

Sensitivity of the frog model to initial conditions

Tobias Johnson* Leonardo T. Rolla†

Abstract

The frog model is an interacting particle system on a graph. Active particles perform independent simple random walks, while sleeping particles remain inert until visited by an active particle. Some number of sleeping particles are placed at each site sampled independently from a certain distribution, and then one particle is activated to begin the process. We show that the recurrence or transience of the model is sensitive not just to the expectation but to the entire distribution. This is in contrast to closely related models like branching random walk and activated random walk.

Keywords: frog model; recurrence; transience; branching random walk; activated random walk.

AMS MSC 2010: 60K35; 60J80; 60J10.

Submitted to ECP on September 10, 2018, final version accepted on April 12, 2019.

Supersedes arXiv:1809.03082.

1 Introduction

The frog model is an interacting particle system on a graph. Initially, one designated site contains an active particle, and all other sites contain some number of sleeping particles, typically sampled independently from a given distribution. Active particles perform simple random walks; these are generally taken to be in discrete time, though it is irrelevant to this paper. When an active particle visits a site, all sleeping particles there are activated. For no deep reason, the particles have come to be called frogs. We represent a frog model as a pair (S, η) where $\eta(v)$ gives the count of sleeping particles on a vertex v , and $S_t(v, i)$ gives the path of the i th particle on vertex v for each $1 \leq i \leq \eta(v)$.

We call a realization of the frog model *recurrent* if the starting site is visited infinitely often by particles and *transient* if not. The first question for the frog model on a given infinite graph is whether it is recurrent or transient. On $\mathbb{Z}^d$, the frog model is recurrent a.s. if the initial configuration $(\eta(v))_{v \in \mathbb{Z}^d}$ is i.i.d., so long as $\eta(v)$ is not deterministically equal to zero [1]. On the other hand, the frog model on the infinite d -ary tree can be either transient a.s. or recurrent a.s., depending on the initial configuration. For example, on the d -ary tree when $(\eta(v))_v$ is i.i.d.-Poisson(μ), the frog model is recurrent or transient depending on whether μ is greater or less than a critical value $\mu_c(d)$ [5, 7]. In [8], the authors give a theorem comparing frog models with different initial conditions on the same graph, which shows that the frog model on the d -ary tree is recurrent if $\eta(v) \succeq_{\text{pgf}} \text{Poisson}(\mu)$ for all v for some $\mu > \mu_c(d)$. The condition $X \succeq_{\text{pgf}} Y$ means that

*College of Staten Island, United States of America. E-mail: tobias.johnson@csi.cuny.edu

†Argentina National Research Council at the University of Buenos Aires, Argentina; NYU-ECNU Institute of Mathematical Sciences at NYU Shanghai, China. E-mail: leorolla@dm.uba.ar

$\mathbf{E}t^X \leq \mathbf{E}t^Y$ for all $t \in (0, 1)$, which roughly speaking requires $\eta(v)$ to have expectation at least μ and be less dispersed than the Poisson distribution.

This raises the question of whether the transience and recurrence of the frog model depends on the entire distribution of each $\eta(v)$ or just on the expectation. This question was posed in [8, Open Question 11] with the conjecture that the entire distribution matters. We confirm this:

Theorem 1.1. *Consider the frog model on the d -ary tree with i.i.d.- π initial conditions. For arbitrarily large μ , there exists a distribution π with mean μ so that the model is a.s. transient.*

Thus, for large enough μ , there exist distributions π with mean μ so that the frog model on the d -ary tree is recurrent or so that it is transient. In §2 we prove Theorem 1.1. In §3 we describe how the argument can be extended to show transience for some initial distributions with infinite mean and (for $d \geq 14$) even when every site contains at least one particle.

For two processes resembling the frog model, the long-time behavior of the model depends only on the mean particle distribution. The first of these is *branching random walk*, essentially the frog model except that particles spawn new particles even when moving to a previously visited site. It is a classical result of Biggins's that for BRW, recurrence vs. transience depends only on the mean number of particles spawned. We explain this in more detail in Remark 2.2.

