

# On Approximating the Stationary Distribution of Time-reversible Markov Chains

**Marco Bressan**

Dipartimento di Informatica, Sapienza Università di Roma, Italy  
bressan@di.uniroma1.it

**Enoch Peserico**

Dipartimento di Ingegneria dell'Informazione, Università di Padova, Italy  
enoch@dei.unipd.it

**Luca Pretto**

Dipartimento di Ingegneria dell'Informazione, Università di Padova, Italy  
pretto@dei.unipd.it

---

## Abstract

Approximating the stationary probability of a state in a Markov chain through Markov chain Monte Carlo techniques is, in general, inefficient. Standard random walk approaches require  $\tilde{O}(\tau/\pi(v))$  operations to approximate the probability  $\pi(v)$  of a state  $v$  in a chain with mixing time  $\tau$ , and even the best available techniques still have complexity  $\tilde{O}(\tau^{1.5}/\pi(v)^{0.5})$ ; and since these complexities depend inversely on  $\pi(v)$ , they can grow beyond any bound in the size of the chain or in its mixing time. In this paper we show that, for time-reversible Markov chains, there exists a simple randomized approximation algorithm that breaks this “small- $\pi(v)$  barrier”.

**2012 ACM Subject Classification** Mathematics of computing → Markov-chain Monte Carlo methods, Theory of computation → Sketching and sampling, Theory of computation → Random walks and Markov chains

**Keywords and phrases** Markov Chains, MCMC Sampling, Large Graph Algorithms, Randomized Algorithms, Sublinear Algorithms

**Digital Object Identifier** 10.4230/LIPIcs.STACS.2018.18

**Related Version** For a full version of this paper see <https://arxiv.org/abs/1801.00196v1>, [7].

**Funding** Marco Bressan is supported in part by the SIR Grant RBSI14Q743. Enoch Peserico is supported in part by Univ. Padova under Project CAEPAE.

## 1 Introduction

We investigate the problem of approximating efficiently a *single* entry of the stationary distribution of an ergodic Markov chain. This problem has two main motivations. First, with the advent of massive-scale data, even complexities linear in the size of the input are often excessive [19]; therefore computing explicitly the entire stationary distribution, e.g. via the power method [10], can be simply infeasible. As an alternative one can then resort to approximating only individual entries of the vector, in exchange for a much lower computational complexity [13, 20]. In fact, if such a complexity is low enough one could efficiently “sketch” the whole vector by quickly getting a fair estimate of its entries. Second, in many practical cases one is really interested in just a few entries at a time. A classic example is that of network centralities, many of which are stationary distributions of an



© Marco Bressan, Enoch Peserico, and Luca Pretto;  
licensed under Creative Commons License CC-BY

35th Symposium on Theoretical Aspects of Computer Science (STACS 2018).

Editors: Rolf Niedermeier and Brigitte Vallée; Article No. 18; pp. 18:1–18:14

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



SYMPOSIUM  
ON THEORETICAL  
ASPECTS  
OF COMPUTER  
SCIENCE

ergodic Markov chain [4]. Indeed, the problem of approximating the Personalized PageRank score of a few nodes in a graph has been repeatedly addressed in the past [5, 6, 16, 15].

In this paper we seek for efficient algorithms for approximating the stationary probability  $\pi(v)$  of some target state  $v$  in the state space of a discrete-time ergodic Markov chain. Besides the motivations above, the problem arises in estimating heat kernels and graph diffusions, testing the conductance of graphs and chains, developing local algorithms, and has applications in machine learning; see [3, 12] for a thorough discussion. We adopt a simple model where with a single operation one can either (i) simulate one step of the chain or (ii) retrieve the transition probability between a pair of states. Although recent research has provided encouraging results, existing algorithms suffer from a crucial bottleneck: to guarantee a small *relative* error in the approximation of  $\pi(v)$ , they incur a cost that grows with  $1/\pi(v)$  itself (basically because estimating  $\pi(v)$  via repeated sampling requires  $1/\pi(v)$  samples). This is a crucial issue since in general there is no lower bound on  $\pi(v)$ ; even worse, if the state space has  $n$  states, then most states have mass  $\pi(v) = O(\frac{1}{n})$ , and one can easily design chains where they have mass exponentially small in  $n$ . In general, then, the cost of existing algorithms can blow up far beyond  $O(n)$  for almost all input states  $v$ . It is thus natural to ask if the dependence of the complexity on  $\pi(v)$  is unavoidable. Unfortunately, one can easily show that  $\Omega(\tau/\pi(v))$  operations can be necessary to estimate  $\pi(v)$  within any constant multiplicative factor if one makes no assumption on the chain (see Appendix 5.2). To drop below this complexity barrier one must then necessarily look at special classes of Markov chains.

We present an algorithm that breaks this “small- $\pi(v)$  barrier” for *time-reversible* Markov chains. Time-reversible chains are a well-known subclass of Markov chains which lie at the heart of the celebrated Metropolis-Hastings algorithm [11] and are equivalent to random walks on weighted undirected graphs [14]. Formally, given any  $\epsilon, \delta > 0$  and any state  $v$  in a time-reversible chain, our algorithm with probability  $1 - \delta$  returns a multiplicative  $(1 \pm \epsilon)$ -approximation of  $\pi(v)$  by using  $\tilde{O}(\tau\|\pi\|^{-1})$  operations, where  $\tau$  is the mixing time of the chain,  $\|\cdot\|$  is the Euclidean norm and  $\tilde{O}(\cdot)$  hides polynomials in  $\epsilon^{-1}$ ,  $\ln(\delta^{-1})$ ,  $\ln(\|\pi\|^{-1})$ . The complexity is independent of  $\pi(v)$ , and for all but a vanishing fraction of states in the chain improves by factors at least  $\sqrt{n}$  or  $\sqrt{\tau}$  over previous algorithms. The heart of our algorithm is a randomized scheme for approximating the sum of a nonnegative vector by sampling its entries with probability proportional to their values. This scheme requires  $\tilde{O}(\|\pi\|^{-1})$  samples if  $\pi$  is the distribution over the vector entries, which generalizes the  $O(\sqrt{n})$  algorithm of [17] and is provably optimal. We prove that our algorithm for estimating  $\pi(v)$  is essentially optimal as a function of  $\tau$ ,  $n$  and  $\|\pi\|$ ; in fact one cannot do better even under a stronger computational model where all transition probabilities to/from all visited states are known. Finally, we show the number of distinct states visited by our algorithm may be further reduced, provided such a number satisfies some concentration hypotheses. This is useful if visiting a new state is expensive (e.g. if states are users in a social network). All our algorithms are simple to implement, require no tuning, and experimentally they appear faster than existing alternatives already for medium-sized chains.

The rest of the paper is organized as follows. Subsection 1.1 pins down definitions and notation; Subsection 1.2 formalizes the problem; Subsection 1.3 discusses related work; Subsection 1.4 summarizes our results. Section 2 presents our vector sum approximation algorithm. Section 3 presents our approximation algorithm for  $\pi(v)$ . The proofs of our lower bounds (Theorem 3 and Theorem 5), the pseudocode of our algorithms, and our experimental results can be found in the full version of this paper [7].

## 1.1 Preliminaries

A discrete-time, finite-state Markov chain is a sequence of random variables  $X_0, X_1, \dots$  taking value over a set of states  $V = \{1, \dots, n\}$ , such that for all  $i \geq 1$  and all  $u_0, \dots, u_i \in V$  with  $\Pr(X_0 = u_0, \dots, X_{i-1} = u_{i-1}) > 0$  we have  $\Pr(X_i = u_i | X_0 = u_0, \dots, X_{i-1} = u_{i-1}) = \Pr(X_i = u_i | X_{i-1} = u_{i-1})$ . Denote by  $\mathbf{P} = [p_{uu'}]$  the transition matrix of the chain, so that  $p_{uu'} = \Pr(X_i = u' | X_{i-1} = u)$ . We assume the chain is *ergodic*, and thus has a limit distribution that is independent from the distribution of  $X_0$ ; the limit distribution then coincides with the *stationary distribution*  $\pi$ . Thus  $\pi$  is the unique distribution vector such that for any distribution vector  $\pi_0$ :

$$\pi = \pi \mathbf{P} = \lim_{t \rightarrow \infty} \pi_0 \mathbf{P}^t \quad (1)$$

We denote by  $\pi(u)$  the stationary probability, or *mass*, of  $u$ , and we always denote by  $v$  the target state whose mass is to be estimated. For any  $V' \subseteq V$  we let  $\pi(V')$  denote  $\sum_{u \in V'} \pi(u)$ . We also assume the chain is *time-reversible*, i.e. that for any pair of states  $u$  and  $u'$  we have:

$$\pi(u)p_{uu'} = \pi(u')p_{u'u} \quad (2)$$

We denote by  $\tau$  the standard  $\frac{1}{4}$ -*mixing time* of the chain. In words,  $\tau$  is the smallest integer such that after  $\tau$  steps the total variation distance between  $\pi$  and the distribution of  $X_\tau$  is bounded by  $\frac{1}{4}$ , irrespective of the initial distribution. Formally,  $\tau := \min\{t : d(t) \leq \frac{1}{4}\}$ , where

$$d(t) := \max_{\pi_0} \|\pi_0 \mathbf{P}^t - \pi\|_{\text{TV}} = \max_{\pi_0} \frac{1}{2} \|\pi_0 \mathbf{P}^t - \pi\|_1 \quad (3)$$

After  $\tau$  steps, the distribution of  $X_t$  converges to  $\pi$  exponentially fast; that is, if  $t = \eta\tau$  with  $\eta \geq 1$ , then  $\|\pi_0 \mathbf{P}^t - \pi\|_{\text{TV}} \leq 2^{-\eta}$ . In the rest of the paper,  $\|\cdot\|$  always denotes the  $\ell^2$  norm. One may refer to [14] for a detailed explanation of the notions recalled here.

