Electron. Commun. Probab. **18** (2013), no. 43, 1–6.
DOI: 10.1214/ECP.v18-2763
ISSN: 1083-589X

ELECTRONIC COMMUNICATIONS in PROBABILITY

Controlled random walk with a target site*

Kenneth S. Alexander†

Abstract

We consider a symmetric simple random walk $\{W_i\}$ on $\mathbb{Z}^d, d = 1, 2$, in which the walker may choose to stand still for a limited time. The time horizon is n, the maximum consecutive time steps which can be spent standing still is m_n and the goal is to maximize $P(W_n = 0)$. We show that for $d = 1$, if $m_n \gg (\log n)^{2+\gamma}$ for some $\gamma > 0$, there is a strategy for each n yielding $P(W_n = 0) \to 1$. For $d = 2$, if $m_n \gg n^\epsilon$ for some $\epsilon > 0$ then there are strategies yielding $\liminf_n P(W_n = 0) > 0$.

Keywords: random walk; stochastic control.
AMS MSC 2010: Primary 60G50, Secondary 93E20.
Submitted to ECP on April 25, 2013, final version accepted on May 7, 2013.

1 Introduction.

We consider a process $\{W_i, 0 \le i \le n\}$ on $\mathbb{Z}^d$ ($d = 1$ or 2) in which $W_0 = 0$ and each step $W_i - W_{i-1}$ either is 0 (i.e. standing still) or is a step $\pm e_i$ of symmetric simple random walk (SSRW), with e_i being the ith unit coordinate vector. The choice between standing still and SSRW step is determined by a strategy. Formally, a *strategy* (or *n-strategy*) is a mapping $\delta_n : \mathcal{T}_n \to \{0, 1\}$ defined on the space

$$\mathcal{T}_n = \Big\{\{(i, w_i), 0 \le i \le j\} : 0 \le j \le n-1, |w_i - w_{i-1}| \le 1 \text{ for all } i\Big\}$$

of all space-time trajectories of length less than n; here the values 0 and 1 for δ_n correspond to standing still and taking a SSRW step, respectively. Thus the choice of whether to take a step at a time i depends on the trajectory up to time $i-1$. Letting $\{Y_i, i \ge 1\}$ be SSRW, we then construct the controlled random walk process iteratively from the strategy by

$$W_0 = 0, \quad W_i - W_{i-1} = \Delta_i(Y_i - Y_{i-1}) \quad i \le n,$$

where

$$\Delta_i = \delta_n\Bigg(\{(k, W_k) : 0 \le k \le i-1\}\Bigg).$$

We define the time since the last SSRW step to be

$$J_i = \min\{j \in [0, i-1] : W_{i-j} \ne W_{i-j-1}\},$$

with $J_i = i$ if the set in this definition is empty. We designate a maximum number $m_n - 1$ of consecutive steps standing still and say that a strategy is *admissible* if $J_i \le m_n - 1$

*This research was supported by NSF grant DMS-0804934.
†University of Southern California, USA. E-mail: alexandr@usc.edu

a.s. for all $i \le n$. A strategy is *Markov* if the value of $\delta_n\big(\{(k,W_k) : 0 \le k \le i-1\}\big)$ depends only on the pair (W_i, J_i). It is easily seen that if the strategy is Markov, then $\{(W_i, J_i) : i \ge 0\}$ is a Markov process. We write $\mathcal{A}_n$ for the set of all admissible strategies and $\mathcal{A}_n^*$ for the set of all admissible Markov strategies. When the dependence on the strategy needs to be made clear, we write $P(\cdot;\delta_n)$ for probability when strategy δ_n is used, but generally we suppress the δ_n in the notation.

We are interested in steering the process toward a target site, and specifically in the behavior of

$$\hat{P}(W_n = 0) = \sup_{\delta_n \in \mathcal{A}_n} P(W_n = 0; \delta_n) \tag{1.1}$$

as $n \to \infty$. The word "steering" is a bit misleading here, as the process never has a drift. For SSRW, of course $P(W_n = 0) \sim Cn^{-d/2}$. (Here and throughout the paper, C and $C_1, C_2, \dots$ are generic constants, and $a_n \sim b_n$ means the ratio converges to 1.) A simple strategy to increase this is to minimize steps, i.e. always stand still as long as allowed, taking only n/m_n steps, yielding $P(W_n = 0) \asymp (m_n/n)^{d/2}$, where we use $a_n \asymp b_n$ to mean the ratio is bounded away from 0 and ∞. A slightly more sophisticated strategy is to minimize steps until time $n - m_n$ then maximize steps (i.e. never stand still) until $\{W_i\}$ hits 0 (if it does), and then stand still until time n. For $d = 1$, conditionally on $\{|W_{n-m_n}| \le m_n^{1/2}\}$ this strategy has a success probability of order 1, so the unconditional success probability is of order $\min(m_n/n^{1/2}, 1)$; in particular it is bounded away from 0 if $m_n \ge Cn^{1/2}$.

