Wesleyan University The Honors College

Once-Reinforced Random Walks on $\mathbb{Z}$ and The Doubly Infinite Ladder

by

Rebecca Feiden

Class of 2008

A thesis submitted to the
faculty of Wesleyan University
in partial fulfillment of the requirements for the
Degree of Bachelor of Arts
with Departmental Honors in Mathematics

Middletown, Connecticut April, 2008

Abstract

The once-reinforced random walk is a type of edge-reinforced random walk in which edges have two possible weights. If an edge is familiar to the walker it has weight β and if it is unfamiliar it has weight 1. This paper proves that the once-reinforced random walk on $\mathbb{Z}$ is recurrent for $\beta > 0$ and recurrent on a doubly infinite ladder, $\{\mathbb{Z} \times \{1, 2\}\}$, for $\beta > 0$.

Acknowledgments

Many thanks to my friends, family, and professors who have supported me throughout this project. In particular, I would like to thank Professor Keane for his patience, encouragement, and guidance. I have been honored to work with such a creative and dedicated advisor. In addition, I would like to thank Sam Gottlieb for her willingness to listen and discuss.

Contents

Ch. 1: Introduction

In this paper I shall discuss once-reinforced random walks on $\mathbb{Z}$ and the doubly infinite ladder, $\{\mathbb{Z} \times \{1,2\}\}$. I begin by providing a summary of the research on simple random walks and some of the many variations and applications of random walks that have been studied.

A simple random walk is a sequence of sums of independent identically distributed random variables. In 1921, George Pólya proved an important theorem about simple random walks on $\mathbb{Z}^d = \underbrace{\{\mathbb{Z} \times \mathbb{Z} \times \cdots \times \mathbb{Z}\}}_{d}$. The graph of $\mathbb{Z}^d$ is made up of vertices

$$V = \{(a_1, a_2, \ldots, a_d)\colon a_i \in \mathbb{Z} \ \forall \ i\}$$

and edges

$$\begin{aligned} E \quad = \quad & \big\{\{(a_1, a_2, \ldots, a_d), (b_1, b_2, \ldots, b_d)\}\colon (a_1, a_2, \ldots, a_d), (b_1, b_2, \ldots, b_d) \in V \\ & \text{and } |(a_1, a_2, \ldots, a_d) - (b_1, b_2, \ldots, b_d)| = 1\big\}. \end{aligned}$$

The sum of the random variables, $S_n = (x_0 + x_1 + \cdots + x_n)$ describes the position of a walker on the graph at each time $n \geq 0$. Intuitively, we imagine a walker that begins at some vertex, $v_0 \in V$. At each time $n \geq 0$, the walker moves one step along any of the $2d$ edges connected to v_n. The probability of choosing each of these edges is always $\frac{1}{2d}$. Pólya proved that a random walk is *recurrent*; that is, the walker returns to its starting position at some time $t < \infty$ with probability 1, on $\mathbb{Z}^d$ for $d = 1, 2$. However, for $d \geq 3$ the walk is *transient*, meaning that there is a positive probability of escape [1]. Since Pólya's theorem, several variations of the simple random walk on $\mathbb{Z}^d$ have been studied. One of the major areas of research has involved random walks with different probability schemes.

In 1987, Persi Diaconis described the edge-reinforced random walk. Diaconis defined the *weight* of an edge as one plus the number of times that the edge has been crossed from either direction. At each time $t \geq 0$, the probability of crossing some edge $e \in E$ is equal to the weight of edge e divided by the sum of the weights of all edges connected to vertex v_t. So at time $t = 0$ all edges have weight 1 and there is an equal probability of crossing any of the $2d$ edges connected to vertex v_0. Diaconis proved that for a finite graph $G = \{V, E\}$, an edge-reinforced random walk is a mixture of Markov chains. While many attempts have been made to determine whether an edge-reinforced random walk on $\mathbb{Z}^d$ is recurrent, the question remains open [2].

Another type of reinforced random walk first proposed by Diaconis is the vertex-reinforced random walk. It is defined similarly to the edge-reinforced random walk except that instead of edges carrying weights, vertices are assigned weights for each time $t \geq 0$ depending on how many times the vertex has been visited. Pierre Tarrès, with the help of earlier proofs by Robin Pemantle and Stanislav Volkov, showed that an vertex-reinforced random walk on $\mathbb{Z}$ will eventually get stuck within five adjacent vertices [3].

Besides reinforcement schemes, several other types of biased random walks have been studied. For example, in a persistent random walk, the walker's next step is dependent on the direction of its previous step. This process, like the edge-reinforced and vertex-reinforced random walks, requires knowledge of past events. With reinforced random walks, the walker's decision at time t is dependent on all events leading up to time t, but in a persistent random walk, the walker's decision is only dependent on what happened at time $t-1$. Ultimately, the path of the walker and the possibility of recurrence in a persistent random walk is dependent on some probability,

$$Q = P(\text{direction of motion at time } t = \text{direction of movement at time t-1})$$

If $Q = 1$ or $Q = 0$ the walk is not random at all, but completely known [4].

The self-avoiding random walk is another type of biased random walk. This process is a simple random walk except that the walker avoids vertices that have already been visited. Paul Flory used this as a model to study the size of randomly coiled linear polymers [4]. Polymers are made up of monomers. If a polymer is made up of monomers that, because of their chemical make up can only make two connections, it is a linear polymer [5]. Because polymers generally avoid intersections, the self-avoiding random walk serves as a good model for the random shape of a coiled linear polymer. Ultimately, Flory was able to determine a relationship between the size of a randomly coiled polymer and the expected number of monomers contained in it [4]. Many other physicists and chemists have used self-avoiding random walks to continue to study the properties of polymers [5].

Polymer shapes are the not the only situation for which random walks provide a good model. As it turns out, random walks occur throughout nature. Brownian motion describes the movement of particles suspended in gas or liquid [6]. Under certain conditions, a random walk on a graph is an accurate model of this type of motion [7]. Simple random walks also model many types of electric networks. In this case, the current is analogous to the random walker and resistors create probabilities of moving through different wires or edges [1]. Recently, reinforced random walks have been used to model the growth of tumors [8] and self-organization and learning behaviors [3]. As we continue to study and describe the properties of random walks, we will be able to more accurately model and understand these and other natural events.

In my study of random walks I decided to focus on a simplified version of the edge-reinforced random walk, called the once-reinforced random walk. The once-reinforced random walk is similar to the edge-reinforced random walk in that at each time $t > 0$ the walker's move depends on what has happened up to time t. While the weight of each edge in an edge-reinforced random walk depends on how many times the edge has been crossed, the weight of an edge in a once-reinforced random walk only depends on whether the edge has ever been crossed. We define the weight of edge $e \in E$ to be β if it has been crossed and 1 if it has not been crossed. Once an edge has weight β it will remain β forever. I will begin by proving that a once-reinforced random walk on $\mathbb{Z}$ is recurrent for all $\beta > 0$. In an effort to begin to understand once-reinforced random walks on $\mathbb{Z}^2$, I will then describe the once-reinforced random walk on a doubly infinite ladder, $\{\mathbb{Z} \times \{1, 2\}\}$. I will prove that the once-reinforced random walk on the doubly infinite ladder is recurrent for all $\beta \geq 1$. This result was originally proved by Thomas Sellke [9]. My proof, though similar to Sellke's has one key difference. I allow for the weight of the vertical edges to change, resulting in a compensator that has five possible values instead of three possible values. This is helpful in understanding reinforced random walks on graphs with a wider range of y-values.

