Filling a gap from last time

Back in Lec07, I wrote "... Also note that $V_{ij}(t)$ differs from $p_{ij}\ V_i(t)$ by at most a constant C_{ij} that does not depend on *t,*" where for each state s_i, $V_i(t)$ is the number of times that the rotor-walk has been in state s_i in the first t time-steps, and $V_{ij}(t)$ is the number of times (in the first t time-steps) that the rotor-walk has been in state s_i and then gone immediately to state s_j. Now I want to explain that.

Suppose first that p_{ij} is $\frac{1}{2}$. Then every time $V_i(t)$ increases by 2, $V_{ij}(t)$ goes up by 1, so that the difference $V_{ij}(t) - \frac{1}{2}V_i(t)$ doesn't change at all. This implies that $V_{ij}(t) - \frac{1}{2}V_i(t)$ takes on just two different values, in alternation, as t goes to infinity: one value for $V_i(t)$ odd, and one value for $V_i(t)$ even.

More generally, suppose that the rotor at i has period m. Then every time $V_i(t)$ increases by m, $V_{ij}(t)$ goes up by $p_{ij}\ m$, so that $V_{ij}(t) - p_{ij}\ V_i(t)$ doesn't change at all. So $V_{ij}(t) - p_{ij}\ V_i(t)$ takes on at most m different values, repeating with period m. In particular, the values taken on by $V_{ij}(t) - p_{ij}\ V_i(t)$ are bounded.

Random walks in one dimension

Walk on ℕ

Suppose we do a biased random walk on ℕ = {0,1,2,...} with absorption at 0, where a walker at position $k > 0$ moves to position $k+1$ with probability p and moves to position $k-1$ with probability $q = 1 - p$.

(This is a Markov chain with infinitely many states: $p_{00}=1$ (state 0 is an absorbing state), and for all i >0,

$p_{ij} = p$ if $j = i+1$, $p_{ij} = q$ if $j = i-1$, and $p_{ij} = 0$ otherwise.)

What is the probability that a walker who starts at k eventually gets absorbed at 0?

Let p_k be this absorption probability. We have the relation

(1) $\quad p_k = p\, p_{k+1} + q\, p_{k-1}$

for all $k > 0$, i.e., $p_{k+1} = [p_k - q\, p_{k-1}]/p$, so we can compute p_2 from p_1 and p_0, compute p_3 from p_2 and p_1, etc. Hence it's enough to know p_0 and p_1.

Clearly $p_0 = 1$; what about p_1?

Claim: $p_2 = p_1^2$.

Proof: Let $P(x,y)$ be the probability, if you start at x, that you eventually hit y. Then
$p_2 = P(2,0) = P(2,1)\,P(1,0) = P(1,0)\,P(1,0) = p_1^2$.

Remark #1: The second "=" is a consequence of two facts: in order to get from 2 to 0, the walker must get to 1 first; and once the walker has arrived at 1 for the first time, the chance of eventually getting to 0 is P(1,0), since the walk has no memory.

Remark #2: The third "=" is a consequence of the fact that our transition probabilities are translation-invari - ant: that is, $P(i+1,j+1) = P(i,j)$ for all i,j because $p_{i+1,j+1} = p_{i,j}$ for all i,j.

Likewise, $p_k = p_1^k$ for all $k > 1$.

Consequence: Plugging $p_2 = p_1^2$
into
$$p_1 = p\, p_2 + q\, p_0$$
(the case $k = 1$ of (1)) and recalling that $p_0 = 1$, we get
$$p_1 = p\, p_2 + q\, p_0 = p\, p_1^2 + q\, 1,$$
i.e.,
$$p\, p_1^2 + (-1)\, p_1 + (1-p) = 0$$
i.e., $(p\, p_1 - (1-p))(p_1 - 1) = 0$.

This quadratic equation in p_1 has two roots: p_1=1 and p_1=$(1-p)/p = q/p$.

Which root is p_1 equal to?

Leftward bias

If $p < \frac{1}{2}$ (that is, if our walk has leftward drift), then we can't have $p_1 = q/p > 1$, since p_1 is a probability.

Hence $p_1 = 1$, and it follows that $p_2 = 1$, etc.

