Random recursion

Matthew Junge

A dissertation submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

University of Washington

2016

Reading Committee:

Christopher Hoffman, Chair

Krzysztof Burdzy

Zhen-Qing Chen

Program Authorized to Offer Degree:
Mathematics

University of Washington

Abstract

Random recursion

Matthew Junge

Chair of the Supervisory Committee:
Professor Christopher Hoffman
Mathematics

We study the limiting behavior of three stochastic processes. Two are interacting particle systems, the frog model and coalescing random walk. We work out transience and recurrence properties on various graphs. The last is an interval splitting algorithm, which is shown to be equidistributed in the limit. Many of the proofs hinge on recursive equations of random variables.

ACKNOWLEDGMENTS

It is humbling to be here after nearly failing calculus, and actually failing many prelims. This document would not exist without the support of my family and friends. Both Chris Hoffman and Itai Benjamini are inspirations for my approach to mathematics. Many conversations with Shirshendu Ganguly have been invaluable. Collaborating with Eric Foxall has been fortuitous and gratifying. Toby Johnson helped me believe I could be excellent, and continues to be an immeasurably positive influence.

Table of Contents

Introduction 1

Transience and recurrence for the frog model on trees 10

The frog model is a growing system of random walks where a particle is added whenever a new site is visited. A longstanding open question is how often the root is visited on the infinite d-ary tree. We prove the model undergoes a phase transition, finding it recurrent for $d = 2$ and transient for $d \geq 5$. Simulations suggest strong recurrence for $d = 2$, weak recurrence for $d = 3$, and transience for $d \geq 4$. Additionally, we prove a 0-1 law for all d-ary trees, and we exhibit a graph on which a 0-1 law does not hold. Joint with Christopher Hoffman and Tobias Johnson.

From transience to recurrence with Poisson tree frogs 33

On a d-ary tree one particle is awake at the root and i.i.d. Poisson many particles are sleeping at every other vertex. Particles that are awake perform simple random walks, awakening any sleeping particles they encounter. We prove that there is a phase transition between transience and recurrence as the initial density of particles increases, and we give the order of the transition up to a logarithmic factor. Joint with Christopher Hoffman and Tobias Johnson.

Stochastic orders and the frog model 46

The frog model *starts with one active particle at the root of a graph and $\eta(v)$ dormant particles at all nonroot vertices v. Active particles follow independent random paths, waking all inactive particles they encounter. We prove that certain frog model statistics are monotone in the initial configuration with respect to stochastic dominance in the increasing concave order. As a consequence, a frog model recurrent with i.i.d.-η particles per site is also recurrent with deterministically $\lceil \mathbf{E}\eta \rceil$ particles per site. We deduce recurrence for the infinite d-ary tree with simple random walk paths and $k > \mu_c(d)$ frogs at each site, with $\mu_c(d)$ the threshold for recurrence of a frog model with i.i.d. Poisson frogs per site. We end with a proof that $\mu_c(d)$ is of order d, removing a logarithmic term from the previous upper bound. Joint with Tobias Johnson.*

Site recurrence for coalescing random walk 67

Begin continuous time random walks from every vertex of a graph and have particles coalesce when they collide. We use a duality relation with the voter model to prove the process is site recurrent on bounded degree graphs, and for Galton-Watson trees whose offspring distribution has exponential tail. We prove bounds on the occupation probability of a site, as well as a general 0-1 law. Similar conclusions hold for a coalescing process on trees where particles do not backtrack. Joint with Itai Benjamini, Eric Foxall, Ori Gurel-Gurevich, and Harry Kesten.

Choices, intervals and equidistribution 78

We give a sufficient condition for a random sequence in [0,1] generated by a Ψ-process to be equidistributed. The condition is met by the canonical example – the max-*2 process – where the nth term is whichever of two uniformly placed points falls in the larger gap formed by the previous $n-1$ points. This solves an open problem from Itai Benjamini, Pascal Maillard and Elliot Paquette. We also deduce equidistribution for more general Ψ-processes. This includes an interpolation of the* min-*2 and* max-*2 processes that is biased towards* min-*2.*

Introduction

Random walk—the process that moves randomly along adjacent edges of a graph—has applications to virtually all of the sciences. Indeed, many physical, biological, chemical and financial systems are driven by small, random fluctuations. Of particular interest to mathematicians is the long term behavior. For example, the classical Pólya's recurrence theorem says that a random walk on $\mathbb{Z}^d$ visits the origin infinitely often when $d \leq 2$, but only finitely many times for larger d. One part of my research looks for similar phase transitions for multiple interacting random walks. Another part of my work studies how random decimals spread throughout $[0, 1]$, and asks questions inspired by Weyl's equidistribution theorem.

My work is contained in: [HJJ15b, HJJ15a, JJ16, Jun15, BFG$^+$15]. This includes solutions to two widely disseminated open problems, and also generalizes theorems for well-known models. These results have garnered some attention and are cited in [KZ15, HW15, GNR15, MP14]. [HJJ15a] is to appear in the Annals of Applied Probability and [Jun15] is published in the Electronic Journal of Probability. The papers [HJJ15b, JJ16, BFG$^+$15] are pending review at Annals of Probability, Probability Theory and Related Fields, and Electronic Communications in Probability, respectively.

The frog model. The *frog model* is an interacting particle system that can model the spread of a virus. Initially there is one awake particle at the root of a graph, and one dormant particle at every other vertex. Awake particles perform random walk, waking any dormant particles they visit. Rick Durrett is credited with the name frog model, which zoomorphisizes the chaotic way the set of walkers grows.

On $\mathbb{Z}^d$ the growth of the frog model is well understood. The articles [AMP02], [AMPR01] and [RS04] successfully apply the sub-additive ergodic theorem, proving that the activated sites converge to a limiting shape. The primary obstacle to applying Kingman's ergodic theory is estimating the probability the root is occupied at time t. A related measurement is the limiting number of visits to the root. Call the model *recurrent* if the root is visited by infinitely many frogs and *transient* if the root is visited finitely many times. The transience/recurrence behavior of the model relates to the occupation probability. For example, these probabilities are summable in the transient case and divergent in the recurrent case. Visits to the root was first studied on $\mathbb{Z}^d$ where the model is recurrent for all $d \geq 1$ ([TW99]). Compare to the change as d increases in Pólya's recurrence theorem.

In my project ([HJJ15b, HJJ15a]) with Christopher Hoffman and Tobias Johnson we study the frog model on infinite d-ary trees. How often the root is visited on trees was one of longest standing and most basic questions surrounding the model (asked in [GS09, Pop03, AMP02]). We prove a phase transition occurs as d increases; recall, this does not happen on $\mathbb{Z}^d$.

Theorem 1 (Hoffman, Johnson, J.)**.** *The frog model on the d-ary tree is recurrent for $d = 2$ and transient for $d \geq 5$.*

Random recursion is vital to the proof of recurrence for $d = 2$. We are able to simplify the model, and then express the visits to the root, V, as an equation $V \sim f(V_1, V_2)$ for two idendtically distributed copies of V. We then prove that $V \equiv \infty$ is the only random variable that can satisfy the equation.

The degree, d, of the tree is a coarse parameter. Working with some continuous parameter would give us more control. In [HJJ15a] we study the frog model on the d-ary tree with a random number of frogs at each site. Something similar occurs in [Pop01]. By increasing a parameter, α, the frog model on $\mathbb{Z}^d$ with $\text{Ber}(\alpha \|x\|^{-2})$ frogs at each vertex x switches from transient to recurrent. In our paper [HJJ15a] we vary the mean density of frogs on a fixed d-ary tree and observe a phase transition.

Theorem 2 (Hoffman, Johnson, J.)**.** *There exists $\mu_d > 0$ such that the frog model with one awake frog at the root and* Poisson(μ) *frogs independently placed at each site is transient for $\mu < \mu_d$ and recurrent for $\mu > \mu_d$. Moreover, $\mu_d = \Theta(d)$.*

One shortcoming of this theorem was that we could not prove recurrence on, say, a 5-ary tree with a large deterministic number of frogs at each site. A recent project with Toby Johnson ([JJ16]) resolves this. We look at a more general frog model. It starts with one active particle at the root of a graph and $\eta(v)$ dormant particles at all nonroot vertices v. Active particles follow independent random paths (not necessarily simple random walk paths), waking all inactive particles they encounter. We prove that certain frog model statistics—like visits to the root—are monotone in the initial configuration with respect to stochastic dominance in the increasing concave order; $X \preceq_{\text{icv}} Y$ iff $\varphi(X) \leq \varphi(Y)$ for all increasing concave functions φ.

It is a consequence of Jenson's inequality that the constant functions are maximal in this stochastic order. Thus, a frog model recurrent with i.i.d.-η particles per site is also recurrent with deterministically $\lceil \mathbf{E}\eta \rceil$ particles per site. We deduce recurrence for the infinite d-ary tree with simple random walk paths and $k > \mu_c(d)$ frogs at each site, with $\mu_c(d)$ the threshold for recurrence of a frog model with i.i.d. Poisson frogs per site. Our main theorem is as follows.

Theorem 3. *Assume that the frog paths S and counts $\eta(v)$ and $\eta'(v)$ are mutually independent for all v and i, and that the paths at a particular vertex v are identically distributed for all i. Let f be an icv statistic of the frog model. If $\eta(v) \preceq_{icv} \eta'(v)$ for all v, then $f(\eta, S) \preceq_{icv} f(\eta', S)$.*

Though the definition is a bit technical, the statistic to keep in mind is f being the number of visits to the root. This allows us to deduce the following desired corollary for d-ary trees with deterministically many frogs per site.

Corollary 4. *The frog model on $\mathbb{T}_d$ with $k > \mu_c(d)$ frogs per site is almost surely recurrent.*

.

Coalescing random walk. Start with one particle at each vertex of a graph and have particles perform random walk. Upon colliding, they bind together and proceed as one particle. This is called *coalescing random walk*. The study of such systems began in the 1970's with the paper of Erdős and Ney [EN74], and to this day, variations continue to find new applications. For example, random coalescence involving multiple types of particles is used to model chemical reactions ([Hol83, BL88, vdBK00]). It's connection to the voter model first was noted in [BG80]. The recent paper [BL15] also exploits this relationship. Computer scientists are interested in the process on finite graphs ([Cox89, CEOR12]).

Early studies of coalescing random walk were on $\mathbb{Z}^d$. [Gri78] shows that coalescing random walk is *site recurrent* for all d, meaning that each site is In [BFG$^+$15] we revisit this problem and prove a general criterion for recurrence. This criterion is met by a large family of graphs:

Theorem 5 (Benjamini, Foxall, Gurel-Gurevich, J., Kesten)**.** *Coalescing random walk is site recurrent on any bounded degree graph.*

Our proof uses duality with the voter model and martingale techniques. The random recursion in this setting is describing the survival time of cluster whose expansion and regression rates depend explicitly on its current size. We are ably to analyze this with a random walk. Theorem 5 resolves the question of recurrence for coalescing random walk in many settings. We are also interested in understanding similar, but less random processes. Consider coalescing random walk where particles follow non-backtracking paths. A lack of symmetry makes it difficult to apply known techniques. In [BFG$^+$15] we use duality and martingales to prove:

Theorem 6 (Benjamini, Foxall, Gurel-Gorevich, J., Kesten)**.** *Non-backtracking coalescing random walk on bounded degree trees spends infinite time at the root.*

Interval splitting. A sequence in $[0,1]$ is *equidistributed* if the limiting proportion of points in each subinterval is equal to the subinterval's length. A century ago Weyl proved that $\{\beta n \mod 1\}_{n\geq 1}$ is equidistributed for any irrational number β (see [Wey10]). Since then connections have been found in ergodic theory, number theory, complex analysis and computer science ([BM72], [Vau77], [FSZ09], [CKK+07]).

[MP14] introduces a family of interval splitting procedures called Ψ-processes. The canonical example is the *max-2 process*. We obtain the nth term of the random sequence by placing two candidate points uniformly in $[0,1]$ and keeping whichever lies in the larger subinterval formed by the previous $n-1$ points.

A discrete analogue of the max-2 process appears in [ABKU99] where n balls are randomly placed into n bins. For each ball two bins are selected uniformly and the ball is placed in the bin with fewer balls. They find that the most-filled bin has exponentially fewer balls then if they were instead placed uniformly. This phenomenon has been dubbed "the power of two choices," and is studied in more detail in [MRS00] and [LM05]. Analogously, the max-2 process should spread points more evenly than uniformly adding points. Despite our intuition this is difficult to formalize, and equidistribution was the primary open problem from [MP14]. In [Jun15] I resolve this.

Theorem 7 (J.)**.** *The* max-2 *process is equidistributed almost surely.*

The random recursion in this solution involves conditioning on the past of the size-biased empirical distribution of interval lengths, and using this to describe the future evolution of the process.

References

[ABKU99] Yossi Azar, Andrei Z. Broder, Anna R. Karlin, and Eli Upfal. Balanced allocations. *SIAM J. Comput.*, 29(1):180–200, September 1999.

[AMP02] O. S. M. Alves, F. P. Machado, and S. Yu. Popov. The shape theorem for the frog model. *Ann. Appl. Probab.*, 12(2):533–546, 2002.

[AMPR01] 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(4):525–539, 2001.

[BFG+15] Itai Benjamini, Eric Foxall, Ori Gurel-Gurevich, Matthew Junge, and Harry Kesten. Site recurrence for coalescing random walk. *ArXiv e-prints:* 1510.04721, October 2015.

[BG80] Maury Bramson and David Griffeath. Asymptotics for interacting particle systems on $\mathbb{Z}^d$. *Zeitschrift fr Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 53(2):183–196, 1980.

[BL88] Maury Bramson and Joel L. Lebowitz. Asymptotic behavior of densities in diffusion-dominated annihilation reactions. *Phys. Rev. Lett.*, 61:2397–2400, Nov 1988.

[BL15] M. Balázs and A. László Nagy. Dependent augmented Branching Annihilating Random Walk. *ArXiv e-prints*, January 2015.

[BM72] J. R. Blum and V. J. Mizel. A generalized Weyl equidistribution theorem for operators, with applications. *Transactions of the American Mathematical Society*, 165(2):291–307, 06 1972.

[CEOR12] Colin Cooper, Robert Elsässer, Hirotaka Ono, and Tomasz Radzik. Coalescing random walks and voting on graphs. *CoRR*, abs/1204.4106, 2012.

[CKK+07] Jacek Cichoń, Marek Klonowski, Lukasz Krzywiecki, Bartlomiej Różański, and Pawel Zieliński. Random subsets of the interval and p2p protocols. In Moses Charikar, Klaus Jansen, Omer Reingold, and José D.P. Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, volume 4627 of *Lecture Notes in Computer Science*, pages 409–421. Springer Berlin Heidelberg, 2007.

[Cox89] J. T. Cox. Coalescing random walks and voter model consensus times on the torus in $\mathbb{Z}^d$. *Ann. Probab.*, 17(4):1333–1366, 10 1989.

[EN74] P. Erdos and P. Ney. Some problems on random intervals and annihilating particles. *Ann. Probab.*, 2(5):828–839, 10 1974.

[FSZ09] Kevin Ford, K. Soundararajan, and Alexandru Zaharescu. On the distribution of imaginary parts of zeros of the Riemann zeta function, ii. *Mathematische Annalen*, 343(3):487–505, 2009.

[GNR15] Arka P. Ghosh, Steven Noren, and Alexander Roitershtein. On the range of the transient frog model on $\mathbb{Z}$. available at arXiv:1502.02738, 2015.

[Gri78] David Griffeath. Annihilating and coalescing random walks on z^d. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 46(1):55–65, 1978.

[GS09] Nina Gantert and Philipp Schmidt. Recurrence for the frog model with drift on $\mathbb{Z}$. *Markov Process. Related Fields*, 15(1):51–58, 2009.

[HJJ15a] Christopher Hoffman, Tobias Johnson, and Matthew Junge. From transience to recurrence with Poisson tree frogs. *ArXiv e-prints:* 1501.05874, January 2015. *To appear in Annals of Applied Probability.*

[HJJ15b] Christopher Hoffman, Tobias Johnson, and Matthew Junge. Recurrence and transience for the frog model on trees. *ArXiv e-prints:* 1404.6238, July 2015.

[Hol83] Richard Holley. Two types of mutually annihilating particles. *Advances in Applied Probability*, 15(1):133–148, 1983.

[HW15] T. Höfelsauer and F. Weidner. The speed of frogs with drift on $\mathbb{Z}$. *ArXiv e-prints*, May 2015.

[JJ16] T. Johnson and M. Junge. Stochastic orders and the frog model. *ArXiv e-prints*, February 2016.

[Jun15] Matthew Junge. Choices, intervals and equidistribution. *Electronic Journal of Probability.*, 20:no. 97, 1–18, 2015.

[KZ15] E. Kosygina and M. P. W. Zerner. A zero-one law for recurrence and transience of frog processes. *ArXiv e-prints*, August 2015.

[LM05] Malwina J. Luczak and Colin McDiarmid. On the power of two choices: Balls and bins in continuous time. *The Annals of Applied Probability*, 15(3):1733–1764, 08 2005.

[MP14] P. Maillard and E. Paquette. Choices and intervals. *ArXiv e-prints*, 2014. To appear in Israel Journal of Mathematics.

[MRS00] Michael Mitzenmacher, Andra W. Richa, and Ramesh Sitaraman. The power of two random choices: A survey of techniques and results. In *in Handbook of Randomized Computing*, pages 255–312. Kluwer, 2000.

[Pop01] S.Yu. Popov. Frogs in random environment. *Journal of Statistical Physics*, 102(1-2):191–201, 2001.

[Pop03] Serguei Yu. Popov. Frogs and some other interacting random walks models. In *Discrete random walks (Paris, 2003)*, Discrete Math. Theor. Comput. Sci. Proc., AC, pages 277–288 (electronic). Assoc. Discrete Math. Theor. Comput. Sci., Nancy, 2003.

[RS04] Alejandro F. Ramírez and Vladas Sidoravicius. Asymptotic behavior of a stochastic combustion growth process. *J. Eur. Math. Soc. (JEMS)*, 6(3):293–334, 2004.

[TW99] András Telcs and Nicholas C. Wormald. Branching and tree indexed random walks on fractals. *J. Appl. Probab.*, 36(4):999–1011, 1999.

[Vau77] R. C. Vaughan. On the distribution of p modulo 1. *Mathematika*, 24:135–141, 12 1977.

[vdBK00] J. van den Berg and Harry Kesten. Asymptotic density in a coalescing random walk model. *Ann. Probab.*, 28(1):303–352, 01 2000.

[Wey10] Hermann Weyl. Über die gibbs'sche erscheinung und verwandte konvergenzphänomene. *Rendiconti del Circolo Matematico di Palermo*, 30(1):377–407, 1910.

RECURRENCE AND TRANSIENCE FOR THE FROG MODEL ON TREES

CHRISTOPHER HOFFMAN, TOBIAS JOHNSON, AND MATTHEW JUNGE

Abstract. The frog model is a growing system of random walks where a particle is added whenever a new site is visited. A longstanding open question is how often the root is visited on the infinite d-ary tree. We prove the model undergoes a phase transition, finding it recurrent for $d = 2$ and transient for $d \geq 5$. Simulations suggest strong recurrence for $d = 2$, weak recurrence for $d = 3$, and transience for $d \geq 4$. Additionally, we prove a 0-1 law for all d-ary trees, and we exhibit a graph on which a 0-1 law does not hold.

1. Introduction

The frog model is a system of interacting random walks on a given rooted graph. Initially, the graph contains one particle at the root and some configuration of sleeping particles on its vertices; unless otherwise stated, we will assume an initial condition of one sleeping particle per vertex. The particle at the root starts out awake and performs a simple nearest-neighbor random walk in discrete time. Whenever a vertex with sleeping particles is first visited, all the particles at the site wake up and begin their own independent random walks, waking particles as they visit them. A formal definition of the frog model is in [AMP02a], and a nice survey of variations is in [Pop03]. Traditionally, particles are referred to as frogs, a practice we will continue here.

One of the most basic questions about the frog model on an infinite graph is whether it is recurrent or transient. Telcs and Wormald determined that the frog model was recurrent on $\mathbb{Z}^d$ for any d, the first published result on the frog model [TW99]. On an infinite d-ary tree, this question is more difficult. It was first posed in [AMP02b]. It was asked again in [Pop03] and in [GS09], which pointed out that the answer was unknown even on a binary tree.

Our main result in this paper is that the frog model is recurrent on the binary tree but transient on the d-ary tree for $d \geq 5$, demonstrating a phase transition not found on $\mathbb{Z}^d$. We use martingale techniques for the transience results, though pushing this result down to $d = 5$ is more complicated and requires computer assistance. Our proof of recurrence on the binary tree uses a bootstrapping argument in which we iteratively assume that the number of visits to the root is stochastically large and then prove it even larger; the argument seems novel to us.

Background on the frog model. It is common to use the frog model as a model for the spread of rumors or infections, thinking of awakened frogs as informed or infected agents. See [DG99] for an overview and [KPV04] for more tailored discussion. Another perspective on the frog model is as a conservative lattice gas model with the reaction $A + B \to 2A$. Here A represents an an active particle and B an inert particle. Active particles disperse throughout the graph, igniting any inert particles they contact. Several papers taking this perspective study a process identical to the frog model except that particles move in continuous rather than discrete time [RS04, CQR09, BR10]. This process and its variants have also seen much study by physicists; see the references in [CQR09] and [BR10]. Our results in this paper depend only on the paths of the frogs and not on the timing of their jumps, and so they apply equally well to this continuous-time process.

In the larger mathematical context, the frog model is part of a family of self-interacting random walks which have proven quite difficult to analyze. ([Pem07] provides a nice survey of this family.) In recent years progress on a few self-interacting random walks has generated considerable interest. One of these close relatives is activated random walk, which is touched on in [KS06] and studied in depth in [DRS10, RS12, ST14]. Activated random walk can be described as a frog model where frogs fall back asleep at some given rate. Another related process is excited random walk [BW03]. This walk has a bias the first time it visits a site but is unbiased each subsequent time that it returns. The frog model can be thought of as an "excited" branching process, which branches at a site v only the first time the process visits v.

Initial interest in the frog model was on the graph $\mathbb{Z}^d$. For any d, it was shown that the process is recurrent [TW99] and that the set of visited vertices grows linearly and when rescaled converges to a limiting shape [AMP02a]. A similar shape theorem was proven independently in [RS04] for the process with continuous-time particles. Both shape theorems rely on the subadditive ergodic theorem. A technical difficulty that arises is proving the expected time to wake a given frog is finite. Thus, measuring recurrence on a given graph is the first step in understanding the long time behavior of the model. Other results for the frog model on $\mathbb{Z}^d$ include determining the decay of initially sleeping frogs necessary to make the model transient [Pop01]. Later this question was studied on the integers with drift [GS09]. A finer analysis in this setting recently appeared in [DP14] and [GNR15].

Our main interest in this paper is in recurrence and transience on $\mathbb{T}_d$, the infinite rooted d-ary tree. We denote the root by $\varnothing$. We also study aspects of the process on $\mathbb{T}_d^{\mathrm{hom}}$, the infinite homogeneous degree $(d+1)$-tree.

Some attention has already been given to a relative of our model on $\mathbb{T}_d^{\mathrm{hom}}$ in which awake frogs die after independently taking a geometrically distributed number of jumps. In [AMP02b] and [LMP05], the authors prove a phase transition for survival. Depending on the parameter, there will either be frogs alive at all times with positive probability or the process will die out almost surely. We study the model in which frogs jump perpetually—a fundamentally different problem, since it switches the emphasis from the local to the global behavior of the model.

Statement and discussion of results. For a given rooted graph we call a realization of the model recurrent if the root is visited infinitely many times and transient if it is visited finitely often. Our main theorem covers the d-ary tree for all but two degrees:

Theorem 1.

(i) *The frog model on $\mathbb{T}_2$ is almost surely recurrent.*
(ii) *The frog model on $\mathbb{T}_d$ for $d \geq 5$ is almost surely transient.*

We also make a conjecture on these two unsolved degrees based on fairly convincing evidence from simulations, presented in Section 5.

Conjecture 2. *The frog model on $\mathbb{T}_d$ is recurrent a.s. for $d = 3$ and transient a.s. for $d = 4$.*

The simulations suggest the possibility of a three-phase transition as d increases. We call the model strongly recurrent if the probability the root is occupied is bounded away from zero for all time. We call it weakly recurrent if it is recurrent with positive probability, but the probability that the root is occupied decays to zero.

Open Question 3. *Is the frog model strongly recurrent on $\mathbb{T}_2$ but weakly recurrent on $\mathbb{T}_3$?*

Such a transition would give information about the time to wake all children of the root. For instance, strong recurrence on $\mathbb{T}_2$ would imply that this time has finite expectation and an exponential tail. Should $\mathbb{T}_3$ exhibit a weak recurrence phase, then a tantalizing problem would be to estimate the decay of the occupation time at the root.

The recurrence of the frog model on the binary tree established in Theorem 1 (i) is the flagship result of this article. The proof goes by coupling the frog model with a process in which the root is visited less often. Let V be the number of visits to the root in this restricted model. The payoff is (2), a recursive distributional equation (RDE, see [AB05]) relating V and two independent copies of itself. We find that $V \sim \delta_\infty$ is the unique solution. Thus, the original model is recurrent.

The proof that $V \sim \delta_\infty$ is the unique solution of the RDE uses a seemingly novel bootstrapping argument. We assume that V dominates a Poisson and show that in fact, V dominates a Poisson with slightly larger mean. It follows by repeating this argument that V dominates any Poisson. One obstacle to making this work is that using the typical definition of stochastic dominance, we cannot establish a base case for the argument. To get around this, we instead use a weaker stochastic ordering defined in terms of generating functions, under which the argument holds even when starting with the trivial base case of V dominating the distribution $\mathrm{Poi}(0)$. The situation is different for the frog model with initial conditions of $\mathrm{Poi}(\mu)$ frogs per site. In this setting, we use the usual notion of stochastic dominance and a related bootstrapping argument to prove a phase transation from transience to recurrence on any d-ary tree as μ increases [HJJ15].

We believe that these ideas are more widely applicable. Aldous and Bandyopadhyay study RDEs in general in [AB05]. Aother example of analyzing an RDE through an induced relation of generating functions can be found in [Liu98]. The RDE (2) in this paper is specific to our setting and much more complicated than the RDEs analyzed in either of these sources. Still, we think that our argument can be applied to other RDEs, including ones derived from similar interacting particle systems like activated random walk and the frog model with death.

For the transience part of Theorem 1, the idea is to dominate the frog model by a branching process. At the beginning of Section 3, we show in a few lines that a doubling branching random walk is transient on the 15-ary tree. A simple refinement in Proposition 18 improves this to $d \geq 6$. The case $d = 5$ uses a branching random walk with 27 particle types. This is significantly more complicated, and computing the transition probabilities requires computer assistance. Conceptually our approach could extend to a computer-assisted proof for transience when $d = 4$, but the demands of this theoretical program seem well beyond current processing power.

We present two other results besides Theorem 1. The first is a 0-1 law for transience and recurrence of the frog model on a d-ary tree that applies under more general initial conditions than one frog per site. For a given distribution ν on the nonnegative integers, we consider the frog model on a d-ary tree with the number of sleeping frogs on each vertex other than the root drawn independently from ν. The root initially contains one frog, which begins its life awake. Recall that when a site is visited for the first time, all sleeping frogs at that site are awoken. We refer to this as the frog model with i.i.d.-ν initial conditions. When $\nu = \delta_1$, this is the usual one-per-site frog model. This theorem complements the 0-1 law for recurrence proven in [GS09] in a frog model on $\mathbb{Z}$ with drift. It also plays an important role in [HJJ15], where we use it to show that the probability of recurrence for the frog model on a d-ary tree with i.i.d.-Poi(μ) initial conditions jumps abruptly from 0 to 1 as μ increases.

Theorem 4. *The frog model on $\mathbb{T}_d$ for any d and any i.i.d. initial conditions is recurrent with probability 0 or 1.*

Our final related result is that in contrast to the 0-1 law on $\mathbb{T}_d$, there is a graph on which the frog model is recurrent with probability strictly between 0 and 1.

Theorem 5. *Let G be the graph formed by merging the root of $\mathbb{T}_6$ and the origin of $\mathbb{Z}$ into one vertex. The frog model on G has probability $0 < p < 1$ of being recurrent.*

We remark that [Pop01] exhibits a frog model without a 0-1 law on $\mathbb{Z}^d$. In their example the initial distribution of frogs decays in the distance from the origin.

A few of our proofs would be simplified by changing the setting from d-ary to homogeneous trees. However, we are interested in applying these results to finite trees, and the infinite d-ary tree is more natural to work with from that perspective. In any event, the techniques underlying our theorems can all be cleanly modified to prove similar statements about the homogeneous tree.

2. Recurrence for the binary tree

An outline of our proof is as follows. We start by a defining a process that we call the self-similar frog model. A consequence of Proposition 7 is that the number of visits to the root in this model is smaller than in the original one. Thus it suffices to prove the self-similar frog model recurrent. To do this, we define the random variable V to be the number of returns to the root and set $f(x) = \mathbf{E}x^V$, the generating function of V. The self-similarity of our model established in Proposition 6 allows us to show in Proposition 9 that the generating function satisfies the relation $f = \mathcal{A}f$ for an explicit operator $\mathcal{A}$. In Lemma 10, we show that $\mathcal{A}$ is monotone on a large class of functions. Combining this with $f \leq 1$ on $[0,1]$, we get

$$f = \mathcal{A}^n f \leq \mathcal{A}^n 1,$$

and in Lemma 14, we prove that this converges to 0 as $n \to \infty$. This implies that $f \equiv 0$ and $V = \infty$ a.s.

This proof can be interpreted as an argument about stochastic orders. One can define a stochastic order by saying that if X and Y are nonnegative integer-valued random variables and $\mathbf{E}t^X \geq \mathbf{E}t^Y$ for $t \in (0,1)$, then X is *smaller in the probability generating function order* than Y. This order and an equivalent one called the Laplace transform order are discussed in in [SS07, Section 5.A]. From this perspective, each application of the operator $\mathcal{A}$ shows that the distribution of V is slightly larger in this stochastic order.

2.1. The non-backtracking frog model.

We will define the *non-backtracking frog model*, in which frogs move as random non-backtracking walks stopped at the root. More formally, we define the random non-backtracking walk $(X_n,\, n \geq 0)$ as a process taking values in $\mathbb{T}_d$, with $X_0 = x_0$. On its first step, the walk moves to a uniformly random neighbor of x_0. At every subsequent step, it chooses uniformly from its neighbors other than the one from which it arrived. Let $T = \inf\{n \geq 1\colon X_n = \varnothing\}$, taking this to be ∞ if the walk never visits $\varnothing$. Define the non-backtracking frog model by changing the frog's paths in the definition of the frog model from simple random walks to the stopped non-backtracking walks given by $(X_{n\wedge T},\, n \geq 0)$. Notice that the initial frog is never stopped in this model, and only one child of the root is ever visited. Call this child $\varnothing'$.

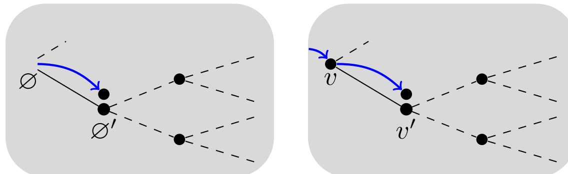

FIGURE 1. Conditional on v being visited, V and $V_{v'}$ are identically distributed in the self-similar model.

2.2. **The self-similar frog model.** We now make a further alteration to the non-backtracking frog model. When a frog x in this model is woken, if its first step away from the root takes it to a vertex that has already been visited by another frog, we stop x after this step. If multiple frogs visit a vertex for the first time in the same step, we stop all but one of them. We call the resulting model the *self-similar frog model*.

Let $V = V_{\varnothing'}$ be the number of visits to the root in the self-similar frog model. Note that only frogs initially sleeping in $\mathbb{T}_d(\varnothing')$, the subtree rooted at $\varnothing'$, have a chance of visiting the root. Suppose that vertex v is visited by a frog. Conditional on this, let v' be the child of v that the waking frog moves to next, and define $V_{v'}$ as the number of visits to v from frogs in $\mathbb{T}_d(v')$, the subtree rooted at v'. The following proposition explains how the self-similar frog model earns its name.

Proposition 6. *The distribution of $V_{v'}$ conditional on some frog visiting v and moving next to v' is equal to the (unconditioned) distribution of V.*

Proof. Let x be the frog that wakes vertex v and moves from there to v'. Besides x, all frogs that start outside of $\mathbb{T}_d(v')$ get stopped when they try to enter this subtree. Thus, from the time that v' is woken on, if we consider the model restricted to $\{v\} \cup \mathbb{T}_d(v')$, it looks identical to the original self-similar frog model (see Figure 1).

To turn this into a precise statement, consider the model from the time x visits v on. Ignore the frog initially at v. Freeze frogs when they visit v from $\mathbb{T}_d(v')$. Since no frogs ever enter $\mathbb{T}_d(v')$, the process depends only on the frogs initially in $\mathbb{T}_d(v')$ and the initial frog x, and relabeling vertices $\{v\} \cup \mathbb{T}_d(v')$ as $\{\varnothing\} \cup \mathbb{T}_d(\varnothing')$ in the obvious way produces a process identically distributed as the original self-similar frog model. Thus V and $V_{v'}$ are functionals of identically distributed processes. □

2.3. **Coupling the models.** Suppose we wanted to couple a simple and a non-backtracking random walk starting from a vertex v on the homogeneous tree $\mathbb{T}_d^{\text{hom}}$. Almost surely, there is a unique geodesic from v to infinity that intersects the walk infinitely many times, obtained by trimming away the backtracking portions from the walk. By symmetry, this geodesic is a uniformly random non-backtracking walk on $\mathbb{T}_d^{\text{hom}}$, coupled so that its path is a subset of the simple random walk's path. If we were working on $\mathbb{T}_d^{\text{hom}}$ and not $\mathbb{T}_d$, we could couple the non-backtracking and usual frog models as desired by coupling each frog in this way. To address the asymmetry of $\mathbb{T}_d$ at its root, our coupling of non-backtracking and normal frogs on $\mathbb{T}_d$ will involve an intermediate coupling with a random walk on $\mathbb{T}_d^{\text{hom}}$.

Proposition 7. *There is a coupling of the non-backtracking, the self-similar and the usual frog models so that the path of every non-backtracking (self-similar) frog path is a subset of the path of the corresponding frog in the usual model.*

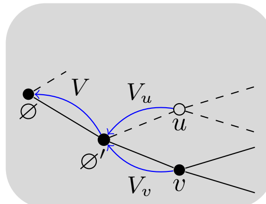

FIGURE 2. V is the total number of visits to $\varnothing$ in the self-similar process, V_v and V_u are the number of visits to $\varnothing'$ from frogs originally in $\mathbb{T}_2(v)$ and $\mathbb{T}_2(u)$, respectively. In the self-similar model V, V_v, and $V_u \mid \{u \text{ is visited}\}$ are identically distributed

Proof. First, we couple a non-backtracking walk to a simple random walk not on $\mathbb{T}_d$, but on $\mathbb{T}_d^{\mathrm{hom}}$. Let $(Y_n,\, n \geq 0)$ be a simple random walk on $\mathbb{T}_d^{\mathrm{hom}}$ starting at x_0. This random walk diverges almost surely to infinity, and there is a unique geodesic from x_0 to the path's limit. Let $(X_n, n \geq 0)$ be the path of this geodesic. By the symmetry of $\mathbb{T}_d^{\mathrm{hom}}$, the process (X_n) is a random non-backtracking walk from x_0.

Next, we consider $\mathbb{T}_d$ as a subset of $\mathbb{T}_d^{\mathrm{hom}}$ and define a new random walk $(Z_n,\, n \geq 0)$ by modifying (Y_n) as follows. First, delete all excursions of (Y_n) away from $\mathbb{T}_d$. This might leave the walk sitting at the root for consecutive steps; if so, we replace all consecutive occurrences of the root by a single one. This results in either an infinite path on $\mathbb{T}_d$ or a finite path on $\mathbb{T}_d$ truncated at a visit to the root. In the second case, we extend the path by tacking on an independent simple random walk to its end. It follows from the independence of excursions in simple random walk that the resulting process (Z_n) is a simple random walk on $\mathbb{T}_d$.

Let T be the first time past 0 that (X_n) hits the root, or ∞ if it never does. By our construction, $\{X_0, \ldots, X_T\} \subseteq \{Z_n,\, n \geq 0\}$. Thus we have coupled the stopped non-backtracking walks and simple random walks on $\mathbb{T}_d$. Coupling each frog in the non-backtracking frog model to the corresponding frog in the usual model gives the desired coupling between the non-backtracking and usual frog models. As the self-similar model is obtained by stopping frogs in the non-backtracking model, we obtain a coupling for it as well. ☐

2.4. **Generating function recursion.** We now apply the self-similarity described in Proposition 6 to obtain a relation satisfied by the generating function for the number of visits to the root in the self-similar model.

Definition 8. *Let V be the number of visits to the root in the self-similar frog model on $\mathbb{T}_2$. Define $f\colon [0,1] \to [0,1]$ by $f(x) = \mathbf{E}x^V$ with the convention that if $V = \infty$ a.s. then $f(1) = 0$.*

Proposition 9. *Define $\mathcal{A}$, an operator on functions on $[0,1]$, by*

$$\mathcal{A}g(x) = \frac{x+2}{3} g\Big(\frac{x+1}{2}\Big)^2 + \frac{x+1}{3} g\Big(\frac{x}{2}\Big)\bigg(1 - g\Big(\frac{x+1}{2}\Big)\bigg). \tag{1}$$

The generating function f satisfies $f = \mathcal{A}f$.

Proof. If $\mathbf{P}[V = \infty] = 1$ then $f \equiv 0$. This is easily checked to be a fixed point of $\mathcal{A}$. So, for the remainder of the argument suppose that $\mathbf{P}[V = \infty] < 1$.

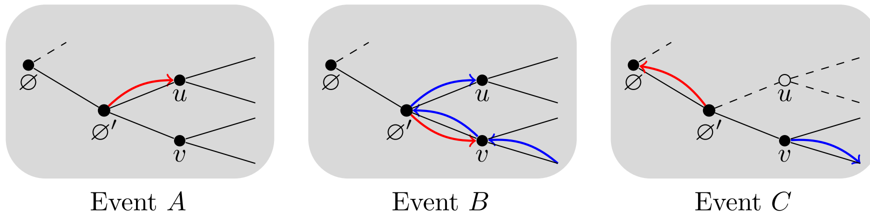

Figure 3. Outcomes that would result in events A, B and C, respectively. The path of the frog at $\varnothing'$ is red and the path of a frog from the subtree $\mathbb{T}(v)$ is blue.

The frog at the root in the self-similar model follows a non-backtracking path and visits one of its children and then one of this child's children; call these vertices $\varnothing'$ and v, respectively. We label the yet to be visited child u (see Figure 2). Define V_v and V_u to be the number of frogs which visit $\varnothing'$ that were originally sleeping in $\mathbb{T}_2(v)$ and $\mathbb{T}_2(u)$, respectively.

Proposition 6 guarantees that, since v has been visited, the random variable V_v is distributed identically to V. Conditionally on u being visited, the random variable V_u is also distributed identically to V. In fact, because frogs outside of $\mathbb{T}_2(u)$ affect $\mathbb{T}_2(u)$ only by determining whether or not u is visited, we can express V_u as $V_u = \mathbf{1}\{u \text{ is visited}\}V'$, where V' is distributed as V and is independent of V_v. This yields a description of V in terms of a pair of independent copies of itself:

$$V = \underbrace{\mathbf{1}\{\text{frog at } \varnothing' \text{ visits } \varnothing\}}_{\text{term 1}} + \underbrace{\mathbf{1}\{u \text{ is visited}\}\operatorname{Bin}(V', \tfrac{1}{2})}_{\text{term 2}} + \underbrace{\operatorname{Bin}(V_v, \tfrac{1}{2})}_{\text{term 3}}. \tag{2}$$

Term 1 accounts for a possible visit to $\varnothing$ by the frog started at $\varnothing'$. The conditional binomial distributions in terms 2 and 3 arise because each frog that visits $\varnothing'$ from u or v has a $\frac{1}{2}$ chance of jumping back to $\varnothing$.

Despite the independence between V_v and V', the three terms are dependent. For example, if term 1 is zero, then term 2 is more likely to be nonzero, since the frog at $\varnothing'$ not visiting $\varnothing$ makes it more likely to visit u. We unearth the pairwise independence of V_v and V' from (2) by conditioning on the following three disjoint events (see Figure 3):

- A. the frog starting at $\varnothing'$ visits u;
- B. the frog at $\varnothing'$ does not visit u, and a frog returns to $\varnothing'$ through v and visits u;
- C. no frog ever visits u.

Event A occurs with probability $1/3$. Given that k frogs return to $\varnothing'$ through v, the probability of C is $(2/3)2^{-k}$. Since the number of frogs returning to $\varnothing'$ through v is distributed identically to V, the probability of C is $\frac{2}{3}\mathbf{E}\left(\frac{1}{2}\right)^V$, which we call $2q/3$. The probability of event B is $2(1-q)/3$. Note that under our assumption $\mathbf{P}[V = \infty] < 1$ it follows that $0 < q < 1$.

Conditional on event A, B, or C, the terms in (2) are independent. Indeed, conditioning on whether u is visited makes terms 2 and 3 independent, and conditioning further on whether the frog at $\varnothing'$ visits u then makes term 1 independent of the other two. Now, we describe the distributions of each term in (2) conditional on events A, B, and C. For a given random variable X, we use $\operatorname{Bin}(X, p)$ to denote the random variable $\sum_{i=1}^{X} B_i$, where $\{B_i\}_{i\in\mathbb{N}}$ are distributed as Bernoulli(p), independent of each other and of X.

