Nonfixation for Activated Random Walks

Eric Shellef

The Weizmann Institute of Science, 76100 Rehovot, Israel
E-mail address: `shellef@gmail.com`

Abstract. We consider the activated random walk (ARW) model where particles follow the path of a general Markov process on a general graph. We prove that ARW, regardless of sleep rate, dominates a simpler process, multiple source internal aggregation (MSIA), and use this to formulate a deterministic sufficient condition on initial occupations for nonfixation of ARW and similar variants. In particular, on bounded degree graphs, initial occupation density greater than one almost surely implies nonfixation, where independence requirements are weakened to ergodic in the case of Euclidean lattices. Finally, we prove the critical density for the infinite sleep rate ARW is positive for all dimensions.

1. Introduction

Given a graph, the activated random walks (ARW) process is an interacting system of *active* and *sleeping* particles on the vertices of the graph. Active particles perform independent, rate one, random walks, while sleeping particles stay put. A sleeping particle remains sleeping as long as it is alone at a vertex. Active particles fall asleep independently at a rate $\lambda > 0$. However, if two or more particles occupy the same vertex, they will all be active. Thus a sleeping particle at a vertex v will be activated by an active particle that jumps to v, and an active particle will never fall asleep if it is not alone at a vertex. All particles begin as active from some initial occupation state at time $t = 0$. We call a vertex *fixating* if for some finite time onward, no active particle visits the vertex. A vertex is *nonfixating* if no such time exists. A graph is nonfixating for an initial occupation if all vertices are nonfixating almost surely. Otherwise, it is fixating for that initial occupation.

The main theme explored is the connection between initial particle placement and nonfixation. In particular, if the initial occupation of each vertex is an iid r.v. X with mean μ, is there a phase transition from fixation to nonfixation as we increase μ? Several papers treating this question used $\mathbb{Z}^d$ as a setting and took X to have a Poisson distribution with mean μ. An initial result in this framework can be found in Kesten and Sidoravicius (2006) where a model similar to ARW, in

which sleeping particles are not static, was studied in detail. In Remark 5 of the introduction of Kesten and Sidoravicius (2006), ARW itself with "sleepier" initial conditions is treated, and a general upper bound proven on the critical density for global fixation. Specifically, it is shown that for some large initial density $\mu(d)$, there is a positive probability for a system starting with all particles sleeping except for a single active one to never reach a state where all particles are sleeping, regardless of how high the sleep rate is. In Rolla and Sidoravicius (2009), ARW is studied by Rolla and Sidoravicius, again on $\mathbb{Z}^d$ with X Poisson. A monotonicity theorem proved in the same paper, implies that there is at most one phase transition in the density μ for the ARW process on a graph. An explicit upper bound and a non-trivial lower bound on the critical phase transition density in one dimension are given. In Amir and Gurel-Gurevich (2009), Amir and Gurel-Gurevich show that on transitive graphs satisfying the unimodularity condition, if each vertex is fixating, then eventually all particles fall asleep. They proceed to show that a.s. there is no fixation on such graphs for initial occupation density larger than one (compare with Theorem 3.1).

In this paper we consider ARW in a general form, allowing the particles to follow the path of a general Markov process on the graph. Using a dominating coupling of a simpler finite process (Lemma 2.1), we formulate in Theorem 2.3 a deterministic sufficient condition on initial occupations for nonfixation, regardless of sleep rate λ. We use this condition to prove nonfixation in several scenarios. In Theorem 3.1, we prove that on all bounded degree graphs, an initial occupation density greater than one almost surely implies nonfixation, answering a question posed by Rolla and Sidoravicius (Rolla, 2008; Rolla and Sidoravicius, 2009). The iid assumption on the initial occupation of each vertex can be significantly weakened, and in the special case of $\mathbb{Z}^d$, we show that ergodic initial occupation suffices (Theorem 3.4). In the next section, we prove the critical density for the infinite sleep rate ARW is positive for all dimensions. In the appendix we prove a general lemma for simple random walks on a network that may be of independent interest. The lemma gives a lower bound for the average exit time from a set starting at some vertex using the average number of visits to that vertex before exit.

Note that the nonfixating condition (and thus the results of section 3) holds for a large class of ARW variants, since the only assumption used is that if a vertex is occupied by more than one particle, there must be a time when one of the particles hops from that vertex.

2. Notation and Nonfixation Criterion

Let S be a countable graph, let $\eta_0 : S \to \mathbb{N} \cup \{0\}$ be an initial occupation for a ARW process on S, and let $P(\cdot, \cdot)$ be a transition kernel for a Markov process on S. Each particle is assigned an independent rate one Poisson clock, which together with P, determines the evolution of the particle in the ARW when it is active. We denote the triplet (S, η_0, P) by η, and write $\mathbb{P}^{\eta}_{(\lambda)}[\cdot]$ for the law of the ARW process with kernel P on S with initial occupation η_0 and λ a positive or infinite sleep rate. We write $\lambda = \infty$ for the process where a particle is sleeping if and only if it is alone at a vertex.

The question of existence of the process, and ruling out the possibility that an infinite number of particles reach some vertex in a finite time is not treated here.

See Andjel (1982) for conditions and a construction which can easily be adapted to the ARW process.

Fix a vertex $x \in S$. For an integer r, $\mathcal{A}_r = \mathcal{A}_r(x)$ is the event that x is visited by an active particle at least r distinct times.