Another way to see that $p_1 = 1$ in this case is to look at the corresponding walk on $\{0,1,2,...,n\}$:

If the walker on $\mathbb{N}$ never gets absorbed at 0, then for all n, the walker hits n before hitting 0. (We're taking it for granted here that the walker must eventually hit either 0 or n with probability 1; but we already know this for walk on $\{0,1,2,...,n\}$, and the fact for $\mathbb{N}$ follows.)

Hence, for all n, the probability that the walker on $\mathbb{N}$ who starts at 1 never hits 0 is at most the probability that the walker on $\mathbb{N}$ who starts at 1 hits n before 0, which is equal to the probability that the walker on $\{0,1,2,...,n\}$ who starts at 1 hits n before 0, which is

$((q/p)^1 - 1)/((q/p)^n - 1)$, which goes to 0 as $n \to \infty$ since $q/p > 1$. Hence the probability of never hitting 0 on $\mathbb{N}$ is 0, i.e., the probability of eventually hitting 0 is 1.

How long does it take on average for a walker who starts at k to get absorbed at 0?

Let a_k be this mean absorption time. We have the relation

(2) $\quad a_k = 1 + p\, a_{k+1} + q\, a_{k-1}$

for all $k > 0$, i.e., $a_{k+1} = [a_k - q\, a_{k-1} - 1]/p$,
so as before it's enough to know a_0 and a_1.

Clearly $a_0 = 0$; what about a_1? It's not so clear. But we can say something about the relationship between a_1 and a_2, just as we found earlier a relationship between p_1 and p_2, specifically $p_2 = p_1^2$.

What is this relationship? ...
..?..

Claim: $a_2 = 2\ a_1$.

Proof: Write the (random) time $T_{2,0}$ it takes the walker to get from 2 to 0 as a sum of two times: the time $T_{2,1}$ it takes until the walker first hits 1, and the time $T_{1,0}$ it takes after that for the walker to hit 0. Since a Markov chain doesn't remember its past, the random variables $T_{2,1}$ and $T_{1,0}$ are independent, but we actually don't need that fact; all we need is that $\text{Exp}(T_{2,0}) = \text{Exp}(T_{2,1}) + \text{Exp}(T_{1,0})$. Clearly $\text{Exp}(T_{1,0}) = a_1$, and the translation-invariance of the transition probabilities tells us that $\text{Exp}(T_{2,1}) = \text{Exp}(T_{1,0}) = a_1$. So $a_2 =$
$\text{Exp}(T_{2,0}) =$
$a_1 + a_1$.

Likewise $a_k = k\, a_1$ for all $k > 1$.

Consequence:Plugging $a_2 = 2\, a_1$
into

$$a_1 = 1 + p\, a_2 + q\, a_0$$

((2) in the case $k = 1$) and recalling that $a_0 = 0$, we get

$$a_1 = 1 + p\, a_2 + (1\text{-}p)\, a_0 = 1 + 2\, p\, a_1,$$

so

$$(1 - 2p)\, a_1 = 1,$$

and $a_1 = 1/(1\text{-}2\, p)$.

So, when $p < \frac{1}{2}$, we get

$$p_k = 1$$

$$a_k = \frac{k}{1-2\,p} = \frac{k}{q-p}$$

for all $k \geq 0$.

Rightward bias

If $p > \frac{1}{2}$ (that is, if our walk has rightward drift), then $q/p < 1$, so both possibilities $p_1 = 1$ and $p_1 = q/p$ are conceivable.

To see which is correct, look at the corresponding walk on $\{ 0,1,2,...,n \}$.

Before, we showed that the nonabsorption probability $1\text{-}p_1$ is bounded above by $((q/p)^1 - 1)/((q/p)^n - 1)$. Now that $q/p < 1$, this goes to $((q/p) - 1)/(0 - 1) = 1 - q/p$. So $1\text{-}p_1 \leq 1 - q/p$. This tells us that p_1 is at least q/p. But how do we know it isn't 1?

Let Ω be the probability space for the walker starting at 1.