It is important to note that there are several different definitions of recurrence. Pólya's original definition required that the walker visit each vertex of the graph. This turns out to be equivalent to requiring that the walker return to its starting point at some time $t < \infty$ [1]. For a random walk on $\mathbb{Z}$, we assume that the walker beings in positive territory or his first step is from zero to one. If we can prove that the walker returns to zero then we can make a symmetric argument for the case when the walker begins in negative territory. On the doubly infinite ladder the walker can move vertically first. We know, however, that the probability of the walker moving up and down forever converges to zero. Therefore at some time $t \geq 0$ the walker must move either left or right. I will prove that a random walks on the graphs of $\mathbb{N}$ and $\{\mathbb{N} \times \{1, 2\}\}$ are recurrent with the understanding that this implies recurrence on $\mathbb{Z}$ and the doubly infinite ladder.

Ch. 2: Once-Reinforced Random Walk on $\mathbb{Z}$

2.1 Definitions

Let N be a positive integer and $\beta > 0$. Let $G = \{V, E\}$ be a graph with vertices $V = \{0, 1, \dots, N\}$ and edges $E = \Big\{\{0,1\}, \{1,2\}, \dots, \{N-1, N\}\Big\}$. A *path* is a sequence

$$x = (x_0, x_1, \dots)$$

with $x_t \in V$ for each $t \geq 0$. For each $x_t \in \{1, 2, \dots, N-1\}$

$$x_{t+1} = x_t + 1 \text{ or } x_{t+1} = x_t - 1.$$

If $x_t = 0$ or $x_t = N$, then we require that $x_{t+1} = x_t$. We define the stopping time

$$\tau = \min\{t \colon x_t = 0 \text{ or } N\}.$$

Paths can be infinite, that is x_t is defined $\forall\ t \in \mathbb{N}$, or finite, that is $x = (x_0, x_1, \dots, x_n)$ for some $n \in \mathbb{N}$. We call the length of the finite path $\text{length}(x) = n$. Intuitively, we imagine a walker who begins at time $t = 0$ at some position $x_0 \in \{1, 2, \dots, N-1\}$ and at each time $t > 0$ moves one step left or one step right. If at some time $t > 0$ the walker reaches the position $x_t = 0$ or N, then the walker stops and remains at that position forever. Thus, a finite path is a piece of an infinite path that has some length n.

Define

$$\Omega = \{x = (x_0, x_1, \dots) : x_t \text{ is defined } \forall\, t \in \mathbf{N}\}$$

to be the set of infinite paths. Define

$$\Omega^f = \{x \in \Omega : \exists\, t \in \mathbf{N} : x_t = 0 \text{ or } x_t = N\}$$

to be the set of infinite paths in which the walker reaches zero or N. For each $x \in \Omega^f$, $\tau(x) < \infty$. So, Ω^f is a countable set.

A *weight configuration* is a mapping

$$w_t \colon E \to \mathbb{R}^{+*} = \{r \in \mathbb{R} \colon r > 0\}$$

that defines the weight of an edge $e \in E$ for each time $t \geq 0$.

Define

$$x_t^m = \max\{x_i \colon 0 \leq i \leq t\}$$

to be the maximum position reached by the walker up to some time t. With a once-reinforced random walk, we are interested in understanding the impact of the walker being familiar with certain edges of the graph. Therefore, we use the weight configuration to label the edges that the walker knows. At time $t = 0$, the walker is at position $x_0 \in V$ and is familiar with all edges of the graph up to a fixed vertex l, where $0 \leq x_0 \leq l < N$. To designate the boundary between known and unknown territory, we define

$$x_t^* = \begin{cases} l & \text{if } x_t^m < l \\ x_t^m & \text{otherwise.} \end{cases}$$

So, we will consider weight configurations of the form

$$w_t(e) = \begin{cases} \beta & \text{for } e = \{i-1, i\},\ 1 \leq i \leq x_t^* \\ 1 & \text{for } e = \{i-1, i\},\ x_t^* < i \leq N. \end{cases}$$

Intuitively this means that any edge that has been crossed has weight β while all other edges have weight 1. It is clear that $0 \leq x_t \leq x_t^*$ for all $t \geq 0$.

For all times $0 < t < \tau$

$$P(x_{t+1} = x_t + 1) = \frac{w(x_t, x_t + 1)}{w(x_t - 1, x_t) + w(x_t, x_t + 1)}$$

and

$$P(x_{t+1} = x_t - 1) = \frac{w(x_t - 1, x_t)}{w(x_t - 1, x_t) + w(x_t, x_t + 1)}.$$

So, for each time $0 \leq t < \tau$ where $x_t \neq x_t^*$, the walker moves left and right with equal probability. When $x_t = x_t^*$ the walker stays in known territory with probability $\frac{\beta}{1+\beta}$ and goes into unknown territory with probability $\frac{1}{1+\beta}$. We can look at β as the proportional risk of moving into unknown territory. For $\beta > 1$, the walker avoids risk. We call this *positive reinforcement*. For $0 < \beta < 1$ the walker seeks risk. This is called *negative reinforcement*. When $\beta = 1$ there is no reinforcement we have a simple random walk on $\mathbb{Z}$ which is recurrent [1].

For some $x \in \Omega^f$ with $x_0 = k$, stopping time τ and $\beta > 0$, the probability of path x is given by

$$P^{k,l,\beta}(x) = \left(\frac{1}{2}\right)^{\tau(x)-m(x)} \left(\frac{1}{1+\beta}\right)^{m^+(x)} \left(\frac{\beta}{1+\beta}\right)^{m^-(x)}$$

where,

$$\begin{aligned} m^+(x) &= card\{s : 0 \le s < \tau(x),\ x_s = x_s^* \text{ and } x_{s+1} = x_s + 1\} \\ m^-(x) &= card\{s : 0 \le x < \tau(x),\ x_s = x_s^* \text{ and } x_{s+1} = x_s - 1\} \\ m(x) &= m^+(x) + m^-(x). \end{aligned}$$

It can be shown that with probability one, any infinite path must eventually reach zero or N [10]. So,

$$\sum_{x\in\Omega^f} P^{k,l,\beta}(x) = 1.$$

Therefore, $P^{k,l,\beta}$ is a probability measure on the countable set Ω^f.

