Random Walks with Restarts, 3 Examples

Nathan McNew Dartmouth College nathan.g.mcnew.gr@dartmouth.edu

March 3, 2013

Abstract

While infinite graphs generally have infinite expected hitting times, we prove that the hitting times become finite if we allow a random walk to restart at the source vertex at any point in the walk. We then compute explicitly the expected hitting time with restarts for Cayley trees and the hypercube (including a small correction to the original paper from which this material is drawn.) Finally we state a theorem regarding these hitting times on the integer lattice of dimension d, and provide a surprisingly accurate heuristic argument for d=2.

1 Introduction

A natural question when considering a random walk on a graph is to compute the expected time to get from one vertex of the graph to another. This quantity, the hitting time, has been well studied for finite graphs. It is not as useful a statistic, however, for infinite graphs, as the hitting time between vertices is generally infinite. For example on the infinite line, \mathbb{Z} the expected time to get from 0 to 1 is infinite. (The hitting time between vertices in an infinite graph may be finite in one direction if the graph obtained by removing the target vertex leaves the source vertex isolated in a finite component of the graph.)

If, however, we allow the random walk to "restart" at the source vertex we might expect the "hitting time with restarts" to be finite, as it eliminates the possibility a walk will wander to far off "towards inifinity." We can make this idea rigorous, so long as our graph is reasonable:

Theorem 1.1. In any locally finite graph, the expected hitting time from any vertex s to another vertex t is finite.

Proof. Let k be the length of a shortest path from s to t. Since the graph is locally finite there are a finite number of vertices with distance less than k from s. Let d be the maximum degree over all such vertices. Now use the simple (but far from optimal) restart rule: Take k steps from s, then restart if not at t. The probability of selecting the path from s to t above is at least $\left(\frac{1}{d}\right)^k$. Treating this as a Bernoulli trial, we see that the expected hitting time with restarts is at most d^k .

It is shown in [1] that the hitting time with restarts is equal to a quantity called the grade $\gamma(s,t)$, which can also be defined in greater generality, allowing different costs when moving at different places in the graph. The theorem above is proven in much greater generality in [1], Theorem 9.1. In what follows we demonstrate three examples where we can explicitly compute the expected hitting time with restarts.

2 The *n*-dimensional hypercube

We start by considering a finite example. Let Q^n denote the n-dimensional hypercube, where we represent each vertex by a binary sequence $u = (u_1, u_2, \dots u_n)$. Vertices are connected if and only if their binary sequences differ in exactly one position. By the symmetry of the hypercube, we can always assume the target vertex is the origin. Now, define the j-th level of Q^n to be all vertices whose binary representation contains exactly j 1's. We would like to find the expected hitting time, allowing restarts, from a vertex in level k to the origin.

Again exploiting the symmetry of the hypercube, we see that any time our random walk takes us to level k+1 we are better off restarting at our original vertex. We can also use the symmetry to simplify the idea of restarting: Since some rotation of the hypecube fixes the origin and interchanges any two vertices in the same level, we can restart any time we get to level k+1 by simply returning to our prior vertex in level k rather than actually returning to our initial vertex s without changing the problem.

Therefore, we construct the truncated hypercube Q_k^n by removing all vertices of level greater than k from Q_n and replace each of the n-k edges from each vertex of level k to level k+1 with a loop. By the discussion above, a simple random walk on this graph will be equivalent to a random walk with restarts on Q_n .

Let T_j denote the expected time to get from level j to level j-1 in Q_k^n . So the expected hitting time with restarts (equivalently the grade) $\gamma(k,0) = \sum_{j=1}^k T_j$. Getting from level k to level k-1 requires simply not choosing one of the loops out of level k, and hence as a bernouli trial has expected time $T_k = \frac{n}{k}$.

Now for arbitrary 0 < j < k, we must take at least one step, which will either take us to level j-1 or j+1. If we step to j-1, we are done, but if we go to j+1, we must first get back to level j, and then on to level j-1. This happens with probability $\frac{n-j}{n}$, which gives us the recurrence: $T_j = 1 + \frac{n-j}{n}(T_j + T_{j+1})$. The paper, [1], from which this example comes, makes the assertion: It is straightforward to verify that

