

WAITING FOR A BAT TO FLY BY (IN POLYNOMIAL TIME)

ITAI BENJAMINI, GADY KOZMA, LÁSZLÓ LOVÁSZ, DAN ROMIK, AND GÁBOR TARDOS

ABSTRACT. We observe returns of a simple random walk on a finite graph to a fixed node, and would like to infer properties of the graph, in particular properties of the spectrum of the transition matrix. This is not possible in general, but at least the set of eigenvalues can be recovered under fairly general conditions, e.g. when the graph has a node-transitive automorphism group. The main result is that by observing polynomially many returns, it is possible to estimate the spectral gap of such a graph up to a constant factor.

1. INTRODUCTION

A spelunker has an accident in the cave. His lamp goes out, he cannot move, all he can hear is a bat flying by every now and then on its random flight around the cave. What can he learn about the shape of the cave?

In other words: What can we learn about the structure of a finite graph using only information obtained by observing the returns of a random walk on the graph to this node?

Let $G = (V, E)$ be a connected simple graph with $n = |V| > 1$ vertices, and let $r \in V$ be a fixed node. Let $w_0 = r, w_1, w_2, \dots$ be the steps of a simple random walk on G starting from r . Assume that we observe the *return time sequence*, the infinite sequence of (random) times $0 < T_1 < T_2 < \dots$ when the walk visits r . Alternatively this can be described as a sequence a_1, a_2, a_3, \dots of bits, where $a_i = 1$ if the walk is at r at time i , 0 otherwise. Note that $T_2 - T_1, T_3 - T_2, \dots$ are independent samples from the same distribution as T_1 , which we call the *return distribution* of G to r .

We say that a parameter $p(G, r)$ of the graph G and root r can be reconstructed (from the return time sequence), if for every two rooted graphs (G, r) and (G', r') for which the return time sequence has the same distribution, we have $p(G, r) = p(G', r')$.

Which graph parameters can be reconstructed from the return time sequence? There is a trivial way to construct different graphs with the same return sequence: take two isomorphic copies and glue them together at the root. Sometimes it makes sense to assume that we also know the degree $d(r)$ of the root. In this case, we can reconstruct the number of edges through

$$|E| = d(r)E(T_1)/2. \tag{1}$$

If the graph is regular, then we can reconstruct the number of nodes:

$$n = |V| = E(T_1). \tag{2}$$

Another trivial example is to observe if all the numbers T_i are even. This is so if the graph is bipartite, and it happens with probability 0 otherwise.

A natural candidate for a reconstructible quantity is the spectrum of the transition matrix M of the random walk on G . Let $\lambda_1 = 1, \lambda_2, \dots, \lambda_n$ be the eigenvalues of M , arranged in decreasing order. Bipartiteness is equivalent to saying that $\lambda_n = -1$.

We are going to show by a simple example that the spectrum is not reconstructible in general. On the other hand, we show that if λ is an eigenvalue of G which has an eigenvector $v \in \mathbb{R}^V$ such that $v_r \neq 0$, then λ is reconstructible. We note that the *multiplicity* of λ is not necessarily reconstructible.

A special case where the eigenvector condition above is satisfied for all eigenvalues is when G is node-transitive. We don't know whether in this case the multiplicities are reconstructible.

Of particular interest is the issue of *efficient reconstruction*, by which we mean observing a polynomial (or expected polynomial) number of returns. We consider this question in the case of the *spectral gap* $\tau = 1 - \lambda_2$. Assuming the graph is node transitive, we describe a procedure to estimate τ up to a constant factor, using just polynomially many (in n) of the first values of the T_i . We give an example of a graph where the spectral gap cannot be recovered *at all* from observations made at one particular node.

This question was first mentioned, together with other related problems, in [3]. Another related work is that of Feige [4] which presents a randomized space-efficient algorithm that determines whether a graph is connected. His method uses return times of random walks to estimate the size of connected components.