Let $E \subset \Omega$ be the event "the walker never reaches 0", and write E as the intersection of the sets E_2, E_3, ..., where E_n is the event "the walker hits n without hit - ting 0". Since
$E = \bigcap_{n=2}^{\infty} E_n$, and since the sets E_n are nested, we have

$\text{Prob}(E) = \lim_{n\to\infty} \text{Prob}(E_n) =$
$\lim_{n\to\infty} ((q/p)^1 - 1)/((q/p)^n - 1) = 1 - q/p$,
so $p_1 = 1 - \text{Prob}(E) = q/p$.

So, when $p > \frac{1}{2}$, we get

$$p_k = (q/p)^k$$

for all $k \geq 0$.

Note that it makes no sense to talk about a_k in this case, since the time until absorption is infinite with positive probability.

No bias

In the case $p = \frac{1}{2}$, the two candidate values for p_1 (1 and q/p) coincide, so we get $p_1 = 1$ and $p_k = 1$ for all $k \geq 0$.

We can also see this by comparison with the finite- n case. Defining E and E_2, E_3, ... as above, we have P(E_n) = $\frac{1}{n}$, which doesn't go to 0 as fast as $((q/p)^1 - 1)/((q/p)^n - 1)$ did in the case $p < \frac{1}{2}$, but that's okay: we still get Prob(E) = $\lim_{n\to\infty}$ Prob(E_n) = $\lim_{n\to\infty} \frac{1}{n}$ = 0, so $p_1 = 1 - 0 = 1$.

How long does it take on average for the walker to get from 1 to 0?

It can be shown that the expected value of the absorption time is <u>infinite</u>.

One way to prove this is by viewing the case $p = \frac{1}{2}$ as a limiting case of the $p < \frac{1}{2}$ regime, noting that as $p \to \frac{1}{2}$ from below, we get $a_k = \frac{k}{1-2p} \to \infty$.

Another way to prove it is to compare with random walk on $\{0,1,2,...,n\}$, which has an expected time-until absorption (starting from 1) equal to (1)(n-1) (you proved this fact about gambler's ruin on the homework; see the solution to 11.2.26, aka problem I from assignment #2).

Walk on $\mathbb{Z}$

Suppose we do a biased random walk on $\mathbb{Z}$ = {...,-2,-1,0,1,2,...}, where a walker at position $k > 0$ moves to position $k + 1$ with probability p and moves to position $k - 1$ with probability $q = 1 - p$.

Say the walker starts from 0; does the walker get back to 0 again? How often?

Another way to look at this process is to view the walker's location at time n as a sum of n i.i.d. random variables $X_1+X_2+...+X_n$, where each X_k is +1 with probability p and is -1 with probability q.

The Strong Law of Large Numbers, applied to the i.i.d. process, tells us that with probability 1, $(X_1+X_2+...+X_n)/n \to p-q$ (since $p-q$ is the shared expected value of the X_k's). In particular, if $p-q \neq 0$ (i.e., if $p \neq \frac{1}{2}$), then with probability 1, $(X_1+X_2+...+X_n)/n$ will equal 0 only finitely often, i.e., $X_1+X_2+...+X_n$ will equal 0 only finitely often. That is, with probability 1, the random walk on $\mathbb{Z}$ that starts from 0 will visit 0 only finitely many times. (If $p > \frac{1}{2}$, the walk drifts off to the right; if $p < \frac{1}{2}$, the walk drifts off to the left.)

What about $p = \frac{1}{2}$?

We have shown that, starting from 1, the probability that an unbiased random walker will eventually hit 0 is 1. The same is true starting from -1. Hence, if a walker takes a random step starting from 0, regard - less of whether the step is to the right or the left, the probability that the walker will return to 0 is 1.

And when that happens, the probability that the walker will return again is likewise 1. Etc.
So the probability that the walker starting from 0 will visit 0 at least *n* times is 1, for every *n*. It follows that with probability 1, the walker who starts from 0 will visit 0 infinitely often ("i.o." for short).

In fact, with probability 1, <u>every</u> integer in $\mathbb{Z}$ will be visited by the walker infinitely often.

However, the proportion of the time that the walker spends at any one of these states goes to zero in the limit.