$$T_j = \frac{1}{\binom{n-1}{n-j}} \sum_{i=n-j}^n \binom{n}{i}.$$

Plugging this expression into the recurrence to verify we obtain

$$\frac{j}{n} \left(\frac{1}{\binom{n-1}{n-j}} \sum_{i=n-j}^{n} \binom{n}{i} \right) = 1 + \frac{n-j}{n} \left(\frac{1}{\binom{n-1}{n-j-1}} \sum_{i=n-j-1}^{n} \binom{n}{i} \right)
j \left(\sum_{i=n-j}^{n} \binom{n}{i} \right) = n \binom{n-1}{n-j} + (n-j) \left(\frac{j-1}{n-j} \sum_{i=n-j-1}^{n} \binom{n}{i} \right)
\sum_{i=n-j}^{n} \binom{n}{i} = j \binom{n}{n-j} + (j-1) \binom{n}{n-j-1}
\sum_{i=0}^{j-1} \binom{n}{i} = (j-1) \binom{n}{n-j} + (j-1) \binom{n}{n-j-1} = (j-1) \binom{n+1}{n-j+1}.$$

Note that this gives us a closed form expression for the sum of the first j binomial coefficients on level n of Pascal's triangle! Unfortunately, no such closed form expression exists—something has gone wrong. (Consider this a reminder to verify assertions in papers, even if the paper says that it is easy to do so!) We derive the correct solution by writing out the first few terms and recognizing the pattern:

$$T_{k-1} = \frac{n}{k-1} + \frac{n-k+1}{k-1} \left(\frac{n}{k} \right)$$

$$T_{k-2} = \frac{n}{k-2} + \frac{n-k+2}{k-2} \left(\frac{n}{k-1} \right) + \frac{n-k+2}{k-2} \left(\frac{n-k+1}{k-1} \right) \left(\frac{n}{k} \right)$$

$$\vdots$$

$$T_{j} = n \sum_{i=0}^{k-j} \frac{(n-j)(n-j-1)\cdots(n-j-i+1)}{(j)(j+1)\cdots(j+i)} = \frac{n}{j} \sum_{i=0}^{k-j} \frac{(n-j)!j!}{(n-(j+i))!(j+i)!} = \frac{n}{j} \sum_{i=0}^{k-j} \frac{\binom{n}{j+i}}{\binom{n}{j}}$$

$$\gamma(k,0) = \sum_{j=0}^{k} T_{j} = \sum_{j=0}^{k} \frac{n}{j\binom{n}{j}} \sum_{i=0}^{k-j} \binom{n}{j+i}$$

3 Cayley Trees

We now consider an easy, infinite example, Cayley trees. The d-regular Cayley tree, T^d is the unique, infinite tree in which each vertex has degree d. Like the hypercube, this graph is completely symmetric, so the target vertex can be assumed to be the "root." Also, analogously, our restart rule will again be to restart anytime we are further from the root than when we started.

Applying the same analysis as in the hypercube, we see that the hitting time with restarts from level k to the root, $\gamma(k,0)$, is the hitting time on the truncated tree T_k^d , which contains the first k levels of T^d , and where each vertex in level k is given d-1 loops. Now, we know:

Theorem 3.1. In a finite tree (possibly with some loops) the length of any whirling tour from v to t is the expected hitting time from v to t. (See [1] for the proof and relevent definitions.)

In light of this theorem, it suffices to compute a whirling tour on T_k^d . Going from level j+1 to an adjacent vertex u on level j on a such a whirling tour requires itself a whirling tour on the subgraph T_{k-j}^d rooted at u. Let W_j be the length of such a tour.

 T_{k-j}^d has $\sum_{i=0}^{k-j-1} (d-1)^i$ edges and $(d-1)^{k-j}$ loops. Each edge is used twice (except for the final edge) and each loop once, so

$$W_j = 2\left(\sum_{i=0}^{k-j-1} (d-1)^i\right) - 1 + (d-1)^{k-j} = \frac{(d-1)^{k-j}-1}{d-2} - 1 + (d-1)^{k-j}.$$

Summing this over all levels, we obtain

$$\gamma(k,0) = \sum_{i=0}^{k-1} W_i = \sum_{i=0}^{k-1} \left(\frac{(d-1)^{k-j}-1}{d-2} - 1 + (d-1)^{k-j} \right)$$
$$= \frac{(d-1)^{k+2} - (d-1)(d+(d-2)k-1)}{(d-2)^2} \quad (d > 2).$$

4 The infinite grid, \mathbb{Z}^d

Notice that the above solution works only when d > 2. In the case d = 2 we just have a walk on the integers, \mathbb{Z} . It is left as an exercise to the reader to show that the expected time to get from an integer x to the origin is $\gamma(x,0) = |x|(|x|+1)$. We consider the case d > 2.