We call η *nonfixating* at x if $\mathbb{P}^{\eta}_{(\lambda)}[\mathcal{A}_r(x)] = 1$ for every $r \in \mathbb{N}$ and every positive or infinite λ. We say η is nonfixating if it is nonfixating at all its vertices. From here on, we omit λ from the notation as the results hold regardless of λ.

Given a graph S and kernel P, we define the following multiple source internal aggregation (MSIA) process. We begin with a finite number of indistinguishable "explorers" occupying different vertices of S according to an initial occupation $\gamma_0 : S \to \mathbb{N} \cup \{0\}$. If a vertex contains more than one explorer, then one of the explorers at that vertex begins a walk according to kernel P until reaching an unoccupied vertex, where he remains forever. Assume some fixed ordering of the vertices of S. The explorers perform these walks one at a time by this order. This continues until each vertex has at most one explorer. The law of the random subset of occupied vertices at this stage is not dependent on the order in which explorers are chosen, as was shown in Diaconis and Fulton (1991), where this growth model is described. Moreover, the below domination statement is true for any fixed order. Thus when we discuss an MSIA process we implicitly assume some arbitrary ordering of the vertices.

We identify the triplet (S, γ_0, P) by γ and write $\mathcal{P}^{\gamma}[\cdot]$ for the law of the MSIA process on S with kernel P and initial occupation γ_0. Assume we have two triplets $\eta = (S, \eta_0, P)$ and $\gamma = (S, \gamma_0, P)$ such that $\sum_{x \in S} \gamma_0(x)$ is finite and $\gamma_0 \leq \eta_0$, i.e. there are initially no less particles than explorers for each vertex in S. In Lemma 2.1, which uses ideas found in Lawler et al. (1992); Diaconis and Fulton (1991); Rolla and Sidoravicius (2009), we show that the number of visits by an explorer to a vertex x in γ-MSIA is stochastically dominated by the number of visits of particles to x in η-ARW.

The advantage of this coupling technique is that it assumes only a very basic property of the ARW and is thus valid for many ARW variations, see Remark 2.2.

In the context of the MSIA process, we write $\mathcal{A}_r(x)$ for the event that vertex x is visited by at least r explorers.

Lemma 2.1. *Let* $\eta = (S, \eta_0, P)$ *define an ARW, let* $\gamma = (S, \gamma_0, P)$ *with* $\sum_{x \in S} \gamma_0(x) < \infty$ *define an MSIA, and assume* $\gamma_0 \leq \eta_0$. *Then for any integer* r *and* $x \in S$,

$$\mathcal{P}^{\gamma}[\mathcal{A}_r(x)] \leq \mathbb{P}^{\eta}[\mathcal{A}_r(x)].$$

Proof: We look at the ARW defined by η. Recall that regardless of λ (the rate by which particles fall asleep in the ARW process), if two particles or more occupy the same vertex, they will all be active almost surely. The general consequence is as follows. Let x be a vertex that initially has $\eta_0(x)$ particles on it, and let $m(x)$ be the set of distinct times at which a particle hops *to* x. Writing $\alpha(x)$ for the number of distinct times at which a particle hops *from* x, we have that with probability one, for any x, $\alpha(x) \geq |m(x)| + \eta_0(x) - 1$. For each $x \in S$ label the active particles that will hop from x by the time order of the hop, where the particle to hop first is labeled 1, the second to hop 2 and so on. We now implement the γ-MSIA as a marginal of the η-ARW in a way that each move of each explorer is paired to a

unique jump of a distinct particle. Essentially, when an explorer is at a vertex v we assign to him the first particle that jumped from v that wasn't previously assigned to him or an explorer before him. The explorer then follows that particle. The formal proof is below.

Let $\gamma_i(\cdot)$ where $i \in \mathbb{N} \cup \{0\}$, denote the number of explorers at a vertex after the i'th explorer completed his exploration. Next, let the discrete time variable t start at 0 and count the moves of the explorers, accumulating from one explorer to the next. So if the first explorer made τ moves before stopping, the first move of the second explorer is at time τ. Let $\beta_t(x)$ be a counter for the number of distinct times an explorer has jumped from x up to time t. Thus $\beta_0(x) = 0$ for all x. We now describe the algorithm that chooses the path of the $(i+1)$'th explorer starting at a vertex x with $\gamma_i(x) > 1$ until reaching a vertex y with $\gamma_i(y) = 0$. For time t assume the current explorer is at a vertex v with $\gamma_i(v) > 0$. Then for all $x \in S$ set $\beta_{t+1}(x) = \beta_t(x) + \mathbf{1}_{\{x=v\}}$, and move the explorer to the vertex to which the particle labeled $\beta_{t+1}(v)$ jumped, assuming for now that such a particle exists. Continue thus until the explorer reaches a vertex y with $\gamma_i(y) = 0$. Set $\gamma_{i+1}(x) = \gamma_i(x) - 1$, $\gamma_{i+1}(y) = 1$ and $\gamma_{i+1}(v) = \gamma_i(v)$ for $v \notin \{x, y\}$.

It is clear that each explorer move is paired to a unique jump since we increase $\beta_t(x)$ at every move from x. It is left to show that at the end of the process after T time steps $\beta_T(x) \le \alpha(x)$ for all x and that the MSIA is well defined. $\beta_t(x)$ is incremented is two cases. In the first case, an explorer has just hopped to x from another vertex (at time $t-1$). This implies there is a unique time in $m(x)$ corresponding to this explorer. The second case is that the explorer began his walk at x, which happens $\gamma_0(x) - 1$ times. Since $\alpha(x) \ge |m(x)| + \eta_0(x) - 1$ and $\eta_0 \ge \gamma_0$ we are done. $\qquad\square$

Remark 2.2. The only assumption we used from the ARW is that if a vertex is occupied by more than one particle, there must be a time when one of the particles hops from that vertex.

