This is a preprint of an article published in *J. Graph Theory* **29** (1998) 57-62.

For the definitive version of this article, visit http://www.interscience.wiley.com/.

Random walks and the regeneration time

Andrew Beveridge and László Lovász

December 1996

Abstract

Consider a graph G and a random walk on it. We want to stop the random walk at certain times (using an optimal stopping rule) to obtain independent samples from a given distribution ρ on the nodes. For an undirected graph, the expected time between consecutive samples is maximized by a distribution equally divided between two nodes, and so the worst expected time is 1/4 of the maximum commute time between two nodes. In the directed case, the expected time for this distribution is within a factor of 2 of the maximum.

1 Introduction

Consider a random walk on the graph $G = (V, E)$ starting at node i. We can use an optimal stopping rule (see section 2 for the exact definition) to halt a random walk starting at i so that the distribution of the final node is a given distribution ρ. Denote the expected length of this walk by $H(i, \rho)$.

The quantity $H(i, \pi)$ measures the mixing speed of the walk, starting at i. A number of other mixing measures have been introduced and studied (see Aldous [1], Lovász and Winkler [5], [6] and Aldous, Lovász and Winkler [2]). One of these is the reset time $T_{\text{reset}} = \sum_i \pi_i H(i, \pi)$ (here π is the stationary distribution). One of the results in [2] (formulated here for undirected graphs only) states that the reset time is within an absolute constant factor of other mixing measures, like the number of steps after which the distribution of the current node is closer than (say) one percent to the stationary distribution, starting from the worst point.

We consider the more general regeneration time $T_{\text{regen}}(\rho) = \sum_i \rho_i H(i, \rho)$. One can view $T_{\text{regen}}(\rho)$ as follows. We start a random walk from a random node drawn from distribution ρ, and want to stop it (by an optimal stopping rule) so that an independent sample from the same distribution ρ is obtained. We define $T_{\text{regen}} = \max_\rho T_{\text{regen}}(\rho)$, where the maximum is taken over all distributions on G.

It turns out that the regeneration time is in general much larger than the reset time, and in fact it is closely related to the *commute time*. The commute time between two nodes i and j is $H(i,j)+H(j,i)$, and the commute time T_{comm} of the graph is defined as the maximum commute time between any pair of nodes of G.

Our main result is the following.

Theorem 1 *For an undirected graph, the regeneration time is maximized by a distribution equally divided between two nodes, and hence*

$$T_{\text{regen}} = \frac{1}{4}T_{\text{comm}}.$$

Turning to the directed case, in general we can determine the regeneration time within a factor of 2.

Theorem 2 *For a directed graph G,*

$$\frac{1}{4}T_{\text{comm}} \le T_{\text{regen}} \le \frac{n-1}{2n}T_{\text{comm}},$$

and these bounds are tight.

Let us note that the *minimum* regeneration time of distributions is trivial: the regeneration time of a distribution concentrated on a single node is 0.

Finally, we have to point out that all of our results hold without any essential change for finite Markov chains (where undirected graphs correspond to time-reversible chains).

2 Preliminaries

The *support* of a distribution ρ on V, denoted S_ρ, is the set of nodes i such that $\rho_i > 0$.

Given a directed, connected graph $G=(V,E)$ with transition matrix $M=(p_{ij})$ a *random walk* on G is a sequence of nodes $(w_0, w_1, \ldots, w_t, \ldots)$ such that the probability that w_{t+1}=j is $p_{w_t j}$. For an undirected graph G, a random walk has transition probabilities $p_{ij} = 1/d(i)$ if $ij \in E$ and 0 otherwise where $d(i)$ is the degree of i. For irreducible, aperiodic G, as time t tends to infinity the distribution of the tth state tends to the so called stationary distribution π; for a random walk on an undirected graph G, $\pi_i = \frac{d(i)}{2|E|}$. The *hitting time* $H(i,j)$ is the expected length of a random walk from i to j. The cycle reversing identity of [3] states that for an undirected graph,

$$H(i,j)+H(j,k)+H(k,i) = H(i,k)+H(k,j)+H(j,i) \tag{1}$$

for any three nodes i, j and k. The analogous result holds for more than three nodes with an identical proof.

