Reversal of Markov Chains and the Forget Time

László Lovász[1] and Peter Winkler[2]

(version of June 10, 1996)

Abstract

We study three quantities that can each be viewed as the time needed for a finite irreducible Markov chain to "forget" where it started. One of these is the *mixing time*, the minimum mean length of a stopping rule that yields the stationary distribution from the worst starting state. A second is the *forget time*, the minimum mean length of any stopping rule that yields the same distribution from any starting state. The third is the *reset time*, the minimum expected time between independent samples from the stationary distribution.

Our main results state that the mixing time of a chain is equal to the mixing time of the time-reversed chain, while the forget time of a chain is equal to the *reset* time of the reverse chain. In particular, the forget time and the reset time of a time-reversible chain are equal. Moreover, the mixing time lies between absolute constant multiples of the sum of the forget time and the reset time.

We also derive an explicit formula for the forget time, in terms of the "access times" introduced in [11]. This enables us to relate the forget and reset times to other mixing measures of the chain.

1 Introduction

Consider a finite, irreducible, discrete-time Markov chain. How long does it take to "forget" the state in which the chain was started? The following three quantities can each be interpreted as the answer to this question.

The *mixing time* is defined as the least number $T_{\rm mix}$ such that for every starting state, there is a stopping rule that stops the chain in expected time at most $T_{\rm mix}$, such that the distribution of the final state is the stationary distribution. This quantity plays an important role in the analysis of many randomized algorithms that use Markov chains to generate a random sample point from (approximately) the stationary distribution.

[1]Dept. of Computer Science, Yale University, New Haven, CT 06510

[2]Bell Laboratories 2C-379, Murray Hill NJ 07974

The *forget time* is the smallest number T_{forget} such that for some distribution μ (not necessarily stationary) and every starting state, there exists a stopping rule that stops the chain at a state drawn from μ in expected time at most T_{forget}.

The third quantity, which we call *reset time*, will turn out to be closely related to the forget time. If the chain is started at a state drawn from the stationary distribution, how long must the chain be run to get another stationary state, independent of the first? The minimum mean time for a stopping rule for this task is denoted by T_{reset}.

For our purposes, a "stopping rule" is a randomized algorithm for stopping a Markov chain, according to what has been observed so far; a more precise definition is given below in Section 3. A rule which stops at some particular distribution (usually the stationary distribution) and whose expected number of steps is minimal subject to that condition, is said to be *optimal.*

One of our main results is the following

Theorem 1 *If the Markov chain is time-reversible, then* $T_{\text{forget}} = T_{\text{reset}}$.

To generalize this identity to all Markov chains, we need the following notion. Given a Markov chain with transition probabilities p_{ij}, and stationary distribution π, consider the *reverse chain* with transition probabilities $\hat{p}_{ij} = \pi_j p_{ji}/\pi_i$. In what follows, the "hat" over a symbol means that it concerns this reverse chain. It is clear that this chain has the same stationary distribution. We prove the following connections between our mixing quantities defined for the forward and the reverse chain:

Theorem 2 *For every finite Markov chain,* $T_{\text{mix}} = \hat{T}_{\text{mix}}$.

Theorem 3 *For every finite Markov chain,* $T_{\text{forget}} = \hat{T}_{\text{reset}}$ *and* $T_{\text{reset}} = \hat{T}_{\text{forget}}$.

Of these, our proof of Theorem 3 is harder. It is possible that this theorem also has a direct probabilistic proof, but at this time we can only offer a derivation through a more complicated, explicit formula for the forget time (Theorem 4). We also show that the best target distribution μ in the definition of the forget time is uniquely determined, and find an explicit formula for it.

Sections 4 and 7 contain a number of further facts relating reverse chains that follow from our techniques. In particular, we show that *the worst starting states for mixing are also worst for forgetting*, and also relate the best starting states for mixing to the forget distribution.

The mixing time is in general different from the forget (or reset) times even for time-reversible chains, but there are quite tight inequalities relating them. It is clear that

$$T_{\text{forget}} \le T_{\text{mix}}, \qquad \text{and} \qquad T_{\text{reset}} \le T_{\text{mix}} .$$

The following inequality is noted in [5]:

$$T_{\text{mix}} \le 2(T_{\text{reset}} + T_{\text{forget}}) . \tag{1}$$

Thus it follows that $\max\{T_{\text{forget}}, T_{\text{reset}}\}$ is between T_{mix} and $T_{\text{mix}}/4$. We are going to show by an example that $\min\{T_{\text{forget}}, T_{\text{reset}}\}/T_{\text{mix}}$ can be arbitrarily close to 0.

These facts relate our results to studies important for the analysis of certain randomized algorithms for sampling. Aldous [3], [1] proved the surprising and powerful result that for time-reversible chains, essentially all reasonable "mixing parameters" are within absolute constant factors of each other. These parameters include, besides T_{mix}, the times T_{tv} and T_{fill}. T_{tv} is defined as the least t such that for every starting state, there is a stopping rule that stops the chain in expected time at most t, such that the distribution of the final state is closer than (say) 1/4 to the stationary distribution in total variation distance; T_{fill} is defined analogously, except that the requirement for the final distribution is that the probability of each state i be at least 3/4 of its stationary probability.