We introduce some more notation used in the paper. For a set $A \subset S$, let $\eta_0|_A : S \to \mathbb{N} \cup \{0\}$ be equal to η_0 on the set A and 0 outside of it.

When P is implicit, let $X(t)$ denote the discrete time Markov chain according to kernel P.

For a set $A \subset S$, write $\tau_A = \{\inf t \ge 0 : X(t) \in A\}$ for the random time it takes the walk to hit A. If A contains a single vertex x, write τ_x for $\tau_{\{x\}}$.

Let $\{A_m\}_{m \in \mathbb{N}}$ be a rising sequence of sets containing $\mathbf{0}$. When there is no ambiguity, we write τ_m for $\tau_{A_m^c}$ - the first exit time of A_m, and η_m for $\eta_0|_{A_m}$.

Write $p_x^m = P_x[\tau_{\mathbf{0}} < \tau_m]$ for the probability a random walk on S according to the kernel P and starting at $x \in A_m$ hits $\mathbf{0}$ before exiting A_m. Note we use the same P for the kernel and the probability on $X(t)$.

Let $\Lambda_m = \sum_{x \in A_m} p_x^m$ and let $\Omega_m = \sum_{x \in A_m} \eta_0(x) p_x^m$. Write $\Delta_m = \Omega_m - \Lambda_m$ for the difference.

Theorem 2.3. *Given a triplet $\eta = (S, \eta_0, P)$, η is nonfixating at $\mathbf{0} \in S$ if there exists a sequence of sets $\{A_m\}_{m \in \mathbb{N}}$ containing $\mathbf{0}$ such that*

$$\lim_{m \to \infty} \frac{\Delta_m}{\Lambda_m^{1/2}} = +\infty. \tag{2.1}$$

Proof: Since $\Lambda_m \geq p_{\mathbf{0}}^m = 1$, (2.1) implies that $\Delta_m \to \infty$. Writing $\mathcal{P}^m$ for $\mathcal{P}^{(S,\eta_m,P)}$, by Lemma 2.1 it is thus enough to show that for $r(m) = \lfloor \Delta_m/3 \rfloor$,

$$\lim_{m \to \infty} \mathcal{P}^m \left[\mathcal{A}_{r(m)} \right] = 1. \tag{2.2}$$

That is, we show that the probability $\mathbf{0}$ will be visited at least $\lfloor \Delta_m/3 \rfloor$ times by an explorer in the MSIA process defined by (S, η_m, P) goes to one with m.

The condition in (2.1) also implies the stronger condition that

$$\lim_{m \to \infty} \frac{\Delta_m}{\Omega_m^{1/2}} = \infty. \tag{2.3}$$

To see this, assume first that $\Delta_m > \Lambda_m$ in which case $\Delta_m^2/\Omega_m = \Delta_m^2/(\Delta_m + \Lambda_m) > \Delta_m/2$. Otherwise $\Delta_m^2/\Omega_m = \left(\Delta_m^2/\Lambda_m \right) \cdot \Lambda_m/(\Lambda_m + \Delta_m) \geq \Delta_m^2/2\Lambda_m$. By assumption, both Δ_m and Δ_m^2/Λ_m go to infinity with m, and we get (2.3).

Fix $k > 0$ and choose m large enough so that $\Delta_m > 3k\Omega_m^{1/2}$. Using the idea from the original IDLA paper Lawler et al. (1992), we let the explorers positioned by η_m start the walks one at a time but assume that once an explorer reaches an unoccupied vertex and remains there, his ghost continues to walk forever. Thus to each explorer we associate a *walk* that begins together with the explorer but continues indefinitely. Let $W = W(m)$ be the number of walks that visit $\mathbf{0}$ before exiting A_m, and let $L = L(m)$ be the number of walks that visit $\mathbf{0}$ before exiting A_m, but do this as ghosts (i.e. after stopping in the original model). $W - L$ thus counts the explorers that visit $\mathbf{0}$ before stopping and before exiting A_m. Letting $F(m)$ be the event $\{W - L < \Delta_m/3\}$, to prove (2.2) it is enough to show that the probability of $F(m)$ goes to zero with m. We write $\mathcal{P}[\cdot] = \mathcal{P}(m)[\cdot]$ for the law of the MSIA with ghosts.

W is a sum of independent variables with mean Ω_m. $E[L]$ is hard to calculate, but note that each ghost that contributes to L can be tied to the unique point at which it turns from an explorer into a ghost. Thus, by the Markov property, if we start an independent walk from each vertex in A_m and let $\hat{L}$ be the number of such walks that hit $\mathbf{0}$ before exiting A_m, we have $\mathcal{P}[L \geq a] \leq \mathcal{P}[\hat{L} \geq a]$.