2. EXAMPLES


Example 1. Consider the two trees in Figure 1. The distribution of the return time to the root is the same in both trees (see later). The eigenvalues of the tree on the left are

$$1, \sqrt{3}/2, \sqrt{6}/4, 0, 0, 0, 0, 0, -\sqrt{6}/4, -\sqrt{3}/2, -1,$$

while the eigenvalues of the tree on the right are

$$1, \sqrt{3}/2, \sqrt{3}/2, \sqrt{6}/4, 0, 0, 0, -\sqrt{6}/4, -\sqrt{3}/2, -\sqrt{3}/2, -1.$$

Note that the eigenvalues are the same, but their multiplicities are different.



$$h_{G_1}(x) = h_{G_2}(x) = \frac{1 + h(x)}{1 + (1 - x)h(x)}, \text{ where } h(x) = \frac{12 - 3x}{4 - 3x}$$

FIGURE 1. Two trees with the same return times but different spectra

Example 2. Let T be a tree in which all internal nodes have degree $d+1$ and which has a “root” r such that all leaves are at distance h from the root. We construct a graph G by adding a d -regular graph on the leaves.

For a fixed h and d , all graphs obtained this way are $(d+1)$ -regular graphs, and the distribution of the return time to the root is the same in all such graphs. On the other hand, graphs obtained this way can have very different properties. If we add an expander on the leaves, the graph G will be an expander. (Recall that G is a c -expander iff $|\partial S| > c|S|$ for every non empty set of vertices S with $|S| < |G|/2$. For background on expanders and spectral gap see e.g. [5].) If we connect “twin” leaves to each other, and also match up “cousins” to get d new edges at each node, then for $h > 2$ the root will be a cutpoint. For expanders, the eigenvalue gap $\lambda_1 - \lambda_2$ is bounded from below by a positive function of d , while for the graphs with cutpoints in the middle the eigenvalue gap tends to 0 as $h \rightarrow \infty$.

3. PREPARATION: SOME ALGEBRA AND GENERATING FUNCTIONS

3.1. Return probabilities and eigenvalues. Denote by $P_k(x, y)$ the probability that a simple random walk on G starting at $x \in V$ will be at $y \in V$ at time k . Clearly

$$P_k(x, y) = e_x^\top M^k e_y. \quad (3)$$

Here M , the transition matrix of the random walk on G , is not symmetric, but we can consider the symmetrized matrix $N = DM D^{-1}$, where D is a diagonal matrix with the positive numbers $\sqrt{d(i)}$ in the diagonal. The matrix N has the same eigenvalues as M , and so we have

$$P_k(r, r) = \sum_{i=1}^n f_i(r)^2 \lambda_i^k, \quad (4)$$

where f_1, f_2, \dots, f_n is an orthonormal basis of eigenfunctions of N corresponding to the eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_n$.

We note that if the graph is node-transitive, then the value $P_k(r, r)$ is the same for all r , and hence by averaging (4) we get the simpler formula

$$P_k(r, r) = \frac{1}{n} \text{trace}(M^k) = \frac{1}{n} \sum_{i=1}^n \lambda_i^k. \quad (5)$$

At some point, it will be convenient to consider the *lazy version* of our chain, i.e., the Markov chain with transition matrix $M' = (1/2)(I + M)$ (before doing a step, we flip a coin to decide if we want to move at all). The observer can easily pretend that he or she is watching the lazy version of the chain: after each step, he flips a coin in quick succession until he tosses a head, and advances his watch by the number of coinflips. The distribution after k lazy steps is easy to compute from (3):

$$P'_k(x, y) = 2^{-k} e_x^\top (I + M)^k e_y = 2^{-k} \sum_{j=0}^k \binom{k}{j} e_x^\top M^j e_y = 2^{-k} \sum_{j=0}^k \binom{k}{j} P_j(x, y). \quad (6)$$

The main advantage of the lazy chain is that its eigenvalues are nonnegative. Furthermore, for a lazy chain we have

$$\lambda_2 + \dots + \lambda_n = \text{trace}(M) - 1 = \frac{n}{2} - 1.$$

3.2. The generating function of return times. Let us introduce the generating function

$$F(t) = \sum_{k=0}^{\infty} P_k(r, r)t^k = \sum_{i=1}^n f_i(r)^2 \frac{1}{1 - t\lambda_i}. \quad (7)$$

There are several other useful expressions for $F(t)$; for example, we get from (3) that

$$F(t) = e_r^\top (I - tM)^{-1} e_r,$$

and expressing this in terms of determinants, we get

$$F(t) = \frac{\det(I' - tM')}{\det(I - tM)}, \quad (8)$$

where M' is the matrix obtained from M by deleting the row and column corresponding to the root, and I' is the $(n-1) \times (n-1)$ identity matrix.