Unless necessary, we drop multiplicative factors depending only on  $\epsilon, \delta$  (see below) from the asymptotic complexity notation. Furthermore, we use the tilde notation to hide polylogarithmic factors, i.e. we denote  $O(f \cdot \text{poly}(\log(f)))$  by  $\tilde{O}(f)$ .

## 1.2 Problem formulation

Consider now a discrete-time, finite-state, time-reversible, ergodic Markov chain on  $n$  states. The chain is initially unknown and can be accessed via two operations (also called queries):

*step()*: accepts in input a state  $u$ , and returns state  $u'$  with probability  $p_{uu'}$

*probe()*: accepts in input a pair of states  $u, u'$ , and returns  $p_{uu'}$

These queries are the *de facto* model of previous work. *step()* is used in [5, 12, 6, 16, 15, 3] to simulate the walk, assuming each step costs  $O(1)$ . *probe()* is used in [16, 15, 3] to access the elements of the transition matrix, assuming again one access costs  $O(1)$ . Here, too, we assume *step()* and *probe()* as well as all standard operations (arithmetics, memory access, ...) cost  $O(1)$ . This includes set insertion and set membership testing; in case their complexity is  $\omega(1)$ , our bounds can be adapted correspondingly. The problem can now be formalized as follows. The algorithm is given in input a triple  $(v, \epsilon, \delta)$  where  $v$  is a state in the state space of the chain and  $\epsilon, \delta$  are two reals in  $(0, 1)$ . It must output a value  $\hat{\pi}(v)$  such that, with probability  $1 - \delta$ , it holds  $(1 - \epsilon)\pi(v) \leq \hat{\pi}(v) \leq (1 + \epsilon)\pi(v)$ . The complexity of the algorithm is counted by the total number of operations it performs. Obviously we seek for an algorithm of minimal complexity.

A final remark. We say state  $u$  has been *visited* if  $u = v$  or if  $u$  has been returned by a  $step()$  call. In line with previous work, we adopt the following “locality” constraint: the algorithm can invoke  $probe()$  and  $step()$  only on visited states.

### 1.3 Related work

Two recent works address precisely the problem of estimating  $\pi(v)$  in Markov chains. The key differences with our paper are that they consider general (i.e. not necessarily time-reversible) chains, and that we aim at a small *relative* error for *any*  $\pi(v)$  and not only for large  $\pi(v)$ .

- [12] gives a local approximation algorithm based on estimating return times via truncated random walks. Given any  $\Delta > 0$ , if  $\pi(v) \geq \Delta$  the algorithm with probability  $1 - \delta$  outputs a multiplicative  $\epsilon Z(v)$ -approximation of  $\pi(v)$ , where  $Z(v)$  is a “local mixing time” that depends on the structure of the chain. The cost is  $\tilde{O}(\ln(1/\delta)/\epsilon^3 \Delta)$   $step()$  calls. If one wants a multiplicative  $(1 \pm \epsilon)$ -approximation of  $\pi(v)$  for a generic  $v$ , the cost becomes  $\tilde{O}(\tau/\pi(v))$   $step()$  calls since one must wait for the walks to hit  $v$  after having mixed.
- [3] gives an algorithm to approximate  $\ell$ -step transition probabilities based on coupling a local exploration of the transition matrix  $\mathbf{P}$  with simulated random walks. Given any  $\Delta > 0$ , if the probability to be estimated is  $\geq \Delta$  then with probability  $1 - \delta$  the algorithm gives a multiplicative  $(1 \pm \epsilon)$ -approximation of it at an expected cost of  $\tilde{O}(\ell^{1.5} \sqrt{d} \ln(1/\delta) / \epsilon \Delta^{0.5})$  calls to both  $step()$  and  $probe()$ , for a uniform random choice of  $v$  in the chain, where  $d$  is the density of  $\mathbf{P}$ . To estimate  $\pi(v)$  for a generic  $v$  one must set  $\ell = \tau$  and  $\Delta = \pi(v)$ , and since if the chain is irreducible then  $d = \Omega(1)$ , the bound stays at  $\tilde{O}(\tau^{1.5}/\pi(v)^{0.5})$ . This does not contradict our lower bound of Appendix 5.2, since their model allows for probing transition probabilities even between unvisited states.

Similar results are known for specific Markov chains, and in particular for PageRank (note that in PageRank  $\tau = O(1)$ ). [5, 6] give an algorithm for approximating the PageRank  $\pi(v)$  of the nodes  $v$  having  $\pi(v) \geq \Delta$ , at the cost of  $\tilde{O}(1/\Delta)$   $step()$  calls; again, if one desires a multiplicative  $(1 \pm \epsilon)$ -approximation of  $\pi(v)$ , the cost becomes  $\tilde{O}(1/\pi(v))$ . [16] gives an algorithm, with techniques similar to [3], for estimating the Personalized PageRank  $\pi(v)$  of a node  $v$ ; if one aims at a multiplicative  $(1 \pm \epsilon)$ -approximation of  $\pi(v)$ , the algorithm makes  $\tilde{O}(d^{0.5}/\pi(v)^{0.5})$   $step()$  and  $probe()$  calls where  $d$  is the average degree of the graph. Similar bounds can be found in [15] for Personalized PageRank on undirected graphs.

Summarizing, existing algorithms require either  $\tilde{O}(\tau/\pi(v))$  or  $\tilde{O}(\tau^{1.5}/\pi(v)^{0.5})$   $step()$  and  $probe()$  calls to ensure a  $(1 \pm \epsilon)$ -approximation of  $\pi(v)$  for a generic state  $v$ . Note that the complexity and approximation guarantees of these algorithms depend on knowledge of  $\tau$ ; our algorithms are no exception, and we prove our bounds as a function of  $\tau$ .

Finally, for the problem of estimating the sum of a nonnegative  $n$ -entry vector  $\mathbf{x}$  by sampling its entries  $x_i$ , with probability  $\pi_i = x_i / \sum_i x_i$ , the only algorithm existing to date is that of [17]. That algorithm takes  $O(\sqrt{n})$  samples independently of  $\pi$ , while ours needs  $O(\sqrt{n})$  samples only in the worst case, i.e. if  $\pi$  is (essentially) the uniform distribution.

### 1.4 Our results

Our first contribution is SumApprox, a randomized algorithm for estimating the sum  $\gamma$  of a nonnegative vector  $\mathbf{x}$ , assuming one can sample its entries according to the probability distribution  $\pi = \mathbf{x}/\gamma$ . Formally, we prove:

► **Theorem 1.** *Given any  $\delta, \epsilon \in (0, 1)$ ,  $SumApprox(\epsilon, \delta)$  with probability at least  $1 - \delta$  returns a multiplicative  $(1 \pm \epsilon)$ -approximation of  $\gamma$  by taking  $O(\|\pi\|^{-1} \epsilon^{-3} (\ln \frac{1}{\delta})^{3/2})$  samples.*

SumApprox is extremely simple, yet it improves on the state-of-the-art  $O(\sqrt{n})$  algorithm of [17]. We prove  $\Omega(\|\pi\|^{-1})$  samples are necessary, too, to get a fair estimate of  $\gamma$ .

We then employ SumApprox to build MassApprox, a randomized algorithm for approximating  $\pi(v)$ . Random-walk-based sampling and time reversibility are the ingredients that allow one to make the connection. We prove:

► **Theorem 2.** *Given any  $\delta, \epsilon \in (0, 1)$  and any state  $v$  in a time-reversible Markov chain,  $\text{MassApprox}(\epsilon, \delta, v)$  with probability  $(1 - \delta)$  returns a multiplicative  $(1 \pm \epsilon)$  approximation of  $\pi(v)$  using  $\tilde{O}(\tau\|\pi\|^{-1}\epsilon^{-3}(\ln \frac{1}{\delta})^{3/2}) = \tilde{O}(\tau\|\pi\|^{-1})$  elementary operations and calls to  $\text{step}()$  and  $\text{probe}()$ .*

Previous algorithms work also for general (i.e. non-reversible) chains; but on the  $n - o(n)$  states with mass  $\pi(v) = O(1/n)$ , their complexity becomes at least  $\tilde{O}(\tau n)$  [12] or  $\tilde{O}(\tau^{1.5}\sqrt{n})$  [3]. In fact,  $\pi(v)$  can be arbitrarily small (even exponentially small in  $n$  and  $\tau$ ) for almost all states in the chain, so for almost all states the complexity of previous algorithms blow up while that of MassApprox remains unchanged: since  $\|\pi\|^{-1} \leq \sqrt{n}$  for any  $\pi$ , the complexity of MassApprox is at most  $\tilde{O}(\tau\sqrt{n})$ .