This brings up two questions. For which $\{m_n\}$ is $\hat{P}(W_n = 0)$ bounded away from 0? And are there nontrivial $\{m_n\}$ for which $\hat{P}(W_n = 0) \to 1$? Our two main theorems give some answers.

Theorem 1.1. *Suppose $d = 1$ and $m_n \gg (\log n)^{2+\gamma}$ for some $\gamma > 0$.. Then $\hat{P}(W_n = 0) \to 1$ as $n \to \infty$.*

In one dimension, the "more sophisticated" strategy described above relies on the fact that a SSRW started at distance k from 0 has a probability of order one to hit 0 by time k^2. In two dimensions, this probability is only of order $(\log k)^{-1}$ so it is more difficult to construct a strategy based on waiting for the RW to return to 0 after it has wandered away. Nonetheless we have the following.

Theorem 1.2. *Suppose $d = 2$ and $m_n \gg n^\epsilon$, for some $\epsilon > 0$. Then*

$$\liminf_n \hat{P}(W_n = 0) > 0.$$

In (1.1) one could consider only Markov strategies, i.e. take the sup over $\mathcal{A}_n^*$. Since the underlying SSRW is Markov, one cannot actually do better with a non-Markov strategy, so the two sups are the same. Allowing non-Markov strategies simply lets us use ones that can be concisely described and analyzed.

Remark 1.3. *One could consider alternate ways of slowing down the controlled RW, in place of standing still. For example, one could allow a choice between a standard SSRW step and a delayed SSRW step, the latter meaning we stand still with probability $1 - \frac{1}{m_n}$ and take a standard SSRW step of $\pm e_i$ with probability $1/m_n$, with no limit on how many consecutive times we choose the delayed SSRW step. A brief examination of the proofs shows that both theorem statements above remain valid in this case.*

Remark 1.4. *The continuous analog of the problem we consider is a diffusion ξ_t in which the drift is always 0 and one can control the diffusivity $\sigma(x,t)$, but constrained to an interval $[\sigma_1, \sigma_2]$, on a time interval $[0,T]$, where $\sigma_1 > 0$. We take $[\sigma_1, \sigma_2] = [\epsilon_T, 1]$ and ask, how slowly can we have $\epsilon_T \to 0$ as $T \to \infty$ and still have $\hat{P}(|\xi_T| < 1) \to$*

1 *or* $\liminf_T P(|\xi_T| < 1) > 0$*? Note that* ϵ_T *is the analog of* $1/m_n$*. McNamara [3] considered closely related questions for a one-dimensional diffusion; see also [2] for another variant. He proved that there are constants* $K = K(\sigma_1, \sigma_2)$ *and* $\beta = \beta(\sigma_1/\sigma_2)$ *such that*

$$\hat{P}^x \left(|\xi_T| \le \frac{\delta}{2} \right) \le K \left(\frac{\delta}{\sqrt{T}} \right)^{\beta} \quad \textit{for all } x \in \mathbb{R}, \delta > 0, \tag{1.2}$$

while in the other direction, for each $h > 0$ *there exists* $\Delta > 0$ *such that*

$$\hat{P}^0 \left(|\xi_T| \le \frac{\delta}{2} \right) > K \left(\frac{\delta}{\sqrt{T}} \right)^{\beta+h} \quad \textit{for all } \delta \in (0, \Delta). \tag{1.3}$$