The other process with contrasting behavior is *activated random walk*. Particles in this model move with the same dynamics as the frog model except that active particles have some probability of falling back asleep at each step. This process is usually considered with all particles starting awake and moving in continuous time. Both ARW and the frog model have been of significant interest in physics. The frog model and the broader family of $A + B \rightarrow 2B$ models that it falls into have been viewed as stochastic combustion models. They have also been investigated as part of the general study of propagating fronts; see Section 2.6.1 (i)–(v) in [9] for a survey of the physics literature. The interest in ARW comes from a connection with the phenomenon of self-organized criticality. For a given sleep rate, ARW has been shown on many graphs to undergo an absorbing-state phase transition; see the introductions of [2, 10, 12, 13]. In [11], it is shown that the phase transition for this model on a regular tree with all particles initially active depends only on the initial density of particles, not on the distribution of particles at each site (it is stated for $\mathbb{Z}^d$ but the argument also works for unimodular graphs). On the other hand, by Theorem 1.1 the same model starting with all but one particle sleeping can fixate with arbitrarily high densities. This suggests that the question of fixation vs. activity for the ARW is more delicate than it may seem, at least outside the setup of graphs with polynomial growth.

2 Proofs

We consider the $(d + 1)$ -regular tree, denoted $\mathbf{T}_d$, and in the end we translate the result from $\mathbf{T}_d$ to the d -ary tree. Let $\rho \in \mathbf{T}_d$ be the designated vertex starting with one active particle, which we will call the root. Fix a choice of μ , and let

$$\pi_N := \frac{\mu}{N} \delta_N + \left(1 - \frac{\mu}{N}\right) \delta_0.$$

From now on, we consider the frog model on $\mathbf{T}_d$ with i.i.d.- π_N initial conditions. The initial configuration can be thought of as a sea of empty sites, with islands on which N particles sleep. Our goal is to show that for large enough N , this frog model is transient.

The typical way of proving transience for the frog model on a tree is to add extra particles to the model so that a particle *always* wakes up particles when it moves. The




Figure 1: A neighborhood of the root ρ in $\mathbf{T}_2$, organized by levels.

resulting process is a branching random walk, which can then be shown to be transient when the distribution of sleeping particles is sufficiently small. See [5, Proposition 15] for the basic example of this argument, or see [6, Section 3.2] for a more elaborate one. The recurrence or transience of the branching random walk produced by this argument depends only on the expected number of particles per site (see Remark 2.2). Thus, we will need a different approach, since for large values of μ this branching random walk will always be recurrent.

Our argument instead uses a different branching process not indexed by time in the usual way. Rather than track individual particles and branch at each site they visit, we branch at each site visited over all time by the batch of particles starting from the same island. The advantage is that if many particles starting from the same island visit the same site, the process branches only once at that site.

We define a function $\ell: \mathbf{T}_d \rightarrow \mathbb{Z}$ where $\ell(v)$ represents the “level” of vertex v as follows. We set $\ell(\rho) = 0$, and then specify that for each $v \in \mathbf{T}_d$, we have $\ell(u) = \ell(v) - 1$ for one neighbor u of v and $\ell(u) = \ell(v) + 1$ for all other neighbors u (see Figure 1). Note that $\ell(v)$ is not the same as the distance between v and ρ .¹ Fix $\lambda > 0$. We define the weight function

$$w_\lambda(A) = \sum_{v \in A} \lambda^{\ell(v)}, \quad A \subseteq \mathbf{T}_d.$$

Most of our work will be in the following lemma.

Lemma 2.1. *Suppose a random number of particles distributed as π_N start at ρ and perform independent random walks on $\mathbf{T}_d$. Let A be the random (and possibly empty) set of sites visited by these particles. Then, for $\lambda = \frac{1}{\sqrt{d}}$ and N sufficiently large, $\mathbb{E}w_\lambda(A) < 1$.*

Choosing $\frac{1}{d} < \lambda < 1$ would also work but the above choice simplifies the computations, so we fix $\lambda = \frac{1}{\sqrt{d}}$. This choice also asymptotically minimizes the bound we prove as $N \rightarrow \infty$. Before we prove this estimate, we show how it implies Theorem 1.1.