2.2 An algebraic proof of recurrence

Consider the probability

$$\begin{aligned} P(k,l) &= P(\text{the walker reaching N before 0} \mid x_0 = k \text{ and } x_0^* = l) \\ &= \sum_{[x\in\Omega^f : \tau(x)=N]} P^{k,l,\beta}(x). \end{aligned}$$

Then, $P(k,l)$ is well defined and by standard arguments the walker will reach either zero or N in finite time with probability one [10]. There are two cases:

Case 1 If $0 < k < l$ then

$$P(k,l) = \frac{P(k+1,l) + P(k-1,l)}{2}. \tag{2.1}$$

Case 2 If $k = l$ then

$$P(l,l) = \left(\frac{\beta}{1+\beta}\right) P(l-1,l) + \left(\frac{1}{1+\beta}\right) P(l+1,l+1). \tag{2.2}$$

At the boundary,

$$P(0,l) = 0 \text{ and } P(N,N) = 1. \qquad (0 \leq l \leq N)$$

Rearranging equation, 2.1 we have

$$\begin{aligned} P(k+1,l) - P(k,l) &= P(k,l) - P(k-1,l) \\ &= P(k-1,l) - P(k-2,l) \\ &= P(k-2,l) - P(k-3,l) \\ &\vdots \\ &= P(1,l) - P(0,l). \end{aligned}$$

By applying the boundary condition, $P(0,l) = 0$, we have

$$P(k+1,l) - P(k,l) = P(1,l). \tag{2.3}$$

So, $P(1,l)$ is the difference in probability, between two adjacent points, of reaching N before zero. Since there are l steps between zero and l, we can write case 2 as

$$P(l,l) = lP(1,l). \tag{2.4}$$

Similarly,

$$P(l-1,l) = (l-1)P(1,l) \tag{2.5}$$

and,

$$P(k,l) = kP(1,l). \tag{2.6}$$

By combining equations 2.4 and 2.5 we have

$$P(l-1,l) = \left(\frac{l-1}{l}\right) P(l,l). \tag{2.7}$$

Combining equations 2.2 and 2.7 we get

$$P(l,l) = \left(\frac{\beta}{1+\beta}\right)\left(\frac{l-1}{l}\right) P(l,l) + \left(\frac{1}{1+\beta}\right) P(l+1,l+1).$$

Therefore,

$$\begin{aligned} P(l+1,l+1) &= P(l,l)\left[(1+\beta) - \frac{\beta(l-1)}{l}\right] \\ &= P(l,l)\left[\frac{l+\beta}{l}\right]. \end{aligned} \tag{2.8}$$

Consider the case where the walker is already at N and apply the boundary condition $P(N,N) = 1$.

$$\begin{aligned}
1 = P(N,N) &= P(N-1,N-1)\left[\frac{N-1+\beta}{N-1}\right] \\
&= P(N-2,N-2)\left[\frac{N-2+\beta}{N-2}\right] \times \left[\frac{N-1+\beta}{N-1}\right] \\
&\vdots \\
&= P(1,1)\left[\frac{1+\beta}{1}\right] \times \left[\frac{2+\beta}{2}\right] \times \cdots \times \left[\frac{N-1+\beta}{N-1}\right]
\end{aligned}$$

So,

$$\left[\prod_{m=1}^{N-1} \frac{m+\beta}{m}\right] P(1,1) = 1.$$

Therefore,

$$P(1,1) = \prod_{m=1}^{N-1} \frac{m}{m+\beta}. \tag{2.9}$$

The event that the walker never returns to zero is contained in the event that the walker reaches N before zero. Therefore, the probability that the walker never returns to zero must be less than or equal to the probability that the walker reaches N before zero.

Now, we evaluate the limit of equation 2.9 as N goes to infinity to examine recurrence on an infinite line.

$$\lim_{N\to\infty} \prod_{m=1}^{N-1} \frac{m}{m+\beta} = \lim_{N\to\infty} \frac{1}{\prod_{m=1}^{N-1}(1+\frac{\beta}{m})}$$

Since $\sum_{m=1}^{\infty} \frac{1}{m} \to \infty$, we know that $\prod_{m=1}^{\infty}(1 + \frac{\beta}{m}) \to +\infty$. Thus,

$$\lim_{N\to\infty} \prod_{m=1}^{N-1} \frac{m}{m+\beta} = 0. \tag{2.10}$$

So, on an infinite line, with $k = l = 1$, the probability of never reaching position $x = 0$ is zero. Now we will prove the same result for any k and l where $0 \leq k \leq l < N$.

From equation 2.8 we have

$$\begin{aligned}
P(l,l) &= P(l-1,l-1)\left[\frac{l-1+\beta}{l-1}\right] \\
&= P(l-2,l-2)\left[\frac{l-2+\beta}{l-2}\right] \times \left[\frac{l-1+\beta}{l-1}\right] \\
&\vdots \\
&= P(1,1)\left[\frac{l+\beta}{1}\right] \times \cdots \times \left[\frac{l-2+\beta}{l-2}\right] \times \left[\frac{l-1+\beta}{l-1}\right].
\end{aligned}$$

So,

$$P(l,l) = P(1,1)\left[\prod_{m=1}^{l-1} \frac{m+\beta}{m}\right]. \tag{2.11}$$

Combining equations 2.9 and 2.11 we get

$$\begin{aligned}
P(l,l) &= \left[\prod_{m=1}^{N-1} \frac{m}{m+\beta}\right]\left[\prod_{m=1}^{l-1} \frac{m+\beta}{m}\right] \\
&= \left[\prod_{m=l}^{N-1} \frac{m}{m+\beta}\right].
\end{aligned} \tag{2.12}$$

Combining equations 2.4 and 2.6 we get

$$P(k,l) = \frac{k}{l} P(l,l).$$

Therefore,

$$P(k,l) = \frac{k}{l} \left[\prod_{m=l}^{N-1} \frac{m}{m+\beta} \right]. \tag{2.13}$$

As shown earlier,

$$\lim_{N\to\infty} \prod_{m=l}^{N-1} \frac{m}{m+\beta} = 0.$$

So $P(k,l) = 0$. The probability of never reaching zero is less than or equal to the probability of reaching N before zero. So, we know that for an infinite line, the walker must return to position $x = 0$ at some time $0 < t < \infty$. Thus, we have proved the following theorem.

Theorem 2.1

A once reinforced random walk on $\mathbb{Z}$ is recurrent for any reinforcement $\beta > 0$, starting position k, and initial familiarity l where $0 \leq |k| \leq |l| < N$.

2.3 A proof using martingales

Define an indicator function:

$$R_n = \begin{cases} 1 & \text{if } x_n = x_n^* \qquad (n > 0) \\ 0 & \text{if } x_n < x_n^*. \end{cases}$$

Define the *compensator* for each time $0 \leq k < \tau$

$$C_k = \left(\frac{1-\beta}{1+\beta}\right) R_{k-1} \qquad (k \geq 1).$$

The compensator, C_n is $\mathcal{F}_{n-1}$ measurable so it is a previsible process. For each time $0 < t \leq \tau$, the compensator is the expected change in the walkers position. When $x_t < x_t^*$ there is equal probability of moving the same distance in each direction. So for all t where $x_t < x_t^*$, the compensator is zero. However, when $x_t = x_t^*$ the walker will be more likely to move either left or right depending on the value of β. Define the *total compensator* to be the compensator summed over all times up to τ.

$$C = \sum_{k=1}^{\tau} C_k.$$

Let $B_n = x_n - \sum_{k=1}^{n} C_k$ for $n \geq 0$. B_n is a martingale, x_0 at zero. By martingale theory, $E(B_\tau) = E(B_0) = x_0$. So,

$$\begin{aligned} x_0 &= E\left(x_\tau - \sum_{k=1}^{\tau} C_k\right) = E(x_\tau) - E\left(\sum_{k=1}^{\tau} C_k\right) \qquad (2.14) \\ &= P(X_\tau = N) \times N + P(X_\tau = 0) \times 0 - E\left(\sum_{k=1}^{\tau} C_k\right) \end{aligned}$$

Let $P(x_\tau = N) = \alpha$. So, $P(x_\tau = 0) = 1 - \alpha$. Therefore,

$$x_0 = \alpha N - E(C). \qquad (2.15)$$

There are three cases:

Case 1, $\beta > 1$, positive reinforcement

Then, $C_k \leq 0 \; \forall \; k$. Therefore, $C \leq 0$. From equation 2.14, $x_0 = E(x_\tau) - C$. Since $E(x_0) = x_0$, then

$$E(x_0) \geq E(x_\tau).$$

Therefore, x is a supermartingale [10]. This means that on average x decreases. So the walker must return to position $x = 0$ and its starting position. So, for $\beta > 1$, a once-reinforced random walk on $\mathbb{Z}$ is recurrent.

