Electronic Journal of Probability

Vol. 11 (2006), Paper no. 11, pages 301–310.

Journal URL
http://www.math.washington.edu/~ejpecp/

Recurrence of Reinforced Random Walk On a Ladder

Thomas Sellke
Department of Statistics Purdue University
1399 Math Sciences Bldg.
West Lafayette, IN 47907–1399

Abstract

Consider reinforced random walk on a graph that looks like a doubly infinite ladder. All edges have initial weight 1, and the reinforcement convention is to add $\delta > 0$ to the weight of an edge upon first crossing, with no reinforcement thereafter. This paper proves recurrence for all $\delta > 0$. In so doing, we introduce a more general class of processes, termed multiple-level reinforced random walks.

Editor's Note

A draft of this paper was written in 1994. The paper is one of the first to make any progress on this type of reinforcement problem. It has motivated a substantial number of new and sometimes quite difficult studies of reinforcement models in pure and applied probability. The persistence of interest in models related to this has caused the original unpublished manuscript to be frequently cited, despite its lack of availability and the presence of errors. The opportunity to rectify this situation has led us to the somewhat unusual step of publishing a result that may have already entered the mathematical folklore.

Key words: Reinforced random walk, learning, Markov, multiple-level, martingale.

AMS 2000 Subject Classification: Primary: 60J10, 60G50.

Submitted to EJP on September 24, 2001. Final version accepted on January 17, 2006.

1 Introduction and summary

Coppersmith and Diaconis (1987) have initiated the study of a class of processes called reinforced random walks; see also Diaconis (1988). Take a graph with initial weights assigned to the edges. Then define a discrete-time, nearest-neighbor random walk on the vertices of this graph as follows. At each stage, the (conditional, given the past) probability of transition from the current vertex to an adjacent vertex is proportional to the weight currently assigned to the connecting edge. The random walk always jumps, so these conditional transition probabilities sum to 1. The weight of an edge can increase when the edge is crossed, with the amounts of increase depending on the reinforcement convention. The convention most studied by Coppersmith and Diaconis (1987) is to always add $+1$ to the weight of an edge each time it is crossed. In this setting, they show that a reinforced random walk on a finite graph is a mixture of stationary Markov random walks. The mixing measure is given explicitly in terms of the "loops" of the graph.

Pemantle (1988) has studied reinforced random walks with the Coppersmith-Diaconis reinforcing convention on infinite acyclic graphs. Davis (1989) obtained results for nearest-neighbor reinforced random walks on the integers $\mathbb{Z}$ with very general reinforcement schemes.

Consider nearest-neighbor reinforced random walk on the lattice $\mathbb{Z}^2$ of points in $\mathbb{R}^2$ with integer coordinates. All edges between neighboring points are assigned initial weight 1. It seems plausible (perhaps even "obvious") that any spatially homogeneous reinforcement scheme for which the process cannot get stuck forever on a finite set of points will be recurrent, that is, will visit each point of the lattice infinitely often. However, determining whether any such reinforcement scheme leads to recurrence of reinforced random walk on $\mathbb{Z}^2$ has remained open for fifteen years.

Michael Keane (2002) has proposed the following simpler variant: consider nearest-neighbor reinforced random walk on the points of $\mathbb{Z}^2$ with y coordinate 1 or 2 (and starting at $(0, 1)$, say). If one draws in the edges between nearest neighbors, one of course gets an infinite horizontal ladder. Again, all initial weights are taken to be $+1$. For the reinforcement scheme, Keane suggested that edges be reinforced by $\delta = 1$ the first time they are crossed and then never again. This kind of reinforcement is now known as *once-reinforced* random walk; see, e.g., Durrett, Kesten and Limic (2002). This main result of this paper is that Keane's one-time reinforced random walk on a ladder is recurrent for any positive reinforcement parameter δ.

2 Notation and results

Let $(\Omega, \{\mathcal{F}_n\}_{n\geq 0}, \mathbb{P})$ be a standard filtration. Let $d \geq 2$ be an integer and let $\{(X_n, Y_n)\}_{n\geq 0}$ be random variables in $\mathbb{Z} \times \{1, \ldots, d\}$ that are adapted (that is, $(X_n, Y_n) \in \mathcal{F}_n$) and satisfy

(i) $(X_0, Y_0) = (0, 1)$.