It will be convenient to do a little algebraic manipulation. The reciprocal of this function is also an interesting generating function:

$$\frac{1}{F(t)} = 1 - \sum_{k=1}^{\infty} s_k t^k, \quad (9)$$

where $s_k = P(T_1 = k)$ is the probability that the first return to the root occurs at the k -th step. This function has a zero at $t = 1$, so it makes sense to divide by $1 - t$, to get the analytic function

$$\frac{1}{(1-t)F(t)} = \sum_{k=0}^{\infty} z_k t^k, \quad (10)$$

where

$$z_k = 1 - \sum_{j \leq k} s_j = \sum_{j > k} s_j$$

is the probability that the random walk does not return to the root during the first k steps.

4. RECONSTRUCTING NONDEGENERATE EIGENVALUES

It is these formulas which form the basis of learning about the spectrum of G from the visiting times of the random walk at x , since $P_k(r, r)$ is determined by the distribution of return times, and can be easily estimated from the visiting times (see section 6). We call an eigenvalue of M *nondegenerate* if at least one of the corresponding eigenfunctions $f_i(x)$ satisfies $f_i(r) \neq 0$. One can see from (4) that the nonzero nondegenerate eigenvalues are determined by the distribution of return times. Using $\sum_{i=1}^n f_i(r)^2 = 1$ for the orthonormal basis f_i we conclude that whether zero is a nondegenerate eigenvalue of M is also determined. The return time distribution determines $F(t)$ and this can also be used to find the nondegenerate eigenvalues: the poles of $F(t)$ are exactly the reciprocals of the nonzero, nondegenerate eigenvalues of M . Zero is a nondegenerate eigenvalue if and only if $\lim_{t \rightarrow \infty} F(t) > 0$. Then we get

Proposition 1. *If two rooted graphs have the same return time distribution, then they have the same nondegenerate eigenvalues.*

Let us remark that if G has a node-transitive automorphism group, then *every eigenvalue of M is nondegenerate*. Indeed, every eigenvalue has an eigenvector, which does not vanish at some node; by node-transitivity, it also has an eigenvector that does not vanish at the root.

Let us also remark that the *multiplicity* of a nondegenerate eigenvalue is not uniquely determined: 0 is a nondegenerate eigenvalue of both trees in Example 1, but it has different multiplicities in the two. Furthermore, degenerate eigenvalues are not determined by the return times: the second largest eigenvalues of the transition matrices of the two $(d+1)$ -regular graphs constructed in Example 2 are different. It follows from Proposition 1 that at least for the second graph, the second largest eigenvalue is degenerate.

5. TREES

We want to put Example 1 in broader context. For trees, we can simplify the generating function a bit: Since trees are bipartite, we have $z_{2k} = z_{2k+1}$, and hence it makes sense to divide by $t+1$ and then substitute $x = t^2$. It will be convenient to scale by the degree of the root, and to work with the function

$$h_G(x) = d(r) \sum_{k=0}^{\infty} z_{2k} x^k = \frac{d(r)}{(1-x)F(\sqrt{x})}. \quad (11)$$

It is easy to see that we did not lose any information here: we have $h_{G_1}(x) = h_{G_2}(x)$ for two trees G_1 and G_2 if and only if they have the same return time distribution and their roots have the same degree.