Case 2, $\beta = 1$, no reinforcement

This is a simple random walk which, as previously mentioned, is recurrent on $\mathbb{Z}$.

Case 3, $0 < \beta < 1$, negative reinforcement

Then $C_k \geq 0 \; \forall \; k$. Therefore, $C \geq 0$. Rearranging equation 2.15 we have

$$\alpha = \frac{x_0}{N} + \left(\frac{1}{N}\right) E(C)$$

We know that at time $0 \leq n < \tau$, if $x_n = x_n^*$, $P(x_{n+1} = X_n + 1) = \frac{1}{1+\beta}$. Therefore, for any vertex $v \geq x_0^*$, the expected number of times that v must be reached before the walker enters unfamiliar territory is $1 + \beta$. Each time n that $x_n = x_n^*$, $C_n = \frac{1-\beta}{1+\beta}$ is added to the total compensator. So the expected total compensation for each vertex $v \geq x_0^*$ is $(1+\beta)C_n = (1-\beta)$. There are at most $N - x_0^*$ vertices that can at some time be the maximum, so we see that

$$E(C) \leq (N - x_0^*)(1 - \beta).$$

Therefore,

$$\begin{aligned}\alpha &\leq \frac{x_0}{N} + \left(\frac{1}{N}\right)(N - x_0^*)(1-\beta) \\ &\leq 1 - \beta + \frac{x_0}{N} - \frac{x_0^*}{N} + \frac{x_0^*\beta}{N}\end{aligned}$$

For large N, we have

$$\alpha \leq 1 - \beta.$$

So $\alpha < 1$. Assume now that the walker reaches N. We can specify another integer $\tilde{N} \gg N$. The probability of reaching $\tilde{N}$ starting at x_0 would be $\alpha^2 \leq (1-\beta)^2$. We continue designating larger and larger $\tilde{N}$. Since $\lim_{N\to\infty}\alpha^N = 0$ for $0 \leq \alpha < 1$, as N goes to infinity, α goes to zero. So for an infinite line, the probability α of the walker reaching some very distant position N before 0 converges to zero. It can be shown then, that at some time $t < \infty$, the walker must return to 0 [10]. Again we have proved Theorem 2.1.

Ch. 3: Once-Reinforced Random Walk On The Doubly Infinite Ladder

3.1 Definitions

Let N be a positive integer and $\beta > 0$. Define a graph $G = \{V, E\}$ with vertices

$$V = \{(a,b)\colon a \in \{0,1,\ldots,N\} \text{ and } b \in \{1,2\}\}$$

and edges

$$\begin{aligned} E &= \Big\{\{(a,b),(c,d)\}\colon (a,b) \text{ and } (c,d) \in V \text{ and either} \\ &\quad b = d \text{ and } |a-c| = 1 \text{ or } a = c \text{ and } |b-d| = 1\Big\}. \end{aligned}$$

A *path* is a sequence

$$(x,y) = ((x_0,y_0),(x_1,y_1),\ldots)$$

with $(x_t,y_t) \in V$ for each $t \geq 0$ such that if $x_t \in \{1,2,\ldots,N-1\}$ then

$$|(x_t,y_t) - (x_{t+1},y_{t+1})| = 1$$

and if $x_t = 0$ or $x_t = N$, then we require that $(x_{t+1},y_{t+1}) = (x_t,y_t)$. Define

$$\tau = \min\{t \ :\ x_t = 0 \text{ or } N\}$$

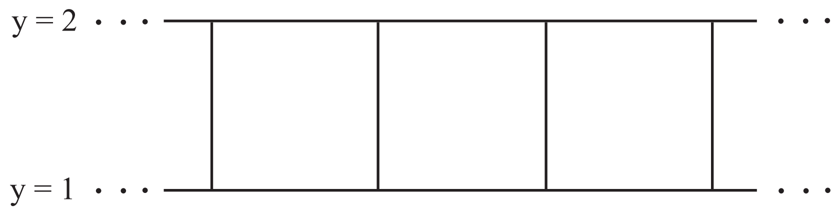

Figure 3.1: The doubly infinite ladder with $x \in \mathbb{Z}$ and $y \in \{1, 2\}$. Here we focus on a section of the doubly infinite ladder with $x \in \{0, 1, \ldots, N\}$.

to be a stopping time. Intuitively, we imagine a walker beginning at some vertex $(x_0, y_0) \in V$ at time $t = 0$. At each time $t > 0$, if $y_t = 1$ the walker moves one step left, right or up, and if $y_t = 2$ the walker moves one step left, right or down. If at some time $t > 0$, $x_t = 0$ or N, the walker stops and remains there forever. As with the a random walk on $\mathbb{Z}$, paths may be infinite or finite. We call the length of a finite path length$(x, y) = n$ and we can see that a finite path is a part of an infinite path. As with the random walk on $\mathbb{Z}$, we define

$$\Omega = \{(x, y) = \big((x_0, y_0), (x_1, y_1), \ldots\big) \colon \forall\, t \geq 0\ \exists\ (x_t, y_t)\}$$

to be the set of all infinite paths. Define

$$\Omega^f = \{(x, y) \in \Omega \colon \exists\ t \colon x_t = 0 \text{ or } N\}$$

to be the set of infinite paths in which the walker reaches the horizontal position zero or N. For each $(x, y) \in \Omega^f$, $\tau(x, y) < \infty$. There are a countable number of paths for which each $\tau \in \mathbb{N}$ is the stopping time. Since the set of all τ is countably infinite, the total number of paths must also be countably infinite. Define $P^{(x_0,y_0),\beta}(x)$ to be the probability of some path x. It can be shown that with probability one, any infinite path must eventually reach zero [10]. Therefore,

$$\sum_{x \in \Omega^f} P^{(x_0,y_0),\beta}(x) = 1.$$

So, $P^{(x_0,y_0),\beta}$ is a probability measure on the countable set Ω^f.

With a once reinforced random walk we are interested in the impact of a walker being familiar with certain edges of the graph. In order to keep track of edges that have been crossed, and are therefore familiar, we define a *weight configuration* to be a mapping

$$w_t \,:\; E \longrightarrow \mathbb{R}^{+*} = \{r \in \mathbb{R} \colon r > 0\}.$$

We call $w_t(e) \in \mathbb{R}^{+*}$ the weight of edge $e \in E$ at time $t \geq 0$. We separate the edges into two categories, vertical edges, E_V and horizontal edges, E_H. Note, $E = E_V + E_H$ and E_V and E_H are defined explicitly as follows:

$$E_V = \big\{\{(a,b),(c,d)\} \colon (a,b) \text{ and } (c,d) \in V, |b-d| = 1, \text{ and } a = c\big\};$$

$$E_H = \big\{\{(a,b),(c,d)\} \colon (a,b) \text{ and } (c,d) \in V, |a-c| = 1, \text{ and } b = d\big\}.$$

For all edges, $\{(a,b),(c,b)\} \in E_H$

$$w_t\{(a,b),(c,b)\} = \begin{cases} \beta & \text{if } \exists\, n \leq t \colon \{x_n, x_{n+1}\} = \{a,c\} \text{ or } \{c,a\} \\ 1 & \text{otherwise.} \end{cases}$$

For all edges, $\{(a,b),(a,c)\} \in E_V$

$$w_t\{(a,b),(a,c)\} = \begin{cases} \beta & \text{if } \exists\, n \leq t \colon \{y_n, y_{n+1}\} = \{b,c\} \text{ or } \{c,b\} \\ 1 & \text{otherwise.} \end{cases}$$

So intuitively, an edge has weight 1 if it is unfamiliar and β if it is familiar. As with the random walk on $\mathbb{Z}$, β as the proportional risk of moving into unknown territory. If $\beta > 1$ we have positive reinforcement. For $0 < \beta < 1$ we have negative reinforcement.

