MULTIPLE PARTICLE EDGE REINFORCED RANDOM WALKS ON $\mathbb{Z}$

ELIZABETH DEYOUNG AND JONATHAN HANSELMAN

ADVISOR: YEVGENIY KOVCHEGOV
OREGON STATE UNIVERSITY

ABSTRACT. It is known that single particle edge reinforced random walks (ERRWs) on certain acyclic graphs are recurrent, i.e. they return to the starting point infinitely often almost surely. It has also been shown that a two particle ERRW on $\mathbb{Z}$ is recurrent, i.e. the particles meet infinitely often almost surely. We extend this result and prove that for any number of particles k, the k particle ERRW on $\mathbb{Z}$ is recurrent. Although the multiparticle ERRW is a non-exchangeable process, we couple it to a similar exchangeable process, allowing us to use familiar results and methods to complete the proof. We extend our result to the infinite binary tree and show that a multi particle ERRW is recurrent for the range of initial weights in which the single particle ERRW is known to be recurrent.

1. INTRODUCTION

The concept of an edge reinforced random walk (ERRW) was first introduced by Coppersmith and Diaconis in 1987 [1]. A particle walks on a weighted graph G, where each edge has initial weight a. The particle moves randomly about the graph, choosing at each step one of the edges leaving from the particle's position with probabilities proportional to the weights of the edges. Whenever the particle traverses an edge, the weight of that edge is increased by one, reinforcing the likelihood of the particle traversing that edge again in the future.

1.1. **Acyclic ERRWs.**

It is well known that for acyclic graphs, the ERRW can be modeled using Polya's urns [1]. At each vertex, we place an urn containing several colored marbles, with each color corresponding to an edge connected to that vertex. For the vertex at which the particle starts, we initially have a marbles of each color, corresponding to the initial weight of a on each edge. The first move is chosen by randomly drawing a marble from the urn, and moving along the edge corresponding to the resulting color. This is equivalent to choosing based on edge weights. Once a marble is drawn, it is put back into the urn, along with two more marbles of the same color. Because the graph is acyclic, if the particle ever returns to that vertex, it will do so by retraversing that same edge. Thus when the next decision needs to be made at that vertex, the weight of the chosen edge will have increased by exactly two, and a drawing from the urn can again be used to determine the particle's next move with the correct probabilities. All other vertices behave in the same way, except that the urns start out with a marbles of each color, plus one extra marble of the color corresponding to the edge that goes toward the particle's starting position: the first time the particle arrives at that vertex, it will do so by that edge, and the weight will already be $a+1$.

Date: August 14, 2008.
This work was done during the Summer 2008 REU program in Mathematics at Oregon State University.

Perhaps the simplest example of this model is when the graph is the integer line, $\mathbb{Z}$. A particle starts at 0, and at each step can move either one unit left or one unit right, based on the weights of the edges (which connect adjacent integers). Assume for this example that $a = 1$. We can describe the system by a Polya urn at each integer containing red and blue marbles. Blue corresponds to moving right and red corresponds to moving left. The urn at the origin initially has one blue and one red marble, while any urn to the right of the origin starts with two red, one blue, and any urn to the left of the origin starts with two blue, one red (see Figure 1.1). At each step the random walker draws a marble from the appropriate urn to decide which direction to go, and then adds two more of the same color to the urn.

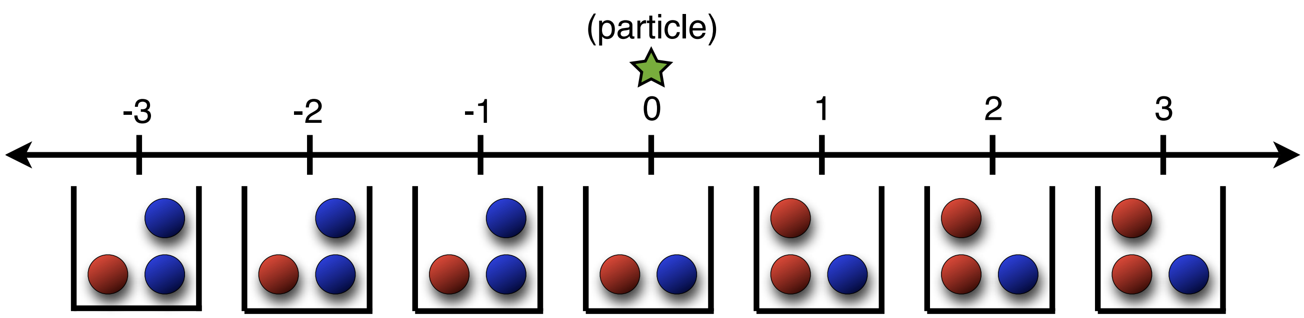

FIGURE 1. Initial set up for Polya urns with initial weight $a = 1$

The connection between acyclic ERRWs and Polya's urns is a powerful tool because Polya's urns have several nice and well understood properties. In particular, Polya's urns are *exchangeable*. That is, the results of any sequence of drawings can be permuted without changing the probability of the overall sequence. For example, the chance of drawing three red marbles and then seven blue marbles is the same as the probability of drawing seven blue marbles and then three red marbles, or of drawing four blue, three red, three blue. The probability of a string of colors being drawn depends only on the number of each color drawn and not on the order they are drawn in. Because of this property, the specific but powerful de Finetti Theorem can be applied. For a detailed discussion of how exchangeability and the de Finetti Theorem applies to ERRWs, see Coppersmith and Diaconis [1]. Consider, for example, the ERRW on $\mathbb{Z}$. The sequence of decisions at a particular vertex can be treated as a sequence of Polya's urn drawings from the urn at that vertex. This sequence can be predrawn, independently of the urns at other vertices. In the urn at each vertex v, let the number of red marbles and number of blue marbles after n drawings be denoted $R_n(v)$ and $B_n(v)$, respectively. Then the proportion of red marbles in the urn after n drawings is

$$\rho_n(v) = \frac{R_n(v)}{R_n(v) + B_n(v)}.$$

This ratio converges to some limiting fraction,

$$\rho_\infty(v) = \lim_{n\to\infty} \rho_n(v).$$

After sufficient time has passed, the drawings at vertex v are effectively independent Bernoulli trials with constant probability $\rho_\infty(v)$ of drawing a red marble. Thus the ERRW behaves like a random walk in a random environment (RWRE). Due to exchangeability, we know that the limiting fraction $\rho_\infty(v)$ is distributed as a beta distribution (in the case of a random walk on $\mathbb{Z}$, or more generally as a Dirichlet distribution), with parameters based on the initial marbles in the urn at vertex v. This fact, along with results from the study of RWRE's, allows one to prove, for instance, that an ERRW on $\mathbb{Z}$ is recurrent; that is, the particle returns to the origin with probability 1. Again, for the details of this argument, see [1]. Further, Pemantle has applied a similar method to prove recurrence and non-recurrence under certain conditions on the infinite binary tree [5], also an acyclic graph.

1.2. **Non-exchangeable ERRWs.** Significant study has been done on ERRWs on cyclic graphs. However, these random walks are harder to work with due to the absence of exchangeability. A cyclic graph can not be modeled by a Polya's urn at each vertex as with acyclic graphs. When the particle leaves a vertex, it is no longer required to return to that vertex by the same edge, so adding two marbles of the selected color to the urn would not be accurate. Despite this difficulty several notable results are known for ERRWs on cyclic graphs. Coppersmith and Diaconis found results for some finite cyclic graphs [1]. Merkl and Rolles proved recurrence for and ERRW on the infinite ladder $(\mathbb{Z}\times\{0,1\})$ for certain initial weights [3]. Rolles further showed recurrence on the finite width infinite ladder $(\mathbb{Z}\times\{0,\dots,d\})$ for any d, provided the initial weight is large enough [6]. However, the recurrence of a single particle ERRW on $\mathbb{Z}^2$ still remains an open problem.

There is another type of ERRW on which exchangeability fails: acyclic graphs with multiple particles. It is this class of ERRWs with which this paper is concerned. With multiple particles, the system can not be modeled with Polya's urns at each vertex. If particle 1 is initially at vertex v and leaves via the edge e_1, then in the Polya's urn model it would add two extra marbles to the urn. When particle 1 returns to vertex v, it is still guaranteed to do so by edge e_1, since the graph is acyclic. However, before that happens, particle 2 might arrive at vertex v by some other edge, say e_2. The correct edge weights would be $a+1$ for both e_1 and e_2 (where a is the initial weight), but the urn would have $a+2$ marbles of the color corresponding to e_1 and only a marbles of the color corresponding to e_2.