For instance, if the walker starts at 0, the expected number of visits to 0 by time n grows like the square root of n.

To see why this makes intuitive sense, note that at time n, the variance of the position of the walker is $\text{Var}(X_1+X_2+...+X_n) = C\,n$ for suitable C. The distribution of values of the position of the walker is binomial, and hence very close to Gaussian, with standard deviation $\sqrt{C\,n}$. So, while the value $X_1+X_2+...+X_n=0$ is the most likely, the other $2\sqrt{C\,n}$ values between $+\sqrt{C\,n}$ and $-\sqrt{C\,n}$ aren't that much more unlikely, and their total probability is about .68, so each of the values has probability on the order of $1/\sqrt{n}$.

I haven't shown you random or quasirandom simulations of random walk; this is something someone could do for a term project.

Recurrence and transience

A queueing model

(material taken from the excellent book *Probability and Computing: Randomized Algorithms and Probabilistic Analysis* by Michael Mitzenmacher and Eli Upfal, pp. 173-174)

Customers wait in line at an ATM. At each discrete time step, exactly one of the following occurs: a new customer joins the end of the queue (so the size of the queue increases by 1), the person at the head of the queue gets served (so the size of the queue decreases by 1), or nothing happens (so the size of the queue doesn't change). We assume that new customers will not join the queue if it contains n customers; other than that, we assume that the rate at which customers join the queue, and the rate at which customers leave the queue, does not depend on the size of the current queue.

(Note: It is more customary to model queues in continuous time; but if we divide time into very small increments, then we can approximate the continuous-time queueing process by a discrete-time queueing process. Moreover, the assumption that no time intervals contains both an arrival and a departure, or two arrivals, or two departures, makes sense if the time increments are sufficiently small, since the probability of such an event is low. But we'll say more about continuous-queues later in the semester.)

The size of the queue after k time-steps (k = 1,2,3,...) is an integer-valued random process that can be viewed as a Markov chain whose states are 0,1,2,...,n, with transition probabilities $p_{i,j}$ of the form

$$p_{i,i+1} = \lambda \text{ if } i < n$$

$$p_{i,i-1} = \mu \text{ if } i > 0$$

$$p_{i,i} = \begin{cases} 1-\lambda & \text{if } i = 0 \\ 1-\lambda-\mu & \text{if } 1 \le i \le n-1 \\ 1-\mu & \text{if } i = n \end{cases}$$

($p_{i,j}$ is 0 otherwise)

where λ is the probability that someone joins the queue (if the current queue length is less than n) and μ is the probability that someone leaves the queue (if the current queue length is greater than 0).

The chain is ergodic, so there is a unique stationary distribution **w**. We use **w** = **wP** to write

$w_0 = (1\text{-}\lambda)w_0 + \mu w_1,$

$w_i = \lambda w_{i-1} + (1\text{-}\lambda\text{-}\mu)w_i + \mu w_{i+1}$ (for $1 \le i \le n$-1),

$w_n = \lambda w_{n-1} + (1\text{-}\mu)w_n.$

The (non-stochastic) vector $w_i = (\lambda/\mu)^i$ is a solution (check:

In[40]:=
```
(L / M)^i == L (L / M)^(i - 1) + (1 - L - M) (L / M)^i + M (L / M)^(i + 1)
```

Out[40]= $\left(\frac{L}{M}\right)^i = L\left(\frac{L}{M}\right)^{i-1} + (-L - M + 1)\left(\frac{L}{M}\right)^i + M\left(\frac{L}{M}\right)^{i+1}$

In[41]:=
```
Simplify[%]
```

Out[41]= True

which works), so by rescaling to make the entries sum to 1, we get the stationary distribution

$$w_i = (\lambda/\mu)^i / Z$$

with

$$Z = \sum_{i=0}^{n} (\lambda/\mu)^i .$$

This is just a random walk model on {0,1,2,...,n} with leftward, rightward, and holding steps, each occurring with respectively probability μ, λ, and 1-λ-μ, with partially reflecting boundaries at 0 and n.

