Exact Mixing in an Unknown Markov Chain

László Lovász and **Peter Winkler**

Submitted: May 2, 1995; Accepted: May 27, 1995.

Abstract

We give a simple stopping rule which will stop an unknown, irreducible n-state Markov chain at a state whose probability distribution is exactly the stationary distribution of the chain. The expected stopping time of the rule is bounded by a polynomial in the maximum mean hitting time of the chain. Our stopping rule can be made deterministic unless the chain itself has no random transitions.

Mathematics Subject Classification: 60J10

Key Words and Phrases: Markov chain, stopping time, mixing, stationary distribution

1 Introduction

Suppose a Markov process $s_0, s_1, s_2, \ldots$ on the state space $\{1, 2, \ldots, n\}$ is observed, with no knowledge either of the transition probabilities or the distribution of s_0. Unless the process is reducible (some states inaccessible from others) or periodic, the probability distribution of the state s_m will be approximately equal to the stationary distribution π of the process, for m sufficiently large.

In fact, this approach to sampling from a state space according to the stationary distribution is the basis for numerous recent estimation algorithms (see, e.g., [1], [16], [17]). Typically the initial state is fixed, the process is reversible (representable as a random walk on a graph) and some bound is obtained for the "mixing time" m. The payoff has been polynomial time randomized approximation algorithms for counting combinatorial objects such as matchings [17, 10], linear extensions [18], and Eulerian orientations [20]; estimating the volume of a convex body [16, 19]; and for Monte Carlo integration [6].

There is no *a priori* reason why a state must be sampled at a fixed number of steps. If the transition probabilities are known, a stopping rule which "looks where it is going" is capable of reaching the stationary distribution rapidly and exactly; in [5] a construction is given for intelligent stopping rules that achieve any target distribution in both the minimum expected number of steps and the minimum maximum number of steps. Several formulas and bounds are given for the number of steps $T_{\rm mix}$ required by an optimal stopping rule (starting from the worst state).

When the transition probabilities are not known, an intelligent stopping rule can be used to examine the process and then estimate how many steps must be taken to approximate the stationary distribution. Using this approach Aldous [3] comes within total variation ϵ of the stationary distribution in time polynomial in $1/\epsilon$ and linear in the maximum hitting time of the chain.

Since it is obviously impossible to *compute* the stationary distribution of an unknown chain exactly, it seems a bit surprising that one can *achieve* it exactly. Nonetheless that is what is done in Asmussen et al. [7]. However, the algorithm employed there is complex and requires perfect generation of random variables with certain exponential distributions. The expected number of steps required appears to be super-polynomial in the maximum hitting time, although no bound or estimate is given in the paper.

It turns out, however, that there is a simple, combinatorial stopping rule which can reach the stationary distribution exactly, in any irreducible, n-state Markov chain; the rule requires only coin-flips for its randomization and can even be made deterministic unless the chain itself is completely deterministic. The expected stopping time of the randomized rule is bounded by a polynomial (namely, $6h^4$) in the maximum hitting time of the chain.

We point out that this time bound is not good enough for the randomized algorithms mentioned above, since in them the approximately stationary distribution is achieved in a time $O(T_{\rm mix})$, which is typically polylogarithmic in h. But this shortcoming of our algorithm cannot be fixed; we will show that mixing in an unknown Markov chain cannot be achieved in time less than h.

2 Notation and Preliminaries

In what follows $M = \{p_{ij}\}$ is the transition matrix for an irreducible Markov chain on the state space $S = \{1, 2, \ldots, n\}$. Let $\pi = (\pi_1, \ldots, \pi_n)$ be the stationary distribution of the chain, so that $\pi^{\mathsf T} M = \pi^{\mathsf T}$.

Following the notation of Aldous (see e.g. [1]), we let T_j be the number of steps before first arrival at state j, with $\mathrm{E}_i T_j$ being the expected value of T_j when the process is begun in state

i. Then what we have been calling the "maximum hitting time" is $\max_{i,j\in S} \mathrm{E}_i T_j$ and will be denoted here by the letter h. The maximum hitting time is a lower bound on the *cover time*, which is the expected number of steps before all states are visited, maximized over all starting states.