(ii) $\{X_n\}_{n=0}^{\infty}$ is a nearest neighbor random motion on $\mathbb{Z}$, i.e., $P\{|X_{n+1} - X_n| = 1\} = 1 \ \ \forall\, n$.

(iii) $P\{X_{n+1} = X_n + 1 | \mathcal{F}_n\} = \frac{W_n(X_n,Y_n)}{W_n(X_n,Y_n)+W_n(X_n-1,Y_n)}$.

Here $W_n(x, y) \in \mathcal{F}_n$ is the weight at time n of the horizontal edge to the right of (x, y), and is defined in our model by $W_n = 1 + \delta \cdot R(x, y, n)$ and $R(x, y, n)$ is the indicator function of the event that the edge to the right of (x, y) has been crossed by time n:

$$R(x, y, n) = I\left\{\exists j < n : \{X_j, X_{j+1}\} = \{x, x+1\} \text{ and } Y_j = y\right\} . \tag{2.1}$$

For lack of a better name, the process just described will be called MLRRW, standing for *multiple-level reinforcing random walk*. The way to think about it is that we first move horizontally from (X_n, Y_n) according to the rules of reinforced random walk, and then we can move vertically in an arbitrary way before the next horizontal move. To make Keane's reinforced random walk on a ladder into an MLRRW, take $\mathcal{F}_n$ to be the σ-field generated by the process up to just before the $(n+1)^{st}$ horizontal step, together with the knowledge that the next step will be horizontal. It is easy to show that the conditions for MLRRW are satisfied by this choice of $\{\mathcal{F}_n\}_{n=0}^{\infty}$.

Let $p = (1+\delta)/(2+\delta)$ and $q = 1 - p = 1/(2+\delta)$. Note that p is the probability of crossing the reinforced edge when the choice is between one reinforced edge and one unreinforced edge. The notation $I(A)$ is used for the indicator function of the event A as in (2.1) above. A sum $\sum_{i=j}^{m} a_i$ will be taken to equal 0 when $m < j$. The main result of the paper is:

Theorem 2.1 *For an MLRRW with $d = 2$, one has $X_n = 0$ infinitely often, almost surely. Consequently, Keane's once-reinforced random walk on a ladder is recurrent.*

3 Outline of the argument and an easy weaker result

The idea of the argument is this. If $\{X_n\}$ were a martingale, it would be forced to return to zero infinitely often. While it is not a martingale, we may get an upper bound on the total amount of compensation per horizontal distance that can be required to make it into a martingale. The following result, while not strong enough to imply Theorem 2.1, is a useful preliminary result and gives a flavor of the argument. Define

$$C_n = \mathbb{E}(X_{n+1} - X_n \,|\, \mathcal{F}_n) = \frac{\delta}{2+\delta} \mathrm{sgn}\,(R(X_n, Y_n, n) - R(X_n - 1, Y_n, n)) \tag{3.2}$$

to be the compensator of the increment $X_{n+1} - X_n$ so that $X_n - \sum_{i=1}^{n-1} C_i$ is an $\mathcal{F}_n$-martingale.

Proposition 3.1 *For any x and y, and any stopping times σ_1, σ_2,*

$$\mathbb{E} \sum_{n=\sigma_1}^{\sigma_2 - 1} C_n I\{(X_n, Y_n) = (x, y)\} \leq \delta\,.$$

PROOF: Let $\tau_0 \leq \infty$ denote the least time $n \geq \sigma_1$ that $(X_n, Y_n) = (x, y)$ and $R(x-1, y, n) = 1$, let τ_j be the least $n > \tau_{j-1}$ for which $(X_n, Y_n) = (x, y)$, and let T' be the least n for which $R(x, y, n) = 1$. When $n \geq \sigma_1$, we may bound $C_n I\{(X_n, Y_n) = (x, y)\}$ above by zero unless $n = \tau_j < T'$ for some j, in which case $C_n \leq \delta/(2+\delta)$. Evidently,

$$\mathbb{P}(\tau_{j+1} < T') \leq \mathbb{E}\left[\mathbb{P}(\tau_j < T')\mathbb{P}(X_{\tau_j + 1} = x - 1 \,|\, \mathcal{F}_{\tau_j})\right] \leq \frac{1+\delta}{2+\delta} \mathbb{P}(\tau_j < T')$$