The case of non-reversible chains is only a little more complicated. It is shown in [5] that most "mixing parameters" of a Markov chain fall into one of two groups, and parameters in the same group differ only by a constant factor. One group contains the time T_{tv} (along with other mixing parameters). The other group contains the *mixing time* T_{mix}, and also the time T_{fill}. It is also shown in [5] that $T_{\text{mix}} \le (\log(1/\pi_0))T_{\text{tv}}$, where $\pi_0 = \min_i \pi_i$ (the factor $\log 1/(\pi_0)$ is best possible up to a constant).

The forget time, which may also be considered as a mixing parameter, fits this pattern. In the time-reversible case, $T_{\text{forget}} = T_{\text{reset}}$ is between T_{mix} and $T_{\text{mix}}/4$, by inequality (1). It is proved in [5] that in the general case, T_{forget} is within a constant factor of T_{tv}. On the other hand, the reset time, which seems to be conceptually closer to the mixing time, belongs to a new group. In fact, it follows that the "mixing time group" is self-reverse, but the "forget time group" is not.

The organization of this paper is the following: in Section 2, we use several examples to illustrate mixing, forgetting, and chain reversal; we shall return to one of the examples in Section 8, to show some implications of our results. In Section 3, we recall from [11, 12] a calculus of "access times between distributions" and properties of optimal stopping rules. Section 4 contains some elementary facts about reversed Markov chains. Section 5 collects some known and (perhaps) not-quite-known relations between hitting times and the transition matrix. Section 6 contains the main formula for the forget time, and a linear programming formulation of this quantity. Section 7 contains several other relations between mixing parameters, and Section 8 returns to one of the examples.

2 Examples

The following examples should help to clarify the mixing parameters described above. We will return later to Example 4, the "winning streak," to show how our results verify or disprove some natural guesses concerning optimal mixing and forgetting.

Example 1 Consider a path on n nodes $0, 1, \ldots, n-1$, and assume, for the sake of simplicity, that n is odd. Consider a simple random walk on the nodes. The stationary distribution is proportional to the degrees, i.e., $\pi_i = 1/(2(n-1))$ for $i = 0$ and $i = n-1$, and $\pi_i = 1/(n-1)$ otherwise. To generate a stationary state, we may use a very simple brute-force approach: pick the state in advance and walk until it is hit. The worst starting states are 0 and $n-1$, and the average number of steps it takes to reach node i from 0 is i^2. Hence this naive stopping rule takes

$$\frac{1}{2(n-1)} \cdot 0 + \frac{1}{n-1} \cdot 1 + \frac{1}{n-1} \cdot 4 + \cdots$$
$$\cdots + \frac{1}{n-1} \cdot (n-2)^2 + \frac{1}{2(n-1)} \cdot (n-1)^2 = \frac{2n^2 - 4n + 3}{6}$$

steps.

It turns out that this rule is optimal (subject to generating the stationary distribution) when the walk is begun at 0 or $n-1$. Optimal rules for mixing when starting from another node of the path are more complicated and not described here, but if we apply one of them and average over all starting nodes according to the stationary distribution (which means that given a stationary node, we want to generate another stationary node independent from the first) we get a mean stopping time of $(n-1)^2/4$.

If we want to forget the starting node, it turns out that the best choice is to walk until we hit the midpoint of the path. This takes on the average $(n-1)^2/4$ steps, starting from an endpoint. The results in this paper easily imply that this simple "forgetting rule" is optimal.

Example 2 Let us add the edge between 0 and $n-1$ to make a cycle; now the stationary distribution is uniform. Starting from 0, we can describe a stopping rule as follows: select a random node i uniformly, and walk until either i or $n-i$ is reached. This rule (which is optimal, as is easily shown using the results in Section 3) takes an expected number of $(n^2+2)/12$ steps if n is even and $(n^2-1)/12$ steps if n is odd.

Trivially this value is the same for all starting nodes, and hence $T_{\text{reset}} = T_{\text{mix}}$. The chain is time-reversible, so it follows that also $T_{\text{forget}} = T_{\text{mix}}$. (But in fact it follows from the uniqueness statement in Theorem 4 that for a chain with a state-transitive automorphism group, the best way to forget is to go to the uniform distribution, so $T_{\text{forget}} = T_{\text{mix}}$.)

Example 3 Next, we add a new node n and connect it to all nodes on the cycle, to get a wheel. Assume that n is odd and replace each edge on the rim by $(n-1)/2$ parallel edges, so that we get a regular graph, and the stationary distribution is again uniform. The interesting point in this example is that if we start from the center, in less than one move we can reach the stationary distribution (make one move with probability $\frac{n}{n+1}$). Starting from any other node, we can walk until we hit the center and then (with probability $\frac{n}{n+1}$) make one more move, taking, on the average, $n + \frac{n}{n+1}$ moves. It turns out that this rule is almost optimal, hence the reset time is only slightly less than the mixing time.