Proof of Theorem 1.1 assuming Lemma 2.1. To prove a.s. transience, it is enough to do so assuming $\eta(\rho)$ is also random and distributed as π_N like the other sites. Define $\mathcal{B}_0 = \{\rho\}$. Define $\mathcal{B}_{n+1}$ inductively as the set of sites outside $\mathcal{B}_0 \cup \dots \cup \mathcal{B}_n$ that are visited at any time by a particle originating at a site in $\mathcal{B}_n$. The set $\bigcup_{n=0}^\infty \mathcal{B}_n$ then consists of all

¹This artificial introduction of levels is the same as giving an arbitrary genealogy to the tree, so that the level decreases towards ancestors and increases towards descendants. In particular, the distance from ρ to v equals $\pm\ell(v)$ if and only if v is a direct descendant or ancestor of ρ , respectively.

sites ever visited by the process. We will show that this set does not encompass all of $\mathbf{T}_d$, from which we can conclude that the model is transient. Although the sets $\mathcal{B}_n$ will be infinite when $\eta(\rho) \geq 1$, their weight is finite, and using Lemma 2.1 recursively we show that they are in fact small.

Take N large enough that $\alpha := \mathbf{E}w_\lambda(A) < 1$ for the random set A defined in Lemma 2.1. Suppose $v \in \mathcal{B}_n$ for some n . The set of sites visited by the particles originating at v is distributed identically to A , except that it is shifted from starting at v instead of the root. Thus, the expected weight of the sites visited by the particles originating at v is $\lambda^{\ell(v)}\mathbf{E}w_\lambda(A)$. Hence,

$$\mathbf{E}[w_\lambda(\mathcal{B}_{n+1}) \mid \mathcal{B}_n] \leq \sum_{v \in \mathcal{B}_n} \lambda^{\ell(v)} \mathbf{E}w_\lambda(A) = \alpha w_\lambda(\mathcal{B}_n).$$

Since $w_\lambda(\mathcal{B}_0) = 1$, we get $\mathbf{E}[w_\lambda(\mathcal{B}_n)] \leq \alpha^n$ and finally

$$\mathbf{E}\left[w_\lambda\left(\bigcup_{n=0}^{\infty} \mathcal{B}_n\right)\right] = \mathbf{E}\left[\sum_{n=0}^{\infty} w_\lambda(\mathcal{B}_n)\right] \leq \sum_{n=0}^{\infty} \alpha^n < \infty.$$

Therefore $w_\lambda(\bigcup_{n=0}^{\infty} \mathcal{B}_n) < \infty$ a.s. Since $\mathbf{T}_d$ has infinite weight, we can conclude that with probability 1, not every site is visited. Moreover, only finitely many sites at each level are ever visited. We can also conclude that no site is visited infinitely often, since if one were, then a.s. all vertices would be visited. Hence level 0 is visited finitely many times, which by [6, Corollary 16] implies transience for the equivalent frog model on the d -ary tree. $\square$

In what follows, each appearance of C denotes a different positive finite constant which depends on d and whose actual value (sometimes easy to find) is irrelevant.

Proof of Lemma 2.1. Our strategy is to sum $\mathbf{P}[v \in A]\lambda^{\ell(v)}$ over all $v \in \mathbf{T}_d$. Let $\varphi(j, k)$ be the number of vertices in $\mathbf{T}_d$ at level j that have distance k from the root. For $j \geq 1$,

$$\begin{aligned} \varphi(j, j) &= d^j, \\ \varphi(j, j+2i) &= (d-1)d^{j+i-1} \leq d^{j+i} \text{ for } i \geq 1, \end{aligned}$$

and for $j \leq 0$,

$$\begin{aligned} \varphi(j, |j|) &= 1, \\ \varphi(j, |j|+2i) &= (d-1)d^{i-1} \leq d^i \text{ for } i \geq 1. \end{aligned}$$

If $k \neq |j| + 2k$ for some $i \geq 0$, then $\varphi(j, k) = 0$.