so we see inductively that

$$\mathbb{P}(\tau_j < T') \leq \left(\frac{1+\delta}{2+\delta}\right)^j .$$

Then

$$\begin{aligned}
\mathbb{E} \sum_{n=\sigma_1}^{\sigma_2 - 1} C_n I\{(X_n, Y_n) = (x, y)\} &\leq \frac{\delta}{2+\delta} \sum_{j=0}^{\infty} \mathbb{P}(\tau_j < T' \wedge \sigma_2) \\
&\leq \frac{\delta}{2+\delta} \sum_{j=0}^{\infty} \left(\frac{1+\delta}{2+\delta}\right)^j \\
&= \delta\,,
\end{aligned}$$

proving the proposition. $\square$

Corollary 3.2 *Define*

$$C = C(x, \sigma_1, \sigma_2) = \sum_{n=\sigma_1}^{\sigma_2 - 1} C_n I\{X_n = x\} \tag{3.3}$$

to be the total compensation occurring at sites (x, y), summed over y. Then for all $x > 0$,

$$\mathbb{E}C \leq (d-1)\delta\,.$$

PROOF: Condition on the first time $\tau \geq \sigma_1$ that $X_n = x$. If $\tau = \infty$ then of course $C(x) = 0$. If not, then

$$\sum_{n=\sigma_1}^{\sigma_2 - 1} C_n I\{(X_n, Y_n) = (x, Y_\tau)\} \leq 0$$

while for every $y \neq Y_\tau$, the previous lemma with $\sigma_1 = \tau$ gives

$$\mathbb{E}\left[\sum_{n=\sigma_1}^{\sigma_2 - 1} C_n I\{(X_n, Y_n) = (x, y)\} \,|\, \mathcal{F}_\tau\right] \leq \delta$$

and taking expectations with respect to $\mathcal{F}_{\sigma_1}$ then proves the corollary. $\square$.

This already allows us to prove the following weaker recurrence result.

Theorem 3.3 *In a MLRRW with d levels, the condition $\delta < (d-1)^{-1}$ is sufficient for recurrence.*

PROOF: Fix any $M > 0$ and any stopping time τ and let $T = T(\tau, M)$ be the least $n \geq \tau$ for which $|X_n| = 0$ or $|X_n| = M$. We first show

$$\mathbb{P}(|X_T| = M \,|\, \mathcal{F}_\tau) \leq \frac{|X_\tau|}{M} + (d-1)\delta \tag{3.4}$$

on the event that τ is finite and $|X_\tau| < M$. Assume without loss of generality that $X_\tau > 0$ (the argument for $X_\tau < 0$ is similar and the case $X_\tau = 0$ is automatic). Since

$$\left\{X_{n \wedge T} - \sum_{i=\tau}^{n \wedge T - 1} C_i\right\}_{n \geq \tau}$$

is an L^2-bounded martingale,

$$\begin{aligned}
\mathbb{E}(X_T - X_\tau \,|\, \mathcal{F}_\tau) &= \mathbb{E}\left(\sum_{i=\tau}^{T-1} C_i \,|\, \mathcal{F}_\tau\right) \\
&= \sum_{x=1}^{M-1} \mathbb{E}\left(\sum_{i=\tau}^{T-1} C_i I\{X_i = x\} \,|\, \mathcal{F}_\tau\right) \\
&\leq (M-1)(d-1)\delta
\end{aligned}$$

by Corollary 3.2, proving (3.4).

The theorem now follows directly. Take $\tau = \tau^{(k)}$ to be the least time n that $|X_n| = 2^k$ and take $M = A2^k$ with A sufficiently large so that $(d-1)\delta + A^{-1} = 1 - \epsilon < 1$. The conditional probability given $\mathcal{F}_{\tau^{(k)}}$ of returning to zero between times $\tau^{(k)}$ and $\tau^{(k+1)}$ is at least ϵ, so the theorem follows from the conditional Borel-Cantelli lemma. $\square$

4 Recurrence when $d = 2$ and improvements for $d > 2$

The key lemma will be the following strengthening of Corollary 3.2.