We are interested in multi-particle ERRWs because, though they are not exchangeable, they are closely related to single particle ERRWs, which *are* exchangeable. Thus we can attempt to use the powerful tools at our disposal for exchangeable processes to solve non-exchangeable problems. Perhaps the simplest system of this type, the two particle ERRW on $\mathbb{Z}$, is explored in [2]. It is proved that the two particles are recurrent, that is they meet each other infinitely often with probability 1. This is done by describing the movement of the particles using modified Polya's urns: the urn at each vertex contains red and blue marbles, but each urn also contains one "magic marble," which becomes blue if the rightmost particle is at the vertex and red if the leftmost particle is there. This single switching marble is enough to correct the discrepancies between a normal Polya's urn and the actual edge weights until the two particles meet. Thus this modified Polya's urn accurately describes the ERRW. It is no longer an exchangeable process, but one can construct a right limiting particle, which behaves like the right particle would if the magic marble were simply blue. This particle is guaranteed to be to the right of the right particle, and it is controlled by a regular, exchangeable Polya's urn. Likewise, a left limiting particle can be constructed that is always to the left of the left particle, and is controlled by a regular Polya's urn. It can be shown that these two limiting particles meet with probability 1, so the two real particles, sandwiched in between, must also meet.

1.3. **Goals of this paper.** The aim of this paper is to generalize the method used in [2] for the two particle ERRW on $\mathbb{Z}$ and extend the result to any number of particles. In addition to proving the recurrence of the k particle ERRW on $\mathbb{Z}$, our method of coupling non-exchangeable processes to closely related exchangeable processes may prove valuable in solving other problems. The principle theorem of this paper will be the following:

Theorem 1.1. *Let $x_1(t),\ldots,x_n(t)$ be the positions of n particles in and edge reinforced random walk on $\mathbb{Z}$. Then all n particles meet recurrently, for any $\in \mathbb{Z}^+$. That is,*

$$Prob(x_1(t) = \cdots = x_n(t) \text{ for infinitely many } t > 1) = 1$$

2. Method of Proof

In our proof of Theorem 1, we will not make use of the magic marbles or limiting particles that appear in [2]. Nonetheless, our method of proof is motivated by that approach. The underlying principles are exactly the same, but we state them without introducing magic marbles or limiting particles. These constructions are excellent for understanding the intuition behind the argument, and in fact they can be extended to prove Theorem 1. This alternate proof is included in Appendix A. However, extending magic marbles to many particles gets somewhat messy. We are able to avoid using magic marbles in our more general proof by noting that describing the movement of the particles exactly is not necessary. What really matters is that a given particle in a multi-particle ERRW is more inward drifting than a a single-particle ERRW would be. Our proof still makes use of a modified Polya's urn, but instead of magic marbles, we simply note that the drawings for the actual particle are skewed to one side with respect to the drawings for a single-particle ERRW. This skew can be thought of as removing and adding marbles from the urn temporarily, which in the two particle case is completely equivalent to the magic marble. However, our model pays no attention to how big the skew is at any given time, only that a skew exists. This aspect of the proof is directly motivated by the magic marbles used in [2], but it is done in a less specific way that is easier to generalize to many particles.

Another problem with extending the method in [2] is the use of strict limiting particles. [2] introduces an actual right (left) limiting particle that moves independently of the right (left) particle, except when the two are together, thus guaranteeing that the limiting particles are always on the outside of the real particles. Not only is this method difficult to define without magic marbles, but the right limiting particle, as defined in [2], is guaranteed to be to the right of the right-starting particle, but not to the right of the left-starting particle. This works in the two particle system, because if the left and right particles switch places they must have met somewhere. However, for three or more particles, this is not true; particles can pass each other without all of them meeting at one place. This problem can be fixed by adding more limiting particles (one on each side of each particle). However, the motivation of these limiting particles is really to show that a particle in a multi-particle ERRW is more inward drifting than a single particle ERRW would be. In our proof, we able to couple the distributions obtained from our skewed Polya's urn and a normal Polya's urn to show this same property without introducing limiting particles. Again, the motivation is the same, but the method is more general and easier to apply to any number of particles.

The basic outline of our proof is as follows: We describe particle movement in a k particle ERRW on $\mathbb{Z}$ in terms of modified Polya's urns. We find the distributions for the limiting ratio of marbles in these urns, or rather how those distributions compare to the familiar beta distribution obtained from a single-particle ERRW. Coupling these two distributions allows us to apply a result from RWREs, since after sufficient time this is how the ERRW behaves. This proves that all k particles return to a given finite region together infinitely often. Finally, we show that each time this happens, there is a nonzero probability of the particles meeting, and so the walk is recurrent.

3. Proof of Theorem 1

Recall that if a Polya's urn initially has R_0 red marbles and B_0 blue marbles, and d marbles are added on each step, then the limiting fraction ρ_∞ of red marbles in the urn is distributed as the beta distribution $\beta_{\frac{R_0}{d},\frac{B_0}{d}}$, and the limiting fraction σ_∞ of blue marbles in the urn is distributed as $\beta_{\frac{B_0}{d},\frac{R_0}{d}}$. The probability distribution function (pdf) for a beta distribution with parameters α and β, defined for $0 \le x \le 1$, is

$$\beta_{\alpha,\beta}(x) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha,\beta)}, \tag{1}$$

where $B(\alpha,\beta)$ denotes the beta function:

$$B(\alpha,\beta) = \int_0^1 x^{\alpha-1}(1-x)^{\beta-1}dx = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}.$$

The cumulative distribution function (cdf) with the same parameters, defined for $0 \le t \le 1$, is

$$I_{\alpha,\beta}(t) = \int_0^t \beta_{\alpha,\beta}(x)dx = \frac{\int_0^t x^{\alpha-1}(1-x)^{\beta-1}dx}{\int_0^1 x^{\alpha-1}(1-x)^{\beta-1}dx} \tag{2}$$

This cdf has a property that will be useful in our proof: for a given x, it is a non-increasing function of α and a non-decreasing function of β.

Proposition 3.1. *Let $\alpha' \ge \alpha$ and $\beta' \le \beta$. Then for all $t \in [0,1]$, $I_{\alpha',\beta'}(t) \le I_{\alpha,\beta}(t)$.*

Proof.

$$\begin{aligned}
\frac{1}{I_{\alpha',\beta'}(t)} &= \frac{\int_0^1 x^{\alpha'-1}(1-x)^{\beta'-1}dx}{\int_0^t x^{\alpha'-1}(1-x)^{\beta'-1}dx} = \frac{\int_0^t x^{\alpha'-1}(1-x)^{\beta'-1}dx + \int_t^1 x^{\alpha'-1}(1-x)^{\beta'-1}dx}{\int_0^t x^{\alpha'-1}(1-x)^{\beta'-1}dx} \\
&= 1 + \frac{\int_t^1 x^{\alpha'-1}(1-x)^{\beta'-1}dx}{\int_0^t x^{\alpha-1}(1-x)^{\beta-1}dx} = 1 + \frac{\int_t^1 \frac{x^{\alpha'-\alpha}}{(1-x)^{\beta-\beta'}} x^{\alpha-1}(1-x)^{\beta-1}dx}{\int_0^t \frac{x^{\alpha'-\alpha}}{(1-x)^{\beta-\beta'}} x^{\alpha-1}(1-x)^{\beta-1}dx}
\end{aligned}$$

Now, $x^{\alpha-1}(1-x)^{\beta-1}$ is nonnegative for all $x \in [0,1]$, and since $\alpha'-\alpha$ and $\beta-\beta'$ are both nonnegative,

$$\frac{x^{\alpha'-\alpha}}{(1-x)^{\beta-\beta'}} \le \frac{t^{\alpha'-\alpha}}{(1-t)^{\beta-\beta'}} \qquad \forall x \in [0,t]$$

$$\frac{x^{\alpha'-\alpha}}{(1-x)^{\beta-\beta'}} \geq \frac{t^{\alpha'-\alpha}}{(1-t)^{\beta-\beta'}} \qquad \forall x \in [t,1]$$