The above combinatorial terms will be controlled using hitting probabilities. For any $v \in \mathbf{T}_d$ at distance k from the root,

$$\mathbf{P}[\text{a random walk starting at } \rho \text{ ever visits } v] = d^{-k}. \quad (2.1)$$

Let $p_N(k)$ be the probability that at least one of N independent random walks starting at the root eventually hits a given vertex at distance k from the root in $\mathbf{T}_d$. By a union bound,

$$p_N(k) \leq d^{-k} N. \quad (2.2)$$

We have $v \in A$ if $\eta(\rho) = N$ and one of the N random walks hits v . Assuming $N = d^m$,

$$\mathbf{P}[v \in A] = \frac{\mu}{N} p_N(k) \leq \mu d^{-m} \min\{d^{-k} N, 1\} = \mu \min\{d^{-k}, d^{-m}\}.$$

Now let $\mathcal{S} = \mathbb{Z} \times \{0, 1, \dots\}$. We split the vertices according to $(j, i) \in \mathcal{S}$, since each $v \in \mathbf{T}_d$ is at some level j and some distance $|j| + 2i$ from the root. Our goal is to bound

$$\mathbf{E}w_\lambda(A) = \sum_{v \in \mathbf{T}_d} \mathbf{P}[v \in A] \lambda^{\ell(v)} \leq \sum_{(j,i) \in \mathcal{S}} \mu \lambda^j \varphi(j, |j| + 2i) \min\{d^{-|j|-2i}, d^{-m}\}. \quad (2.3)$$

We break the set $\mathcal{S}$ into six parts:

$$\begin{aligned} \mathcal{S}_1^+ &= \{(j, i) \in \mathcal{S}: 1 \leq j \leq m, 2i \leq m - j\}, \\ \mathcal{S}_2^+ &= \{(j, i) \in \mathcal{S}: 1 \leq j \leq m, 2i > m - j\}, \\ \mathcal{S}_3^+ &= \{(j, i) \in \mathcal{S}: j > m\}, \\ \mathcal{S}_1^- &= \{(-j, i) \in \mathcal{S}: 0 \leq j \leq m, 2i \leq m - j\}, \\ \mathcal{S}_2^- &= \{(-j, i) \in \mathcal{S}: 0 \leq j \leq m, 2i > m - j\}, \\ \mathcal{S}_3^- &= \{(-j, i) \in \mathcal{S}: j > m\}. \end{aligned}$$

The sets $\mathcal{S}_3^\pm$ are the easiest to estimate. Let $e_{j,i} = \lambda^j \varphi(j, |j| + 2i) \min\{d^{-|j|-2i}, d^{-m}\}$. Then

$$\sum_{(j,i) \in \mathcal{S}_3^+} e_{j,i} \leq \sum_{j > m} \lambda^j \sum_{i \geq 0} d^{j+i} d^{-j-2i} = C \sum_{j > m} \lambda^j$$

and

$$\sum_{(-j,i) \in \mathcal{S}_3^-} e_{-j,i} \leq \sum_{j > m} \lambda^{-j} \sum_{i \geq 0} d^i d^{-j-2i} = C \sum_{j > m} (\lambda d)^{-j}.$$

Note that both vanish as $m \rightarrow \infty$. For $\mathcal{S}_1^\pm$, we have

$$\begin{aligned} \sum_{(j,i) \in \mathcal{S}_1^+} e_{j,i} &\leq \sum_{1 \leq j \leq m} \lambda^j \sum_{0 \leq 2i \leq m-j} d^{j+i} d^{-m} = d^{-m} \sum_{1 \leq j \leq m} (\lambda d)^j \sum_{i=0}^{\lfloor \frac{m-j}{2} \rfloor} d^i \\ &\leq C d^{-m} \sum_{1 \leq j \leq m} (\lambda d)^j d^{\frac{m-j}{2}} = C d^{-\frac{m}{2}} \sum_{1 \leq j \leq m} (\lambda \sqrt{d})^j = C d^{-\frac{m}{2}} m, \end{aligned}$$