Let V^* be the space of finite walks on V, i.e. the set of finite strings $w = (w_0, w_1, w_2, \ldots, w_t)$, $w_i \in V$ and $w_i w_{i+1} \in E$. For a given initial distribution σ the probability of w being the walk after t steps is

$$\Pr(w) = \sigma_{w_0} \prod_{i=0}^{t-1} p_{w_i, w_{i+1}}.$$

A *stopping rule* Γ is a map from V^* to $[0, 1]$ such that $\Gamma(w)$ is the probability of continuing given that w is the walk so far observed. We assume that with probability 1 the rule stops the walk in a finite number of steps.

Given another distribution τ on V, we define the *access time* $H(\sigma, \tau)$ as the minimum expected length of a stopping rule Γ that produces τ when started at σ. We say Γ is *optimal* if it achieves this minimum. Optimal stopping rules exist for any pair σ, τ of distributions (see e.g. [6]). When σ and τ are concentrated on single states i and j respectively (we write $\sigma = i, \tau = j$), then the access time $H(i, j)$ is the (expected) hitting time from i to j. In this instance, the only optimal stopping rule is "walk until you hit j." Whenever our target is concentrated on a node, the analogous rule is optimal for any starting distribution: $H(\sigma, j) = \sum_i \sigma_i H(i, j)$. However, equality does not usually hold if our target is not a singleton.

Given a stopping rule Γ from σ to τ, for each $i \in V$ we define its *exit frequency* $x_i(\Gamma)$ to be the expected number of times the walk leaves state i before halting. Exit frequencies are fundamental to virtually all access time results. A key observation due to Pitman [7] is that exit frequencies satisfy

$$\sum_j p_{ji} x_j(\Gamma) - x_i(\Gamma) = \tau_i - \sigma_i. \tag{2}$$

It follows from this "conservation equation" that the exit frequencies for two rules from σ to τ differ by $K\pi_i$ where K is the difference between the expected lengths of these rules. Hence the distributions σ and τ uniquely determine the exit frequencies for an optimal stopping rule between them and we denote these optimal exit frequencies by $x_i(\sigma, \tau)$. Moreover, Lovász and Winkler [6] prove that a stopping rule Γ is optimal if and only if there exists a *halting state* k such that $x_k(\Gamma) = 0$.

Any three distributions ρ, σ and τ satisfy the "triangle inequality"

$$H(\rho, \tau) \le H(\rho, \sigma) + H(\sigma, \tau) \tag{3}$$

with equality holding if and only if there is a k that is a halting state from ρ to σ and simultaneously from σ to τ. In particular, $H(\sigma,j) \leq H(\sigma,\tau) + H(\tau,j)$ and equality holds if and only if j is a halting state for the σ,τ-walk. Hence

$$H(\sigma,\tau) = \max_j (H(\sigma,j) - H(\tau,j)).$$

3 Proofs

Consider a directed graph G and a distribution ρ concentrated on two nodes i, j. Then

$$\sum_k \rho_k H(k,\rho) = \rho_i \rho_j (H(i,j) + H(j,i)) = \rho_i (1-\rho_i)(H(i,j) + H(j,i)).$$

Clearly this quantity is maximized when ρ is equally divided between the two nodes which maximize the commute time, showing that $T_{\text{regen}} \geq \frac{1}{4} T_{\text{comm}}$. This proves the lower bound in theorem 2 (and of course also in theorem 1).

To complete the proof of theorem 1 we need some lemmas.