Here $\hat{P}^x$ *denotes probability (maximized over the allowed controls) for a process started at* x*. Numerical evidence was given that* β *is approximately proportional to* σ_1/σ_2*. If we assume this to be true and consider* ϵ_T *of order* $1/\log T$*, we get* β *of order* $1/\log T$ *as well. If we could take* h *also of this order, then the right side of* (1.3) *would be bounded away from 0 in* T*, as desired. The problem is that one cannot get a useful result from* (1.3) *if one takes* h *depending on* T*, since then* δ *must also depend on* T*. Nonetheless we may observe that* $\epsilon_T \le C/\log T$ *corresponds to* $m_n \ge C \log n$ *in our discrete problem. Further, still assuming* β *proportional to* σ_1/σ_2*, we see that since we want the left side of* (1.2) *bounded away from 0, we cannot allow* $\epsilon_T \gg 1/\log T$*, suggesting that we perhaps cannot do better than requiring* $m_n \ge C \log n$ *in Theorem 1.1. But we know of no results giving lower bounds for the rate at which* m_n *must grow to give* $\hat{P}(W_n = 0) \to 1$*.*

2 Proof of Theorem 1.1

Throughout the paper we will make use of various quantities which approach infinity as $n \to \infty$. For $\lambda > 0$ let $u_n(\lambda) = \max\{k : m_n^{1+k\lambda} \le n\}$; note that $u_n(\lambda)$ is of order $\log n/\log m_n$ for all $\lambda > 0$. Also, the hypothesis that $m_n \gg (\log n)^{2+\gamma}$ for some $\gamma > 0$ is equivalent to

$$m_n \gg \left(\frac{(\log n)^2}{\log\log n} \right)^{1/(1-\eta)} \quad \text{for some } \eta \in (0,1). \tag{2.1}$$

Fixing such an η and writing u_n for $u_n(\eta)$, we observe that (2.1) is equivalent to

$$m_n^{1-\eta} \log m_n \gg (\log n)^2,$$

hence also to

$$m_n^{1+\eta\theta(m_n^{1-\eta}/\log m_n)^{1/2}} \ge n \quad \text{for all large } n\text{, for every } \theta > 0,$$

and therefore finally to

$$u_n = o\left(\left(\frac{m_n^{1-\eta}}{\log m_n} \right)^{1/2} \right). \tag{2.2}$$

Fix n, write m for m_n. and define

$$\epsilon_m = \frac{\log m}{m^{1-\eta}}, \quad \text{so} \quad \epsilon_m = o\left(\frac{1}{u_n^2} \right) \quad \text{as } n \to \infty.$$

We define a sequence of windows $\{t_k\} \times I_k, 1 \le k \le u_n$, in space-time, with size decreasing as k increases and the target $(n, 0)$ is approached; we then construct a strategy which makes the space-time trajectory of the process pass through all of these windows, with high probability. Specifically, let

$$t_k = n - m^{1+\eta(u_n - k)}, \quad 2 \le k \le u_n, \quad \text{and } t_0 = 0, \quad t_1 = \frac{t_2}{2}, \quad t_{u_n+1} = n,$$

$$N_k = t_k - t_{k-1}, \quad h_{mk} = (\epsilon_m N_{k+1})^{1/2}.$$

Let $I_0 = \{0\}$,

$$I_k = [-h_{mk}, h_{mk}], \quad 1 \le k \le u_n, \quad \text{and } I_{u_n+1} = \{0\}.$$

We want to find a strategy, and choice of ϵ_m, for which we can show

$$\max_{k \le u_n+1} P(W_{t_k} \notin I_k \mid W_{t_{k-1}} \in I_{k-1}) = o\left(\frac{1}{u_n}\right) \quad \text{as } n \to \infty. \tag{2.3}$$

From (2.2) and (2.3) it follows that

$$P(W_n = 0) \ge \prod_{k=1}^{u_n+1} P(W_{t_k} \in I_k \mid W_{t_{k-1}} \in I_{k-1}) \ge \left(1 - o\left(\frac{1}{u_n}\right)\right)^{u_n+1} = 1 - o(1) \tag{2.4}$$

as $n \to \infty$, proving the theorem.

For (2.3), fix $1 \le k \le u_n + 1$ and let $\tau_0 = \tau_0^{(k)} = \min\{t > t_{k-1} : W_t = 0\}$. Fix $x \in I_{k-1}$ and observe that

$$\begin{aligned} P(W_{t_k} \notin I_k \mid W_{t_{k-1}} = x) &\le P\left(\tau_0 > t_k \middle| W_{t_{k-1}} = x\right) \\ &\quad + P\left(\tau_0 \le t_k, |W_{t_k}| > h_{mk} \middle| W_{t_{k-1}} = x\right). \end{aligned} \tag{2.5}$$

To fully specify, and then bound, these probabilities we need to designate a strategy; we do so by describing how the strategy works between t_{k-1} and t_k, for general k. For $k \le u_n$, we begin by taking all SSRW steps (i.e. no standing still) from time t_{k-1} until time $\tau_0 \wedge t_k$. If $\tau_0 > t_k$, we deem the strategy to have failed and we continue in an arbitrary manner, say all SSRW steps. If $\tau_0 \le t_k$, we continue from time τ_0 to t_k by always standing still for the maximum allowed period of time m during the interval $(\tau_0, t_k]$, that is, we take an SSRW step every mth time step. The last standing period is truncated if it would otherwise go beyond time t_k.