recalling that we set $\lambda = 1/\sqrt{d}$. Similarly,

$$\begin{aligned} \sum_{(-j,i) \in \mathcal{S}_1^-} e_{-j,i} &\leq \sum_{0 \leq j \leq m} \lambda^{-j} \sum_{0 \leq 2i \leq m-j} d^i d^{-m} = d^{-m} \sum_{0 \leq j \leq m} \lambda^{-j} \sum_{i=0}^{\lfloor \frac{m-j}{2} \rfloor} d^i \\ &\leq C d^{-m} \sum_{0 \leq j \leq m} \lambda^{-j} d^{\frac{m-j}{2}} = C d^{-\frac{m}{2}} \sum_{0 \leq j \leq m} (\lambda \sqrt{d})^{-j} = C d^{-\frac{m}{2}} (m+1). \end{aligned}$$

Both expressions again vanish as $m \rightarrow \infty$. Finally,

$$\begin{aligned} \sum_{(j,i) \in \mathcal{S}_2^+} e_{j,i} &\leq \sum_{1 \leq j \leq m} \lambda^j \sum_{2i > m-j} d^{j+i} d^{-j-2i} = \sum_{1 \leq j \leq m} \lambda^j \sum_{2i > m-j} d^{-i} \\ &\leq C \sum_{1 \leq j \leq m} \lambda^j d^{-\frac{m-j}{2}} = C d^{-\frac{m}{2}} \sum_{1 \leq j \leq m} (\lambda \sqrt{d})^j = C d^{-\frac{m}{2}} m \end{aligned}$$

and

$$\begin{aligned} \sum_{(-j,i) \in \mathcal{S}_2^-} e_{-j,i} &\leq \sum_{0 \leq j \leq m} \lambda^{-j} \sum_{2i > m-j} d^i d^{-j-2i} = \sum_{0 \leq j \leq m} \lambda^{-j} d^{-j} \sum_{2i > m-j} d^{-i} \\ &\leq C \sum_{0 \leq j \leq m} \lambda^{-j} d^{-j} d^{-\frac{m-j}{2}} = C d^{-\frac{m}{2}} \sum_{0 \leq j \leq m} (\lambda \sqrt{d})^{-j} = C d^{-\frac{m}{2}} (m+1), \end{aligned}$$

both vanishing like the other terms. Therefore,

$$\mathbf{E}w_\lambda(A) \leq \mu \sum_{(j,i) \in S} e_{j,i}$$

can be made less than 1 by choosing m large and $N = d^m$, which proves the lemma. $\square$

Remark 2.2. It is a classical result that the recurrence or transience of a branching random walk on the integers depends only on the expected offspring distribution. We explain now how this result extends to symmetric nearest-neighbor branching random walk on a regular tree. To precisely define the BRW on the integers, let Z be a point process on the integers. The initial generation of the BRW is a single particle at 0. To obtain generation $n+1$, each particle in generation n places new particles with positions given by an independent copy of Z translated by the particle's position. We assume that Z always contains at least one point so that the process survives a.s., and that it contains a point in the negative integers with positive probability.

To give Biggins's criterion for transience, let $m(\lambda) = \mathbf{E} \sum_{x \in Z} e^{-\lambda x}$, where the sum is over the atoms of Z . The BRW is transient in the positive direction if and only if there exists $\lambda > 0$ so that $m(\lambda) \leq 1$. See [4, Theorem 3] for a proof, though the result goes back to [3]. One can also of course test for transience in the negative direction by flipping Z across 0, and it turns out that transience in the negative and positive directions and recurrence are the only possibilities. Note that $m(\lambda)$ depends only on the expected number of atoms in Z at each integer.

Now, consider a BRW on a regular tree starting with a single particle at the root, where each particle reproduces independently by placing particles relative to itself sampled from some distribution. Assume that this distribution places particles only at neighbors of the root and is invariant under tree automorphisms fixing the root. Projecting each particle by ℓ yields a BRW on the integers whose transience is determined by the criterion above. Transience of the BRW on the integers implies transience of the BRW on the tree. Recurrence of the BRW on the integers implies that level 0 of the tree is visited infinitely often. By this, invariance, and independence of particles, each particle in generation 1 almost surely has a descendant that visits level 0. For a generation 1 particle at level 1, this implies that one of its descendant visits the root again. By invariance under tree automorphisms, this is also true of a particle at level -1 . Thus the root is visited infinitely often a.s.