- Conditional on A, term 1 is 0 and terms 2 and 3 are distributed as independent $\mathrm{Bin}(V, 1/2)$.
- Conditional on B, term 1 is Bernoulli(1/2), term 2 is $\mathrm{Bin}(V, 1/2)$, and term 3 is $\mathrm{Bin}(V, 1/2)$ conditional on being strictly less than V (since at least one frog will visit u and not move to $\varnothing$).
- Conditional on C, term 1 is Bernoulli(1/2), term 2 is 0, and term 3 is $\mathrm{Bin}(V, 1/2)$ conditional on being equal to V (since every frog counted by V_v will return to $\varnothing$).

To summarize, let X' and X be distributed as $\mathrm{Bin}(V, 1/2)$. Let Y be distributed as $\mathrm{Bin}(V, 1/2)$ conditional on $\mathrm{Bin}(V, 1/2) < V$. Let Z be distributed as $\mathrm{Bin}(V, 1/2)$ conditional on $\mathrm{Bin}(V, 1/2) = V$. Let $I \sim$ Bernoulli(1/2). Take all of these to be independent. Conditioning on events A, B, and C, equation (2) yields

$$V \overset{d}{=} \begin{cases} X' + X & \text{with probability } 1/3, \\ I + X' + Y & \text{with probability } 2(1-q)/3, \\ I + Z & \text{with probability } 2q/3. \end{cases} \tag{3}$$

From this description of the distribution of V,

$$\begin{aligned} \mathbf{E}x^V &= \frac{1}{3}\mathbf{E}x^{X'+X} + \frac{2(1-q)}{3}\mathbf{E}x^{I+X'+Y} + \frac{2q}{3}\mathbf{E}x^{I+Z} \\ &= \frac{1}{3}\mathbf{E}x^{X'}\mathbf{E}x^X + \frac{2(1-q)}{3}\mathbf{E}x^I\mathbf{E}x^{X'}\mathbf{E}x^Y + \frac{2q}{3}\mathbf{E}x^I\mathbf{E}x^Z. \end{aligned} \tag{4}$$

Recall that a Bernoulli(p) random variable has generating function $px + 1 - p$ and that a random sum of i.i.d. random variables, $\sum_1^N X_i$, has generating function $g_N(g_{X_1}(x))$, where g_N and g_{X_1} are the generating functions of N and X_1. From these facts,

$$\begin{aligned} \mathbf{E}x^I &= \frac{x+1}{2}, \\ \mathbf{E}x^{X'} = \mathbf{E}x^X &= f\left(\frac{x+1}{2}\right). \end{aligned}$$

The generating functions $\mathbf{E}x^Y$ and $\mathbf{E}x^Z$ are a bit more complicated. The random variable Y is distributed as X conditional on $X < V$. Using the basic formula for conditional probability,

$$\begin{aligned} \mathbf{P}[Y = k] = \mathbf{P}[X = k \mid X < V] &= \frac{\mathbf{P}[X = k \text{ and } X < V]}{\mathbf{P}[X < V]} \\ &= \frac{\mathbf{P}[X = k] - \mathbf{P}[X = V = k]}{1-q} \\ &= \frac{\mathbf{P}[X = k] - 2^{-k}\mathbf{P}[V = k]}{1-q}. \end{aligned}$$

Thus, the probability generating function of Y is

$$\begin{aligned} \mathbf{E}x^Y &= \frac{1}{1-q}\sum_{k=0}^{\infty} x^k \bigl(\mathbf{P}[X = k] - 2^{-k}\mathbf{P}[v = k]\bigr) \\ &= \frac{1}{1-q}\mathbf{E}\Bigl[x^X - \Bigl(\frac{x}{2}\Bigr)^V\Bigr] \\ &= \frac{1}{1-q}\biggl(f\Bigl(\frac{x+1}{2}\Bigr) - f\Bigl(\frac{x}{2}\Bigr)\biggr). \end{aligned} \tag{5}$$

In (5) we are making use of the general fact that $\sum(a_n - b_n) = \sum a_n - \sum b_n$ so long as each sum is finite. Similarly,

$$\mathbf{P}[Z = k] = \mathbf{P}[X = k \mid X = V] = \frac{2^{-k}\mathbf{P}[V = k]}{q},$$

and so

$$\mathbf{E}x^Z = \frac{1}{q}\sum_{k=0}^{\infty} x^k 2^{-k}\mathbf{P}[V = k] = \frac{1}{q} f\Big(\frac{x}{2}\Big).$$

Using all of these generating functions and (4)

$$\begin{aligned} f(x) &= \frac{1}{3}\mathbf{E}x^{X'}\mathbf{E}x^X + \frac{2(1-q)}{3}\mathbf{E}x^I\mathbf{E}x^{X'}\mathbf{E}x^Y + \frac{2q}{3}\mathbf{E}x^I\mathbf{E}x^Z \\ &= \frac{1}{3} f\Big(\frac{x+1}{2}\Big)^2 + \frac{2(1-q)}{3}\bigg(\frac{x+1}{2} f\Big(\frac{x+1}{2}\Big)\frac{1}{1-q}\Big(f\Big(\frac{x+1}{2}\Big) - f\Big(\frac{x}{2}\Big)\Big)\bigg) \\ &\qquad\qquad + \frac{2q}{3}\bigg(\frac{x+1}{2q} f\Big(\frac{x}{2}\Big)\bigg) \\ &= \frac{x+2}{3} f\Big(\frac{x+1}{2}\Big)^2 - \frac{x+1}{3} f\Big(\frac{x+1}{2}\Big) f\Big(\frac{x}{2}\Big) + \frac{x+1}{3} f\Big(\frac{x}{2}\Big) = \mathcal{A}f(x), \end{aligned}$$

which establishes our claim. □

2.5. **Proving recurrence.** We have reduced the problem to understanding the properties of the operator $\mathcal{A}$ defined in (1). In Lemma 10, we prove that $\mathcal{A}$ is monotonic for functions belonging to the set $\mathscr{S} = \{g\colon [0,1] \to [0,1],\ \text{nondecreasing}\}$. In Lemma 11, we show that $\mathcal{A}$ maps $\mathscr{S}$ into itself, so that we can apply Lemma 10 after applying $\mathcal{A}$ iteratively. Finally, we show in Lemmas 12 and 14 that $\mathcal{A}^n 1 \to 0$. Starting at the conclusion of Proposition 9 (that the generating function f is a fixed point of $\mathcal{A}$), we will then apply these results to show that $f \equiv 0$, thus proving that the number of visits to the root in the self-similar frog model is a.s. infinite.

Lemma 10. *Let $g, h \in \mathscr{S}$. If $g \le h$, then $\mathcal{A}g \le \mathcal{A}h$.*

Proof. For $0 \le t \le 1$ define the interpolation between g and h by

$$i_t(x) = (1-t)\cdot g(x) + t \cdot h(x).$$

Since $\mathcal{A}i_0 = \mathcal{A}g$ and $\mathcal{A}i_1 = \mathcal{A}h$ it suffices to prove that $\frac{d}{dt}\mathcal{A}i_t(x) \ge 0$. Fix x and set $a = i_t\big(\frac{x+1}{2}\big)$ and $b = i_t\big(\frac{x}{2}\big)$ so that

$$\mathcal{A}i_t(x) = \frac{2+x}{3}a^2 + \frac{1+x}{3}b(1-a).$$

Define $s(a,b) = \mathcal{A}i_t(x)$. The chain rule implies

$$\frac{d}{dt}\mathcal{A}i_t(x) = \frac{\partial}{\partial a}s(a,b)\frac{da}{dt} + \frac{\partial}{\partial b}s(a,b)\frac{db}{dt}.$$

To prove $\frac{d}{dt}\mathcal{A}i_t \ge 0$ it suffices to prove each term in the above formula is nonnegative.

- The assumption that $g \le h$ implies that $\frac{d}{dt}i_t(x) = h(x) - g(x) \ge 0$ for all t and x. In particular, this implies $\frac{da}{dt}, \frac{db}{dt} \ge 0$.

- First we compute the partials
$$\frac{\partial}{\partial a}s(a,b) = 2a\frac{2+x}{3} - b\frac{1+x}{3} \qquad \text{and} \qquad \frac{\partial}{\partial b}s(a,b) = (1-a)\frac{1+x}{3}.$$
As g and h are nondecreasing, i_t is also nondecreasing in x for any fixed t. Hence $b \le a$. Along with the bound $a \le 1$ this immediately implies both partials are positive. □

Lemma 11. *If $g \in \mathscr{S}$, then $\mathcal{A}g \in \mathscr{S}$.*

Proof. All summands in (1) are nonnegative when $g(x) \le 1$, which implies that $\mathcal{A}g \ge 0$. By the previous lemma, $\mathcal{A}g \le \mathcal{A}1 \le 1$. We can conclude then that $0 \le \mathcal{A}g \le 1$. To see that $\mathcal{A}g$ is nondecreasing, suppose that $x \le y$, and let $a = g\big(\frac{y+1}{2}\big) - g\big(\frac{x+1}{2}\big)$. Then we have
$$\begin{aligned}\mathcal{A}g(y) &\ge \frac{x+2}{3}g\Big(\frac{x+1}{2}\Big)g\Big(\frac{y+1}{2}\Big) + \frac{x+1}{3}g\Big(\frac{x}{2}\Big)\bigg(1 - g\Big(\frac{y+1}{2}\Big)\bigg)\\ &= \mathcal{A}g(x) + \bigg(\frac{x+2}{3}g\Big(\frac{x+1}{2}\Big) - \frac{x+1}{3}g\Big(\frac{x}{2}\Big)\bigg)a \ge \mathcal{A}g(x).\end{aligned}$$
□

We now analyze the behavior of $\mathcal{A}$ on the family of generating functions for Poisson random variables. Recall that the generating function of $\mathrm{Poi}(a)$ is $e^{a(x-1)}$.

Lemma 12. *Define $g_a(x) = e^{a(x-1)}$ for all $a \ge 0$. For all $x \in [0,1]$,*
$$\mathcal{A}g_a(x) \le g_{a+c_a}(x),$$
where
$$c_a = \begin{cases} \frac{1}{3}e^{-2} & 0 \le a \le 4, \\ \frac{1}{3}e^{-a/2} & a \ge 4. \end{cases} \tag{6}$$

Proof. Applying the operator $\mathcal{A}$, we have
$$\begin{aligned}\mathcal{A}g_a(x) &= \frac{x+2}{3}e^{a(x-1)} + \frac{x+1}{3}e^{ax/2-a}\Big(1 - e^{a(x-1)/2}\Big)\\ &= g_a(x)r_{a/2}(x),\end{aligned} \tag{7}$$
where
$$r_b(x) = \frac{2+x}{3} + \frac{1+x}{3}\left(e^{-bx} - e^{-b}\right).$$
Note that $g_a(x)g_b(x) = g_{a+b}(x)$. It thus suffices to establish

Claim. *For $x \in [0,1]$, we have $r_b(x) \le g_{c_{2b}}(x)$.*

Proof of claim. We drop subscripts and let $r(x) = r_b(x)$ and $c = c_{2b}$. Calculus and a little algebra show that
$$r'(x) = \frac{1}{3}\left(1 - e^{-b} + e^{-bx}(-bx - b + 1)\right)$$
and
$$r''(x) = \frac{1}{3}e^{-bx}\left(b^2(x+1) - 2b\right).$$
We break the proof up into cases.

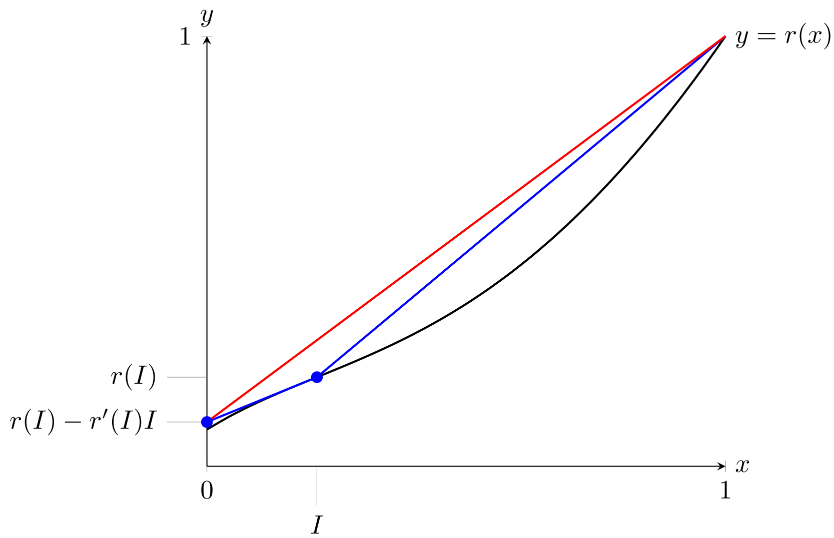

FIGURE 4. Above the graph of $y = r(x)$ sits the secant line from $(I, r(I))$ to $(1, 1)$ and the tangent line to $r(x)$ at $x = I$, both depicted in blue. Above them in red is the line $y = 1 + (1 - r(I) + Ir'(I))(x - 1)$.

- If $b \leq 1$ then $r(x)$ is concave down on $[0, 1]$ and the graph of $r(x)$ lies below its tangent line at $x = 1$. Thus
$$\begin{aligned} r(x) \leq 1 + r'(1)(x-1) &= 1 + \frac{1}{3}\big[1 - 2be^{-b}\big](x-1) \\ &\leq \exp\left[\frac{1}{3}\big(1 - 2be^{-b}\big)(x-1)\right]. \end{aligned}$$
It is easily verified that $\frac{1}{3}\big(1 - 2be^{-b}\big) \geq \frac{1}{3}\big(1 - 2e^{-1}\big) \geq e^{-2}/3$ for $b \leq 1$ and hence that $r(x) \leq g_c(x)$.
- If $b \geq 2$ then $r(x)$ is concave up on $[0, 1]$ and the graph of $r(x)$ lies below the secant line between $(0, r(0))$ and $(1, r(1))$. Thus as $r(1) = 1$ we have
$$\begin{aligned} r(x) \leq 1 + (1 - r(0))(x-1) &= 1 + \frac{1}{3}e^{-b}(x-1) \\ &\leq \exp\left[\frac{1}{3}e^{-b}(x-1)\right] = g_c(x). \end{aligned}$$
- If $1 < b < 2$ then there is a unique inflection point at $I = \frac{2}{b} - 1$ where r switches from concave down to concave up. Since r is concave up on $[I, 1]$, the graph of r lies below the line connecting $(1, 1)$ to $(I, r(I))$. Since r is concave down on $[0, I]$, to the left of I the graph of r lies below its tangent line at $(I, r(I))$. Thus the line segment from $(I, r(I))$ to $(0, r(I) - r'(I)I)$ lies above r, as in Figure 4. Therefore r lies below the line between (1,1) and $(0, r(I) - r'(I)I)$, and
$$r(x) \leq 1 + (1 - r(I) + Ir'(I))(x-1). \tag{8}$$

Next, we evaluate

$$1 - r(I) + Ir'(I) = 1 - \frac{1}{3}\biggl(2 + \Bigl(\frac{4}{b} - 1\Bigr)e^{b-2} - e^{-b}\biggr) \tag{9}$$

and try to bound this expression from below for $b \in (1,2)$. We proceed as calculus students, looking for critical points in this interval. The derivative with respect to b is

$$-\frac{1}{3}\biggl(\Bigl(\frac{4}{b} - 1 - \frac{4}{b^2}\Bigr)e^{b-2} + e^{-b}\biggr),$$

and a bit of algebra shows that the zeros of this expression are the solutions to

$$e^{2(b-1)}\left(\frac{2-b}{b}\right)^2 = 1.$$

Taking logarithms, we are interested in solutions to

$$b - 1 + \log(2-b) - \log b = 0.$$

on $(1,2)$. On this interval we can replace the logarithms with their power series expansions around 1 to rewrite the left-hand side as

$$b - 1 + 2\biggl(\frac{(b-1)^2}{2} + \frac{(b-1)^4}{4} + \frac{(b-1)^6}{6} + \cdots\biggr),$$

which is strictly positive for $b \in (1,2)$. Thus (9) has no critical values on $(1,2)$, and its minimum on $[1,2]$ is $e^{-2}/3$, occurring at $b = 2$. Applying this to (8), we have shown that

$$r(x) \le 1 + \frac{1}{3}e^{-2}(x-1) \le \exp\biggl[\frac{1}{3}e^{-2}(x-1)\biggr] = g_c(x).$$

This concludes the proof of both the claim and the lemma. □

Remark 13. Though the preceding lemma was an exercise in calculus, it has a probabilistic intepretation. If we think of $\mathcal{A}$ as acting directly on distributions instead of on their generating functions, this lemma shows that the result of applying $\mathcal{A}$ to $\mathrm{Poi}(a)$ is larger than $\mathrm{Poi}(a+c_a)$ in the probability generating function stochastic order described at the beginning of Section 2. The reason that $\mathcal{A}g_a$ simplifies so nicely in (7) is the Poisson thinning property, and the fact that $g_a(x)g_b(x) = g_{a+b}(x)$ is just the statement that the sum of independent Poissons is Poisson. There is a temptation to interpret $\mathcal{A}g_a(x) = g_a(x)r_{a/2}(x)$ as saying that the distribution resulting from applying $\mathcal{A}$ to $\mathrm{Poi}(a)$ is a convolution of $\mathrm{Poi}(a)$ and another distribution, but $r_{a/2}(x)$ is not monotone in x and hence not the generating function of a probability distribution.

Lemma 14. *For $x \in [0,1)$,*

$$\lim_{n\to\infty} \mathcal{A}^n g_0(x) = 0.$$

Proof. Define the sequence a_n by $a_0 = 0$ and $a_{n+1} = a_n + c_{a_n}$. By Lemmas 10, 11, and 12,

$$\mathcal{A}^n g_0(x) \le g_{a_n}(x) = e^{a_n(x-1)}.$$

We need to show that $a_n \to \infty$ as $n \to \infty$. Suppose this does not hold. Since the sequence is increasing, $a_n \to a$ for some constant a. Looking back at (6), this implies that c_{a_n} converges to a strictly positive limit. We can then choose n sufficiently large that $a_n + c_{a_n} > a$, a contradiction. □

Proof of Theorem 1 *(i).* Let f be the generating function $f(x) = \mathbf{E}x^V$ with V the number of visits to the root in the self-similar model frog model on the binary tree. By Proposition 9 we know that f satisfies the recursion relation $\mathcal{A}f = f$. Since f is a probability generating function, it satisfies $f(x) \leq 1 = g_0(x)$ for $x \in [0,1]$. Proposition 9 and Lemmas 10 and 11 imply $f(x) \leq \mathcal{A}^n g_0(x)$ for all n. By Lemma 14, f is identically zero on $[0,1)$. Thus the probability of any finite number of returns to the root is 0. This implies there are a.s. infinitely many returns to the root in the self-similar model. By the coupling in Proposition 7 each return in the self-similar model corresponds to a distinct return in the frog model. So, the frog model on the binary tree is a.s. recurrent. □

3. Transience for $d \geq 5$

The non-backtracking model was useful in the previous section because it was dominated by the usual frog model but was still recurrent. To prove transience we instead seek processes that dominate the frog model and can be proven transient. For example, consider a branching random walk on $\mathbb{T}_d$ whose particles split in two at every step. Using a union bound and asymptotics for the number of Dyck paths, one can show that the probability that any of the 2^{2n} particles at time $2n$ are at the root is $O(n^{-3/2})$ when $d \geq 15$, and hence the branching random walk visits the root finitely many times. As this walk can be naturally coupled to the frog model so that every awake frog has a corresponding particle, this proves that the frog model is transient for $d \geq 15$.

In this section, we will present a series of refinements to this argument to ultimately prove Theorem 1 (ii). In Proposition 18, we use a branching random walk on the integers and martingale techniques to prove transience for $d \geq 6$. We use this argument as a base for our proofs of Proposition 19, transience on the deterministic tree which alternates between five and six children, and Theorem 1 (ii), transience for $d \geq 5$. Both proofs use a multitype branching random walk. We included Proposition 19 because its calculations can easily be done by hand. In Theorem 1 (ii), on the other hand, we use a branching random walk with 27 types. The necessary calculations are intractable by hand, but they take only a few seconds on a computer. To get started we first address some difficulties that arise from reflection at the root. In doing so we also prove the 0-1 law described in Theorem 4.

3.1. **Couplings and 0-1 law.**

We will need to consider the frog model on several modifications of a rooted tree. We can handle these special cases all at once by working in a more general setting. Let Λ be any infinite rooted graph and H any graph. Enumerate finitely or countably many copies of Λ by $\Lambda(i)$, and form a graph G by adding an edge from the root of each $\Lambda(i)$ into H. Our next lemma shows that regardless of the number of sleeping frogs placed on H, a frog model is less transient on G than on Λ. Our motivation is the case when $\Lambda = \mathbb{T}_d$, as in Corollaries 16 and 17.

Lemma 15. *We consider two frog models. The first is on Λ with i.i.d.-ν initial conditions, for any measure ν on the nonnegative integers. The second is on G with the following initial conditions: one initially active frog at the root of $\Lambda(1)$; i.i.d.-ν sleeping frogs at all other vertices of $\bigcup_i \Lambda(i)$; and any configuration of sleeping frogs in H. Assume H is such that a random walk on G a.s. escapes H.*

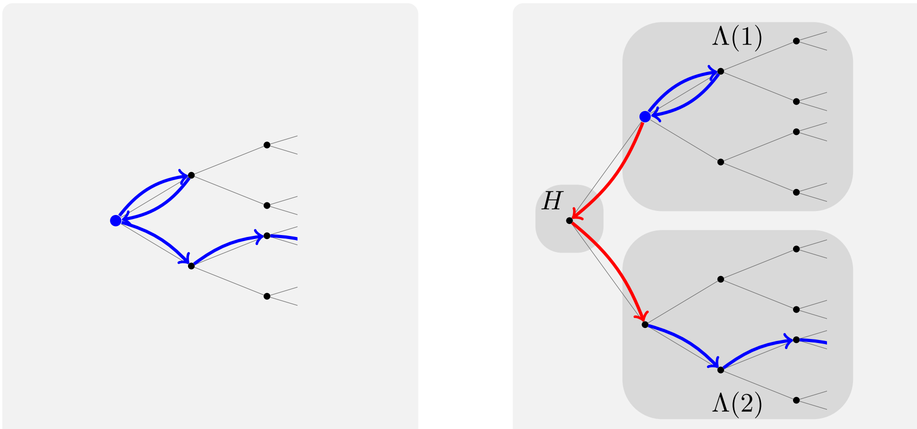

Graph Λ.

Graph G, with subgraphs H, $\Lambda(1)$, and $\Lambda(2)$ indicated.

FIGURE 5. The paths of frog x in Λ and frog x' in G. Frog x follows the blue steps of x and ignores the red steps.

Let V_G be the number of times the root of any $\Lambda(i)$ is visited in the frog model on G, not counting steps from H to a root. Let V_Λ be the the number of times the root is visited in the frog model on Λ. Then the two frog models can be coupled so that $V_G \geq V_\Lambda$.

Proof. Have the frog x, awake at the root $\varnothing \in \Lambda$, mime the frog x' that starts at the root of $\Lambda(1)$. As depicted in Figure 5, whenever x' enters H the frog x pauses at $\varnothing$; when x' re-enters any $\Lambda(i)$, the frog x begins following x' again.

When x visits a vertex that has yet to be visited, so will x'. Couple the number of sleeping frogs at the vertices occupied by x and x', and couple the newly awoken frogs to each other as with x and x'. In this way, x and all descendants on Λ perform simple random walks coupled to a frog on some $\Lambda(i)$. Thus, every visit to the root in Λ corresponds to a visit to level 0 in G, showing that $V_G \geq V_\Lambda$ under this coupling. □

We give two corollaries. The first will help us prove our transience results, and the second will help us prove a 0-1 law for transience and recurrence.

Corollary 16. *Consider the frog model on the $(d+1)$-homogeneous tree $\mathbb{T}_d^{hom}$ starting with a single active frog at the root, and with no sleeping frog at direct ancestors of the root. If level 0 is almost surely visited finitely many times in this model, then the frog model on $\mathbb{T}_d$ is almost surely transient.*

Proof. Let $G = \mathbb{T}_d^{\mathrm{hom}}$, thinking of it as countably many copies of $\mathbb{T}_d$ each joined at its root to a leaf of the infinite graph consisting of all the negative levels of G. The statement then follows immediately from Lemma 15. □

Corollary 17. *Run the frog model on $\mathbb{T}_d$, starting with an active frog not at the root but at level k. Assume that there are no sleeping frogs at levels $0, \ldots, k-1$ and i.i.d.-ν sleeping frogs at level k and beyond, with the exception of the location of the initial frog. The probability that the root is visited infinitely often in this model is at least the probability that the root is visited infinitely often in the usual frog model on $\mathbb{T}_d$ with i.i.d.-ν initial conditions.*

Proof. To set up our alternate frog model, let $G = \mathbb{T}_d$, thinking of it as d^k copies of $\mathbb{T}_d$ joined by a graph consisting of levels 0 to $k-1$ of the original graph. Let p be the probability that the root is visited infinitely often in the usual frog model. Let Y be the number of visits from level $k+1$ to k in the alternate model. It follows immediately from Lemma 15 that $\mathbf{P}[Y = \infty] \geq p$.

Let X be the number of visits to the root. We would like to show that

$$\mathbf{P}[Y = \infty,\, X < \infty] = 0, \tag{10}$$

thus proving that $\mathbf{P}[X = \infty] \geq p$. Call it a *dash* if a frog moves from level $k+1$ to a vertex v at level k, walks directly to the root, and then walks directly back to v. Let X' be the total number of dashes that occur. Conditional on a frog stepping from level $k+1$ to k, it makes a dash independently of all other frogs, since the model has no sleeping frogs at levels 0 to $k-1$. Whether or not it makes a dash is also independent of its own future number of visits from level $k+1$ to k and of dashes. Thus, at every visit from level $k+1$ to k, there is an independent $1/d(d+1)^{2k-1}$ chance of a dash, showing that

$$\mathbf{P}[Y = \infty,\, X' < \infty] = 0.$$

Since $X' < X$, this shows (10) and completes the proof. □

We are now ready to prove the 0-1 law.

Proof of Theorem 4. Suppose the probability that the root is visited infinitely often in the frog model on $\mathbb{T}_d$ with i.i.d.-ν initial conditions is $p > 0$. We wish to show that $p = 1$. The idea of the proof is to turn this statement into a more finite event, and then show that there are infinitely many independent opportunities for this event to occur. To this end, fix a constant N. We will show that at least N frogs visit the root with probability 1.

Claim. *For any k and N, there is a constant $K = K(k, N)$ such that the following statement holds. Consider the frog model on $\mathbb{T}_d$ starting with a frog at level k, with i.i.d.-ν sleeping frogs at levels $k, k+1, \ldots, K-1$ with the exception of the vertex of the initial frog, and with no sleeping frogs outside of this range. With probability at least $p/2$, this process makes at least N visits to the root.*

Proof. Consider the frog process with no sleeping frogs below level k, as in Corollary 17. Let E_K be the event that there are at least N visits to the root by frogs that are woken without the help of any frogs at level K or beyond. As $K \to \infty$, the event E_K converges upward to the event that there are at least N visits to the root by any active frog, which occurs with at least probability p by Corollary 17. Thus, for sufficiently large K, we have $\mathbf{P}[E_K] \geq p/2$. □

Now, we can find infinitely many independent events with probability $p/2$, each implying N visits to the root. Let $k_0, k_0' = 0$, and inductively choose k_i, k_i' as follows. Let $k_i' = K(k_{i-1}, N)$ from the claim. Let k_i be the level of the first frog that wakes up at level k_i' or beyond (assuming that ν is not a point mass at 0, there will be such a frog). Now, imagine a frog process starting with this frog, with no sleeping frogs below level k_i or at level k_{i+1} or beyond. These processes can all be embedded into the original frog process on $\mathbb{T}_d$, and each one independently has a $p/2$ chance of visiting the root at least N times, by the claim. Thus, the root is visited at least N times almost surely, for arbitrary N. □

3.2. **Proving transience.** Consider the branching random walk where each particle gives birth at each step either to one child to its left or to two children to its right. Formally, we define this as a sequence of point processes. Start with ξ_0 as a single particle at 0. With probability $1/(d+1)$, the point process ξ_1 consists of a single particle at -1; with probability $d/(d+1)$, it consists of two particles at 1. After this, each particle in ξ_n produces children in ξ_{n+1} in the same way relative to its position, independently of all other particles. We will use this branching random walk to prove the frog model transient for $d \geq 6$ and closely related processes to extend this down to $d = 5$.

Proposition 18. *For $d \geq 6$, the frog model on $\mathbb{T}_d$ is almost surely transient.*

Proof. Consider the frog model on $\mathbb{T}_d^{\mathrm{hom}}$, starting with no sleeping frogs at direct ancestors of the root, as in Corollary 16. When a frog jumps backward in this process, it never spawns a new frog, and when it moves forward, it sometimes does. Thus, the projection of this frog model onto the integers can be coupled with $(\xi_n,\, n \geq 0)$ so that every frog has a corresponding particle. By Corollary 16, proving that ξ_n visits 0 finitely many times a.s. proves that the frog model on $\mathbb{T}_d$ is transient a.s.

To determine the behavior of ξ_n, we define a weight function w on point process configurations. We refer to the position of a particle i in a point process configuration by $P(i)$ and define

$$w(\xi) = \sum_{i \in \xi} e^{-\theta P(i)}, \tag{11}$$

with θ to be chosen later. Letting $\mu = \mathbf{E} w(\xi_1)$ we have

$$\mathbf{E}[w(\xi_{n+1}) \mid \xi_n] = \mu w(\xi_n),$$

and so the sequence $w(\xi_n)/\mu^n$ is a martingale. As it is positive, it converges almost surely. When $\mu < 1$ this means $w(\xi_n) \to 0$. If a particle in ξ_n occupies the origin then $w(\xi_n) \geq 1$, and so infinitely many visits to the origin prevents $w(\xi_n)$ from converging. Hence, $\mu < 1$ implies the a.s. transience of ξ_n. (In fact, this holds when $\mu = 1$ as well, though we will not need this.) It then suffices to show that there exists θ making $\mu < 1$. We compute

$$\mu = \tfrac{1}{d+1} e^{\theta} + 2 \tfrac{d}{d+1} e^{-\theta}.$$

This is minimized by setting $\theta = \log(2d)/2$, which makes $\mu = 2\sqrt{2d}/(d+1)$. A bit of algebra shows that $\mu < 1$ when $d > 3 + 2\sqrt{2} \approx 5.83$. ☐

By using a multitype branching process, we can extend this proof to show transience for $\mathbb{T}_5$. Before we do so, we will show how it works in a setting where humans can do the math without much assistance.

Proposition 19. *Let $\mathbb{T}_{5,6}$ be the tree whose levels alternate between vertices with 5 children and vertices with 6 children, starting with the root having either 5 or 6 children. The frog model on this tree is transient a.s.*

Proof. Let $\mathbb{T}_{5,6}^{\mathrm{hom}}$ be the five-six children alternating homogeneous tree which contains $\mathbb{T}_{5,6}$ and place a sleeping frog at each vertex except for direct ancestors of the root of $\mathbb{T}_{5,6}$. Lemma 15 implies that it suffices to prove transience of this frog model on $\mathbb{T}_{5,6}^{\mathrm{hom}}$.

First note that a frog at a vertex with five children has different probabilities of moving forwards or backwards than a frog at a vertex with six children. By design the tree deterministically alternates, so a frog also alternates between each state.

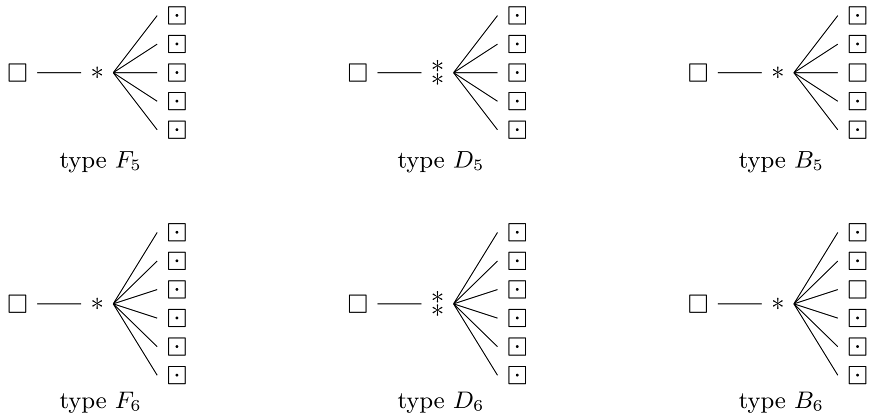

FIGURE 6. A depiction of the six particle types from the proof of Proposition 19. Each asterisk is a frog represented by the particle. The symbol $\boxdot$ signifies a vertex with a sleeping frog, and the symbol $\square$ represents a vertex with no sleeping frog.

When a frog moves backwards there is chance it immediately jumps forward to the same vertex, which will never spawn a new frog. Similarly, when two frogs occupy the same site there is a chance both jump forward to the same vertex, spawning at most one frog, not two. The idea is to introduce additional particle types that act like frogs in these more advantageous states.

Consider a multitype branching random walk on $\mathbb{Z}$ with six particle types, F_5, D_5, B_5, F_6, D_6, and B_6. The subscript accounts for whether a frog is at a vertex with 5 or 6 children. B particles represent frogs that have just stepped backward. D particles represent two frogs at once, the waker and wakee at a vertex where a frog has just woken up. Last, F particles represent single frogs with sleeping frogs present at all children. A visual depiction of these particle types is provided in Figure 6, and the distribution of children for each particle type is defined in Figure 7.

Let ζ_n be the branching random walk in which particles reproduce independently with the given child distributions. These distributions are chosen to match how the projections of frogs on the integers behave. Ignoring for a moment whether a frog is at a site with five or six children, when a frog jumps back it becomes of type B and when a new frog wakes it and its waker consolidate into a type D particle. Any extra frogs become a type F particle. These particles then reproduce independently on a "fresh" tree configured so that the particles always generate at least as many frogs as the projection of the actual frog model. For this reason we can couple the integer projection of the frog model on $\mathbb{T}_{5,6}^{\mathrm{hom}}$ with ζ_n so that the particles representing awake frogs are a subset of ζ_n. It therefore suffices to prove that ζ_n is transient.

To analyze ζ_n, we use a generalization of the martingale from Proposition 18 to the multitype setting, introduced in [Big76]. Let $\zeta_n = \sum_i \zeta_n^i$, where i ranges over the six particle types and ζ_n^i denotes the restriction of ζ_n to particles of type i. Recalling the weight function w given by (11), we define a matrix $\Phi(\theta)$ by

$$\Phi_{ij}(\theta) = \mathbf{E}_i\Big[w\big(\zeta_1^j\big)\Big].$$

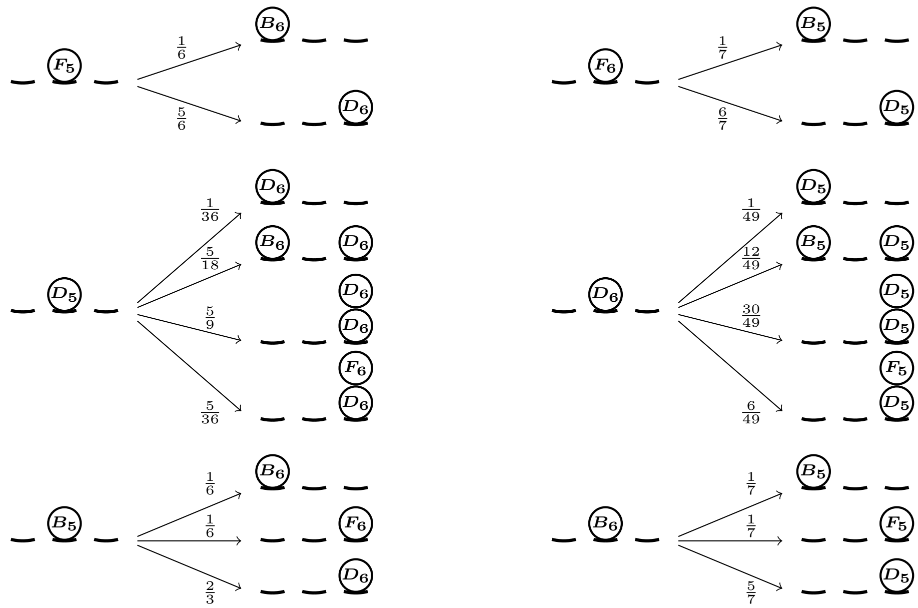

FIGURE 7. The distribution of children for each particle type in the proof of Proposition 19.

Here, we use $\mathbf{E}_i$ to denote expectation when ζ_0 is a single particle at the origin of type i. Let w_n denote a row vector whose ith entry is $w\bigl(\zeta_n^i\bigr)$. Then

$$\mathbf{E}[w_{n+1} \mid \zeta_n] = w_n \Phi(\theta). \tag{12}$$

Thus, for any eigenvalue λ and associated right eigenvector v of $\Phi(\theta)$,

$$\mathbf{E}[w_{n+1}v \mid \zeta_n] = w_n \Phi(\theta) v = \lambda w_n v,$$

and so $w_n v/\lambda^n$ is a martingale.

Since $\Phi(\theta)$ is a nonnegative irreducible matrix, there is a positive eigenvalue $\phi(\theta)$ equal to the spectral radius of $\Phi(\theta)$ by the Perron–Frobenius theorem. The eigenvector $v(\theta)$ associated with $\phi(\theta)$ has strictly positive entries. We then have a positive martingale $w_n v(\theta)/\phi(\theta)^n$. If $\phi(\theta) < 1$, then it follows as in Proposition 18 that the branching random walk visits 0 finitely often, thus proving that the frog model is almost surely transient.

All that remains is to find some value of θ such that $\phi(\theta) < 1$. Ordering the rows and columns F_5, D_5, B_5, F_6, D_6, B_6 and reading off $\mathbf{E}_i\bigl[w(\zeta_i^j)\bigr]$ from Figure 7,

$$\Phi(\theta) = \begin{bmatrix} 0 & 0 & 0 & 0 & \frac{5}{6}e^{-\theta} & \frac{1}{6}e^{\theta} \\ 0 & 0 & 0 & \frac{5}{36}e^{-\theta} & \frac{1}{36}e^{\theta} + \frac{55}{36}e^{-\theta} & \frac{5}{18}e^{\theta} \\ 0 & 0 & 0 & \frac{1}{6}e^{-\theta} & \frac{2}{3}e^{-\theta} & \frac{1}{6}e^{\theta} \\ 0 & \frac{6}{7}e^{-\theta} & \frac{1}{7}e^{\theta} & 0 & 0 & 0 \\ \frac{6}{49}e^{-\theta} & \frac{1}{49}e^{\theta} + \frac{78}{49}e^{-\theta} & \frac{12}{49}e^{\theta} & 0 & 0 & 0 \\ \frac{1}{7}e^{-\theta} & \frac{5}{7}e^{-\theta} & \frac{1}{7}e^{\theta} & 0 & 0 & 0 \end{bmatrix}.$$

Computing the eigenvalues of this matrix numerically, one can confirm that there exists θ with $\phi(\theta) < 1$; for example, $\phi(\log 3) \approx 0.9937$. To be completely certain that this is not an artifact of rounding, we will justify that $\phi(\log 3) < 1$ without using floating-point arithmetic. Observe that $\Phi(\log 3)$ has rational entries. Using the computer algebra system SAGE, we calculated $(\Phi(\log 3))^{66}$ using exact arithmetic, and we found that its largest row sum was less than 1. (The only significance of the 66th power is that it is the lowest one for which this is true.) This implies that all eigenvalues of $(\Phi(\log 3))^{66}$ are less than 1, which implies that all eigenvalues of $\Phi(\log 3)$ are less than 1 as well. The source code accompanying this paper includes this matrix and has instructions so that readers can easily check these claims. $\square$

Remark 20. We chose to include this proof to illustrate the technique we use to prove Theorem 1 (ii). Furthermore, this provides an example of proving the frog model is transient on an interpolation between different degree trees. This is relevant because the sharpest proof of Conjecture 2 would find exactly where the phase transition occurs between recurrence and transience on $\mathbb{T}_d$, perhaps between $d = 3$ and $d = 4$. Last, a natural generalization is a frog model on Galton-Watson trees. Our argument depends on the deterministic structure of $\mathbb{T}_{5,6}$ and we do not see an obvious way to generalize it.

Having proven transience for the frog model on $\mathbb{T}_d$ with $d \geq 6$ and on $\mathbb{T}_{5,6}$, we present our final refinement to prove the $\mathbb{T}_5$ case. The proof is essentially the same as the previous one, but with more particle types and a more difficult calculation.

Proof of Theorem 1 *(ii).* We define a particle type $P(a, b, c)$, for $a \geq 1$ and $b, c \geq 0$. A particle of type $P(a, b, c)$ represents a frogs on one vertex. There are no sleeping frogs on at least b of the vertex's children and on at least c of the vertex's siblings. In this scheme, the F types from the previous proof would translate to $P(1, 0, 0)$, the D types would translate to $P(2, 0, 0)$, and the B types would translate to $P(1, 1, 0)$.