Next, we show that MassApprox is optimal as a function of  $\tau$ ,  $n$  and  $\|\pi\|$ , up to small factors. In fact, no algorithm can perform better even if equipped with an operation  $\text{neigh}(u)$  that returns all incoming and outgoing transition probabilities of  $u$ . Formally, we prove:

► **Theorem 3.** *For any function  $\nu(n) \in \Omega(1/\sqrt{n}) \cap O(1)$  there is a family of time-reversible chains on  $n$  states where (a)  $\|\pi\| = \Theta(\nu(n))$ , and (b) there is a target state  $v$  such that, to estimate its mass  $\pi(v)$  within any constant multiplicative factor with constant probability, any algorithm requires  $\Omega(\tau\|\pi\|^{-1}/\ln n)$   $\text{neigh}()$  calls where  $\tau$  is the mixing time of the chain.*

Although bounding time complexity is our primary goal, in some scenarios one wants to bound the *footprint*, i.e. the number of distinct states visited. Obviously, the footprint of MassApprox is bounded by its complexity (Theorem 2). We give a second algorithm, FullMassApprox, whose footprint can be smaller than that of MassApprox depending on  $\tau$ ,  $n$ , and  $\|\pi\|$ . More precisely, we prove a footprint bound that is conditional on the concentration of the footprint itself (see Subsection 3.1 for the intuition behind it).

► **Theorem 4.** *Let  $N_{v,T}$  be the number of distinct states visited by a random walk of  $T$  steps starting from  $v$ . Assume for a function  $\bar{\tau}$  of the chain we have  $\Pr[N_{v,T} \notin \Theta(\mathbb{E}[N_{v,T}])] = o(\frac{\bar{\tau}}{\mathbb{E}[N_{v,T}]})$ . Then, given any  $\delta, \epsilon \in (0, 1)$ , with probability  $(1 - \delta)$  one can obtain a multiplicative  $(1 \pm \epsilon)$ -approximation of  $\pi(v)$  by visiting  $O(f(\epsilon, \delta)(\tau \ln n + \sqrt{\bar{\tau}n}))$  distinct states.*

If in Theorem 4 we have  $\bar{\tau} = \tau$ , then FullMassApprox is essentially optimal too. Formally:

► **Theorem 5.** *For any function  $\tau(n) \in \Omega(\ln n) \cap O(n)$  there is a family of time-reversible chains on  $n$  states where (a) the mixing time is  $\tau = \Theta(\tau(n))$ , and (b) there is a target state  $v$  such that, to estimate its mass  $\pi(v)$  within any constant multiplicative factor with constant probability, any algorithm requires  $\Omega(\sqrt{\tau n / \ln n})$   $\text{neigh}()$  calls.*

## 2 Estimating sum by weighted sampling

In this section we analyse the following problem. We are given a vector of nonnegative reals  $\gamma_u$  indexed by the elements  $u$  of a set  $V$ . The vector is unknown, including its length, but we can draw samples from  $V$  according to the distribution  $\pi$  where  $u$  has probability  $\gamma_u / \sum_{u \in V} \gamma_u$ . The goal is to approximate the vector sum  $\gamma = \sum_{u \in V} \gamma_u$ . We describe a

simple randomized algorithm, SumApprox, which proceeds by repeatedly drawing samples and checking for repeats (i.e. a draw that yields an element already drawn before). The key intuition is the following: at any instant, if  $S \subseteq V$  is the subset of elements drawn so far, then the next draw is a repeat with probability  $\sum_{u \in S} \gamma_u / \gamma$ . By drawing a sequence of samples we can thus flip a sequence of binary random variables, each one telling if a draw is a repeat, whose expectation is known save for the factor  $1/\gamma$ . If the sum of these random variables is sufficiently close to its expectation, one can then get a good approximation of  $\gamma$  by simply computing a ratio. The code of SumApprox is listed below.

---

**Algorithm** SumApprox( $\epsilon, \delta$ ).

---

```

1:  $S \leftarrow \emptyset$                                 ▷ distinct elements drawn so far
2:  $w_S \leftarrow 0$                                ▷  $\sum_{u \in S} \gamma_u$  for the current  $S$ 
3:  $w \leftarrow 0$                                  ▷ cumulative sum of  $\sum_{u \in S} \gamma_u$  so far
4:  $r \leftarrow 0$                                  ▷ number of repeats so far
5:  $k_{\epsilon, \delta} \leftarrow \lceil \frac{2+4\epsilon}{\epsilon^2} \ln \frac{3}{\delta} \rceil$     ▷ halting threshold on the number of repeats
6: while  $r < k_{\epsilon, \delta}$  do
7:    $w \leftarrow w + w_S$ 
8:    $(u, \gamma_u) \leftarrow$  sample drawn from distribution  $\pi$ 
9:   if  $u \in S$  then                             ▷ detect collision
10:     $r \leftarrow r + 1$ 
11:   else
12:     $S \leftarrow S \cup \{u\}$ 
13:     $w_S \leftarrow w_S + \gamma_u$ 
14: return  $w/r$                                    ▷ estimate of  $\gamma$ 

```

---

We prove:

► **Theorem 6.** *SumApprox( $\epsilon, \delta$ ) with probability at least  $1 - \frac{2\delta}{3}$  returns an estimate  $\hat{\gamma}$  such that  $|\hat{\gamma} - \gamma| < \epsilon\gamma$ .*

**Proof.** We make use of a martingale tail inequality originally from [9] and stated (and proved) in the following form as Theorem 2.2 of [1], p. 1476:

► **Theorem 7** ([1], Theorem 2.2). *Let  $(Z_0, Z_1, \dots)$  be a martingale with respect to the filter  $(\mathcal{F}_i)$ . Suppose that  $Z_{i+1} - Z_i \leq M$  for all  $i$ , and write  $V_t = \sum_{i=1}^t \text{Var}(Z_i | \mathcal{F}_{i-1})$ . Then for any  $z, v > 0$  we have*

$$\Pr[Z_t \geq Z_0 + z, V_t \leq v \text{ for some } t] \leq \exp\left[-\frac{z^2}{2(v + Mz)}\right]$$

Let us plug into the formula of Theorem 7 the appropriate quantities from SumApprox:

- Let  $X_i$  be the  $(i+1)^{\text{th}}$  sample (i.e.  $(X_i, \gamma_{X_i})$  is the pair  $(u, \gamma_u)$  drawn at the  $(i+1)^{\text{th}}$  invocation of line 8).
- Let  $\mathcal{F}_i$  be the event space generated by  $X_0, \dots, X_i$ , so that for any random variable  $Y$ , with  $\mathbb{E}[Y | \mathcal{F}_i]$  we mean  $\mathbb{E}[Y | X_0, \dots, X_i]$  and with  $\text{Var}[Y | \mathcal{F}_i]$  we mean  $\text{Var}[Y | X_0, \dots, X_i]$ .
- Let  $\chi_i = \mathbb{1}[X_i \in \bigcup_{j=0}^{i-1} \{X_j\}]$  be the indicator variable of a repeat on the  $(i+1)^{\text{th}}$  sample.
- Let  $P_i = \sum_{u \in \bigcup_{j=0}^{i-1} \{X_j\}} \frac{\gamma_u}{\gamma}$  be the probability of a repeat on the  $(i+1)^{\text{th}}$  sample as a function of all the (distinct) samples up to the  $i^{\text{th}}$ , i.e.  $P_i = \mathbb{E}[\chi_i | \mathcal{F}_{i-1}] \leq 1$ .
- Let  $Z_i = \sum_{j=0}^i (\chi_j - P_j)$ ; it is easy to see that  $(Z_i)_{i \geq 0}$  is a martingale with respect to the filter  $(\mathcal{F}_i)_{i \geq 0}$ , since  $Z_i$  is obtained by adding to  $Z_{i-1}$  the indicator variable  $\chi_i$  and subtracting  $P_i$  i.e. its expectation in  $\mathcal{F}_{i-1}$ . More formally,  $\mathbb{E}[Z_i | \mathcal{F}_{i-1}] = \mathbb{E}[Z_{i-1} +$

$\chi_i - P_i | \mathcal{F}_{i-1}]$ , and since  $Z_{i-1}$  and  $P_i$  are completely determined by  $X_0, \dots, X_{i-1}$ , the right-hand term is simply  $Z_{i-1} + (\mathbb{E}[\chi_i | \mathcal{F}_{i-1}] - P_i) = Z_{i-1}$ . Note also that  $Z_0 = 0$ .

■ Let  $M = 1$ , noting that  $|Z_{i+1} - Z_i| = |\chi_{i+1} - P_{i+1}| \leq 1$  for all  $i$ .