3 Generalizations of the argument

We now sketch two extensions of the previous argument. First and very simple is to show transience on $\mathbf{T}_d$ when the particles per site has infinite mean. Second and more delicate is to show transience on $\mathbf{T}_d$ with $d \geq 14$ when the number of particles not only has infinite mean but is a.s. at least 1.

Fix some $\mu > 0$. In the proof of Lemma 2.1 we have in fact shown that $\mathbf{E}_{\pi_N} w_\lambda(A)$ can be made arbitrarily small as $N \rightarrow \infty$, where we use $\mathbf{E}_{\pi_N}$ to denote the expectation assuming the distribution of particles is π_N . Choose N_n so that $\mathbf{E}_{\pi_{N_n}} w_\lambda(A) < 2^{-n}$. Let X_n be distributed as π_{N_n} , let $X = \sum_{n=1}^{\infty} X_n$, and let π be the distribution of X . Then

$$\mathbf{E}_\pi w_\lambda(A) \leq \sum_n \mathbf{E}_{\pi_{N_n}} w_\lambda(A) < 1, \quad (3.1)$$

and the proof of Theorem 1.1 carries on without any change to show transience with particle counts given by π , even though π has expectation $\sum_n \mu = \infty$.

Now suppose $d \geq 14$. We start by showing that the model is transient when the initial distribution is given by π_N plus an extra particle at each site. We combine the approach

introduced in §2 with the standard one based on branching random walks. First take $\zeta(v)$ independent with distribution π_N , and consider the configuration η given by $\eta(\rho) = \zeta(\rho)$ and $\eta(v) = \zeta(v) + 1$ for $v \neq \rho$ (as before, altering the distribution at the root does not affect the property of a.s. transience). We assign the guaranteed particle at each site *type 1* and the $\zeta(v)$ particles *type 2*. Start with $\mathcal{B}_0 = \{\rho\}$. Define $\mathcal{B}_{n+1}$ inductively as follows. Launch the type 2 particles from each site $v \in \mathcal{B}_n$, and allow only the type 1 particles to wake. Take $\mathcal{B}_{n+1}$ as the set of sites outside of $\mathcal{B}_0 \cup \dots \cup \mathcal{B}_n$ that are visited. As before, the set $\cup_{n=0}^\infty \mathcal{B}_n$ consists of all sites ever visited in the frog model. Moreover, $\mathbf{E}w_\lambda(\mathcal{B}_n) \leq \alpha^n$ where $\alpha := \mathbf{E}w_\lambda(\mathcal{B}_1)$.

So again it suffices to show that $\alpha < 1$. Instead of the hitting probability (2.1), we use the following estimate based on BRWs:

Lemma 3.1. *Let $d \geq 6$, and let $v \in \mathbf{T}_d$ be an arbitrary vertex at distance k from the root. Run the frog model with one sleeping frog per site starting with the particle at the root active. Then the probability of ever hitting v is at most*

$$\frac{d+1}{d+1-\sqrt{8d}} \left(\frac{4}{d+1} \right)^k.$$

For $d \geq 14$, we have $4/(d+1) < d^{-1/2}$. By the above lemma, if a single type 2 particle begins at the root and the frog model runs with only type 1 particles allowed to wake, then v at level k is hit with probability less than $Cd^{-\beta k}$ for some $\beta > 1/2$ (constants C now depend on d and β). Let $\varepsilon = 1 - \beta < 1/2$, and let $\beta' = 1 - 2\varepsilon > 0$. Instead of (2.2), we have $p_N(k) \leq Cd^{-\beta k} N$, and we consider the sum over $(j, i) \in \mathcal{S}$ of