In particular, $E[L] \leq E[\hat{L}] = \Lambda_m$ and we have that $E[W] - E[L] \geq \Delta_m$. Thus we upperbound $F(m)$ with the union bound

$$\mathcal{P}\left[\Omega_m - W > \Delta_m/3 \right] + \mathcal{P}\left[L - \Lambda_m > \Delta_m/3 \right]. \tag{2.4}$$

To get a bound on the probability of deviation from the mean, note that each of W and $\hat{L}$ is a sum of independent indicators and thus their variances are bounded by their means. By this and the Chebyshev inequality,

$$\mathcal{P}\left[\Omega_m - W > \Delta_m/3 \right] \leq \mathcal{P}\left[\Omega_m - W > k\Omega_m^{1/2} \right] \leq \mathcal{P}\left[\Omega_m - W > k\sigma_W \right] \leq k^{-2}$$

and similarly

$$\mathcal{P}\left[L - \Lambda_m > \Delta_m/3 \right] \leq \mathcal{P}\left[\hat{L} - \Lambda_m > \Delta_m/3 \right] \leq k^{-2}.$$

Since k was arbitrary we are done. $\square$

3. Nonfixation for Random Initial Occupation

The above theorem treats the case of a fixed initial occupation. From here on, we take η_0 to be random, using boldcase $\mathbf{P}$ and $\mathbf{E}$ to denote the probability and expectation corresponding to the law of the initial occupation.

For the formulation of below theorem, it is simpler to view S as an infinite connected network, i.e. a graph where each edge $e \in E(S)$ is assigned a conductance $c(e)$. For a vertex $x \in S$, let $\mathcal{E}(x)$ be the edges in S that have x as an endpoint. Let $\pi(x) = \sum_{e \in \mathcal{E}(x)} c(e)$. We call $\pi(x)$ the weight of x. Let P be the transition kernel of a simple random walk (SRW) on the network S. That is, unlike SRW on a graph in which the transition probability from a vertex x to a neighbor is the reciprocal of the degree of x, in SRW on a network, the probability to move from x to a neighbor connected by edge e is $c(e)/\pi(x)$.

We say S is γ-bounded if there is a constant $\gamma > 0$ such that $\gamma < \pi(x), c(e) < \gamma^{-1}$ for each vertex x and edge e of S. A SRW on a γ-bounded network includes diverse examples such as SRW on a bounded degree graph and a large class of bounded range walks on transitive graphs.

Theorem 3.1. *Let S be a γ-bounded network, and let P be the kernel of an SRW on S. Let $\{\eta_0(x)\}_{x \in S}$ be uncorrelated r.v.'s, all with uniformly bounded variance $V < \infty$ and mean $1 + \epsilon$ for some $\epsilon > 0$. Then $\mathbf{P}$-almost surely , $\eta = (S, \eta_0, P)$ is nonfixating.*

Proof: First we show that without any assumptions on S or P, if there is a vertex $\mathbf{0} \in S$ and a sequence of sets $\{A_m\}_{m \in \mathbb{N}}$ containing $\mathbf{0}$ such that $\Lambda_m \to \infty$, then $\eta = (S, \eta_0, P)$ is nonfixating at $\mathbf{0}$, $\mathbf{P}$-almost surely.

We may assume (by taking a subsequence) that $\{A_m\}_{m \in \mathbb{N}}$ is such that $\Lambda_m > m^2$. By assumptions the variance of $\Omega_m = \sum_{x \in A_m} \eta_0(x) p_x^m$ is bounded by $V \sum_{x \in A_m} (p_x^m)^2 \le V \Lambda_m$. Thus by Chebyshev,

$$\mathbf{P}\left[|\mathbf{E}\left[\Omega_m \right] - \Omega_m| > \frac{\epsilon}{2} \Lambda_m \right] < \frac{4V}{\epsilon^2 \Lambda_m} < cm^{-2}.$$

By Borel Cantelli, $\mathbf{P}$-almost surely, for all large enough m, Ω_m is not far from its mean and we have for all large enough m,

$$\Delta_m = \Omega_m - \Lambda_m > \frac{\epsilon}{2} \Lambda_m. \tag{3.1}$$

Since $\Lambda_m \to \infty$ the condition in Theorem 2.3 holds and we are done.

Back to our γ-bounded network S, by above it is enough to show that for an arbitrary vertex $\mathbf{0} \in S$, and $A_m = \{x \in S : d_S(x, \mathbf{0}) < m\}$, where $d_S(\cdot, \cdot)$ is graph distance, we have $\Lambda_m \to \infty$. For $x, y \in S$ and $m \in \mathbb{N}$, define the Green's function as

$$G_m(x, y) = \mathbb{E}_x \left[\sum_{t=0}^{\tau_m - 1} 1_{\{X(t) = y\}} \right]$$

where $X(t)$ is our discrete time SRW as defined above. By standard Markov chain theory,

$$p_x^m = G_m(x, \mathbf{0})/G_m(\mathbf{0}, \mathbf{0}).$$

Next, for a simple random walk on a network (see e.g. chapter 2 in Lyons and Peres, 2010),

$$G_m(x, \mathbf{0})\pi(x) = G_m(\mathbf{0}, x)\pi(\mathbf{0}).$$

By our assumption that $\pi(\cdot)$ is uniformly bounded away from zero and infinity on S, we may sum over A_m to get that for some $c > 0$,

$$\Lambda_m = \sum_{x \in A_m} p_x^m \ge cG_m(\mathbf{0}, \mathbf{0})^{-1} E\left[\tau_m \right]. \tag{3.2}$$

Using the notation of Lemma 6.1, if we take $Z = S \backslash A_m$ and $x = \mathbf{0}$, we have $G_Z = G_m(\mathbf{0}, \mathbf{0})$. Note that for any $m > 1$, A_m contains all the neighbors of $\mathbf{0}$ and the conditions of the lemma are satisfied. By (6.1) we have

$$G_m(\mathbf{0}, \mathbf{0}) < (E[\tau_m] \log G_m(\mathbf{0}, \mathbf{0})/k(\gamma))^{\frac{1}{2}}.$$