Finally, note that  $\text{Var}(Z_j | \mathcal{F}_{j-1}) = \text{Var}(\chi_j | \mathcal{F}_{j-1})$  (as  $Z_j = Z_{j-1} + \chi_j - P_j$  and, again,  $Z_{j-1}$  and  $P_j$  are completely determined by  $X_0, \dots, X_{j-1}$ ). Since  $\text{Var}(\chi_j | \mathcal{F}_{j-1}) = P_j(1 - P_j) \leq P_j$ , we have  $V_i = \sum_{j=1}^i \text{Var}(Z_j | \mathcal{F}_{j-1}) \leq \sum_{j=1}^i P_j$ . Theorem 7 then yields the following:

► **Lemma 8.** *For all  $z, v > 0$  we have*

$$\Pr[Z_i \geq z, \sum_{j=1}^i P_j \leq v \text{ for some } i] \leq \exp\left[-\frac{z^2}{2(v+z)}\right] \quad (4)$$

Recall now SumApprox. Note that  $\sum_{j=1}^i P_j$  and  $Z_i$  are respectively the value of  $\frac{w}{\gamma}$  and of  $r - \frac{w}{\gamma}$  just after the *while* loop has been executed for the  $(i+1)$ -th time. Note also that, when SumApprox returns,  $r = k_{\epsilon, \delta}$ . Therefore the event that, when SumApprox returns,  $\frac{w}{r} \leq \gamma(1 - \epsilon)$  i.e.  $\frac{w}{\gamma} \leq r(1 - \epsilon) \leq (1 - \epsilon)k_{\epsilon, \delta}$  corresponds to the event that  $Z_i \geq \epsilon r = \epsilon k_{\epsilon, \delta}$  and  $\sum_{j=1}^i P_j \leq (1 - \epsilon)k_{\epsilon, \delta}$ . Invoking Lemma 8 with  $z = \epsilon k_{\epsilon, \delta}$  and  $v = (1 - \epsilon)k_{\epsilon, \delta}$ :

$$\Pr\left[\frac{w}{r} \leq \gamma(1 - \epsilon)\right] \leq \exp\left[-\frac{\epsilon^2 k_{\epsilon, \delta}^2}{2(\epsilon k_{\epsilon, \delta} + (1 - \epsilon)k_{\epsilon, \delta})}\right] = \exp\left[-\frac{\epsilon^2 k_{\epsilon, \delta}}{2}\right] \quad (5)$$

which is smaller than  $\delta/3$  since clearly  $k_{\epsilon, \delta} > \frac{2}{\epsilon^2} \ln \frac{3}{\delta}$ . Consider instead the event that, when SumApprox returns,  $\frac{w}{r} \geq \gamma(1 + \epsilon)$  i.e.  $\frac{w}{\gamma} \geq r(1 + \epsilon) = k_{\epsilon, \delta}(1 + \epsilon)$ . This is the event that  $Z_i \leq -\epsilon k_{\epsilon, \delta}$ , or equivalently  $-Z_i \geq \epsilon k_{\epsilon, \delta}$ . Note that Lemma 8 still holds if we replace  $Z_i$  with  $-Z_i$ , as  $(-Z_i)_{i \geq 0}$  too is obviously a martingale with respect to the filter  $(\mathcal{F}_i)_{i \geq 0}$ , with  $-Z_0 = 0$ . Let then  $i_0 \leq i$  be the smallest time such that  $-Z_{i_0} \geq \epsilon k_{\epsilon, \delta}$ . Since  $|Z_j - Z_{j-1}| \leq 1$ , it must be  $-Z_{i_0} < \epsilon k_{\epsilon, \delta} + 1$ . Also, since  $\sum_{j=0}^i \chi_j$  is nondecreasing with  $i$ , then  $\sum_{j=0}^{i_0} \chi_j \leq k_{\epsilon, \delta}$ . It follows that  $\sum_{j=1}^{i_0} P_j = -Z_{i_0} + \sum_{j=0}^{i_0} \chi_j \leq \epsilon k_{\epsilon, \delta} + 1 + k_{\epsilon, \delta} = (1 + \epsilon)k_{\epsilon, \delta} + 1$ . Invoking again Lemma 8 with  $z = \epsilon k_{\epsilon, \delta}$  and  $v = (1 + \epsilon)k_{\epsilon, \delta} + 1$ , we obtain:

$$\Pr\left[\frac{w}{r} \geq \gamma(1 + \epsilon)\right] \leq \exp\left[-\frac{\epsilon^2 k_{\epsilon, \delta}^2}{2((1 + 2\epsilon)k_{\epsilon, \delta} + 1)}\right] \quad (6)$$

Note that  $\frac{1}{k_{\epsilon, \delta}} < \frac{\epsilon^2}{2+4.4\epsilon} < 0.2\epsilon$  since  $\epsilon \leq 1$ ; so  $2((1 + 2\epsilon) + \frac{1}{k_{\epsilon, \delta}}) < 2 + 4.4\epsilon$ , and since  $k_{\epsilon, \delta} \geq \frac{2+4.4\epsilon}{\epsilon^2} \ln \frac{3}{\delta}$  the right-hand term is at most  $\frac{\delta}{3}$ . Finally, by a simple union bound the probability that  $|\hat{\gamma} - \gamma| \geq \epsilon\gamma$  is at most  $2\frac{\delta}{3}$ , and the proof of Theorem 6 is complete. ◀

► **Theorem 9.** *SumApprox( $\epsilon, \delta$ ) draws at most  $\lceil 45\|\pi\|^{-1}\epsilon^{-3}(\ln \frac{3}{\delta})^{3/2} \rceil$  samples with probability at least  $1 - \frac{\delta}{3}$ .*

**Proof.** We show that the probability that  $s = \lceil 45\|\pi\|^{-1}\epsilon^{-3}(\ln \frac{3}{\delta})^{3/2} \rceil$  draws yield less than  $k_{\epsilon, \delta}$  repeats is less than  $\frac{\delta}{3}$ . Let  $\bar{p} = \frac{5}{18}\|\pi\|\epsilon(\ln \frac{3}{\delta})^{-1/2}$ . We consider two cases.

**Case 1:**  $\exists u \in V$  with  $\pi(u) > \bar{p}$ . Let then  $C_u^s$  be the random variable counting the number of times  $u$  appears in  $s$  draws. Since if  $C_u^s > k_{\epsilon, \delta}$  then  $u$  causes at least  $k_{\epsilon, \delta}$  repeats, the probability that SumApprox needs more than  $s$  draws is upper bounded by  $\Pr[C_u^s \leq k_{\epsilon, \delta}]$ . Now  $\mathbb{E}[C_u^s] = s\pi(u) > s\bar{p} > 45\frac{5}{18}\frac{1}{\epsilon^2} \ln \frac{3}{\delta} = \frac{12.5}{\epsilon^2} \ln \frac{3}{\delta} \geq 1.7(\frac{6.4}{\epsilon^2} \ln \frac{3}{\delta} + 1) \geq 1.7\lceil \frac{2+4.4\epsilon}{\epsilon^2} \ln \frac{3}{\delta} \rceil = 1.7k_{\epsilon, \delta}$ , therefore  $C_u^s \leq k_{\epsilon, \delta}$  implies  $C_u^s < \frac{1}{1.7}\mathbb{E}[C_u^s] < (1 - 0.41)\mathbb{E}[C_u^s]$ . Since  $C_u^s$  is a sum of independent binary random variables, the bounds of Appendix 5.1 give  $\Pr[C_u^s \leq k_{\epsilon, \delta}] < \exp(-\frac{1}{2}0.41^2\mathbb{E}[C_u^s]) < \exp(-0.5 \cdot 0.41^2 \cdot \frac{12.5}{\epsilon^2} \ln \frac{3}{\delta}) < \exp(-1.05 \ln \frac{3}{\delta}) < \frac{\delta}{3}$ .