For $k = u_n + 1$, our strategy during $(t_{u_n}, t_{u_n+1}] = (n - m, n]$ is to maximize steps until the the time $\tau_0 = \tau_0^{(u_n+1)}$ when the process first hits 0 (if $\tau_0 \le n$), then stand still until time n.

Note that the failure of the Markov property for our strategy comes from the fact that standing still vs. taking a step depends on whether the stopping time τ_0 has occurred yet, not just on the current location of the process.

We now bound the first term on the right side of (2.5). From the Reflection Principle we have (for ℓ, h_{mk} of opposite even-odd parity):

$$P(\tau_0^{(k)} - t_{k-1} > l \mid W_{t_{k-1}} = x) = P(-|x| < Y_l < |x|), \tag{2.6}$$

and hence for every $\epsilon > 0$,

$$P(\tau_0^{(k)} - t_{k-1} > yx^2 \mid W_{t_{k-1}} = x) \sim 2\left(\Phi\left(\frac{1}{\sqrt{y}}\right) - \Phi(0)\right)$$
$$\text{as } |x| \to \infty, \text{ uniformly over } y \in [\epsilon, \infty). \tag{2.7}$$

The left side of (2.6) is a nondecreasing function of $|x|$, so for all $x \in I_k$, by (2.7)

$$P\left(\tau_0 > t_k \mid W_{t_{k-1}} = x\right) \le P\left(\tau_0 - t_{k-1} > N_k \mid W_{t_{k-1}} = h_{mk}\right) \le C\sqrt{\epsilon_m}. \tag{2.8}$$

Turning to the second term on the right side of (2.5), it is 0 for $k = u_n + 1$ so we consider $k \le u_n$. We can condition also on τ_0 as follows: for $t \in (t_{k-1}, t_k]$, using

Hoeffding's Inequality [1],

$$\begin{aligned} P\left(|W_{t_k}| > h_{mk} \,\middle|\, \tau_0 = t, W_{t_{k-1}} = x\right) &= P\left(|W_{t_k}| > h_{mk} \,\middle|\, W_t = 0\right) \\ &= P\left(|Y_{(t_k - t)/m}| > h_{mk}\right) \\ &\le e^{-h_{mk}^2 m / 2N_k} \\ &= e^{-m\epsilon_m N_{k+1}/2N_k}. \end{aligned} \tag{2.9}$$

For $k \ge 1$ we have $N_{k+1}/N_k \ge m^{-\eta}$ (with equality for $k \ge 3$), so (2.9) says that

$$P\left(|W_{t_k}| > h_{mk} \,\middle|\, \tau_0 = t, W_{t_{k-1}} = x\right) \le m^{-1/2} \le \sqrt{\epsilon_m}. \tag{2.10}$$

Since $x \in I_{k-1}$ is arbitrary, combining (2.5), (2.8) and (2.10) yields that for all $1 \le k \le u_n + 1$,

$$P(W_{t_k} \notin I_k \mid W_{t_{k-1}} \in I_{k-1}) \le C\sqrt{\epsilon_m} = o\left(\frac{1}{u_n}\right) \quad \text{as } n \to \infty, \tag{2.11}$$

proving (2.3) and thereby proving Theorem 1.1.

3 Proof of Theorem 1.2

We keep the same definition of $u_n(\lambda)$ and note that now our hypothesis on m_n is equivalent to the statement that $\{u_n\}$ is bounded, say $u_n \le u < \infty$ for all n. We keep the same formula for t_k but with η replaced by κ, determined as follows. Choose