Since $a > \log a$ for any positive a, again from (6.1) we get $G_m(\mathbf{0}, \mathbf{0}) < E[\tau_m]/k(\gamma)$ and plugging into above we have

$$G_m(\mathbf{0}, \mathbf{0}) \leq \left(E[\tau_m] \log \frac{E[\tau_m]}{k}/k \right)^{1/2}.$$

Since $E[\tau_m] \to \infty$ we get that the right hand side of (3.2) goes to infinity and are finished. $\qquad\square$

Remark 3.2. Note that some uniform bound on weights and edges of S is needed to get nonfixation with the conditions of above theorem. As an example for this, we could take $\mathbb{N}$ as the graph and set the conductances such that the the probability for a SRW to ever hit 1 is summable on $\mathbb{N}$. Thus any initial occupation distributed like η_0 in above theorem would fixate at 1 a.s. as can be seen by bounding the expected number of particles to hit 1.

Remark 3.3. If $\{\eta_0(x)\}_{x \in S}$ are iid with mean greater than one, we don't need the finite variance assumption since for some $M < \infty$, $\eta_0(x) \wedge M$ has mean greater than one and all moments, thus η_0 is nonfixating by Theorem 3.1.

Theorem 3.4. *Let $S = \mathbb{Z}^d$ and let P be the transition kernel of a simple random walk on S. Let $\{\eta_0(x)\}_{x \in \mathbb{Z}^d}$ be distributed such that the action of the group of translations is ergodic and such that $\mu = \mathbf{E}[\eta_0(\mathbf{0})] = 1 + \epsilon$ for some $\epsilon > 0$. Then $\mathbf{P}$-almost surely, $\eta = (S, \eta_0, P)$ is nonfixating.*

Proof: Let $A_m = \{x \in \mathbb{Z}^d : \|x\|_2 < m\}$. As in Theorem 3.1, $p_x^m = G_m(x, \mathbf{0})/G_m(\mathbf{0}, \mathbf{0})$. Since $\mathbb{Z}^d$ is regular, the Green function is symmetric and we get $\Lambda_m = E[\tau_m]/G_m(\mathbf{0}, \mathbf{0})$ which tends to infinity by Lemma 6.1 as shown in Theorem 3.1. Thus by Theorem 2.3 it is enough to show that $\mathbf{P}$-almost surely,

$$\Omega_m/\Lambda_m \to \mu. \tag{3.3}$$

We prove for $d > 2$. The same proof works for $d = 2$ using the estimate $G_m(\mathbf{0}, \mathbf{0}) = \frac{2}{\pi} \ln n + O(n^{-1})$ (see Proposition 1.6.7 of Lawler, 1991), and for $d = 1$ where $G_m(\mathbf{0}, \mathbf{0}) = m$. By the optional stopping theorem with the martingale $\|X(t)\|^2 - t$, $\mathbb{E}[\tau_m] = m^2 + o(m^2)$. Second, $G_m(\mathbf{0}, \mathbf{0})^{-1} \to \sigma$ where $\sigma(d)$ is the escape probability from $\mathbf{0}$. σ is positive since $d > 2$. Note that in the $d = 1, 2$ cases, $G_m(\mathbf{0}, \mathbf{0})^{-1}$ needs to be multiplied by m and $\ln m$ respectively in order to converge to a positive value. Continuing, we have

$$\Lambda_m/m^2 \to \sigma > 0. \tag{3.4}$$

Next, define the average $Q_m = |A_m|^{-1} \sum_{x \in A_m} \eta_0(x)$. By Theorem 1.2 in Lindenstrauss (2001), since $\mathbb{Z}^d$ is amenable and A_m is a tempered Følner sequence, we have the following law of large numbers

$$Q_m \xrightarrow{\mathbf{P}-a.e.} \mu. \tag{3.5}$$

For $x \in \mathbb{Z}^d$, write $\lceil x \rceil$ for the smallest integer strictly greater than the Euclidean norm of x. Thus $\lceil \mathbf{0} \rceil = 1$. Let $c_d = \frac{2}{d-2}\sigma\omega_d^{-1}$ where ω_d is the volume of the unit sphere in $\mathbb{R}^d$. For $k \in \mathbb{N}$ let

$$q_k^m = c_d \left[k^{2-d} - m^{2-d} \right].$$

and for $x \in \mathbb{Z}^d$ let $q_x^m = q_{\lceil x \rceil}^m$. By Proposition 1.5.9 in Lawler (1991),

$$p_x^m = q_x^m + O(\lceil x \rceil^{1-d}).$$

Summing on level sets, and using that for $k \in \mathbb{N}$, $k^{1-d} - (k+1)^{1-d} < ck^{-d}$, we get

$$\sum_{x \in A_m} \eta_0(x) \left[p_x^m - q_x^m \right] \quad < \quad cQ_m \left| A_m \right| m^{1-d} + c \sum_{k=1}^{m-1} Q_k \left| A_k \right| k^{-d}$$

$$< \quad c \left[Q_m m + \sum_{k=1}^{m-1} Q_k \right].$$

Since Λ_m is order of m^2 then by (3.5)

$$\sum_{x \in A_m} \eta_0(x) \left[p_x^m - q_x^m \right] / \Lambda_m \xrightarrow{a.e.} 0,$$

and to calculate the limit in (3.3), we can replace p_x^m in $\Omega_m = \sum_{x \in A_m} \eta_0(x) p_x^m$ by q_x^m. Since $q_m^m = 0$, we can write

$$\sum_{x \in A_m} \eta_0(x) q_x^m \quad = \quad \sum_{k=1}^{m-1} Q_k \left| A_k \right| \left[q_k^m - q_{k+1}^m \right]$$