**Case 2:**  $\pi(u) \leq \bar{p}$  for all  $u \in V$ . Let then  $\bar{s} = \lceil \bar{p}^{-1} \rceil$ , let  $\bar{S}$  be the set of distinct elements in the first  $\bar{s}$  draws, and let  $w(\bar{s}) = \sum_{u \in \bar{S}} \pi(u)$ . First we show that  $\mathbb{E}[w(\bar{s})] \geq \frac{4}{9} \bar{s} \|\pi\|^2$ . Write  $\mathbb{E}[w(\bar{s})] = \sum_{u \in V} \pi(u) (1 - (1 - \pi(u))^{\bar{s}})$ . Since for all  $x \in [0, 1]$  and  $k \geq 1$  it holds  $(1 - x)^k \leq (1 + kx)^{-1}$ , by setting  $x = \pi(u)$  and  $k = \bar{s}$  we obtain  $1 - (1 - \pi(u))^{\bar{s}} \geq 1 - (1 + \bar{s}\pi(u))^{-1} = \bar{s}\pi(u)(1 + \bar{s}\pi(u))^{-1}$ . Moreover note that  $\bar{p}^{-1} \geq \frac{18}{5} = 3.6$  and thus  $\lceil \bar{p}^{-1} \rceil \leq \frac{5}{4} \bar{p}^{-1}$ . Therefore  $\bar{s}\pi(u) \leq \lceil \bar{p}^{-1} \rceil \bar{p} \leq \frac{5}{4}$  for all  $u$ , and thus  $\bar{s}\pi(u)(1 + \bar{s}\pi(u))^{-1} \geq \bar{s}\pi(u) \frac{1}{1 + \frac{5}{4}} = \frac{4}{9} \bar{s}\pi(u)$ . Therefore  $\mathbb{E}[w(\bar{s})] \geq \frac{4}{9} \bar{s} \sum_{u \in V} \pi(u)^2 = \frac{4}{9} \bar{s} \|\pi\|^2$ . Now we consider two cases. First, suppose the event  $w(\bar{s}) \geq 0.4 \mathbb{E}[w(\bar{s})]$  takes place. For  $i = \bar{s} + 1, \dots, s$  let  $\chi_i$  be the indicator random variable of the event that the  $i$ -th draw is an element of  $\bar{S}$ , and let  $C_s = \sum_{i=\bar{s}+1}^s \chi_i$ . Clearly SumApprox witnesses at least  $C_s$  repeats in the last  $s - \bar{s}$  draws, and thus overall. We shall then bound  $\Pr[C_s < k_{\epsilon, \delta}]$ . First, since by hypothesis the total mass of  $\bar{S}$  is  $w(\bar{s}) \geq 0.4 \mathbb{E}[w(\bar{s})]$ , we also have  $\mathbb{E}[\chi_i] \geq 0.4 \mathbb{E}[w(\bar{s})] \geq \frac{1}{9} \bar{s} \|\pi\|^2$ . Therefore  $\mathbb{E}[C_s] = \sum_{i=\bar{s}+1}^s \mathbb{E}[\chi_i] \geq \frac{1}{9} \bar{s} (s - \bar{s}) \|\pi\|^2$ . Now note that  $s - \bar{s} > 10\bar{s}$ , therefore  $\mathbb{E}[C_s] \geq \frac{16}{9} \bar{s}^2 \|\pi\|^2$ . Finally, since  $\bar{s} = \lceil \bar{p}^{-1} \rceil \geq \frac{18}{5} \|\pi\|^{-1} \epsilon^{-1} (\ln \frac{3}{\delta})^{1/2}$ , it holds  $\mathbb{E}[C_s] \geq (\frac{18}{5})^2 \frac{16}{9} \frac{1}{\epsilon^2} \ln \frac{3}{\delta} > 23 \frac{1}{\epsilon^2} \ln \frac{3}{\delta} > 3.14 k_{\epsilon, \delta}$ . It follows that the event  $C_s < k_{\epsilon, \delta}$  implies  $C_s < \frac{1}{3.14} \mathbb{E}[C_s] < (1 - 0.68) \mathbb{E}[C_s]$ . By the concentration bounds of Appendix 5.1, the probability of the latter is  $\Pr[C_s < k_{\epsilon, \delta}] \leq \exp(-\frac{1}{2} 0.68^2 \mathbb{E}[C_s]) < \exp(-\frac{1}{2} 0.68^2 23 \frac{1}{\epsilon^2} \ln \frac{3}{\delta}) < \exp(-5 \ln \frac{3}{\delta}) < \frac{\delta}{243}$ . The second case corresponds to the event  $w(\bar{s}) < 0.4 \mathbb{E}[w(\bar{s})] = (1 - 0.6) \mathbb{E}[w(\bar{s})]$ , of which we shall bound the probability. Let  $\chi_u^{\bar{s}}$  be the indicator variable of the event  $u \in \bar{S}$ , so  $w(\bar{s}) = \sum_{u \in V} \chi_u^{\bar{s}} \pi(u)$ . Since  $\pi(u) \leq \bar{p}$  for all  $u$ , we can write  $w(\bar{s}) = \bar{p} \sum_{u \in V} \chi_u^{\bar{s}} \bar{p}^{-1} \pi(u)$  so that the coefficients  $\bar{p}^{-1} \pi(u)$  are in  $[0, 1]$ . Clearly, the  $\chi_u^{\bar{s}}$  are non-positively correlated. We can thus apply the bounds of Appendix 5.1 and get  $\Pr[w(\bar{s}) < 0.4 \mathbb{E}[w(\bar{s})]] \leq \exp(-0.5 \cdot 0.6^2 \bar{p}^{-1} \mathbb{E}[w(\bar{s})])$ . By replacing the definitions and bounds for  $\mathbb{E}[w(\bar{s})]$ ,  $\bar{s}$  and  $\bar{p}^{-1}$  from above, we get  $\Pr[w(\bar{s}) < 0.4 \mathbb{E}[w(\bar{s})]] < \exp(-2.88 \ln \frac{3}{\delta}) < \frac{\delta}{23}$ . Again by a union bound, the probability that SumApprox draws more than  $s$  samples is less than  $\frac{\delta}{243} + \frac{\delta}{23} < \frac{\delta}{3}$ .  $\blacktriangleleft$

We remark that the previous existing algorithm for the sum estimation problem [17] needs knowledge of  $n = |V|$  and uses  $O(\sqrt{n} \epsilon^{-7/2} \log(n) (\log \frac{1}{\delta} + \log \frac{1}{\epsilon} + \log \log n))$  samples. SumApprox is simpler, oblivious to  $n$ , and gives more general bounds. It is also asymptotically faster unless  $\pi$  is (essentially) the uniform distribution.

Finally, we show that SumApprox is essentially optimal, by proving  $\Omega(\|\pi\|^{-1})$  samples are in general necessary to estimate  $\gamma$  even if  $n$  is known in advance. This extends to arbitrary distributions the  $\Omega(\sqrt{n})$  lower bound given by [17] for the uniform distribution.

**► Theorem 10.** *For any function  $\nu(n) \in \Omega(n^{-\frac{1}{2}}) \cap O(1)$  there exist vectors  $\mathbf{x} = \gamma\pi = (\gamma_1, \dots, \gamma_n)$  with  $\|\pi\| = \Theta(\nu(n))$  such that  $\Omega(\|\pi\|^{-1})$  samples are necessary to estimate  $\gamma$  within constant multiplicative factors with constant probability, even if  $n$  is known.*

**Proof.** Let  $k \in \Theta(\nu(n)^{-2})$  with  $1 \leq k \leq \frac{n}{2}$ . Consider the two vectors  $\mathbf{x} = (\gamma_1, \dots, \gamma_n)$  and  $\mathbf{x}' = (\gamma'_1, \dots, \gamma'_n)$  defined as follows:

$$\begin{aligned} \gamma_j &= 1 : j \leq k, & \gamma_j &= \sqrt{k}/n : j > k \\ \gamma'_j &= 1 : j \leq 2k & \gamma'_j &= \sqrt{k}/n : j > 2k \end{aligned}$$

Now let  $\gamma = \sum_{i=1}^n \gamma_i$  and  $\gamma' = \sum_{i=1}^n \gamma'_i$ . One can check that  $\gamma \leq 2k$  and  $|\gamma - \gamma'| \geq \frac{k}{2}$ . Hence, to obtain an estimate  $\hat{\gamma}$  of  $\gamma$  with  $\hat{\gamma} \leq \frac{5}{4} \gamma$ , one must distinguish  $\mathbf{x}$  from  $\mathbf{x}'$ . Note that the norm of  $\pi = \mathbf{x}/\gamma$  is in  $\Theta(1/\sqrt{k}) = \Theta(\nu(n))$ , as requested. Now, for each one of  $\mathbf{x}$  and  $\mathbf{x}'$  in turn, pick a permutation of  $\{1, \dots, n\}$  uniformly at random and apply it to the entries of the vector. Suppose then we sample  $o(\|\pi\|^{-1}) = o(\sqrt{k})$  entries from  $\mathbf{x}$ . We shall see that, with probability  $1 - o(1)$ , we cannot distinguish  $\mathbf{x}$  from  $\mathbf{x}'$ . First, note that the total mass



of the entries with value  $\sqrt{k}/n$  is at most  $1/\sqrt{k}$ . Hence the probability of drawing *any* of those entries with  $o(\sqrt{k})$  samples is  $o(1)$ , and we can assume all draws yield entries having value 1. Since there are  $O(k)$  such entries in total, the probability of witnessing any repeat is also  $o(1)$ , and we can assume no repeat is witnessed. Furthermore, because of the random permutation, the indices of samples are distributed uniformly over  $\{1, \dots, n\}$  (recall that we actually sample from the index set  $\{1, \dots, n\}$ , so we could use the distribution of the indices to distinguish  $\mathbf{x}$  from  $\mathbf{x}'$ ). The same argument applies to  $\mathbf{x}'$ , so drawing  $o(\sqrt{k})$  samples from  $\mathbf{x}'$  yields exactly the same distribution and the two vectors are indistinguishable. To adapt the construction to larger approximation factors, set  $\gamma'_j = 1 : j \leq \eta k$  for  $\eta$  large enough. ◀

### 3 Approximating the stationary distribution

In this section we address the problem of approximating  $\pi(v)$ . Such a problem can in fact be reduced to the sum estimation problem of Section 2 by drawing states via random walks. The crux is determining how long the walks must be in order for the samples to come from a distribution close enough to  $\pi$ , so that the approximation guarantees of SumApprox transfer directly to our estimate of  $\pi(v)$ .