Thus,

$$\begin{aligned}
\frac{1}{I_{\alpha',\beta'}(t)} &= 1+\frac{\int_t^1 \frac{x^{\alpha'-\alpha}}{(1-x)^{\beta-\beta'}} x^{\alpha-1}(1-x)^{\beta-1}dx}{\int_0^t \frac{x^{\alpha'-\alpha}}{(1-x)^{\beta-\beta'}} x^{\alpha-1}(1-x)^{\beta-1}dx} \geq 1+\frac{\int_t^1 \frac{t^{\alpha'-\alpha}}{(1-t)^{\beta-\beta'}} x^{\alpha-1}(1-x)^{\beta-1}dx}{\int_0^t \frac{t^{\alpha'-\alpha}}{(1-t)^{\beta-\beta'}} x^{\alpha-1}(1-x)^{\beta-1}dx} \\
&= 1+\frac{\frac{t^{\alpha'-\alpha}}{(1-t)^{\beta-\beta'}} \int_t^1 x^{\alpha-1}(1-x)^{\beta-1}dx}{\frac{t^{\alpha'-\alpha}}{(1-t)^{\beta-\beta'}} \int_0^t x^{\alpha-1}(1-x)^{\beta-1}dx} = 1+\frac{\int_t^1 x^{\alpha-1}(1-x)^{\beta-1}dx}{\int_0^t x^{\alpha-1}(1-x)^{\beta-1}dx} \\
&= \frac{\int_0^1 x^{\alpha-1}(1-x)^{\beta-1}dx}{\int_0^t x^{\alpha-1}(1-x)^{\beta-1}dx} = \frac{1}{I_{\alpha,\beta}(t)}
\end{aligned}$$

$$\therefore I_{\alpha',\beta'}(t) \leq I_{\alpha,\beta}(t)$$

□

Now consider the *left skewed Polya's urn* in which at each drawing, the probability may be skewed slightly toward drawing red. More precisely, at the $(n+1)$th drawing, there are some deviations i_n and j_n s.t. $0 \leq i_n \leq k-1$ and $0 \leq j_n \leq k-1$. The probability of drawing a red marble from the urn is no longer $\frac{R_n}{R_n+B_n}$, but the greater probability $\frac{R_n+i_n}{(R_n+i_n)+(B_n-j_n)}$. This can be thought of in the following way: at the time of the drawing, we *temporarily* add i_n red marbles to the urn and remove j_n blue marbles. After the drawing, the number of red and blue marbles in the urn is returned to normal, but the deviations affect the probability of drawing red on that drawing.

With the above description, it is clear that the left skewed Polya's urn can describe the k particle ERRW on $\mathbb{Z}$ for vertices to the right of r_0. At each of these vertices, we place an urn initially containing $a+1$ red marbles and a blue marbles (where a is the initial weight), just as in a single particle ERRW. The decisions made at each vertex are described by a series of urn drawings, where two marbles of the drawn color are added each time. However, special adjustments must be made depending on which particle is doing the drawing. If the rightmost particle is at the vertex in question, then $i_n = j_n = 0$; the marbles accurately represent the weights in the ERRW. However, if the second particle from the right is at the vertex, then the marbles in the urn do not accurately represent. There is one too many blue marbles, since the last time the rightmost particle left that vertex to the right it added two blue marbles, but only increased the weight of the right leaving edge by 1. There are also too few red marbles, since the particle there now increased the weight of the left leaving edge by 1, but no red marbles were added. So, to accurately describe the particle's movement with the urn, we must temporarily remove one blue marble and add one red marble. If the third and fourth particles from the right are simultaneously at the vertex, we must remove two blue marbles (because the right two particles have left to the right and not come back) and add three red marbles (because the second, third, and fourth particles from the right have all arrived at that vertex from the left). In general, to make an adjusted urn drawing at a vertex, we must remove one blue marble for each particle currently to the right of that vertex, and add one red marble for each particle at or to the right of that vertex, except for the rightmost particle. We at most have to remove $k-1$ blue marbles and add $k-1$ red marbles. Keeping track of exactly how many marbles need to be added or removed would be difficult, but all that is important, we will see, is that any adjustment made increases the probability of drawing a red marble and moving left. Thus in the left skewed Polya's urn, the probability of drawing a red marble is always greater than or equal to the proportion of red marbles in the urn.

As in the case of a simple Polya's urn, the proportion of red marbles in the skewed Polya's urn after n drawings, ρ_n, converges to a limiting fraction ρ_∞, and the proportion of blue marbles after n drawings, σ_n, converges to a limiting fraction $\sigma_\infty = 1-\rho_\infty$. This may not seem as obvious for the skewed Polya's urn, since the probability of drawing blue, and thus adding more blue marbles, is not equal to the fraction of blue marbles in the urn. However, since the deviations i_n and j_n are bounded by k, this discrepancy is negligible after enough time has passed. If the total number of marbles in the urn is very large compared to k, then removing k blue marbles and adding k red marbles does not make a significant difference in the probability of drawing blue, and furthermore this difference goes to zero as n goes to infinity. So, the proportion of blue marbles in the left skewed Polya's urn still converges to a limiting fraction σ_∞. We expect this limiting fraction to be distributed between 0 and 1 with some probability distribution. For the simple Polya's urn, we can use exchangeability to show that the limiting fraction is distributed as a beta distribution with parameters $\frac{B_0}{2}$ and $\frac{R_0}{2}$, but in the absence of exchangeability we have no way of calculating the distribution of σ_∞. However, since in the left skewed Polya's urn the probability of drawing blue is always less than or equal to the proportion of blue marbles in the urn, it makes sense that the distribution would be somehow shifted to the left from the distribution of limiting fractions for the simple Polya's urn with the same initial marbles. This notion is indeed true, and is formalized in the following Lemma.

Lemma 3.2. *Let the distribution of the limiting ratio σ_∞ of blue marbles in the left skewed Polya's urn with parameters α and β (corresponding to the urn initially having 2α blue marbles and 2β red marbles) have the probability density function $f_{\alpha,\beta}(x)$. Let the corresponding cumulative distribution function be $C_{\alpha,\beta}(t) = \int_0^t f_{\alpha,\beta}(x)dx$. Then $C_{\alpha,\beta}(t) \geq I_{\alpha,\beta}(t)$ $\forall t \in [0,1]$.*

Proof. Essentially we prove this by showing that the effect of taking into account the "skew" on any one particular drawing is to shift the distribution to the left (shift the cumulative distribution function up). Thus when all the skews are taken into account, the cumulative distribution function is shifted up. More formally, we define a sequence of functions $C_{\alpha,\beta,n}(t)$ that converges pointwise to $C_{\alpha,\beta}(t)$ and such that

$$I_{\alpha,\beta} \leq C_{\alpha,\beta,1} \leq C_{\alpha,\beta,2} \leq \cdots$$

Thus, it is clear that $I_{\alpha,\beta} \leq C_{\alpha,\beta}$

Imagine a simple Polya's urn starting with $B_0 = 2\alpha$ blue marbles and $R_0 = 2\beta$ red marbles, where two marbles are added after each drawing. In this case, the limiting ratio σ_∞ of blue marbles is distributed as $\beta_{\alpha,\beta}(x)$, and the distribution has a cdf of $I_{\alpha,\beta}(t)$. Now imagine altering this urn so that the first drawing has a red-biased skewed, but all other drawings are done normally. That is, for the first drawing, there is some probability $p_1 \geq \frac{R_0}{R_0+B_0}$ of drawing red, but for all other drawings the probability of drawing red is still $\frac{R_n}{R_n+B_n}$. Let the cumulative distribution function for the limiting ratio σ_∞ in this case be $C_{\alpha,\beta,1}(t)$. This distribution can be calculated, since there are only two possible outcomes for the first drawing, and after that the urn behaves normally. Let y_1

be the result of the first drawing, which can be either r (red) or b (blue). Then

$$P(\sigma_\infty \le t) = P(\sigma_\infty \le t | y_1 = r)P(y_1 = r) + P(\sigma_\infty \le t | y_1 = b)P(y_1 = b). \tag{3}$$

Note that in the case of a simple Polya's urn,

$$P(\sigma_\infty \le t) = \frac{R_0}{R_0 + B_0} P(\sigma_\infty \le t | y_1 = r) + \frac{B_0}{R_0 + B_0} P(\sigma_\infty \le t | y_1 = b)$$

and so,

$$I_{\alpha,\beta}(t) = \frac{\beta}{\alpha+\beta} I_{\alpha,\beta+1}(t) + \frac{\alpha}{\alpha+\beta} I_{\alpha+1,\beta}(t) \tag{4}$$

Now when we add the effect of the red-biased skew in the first drawing and use (3), we have

$$C_{\alpha,\beta,1}(t) = p_1 I_{\alpha,\beta+1}(t) + (1 - p_1) I_{\alpha+1,\beta}(t) \tag{5}$$

Subtracting (4) from (5) gives

$$C_{\alpha,\beta,1}(t) - I_{\alpha,\beta}(t) = \left(p_1 - \frac{\beta}{\alpha+\beta}\right)\left(I_{\alpha,\beta+1}(t) - I_{\alpha+1,\beta}(t)\right) \tag{6}$$

The first term on the right of (6) is positive, and by Proposition 1 the second term is positive too. Therefore $I_{\alpha,\beta}(t) \le C_{\alpha,\beta,1}(t)$.