For all times $0 \leq t < \tau$ where $y_t = 1$ define

$$\tilde{w}_t^1 = w_t\{(x_t-1,1),(x_t,1)\} + w_t\{(x_t,1),(x_t+1,1)\} + w_t\{(x_t,1),(x_t,2)\}.$$

So,

$$\begin{aligned} P(x_{t+1} &= x_t + 1 | y_t = 1) = \frac{w_t\{(x_t,1),(x_t+1,1)\}}{\tilde{w}_t^1}; \\ P(x_{t+1} &= x_t - 1 | y_t = 1) = \frac{w_t\{(x_t-1,1),(x_t,1)\}}{\tilde{w}_t^1}; \\ P(y_{t+1} &= y_t + 1 | y_t = 1) = \frac{w_t\{(x_t,1),(x_t,2)\}}{\tilde{w}_t^1}. \end{aligned}$$

For all times $0 \leq t < \tau$ where $y_t = 2$ define

$$\tilde{w}_t^2 = w_t\{(x_t-1,2),(x_t,2)\} + w_t\{(x_t,2),(x_t+1,2)\} + w_t\{(x_t,2),(x_t,1)\}.$$

So,

$$
\begin{aligned}
P(x_{t+1} &= x_t + 1 | y_t = 2) = \frac{w_t\{(x_t, 2), (x_t + 1, 2)\}}{\tilde{w}_t^2};\\
P(x_{t+1} &= x_t - 1 | y_t = 2) = \frac{w_t\{(x_t - 1, 2), (x_t, 2)\}}{\tilde{w}_t^2};\\
P(y_{t+1} &= y_t + 1 | y_t = 2) = \frac{w_t\{(x_t, 2), (x_t, 1)\}}{\tilde{w}_t^2}.
\end{aligned}
$$

Define the *horizontal compensator* for each time $0 < k \leq \tau$

$$C_n = 1 \times P(x_{n+1} = x_n + 1) + (-1) \times P(x_{n+1} = x_n - 1)$$

to be the expected change in horizontal position at time k. The compensator, C_n is $\mathcal{F}_{n-1}$ measurable so it is a previsible process.

Define the *total horizontal compensator* to be the compensator summed over all times up to τ.

$$C = \sum_{k=1}^{\tau} C_k$$

Let $Z_n = x_n - \sum_{k=1}^{n} C_k$ for $n \geq 0$. Z_n is a martingale, x_0 at time zero.

Let $\tilde{x}$ be a designated horizontal coordinate such that $x_0 < \tilde{x} < N$. Define

$$
\begin{aligned}
\sigma_1 &= \min\{t\colon w_t\{(\tilde{x} - 1, 1), (\tilde{x}, 1)\} = \beta\}\\
\sigma_2 &= \min\{t\colon w_t\{(\tilde{x} - 1, 2), (\tilde{x}, 2)\} = \beta\}\\
\mu_1 &= \min\{t\colon w_t\{(\tilde{x}, 1), (\tilde{x} + 1, 1)\} = \beta\}\\
\mu_2 &= \min\{t\colon w_t\{(\tilde{x}, 2), (\tilde{x} + 1, 2)\} = \beta\}\\
\theta &= \min\{t\colon w_t\{(\tilde{x}, 1), (\tilde{x}, 2)\} = \beta\}
\end{aligned}
$$

We know that if $\tilde{x}$ is not reached, the total compensation for the horizontal position $\tilde{x}$ is zero. In addition, this means that the walker never reaches N and at some finite time must return to zero [10]. In this case, the random walk is recurrent. Therefore, we will focus on the situation where the walker reaches almost all horizontal positions. We assume that not only is $\tilde{x}$ reached, but that $\tilde{x}+1$ is reached as well. So, at least one of μ_1 and $\mu_2 < \infty$. Define

$$\Omega^{\tilde{x}} = \{\omega \in \Omega^f : \exists\, t, s < \infty : x_t = \tilde{x} \text{ and } x_s = \tilde{x}+1\}$$

to be the set of all paths that reach horizontal positions $\tilde{x}$ and $\tilde{x}+1$ at some times $t, s < \infty$. Define the following events:

$$\begin{aligned}
A &= \{\omega \in \Omega^{\tilde{x}} : \sigma_1 < \sigma_2, \mu_1, \mu_2, \theta\} \\
B &= \{\omega \in \Omega^{\tilde{x}} : \sigma_2 < \sigma_1, \mu_1, \mu_2, \theta\} \\
C &= \{\omega \in \Omega^{\tilde{x}} : \sigma_1 < \mu_1, \theta\} \\
F &= \{\omega \in \Omega^{\tilde{x}} : \sigma_2 < \mu_2, \theta\} \\
G &= \{\omega \in \Omega^{\tilde{x}} : \mu_2 < \sigma_2, \theta\} \\
H &= \{\omega \in \Omega^{\tilde{x}} : \theta < \sigma_2, \mu_2\} \\
J &= \{\omega \in \Omega^{\tilde{x}} : \sigma_2, \mu_2, \theta = \infty\} \\
K &= \{\omega \in \Omega^{\tilde{x}} : \omega \in A, \omega \notin G, H, J\}
\end{aligned}$$

If horizontal position $\tilde{x}$ is reached at all, then we can see that either event A or event B occurs. These two events are disjoint so

$$\Omega^{\tilde{x}} = A \cup B.$$

Also note that events F, G and H are disjoint and G, H, J and $K \subset A$.

We are now prepared to show that the once-reinforced random walk is recurrent on the doubly infinite ladder.