Consider a random walk of length  $t+1$  that starts at  $v$ . Obviously we can simulate such a walk by setting  $u_0 = v$  and then invoking  $step(u_i)$  to obtain the state  $u_{i+1}$ , for  $i = 0, \dots, t-1$ . Crucially, using the time-reversibility of the chain, for any visited state  $u$  we can obtain the ratio  $\gamma_u$  between  $\pi(u)$  and  $\pi(v)$  using  $O(1)$  operations. Formally, let  $\gamma_v = \pi(v)/\pi(v) = 1$ , and in general let  $\gamma_u = \pi(u)/\pi(v)$ . Note that:

$$\gamma_{u_{i+1}} = \frac{\pi(u_{i+1})}{\pi(v)} = \frac{\pi(u_{i+1})}{\pi(u_i)} \cdot \frac{\pi(u_i)}{\pi(v)} = \frac{\pi(u_{i+1})}{\pi(u_i)} \cdot \gamma_{u_i} \quad (7)$$

The time-reversibility of the chain (see Equation 2) implies  $\frac{\pi(u_{i+1})}{\pi(u_i)} = \frac{p^{u_i, u_{i+1}}}{p^{u_{i+1}, u_i}} = \frac{probe(u_i, u_{i+1})}{probe(u_{i+1}, u_i)}$ , hence we can compute  $\gamma_{u_{i+1}}$  with  $O(1)$  operations if we know  $\gamma_{u_i}$ . But then we can keep track of  $\gamma_u$  for any  $u$  visited so far, starting with  $\gamma_{u_0} = 1$  and computing  $\gamma_{u_{i+1}}$  by Equation 7 the first time  $u_{i+1}$  is visited.

Suppose now to pick  $t$  large enough so that the chain reaches its stationary distribution, i.e.  $u_t \sim \pi$  irrespective of  $v$ . One is then drawing state  $u$ , as well as its associate weight  $\gamma_u$ , with probability  $\pi(u)$ . Now if we let  $\gamma = \sum_{u \in V} \gamma_u$ , then  $\pi(u) = \gamma_u \gamma^{-1}$  and in particular  $\pi(v) = \gamma^{-1}$ . Therefore approximating  $\pi(v)$  amounts to approximating  $\gamma$ ; more formally, for any  $\epsilon \in (0, 1)$ , if  $\hat{\gamma}$  is a  $(1 \pm \frac{\epsilon}{2})$ -approximation of  $\gamma$  then  $\hat{\gamma}^{-1}$  is a  $(1 \pm \epsilon)$ -approximation of  $\pi(v)$ . We can therefore reduce to the sum approximation problem of Section 3: compute with probability  $(1 - \delta)$  a  $(1 \pm \epsilon)$ -approximation of  $\gamma$ , assuming we can draw pairs  $(u, \gamma_u)$  according to  $\pi$ . The only problem is that by simulating the chain we can only come close to (but not exactly on) the stationary distribution  $\pi$ . We must then tie the approximation guarantees of SumApprox to the length  $t$  of the random walks, or better to the distance  $\|\pi' - \pi\|_{TV}$  between  $\pi$  and the distribution  $\pi'$  from which  $u_t$  is drawn. Formally, we show:

► **Lemma 11.** *There exists some constant  $c > 0$  such that the following holds. Choose any  $\delta, \epsilon \in (0, 1)$ , and suppose we draw the pairs  $(u, \gamma_u)$  from a distribution  $\pi'$  such that  $\|\pi - \pi'\|_{TV} \leq \left(\frac{\epsilon \|\pi\|}{\ln(3/\delta)}\right)^c$ . Then SumApprox( $\frac{\epsilon}{2}, \delta$ ) with probability at least  $1 - \delta$  returns a multiplicative  $(1 \pm \epsilon)$ -approximation of  $\gamma$  by taking at most  $\lceil 720 \|\pi\|^{-1} \epsilon^{-3} (\ln \frac{3}{\delta})^{3/2} \rceil$  samples.*

**Proof.** Let us start with the bound on the number of samples. Recall the proof of Theorem 9, and note that the whole argument depends on  $\pi$  but not on the values  $\gamma_u$ . Indeed,  $\pi$  alone determines the probability of repeats and thus controls the distribution of the number of

## 18:10 On Approximating the Stationary Distribution of Time-reversible Markov Chains

samples drawn by SumApprox. Hence, by Theorem 9 SumApprox( $\frac{\epsilon}{2}, \delta$ ) takes more than  $\lceil 45\|\pi'\|^{-1}8\epsilon^{-3}(\ln\frac{3}{\delta})^{3/2} \rceil = \lceil 360\|\pi'\|^{-1}\epsilon^{-3}(\ln\frac{3}{\delta})^{3/2} \rceil$  samples with probability less than  $\frac{\delta}{3}$ . Now  $\|\pi - \pi'\| \leq 2\|\pi - \pi'\|_{\text{TV}} \leq 2\left(\frac{\epsilon\|\pi\|}{\ln(3/\delta)}\right)^c \leq \|\pi\|2(\ln 3)^{-c}$ , which for  $c \geq 15$  is bounded by  $\frac{1}{2}\|\pi\|$ . Then, since  $\|\pi'\| \geq \|\pi\| - \|\pi - \pi'\|$ , we have  $\|\pi'\|^{-1} \leq 2\|\pi\|^{-1}$  and the bound above is in turn bounded by  $\lceil 720\|\pi\|^{-1}\epsilon^{-3}(\ln\frac{3}{\delta})^{3/2} \rceil$ .

Let us now see the approximation guarantees. Recall the proof of Theorem 6. We want to show again that  $\Pr[\frac{w(s)}{r} - \gamma \geq \frac{\epsilon}{2}\gamma] \leq \frac{2\delta}{3}$ . However, now the samples are drawn according to  $\pi'$  instead of  $\pi$ . Let then  $P'_j = \sum_{u \in \cup_{h=0}^{j-1} \{X_h\}} \pi'(u)$  and  $Z'_i = \sum_{j=0}^i (\chi_j - P'_j)$ ; in a nutshell,  $P'_j$  and  $Z'_i$  are the analogous of  $P_j$  and  $Z_i$  under  $\pi'$ . It is immediate to check that Lemma 8 holds with  $Z'_i$  and  $P'_j$  in place of  $Z_i$  and  $P_j$ . Let now  $w'(i) = \gamma \sum_{j=1}^i P'_j$ . Note that  $\sum_{j=1}^i P'_j$  and  $Z'_i$  are respectively the value of  $\frac{w'(i)}{\gamma}$  and of  $r - \frac{w'(i)}{\gamma}$  just after line 9 has been executed for the  $(i+1)$ -th time. Therefore, the argument following Lemma 8 holds if we put  $w'(i)$  in place of the value taken by  $w$  after the  $(i+1)$ -th execution of line 9. Hence the same bounds hold, and SumApprox( $\frac{\epsilon}{2}, \delta$ ) ensures  $\Pr[\frac{w'(s)}{r} - \gamma \geq \frac{\epsilon}{2}\gamma] \leq \frac{\delta}{3}$  where  $s$  is the total number of draws. Now note that SumApprox does not return  $\frac{w'(s)}{r}$ , but  $\frac{w(s)}{r}$  where  $w(i) = \gamma \sum_{j=1}^i P_j$  is the value of  $w$  in SumApprox after line 9 has been executed for the  $(i+1)$ -th time. We shall now make  $|\frac{w(s)}{r} - \frac{w'(s)}{r}| \leq \frac{\epsilon}{2}\gamma$ ; by the triangle inequality we will then be done. First of all, by the definition of  $w(s)$  and  $w'(s)$  we have

$$\left| \frac{w(s)}{r} - \frac{w'(s)}{r} \right| = \gamma r^{-1} \left| \sum_{j=1}^s P_j - \sum_{j=1}^s P'_j \right| \leq \gamma r^{-1} \sum_{j=1}^s |P_j - P'_j| \quad (8)$$

Now note that  $|P_j - P'_j| \leq \|\pi - \pi'\|_{\text{TV}}$ , since  $P_j$  and  $P'_j$  are the probability of the same event under respectively  $\pi$  and  $\pi'$ . Therefore the right-hand side of Equation 8 is bounded by  $\gamma r^{-1} s \|\pi - \pi'\|_{\text{TV}}$ . Now, when SumApprox( $\frac{\epsilon}{2}, \delta$ ) terminates  $r = k_{\frac{\epsilon}{2}, \delta} \geq 4\frac{2+2.2\epsilon}{\epsilon^2} \ln\frac{3}{\delta}$ , and by hypothesis  $\|\pi - \pi'\|_{\text{TV}} \leq \left(\frac{\epsilon\|\pi\|}{\ln(3/\delta)}\right)^c$ . Therefore:

$$\left| \frac{w(s)}{r} - \frac{w'(s)}{r} \right| \leq \gamma s \frac{\epsilon^2}{4(2+2.2\epsilon)\ln(\frac{3}{\delta})} \left(\frac{\epsilon\|\pi\|}{\ln(\frac{3}{\delta})}\right)^c \leq \gamma s \|\pi\| \frac{\epsilon^{3+c}}{8\ln(\frac{3}{\delta})^{1+c}} \quad (9)$$