We think of a stopping rule as a (possibly randomized) algorithm which, based on states so far seen, decides at each step whether to stop the Markov chain. Since we are interested in stopping rules that work for an unknown chain, the rule must decide when to stop based on the *pattern* of the states visited. This implies that such a rule needs substantial time; for example, we cannot rely on repetitions before n steps. (The "time" taken by a stopping rule is merely the expected number of steps before stopping, and has nothing to do with the computational complexity of the algorithm itself. However, our algorithm will only use polynomial time computations.) In fact, we show that the cover time is a lower bound on the expected number of steps. This follows immediately from the next observation.

Proposition 1 *Let the number n of states be fixed. Consider any stopping rule that decides when to stop based on the pattern of the states seen before, and assume that for every Markov chain on n states, the distribution of the state where it stops is the stationary distribution. Then it never stops without visiting all nodes.*

Proof. Consider any Markov chain M on n states, and consider a walk $(v_0, \ldots, v_t)$ that is stopped before seeing all states, and let j be state not visited. We replace j by a nearly absorbing state as follows. Construct a new Markov chain M' by replacing p_{ji} by δp_{ji} for all $i \neq j$ and p_{ii} by $1 - \delta(1 - p_{jj})$, where δ is very small. The stationary distribution of the new chain is $\pi'_i = \delta\pi_i/(\pi_i + \delta - \delta\pi_i)$ for $i \neq j$ and $\pi'_j = \pi_j/(\pi_j + \delta - \delta\pi_j)$. The walk $(v_0, \ldots, v_t)$ has the same probability in the old chain as in the new, and hence it must not exceed $\pi'(v_t)$, which tends to 0 as $t \to \infty$. This is a contradiction. $\square$

The same argument holds if we assume only that the probability of stopping at any state is at most some constant times its stationary probability.

3 Random Trees

Definition. Let $j \in S$. A *j-assignment* is a function $A_j : S \setminus \{j\} \longrightarrow S$. The *weight* $w(A_j)$ is defined by

$$w(A_j) := \prod_{i \neq j} p(i, A_j(i)) \ .$$

We may, for example, define a j-assignment A_j^t by "first exit after time t", that is, $A_j^t(i) = w_{k+1}$ where $k = \min\{t' : t' \geq t \text{ and} w_{t'} = j\}$. Then we can interpret $w(A_j)$ as the probability $\Pr(A_j^t = A_j)$ that a particular assignment A_j occurs in this construction, since all the exits are independent.

A j-assignment A_j defines a directed graph on S by placing an arc from i to $A_j(i)$ for each $i \neq j$; we say that A_j is a *j-tree* if this graph is a tree, necessarily an in-directed tree rooted at j. We denote by Υ_j the set of all j-trees on S. The following "random tree lemma" (which can be verified by straightforward substitution) has been, according to Aldous [2], frequently rediscovered; the earliest explicit reference we know of is [15], but it also follows easily from Tutte's matrix-tree theorem (see e.g [8]).

Lemma 1 *For any state $j \in S$,*

$$\pi_j = w(\Upsilon_j) / \sum_{i \in S} w(\Upsilon_i)$$

where $w(\Upsilon_i) := \sum_{A \in \Upsilon_i} w(A)$.

Remark. It may be instructive to describe the following construction related to the lemma. Run the Markov process given by M from $-\infty$ to $+\infty$ and for each time t, define a k-assignment A^t by *last prior exit*, where k is the state of the chain at time t. In other words, for each $i \neq k$, if t_i is the last time before t at which the chain is in state i, then $A^t(i)$ is defined to be the state of the chain at time $t_i + 1$. Note that A^t must be a tree, rooted at k, since all the arcs are oriented forward in time. Furthermore, A^{t+1} depends only on A^t and the state at time $t+1$, so we now have a stationary Markov process on trees.