Lemma 1 *For an undirected graph G and a distribution ρ on G, assume there is a sequence of nodes $(i_1, i_2, \ldots, i_r)$ such that for all k, i_{k+1} is a halting state for the (i_k, ρ)-walk. Then i_{k-1} is also halting for the (i_k, ρ)-walk (here $i_0 = i_r$ and $i_{r+1} = i_1$).*

Proof. If i^* is a halting state for the (i,ρ)-walk, then $H(i,i^*) - H(\rho,i^*) \geq H(i,j) - H(\rho,j)$ for all j with equality holding if and only if j is also a ρ-halting state for i. Therefore

$$\sum_{k=1}^{r} (H(i_k, i_{k+1}) - H(\rho, i_{k+1})) \geq \sum_{k=1}^{r} (H(i_k, i_{k-1}) - H(\rho, i_{k-1})).$$

Adding $\sum_k H(\rho, i_k)$ to both sides yields $\sum_k H(i_k, i_{k+1}) \geq \sum_k H(i_k, i_{k-1})$. This is satisfied with equality by the cycle reversing identity (1) and the lemma follows. □

Corollary 1 *For every distribution ρ on an undirected graph G, there exist nodes i, j such that i is ρ-halting for j and vice versa. If ρ is not a singleton, such indices i and j exist in the support of ρ.*

Proof. If the distribution is concentrated on a single node $\rho = j$ then every state is halting from j and vice versa. For a non-singleton ρ, consider the sequence $i_1, i_2, \ldots, i_r$ where i_{k+1} is ρ-halting for i_k and $\rho_{i_{k+1}} > 0$. Since G is finite, eventually a node must repeat. We may assume $i_{r+1} = i_1$ is the first repeat of a node in this sequence. Since ρ is not a singleton, $r \geq 2$.

Hence $\{i_1, i_2, \ldots, i_r\}$ satisfies the conditions of lemma 1 so i_k and i_{k+1} are mutually ρ-halting for every k. □

Corollary 1 generalizes the observation of Lovász and Winkler [5] that for an undirected graph there are two nodes which are mutually π-halting (in fact, each of these nodes achieves $\max_i H(i, \pi)$).

Let ρ be any distribution on V. For $T = \{s_1, s_2, \ldots, s_r\} \subset S_\rho$, let σ^j denote the distribution achieved by walking from s_j until you hit a node in $S_\rho \backslash T$. Let ρ' denote the distribution with $\rho'_i = \rho_i + \sum_{j \in T} \rho_{s_j} \sigma_i^j$ for $i \in S_\rho \setminus T$ and zero elsewhere. An optimal stopping rule from ρ to ρ' is given by "walk until you hit a node in $S_\rho \setminus T$." We refer to this process as *dispersing* the set T.

Lemma 2 *Given a (directed or undirected) graph G and a distribution ρ, assume that $T = \{s_1, s_2, \ldots, s_r\}$ is a subset of S_ρ such that every node in $S_\rho \setminus T$ has a ρ-halting state in $S_\rho \setminus T$. If ρ' is the distribution achieved by dispersing T then $\sum_k \rho_k H(k, \rho) \leq \sum_k \rho'_k H(k, \rho')$.*

Proof. Using the notation above, we have $\rho'_i = \rho_i + \sum_{j \in T} \rho_{s_j} \sigma_i^j$ for $i \in S_\rho \setminus T$ and zero elsewhere. For $i \in S_\rho \setminus T$ let $i^* \in S_\rho \setminus T$ be a ρ-halting state. Then $x_{i^*}(i, \rho) + x_{i^*}(\rho, \rho') = 0$ so that $H(i, \rho') = H(i, \rho) + H(\rho, \rho')$. Therefore

$$\begin{aligned}
\sum_k \rho'_k H(k, \rho') &= \sum_k \rho'_k (H(k, \rho) + H(\rho, \rho')) = \sum_k \rho_k H(k, \rho) + \sum_k (\rho'_k - \rho_k) H(k, \rho) + H(\rho, \rho') \\
&= \sum_k \rho_k H(k, \rho) + \sum_j \rho_{s_j} \left(H(s, \sigma^j) + \sum_k \sigma_k^j H(k, \rho) - H(s_j, \rho) \right) \\
&\geq \sum_k \rho_k H(k, \rho) + \sum_j \rho_{s_j} (H(s_j, \sigma^j) + H(\sigma^j, \rho) - H(s_j, \rho)) \geq \sum_k \rho_k H(k, \rho)
\end{aligned}$$

where the final inequality follows from the triangle inequality (3). □

Proof of Theorem 1. It is enough to show that for an undirected graph G the regeneration time is achieved by a distribution supported by two nodes.