Lemma 4.1 *Let $2 \leq x < M$ and $0 \leq k < d$ be fixed. Let $\tau = \tau(x)$ be the least n for which $X_n = x$, and define T to be the least $n \geq \tau$ for which $X_T = 0$ or $X_T = x + 1$. Then, on the event that $\sum_{y=1}^{d} R(x-1, y, \tau) = d - k$,*

$$\mathbb{E}(C \,|\, \mathcal{F}_\tau) \leq \delta \cdot (k - 1 + \mathbb{P}(X_T = 0 \,|\, \mathcal{F}_\tau)) \, ,$$

where $C(x, \tau, T)$ is defined as in (3.3). In other words, the accumulation of compensation at horizontal position x, from the time x is hit until the time $x + 1$ or zero is hit, can be no more than -1, plus the number of unreinforced edges immediately to the left of position x when x was hit, plus the probability that the walk will return to zero before ever reaching $x + 1$.

PROOF: The proof is by induction on k. First assume $k = 0$. Let T' be the least $n > \tau$ for which $X_n = x + 1$ or $X_n = 0$. Define $\tau_0 = \tau$ and $\tau_j = \inf\{n > \tau_j : X_n = x\}$. We may compute

$$\begin{aligned}
1 - \mathbb{P}(X_T = 0 \,|\, \mathcal{F}_\tau) &\leq \mathbb{P}(X_{T'} = x + 1 \,|\, \mathcal{F}_\tau) \\
&= \sum_{j \geq 0} \mathbb{P}(\tau_j < T' = \tau_j + 1 \,|\, \mathcal{F}_\tau) \\
&= \frac{1}{2 + \delta} \sum_{j \geq 0} \mathbb{P}(\tau_j < T' \,|\, \mathcal{F}_\tau) \, .
\end{aligned}$$

On the other hand,

$$\mathbb{E}(C \,|\, \mathcal{F}_\tau) = \sum_{j \geq 0} -\frac{\delta}{2 + \delta} \mathbb{P}(\tau_j < T' \,|\, \mathcal{F}_\tau)$$

and these two together yield

$$\mathbb{E}(C \,|\, \mathcal{F}_\tau) \leq -\delta(1 - \mathbb{P}(X_T = 0 \,|\, \mathcal{F}_\tau))$$

which proves the lemma in the case $k = 0$.

Now assume for induction that the lemma is true when k is replaced by $k - 1$. There are two cases in the induction step. The first case is when at time τ, the random walker sees an unreinforced edge to the left, that is, $R(X_\tau - 1, Y_\tau, \tau) = 0$. The next paragraph refers to inequalities holding on this event.

We may split C into three components, depending on whether the next move is to the right or the next move is to the left and the walk does or does not return to horizontal position x before zero. Formally, write $C = C(I_1 + I_2 + I_3)$ where

$$\begin{aligned} I_1 &= I\{X_{\tau+1} = x+1\}\,; \\ I_2 &= I\{X_{\tau+1} = x-1, \tau_1 < T'\}\,; \\ I_3 &= I\{X_{\tau+1} = x-1, \tau_1 > T'\}\,. \end{aligned}$$

Clearly $CI_3 = 0$ since then the walk is at x only once before time T, at which time the compensator is zero. For the first piece, let $I_4 = I\{X_{\tau+1} = x+1, \tau_1 < T\}$, and observe that $I_4 \in \mathcal{F}_{\tau_1}$ while, by Corollary 3.2, $\mathbb{E}(C \,|\, \mathcal{F}_{\tau_1}) \le \delta k$. Thus

$$\begin{aligned} \mathbb{E}(CI_1 \,|\, \mathcal{F}_\tau) &= \mathbb{E}(CI_4 \,|\, \mathcal{F}_\tau) \\ &= \mathbb{E}\left[\mathbb{E}\left(CI_4 \,|\, \mathcal{F}_{\tau_1}\right) \,|\, \mathcal{F}_\tau\right] \\ &\le \mathbb{E}\left[I_4 \delta k \,|\, \mathcal{F}_\tau\right] \\ &\le \frac{1}{2}\delta k\,. \end{aligned}$$