Suppose now that the probability distribution of the tree observed at time t is given by $\Pr(A^t) = cw(A^t)$, where c is (necessarily) the reciprocal of the sum of the weights of all trees on the state space S. If a certain fixed tree A rooted at k is to occur at time $t+1$, then its predecessor, the tree A^t at time t, must be constructible from A by adding the arc $k \to i$ for some i, and then removing exactly the arc $i \to j$ where $j = j(i)$ is the last state before k in the path from i to k in A. For such an A^t the *a priori* probability of achieving A at the next step is just $p_{j,k}$, thus the total probability of seeing A at time $t+1$ is

$$\sum_{i \in S} \left[p_{j(i),k} \left(cw(A) \cdot \frac{p_{k,i}}{p_{j(i),k}} \right) \right] = cw(A) \ .$$

It follows that $cw(\cdot)$ is the stationary distribution for our tree process, but of course the stationary distribution for the roots is just π so we have that π_i is proportional to $w(\Upsilon_i)$.

Aldous [2] and Broder [11] use a closely related construction to design an elegant algorithm to generate a random spanning tree of a graph.

Lemma 1 already provides a stopping rule, described below, that attains the stationary distribution. In contrast to the procedure described above, the stopping rule constructs a random j-assignment by looking *forward* in time; then, as previously noted, the probability of a given assignment is exactly its weight, independent of the starting state. The price we pay is that the assignment is no longer necessarily a tree.

1. Choose a state j uniformly from S, and set current time equal to 0.
2. For each $i \in S \setminus \{j\}$ let t_i be the least $t \geq 0$ at which the chain is in state i, and set $A_j(i)$ to be the state of the chain at time $t_i + 1$.
3. By the time every state $i \in S \setminus \{j\}$ has been exited, we will know whether the resulting assignment A_j is a tree. If it is, we continue until the chain reaches j and then stop; if not, we repeat step 1.

Since the chain is irreducible, step 2 is finite with probability 1 and there must be some tree assignment which is eventually reached, say at iteration k. Letting T_i be the tree assignment constructed at that time, we have that Pr(the rule stops at j) = $\Pr(i = j) = \Pr(T_i \in \Upsilon_j \mid T_i$ is a tree assignment$) = \pi_j$. Unfortunately it may be the case that $\Pr(A_j$ is a tree) is exponentially small in n, even when the Markov chain has no small positive transition probabilities. For example, in a simple random walk on an n-cycle, where $p_{i,i+1} = p_{i+1,i} = 1/2$ for $i = 0, \ldots, n-1$ mod n, our stopping rule takes more than 2^n steps on average while the maximum expected time to hit a given state is only $n^2/4$.

To speed up the stopping rule, we make use of the fact that for an independent stochastic process (i.e. a Markov chain whose transition matrix has identical rows) the probability that a random assignment is a tree is fairly high—in fact, surprisingly, it depends only on n. The following lemma has appeared in many guises and is deducible, for example, from Theorem 37 of Chapter XIV in Bollobás [9]; we give an independent proof.

Lemma 2 *Let $X_1, \ldots, X_n$ be i.i.d. random variables with values in S. Define an assignment A_j by choosing $j \in S = \{1, 2, \ldots, n\}$ uniformly at random, then setting $A_j(i) = X_i$ for $i \neq j$. Then $Pr(A_j \in \Upsilon_j) = 1/n$.*

Proof. Let $m_1, \ldots, m_n$ be non-negative integers which sum to $n-1$. We may build an in-directed tree in which vertex i has in-degree m_i as follows: assuming that the in-neighbor sets $N^{\text{in}}(1), \ldots, N^{\text{in}}(k-1)$ have already been chosen, we select $N^{\text{in}}(k)$ from

$$S \setminus \cup_{i=1}^{k-1} N^{\text{in}}(i) \setminus \{j\}$$

where j is the root (possibly k itself) of the component currently containing k. It follows that the number of such trees is

$$\binom{n-1}{m_1} \cdot \binom{n-m_1-1}{m_2} \cdot \binom{n-m_1-m_2-1}{m_3} \cdot \cdots \cdot \binom{m_n}{m_n} = \binom{n-1}{m_1, m_2, \ldots, m_n} .$$

Since the weight of such a tree is $\prod_{i=1}^n p_i^{m_i}$ where $p_i = \Pr(X = i)$, we have that the sum of the weights of all the in-directed trees is

$$\sum_{m_1+\cdots+m_n=n-1} \binom{n-1}{m_1, m_2, \ldots, m_n} \prod_{i=1}^n p_i^{m_i} = (p_1 + \cdots + p_n)^{n-1} = 1$$

and thus the desired probability is

$$\frac{1}{n} \sum_{j=1}^{n} \Pr(A_j \in \Upsilon_j) \;=\; \frac{1}{n} \,.$$

□

4 A Randomized Stopping Rule

To make use of Lemma 2 we need to replace the transition matrix M by a new matrix N having the same stationary distribution but which represents a nearly independent process; in other words, the rows of N should be similar to one another (and therefore to the stationary vector π).