$$e_{j,i} = C\lambda^j \varphi(j, |j| + 2i) \min\{d^{-\beta(|j|+2i)}, d^{-m}\}$$

which is the same as before with an extra β . The estimates in $\mathcal{S}_1^\pm$ of course do not change. The estimate in $\mathcal{S}_2^+$ and $\mathcal{S}_3^+$ becomes $e_{j,i} \leq C(\lambda d^\varepsilon)^j d^{-\beta' i}$. The estimate in $\mathcal{S}_2^-$ and $\mathcal{S}_3^-$ becomes $e_{-j,i} \leq C(\lambda d^\beta)^{-j} d^{-\beta' i}$. Since these are still summable over $\mathcal{S}$, the sum over $\mathcal{S}_3^\pm$ vanishes for large m as before. Carrying the same computations as in the proof of Lemma 2.1 we get as upper bounds $Cd^{-\beta' \frac{m}{2}}$ for the sum over $\mathcal{S}_2^+$ and $Cd^{-\beta' \frac{m}{2}}(m+1)$ for the sum over $\mathcal{S}_2^-$. Both vanish for large m , and therefore α can be made arbitrarily small. This proves that the model is transient with distribution π_N plus an extra particle for large enough N . Using the same idea as in (3.1), we can maintain $\alpha < 1$ while replacing π_N with a distribution with infinite expectation, and the proof goes on as before.

Proof of Lemma 3.1. Consider the following BRW. One particle starts at the root $\rho \in \mathbf{T}_d$. At every step, each particle at u chooses a neighbor u' of u uniformly at random. If u' lies on the path from ρ to u (including ρ itself), the particle produces one offspring at u' ; if not, it produces two offspring at u' . Since a particle moving towards the root in the frog model never wakes any particles, this BRW dominates the frog model with one sleeping particle per site, in the sense that it can be coupled with it so that every active particle in the frog model at time n also exists in the BRW. Thus it suffices to bound the probability that this BRW hits v .

Assign a particle at distance k from the root to have weight $e^{-\theta k}$, for θ to be chosen later. Let W_n be the total weight of all particles in the BRW at time n . Define

$$m = \frac{1}{d+1} e^\theta + \frac{2d}{d+1} e^{-\theta}.$$

We claim that

$$\mathbf{E}[W_{n+1} \mid W_n] \leq mW_n. \quad (3.2)$$

Indeed, consider a particle at distance k from the root at time n , which has weight $e^{-\theta k}$. If $k \geq 1$, then at its next step it produces a single particle with weight $e^{-\theta(k-1)}$ with probability $1/(d+1)$, or it produces two particles each with weight $e^{-\theta(k+1)}$ with probability $d/(d+1)$. Thus the expected weight from the offspring is exactly m times the weight of the vertex. If $k = 0$, then the particle at its next step deterministically produces two offspring with weight $e^{-\theta}$, which collectively have weight $2e^{-\theta} < m$. Thus, the expected weight of the offspring of any particle is bounded by m times the weight of the particle, which proves (3.2). Since $W_0 = 1$, this shows that $\mathbf{E}W_n \leq m^n$. Set $\theta = \log(2d)/2$ to optimize m . This gives $m = \sqrt{8d}/(d+1)$. Under our assumption $d \geq 6$, we have $m < 1$.

Now, let X_k denote the number of vertices at distance k from the root that are ever visited by the BRW. A particle at distance k from the root has weight $e^{-\theta k}$ at that time, and the first time that there can be a particle at distance k from the root is time k . Thus, $e^{-\theta k} X_k \leq \sum_{n=k}^{\infty} W_n$. Taking expectations,

$$\mathbf{E}X_k \leq e^{\theta k} \frac{m^k}{1-m} = \frac{1}{1-m} \left(\frac{4d}{d+1} \right)^k.$$

Since $\mathbf{E}X_k$ is the sum of the probabilities of each the $(d+1)d^{k-1}$ vertices at distance k from the root being visited,

$$\mathbf{P}[v \text{ is visited}] = \frac{1}{(d+1)d^{k-1}} \mathbf{E}X_k \leq \frac{1}{1-m} \left(\frac{4}{d+1} \right)^k. \quad \square$$