For a rooted tree with a single edge, $h_G(x) = 1$. If a rooted tree G is obtained by gluing together the roots of two rooted trees G_1 and G_2 , then

$$h_G(x) = h_{G_1}(x) + h_{G_2}(x). \quad (12)$$

This is easily seen by conditioning on which tree the random walk starts in. Furthermore, if we attach a new leaf r' to the root r of a tree G and make this the root to get a new rooted tree G' , then

$$h_{G'}(x) = \frac{1 + h_G(x)}{1 + (1-x)h_G(x)}. \quad (13)$$

To see this, consider a walk on G' starting at r' , and the probability z'_{2k} that it does not return to r' in the first $2k$ steps. Let z_{2k} denote the corresponding probability in the random walk in G starting from r . Let d be the degree of r in G . The first step of the random walk in G' leads to r . If the walk does not return to r' for $2k$ steps for some $k \geq 1$, then the second step has to use a different edge, this happens with probability $d/(d+1)$. We can view the walk now as a random walk on G until it returns to r . The probability that this happens after $2j$ steps is $z_{2j-2} - z_{2j}$. If $j \geq k$ then the walk will certainly not return to r' in the first $2k$ steps. If $j < k$, then we can think of the situation as just having made a step from r' , and so the probability that we don't return to r' in the next $2k - 2j - 1$ steps is z'_{2k-2j} . Hence we get the equation

$$z'_{2k} = \frac{d(r)}{d(r)+1} \left(z_{2k-2} + \sum_{j=1}^{k-1} (z_{2j-2} - z_{2j}) z'_{2k-2j} \right).$$

Multiplying by x^k and summing over all $k \geq 0$, we get (13).

Formulas (12) and (13) imply that h_G is a rational function with integral coefficients. They also provide us with a fast way to compute h_G , and through this, to verify that the two trees in Example 1 have the same return distribution. But we can get more, a way to generate many such pairs.

Suppose that we find a linear dependence between functions h_G for various trees G . This can be written as

$$a_1 h_{G_1} + \cdots + a_k h_{G_k} = b_1 h_{G'_1} + \cdots + b_m h_{G'_m}$$

with some positive integers $a_1, \dots, a_k, b_1, \dots, b_m$. Now if we glue together the roots of a_1 copies of G_1, \dots, a_k copies of G_k to get G , and the roots of b_1 copies of G'_1, \dots, b_m copies of G'_m to get G' , then by (12) we'll have

$$h_G(x) = h_{G'}(x).$$

We can add a new root to both if we prefer to have an example rooted at a leaf.

Obviously, we only need to look for trees rooted at leaves. To find such linear dependencies, it is natural to find trees for which $h_G(x)$ is “simple”, namely the ratio of two linear functions, and then find three with a common denominator. A general example is a tree $G = G_{a,b}$ of height 3, where the neighbor of the root has degree a and has $a - 1$ neighbors of degree b . We can allow the degenerate cases $b = 1$ (when G is a star rooted at a leaf) and $a = 1$ (when G is a single edge). It is easy to compute that

$$h_G(x) = \frac{ab - (b - 1)x}{ab - (ab - 1)x}.$$

So if we fix a k which is not a prime, and consider trees $G = G_{a,b}$ with $ab = k$, they all have the same denominator $k - (k - 1)x$, and so for any three of them their functions h_G will be linearly dependent. The simplest choice is $k = 4$, when we get the trees $G_{1,4}$ (a single edge), $G_{2,2}$ (a path of length 3) and $G_{4,1}$ (a 4-star). Simple computation shows that

$$h_{G_{1,4}} - 3h_{G_{2,2}} + 2h_{G_{4,1}} = 0.$$

Gluing these together as described above, and adding a new root for good measure, gives the two trees in Example 1.

Using (8) and (11), it is not hard to see that the roots of the numerator of $h_G(x)$ are the squared reciprocals of the nondegenerate non zero eigenvalues of G , except for the trivial nondegenerate eigenvalues ± 1 . The multiplicities, as we have seen, are not necessarily determined by h_G .

Remark. In the special trees constructed above, the squareroots of the root of the denominator are exactly the degenerate eigenvalues of G . We don't know if this is always so. An interesting open question seems to be whether the degenerate eigenvalues are reconstructible for trees. Several other questions concerning the functions $h_G(x)$ arise: Are the trees above the only trees for which h_G has linear numerator and denominator? Are there interesting trees for which $h_G(x)$ is the ratio of quadratic polynomials? What can be said about $h_G(x)$ for trees of depth 4?