An obvious candidate for N is M^t, for t some polynomial in n and the maximum hitting time h, and in fact this choice will suffice for reversible Markov chains. In general, however, "mixing time" may be exponentially larger than both n and h. For example, suppose $p_{i,i+1} = 1$ for $i = 1, \ldots, n-1$, $p_{n,1} = 1 - 2^{-n}$, $p_{n,n} = 2^{-n}$ and all other transitions are forbidden. Then h is only about n but the state of the chain t steps after being at state j is $j + t \pmod n$ with high probability for fixed $t < 2^n$.

Instead we take N to be an average of the matrices M^k for k between 1 and some sufficiently large bound t.

Lemma 3 *Let M be the transition matrix for an n-state irreducible Markov chain $(v^0, v^1, \ldots)$ with stationary distribution π and maximum hitting time h, and let $t \geq 1$. Let Z be chosen uniformly from $\{1, \ldots, t\}$. Then for every state j,*

$$\Pr(v^Z = j) \geq (1 - \frac{h}{t})\pi_j.$$

Proof. Let s be any positive integer, and let Y_j^s be a random variable which counts the number of hits of state j in the next s steps of the chain M. Again using Aldous' notation, we let $\mathrm{E}_\sigma Y_j^s$ be the expected value of Y_j^s when the chain is begun in a state drawn from the distribution σ; if σ is concentrated at i we just write $\mathrm{E}_i Y_j^s$.

For any i and s, we have $\mathrm{E}_i Y_j^s \leq 1 + \mathrm{E}_j Y_j^s$ (by waiting for the first occurrence of j) and thus in particular, $\pi_j = \frac{1}{s}\mathrm{E}_\pi Y_j^s \leq \frac{1}{s}(1 + \mathrm{E}_j Y_j^s)$.

Fix i and j and let q_s be the probability that, when started at state i, the first occurrence of state j is at step s. By definition of N, we have

$$\begin{aligned}
\Pr(v^Z = j) &= \frac{1}{t}\mathrm{E}_i Y_j^t = \frac{1}{t}\sum_{s=1}^{t} q_s(1 + \mathrm{E}_j Y_j^{t-s}) \\
&\geq \frac{1}{t}\sum_{s=1}^{t} q_s(\pi_j(t-s)) \\
&= \pi_j - \frac{\pi_j}{t}\sum_{s=1}^{t} s q_s \\
&\geq \pi_j - \frac{\pi_j}{t}\mathrm{E}_i T_j \\
&\geq \pi_j - \frac{\pi_j}{t} h
\end{aligned}$$

as desired. □

Below we shall need that $N_{ij} \geq (1-(1/n))\pi_j$ for all i and j. We can achieve this by choosing $t = \lceil nh \rceil$. This is good enough if we are only interested in polynomiality, but the time bounds we get this way are too pessimistic on two counts. We could apply the "multiplicativity property" in Aldous [4] to show that the factor n could be replaced by $\log n$, and results from [5] to show that h can be replaced by the mixing time T_{mix}.

More exactly, let $M = \lceil \log n \rceil$ and $s = 8\lceil T_{\mathrm{mix}} \rceil$, and let Z be the sum of M independent random variables $Y_1, \ldots, Y_M$, each distributed uniformly in $\{0, \ldots, s-1\}$. Then results from [5] imply that for any starting point,

$$\Pr(v^Z = j) \geq \left(1 - \frac{1}{n}\right)\pi_j.$$

To get around the difficulty that the maximum hitting time h is not known, we start with $t = 1$ and double t until we are successful in constructing a tree; for each t we construct $3n$ assignments (the proof below uses $3 > e$). Altogether our randomized stopping rule Θ runs as follows:

