

Beyond Random Walk and Metropolis-Hastings Samplers: Why You Should Not Backtrack for Unbiased Graph Sampling

Chul-Ho Lee Xin Xu Do Young Eun
Department of Electrical and Computer Engineering
North Carolina State University
Raleigh, NC 27695
{clee4,xxu5,dyeun}@ncsu.edu

ABSTRACT

Graph sampling via crawling has been actively considered as a generic and important tool for collecting uniform node samples so as to consistently estimate and uncover various characteristics of complex networks. The so-called simple random walk with re-weighting (SRW-rw) and Metropolis-Hastings (MH) algorithm have been popular in the literature for such unbiased graph sampling. However, an unavoidable downside of their core random walks – slow diffusion over the space, can cause poor estimation accuracy. In this paper, we propose non-backtracking random walk with re-weighting (NBRW-rw) and MH algorithm with delayed acceptance (MHDA) which are *theoretically guaranteed* to achieve, at *almost no additional cost*, not only unbiased graph sampling but also higher efficiency (smaller asymptotic variance of the resulting unbiased estimators) than the SRW-rw and the MH algorithm, respectively. In particular, a remarkable feature of the MHDA is its applicability for *any* non-uniform node sampling like the MH algorithm, but ensuring better sampling efficiency than the MH algorithm. We also provide simulation results to confirm our theoretical findings.

Categories and Subject Descriptors: G.3 [Probability and Statistics]: Probabilistic algorithms, Stochastic processes

General Terms: Theory, Algorithms, Performance

Keywords: unbiased graph sampling, random walks, non-reversible Markov chains, asymptotic variance

1. INTRODUCTION

Estimating various nodal and topological properties of complex networks such as online social networks (OSNs), peer-to-peer (P2P) networks, and the world wide web (WWW) has recently attracted much attention from research community because of their ever-increasing popularity and importance in our daily life. However, the estimation of network characteristics is a non-trivial task, as these networks are

typically too large to measure, making a complete picture of the network hard to obtain and even its size unknown. It is thus infeasible to perform ‘independence sampling’ which obtains uniform node samples (for unbiased estimation) directly and independently from such a large, unknown network. Instead, graph crawling techniques – graph sampling via crawling, have been widely used for that purpose. In particular, random walk-based graph sampling methods (or Markov chain samplers) have become popular, as they are simple and implementable in a distributed fashion and also able to provide *unbiased* graph sampling, unlike the breadth-first-search (BFS) and its variants leading to unknown bias [12, 20].

In the literature, the most popular random walk-based graph sampling methods are the so-called simple random walk with re-weighting (SRW-rw) [29, 12] and Metropolis-Hastings (MH) algorithm [25, 16, 34, 29, 12, 15]. The former launches a simple random walk (SRW) over a graph \mathcal{G} , which moves from a node to one of its neighbors chosen uniformly at random, to collect random node samples, followed by a re-weighting process in order to eliminate the bias caused by the non-uniform stationary distribution of the SRW. The other method is to rely on a Metropolis-Hastings random walk (MHRW) crawling over \mathcal{G} – a random walk achieving a uniform distribution constructed by the famous MH algorithm [25, 16], to obtain uniform node samples.

Motivation and Contributions: While the SRW-rw and MH algorithm ensure unbiased graph sampling, the core components – SRW and MHRW, suffer from their slow diffusion over the space, which can in turn lead to poor estimation accuracy. In particular, their fully random nature in selecting the next node, when making a transition, often cause them to go back to the previous node from where they just came. This produces many duplicate samples for a short to moderate time span, thereby reducing estimation accuracy. It is apparently desirable to avoid such backtracking transitions whenever possible, so as to steer them toward ‘unvisited’ places (or to obtain new node samples), as long as such a modification does not affect the unbiased estimation.

However, it is still uncertain how to achieve this *at almost no additional cost* and whether it really results in better estimation accuracy. We provide affirmative answers for these questions. Specifically, we propose non-backtracking random walk with re-weighting (NBRW-rw) and MH algorithm with delayed acceptance (MHDA), and prove that each of them guarantees not only *unbiased* graph sampling but also *higher efficiency* (smaller asymptotic variance of the estimators) than the SRW-rw and the MH algorithm, respec-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGMETRICS’12, June 11–15, 2012, London, England, UK.
Copyright 2012 ACM 978-1-4503-1097-0/12/06 ...\$10.00.

tively. A notable feature of our MHDA is its generic purpose: the MHDA is theoretically guaranteed to enhance the standard MH algorithm for constructing a random walk or a Markov chain with *any* arbitrarily given stationary distribution under the constraints of graph structure. Thus, the MHDA is applied, as ‘a special case’, to construct a random walk crawling over a graph \mathcal{G} achieving a uniform stationary distribution, leading to higher efficiency than the MHRW while ensuring the unbiased estimation. To the best of our knowledge, this is the first theoretical result to improve, with proven guarantee, *both* SRW-rw and the MH algorithm for unbiased graph sampling.