Now suppose there is no upper limit n to the size of the queue. This is an example of a Markov chain with infinite state space; it is ergodic because it is possible to get from any state to any other. It can be shown that for an ergodic Markov chain on an infinite state space, stationary probability measure may or may not exist, but if it exists it is unique. In our particular case, a stationary probability measure is a vector **w** =

$(w_0, w_1, w_2, \ldots)$ consisting of infinitely many non-negative real numbers summing to 1, such that

$$w_0 = (1-\lambda) w_0 + \mu w_1,$$
$$w_i = \lambda w_{i-1} + (1-\lambda-\mu) w_i + \mu w_{i+1} \text{ (for } i \geq 1).$$

It can be checked that one vector satisfying these equations is **w** with $w_i = (\lambda/\mu)^i$.

If $\lambda < \mu$, then $\Sigma_{i \geq 0} (\lambda/\mu)^i$ converges to the finite number $Z = 1/(1 - \lambda/\mu)$, so the unique stationary probability measure is

$$w_i = (\lambda/\mu)^i (1 - \lambda/\mu).$$

If $\lambda > \mu$, then the rate at which customers arrive is higher than the rate at which they leave, and the queue becomes arbitrarily long as time passes (indeed, its length after k steps is roughly $(\lambda-\mu)k$); there is no stationary distribution on queue-length.

If $\lambda = \mu$, there is no stationary distribution; if we define M_n as the maximum queue length from time 1 to time n, one finds that M_n drifts upward to infinity. However, unlike the case where $\lambda > \mu$ (where the queue length exhibited linear growth), we have $M_n \approx \sqrt{n}$. Furthermore, there will (with probability 1) be infinitely many values of k for which the queue length at time k equals 0, even though as n goes to infinity, the proportion of the time between time 0 and time n that the queue has been empty goes to zero (with probability 1).

The recurrence trichotomy

Think of a Markov chain with state-space S as a sequence of random variables: a random variable X_1 taking its values in S, another S-valued random variable X_2 for which

$$\text{Prob}(X_2 = j \mid X_1 = i) = p_{i,j},$$

another S-valued random variable X_3 for which

$$\text{Prob}(X_3 = k \mid X_1 = i \text{ and } X_2 = j) = p_{j,k},$$

etc.

If S is countably infinite then the entries $p_{i,j}$ do not form an ordinary matrix, but as long as each entry $p_{i,j}$ is non-negative and each "row-sum" $\sum_j p_{i,j}$ is 1, we can (at least in principle) simulate the Markov chain, starting with X_1 equal to some particular state i (or some specified random distribution on states). As usual, we find it helpful to speak of a "walker" moving from one state in S to another.

Typical examples of infinite state spaces are $\mathbb{N}$, $\mathbb{Z}$, $\mathbb{Z}^2$,...

Assume that our transition matrix **P** is "ergodic" in the sense that it is possible to get from any state to any other in some finite number of steps. Then there are two cases:

(1) For every i and j, the expected number of times a walker who starts at i visits j is infinite.

(2) For every i and j, the expected number of times a walker who starts at i visits j is finite.

In case (1), it can be shown (proof omitted) that with probability 1, a walker who starts at *i* will visit *j* infinitely many times, whereas in case (2), it can be shown that with probability 1, a walker who starts at *i* will visit *j* finitely many times.

Case (1) bifurcates, so we get three cases:

(1a) For every *i* and *j*, a walker who starts at *i* will (with probability 1) visit *j* infinitely often, and indeed will visit *j* a positive proportion of the time.

(1b) For every *i* and *j*, a walker who starts at *i* will (with probability 1) visit *j* infinitely often, but asymptotically will visit *j* a vanishingly small proportion of the time.

(2) For every *i* and *j*, a walker who starts at *i* will (with probability 1) visit *j* only finitely often.

Case (1a) is called the <u>positive recurrent</u> case. An example is the queue model in the case $\lambda < \mu$ (aka leftward-biased walk on $\mathbb{N}$).
Case (1b) is called the <u>null recurrent</u> case. An example is the queue model in the case $\lambda = \mu$ (aka unbiased walk on $\mathbb{N}$ with a semireflecting barrier at 0).
Case (2) is called the <u>transient</u> case. An example is the queue model in the case $\lambda > \mu$ (aka rightward-biased walk on $\mathbb{N}$).