3.2 Positive Reinforcement

As with the once-reinforced random walk on $\mathbb{Z}$ we will try to calculate the maximum expected total compensation. Again, we would like this value to be negative in order to show that the walker is, on average, pulled toward zero. We know that in order to produce a positive compensator at some vertex, it is necessary for $\mu < \sigma$. As we can see in figure 3.2, there are several scenarios in which the total compensation at horizontal position $\tilde{x}$ will be negative. However, there are some situations in which the total compensation at $\tilde{x}$ is positive. The idea is to calculate the maximum value of the compensation under each scenario and then prove that due to the probability of each of the scenarios, the total expected compensation must be negative. In all cases, the compensator is calculated by finding the expected maximum number of times that a vertex must be reached before a new edge is crossed and then multiplying by the compensator for each of these times. In some cases, we define new probabilities and make claims about their relationships. Though proofs of these claims are not given, the intuitive arguments

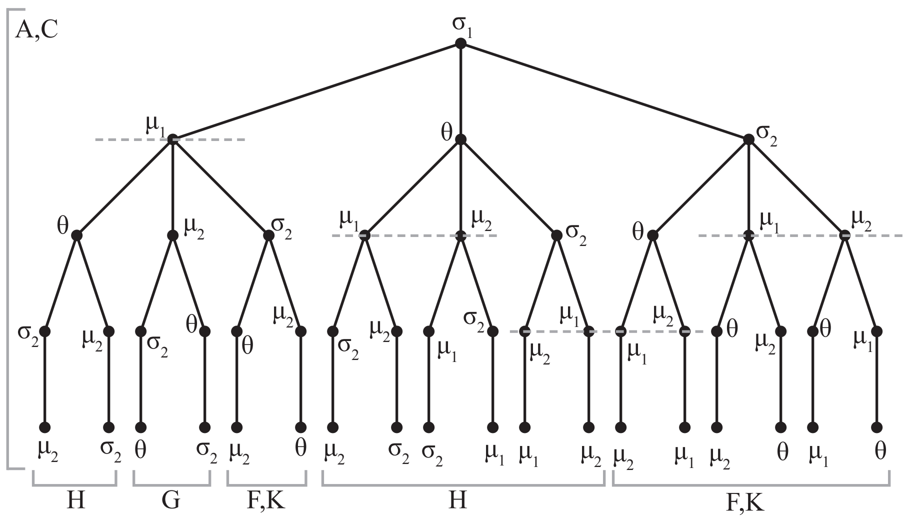

Figure 3.2: All of the possible sequences of stopping times in event A. The dotted lines show the times at which $\tilde{x}$ and $\tilde{x}+1$ have been reached. It is possible for a sequence to stop at any time after the dotted line.

validate their use.

The ladder is symmetric and we assume that $P(y_0 = 1) = P(y_0 = 2) = \frac{1}{2}$. So for any argument made for event A, we can make a symmetric argument for event B. We assume, without loss of generality, that $\sigma_1 < \sigma_2, \mu_1, \mu_2, \theta$, that is event A occurs. This means that the total compensation at vertex $(\tilde{x}, 1)$ must be negative. The only way to produce a positive compensator at vertex $(\tilde{x}, 2)$ is for $\mu_2 < \sigma_2$. This is only possible if either event G or H occurs. We will therefore calculate the expected compensation for horizontal position $\tilde{x}$ for all possible events contained in A.

Case 1 If event G occurs

Then $\sigma_1 < \mu_1 < \mu_2$. At each time t where $\sigma_1 \leq t < \mu_1$ and $(x_t, y_t) = (\tilde{x}, 1)$

$$P((x_{t+1}, y_{t+1}) \neq (x_{t-1}, y_{t-1})) = \frac{2}{2+\beta}$$

and

$$C_t = \frac{1-\beta}{2+\beta}.$$

Therefore, the expected number of times that $(\tilde{x}, 1)$ must be reached before $w\{(\tilde{x}, 1), (\tilde{x}+1, 1)\} = \beta$ is $\frac{2+\beta}{2}$. So, the expected total compensation for vertex $(\tilde{x}, 2)$ is

$$E(C^{(\tilde{x},1)}|G) = \frac{1-\beta}{2}.$$

We know that at time μ_2, $\frac{\beta-1}{2+\beta}$ must be added to the compensator. Since by time μ_2, $(\tilde{x}+1)$ has been reached, it is possible that vertex $(\tilde{x}, 1)$ is never visited again. If $(\tilde{x}, 1)$ is visited again after time μ_2 then, at each time s where $\sigma_1, \mu_1, \mu_2 < s < \sigma_2, \theta$ and $(x_s, y_s) = (\tilde{x}, 2)$

$$P((x_{s+1}, y_{s+1}) \neq (x_{s-1}, y_{s-1})) = \frac{2+\beta}{2}$$

and

$$C_s = \frac{\beta-1}{2+\beta}.$$

It is possible for edges $\{(\tilde{x}, 1), (\tilde{x}, 2)\}$ and $\{(\tilde{x}-1, 2), (\tilde{x}, 2)\}$ to be crossed from another vertex to $(\tilde{x}, 2)$. So the maximum expected wait time for either $w\{(\tilde{x}, 1), (\tilde{x}, 2)\} = \beta$ or $w\{(\tilde{x}-1, 2), (\tilde{x}, 2)\} = \beta$ is $\frac{2+\beta}{2}$. Therefore, at most $\frac{\beta-1}{2} - \frac{\beta-1}{2+\beta}$ is added to the total compensator for $\tilde{x}$. We subtract $\frac{\beta-1}{2+\beta}$ because it has already been added at time μ_2.

If $\sigma_1, \mu_1, , \mu_2, \theta < \sigma_2$ and $\sigma_2 < \infty$ then at each time r where $\theta \leq r < \sigma_2$ and $(x_r, y_r) = (\tilde{x}, 2)$

$$P((x_{r+1}, y_{r+1}) = (\tilde{x} - 1, 2)) = \frac{1}{1 + 2\beta}$$

and

$$C_r = \frac{\beta - 1}{1 + 2\beta}.$$

So, the maximum expected wait time for $w\{(\tilde{x} - 1, 2), (\tilde{x}, 2)\} = \beta$ is $1 + 2\beta$. Therefore, at most $\beta - 1$ is added to the total compensator for $\tilde{x}$. Now, we calculate the expected total compensation for horizontal position $\tilde{x}$ for the case where G occurs.

$$E(C^{\tilde{x}}|G) \leq \frac{1 - \beta}{2} + \frac{\beta - 1}{2 + \beta} + \ell \left(\frac{\beta - 1}{2} - \frac{\beta - 1}{2 + \beta} \right) + av(\beta - 1)$$

where

$$\begin{aligned}
\ell &= P(\text{at least one of } \sigma_2 \text{ and } \theta < \infty \mid \mu_2 < \sigma_2, \theta) \\
a &= P(\sigma_2 < \infty \mid \mu_2 < \sigma_2, \theta) \\
v &= P(\theta < \sigma_2 \mid \mu_2 < \sigma_2, \theta).
\end{aligned}$$

Case 2 If event H occurs

Then $\sigma_1 < \theta$. Note that it is possible for $\mu_1 < \theta$ or $\theta < \mu_1$. At each time t where $\sigma_1 \leq t < \mu_1, \theta$ and $(x_t, y_t) = (\tilde{x}, 1)$

$$P((x_{t+1}, y_{t+1}) \neq (x_{t-1}, y_{t-1})) = \frac{2}{2 + \beta}$$

and

$$C_t = \frac{1-\beta}{2+\beta}.$$

Therefore, the expected number of times that $(\tilde{x}, 1)$ must be visited before $w\{(\tilde{x}, 1), (\tilde{x}, 2)\} = \beta$ or $w\{(\tilde{x}, 1), (\tilde{x}+1, 1)\} = \beta$ is $\frac{2+\beta}{2}$. At time μ_2 or θ, whichever comes first, the expected total compensation is $\frac{1-\beta}{2}$.

If $\theta < \mu_1$ and $\mu_1 < \infty$, then at each time s where $\theta \leq s < \mu_1$ and $(x_s, y_s) = (\tilde{x}, 1)$

$$P((x_{s+1}, y_{s+1}) = (\tilde{x}+1, 1)) = \frac{1}{1+2\beta}$$

and

$$C_s = \frac{1-\beta}{1+2\beta}.$$

Therefore, the expected number of times that $(\tilde{x}, 1)$ must be visited before $w\{(\tilde{x}, 1), (\tilde{x}+1, 1)\} = \beta$ is $1+2\beta$. So, $1-\beta$ is added to the compensator.