Example 4 Let us consider the "winning streak" chain, where states are $0, \ldots, n-1$ and all transition probabilities are zero except $p_{i,0} = 1/2$ for all $0 \le i \le n-1$, $p_{i,i+1} = 1/2$ for $0 \le i < n-1$, and $p_{n-1,n-1} = 1/2$. Here the stationary distribution is given by $\pi_i = 2^{-i-1}$ $(0 \le i < n-1)$ and $\pi_{n-1} = 2^{-(n-1)}$.

At every state, we have probability 1/2 of "falling back" to 0. If we want to forget where we started, then a natural idea is to wait until we get to 0; at worst this takes an average of 2 moves. A computation based on Theorem 4 (see Section 8) will show that this strategy is not quite optimal.

To get a stationary state takes much more time. If we start at any state, we have the following strategy: walk $n-1$ steps and stop. It is not difficult to see that this rule leads to the stationary distribution from any starting point (we end up at i if and only if the number of tailing "forward" moves is i). This stopping rule will turn out optimal if we start at 0 or $n-1$. Starting from other nodes, this strategy will not be optimal any more, but will be close to it.

Now look at the reverse of the winning streak chain. Here the transition probabilities are $p_{i,i-1} = 1$ for $1 \le i < n-1$, $p_{n-1,n-2} = p_{n-1,n-1} = 1/2$, $p_{0,i} = 2^{-i-1}$ for $0 \le i \le n-2$ and $p_{0,n-1} = 2^{-(n-1)}$ (and 0 otherwise). A step from state 0 goes directly to the stationary distribution. From any other state i, we can walk to 0 and then make one more move, which takes $i+1$ moves if $i < n-1$, and $n+1$ steps (on the average) if we start in the last state. This strategy will turn out optimal except if we start at the last state. Starting from the last state, the following strategy is better: walk for $n-1$ steps. So the mixing time is at most $n-1$; it will turn out to be exactly $n-1$.

If we want to generate a stationary state independent from the (stationary) starting state, we can follow the rule described above; this takes, on the average,

$$\frac{1}{2} \cdot 1 + \frac{1}{4} \cdot 2 + \frac{1}{8} \cdot 3 + \cdots + \frac{1}{2^{n-1}}(n-1) + \frac{1}{2^{n-1}}(n-1) = 2 - \frac{1}{2^{n-2}}$$

steps. So this is an upper bound on the reset time; again, it will turn out to be the exact value. To forget, we can again walk to node 0; in this chain, this will take $n-1$ steps from the worst starting node. It will turn out that this rule is nearly optimal.

Example 5 A classic example of mixing occurs in shuffling a deck of playing cards. Here the states of the chain are the 52! permutations of the deck; in the Gilbert–Shannon–Reeds model ([10], [14]) of the riffle (or "dovetail") shuffle the *reverse* of a shuffle is described elegantly as follows. Each card is randomly, uniformly, and independently assigned a bit from $\{0, 1\}$. All cards assigned the same bit are kept in the same order, but all cards with "0" are placed above all cards with "1".

In a series of t such "unshuffles" each card receives a string of t bits which we interpret as a binary number; the final order of the deck is then determined by the natural order of these numbers, with all ties broken in favor of the initial order. If there are no ties then the final order is perfectly uniformly random, thus the following stopping rule yield perfect mixing: unshuffle repeatedly until all 52 binary numbers are distinct.

In fact, this rule is optimal for mixing in the reverse riffle-shuffle chain (because the inverse of the starting permutation is a "halting state" (see Section 3; it cannot be achieved unless the binary labels are, indeed, all distinct). A calculation based on the "birthday problem" shows that the expected time for this rule is about $11\frac{3}{4}$ unshuffles, so by Theorem 2 we have for the original riffle shuffle chain that:

$$T_{\mathrm{mix}} = \hat{T}_{\mathrm{mix}} \approx 11\frac{3}{4} \ .$$

It is interesting to compare this figure with the famous 7 riffle shuffles recommended by Bayer and Diaconis [6]. Their need for a fixed-time stopping rule seems to be much more than compensated by sufficiency of *approximate* mixing. In fact Doyle [8] has devised a game which he calls "New Age Solitaire" which has a winning probability of 50% for a random deck but about 80% for a 7-times-riffle-shuffled fresh deck of Bicycle brand playing cards.

Fortunately Doyle's game is not one which is likely to replace standard forms of solitaire. It is clear that if the public really could be persuaded to shuffle as many as 7 times, then perceived differences (well known to tournament bridge players) between hand-shuffled and computer-shuffled decks would disappear.

3 Access times between distributions

Let $w^0, w^1, \ldots, w^k, \ldots$ be a finite irreducible Markov chain with state space V and transition probabilities p_{ij}. We denote by π its stationary distribution. Let σ^k denote the distribution of w^k; we often denote the starting distribution σ^0 by σ.

A *stopping rule* is a map Γ from V^* (the set of finite strings of states) to $[0,1]$: for $w = w^0 w^1 \ldots w^t$, the value of $\Gamma(w)$ is the probability of continuing given that w is the walk so far observed. (Each such stop-or-go decision is made independently.) It is useful sometimes to regard a stopping rule Γ as a stopping time, i.e., a random variable whose value is the actual number of steps made before stopping (so we stop at w^Γ). We assume that with probability 1, the rule stops the chain, i.e., the stopping time is finite. We denote the distribution of the final state by σ^Γ. A stopping rule Γ for which $\sigma^\Gamma = \tau$ is also called a stopping rule from σ to τ.