In case (1a), there is a stationary probability measure on S; in cases (1b) and (2), there is not.

Unbiased walk on $\mathbb{Z}^d$

It's easy to see that n steps of unbiased random walk on $\mathbb{Z}$ starting from 0 gives a distribution on $\{-n,-n+2,...,n-2,n\}$ that, if you shift values to the right by n and divide by 2, becomes the Binomial(n, $\frac{1}{2}$) distribution. So the position of the walker at time n is approximated by a normal distribution with mean 0 and standard deviation $\sqrt{n}$, and the probability p_n that the walker is actually at 0 after n steps is on the order of $1/\sqrt{n}$. Note that the sum $\sum_{n=1}^{\infty} 1/\sqrt{n}$ diverges, so the expected number of visits to 0 is ∞, and indeed, with probability 1, the walker will visit 0 infinitely many times.

What about random walk on $\mathbb{Z}^2$? That is, from (i,j) in $\mathbb{Z}^2$ we move with equal probability to one of the four neighbors $(i-1,j)$, $(i+1,j)$, $(i,j-1)$, $(i,j+1)$.

Trick: Use instead the four neighbors $(i\pm 1,j\pm 1)$. This just amount to rotating the walk by 45 degrees and scaling by $\sqrt{2}$.

The reason this trick is smart is that the rotated and rescaled walk gives a path in $\mathbb{Z}^2$ such that the x- and y-coordinates of the points in the path are just doing random walk in $\mathbb{Z}$, independently of each other.

Hence, for the rotated and rescaled walk, the probability that a walker who starts at (0,0) is back at (0,0) after n steps is p_n^2, where p_n is the probability that a walker in $\mathbb{Z}$ who starts at 0 is back at 0 after n steps. Since $p_n \approx 1/\sqrt{n}$, we have $p_n^2 \approx 1/n$. So the expected number of visits to (0,0) is $\sum_{n=1}^{\infty} 1/n$, which is infinite, and we can see that this walk (like the walk in 1 dimension) is null-recurrent.

This rotate-and-rescale trick doesn't work in more than two dimensions. E.g., in three dimensions, we want to study a walk where from (i,j,k) in $\mathbb{Z}^3$ we move with equal probability to one of the <u>six</u> neighbors

$(i-1, j, k)$, $(i+1, j, k)$, $(i, j-1, k)$, $(i, j+1, k)$, $(i, j, k-1)$, $(i, j, k+1)$. However, the trick applies to a walk where from (i, j, k) in $\mathbb{Z}^3$ we move with equal probability to one of the <u>eight</u> neighbors $(i \pm 1, j \pm 1, k \pm 1)$. These two walks are genuinely different (they don't just differ by rotation and scaling). However, they do turn out to have something in common: they are both **transient**. We can see that the eight-neighbor walk in $\mathbb{Z}^3$ is transient by using the same method as before: since each of the three coordinates is doing a one-dimensional random walk independently of the other two, the probability that a walker who starts at (0,0,0) is back at (0,0,0) after n steps is
$p_n^3 \approx (1/\sqrt{n})^3$, and $\sum_{n=1}^{\infty} n^{-3/2}$, unlike $\sum_{n=1}^{\infty} n^{-1}$ and $\sum_{n=1}^{\infty} n^{-1/2}$, is finite. So this walk is transient. And the same is true for six-neighbor random walk in $\mathbb{Z}^3$ (proof omitted).

Summary:

In $\mathbb{Z}^d$ with d=1 or d=2, random walk is null-recurrent, and the probability that you will eventually return to the origin, given that you're currently at the origin, is 1.

In $\mathbb{Z}^d$ with $d > 2$, random walk is transient, and the probability that you will eventually return to the origin, given that you're currently at the origin, is some probability p strictly between 0 and 1. In this case, the expected number of returns to the origin, if you start at the origin, is
$\sum_{n=1}^{\infty}$ Prob(the number of visits is at least n)
$= \sum_{n=1}^{\infty} p^n = p/(1-p)$.

Polya: "A drunk man will eventually return home but a drunk bird will lose its way in space."