Finally, recall from above that with probability  $1 - \frac{\delta}{3}$  we have  $s \leq \lceil 720\|\pi\|^{-1}\epsilon^{-3}(\ln\frac{3}{\delta})^{3/2} \rceil$ . In this case the equation above yields  $|\frac{w(s)}{r} - \frac{w'(s)}{r}| \leq \gamma \cdot 721\epsilon^c \ln(\frac{3}{\delta})^{0.5-c}$ , which is smaller than  $\frac{\epsilon}{2}\gamma$  for  $c \geq \frac{1}{2} + \frac{\ln 1442}{\ln \ln 3} \approx 78$ . A simple union bound completes the proof.  $\blacktriangleleft$

We are now ready to prove Theorem 2. Pick  $t = \tau c \ln(\|\pi\|^{-1}\epsilon^{-1} \ln\frac{3}{\delta}) / \ln 2$ , where  $c$  is the constant of Lemma 11 and  $\tau$  is the mixing time of the chain. Simulate the walk for  $t$  steps starting from  $v$ , and let  $\pi'$  be the distribution of the final state. By the properties of the mixing time (see Section 1.1) it holds  $\|\pi - \pi'\|_{\text{TV}} \leq 2^{-c \ln(\|\pi\|^{-1}\epsilon^{-1} \ln\frac{3}{\delta}) / \ln 2} \leq \left(\frac{\epsilon\|\pi\|}{\ln(3/\delta)}\right)^c$  and therefore by Lemma 11 we obtain a  $(1 \pm \epsilon)$  approximation of  $\gamma$ . By choosing  $\epsilon$  small enough we can obtain a  $(1 \pm \epsilon')$  approximation of  $\pi(v)$  for any desired  $\epsilon'$ . The total number of operations performed is clearly bounded by  $t = \tau c \ln(\|\pi\|^{-1}\epsilon^{-1} \ln\frac{3}{\delta}) / \ln 2$  times the number of samples taken by SumApprox, and by substituting this value in the bound of Theorem 1 we obtain Theorem 2.

### 3.1 Reducing the footprint

In this section we describe FullMassApprox, the algorithm behind the bounds of Theorem 5. FullMassApprox is derived from MassApprox as follows. First, instead of performing a new walk of length  $t$  from  $v$  for each sample, the algorithm performs one long random walk of

length  $T$  and takes one sample every  $t$  steps. The correctness guarantees do not change, since although the samples do not come all from the same distribution, they are still drawn from a distribution sufficiently close to  $\pi$ . Second, after checking if the current draw yields a repeat, the algorithm includes in the set  $S$  not only the draw but also all other states visited so far. Again, this does not affect the guarantees, since we do not need the set  $S$  to be built on independent samples. However, this makes the mass of  $S$  grow potentially faster, so we can hope to get more repeats and decrease the total number of samples.

**The concentration hypothesis.** Before continuing to the proof of Theorem 5, let us provide some intuition behind the concentration hypothesis. Suppose the walk runs for  $T = k\bar{\tau}$  steps for some  $\bar{\tau} = \tau \text{poly}(\log(\|\pi\|^{-1}))$ . Such a process can be seen as a coupon collector over  $k$  rounds, where a subset of at most  $\bar{\tau}$  states is collected (i.e. visited) at each round. Now, if we pick  $\bar{\tau}' \leq \bar{\tau}$  with  $\bar{\tau}' = \tau \text{poly}(\log(\|\pi\|^{-1}))$ , then in each round the  $\bar{\tau} - \bar{\tau}'$  central steps are essentially independent of other rounds (more formally, the correlation is  $O(\text{poly}(n)^{-1})$ ). Each round is then in large part independent of the others; the issue is that the states visited *within* a single round are correlated. Such a correlation is responsible for the factor  $\bar{\tau}$  in the concentration hypothesis and amounts for the (intuitive) fact that conditioning on the outcome of one step of the walk does not affect the distribution of those steps that are more than  $\bar{\tau}$  steps away. We note that the concentration bounds of [8] give  $\Pr[\sum_{i=1}^T f_i \notin (1 \pm \bar{\epsilon})\mathbb{E}[\sum_{i=1}^T f_i]] < 2 \exp -\Omega(\bar{\epsilon}^2 \mathbb{E}[\sum_{i=1}^T f_i]/\tau)$  where  $f_i \in [0, 1]$  is a function of state  $X_i$ ; however we could not use them to prove the concentration hypothesis of Theorem 4. Let us now delve into the proof.

**Proof.** Observe the random walk performed by FullMassApprox. Clearly if  $\bar{\tau} = \Omega(n)$  then the walk visits  $O(\tau \ln n + \sqrt{\bar{\tau}n})$  distinct states, and the theorem holds unconditionally. Let us then assume  $\bar{\tau} = o(n)$ . We disregard the first  $T_0 = \Theta(\tau \ln n)$  steps of the walk, which of course yield at most  $T_0$  distinct states, and focus on the last  $T$  steps, which we denote by  $X_1, \dots, X_T$  (one may thus plug  $T + T_0$  in place of  $T$  in the concentration hypothesis). Let  $\pi_i$  denote the distribution of state  $X_i$ ,  $i = 1, \dots, T$ . Since  $T_0 = \Theta(\tau \ln n)$ , then we can make  $\|\pi_i - \pi\|_{\text{TV}} \leq \frac{1}{\text{poly}(n)}$ . One can adapt the proof of Lemma 11 to FullMassApprox, using the hypothesis  $\|\pi - \pi_i\|_{\text{TV}} \leq (\frac{\epsilon \|\pi\|}{\ln(3/\delta)})^c$  for all  $i \geq 1$ . This changes the bounds of the lemma only by constant multiplicative factors. We can thus focus on proving the bound on the number of states visited by the walk. In the analysis we assume  $X_i \sim \pi$ , but again the same asymptotic bounds hold if  $\|\pi_i - \pi\|_{\text{TV}} \leq \frac{1}{\text{poly}(n)}$ . Let  $S_{v,t} = \cup_{i=1}^t \{X_i\}$ , let  $N_{v,t} = |S_{v,t}|$ , and let  $M_{v,t} = \sum_{u \in S_{v,t}} \pi(u)$ . For brevity we simply write  $S_t, N_t, M_t$ .

The crux is to show that  $M_t$ , the aggregate mass of  $S_t$ , grows basically as  $N_t^2/t$ . Formally we prove that, for any  $\epsilon, \delta, q > 0$ , if  $\Pr[N_t \geq q] \geq 1 - \delta$  then  $\Pr[M_t \geq q^2 \frac{\epsilon}{4tn}] \geq 1 - \epsilon - \delta$ . First, for any  $\lambda > 0$  let  $V_\lambda = \{u \in V : \pi(u) < \frac{\lambda}{n}\}$ . Clearly  $\Pr[X_i \in V_\lambda] = \sum_{u \in V_\lambda} \pi(u) < \lambda$ . Therefore the number of steps  $J_t(\lambda)$  the chain was on a state of  $V_\lambda$  satisfies  $\mathbb{E}[J_t(\lambda)] < t\lambda$ . Now, by Markov's inequality  $\Pr[J_t(\lambda) > \frac{q}{2}] < \frac{2t\lambda}{q}$ , and setting  $\lambda = \epsilon \frac{q}{2}t$  we obtain  $\Pr[J_t(\lambda) > \frac{q}{2}] < \epsilon$ . Since by hypothesis  $\Pr[N_t < q] < \delta$ , by a union bound we get  $\Pr[N_t \geq q, J_t(\lambda) < \frac{q}{2}] \geq 1 - \delta - \epsilon$ . But if  $N_t \geq q$  and  $J_t(\lambda) < \frac{q}{2}$  then  $S_t$  contains at least  $\frac{q}{2}$  distinct states with individual mass at least  $\frac{\epsilon q}{2tn}$ , and thus  $M_t \geq \frac{q}{2} \frac{\epsilon q}{2tn} = q^2 \frac{\epsilon}{4tn}$ .

Now choose  $t$  such that  $\mathbb{E}[N_t] = \Omega(\sqrt{n\bar{\tau}})$ ; note that  $\mathbb{E}[N_t] = \Omega(\bar{\tau})$  since  $\bar{\tau} = o(n)$ . By plugging  $\mathbb{E}[N_t]$  into the concentration bound for  $N_t$  we can then make  $\Pr[N_t < (1 - \bar{\epsilon})\mathbb{E}[N_t]]$  arbitrarily small for any  $\bar{\epsilon} > 0$ . Let then  $q = (1 - \bar{\epsilon})\mathbb{E}[N_t]$ . By the bounds of the previous paragraph, for any  $\delta > 0$  with probability  $1 - \delta - \bar{\epsilon}$  we have  $M_t \geq q^2 \frac{\bar{\epsilon}}{4tn} = \Omega(n\bar{\tau}) \frac{\bar{\epsilon}}{4tn} = \Omega(\frac{\bar{\tau}}{t})$ . Conditioned on the event that  $M_t = \Omega(\frac{\bar{\tau}}{t})$ , any sample drawn after  $t$  steps is a repeat with

probability  $\Omega(\frac{\tau}{t})$ . If we then draw  $\Theta(\frac{t}{\tau})$  samples, which require  $\Theta(t)$  steps, we witness an expected  $\Omega(1)$  samples, which can be made larger than  $k_{\epsilon,\delta}$  by appropriately increasing  $t$ . Again by the concentration bounds on  $N_t$ , the total number of states visited can be made  $O(2\mathbb{E}[Q_t]) = O(\sqrt{n\tau}) = \tilde{O}(\sqrt{n\tau})$  with probability arbitrarily close to 1 by appropriately increasing  $t$ . ◀

## 4 Conclusions

We have given improved, optimal algorithms for approximating the stationary probability of a given state in a time-reversible Markov chain, and for approximating the sum of nonnegative real vectors by weighted sampling. Although time-reversible chains are of clear relevance, extending our results to other classes of Markov chains is an intriguing open question. We have also shown that the footprint of our algorithms in terms of number of distinct states visited is tied to the concentration of the number of distinct states visited by the chain; investigating such a concentration is thus an obvious line of future research.