We know the compensator for $(\tilde{x}, 2)$ will be zero until either $w\{(\tilde{x}-1, 2), (\tilde{x}, 2)\} = \beta$ or $w\{(\tilde{x}, 2), (\tilde{x}+1, 2)\} = \beta$. After time θ, if the walker returns to $(\tilde{x}, 2)$ there are two possibilities.

If $\mu_2 < \sigma_2$ then at each time r where $\mu_2 \leq r < \sigma_2$ and $(x_r, y_r) = (\tilde{x}, 2)$

$$P((x_{r+1}, x_{r+1}) = (\tilde{x}-1, 2)) = \frac{1}{1+2\beta}$$

and

$$C_r = \frac{\beta-1}{1+2\beta}.$$

Therefore, the expected wait time for $w\{(\tilde{x}-1, 2), (\tilde{x}, 2)\} = \beta$ is $1+2\beta$. So $\beta-1$ must be added to the total compensator for $\tilde{x}$.

If $\sigma_2 < \mu_2$ then at each time n where $\sigma_2 \leq n < \mu_2$ and $(x_n, y_n) = (\tilde{x}, 2)$

$$P((x_{n+1}, x_{n+1}) = (\tilde{x}+1, 2)) = \frac{1}{1+2\beta}$$

and

$$C_n = \frac{1-\beta}{1+2\beta}.$$

Therefore, the expected wait time for $w\{(\tilde{x}, 2), (\tilde{x}+1, 2)\} = \beta$ is $1+2\beta$. So $1-\beta$ must be added to the compensator. Now we calculate the expected total compensation for horizontal position $\tilde{x}$ for the case where event H occurs.

$$E(C^{\tilde{x}}|H) = \frac{1-\beta}{2} + \bar{v}\bar{a}(\beta-1) + b(\beta-1) + c(1-\beta)$$

where

$$\begin{aligned}
\bar{a} &= P(\mu_1 < \infty \mid \sigma_1 < \mu_1, \theta) \\
\bar{v} &= P(\theta < \mu_1 \mid \sigma_1 < \mu_1, \theta) \\
b &= P(\mu_2 < \sigma_2 \mid \theta < \sigma_2, \mu_2 \text{ and } \sigma_2, \mu_2 < \infty) \\
c &= P(\sigma_2 < \mu_2 \mid \theta < \sigma_2, \mu_2 \text{ and } \sigma_2, \mu_2 < \infty).
\end{aligned}$$

Case 3 If event J occurs

Then $\sigma_1 < \mu_1$ and $\theta = \sigma_2 = \mu_2 = \infty$. Then at time t where $\sigma_1 \leq t < \mu_1$ and $(x_t, y_t) = (\tilde{x}, 1)$

$$P((x_{t+1}, y_{t+1}) \neq (x_{t-1}, y_{t-1})) = \frac{2}{2+\beta}$$

and

$$C_t = \frac{1-\beta}{2+\beta}.$$

So the expected wait time for $w\{(\tilde{x},1),(\tilde{x}+1,1)\} = \beta$ is $\frac{2+\beta}{2}$. So the expected total compensation for $\tilde{x}$ for the case where event J occurs is

$$E(C^{\tilde{x}}|J) = \frac{1-\beta}{2}.$$

Case 4 If event K occurs

Then events C and F occur and $\sigma_1 < \sigma_2$. In order for horizontal position $\tilde{x}+1$ to be reached, at least one of μ_1 and $\mu_2 < \infty$. So, at each time t where $\sigma_1 \leq t < \mu_1, \theta$ and $(x_t, y_t) = (\tilde{x}, 1)$ and time s where $\sigma_2 \leq s < \mu_2$ and $(x_s, y_s) = (\tilde{x}, 2)$

$$P((x_{s+1,t+1}, y_{s+1,t+1}) \neq (x_{s-1,t-1}, y_{s-1,t-1})) = \frac{2}{2+\beta}$$

and

$$C_{s,t} = \frac{1-\beta}{2+\beta}.$$

So the expected number of times that $(\tilde{x},1)$ or $(\tilde{x},2)$ must be visited before the walker enters unfamiliar territory is $\frac{2+\beta}{2}$. Both vertices $(\tilde{x},1)$ and $(\tilde{x},2)$ must be reached but only one of vertices $(\tilde{x}+1,1)$ and $(\tilde{x}+1,2)$ must be visited. So, at most, the total compensation for $\tilde{x}$ is $\frac{1-\beta}{2} + \frac{1-\beta}{2+\beta}$.

It is possible for $\theta < \mu_1, \mu_2$. If this is the case, at each time before μ_1 and μ_2 that the walker is in a position to cross edge $\{(\tilde{x},1),(\tilde{x},2)\}$, the probability of doing so is $\frac{1}{2}$. If this happens, the expected wait time for the walker to enter new territory at each of the vertices is $1 + 2\beta$. Each time that walker is at one of

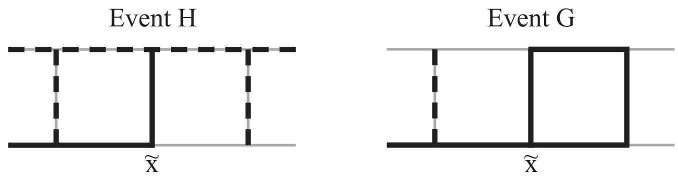

Figure 3.3: Events G and H under the condition that the walker crosses the vertical edge before the horizontal edge. The gray lines show the graph, the solid black lines represent edges that have been crossed, and the dashed black lines represent edges that may have been crossed.

these vertices after θ but before μ_1 and μ_2, the compensation is $\frac{1-\beta}{1+2\beta}$. So, $1-\beta$ is added to the compensator. There are many possible sequences of the stopping times from this point on. However, all of these sequences can only make the compensator more negative. Therefore, we set an upper bound for the expected total compensator for $\tilde{x}$ for the case where none of the events G,H or J occur:

$$E(C^{\tilde{x}}|K) \leq \frac{1-\beta}{2} + \frac{1-\beta}{2+\beta} + \frac{1}{2}(1-\beta).$$

Now we calculate the total expected compensation for horizontal position $\tilde{x}$ for the case where event A occurs:

$$\begin{aligned} E(C^{\tilde{x}}|A) \quad \leq \quad & P(G)\left[\frac{1-\beta}{2} + \frac{\beta-1}{2+\beta} + \ell\left(\frac{\beta-1}{2} - \frac{\beta-1}{2+\beta}\right) + va(\beta-1)\right] + \\ & P(H)\left[\frac{1-\beta}{2} + \bar{v}\bar{a}(1-\beta) + b(\beta-1) + c(1-\beta)\right] + \\ & P(J)\left[\frac{1-\beta}{2}\right] + P(K)\left[1-\beta + \frac{1-\beta}{2+\beta}\right]. \end{aligned} \tag{3.1}$$

As we can see in figure 3.3, there is some symmetry between events G and H.