We still have translational symmetry on the grid \mathbb{Z}^d , so we can assume the target is the origin, however we don't have complete (rotational) symmetry, so the hitting time cannot be dependent solely on the number of steps from the origin (the l_1 norm.) We can, however, get an asymptotic estimate. [2]

Theorem 4.1 (Janson and Peres, 2010). For a random walk with restarts on \mathbb{Z}^d ,

$$\gamma(x,0) = 2|x|^2 \log|x| + (2\gamma_e + 3\log(2) - 1)|x|^2 + O(|x|\log|x|)$$

when d=2, where |x| is the Euclidean distance from the origin and γ_e is the Euler-Mascheroni Constant, $\gamma_e:=\lim_{N\to\infty}\sum_{i=1}^N-\log N$. For $d\geq 3$.

$$\gamma(x,0) = \frac{\omega_d}{p_d} |x|^d + O(|x|^{d-1})$$

where ω_d is the volume of the unit ball in \mathbb{R}^d , and p_d probability that a random walk in \mathbb{Z}^d never returns to 0.

The proof is complicated, and involves submartingales and the optional sampling theorem to bound the grade using bounds for the local variance of a harmonic function, and various results on the potential kernel (for \mathbb{Z}^2 .)

Instead, we present a heuristic argument for the case d = 2. Define as ulual the jth level of the grid to be the points j steps from the origin, and bserve that the 2d grid is nearly symmetric, except on the axes. Each non-axis point has 2 edges to level j + 1 and 2 edges to j - 1. There are a total of 16j edges: 8j + 4 to level j + 1 and 8j - 4 to level j - 1.

So, we ignore the asymmetry, and blindly apply the method we've used before to calculate $\gamma(k,0)$ from an "arbitrary" point k steps from the origin. Truncate the grid at level k as usual, replacing each outgoing edge with a loop, and let S_j be the expected time to get from level j to level j-1.

Now, applying the same analysis as in the hypercube we find that $S_k = \frac{16k}{8k-4} \approx 2$ and $S_j = 1 + \frac{8j+4}{16j}(S_j + S_{j+1})$

It is straightforward to verify that: (Remember the lesson above?)

$$S_j = 2\left(\frac{k^2 - j^2 + k + j - 1}{2j - 1}\right)$$

So,

$$\gamma(k,0) = \sum_{j=1}^{k} S_j = \sum_{j=1}^{k} 2\left(\frac{k^2 - j^2 + k + j - 1}{2j - 1}\right)$$

$$\approx 2\int_{j=1}^{k} \left(\frac{k^2 - j^2 + k + j - 1}{2j - 1}\right) dj$$

$$= \frac{1}{4}((4k(k+1) - 3)\log(2k - 1) - 2(k-1)k)$$

$$= k^2\log(2k) - \frac{1}{2}k^2 + O(k\log k).$$

So we see the $|x|^2 \log |x|$ and $|x|^2$ terms appearing as in the theorem. Our heuristic is actually better than it first appears, however. As a final observation, note that as long as we stay away from the axes, our symmetry assumption isn't so bad. Take x = (a, a) to be a point on the diagonal and recall that for such an x (at level k) $k = ||x||_1 = \sqrt{2}|x|$. Plugging this in our estimate, we find that

$$k^{2} \log(2k) - \frac{1}{2}k^{2} \to (\sqrt{2}|x|)^{2} \log(2(\sqrt{2}|x|)) - \frac{1}{2}(\sqrt{2}|x|)^{2}$$

$$= 2|x|^{2} (\log|x| + \log(2^{3/2})) - |x|^{2}$$

$$= 2|x|^{2} \log|x| + (3\log 2 - 1)|x|^{2}.$$

This agrees exactly with the result from the theorem, with the exception of the factor of $2\gamma_e$. But, we approximated a sum with an integral, and obtained a log. So, recalling the definition of γ_e , it would make sense that replacing the $\log |x|$ with a $\log |x| + \gamma_e$ would give us a better assymptotic estimate. Doing so gives us

$$= 2|x|^2 \log|x| + (2\gamma_e + 3\log 2 - 1)|x|^2,$$

the same expression given in the theorem above.

References

- [1] Dumitriu, Ioana and Tetali, Prasad and Winkler, Peter (2003), On playing golf with two balls, SIAM Journal on Discrete Mathematics 16, 1604–615.
- [2] Janson, Svante and Peres, Yuval (2012), SIAM Journal on Discrete Mathematics 26, 537–547.