In the same way, we can consider the Polya's urn with red-biased skew on the first n drawings, but which behaves as a normal Polya's urn for all drawings after that. Define $C_{\alpha,\beta,n}$ to be the cdf for the distribution of σ_∞ in this case. For each n, there are a finite number of states the urn can be in after n drawings, so the cdf can also be expressed as a finite weighted average of cdf's for the normal Polya's urns starting at those states. That is,

$$C_{\alpha,\beta,n}(t) = \sum_{m=0}^{n} p(\alpha+m, \beta+n-m) I_{\alpha+m,\beta+n-m}(t), \tag{7}$$

where $p(\alpha+m, \beta+n-m)$ represents the probability of getting to the state where the urn has $2\alpha+2m$ blue marbles and $2\beta+2(n-m)$ red marbles after n drawings. Adding the effect of a red-biased skew on the $(n+1)$st drawing is the same as replacing each of the normal Polya's urns for all the possible states after n drawings with urns with a skewed first drawing. For all n,

$$\begin{aligned} C_{\alpha,\beta,n+1}(t) &= \sum_{m=0}^{n} p(\alpha+m, \beta+n-m) C_{\alpha+m,\beta+n-m,1}(t) \\ &\ge \sum_{m=0}^{n} p(\alpha+m, \beta+n-m) I_{\alpha+m,\beta+n-m}(t) = C_{\alpha,\beta,n}(t) \end{aligned}$$

By induction, $I_{\alpha,\beta} \le C_{\alpha,\beta,1} \le C_{\alpha,\beta,2} \le \cdots$. Now $C_{\alpha,\beta}(t)$ is the distribution of σ_∞ taking into account the right-biased skew of all drawings. Also, the effect of the skew in the probability goes to zero for large n, since the number of marbles that are changed on each drawing is bounded by k. Thus it is clear that for any $t \in [0,1]$,

$$\lim_{n\to\infty} C_{\alpha,\beta,n} = C_{\alpha,\beta}(t)$$

$\therefore I_{\alpha,\beta}(t) \le C_{\alpha,\beta}(t)$ for all $t \in [0,1]$. □

We have shown that for vertices to the right of r_0, the k particle ERRW on $\mathbb{Z}$ is described by a left skewed Polya's urn initially containing $a+1$ red marbles and a blue marbles at each vertex. Furthermore, since the ratio of blue marbles in the urn at vertex v approaches a limiting fraction $\sigma_\infty(v)$, after sufficient time a particle at v effectively jumps right with a constant probability $\sigma_\infty(v)$. Thus after sufficient time, the ERRW can be treated as a RWRE in the region to the right of r_0. Finally, while we can not calculate the distribution of $\sigma_\infty(v)$ at each vertex directly, we have proved that the distribution is shifted to left from $\beta_{\frac{a}{2},\frac{a+1}{2}}$ in the sense that the cdf is greater for all $t \in [0,1]$.

We can also define a *right skewed Polya's urn*, in which at each drawing the probability of drawing a red marble is less than or equal to the proportion of red marbles in the urn. By symmetry, the same argument used above shows that the k particle ERRW on $\mathbb{Z}$ in the region to the left of l_0 can be described by a right skewed Polya's urn initially containing one red marble and two blue marbles at each vertex, and also that the distribution of the limiting ratio ρ_∞ of red marbles in the urn is shifted to the left of $\beta_{\frac{a}{2},\frac{a+1}{2}}$ in the sense that the cdf is greater than or equal to $I_{\frac{a}{2},\frac{a+1}{2}}(t)$ for all $t \in [0,1]$.

So at least outside of the finite region between l_0 and r_0, the k particle ERRW on $\mathbb{Z}$ behaves as a RWRE. In this RWRE, the k particles move independently, and can be treated separately. This also allows us to use the following result from the study of RWREs.

Lemma 3.3. *For each $1 \le i \le l$, let $p_i(1), p_i(2), \dots$ be independent random variables defined on (0, 1) with*

$$\prod_{m=1}^{n} \frac{(1-p_i(m))}{p_i(m)} \to f(n) \ge e^{\mu_i n} \text{ as } n \to \infty \text{ for some } \mu_i > 0 \tag{8}$$

Also let $p_i(0) = 1$ for all $1 \le i \le l$. If $p_i(0), p_i(1), p_i(2), \dots$ are the forward rates for the birth-and-death chain Z_t^i, then the l-dimensional RWRE $X_t = (Z_t^1, \dots, Z_t^l)$ returns to zero infinitely often.

Proof. Suppose we have $X_t = (x_1, ..., x_l) \in \mathbb{Z}_1^+ \times ... \times \mathbb{Z}_l^+$. X_t will move in a particular dimension with probability $\frac{1}{l}$. We therefore have:

$$P[X_{t+1} = (x_1, ..., x_i + 1, ..., x_l) | X_t = (x_1, ..., x_l)] = \frac{p_i(x_i)}{l}$$
$$P[X_{t+1} = (x_1, ..., x_i - 1, ..., x_l) | X_t = (x_1, ..., x_l)] = \frac{1 - p_i(x_i)}{l}$$

To prove recurrence, we will generalize the argument used in [4] to multiple-dimensions using Lyapunov functions [1] that are supermartingale [2].

Let $\nu_i \in (0, \mu_i)$ and define a Lyapunov function ψ where for each i we have:

$$\psi_i(n) = A_i + 1 + \sum_{m=1}^{n-1} \frac{(1-p_i(1)) \cdot \ldots \cdot (1-p_i(m))}{p_i(1) \cdot \ldots \cdot p_i(m)} e^{-\nu_i m} \tag{9}$$

We define ψ as follows:

$$\psi(X_t) = \sum_{i=1}^{l} \psi_i(x_i) \tag{10}$$

[1] A Lyapunov function is a function ϕ such that $\lim_{x \to \infty} \phi(x) = \infty$.
[2] $\phi(X_t) \ge E[\phi(X_{t+1}) | X_t]$

It is fairly easy to check that $\psi(X_n)$ is now supermartingale for points not on the axes, which gives us

$$\psi(X_t) \geq E[\psi(X_{t+1})|X_t = (x_1,...,x_l)] \qquad \text{where } x_i \neq 0 \text{ for all } 1 \leq i \leq l$$

We claim that $\psi(X_t)$ is supermartingale everywhere outside some ball around the origin. For $X_t = (x_1,...,x_l)$, the worst case will be when $x_i = 0$ for all but one i between 1 and l. Without loss of generality, assume $x_1 = n \neq 0$ and $x_i = 0$ for all $1 < i \leq l$. The conditional expectation of $\psi(X_{t+1})$ is as follows:

$$E[\psi(X_{t+1})|X_t = (n,0,...,0)] = \psi(X_t) + \frac{l-1}{l} - \frac{p_1(n)}{l} \cdot \frac{(1-p_1(1)) \cdot \ldots \cdot (1-p_1(n))}{p_1(1) \cdot \ldots \cdot p_1(n)} (e^{-\nu_1(n-1)} - e^{-\nu_1 n})$$

To show that this expectation is less then $\psi(X_t)$, we need to show that for large enough n,

$$\frac{l-1}{l} - \frac{p_1(n)}{l} \cdot \frac{(1-p_1(1)) \cdot \ldots \cdot (1-p_1(n))}{p_1(1) \cdot \ldots \cdot p_1(n)} (e^{-\nu_1(n-1)} - e^{-\nu_1 n}) < 0 \tag{11}$$

By equation 8, we have that $\frac{(1-p_1(1)) \cdot \ldots \cdot (1-p_1(n))}{p_1(1) \cdot \ldots \cdot p_1(n)} \geq e^{\mu_1 n}$ for large enough n. To show that statement 11 is true, we show that for any constant $c > 0$, there a.s. exists $N > 0$ such that

$$c p_1(n) e^{(\mu_1 - \nu_1)n} \geq 1 \qquad \text{for all } n > N \tag{12}$$

Note that $P[p_1(n) < e^{-(\mu_1-\nu_1)\frac{n}{2}}/c]$ decreases exponentially. This implies that

$$\prod_{n=1}^{\infty} P[p_1(n) > e^{-(\mu_1-\nu_1)\frac{n}{2}}/c] \neq 0$$

By making $c = \frac{e^{\nu_1}-1}{2(l-1)}$, we know we can find N such that statement 12 holds. We can therefore show that statement 11 is true.