We use 27 of these particles, $P(a, b, c)$ with $1 \leq a \leq 3$ and $0 \leq b, c \leq 2$. For particle type $P(a, b, c)$, consider the frog model on the homogeneous tree, starting with a frogs at the root. As usual, remove the sleeping frogs from direct ancestors of the root. Also remove the sleeping frogs from b children of the root and from c siblings. From each of these 27 initial states, we compute all possible states to which the frog model could transition in two steps, along with the exact probabilities of doing so. We then represent each of these final states as a collection of particles of the 27 types, at levels -2, 0, and 2 on the tree. In this way, we determine child distributions for each particle type, as in Figure 7. There is a slight ambiguity in how to do this, as a state of frogs can be represented in more than one way by these particle types. For example, four frogs on one vertex with one sibling vertex with no sleeping frog could be represented as two particles of type $P(2, 0, 1)$, or as one of type $P(3, 0, 1)$ and one of type $P(1, 0, 1)$. We always chose particles greedily, opting for as many 3-frog particles as possible. Whatever choice we make here, our branching random walk will still dominate the frog model, since when we assign new particles we "reset" the tree below them so the particles wake at least as many frogs as their counterpart in the frog model.

As in Proposition 19, it suffices to compute the matrix $\Phi(\theta)$ and show that for some choice of θ, its highest eigenvalue is less than one. Our attached source code computes $\Phi(\theta)$ exactly. We include additional documentation there explaining how we performed this calculation and describing the steps we took to make sure it was trustworthy. To avoid rounding issues, we proceeded as with Proposition 19. We exactly computed $(\Phi(\log 3))^{1024}$ by succesively squaring the matrix ten times, and we then checked that all of its row sums were less than 1. (There is no significance to the value $\log 3$; it just happens to work.) Thus, all eigenvalues of $(\Phi(\log 3))^{1024}$ are less than 1, implying that all eigenvalues of $\Phi(\log 3)$ are less than 1 as well. ☐

4. A frog model without a 0-1 law

We obtain a graph on which the frog model does not satisfy a 0-1 law by combining the transient graph $\mathbb{T}_6$ with the recurrent graph $\mathbb{Z}$ and proving that there is a positive probability that the frogs in each do not interact much.

To this end, we first prove two lemmas. The first shows there is a positive probability that the rightmost frog on the $\mathbb{Z}$ part of the combined graph escapes to ∞ while avoiding 0. This is necessary to rule out the possibility that too many frogs from $\mathbb{Z}$ get lost forever in $\mathbb{T}_6$. The second lemma proves there is a positive probability a frog model on $\mathbb{T}_6$ never returns to the origin, thus establishing a chance that the frog model on G gets lost in the transience of $\mathbb{T}_6$.

Lemma 21. *Consider the frog model on G, the graph formed by merging the root of $\mathbb{T}_6$ and the origin of $\mathbb{Z}$. With positive probability, the frogs starting at $1, 2, \dots$ in $\mathbb{Z}$ all wake up.*

Proof. Let

$$\delta_n = \frac{1}{8}\prod_{k=1}^{n-1}\left(1 - \frac{1}{(k+1)^2}\right),$$

taking $\delta_1 = 1/8$. We will show by induction that the frogs at $1, \dots, n$ wake up with probability at least δ_n. When $n = 1$, this holds because the initial frog moves right on its first step with probability 1/8. Now, assume the statement for n. Condition on the frogs at $1, \dots, n$ being woken. From the time when the frog at n is woken on, the two frogs there are independent random walkers, and at least one of them reaches $n+1$ before 0 with probability $1-1/(n+1)^2$ by a standard martingale argument. Thus frog $n+1$ is woken with probability at least $\delta_n\bigl(1 - 1/(n+1)^2\bigr) = \delta_{n+1}$, completing the induction.

Taking a limit of increasing events, the probability of the frogs at $1, 2, \dots$ all waking is at least $\lim_{n\to\infty}\delta_n > 0$. ☐

Lemma 22. *Let p' be the probability that the root is never visited past the initial frog's first move in the frog model on $\mathbb{T}_6$. It holds that $p' > 0$.*

Proof. As in Proposition 18, consider the frog model on $\mathbb{T}_6^{\mathrm{hom}}$, starting with no sleeping frogs at direct ancestors of the root. By Lemma 15, following the reasoning of Corollary 16, there is a coupling so that the number of visits to level 0 in the frog model on $\mathbb{T}_6^{\mathrm{hom}}$ is at least the number of visits to the root in the frog model on $\mathbb{T}_6$.

Now, recall from Proposition 18 the point process ξ, a branching random walk on $\mathbb{Z}$ in which particles split whenever they move in the positive direction. This process dominates the projection of the frog model on $\mathbb{T}_6^{\mathrm{hom}}$ onto the integers. Putting this all together, it suffices to show that with positive probability, ξ_n avoids 0 for all $n \geq 1$.

Suppose not, so ξ a.s. revisits 0. Since particles in ξ reproduce independently, this implies that ξ returns to the origin infinitely often. This is a contradiction, as we showed the opposite in proving Proposition 18. ☐

Proof of Theorem 5. In two steps, we bound the probability p of recurrence on $G = (\mathbb{Z} \cup \mathbb{T}_6)/\{0 \sim \varnothing\}$:

($p > 0$): The probability is 0 that any frog starting in $\mathbb{Z}$ wakes but fails to visit 0, by the recurrence of simple random walk on $\mathbb{Z}$. All frogs at $1, 2, \dots$ wake up with positive probability by Lemma 21, and on this event they therefore all visit 0.

($p < 1$): With probability 6/8 the first jump of the frog at $0 \in G$ will be into $\mathbb{T}_6$. Conditional on this, Lemma 22 guarantees a frog model in this configuration will never again visit the origin with probability $p' > 0$. Therefore, $1 - p \geq \frac{6}{8}p'$. □

5. Conjectures

Simulations suggest that for $d = 3$ the frog model on $\mathbb{T}_d$ is recurrent a.s., while for $d = 4$ the model is transient a.s. Our approach was to consider the frog model with the addition of stunning fences at each depth. When a frog jumps on a fence for the first time, it is stunned and stops moving. When all frogs are stunned at depth k, the fence turns off, and frogs resume their motion until they reach depth $k+1$ and are stunned again. Let $A_{d,k}$ be the number of stunned frogs that pile up on the fence at depth k before it turns off. We then examined the growth of $A_{d,k}$ in k for different choices of d. (The more obvious approach of directly simulating the frog model and counting visits to the root does not yield any obvious conclusions, as the rapid growth of the frog model makes it impossible to simulate very far.)

Martingale techniques tell us that the probability a frog at distance k from the root reaches the root before visiting depth $k+1$ is greater than cd^{-k} for some $c > 0$ independent of k. It follows that

$$\mathbf{E}[\text{visits to root between } k\text{th and } (k+1)\text{th stunnings}] \geq cd^{-k}\mathbf{E}[A_{d,k}].$$

So, if $\sum_{k=1}^{\infty} d^{-k}\mathbf{E}[A_{d,k}] = \infty$, then the expected number of visits to the root is infinite, which strongly suggests the model is recurrent.

This occurs if $kd^{-k}\mathbf{E}[A_{d,k}]$ is bounded from below. The data in Figure 8 summarizes the behavior of $kd^{-k}\mathbf{E}[A_{d,k}]$ to the maximum k we could easily simulate, $k = 18$. For $d = 4$, the slow growth of $A_{d,k}$ makes us suspect that the model is transient. The plot for $d = 2$ confirms Theorem 1 (i). Interestingly, $d = 3$ appears to be recurrent but very near criticality. The different growth for $d = 3$ is grounds for further investigation: it is possible $d = 2$ and $d = 3$ exhibit different forms of recurrence.

For $d = 2$ the simulated values of $k2^{-k}\mathbf{E}[A_{2,k}]$ appear to be growing linearly. This suggests a constant expected number of returns between each successive stunning. As the average number of steps for an individual frog between stunnings is constant, this could indicate that the average time between returns is also bounded away from infinity. If this is the case then the probability that there is a frog at the origin at time t would be bounded away from 0 as t gets large. However, for $d = 3$ it appears that $k3^{-k}\mathbf{E}[A_{3,k}]$ is sublinear. This might indicate that the average time between returns is unbounded and the probability of a frog occupying the origin at time t is approaching 0 as t approaches infinity. This leads us to ask the following:

Open Question 3. *Is the frog model on strongly recurrent on $\mathbb{T}_2$ but only weakly recurrent on $\mathbb{T}_3$?*

Such a result would have analogues with other interacting particle systems on trees. For example percolation on $\mathbb{T}_6 \times \mathbb{Z}$ has a three phases: no infinite components, infinitely many infinite components and a unique infinite component [GN90]. Similarly the contact process on trees can have strongly recurrent, weakly recurrent, and extinction phases [Pem92].

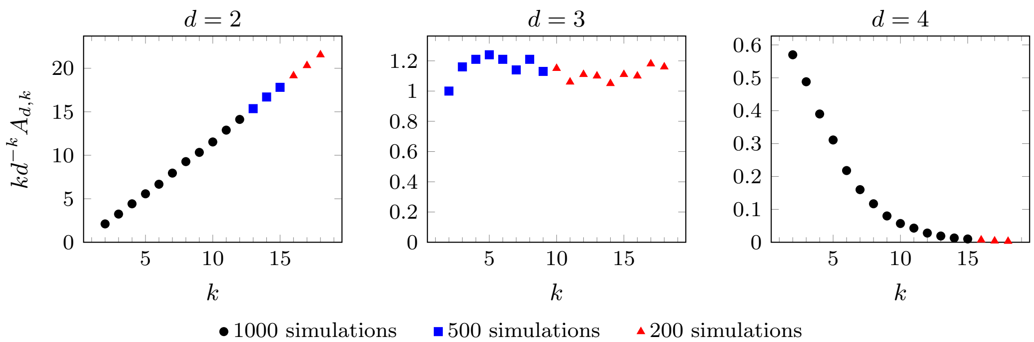

FIGURE 8. Plots of simulated values of $kd^{-k}A_{d,k}$ against k for $d = 2, 3, 4$. The number of simulations used in each estimate is shown in the chart.

Acknowledgments. We would like to thank Shirshendu Ganguly for his suggestions throughout the project. Christopher Fowler helped with a calculation, Avi Levy asked a question which led to the inclusion of Lemma 11 and James Morrow helped address potential concerns about roundoff error. Thanks to Soumik Pal who pointed out for large d the dynamics should be simpler—this remark sparked our study of transience. We thank Robin Pemantle for directing our attention to [AB05]. We also thank Nina Gantert who pointed out an inaccuracy in a previous version.

The first author was partially supported by NSF grant DMS-1308645 and NSA grant H98230-13-1-0827, the second author by NSF CAREER award DMS-0847661 and NSF grant DMS-1401479, and the third author by NSF RTG grant 0838212.

References

[AB05] David J. Aldous and Antar Bandyopadhyay, *A survey of max-type recursive distributional equations*, Ann. Appl. Probab. **15** (2005), no. 2, 1047–1110. MR 2134098 (2007e:60010)

[AMP02a] O. S. M. Alves, F. P. Machado, and S. Yu. Popov, *The shape theorem for the frog model*, Ann. Appl. Probab. **12** (2002), no. 2, 533–546. MR 1910638 (2003c:60159)

[AMP02b] Oswaldo Alves, Fabio Machado, and Serguei Popov, *Phase transition for the frog model*, Electron. J. Probab. **7** (2002), no. 16, 1–21, `http://ejp.ejpecp.org/article/view/115`.

[Big76] J. D. Biggins, *The first- and last-birth problems for a multitype age-dependent branching process*, Advances in Appl. Probability **8** (1976), no. 3, 446–459. MR 0420890 (54 #8901)

[BR10] Jean Bérard and Alejandro F. Ramírez, *Large deviations of the front in a one-dimensional model of $X + Y \to 2X$*, Ann. Probab. **38** (2010), no. 3, 955–1018. MR 2674992 (2011e:60219)

[BW03] Itai Benjamini and David B Wilson, *Excited random walk*, Electron. Comm. Probab **8** (2003), no. 9, 86–92.

[CQR09] Francis Comets, Jeremy Quastel, and Alejandro F. Ramírez, *Fluctuations of the front in a one dimensional model of $X + Y \to 2X$*, Trans. Amer. Math. Soc. **361** (2009), no. 11, 6165–6189. MR 2529928 (2010i:60281)

[DG99] D. J. Daley and J. Gani, *Epidemic modeling: an introduction*, Cambridge Studies in Mathematical Biology, vol. 15, Cambridge University Press, Cambridge, 1999. MR 1688203 (2000e:92042)

[DP14] Christian Döbler and Lorenz Pfeifroth, *Recurrence for the frog model with drift on $\mathbb{Z}^d$*, Electron. Commun. Probab. **19** (2014), no. 79, 13. MR 3283610

[DRS10] Ronald Dickman, Leonardo T. Rolla, and Vladas Sidoravicius, *Activated random walkers: facts, conjectures and challenges*, J. Stat. Phys. **138** (2010), no. 1-3, 126–142. MR 2594894 (2011h:82047)

[GN90] G. R. Grimmett and C. M. Newman, *Percolation in $\infty + 1$ dimensions*, Disorder in physical systems, Oxford Sci. Publ., Oxford Univ. Press, New York, 1990, pp. 167–190. MR 1064560 (92a:60207)

[GNR15] Arka P. Ghosh, Steven Noren, and Alexander Roitershtein, *On the range of the transient frog model on $\mathbb{Z}$*, arXiv:1502.02738, 2015.

[GS09] N. Gantert and P. Schmidt, *Recurrence for the frog model with drift on $\mathbb{Z}$*, Markov Process. Related Fields **15** (2009), no. 1, 51–58, `http://wwwmath.uni-muenster.de/statistik/gantert/frogs.pdf`. MR 2509423 (2010g:60170)

[HJJ15] Chris Hoffman, Tobias Johnson, and Matthew Junge, *From transience to recurrence with Poisson tree frogs*, arXiv:1501.05874, 2015.

[KPV04] Irina Kurkova, Serguei Popov, and M. Vachkovskaia, *On infection spreading and competition between independent random walks*, Electron. J. Probab. **9** (2004), no. 11, 293–315.

[KS06] Harry Kesten and Vladas Sidoravicius, *A phase transition in a model for the spread of an infection*, Illinois J. Math. **50** (2006), no. 1-4, 547–634. MR 2247840 (2007m:60298)

[Liu98] Quansheng Liu, *Fixed points of a generalized smoothing transformation and applications to the branching random walk*, Adv. in Appl. Probab. **30** (1998), no. 1, 85–112. MR 1618888 (99f:60151)

[LMP05] Élcio Lebensztayn, Fábio Machado, and Serguei Popov, *An improved upper bound for the critical probability of the frog model on homogeneous trees*, Journal of Statistical Physics **119** (2005), no. 1-2, 331–345 (English), `http://dx.doi.org/10.1007/s10955-004-2051-8`.

[Pem92] Robin Pemantle, *The contact process on trees*, Ann. Probab. **20** (1992), no. 4, 2089–2116. MR 1188054 (94d:60155)

[Pem07] ———, *A survey of random processes with reinforcement*, Probab. Surv. **4** (2007), 1–79. MR 2282181 (2007k:60230)

[Pop01] S.Yu. Popov, *Frogs in random environment*, Journal of Statistical Physics **102** (2001), no. 1-2, 191–201 (English).

[Pop03] Serguei Yu. Popov, *Frogs and some other interacting random walks models*, Discrete random walks (Paris, 2003), Discrete Math. Theor. Comput. Sci. Proc., AC, Assoc. Discrete Math. Theor. Comput. Sci., Nancy, 2003, pp. 277–288 (electronic). MR 2042394

[RS04] Alejandro F. Ramírez and Vladas Sidoravicius, *Asymptotic behavior of a stochastic combustion growth process*, Journal of the European Mathematical Society **6** (2004), no. 3, 293–334.

[RS12] 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

[SS07] Moshe Shaked and J. George Shanthikumar, *Stochastic orders*, Springer Series in Statistics, Springer, New York, 2007. MR 2265633 (2008g:60005)

[ST14] Vladas Sidoravicius and Augusto Teixeira, *Absorbing-state transition for Stochastic Sandpiles and Activated Random Walks*, arXiv:1412.7098, 2014.

[TW99] András Telcs and Nicholas C. Wormald, *Branching and tree indexed random walks on fractals*, Journal of Applied Probability **36** (1999), no. 4, 999–1011.

Department of Mathematics, University of Washington
E-mail address: `hoffman@math.washington.edu`

Department of Mathematics, University of Southern California
E-mail address: `tobias.johnson@usc.edu`

Department of Mathematics, University of Washington
E-mail address: `jungem@math.washington.edu`

FROM TRANSIENCE TO RECURRENCE WITH POISSON TREE FROGS

CHRISTOPHER HOFFMAN, TOBIAS JOHNSON, AND MATTHEW JUNGE

Abstract. Consider the following interacting particle system on the d-ary tree, known as the *frog model*: Initially, one particle is awake at the root and i.i.d. Poisson many particles are sleeping at every other vertex. Particles that are awake perform simple random walks, awakening any sleeping particles they encounter. We prove that there is a phase transition between transience and recurrence as the initial density of particles increases, and we give the order of the transition up to a logarithmic factor.

1. Introduction

We study a system of branching random walks known as the frog model, and we discover a phase transition as the initial state becomes more saturated with particles. Similar phase transitions have been observed in related models, including activated random walk [DRS10, ST14], reinforced random walk [Pem88], killed branching random walk [AHZ13] and the contact process [Pem92].

The frog model starts with a single particle awake at the root of a graph and sleeping particles at the other vertices. The initial configuration of sleeping particles can be deterministic or random. Particles that are awake perform independent simple random walks in discrete time. When a vertex with sleeping particles is first visited, all of the particles at the site wake up and each begins its own walk. The name "frog model" was coined in 1996 by Rick Durrett; we continue the zoomorphism and refer to the particles as frogs. As with other interacting particle systems, the frog model is often motivated as a model for the spread of a rumor or infection (see [AMP02a], for instance). It and its variants have also found interest as models of combustion [RS04, CQR09, RS12], generally with particles moving in continuous time.

We call a realization of the frog model *recurrent* if the root is visited infinitely often by frogs and *transient* if not. Even if each individual frog is transient, the aggregate of visits to the root can still be infinite. For this reason, the transience or recurrence of the frog model gives a measurement of its growth, and the question of transience or recurrence for the frog model on a given graph is one of the most fundamental ones.

The first ever published result on the frog model is that it is recurrent on $\mathbb{Z}^d$ with one sleeping frog per site for all d [TW99]. In fact, the frog model on $\mathbb{Z}^d$ is recurrent for any i.i.d. initial configuration of sleeping frogs [AMPR01]. It is natural to wonder if a sparser configuration changes the behavior. [Pop01] exhibits a threshold at which a frog model with Bernoulli($\alpha\|x\|^{-2}$) frogs at each $x \in \mathbb{Z}^d$ switches from transience to recurrence. A similar

Date: January 23, 2015.

2010 *Mathematics Subject Classification.* 60K35, 60J80, 60J10.

Key words and phrases. Frog model, transience, recurrence, phase transition.

The first author was partially supported by NSF grant DMS-1308645 and NSA grant H98230-13-1-0827 and the second author by NSF grant DMS-1401479.

phenomenon occurs when the walks have a bias in one direction: [GS09] finds that on $\mathbb{Z}$, the model is recurrent if and only if the number of sleeping frogs per site has infinite logarithmic moment. Recently this result was partially extended to $\mathbb{Z}^d$ in [DP14] and worked out in finer detail in [GNR15].

Let $\mathbb{T}_d$ denote the full infinite d-ary tree, in which the root has degree d and all other vertices degree $d+1$. The question of transience or recurrence on $\mathbb{T}_d$ is especially subtle. On one hand, the number of sleeping frogs grows exponentially with the distance from the root. On the other hand, each frog that wakes up has a drift away from the root; its probability of visiting the root shrinks exponentially as the starting vertex of the frog moves outward. The question of whether $\mathbb{T}_d$ is transient for the one-per-site model is posed in [AMP02b] and again in [Pop03] and [GS09]. Surprisingly, the answer depends on the degree of the tree. In [HJJ15], we prove that the one-per-site frog model is recurrent on the binary tree and transient on d-ary trees with $d \geq 5$.

We conjecture that the one-per-site frog model is recurrent for $d = 3$ and transient for $d = 4$. While we would like to pin this down and complete the picture of transience and recurrence for the one-per-site frog model on trees, we believe that the most interesting aspect of this work is that the frog model on trees is teetering on the edge between recurrence and transience. The point of this paper is to demonstrate this more precisely. We consider the frog model on $\mathbb{T}_d$ with i.i.d. $\mathrm{Poi}(\mu)$ sleeping frogs at each site. Our result is a phase transition between recurrence and transience as μ varies:

Theorem 1. *Consider the frog model on a d-ary tree with* $\mathrm{Poi}(\mu)$ *sleeping frogs per site. For all* $d \geq 2$*, there exists a critical value* $\mu_c(d) > 0$ *such that the model is recurrent a.s. if* $\mu > \mu_c(d)$ *and transient a.s. if* $\mu < \mu_c(d)$*. The critical value satisfies*

$$Cd < \mu_c(d) < C' d \log d$$

for some constants C *and* C'*.*

Proof. By a straightforward coupling, the probability of recurrence is monotone in μ. By [HJJ15, Theorem 4], the probability of recurrence is either 0 or 1. The theorem is then an immediate consequence of Propositions 6 and 15, where we prove recurrence and transience, respectively. ☐

Contrast our result with the frog model on $\mathbb{Z}^d$, which is is recurrent for any i.i.d. configuration of sleeping frogs [AMPR01]. To show the existence of the recurrence phase, we consider a restricted process that lets us take advantage of the recursive structure of $\mathbb{T}_d$. We then use a bootstrapping argument, showing that the number of returns to the root is stochastically larger and larger at each step. We establish the transience phase essentially by dominating the model with a branching random walk, using a similar argument as in [HJJ15]. As in that paper, the most difficult part is recurrence. Our result is an advance in that we are able to show recurrence on any d-ary tree with enough sleeping frogs. In [HJJ15], we prove recurrence only for $d = 2$, and the proof does not apply to a general choice of d; even extending it to $d = 3$ seems difficult. The argument here relies on having Poisson many sleeping frogs at each site, however, and thus neither result implies the other. A more detailed comparison between the recurrence proofs in the two papers is in Section 2.4.

Further Questions. A nice general survey on the frog model can be found in [Pop03]. Here we pose four questions specifically related to the frog model on trees.

The question most directly related to our paper is to better estimate the critical value $\mu_c(d)$. We are interested in both the asymptotic behavior and precise values for small d.

Open Question 2. *What is the correct order of $\mu_c(d)$ as $d \to \infty$? Also, what is the value of $\mu_c(d)$ for small d?*

We suspect that $\mu_c(d) = \Theta(d)$. As for the second question, the best bounds we can prove for $d = 2$ are $.125 \le \mu_c(2) \le 1.13$ (see Section 2.4).

As a start at considering the frog model on less regular graphs, we would like to know if the analogue of our result holds on Galton-Watson trees.

Open Question 3. *Consider a frog model with* $\mathrm{Poi}(\mu)$ *frogs at each site of an infinite Galton-Watson tree. As μ varies, does a phase transition occur between transience and recurrence?*

We are also interested in the relationship between the frog model and the degree distribution of the tree.

Open Question 4. *Does the recurrence of the frog model on a Galton-Watson tree depend on the entire degree distribution or just the maximal degree? Concretely, consider a one-per-site frog model on a Galton-Watson tree where each vertex has probability p of having two children and probability $1-p$ of having five children.* [HJJ15, Theorem 1] *implies that this is recurrent when $p = 1$ and transient when $p = 0$. Is it recurrent for any $p < 1$?*

This dependence on the maximal degree of the tree alone is seen in the contact process (see [Pem92] and [PS01, Proposition 2.5]).

Our next question comes from Itai Benjamini and concerns the frog model on finite trees. Define the *cover time* to be the expected time for every frog to wake up in a one-per-site frog model on the full d-ary tree with height n. We call this the cover time since it is equivalent to the time when every site is visited. A naive bound on the cover time is $O(n^2 d^n)$, the expected time for a single random walk to visit every site, as shown in [Ald91]. We have an unpublished proof improving this to $O(n^5(d/\sqrt{2})^n)$, but we suspect the correct value is polynomial.

Open Question 5. *Is the cover time for the one-per-site frog model on a d-ary tree of height n polynomial in n?*

Possibly the cover time on finite trees relates to the recurrence and transience properties on the corresponding infinite tree. For instance, it would be exciting to see that the cover time is polynomial in the height of the tree for $d = 2$ but exponential for higher d. This would be reminiscent of the contact process, which behaves similarly on finite lattices and trees as on their infinite counterparts [DL88, DS88, DST89, CMMV14].

2. Recurrence

We start with a sketch. Let ν' be the law of the number of visits to the root in the frog model with $\mathrm{Poi}(\mu)$ frogs at each site. To get some regularity, we restrict the motion of awakened frogs to the non-backtracking component of their ranges. Call this the *non-backtracking frog model* (more details are in Section 2.1) and let ν be the law of the number of visits to the root in this model. A coupling argument in Proposition 7 confirms the intuition that

$$\nu \preceq \nu'. \tag{1}$$

Here $\preceq$ denotes *stochastic dominance*, i.e. $\nu([x,\infty)) \le \nu'([x,\infty))$ for all x.

In Section 2.2 we define an operator $\mathcal{A}$ under which the image of ν has an interpretation in an even more restricted frog model. First a bit of notation (see Figure 1). Say the initial

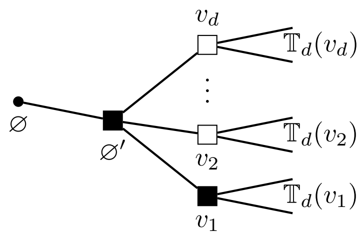

FIGURE 1. The frog from $\varnothing$ visits $\varnothing'$ and v_1. Suppose at most one frog in the non-backtracking frog model is allowed to enter each $\mathbb{T}_d(v_i)$ and only frogs woken at $\varnothing'$ and emerging from $\mathbb{T}_d(v_1)$ can enter other subtrees. We see in Lemma 9 that the number of visits to $\varnothing$ is stochastically fewer than ν and is distributed as $\mathcal{A}\nu$.

non-backtracking frog moves down the tree from the root $\varnothing$ to $\varnothing'$ and then to v_1. Let $v_2, \dots, v_d$ be the other children of $\varnothing'$ and let $\mathbb{T}_d(v_i)$ denote the subtree rooted at v_i. The measure $\mathcal{A}\nu$ is the law of the number of visits to the root in the non-backtracking frog model with two further restrictions:

(i) At most one frog can enter $\mathbb{T}_d(v_i)$ for each $1 \leq i \leq d$.
(ii) Only frogs woken at $\varnothing'$ and those emerging from $\mathbb{T}_d(v_1)$ can enter the other $\mathbb{T}_d(v_i)$.

The advantage of (i) is that it makes the number of frogs emerging from the activated subtrees i.i.d. random variables. The advantage of (ii) is that it simplifies which subtrees become activated (see Lemma 8). Intuitively, these restrictions reduce the number of visits to the root. This is made rigorous in Lemma 9 where we prove that

$$\mathcal{A}\nu \preceq \nu. \tag{2}$$

We stress that this a special property of ν. In fact, the essence of our argument is to show that when μ is large enough, (2) can hold only if $\nu = \delta_\infty$.

Section 2.3 explores properties of $\mathcal{A}$. In Lemma 10 we show that $\mathcal{A}$ is monotonic, meaning that for two probability measures π_1 and π_2,

$$\text{if } \pi_1 \preceq \pi_2, \text{ then } \mathcal{A}\pi_1 \preceq \mathcal{A}\pi_2. \tag{3}$$

Lemma 11 shows that $\mathcal{A}$ acts nicely on the Poisson distribution. In fact, by writing the Poisson distribution in a nonstandard way (see Lemma 13), we can compare $\mathcal{A}\,\mathrm{Poi}(\lambda)$ with $\mathrm{Poi}(\lambda+\epsilon)$. We carry this out in Proposition 14, where we show that when $\mu \geq 2(d+1)\log d$, there exists ϵ such that

$$\mathrm{Poi}(\lambda+\epsilon) \preceq \mathcal{A}\,\mathrm{Poi}(\lambda) \tag{4}$$

for all $\lambda \geq 0$. This is where the value of μ plays a role. Proving (4) reduces to comparing two binomial distributions with parameters depending on μ.

Now we explain how (1), (2), (3) and (4) imply the recurrence part of Theorem 1.

Proposition 6. *If $\mu > 2(d+1)\log d$, then the frog model is recurrent a.s. on the d-ary tree with an initial configuration of* $\mathrm{Poi}(\mu)$ *sleeping frogs per vertex.*

Proof. By (1) it suffices to prove that ν is a point mass at infinity. From (2) we have

$$\mathrm{Poi}(0) \preceq \mathcal{A}\nu \preceq \nu.$$

Statement (3) implies this relation is preserved under iterations of $\mathcal{A}$. Moreover, (4) lets us increase the Poisson term by ϵ with each iteration. In symbols this says that for all $n \geq 1$,

$$\mathrm{Poi}(\epsilon n) \preceq \mathcal{A}^n \nu \preceq \mathcal{A}^{n-1} \nu \preceq \cdots \preceq \mathcal{A}\nu \preceq \nu.$$

Taking $n \to \infty$ implies that ν is a point mass at infinity, and so the frog model is recurrent almost surely. ☐

In the rest of this section, we will carry out this plan and prove statements (1)–(4). First, we give some notation. Recall that $\preceq$ denotes stochastic domination. We also use the notation $X \preceq Y$ to indicate that the law of X is stochastically dominated by the law of Y. An equivalent condition to stochastic dominance is that $\pi_1 \preceq \pi_2$ if and only if there exists a coupling (X, Y) with $X \sim \pi_1$, $Y \sim \pi_2$, and $X \leq Y$ a.s. A thorough reference on stochastic domination is [SS07].

For a nonnegative random variable N, we use $\mathrm{Poi}(N)$ to denote a mixture of Poisson distributions; when we write $X \sim \mathrm{Poi}(N)$, we mean that X is coupled with N such that the distribution of X conditional on $N = n$ is $\mathrm{Poi}(n)$. If $N \sim \pi$, we also use $\mathrm{Poi}(\pi)$ to denote the same Poisson mixture. We similarly use the notations $\mathrm{Bin}(N, p)$ and $\mathrm{Bin}(\pi, p)$.

2.1. **The non-backtracking frog model.** A *random non-backtracking walk* on $\mathbb{T}_d$ starting at a vertex x_0 moves in its first step to a uniformly random neighbor of x_0. In all subsequent steps it moves to a vertex chosen uniformly from all its neighbors except for the one it just arrived from.

Suppose that $(S_n, n \geq 0)$ is a random non-backtracking walk starting from x_0, stopped if it arrives at the root at step 1 or beyond. (If x_0 is the root, then it is never stopped.) Define the *non-backtracking frog model* just as the usual frog model, except that the motion of a frog waking at x_0 is an independent copy of (S_n), rather than a simple random walk. The advantage is that when a non-backtracking frog moves away from the root, it will forever remain in the just-entered subtree. This gives the model more self-similarity. As shown in [HJJ15, Proposition 7], (S_n) can be coupled with a simple random walk on $\mathbb{T}_d$ starting from x_0 so that its path is a subset of the simple random walk's. This lets us relate the non-backtracking and usual frog models, proving (1):

Proposition 7. *Let ν and ν' be the laws of the number of returns to the root in the non-backtracking and usual frog models on $\mathbb{T}_d$, respectively, both with* $\mathrm{Poi}(\mu)$ *sleeping frogs per vertex. Then $\nu \preceq \nu'$.*

Proof. It suffices to show that we can couple the two models so that at least as many frogs visit the root in the usual model as in the non-backtracking model. We construct the coupling as follows. For each vertex $v \in \mathbb{T}_d$, make the number of sleeping frogs on v identical in the two models. Make each frog's path in the non-backtracking model a subset of the corresponding frog's path in the usual model as previously described. Thus any frog woken in the non-backtracking model is also woken in the usual model, and any visit to the root in the non-backtracking model corresponds to a visit in the usual model. ☐

For the remainder of this section we only consider the non-backtracking frog model. We record an observation: Suppose the initial frog in the non-backtracking model steps from the root $\varnothing$ to a child $\varnothing'$. Since frogs are stopped at the root, no other child of the root besides $\varnothing'$ is ever visited, and all action occurs in the subtree rooted at $\varnothing'$.

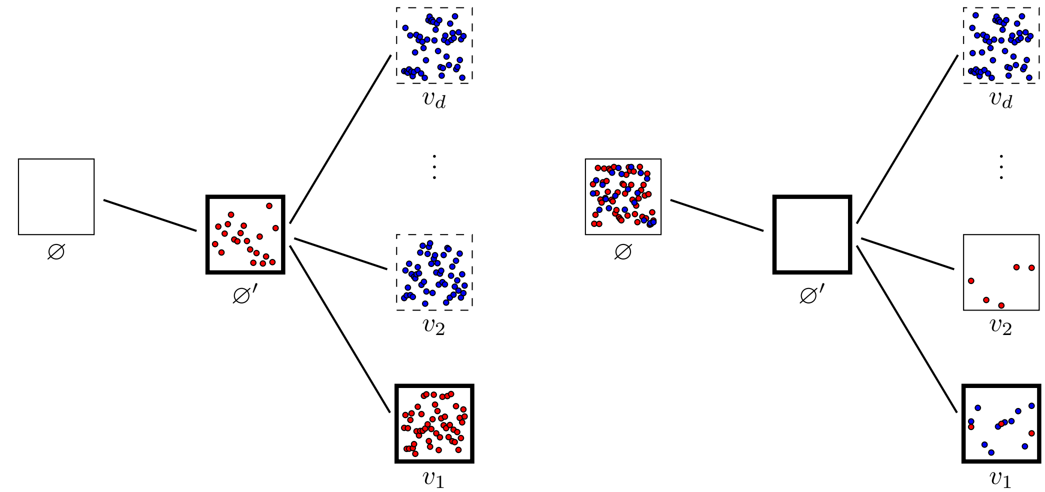

(a) **Initial state:** particles at $\varnothing'$ and v_1 will move first and possibly release a second wave of particles from $v_2, \dots, v_d$.

(b) **Terminal state:** $\#\{\text{particles at } \varnothing\} \sim \mathcal{A}\pi$.

FIGURE 2. An interacting particle system related to the frog model. Initially, the number of active particles at $\varnothing'$ is distributed as $\text{Poi}(\mu)$, and the number of active particles at v_1 is distributed according to some probability measure π. Active particles take random non-backtracking steps until reaching a leaf. For each $2 \leq i \leq d$, if any of these particles reach v_i, then a new π-distributed batch of particles is released at v_i. These second-wave particles do not activate other vertices.

2.2. **Formal definition of $\mathcal{A}$.** Fix a probability measure π on the nonnegative integers. We will define $\mathcal{A}\pi$ to be the probability measure for the number of particles ending at $\varnothing$ (see Figure 2) in the random system of non-backtracking particles described below.

The setting for the particle system is a star graph, consisting of a central vertex connected to $d+1$ leaf vertices. In a slight abuse of notation, we reuse the vertex names from Figure 1, calling the central vertex $\varnothing'$ and the leaves $\varnothing$ and $v_1, \dots, v_d$. Let $X \sim \text{Poi}(\mu)$ and $X_1, \dots, X_d \sim \pi$, all independent. Place X particles at $\varnothing'$ and X_i particles at each v_i. Each particle if activated will perform an independent random non-backtracking walk until it halts at a leaf.

Initially, only the particles at $\varnothing'$ and at v_1 are active. If one of these first-wave particles lands at v_i for $i \geq 2$, then the particles there are activated and begin independent non-backtracking random walks until reaching a leaf. These second-wave particles do not activate other particles; only the first-wave particles have that power. The number of particles that finish at $\varnothing$ in this system is a random variable, and we define $\mathcal{A}\pi$ as its law. With these dynamics, we can summarize the system as follows:

- Particles at $\varnothing'$ move to one of $\{\varnothing, v_1, \dots, v_d\}$ each with probability $1/(d+1)$.
- Particles at v_1 move to one of $\{\varnothing, v_2, \dots, v_d\}$ each with probability $1/d$.
- If a first-wave particle visits v_i, the particles at v_i move to $\varnothing$ with probability $1/d$.

For $2 \leq i \leq d$, let E_i be the event that a first-wave particle ends at v_i. The following lemma follows from the definition of $\mathcal{A}$. Informally it says that conditional on how many of the events $E_2, \dots, E_d$ occur, the number of second-wave particles ending at $\varnothing$ is a sum of independent thinned copies of π.

Lemma 8. *Conditional on $\sum_{i=2}^{d} \mathbf{1}_{E_i} = u$, the number of second-wave particles ending at $\varnothing$ is distributed as the sum of u independent $\mathrm{Bin}(\pi, 1/d)$-distributed random variables.*

Proof. If E_i occurs then by definition a π-distributed batch of particles is released at v_i. With probability $1/d$ each released particle halts at $\varnothing$. As particles move independently the total number is distributed as $\mathrm{Bin}(\pi, 1/d)$. Since the second-wave particles cannot wake other sites, the total number of particles to arrive is distributed as claimed. □

Now, we show the connection between this operator and the frog model.

Lemma 9. *Let ν be the distribution of number of returns to the root in the non-backtracking frog model on the d-ary tree with sleeping frog distribution $\mathrm{Poi}(\mu)$. Then $\mathcal{A}\nu \preceq \nu$.*

Proof. Let $\mathbb{T}_d(x)$ denote the subtree of $\mathbb{T}_d$ rooted at a given vertex x. Recall that no children of the root other than $\varnothing'$, the child visited by the initial frog, are ever visited. In light of this, it will be helpful to think of the non-backtracking frog model as taking place on $\varnothing \cup \mathbb{T}_d(\varnothing')$ rather than on all of $\mathbb{T}_d$.

We say that the frogs sleeping on some vertex $v \in \mathbb{T}_d(v_1)$ *wake within* $\mathbb{T}_d(v_1)$ if there exists a chain of vertices $x_1, \dots, x_m = v$ all in $\mathbb{T}_d(v_1)$ such that the initial frog starting from the root visits x_1, a frog starting at x_1 visits x_2, and so on. More simply, a frog is woken within $\mathbb{T}_d(v_1)$ if it would have been woken even if there were no frogs sleeping on any vertices outside of $\mathbb{T}_d(v_1)$.

We define some random variables counting frogs that might possibly visit the root. Let $X \sim \mathrm{Poi}(\mu)$ be the number of frogs sleeping on $\varnothing'$, which are woken by the initial frog. Let X_1 be the number of frogs waking within $\mathbb{T}_d(v_1)$ that visit $\varnothing'$. We claim that X_1 is distributed as ν. Indeed, when we consider frogs as waking only if they wake within $\mathbb{T}_d(v_1)$ and relabel the vertices $\{\varnothing'\} \cup \mathbb{T}_d(v_1)$ as $\{\varnothing\} \cup \mathbb{T}_d(\varnothing')$, we see a process identical in law to the original non-backtracking frog model. Call the frogs counted by X and X_1 the *first-wave frogs*.

For each $2 \leq i \leq d$, let E_i be the event that some of the frogs counted by X or X_1 move to v_i. Conditional on E_i, arbitrarily choose one of these frogs that visits v_i and call it f. We say that the frogs at v are woken within $\mathbb{T}_d(v_i)$ if there exists a chain of vertices $x_1, \dots, x_m = v$ in $\mathbb{T}_d(v_i)$ such that f visits x_1, a frog starting at x_1 visits x_2, and so on. Let X_i be the number of frogs waking within $\mathbb{T}_d(v_i)$ that visit $\varnothing'$. By the same argument showing that $X_1 \sim \nu$, the distribution of X_i conditional on E_i is also ν. Furthermore, for any $\{i_1, \dots, i_k\} \subseteq \{2, \dots, d\}$, the random variables $X_{i_1}, \dots, X_{i_k}$ are conditionally independent given $E_{i_1}, \dots, E_{i_k}$, since each X_i is determined solely by the paths of the frogs sleeping in $\mathbb{T}_d(v_i)$. We call the frogs counted by $X_2, \dots, X_d$ the *second-wave frogs*.

The first- and second-wave frogs all visit $\varnothing'$. We define V'' as the number of these that move from there to $\varnothing$.

Claim. *$V'' \sim \mathcal{A}\nu$.*

Proof. Our strategy is to show that the first-wave frogs behave identically as the first-wave particles, and then to show that the second-wave frogs conditional on the behavior of first-wave frogs behave the same as the second-wave particles conditional on the behavior of the first-wave particles.

For the first of these claims, consider the first-wave frogs, counted by X and X_1. Observe that X and X_1 are independent with $X \sim \mathrm{Poi}(\mu)$ and $X_1 \sim \nu$, just as in the particle system defining $\mathcal{A}\nu$. The frogs counted by X move from $\varnothing'$ independently to a random choice out of $\varnothing, v_1, \ldots, v_d$, and the frogs counted by X_1 move from $\varnothing'$ independently to a random choice out of $\varnothing, v_2, \ldots, v_d$, also matching the particle system. Thus, the locations of the first-wave frogs one step after leaving $\varnothing'$ are distributed identically to the ending locations of the first-wave particles.