6. EFFECTIVE RECONSTRUCTION

In the previous section, we assumed that the exact distribution of the return time is known, which is the same as saying that we can observe the random walk forever. In this section we are concerned with determining quantities after observing a polynomial number of returns.

6.1. Estimating return probabilities. We show that we can estimate $P_k(r, r)$, the probability that the random walk starting from r is at r at time k , from the observation of polynomially many return times. Fix k and observe the returns T_1, T_2, \dots until the first T_{i_1} with $T_{i_1} \geq k$; call this period an *experiment*. Call the experiment *successful* if $T_{i_1} = k$. The probability that an experiment is successful is $P_k(r, r)$. Note that observing the next k steps and then until the first return (i.e., $T_{i_1+1}, \dots, T_{i_2}$ with the smallest i_2 such that $T_{i_2} \geq T_{i_1} + k$) is an independent experiment.

Continuing in this way, we obtain a sequence of independent events A_1, A_2, \dots with the same probability $p = P_k(r, r)$, and we want to estimate p . Let ε, δ be given positive numbers, and set $m = \varepsilon^{-2}\delta^{-1}$. By Chebyshev's inequality, after observing m of these events, the relative frequency of their occurrence is closer than ε to its mean, p , with probability at least

$$1 - \frac{p(1-p)}{m\varepsilon^2} > 1 - \delta.$$

The amount of time a particular experiment takes is a random variable, whose expectation is k plus the time it takes to get back to r after k steps. This can be bounded by the maximum hitting time between nodes, which is $O(n^3)$. Summing up,

Proposition 2. *In an expected time of $O((k + n^3)\varepsilon^{-2}\delta^{-1})$ we can compute an estimate of $P_k(r, r)$ which is within an additive error of ε with probability at least $1 - \delta$.*

6.2. Reconstructing the eigenvalue gap. We restrict our attention to node-transitive graphs, in which case we can use the trace formula (5). We can use (2) to reconstruct the number of nodes n . Furthermore, we assume that the chain is lazy, so that its eigenvalues are nonnegative, and their sum is $n/2$.

For a lazy chain, $P_k(r, r)$ tends to $1/n$ monotone decreasing. Furthermore, (5) implies that setting

$$q_k = P_k(r, r) - \frac{1}{n},$$

we have

$$nq_{k+1} = \sum_{i=2}^n \lambda_i^{k+1} \geq \frac{1}{n-1} \left(\sum_{i=2}^n \lambda_i \right) \left(\sum_{i=2}^n \lambda_i^k \right) = \frac{1}{n-1} (\text{trace}(M) - 1) nq_k,$$

and hence

$$q_{k+1} \geq \frac{1}{3} q_k \tag{14}$$

for $n \geq 4$ (which we assume without loss of generality).

We can try to compute recursively $\lambda_1 = 1$ and

$$\lambda_i = \lim_{k \rightarrow \infty} \left[P_k(r, r) - \sum_{j=1}^{i-1} \frac{\lambda_j^k}{n} \right]^{1/k}.$$

This, however, does not seem to give an effective means of estimating λ_i in polynomial time. But to estimate at least the eigenvalue gap $\tau = 1 - \lambda_2$ we can use the following fact.

Lemma 1. *We have*

$$\left(1 + \frac{\ln n}{\ln q_k}\right) (1 - q_k^{1/k}) \leq \tau \leq 1 - q_k^{1/k}. \quad (15)$$

Proof. From (5),

$$P_k(r, r) = \frac{1}{n} + \sum_{i=2}^n \frac{\lambda_i^k}{n},$$

and hence

$$\frac{\lambda_2^k}{n} \leq \sum_{i=2}^n \frac{\lambda_i^k}{n} = q_k \leq \lambda_2^k.$$