$$\theta \in \left(\frac{1-\epsilon}{2}, \frac{1}{2}\right),$$

so $1 - 2\theta < \epsilon$, then choose κ small enough so

$$0 < \frac{1-2\theta}{1-2\kappa\theta} < \epsilon. \tag{3.1}$$

We then write u_n for $u_n(\kappa)$. For our windows, in place of the interval I_k we have the square $Q_k = [-N_{k+1}^\theta, N_{k+1}^\theta]^2$. To distinguish dimensions clearly, we now write $\{Y_i^{(d)}\}$ for d-dimensional SSRW, $d = 1, 2$. We use the same strategy as in one dimension: in each interval $(t_{k-1}, t_k]$, take an SSRW step every time step until time $\tau_0^{(k)}$, then an SSRW step every mth time step from time $\tau_0^{(k)}$ to t_k. In place of (2.3), we will need that for some $C > 0$,

$$P(W_{t_k} \in Q_k \mid W_{t_{k-1}} \in Q_{k-1}) \ge C \quad \text{for all } k\text{, for all large } n. \tag{3.2}$$

Let $L_{[a,b]}$ be the number of visits to 0 by the SSRW $\{Y_i^{(2)}\}$ during the time interval $[a, b]$. In comparison to (2.8), we have for all $x \in Q_{k-1}$:

$$\begin{aligned} P\left(\tau_0 \le t_k \,\middle|\, W_{t_{k-1}} = x\right) &= P\left(\tau_0 \le t_k \,\middle|\, Y_{t_{k-1}}^{(2)} = x\right) \\ &\ge \frac{E(L_{(t_{k-1}, t_k]} \mid Y_{t_{k-1}}^{(2)} = x)}{E(L_{(t_{k-1}, t_k]} \mid \tau_0 \le t_k, Y_{t_{k-1}}^{(2)} = x)} \end{aligned} \tag{3.3}$$

Since $|x| \le CN_k^\theta$ with $\theta < \frac{1}{2}$, the numerator in (3.3) is equal to

$$\sum_{i=1}^{N_k} P(Y_i = 0 \mid Y_0^{(2)} = x) \ge C \sum_{i=|x|^2}^{N_k} \frac{1}{i} \ge C_1 \log N_k. \tag{3.4}$$

Here C_1 depends on θ. The denominator in (3.3) is bounded above by

$$E(L_{(t_{k-1},t_k]} \mid Y_0^{(2)}=0) \leq C \sum_{i=1}^{N_k} \frac{1}{i} \leq C_2 \log N_k. \tag{3.5}$$

Therefore we have the analog of (2.8):

$$P\left(\tau_0 \leq t_k \;\middle|\; W_{t_{k-1}}=x\right) \geq \frac{C_1}{C_2} > 0 \quad \text{for all } k \leq u_n, x \in Q_{k-1}. \tag{3.6}$$

In comparison to (2.9), we have for $x \in Q_k$ and $t \in (t_{k-1}, t_k]$, using again Hoeffding's Inequality [1]:

$$\begin{aligned} P\left(W_{t_k} \notin Q_k \;\middle|\; \tau_0 = t, W_{t_{k-1}}=x\right) &= P\left(W_{t_k} \notin Q_k \;\middle|\; W_t = 0\right) \\ &= P\left(Y_{(t_k-t)/m} \notin Q_k \mid Y_0^{(2)}=0\right) \\ &\leq 2 \max_{j \leq N_k/m} P(|Y_j^{(1)}| > N_{k+1}^{\theta}) \\ &\leq 4e^{-mN_{k+1}^{2\theta}/2N_k}. \end{aligned} \tag{3.7}$$

For $k \geq 1$ we have $N_{k+1} \geq m^{-\kappa} N_k$ (with equality for $k \geq 3$), so by (3.1),

$$\frac{mN_{k+1}^{2\theta}}{N_k} \geq \frac{m^{1-2\kappa\theta}}{N_k^{1-2\theta}} \geq \frac{m^{1-2\kappa\theta}}{n^{1-2\theta}} \to \infty \text{ as } n \to \infty. \tag{3.8}$$

Combining (3.6), (3.7) and (3.8) we see that (3.2) holds. Since $\{u_n\}$ is bounded, it follows as in (2.4) that $\liminf_n P(W_n = 0) > 0$.

References

[1] Hoeffding, W.: Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.* **58**, (1963),13–30. MR-0144363

[2] McNamara, J. M.: Optimal control of the diffusion coefficient of a simple diffusion process. *Math. Oper. Res.* **8**, (1983), 373–380. MR-0716119

[3] McNamara, J. M.: A regularity condition on the transition probability of a diffusion process. *Stochastics* **15**, (1985), 161–182. MR-0869198

Acknowledgments. The author would like to thank Ananda Weerasinghe for helpful conversations.