Related Work: Very recently, there have been a few attempts to improve the estimation accuracy against the SRW-rw (not the MH algorithm) through multiple dependent random walks [30], a random walk on a weighted graph (with a priori estimate of network information) [20], and the addition of random jumps (to anywhere in the graph) [5]. The corresponding Markov chains are time-reversible, whereas the main kernel of our proposed methods is transforming ‘any’ reversible Markov chain to its related non-reversible chain which avoids backtracking transitions and also achieves the same stationary distribution. Thus, our work is complementary to their approaches.

On the other hand, there is a body of research works across many disciplines for speeding up a random walk, or Markov chain, on a graph \mathcal{G} in terms of its mixing time, hitting time, and/or cover time. The fastest mixing (reversible) Markov chain on a graph is obtained in [8] with complete knowledge of entire graph. [9, 10] showed that certain ‘lifted’ (non-reversible) Markov chains converge to their stationary distributions faster than their related reversible chain, and [19, 23] subsequently applied this idea to design a fast and efficient average consensus algorithm. It is, however, still unknown how to construct such a ‘lifted’ Markov chain in a distributed or decentralized manner for a general graph.

[3, 7, 17] recently undertook speeding up a SRW based only on local information, but did not provide any direct implication to the unbiased graph sampling. As the MH algorithm is the most popular method of Markov Chain Monte Carlo (MCMC) simulations or samplers, it has been an active research topic to improve the MH algorithm in terms of the sampler performance (asymptotic variance) in the MCMC literature (e.g., [26, 14, 35]). However, most works toward more efficient MCMC samplers (including [26, 14, 35]) do not take into account graph-topological constraints in that transition from node i to $j \neq i$ is allowed only when they are neighbors of each other, and thus cannot be directly applicable to unbiased graph sampling.

Organization: The rest of the paper is organized as follows. We first provide an in-depth overview on generic Markov chain samplers for unbiased graph sampling in Section 2, and then briefly review the SRW-rw and MH algorithm in Section 3. In Section 4, we present a general recipe for the transformation of a time-reversible Markov to its related non-reversible Markov chain, which forms a common building block for our proposed NBRW-rw and MHDA. We then explain the details of the NBRW-rw and MHDA, and provide relevant analysis. In Section 5, we provide simulation results obtained based on real graphs to support our theoretical findings. We finally conclude in Section 6.

2. PRELIMINARIES

2.1 Unbiased Graph Sampling

Consider a *connected, undirected* graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ with a set of nodes (vertices) $\mathcal{N} = \{1, 2, \dots, n\}$ and a set of edges \mathcal{E} . We assume that $3 \leq |\mathcal{N}| = n < \infty$. We also assume that the graph \mathcal{G} has no self-loops and no multi-edges. Let $N(i) \triangleq \{j \in \mathcal{N} : (i, j) \in \mathcal{E}\}$ be the set of neighbors of node $i \in \mathcal{N}$, and $d(i) \triangleq |N(i)|$ be the degree of node i .

Unbiased graph (or node) sampling, via crawling, is to consistently estimate nodal or topological properties of a target graph \mathcal{G}^* (e.g., an overlay network or an OSN) based upon *uniform* node samples obtained by a random walk (or possibly multiple random walks) crawling over the graph \mathcal{G} . The goal here is to unbiasedly estimate a proportion of the nodes with a specific characteristic. Thus, the unbiased, uniform graph sampling is, in principle, developing a random walk-based ‘estimator’ or a Markov chain sampler for the expectation of any given, desired function f with respect to a uniform distribution, i.e., $\mathbb{E}_{\mathbf{u}}(f) \triangleq \sum_{i \in \mathcal{N}} f(i)/n$, where $\mathbf{u} \triangleq [u(1), u(2), \dots, u(n)] = [1/n, 1/n, \dots, 1/n]$. Note that a nodal (or topological) characteristic of interest can be specified by properly choosing a function f . For example, for a target graph \mathcal{G} , if one is interested in estimating its degree distribution (say, $\mathbb{P}\{D_{\mathcal{G}} = d\}, d = 1, 2, \dots, n-1$), then choose a function f such that $f(i) = \mathbf{1}_{\{d(i)=d\}}$ for $i \in \mathcal{N}$, i.e., $f(i) = 1$ if $d(i) = d$, and $f(i) = 0$ otherwise.

We below review a basic Markov chain theory which serves as the mathematical foundation for unbiased graph sampling via a (Markovian) random walk crawling over a graph \mathcal{G} . Define a random walk or a finite discrete-time Markov chain $\{X_t \in \mathcal{N}\}_{t \geq 0}$ on the nodes of the graph \mathcal{G} with its transition matrix $\mathbf{P} \triangleq \{P(i, j)\}_{i, j \in \mathcal{N}}$ in which $P(i, j) = \mathbb{P}\{X_{t+1} = j | X_t = i\}, i, j \in \mathcal{N}$, and $\sum_j P(i, j) = 1$ for all i . Each edge $(i, j) \in \mathcal{E}$ is associated with a transition probability $P(i, j) \geq 0$ with which the chain (or random walk) makes a transition from node i to node j . We allow the chain to include self-transitions, i.e., $P(i, i) > 0$ for some i , although \mathcal{G} has no self-loops. Clearly, $P(i, j) = 0$ for all $(i, j) \notin \mathcal{E}$ ($i \neq j$). We then assume that the Markov chain $\{X_t\}$ is irreducible, i.e., every node in \mathcal{N} is reachable in finite time with positive probability, such that the chain has a unique stationary distribution $\boldsymbol{\pi} \triangleq [\pi(1), \pi(2), \dots, \pi(n)]$.