$$= \quad \sum_{k=1}^{m-1} Q_k \omega_d \left[k^d + o(k^d) \right] c_d (d-2) \left[k^{1-d} + o(k^{1-d}) \right]$$

$$= \quad 2\sigma \sum_{k=1}^{m-1} Q_k \left[k + o(k) \right].$$

Applying (3.5) we get that

$$\sum_{x \in A_m} \eta_0(x) q_x^m / \sigma m^2 \xrightarrow{a.e.} \mu.$$

$\square$

4. Fixation for infinite sleep rate

In this section we prove that for the ARW process with $\lambda = \infty$, if a deterministic initial occupation satisfies a "density bound" condition around a fixed vertex, only a finite number of distinct particles ever visit this vertex, almost surely. We then use this to show there are nontrivial iid distributions on $\mathbb{Z}^d$ (including Poisson iid) that satisfy this condition, hence showing together with Theorem 3.1 and the monotonicity result in Rolla and Sidoravicius (2009), that the critical density for $\lambda = \infty$ is positive for all dimensions.

Let $\eta = (S, \eta_0, P)$ define an ARW process. Fix $\mathbf{0} \in S$. For a set of vertices $A \subset S$, let the weight of A be $w(A) = \sum_{x \in A} \eta_0(x)$. Let $W(n)$ denote the maximal weight among all connected vertex sets of size n that include $\mathbf{0}$. Let

$$\mathcal{A}(S, \eta_0) = \sup \{ n \in \mathbb{N} : W(n) \geq n \}.$$

For the process defined by η, let $\mathcal{C}(T)$ be the random set of vertices visited by any one of the particles by time $T \in [0, \infty]$. Let $\mathcal{C}_{\mathbf{0}}(T)$ be the (possibly empty) connected component of $\mathcal{C}(T)$ containing $\mathbf{0}$.

Theorem 4.1. *For the ARW process with infinite sleep rate on $\eta = (S, \eta_0, P)$, where P is such that only nearest neighbor moves are allowed,*

$$\mathbb{P}^{\eta}_{(\infty)} [|\mathcal{C}_{\mathbf{0}}(\infty)| \leq \mathcal{A}(S, \eta_0)] = 1. \tag{4.1}$$

In particular, since any particle that visits $\mathbf{0}$ is in $\mathcal{C}_{\mathbf{0}}(\infty)$, and we assume η_0 is locally finite, if $\mathcal{A} < \infty$ then only a finite number of distinct particles ever visit $\mathbf{0}$ almost surely.

Proof: Since $\lambda = \infty$, every vertex of $\mathcal{C}_{\mathbf{0}}(T)$ is occupied by a distinct particle at time T. Next, note that since P only allows nearest neighbor jumps, any particle in $\mathcal{C}_{\mathbf{0}}(T)$ at time T was in $\mathcal{C}_{\mathbf{0}}(T)$ at time 0. Hence w.p. 1,

$$w(\mathcal{C}_{\mathbf{0}}(T)) \geq |\mathcal{C}_{\mathbf{0}}(T)|.$$

So if $\mathcal{C}_{\mathbf{0}}(T)$ is finite then by definition of $\mathcal{A}$, $|\mathcal{C}_{\mathbf{0}}(T)| \leq \mathcal{A}$ w.p. 1. We assume $\mathcal{A} < \infty$ as otherwise (4.1) holds trivially. We calculate the probability for the local event $\{|\mathcal{C}_{\mathbf{0}}(T)| \leq \mathcal{A}\}$ as a limit of its probability with initial occupation η_0 replace by $\eta_m = \eta_0|_{A_m}$ where $A_m \nearrow S$ are finite sets increasing to S. As always, we assume that η_0 and P generate a well defined ARW process on S and any local event measurable up to a finite time can be finitely approximated. For all m and all $T < \infty$

$$\mathbb{P}^{(S, \eta_m, P)}_{(\infty)} [|\mathcal{C}_{\mathbf{0}}(T)| \leq \mathcal{A}(S, \eta_0)] = 1.$$

Thus we get (4.1) for $\mathcal{C}_{\mathbf{0}}(T)$ for any $T < \infty$ and hence for $\mathcal{C}_{\mathbf{0}}(\infty)$ by monotonic convergence. $\square$

Remark 4.2. The combinatorical nature of the proof allows extension of the result from ARW to a controllable process in which an adversary attempts to bring infinitely many particles to $\mathbf{0}$ while observing the rule that a particle may be moved only if it is not alone at a vertex. Let $D(S, \eta_0) = \limsup_{n \to \infty} W(n)/n$. In the setting of the controllable process, $D < 1$ implies finitely many visits while $D > 1$ allows for infinitely many visits, thus in a sense the condition is sharp.

Below we prove there exist nontrivial iid distribution on $\mathbb{Z}^d$ for which $\mathcal{A}(\mathbb{Z}^d, \eta_0)$ is finite almost surely. For one dimension fixation in this setting was already known and proven in Rolla and Sidoravicius (2009).