The "Goldbug walk"

Consider random walk on {-1,0,1,2,3,...} with absorption at {-1,0} in which $p_{i,j}$ is $\frac{1}{2}$ if j equals i-2 or i+1 and is 0 otherwise (as in the "1-D Walk" mode of the rotor-router applet).

This random walk is biased to the left, so with probability 1, the walk will get absorbed in {-1,0}.

More rigorously: View the walker's location at time n (prior to absorption) as a sum of n i.i.d. random variables $X_1+X_2+...+X_n$, where each X_k is +1 with probability $\frac{1}{2}$ and is -2 with probability $\frac{1}{2}$. The Strong Law of

Large Numbers, applied to the i.i.d. process, tells us that with probability 1, $(X_1+X_2+\ldots+X_n)/n \rightarrow -\frac{1}{2}$ (since $-\frac{1}{2}$ is the shared expected value of the X_k's). In particular, the probability that $(X_1+X_2+\ldots+X_n)/n$ will stay positive forever is 0. So absorption occurs with probability 1.

Hence with probability 1, the walker will be absorbed at either -1 or 0, regardless of the walker's starting location.

What is the probability that a walker who starts at 1 eventually gets absorbed at -1?

Call this probability p.

We know that $p > 0$ (why? ...

..?..
because the walker could go from 1 to -1 in a single move), and we also know that $p<1$ (why? ...
..?..
because the walker could go from 1 to 2 to 0).

It'll be handy to think of p as the probability that a random walker who starts at 1 will hit -1 before hitting 0; and we will need to know that p is also the probability that a random walker who starts at 2 will hit 0 before hitting 1 (using the translation-invariance of the transition probabilities for the random walk).

I claim that

$$p = \frac{1}{2} + \frac{1}{2}\,((p)(0)+(1-p)(p)).$$

The first term in the RHS corresponds to the possibility that the walker who starts at 1 will jump to -1 immediately. If this doesn't happen, then the walker jumps from 1 to 2. Let's look at what happens after the walker continues onward from 2. We know that with probability 1, the walker who continues from 2 will eventually hit {0,1}, and indeed, we know that when this happens for the first time, the respective probabilities of the walker being at 0 and 1 are p and 1-p (see the discussion of translation-invariance in the preceding paragraph). In the former case, the game is over; the walker has hit 0 without first hitting 1, so the walker's chance of "winning" (i.e. hitting -1 before 0) is zero. But, in the latter case, the walker is back where he started, and his chance of winning is p, just as it was at the start.

[Go over this analysis a second time, to make sure every one's got it! There'll be a homework problem requiring this style of analysis.]

We can simplify the equation:

$$p = \frac{1}{2} + \frac{1}{2}((p)(0)+(1-p)(p))$$

$$2p = 1 + (1-p)(p)$$

$$2p = 1 + p - p^2$$

$$p^2 + p - 1 = 0$$

$$p = (-1 \pm \sqrt{5})/2$$

Since $p > 0$, we must have $p = (-1+\sqrt{5})/2$ (the reciprocal of the golden ratio; hence Kleber's coinage "goldbug").

Thus, a particle added to the system at 1
escapes to infinity with probability 0,
gets absorbed at -1 with probability $p = (-1+\sqrt{5})/2 \approx$

.618,

and gets absorbed at 0 with probability $1-p=p^2\approx.382$.

Read the first four pages of

http://people.brandeis.edu/~kleber/Papers/rotor.pdf

for a nice description of what's going on with the rotor-router simulation of the walk.

If we let $N(t)$ be the number of times, during the first t runs, that the rotor-walk (with all rotors initially pointing rightward) started from 1 gets absorbed at -1 (rather than 0), then it can be shown that $N(t)$ equals $[p(t+1)]$, the greatest integer less than or equal to $p(t+1)$. From this it follows that the difference between $N(t)/t$ (the empirical rotor-router estimate of the probability of absorption at -1) and p (the exact value of this probability) falls off like $1/t$. (Recall that for true random simulation, the corresponding discrepancy drops off only like $1/\sqrt{t}$.)

Can we do better than $O(1/t)$? Strange but true: We can! Stay tuned...