Finally, on the event I_2 there are at most $k-1$ unreinforced edges just to the left of x at time τ_1 so that the induction hypothesis implies that

$$\begin{aligned} \mathbb{E}(CI_2 \,|\, \mathcal{F}_\tau) &= \mathbb{E}\left[\mathbb{E}\left(CI_2 \,|\, \mathcal{F}_{\tau_1}\right) \,|\, \mathcal{F}_\tau\right] \\ &\le \mathbb{E}\left[I_2 \delta\left(k - 1 + \mathbb{P}(X_T = 0 \,|\, \mathcal{F}_{\tau_1})\right) \,|\, \mathcal{F}_\tau\right] \\ &\le \frac{1}{2}\delta(k-2) + \mathbb{E}\left[\mathbb{P}(X_T = 0 \,|\, \mathcal{F}_{\tau_1}) \,|\, \mathcal{F}_\tau\right] \\ &= \frac{1}{2}\delta(k-2) + \mathbb{P}(X_T = 0 \,|\, \mathcal{F}_\tau)\,. \end{aligned}$$

Putting together the three estimates gives

$$\mathbb{E}(C \,|\, \mathcal{F}_\tau)\delta\left(k - 1 + \mathbb{P}(X_T = 0 \,|\, \mathcal{F}_\tau)\right)\,,$$

finishing the case where the walker sees an unreinforced edge to the left.

Finally, we remove the assumption of an unreinforced edge to the left. Let $\tau' \ge \tau$ be the least n for which $X_n = x+1$ or $X_n = x$ and $R(x-1, Y_n, n) = 0$. Then $C_n I\{X_n = x\} \le 0$ for $n < \tau'$, whence

$$E(C(x,\tau) \,|\, \mathcal{F}_\tau) \le \mathbb{E}(C(x,\tau') \,|\, \mathcal{F}_\tau)\,.$$

But $\mathbb{E}(C(x,\tau') \,|\, \mathcal{F}_{\tau'})$ is bounded above by $\delta(k - 1 + \mathbb{P}(X_T = 0 \,|\, \mathcal{F}_{\tau'}))$ on the event $\{X_{\tau'} = x\}$ (this is the case in the previous paragraph), and by $(k-1)\delta$ if $X_{\tau'} = x+1$ (Corollary 3.2). Removing the conditioning on $\mathcal{F}_{\tau'}$, we see that $\mathbb{E}(C \,|\, \mathcal{F}_\tau) \le \delta(k - 1 + \mathbb{P}(X_T = 0 \,|\, \mathcal{F}_\tau))$ in either case, which completes the induction and the proof of the lemma. $\square$

As a Corollary we get the following strengthening of (3.4):

Corollary 4.2 *For $2 \leq m < M$, let τ be the least n for which $X_n = m$ and let T be the least $n > \tau$ for which $X_n = 0$ or $X_n = M$. Then*

$$\mathbb{P}(|X_T| = M \mid \mathcal{F}_\tau) \leq 1 - \frac{(M-m)(1-(d-2)\delta) - m(d-1)\delta}{M + (M-m)\delta} .$$

PROOF: As in the proof of (3.4), the quantity

$$\left\{ X_{n \wedge T} - \sum_{i=\tau}^{n \wedge T - 1} C_i \right\}_{n \geq \tau}$$

is a martingale, and X_T is either 0 or M, so

$$\mathbb{P}(X_T = M \mid \mathcal{F}_\tau) = \frac{1}{M}\mathbb{E}(X_T \mid \mathcal{F}_\tau) = \frac{1}{M}\left[m + \mathbb{E}\left(\sum_{n=\tau}^{T-1} C_n \mid \mathcal{F}_\tau \right) \right] .$$

But

$$\begin{aligned} \mathbb{E}\left(\sum_{n=\tau}^{T-1} C_i \mid \mathcal{F}_\tau \right) &= \sum_{x=1}^{m} \mathbb{E}\left(\sum_{n=\tau}^{T-1} C_n I\{X_n = x\} \mid \mathcal{F}_\tau \right) + \sum_{x=m+1}^{M-1} \mathbb{E}\left(\sum_{n=\tau}^{T-1} C_n I\{X_n = x\} \mid \mathcal{F}_\tau \right) \\ &\leq \delta\left[m(d-1) + (M-m)(d-2+\mathbb{P}(X_T = 0 \mid \mathcal{F}_\tau)) \right] \end{aligned}$$

where Corollary 3.2 and Lemma 4.1 were used to bound the two summations. Thus

$$\mathbb{P}(X_T = M \mid \mathcal{F}_\tau) \leq \frac{1}{M}\left[m + \delta(m(d-1) + (M-m)(d-2+\mathbb{P}(X_T = 0 \mid \mathcal{F}_\tau))) \right] . \tag{4.5}$$