For $t = 1, 2, 4, 8, \ldots$ do

    For $k = 1, 2, 3, \ldots, 3n$ do

        Choose a state j uniformly from S

        Put $U = \{j\}$

        Do until $U = S$

            Proceed until a state $i \notin U$ is reached

            Choose a random number m uniformly from $1, \ldots, t$

            Proceed m steps and designate the current state as $A_j(i)$

            Update $U \leftarrow U \cup \{i\}$

            End

        If the assignment A_j is a tree, proceed to state j and STOP

    Next k

Next t

Theorem 1 *Stopping Rule Θ runs in expected number of steps polynomial in h, and stops at state j with probability exactly π_j.*

Remark. The proof below gives that the expected number of steps is $O(n^2h^2) = O(h^4)$. Using the bounds mentioned after Lemma 3 we get the tighter bound $O(hT_{\rm mix} n \log n) = O(h^2 n \log n) = O(h^3 \log n)$ by the same argument.

Proof. For each fixed t and k, the expected number of steps taken by the assignment construction is no more than $3nh(t+1)/2$, hence before t reaches nh the algorithm expects to take fewer than $3n^2h^2$ steps. Afterwards the probability of "success" (achieving a tree and stopping) for given t and k is at least

$$\frac{1}{n}\left(1-\frac{1}{n}\right)^{n-1} > \frac{1}{en}$$

on account of Lemmas 2 and 3, since each factor in the expression for the weight of any assignment (in particular, any tree) is short by at most a factor of $1-1/n$ of the corresponding stationary probability.

It follows that for fixed $t \geq nh$ the success probability is at least

$$1-\left(1-\frac{1}{en}\right)^{3n} > 1-e^{-3/e} > 2/3\ .$$

Setting t_0 equal to the first value of t above nh, and letting m be such that the algorithm stops at $t = 2^m t_0$, we have that the expected total number of steps is less than

$$3n^2h^2 + \sum_{m=0}^{\infty}(2/3)(1/3)^m 2^m (3nht_0/2) < 6n^2h^2\ .$$

It remains only to argue that Pr(the rule stops at state i) $= \pi_i$, but this follows from previous remarks plus the fact that the stationary distribution for $N = \frac{1}{t}\sum_{k=1}^{t} M^k$ is the same as for M. □

As an example, suppose M has only two states a and b, with transition probabilities $p_{a,b} = p > 0$ and $p_{b,a} = q > 0$. We may achieve the stationary distribution $\pi = (q/(p+q), p/(p+q))$ by the following procedure: flip a fair coin; if "heads" wait for the first exit from state a and if "tails", the first exit from b. If the exit is to the opposite state, stop right there; else flip again. After 6 unsuccessful flips, repeat but take 1- or 2-step exits with equal probability; then 1-, 2-, 3- or 4-step etc.

This two-state algorithm can be generalized to an n-state stopping rule by recursion, giving another solution to our problem (with about the same bound on expected number of steps).

5 Derandomization

The randomization required for Stopping Rule Θ is easily accomplished by coin flips, since we need only uniform choices between 1 to n and between 1 to t, with n and t both known. But coin flips can be done using the Markov process itself as long as there is some transition probability $p_{i,j}$ which lies in the open interval (0,1). (Otherwise there is no hope, as we cannot reach a random state in a deterministic process without outside randomization.) The technique is "von Neumann's trick" for getting a fair decision from a bent coin.

To obtain from Θ a deterministic stopping rule Δ, we observe the process for a while and make a list L of states i with corresponding sets $U_i \subset S$ such that

$$\pi_i \left(\sum_{j \in U_i} p_{i,j} \right) \left(1 - \sum_{j \in U_i} p_{i,j} \right)$$

is about as high as possible.

Then we proceed with Θ but when a coin flip is needed, we wait until some state in L occurs. Suppose this happens at time t_1 with state i; we then proceed to the next occurrence of i, say at time t_2, and we take one further step. We now check whether we were in U_i at exactly one of the two times $t_1 + 1$ and $t_2 + 1$. If so we have made a successful coin flip, the result being "heads" if our state at time $t_1 + 1$ is in U_1 and "tails" otherwise.