Corollary 4.3. *Let $S = \mathbb{Z}$ and let $\{\eta_0(x)\}_{x \in \mathbb{Z}}$ be an ergodic distribution on $\mathbb{N} \cup \{0\}$ with mean less than one. Let P be a transition kernel on S allowing only nearest neighbor jumps. Then for the ARW process with $\lambda = \infty$, $\mathbf{P}$-almost surely, $\eta = (S, \eta_0, P)$ is fixating.*

Proof: By the law of large numbers, $W(n) < n$, for all large enough n almost surely. Thus $\mathbf{P}$-almost surely, $\mathcal{A}(\mathbb{Z}^d, \eta_0)$ is finite and we have fixation. $\square$

Corollary 4.4. *Let $S = \mathbb{Z}^d$ for $d \geq 2$, and let $\mathcal{D}$ be a probability mass function on $\mathbb{N} \cup \{0\}$ such that for some constant $c > 0$,*

$$\sum_{n=0}^{\infty} \left(1 - \sum_{i=0}^{n} \mathcal{D}(i) \right)^{1/d} < c.$$

Let P be a transition kernel on S allowing only nearest neighbor jumps. Set the initial occupation $\{\eta_0(x)\}_{x \in \mathbb{Z}^d}$ to be iid r.v.'s with distribution $\mathcal{D}$. Then for the ARW process with $\lambda = \infty$, $\mathbf{P}$-almost surely, $\eta = (S, \eta_0, P)$ is fixating. In particular, when $\lambda = \infty$, the critical density for the Poisson iid initial occupation is positive for all dimensions.

Proof: Theorem 1 in Martin (2002) (see also Cox et al., 1993) implies that with some $c > 0$, for any distribution satisfying above condition, we almost surely have $W(n) < n$ for all large enough n. Thus $\mathbf{P}$-almost surely, $\mathcal{A}(\mathbb{Z}^d, \eta_0)$ is finite and we have fixation.

By Proposition 8.1 in Martin (2002) for some c',

$$\sum_{n=0}^{\infty} \left(1 - \sum_{i=0}^{n} \mathcal{D}(i) \right)^{1/d} < c' \mathbf{E} \left[\mathcal{D}^{d+1} \right].$$

For a Poisson distribution with intensity γ, the right hand side of above tends to 0 as γ approaches 0 and we are done. □

5. Further Remarks and Questions

(1) There are many variations which can be analyzed using the above framework. For example, given a function $f : S \to \mathbb{N}$ from a graph to the naturals, let $\mathrm{ARW}(f)$, be the ARW process where there must be more than $f(x)$ particles at a vertex x to know they are all active. Thus the usual process is $\mathrm{ARW}(\mathbf{1})$. Lemma 2.1 and Theorem 2.3 can be modified by redefining MSIA and Λ_m according to f, and similar theorems on nonfixation for fixed or random initial occupations can be proved.

(2) What can be shown if the random walks of particles are dependent? (e.g. an exclusion process) A review of the proof of Theorem 2.3 shows we use a property of independent indicators that could arise in weakly dependent situations as well. Namely that the variance of W and $\hat{L}$ is the same order as their mean. A different relationship between these quantities could imply a similar theorem with an appropriate update of (2.1).

(3) Is there an example of a graph which has nonfixation for iid initial occupation with mean smaller than one for infinite sleep rate? For fixed finite sleep rate, can one find the exact critical density for some graph with iid initial occupation? What are achievable values for critical density for Cayley graphs? For general graphs?

(4) The model is not an attractive particle system and thus it is not clear whether there exist stationary distributions for nonfixating scenarios or what could be possible candidates. Can one find a graph with a stationary distribution for the ARW?

(5) When fixation occurs on a graph, what can we say about the speed with which vertex fixation spreads?

For more questions, conjectures and numerical results, see Dickman et al. (2010).
Acknowledgment: Thanks to Gady Kozma and Vladas Sidoravicius for telling me about this problem and for very helpful discussions and ideas. Thanks to Itai Benjamini for suggesting some more directions and to a referee for a careful reading which helped weed out problems in the original manuscript.

6. Appendix

Here we prove a general lemma for simple random walks on a network relating the average exit time from a set starting at some vertex to the average number of visits to that vertex before exit. Let S be a network, let $Z \subset S$ be a set of vertices and let τ_Z be the first hitting time of Z for $X(t)$ the discrete time simple random walk on S. For $x \in S$ we set $G_Z(x) = E_x \left[\sum_{t=0}^{\tau_Z} \mathbf{1}_{\{X(t)=x\}} \right]$, the expected number of visits to x of a walk starting at x before τ_Z. Let $B_r(x) = \{v \in S : d_S(v,x) < r\}$ and let $\partial B_r(x) = \{v \in S : d_S(v,x) = r\}$ where $d_S(\cdot,\cdot)$ is graph distance in S.

Lemma 6.1. *Assume S is an infinite γ-bounded connected network (i.e. there is a $\gamma > 0$ for which $\gamma < c(e), \pi(x) < \gamma^{-1}$ for every vertex x and edge e in S). Then there is a $k = k(\gamma) > 0$ such that for any $x \in S$ and $Z \subset S$ where $B_2(x) \cap Z = \emptyset$.*

$$E_x\left[\tau_Z\right] > kG_Z^2 / \log G_Z. \tag{6.1}$$

Proof: Fix $x \in S$, set $T_0 = 0$, and define for each $i \in \mathbb{N}$ the r.v.'s

$$T_i = \inf \left\{ t > T_{i-1} : X(t) = x \right\}.$$

Let $i^* = \inf \{i : T_i = \infty\}$. For $1 \le i \le i^*$ let $\rho_i = T_i - T_{i-1}$. We show there are positive constants k_1, k_2 dependent only on γ such that

$$P\left[\rho_1 \ge k_1 r^2 / \log r\right] \ge \frac{k_2}{r}. \tag{6.2}$$

By electrical network interpretation (see e.g. Lyons and Peres, 2010), the probability for a walk beginning at x to hit $\partial B_r(x)$ before returning to x is $C_{eff}(r)/\pi(x)$, where $C_{eff}(r)$ is the effective conductance from x to $\partial B_r(x)$. Since S is infinite and connected, for any r there is a connected path of r edges from x to $\partial B_r(x)$. By the monotonicity principle, $C_{eff}(r)$ is at least the conductance on this path, which is γr^{-1}. Thus the probability to hit some $y \in \partial B_r(x)$ before returning to x is at least $\gamma^2 r^{-1}$.