If the support set S_ρ contains more than two nodes, then by Corollary 1 there exits a pair of nodes $i, j \in S_\rho$ such that i is ρ-halting for j and vice versa. Let $T = S_\rho \setminus \{i, j\}$ and disperse T to get a distribution ρ' concentrated on i and j. By lemma 2, $\sum_k \rho_k H(i, \rho) \leq \sum_k \rho'_k H(k, \rho')$.
□

Proof of Theorem 2. We have proved the lower bound for any graph G already. (Theorem

1 shows that this lower bound is tight.) The upper bound follows from the naive rule to ρ:

$$\sum_i \rho_i H(i,\rho) \le \sum_{i,j} \rho_i\rho_j H(i,j) = \sum_{i<j} \rho_i\rho_j (H(i,j)+H(j,i)) \le \sum_{i<j} \rho_i\rho_j T_{\rm comm}.$$

The coefficient $\sum_{i<j} \rho_i\rho_j \le (n-1)/2n$ with equality if and only if $\rho_i = 1/n$ for all i.

The following example shows that the upper bound is best possible as well. Consider the directed cycle on n nodes. The stationary distribution π is the uniform distribution. For every pair of nodes i,j, $H(i,j)+H(j,i) = n$ and the naive walk from i to π is optimal for all i. Hence taking $\rho = \pi$ gives equality at each step in the proof of the upper bound, showing that the regeneration time of the directed cycle is $T_{\rm comm}(n-1)/2n = (n-1)/2$. □

If G is a directed cycle with loops, then an identical argument shows that $T_{\rm regen}$ of G achieves the upper bound of theorem 2. The following lemma shows that the converse holds as well: the family of directed cycles (with or without loops) is the unique set of graphs achieving $T_{\rm regen} = \frac{n-1}{2n} T_{\rm comm}$.

Lemma 3 *Let G be a graph and $\rho > 0$, a distribution on V. Assume that the naive rule is optimal from each node to ρ. Then G is a directed cycle with loops.*

Proof. Assume that there exists a node i with outdegree at least 2. By assumption the naive rule from i to ρ is optimal, hence it has a halting state z. Let P be a shortest path from i to z. Clearly i has an outneighbor j that is not on P. Now the naive rule selects j with positive probability, and the random walk traverses P with positive probability, and in this case it does not halt at z, a contradiction. □

References

[1] D.J. Aldous, Some inequalities for reversible Markov chains, *J. London Math. Soc.* **25** (1982), 564–576.

[2] D.J. Aldous, L. Lovász and P. Winkler, Inequalities for Markov Chain Mixing Times, to appear in: *Stochastic Processes and their Applications*

[3] D. Coppersmith, P. Tetali, and P. Winkler, Collisions among Random Walks on a Graph, *SIAM J. on Discrete Mathematics* **6** No. 3 (1993), 363-374.

[4] L. Lovász, Random Walks on Graphs: a Survey, *Combinatorics, Paul Erdös is Eighty*, Vol II (eds. D. Miklós, V.T. Sós and T. Szönyi) J. Bolyai Math. Soc. (1996), 353–397.

[5] L. Lovász and P. Winkler, Reversal of Markov Chains and the Forget Time, to appear in *Combinatorics, Probability and Computing*

[6] L. Lovász and P. Winkler, Efficient Stopping Rules for Markov Chains, *Proceedings of the 27th ACM Symposium on the Theory of Computing,* (1995) 76–82.

[7] J.W. Pitman, Occupation measures for Markov chains, *Adv. Appl. Prob.* **9** (1977), 69-86.