Thus

$$1 - (nq_k)^{1/k} \leq \tau \leq 1 - q_k^{1/k}.$$

Using the elementary inequality

$$\frac{1-x}{1-y} \leq \frac{\ln x}{\ln y}$$

valid for $0 < x < y < 1$, (15) follows. Simply take $x = q_k^{1/k}$, $y = (nq_k)^{1/k}$. Note that $nq_k > 1$ makes the lower bound in (15) trivially true. \square

Let $c > 1$. It follows that if we find an integer $k > 0$ such that $q_k < 1/n^c$, then $1 - q_k^{1/k}$ is an estimate for the eigenvalue gap τ which is within a factor of $1/(1 - 1/c)$ to the true value. So, if we want to estimate τ to within a factor of $1 \pm \varepsilon$, for some $0 < \varepsilon < 1$, then picking $c = 5/\varepsilon$ we get that $1 - q_k^{1/k}$ estimates τ to within a factor of $1 \pm \varepsilon/4$. But of course we don't know q_k exactly, only with an additive error: by proposition 2, we can estimate q_k in polynomial time with an additive error less than (say) ε/n^c , with high probability.

It is known (see, e.g., [2, Chapter 4]) that the eigenvalue gap of a connected node-transitive graph with n nodes is at least $1/n^2$. So we get that for $k \geq K_0 = (c+1)n^2 \ln n$,

$$q_k \leq n \left(1 - \frac{1}{n^2}\right)^k < ne^{-k/n^2} \leq \frac{1}{n^c}.$$

Applying proposition 2, we can compute an approximation Q_k of q_k that is within an additive error of $\varepsilon/(8n^c)$ with probability $\geq 1 - \delta/(\log_2 K_0)$. By binary search, we can find a k in the interval $[0, K_0]$ for which

$$Q_k \leq 1/n^c \text{ but } Q_{k-1} > 1/n^c. \quad (16)$$

Proposition 3. *Let $0 < \varepsilon < 1$, and let $0 < \delta < 1$. Let Q_k be as above and let k be defined by (16). Then $1 - Q_k^{1/k}$ is an estimate of the spectral gap τ that is within a factor of $1 \pm \varepsilon$ of τ with probability at least $1 - \delta$.*

Proof. With large probability, we have

$$|q_m - Q_m| < \frac{\varepsilon}{8n^c}$$

for all m for which we compute Q_m , in particular for $m = k - 1$ and $m = k$. Using (14),

$$q_k \geq \frac{1}{3}q_{k-1} \geq \frac{1}{3} \left(Q_{k-1} - \frac{\varepsilon}{8n^c} \right) \geq \frac{1}{4n^c},$$

and also

$$Q_k \geq q_k - \frac{\varepsilon}{8n^c} \geq \left(1 - \frac{\varepsilon}{2}\right)q_k. \quad (17)$$

Similarly,

$$Q_k \leq \left(1 + \frac{\varepsilon}{2}\right)q_k. \quad (18)$$

We claim that

$$1 - \frac{\varepsilon}{2} \leq \frac{1 - Q_k^{1/k}}{1 - q_k^{1/k}} \leq 1 + \frac{\varepsilon}{2}. \quad (19)$$

To show the upper bound, we may assume that $Q_k \leq q_k$. Then using (18),

$$\frac{1 - Q_k^{1/k}}{1 - q_k^{1/k}} \leq \frac{\ln Q_k}{\ln q_k} \leq \frac{\ln\left(\left(1 - \frac{\varepsilon}{2}\right)q_k\right)}{\ln q_k} = 1 + \frac{\ln\left(1 - \frac{\varepsilon}{2}\right)}{\ln q_k} < 1 - \ln\left(1 - \frac{\varepsilon}{2}\right) \leq 1 + \frac{\varepsilon}{2}.$$

The lower bound in (19) follows similarly. Hence by Lemma 1,

$$\tau \geq 1 - q_k^{1/k} \geq (1 - \varepsilon)(1 - Q_k^{1/k}),$$

and

$$\begin{aligned} \tau &\leq 1 - \left(\frac{q_k}{n}\right)^{1/k} \leq \left(1 + \frac{\ln n}{\ln(1/q_k)}\right) (1 - q_k^{1/k}) \\ &\leq \left(1 + \frac{1}{c}\right) \left(1 + \frac{\varepsilon}{2}\right) (1 - Q_k^{1/k}) \leq (1 + \varepsilon)(1 - Q_k^{1/k}). \end{aligned}$$