Letting $r = \mathbb{P}(X_T = 0 \mid \mathcal{F}_\tau) = 1 - \mathbb{P}(X_t = M \mid \mathcal{F}_\tau)$

$$M - Mr \leq m + \delta[m(d-1) + (M-m)(d-2+r)]$$

and solving for r gives

$$r \geq \frac{M - m - \delta(m(d-1) + (M-m)(d-2))}{M + \delta(M-m)}$$

which is equivalent to the conclusion of the corollary. □

PROOF OF THEOREM 2.1 The argument is exactly the same as the derivation of Theorem 3.3 from (3.4). □

Another immediate consequence of Lemma 4.1 is that when $d > 2$, we may strengthen Theorem 3.3 replacing $(d-1)^{-1}$ by $(d-2)^{-1}$:

Theorem 4.3 *If $\delta < (d-2)^{-1}$ then MLRRW with d levels is recurrent.* □

5 Further remarks

It is expected that Keane's once edge-reinforced walk on a ladder is recurrent for any d and δ. However, Theorem 4.3 is sharp in the sense that a MLRRW with d levels can be transient for any $\delta > (d-2)^{-1}$; thus the freedom to take arbitrary vertical jumps does seem to alter the critical value of δ. To see that Theorem 4.3 is sharp, define an MLRRW by choosing Y_i at each stage to make C_i positive whenever possible. If C_i cannot be made positive, then let it be zero if possible, giving preference to sites where the edges to either side are already reinforced. The proof of transience is similar to arguments to be found in Sellke (1993). The gist is as follows. First of all, it is easy to show that X_n is transient if and only if X_n^+, the positive part of X_n, is transient. (Note that C_i is never negative when X_i is negative.) So consider the process X_n^+. A zero-one law argument shows that X_n^+ is either almost surely transient or almost surely recurrent. If X_n^+ were almost surely recurrent, we could find an M large enough so that, for the overwhelming majority of x values between 0 and M, the probability is near 1 that all horizontal edges at x are reinforced before X_n^+ hits M. One then shows that, for the randomized enlightened greedy algorithm, the expected cumulative bias at a positive x is $(d-2)\delta$ if X_n^+ visits x often enough. Consequently, the expected cumulative bias accumulated by X_n^+ by the time T_M that M is finally hit can be shown to be greater than M. But this would imply $E(X_{T_M}) > M$, which contradicts $X_{T_M} \equiv M$.

In the critical case $\delta = (d-2)^{-1}$, this algorithm can be shown to produce recurrence, again by arguments similar to those in Sellke (1993).

Acknowledgement

I am very grateful to Burgess Davis for telling me about this problem and for useful conversations about it.

References

[1] Coppersmith, D. and Diaconis, P. (1987). Random walk with reinforcement. *Unpublished manuscript.*

[2] Davis, B. (1990). Reinforced random walk. *Probab. Th. Rel. Fields* **84**, 203–229. MR1030727

[3] Diaconis, P. (1988). Recent progress on de Finetti's notions of exchangeability. *Bayesian Statistics* **3**, 111–125, J.M. Bernardo, M.H. DeGroot, D.V. Lindley, and A.F.M. Smith, Eds. Oxford University Press. MR1008047

[4] Durrett, R., Kesten, H. and Limic, V. (2002). Once edge-reinforced random walk on a tree. *Prob. Theory. Rel. Fields* **122**, 567–592. MR1902191

[5] Keane, M. (2002). Lecture at the University of Maryland, February 19, 2002.

[6] Neveu, J. (1965). *Mathematical Foundations of the Calculus of Probabilities*, Holden-day, San Francisco. MR198505

[7] Pemantle, R. (1988). Phase transition in reinforced random walk and RWRE on trees. *Ann. Probab.* **16**, 1229–1241. MR942765

[8] Sellke, T. (1993). Nearest-neighbor random walk in a changing environment. *Unpublished manuscript.*