Next, let $y \in \partial B_r(x)$. By the Carne-Varopoulos upper bound (see Varopoulos, 1985),

$$P\left[X(t) = x | X(0) = y\right] \le 2 \left(\pi(y)/\pi(x)\right)^{1/2} \exp\left(-\frac{r^2}{2t}\right)$$

and thus, for some $k_1(\gamma), k_2(\gamma) > 0$ and all $r > 1$, the probability that a walk starting at $y \in \partial B_r(x)$ does not hit x in the next $\lfloor k_1 r^2 / \log r \rfloor$ steps, by union bound, is greater than k_2. Together with our lower bound on the probability that we arrive at such a $y \in \partial B_r(x)$, we get (6.2).

Next, let $g = \inf \{i : T_i > \tau_Z\}$ be the number of visits of $X(t)$ to x before hitting Z, including $t = 0$. g is a geometric random variable with mean $G = G_Z$. Let $\alpha = \frac{1}{2} \ln (4/3)$ and note that since there is a constant in (6.1) and $G \ge 1 + \gamma^4$, as

$P[X(2) = x|X(0) = x] \geq \gamma^4$, we can assume that $G > 2$. Thus

$$\begin{aligned} P\left[g \geq \alpha G\right] &\geq \left(1 - G^{-1}\right)^{\alpha G} \\ &\geq \left(1 - G^{-1}\right)^{2\alpha(G-1)} \\ &\geq e^{-2\alpha} = 3/4 \end{aligned}$$

We further assume $G > \frac{2}{\alpha} \vee \frac{16}{k_2}$ so that $\alpha G - 1 > \alpha G/2$ and $G\frac{k_2}{16} > 1$. Let A be the event that there is an $1 \leq i \leq (\alpha G - 1) \wedge i^*$ such that $\rho_i > k_1 \left(\frac{k_2}{16}G\right)^2 / \log\left(\frac{k_2}{16}G\right)$. Note that $i^* \leq \alpha G - 1$ implies A. Thus by (6.2) and the independence of consecutive excursions from x,

$$P\left[A^c\right] \leq \left(1 - \frac{16}{G}\right)^{\alpha G/2} \leq e^{-8\alpha}$$

which is smaller than $1/4$.

Thus

$$P\left[\{g \geq \alpha G\}, A\right] \geq \frac{1}{2}.$$

This implies the lemma since $\tau_Z > \sum_{i=1}^{g-1} \rho_i$, and for $k = k_1 k_2^2/256$, $\sum_{i=1}^{g-1} \rho_i > kG^2/\log G$. $\qquad\square$

References

G. Amir and O. Gurel-Gurevich. On fixation of activated random walks. *ArXiv Mathematics e-prints* (2009). arXiv:math/0910.3730v1.

E. D. Andjel. Invariant measures for the zero range processes. *Ann. Probab.* **10** (3), 525–547 (1982). MR659526.

J. T. Cox, A. Gandolfi, P. S. Griffin and Harry Kesten. Greedy lattice animals. I. Upper bounds. *Ann. Appl. Probab.* **3** (4), 1151–1169 (1993). MR1241039.

P. Diaconis and W. Fulton. A growth model, a game, an algebra, Lagrange inversion, and characteristic classes. *Rend. Sem. Mat. Univ. Politec. Torino* **49** (1), 95–119 (1993) (1991). Commutative algebra and algebraic geometry, II (Italian) (Turin, 1990); MR1218674.

R. Dickman, L. T. Rolla and V. Sidoravicius. Activated random walkers: Facts, conjectures and challenges. *J. Statist. Phys.* **138** (1-3), 126–142 (2010).

H. Kesten and V. Sidoravicius. A phase transition in a model for the spread of an infection. *Illinois J. Math.* **50** (1-4), 547–634 (electronic) (2006). MR2247840.

G. F. Lawler. *Intersections of random walks.* Probability and its Applications. Birkhäuser Boston Inc., Boston, MA (1991). ISBN 0-8176-3557-2. MR1117680.

G. F. Lawler, M. Bramson and D. Griffeath. Internal diffusion limited aggregation. *Ann. Probab.* **20** (4), 2117–2140 (1992). MR1188055.

Elon Lindenstrauss. Pointwise theorems for amenable groups. *Invent. Math.* **146** (2), 259–295 (2001). MR1865397.

R. Lyons and Y. Peres. *Probability on trees and networks.* Cambridge University Press (2010). Book in progress; prbtree.html.

J. B. Martin. Linear growth for greedy lattice animals. *Stochastic Process. Appl.* **98** (1), 43–66 (2002). MR1884923.

L.T. Rolla. Generalized hammersley process and phase transition for activated random walk models (2008).

L.T. Rolla and V. Sidoravicius. Absorbing-state phase transition for stochastic sandpiles and activated random walks. *ArXiv Mathematics e-prints* (2009). arXiv:math/0908.1152.

N. Th. Varopoulos. Long range estimates for Markov chains. *Bull. Sci. Math. (2)* **109** (3), 225–252 (1985). MR822826.