□

7. CONCLUDING REMARKS

1. We can estimate for every node-transitive graph, by similar means, the value $1 - \max(\lambda_2, |\lambda_n|)$, which governs the mixing time of the chain. The trick is to consider the matrix M^2 instead of M , i.e., observe the chain only every other step. A little care is in order, since this new chain is not connected if G is bipartite. We have to start by observing if the graph is bipartite and if so return 0 as $\lambda_n = -1$. As mentioned in the Introduction, whether G is bipartite can be decided by checking if all return times are even. Clearly, considering only a polynomial number of return times introduces a negligible error probability.

2. The second moment of the first return time also has some more direct meaning. Let $H(\pi, r)$ denote the expected number of steps before a random walk starting from the stationary distribution π hits the root r . Then it is not hard to show using that the walk is close to stationary at a far away time that

$$H(\pi, r) = \frac{\mathbf{E}(T_1^2)}{2\mathbf{E}(T_1)} - \frac{1}{2}.$$

To see this consider a random walk from r with return times T_1, T_2 , etc. Let t be large enough and fixed and let s be distributed uniformly among the integers in $t < s \leq 2t$. The walk after s steps gets us very close to the stationary distribution. So $H(\pi, r)$ is close to the expected time it takes to return to r after the first s steps. The contribution to this expectation from the return T_i with $t \leq T_{i-1} < T_i \leq 2t$ is

exactly $\binom{T_i - T_{i-1}}{2}/t$. The contribution of the first returns after step t and after step $2t$ can be neglected (t is large). Here $T_i - T_{i-1}$ is distributed independently as T_1 so each return contributes approximately $\mathbf{E}(\binom{T_1}{2})/t$ to the expectation and we have an expected number of approximately $t/\mathbf{E}(T_1)$ such returns in the interval. This yields an estimate of $H(\pi, r) \approx \mathbf{E}(\binom{T_1}{2})/\mathbf{E}(T_1)$. The error of this approximation goes to 0 as t grows, so it has to be exact yielding the claimed formula.

It is not clear whether any of the higher moments have any direct combinatorial significance.

3. Here are a couple of related problems.

Problem 1: Let G be a connected graph of size n . We label the vertices randomly by $m(n)$ colors and observe the colors as they are visited by a simple random walk: after each step, the walker tells you “now I’m at red”, “now at blue”, and so on. How many colors are needed in order to recover the shape of G almost surely from this sequence of colors?

Problem 2: Consider an n -node connected graph. Take n particles labeled $1, \dots, n$. In a configuration, there is one particle at each node. The interchange process introduced in [1] is the following continuous time Markov chain on configurations: For each edge (i, j) at rate 1 the particles at i and j interchanged. Assume you observed the restriction of the interchange process to a fixed node, what graph properties can be recovered? Obviously you get more information than in the case discussed in the paper, which corresponds to noticing only one of the particles. But is it really possible to use this information to discover more about the graph?

REFERENCES

- [1] D. Aldous, My favorite 6 open problems in mathematical probability. <http://stat-www.berkeley.edu/users/aldous>
- [2] D. Aldous, J. Fill, Reversible Markov Chains and Random Walks on Graphs. <http://stat-www.berkeley.edu/users/aldous/RWG/book.html>
- [3] I. Benjamini and L. Lovász, Global Information from Local Observation. FOCS 2002: 701-710
- [4] U. Feige, A Fast Randomized LOGSPACE Algorithm for Graph Connectivity. Theoretical Computer Science 169, 147–160, 1996.
- [5] A. Lubotzky, Cayley graphs: eigenvalues, expanders and random walks. Surveys in combinatorics, 1995 (Stirling), 155–189, London Math. Soc. Lecture Note Ser., 218, Cambridge Univ. Press, Cambridge, 1995