Now, condition on some arrangement of the first-wave frogs on $\varnothing, v_1, \ldots, v_d$ one step after leaving $\varnothing$. Suppose that u out of the vertices $v_2, \ldots, v_d$ are occupied by first-wave frogs in this arrangement. The number of second-wave frogs visiting $\varnothing'$ conditional on this arrangement of first-wave frogs is a sum of u independent copies of ν. Each second-wave frog that visits $\varnothing'$ has an independent $1/d$ chance of moving next to $\varnothing$. Thus, the number of second-wave frogs that visit $\varnothing$ is the sum of u independent copies of $\mathrm{Bin}(\nu, 1/d)$. This matches the conditional distribution of second-wave particles ending at $\varnothing$ given in Lemma 8. Thus, the distribution of the number of first- and second-wave frogs visiting $\varnothing$ is the same as the distribution of the number of first- and second-wave particles ending at $\varnothing$, which is by definition $\mathcal{A}\nu$. □

With this claim, the proof of the lemma is almost complete: Let V be the total number of visits to $\varnothing$ in the non-backtracking frog model. Since $V'' \leq V$ with $V'' \sim \mathcal{A}\nu$ and $V \sim \nu$, we have shown that $\mathcal{A}\nu \preceq \nu$. □

2.3. **Properties of $\mathcal{A}$.** We first show (3), monotonicity of $\mathcal{A}$ with respect to stochastic dominance.

Lemma 10. *If $\pi_1 \preceq \pi_2$, then $\mathcal{A}\pi_1 \preceq \mathcal{A}\pi_2$.*

Proof. If $\pi_1 \preceq \pi_2$, then we can couple the two particle systems defining $\mathcal{A}\pi_1$ and $\mathcal{A}\pi_2$ so that the second particle system contains all the same particles as the first, moving identically, as well as additional ones. Thus at least as many particles visit $\varnothing$ in the second system as in the first, and $\mathcal{A}\pi_1 \preceq \mathcal{A}\pi_2$. □

Now, we describe the result of applying $\mathcal{A}$ to a Poisson distribution, whose thinning property simplifies things.

Lemma 11. *The distribution $\mathcal{A}\,\mathrm{Poi}(\lambda)$ is a mixture of Poisson distributions, given by*

$$\mathcal{A}\,\mathrm{Poi}(\lambda) \sim \mathrm{Poi}\biggl(\frac{(U+1)\lambda}{d} + \frac{\mu}{d+1}\biggr), \tag{5}$$

where

$$U \sim \mathrm{Bin}\biggl(d-1,\ 1 - \exp\Bigl(-\frac{\lambda}{d} - \frac{\mu}{d+1}\Bigr)\biggr). \tag{6}$$

Proof. In the particle process defining $\mathcal{A}\,\mathrm{Poi}(\lambda)$, let $Y_{u\to v}$ be the number of particles that start at u and finish at v, for $u \in \{\varnothing', v_1, \ldots v_d\}$ and $v \in \{\varnothing, v_1, \ldots, v_d\}$. Each of the $\mathrm{Poi}(\mu)$ particles starting at $\varnothing'$ moves to a random neighbor. By Poisson thinning, then, the random variables $Y_{\varnothing'\to v}$ for $v \in \{\varnothing, v_1, \ldots, v_d\}$ are independent and distributed as $\mathrm{Poi}(\mu/(d+1))$. Similarly, $Y_{v_1\to v}$ for $v \in \{\varnothing, v_2, \ldots, v_d\}$ are independent and distributed as $\mathrm{Poi}(\lambda/d)$. These two collections of random variables are also independent of each other.

Thus, the number of first-wave particles that move to v_i for each $2 \leq i \leq d$ are independent and distributed as $\mathrm{Poi}\bigl(\lambda/d + \mu/(d+1)\bigr)$. Let U be the number of vertices out of $\{v_2, \ldots, v_d\}$ that are visited. As each vertex has an independent $1 - \exp\bigl(-\lambda/d - \mu/(d+1)\bigr)$ chance of being visited, the distribution of U is as given in (6). And since U is determined by $Y_{\varnothing'\to v_i}$ and $Y_{v_1\to v_i}$ for $i=2,\dots d$, it is independent of $Y_{\varnothing'\to\varnothing}$ and $Y_{v_1\to\varnothing}$.

By Lemma 8 and Poisson thinning, the number of second-wave particles ending at $\varnothing$ is $\mathrm{Poi}(U\lambda/d)$. The number of first-wave particles ending at $\varnothing$ is $Y_{\varnothing'\to\varnothing}+Y_{v_1\to\varnothing}$, independent of U and distributed as $\mathrm{Poi}\big(\lambda/d+\mu/(d+1)\big)$. Summing these together yields yields (5). ☐

We are nearly in a position to establish that $\mathcal{A}^n\,\mathrm{Poi}(0)$ grows without limit as $n\to\infty$. First we need two technical lemmas on the Poisson distribution.

Lemma 12. *Let $\overline{Z}_\lambda$ be distributed as* $\mathrm{Poi}(\lambda)$ *conditioned to be nonzero. If $\lambda_1\le\lambda_2$, then $\overline{Z}_{\lambda_1}\preceq\overline{Z}_{\lambda_2}$.*

Proof. Consider the Radon-Nikodym derivative of the law of $\overline{Z}_{\lambda_2}$ with respect to the law of $\overline{Z}_{\lambda_1}$,

$$r(k)=\frac{\mathbf{P}[\overline{Z}_{\lambda_2}=k]}{\mathbf{P}[\overline{Z}_{\lambda_1}=k]}=\frac{1-e^{-\lambda_1}}{1-e^{-\lambda_2}}e^{\lambda_1-\lambda_2}\left(\frac{\lambda_2}{\lambda_1}\right)^k.$$

The function $r(k)$ is increasing, and it is straightforward to show that this implies that $\overline{Z}_{\lambda_1}\preceq\overline{Z}_{\lambda_2}$ (or see [SS07, Theorem 1.C.1]). ☐

Lemma 13. *Let $\overline{Z}^{(1)}_{\lambda/n},\overline{Z}^{(2)}_{\lambda/n},\dots$ be independent and distributed as* $\mathrm{Poi}(\lambda/n)$ *conditioned to be nonzero. Let M be independent of these and be distributed as* $\mathrm{Bin}(n,1-e^{-\lambda/n})$*, and let*

$$Z=\sum_{i=1}^{M}\overline{Z}^{(i)}_{\lambda/n}.$$

Then Z is distributed as $\mathrm{Poi}(\lambda)$*.*

Proof. Decompose $\mathrm{Poi}(\lambda)$ as a sum of n independent copies of $\mathrm{Poi}(\lambda/n)$. Let M be the number of these that are nonzero, and condition on M to get the desired representation. ☐

Finally, we prove (4).

Proposition 14. *If $\mu>2(d+1)\log d$, then there exists $\epsilon>0$ such that*

$$\mathrm{Poi}(\lambda+\epsilon)\preceq\mathcal{A}\,\mathrm{Poi}(\lambda)$$

for all $\lambda\ge0$.

Proof. Let $X\sim\mathrm{Poi}(\lambda+\epsilon)$ for some $\epsilon>0$ to be chosen later, and let $Y\sim\mathcal{A}\,\mathrm{Poi}(\lambda)$. We start by decomposing X into a sum of Poissons conditioned to be nonzero. For any a, let $\overline{Z}^{(1)}_a,\overline{Z}^{(2)}_a,\dots$ be distributed as $\mathrm{Poi}(a)$ conditioned to be nonzero, and let $Z_a\sim\mathrm{Poi}(a)$ (with no conditioning). Take all these random variables to be independent. By Lemma 13, we can write X as

$$X=Z_{(\lambda+\epsilon)/d}+\sum_{i=1}^{M}\overline{Z}^{(i)}_{(\lambda+\epsilon)/d},\tag{7}$$

where

$$M\sim\mathrm{Bin}\biggl(d-1,\ 1-\exp\Bigl(-\frac{\lambda+\epsilon}{d}\Bigr)\biggr).$$

We now turn to Y, which by Lemma 11 is distributed as

$$\mathrm{Poi}\biggl(\frac{(U+1)\lambda}{d}+m\biggr),\tag{8}$$

where $m = \mu/(d+1)$ and $U \sim \mathrm{Bin}\bigl(d-1,\ 1-\exp(-\lambda/d-m)\bigr)$. Let $Y' \sim \mathrm{Poi}\bigl((U+1)(\lambda+m)/d\bigr)$. For each u, the distribution of Y' conditional on $U=u$ is stochastically dominated by the distribution of Y conditional on $U=u$, simply because $\mathrm{Poi}(a) \preceq \mathrm{Poi}(b)$ when $a \leq b$. It follows that $Y' \preceq Y$. Thus it suffices to show that $X \preceq Y'$. Decomposing Y' by Lemma 13 and using the same notation as before, we can write Y' as

$$Y' = Z_{(\lambda+m)/d} + \sum_{i=1}^{N} \overline{Z}^{(i)}_{(\lambda+m)/d} \tag{9}$$

with

$$N \sim \mathrm{Bin}\biggl(U,\ 1-\exp\Bigl(-\frac{\lambda+m}{d}\Bigr)\biggr).$$

These decompositions allow us to stochastically compare X and Y'. Assume that ϵ is chosen to be smaller than m. We claim that to show that $X \preceq Y'$, it suffices to show that $M \preceq N$. Indeed, we can then couple the random variables on the right-hand sides of (7) and (9) so that

(1) $M \leq N$;
(2) $Z_{(\lambda+\epsilon)/d} \leq Z_{(\lambda+m)/d}$;
(3) $\overline{Z}^{(i)}_{(\lambda+\epsilon)/d} \leq \overline{Z}^{(i)}_{(\lambda+m)/d}$ for each i.

Property (2) is possible because $\mathrm{Poi}(a) \preceq \mathrm{Poi}(b)$ if $a \leq b$, and (3) is possible by Lemma 12. Together, this yields a coupling of X and Y' with $X \leq Y'$

Thus, it only remains to show that $M \preceq N$. Recalling that U is itself binomial, we have

$$\begin{aligned} N &\sim \mathrm{Bin}\biggl(\mathrm{Bin}\Bigl(d-1,\, 1-\exp\Bigl(-\frac{\lambda}{d}-m\Bigr)\Bigr),\ 1-\exp\Bigl(-\frac{\lambda+m}{d}\Bigr)\biggr) \\ &= \mathrm{Bin}\biggl(d-1,\ \Bigl(1-\exp\Bigl(-\frac{\lambda}{d}-m\Bigr)\Bigr)\Bigl(1-\exp\Bigl(-\frac{\lambda+m}{d}\Bigr)\Bigr)\biggr). \end{aligned}$$

Since M and N are both binomial, proving $M \preceq N$ reduces to comparing their parameters. The argument will be complete once we show for some $\epsilon > 0$ and all $\lambda > 0$,

$$1-\exp\Bigl(-\frac{\lambda+\epsilon}{d}\Bigr) \leq \Bigl(1-\exp\Bigl(-\frac{\lambda}{d}-m\Bigr)\Bigr)\Bigl(1-\exp\Bigl(-\frac{\lambda+m}{d}\Bigr)\Bigr). \tag{10}$$

Some basic calculus (see Lemma 16 in the appendix) establishes that for all $d \geq 2$,

$$e^{-2\log d} + e^{-2\log d/d} < 1.$$

Since $m > 2\log d$, we can choose $\epsilon > 0$ such that

$$1 > \exp\Bigl(-\frac{\epsilon}{d}\Bigr) \geq e^{-m} + e^{-m/d}.$$

Multiplying both sides of this inequality by $e^{-\lambda/d}$ gives

$$\exp\Bigl(-\frac{\lambda+\epsilon}{d}\Bigr) \geq \exp\Bigl(-\frac{\lambda}{d}-m\Bigr) + \exp\Bigl(-\frac{\lambda+m}{d}\Bigr).$$

Thus

$$\begin{aligned} 1-\exp\Bigl(-\frac{\lambda+\epsilon}{d}\Bigr) &\leq 1-\exp\Bigl(-\frac{\lambda}{d}-m\Bigr)-\exp\Bigl(-\frac{\lambda+m}{d}\Bigr) \\ &\leq \Bigl(1-\exp\Bigl(-\frac{\lambda}{d}-m\Bigr)\Bigr)\Bigl(1-\exp\Bigl(-\frac{\lambda+m}{d}\Bigr)\Bigr). \end{aligned}$$

Looking back at (10), we have shown that $M \preceq N$. □

We have now proven (1)–(4), completing the proof of Proposition 6.

2.4. **Comparison to one-per-site results.** In [HJJ15], we proved that the frog model on a binary tree with one sleeping frog per site is recurrent. The proof has the same overarching idea as here: We use the self-similarity of the tree to obtain a recursive distributional relationship for the number of returns to the root. We then use this relationship in a bootstrapping argument, assuming that the number of visits to the root is stochastically larger than $\mathrm{Poi}(\lambda)$ and proving that it is in fact stochastically larger than $\mathrm{Poi}(\lambda+\epsilon)$.

The major difference between the two arguments is in the bootstrapping portion. The approach in this paper using traditional stochastic domination fails with the one-per-site frog model. The problem is that the distributions given by successively applying the analogue of the $\mathcal{A}$ operator in the one-per-site model have finite support and hence are never stochastically greater than any Poisson distribution. Our proof in [HJJ15] instead uses an exotic definition of stochastic dominance, where π_1 is dominated by π_2 if the probability generating function of π_1 is greater than the probability generating function of π_2.

This generating function approach works better than the technique in this paper in some ways and worse in others. On one hand, it can handle both deterministic and random initial configurations. When applied to the binary tree with $\mathrm{Poi}(\mu)$ sleeping frogs per site, it gives slightly better results than we obtain in this paper: in unpublished work, we have used it to prove that this model is recurrent for $\mu > 1.13$, better than the $\mu > 6\log 2 \approx 4.16$ given by Proposition 6. (By specializing the argument in Proposition 14 to $d = 2$ and refining it as much as possible, it is possible to improve this to $\mu > 3\log 2$, but we cannot do any better than that using stochastic dominance.)

On the other hand, the generating function approach seems confined to small values of d. It relies on a purely analytic argument that is elementary but difficult. It seems impossible to apply this argument to an arbitrary choice of d. Even for $d = 3$, the generating functions to be analyzed become extremely complicated. The technical advance in this paper is the probabilistic argument we give in Proposition 14, which allows us to work on any d-ary tree.

3. Transience

The main idea of our proof of transience is to consider a *weight function* on the frog model. To analyze the weight function, we bound the frog model by a branching random walk. The weight function is the frog model analogue to a common martingale derived from branching random walk (see [Big77]).

Proposition 15. *If* $\mathbf{E}\eta < \frac{(d-1)^2}{4d}$*, then the frog model with an independent copy of* η *frogs per site on* $\mathbb{T}_d$ *is almost surely transient.*

Proof. Let F_n be the set of frogs awake at time n. For $f \in F_n$, let $|f|$ denote the level of f on the tree (that is, its distance from the root). We define a weight function

$$W_n = \sum_{f \in F_n} e^{-\theta|f|},$$

with θ to be chosen shortly. Let

$$m = \frac{1}{d+1}e^{\theta} + \frac{d}{d+1}\mathbf{E}[\eta+1]e^{-\theta}.$$

Before we explain the meaning of this, we minimize m by setting $\theta = \log\bigl((\mathbf{E}\eta + 1)d\bigr)/2$, making

$$m = \frac{2\sqrt{(\mathbf{E}\eta + 1)d}}{d+1} < 1$$

under our assumption that $\mathbf{E}\eta < \frac{(d-1)^2}{4d}$.

The strategy of the proof now is to show that $W_n \to 0$, and hence that the root eventually stops being visited. The term m gives an upper bound for the expected contribution to W_{n+1} of a frog at time n in the following way: Suppose that at time n, some frog f is at level i of the tree for any $i \geq 1$. With probability $1/(d+1)$, the next jump of f is towards the root, waking no frogs. With probability $d/(d+1)$, the jump is away from the root, possibly waking up an η-distributed number of frogs. Thus, the expected contribution to W_{n+1} from f and any frogs it wakes at time $n+1$ is at most $e^{-\theta i}m$. If f is at the root at time n, then the expected contribution to W_{n+1} from f and the frogs it wakes is at most $\mathbf{E}[\eta + 1]e^{-\theta}$, which is bounded by m given our choice of θ. Therefore

$$\mathbf{E}[W_{n+1} \mid W_n] \leq \sum_{f \in F_n} e^{-\theta|f|}m = mW_n.$$

Thus W_n/m^n is a positive supermartingale. By the martingale convergence theorem, it converges almost surely to a finite limit. Since $m^n \to 0$, we also have $W_n \to 0$ a.s., which implies that eventually no frogs are present at the root. □

4. Appendix

Lemma 16. $x^{-2} + x^{-2/x} < 1$ *for all* $x \geq 2$.

Proof. Let $f(x) = x^{-2} + x^{-2/x}$. First we show the inequality holds on the interval $[2, 8]$. Since x^{-2} is decreasing,

$$f(x) \leq \frac{1}{4} + x^{-2/x}.$$

It is easily checked that the maximum of $x^{-2/x}$ on $[2, 8]$ occurs at $x = 8$ and is less than $\frac{3}{4}$.

Next, we consider $x \geq 8$. L'Hôpital's rule implies that $\lim_{x\to\infty} f(x) = 1$. Thus it suffices to confirm that $f(x)$ is increasing on $[8, \infty)$. We compute

$$f'(x) = 2x^{-(2/x)-2}\Bigl(\log x - x^{(2/x)-1} - 1\Bigr).$$

For $x \geq 8$, it holds that $x^{(2/x)-1} < 1$. Hence

$$f'(x) \geq 2x^{-(2/x)-2}\Bigl(\log x - 2\Bigr),$$

which is positive on $[8, \infty)$ since $\log 8 > 2$. □

References

[AHZ13] Elie Aïdékon, Yueyun Hu, and Olivier Zindy, *The precise tail behavior of the total progeny of a killed branching random walk*, Ann. Probab. **41** (2013), no. 6, 3786–3878. MR 3161464

[Ald91] David J. Aldous, *Random walk covering of some special trees*, J. Math. Anal. Appl. **157** (1991), no. 1, 271–283. MR 1109456 (93b:60147)

[AMP02a] O. S. M. Alves, F. P. Machado, and S. Yu. Popov, *The shape theorem for the frog model*, Ann. Appl. Probab. **12** (2002), no. 2, 533–546. MR 1910638 (2003c:60159)

[AMP02b] Oswaldo Alves, Fabio Machado, and Serguei Popov, *Phase transition for the frog model*, Electron. J. Probab. **7** (2002), no. 16, 1–21.

[AMPR01] 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 (2003f:60171)

[Big77] John D. Biggins, *Martingale convergence in the branching random walk*, J. Appl. Probability **14** (1977), no. 1, 25–37. MR 0433619 (55 #6592)

[CMMV14] Michael Cranston, Thomas Mountford, Jean-Christophe Mourrat, and Daniel Valesin, *The contact process on finite homogeneous trees revisited*, available at arXiv:1403.5927, 2014.

[CQR09] Francis Comets, Jeremy Quastel, and Alejandro F. Ramírez, *Fluctuations of the front in a one dimensional model of $X + Y \to 2X$*, Trans. Amer. Math. Soc. **361** (2009), no. 11, 6165–6189. MR 2529928 (2010i:60281)

[DL88] Richard Durrett and Xiu Fang Liu, *The contact process on a finite set*, Ann. Probab. **16** (1988), no. 3, 1158–1173. MR 942760 (89f:60119)

[DP14] Christian Döbler and Lorenz Pfeifroth, *Recurrence for the frog model with drift on $\mathbb{Z}^d$*, Electron. Commun. Probab. **19** (2014), no. 79, 13. MR 3283610

[DRS10] Ronald Dickman, Leonardo T. Rolla, and Vladas Sidoravicius, *Activated random walkers: facts, conjectures and challenges*, J. Stat. Phys. **138** (2010), no. 1-3, 126–142. MR 2594894 (2011h:82047)

[DS88] Richard Durrett and Roberto H. Schonmann, *The contact process on a finite set. II*, Ann. Probab. **16** (1988), no. 4, 1570–1583. MR 958203 (89m:60254)

[DST89] Richard Durrett, Roberto H. Schonmann, and Nelson I. Tanaka, *The contact process on a finite set. III. The critical case*, Ann. Probab. **17** (1989), no. 4, 1303–1321. MR 1048928 (91d:60251)

[GNR15] Arka P. Ghosh, Steven Noren, and Alexander Roitershtein, *On the range of the transient frog model on $\mathbb{Z}$*, available at arXiv:1502.02738, 2015.

[GS09] Nina Gantert and Philipp Schmidt, *Recurrence for the frog model with drift on $\mathbb{Z}$*, Markov Process. Related Fields **15** (2009), no. 1, 51–58. MR 2509423 (2010g:60170)

[HJJ15] Christopher Hoffman, Tobias Johnson, and Matthew Junge, *Recurrence and transience for the frog model on trees*, available at arXiv:1404.6238, 2015.

[Pem88] Robin Pemantle, *Phase transition in reinforced random walk and RWRE on trees*, Ann. Probab. **16** (1988), no. 3, 1229–1241. MR 942765 (89g:60220)

[Pem92] ———, *The contact process on trees*, Ann. Probab. **20** (1992), no. 4, 2089–2116. MR 1188054 (94d:60155)

[Pop01] Serguei Yu. Popov, *Frogs in random environment*, J. Statist. Phys. **102** (2001), no. 1-2, 191–201. MR 1819703 (2002a:82064)

[Pop03] ———, *Frogs and some other interacting random walks models*, Discrete random walks (Paris, 2003), Discrete Math. Theor. Comput. Sci. Proc., AC, Assoc. Discrete Math. Theor. Comput. Sci., Nancy, 2003, pp. 277–288 (electronic). MR 2042394

[PS01] Robin Pemantle and Alan M. Stacey, *The branching random walk and contact process on Galton-Watson and nonhomogeneous trees*, Ann. Probab. **29** (2001), no. 4, 1563–1590. MR 1880232 (2002m:60193)

[RS04] Alejandro F. Ramírez and Vladas Sidoravicius, *Asymptotic behavior of a stochastic combustion growth process*, J. Eur. Math. Soc. (JEMS) **6** (2004), no. 3, 293–334. MR 2060478 (2005e:60234)

[RS12] 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

[SS07] Moshe Shaked and J. George Shanthikumar, *Stochastic orders*, Springer Series in Statistics, Springer, New York, 2007. MR 2265633 (2008g:60005)

[ST14] Vladas Sidoravicius and Augusto Teixeira, *Absorbing-state transition for Stochastic Sandpiles and Activated Random Walks*, available at arXiv:1412.7098, 2014.

[TW99] András Telcs and Nicholas C. Wormald, *Branching and tree indexed random walks on fractals*, J. Appl. Probab. **36** (1999), no. 4, 999–1011. MR 1742145 (2001m:60199)

Department of Mathematics, University of Washington
E-mail address: hoffman@math.washington.edu

Department of Mathematics, University of Southern California
E-mail address: tobias.johnson@usc.edu

Department of Mathematics, University of Washington
E-mail address: jungem@math.washington.edu

STOCHASTIC ORDERS AND THE FROG MODEL

TOBIAS JOHNSON AND MATTHEW JUNGE

ABSTRACT. The *frog model* starts with one active particle at the root of a graph and $\eta(v)$ dormant particles at all nonroot vertices v. Active particles follow independent random paths, waking all inactive particles they encounter. We prove that certain frog model statistics are monotone in the initial configuration with respect to stochastic dominance in the increasing concave order. As a consequence, a frog model recurrent with i.i.d.-η particles per site is also recurrent with deterministically $\lceil \mathbf{E}\eta \rceil$ particles per site. We deduce recurrence for the infinite d-ary tree with simple random walk paths and $k > \mu_c(d)$ frogs at each site, with $\mu_c(d)$ the threshold for recurrence of a frog model with i.i.d. Poisson frogs per site. We end with a proof that $\mu_c(d)$ is of order d, removing a logarithmic term from the previous upper bound.

1. INTRODUCTION

Let G be a countable collection of vertices, one of which we distinguish as the root and call $\varnothing$. A general frog model (η, S) starts with one active particle at $\varnothing$, and $\eta(v)$ dormant particles at each $v \neq \varnothing$. The ith particle at v starting from its time of activation moves according to the path $S_\bullet(v,i)$, with $S_0(v,i)$ assumed equal to v. When an active particle visits a site containing dormant particles, *all* of the dormant particles activate. The particles move in discrete time, though this will be unimportant since most of the properties of the frog model we consider depend only on the particles' paths and not on the time they make their moves. The particles are traditionally called frogs, and we continue the zoomorphism. Typically, G is a graph, the frog paths $(S_\bullet(v,i))_{v\in G, i\geq 1}$ are independent random walks, the frog counts $(\eta_v)_{v\in G}$ are either deterministic or i.i.d., and $(S_\bullet(v,i))_{v\in G, i\geq 1}$ and $(\eta_v)_{v\in G}$ are independent of each other. We will not belabor an example like the frog model with simple random walk paths on $\mathbb{Z}^d$ and i.i.d.-Poi(μ) frogs per vertex by stating that the frog paths are mutually independent, and that the frog counts and paths are independent.

Our main result is about a class of frog model functionals we call *icv statistics*. The prime example is the number of visits to $\varnothing$ in the frog model (η, S) over all time, which we denote $r(\eta, S)$. A realization of the frog model is called *recurrent* if $r(\eta, S) = \infty$ and *transient* otherwise. In [TW99], the frog model with one sleeping frog per site and simple random walk paths is shown to be recurrent on $\mathbb{Z}^d$ for all d. This is further refined in [Pop01], which exhibits a threshold in α at which a frog model with Bernoulli$(\alpha\|x\|^{-2})$ frogs at each $x \in \mathbb{Z}^d$ switches from transience to recurrence. A similar phenomenon occurs when the walks have a bias in one direction: [GS09] finds that on $\mathbb{Z}$, the model is recurrent if and only if the number of sleeping frogs per site has infinite logarithmic moment. A sufficient condition for

Date: February 13, 2016.

2010 *Mathematics Subject Classification.* 60K35, 60J80, 60J10.

Key words and phrases. Frog model, transience, recurrence, phase transition, stochastic orders, increasing concave order.

The first author was partially supported by NSF grant DMS-1308645 and NSA grant H98230-13-1-0827 and the second author by NSF grant DMS-1401479.

recurrence in this setting on $\mathbb{Z}^d$ was given in [DP14] and improved on in [KZ15]. Our papers [HJJ15b] and [HJJ15a] study the frog model with simple random walk paths on d-ary trees. We prove that the frog model on a d-ary tree switches from transient to recurrent by either fixing d and increasing the density of frogs, or by fixing the density and decreasing d.

Statement of main theorem. Our main result is a comparison theorem relating certain statistics of the frog model, including the number of returns $r(\eta, S)$, when we vary the distribution of the initial configuration η. Our motivation is that while the most convenient setting has Poisson-distributed frog counts, the most basic questions assume a deterministic number of frogs per site. As an example, in [HJJ15a] we showed the existence of a recurrence phase on the d-ary tree with Poisson frogs per site for any $d \geq 2$. This left open the existence of a recurrence phase for initial conditions other than i.i.d. Poisson. For instance, for large enough k, is the frog model recurrent on the d-ary tree with k frogs per site? With our previous tools, we could answer this question only for the case $d = 2$ [HJJ15b], but our comparison theorem tidily transfers the result from Poisson to deterministic initial conditions (see Corollary 4).

If $\eta(v)$ is dominated by $\eta'(v)$ in the usual stochastic order, then we can couple the corresponding frog models and deduce that $f(\eta, S)$ is dominated by $f(\eta', S)$ for any statistic f that is increasing in η. This is not helpful for the problem described above, since we cannot relate a Poisson random variable to the constant k in this stochastic order. We instead turn to a weaker stochastic dominance relation known as the increasing concave order, '$\preceq_{\text{icv}}$' (defined in Section 2). Our main theorem applies to frog model functions we call *icv statistics*, defined in Definition 18. These include the count $r(\eta, S)$ of visits to the root, as well as the total number of activated sites. Our result is that such statistics are monotonic in the initial frog configuration with respect to this weaker stochastic order.

Theorem 1. *Assume that the frog paths $S_\bullet(v,i)$ and counts $\eta(v)$ and $\eta'(v)$ are mutually independent for all v and i, and that the paths $S_\bullet(v,i)$ at a particular vertex v are identically distributed for all i. Let f be an icv statistic of the frog model in the sense of Definition* 18. *If $\eta(v) \preceq_{icv} \eta'(v)$ for all v, then $f(\eta, S) \preceq_{icv} f(\eta', S)$.*

The intuition behind the proof is the following property of the increasing concave order: for an increasing concave function $f\colon \mathbb{R} \to \mathbb{R}$, if $X \preceq_{\text{icv}} Y$, then $f(X) \preceq_{\text{icv}} f(Y)$. While this fact follows immediately from the definition of the increasing concave order (see Section 2), the proof of Theorem 1 requires some argument.

Applications. As we mentioned, our main statistic of interest fits the criteria of Theorem 1.

Lemma 2. *The count $r(\eta, S)$ of visits to $\varnothing$ in the frog model (η, S) is an icv statistic of the frog model.*

This allows us to transfer many recurrence and transience results to different initial conditions. In the increasing concave order, the constant k dominates all mean k random variables. Theorem 1 and Lemma 2 therefore imply the following:

Corollary 3. *Consider the frog model on a graph with mutually independent frog paths and i.i.d. frogs per site with common mean μ. If this is almost surely recurrent, then for any integer $k \geq \mu$, the same frog model with k frogs per site is almost surely recurrent.*

This solves our problem of showing that the frog model on a d-ary tree with deterministically k frogs per site is recurrent for large enough k. In more detail, [HJJ15a, Theorem 1] establishes that on the d-ary tree with i.i.d.-Poi(μ) frogs per site, there is a critical value

$\mu_c(d)$ such that the frog model is recurrent a.s. if $\mu > \mu_c(d)$ and transient a.s. if $\mu < \mu_c(d)$. Corollary 3 thus gives us the desired result:

Corollary 4. *The frog model on $\mathbb{T}_d$ with $k > \mu_c(d)$ frogs per site is almost surely recurrent.*

In light of Corollary 4 it is especially relevant to know the correct order of $\mu_c(d)$. [HJJ15a, Theorem 1] establishes that for some constants C and C', we have $Cd \le \mu_c(d) \le C'd \log d$. In the following theorem, we remove the $\log d$ factor, establishing that $\mu_c(d)$ grows linearly in d. The proof is similar to that in [HJJ15a], but with several technical improvements and simplifications.

Theorem 5. *For all sufficiently large d, it holds that $.24d \le \mu_c(d) \le 3d$.*

Another application of Theorem 1 concerns the transience regime of the d-ary tree. In [HJJ15b, Theorem 1] we show that on $\mathbb{T}_d$ with one frog per site and simple random walk paths, the frog model is transient for $d \ge 5$. An immediate corollary of Theorem 1 is transience for all other mean 1 configurations.

Corollary 6. *For $d \ge 5$, the frog model on $\mathbb{T}_d$ with $\eta(v)$ frogs at each site and $\mathbf{E}\eta(v) \le 1$ for all $v \in \mathbb{T}_d$ is almost surely transient.*

Our next application is to the frog model on $\mathbb{Z}^d$. As mentioned earlier, [Pop01, Theorem 1.1] establishes the existence of a critical parameter $0 < \alpha_c(d) < \infty$ for the frog model with simple random walk paths on $\mathbb{Z}^d$ and initial configuration given by $\eta(x) \sim \mathrm{Ber}(p_x)$ such that

- (i) if $p_x \le \alpha/\|x\|^2$ for $\alpha < \alpha_c(d)$ and all sufficiently large x, then the model is transient with positive probability;
- (ii) if $p_x \ge \alpha/\|x\|^2$ for $\alpha > \alpha_c(d)$ and all sufficiently large x, then the model is transient with probability zero.

Theorem 1 allows us to extend part (i) of this result to non-Bernoulli distributions of sleeping frogs. Other results in [Pop01] like Theorem 1.3 can be similarly extended.

Corollary 7. *For all $\alpha < \alpha_c(d)$ and any $\big(\eta(x),\, x \in \mathbb{Z}^d \setminus \{0\}\big)$ satisfying $\mathbf{E}\eta(x) \le \alpha/\|x\|^2$ for sufficiently large x, the frog model on $\mathbb{Z}^d$ with simple random walk paths and initial configuration η is transient with positive probability.*

Another application is to the frog model with death, explored in [AMP02, FMS04, LMP05], where frogs have an independent chance $1 - p$ of dying at each step. This is a frog model according to our general definition, taking the frog paths to be stopped random walks. In this setting, the statistic of interest has been the total number of sites visited, which undergoes a phase transition on the regular tree from being finite a.s. to being infinite with positive probability as p grows. The model is said to *die out* in the first case and to *survive* in the second. The number of sites visited is an icv statistic, as we show in Lemma 20, and we therefore obtain the following result.

Corollary 8. *Let $\eta'(v) \succeq_{icv} \eta(v)$ be independent random variables indexed by the vertices v of an arbitrary graph G. If the frog model with death on G survives with $\eta(v)$ frogs at each v, then it survives with $\eta'(v)$ frogs at each v.*

Questions. We will give a few open problems on the theme of comparison theorems. We also list many open questions in [HJJ15b, HJJ15a].

Open Question 9. *Does the analogue of Theorem 1 hold in any weaker stochastic orders? For example, does it hold in the probability generating function order, described in Section 2?*

We used the probability generating function order in [HJJ15b] to prove the frog model recurrent on the binary tree with one sleeping frog per site. Many of the results in that paper rely on explicit calculations that do not generalize beyond the binary tree. A positive answer to the above question would provide more conceptual proofs that might help establish recurrence of the frog model on the 3-ary tree with one sleeping frog per site. For example, it would give an immediate proof of the analogue of Lemma 25 for the probability generating function order, which we proved in the binary tree case by a technical calculation that evades probabilistic interpretation [HJJ15b, Lemma 10].

The next question asks how sensitive the recurrence of the frog model is to the distribution of the frog counts. We believe that recurrence depends not just on the mean number of frogs at each vertex, but on the entire distribution.

Open Question 10. *Give an example where $r(\eta, S) = \infty$ a.s. and $r(\eta', S) < \infty$ a.s. with $\mathbf{E}\eta(v) = \mathbf{E}\eta'(v)$ for all v.*

Specifically, we would like to know that with simple random walk paths on the binary tree and i.i.d.-π frogs per vertex with mean 1, the frog model is transient when π is sufficiently unconcentrated.

Finally, we are interested in comparing frog models when the graph is altered rather than the initial configuration. As a concrete question in this vein, we ask if the d-regular tree is the most transient graph in the following sense:

Open Question 11. *Suppose the frog model is transient on a d-regular graph G with simple random walks. Is it necessarily transient on an infinite d-regular tree with simple random walk paths and the same initial conditions?*

Acknowledgments. We thank Chris Hoffman for his general assistance, and we thank Jonathan Hermon for a discussion in 2014 that mentioned Open Question 11.

2. Background material on stochastic orders

Let π_1 and π_2 be probability measures on the extended real numbers $\overline{\mathbb{R}} = (-\infty, \infty]$, and let $X \sim \pi_1$ and $Y \sim \pi_2$. The following three stochastic orders will play a role in this paper:

Standard stochastic order: $\pi_1 \preceq_{\text{st}} \pi_2$ if $\mathbf{E}f(X) \leq \mathbf{E}f(Y)$ for all increasing functions $f\colon \overline{\mathbb{R}} \to \overline{\mathbb{R}}$.

Increasing concave order: $\pi_1 \preceq_{\text{icv}} \pi_2$ if $\mathbf{E}f(X) \leq \mathbf{E}f(Y)$ for all increasing concave functions $f\colon \overline{\mathbb{R}} \to \overline{\mathbb{R}}$.

Probability generating function order: $\pi_1 \preceq_{\text{pgf}} \pi_2$ if $\mathbf{E}t^X \geq \mathbf{E}t^Y$ for all $t \in [0,1]$, with 0^0 interpreted as 1, and t^∞ interpreted as 0 for $t \in [0,1)$. We always assume that π_1 and π_2 are supported on $[0,\infty]$ when using this order.

We use $X \preceq_{\text{st}} Y$, $X \preceq_{\text{st}} \pi_2$, and $\pi_1 \preceq_{\text{st}} Y$ all to mean that $\pi_1 \preceq_{\text{st}} \pi_2$, and we do the corresponding thing with the other two orders. We abbreviate the increasing concave and the probability generating function orders as the icv and pgf orders, respectively.

We have listed these three stochastic orders in decreasing strength. That is,

$$\pi_1 \preceq_{\text{st}} \pi_2 \implies \pi_1 \preceq_{\text{icv}} \pi_2 \implies \pi_1 \preceq_{\text{pgf}} \pi_2. \tag{1}$$

The first implication is obvious. For the second, the map $x \mapsto 1 - t^x$ is an increasing concave function for any $t \in (0,1]$, establishing that $\mathbf{E}t^X \geq \mathbf{E}t^Y$ for $t \in (0,1]$ if $X \preceq_{\text{icv}} Y$. Taking a limit shows that it holds for $t = 0$ as well.

We have defined these orders to allow for random variables that take the value ∞. Though this is not typical, all of stochastic order theory can be extended to allow it in a routine way.

In keeping with the spirit of concavity, we require an increasing concave function f on $\overline{\mathbb{R}}$ to satisfy $f(\infty) = \lim_{x\to\infty} f(x)$. We note that if $X \preceq_{\text{pgf}} Y$, then $\mathbf{P}[X = \infty] \le \mathbf{P}[Y = \infty]$. To see this, note that as $t \nearrow 1$, we have $t^x \to \mathbf{1}\{x < \infty\}$. Thus, by the monotone convergence theorem,

$$\mathbf{E}t^X \to \mathbf{P}[X < \infty] \qquad \text{and} \qquad \mathbf{E}t^Y \to \mathbf{P}[Y < \infty]$$

as $t \nearrow 1$. Now $\mathbf{E}t^X \ge \mathbf{E}t^Y$ implies that $\mathbf{P}[X < \infty] \ge \mathbf{P}[Y < \infty]$. By (1), the conclusion also holds for X, Y taking nonnegative values under the assumption $X \preceq_{\text{st}} Y$ or $X \preceq_{\text{icv}} Y$.

Two useful equivalent conditions for $\pi_1 \preceq_{\text{st}} \pi_2$ are that $\mathbf{P}[X > t] \le \mathbf{P}[Y > t]$ for all t, and that X and Y can be coupled so that $X \le Y$ a.s. An equivalent condition for $X \preceq_{\text{icv}} Y$ is that X and Y can be coupled so that $\mathbf{E}[X \mid Y] \le Y$ a.s. [SS07, Theorem 4.A.5]. While this is less useful than the coupling characterization of the standard order, it gets across that $X \preceq_{\text{icv}} Y$ roughly means that X is smaller and less concentrated than Y. We are not aware of an equivalent definition of the pgf order in terms of couplings. One way to interpret the pgf order probabilistically if X and Y are integer-valued is that $X \preceq_{\text{pgf}} Y$ if the p-thinning of X is more likely than the p-thinning of Y to be zero, for any $p \in [0, 1]$. We never use the pgf order in this paper, but it played a prominent role in our proof of recurrence for the one-per-site frog model on the binary tree in [HJJ15b]. For a thorough reference on stochastic orders, see [SS07].

Theorem 13 provides a necessary and sufficient condition for a Poisson mixture to dominate a Poisson distribution, which we will need in Section 5. The result was first proven in [MSH03] for the standard stochastic order, but the proof given there also works for the icv and pgf orders. We reproduce it here for our readers' convenience. See also [Yu09] for a more general result.

Lemma 12 ([MSH03, Lemma 3.1(b)])**.** *For any positive integer n, the function*

$$h_n(x) = x \sum_{k=0}^{n} \frac{(-\log x)^k}{k!}$$

is increasing and concave on $(0, 1]$.

Proof. We compute

$$h_n'(x) = \sum_{k=0}^{n} \frac{(-\log x)^k}{k!} - \sum_{k=1}^{n} \frac{(-\log x)^{k-1}}{(k-1)!} = \frac{(-\log x)^n}{n!},$$

which is positive and decreasing on $(0, 1]$, showing that $h_n(x)$ is increasing and concave. $\square$

Theorem 13 ([MSH03, Theorem 3.1(b)])**.** *Let $X \sim$ Poi(λ), and let $Y \sim$ Poi(U) for some nonnegative random variable U. Then the following are equivalent:*

(i) $X \preceq_{st} Y$,
(ii) $X \preceq_{icv} Y$,
(iii) $X \preceq_{pgf} Y$,
(iv) $\mathbf{P}[X = 0] \ge \mathbf{P}[Y = 0]$, *and*
(v) $\lambda \le -\log \mathbf{E}e^{-U}$.

Proof. The implications (i) $\Longrightarrow$ (ii) $\Longrightarrow$ (iii) were explained previously. Conditions (iv) and (v) are just restatements of each other, since $\mathbf{P}[X = 0] = e^{-\lambda}$ and $\mathbf{P}[Y = 0] = \mathbf{E}e^{-U}$. Condition (iii) implies (iv) by the definition of the pgf order applied with $t = 0$.