This gives our desired result: $E[\psi(X_{t+1})|X_t = (n,0,...,0)] \leq \psi(X_t)$.

We now have a supermartingale Lyapunov function for our random variables X_t outside some ball around the origin, implying recurrence around the origin[4][2]. □

We apply this lemma with $l = 2k$. We let p_i for $1 \leq i \leq k$ be the limiting right jump probability σ_∞ at each vertex, starting with r_0 as the origin. That is, $p_i(m) = \sigma_\infty(r_0 + m)$ for $m \geq 1$. Thus $p_i(m)$ is distributed as the limiting fraction of blue marbles in the urn at vertex $r_0 + m$, whose distribution is shifted to the left of $\beta_{\frac{a}{2},\frac{a+1}{2}}$. Note that this distribution is independent of i; all particles have the same limiting right jump probability at any given vertex. Now let Z_t^i be the maximum of the ith particle's position minus r_0 at time t, and 0. Thus Z_t^i exactly traces the particle's movement, indexing how many spaces to the right of r_0 the particle is, except when the particle moves to the left of r_0. In this case, Z_t^i is 0 until the particle returns and moves to the right of r_0 again. Z_t^i can only increase from 0, so if $p_i(0) = 1$, then $p_i(0), p_i(1), p_i(2), \ldots$ are the forward rates for the birth-and-death chain Z_t^i. $(Z_t^1, \ldots, Z_t^k)$ being at zero is equivalent to all k particles being at or to the left of vertex r_0. In the same way we let p_{k+i} for $1 \leq i \leq k$ be the limiting *left* jump probability at each vertex to the left of l_0. That is, $p_{k+i}(m) = \rho_\infty(l_0 - m)$ for $m \geq 1$. Now let Z_t^{k+i} be the maximum of l_0 minus the ith particle's position at time t, and 0. Thus Z_t^{k+i} exactly traces the ith particles movement, indexing how many spaces to the left of l_0 the particle is, except when the particle is to the right of l_0, in which case Z_t^{k+i} is 0. If $p_{k+i}(0) = 1$, then $p_{k+i}(0), p_{k+i}(1), p_{k+i}(2), \ldots$ are the forward rates for the birth-and-death chain Z_t^{k+i}. $(Z_t^{k+1}, \ldots, Z_t^{2k})$ being at zero is equivalent to all k

particles being at or to the right of vertex l_0. If $X_t = (Z_t^1, \ldots, Z_t^{2k})$ returns to zero, which is the result of Lemma 3.3, then all k particles are simultaneously to the left of or at r_0 and to the right of or at l_0. Thus in order to prove that all k particles return to the finite region $[l_0, r_0]$ together infinitely often, it only remains to check that the distributions of ρ_∞ and σ_∞ derived from the skewed Polya's urns satisfy the conditions of Lemma 3.3, which we will do by means of coupling.

Lemma 3.4. *The distributions $p_i(0), p_i(1), p_i(2), \ldots$ as defined above in terms of the distributions of $\rho_\infty(v)$ and $\sigma_\infty(v)$ satisfy the conditions of Lemma 3.3.*

Proof. The urns at each vertex can be treated as independent, because they can be pre-drawn separately with any amount of skew on each drawing. In reality, the exact amount of skew at a given drawing depends on the movement of the particles and thus on the other urns, but we do not need to know exactly what the skew is. We just need to know that a skew (potentially zero) exists at every drawing. So all we have to check is that our distributions satisfy condition (8).

While we do not know the distributions of $\rho_\infty(v)$ or $\sigma_\infty(v) = 1 - \rho_\infty(v)$ exactly, by Lemma 3.2 we do know that for vertices to the right of r_0, the distribution of $\sigma_\infty(v)$ is shifted to the left of $\beta_{\frac{a}{2},\frac{a+1}{2}}$, that is, it has a cumulative distribution function $C_{\frac{a}{2},\frac{a+1}{2},v}(t)$ s.t. $C_{\frac{a}{2},\frac{a+1}{2},v}(t) \geq I_{\frac{a}{2},\frac{a+1}{2}}(t)\ \forall t \in [0,1]$. As a result of this, we can couple the random variable $p_i(m)$ for $1 \leq i \leq k$ with a $\beta_{\frac{a}{2},\frac{a+1}{2}}$ distribution for each m. We can pick $p_i(1), p_i(2), \ldots$ and $q_i(1), q_i(2), \ldots$ simultaneously so that $p_i(m)$ is distributed as $\sigma_\infty(r_0 + m)$ (it's cdf is $C_{\frac{a}{2},\frac{a+1}{2},r_0+m}$), $q_i(m)$ is distributed as a $\beta_{\frac{a}{2},\frac{a+1}{2}}$ distribution, and $p_i(m) \leq q_i(m)\ \forall m$. Do this in the following way: Let $u(m)$ be uniformly distributed on [0, 1]. Then let $p_i(m) = C^{-1}_{\frac{a}{2},\frac{a+1}{2},r_0+m}(u(m))$ and let $q_i(m) = I^{-1}_{\frac{a}{2},\frac{a+1}{2}}(u(m))$ (see Figure 2). It is clear that the random variables are distributed with the appropriate distributions, and since the cdf's are non-decreasing functions and $C_{\frac{a}{2},\frac{a+1}{2},r_0+m} \geq I_{\frac{a}{2},\frac{a+1}{2}}$, it is clear that $p_i(m) \leq q_i(m)\ \forall m$.

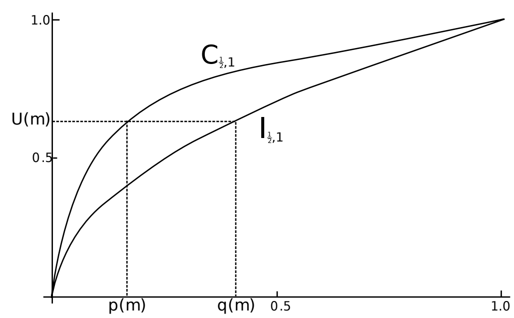

FIGURE 2. Coupling of skewed urn distribution and beta distribution

Now $q_i(1), q_i(2), \ldots$ are i.i.d. beta distributions with parameters $\alpha < \beta$. It is known that for this distribution,

$$\mu_i = E\left[log\left(\frac{1 - q_i(1)}{q_i(1)}\right)\right] > 0.$$

Applying the strong law of large numbers, we see that

$$\prod_{m=1}^{n} \frac{1 - q_i(m)}{q_i(m)} \to e^{\mu_i n} \quad \textit{almost surely as} \quad n \to \infty$$

Since $p_i(m) \leq q_i(m)$ $\forall m$, we have that

$$\prod_{m=1}^{n} \frac{1-p_i(m)}{p_i(m)} \geq \prod_{m=1}^{n} \frac{1-q_i(m)}{q_i(m)},$$

and thus for $1 \leq i \leq k$,

$$\prod_{m=1}^{n} \frac{1-p_i(m)}{p_i(m)} \to f(n) \geq e^{\mu_i n} \quad \textit{almost surely as} \quad n \to \infty.$$

Similarly, for vertices to the left of l_0, we know by Lemma 3.2 that the distribution of $\rho_\infty(v)$ is shifted to the left of $\beta_{\frac{a}{2},\frac{a+1}{2}}(t)$, that is, it has a cumulative distribution function $C_{\frac{a}{2},\frac{a+1}{2},v}(t)$ s.t. $C_{\frac{a}{2},\frac{a+1}{2},v}(t) \geq I_{\frac{a}{2},\frac{a+1}{2}}(t)$ $\forall t \in [0,1]$. Let $v(m)$ be uniformly distributed on [0, 1]. Then for $k+1 \leq i \leq 2k$ let $p_i(m) = C^{-1}_{\frac{a}{2},\frac{a+1}{2},r_0+m}(v(m))$ and let $q_i(m) = I^{-1}_{\frac{a}{2},\frac{a+1}{2}}(v(m))$. Now $p_i(m)$ has the same distribution as $\rho_\infty(l_0-m)$, $q_i(m)$ is distributed as $\beta_{\frac{a}{2},\frac{a+1}{2}}$, and $p_i(m) \leq q_i(m)$ $\forall m$. Again since $q_i(1), q_i(2), \ldots$ are i.i.d. $\beta_{\frac{a}{2},\frac{a+1}{2}}$ distributions with parameters $\alpha < \beta$, it is known that

$$\mu_i = E\left[log\left(\frac{1-q_i(1)}{q_i(1)}\right)\right] > 0.$$

Since $p_i(m) \leq q_i(m)$ $\forall m$, we have that

$$\prod_{m=1}^{n} \frac{1-p_i(m)}{p_i(m)} \geq \prod_{m=1}^{n} \frac{1-q_i(m)}{q_i(m)},$$

and thus for $k+1 \leq i \leq 2k$,

$$\prod_{m=1}^{n} \frac{1-p_i(m)}{p_i(m)} \to f(n) \geq e^{\mu_i n} \quad \textit{almost surely as} \quad n \to \infty.$$

Therefore condition (8) holds for all $1 \leq i \leq 2k$, and Lemma 3.3 applies.