If we hit U_i both times or neither time we try again, waiting for another state in L to occur.

For "most" Markov processes the time to consummate a coin flip will be negligible but if all transition probabilities are close to 0 or 1, or if the only exceptional $p_{i,j}$'s correspond to states i with very low stationary probability, then the derandomization may cost Θ its polynomiality in h. The deterministic stopping rule Δ will, however, be polynomial in h and r where $1/r$ is the stationary frequency of the most common transition $i \to j$ such that $p_{i,j} < p_{i,j'}$ for some j'.

Remark.

A faster (randomized) algorithm for exact mixing in an unknown chain has now been devised by J.G. Propp and D.B. Wilson, using the elegant notion of "coupling from the past." Their stopping rule runs in expected time bounded by a constant times the expected cover time (thus best possible), and will appear in a paper entitled "How to get an exact sample from a generic Markov chain."

Acknowledgments.

The authors are indebted to David Aldous and Eric Denardo for many useful comments and corrections.

Authors' addresses:

Dept. of Computer Science, Yale University, New Haven CT 06510 USA, lovasz@cs.yale.edu

AT&T Bell Laboratories 2D-147, 600 Mountain Ave., Murray Hill NJ 07974 USA, pw@research.att.com

References

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

[2] D.J. Aldous, The random walk construction of uniform spanning trees and uniform labelled trees, *SIAM J. Disc. Math.* **3** (1990), 450–465.

[3] D.J. Aldous, On simulating a Markov chain stationary distribution when transition probabilities are unknown, prepint (1993).

[4] D.J. Aldous, *Reversible Markov Chains and Random Walks on Graphs* (book), to appear.

[5] D.J. Aldous, L. Lovász and P. Winkler, Fast mixing in a Markov chain, in preparation (1995).

[6] D. Applegate and R. Kannan, Sampling and integration of near log-concave functions, *Proc. 23rd ACM Symp. on Theory of Computing* (1991), 156–163.

[7] S. Asmussen, P.W. Glynn and H. Thorisson, Stationary detection in the initial transient problem, *ACM Transactions on Modeling and Computer Simulation* **2** (1992), 130–157.

[8] C. Berge, *Graphs and Hypergraphs*, North-Holland, Amsterdam, 1973.

[9] B. Bollobás, *Random Graphs*, Academic Press, London, 1985.

[10] A.Z. Broder, How hard is it to marry at random? (on the approximation of the permanent), *Proc. 18th ACM Symp. on Theory of Computing* (1986), 50–58.

[11] A.Z. Broder, Generating random spanning trees, *Proc. 30th IEEE Symp. on Found. of Computer Science* (1989), 442–447.

[12] A.K. Chandra, P. Raghavan, W.L. Ruzzo, R. Smolensky, and P. Tiwari, The electrical resistance of a graph captures its commute and cover times, *Proc. 21st ACM Symp. on Theory of Computing* (1989), 574–586.

[13] 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.

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

[15] M.I. Friedlin and A.D. Wentzell, Random perturbations of dynamical systems, *Russian Math. Surveys* (1970) 1–55.

[16] M. Dyer, A. Frieze and R. Kannan, A random polynomial time algorithm for estimating volumes of convex bodies, *Proc. 21st ACM Symp. on Theory of Computing* (1989), 375–381.

[17] M. Jerrum and A. Sinclair, Conductance and the rapid mixing property for Markov chains: the approximation of the permanent resolved, *Proc. 20th ACM Symp. on Theory of Computing* (1988), 235–243.

[18] A. Karzanov and L. Khachiyan, On the conductance of order Markov chains, Technical Report DCS TR 268, Rutgers University, June 1990.

[19] L. Lovász and M. Simonovits, Random walks in a convex body and an improved volume algorithm, *Random Structures and Algorithms* **4** (1993), 359-412.

[20] M. Mihail and P. Winkler, On the number of Eulerian orientations of a graph, *Algorithmica*, to appear.

[21] D.E. Symer, Expanded ergodic Markov chains and cycling systems, senior thesis, Dartmouth College, Hanover NH (1984).

[22] P. Tetali, Random walks and effective resistance of networks, *J. Theor. Prob.* **1** (1991), 101-109.