It remains to prove that (v) implies (i). It suffices to show that $\mathbf{P}[Y \le n] \le \mathbf{P}[X \le n]$ for all nonnegative integers n. We compute

$$\mathbf{P}[Y \le n] = \sum_{k=0}^{n} \mathbf{E}\left[\frac{e^{-U}U^k}{k!}\right] = \mathbf{E}\left[\zeta \sum_{k=0}^{n} \frac{(-\log\zeta)^k}{k!}\right] = \mathbf{E}h_n(\zeta),$$

where $\zeta = e^{-U}$. We have assumed that $\mathbf{E}\zeta = \mathbf{E}e^{-U} \le e^{-\lambda}$. By Lemma 12, the function $h_n(x)$ is increasing and concave on $(0, 1]$, where ζ takes values. Thus

$$\mathbf{E}h_n(\zeta) \le h_n(\mathbf{E}\zeta) \le h_n\bigl(e^{-\lambda}\bigr) = \sum_{k=0}^{n} \frac{e^{-\lambda}\lambda^k}{k!} = \mathbf{P}[X \le n],$$

where we use that h_n is concave to apply Jensen's inequality in the first step, and we use that h_n is increasing in the second step. ☐

The following proposition shows that the maximal real- and integer-valued distributions in the icv order with a given expectation are the distributions that are as concentrated as possible.

Proposition 14.

(a) If $\mathbf{E}X \le c$, then $X \preceq_{icv} c$.

(b) Suppose X takes integer values and $\mathbf{E}X \in [k, k+1]$ for an integer k. Let Y be a random variable taking values k and $k+1$ satisfying $\mathbf{E}X \le \mathbf{E}Y$. Then $X \preceq_{icv} Y$.

Proof. Part (a) follows immediately from Jensen's inequality. For part (b), since $X \preceq_{\text{icv}} Y$ if and only if $X - k \preceq_{\text{icv}} Y - k$, we can assume without loss of generality that $k = 0$. Let φ be an arbitrary increasing concave function; by translating it, we can assume without loss of generality that $\varphi(0) = 0$. Under these assumptions, Y is Bernoulli with mean at least $\mathbf{E}X$, implying that

$$\mathbf{E}\varphi(Y) \ge (\mathbf{E}X)\varphi(1). \tag{2}$$

Define

$$a = \mathbf{E}[X \mid X \le 0], \qquad p = \mathbf{P}[X \le 0],$$
$$b = \mathbf{E}[X \mid X \ge 1], \qquad q = \mathbf{P}[X \ge 1].$$

If $p = 0$ or $q = 0$, then X is deterministic and the result is trivial. Thus we can assume that both conditional expectations above are well defined.

Applying Jensen's inequality,

$$\mathbf{E}\varphi(X) = p\mathbf{E}[\varphi(X) \mid X \le 0] + q\mathbf{E}[\varphi(X) \mid X \ge 1] \le p\varphi(a) + q\varphi(b). \tag{3}$$

As $a \le 0$ and $b \ge 1$, the points $(a, \varphi(a))$ and $(b, \varphi(b))$ lie under the secant line connecting $(0, 0)$ and $(1, \varphi(1))$. Thus $\varphi(a) \le a\varphi(1)$ and $\varphi(b) \le b\varphi(1)$. Applying to this to (3) and combining with (2) gives

$$\mathbf{E}\varphi(X) \le (pa + qb)\varphi(1) = (\mathbf{E}X)\varphi(1) \le \mathbf{E}\varphi(Y).$$ ☐

3. Proof of the comparison theorem

We start by proving Proposition 17, a general result useful for proving domination in the icv stochastic order. We then apply this to the frog model to prove Theorem 1.

3.1. **A method for proving icv domination.** Here is an example of the sort of statement that Proposition 17 is designed to prove:

Example 15. If $Z_1, Z_2, \dots$ are i.i.d. nonnegative random variables and $M \preceq_{\mathrm{icv}} N$, then

$$\sum_{i=1}^{M} Z_i \preceq_{\mathrm{icv}} \sum_{i=1}^{N} Z_i. \tag{4}$$

We give the proof immediately following Proposition 17. Example 15 is not a new result—see Theorem 4.A.9 in [SS07] for a proof. Proposition 17 is modeled on this theorem but is more general. First, we need a measure theory lemma whose technical proof can be safely ignored.

Lemma 16. *Let X and Y be real-valued random variables on a common probability space, and let $f(x, y)$ be a measureable function. Define for each $x \in \mathbb{R}$ the random variable $F(x) = \mathbf{E}[f(x, Y) \mid X]$. Then*

$$F(X) = \mathbf{E}[f(X, Y) \mid X] \text{ a.s.}$$

Proof. Let μ be a regular conditional distribution of Y given X. That is, $\mu(x, \cdot)$ is a probability measure on the Borel sets for all $x \in \mathbb{R}$, the map $x \mapsto \mu(x, B)$ is measurable for all Borel sets B, and

$$\mu(X, B) = \mathbf{P}[Y \in B \mid X] \text{ a.s.}$$

Such a random measure μ exists by [Kal02, Theorem 6.3]. By [Kal02, Theorem 6.4],

$$\mathbf{E}[f(X, Y) \mid X] = \int f(X, y)\mu(X, dy) \text{ a.s.,} \tag{5}$$

and

$$F(x) = \mathbf{E}[f(x, Y) \mid X] = \int f(x, y)\mu(X, dy) \text{ a.s.} \tag{6}$$

Substituting X in place of x on the right hand side of (6), we see that it matches the right hand side of (5), proving that $\mathbf{E}[f(X, Y) \mid X] = F(X)$ a.s. ☐

Now, we state the proposition. The rough idea is that for a sequence of random variables $X_0 \le X_1 \le \cdots$ with a tendency to increase more and more slowly, if $N \preceq_{\mathrm{icv}} N'$ then $X_N \preceq_{\mathrm{icv}} X_{N'}$.

If X and Y are random variables defined on the same probability space, then we say that $X \preceq_{\mathrm{icv}} Y$ conditional on a σ-algebra $\mathcal{F}$ if for any increasing concave function φ,

$$\mathbf{E}[\varphi(X) \mid \mathcal{F}] \le \mathbf{E}[\varphi(Y) \mid \mathcal{F}] \text{ a.s.}$$

We say that $X \preceq_{\mathrm{icv}} Y$ conditional on Z to mean that $X \preceq_{\mathrm{icv}} Y$ conditional on the σ-algebra generated by Z. If $X \preceq_{\mathrm{icv}} Y$ conditional on $\mathcal{F}$, then by taking expectations in the above inequality, we see that $X \preceq_{\mathrm{icv}} Y$ also holds unconditionally.

Proposition 17. *Let $X_0 \le X_1 \le \cdots$ be a sequence of random variables. Suppose that*

$$X_{n+2} - X_{n+1} \preceq_{icv} X_{n+1} - X_n \text{ conditional on } X_n$$

for all $n \ge 0$. Let M and N be independent of $(X_i)_{i \ge 0}$. If $M \preceq_{icv} N$, then $X_M \preceq_{icv} X_N$.

Proof. Let φ be an arbitrary increasing concave function, and define $g(n) = \mathbf{E}\varphi(X_n)$. We claim that $g(n)$ is increasing and concave. Once we prove this, it holds by our assumption $M \preceq_{\mathrm{icv}} N$ that $\mathbf{E}g(M) \leq \mathbf{E}g(N)$. By the independence of M and N from $(X_i)_{i\geq 0}$, we have $\mathbf{E}g(M) = \mathbf{E}\varphi(X_M)$ and $\mathbf{E}g(N) = \mathbf{E}\varphi(X_N)$, showing that $\mathbf{E}\varphi(X_M) \leq \mathbf{E}\varphi(X_N)$, and hence $X_M \preceq_{\mathrm{icv}} X_N$.

Thus, we just need to show that $g(n)$ is increasing and concave. That it is increasing follows from our assumption that X_n is increasing. For the concavity, we need to show that $g(n+1) - g(n)$ is decreasing. First, we note that the function $x \mapsto \varphi(c+x)$ is increasing and concave for any c. By our assumption that $X_{n+1} - X_n \succeq_{\mathrm{icv}} X_{n+2} - X_{n+1}$ conditional on X_n,

$$\mathbf{E}\bigl[\varphi(c + X_{n+1} - X_n) \mid X_n\bigr] \geq \mathbf{E}\bigl[\varphi(c + X_{n+2} - X_{n+1}) \mid X_n\bigr] \text{ a.s.} \tag{7}$$

for any $c \geq 0$. With an eye towards applying Lemma 16, let $f(x, y) = \varphi(x + y)$, and let

$$F_1(x) = \mathbf{E}[f(x, X_{n+1} - X_n) \mid X_n] \qquad \text{and} \qquad F_2(x) = \mathbf{E}[f(x, X_{n+2} - X_{n+1}) \mid X_n],$$

so that $F_1(c)$ and $F_2(c)$ are the left and right hand sides of (7), respectively. By Lemma 16,

$$F_1(X_n) = \mathbf{E}[f(X_n, X_{n+1} - X_n) \mid X_n] = \mathbf{E}\bigl[\varphi(X_{n+1}) \mid X_n\bigr] \text{ a.s.}$$

and

$$F_2(X_n) = \mathbf{E}[f(X_n, X_{n+2} - X_{n+1}) \mid X_n] = \mathbf{E}\bigl[\varphi(X_n + X_{n+2} - X_{n+1}) \mid X_n\bigr] \text{ a.s.}$$

Since $F_1(X_n) \geq F_2(X_n)$ a.s. by (7), we have

$$\mathbf{E}\bigl[\varphi(X_{n+1}) \mid X_n\bigr] \geq \mathbf{E}\bigl[\varphi(X_n + X_{n+2} - X_{n+1}) \mid X_n\bigr] \text{ a.s.},$$

and subtracting $\varphi(X_n)$ from both sides gives

$$\begin{aligned}\mathbf{E}\bigl[\varphi(X_{n+1}) - \varphi(X_n) \mid X_n\bigr] &\geq \mathbf{E}\bigl[\varphi(X_n + X_{n+2} - X_{n+1}) - \varphi(X_n) \mid X_n\bigr] \\ &\geq \mathbf{E}\bigl[\varphi(X_{n+2}) - \varphi(X_{n+1}) \mid X_n\bigr] \text{ a.s.},\end{aligned}$$

with the last line following because $\varphi(x + X_{n+2} - X_{n+1}) - \varphi(x)$ is decreasing in x by the concavity of φ, and $X_{n+1} \geq X_n$. Taking expectations, we have

$$\begin{aligned}g(n+1) - g(n) &= \mathbf{E}\bigl[\varphi(X_{n+1}) - \varphi(X_n)\bigr] \\ &\geq \mathbf{E}\bigl[\varphi(X_{n+2}) - \varphi(X_{n+1})\bigr] = g(n+2) - g(n+1).\end{aligned}$$

Thus $g(n+1) - g(n)$ is decreasing, completing the proof. □

Proof of Example 15. We define $X_n = \sum_{i=1}^{n} Z_i$. Since Z_i is nonnegative for all i, we have $X_0 \leq X_1 \leq \cdots$. By the independence of $(Z_i)_{i\geq 1}$, the conditional stochastic dominance condition reduces to $Z_{n+2} \preceq_{\mathrm{icv}} Z_{n+1}$ unconditionally, which is trivially true since the two random variables are identically distributed. Thus (4) follows from Proposition 17. □

3.2. **Applying Proposition 17 to the frog model.** We now define *icv statistics* of the frog model, the class of statistics covered by Theorem 1. Roughly speaking, we call a functional of the frog model an icv statistic if it increases when a frog is added to the model, but when two frogs are added at the same vertex it increases less than by the separate addition of each of them. The definition also includes a mild continuity assumption. Many counts in the frog model naturally satisfy these assumptions.

Definition 18. Let $\{\eta(v), S_\bullet(v,i);\, v \in G, i \geq 1\}$ be a deterministic collection of frog counts and paths. For any path $P_\bullet$, let $\sigma_{P_\bullet}(\eta, S)$ denote a new frog model with an extra frog of path $P_\bullet$ added at P_0; that is, $\sigma_{P_\bullet}(\eta, S) = (\eta', S')$, where η' is identical to η except that $\eta'(P_0) = \eta(P_0) + 1$, and S' is identical to S except that $S'_\bullet(P_0, \eta(P_0)+1) = P_\bullet$. For any frog model statistic $f(\eta, S)$, define

$$\Delta_{P_\bullet} f(\eta, S) = f(\sigma_{P_\bullet}(\eta, S)) - f(\eta, S),$$

the change in f when a frog with path $P_\bullet$ is added to the model. Also define

$$\Delta_{P_\bullet, P'_\bullet} f(\eta, S) = f(\sigma_{P_\bullet}\sigma_{P'_\bullet}(\eta, S)) - f(\eta, S),$$

the change in f when frogs with paths $P_\bullet$ and $P'_\bullet$ are added.

A function f taking values in the nonnegative extended real numbers is an *icv statistic* of the frog model if for any (η, S) and any two paths $P_\bullet, P'_\bullet$ starting at the same vertex,

(a) $\Delta_{P_\bullet} f(\eta, S) \geq 0$;
(b) $\Delta_{P_\bullet, P'_\bullet} f(\eta, S) \leq \Delta_{P_\bullet} f(\eta, S) + \Delta_{P'_\bullet} f(\eta, S)$;
(c) if $\eta_k(v)$ converges upwards to $\eta(v)$ as $k \to \infty$ for all $v \in G$, then $f(\eta_k, S)$ converges upwards to $f(\eta, S)$.

Two equivalent formulations of part (b) of this definition are

$$\Delta_{P_\bullet}\Delta_{P'_\bullet} f(\eta, S) \leq 0, \tag{8}$$

and

$$\Delta_{P_\bullet} f\big(\sigma_{P'_\bullet}(\eta, S)\big) \leq \Delta_{P_\bullet} f(\eta, S). \tag{9}$$

Part (a) and formulation (8) of part (b) resemble the conditions for a real-valued smooth function to be increasing and concave. We caution the reader, however, that part (b) is required to hold only for two frogs originating at the same vertex, not for two general frogs.

In the next lemma, we show that an icv statistic is monotone in the number of frogs at a *single* vertex with respect to the icv order.

Lemma 19. *Assume all conditions of Theorem 1, and also assume that η and η' have identical distributions at all but one vertex. Then $f(\eta, S) \preceq_{icv} f(\eta', S)$.*

Proof. Let v_0 be the vertex where η and η' differ. Define η_k to be the same as η except that $\eta_k(v_0) = k$. Let $W_k = f(\eta_k, S)$. By our assumptions, $\eta(v_0)$ and $\eta'(v_0)$ are independent of W_k, and hence

$$W_{\eta(v_0)} \sim f(\eta, S) \qquad \text{and} \qquad W_{\eta'(v_0)} \sim f(\eta', S).$$

Our goal is to apply Proposition 17 with $M = \eta(v_0)$ and $N = \eta'(v_0)$ to show that $W_{\eta(v_0)} \preceq_{\text{icv}} W_{\eta'(v_0)}$.

The random variables $(W_i)_{i \geq 0}$ are naturally coupled together, since $S_\bullet(v_0, i)$ for $i \geq 1$ are all defined on a common probability space. The condition that $W_k \leq W_{k+1}$ for all $k \geq 0$ holds by part (a) of Definition 18, as $W_{k+1} - W_k = \Delta_{S_\bullet(v_0, k+1)} f(\eta_k, S)$. To apply Proposition 17, we need to confirm that $W_{k+2} - W_{k+1} \preceq_{\text{icv}} W_{k+1} - W_k$ conditional on W_k for all $k \geq 0$. Let $\mathcal{F}_k$ be the σ-algebra generated by

$$\big\{S_\bullet(v_0, i)\big\}_{1 \leq i \leq k} \quad \text{and} \quad \big\{S_\bullet(v, i),\, \eta(v)\big\}_{v \neq v_0,\, i \geq 1}$$

which represents all the information about the frog model (η_k, S). As W_k is measurable with respect to $\mathcal{F}_k$, it suffices to prove that $W_{k+2} - W_{k+1} \preceq_{\text{icv}} W_{k+1} - W_k$ conditional on $\mathcal{F}_k$.

Let $P_\bullet = S_\bullet(v_0, k+1)$ and $P'_\bullet = S_\bullet(v_0, k+2)$, and with (9) in mind observe that

$$W_{k+1} - W_k = \Delta_{P_\bullet} f(\eta_k, S), \tag{10}$$

and

$$W_{k+2} - W_{k+1} = \Delta_{P'_\bullet} f\bigl(\sigma_{P_\bullet}(\eta_k, S)\bigr). \tag{11}$$

Conditional on $\mathcal{F}_k$, the paths $P_\bullet$ and $P'_\bullet$ are i.i.d. Hence, the conditional distribution of $\Delta_{P'_\bullet} f\bigl(\sigma_{P_\bullet}(\eta_k, S)\bigr)$ on $\mathcal{F}_k$ is unaffected by swapping $P_\bullet$ and $P'_\bullet$. By (11), this shows that conditional on $\mathcal{F}_k$, the random variable $\Delta_{P_\bullet} f\bigl(\sigma_{P'_\bullet}(\eta_k, S)\bigr)$ is distributed identically to $W_{k+2} - W_{k+1}$. Thus, for an arbitrary increasing function ψ,

$$\begin{aligned}\mathbf{E}\bigl[\psi(W_{k+2} - W_{k+1}) \mid \mathcal{F}_k\bigr] &= \mathbf{E}\bigl[\psi\bigl(\Delta_{P_\bullet} f\bigl(\sigma_{P'_\bullet}(\eta_k, S)\bigr)\bigr) \mid \mathcal{F}_k\bigr] \\ &\leq \mathbf{E}\bigl[\psi\bigl(\Delta_{P_\bullet} f(\eta_k, S)\bigr) \mid \mathcal{F}_k\bigr] = \mathbf{E}\bigl[\psi(W_{k+1} - W_k) \mid \mathcal{F}_k\bigr],\end{aligned}$$

with (9) applied in the second step. This proves that $W_{k+2} - W_{k+1} \preceq_{\text{icv}} W_{k+1} - W_k$ conditional on $\mathcal{F}_k$ (in fact, it proves that this holds in the standard stochastic order as well, as we never assumed ψ concave). Proposition 17 now shows that $W_{\eta(v_0)} \preceq_{\text{icv}} W_{\eta'(v_0)}$. ☐

Proof of Theorem 1. By applying Lemma 19 repeatedly, the result holds if η and η' have the same distribution at all but finitely many vertices. To justify the general case, we use the continuity assumption, part (c) of Definition 18, to make a limiting argument. Let $G_1 \subseteq G_2 \subseteq \cdots$ be finite sets of vertices whose union is G. We use $\eta|_{G_k}$ and $\eta'|_{G_k}$ to denote restrictions to G_i; that is, $\eta|_{G_k}(v) = \eta(v)\mathbf{1}\{v \in G_k\}$.

Since $\eta|_{G_k}$ and $\eta'|_{G_k}$ differ at only finitely many vertices, Lemma 19 implies that

$$f(\eta|_{G_k}, S) \preceq_{\text{icv}} f(\eta'|_{G_k}, S). \tag{12}$$

Let φ be an arbitrary increasing concave function on $[0, \infty]$. In proving that $\mathbf{E}\varphi(f(\eta, S)) \leq \mathbf{E}\varphi(f(\eta', S))$, we can assume without loss of generality that φ takes nonnegative values by replacing it with $\varphi(x) - \varphi(0)$. By part (c) of Definition 18,

$$\varphi\bigl(f(\eta|_{G_k}, S)\bigr) \nearrow \varphi\bigl(f(\eta, S)\bigr) \text{ a.s.} \qquad \text{and} \qquad \varphi\bigl(f(\eta'|_{G_k}, S)\bigr) \nearrow \varphi\bigl(f(\eta', S)\bigr) \text{ a.s.}$$

as $k \to \infty$, and this holds even if random variables take the value ∞ with positive probability (recall that φ is assumed continuous at infinity). By the monotone convergence theorem,

$$\mathbf{E}\varphi\bigl(f(\eta|_{G_k}, S)\bigr) \to \mathbf{E}\varphi\bigl(f(\eta, S)\bigr) \qquad \text{and} \qquad \mathbf{E}\varphi\bigl(f(\eta'|_{G_k}, S)\bigr) \to \mathbf{E}\varphi\bigl(f(\eta', S)\bigr)$$

as $k \to \infty$. Since $\mathbf{E}\varphi\bigl(f(\eta|_{G_k}, S)\bigr) \leq \mathbf{E}\varphi\bigl(f(\eta'|_{G_k}, S)\bigr)$ by (12), it holds that $\mathbf{E}\varphi\bigl(f(\eta, S)\bigr) \leq \mathbf{E}\varphi\bigl(f(\eta', S)\bigr)$. ☐

4. Applications of the comparison theorem

To apply Theorem 1 to the frog model, we just need to establish that the frog model functionals we are interested in are icv statistics.

Proof of Lemma 2. Part (a) of Definition 18, that r is increasing under the addition of extra frogs, is obvious. Part (b) follows from a subadditivity property of the frog model: the sites activated when frogs $S_\bullet(v, 1), \ldots, S_\bullet(v, k)$ are initially at vertex v is the union of sites activated in the k frog models that are identical to the original one, but have a single frog $S_\bullet(v, i)$ at v, for $i = 1, \ldots, k$. We explain this in more detail now. Fix η and S and let $P_\bullet$ and $P'_\bullet$ be arbitrary paths starting from some vertex v_0. For any $v \in G$ and $i \geq 1$, let $V_{v,i}$ denote the number of visits to $\varnothing$ from time 1 on by the path $S_\bullet(v, i)$. Also let W and W' denote the number of visits to $\varnothing$ by $P_\bullet$ and $P'_\bullet$, respectively. Define the following indicators for each $v \in G$:

- I_v is an indicator on vertex v being visited by the frog model $\sigma_{P_\bullet}(\eta, S)$ but not by (η, S).
- I'_v is an indicator on v being visited by the frog model $\sigma_{P'_\bullet}(\eta, S)$ but not by (η, S).
- J_v is an indicator on v being visited by the frog model $\sigma_{P_\bullet}\sigma_{P'_\bullet}(\eta, S)$ but not by (η, S).

If v_0 is never visited in (η, S), then

$$\Delta_{P_\bullet} r(\eta, S) = \Delta_{P'_\bullet} r(\eta, S) = \Delta_{P_\bullet, P'_\bullet} r(\eta, S) = 0,$$

and part (b) of Definition 18 holds trivially. Otherwise, we have

$$\begin{aligned}\Delta_{P_\bullet} r(\eta, S) &= \sum_{v,i} I_v V_{v,i} + W,\\ \Delta_{P'_\bullet} r(\eta, S) &= \sum_{v,i} I'_v V_{v,i} + W',\\ \Delta_{P_\bullet, P'_\bullet} r(\eta, S) &= \sum_{v,i} J_v V_{v,i} + W + W'.\end{aligned}$$

To prove that (b) holds, it suffices to show that $J_v \leq I_v + I'_v$ for $v \neq v_0$. Equivalently, we must show that if v is activated by $\sigma_{P_\bullet}\sigma_{P'_\bullet}(\eta, S)$ but not by (η, S), then it is activated either by $\sigma_{P_\bullet}(\eta, S)$ or $\sigma_{P'_\bullet}(\eta, S)$.

To show that this subadditivity property holds, suppose v is activated by $\sigma_{P_\bullet}\sigma_{P'_\bullet}(\eta, S)$ but not by (η, S). Then there is a sequence of frogs that wake each other, starting with the initial frog and ending with a frog that visits v, and this sequence must include either $P_\bullet$ or $P'_\bullet$. If it includes both, all portions of the sequence starting at the first and ending immediately before the second can be eliminated, yielding a sequence of awakenings demonstrating that v is activated by $\sigma_{P_\bullet}(\eta, S)$ or $\sigma_{P'_\bullet}(\eta, S)$. This establishes that part (b) holds.

Finally, we must show the continuity property (c). This holds because any frog woken in (η, S) relies only on a finite sequence of frogs to wake it. More formally, suppose that the components of η_k converge upwards to η as $k \to \infty$. Let K_v and $K_v^{(k)}$ be indicators on v being activated by (η, S) and by (η_k, S), respectively. Then

$$r(\eta_k, S) = \sum_{v,i} K_v^{(k)} V_{v,i} \qquad \text{and} \qquad r(\eta, S) = \sum_{v,i} K_v V_{v,i}.$$

If $K_v = 1$, then eventually η_k contains enough frogs to make $K_v^{(k)} = 1$, by the property of a site being activated by a finite sequence of frogs. Thus $K_v^{(k)} \nearrow K_v$ as $k \to \infty$. By monotone convergence, $r(\eta_k, S) \nearrow r(\eta, S)$. □

Next, we prove that the number of sites visited by the frog model is an icv statistic. This is of interest for frog models with stopped paths, where it is possible for only finitely many frogs to be visited.

Lemma 20. *The number of sites ever visited by the frog model (η, S), which we call $a(\eta, S)$, is an icv statistic.*

Proof. This was nearly shown in the course of proving Lemma 2. Part (a) of Definition 18 is obvious. For part (b), let I_v, I'_v, and J_v be as in the proof of Lemma 2. Then

$$\Delta_{P_\bullet} a(\eta, S) = \sum_{v,i} I_v,$$
$$\Delta_{P'_\bullet} a(\eta, S) = \sum_{v,i} I'_v,$$
$$\Delta_{P_\bullet, P'_\bullet} a(\eta, S) = \sum_{v,i} J_v.$$

We showed in the proof of Lemma 2 that $J_v \leq I_v + I'_v$, which establishes part (b). Part (c) holds by the same argument used in the previous proof. □

Proofs of Corollaries 3, 4, and 6. We apply Theorem 1, Lemma 2, and Proposition 14(a), along with the observation made in Section 2 that $\mathbf{P}[X = \infty] \leq \mathbf{P}[Y = \infty]$ if $X \preceq_{\mathrm{icv}} Y$. □

Proof of Corollary 7. This is proven the same as Corollaries 3, 4, and 6, except that Proposition 14(b) is used instead of Proposition 14(a). □

Proof of Corollary 8. This also has the same proof as Corollaries 3, 4, and 6, except that Lemma 20 replaces Lemma 2. □

5. Critical parameters for d-ary trees

In [HJJ15b, HJJ15a], we prove recurrence for the frog model on a d-ary tree with simple random walk paths under different initial conditions. The techniques in the two papers are not identical, but both are based on recursion and bootstrapping. To set this up, we show that it is enough to establish recurrence for a frog model whose paths are stopped non-backtracking walks, which we call the *self-similar frog model.* Let V be the number of visits to the root in this process. A self-similarity yields a relation between V and a collection of independent copies of V. Such relations are referred to as recursive distributional equations (see [AB05] for further discussion).

In the bootstrap part of the argument, we assume that V is stochastically larger than $\mathrm{Poi}(\lambda)$ for some $\lambda \geq 0$. We then analyze the recursive distributional equation to show that V is in fact stochastically larger than $\mathrm{Poi}(\lambda + \epsilon)$. Iterating this argument starting at $\lambda = 0$, we show that V is larger than $\mathrm{Poi}(\epsilon)$, then larger than $\mathrm{Poi}(2\epsilon)$, and so on, with the conclusion that $V = \infty$ a.s. In [HJJ15a] and here, this argument uses the standard stochastic order, while in [HJJ15b] it uses the pgf order.

The result in [HJJ15a] was recurrence on a d-ary tree with i.i.d.-$\mathrm{Poi}(\mu)$ frogs per vertex for $\mu = \Omega(d \log d)$. Here we improve this by eliminating the $\log d$ factor. With the lower bound from [HJJ15a], this proves that the critical parameter $\mu_c(d)$ is of order d. Section 5.1 defines the self-similar frog model and gives the recursive distributional equation that the law of V satisfies (see Lemma 25). The ideas in this section can be found in [HJJ15b, HJJ15a], but they take some work to extract in the form we need. Though we do our best to avoid duplicating material, when in doubt we have opted for comprehensibility over efficiency.

Section 5.2 uses the set-up of Section 5.1 to make a bootstrapping argument establishing recurrence for $\mu = \Omega(d)$. This argument is a more elaborate version of the one used in [HJJ15a]. When writing that paper, we were unaware of the criterion for stochastic dominance of a Poisson distribution by a Poisson mixture given in Theorem 13, and as a result our argument was more complex than necessary. With this theorem at our disposal, we can prove the sharper result of Theorem 5.

5.1. **The bootstrapping set-up.**

5.1.1. *The self-similar frog model.* The frog model depends only on the range of each frog. This yields rather nice abelian and monotonicity properties. For example, the total number of visits to the root is unaffected by the order frogs wake up in and the rate they reveal vertices in their ranges. Also, trimming the range of frogs can only reduce the number of visits to the root. Applying this observation in combination with the coupling characterization of stochastic dominance, we note the following fact. Recall that $r(\eta, S)$ is the number of visits to the root in the frog model (η, S).

Fact 21. *Consider a collection of frog paths $S = \big(S_\bullet(v,i)\big)_{v\in G, i\geq 1}$ on a graph G. Suppose that another collection of paths $\tilde{S}$ can be coupled with S such that for all i and v, the range of $\tilde{S}_\bullet(v,i)$ is a subset of the range of $S_\bullet(v,i)$. Then $r(\eta,\tilde{S}) \preceq_{st} r(\eta, S)$.*

From now on, let $S = (S_\bullet(v,i),\ v \in \mathbb{T}_d, i \geq 1)$ denote a collection of independent simple random walks with $S_\bullet(v,i)$ started at v, and let the components of $\eta = (\eta(v))_{v\in\mathbb{T}_d}$ be i.i.d.-Poi(μ), independent of S. The first step in studying the frog model (η, S) will be to replace S by a collection of paths T to obtain (η, T), which we call the *self-similar frog model* in reference to a useful property described in Fact 23.

We define T in two steps. First, let $S' = (S'_\bullet(v,i),\ v \in \mathbb{T}_d, i \geq 1)$ denote a collection of independent random non-backtracking walks stopped at $\varnothing$. In more detail, call a random walk a *simple random non-backtracking walk* on an arbitrary graph if it chooses from its neighbors uniformly for its first step, and then in all subsequent steps it chooses uniformly from its current neighbors except the one it just arrived from. We define $S'_\bullet(v,i)$ to be a simple non-backtracking random walk stopped on arrival at $\varnothing$. The walks $S'_\bullet(v,i)$ and $S_\bullet(v,i)$ can be coupled so that the range of the first is a subset of the range of the second by making $S'_\bullet(v,i)$ a stopped, loop-erased version of $S_\bullet(v,i)$. This is proved in detail in [HJJ15b, Proposition 7].

Now we construct T as a modification of S'. Each path $T_\bullet(v,i)$ will be a stopped version of $S'_\bullet(v,i)$. Let v be a nonroot vertex in $\mathbb{T}_d$ with parent u. Suppose that v is visited in the frog model (S',η) for the first time at time j, necessarily by one or more frogs moving from u to v. Select one of these visiting frogs arbitrarily, and stop all of the other ones. (Observe that it is irrelevant which frog is allowed to continue, so long as one views frogs as indistinguishable.) If any frogs move from u to v at subsequent times, stop them at v as well. Do this for all vertices $v \in \mathbb{T}_d$, and let T be the resulting collection of stopped walks. As the range of each $T_\bullet(v,i)$ is a subset of the range of $S'_\bullet(v,i)$, the following fact (also noted in [HJJ15b, Proposition 7]) follows:

Fact 22. *There is a coupling of S and T so that the range of each $T_\bullet(v,i)$ is a subset of the range of $S_\bullet(v,i)$.*

By Facts 21 and 22, we have $r(\eta,T) \preceq_{\mathrm{st}} r(\eta,S)$. We will now work exclusively with the self-similar frog model, (η,T), and prove recurrence for it with sufficiently large μ. Unlike all other frog models considered in this paper, the frog paths T are not independent of each other nor of η, because one frog's motion in (η,T) can cause another frog to be stopped. This is the only form of dependence, however, and frogs that have not been stopped move independently of each other. So, it is not a serious obstacle.

Let $V = r(\eta,T)$. Next, we discuss a self-similarity property of the model and its consequences for V. For any vertex $v \in \mathbb{T}_d$, let $\mathbb{T}_d(v)$ denote the subtree made up of v and its descendants. We call $\mathbb{T}_d(v)$ *activated* in the self-similar frog model if v is ever visited. Let u

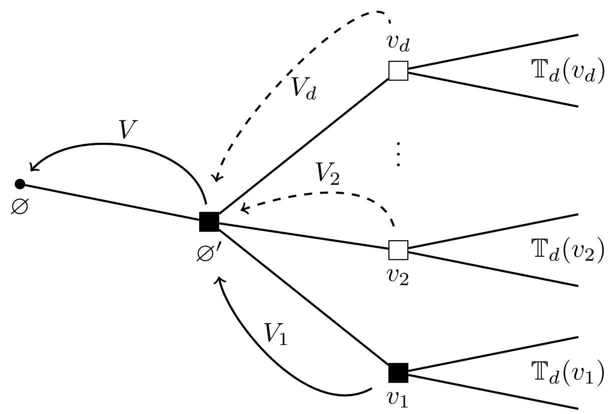

FIGURE 1. In the self-similar frog model on $\mathbb{T}_d$, the initial frog moves from $\varnothing$ to $\varnothing'$ to v_1 and then continues down the tree. The random variables V and V_1, counting the number of frogs moving from $\varnothing'$ to $\varnothing$ and from v_1 to $\varnothing'$, respectively, both have distribution ν (see Lemma 25). For $i \in \{2, \dots, d\}$, the distribution of V_i conditional on a frog entering the subtree $\mathbb{T}_d(v_i)$ is also ν.

be the parent of v. By our construction of T, if $\mathbb{T}_d(v)$ is activated, then there is a unique frog that moves from u to v, entering $\mathbb{T}_d(v)$ and then never leaving it. The frog model viewed starting from the time of activation only at vertices $\{u\} \cup \mathbb{T}_d(v)$ then looks identical to the original self-similar frog model viewed on $\{\varnothing\} \cup \mathbb{T}_d(\varnothing')$. This yields the following fact, proved in more detail in [HJJ15b, Proposition 6].

Fact 23. *Let V' be the number of frogs that move from v to its parent u in the self-similar frog model. The distribution of V' conditional on $\mathbb{T}_d(v)$ being activated is identical to the distribution of V.*

The following observation shows that once a subtree $\mathbb{T}_d(v)$ is activated, the random variable V' defined in the above fact is independent of the frog model outside of $\mathbb{T}_d(v)$.

Fact 24. *Let V' be defined as in Fact 23. Conditional on $\mathbb{T}_d(v)$ being activated, V' depends only on the path of the activator and on $\{T_\bullet(w,i), \eta(w)\colon w \in \mathbb{T}_d(v), i \geq 1\}$.*

We will use Facts 23 and 24 to express V recursively in terms of independent copies of itself, an idea expressed in Figure 1. This relation will be given in terms of an operator we define next.

5.1.2. *The operators $\mathcal{B}$ and $\mathcal{U}$.* Suppose that the initial frog in the self-similar frog model moves from $\varnothing$ to $\varnothing'$ to v_1. Let $v_2, \dots, v_d$ be the remaining children of $\varnothing'$. Observe that since frogs are stopped at $\varnothing$, no children of $\varnothing$ other than $\varnothing'$ are ever visited. The idea of this section is to view the self-similar frog model only at the vertices mentioned above. If a vertex v_i is visited, we close our eyes to $\mathbb{T}_d(v_i)$, thinking of this entire subtree as a black box that eventually emits some frogs from v_i back to $\varnothing'$.

Enacting this view, we now define operators $\mathcal{B}$ and $\mathcal{U}$ on probability measures supported on the extended nonnegative integers. Informally, the operator $\mathcal{B}$ corresponds to the number of visits to the root, and $\mathcal{U}$ corresponds to the number of subtrees $v_1, \ldots, v_d$ that are activated. Let π be a probability measure on the nonnegative integers. To define $\mathcal{B}\pi$ and $\mathcal{U}\pi$, we consider the following frog model. The example to keep in mind is when π is the law of V, in which case the following description matches up with the black box view of the self-similar frog model described above.

Graph: a star graph with center ρ' and leaf vertices $\rho, u_1, \ldots, u_d$ (think of these as paralleling $\varnothing'$ and $\varnothing, v_1, \ldots, v_d$). The root of the graph is ρ.

Sleeping frog counts: all independent, distributed as $\mathrm{Poi}(\mu)$ at ρ' and as π at $u_1, \ldots, u_d$. There is one frog at ρ, as is always true at the root vertex.

Paths: All frogs have independent paths. The initial frog moves deterministically from ρ to ρ' to u_1 and then remains there. All other frogs, if woken, perform simple random non-backtracking walks from their starting points, stopped on arrival at a leaf vertex.

We then define two quantities:

– $\mathcal{B}\pi$ is the distribution of the number of frogs that terminate at ρ.
– $\mathcal{U}\pi$ is the distribution of the final number of $u_1, \ldots, u_d$ that are visited by a frog.

Note that our definition of the initial frog path as deterministic is just for convenience. By symmetry, we would arrive at the same measures $\mathcal{B}\pi$ and $\mathcal{U}\pi$ if it were also defined as a stopped simple random non-backtracking walk.

We mention that $\mathcal{B}$ is closely related to the operators $\mathcal{A}$ defined in [HJJ15b] and [HJJ15a], but differs from both of them. The operator $\mathcal{A}$ in [HJJ15b] is the same as $\mathcal{B}$ in the $d = 2$ case if the initial distribution at ρ' in the definition of $\mathcal{B}$ is changed from $\mathrm{Poi}(\mu)$ to δ_1 (except that $\mathcal{A}$ acts on probability generating functions rather than distributions). The operator $\mathcal{A}$ in [HJJ15a] would be the same as $\mathcal{B}$ if in the frog model defining $\mathcal{B}$, frogs initially at $v_2, \ldots, v_d$ do not wake other frogs.

Now we relate this system back to the frog model.

Lemma 25. *Let ν be the law of $V = r(\eta, T)$, the number of visits to the root in the self-similar frog model on $\mathbb{T}_d$. It holds that $\mathcal{B}\nu = \nu$.*

Proof. Essentially, the frog model on the star graph exactly matches the black box view of the self-similar frog model described at the beginning of Section 5.1.2, and the result then follows from Facts 23 and 24. To make this more formal, we couple the two frog models. We take full advantage of the abelian properties of the frog model by viewing the frogs' motions in a convenient order.

Consider the frog model used to define $\mathcal{B}\nu$ as well as the self-similar frog model. We can couple the initial number of frogs on ρ' to be the same as on $\varnothing'$, and we can couple the first (and only) step of each frog at ρ' with the first step of the corresponding frog at $\varnothing'$.

Let V_i be the number of frogs that ever move from v_i to $\varnothing'$ in the self-similar model, and let U_i be the number of frogs initially at u_i in the star graph model. Noting that $\mathbb{T}_d(v_1)$ is activated by the initial frog, $V_1 \sim \nu$ by Fact 23. By Fact 24, V_1 is independent of all that we have coupled so far (that is, the number and first steps of frogs initially at $\varnothing'$). The random variable U_1 is also independent of all we have coupled so far and is distributed identically to V_1. We can therefore couple U_1 and V_1 to be equal. Next, we couple the second (and final) step of each frog at u_1 with the step of the corresponding frog counted by V_1 after it moves from v_1 to $\varnothing'$.

Let $\mathcal{V}_1$ consist of the indices $i \in \{2, \ldots, d\}$ such that u_i has been visited so far. By the construction of our coupling, we can also describe $\mathcal{V}_1$ as the set of $i \in \{2, \ldots, d\}$ such that

$\mathbb{T}_d(v_i)$ has been activated so far. Furthermore, identically many frogs have returned so far to ρ as to $\varnothing$. By Facts 23 and 24, conditional on the information so far, the random variables $(V_i,\, i \in \mathcal{V}_1)$ are i.i.d.-ν and are independent of the information so far, as are the random variables $(U_1,\, i \in \mathcal{V}_1)$. We can therefore couple these two random vectors to be equal. We then couple the paths of the frogs at these vertices to match up as we did with the frogs at u_1 and v_1.

As above, a vertex u_i is visited for the first time in this second round if and only if $\mathbb{T}_d(v_i)$ is visited for the first time in this second round. Let $\mathcal{V}_2$ be the set of such i. We can repeat the coupling argument of the previous paragraph, maintaining identical numbers of frogs terminating at ρ as at $\varnothing$, until we get an empty $\mathcal{V}_j$ and have counted all returns to ρ and $\varnothing$. Thus, under this coupling, the number of frogs terminating at ρ in the star graph model is the same as the number of frogs terminating at $\varnothing$ in the self-similar model. The first of these counts has distribution $\mathcal{B}\nu$, while the second has distribution ν, showing that the two are equal. □

The next lemma is similar to [HJJ15b, Lemma 10] and [HJJ15a, Lemma 10]. Theorem 1 shows that its analogue holds in the icv order, though we will not make use of this.

Lemma 26. *If $\pi_1 \preceq_{st} \pi_2$, then $\mathcal{B}\pi_1 \preceq_{st} \mathcal{B}\pi_2$.*

Proof. This immediately follows from the coupling definition of stochastic dominance. We couple the frog models defining $\mathcal{B}\pi_1$ and $\mathcal{B}\pi_2$ so that the frogs in the former are a subset of the frogs of the latter model, resulting in more visits to the root. □

Just as in [HJJ15a, Lemma 11], the operator $\mathcal{B}$ applied to a Poisson distribution yields a mixture of Poisson distributions. This is a consequence of the following property, known as *Poisson thinning*: Consider a multinomial distribution with $\mathrm{Poi}(\lambda)$ trials and n-types, each having probability p_k. Then the vector of outcomes is distributed like an independent collection of $\mathrm{Poi}(\lambda p_k)$ random variables.