□

We have now proved that in the k particle ERRW on $\mathbb{Z}$, with probability 1 all k particles will be in the finite region $[l_0, r_0]$ at the same time infinitely often. The only remaining step in the proof of Theorem 1 is to show that this implies recurrent meeting of all k particles.

Lemma 3.5. *In an k-particle ERRW, if all k particles will be in the finite region* $[l_0, r_0]$*at the same time infinitely often, then all k particles will meet recurrently.*

Proof. Recall that at each vertex v within the finite region $[l_0, r_0]$ there is a limiting ratio of red (or blue) marbles in the urn:

$$\rho_\infty(v) = \lim_{n\to\infty} \frac{R_n(v)}{R_n(v)+B_n(v)}$$

$$\sigma_\infty(v) = 1-\rho_\infty(v) = \lim_{n\to\infty} \frac{B_n(v)}{R_n(v)+B_n(v)}$$

$R_n(v)$ and $B_n(v)$ are increasing functions of n, and, as all the particles will return to positions between l_0 and r_0 recurrently, $R_n(v)$ and $B_n(v)$ go to infinity as n goes to infinity for $l_0 < v < r_0$. Thus for any $\varepsilon > 0$, there exists a $N > 0$ such that for all $n > N$,

$$|\rho_\infty(v) - p_n(v)| < \varepsilon$$

where $p_n(v)$ is the probability of moving to the left from position v after n drawings. Assume at least N drawings have occurred, and let d_i be the distance from l_0 to the ith particle (for some arbitrary indexing of the particles).

One possible way the particles can all meet is if the first particle travels directly left to l_0 (without any other particles moving), then the next particle travels directly to l_0, and so on, until all the particles are at l_0. The probability of this event is given by:

$$p = \prod_{i=1}^{k}\left(\prod_{j=1}^{d_i}(p_n(l_0+j))\frac{1}{k}\right) \geq \prod_{i=1}^{k}\left(\prod_{j=1}^{r_0-l_0}(\rho_\infty(l_0+j)-\varepsilon)\frac{1}{k}\right) \equiv p_{min}$$

If $\rho_\infty(v) > 0$ for all $v \in [l_0, r_0]$ (which is true a.s.), then as long as $\varepsilon < \min\{\rho_\infty(v) | l_0 \leq v \leq r_0\}$, p_{min} must be greater than zero.

When all k particles are in $[l_0, r_0]$ after N drawings, they have a probability of at least $p_{min} > 0$ of meeting. Because all particles will be in $[l_0, r_0]$ at the same time infinitely often, the probability of them never meeting is

$$\lim_{t\to\infty}(1-p_{min})^t = 0$$

□

This completes the proof of Theorem 1.

4. Further Applications

While Theorem 1 in itself is an interesting result, the method developed above can easily be applied to multi-particle ERRWs on acyclic graphs other than $\mathbb{Z}$. For example, the method transfers easily to the infinite binary tree. In [5], Pemantle proves that a single particle ERRW on an infinite binary tree is recurrent if the initial weight a is less than $a_0 \approx 0.233$. Using skewed Polya's urns and coupled distributions, we show that in this range of initial weights, a k particle ERRW on the infinite binary tree is recurrent.

Theorem 4.1. *For $a < a_0 \approx 0.233$, the range of initial weights in which a single particle ERRW on an infinite binary tree is recurrent as shown in* [5]*, the k particle ERRW on the infinite binary tree is also recurrent.*

Proof. The proof follows the same outline as the proof of Theorem 1, and uses much of the same machinery. Let y_i be the initial distance from the ith particle to the root for $1 \leq i \leq k$, and let $y = max\{y_i | 1 \leq i \leq k\}$. The set of vertices whose distance from the root is at most y is finite region in which all k particles start, analogous to the interval $[l_0, r_0]$ for the ERRW on $\mathbb{Z}$. At vertices outside of this region, place a Polya's urn which initially contains $a+1$ red marbles, a blue marbles, and a green marbles, where red marbles correspond to rootward edges and blue and green marbles correspond to the remaining two edges. This urn arrangement accurately describes the single particle ERRW on the infinite binary tree. It nearly describes the multi-particle ERRW as well, except that if a particle arrives at a vertex when other particles are at descendants of that vertex, there will be too many blue or green marbles and too few red marbles in the urn. To get the correct weights for a given drawing at vertex v, we must temporarily remove one blue marble for every particle that is at a descendant of v in the direction of the "blue" edge, remove one green marble for every particle that is at a descendant of v in the direction of the "green" edge, and add one red marble for all but one particle that is at v or a descendant of v. When these changes are taken into account, the probability of drawing red is greater than or equal to the original proportion of red marbles in the urn, and the probability of drawing blue or green is less than or equal to the original proportions of blue or green marbles in the urn, respectively. We take this as the definition for the *root skewed Polya's urn*.

Let the limiting distributions of red, blue, and green marbles in the root skewed urn at vertex v be $\rho_\infty(v)$, $\sigma_\infty(v)$, and $\tau_\infty(v)$, respectively. Although there are three marble colors, if we care about the limiting ratio of one color, we can treat the other two colors as one color. For example, for a normal Polya's urn starting with R_0 red marbles, B_0 blue marbles, and G_0 green marbles, where d marbles are added at each drawing, $\rho_\infty(v)$ is distributed as the beta distribution $\beta_{\frac{R_0}{d},\frac{B_0+G_0}{d}}$. Furthermore, Lemma 3.2 can be directly applied to the root skewed Polya's urns defined above. It shows that the distribution of $\rho_\infty(v)$ is shifted to the right and the distributions of $\sigma_\infty(v)$ and $\tau_\infty(v)$ are shifted to the left with respect to the standard (unskewed) distributions. More precisely, if the cdf's are $C_{\rho,v}$, $C_{\sigma,v}$, and $C_{\tau,v}$, then for all t in $[0,1]$

$$C_{\rho,v}(t) \leq I_{\frac{a+1}{2},a}(t) \tag{13}$$

$$C_{\sigma,v}(t) \geq I_{\frac{a}{2},\frac{2a+1}{2}}(t) \tag{14}$$

$$C_{\tau,v}(t) \geq I_{\frac{a}{2},\frac{2a+1}{2}}(t) \tag{15}$$

Because the ratios of marbles in the urn at each vertex converges to a limiting ratio, after sufficient time the ERRW behaves as a RWRE. We can thus apply Lemma 4.2 a slight variation on Lemma 3.3. The proof, included in Appendix A, is almost identical to the proof of Lemma 3.3.

Lemma 4.2. *Let $\mathbb{T}$ denote any tree. For each $1 \leq i \leq l$, let $\phi_i(v)$ for vertices $v \in \mathbb{T}$ be independent random variables defined on (0, 1) such that for any path $v_0, v_1, \ldots$ starting at the root and always moving outward, we have*

$$\prod_{m=2}^{n} \frac{1}{\phi_i(v_m)} \to f(n) \geq e^{\mu_i n} \text{ as } n \to \infty \text{ for some } \mu_i > 0 \tag{16}$$

Also let $\phi_i(0) = 1$ for all $1 \leq i \leq l$. If $\phi_i(v)$ is defined as

$$\phi_i(v) = \frac{\text{prob(transition from parent of } v \text{ to } v)}{\text{prob(transition from parent of } v \text{ to grandparent of } v)}$$

for the birth-and-death chain Z_t^i on $\mathbb{T}$, then the l-dimensional RWRE $X_t = (Z_t^1, \ldots, Z_t^l)$ on $\mathbb{T}^l$ returns to the origin (the root in all l dimensions) infinitely often.