Let σ and τ be two distributions on the set V of states. We define the *access time* $H(\sigma,\tau)$ from σ to τ as the infimum of the mean lengths of all stopping rules that produce the distribution τ, when the chain starts from σ. A stopping rule is *optimal* if it achieves this infimum. In the case when σ and τ are concentrated on a single states s and t respectively (we abuse notation by writing $\sigma = s$ and $\tau = t$) we have that $H(s,t)$ is the *hitting time* from s to t, i.e., the expected number of steps of the Markov chain before state t is reached, when started from state s.

It turns out that the access times are always well-defined (optimal stopping rules exist for any two distributions) and possess many useful algebraic properties. In this paper, we shall need the following facts. First, there is an explicit formula for the access times. If the target distribution is a single state t, then of course there is only one optimal stopping rule (walk until state t is hit), and hence

$$H(\sigma, t) = \sum_i \sigma_i H(i, t) \ . \tag{2}$$

(No similar linearity property holds for the target distribution.) For two arbitrary distributions σ and τ, we have the following formula:

$$H(\sigma, \tau) = \max_j \Big(H(\sigma, j) - H(\tau, j) \Big) \ . \tag{3}$$

Given any stopping rule Γ that produces τ from σ, we define the *exit frequency* x_i $(i \in V)$ for Γ as the expected number of times the walk leaves state i before stopping. Exit frequencies are related to the starting and ending distributions by a simple formula, found in Pitman [13]:

$$\sum_i p_{i,j} x_i - x_j = \tau_j - \sigma_j \ . \tag{4}$$

This equation implies that the exit frequencies are almost determined by σ and τ in the sense that the exit frequencies of two rules from the same starting distribution σ to the the same target distribution τ differ by $K\pi_i$, where K is the difference between the expected lengths of these stopping rules.

In particular, the exit frequencies of an optimal stopping rule from σ to τ are uniquely determined by σ and τ only. They are denoted by $x_i(\sigma, \tau)$.

It is sometimes more convenient to consider the scaled exit frequencies

$$y_i(\sigma, \tau) = \frac{1}{\pi_i} x_i(\sigma, \tau) \ .$$

The following formula (see [11, 12]) expresses exit frequencies in terms of the access times:

$$y_k(\sigma, \tau) = H(\tau, k) - H(\sigma, k) + H(\sigma, \tau) \ . \tag{5}$$

It follows easily from formulas (3) and (5) that for every σ and τ there exists a state k with $x_k(\sigma, \tau) = 0$. Such a state is called a *halting state*, since it has the property that every optimal stopping rule halts if it ever reaches state k. In fact, a stopping rule is optimal if and only if it has a halting state ([12]). The halting states achieve the maximum on the right hand side of (3).

In terms of access times, the mixing time of the chain can be defined by

$$T_{\mathrm{mix}} = \max_{\sigma} H(\sigma, \pi) = \max_{s} H(s, \pi)$$

where s varies over the state space V. A state s maximizing $H(s, \pi)$ is called *mixing-pessimal* or simply *pessimal.*

While this quantity seems to be much less relevant for the applications, we can also say something about

$$T_{\mathrm{bestmix}} = \min_{s} H(s, \pi) \ .$$

A state s minimizing $H(s, \pi)$ is called *mixing-optimal.*

We can also express the forget time by

$$T_{\mathrm{forget}} = \min_{\tau} \max_{\sigma} H(\sigma, \tau) = \min_{\tau} \max_{s} H(s, \tau) \ .$$

We call the distribution τ minimizing the right hand side the *forget distribution* and denote it by μ. (It will turn out that the forget distribution is uniquely determined.) The starting state achieving the maximum for $\tau = \mu$ will be called *forget-pessimal.*

For the reset time, we have

$$T_{\mathrm{reset}} = \sum_{j} \pi_j H(j, \pi) \ ,$$

showing that it can also be viewed as an "average mixing time."

We denote by T_{hit} the average hitting time between two states drawn from the stationary distribution, i.e.,

$$T_{\mathrm{hit}} = \sum_{i} \sum_{j} \pi_i \pi_j H(i, j) \ .$$

Using the notion of access times, we can write this definition as

$$T_{\mathrm{hit}} = \sum_{j} \pi_j H(\pi, j) \ .$$

(This is only formally similar to the definition of T_{reset}; in general, T_{hit} is much larger.)

The very useful "Random Target Lemma" (sometimes called the "right averaging principle" asserts that T_{hit} does not depend on the distribution of the starting point: for every state i,

$$\sum_{j} \pi_j H(i, j) = T_{\mathrm{hit}} \ . \tag{6}$$

4 Reverse chains

Given a Markov chain with transition probabilities p_{ij}, we define the *reverse chain* as the Markov chain on the same set of states, with transition probabilities $\hat{p}_{ij} = \pi_j p_{ji}/\pi_i$. Recall that the "hat" over a symbol means that it concerns this reverse chain.