Lemma 27. *Let U be a random variable distributed as $\mathcal{U}\,\mathrm{Poi}(\lambda)$.*

$$\mathcal{B}\,\mathrm{Poi}(\lambda) = \mathrm{Poi}\left(\frac{\mu}{d+1} + U\frac{\lambda}{d}\right). \tag{13}$$

Proof. In the frog model defining $\mathcal{B}\,\mathrm{Poi}(\lambda)$, the number of frogs at ρ' that move back to ρ is distributed as $\mathrm{Bin}(\mathrm{Poi}(\mu), 1/(d+1))$. By Poisson thinning, this is $\mathrm{Poi}(\mu/(d+1))$. Each visited u_i releases $\mathrm{Poi}(\lambda)$ sleeping frogs. These will take a non-backtracking step back to ρ', then with probability $1/d$ will move to ρ. Thus, each activated u_i sends $\mathrm{Poi}(\lambda/d)$ frogs to ρ. It follows that

$$\mathcal{B}\,\mathrm{Poi}(\lambda) \sim \mathrm{Poi}(\mu/(d+1)) + \textstyle\sum_1^d \mathbf{1}\{u_i \text{ visited}\}\,\mathrm{Poi}(\lambda/d).$$

The above sum is equal to $\sum_1^U \mathrm{Poi}(\lambda/d)$. By Poisson thinning, the $\mathrm{Poi}(\lambda/d)$ terms are independent of U. Applying additivity of Poisson random variables then brings us to the claimed formula. □

5.2. **Carrying out the bootstrap.** First, note that the lower bound in Theorem 5 follows from [HJJ15a, Proposition 15], which is proven by coupling the frog model with a transient branching random walk. Our contribution here is the upper bound.

The idea of the bootstrapping argument is to use Lemma 27 to demonstrate that for some $\delta > 0$, it holds for all $\lambda \geq 0$ that $\mathcal{B}\,\mathrm{Poi}(\lambda) \succeq_{\mathrm{st}} \mathrm{Poi}(\lambda + \delta)$. Lemmas 25 and 26 then combine to show that V is stochastically larger than any Poisson distribution, and hence $V = \infty$ a.s.

Recall that $\mathcal{U}\,\mathrm{Poi}(\lambda)$ is the distribution of the number of vertices $u_1, \ldots, u_d$ visited in the frog model on the star graph defined in Section 5.1.2. Compared to our proof of recurrence for $\mu = \Omega(d \log d)$ in [HJJ15a], the difference is that we give a better lower bound on $\mathcal{U}\,\mathrm{Poi}(\lambda)$. We define a lower bounding random variable $U' \in \{1, \ldots, d\}$ as follows. Consider the frog model used to define $\mathcal{B}\,\mathrm{Poi}(\lambda)$ and $\mathcal{U}\,\mathrm{Poi}(\lambda)$, and observe how many of $u_1, \ldots, u_d$ are visited by the $\mathrm{Poi}(\mu)$ frogs starting at ρ'. If at least $\lceil d/c \rceil$ of these vertices are visited for a yet to be determined constant c, then arbitrarily choose $\lceil d/c \rceil$ of them and allow the frogs activated there the chance to visit the remaining $d - \lceil d/c \rceil$ vertices. If fewer than $\lceil d/c \rceil$ vertices are visited by the frogs at ρ', then we recall that u_1 is guaranteed to be activated by the initial frog, and we just use the frogs at u_1 to try to activate the remaining vertices $u_2, \ldots, u_d$. We define U' as the number of vertices out of $u_1, \ldots, u_d$ activated in the end in this scheme. This is summarized as follows:

Let U_1' be the number of vertices $u_1, \ldots, u_d$ visited by the frogs initially at ρ'.

Case 1: $U_1' \geq \lceil d/c \rceil$

Arbitrarily choose $\lceil d/c \rceil$ of the vertices counted by U_1' and denote them by $\mathcal{V} \subseteq \{u_1, \ldots, u_d\}$. Let U' be the sum of $\lceil d/c \rceil$ and the number of the remaining $d - \lceil d/c \rceil$ subtrees visited by frogs starting in $\mathcal{V}$.

Case 2: $U_1' < \lceil d/c \rceil$

Let U' equal one plus the number of number of vertices $u_2, \ldots, u_d$ visited by frogs returning from u_1.

As U' counts only a subset of the full collection of activated vertices, we have $U' \preceq_{\mathrm{st}} \mathcal{U}\,\mathrm{Poi}(\lambda)$.

Now, we sketch the proof of Theorem 5. Throughout, we will assume that $\mu = C(d+1)$ with C a yet to be determined positive constant. In Lemma 28, we prove that Case 2 occurs with exponentially small probability as d grows. Next, in Lemma 29 we give a very explicit definition of a random variable U'' satisfying $U'' \preceq_{\mathrm{st}} U' \preceq_{\mathrm{st}} \mathcal{U}\,\mathrm{Poi}(\lambda)$. In Lemma 30, we use this lower bound together with Lemma 27 to prove that if $V \succeq_{\mathrm{st}} \mathrm{Poi}(\lambda)$, then $V \succeq_{\mathrm{st}} \mathrm{Poi}(\lambda+\delta)$ for some $\delta > 0$. An iterative argument then implies that $V = \infty$ a.s.

Lemma 28. *Recall that U_1' is the number of vertices $u_1, \ldots, v_d$ visited by the $\mathrm{Poi}(\mu)$ frogs initially at ρ' in the frog model defining $\mathcal{B}\,\mathrm{Poi}(\lambda)$ and $\mathcal{U}\,\mathrm{Poi}(\lambda)$. We have*

$$\mathbf{P}[U_1' < \lceil d/c \rceil] \leq e^{-bd} := p, \tag{14}$$

where $b = 2(1 - e^{-C} - \frac{1}{c})^2$.

Proof. It is a consequence of Poisson thinning that out of the $\mathrm{Poi}(\mu)$ frogs starting at ρ', independently $\mathrm{Poi}(\frac{\mu}{d+1}) = \mathrm{Poi}(C)$ move to each leaf $u_1, \ldots, u_d$. Thus each vertex has an independent $1 - e^{-C}$ chance of having a frog enter it from the ones starting at ρ', showing that $U_1' \sim \mathrm{Bin}(d, 1 - e^{-C})$.

Hoeffding's inequality tailored to a binomial distribution states that $\mathbf{P}[\mathrm{Bin}(n, p) \leq (p - \epsilon)n] \leq \exp(-2\epsilon^2 n)$ (this follows from [Hoe63, eq. (2.3)]). If we apply the inequality to U_1' with $\epsilon = (1 - e^{-C}) - \frac{1}{c}$, we establish (14). ∎

Lemma 29. *Let*

$$U'' \sim \begin{cases} \lceil d/c \rceil + \mathrm{Bin}\Big(d - \lceil d/c \rceil,\, 1 - e^{-\lambda/c}\Big) & \text{with probability } 1 - q, \\ 1 + \mathrm{Bin}\Big(d - 1,\, 1 - e^{-\lambda/d}\Big) & \text{with probability } q, \end{cases} \tag{15}$$

where $q = \mathbf{P}\big[U_1' < \lceil d/c \rceil\big]$. Then $U'' \preceq_{st} U'$.

Proof. Writing $U' \mid E$ to mean U' conditioned on the event E, we claim that

$$U' \mid \{U_1' \geq \lceil d/c \rceil\} \succeq_{\mathrm{st}} \lceil d/c \rceil + \mathrm{Bin}\Big(d - \lceil d/c \rceil,\, 1 - e^{-\lambda/c}\Big), \tag{16}$$

and

$$U' \mid \{U_1' < \lceil d/c \rceil\} \succeq_{\mathrm{st}} 1 + \mathrm{Bin}\Big(d - 1,\, 1 - e^{-\lambda/d}\Big). \tag{17}$$

The lemma then follows because conditional stochastic dominance implies stochastic dominance [SS07, Theorem 1.A.3, (d)].

Thus it just remains to confirm (16) and (17). Suppose $U_1' \geq \lceil d/c \rceil$. Then we are in Case 1, and $U' = \lceil d/c \rceil + U_2'$, where U_2' is the number of vertices in $\{u_1, \ldots, u_d\} \setminus \mathcal{V}$ visited by frogs returning from $\mathcal{V}$. Conditional on $\mathcal{V}$, the counts of frogs proceeding from $\mathcal{V}$ to each of $\{u_1, \ldots, u_d\} \setminus \mathcal{V}$ form a collection of independent $\mathrm{Poi}(\lambda \lceil d/c \rceil / d)$ random variables. Thus each vertex in $\{u_1, \ldots, u_d\} \setminus \mathcal{V}$ has an independent probability of $1 - e^{-\lambda \lceil d/c \rceil / d} \geq 1 - e^{-\lambda/c}$ of being visited by one of these frogs, showing that $U_2' \succeq_{\mathrm{st}} \mathrm{Bin}(d - \lceil d/c \rceil, 1 - e^{-\lambda/c})$ and confirming (16).

Next, suppose that $U_1' < \lceil d/c \rceil$, and Case 2 is in effect. In this case, $U' = 1 + U_2'$, where U_2' is the number of vertices $u_2, \ldots, u_d$ visited by frogs returning from u_1. By the same reasoning as in the previous case, $U_2' \succeq_{\mathrm{st}} \mathrm{Bin}(d-1, 1 - e^{-\lambda/d})$, confirming (17). □

Lemma 30. *Define*

$$h_{C,c} = h_{C,c}(\lambda, d) := \log\Big[\big(e^{-\frac{\lambda}{c} + \frac{\lambda}{d}} + 1 - e^{-\frac{\lambda}{c}}\big)^{d - \lceil d/c \rceil} + p\big(2 - e^{-\frac{\lambda}{d}}\big)^{d-1}\Big],$$

where p is the value defined in (14)*, which depends on C and c. We have*

$$\mathcal{B}\,\mathrm{Poi}(\lambda) \succeq_{st} \mathrm{Poi}\bigg(\lambda + \frac{\mu}{d+1} - h_{C,c}\bigg).$$

Proof. Combining (13) and $U'' \preceq_{\mathrm{st}} \mathcal{U}\,\mathrm{Poi}(\lambda)$, it follows from [SS07, Theorem 1.A.3, (d)] that

$$\mathcal{B}\,\mathrm{Poi}(\lambda) \succeq_{\mathrm{st}} \mathrm{Poi}\bigg(\frac{\mu}{d+1} + U'' \frac{\lambda}{d}\bigg). \tag{18}$$

In light of Theorem 13, we would like to compute $-\log \mathbf{E} e^{-\frac{\lambda}{d} U''}$. Recalling the definition of U'' in (15), we use the fact that $\mathbf{E} x^{\mathrm{Bin}(n,p)} = (1 - p + px)^n$ to compute

$$\begin{aligned}\mathbf{E} e^{-\frac{\lambda}{d} U''} = (1-q) e^{-\frac{\lambda}{d}\lceil d/c \rceil} \Big(e^{-\frac{\lambda}{c}} + \big(1 - e^{-\frac{\lambda}{c}}\big) e^{-\frac{\lambda}{d}}\Big)^{d - \lceil d/c \rceil} \\ + q e^{-\frac{\lambda}{d}} \Big(e^{-\frac{\lambda}{d}} + \big(1 - e^{-\frac{\lambda}{d}}\big) e^{-\frac{\lambda}{d}}\Big)^{d-1}.\end{aligned}$$

Using the bound $q \leq p$ from Lemma 28 and the trivial bound $1 - q \leq 1$ in the first step, and factoring out $e^{-\lambda}$ in the second step,

$$\begin{aligned}\mathbf{E} e^{-\frac{\lambda}{d} U''} &\leq e^{-\frac{\lambda}{d}\lceil d/c \rceil} \Big(e^{-\frac{\lambda}{c}} + \big(1 - e^{-\frac{\lambda}{c}}\big) e^{-\frac{\lambda}{d}}\Big)^{d - \lceil d/c \rceil} \\ &\qquad + p e^{-\frac{\lambda}{d}} \Big(e^{-\frac{\lambda}{d}} + \big(1 - e^{-\frac{\lambda}{d}}\big) e^{-\frac{\lambda}{d}}\Big)^{d-1} \\ &= e^{-\lambda}\Big[\big(e^{-\frac{\lambda}{c} + \frac{\lambda}{d}} + 1 - e^{-\frac{\lambda}{c}}\big)^{d - \lceil d/c \rceil} + p\big(2 - e^{-\frac{\lambda}{d}}\big)^{d-1}\Big].\end{aligned}$$

Thus,

$$-\log \mathbf{E} e^{-\frac{\lambda}{d} U''} = \lambda - h_{C,c}.$$

Using the above calculation and Theorem 13, we deduce that

$$\mathrm{Poi}\left(\frac{\mu}{d+1}+U''\frac{\lambda}{d}\right)\succeq_{\mathrm{st}}\mathrm{Poi}\left(\lambda+\frac{\mu}{d+1}-h_{C,c}\right).$$

Together with (18), this completes the proof. □

Proof of Theorem 5. As we noted before, the lower bound is a consequence of [HJJ15a, Proposition 15], and we just need to establish the upper bound by showing that the frog model on $\mathbb{T}_d$ is almost surely recurrent with i.i.d.-Poi($3d$) frogs per vertex for sufficiently large d. To apply our bootstrapping argument, we seek to show that for some $\delta>0$, it holds for all $\lambda\geq 0$ that $\mathcal{B}\,\mathrm{Poi}(\lambda)\succeq_{\mathrm{st}}\mathrm{Poi}(\lambda+\delta)$. Considering the result of Lemma 30, we need to choose C and c such that $\mu/(d+1)-h_{C,c}(\lambda,d)>\delta$ for all $\lambda\geq 0$ and sufficiently large d. Recalling that $\mu=C(d+1)$, rearranging terms, and exponentiating both sides of the inequality, this is equivalent to showing that for some C, c, δ, and d_0 it holds that

$$\exp(h_{C,c}(\lambda,d))<e^{C-\delta}, \tag{19}$$

on the set $\{(\lambda,d)\colon \lambda\geq 0, d\geq d_0\}$.

Towards proving this, we start with the inequality

$$\begin{aligned}\exp(h_{C,c}(\lambda,d))&=\left(e^{-\frac{\lambda}{c}+\frac{\lambda}{d}}+1-e^{-\frac{\lambda}{c}}\right)^{d-\lceil d/c\rceil}+p\left(2-e^{-\frac{\lambda}{d}}\right)^{d-1}\\&\leq\left(1+e^{-\frac{\lambda}{c}}\left(e^{\frac{\lambda}{d}}-1\right)\right)^{d(1-\frac{1}{c})}+e^{-bd}2^{d-1}\end{aligned} \tag{20}$$

obtained by applying the bounds $2-e^{-\lambda/d}\leq 2$ and $d-\lceil d/c\rceil\leq d(1-1/c)$ and substituting the value of p from (14). Note that b depends on C and c. Now we bound each of the two terms on the right hand side of (20) for the right choice of C, c, and d_0.

Some calculus shows that for any d and c satisfying $d>c$, the first term is maximized in λ when $e^{\lambda/d}=d/(d-c)$. This demonstrates that if $d>c$, then

$$\begin{aligned}\left(1+e^{-\frac{\lambda}{c}}\left(e^{\frac{\lambda}{d}}-1\right)\right)^{d(1-\frac{1}{c})}&\leq\left(1+\left(\frac{d-c}{d}\right)^{d/c}\left(\frac{d}{d-c}-1\right)\right)^{d(1-\frac{1}{c})}\\&=\left(1+\left(\frac{d-c}{d}\right)^{d/c}\frac{c}{d-c}\right)^{d(1-\frac{1}{c})}\\&\leq\left(1+\frac{c}{d-c}\right)^{d(1-\frac{1}{c})}\\&\leq\exp\left(\frac{d(c-1)}{d-c}\right).\end{aligned}$$

Choosing $c=3$, we obtain for any choice $d_0>c$ the following bound holding for all $d\geq d_0$ and $\lambda\geq 0$:

$$\left(1+e^{-\frac{\lambda}{c}}\left(e^{\frac{\lambda}{d}}-1\right)\right)^{d(1-\frac{1}{c})}\leq\exp\left(\frac{c-1}{1-3/d_0}\right). \tag{21}$$

The second term to be bounded, $e^{-bd}2^{d-1}$, vanishes as $d\to\infty$ when $b>\log 2$. Referring back to (14) and doing some algebra, we see that $b>\log 2$ when

$$C>-\log\left(1-\frac{1}{c}-\sqrt{\frac{\log 2}{2}}\right).$$

This inequality is satisified with $c = 3$ and $C = 2.56$. By this and (21), for any $\epsilon > 0$, we can take d_0 sufficiently large that for all $d \geq d_0$,

$$e^{-bd}2^{d-1} \leq \epsilon$$

and

$$\left(1 + e^{-\frac{\lambda}{c}}\left(e^{\frac{\lambda}{d}} - 1\right)\right)^{d(1-\frac{1}{c})} \leq e^{2+\epsilon}.$$

Applying these bounds to (20), we have

$$\exp(h_{C,c}(\lambda, d)) \leq e^{2+\epsilon} + \epsilon < e^{C-\delta}$$

for some choice of $\delta > 0$, confirming (19).

Thus, we have shown that for $d \geq d_0$, if $\mu \geq 2.56(d+1)$ then for all $\lambda \geq 0$,

$$\mathcal{B}\,\mathrm{Poi}(\lambda) \succeq_{\mathrm{st}} \mathrm{Poi}(\lambda + \delta). \tag{22}$$

For the sake of simplicity, we can revise our assumption to $\mu \geq 3d$ for $d \geq d_0$. We are finally ready to bootstrap our way to the conclusion that $r(\eta, T)$, the number of visits to the root in the self-similar frog model, is almost surely infinite given this assumption. Recall that ν is the law of $r(\eta, T)$. As $\nu \succeq_{\mathrm{st}} \mathrm{Poi}(0)$, Lemma 26 shows that $\mathcal{B}\nu \succeq_{\mathrm{st}} \mathcal{B}\,\mathrm{Poi}(0)$, and so $\mathcal{B}\nu \succeq_{\mathrm{st}} \mathrm{Poi}(\epsilon)$ by (22). But ν is a fixed point of $\mathcal{B}$ by Lemma 25, implying that $\nu \succeq_{\mathrm{st}} \mathrm{Poi}(\epsilon)$. Repeating this argument of successively applying Lemma 26, (22), and Lemma 25, we show that $\nu \succeq_{\mathrm{st}} \mathrm{Poi}(2\epsilon)$, and so on. Thus ν is stochastically larger than all Poisson distributions, which implies $\nu = \delta_\infty$. Finally, Facts 21 and 22 imply that $r(\eta, S) \succeq_{\mathrm{st}} \nu$, and we can thus conclude that $r(\eta, S) = \infty$ a.s. when $d \geq d_0$ and $\mu \geq 3d$. □

References

[AB05] David J. Aldous and Antar Bandyopadhyay, *A survey of max-type recursive distributional equations*, Ann. Appl. Probab. **15** (2005), no. 2, 1047–1110. MR 2134098 (2007e:60010)

[AMP02] Oswaldo Alves, Fabio Machado, and Serguei Popov, *Phase transition for the frog model*, Electron. J. Probab. **7** (2002), no. 16, 1–21.

[DP14] Christian Döbler and Lorenz Pfeifroth, *Recurrence for the frog model with drift on $\mathbb{Z}^d$*, Electron. Commun. Probab. **19** (2014), no. 79, 13. MR 3283610

[FMS04] L. R. Fontes, F. P. Machado, and A. Sarkar, *The critical probability for the frog model is not a monotonic function of the graph*, J. Appl. Probab. **41** (2004), no. 1, 292–298. MR 2036292 (2004m:60221)

[GS09] Nina Gantert and Philipp Schmidt, *Recurrence for the frog model with drift on $\mathbb{Z}$*, Markov Process. Related Fields **15** (2009), no. 1, 51–58. MR 2509423 (2010g:60170)

[HJJ15a] Christopher Hoffman, Tobias Johnson, and Matthew Junge, *From transience to recurrence with Poisson tree frogs*, to appear in the *Annals of Applied Probability*, available at arXiv:1501.05874, 2015.

[HJJ15b] ———, *Recurrence and transience for the frog model on trees*, available at arXiv:1404.6238, 2015.

[Hoe63] Wassily Hoeffding, *Probability inequalities for sums of bounded random variables*, J. Amer. Statist. Assoc. **58** (1963), 13–30. MR 0144363 (26 #1908)

[Kal02] Olav Kallenberg, *Foundations of modern probability*, second ed., Probability and its Applications (New York), Springer-Verlag, New York, 2002. MR 1876169 (2002m:60002)

[KZ15] Elena Kosygina and Martin P. W. Zerner, *A zero-one law for recurrence and transience of frog processes*, available at arXiv:1508.01953, 2015.

[LMP05] Élcio Lebensztayn, Fábio P. Machado, and Serguei Popov, *An improved upper bound for the critical probability of the frog model on homogeneous trees*, J. Stat. Phys. **119** (2005), no. 1-2, 331–345. MR 2144514 (2006b:82068)

[MSH03] Neeraj Misra, Harshinder Singh, and E. James Harner, *Stochastic comparisons of Poisson and binomial random variables with their mixtures*, Statist. Probab. Lett. **65** (2003), no. 4, 279–290. MR 2039874 (2005a:60028)

[Pop01] Serguei Yu. Popov, *Frogs in random environment*, J. Statist. Phys. **102** (2001), no. 1-2, 191–201. MR 1819703 (2002a:82064)

[SS07] Moshe Shaked and J. George Shanthikumar, *Stochastic orders*, Springer Series in Statistics, Springer, New York, 2007. MR 2265633 (2008g:60005)

[TW99] András Telcs and Nicholas C. Wormald, *Branching and tree indexed random walks on fractals*, J. Appl. Probab. **36** (1999), no. 4, 999–1011. MR 1742145 (2001m:60199)

[Yu09] Yaming Yu, *Stochastic ordering of exponential family distributions and their mixtures*, J. Appl. Probab. **46** (2009), no. 1, 244–254. MR 2508516 (2010f:60065)

Department of Mathematics, University of Southern California and New York University
E-mail address: tobias.johnson@usc.edu

Department of Mathematics, University of Washington
E-mail address: jungem@math.washington.edu

SITE RECURRENCE FOR COALESCING RANDOM WALK

ITAI BENJAMINI, ERIC FOXALL, ORI GUREL-GUREVICH, MATTHEW JUNGE, AND HARRY KESTEN

ABSTRACT. Begin continuous time random walks from every vertex of a graph and have particles coalesce when they collide. We use a duality relation with the voter model to prove the process is site recurrent on bounded degree graphs, and for Galton-Watson trees whose offspring distribution has exponential tail. We prove bounds on the occupation probability of a site, as well as a general 0-1 law. Similar conclusions hold for a coalescing process on trees where particles do not backtrack.

1. INTRODUCTION

Coalescing random walk (CRW) starts with one particle at each vertex of an undirected graph. Each then performs a continuous time nearest neighbor random walk, jumping according to a mean 1 exponential clock. When particles collide they bind together and proceed as one. Say that CRW is *site recurrent* if every site is almost surely visited infinitely often. If instead this occurs with probability 0, call the process *transient*. Our main tool for proving site recurrence is the following necessary and sufficient condition in terms of the expected occupation time of a vertex.

Proposition 1. *Site recurrence is equivalent to infinite expected occupation time at any vertex. Moreover, CRW is either site recurrent or transient (i.e. it satisfies a 0-1 law).*

Let $p_t(v)$ be the probability a particle is at the vertex v at time t, so that site recurrence is equivalent to divergence of $\int p_t(v)$. We use duality with the voter model to obtain non-integrable lower bounds on the following graphs:

Theorem 2. *CRW is site recurrent on:*

(i) Bounded degree graphs. If the maximum degree is D, then for all vertices v

$$p_t(v) \geq (1 + Dt)^{-1}$$

for all $t \geq 0$.

(ii) Galton-Watson trees with offspring distribution on $\mathbb{Z}^+$ and the probability of k offspring bounded by e^{-ck} for some $c > 0$ and large enough k. Here

$$p_t(v) \geq C(t \log t)^{-1}$$

for all vertices v, large enough t, and some $C > 0$ that depends on c.

Note that there are unbounded degree graphs for which CRW is not site recurrent; even the non-coalescing system of independent random walks is transient on trees with rapidly increasing degree. We are not sure how much the exponential tail hypothesis in (ii) can be weakened. See Further Questions (a) for more discussion. A corollary to Proposition 1 is a general upper bound on the probability that a vertex v is unoccupied on the interval (t, u).

Corollary 3. *Let σ_t be the first time after t that v is occupied by a particle. It holds that*

$$\mathbf{P}(\sigma_t > u) \le \frac{t}{t + \int_t^u p_s(v)ds}.$$

And, for a graph with maximum degree D

$$\mathbf{P}(\sigma_t > u) \le \frac{t}{t + \frac{1}{D}(\log(1+Du) - \log(1+Dt))} = O(1/\log(u)).$$

We also give a universal upper bound for $p_t(v)$ on general graphs. It follows that the occupation probability decays to 0 for any graph. The upper bound is a small modification of an argument in [Gri78], so we also credit David Griffeath.

Proposition 4 (Griffeath)**.** *Let G be a connected, infinite graph. For all vertices v, any $\epsilon > 0$, and large enough t (depending on ϵ) it holds that $p_t(v) \le (1+\epsilon)/(2\sqrt{\pi t})$.*

History. The study of coalescing systems began in the 1970's with the paper of Erdős and Ney [EN74]. The duality relationship to the voter model, which we rely heavily upon, was first observed in [HL75]. Variations of coalescing random walk continue to find new applications. For example, random coalescence involving multiple types of particles, and particle interaction rules, is used to model certain chemical reactions (see [Hol83], [BL88] and [vdBK00]). Non-spatial models such as Kingman coalescence ([Kin82]) find applications in modeling ancestry in biology. A survey of coalescence models can be found in [Ber09]. Arratia [Arr83] looks at site recurrence for discrete time walks, and annihilating systems, both with possibly vacant sites in the starting configuration. CRW is applied to the voter model in [BL15]. Also, it is studied in more generality in [RV15] and [GPTZ15]. Other recent articles have focussed on different settings. Its behavior on finite graphs is of interest to computer scientists. The model on the d-dimensional torus is introduced in [Cox89]. There, they study the expected time for the process to coalesce into a single particle. In [CEOR12] the coalescence time is studied on a variety of finite graphs. Elsewhere, in a continuous spatial setting, recurrence is studied with coalescing diffusions by Cabezas, Rolla and Sidoravicious in [CRS13].

Early results for coalescing random walk focused on the lattices $\mathbb{Z}^d$. In [Gri78] Griffeath shows that both coalescing and annihilating random walk on $\mathbb{Z}^d$ is a.s. *weakly recurrent*, under certain restrictions on the vacant sites in the starting configuration. Weakly recurrent means that each site is occupied infinitely often, but for a decreasing fraction of time. An important ingredient in the proof of recurrence is an estimate for the function p_t, the probability a particle occupies the origin at time t. In [BG80], Bramson and Griffeath study p_t in the coalescing case and compute its asymptotics for every $d \ge 1$. Rather nicely, for $d \ge 3$ it holds that $p_t \sim (\gamma_d t)^{-1}$ with γ_d the probability a random walk on $\mathbb{Z}^d$ never returns to its starting point. The proof of this is computational; later, [vdBK02] Kesten gives a probabilistic argument that revolves around the heuristic $p_t' \approx -\gamma_d p_t^2$.

Overview. The main idea is to obtain information about $p_t(v)$ from a dual voter model. This dual was first applied to CRW in [HL75], and subsequently utilized in [HS79, Gri78, Arr81, Arr83]. In Corollary 6 we deduce that $p_t(v)$ is equivalent to the probability a time-changed nearest neighbor simple random walk avoids 0 up to time t. All of our estimates come from studying this random walk.

Further Questions. We record several questions regarding coalescing and annihilating random walk here:

(a) Can the assumptions on the degree in Theorem 2 (ii) be weakened? We expect that our approach extends (at least) to stationary graphs with finite expected degree.
(b) Suppose G is an infinite unimodular random graph in which each vertex is assigned an infinite trajectory in an ergodic invariant way (see [BC12]). Particles, one from each vertex, start moving along their trajectory in continuous time and annihilate when meeting. Is the resulting process recurrent? *Start by showing it on Euclidean lattices.*
(c) Place ϵ-balls (meteors) in Euclidean space with centers according to a unit intensity Poisson process. At time 0 each chooses a direction uniformly randomly and proceeds along this direction at unit speed (non-random). When two meteors collide, they annihilate. Is the origin a.s. occupied by infinitely many meteors for all $d \geq 1$? This is discussed in more detail in Section 3.
(d) Have particles perform annihilating random walk on a graph where the particle started at x steps according to an exponential clock with mean an independent uniform$(0,1)$ random variable. Is this model on $\mathbb{Z}^d$ still recurrent? If so, what can be said of the limiting speed of the particles visiting the origin? Possibly slower moving particles survive longer, and the average speed of particles visiting the origin decays with time.

Outline. Section 2 starts with the proof of Proposition 1. We also establish, in Lemma 5, that infinite expected occupation time is equivalent to survival of a nearest neighbor random walk in the voter model. Corollary 6 relates this back to p_t. Sections 2.1 and 2.2 contain the proofs of Theorem 2 (i) and (ii), respectively. Section 3 discusses some non-backtracking variants, and contains the proof of site recurrence for a non-backtracking model on bounded degree trees and Galton-Watson trees with exponential tail.

2. Site recurrence for coalescing random walk

Coalescing random walk on a graph, $G = (V, E)$, has a graphical representation as follows: each edge is replaced with two directed edges and an independent Poisson process with unit intensity is placed on each directed edge, indexed by time. When the bell of a Poisson process for the edge (u, w) rings we check if there's a particle at u and if so, we move it to w. If there's already a particle at w, they merge. We denote the process $(\xi_t)_{t\geq 0}$ with $\xi_t \in \{0,1\}^V$ equal to the set of occupied vertices at time t, and occasionally ξ_t^v for the location at time t of the particle that began at v. In this notation we have $p_t(v) = \mathbf{P}(\xi_t(v) = 1)$ is the probability that v is occupied by a particle at time t. Thus, $\int_0^\infty p_t(v)dt$ is the expected occupation time of v.

Proof of Proposition 1. If there is positive probability of infinite occupation time at v, then the expected occupation time is infinite. For the other direction we generalize [Arr83, Lemma 2].

Suppose that $\int_0^\infty p_t(v) = \infty$. For any $t \in [0, \infty)$ let $\sigma_t = \inf\{s \geq t \colon \xi_t(v) = 1\} \in [0, \infty]$, the first time after t that v is occupied by a particle. We wish to establish that

$$\mathbf{P}(\sigma_t < \infty) = 1, \qquad \forall t \geq 0.$$

The basic coupling $\xi_t^A = \{\xi_t^x \colon x \in A\}$ for $A \subseteq V$ has the property that $A \subseteq B \subseteq V$ implies $\xi_t^A \subseteq \xi_t^B \subseteq \xi_t$, so the Markov process $(\xi_t^A \colon A \subseteq V)$, with state space $\{0,1\}^V = \{A \colon A \subseteq V\}$ ordered by set inclusion, is attractive. Define $p_t^A = \mathbf{P}(\xi_t^A(v) = 1)$, and also $I(t, u) =$

$\int_t^u p_s(v)ds$. By assumption,

$$\lim_{u\to\infty} I(t,u) = \infty, \qquad \forall t \geq 0. \tag{1}$$

Let f_{σ_t} be the density function of σ_t and $\mathbf{E}_{\sigma_t}$ denote the expectation taken over all possible realizations of ξ_{σ_t} given that $\sigma_t = r$. Using the strong Markov property we have for $t < u$,

$$\begin{aligned} I(t,u) := \mathbf{E}\int_t^u \mathbf{1}\{\xi_s(v)=1\}ds &= \int_{r=t}^u f_{\sigma_t}(r)\mathbf{E}_{\sigma_t}\left(\mathbf{E}\int_0^{u-r}\mathbf{1}\{\xi_s^A(v)=1\}ds\right)dr \\ &\leq \int_{r=t}^u f_{\sigma_t}(r)\left(\mathbf{E}\int_0^u \mathbf{1}\{\xi_s(v)=1\}ds\right)dr \\ &\leq \mathbf{P}(\sigma_t \leq u)(t + I(t,u)). \end{aligned}$$

Dividing both sides by $t + I(t,u)$ we arrive at the inequality

$$\mathbf{P}(\sigma_t \leq u) \geq I(t,u)/(t + I(0,u)). \tag{2}$$

For fixed t, taking $u \to \infty$ yields $\mathbf{P}(\sigma_t < \infty) \geq 1$ by (1).

We conclude by describing the $0-1$ law. The above argument establishes that if CRW occupies a site for infinite time with positive probability, then it does so with probability 1. As G is assumed to be connected, it follows that all sites are occupied infinitely often with probability 1. Therefore, the process is either site recurrent (recall this is defined as an almost sure event) or transient. □

Proof of Corollary 3. The lower bound on $\mathbf{P}(\sigma_t \leq u)$ at (2) yields an upper bound on the probability v is unoccupied from time t to u:

$$\mathbf{P}(\sigma_t > u) \leq 1 - \frac{I(t,u)}{t + I(t,u)} = \frac{t}{t + \int_t^u p_s(v)ds}. \tag{3}$$

Which is the first part of the corollary. The second part follows by applying the bound on $p_t(v)$ in Theorem 2 (i) and integrating. □

Theorem 2 allows for site recurrence to be deduced by proving $p_t(v)$ is non-integrable. Our approach is to express $p_t(v)$ in another way. Consider the dual process to this model, which is called the *voter model*. In the dual model we start with a partition of the space into clusters, where initially each vertex corresponds to a different cluster. When the bell at (u,w) rings the vertex w is added to the cluster containing u. We denote the process $(\zeta_t^v)_{t\geq 0}$ where ζ_t^v is the set of vertices belonging to the cluster that initially consists of the vertex v. If we run this model in reverse time, from time t to 0, we see the cluster that began at v at time t at time 0 consists of exactly the particles that in the coalescing model are at v at time t. In particular,

$$p_t(v) = \mathbf{P}(\xi_t(v) = 1) = \mathbf{P}(\zeta_t^v \neq \varnothing). \tag{4}$$

The advantage of working with the voter model is that the size of ζ_t^v is a nearest-neighbor symmetric random walk with transition rate depending on the number of boundary edges. Indeed, at any moment there are some directed edges going out of the cluster, and the same number of edges coming in. More precisely, the cluster size is a skip-free process on the integers that moves with rate equal to the size of the current boundary of the cluster. We record this fact in the following lemma. Let $|\cdot|$ denote either the counting measure of a finite set.

Lemma 5. *Define* $\zeta_t^v \subseteq V$ *to be the set of vertices in the cluster of* v *at time* t*. For each* $x \in \zeta_t^v$ *let* $\partial_t(x) = \{(x,y) \in E \colon y \notin \zeta_t^v\}$*. Let* $\sum_{x\in\zeta_t^v} |\partial_t(x)|$ *be the number of edges leading out of* ζ_t^v*. The process has the following properties:*

(i) $|\zeta_0^v| = 1$.
(ii) *Let* $\tau = \inf\{t \colon |\zeta_t| = 0\}$*. For all* $t \geq \tau$ *it holds that* $|\zeta_t^v| = 0$.
(iii) *The process is a martingale that transitions to* $|\zeta_t^v| \pm 1$ *at rate* $\sum_{x\in\zeta_t^v} |\partial_t(x)|$.

Proof. Properties (i) and (ii) follow from the construction of ζ_t^v. We turn our attention to property (iii). For each $x \in \zeta_t^v$, x is removed from ζ_t^v at rate $|\partial_t(x)|$ and each of the $|\partial_t(x)|$ sites in $\partial_t(x)$ is added to ζ_t^v at rate 1. Since the rates balance,

$$\mathbf{E}[|\zeta_t^v| \mid |\zeta_{t-}^v|] = |\zeta_{t-}^v|,$$

which establishes $|\zeta_t^v|$ is a martingale. Summing the rates over $x \in \zeta_t^v$ shows that $|\zeta_t^v|$ transitions to $|\zeta_t| + 1$, and to $|\zeta_t^v| - 1$, each at rate $\sum_{x\in\zeta_t^v} |\partial_t(x)|$. □

This lets us describe p_t in terms of the voter model.

Corollary 6. $p_t(v) = \mathbf{P}(|\zeta_t(v)| > 0)$. *This is the probability a nearest neighbor random walk with transition rate* $\sum_{x\in\zeta_t} |\partial_t(x)|$ *and absorbing state at 0 is yet to reach zero at time* t.

It follows that $p_t \to 0$ on any infinite, connected graph.

Proof of Proposition 4. The transition rate in ζ_t^v is always at least two. By Corollary 6 and a straightforward coupling we have p_t is at most $\tilde{p}_t = \mathbf{P}(X_s > 0,\ \forall s \leq t)$, with X_s a rate-2 continuous time simple symmetric random walk started at 1. Using the reflection principle together with the local central limit theorem, $\tilde{p}_t \sim 1/(2\sqrt{\pi t})$ as $t \to \infty$, and the result follows. □

Remark 7. For coalescing walk on $\mathbb{Z}$ with nearest neighbour connections, since ζ_t^v is always of the form $\{x, x+1, ..., x+k\}$ for some $x \in \mathbb{Z}, k \in \mathbb{Z}^+$, its transition rate is exactly $2 \cdot 2 = 4$, so the above inequality is an equality, and gives the exact asymptotics $p_t \sim 1/(2\sqrt{\pi t})$, as observed in [BG80]. Compared to [BG80] there is an extra factor of 1/2; our convention differs from theirs in that the transition of a particle at v is equal to $\deg v$ and not 1, since in our case $\deg v$ is allowed to vary.

2.1. **Site recurrence for bounded degree graphs.** Now we turn our attention to proving site recurrence on general graphs. Define $\tau_v = \inf\{t \colon \zeta_t^v = \varnothing\}$. Integrating over t in the duality relation (4) we find

$$\mathbf{E}\tau_v = \int_0^\infty \mathbf{P}(|\zeta_t^v| > 0)dt = \int_0^\infty \mathbf{P}(\xi_t(v) = 1)dt = \int_0^\infty p_t(v)dt.$$

So, proving site recurrence is equivalent to showing that the first hitting time of 0 for the simple random walk $|\zeta_t^v|$ (with random and time-varying transition rate) has infinite expectation. We start with the case when G has bounded degree.

Proof of Theorem 2 *(i).* Let $v \in V$ with the maximum degree of vertices in G bounded by D. Lemma 5 establishes that the transition rate of ζ_t^v is less than or equal to $D|\zeta_t^v|$. Let W_t be a continuous time nearest-neighbour random walk on $\mathbb{Z}^+ \cup \{0\}$ with $W_0 = 1$. The walk transitions from $k \in \mathbb{Z}^+ \cup \{0\}$ to $k \pm 1$ each at rate Dk, and is absorbed at 0. Letting $\theta(t) = \mathbf{P}(W_t > 0)$, if follows from Corollary 6 and a straightforward coupling of $|\zeta_t^v|$ with W_t that $p_t(v) \geq \theta(t)$, so it suffices to control $\theta(t)$.

The process W_t can be interpreted as the number of particles in a branching process in which each particle independently dies, or gives birth to a single offspring, each at rate D. Let $\rho(t) = \mathbf{P}(W_t = 0)$. By independence, for each $k \geq 0$ we have

$$\mathbf{P}(W_{t+h} = 0 \mid W_h = k) = \rho(t)^k. \tag{5}$$

Recalling that $W_0 = 1$ then conditioning on W_h for small $h > 0$,

$$\rho(t+h) = \mathbf{P}(W_{t+h} = 0) = \sum_{k \geq 0} \mathbf{P}(W_{t+h} = 0 \mid W_h = k)\mathbf{P}(W_h = k).$$

The event that two or more transitions happens on the interval $[0, h]$ is contained in the event that a rate $2D$ exponential clock rings, then a rate $4D$ exponential clock rings (i.e. we go from 1 to 2 particles then another transition happens). The probability of this is bounded by the probability that $X + Y < h$ for X and Y rate $4D$ exponential random variables, and has density $f_{X+Y}(t) = \lambda^2 te^{-\lambda t}$ with $\lambda = 1/4D$. Integrating on $[0, h]$, then taking the Taylor expansion we have $\mathbf{P}(X + Y < h) = O(h^2)$. Since we will be dividing by h and letting it tend to 0, we can combine all of the events that occur with two or more transitions into an $O(h^2)$ term. Using the expression (5) this lets us write $\rho(t + h)$ as

$$\rho(t+h) = \underbrace{1 \cdot Dh}_{\text{dies out}} + \underbrace{\rho(t)(1-2Dh)}_{\text{no change}} + \underbrace{\rho(t)^2 Dh}_{\text{increases by 1}} + O(h^2).$$

Subtracting $\rho(t)$, dividing by h and taking $h \downarrow 0$ this converges to the equation $\rho' = D(1-\rho)^2$. So, for the survival probability $\theta(t) = 1 - \rho(t)$ we find $\theta' = -\theta^2$, with $\theta(0) = 1$, whose unique solution is $\theta(t) = 1/(1 + Dt)$. □

2.2. **Site recurrence for Galton-Watson trees.** A more general upper bound on the transition rate is

$$|\text{maximum exposed degree}| \cdot |\zeta_t|.$$

Our hypothesis that the offspring distribution of our Galton-Watson tree has exponential tail guarantees that the maximum exposed degree is asymptotically bounded by log(number of steps). Ultimately this lets us compare with the divergent integral $\int_{t_0}^{\infty}(t \log t)^{-1}dt$. This is made rigorous below.