We apply this lemma with $\mathbb{T}$ being the infinite binary tree, $l = k$, and each Z_t^i representing the movements of one particle in the k particle ERRW on $\mathbb{T}$. The challenge is to show that condition (16) holds. Consider vertices v more than d steps from the root. Here the distributions of $\phi_i(v)$ are determined by root skewed Polya's urns. If we ignore the skew, the distributions would be i.i.d.

beta distributions, equivalent to the distributions for a single particle ERRW. Pemantle proved that for initial weights $a < a_0$, this distribution satisfies

$$E[ln(\phi)] < 0 \text{ almost surely.}$$

Let $E\left[ln\left(\frac{1}{\phi(v)}\right)\right] = \mu_i > 0$ for vertices v more than y steps from the root. For any path $v_0, v_1, \ldots$ moving outward from the root, by the strong law of large numbers,

$$\sum_{m=y}^{n} ln\left(\frac{1}{\phi(v_m)}\right) \to (n-y)\mu_i \text{ a.s. as } n \to \infty$$

Since there are finitely many vertices within the first y steps and $ln\left(\frac{1}{\phi(v)}\right)$ is finite a.s. for these vertices, they can be included without affecting the above convergence:

$$\sum_{m=1}^{n} ln\left(\frac{1}{\phi(v_m)}\right) \to n\mu_i \text{ a.s. as } n \to \infty$$

Exponentiating each side gives

$$\prod_{m=1}^{n} \frac{1}{\phi(v_m)} \to e^{n\mu_i} \text{ a.s. as } n \to \infty$$

Thus for the unskewed distributions, condition (16) holds. In complete analogy to Lemma 3.4, we can use (13)-(15) to couple the real distributions for $\rho_\infty(v)$, $\sigma_\infty(v)$, and $\tau_\infty(v)$ to the unskewed distributions. At any vertex, the probability of jumping to a parent will increase, and the probability of jumping to any child will decrease, so $\phi_i(v)$ will decrease for every v at least d steps from the root. For these vertices, the terms in the product above will increase. We can again ignore the effect of the finitely many vertices within d steps of the root, and thus condition (16) holds, and the Lemma applies. The Lemma states that the random walk $X_t = (Z_t^1, \ldots, Z_t^k)$ on $\mathbb{T}^k$ returns to the origin, that is, all k random walkers return to the root at the same time, infinitely often with probability 1.

□

5. Conclusion

We have successfully extended both the result and method of [2] to multiple particle ERRWs on $\mathbb{Z}$ for any finite number of particles k, proving that such random walks are a.s. recurrent. Furthermore, the tool developed in this generalized method of proof turned out to be more general than originally hoped. Using results from [5], we easily proved recurrence for the k particle ERRW on the infinite binary tree, given certain initial weights. In fact, it seems that by this method, if one can prove recurrence for a single particle ERRW on any tree, then the recurrence of a multi-particle ERRW on the same tree follows easily. In some sense, having more particles seems to make an ERRW more likely, or at least as likely, to be recurrent. This is an interesting result, because it is not true for simple random walks, where for instance on $\mathbb{Z}$, the one, two, and three particle random walks are recurrent, but the k particle random walk for $k \geq 4$ is not recurrent. An interesting question is, are there any trees and initial weights for which a single particle ERRW is not a.s. recurrent, but a k particle ERRW is for sufficiently large k? Further explorations might also include finding different applications of our general method, specifically the coupling of random variables in nonexchangeable processes to better understood random variables in exchangeable processes.

APPENDIX A. PROOF OF LEMMA 4.2

Proof. For any vertex v, denote the parent of v by $\tilde{v}$ and the two children of v by v' and v''. Let $p_i(v \to v')$ denote the probability of the ith particle jumping from v to v', etc. Thus for any v,

$$\phi_i(v') = \frac{p_i(v \to v')}{p_i(v \to \tilde{v})}$$

Suppose we have $X_t = (x_1, ..., x_l) \in \mathbb{T}_1 \times ... \times \mathbb{T}_l$. X_t will move in a particular dimension with probability $\frac{1}{l}$. We therefore have:

$$P[X_{t+1} = (x_1, ..., x_i', ..., x_l) | X_t = (x_1, ..., x_l)] = \frac{p_i(x_i \to x_i')}{l}$$

$$P[X_{t+1} = (x_1, ..., x_i'', ..., x_l) | X_t = (x_1, ..., x_l)] = \frac{p_i(x_i \to x_i'')}{l}$$

$$P[X_{t+1} = (x_1, ..., \tilde{x}_i, ..., x_l) | X_t = (x_1, ..., x_l)] = \frac{p_i(x_i \to \tilde{x}_i)}{l}$$

To prove recurrence, we will generalize the argument used in [4] to multiple-dimensions using Lyapunov functions [3] that are supermartingale [4].

Let $\nu_i \in (0, \mu_i)$ and define a Lyapunov function ψ where for each i we have:

$$\psi_i(v) = A_i + 1 + \sum_{m=2}^{n} \frac{1}{\phi_i(v_2) \cdot \ldots \cdot \phi_i(v_m)} e^{-\nu_i(m-1)} \tag{17}$$

where $v_0, v_1, ..., v_n = v$ is the path from the root v_0 to v. We define ψ as follows:

$$\psi(X_t) = \sum_{i=1}^{l} \psi_i(x_i) \tag{18}$$

It is fairly easy to check that $\psi(X_t)$ is now supermartingale for points not on the axes, which gives us

$$\psi(X_t) \geq E[\psi(X_{t+1}) | X_t = (x_1, ..., x_l)] \qquad \text{where } x_i \text{ is not the root for all } 1 \leq i \leq l$$

We claim that $\psi(X_t)$ is supermartingale everywhere outside some ball around the origin. For $X_t = (x_1, ..., x_l)$, the worst case will be when x_i is the root vertex for all but one i between 1 and l. Without loss of generality, assume x_1 is not the root and x_i is for all $1 < i \leq l$. The conditional

[3] A Lyapunov function is a function ϕ such that $\lim_{x \to \infty} \phi(x) = \infty$.
[4] $\phi(X_n) \geq E[\phi(X_{n+1}) | X_n]$

expectation of $\psi(X_{t+1})$ is as follows:

$$\begin{aligned}E[\psi(X_{t+1})|X_t=(x_1,0,...,0)] = \psi(X_t) &+ \frac{l-1}{l} - \frac{p_1(x_1\to\tilde{x_1})}{l}\cdot\frac{1}{\phi_1(v_2)\cdots\phi_1(v_n)}e^{-\nu_1(n-1)}\\ &+ \frac{p_1(x_1\to x_1')}{l}\cdot\frac{1}{\phi_1(v_2)\cdots\phi_1(v_n)\phi(x_1')}e^{-\nu_1 n}\\ &+ \frac{p_1(x_1\to x_1'')}{l}\cdot\frac{1}{\phi_1(v_2)\cdots\phi_1(v_n)\phi(x_1'')}e^{-\nu_1 n}\\ = \psi(X_t) &+ \frac{l-1}{l} - \frac{p_1(x_1\to\tilde{x_1})}{l\cdot\phi_1(v_2)\cdots\phi_1(v_n)}(e^{-\nu_1(n-1)}-e^{-\nu_1 n})\end{aligned}$$

where $v_0, v_1, ..., v_n = x_1$ is the path from the root v_0 to x_1. To show that this expectation is less then $\psi(X_t)$, we need to show that for large enough n,

$$\frac{l-1}{l} - \frac{p_1(x_1\to\tilde{x_1})}{l\cdot\phi_1(v_2)\cdots\phi_1(v_n)}e^{-\nu_1 n}(e^{\nu_1}-1) < 0 \tag{19}$$

By equation 16, we have that $\frac{1}{\phi_1(v_2)\cdots\phi_1(v_n)} \geq e^{\mu_1 n}$ for large enough n. To show that statement 19 is true, we show that for any constant $c > 0$, there a.s. exists $N > 0$ such that

$$cp_1(v_n\to\tilde{v_n})e^{(\mu_1-\nu_1)n} \geq 1 \qquad \text{for all } n > N \tag{20}$$

Note that $P[p_1(v_n\to\tilde{v_n}) < e^{-(\mu_1-\nu_1)\frac{n}{2}}/c]$ decreases exponentially. This implies that

$$\prod_{n=1}^{\infty} P[p_1(v_n\to\tilde{v_n}) > e^{-(\mu_1-\nu_1)\frac{n}{2}}/c] \neq 0$$

By making $c = \frac{e^{\nu_1-1}}{2(l-1)}$, we know we can find N such that statement 20 holds. We can therefore show that statement 19 is true.