We start with the trivial but important remark that the reverse chain has the same stationary distribution as the original. This follows by straightforward calculation.

The key to the proof of many properties of the reverse chain is the following general "duality formula."

Lemma 1 *Let σ, τ, ϕ, ρ be four distributions on the states. Then*

$$\sum_i (\tau_i - \sigma_i)\hat{y}_i(\phi, \rho) = \sum_i (\rho_i - \phi_i) y_i(\sigma, \tau) .$$

Proof. Let $v_0, v_1, \ldots$ be a walk in the forward chain started from σ, and consider the random variables

$$\mathbf{Y}_t = \hat{y}_{v_{t+1}}(\phi, \rho) - \hat{y}_{v_t}(\phi, \rho) + (\phi_{v_t} - \rho_{v_t})/\pi_{v_t} .$$

¿From the conservation equation (4) for exit frequencies applied to the reverse chain, we have

$$\hat{y}_i(\phi, \rho) - (\phi_i - \rho_i)/\pi_i = \frac{1}{\pi_i} \sum_j \hat{p}_{ji} \hat{x}_j(\phi, \rho)$$

$$= \sum_j \frac{p_{ij}}{\pi_j} \hat{x}_j(\phi, \rho) = \sum_j p_{ij} \hat{y}_j(\phi, \rho) ,$$

and therefore

$$\mathrm{E}(\mathbf{Y}_t | v_0 = i_0, \ldots, v_{t-1} = i_{t-1}, v_t = i) = \sum_j p_{ij} \hat{y}_j(\phi, \rho) - \sum_j p_{ij} \hat{y}_j(\phi, \rho) = 0$$

for each i, thus $\mathrm{E}(\mathbf{Y}_t) = 0$ *a priori*. In particular, if $\mathbf{T}$ is the number of steps taken by some optimal stopping rule from σ to τ, then the sum $\sum_{t=0}^{\mathbf{T}-1} \mathbf{Y}_t$ has zero expectation; but

$$\sum_{t=0}^{\mathbf{T}-1} \mathbf{Y}_t = \hat{y}_{v_{\mathbf{T}}}(\phi, \rho) - \hat{y}_{v_0}(\phi, \rho) + \sum_i \left(\frac{\phi_i - \rho_i}{\pi_i} \right) \mathbf{X}_i$$

where $\mathbf{X}_i$ is the number of times state i is exited. Taking expectation, the lemma follows. □

We derive some corollaries of this lemma.

Corollary 1 *For every state j, we have*

$$H(\pi, j) = \hat{H}(\pi, j) \ .$$

Hence also $\hat{T}_{\rm hit} = T_{\rm hit}$.

Proof. Choose $\sigma = \phi = \pi$ and $\tau = \rho = j$ in Lemma 1. $\square$

Now, our Theorem 2 above:

Corollary 2 *For every Markov chain, we have*

$$\hat{T}_{\rm mix} = T_{\rm mix} \ .$$

Moreover, if s is a pessimal state for the first chain and z is a halting state from s to π, then z is a pessimal state for the reverse chain and s is a halting state from z to π in the dual chain.

Remark. For time-reversible chains, the first assertion does not say anything, but the second translates into an interesting fact: *every time-reversible chain has at least two mixing-pessimal starting states.*

Proof. Let s be a pessimal starting state for the primal chain and let z be a halting state from s to π. Apply the lemma with $\sigma = s$, $\phi = z$, and $\tau = \rho = \pi$. Then we get

$$\hat{H}(z, \pi) - \hat{y}_s(z, \pi) = H(s, \pi) - y_z(s, \pi) = H(s, \pi) \ .$$

Hence

$$\hat{T}_{\rm mix} \geq \hat{H}(z, \pi) \geq H(s, \pi) = T_{\rm mix} \ .$$

The reverse inequality follows similarly. $\square$

Corollary 3 *For every two states i and j, we have*

$$\hat{H}(i, j) = H(j, i) - H(\pi, i) + H(\pi, j) \ .$$

(A consequence worth mentioning—although not used in this paper—is the following: if the Markov chain has a state-transitive automorphism group, then for all i and j, $\hat{H}(i, j) = H(j, i)$.)

Proof. Apply the lemma with $\sigma = j$, $\tau = \pi$, $\phi = i$ and $\rho = j$. Then we get

$$\hat{H}(i,j) = y_j(j,\pi) - y_i(j,\pi) \ . \tag{7}$$

Now (5) implies the assertion. □

Corollary 4 *Let s be any state. A state z is a halting state from s to π if and only if $\hat{H}(z,s) \geq \hat{H}(u,s)$ for every state u. In this case,*

$$\hat{H}(z,s) = H(s,\pi) + H(\pi,s) \ .$$

Proof. By (7),

$$\hat{H}(u,s) = y_s(s,\pi) - y_u(s,\pi) \ .$$

It follows that the left hand side is maximized when u is a halting state z from s to π. Note that then

$$\hat{H}(z,s) = y_s(s,\pi) \ . \tag{8}$$

Now apply the lemma with $\sigma = \rho = s$ and $\tau = \phi = \pi$. Then we get

$$\hat{H}(\pi,s) = y_s(s,\pi) - H(s,\pi) \ .$$

Combining with Corollary 1 and (8), we get the second statement of the corollary.