Proof of Theorem 2 *(ii).* Again, by Corollary 6 it suffices to prove that ζ_t^v has infinite expected survival time. For convenience we denote ζ_t^v by ζ_t. Let $H_t = \cup_{s \leq t}\zeta_t$ be the vertices visited up to time t. Define the random times $0 = t_1 < t_2 < ...$ as when a vertex is added to H_t, and list them as $v_1, v_2, ...$ in the order they are discovered, with v_1 being the root (ρ). The transition rate of $|\zeta_t| \to |\zeta_t| \pm 1$ is at most $M_t|\zeta_t|$ where

$$M_t = \sup\{\deg v \colon v \in H_t\}.$$

So, our first goal is to construct the voter model in such a way that M_t can be easily controlled. A simple way to do this is to construct G "on the fly." That is, let $(X_i)_{i \geq 1}$ be an i.i.d. sequence of copies of the offspring distribution, and at time t_i, sample X_i to determine the offspring distribution of v_i, which is then fixed for all $t > t_i$. This does not disturb the sample path distribution of ζ_t, and has the advantage that the quantity

$$D_k = \max_{i \leq k} \deg v_i$$

is equal to $\max_{i \leq k} X_i$ where X_i is a *fixed* (as opposed to being a randomly indexed) i.i.d. sequence. Since, by assumption, $\mathbf{P}(X_i > x) \leq e^{-cx}$ for some $c > 0$ and large enough x, a union bound gives $\mathbf{P}(D_k > x) \leq ke^{-cx}$, and setting $x = (3/c)\log k$,

$$\mathbf{P}(D_k > (3/c)\log k) \leq k^{-2} \tag{6}$$

for large enough k.

Now, let $0 = t_0 < t_1 < t_2 < ...$ denote the jump times of ζ_t. In what follows we will want the set of jump times to be infinite, so if $\zeta_{t_i} = 0$ (i.e. the cluster dies out), just include jumps back to 0 at rate 1. Since $|\xi_i|$ is a martingale with $\mathbf{E}|\zeta_{t_i}| = |\zeta_0| = 1$ for all i, Doob's martingale inequality implies that for all $n > 0$

$$\mathbf{P}(\sup_i |\zeta_{t_i}| > bn) \leq (bn)^{-1} \tag{7}$$

for any $b > 0$. Clearly M_{t_i} is nondecreasing in i and $M_{t_i} \leq D_{1+i}$, since vertices are exposed one at a time. Thus, the transition rate of $|\zeta_{t_i}|$ is at most $D_{i+1}|\zeta_{t_i}|$. Combining these observations with (6) and (7), we find that with probability at least $1 - (bn)^{-1} - n^{-2}$, for $t \leq t_{n^2}$ the transition rate in ζ_t is at most

$$(3/c)\log(n^2+1)bn \leq (6b/c)n\log(n+1). \tag{8}$$

Let $m_n = m_n(b,c) = (6b/c)n\log(n+1)$. A quick summary: with high probability the first n^2 transitions happen at rate no more than m_n. Equivalently, the time t_{n^2} is bounded below by the sum of n^2 independent exponentials with rate m_n. This is an Erlang distribution, X_{n^2}, with shape parameter n^2 and rate m_n. Thus, we have

$$\text{mean: } \mu_n = \frac{n^2}{m_n}, \text{ and variance: } \sigma_n = \frac{\mu_n}{m_n}.$$

Chebyshev's inequality guarantees that

$$\mathbf{P}(|X_n - \mu_n| \geq \mu_n/2) \leq \frac{\mu_n/m_n}{(\mu_n/2)^2} = \frac{4}{\mu_n m_n} = \frac{4}{n^2}.$$

One side of the above estimate is

$$\mathbf{P}(X_{n^2} \leq \mu_n/2) \leq 4n^{-2}. \tag{9}$$

By comparison, and using (8) and (9), we have

$$\mathbf{P}(t_{n^2} \geq n^2/(2m_n)) \geq 1 - (bn)^{-1} - n^{-2} - 4n^{-2}.$$

From the well-known first passage distribution for random walk, for the random variable $N = \inf\{n\colon |\zeta_n| = 0\}$ we have $P(N > n^2) \geq c/n$ for some possibly smaller $c > 0$. Note that although m_n depends on c, shrinking c does not affect the estimate. Letting $\tau = \inf\{t\colon |\zeta_t| = 0\}$ as before, and letting $a_n = \mu_n/2 = cn/(12b\log(n+1))$, for $a_{n-1} < t < a_n$,

$$t\log t \geq \frac{c(n-1)}{12b\log n}(\log(n-1) - \log(12b/c) - \log\log n) \geq \frac{cn}{24b}$$

i.e., $n \leq (24b/c)t\log t$ for n large enough. For the same t, then,

$$\mathbf{P}(\tau > t) = \mathbf{P}(|\zeta_t| > 0) \geq \mathbf{P}(|\zeta_{t_{n^2}}| > 0, t_{n^2} > a_n).$$

Since the survival time of the cluster is independent of the rate at which it jumps we have

$$\mathbf{P}(\tau > t) \geq c/n(1 - (bn)^{-1} - O(n^{-2})), \qquad a_{n-1} < t < a_n.$$

The right side is at least $c/(2n)$ for n greater than some n_0. Letting $t_0 = a_{n_0}$ and using the upper bound on n,

$$\mathbf{E}\tau = \int_0^\infty \mathbf{P}(\tau > t)dt \geq \frac{c^2}{48b}\int_{t_0}^\infty \frac{1}{t\log t}dt = \infty.$$

□

3. Non-backtracking coalescing random walk on trees

We are also interested in understanding similar, but less random processes. A lack of symmetry in these settings makes it difficult to apply known techniques. We are hopeful that progress will lead to new ideas.

The *non-backtracking coalescing random walk* is defined in the same way as the coalescing random walk with particles instead performing non-backtracking random walk. More precisely, the state of a particle is specified by a vertex-edge pair $(u, \{u, v\})$, and when an edge clock rings at a directed edge (u, w), the particle moves to w if and only if $w \neq v$. If the particle moves from u to w, its state is updated to $(w, \{w, u\})$, so that its next jump cannot be back to u. It will be convenient to assume that each particle is initialized with a uniformly chosen edge along which it cannot move, that is, the particle initially at v has state $(v, \{v, u\})$ where u is a uniform random neighbour of v. With particles coalescing there is ambiguity about whose path to remember. There are several well-defined ways to assign priority. On a rooted tree, we analyze the special case where we always remember the path of particles moving towards the root. With the model defined in this way we do not quite have a voter model dual, but a closely related process does. Analogous to Proposition 1 we prove a necessary and sufficient condition for site recurrence.

Proposition 8. *Consider coalescing non-backtracking random walk on a rooted tree with priority given to particles moving towards the root. The process is site recurrent at the root if and only if the expected survival time of the cluster in a certain voter model is infinite.*

We can deduce site recurrence on bounded degree trees and some trees with unbounded degree.

Theorem 9. *The process from Proposition* 8 *is site recurrent at the root of either a bounded degree tree or a Galton-Watson tree whose offspring distribution is as in Theorem* 2 *(ii).*

Non-backtracking removes a vital symmetry from the argument. The proof goes by, once again, constructing a dual voter process and showing the cluster of the root survives for an infinite expected amount of time. Our "priority to the root" rule is hand-picked to preserve monotonicity and the existence of a dual voter model. Neither property exists in other equally natural non-backtracking models. Further progress in these different settings will likely require a new approach. Consider the following conjecture:

Conjecture 10. *Non-backtracking coalescing random walk with any priority scheme is site recurrent on bounded degree trees.*

The inspiration for studying non-backtracking processes comes from the following *meteor model* on $\mathbb{R}^d$. Place ϵ-balls in Euclidean space with centers according to a unit intensity Poisson process. At time 0 each chooses a direction uniformly randomly and proceeds along this direction at unit speed (non-random). When two meteors collide, they annihilate.

Conjecture 11. *The origin a.s. is occupied by infinitely many meteors for all* $d \geq 1$.

This problem appears quite difficult. It could be discretized to an annihilating system of random walks by uniformly assigning each particle a geodesics to ∞ from which it never deviates and steps along according to a Poisson clock. The integer lattice is a natural graph to start with. Or, perhaps hyperbolic space—in which random walk paths stay within a logarithmic neighborhood of a geodesic—would be a more tractable place to study this problem.

3.1. **Recovering a dual.** Since the graph is a tree, a particle either moves towards the root, or away, at each jump. Once it has moved away for the first time, at subsequent jumps it must always move away, since the only way back towards the root requires backtracking. To simplify matters, we suppose that at each vertex, the initial forbidden edge is chosen uniformly from the edges that lead away from the root.

Since coalescing particles may have different histories we must decide which one to remember. Therefore, upon collision we define the following rule for annihilating exactly one of the two colliding particles; in the original setting with memoryless walks, any such rule leads to the coalescing model.

- If one particle is moving towards the root, and the other particle is moving away, annihilate the particle moving away.
- Otherwise, annihilate the particle currently occupying the site (i.e. keep the particle that just moved).

As currently stated, this process depends on past information. Running the process in reverse would require information about the future. Thus, it does not have a dual voter model. Still, a simple observation yields a related model that does have a dual.

Let X_T be the occupation time of the root up to time T. Notice that particles moving away from the root are inert; based on the rule above they cannot block upward moving particles, and as noted before they cannot revisit the root. Thus, X_T is unchanged if we delete particles the instant they turn away from the root. This modification gives the following model; for a vertex v, let $d_v = \deg v$.

- Suppose v is not the root.
  - A particle at v moves towards the root at rate 1.
  - A particle at v is deleted at rate $d_v - 2$.
- Suppose $v = \rho$ is the root. A particle at ρ is deleted at rate $d_\rho - 1$.

We can think of this model as follows: each particle attempts to travel up to the root, coalescing with other particles upon collision, and particles (or coalesced collections of particles) are instantaneously zapped out of existence at some rate that depends on their present location.

We can simplify the description somewhat by introducing a single absorbing vertex $\mathfrak{a}$ and considering the process on $V \cup \{\mathfrak{a}\}$, where every vertex has a directed edge pointing to v which rings at rate described below.

Particles at $\mathfrak{a}$ do not move. The transitions for particles at $v \in V$ are as follows.

- Move towards the root at rate 1.
- Move directly to $\mathfrak{a}$ at rate $d_v - 2$.

A graphical representation of this model can be obtained by placing an independent Poisson process with rate 1 at each upward directed edge, and with rate $d_v - 2$ at each vertex v.

The model enjoys the same monotonicity as the coalescing random walks – resetting to the initially full configuration maximizes the probability of occupying the root in the future. Then, as for the coalescing random walks, there is a dual voter model. In this case, deletion of a particle at v corresponds to the addition of v to the cluster of $\mathfrak{a}$. Note that since the direction of motion is reversed in the voter model, clusters on the tree must expand away from the root. Altogether, the voter model has the following transitions. A *down-going* directed edge (w, u) is an edge directed away from $\mathfrak{a}$. These are the rules we use in the proof of Lemma 12. We box it for emphasis:

- Along each down-going edge (w, u), at rate 1, u is added to the cluster containing w.
- At each vertex v on the tree, at rate $d_v - 2$, v is added to the cluster of $\mathfrak{a}$.

The existence of a dual lets us establish Proposition 8.

Proof of Proposition 8. We have established that X_T in the non-backtracking coalescing random walk has the same distribution as in the simpler model. Since the simpler model has monotonicity and a voter model dual, the same argument used to prove Proposition 1 gives the desired equivalence. □

3.2. **Site recurrence.** Now we can turn our attention to proving infinite expected survival time of the root cluster in the voter model. The voter model from Section 3.1 also has the martingale property. As before, let ζ_t^v denote the cluster that began at v.

Lemma 12. *Consider the voter model $\zeta_t^v \in V \cup \{\mathfrak{a}\}$ described above. For each $v \in V$, so long as $|\zeta_t^v| < \infty$, the size of the cluster ζ_t^v is a martingale. It transitions to $|\zeta_t| \pm 1$ at rate $\sum_{w \in \zeta_t}(d_w - 1) - 2|\{(w, y) \colon w, u \in \zeta_t^v\}|$.*

Proof. Fix a vertex $v \in V$ for which we will consider the cluster $\zeta_t = \zeta_t^v$. Let $r_t^+ = r_t^+(v)$ and $r_t^- = r_t^-(v)$ denote the rate at which $|\zeta_t| \to |\zeta_t| \pm 1$, respectively. Note that the transition rules prohibit $\mathfrak{a} \in \zeta_t$. Moreover, ζ_t is unchanged if we assume that initially, all vertices but v belong to the cluster of $\mathfrak{a}$. Therefore, it is enough to check that for any finite W, if $\zeta_t = W$ and $\zeta_t^{\mathfrak{a}} = V \cup \{a\} \setminus W$ then $r_t^+ - r_t^- = 0$. For a vertex $w \neq \mathfrak{a}$ let $\hat{w}$ denote its unique parent vertex, i.e., the unique vertex such that there is a down-going edge to w, and let $\mathfrak{o}(w)$ denote the set of childs vertices. From the transition rules it follows that

$$r_t^+ = \sum_{w \in W}(d_w - 1) - |\{u \in \mathfrak{o}(w) \colon u \in W\}|$$

and

$$r_t^- = \sum_{w \in W}(d_w - 2) + \mathbf{1}(\hat{w} \notin W) = \sum_{w \in W}(d_w - 1) - \mathbf{1}(\hat{w} \in W)$$

and since $\sum_{w \in W} |\{u \in \mathfrak{o}(w) : u \in W\}|$ and $\sum_{w \in W} \mathbf{1}(\hat{w} \in W)$ are both equal to $|\{(w, u) \colon u, w \in W\}|$,

$$r_t^+ - r_t^- = -\sum_{w \in W} |\{u \in \mathfrak{o}(w) \colon u \in W\}| + \sum_{w \in W} \mathbf{1}(\hat{w} \in W) = 0.$$

□

Proof of Theorem 9. With Lemma 12 we can bound the transition rate of $|\zeta_t^\rho|$ by $\sum_{v \in \zeta_t^\rho} d_v$. To prove the part of Theorem 9 concerning bounded degree trees we can follow the same approach as Theorem 2 (i); again we use the transition rate bound $D \cdot |\zeta_t^\rho|$. Similarly, we can use the same technique as the proof of Theorem 2 (ii) to deduce site recurrence for Galton-Watson trees. □

References

[Arr81] Richard Arratia, *Limiting point processes for rescalings of coalescing and annihilating random walks on z^d*, Ann. Probab. **9** (1981), no. 6, 909–936.

[Arr83] ______, *Site recurrence for annihilating random walks on $\mathbb{Z}^d$*, Ann. Probab. **11** (1983), no. 3, 706–713.

[BC12] Itai Benjamini and Nicolas Curien, *Ergodic theory on stationary random graphs*, Electron. J. Probab. **17** (2012), no. 93, 1–20.

[Ber09] Nathanaël Berestycki, *Recent progress in coalescent theory*, Ensaios Matematicos **16** (2009), no. 1.

[BG80] Maury Bramson and David Griffeath, *Asymptotics for interacting particle systems on* $\mathbb{Z}^d$, Zeitschrift fr Wahrscheinlichkeitstheorie und Verwandte Gebiete **53** (1980), no. 2, 183–196 (English).

[BL88] Maury Bramson and Joel L. Lebowitz, *Asymptotic behavior of densities in diffusion-dominated annihilation reactions*, Phys. Rev. Lett. **61** (1988), 2397–2400.

[BL15] M. Balázs and A. László Nagy, *Dependent augmented Branching Annihilating Random Walk*, ArXiv e-prints (2015).

[CEOR12] Colin Cooper, Robert Elsässer, Hirotaka Ono, and Tomasz Radzik, *Coalescing random walks and voting on graphs*, CoRR **abs/1204.4106** (2012).

[Cox89] J. T. Cox, *Coalescing random walks and voter model consensus times on the torus in* $\mathbb{Z}^d$, Ann. Probab. **17** (1989), no. 4, 1333–1366.

[CRS13] M. Cabezas, L. T. Rolla, and V. Sidoravicius, *Recurrence and Density Decay for Diffusion-Limited Annihilating Systems*, ArXiv e-prints (2013).

[EN74] P. Erdos and P. Ney, *Some problems on random intervals and annihilating particles*, Ann. Probab. **2** (1974), no. 5, 828–839.

[GPTZ15] B. Garrod, M. Poplavskyi, R. Tribe, and O. Zaboronski, *Interacting particle systems on Z as Pfaffian point processes I - annihilating and coalescing random walks*, ArXiv e-prints (2015).

[Gri78] David Griffeath, *Annihilating and coalescing random walks on d*, Zeitschrift fr Wahrscheinlichkeitstheorie und Verwandte Gebiete **46** (1978), no. 1, 55–65 (English).

[HL75] Richard A. Holley and Thomas M. Liggett, *Ergodic theorems for weakly interacting infinite systems and the voter model*, Ann. Probab. **3** (1975), no. 4, 643–663.

[Hol83] Richard Holley, *Two types of mutually annihilating particles*, Advances in Applied Probability **15** (1983), no. 1, 133–148 (English).

[HS79] R. Holley and D. W. Stroock, *Central limit phenomena of various interacting systems*, Annals of Mathematics **110** (1979), no. 2, pp. 333–393 (English).

[Kin82] J.F.C. Kingman, *The coalescent*, Stochastic Processes and their Applications **13** (1982), no. 3, 235 – 248.

[RV15] B. Rath and D. Valesin, *Percolation on the stationary distributions of the voter model*, ArXiv e-prints (2015).

[vdBK00] J. van den Berg and Harry Kesten, *Asymptotic density in a coalescing random walk model*, Ann. Probab. **28** (2000), no. 1, 303–352.

[vdBK02] ———, *Randomly coalescing random walk in dimension* ≥ 3, 1–45 (English).

Department of Mathematics, Weizmann Institute
E-mail address: `itai.benjamini@weizmann.ac.il`

Department of Mathematical and Statistical Sciences, Arizona State University
E-mail address: `Eric.Foxall@asu.edu`

Einstein Institute of Mathematics, Hebrew University of Isreal
E-mail address: `origurel@math.huji.ac.il`

Department of Mathematics, University of Washington
E-mail address: `jungem@math.washington.edu`

Department of Mathematics, Cornell University
E-mail address: `kesten@math.cornell.edu`

CHOICES, INTERVALS AND EQUIDISTRIBUTION

MATTHEW JUNGE

Abstract. We give a sufficient condition for a random sequence in [0,1] generated by a Ψ-process to be equidistributed. The condition is met by the canonical example – the max-2 process – where the nth term is whichever of two uniformly placed points falls in the larger gap formed by the previous $n-1$ points. This solves an open problem from Itai Benjamini, Pascal Maillard and Elliot Paquette. We also deduce equidistribution for more general Ψ-processes. This includes an interpolation of the min-2 and max-2 processes that is biased towards min-2.

1. Introduction

A sequence in $[0,1]$ is *equidistributed* if the limiting proportion of points in each subinterval is equal to the subinterval's length. Over a century ago Weyl proved that $\{\beta n \mod 1\}_{n\geq 1}$ is equidistributed for any irrational number β (see [Wey10]). Since then connections have been found in ergodic theory, number theory, complex analysis and computer science ([BM72], [Vau77], [FSZ09], [CKK$^+$07]). See [KN06] for an overview.

Not long after Weyl's Theorem, attention turned to equidistribution of random sequences. One way to obtain a random sequence in $[0,1]$ is to independently choose points uniformly. Call the resulting sequence the *uniform process*. The strong law of large numbers guarantees this is equidistributed almost surely.

Another random process known to equidistribute points is the *Kakutani interval splitting procedure* (introduced in [Kak76]), where at each step a point is added uniformly to the current largest subinterval. Almost sure equidistribution is proven in [Zwe78] and [Loo78] using stopping times. Because points are placed in the largest gaps they ought to spread more evenly than the uniform process. Indeed, [Pyk80] proves the size of the largest interval is asymptotic to $2/n$; the same order as the average interval. Compare to $\log n/n$ in the uniform process (see [Dar53]).

[MP14] introduces a family of interval splitting processes that exhibit a wider range of behavior. The canonical example is the *max-2 process*. The dynamics are as follows:

- Partition $[0,1]$ into subintervals by placing finitely many points in any manner.
- At each step sample two points uniformly from $[0,1]$. Each lies in a subinterval formed by the previous configuration.
- Keep the point contained in the larger subinterval and disregard the other point. Break a tie by flipping a fair coin.

A discrete analogue of the max-2 process appears in [ABKU99] where n balls are placed into n bins. For each ball two bins are selected uniformly and the ball is placed in the bin with fewer balls. They find that the most-filled bin has $\approx \log_2 \log n$ balls; significantly less then $\approx \log n/\log\log n$ if the balls were instead placed uniformly. This is studied in more detail in [MRS00] and [LM05].

In the max-2 process choosing the larger gap should spread points more evenly. Despite our intuition this is difficult to formalize, and equidistribution was a primary open problem from [MP14]. The natural counterpart is the *min-2 process* where the point contained in the smaller subinterval is kept. Unlike the previous processes, points are prone to clump together. It is natural to also define the max-k and min-k processes; in these the max or (resp.) min of k candidate points is selected at each step.

Before we can state the theorem we describe a more general splitting procedure known as a Ψ-*process* (introduced in [MP14]). For technical convenience we will assume that points arrive according to a Poisson process with intensity e^t. Suppose at time t that N_t points have arrived and we have interval lenghts $I_1^{(t)}, I_2^{(t)}, \ldots, I_{N_t+1}^{(t)}$. Define the size-biased empirical distribution function

$$\widetilde{A}_t(x) = \sum_{i=1}^{N_t+1} I_i^{(t)} \mathbf{1}\{I_i^{(t)} \leq x\}.$$

This function is now defined to evolve according to Markovian dynamics as follows. Let us say that the next point arrives at time $s > t$, for the N_s-th step (with $N_s = N_t + 1$) we choose an interval at random, with length $\ell_s = \widetilde{A}_{s^-}^{-1}(u)$, where u is sampled from a law on $(0, 1]$ whose distribution function we denote by Ψ. This randomly chosen interval is now subdivided into two pieces at a point chosen uniformly inside the interval. This produces a new sequence of interval lengths $I_1^{(s)}, \ldots I_2^{(s)}, \ldots, I_{N_s}^{(s)}$ and the process is repeated. Note that $\widetilde{A}_t(x)$ is constant (in t) between point arrivals. We remark that the max-k, uniform and min-k processes are Ψ-processes with $\Psi(u) = u^k, u$, and $1 - (1 - u)^k$, respectively.

We abbreviate a few common assumptions for Ψ:

- (C) Ψ is continuous.
- (C^1) Ψ is continuously differentiable.
- (C^2) Ψ is twice continuously differentiable.
- (D) There exist $c > 0$ and $\kappa_\Psi \in [1, \infty)$, such that $1 - \Psi(u) \geq c(1 - u)^{\kappa_\Psi}$ for all $u \in (0, 1)$.

Set $A_t(x) = \widetilde{A}_t(e^{-t}x)$. The main theorem of [MP14] proves that, when (C) and (D) hold, $A_t(x)$ converges pointwise to a (deterministic) continuously differentiable distribution function $F^\Psi(x)$. For future theorem statements we note that (C^1) and (C^2) both imply (D).

Here we study $\widetilde{A}_t^\alpha$, the restriction of $\widetilde{A}_t$ to the N_t^α subintervals contained in $[0, \alpha]$. We find conditions on Ψ that guarantee pointwise convergence $A_t^\alpha \to \alpha F^\Psi$, where $A_t^\alpha(x) = \widetilde{A}_t^\alpha(e^{-t}x)$ and αF^Ψ denotes the map $x \mapsto \alpha \cdot F^\Psi(x)$. When this holds the subinterval lengths in $[0, \alpha]$ evolve to look the same as those in all of [0,1]. This sameness is enough to deduce equidistribution.

Theorem 1. *Let $\psi = \Psi'$. If Ψ satisfies* (C^2) *and for some $\delta \in (0, 1]$ and all $z \geq 0$*

$$|z\psi'(F^\Psi(z))(F^\Psi)'(z) - \psi(F^\Psi(z))| \leq (2 - \delta)\psi(F^\Psi(z)), \tag{1}$$

then the Ψ-process is equidistributed a.s.

The condition (1) arises from a technical computation (see the proof Proposition 4) used to show that a family of processes containing $(A_t^\alpha)_{t \geq 0}$ contract in a certain norm. We stress that it is not at all obvious which Ψ and F^Ψ should satisfy this condition. Our only tools are the properties of F^Ψ established in [MP14]. Most importantly, it satisfies the integro-differential equation (see [MP14, Lemma 3.5]):

$$(F^\Psi)'(z) = z \int_z^\infty \frac{1}{y} d\Psi(F^\Psi(y)), \tag{2}$$

and the differential equation (see [MP14, Proposition 8.1]):

$$z(F^{\Psi})''(z) - (F^{\Psi})'(z) + z\psi(F^{\Psi}(z))(F^{\Psi})'(z) = 0. \tag{3}$$

Remarkably, this is enough information to deduce (1) holds for the max-2 process, an interpolation of max-2 and min-2 processes that is biased towards min-2, and arbitrary interpolations of max-k, uniform and min-k processes that place enough weight on the uniform process.

Corollary 2. *The following are equidistributed a.s.*

(1) The max-2 *process.*
(2) The interpolation that is 60%-min-2 *and* 40%-max-2*;* $\Psi(u) = .6(1-(1-u)^2) + .4u^2$.
(3) The interpolation of max-k*, uniform and* min-k *processes given by a probability measure* $\mathbf{p} = (p_k)_{k\neq -1,0}$ *on* $\mathbb{Z}\setminus\{-1,0\}$*, that satisfes* $\sum_{k\geq 2} k(k-1)[p_k + p_{-k}] \leq 1/2$;

$$\Psi(u) = p_1 u + \sum_{k\geq 2} p_k u^k + p_{-k}(1-(1-u)^k).$$

For example, this includes the interpolations
(a) $(1/k^2)$%-min-k *for a single fixed* k *and otherwise* uniform*.*
(b) 99.95%-uniform *and* (5^{-k})%-min-k *for all* $k = 2, 3, \ldots$*.*

The reason our approach works for only certain Ψ is unclear. Numerical methods indicate the inequality fails for other processes, suggesting a different approach is needed. This is surprising since processes which ought to better equidistribute points, like a max-3 process, do not meet our criterium. Nonetheless, we conjecture that all max-k and min-k processes are equidistributed. The properties established in Proposition 5 are an important step in exploring this for max, min and more general Ψ-processes. The rate of convergence to a uniform placement of points and also the asymptotic size of the largest interval are other important open problems. More thorough discussion can be found in [MP14].

Overview. This article is organized to quickly arrive at the proof of Theorem 1. In Section 2 we describe the evolution of intervals in $[0,\alpha]$ and give the major definitions. In Section 3 we state without proof Proposition 4 and Proposition 5. The first proposition describes the importance of (1) holding. The second shows that A_t^{α} has similar properties as those needed of A_t to deduce convergence in [MP14]. We then use this to establish Theorem 1. Section 4 contains the proofs for the previous section. Finally, in Section 5 we prove Corollary 2 by showing that various interpolations satisfy (1).

2. Subintervals in $[0,\alpha]$

We start with a formal definition for a process to be equidistributed. Suppose n_0 points are initially placed. After n iterations of an interval splitting process let N_n^{α} be the number of the first $n_0 + n$ points smaller than α. We say a sequence is *equidistributed* if $n^{-1}N_n^{\alpha} \to \alpha$ for all $\alpha \in [0,1]$. It is convenient to work in continuous time. Following [MP14] we have points arrive as a Poisson process with intensity e^t. Formal details are in Proposition 4. So, in continuous time equidistribution is equivalent to $e^{-t}N_t^{\alpha} \to \alpha$ for all $\alpha \in [0,1]$.

2.1. Describing $\widetilde{\mathbf{A}}_t^{\alpha}$.

Fix $\alpha \in [0,1]$. We use the convention that a bold face letter represents a process indexed by time (i.e. $\widetilde{\mathbf{A}} = (\widetilde{A}_t)_{t\geq 0}$). Define the joint processes $(\widetilde{\mathbf{A}}^{\alpha}, \widetilde{\mathbf{A}}^{\alpha+}, \widetilde{\mathbf{A}})$ to be the size-biased empirical distributions of interval lengths contained in $[0,\alpha]$, $[\alpha,1]$ and $[0,1]$, respectively. Formally, letting $I_1^{\alpha,(t)},\dots,I_{N_t^\alpha}^{\alpha,(t)}$ be the lengths of subintervals contained in $[0,\alpha]$ at time t we define

$$\widetilde{A}_t^\alpha(x)=\sum_{j=1}^{N_t^\alpha} I_j^{\alpha,(t)}\cdot \mathbf{1}\{I_j^{\alpha,(t)}\le x\},$$

and similarly for $\widetilde{A}_t^{\alpha+}$ and $\widetilde{A}_t$. The spark for the refined analysis comes from the relation

$$\widetilde{A}_t^\alpha(x)+\widetilde{A}_t^{\alpha+}(x)=\widetilde{A}_t(x),\qquad \forall t,x\ge 0. \tag{4}$$

To ensure that no intervals are double counted assume the initial set of points placed in $[0,1]$ always contains $\{\alpha\}$. This assumption is only for convenience. Our proof could be adapted to omit it by running the process until two points $\alpha_1\le\alpha\le\alpha_2$ land sufficiently close to α, and then using the bound $N_t^{\alpha_1}\le N_t^\alpha\le N_t^{\alpha_2}$. We further remark that the same reasoning extends our theorems to the unit circle.

In [MP14, Section 2] the authors prove that

$$\widetilde{A}_t(x)=\widetilde{A}_0(x)+\int_0^t e^s x^2\int_x^\infty \frac{\psi(\widetilde{A}_s(z))}{z}d\widetilde{A}_s(z)+\widetilde{M}_t$$

for some martingale $\widetilde{M}_t$. The following proposition shows that $\widetilde{A}_t^\alpha$ satisfies a similar equation.

Proposition 3. *Let $\psi=\Psi'$. For any Ψ-process satisfying* (C^1)*, the joint processes* $(\widetilde{\mathbf{A}}^\alpha,\widetilde{\mathbf{A}}^{\alpha+},\widetilde{\mathbf{A}})$ *satisfy the equation*

$$\widetilde{A}_t^\alpha(x)=\widetilde{A}_0^\alpha(x)+\int_0^t e^s x^2\int_x^\infty \frac{\psi(\widetilde{A}_s(z))}{z}d\widetilde{A}_s^\alpha(z)ds+\widetilde{M}_t^\alpha(x),$$

with $\widetilde{M}_t^\alpha$ a martingale.

Proof. We first build up some necessary definitions. Let Ψ be a continuously differentiable distribution function. Define a Poisson random measure $\prod$ on $[0,\infty)\times[0,1]^2$ with intensity $e^t dt\otimes d\Psi(u)\otimes dv$. Set $\ell_t(u)=\widetilde{A}_{t^-}^{-1}(u)$. We use the function $h(v,\ell,x)=v\mathbf{1}\{\ell v\le x\}+(1-v)\mathbf{1}\{\ell(1-v)\le x\})$ to "cut" our sampled interval by v.

We need to detect whether the sampled interval belongs to $[0,\alpha]$. We use the function $g_t^\alpha(\ell_t(u))=\mathbf{1}\{\ell_t(u)\subset[0,\alpha]\}$. The function g_t^α can be constructed rigorously by assuming all of the subintervals have different lengths, and putting a point mass on each length of subintervals in $[0,\alpha]$. This is a harmless simplification; even for starting configurations with same-length subintervals we know that (when $\Psi\in C^1$) after an a.s. finite time a point will be added to each interval. Once this happens all of the subintervals are of different lengths a.s. and will continue to be of different lengths a.s.

We combine all of this to define

$$\widetilde{B}^\alpha(s,u,v,x)=\ell_s(u)\mathbf{1}\{\ell_s(u)>x\}g_t^\alpha(\ell_s(u))h(v,\ell_s(y)),$$

so that $\widetilde{A}_t^\alpha(x)=\widetilde{A}_0^\alpha(x)+\sum_{(s,u,v,x)\in\Pi,s\le t}\widetilde{B}^\alpha(s,u,v,x)$.

Looking to obtain the semimartingale decomposition of $\widetilde{A}_t^\alpha(x)$ we integrate $B(t,u,v,x)$. Note that $\int_0^1 h(v,\ell,x)dv=(x/\ell)^2$. We then write

$$\begin{aligned}\int\int \widetilde{B}^\alpha(t,u,v,x)dvd\Psi(u) &= \int_0^1 \ell_t(u)\mathbf{1}\{\ell_t(u)>x\}g_t^\alpha(\ell_s(u))(x/\ell_t(u))^2 d\Psi(u)\\ &= x^2\int_0^1 \frac{1}{\ell_t(u)}\mathbf{1}\{\ell_t(u)>x\}g_t^\alpha(\ell_t(u))d\Psi(u)\\ &= x^2\int_x^\infty \frac{1}{z}g_t^\alpha(z)d\Psi(\widetilde{A}_{t^-}(z)).\end{aligned}$$

The last line follows from the fact that for a bounded Borel function, f,

$$\int_0^1 f(\ell_t(u))d\Psi(u) = \int_0^\infty f(z)d\Psi(\widetilde{A}_{t^-}(z)).$$

Recall that Ψ is assumed to be C^1, and that the indicator function g_t^α is zero unless the selected interval belongs to $[0,\alpha]$. This lets us write

$$g_t^\alpha(z)d\Psi(\widetilde{A}_{t^-}(z)) = \psi(\widetilde{A}_{t^-}(z))d\widetilde{A}_{t^-}^\alpha(z).$$

We now rewrite the integral of $\widetilde{B}_t^\alpha$ as

$$\int\int \widetilde{B}^\alpha(t,u,v,x)dvd\Psi(u) = x^2\int_x^\infty \frac{\psi(\widetilde{A}_{t^-}(z))}{z}d\widetilde{A}_{t^-}^\alpha(z).$$

Integrate this from 0 to t and we arrive at the claimed decomposition of $\widetilde{A}_t^\alpha(x)$. □

2.2. **Definitions and notation.** What follows are the essential facts and notation for understanding the proof of Theorem 1. Let non-tilde processes represent the original process scaled by e^{-t} (i.e. $A_t(x) = \widetilde{A}_t(e^{-t}x)$). In light of Proposition 3, a change of variables gives the relationship

$$\mathbf{A}^\alpha = \mathscr{C}(\mathbf{A}^\alpha, \mathbf{A}) + \mathbf{M}^\alpha, \tag{5}$$

where $\mathscr{C}\colon \mathcal{X}\times\mathcal{X}\to C([0,\infty), L^1_{\text{loc}})$ is defined by

$$\mathscr{C}(\mathbf{F},\mathbf{G})_t(x) = F_0(e^{-t}x) + \int_0^t (e^{s-t}x)^2\int_{e^{s-t}x}^\infty \frac{\psi(G_s(z))}{z}dF_s(z)ds.$$

Here $\mathcal{X} = \mathcal{B}([0,\infty),\mathcal{D})$ where $\mathcal{D} = \{F\colon [0,\infty)\to[0,1], \text{cádlág, increasing}\}$. The set $\mathcal{X}$ is a subspace of the space $\mathcal{B}([0,\infty), L^1_{\text{loc}})$ of measurable maps from $[0,\infty)$ to L^1_{loc} with the topology of locally uniform convergence, which we denote by the symbol $\xrightarrow{\mathcal{X}}$.

We say that a family of functions $(\mathbf{F}^{(n)})_{n\in\mathbb{N}}$ in $\mathcal{X}$ is *asymptotically equicontinuous* if for every compact $K\subset[0,\infty)$,

$$\lim_{\delta\to 0}\lim_{n\to\infty}\sup_{\substack{s,t\geq 0\\ |s-t|\leq\delta}}\int_K |F_s^{(n)}(x) - F_s^{(n)}(x)|dx = 0.$$

A family of distributions $(F_t)_{t\geq 0}$ is *tight* if for all $\epsilon>0$ there exists N such that $F_t(N)\geq 1-\epsilon$ for all $t\geq 0$.

We will use $\hat{F}$ and F^Ψ interchangeably to denote the a.s. pointwise limiting distribution of A_t from [MP14, Theorem 1.1]. Also define the stationary distribution $\hat{\mathbf{F}}^*$ so that $\hat{F}_t^* = \hat{F}$ for all $t\geq 0$. With the convergence $A_t\to\hat{F}$ in mind, we consider the operator

$$\mathscr{C}^*(\mathbf{F})_t = \mathscr{C}(\mathbf{F},\hat{\mathbf{F}}^*)_t = F_0(e^{-t}x) + \int_0^t (e^{s-t}x)^2\int_{e^{s-t}x}^\infty \frac{\psi(\hat{F}(z))}{z}dF_s(z)ds.$$

We will see in the proof of Theorem 1 that the limiting distribution of A_t^α belongs to the set of fixed points

$$\mathfrak{F}^\alpha = \{\mathbf{F} \in \mathcal{X}_1 \colon \mathbf{F} = \mathscr{C}^*(\mathbf{F}), F_t(+\infty) = \alpha \text{ and } (\tfrac{1}{\alpha}F_t)_{t\geq 0} \text{ tight}\}.$$

Here $\mathcal{X}_1 = \mathcal{B}([0,\infty), \{F \in \mathcal{D}\colon \|F\|_{x^{-2}} \leq 1\})$, where $\|\cdot\|_{x^{-2}}$ is the case $\delta = 1$ of the following family of norms on $L^1_{\text{loc}}([0,\infty))$:

$$\|f\|_{x^{-1-\delta}} = \int_0^\infty x^{-1-\delta}|f(x)|dx, \quad \delta \in (0,1]. \tag{6}$$

The norm used exclusively in [MP14] is $\|f\|_{x^{-2}} = \int_0^\infty x^{-2}|f(x)|dx$. This extra δ of freedom lets us prove the interpolation between min-2 and max-2 is equidistributed. The effect of working in this norm is the appearance of the $(2-\delta)$ term in (1).

We remark that $\|\cdot\|_{x^{-2}}$ does have special significance. A key property (see Proposition 5 (I)) is that $\|\widetilde{A}_t^\alpha\|_{x^{-2}} = e^{-t}N_t^\alpha$. Thus, we can recover the number of points added to the interval $[0,\alpha]$, which is the fundamental quantity for proving equidistribution.

3. Proof of Theorem 1

We delay the proofs of the following two propositions until the next section. Our goal is to make transparent the necessary ingredients for proving Theorem 1. The first proposition describes the benefit of when a Ψ-process satisfies (1).

Proposition 4. *If Ψ satisfies* (C^1) *and there exists $\delta \in (0,1]$ such that* (1) *holds for all $z \geq 0$, then*

$$\|F_t - \alpha\hat{F}\|_{x^{-1-\delta}} \leq 2(1+\delta^{-1})e^{-\delta t}$$

for all $\mathbf{F} \in \mathfrak{F}^\alpha$.

We will also need several general properties of $\mathbf{A}^\alpha$.

Proposition 5. *The following hold for any Ψ satisfying* (C^2)*:*

- *(I)* $\|A_t^\alpha\|_{x^{-2}} = e^{-t}N_t^\alpha$ *and* $\|\alpha\hat{F}\|_{x^{-2}} = \alpha$.
- *(II) The collection of distribution functions* $(\frac{1}{\alpha}A_t^\alpha)_{t\geq 0}$ *is tight.*
- *(III) The family* $(\mathbf{A}^{\alpha,(n)})$ *defined by* $A_t^{\alpha,(n)} = A_{t+n}^\alpha$ *is asymptotically equicontinuous.*
- *(IV)* $\mathbf{M}^{\alpha,(n)} \xrightarrow{\mathcal{X}} 0$ *as* $n \to \infty$, *where* $M_t^{\alpha,(n)}(x) = M_{t+n}^\alpha(x) - M_n^\alpha(e^{-t}x)$ *for every* $t \geq 0$.
- *(V) Suppose additionally that* $\sup_{z\geq 0} z\hat{F}'(z) < \infty$ *(discussion of this hypothesis appears in Lemma* 6*). Define* $\mathbf{A}^{(n)}$ *by* $A_t^{(n)} = A_{t+n}$. *If* $\mathbf{F}^{(n)} \xrightarrow{\mathcal{X}} \mathbf{F}$ *then* $\mathscr{C}(\mathbf{F}^{(n)}, \mathbf{A}^{(n)}) \xrightarrow{\mathcal{X}} \mathscr{C}^*(\mathbf{F})$.

Proof of Theorem 1. All statements are meant to hold almost surely. Also we abbreviate items from Proposition 5 as a roman numeral. In the continuous process points are added as a Poisson process with intensity $e^t dt$. So, it suffices to show $e^{-t}N_t^\alpha \to \alpha$.