References

- [1] O. S. M. Alves, F. P. Machado, S. Yu. Popov, and K. Ravishankar, *The shape theorem for the frog model with random initial configuration*, Markov Process. Related Fields **7** (2001), no. 4, 525–539. MR-1893139
- [2] Riddhipratim Basu, Shirshendu Ganguly, and Christopher Hoffman, *Non-fixation for conservative stochastic dynamics on the line*, Comm. Math. Phys. **358** (2018), no. 3, 1151–1185. MR-3778354
- [3] J. D. Biggins, *Martingale convergence in the branching random walk*, J. Appl. Probability **14** (1977), no. 1, 25–37. MR-0433619
- [4] J. D. Biggins, *Lindley-type equations in the branching random walk*, Stochastic Process. Appl. **75** (1998), no. 1, 105–133. MR-1629030
- [5] Christopher Hoffman, Tobias Johnson, and Matthew Junge, *From transience to recurrence with Poisson tree frogs*, Ann. Appl. Probab. **26** (2016), no. 3, 1620–1635. MR-3513600
- [6] Christopher Hoffman, Tobias Johnson, and Matthew Junge, *Recurrence and transience for the frog model on trees*, Ann. Probab. **45** (2017), no. 5, 2826–2854. MR-3706732
- [7] Tobias Johnson and Matthew Junge, *The critical density for the frog model is the degree of the tree*, Electron. Commun. Probab. **21** (2016), Paper No. 82, 12. MR-3580451
- [8] Tobias Johnson and Matthew Junge, *Stochastic orders and the frog model*, Ann. Inst. Henri Poincaré Probab. Stat. **54** (2018), no. 2, 1013–1030. MR-3795075
- [9] Debabrata Panja, *Effects of fluctuations on propagating fronts*, Physics Reports **393** (2004), no. 2, 87–174.
- [10] Leonardo T. Rolla and Vladas Sidoravicius, *Absorbing-state phase transition for driven-dissipative stochastic dynamics on $\mathbb{Z}$* , Invent. Math. **188** (2012), no. 1, 127–150. MR-2897694
- [11] Leonardo T. Rolla, Vladas Sidoravicius, and Olivier Zindy, *Universality and sharpness in absorbing-state phase transitions*, Ann. Henri Poincaré **20** (2019), no. 6, 1823–1835.
- [12] Vladas Sidoravicius and Augusto Teixeira, *Absorbing-state transition for stochastic sandpiles and activated random walks*, Electron. J. Probab. **22** (2017), Paper No. 33, 35. MR-3646059
- [13] Alexandre Stauffer and Lorenzo Taggi, *Critical density of activated random walks on transitive graphs*, Ann. Probab. **46** (2018), no. 4, 2190–2220. MR-3813989

Acknowledgments. T.J. thanks the support and hospitality of the NYU-ECNU Institute of Mathematical Sciences at NYU Shanghai. L.R. received support from grants UBACYT-2017 Mod-I 20020160100147BA and PICT 2015-3154. T.J. received support from NSF grants DMS-1401479 and DMS-1811952 and PSC-CUNY award #61540-00 49.

Electronic Journal of Probability

Electronic Communications in Probability

Advantages of publishing in EJP-ECP

- Very high standards
- Free for authors, free for readers
- Quick publication (no backlog)
- Secure publication (LOCKSS¹)
- Easy interface (EJMS²)

Economical model of EJP-ECP

- Non profit, sponsored by IMS³, BS⁴, ProjectEuclid⁵
- Purely electronic

Help keep the journal free and vigorous

- Donate to the IMS open access fund⁶ (click here to donate!)
- Submit your best articles to EJP-ECP
- Choose EJP-ECP over for-profit journals

¹LOCKSS: Lots of Copies Keep Stuff Safe <http://www.lockss.org/>

²EJMS: Electronic Journal Management System <http://www.vtex.lt/en/ejms.html>

³IMS: Institute of Mathematical Statistics <http://www.imstat.org/>

⁴BS: Bernoulli Society <http://www.bernoulli-society.org/>

⁵Project Euclid: <https://projecteuclid.org/>

⁶IMS Open Access Fund: <http://www.imstat.org/publications/open.htm>