Remark. Several results about reversible Markov chains can be extended by using the notion of reverse chain. For example, the "cycle reversing" identity of [7] can be generalized (with a virtually identical proof) as

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

for any three states i, j and k. (From here we could get another way of deriving, among others, corollary 3 above.)

5 Matrix formulas for hitting and mixing times

We describe some useful formulas connecting the transition matrix M with the matrix H whose entry in position (i,j) is the hitting time $H(i,j)$. So in particular the diagonal of H is 0. We denote by R the diagonal matrix with the expected return time $1/\pi_i$ in the i-th position of the diagonal.

It is easy to see that these matrices satisfy the equation

$$(I - M)H = J - R \ . \tag{9}$$

where J is the $n \times n$ matrix of all 1's. Unfortunately, the matrix $I - M$ is singular, and so (9) does not uniquely determine H. But the only left eigenvector of $I - M$ with eigenvalue 0 is $\pi^\top$, and the only right eigenvector with eigenvalue 0 is $\mathbf{1}$, hence $I - M + \mathbf{1}\pi^\top$ is nonsingular. Moreover, we have

$$(I - M + \mathbf{1}\pi^\top)(I - \mathbf{1}\pi^\top) = I - M$$

(using that $M\mathbf{1} = \mathbf{1}$ and $\pi^\top\mathbf{1} = \mathbf{1}$). Hence

$$(I - M + \mathbf{1}\pi^\top)(I - \mathbf{1}\pi^\top)H = (I - M)H = J - R$$

and thus

$$H = (I - M + \mathbf{1}\pi^\top)^{-1}(J - R) + \mathbf{1}\pi^\top H = G + \mathbf{1}\pi^\top H \ . \tag{10}$$

For convenience, let us denote the matrix $(I - M + \mathbf{1}\pi^\top)^{-1}(J - R) = J - (I - M + \mathbf{1}\pi^\top)^{-1}R$ by G; this matrix turns out to carry a lot of information about combinatorial properties of the random walk. It is not difficult to see that for time-reversible walks, G is a symmetric matrix.

Equation (10) is still not the desired formula for H, since H also occurs on the right hand side. But we can determine $\pi^\top H$ by looking at the diagonal: $0 = G_{ii} + (\pi^\top H)_i$ and hence

$$(\pi^\top H)_i = -G_{ii} \ .$$

So the negative of the i-th diagonal entry of G gives the expected time $H(\pi, i)$ needed to hit state i, starting from the uniform distribution. But then we can express H purely from the matrix G. In fact, (10) implies that

$$H_{ij} = G_{ij} - G_{jj} \ .$$

Now consider the mixing time. Using (3), we get

$$H(s, \pi) = \max_t(H(s,t) - H(\pi,t)) = \max_t(G_{st} - G_{tt} + G_{tt}) = \max_t G_{s,t} \ .$$

Hence the mixing time is the largest entry of G.

All these formulas also relate nicely to the reverse chain. Clearly this has transition matrix $\hat{M} = RM^\top R^{-1}$, and simple substitution gives $\hat{G} = G^\top$. We could use this fact to derive the results of the previous section in an algebraic way.

6 Forget time and reset time

Now we turn to the forget time. First, we describe a reformulation of $T_{\rm forget}$ as the optimum value of a linear program, and derive a few consequences. Let us rewrite the definition, using equation (3):

$$T_{\rm forget} = \min_\tau \max_s H(s,\tau) = \min_\tau \max_s \max_j (H(s,j) - H(\tau,j)) = \min_\tau \max_j (H(j',j) - H(\tau,j)) ,$$

where j' is an "antipode" of j, i.e., a state for which $H(j',j)$ is maximal among all values $H(k,j)$. We can write this expression as a linear program:

$$T_{\rm forget} = \begin{array}{lll} \text{minimize} & t & \\ \text{subject to} & \tau_i \geq 0 & (i \in V), \\ & \sum_i \tau_i = 1, & \\ & t + \sum_i \tau_i H(i,j) \geq H(j',j) & (j \in V). \end{array} \tag{11}$$

Let us formulate the dual program; we have a variable r for the equation and a variable ρ_j for each $j \in V$, and then the dual program is:

$$T_{\rm forget} = \begin{array}{lll} \text{maximize} & r + \sum_j \rho_j H(j',j) & \\ \text{subject to} & \rho_j \geq 0 & (j \in V), \\ & \sum_j \rho_j = 1 & \\ & r + \sum_j \rho_j H(i,j) \leq 0 & (i \in V). \end{array} \tag{12}$$

So ρ is a probability distribution; moreover, the best value of r is clearly

$$r = \min_i \left(-\sum_j \rho_j H(i,j) \right) , \tag{13}$$

and so

$$T_{\rm forget} = \max_\rho \min_i \sum_j \rho_j (H(j',j) - H(i,j)) . \tag{14}$$

Any probability distribution ρ gives a lower bound on $T_{\rm forget}$. An attractive choice is $\rho = \pi$, in which case we can use the Random Target Lemma (6) to see that $r = -T_{\rm hit}$ and the minimum in (13) is attained for all i. Hence

$$T_{\rm forget} \geq \sum_j \pi_j (H(j',j) - H(\pi,j)) . \tag{15}$$

The main result of this paper is that equality holds in (15).