Rotor-routing on infinite graphs

Prologue

For the Goldbug walk, the set of sites is infinite, "but not very". Specifically, if you do random Goldbug walk, the probability of visiting infinitely many sites is 0; and if you do rotor-router Goldbug walk with lots of bugs, only one of the bugs can wander off to infinity. That's because the walk is biased to the left.

But what about rightward-biased walks on {0,1,2,3,...} in which you expect to wander off to infinity lots of times?

A rightward-biased walk on ℕ

Example: $p_{0,0} = 1$ (i.e., 0 is an absorbing site), and $p_{i,i-1} = \frac{1}{3}$ and $p_{i,i+1} = \frac{2}{3}$ for all $i > 0$. If we start at 1, the probability p that we'll eventually hit 0 is 1/2 (recall our earlier analysis of biased random walk). Can we do this with rotor-routing?

Yes: At each site k with $k > 0$, send
the 1st particle to $k+1$,
the 2nd to $k-1$,
the 3rd to $k+1$,
the 4th to $k+1$,
the 5th to $k-1$,
the 6th to $k+1$,
etc.:

1 →2→3→4→5→6→... (escape)
1→0 (absorption)
1→2→1→2→3→2→3→4→3→... (escape)
1→0 (absorption)
1→2→1→2→3→2→3→4→3→... (escape)
...

Why does it make sense to play this game?

What does it mean to say "the state of the rotors at time infinity"?
This only makes sense if each site has been visited a finite number of times (for then we can define the state of the rotor at time infinity as the state of the rotor after the last time it changed and forever after).
But how do we know each site gets visited only a finite number of times as the particle wanders off to infinity?
Insight: if there were some site that got visited infinitely often, then so would each of its neighbors, and so would each of their neighbors; but then 0 would get visited, so the particle wouldn't wander off to infinity after all!

Why does the game give the right answer?

Once again, the trick is to assign a numerical value to the rotor-configuration. See http://jamespropp.org/584/bugs.pdf for details.

Rotor-router walk in $\mathbb{Z}^2$

Consider random walk on $\{(i,j): i, j \text{ in } \mathbb{Z}\}$ with $p_{x,y} = \frac{1}{4}$ if $y-x=(1,0)$, $(0,1)$, $(-1,0)$, or $(0,-1)$ and $p_{x,y} = 0$ otherwise (two-dimensional random walk as discussed above). As we saw above, this random walk is recurrent, so that with probability 1, a walker who starts at (0,0) will return to (0,0), and indeed return infinitely often, and indeed visit every site in $\mathbb{Z}^2$ infinitely often.

What is the probability that a random walker who starts at (0,0) will visit (1,0) before returning to (0,0)? The answer is $\frac{1}{2}$, and we will see an (incomplete) proof of this in a few weeks.

What is the probability that a random walker who starts at (0,0) will visit (1,1) before returning to (0,0)? This problem is much harder; it turns out that the answer is $\frac{\pi}{8}$.

We can quasirandomize this walk by using rotors at the sites in $\mathbb{Z}^2$ to decide where the walker goes next. See the "2-D Walk" mode of the Canary-Wong applet .)

If we let $N(t)$ be the number of times, during the first t runs, that the rotor-walk (with all rotors initially pointing rightward) started from (0,0) gets absorbed at (1,1) (rather than (0,0)), then it can be shown that $N(t)/t$ (the empirical rotor-router estimate of the probability of absorption at (1,1)) approaches $\frac{\pi}{8}$ (the exact value of this probability), and that the difference falls off at least as fast as $(\log t)/t$.

Can we do better than $O((\log t)/t)$? I don't know; I have a method that (for t up to 100,000) satisfies $|N(t)/t - \frac{\pi}{8}| < 1/t^{1.5}$, but I have no proof that it works for larger t, nor any heuristic explanation of why the exponent should be about 1.5 and not something else (larger or smaller).

One reason why we might expect to do better than (log t)/ t is the fact that the behavior of the hitting sequence (i.e., the sequence whose nth term is 1 or 0 according to whether the nth stage results in absorption at (1,1) or (0,0)), although not periodic, is "almost-periodic", and even appears to have some periodic subsequences.