---

## References

- 1 Noga Alon, Ori Gurel-Gurevich, and Eyal Lubetzky. Choice-memory tradeoff in allocations. *The Annals of Applied Probability*, 20(4):1470–1511, 2010.
- 2 Anne Auger and Benjamin Doerr, editors. *Theory of Randomized Search Heuristics: Foundations and Recent Developments*, volume 1. World Scientific Publishing Co., Inc., 2011.
- 3 Siddhartha Banerjee and Peter Lofgren. Fast bidirectional probability estimation in Markov models. In *Proc. of NIPS*, pages 1423–1431, 2015.
- 4 Phillip Bonacich and Paulette Lloyd. Eigenvector-like measures of centrality for asymmetric relations. *Social Networks*, 23(3):191–201, 2001.
- 5 Christian Borgs, Michael Brautbar, Jennifer T. Chayes, and Shang-Hua Teng. A sub-linear time algorithm for pagerank computations. In Anthony Bonato and Jeannette C. M. Janssen, editors, *Algorithms and Models for the Web Graph - 9th International Workshop, WAW 2012, Halifax, NS, Canada, June 22-23, 2012. Proceedings*, volume 7323 of *Lecture Notes in Computer Science*, pages 41–53. Springer, 2012. doi:10.1007/978-3-642-30541-2\_4.
- 6 Christian Borgs, Michael Brautbar, Jennifer T. Chayes, and Shang-Hua Teng. Multiscale matrix sampling and sublinear-time PageRank computation. *Internet Mathematics*, 10(1-2):20–48, 2014.
- 7 Marco Bressan, Enoch Peserico, and Luca Pretto. On approximating the stationary distribution of time-reversible Markov chains. *CoRR*, abs/1801.00196, 2018.
- 8 Kai-Min Chung, Henry Lam, Zhenming Liu, and Michael Mitzenmacher. Chernoff-Hoeffding bounds for Markov chains: Generalized and simplified. In *Proc. of STACS*, pages 124–135, 2012.
- 9 David A. Freedman. On tail probabilities for martingales. *The Annals of Probability*, 3(1):100–118, 1975.
- 10 Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. Matrix Computations. Johns Hopkins University Press, 2012.
- 11 Wilfred K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.
- 12 Christina E. Lee, Asuman Ozdaglar, and Devavrat Shah. Computing the stationary distribution, locally. In *Proc. of NIPS*, pages 1376–1384, 2013.
- 13 Christina E. Lee, Asuman E. Ozdaglar, and Devavrat Shah. Solving systems of linear equations: Locally and asynchronously. *CoRR*, abs/1411.2647, 2014.

- 14 David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov chains and mixing times*. American Mathematical Society, 2009.
- 15 Peter Lofgren, Siddhartha Banerjee, and Ashish Goel. Bidirectional PageRank estimation: from average-case to worst-case. In *Proc. of WAW*, pages 164–176, 2015.
- 16 Peter A. Lofgren, Siddhartha Banerjee, Ashish Goel, and C. Seshadhri. FAST-PPR: Scaling personalized PageRank estimation for large graphs. In *Proc. of ACM KDD*, pages 1436–1445, 2014.
- 17 Rajeev Motwani, Rina Panigrahy, and Ying Xu. Estimating sum by weighted sampling. In *Proc. of ICALP*, pages 53–64, 2007.
- 18 Alessandro Panconesi and Aravind Srinivasan. Randomized distributed edge coloring via an extension of the Chernoff–Hoeffding bounds. *SIAM Journal on Computing*, 26(2):350–368, 1997.
- 19 Ronitt Rubinfeld and Asaf Shapira. Sublinear time algorithms. *SIAM Journal on Discrete Mathematics*, 25(4):1562–1588, 2011.
- 20 Nitin Shyamkumar, Siddhartha Banerjee, and Peter Lofgren. Sublinear estimation of a single element in sparse linear systems. In *2016 Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 856–860, 2016.

## 5 Appendix

### 5.1 Probability bounds

This appendix provides Chernoff-type probability bounds that are repeatedly used in our analysis; these bounds can be found in e.g. [2], and can be derived from [18].

Let  $X_1, \dots, X_n$  be binary random variables. We say that  $X_1, \dots, X_n$  are non-positively correlated if for all  $I \subseteq \{1, \dots, n\}$  we have:

$$\Pr[\forall i \in I : X_i = 0] \leq \prod_{i \in I} \Pr[X_i = 0] \quad (10)$$

$$\Pr[\forall i \in I : X_i = 1] \leq \prod_{i \in I} \Pr[X_i = 1] \quad (11)$$

The following lemma holds:

► **Lemma 12.** *Let  $X_1, \dots, X_n$  be independent or, more generally, non-positively correlated binary random variables. Let  $a_1, \dots, a_n \in [0, 1]$  and  $X = \sum_{i=1}^n a_i X_i$ . Then, for any  $\epsilon > 0$ , we have:*

$$\Pr[X < (1 - \epsilon)\mathbb{E}[X]] < e^{-\frac{\epsilon^2}{2}\mathbb{E}[X]} \quad (12)$$

$$\Pr[X > (1 + \epsilon)\mathbb{E}[X]] < e^{-\frac{\epsilon^2}{2+\epsilon}\mathbb{E}[X]} \quad (13)$$

Note that Lemma 12 applies if  $X_1, \dots, X_n$  are indicator variables of mutually disjoint events, or if they can be partitioned into independent families  $\{X_1, \dots, X_{i_1}\}, \{X_{i_1+1}, \dots, X_{i_2}\}, \dots$  of such variables.

### 5.2 A lower bound for non-time-reversible Markov chains

► **Lemma 13.** *For any functions  $\tau(n) = \omega(1)$  and  $p(n) = o(\frac{1}{n})$  there exists a family of ergodic non-time-reversible Markov chains on  $n$  states having mixing time  $\tau = \Theta(\tau(n))$ , and containing a state  $v$  with  $\pi(v) = \Theta(p(n))$  such that any algorithm needs  $\Omega(\frac{\tau}{\pi(v)})$  calls to  $\text{step}()$  to estimate  $\pi(v)$  within constant multiplicative factors with constant probability.*

## 18:14 On Approximating the Stationary Distribution of Time-reversible Markov Chains

**Proof.** Consider a chain with state space  $\{u\} \cup \{u_1, \dots, u_{n-1}\}$  and the following transition probabilities (we assume  $n$  large enough to set in  $[0, 1]$  any quantity where needed). For  $u$ , set  $p_{uu} = 1 - \frac{(n-1)p(n)}{\tau(n)}$ , and  $p_{uu_i} = \frac{p(n)}{\tau(n)}$  for all  $i = 1, \dots, n-1$ . For all  $i = 1, \dots, n-1$ , set  $p_{u_i u_i} = 1 - \frac{1}{\tau(n)}$  and  $p_{u_i u} = \frac{1}{\tau(n)}$ . The chain is clearly ergodic. Note that  $\frac{(n-1)p(n)}{\tau(n)} = o(\frac{1}{\tau(n)})$  and therefore the expected time to leave  $u$  is asymptotically larger than the expected time to leave any of the  $u_i$ . One can then check that (i)  $\pi(u_i) = \Theta(p(n))$ , and (ii) the mixing time is  $\tau = \Theta(\tau(n))$  (essentially, the expected time to leave the  $u_i$ ). Pick any  $u_i$  as target state  $v$ . Suppose now to alter the chain as follows: pick some  $u_j \neq v$  and set  $p_{u_j v} = 1$ . The new stationary probability of  $v$  would then be roughly  $2\pi(v)$ . However one cannot distinguish between the two chains with constant probability with less than  $\Omega(\frac{\tau}{\pi(v)})$  *step()* calls. Indeed, to distinguish between them one must at least visit  $u_j$  (and then perform e.g. *probe*( $u_j, v$ )). Since  $u$  is the only state leading to  $u_j$  with positive probability, one must invoke *step*( $u$ ) until it returns  $u_j$ . But  $p_{uu_j} = \frac{p(n)}{\tau(n)}$ , hence one needs  $\Omega(\frac{\tau(n)}{p(n)}) = \Omega(\frac{\tau}{\pi(v)})$  calls in expectation. The construction can be adapted to any constant approximation factor by adding more transitions towards  $v$ . ◀