Theorem 4 *For every Markov chain,*

$$T_{\rm forget} = \sum_j \pi_j (H(j',j) - H(\pi,j)) .$$

Moreover, the forget distribution μ is uniquely determined and is given by the formula

$$\mu_j = \pi_j(1 - H(j', j) + H(\pi, j)) + \sum_m p_{mj}\pi_m(H(m', m) - H(\pi, m)) \ . \tag{16}$$

Remark: Using the reverse chain and Corollaries 1 and 4, we can write these formulas in the following neater forms:

$$T_{\text{forget}} = \sum_j \pi_j \hat{H}(j, \pi) = \hat{T}_{\text{reset}} \ ,$$

and

$$\mu_j = \pi_j \left(1 - \hat{H}(j, \pi) + \sum_m \hat{p}_{jm}\hat{H}(m, \pi)\right) \ . \tag{17}$$

In particular, Theorem 4 will imply Theorem 3 and hence also Theorem 1.

Proof. It suffices to show that $\rho = \pi$ and $r = -T_{\text{hit}}$ form an optimal solution of (12). For this, it suffices to exhibit a solution (μ, t) of the primal such that the complementary slackness conditions hold:

$$\mu_i > 0 \quad \Rightarrow \quad -T_{\text{hit}} + \sum_j \pi_j H(i, j) = 0 \ , \tag{18}$$

and

$$\pi_j > 0 \quad \Rightarrow \quad t + \sum_i \mu_i H(i, j) = H(j', j) \ . \tag{19}$$

We choose $t = \sum_j \pi_j(H(j', j) - H(\pi, j))$ (recall that we want equality in (15)). To choose the right μ, observe the first set of conditions is fulfilled by the Random Target Lemma (6), independently of the choice of μ. Since $\pi > 0$, the second set applies for every j. This gives $n = |V|$ linear equations on μ (and we have one more from the original system, namely $\sum_i \mu_i = 1$)). We show that this (seemingly overdetermined) system has a unique solution, which is non-negative, and hence is a solution of the linear program (11). This will prove the theorem.

Let $h_j = H(j', j)$. Then (19) can be re-written as

$$H^{\mathsf{T}}\mu = h - t\mathbf{1} \ . \tag{20}$$

and we also must have

$$\mathbf{1}^{\mathsf{T}}\mu = 1 \ . \tag{21}$$

We know from (10) that

$$H = J - (I - M + \mathbf{1}\pi^{\mathsf{T}})^{-1}R + \mathbf{1}\pi^{\mathsf{T}}H \ ,$$

and hence

$$H^\top \mu = J\mu - R(I - M^\top + \pi \mathbf{1}^\top)^{-1}\mu + H^\top \pi \mathbf{1}^\top \mu = \mathbf{1} - R(I - M^\top + \pi \mathbf{1}^\top)^{-1}\mu + H^\top \pi \ .$$

Equating the two expressions for $H^\top \mu$ and rearranging, we can express μ:

$$\mu = (I - M^\top + \pi \mathbf{1}^\top) R^{-1}((1+t)\mathbf{1} + H^\top \pi - h) = \pi + (I - M^\top) R^{-1}(H^\top \pi - h) \ . \qquad (22)$$

Substitution shows that this μ satisfies (20) and (21). The fact that it is non-negative follows by noticing that (22) is equivalent to equation (16), which in turn is equivalent to (17), where the non-negativity of the right hand side is trivial:

$$\hat{H}(j, \pi) \leq 1 + \sum_m \hat{p}_{jm} \hat{H}(m, \pi) \ ,$$

since making one step from j and then following an optimal rule to π is a (not necessarily optimal) stopping rule from j to π.

Thus we have a feasible solution of (11), which satisfies, together with (r, π), the complementary slackness conditions. Our argument also proves the uniqueness of the optimizing distribution, as well as the optimality of π as a dual solution. □

7 Optimal and pessimal starting states

We continue with some further links between the forget time and mixing-optimal and mixing-pessimal states.

Lemma 2 *The scaled exit frequencies from the stationary distribution π to the forget distribution μ, and vice versa, can be expressed by the formulas*

$$y_i(\pi, \mu) = \hat{H}(i, \pi) - \hat{T}_{\mathrm{bestmix}} \ ,$$

and

$$y_i(\mu, \pi) = T_{\mathrm{mix}} - \hat{H}(i, \pi) \ .$$

Proof. By equation (17), we have

$$\mu_i = \pi_i \left(1 + \sum_j \hat{p}_{ij} \hat{H}(j, \pi) - \hat{H}(i, \pi) \right)$$

and hence

$$\sum_j p_{ji}\pi_j \hat{H}(j,\pi) - \pi_i \hat{H}(i,\pi) = \mu_i - \pi_i \ . \tag{23}$$

This equation means that the scaled exit frequencies of any stopping rule from π to μ differ from the numbers $y_i = \hat{H}(i,\pi)$ by an additive constant, and so the numbers $\hat{H}(i,\pi) - \hat{T}_{\rm bestmix}$ are the scaled exit frequencies of an *optimal* stopping rule.