The probabilities of edge $\{(\tilde{x},1),(\tilde{x}+1,1)\}$ in event H and edge $\{(\tilde{x}-1,2),(\tilde{x},2)\}$ in event G being crossed after the vertical edges are crossed is dependent on the probability that the vertical edge is crossed first and the probability that the horizontal edge is ever crossed. We require that the walker reaches $\tilde{x}+1$. So, in event H, even though edge $\{(\tilde{x},2),(\tilde{x}+1,2)\}$ may not have been crossed, we know that one of the horizontal edges between $\tilde{x}$ and $\tilde{x}+1$ must be crossed at some time. I claim that due to symmetry, $va = \bar{v}\bar{a}$.

In figure 3.4 we see that in event H at time t when the walker is at vertex $(\tilde{x},2)$, the probability of going left is equal to the probability of going right. However, the walker must have already crossed edge $\{(\tilde{x}-1,1),(\tilde{x},1)\}$ but may not have crossed edge $\{(\tilde{x},1),(\tilde{x}+1,1)\}$. So I claim that $c \geq b$. Equation 3.1 simplifies to

$$\begin{aligned} E(C^{\tilde{x}}|A) \;\leq\;& \big(P(H)-P(G)\big)(1-\beta)va + P(H)(c-b)(1-\beta) + P(J)\left[\frac{1-\beta}{2}\right] + \\ & (P(K)\left[1-\beta+\frac{1-\beta}{2+\beta}\right] + \\ & P(G)\left[\frac{1-\beta}{2} - \frac{1-\beta}{2+\beta} - \ell\frac{1-\beta}{2} + \ell\frac{1-\beta}{2+\beta}\right] \end{aligned}$$

Since we want the compensator to be less than or equal to zero and $1-\beta$ is negative, we must prove that

$$\begin{aligned} 0 \;\leq\;& \big(P(H)-P(G)\big)va + P(H)(c-b) + P(J)\frac{1}{2} + P(K)(1+\frac{1}{2+\beta}) + \\ & P(G)\left[\frac{1}{2} - \frac{1}{2+\beta}\right](1-\ell) \end{aligned}$$

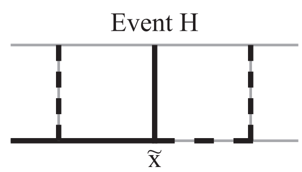

Figure 3.4: Event H at time θ. The gray lines show the graph, the solid black lines represent edges that have been crossed, and the dashed black lines represent edges that may have been crossed.

Each part of the sum must be positive except for $\big(P(H) - P(G)\big)va$ in the case where $P(G) > P(H)$. In order for $E(C^{\tilde{x}}|A) \geq 0$,

$$\begin{aligned} P(G)\left[va - \left(\frac{1}{2} - \frac{1}{2+\beta}\right)(1-\ell)\right] \quad &\geq \quad P(H)va + P(H)(c-b) + \\ &\quad\quad \frac{P(J)}{2} + P(K)\big(1 + \frac{1}{2+\beta}\big). \end{aligned}$$

There are two cases:

Case 1 If $va < \left(\frac{1}{2} - \frac{1}{2+\beta}\right)(1-\ell)$

Then,

$$P(G) \leq \frac{P(H)va + P(H)(c-b) + \frac{P(J)}{2} + P(K)\big(1 + \frac{1}{2+\beta}\big)}{va - \left(\frac{1}{2} - \frac{1}{2+\beta}\right)(1-\ell)}.$$

So, in order for $E(C^{\tilde{x}}|A)$ to be positive, $P(G)$ would have to be negative. Of course this isn't possible because $P(G)$ is between zero and one. Therefore the expected value of the total compensator when event A occurs is

$$E(C^{\tilde{x}}|A) < 0.$$

Case 2 If $va > \left(\frac{1}{2} - \frac{1}{2+\beta}\right)(1-\ell)$

Then,

$$P(G) \geq \frac{P(H)va + P(H)(c-b) + \frac{P(J)}{2} + P(K)\big(1 + \frac{1}{2+\beta}\big)}{va - \left(\frac{1}{2} - \frac{1}{2+\beta}\right)(1-\ell)}.$$

Since $va < 1$ and $\left(\frac{1}{2} - \frac{1}{2+\beta}\right)(1-\ell) < 1$, this would require that

$$P(G) > P(H)va + P(H)(c-b) + \frac{P(J)}{2} + P(K)\big(1 + \frac{1}{2+\beta}\big).$$

We can see in figure 3.2 that there are very few sequences of stopping times in event G while there are many more in sequences for events H and K. Since $\tilde{x}$ is a fixed horizontal position, there are only a finite number of paths for each of these sequences, and so I claim that

$$P(G) \leq P(H)va + P(H)(c-b) + \frac{P(J)}{2} + P(K)\big(1 + \frac{1}{2+\beta}\big).$$

Therefore, the total expected compensation for event A is bounded above by zero. If we were to assume that event B occurred, we could make a symmetrical argument. Therefore, the expected total compensation at vertex $\tilde{x}$ is bounded above by zero. In looking at the entire ladder we see that some vertices may be reached from the right first, but the number of vertices reached from the left will be greater. Thus the total compensation summed over all of the vertices from x_0 to N is

$$C \leq 0.$$

As defined previously, $Z_n = x_n - \sum_{k=1}^{n} C_k$ is a martingale. By martingale theory, $E(Z_\tau) = E(Z_0) = x_0$. So,

$$x_0 = E(x_\tau) - E\left(\sum_{k=1}^{\tau} C_k\right)$$

Therefore, since $E(x_0) = x_0$

$$E(x_0) = E(x_\tau) - C$$

Since $C \leq 0$,

$$E(x_0) \geq E(x_\tau)$$

Following the argument for the once-reinforced random walk on $\mathbb{Z}$ we can prove the following theorem.

Theorem 3.1

A once-reinforced random walk on the doubly infinite ladder is recurrent for any reinforcement $\beta > 1$, and starting position (x_0, y_0).

Bibliography

[1] Doyle, Peter G. and Snell, J. Laurie. *Random Walks and Electric Networks.* United States of America: The Mathematical Association of America, 1984.

[2] Pemantle, Robin. "A Survey of Random Processes With Reinforcement." *Probability Surveys* 2007: 1-79.

[3] Tarrès, Pierre. "Vertex-Reinforced Random Walk On $\mathbb{Z}$ Eventually Gets Stuck On Five Points." *The Annals of Probability.* 2004: pp. 2650-2701.

[4] Rudnick, Joseph and Gaspari, George. *Elements of the Random Walk.* Cambridge, United Kingdom: Cambridge University Press, 2004.

[5] van der Hofstan, R.W. *One-dimensional random polymers.* Amsterdam, The Netherlands: Stichting Mathematisch Centrum, 1998.

[6] Krylov, N. V. *Introduction to the Theory of Diffusion Processes.* The United State of America: American Mathematical Society, 1995.

[7] Kesten, Harry "Frank Spitzer's Work on Random Walk and Brownian Motion." *The Annals of Probability.* Apr. 1993: pp. 593-607.

[8] Sleeman, B.D. and Wallis, I.P. "Tumour induced angiogenesis as a reinforced random walk: modelling capillary network formation without endothelial cell proliferation." *Mathematical Computer Modelling.* 2002: 339–358.

[9] Sellke, Thomas. "Recurrence of reinforced random walk on a ladder." *Electronic Journal of Probability*. 2006: 301–310.

[10] Williams, David. *Probability with Martingales*. New York: Cambridge University Press, 1991.