This gives our desired result: $E[\psi(X_{t+1})|X_t=(x_1,0,...,0)] \leq \psi(X_t)$.

We now have a supermartingale Lyapunov function for our random variables X_t outside some ball around the origin, implying recurrence around the origin[4][**?**]. □

Appendix B. Generalization of Kovchegov's Method

We present an alternate proof of Theorem 1.1 using magic marbles, which more directly extends the method in [2].

B.1. **Setup and particle movement.** In a k-particle system, the particles are differentiated by their positions relative to one another. The starting position of the left most particle is l_0 and the starting position of the right most particle is r_0. Each urn contains $k-1$ magic marbles, which turn blue or red depending on which particle(s) are at that urn. Let b_i be the ith particle from the left. When b_i is at an urn, $i-1$ magic marbles turn blue, and $n-i$ magic marbles turn red. For example, when b_1, the left most particle is at an urn, all the $n-1$ magic marbles turn red. A particle moves by choosing either a red or blue marble from the urn. The color dictates its movement: red means it moves left, blue means it moves right. It then replaces the marble and adds two more of the same color.

$R_0(v)$ and $B_0(v)$ are the initial numbers of red and blue marbles respectively in the urn at position v. The initial urn setup will be as follows for a k-particle system, where b_i^0 is the starting position of particle b_i, and a is the initial weight on each edge:

$$(R_0(v), B_0(v)) = \begin{cases} (a-(k), \quad a+1) & \text{if } v < l_0 \\ (a-(k-i), \quad a-i+1) & \text{if } v = b_i^0 \\ (a-(k-i)+1, \quad a-i+1) & \text{if } b_i^0 < v < b_{i+1}^0 \\ (a+1, \quad a-k) & \text{if } r_0 < v \end{cases}$$

This is assuming that no two particles begin at the same location. Ultimately, we only care about the initial distribution of marbles in the urns to the left of l_0 and to the right of r_0, and these remain the same no matter how many particles share a location.

Until any two particles meet, it is easy to see that the above representation of marbles in the urns correctly represents an ERRW.

B.2. **Limiting particles.** As in [2], we introduce left and right limiting particles. However, instead of just one pair for the entire system, there is a pair for each particle. The limiting particles stay with a particular particle regardless of that particle's position relative to the others.

We will define two concepts: the magic family and magic marbles' offspring. The magic family is composed of all the magic marbles and their offspring. The magic marbles' offspring are marbles that a particle added as a consequence of choosing a magic marble or a magic marble's offspring.

The right limiting particles move according to a limiting probability that is determined by making all marbles in the magic family blue. Similarly, the left limiting particles move according to a limiting probability that is determined by making all marbles in the magic family red. The limiting particles view each urn as a two color urn whose limiting probability can be described by a beta distribution [2].

As a consequence, the right limiting particles' limiting probabilities of moving to the right are beta distributions with parameters $\frac{(B_0+M_0)}{2}$ and $\frac{R_0}{2}$ where R_0, B_0, and M_0 are the initial numbers of red, blue, and magic marbles respectively in an urn. At all positions to the right of r_0, the parameters will be $\frac{a}{2}$ and $\frac{a+1}{2}$. This is also true for the left movement of the left limiting particle to the left of l_0.

Following is a description of the movement of the right limiting particles, which can be extended to the left limiting particles. The right limiting particles all begin at r_0, and may not move to the left of r_0. If a particle and its right limiting particle are at the same location, the right limiting particle moves according to the particle's draw. If the particle draws a red marble not in the magic family, both the particle and the right limiting particle move left. If the particle draws a red marble in the in the magic family, the particle moves left, but the right limiting particle moves right. If the particle draws any kind of blue marble, both the particle and the right limiting particle move right. This corresponds to the right limiting particle's perception that magic family marbles are blue. If the right limiting particle is not at the same position as the particle it limits, it moves according to the limiting probabilities described above.

B.3. **Multiple particles at the same position.** Our description of the system appears to fall apart when more than one particle is at a single location. Not only is it unclear how many magic marbles turn red or blue, but the newly arrived particle has increased an edge-weight by one. This weight increase is not represented in the number of marbles in the urn, as marbles are only added when a particle leaves.

Because of the limiting particles, we can't increase the number of marbles when another particle arrives, as this would not correspond to a Polya's urn. We do, however, need to represent the weights correctly for the actual particles. To solve this problem, we can make the marbles in the urn represent the correct weights in the following way:

(1) If b_i through b_{i+m} are at an urn, $(k-(i+m))$ magic marbles turn red, and $i-1$ marbles turn blue. [5] Note that $m+1$ particles are at the urn, and m magic marbles remain unassigned.
(2) The m left over marbles will each be a fraction red and the rest blue in such a way that the total amount of red to blue will correspond to the ratio of edge-weights. [6]
(3) The fraction of the magic marbles that will be turned blue or red can be determined in the following way:
   - The difference between the correct edge weights and red and blue marbles in the urn (ignoring the m unassigned magic marbles) always happens to be m. This fact is easy to check with some simple algebra.
     For example, if the number of red marbles (including the assigned magic marbles) is R_n and the number of blue marbles (including the assigned magic marbles) is B_n, then the actual weight on the left edge is R_n+m and the actual weight on the right edge is B_n+m.
   - If x is the fraction of an unassigned magic marble that will turn red, we have
     $$\frac{R_n+m}{B_n+m}=\frac{R_n+mx}{B_n+m(1-x)}$$
     We have m because this is the number of unassigned magic marbles. Solving this equation gives
     $$x=\frac{B_n+m}{R_n+B_n+2m} \qquad \text{and} \qquad (1-x)=\frac{R_n+m}{R_n+B_n+2m}$$
(4) Making $\frac{B_n+m}{R_n+B_n+2m}$ of each unassigned magic marble red and $\frac{R_n+m}{R_n+B_n+2m}$ of each unassigned magic marble blue, we have attained the correct ratio of red to blue to match the ratio of the left and right edge-weights without increasing the number of marbles.

B.4. **Conclusion.** There are a few things to note about this method. We have correctly defined the particles' movements to correspond exactly to an ERRW. Also, each particle's right limiting particle must stay to the right of r_0, and each urn to the right of r_0 begins with a leftward inclination according to the right limiting particles: $a+1$ red marbles and a blue marbles. Similarly, each particle's left limiting particle must stay to the left of l_0, and each urn to the left of l_0 begins with a rightward inclination according to the left limiting particles: $a+1$ blue marbles and a red marbles.

[5] In other words, the magic marbles in common that b_i and b_{i+m} would turn red and the magic marbles in common that b_i and b_{i+m} would turn blue remain red or blue.

[6] This can be more easily understood by thinking of pounds of sand instead of marbles. A blue marble would correspond to a pound of blue sand. A magic marble would correspond to a pound of magic sand. If a particle chooses a red or blue grain of sand, it adds two pounds of the corresponding color of sand. So when a fraction of a pound of magic sand is made red, and the rest of the pound is made blue, choosing a magic red grain of sand still corresponds to adding two pounds of magic family red sand.

For the proof of recurrence for this method, we can use Lemma 3.3. Instead of using the particles themselves in the application of Lemma 3.3 as in the more general method, we use the limiting particles. Because the limiting distributions for the limiting particles to the right of r_0 and to the left of l_0 are beta distributions, we do not need to use coupling when applying this lemma.

REFERENCES

[1] COPPERSMITH, D., DIACONIS, P. (1987). *Unpublished manuscript*.

[2] KOVCHEGOV, Y. (2008). Multi-particle processes with reinforcements. *Journal of Theoretical Probability*. 21(2), 437-448.

[3] MERKL, F., AND ROLLES, S.W.W. (2005). Edge-reinforced random walk on a ladder. *The Annals of Probability*. 33(6), 2051-2093.

[4] PAKES, A.G. (1969). Some conditions for ergodicity and recurrence of Markov chains. *Operational Research*. 17, 1058-1061.

[5] PEMANTLE, R. (1988). Phase transition in reinforced random walk and RWRE on trees. *The Annals of Probability*. 16(3), 1229-1241.

[6] ROLLES, S.W.W. (2006). On the recurrence of edge-reinforced random walk on $\mathbb{Z} \times G$. *Probability Theory and Related Fields*. 135(2), 216-264.

UNIVERSITY OF CHICAGO
E-mail address: lizzied@uchicago.edu

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
E-mail address: jhansel@mit.edu