Equation (23) can also be written in the form

$$\sum_j p_{ji}\pi_j (T_{\rm mix} - \hat{H}(j,\pi)) - \pi_i (T_{\rm mix} - \hat{H}(i,\pi)) = \pi_i - \mu_i \ .$$

Hence the second equation in the lemma follows. □

The previous lemma has the following corollaries.

Corollary 5 *A state z is a halting state from π to the "forget distribution" μ if and only if it is a mixing-optimal state for the reverse chain. Moreover, for the reverse mixing time from z we have*

$$\hat{T}_{\rm bestmix} = T_{\rm forget} - H(\pi,\mu) \ .$$

Corollary 6 *A state z is a halting state from μ to π if and only if it is a mixing-pessimal state for the reverse chain. Moreover, for the mixing time we have*

$$T_{\rm mix} = T_{\rm forget} + H(\mu,\pi) \ .$$

So from a mixing-pessimal state, a best stopping rule to get to π is to follow an optimal stopping rule to the forget distribution μ, and then follow a mixing-optimal rule from there. It also follows that

Corollary 7 *Every mixing-pessimal state is forget-pessimal.*

8 The "winning streak" revisited

We conclude with a discussion of one of our introductory examples, the "winning streak" chain (Example 4). The stopping rules from each fixed state to π exhibited for the reverse chain are optimal, since each has a halting state. Hence

$$\hat{T}_{\rm mix} = T_{\rm mix} = n-1 \ .$$

To compute the forget time for the forward chain, we use Theorem 4. It can easily be checked that $H(0,i) = 2^{i+1} - 2$ for all i and $H(j,i) = 2^{i+1}$ for all $i < j \le n-1$. It is clear that a state i' which maximizes $H(i',i)$ is $i+1 \pmod n$. Using the Random Target Lemma we have that for any state i,

$$T_{\text{hit}} = \sum_{j=0}^{n-1} \pi_j H(i,j) = \sum_{j=0}^{n-1} \pi_j H(0,j) = n-1$$

and

$$\sum_{j=0}^{n-1} \pi_j H(j',j) = n + 1 - 2^{-(n-2)} \ ,$$

so that by Theorem 4,

$$T_{\text{forget}} = 2 - 2^{-(n-2)} \ .$$

This is indeed the value for $\hat{T}_{\text{reset}}$ calculated in the introduction. The forget distribution μ is given by

$$\mu_i = \begin{cases} 1 - 2^{-(n-1)}, & \text{if } i = 0, \\ 2^{-(n-1)}, & \text{if } i = n-1, \\ 0, & \text{otherwise.} \end{cases}$$

In fact, $H(i,\mu) = 2 - 2^{-(n-2)}$ for all i (walk until state 0 is hit or $n-1$ steps have been made).

A more tedious calculation shows that

$$\hat{T}_{\text{forget}} = n - k + 1 - \frac{n-k}{2^k} - \frac{1}{2^{n-1}} \ ,$$

where k is the largest integer with $2^k \le n-k+1$ (so $k \sim \log_2 n$; $H(i,\mu)$ behaves differently for $i < k$ and $i \ge k$). This value is strictly smaller than T_{mix}, but still asymptotically n.

Acknowledgement. We are grateful to Andy Beveridge and Jim Fill for their careful reading of this manuscript, which uncovered many errors and ambiguities.

References

[1] D.J. Aldous and J.A. Fill, *Reversible Markov Chains and Random Walks on Graphs* (book in preparation).

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

[3] D.J. Aldous, Applications of random walks on graphs, preprint (1989).

[4] D.J. Aldous and P. Diaconis, Shuffling cards and stopping times, *Amer. Math. Monthly* **93** No. 5 (1986), 333–348.

[5] D.J. Aldous, L. Lovász and P. Winkler, Mixing times for uniformly ergodic Markov chains, preprint.

[6] D. Bayer and P. Diaconis, Trailing the dovetail shuffle to its liar, *Ann. Appl. Prob.* **2** (1992), 294–313.

[7] 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.

[8] P. Doyle, seminar at Center for Elementary Studies, Pinecrest Lodge, Sonora CA (1995).

[9] P.G. Doyle and J.L. Snell, *Random Walks and Electric Networks*, Mathematical Assoc. of America, Washington, DC 1984.

[10] E. Gilbert, Theory of shuffling, Bell Laboratories Technical Memorandum (1955).

[11] L. Lovász and P. Winkler, Efficient stopping Rules for Markov Chains, *Proceedings of the 1995 ACM Symposium on the Theory of Computing,* 76–82.

[12] L. Lovász and P. Winkler, Mixing of random walks and other diffusions on a graph, *Surveys in Combinatorics, 1995*, P. Rowlinson, ed., London Math. Soc. Lecture Note Series 218, Cambridge U. Press (1995), 119–154.

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

[14] J. Reeds, unpublished manuscript (1981).