By (II), (III) and the version of the Arzelá-Ascoli theorem in [MP14, Lemma 7.3] we may choose a sequence $(\mathbf{A}^{\alpha,(n_k)})$ which converges to a family of (scaled by α) distributions $\mathbf{F}^{\alpha,(\infty)}$ with $F_t^{\alpha,(\infty))}(+\infty) = \alpha$ for every $t \geq 0$. Taking limits in the formula at (5) we obtain

$$\mathscr{C}(\mathbf{A}^{\alpha,(n_k)}, \mathbf{A}^{(n_k)}) + \mathbf{M}^{\alpha,(n_k)} \xrightarrow{\mathcal{X}} \mathbf{F}^{\alpha,(\infty)}.$$

By (IV) and (V) we have

$$\mathscr{C}(\mathbf{A}^{\alpha,(n_k)}, \mathbf{A}^{(n_k)}) \xrightarrow{\mathcal{X}} \mathscr{C}^*(\mathbf{F}^{\alpha,(\infty)}).$$

Thus, $\mathbf{F}^{\alpha,(\infty)} \in \mathfrak{F}^\alpha$. Since we are assuming (1) holds, Proposition 4 implies that $\|F_t^{\alpha,(\infty)} - \alpha\hat{F}\|_{x^{-1-\delta}} \leq (2+\delta^{-1})e^{-\delta t}$. A similar argument as the conclusion of the proof of [MP14, Theorem 7.1] gives almost sure pointwise convergence $A_t^\alpha \to \alpha\hat{F}$. [MP14, Theorem 1.1] states that $A_t \to \hat{F}$ pointwise. We can then deduce from (4) that $A_t^{\alpha+} \to (1-\alpha)\hat{F}$. Combining pointwise convergence, (4) and Fatou's lemma we deduce that $\|A_t^\alpha\|_{x^{-2}} \to \|\alpha\hat{F}\|_{x^{-2}}$. Indeed,

$$\begin{aligned}\liminf \|A_t^\alpha\|_{x^{-2}} &\geq \|\alpha\hat{F}\|_{x^{-2}},\\ \limsup \|A_t^\alpha\|_{x^{-2}} &= 1 - \liminf \|A_t^{\alpha+}\|_{x^{-2}} \leq 1-(1-\alpha) = \|\alpha\hat{F}\|_{x^{-2}}.\end{aligned}$$

This finishes the proof since (I) states that $\|A_t^\alpha\|_{x^{-2}} = e^{-t}N_t^\alpha$ and $\|\alpha\hat{F}\|_{x^{-2}} = \alpha$. □

4. Proof of Proposition 4 and Proposition 5

4.1. **Proposition 4.** The proof of Proposition 4 proceeds analogously to [MP14, Lemma 4.1 and Proposition 3.4]. A significant difference is that they apply integration by parts to

$$\frac{1}{z}d\Psi(\widetilde{F}_s(z)),$$

whereas our operator $\mathscr{C}^*$ requires applying integration by parts to

$$\frac{\psi(\hat{F}(z))}{z}d\widetilde{F}_s(z).$$

The requirement at (1) arises from the extra term $\psi(\hat{F}(z))$. Also, note that we work in the norm $\|\cdot\|_{x^{-1-\delta}}$ to obtain the constant $(2-\delta)$ in (1).

Proof of Proposition 4. Let $\mathbf{F} \in \mathfrak{F}^\alpha$. We consider the rescaled processes $\widetilde{F}_t(x) = F(e^t x)$, $\widetilde{F}_t^\Psi(x) = \hat{F}(e^t x)$. It then holds that $\widetilde{\mathbf{F}} = \widetilde{\mathscr{C}}(\widetilde{\mathbf{F}})$ where

$$\widetilde{\mathscr{C}}(\widetilde{\mathbf{F}})_t(x) = \widetilde{F}_0(x) + \int_0^t e^s x^2 \int_x^\infty \frac{\psi(\hat{F}(z))}{z}d\widetilde{F}_s(z)ds.$$

Our goal is to prove the distance between $\widetilde{\mathbf{F}}$ and $\alpha\widetilde{\hat{\mathbf{F}}}^*$ is decreasing in t:

$$\partial_t\|\widetilde{F}_t - \alpha\widetilde{F}_t^\Psi\|_{x^{-1-\delta}} = \int_0^\infty x^{-1-\delta}\partial_t|\widetilde{F}_t(x) - \alpha\widetilde{F}_t^\Psi(x)|dx \leq 0. \tag{7}$$

We start by differentiating under the integral sign

$$\partial_t\widetilde{\mathscr{C}}(\widetilde{\mathbf{F}})_t(x) = e^t x^2 \int_x^\infty \frac{\psi(\hat{F}(z))}{z}d\widetilde{F}_t(z)$$

to write for each $x \geq 0$ the dynamics for the difference $\widetilde{F}_t(x) - \alpha\widetilde{F}_t^\Psi(x)$ as

$$\partial_t(\widetilde{F}_t(x) - \alpha\widetilde{F}_t^\Psi(x)) = e^t x^2 I_t(x),$$

$$I_t(x) = \int_x^\infty \frac{\psi(\hat{F}(z))}{z}\partial_z(\widetilde{F}_t(z) - \alpha\widetilde{F}_t^\Psi(z))dz.$$

Multiply both sides by $\mathrm{sgn}(\widetilde{F}_t - \alpha\widetilde{F}_t^\Psi)$ to obtain

$$e^{-t}\partial_t|\widetilde{F}_t(x) - \alpha\widetilde{F}_t^\Psi(x)| = x^2\begin{cases}\mathrm{sgn}(\widetilde{F}_t(x) - \alpha\widetilde{F}_t^\Psi(x))I_t(x), & \widetilde{F}_t(x) \neq \alpha\widetilde{F}_t^\Psi(x)\\ 0, & \widetilde{F}_t(x) = \alpha\widetilde{F}_t^\Psi(x)\end{cases}.$$

Let $\hat{f}(z) = z\psi'(\hat{F}(z))\hat{F}'(z) - \psi(\hat{F}(z))$. An application of integration by parts to the integral gives

$$I_t(x) = -\frac{\psi(\hat{F}(x))}{x}(\widetilde{F}_t(x) - \alpha\widetilde{F}_t^{\Psi}(x)) + \int_x^\infty \frac{\hat{f}(z)}{z^2}(\widetilde{F}_t(z) - \alpha\widetilde{F}_t^{\Psi}(z))dz.$$

The previous two equations therefore yield

$$e^{-t}\partial_t|\widetilde{F}_t(x) - \alpha\widetilde{F}_t^{\Psi}(x)| \leq -x\psi(\hat{F}(x))|\widetilde{F}_t(x) - \alpha\widetilde{F}_t^{\Psi}(x)| + x^2\int_x^\infty |\hat{f}(z)|\frac{|\widetilde{F}_t(z) - \alpha\widetilde{F}_t^{\Psi}(z)|}{z^2}dz.$$

We next multiply both sides by $x^{-1-\delta}$ and integrate with respect to x from 0 to infinity to obtain the bound

$$\begin{aligned} e^{-t}\int_0^\infty x^{-1-\delta}\partial_t|\widetilde{F}_t(x) - \alpha\widetilde{F}_t^{\Psi}(x)|dx \leq \int_0^\infty -\psi(\hat{F}(x))\frac{|\widetilde{F}_t(x) - \alpha\widetilde{F}_t^{\Psi}(x)|}{x^\delta}dx \\ + \int_0^\infty x^{1-\delta}\int_x^\infty |\hat{f}(z)|\frac{|\widetilde{F}_t(z) - \alpha\widetilde{F}_t^{\Psi}(z)|}{z^2}dzdx. \end{aligned}$$

An application of Fubini's theorem lets us rewrite the second integral as

$$\begin{aligned} \int_0^\infty x^{1-\delta}\int_x^\infty |\hat{f}(z)|\frac{|\widetilde{F}_t(z) - \alpha\widetilde{F}_t^{\Psi}(z)|}{z^2}dzdx &= \int_0^\infty |\hat{f}(z)|\frac{|\widetilde{F}_t(z) - \alpha\widetilde{F}_t^{\Psi}(z)|}{z^2}\int_0^z x^{1-\delta}dxdz \\ &= \int_0^\infty (2-\delta)^{-1}|\hat{f}(z)|\frac{|\widetilde{F}_t(z) - \alpha\widetilde{F}_t^{\Psi}(z)|}{z^\delta}dz. \end{aligned}$$

Hence we can combine the integrals to obtain the bound

$$e^{-t}\int_0^\infty x^{-2}\partial_t|\widetilde{F}_t(x) - \alpha\widetilde{F}_t^{\Psi}(x)|dx \leq \int_0^\infty \Big((2-\delta)^{-1}|\hat{f}(z)| - \psi(\hat{F}(z))\Big)\frac{|\widetilde{F}_t(z) - \alpha\widetilde{F}_t^{\Psi}(z)|}{z^\delta}dz.$$

Our hypothesis (1) guarantees that the term inside the integral:

$$(2-\delta)^{-1}|\hat{f}(z)| - \psi(\hat{F}(z)) \leq 0.$$

Therefore (7) holds. This establishes that

$$\|\widetilde{F}_t - \alpha\widetilde{F}_t^{\Psi}\|_{x^{-1-\delta}} \leq \|\widetilde{F}_0 - \alpha\widetilde{F}_0^{\Psi}\|_{x^{-1-\delta}} = \|F_0 - \alpha\hat{F}\|_{x^{-1-\delta}}. \tag{8}$$

A change of variables $x = e^{-t}z$ gives

$$\begin{aligned} \|F_t - \alpha\hat{F}\|_{x^{-1-\delta}} &= \int_0^\infty x^{-1-\delta}|F_t(x) - \alpha\hat{F}(x)|dx \\ &= e^{-\delta t}\int_0^\infty z^{-1-\delta}|\widetilde{F}_t(z) - \alpha\widetilde{F}_t^{\Psi}(z)|dz \\ &= e^{-\delta t}\|\widetilde{F}_t - \alpha\widetilde{F}_t^{\Psi}\|_{x^{-1-\delta}} \\ &\leq e^{-\delta t}\|F_0 - \alpha\hat{F}\|_{x^{-1-\delta}}, \end{aligned} \tag{9}$$

where at the last line we apply (8).

It remains to prove that $\|F_0 - \alpha\hat{F}\|_{x^{-1-\delta}} \leq C$, for some $C > 0$. By assumption, $\mathbf{F} \in \mathcal{X}_1$ and therefore $\|F_0\|_{x^{-2}} \leq 1$. As $0 \leq F_0(x) \leq 1$ we can break up the integral and use integrability of $x^{-1-\delta}\mathbf{1}\{x > 1\}$:

$$\int_0^\infty x^{-1-\delta}F_0(x)dx \leq \int_0^1 x^{-2}F_0(x)dx + \int_1^\infty x^{-1-\delta}dx \leq \|F_0\|_{x^{-2}} + \delta^{-1} \leq 1 + \delta^{-1}.$$

Similarly, $\|\alpha\hat{F}\|_{x^{-1-\delta}} \leq 1+\delta^{-1}$. Apply the triangle inequality to conclude $\|F_0-\alpha\hat{F}\|_{x^{-1-\delta}} \leq \|F_0\|_{x^{-1-\delta}} + \|\alpha\hat{F}\|_{x^{-1-\delta}} \leq 2(1+\delta^{-1})$. □

4.2. **Proposition 5.** In Proposition 5 we prove that A_t^α and A_t have similar properties. Each statement requires some manipulation. Fortunately [MP14] contains much of the heavy-lifting. We make one remark concerning the proof of (V). In [MP14] they prove continuity of an operator $\mathscr{S}^\Psi$ with domain $\mathcal{X}$. Our operator $\mathscr{C}$ has domain $\mathcal{X}\times\mathcal{X}$. This makes the proof more involved, and also restricts us to proving continuity in sequences of the form $(\mathbf{F}^{(n)}, \mathbf{A}^{(n)})$.

Proof of (I). The equality $\|\alpha\hat{F}\|_{x^{-2}} = \alpha$ is [MP14, Lemma 3.5]. For the other equality, take $I_j^{\alpha,(t)}$ to be the length of an interval in $[0,\alpha]$. Define the measure $\mu_t^\alpha = e^{-t}\sum_1^{N_t^\alpha} \delta_{e^t I_j^{\alpha,(t)}}$. This gives μ_t^α is the empirical distribution of rescaled interval lengths. We can then write

$$A_t^\alpha(x) = \int_0^x y\mu_t(dy).$$

Applying Fubini's theorem shows that

$$\|A_t^\alpha\|_{x^{-2}} = \int_0^\infty x^{-2}\int_0^x y\mu_t^\alpha(dy)dx = \int_0^\infty \mu_t^\alpha(dy) = e^{-t}N_t^\alpha.$$

□

Proof of (II). Recall that a family of distributions $(F_t)_{t\geq 0}$ is *tight* if for all $\epsilon > 0$ there exists N such that $F_t(N) \geq 1-\epsilon$ for all $t \geq 0$. [MP14, Proposition 6.3] implies $(A_t)_{t\geq 0}$ is tight. Fix $\epsilon > 0$ and let N be such that $A_t(N) \geq 1-\alpha\epsilon$ for all $t \geq 0$. The relationship at (4) ensures $A_t^\alpha(N) + A_t^{\alpha+}(N) \geq 1-\alpha\epsilon$. As $A_t^\alpha \leq \alpha$ and $A_t^{\alpha+} \leq 1-\alpha$, this inequality could only hold if $A_t^\alpha(N) \geq \alpha - \alpha\epsilon$ for all $t \geq 0$. Hence, $(\frac{1}{\alpha}A_t^\alpha)_{t\geq 0}$ is tight. □

Proof of (III). Recall, that a family of functions $(\mathbf{F}^{(n)})_{n\in\mathbb{N}}$ in $\mathcal{X}$ is asymptotically equicontinuous if for every compact $K \subset [0,\infty)$,

$$\lim_{\delta\to 0}\lim_{n\to\infty}\sup_{\substack{s,t\geq 0\\ |s-t|\leq\delta}}\int_K |F_s^{(n)}(x) - F_s^{(n)}(x)|dx = 0.$$

The proof is similar to [MP14, Lemma 7.5]. The idea is that it suffices to show the existence of a $\delta_0 > 0$ and constant C so that for every $0 < \delta_1 < \delta_0$ there exists almost surely a $T_{\delta_1} < \infty$ so that

$$\sup_{t\geq T_{\delta_1}, 0\leq\delta\leq\delta_1}\int_0^\infty \frac{|A_{t+\delta}^\alpha(x) - A_t^\alpha(x)|}{x^2}dx \leq C\delta_1. \tag{10}$$

This is sufficient since we for any $\delta_1 > 0$ and any $M > 0$, almost surely

$$\lim_{n\to\infty}\sup_{s,t\geq 0,|s-t|\leq\delta_1}\int_0^M |A_s^{\alpha,(n)}(x) - A_t^{\alpha,(n)}(x)|dx \leq \sup_{t\geq T_{\delta_1},0\leq\delta\leq\delta_1}\int_0^\infty \frac{|A_{t+\delta}^\alpha(x) - A_t^\alpha(x)|}{x^2}dx$$
$$\leq M^2 C\delta_1.$$

As this holds jointly with probability 1 for a countable sequence of δ_1 going to 0 and $M \in \mathbb{N}$, the asymptotic equicontinuity of $(A^{(n)})_{n\geq 0}$ follows.

The formula at (10) follows from the fact that $\widetilde{A}_t^\alpha$ satisfies the monotonicity condition, for any $\delta > 0$,

$$\widetilde{A}_t^\alpha(x) \leq \widetilde{A}_{t+\delta}^\alpha(e^{-\delta}x) \leq \widetilde{A}_{t+\delta}^\alpha(x). \tag{11}$$

Another necessary fact is that number of points kept in $[0,\alpha]$ from time t to $t+\delta$ is bounded by the number of points added to $[0,1]$ in that same time interval. Formally, for any $\delta > 0$ we have $N^\alpha_{t+\delta} - N^\alpha_t \le N^1_{t+\delta} - N^1_t$. This lets us deduce the equivalent for N^α_t as for N_t in [MP14, Lemma 7.6]. Namely, that there is a $\delta > 0$ so that for every $0 < \delta < \delta_0$ there exists almost surely a $T_\delta < \infty$ so that

$$\sup_{t \ge T_\delta} N^\alpha_{t+\delta} - N^\alpha_t \le 2\delta e^t.$$

The argument finishes by using the formula from Proposition 5 (I) for N^α_t in terms of $\|A^\alpha_t\|_{x^{-2}}$. See the proof of [MP14, Lemma 7.5] for further details. □

Proof of (IV). The proof is similar to the decay of the noise subsection in [MP14, Section 7]. The idea is to bound the martingale $\mathbf{M}^\alpha$ by computing various moments of the underlying process $\mathbf{B}^\alpha$. We can use the same bounds as in [MP14] because points are added to $[0,\alpha]$ no faster than to $[0,1]$. This ensures that $B^\alpha(s,u,v,x) \le B(s,u,v,x)$. Here $B(s,u,v,x)$ is the function defined at [MP14, (3)]. □

Proof of (V). Suppose that $\mathbf{F}^{(n)} \xrightarrow{\mathcal{X}} \mathbf{F}$. An equivalent notion of convergence in the topology of local uniform convergence is that $\mathbf{F}^{(n)} \xrightarrow{\mathcal{X}} \mathbf{F}$ if and only if for all compact $K \subset [0,\infty)$

$$\lim_{n\to\infty} \sup_{0 \le s \le t} \int_K |F^{(n)}_s(x) - F_s(x)| dx = 0.$$

[MP14, Theorem 7.1] implies $\mathbf{A}^{(n)} \xrightarrow{\mathcal{X}} \mathbf{F}^*$. Thus it suffices to prove for any fixed $T > 0$ and $K > 0$

$$\int_0^K |\mathscr{C}(\mathbf{F},\mathbf{F}^*)_t(x) - \mathscr{C}(\mathbf{F}^{(n)},\mathbf{A}^{(n)})_t(x)| dx \to 0 \tag{12}$$

uniformly for $t \le T$. For fixed n we can write

$$\mathscr{C}(\mathbf{F}^{(n)},\mathbf{A}^{(n)})_t(x) = F^{(n)}_0(x) + \int_0^t (e^{s-t}x)^2 \int_{e^{s-t}x}^\infty \frac{\psi(A^{(n)}_s(z))}{z} dF^{(n)}_s(z) ds.$$

If we write $\psi(A^{(n)}_s(z)) = \psi(\hat{F}(z)) + \psi(A^{(n)}_s(z)) - \psi(\hat{F}(z))$ the above becomes

$$\mathscr{C}(\mathbf{F}^{(n)},\mathbf{A}^{(n)})_t(x) = \mathscr{C}(\mathbf{F}^{(n)},\mathbf{F}^*)_t(x) + \int_0^t (e^{s-t}x)^2 \int_{e^{s-t}x}^\infty \frac{\psi(A^{(n)}_s(z)) - \psi(\hat{F}(z))}{z} dF^{(n)}_s(z) ds.$$

We can then bound the left side of (12) by

$$\int_0^K |\mathscr{C}(\mathbf{F},\mathbf{F}^*)_t(x) - \mathscr{C}(\mathbf{F}^{(n)},\mathbf{F}^*)_t(x)| dx \tag{13}$$

$$+ \int_0^K \int_0^t (e^{s-t}x)^2 \int_{e^{s-t}x}^\infty \frac{|\psi(A^{(n)}_s(z)) - \psi(\hat{F}(z))|}{z} dF^{(n)}_s(z) ds dx. \tag{14}$$

It suffices to show that as $n \to \infty$ each summand converges to zero uniformly for $t \le T$.

First summand. Start by bounding the summand at (13) by

$$\int_0^K |F_0(e^{-t}x) - F_0^{(n)}(e^{-t}x)|dx + \int_0^K \int_0^t (e^{s-t}x)^2 \left| \int_{e^{s-t}x}^\infty \frac{\psi(\hat{F}(z))}{z} d(F_s(z) - F_s^{(n)}(z)) \right| ds dx.$$

The first quantity goes to zero uniformly for $t \leq T$ by the definition of $\mathbf{F}^{(n)} \xrightarrow{\mathcal{X}} \mathbf{F}$ since a change of variables gives

$$\int_0^K |F_0(e^{-t}x) - F_0^{(n)}(e^{-t}x)|dx \leq e^t \int_0^K |F_0(x) - F_0^{(n)}(x)|dx.$$

Expand the interior of the second quantity with integration by parts and take the absolute value signs inside to bound it by

$$\underbrace{\frac{\psi(\hat{F}(e^{s-t}x))}{e^{s-t}x} |F_s(e^{s-t}x)dx - F_s^{(n)}(e^{s-t}x)|}_{\text{term one}} + \underbrace{\int_{e^{s-t}x}^\infty \left| \frac{d}{dz} \frac{\psi(\hat{F}(z))}{z} \right| |F_s(z) - F_s^{(n)}(z)| dz dx}_{\text{term two}}.$$

Multiply term one by $(e^{s-t}x)^2$ and integrate so it becomes

$$\int_0^K \int_0^t (e^{s-t}x)\psi(\hat{F}(e^{s-t}x))|F_s(e^{s-t}x) - F_s^{(n)}(e^{s-t}x)| ds dx.$$

Since $\hat{F}$ is a distribution function and ψ is continuous we have $(\psi \circ \hat{F})(u) \leq \sup_{u \in [0,1]} \hat{\psi}(u) < D < \infty$ for some constant D. Thus, the above is bounded by

$$D \int_0^K \int_0^t (e^{s-t}x)|F_s(e^{s-t}x) - F_s^{(n)}(e^{s-t}x)|dx.$$

The above goes to zero by the definition of $\mathbf{F}^{(n)} \xrightarrow{\mathcal{X}} \mathbf{F}$. As for term two, we differentiate to rewrite it as

$$\int_{e^{s-t}x}^\infty \frac{|z\psi'(\hat{F}(z))\hat{F}'(z) - \psi(\hat{F}(z))|}{z^2} |F_s(z) - F_s^{(n)}(z)| dz. \tag{15}$$

Our additional hypothesis is that $z\hat{F}'(z)$ is bounded. Since the range of $\hat{F}$ is contained in the compact interval $[0,1]$ and $\Psi \in C^2$ we have $\psi \circ \hat{F}$ and $\psi' \circ \hat{F}$ are also bounded. Therefore, $C = \sup_{0 \leq z \leq \infty} |z\hat{F}'(z)\psi'(\hat{F}(z)) - \psi(\hat{F}(z))| < \infty$. It follows that (15) is less than

$$C \int_{e^{s-t}x}^\infty \frac{1}{z^2} |F_s(z) - F_s^{(n)}(z)| dz. \tag{16}$$

Finally we are in the position of I_2 from [MP14, Lemma 3.3] and can conclude that (16) goes to zero uniformly for $t \leq T$.

Second summand. Fix $M > 0$ and for any function $f : [0, \infty) \to [0, 1]$ define $f^M = f|_{[0,M]}$ to be the restriction to the domain $[0, M]$. We have in [MP14, Theorem 7.1] that A^M converges pointwise to $\hat{F}^M$. Observe that each A_t^M is an increasing function with compact domain, and $\hat{F}^M$ is continuous by [MP14, Lemma 3.5]. Together these imply (see [Rud76, exercise 7.13]) that for any $\epsilon > 0$ there exists t_ϵ such that for all $z \in [0, M]$

$$\sup_{t \geq t_\epsilon} |A_t^M(z) - \hat{F}_t^M(z)| < \epsilon.$$

Because the functions $A_t^{(n)}$ are translates of A_t it follows that for all $n > t_\epsilon$ we have

$$\sup_{t\geq 0} |A_t^{(n),M}(z) - \hat{F}_t^M(z)| \leq \sup_{t\geq t_\epsilon} |A_t^M(z) - \hat{F}_t^M(z)| < \epsilon.$$

As the functions $A_t^{(n)}$ and $\hat{F}$ are supported on $[0,1]$, we have their compositions with ψ are uniformly continuous. We conclude that there exists n_0 such that for all $z \in [0,M]$

$$\sup_{t\geq 0} |\psi(A_t^{(n)}(z)) - \psi(\hat{F}(z))| < \epsilon, \qquad \text{for } n \geq n_0. \tag{17}$$

We truncate the integral then apply (17) to bound the absolute value of (14) by

$$\epsilon \int_0^K \int_0^t (e^{s-t}x)^2 \int_{e^{s-t}x}^M \frac{1}{z} dF_s^{(n)}(z) ds dx \tag{18}$$

$$+ \int_0^K \int_0^t (e^{s-t}x)^2 \int_M^\infty \frac{|\psi(A_s^{(n)}(z)) - \psi(\hat{F}(z))|}{z} dF_s^{(n)}(z) ds dx. \tag{19}$$

We can use the fact that $F_s^{(n)}(z) \leq 1$ and bound the inside integral of (18) by

$$\frac{1}{e^{s-tx}} \int_{e^{s-t}x}^M dF_s^{(n)}(z) \leq \frac{2}{e^{s-t}x}.$$

Thus (18) is bounded by

$$\epsilon \int_0^K \int_0^t 2e^{s-t} x ds dx \leq \epsilon(1-e^{-t})K^2 \leq \epsilon K^2.$$

As K is fixed, this can be made arbitrarily small.

Lastly we consider (19). Since $\sup_{u\geq 0} \psi(u) = D < \infty$ we use similar estimates as in (18) and start with the bound

$$\begin{aligned} \int_0^K \int_0^t (e^{s-t}x)^2 \int_M^\infty \frac{|\psi(A_s^{(n)}(z)) - \psi(\hat{F}(z))|}{z} dF_s^{(n)}(z) ds dx \\ \leq 4D \int_0^K \int_0^t (e^{s-t}x)^2 \frac{1}{M} ds dx \\ \leq \frac{4DK^3(1-e^{-2t})}{6M}. \end{aligned}$$

Since M can be made arbitrarily large, this can be made as small as we like. Therefore, the absolute value of (14) can be bounded by any $\epsilon > 0$ uniformly for $t \leq T$. □

Lemma 6. *If* Ψ *satisfies* (C^2) *and either* $\psi(1) > 0$ *or* $\Psi(u) = 1-(1-u)^k$ *for some positive integer* k *then* $\sup_{z\geq 0} z\hat{F}'(z) < \infty$.

Proof. [MP14, Proposition 8.2] states that when $\psi(1) > 0$ it holds that $\hat{F}'(x) \leq Ce^{-ax}$ for some constants $C, a > 0$. Additionally, for the min-k process ($\Psi(u) = 1-(1-u)^k$) it is shown in [MP14, Proposition 8.4] that $\hat{F}'(x) \leq C_k x^{-1-\epsilon_k}$ for some $C_k, \epsilon_k > 0$. Note that $\sup_{k\geq 0} C_k < \infty$ and $\epsilon_k \to 0$. □

Corollary 7. *From Lemma* 6 $z\hat{F}'(z)$ *is bounded for all interpolations of the* max-k *and* min-k *processes.*

We remark that it appears boundedness of $z\hat{F}'(z)$ does not necessarily hold for general Ψ. At the very least it does not obviously follow from (2) or (3).

5. Proving Corollary 2

For this entire section we will let F denote F^Ψ. To establish (1), we rely almost entirely on (2) and (3). For convenience we rerecord them here:

$$|z\psi'(F(z))F'(z) - \psi(F(z))| \le (2-\delta)\psi(F(z)), \qquad \delta \in (0,1], \tag{1}$$

$$F'(z) = z\int_z^\infty \frac{1}{y}\,d\Psi(F(y)), \tag{2}$$

$$zF''(z) - F'(z) + z\psi(F(z))F'(z) = 0. \tag{3}$$

We start with the proof of Corollary 2. It follows from a sequence of lemmas.

Proof of Corollary 2. First off we need the conclusion of Corollary 7 to guarantee Proposition 5 (V) holds for the interpolations we consider. Equidistribution for the max-2 process then follows from Lemma 8 by taking $p_2 = 1$. The fact that the interpolation that is 60%-min-2 satisfies (1) follows by taking $p_{-2} = .6$ in Lemma 11. Part three (for general interpolations) follows from Lemma 15. □

Now we give the proofs of the necessary lemmas. We break this up into two sections: one for interpolations of max-2 and min-2 processes and the other for general interpolations.

5.1. **Interpolations of** min-2 **and** max-2**.** Fix $p_{-2}, p_2 \in [0,1]$ with $p_2 + p_{-2} = 1$. We will work exclusively in this subsection with Ψ that are interpolations of the min-2 and max-2 process. Thus,

$$\begin{aligned}
\Psi(u) &= p_2u^2 + p_{-2}(1-(1-u)^2),\\
\psi(u) &= 2p_2u + 2p_{-2}(1-u),\\
\psi'(u) &= 2p_2 - 2p_{-2},
\end{aligned}$$

This is the distribution function (and derivatives) for an interpolation where at each step we add a point according the min-2 process with probability p_{-2} and according to the max-2 process with probability p_2.

Our first lemma establishes (1) holds so long as $p_{-2} \le p_2$. Note that the case $p_2 = 1$ is the max-2 process.

Lemma 8. *If* $p_{-2} \le p_2$ *then* (1) *holds.*

Proof. Dropping the constant $2-\delta$ from the right side of (1) it suffices to prove that

$$|\psi(F(z)) - z\psi'(F(z))F'(z)| \le \psi(F(z)). \tag{20}$$

We break into two cases:

- First suppose $\psi(F(z)) \ge z\psi'(F(z))F'(z)$ so that (20) reduces to proving that
$$-z\psi'(F(z))F'(z) \le 0.$$
As F is increasing we know $F'(z) \ge 0$. The hypothesis $p_{-2} \le p_2$ guarantees that $\psi'(F(z)) \ge 0$. Thus, the inequality is satisfied.
- Next, suppose $\psi(F(z)) \le z\psi'(F(z))F'(z)$. Rearranging (20) we seek to show
$$2(p_2 - p_{-2})zF'(z) \le 2\psi(F(z)).$$

Note that both sides are zero at $z = 0$. By the fundamental theorem of calculus it then suffices to prove the above inequality holds for the derivatives. Differentiating and again using the fact that $\psi'(F(z)) = 2(p_2 - p_{-2})$ reduces the problem to establishing

$$2(p_2 - p_{-2})(zF''(z) + F'(z)) \leq 4(p_2 - p_{-2})F'(z).$$

After some algebra this is equivalent to

$$zF''(z) \leq F'(z). \tag{21}$$

From (3) we know that $zF''(z) = F'(z) - z\psi(F(z))F'(z)$. Substitute this into (21) and we have a sufficient condition is that

$$F'(z) - 2z\psi(F(z))F'(z) \leq F'(z).$$

This holds as $F'(z)$ and $\psi(F(z))$ are nonnegative.

□

To prove (1) holds when $p_{-2} > p_2$ requires a different analysis of the differential equation at (3). Lemma 10 shows $zF'(z)$ can be bounded in terms of p_2.

Lemma 9. *If $p_{-2} > p_2$ then $\lim_{\epsilon \to 0} F'(\epsilon)/\epsilon \leq 2$.*

Proof. Starting from the formula at (2) then integrating by parts gives

$$\lim_{\epsilon \to 0} \frac{F'(\epsilon)}{\epsilon} = \int_0^\infty \frac{1}{y} d\Psi(F(y)) = \|\Psi \circ F\|_{x^{-2}}. \tag{22}$$

Plugging into Ψ we have

$$\begin{aligned} \Psi(F(y)) &= p_2 F(y)^2 + p_{-2}(1 - (1 - F(y))^2) \\ &= F(y)[(p_2 - p_{-2})F(y) + 2p_{-2}]. \end{aligned} \tag{23}$$

The hypothesis $p_2 < p_{-2}$ means an upper bound for the above is

$$\Psi(F(y)) \leq 2p_{-2}F(y) \leq 2F(y). \tag{24}$$

Proposition 5 (I) implies that $\|F\|_{x^{-2}} = 1$. It follows from (22) and (24) that

$$\lim_{\epsilon \to 0} \frac{F'(\epsilon)}{\epsilon} \leq 2\|F\|_{x^{-2}} = 2.$$

□

Lemma 10. *It $p_{-2} > p_2$ then*

$$zF'(z) \leq 2(p_2 e)^{-2}.$$

Proof. Integrate (3) as in [MP14, Proposition 8.1] so that for any $\epsilon > 0$

$$F'(z) = \frac{F'(\epsilon)}{\epsilon} z \exp\left(-\int_\epsilon^z \psi(F(y))dy\right).$$

Taking $\epsilon \to 0$ and applying Lemma 9 gives

$$F'(z) \leq 2z \exp\left(-\int_0^z \psi(F(y))dy\right). \tag{25}$$

We observe that $\psi(F(y)) = 2p_2F(y) + 2p_{-2}(1 - F(y))$. Since we are assuming $p_{-2} > p_2$ and know that $F(y) \leq 1$ we obtain a lower bound by evaluating at $\psi(1)$:

$$\psi(F(y)) \geq \psi(1) = 2p_2. \tag{26}$$

Applying this to (25) and multiplying by z gives

$$zF'(z) \leq 2z^2 e^{-2p_2 z}.$$

The maximum of $z^2 e^{-2p_2 z}$ is at $z = 1/p_2$. Plug this in above to obtain the claimed bound. □

Lemma 11. *If $p_2 < p_{-2} \leq .6$ then* (1) *holds.*

Proof. Using the triangle inequality on the left side of (1) it suffices to find δ such that for all $z \geq 0$

$$|z\psi'(F(z))F'(z)| + |\psi(F(z))| \leq (2-\delta)\psi(F(z)). \tag{27}$$

Because F is a distribution function, we know that $F' \geq 0$. Also, note that

$$(2-\delta)\psi(F(z)) - |\psi(F(z))| \leq (1-\delta)\psi(F(z)).$$

Thus, to establish (27) it is enough to prove

$$zF'(z) \leq \frac{(1-\delta)\psi(F(z))}{|\psi'(F(z))|}, \qquad \text{for } z \geq 0.$$

We have from (26) that $\psi(u) \geq 2p_2$ and can compute $|\psi'(u)| = 2|p_2 - p_{-2}|$. It then suffices to prove

$$zF'(z) \leq \frac{p_2(1-\delta)}{|p_2 - p_{-2}|}.$$

By Lemma 10 it suffices to choose δ, p_{-2} and p_2 so that

$$2\,(p_2 e)^{-2} \leq \frac{p_2(1-\delta)}{|p_2 - p_{-2}|}.$$

Combining with our hypotheses we have the following system of constraints

$$\begin{aligned} 2e^{-2}|p_2 - p_{-2}| &\leq (1-\delta)(p_2)^3, \\ p_2 + p_{-2} &= 1, \\ p_2 &< p_{-2}, \\ 0 < \delta &\leq 1. \end{aligned}$$

Take $\delta \to 0$ and use the fact that p_{-2} is assumed to be larger than p_2, and the solution must be strictly smaller than the real root of the cubic

$$\frac{2}{e^2}(p_{-2} - (1-p_{-2})) = (1-p_{-2})^3.$$

This is approximately .61, thus $p_{-2} \leq .6$ lies in the solution set. □

Remark 12. The bound $p_{-2} \leq .6$ could be optimized further in the preceding lemmas, but the gain would be marginal. Something like $p_{-2} \leq .68$ is the best that comes out of optimizing our argument. We sacrifice this marginal gain for the sake of clarity.

5.2. **General interpolations.** We will reprove versions of the previous three lemmas for more general interpolations. Let $\mathbf{p} = (p_k)_{k\neq -1,0}$ be a probability measure on $\mathbf{Z} \setminus \{-1, 0\}$. In this subsection we consider the interpolations

$$\Psi(u) = p_1 u + \textstyle\sum_{k\geq 2} p_k u^k + p_{-k}(1-(1-u)^k).$$

Define $C_{\mathbf{p}} = \sum_{k\geq 2} k(k-1)(p_k + p_{-k})$. This constant arises because $\sup_{u\geq 0} |\psi'(u)| \leq C_{\mathbf{p}}$. First we give a bound on F' that holds for any Ψ-process.

Lemma 13. *Let Ψ satisfy (C) and (D). For all $z \geq 0$ it holds that $F'(z) \leq 1$.*

Proof. This follows from a simple bound on (2):

$$\begin{aligned} F'(z) = z\int_z^\infty \frac{\psi(F(y))}{y} F'(y)dy &\leq z \cdot \frac{1}{z}\int_z^\infty \psi(F(y))F'(y)dy \\ &= \Psi(1) - \Psi(F(z)). \end{aligned}$$

Since $\Psi(1) = 1$ we conclude that $F'(z) \leq 1$. □

Now let us return to the setting where Ψ is an interpolation of max-k, uniform and min-k processes given by $\mathbf{p}$.

Lemma 14. *Suppose that $p_1 > 0$. It holds that*

$$zF'(z) \leq \frac{2e^{-1}}{(p_1)^2}.$$

Proof. Integrate (3) as in [MP14, Proposition 8.1] so that for any $\epsilon > 0$

$$F'(z) = \frac{F'(\epsilon)}{\epsilon} z \exp\left(-\int_\epsilon^z \psi(F(y))dy\right).$$

Taking $\epsilon = 1$ and applying Lemma 13 gives

$$F'(z) \leq z \exp\left(-\int_1^z \psi(F(y))dy\right). \tag{28}$$

Notice that

$$\psi(u) = p_1 + \textstyle\sum_{k\geq 2} k[p_k u^{k-1} + p_{-k}(1-u)^{k-1}] \geq p_1. \tag{29}$$

Apply this to (28) then multiply by z to obtain the bound

$$zF'(z) \leq e^{p_1} z^2 e^{-p_1 z}$$

The maximum of $z^2 e^{-p_1 z}$ is at $z = 2/p_1$. Plug this in above to obtain the claimed bound. □

Lemma 15. *If $C_{\mathbf{p}} \leq \frac{1}{2}$ then (1) holds.*

Proof. As in Lemma 11 it suffices to show for some $\delta \in (0, 1]$ and all $z \geq 0$

$$zF'(z) \leq \frac{(1-\delta)\psi(F(z))}{|\psi'(F(z))|}.$$

We have from (29) that $\psi(u) \geq p_1$ and can compute

$$|\psi'(u)| \leq \sum_{k\geq 2} k(k-1)|u^{k-2} - (1-u)^{k-2}| \leq C_{\mathbf{p}}.$$

It then suffices to prove

$$zF'(z) \leq \frac{(1-\delta)p_1}{C_{\mathbf{p}}}. \tag{30}$$

By Lemma 14 and the hypothesis $C_{\mathbf{p}} \leq 1/2$ it suffices to choose the p_k so that

$$\frac{2e^{-1}}{(p_1)^2} \leq 2(1-\delta)p_1.$$

Rewriting and letting $\delta \to 0$ we require that $e^{-1/3} < p_1$. It is easy to verify (by just checking the case $p_k = 0$ for $k \neq 1, 2$) that we must have $\sum_{k \neq 1} p_k < 1/4$ in order to satisfy $C_{\mathbf{p}} < 1/2$. Thus, $p_1 > 3/4$. Since $e^{-1/3} \approx .71 < 3/4 = p_1$ the above displayed inequality holds. □

References

[ABKU99] Yossi Azar, Andrei Z. Broder, Anna R. Karlin, and Eli Upfal, *Balanced allocations*, SIAM J. Comput. **29** (1999), no. 1, 180–200.

[BM72] J. R. Blum and V. J. Mizel, *A generalized weyl equidistribution theorem for operators, with applications*, Transactions of the American Mathematical Society **165** (1972), no. 2, 291–307.

[CKK+07] Jacek Cicho, Marek Klonowski, ukasz Krzywiecki, Bartomiej Raski, and Pawe Zieliski, *Random subsets of the interval and p2p protocols*, Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (Moses Charikar, Klaus Jansen, Omer Reingold, and JosD.P. Rolim, eds.), Lecture Notes in Computer Science, vol. 4627, Springer Berlin Heidelberg, 2007, pp. 409–421 (English).

[Dar53] D. A. Darling, *On a class of problems related to the random division of an interval*, The Annals of Mathematical Statistics **24** (1953), no. 2, 239–253.

[FSZ09] Kevin Ford, K. Soundararajan, and Alexandru Zaharescu, *On the distribution of imaginary parts of zeros of the riemann zeta function, ii*, Mathematische Annalen **343** (2009), no. 3, 487–505 (English).

[Kak76] S. Kakutani, *A problem of equidistribution on the unit interval [0, 1]*, Measure Theory (Alexandra Bellow and Dietrich Klzow, eds.), Lecture Notes in Mathematics, vol. 541, Springer Berlin Heidelberg, 1976, pp. 369–375 (English).

[KN06] L. Kuipers and H. Niederreiter, *Uniform distribution of sequences*, Dover Books on Mathematics, Dover Publications, 2006.

[LM05] Malwina J. Luczak and Colin McDiarmid, *On the power of two choices: Balls and bins in continuous time*, The Annals of Applied Probability **15** (2005), no. 3, 1733–1764.

[Loo78] J.-C. Lootgieter, *Sur la rpartition des suites de Kakutani (ii)*, Annales de l'institut Henri Poincar (B) Probabilits et Statistiques **14** (1978), no. 3, 279–302 (French).

[MP14] P. Maillard and E. Paquette, *Choices and intervals*, ArXiv e-prints (2014), To appear in Israel Journal of Mathematics.

[MRS00] Michael Mitzenmacher, Andra W. Richa, and Ramesh Sitaraman, *The power of two random choices: A survey of techniques and results*, in Handbook of Randomized Computing, Kluwer, 2000, pp. 255–312.

[Pyk80] Ronald Pyke, *The asymptotic behavior of spacings under Kakutani's model for interval subdivision*, The Annals of Probability **8** (1980), no. 1, 157–163.

[Rud76] W. Rudin, *Principles of mathematical analysis*, third ed., McGraw-Hill, New York, 1976.

[Vau77] R. C. Vaughan, *On the distribution of p modulo 1*, Mathematika **24** (1977), 135–141.

[Wey10] Hermann Weyl, *Über die gibbs'sche erscheinung und verwandte konvergenzphänomene*, Rendiconti del Circolo Matematico di Palermo **30** (1910), no. 1, 377–407 (Italian).

[Zwe78] W. R. Van Zwet, *A proof of Kakutani's conjecture on random subdivision of longest intervals*, The Annals of Probability **6** (1978), no. 1, pp. 133–137 (English).

Department of Mathematics, University of Washington
E-mail address: jungem@math.washington.edu