For any function $f : \mathcal{N} \rightarrow \mathbb{R}$, define an estimator

$$\hat{\mu}_t(f) \triangleq \frac{1}{t} \sum_{s=1}^t f(X_s) \quad (1)$$

for the expectation of the function f with respect to $\boldsymbol{\pi}$ which is given by $\mathbb{E}_{\boldsymbol{\pi}}(f) \triangleq \sum_{i \in \mathcal{N}} f(i)\pi(i)$. Then, the following Strong Law of Large Numbers (SLLN) (a.k.a., ergodic theorem) has been a fundamental basis for most of the random walk-based graph sampling methods in the literature [34, 29, 12, 15, 30, 5, 20], and more generally, MCMC samplers [28, 26, 35, 18, 31, 27].

THEOREM 1. [18, 31] *Suppose that $\{X_t\}$ is a finite, irreducible Markov chain with its stationary distribution $\boldsymbol{\pi}$. Then, for any initial distribution $\mathbb{P}\{X_0 = i\}, i \in \mathcal{N}$, as $t \rightarrow \infty$,*

$$\hat{\mu}_t(f) \rightarrow \mathbb{E}_{\boldsymbol{\pi}}(f) \text{ almost surely (a.s.)}$$

for any function f with $\mathbb{E}_{\boldsymbol{\pi}}(|f|) < \infty$. \square

*A target graph for sampling may be time-varying due to node join/leave, which is beyond the scope of this paper.

The SLLN ensures that the estimator $\hat{\mu}_t(f)$ based on any finite, irreducible Markov chain with the same π can serve as a valid and unbiased approximation of $\mathbb{E}_\pi(f)$.

2.2 Central Limit Theorem and Asymptotic Variance

For a given graph \mathcal{G} , there are potentially many (finite) irreducible Markov chains (or different random walks) preserving the same stationary distribution π , all of which can be used to obtain asymptotically unbiased estimates of $\mathbb{E}_\pi(f)$, and also of $\mathbb{E}_u(f)$, together with proper re-weighting if $\pi \neq u$. One important question would then be how to compare these ‘competing’ Markov chains, or rather, which one is ‘better’ or more efficient than the others as a Markov chain sampler for unbiased graph sampling.

Mixing time [8, 22] can perhaps be a criterion to compare several irreducible, aperiodic Markov chains, all with the same stationary distribution. If the speed of convergence to the stationary distribution is a primary concern, then the mixing time is surely the right metric to compare different Markov chains with the same stationary distribution. However, this is not the case for the unbiased graph sampling. Random walk-based graph sampling methods typically adopt an initial burn-in period over which (initial) sampled values are discarded to get rid of the dependence on the initial position of a random walk [12]. After such a burn-in period, the Markov chain (or random walk) will be close to its stationary regime (well mixed), but many samples are still yet to be obtained from this point onward. Therefore, the primary concern should be, instead, the *efficiency* of the estimator $\hat{\mu}_t(f)$ in deciding how many random samples are required to achieve a certain accuracy of $\hat{\mu}_t(f)$ in regard to $\mathbb{E}_\pi(f)$ (and eventually to $\mathbb{E}_u(f)$ after proper re-weighting if necessary).

To that end, we define by $\sigma^2(f)$ the asymptotic variance of the estimator $\hat{\mu}_t(f)$ based on an irreducible Markov chain $\{X_t\}$ with its stationary distribution π , which is given by

$$\begin{aligned} \sigma^2(f) &\triangleq \lim_{t \rightarrow \infty} t \cdot \text{Var}(\hat{\mu}_t(f)) \\ &= \lim_{t \rightarrow \infty} \frac{1}{t} \mathbb{E} \left\{ \left[\sum_{s=1}^t (f(X_s) - \mathbb{E}_\pi(f)) \right]^2 \right\} \end{aligned} \quad (2)$$

for any function f with $\mathbb{E}_\pi(f^2) < \infty$, where the initial position (state) X_0 is drawn from the stationary distribution π , i.e., $X_0 \sim \pi$. Note that the asymptotic variance $\sigma^2(f)$ is, in fact, *independent* of the distribution of the initial state X_0 [28, 31]. Not only for the well-known case with *i.i.d.* samples (or random variables), the Central Limit Theorem (CLT) holds also for Markov chains, as given below.

THEOREM 2. [18, 31] *For a finite, irreducible Markov chain $\{X_t\}$ with its stationary distribution π ,*

$$\sqrt{t} \cdot [\hat{\mu}_t(f) - \mathbb{E}_\pi(f)] \xrightarrow{d} N(0, \sigma^2(f)), \text{ as } t \rightarrow \infty,$$

for any function f with $\mathbb{E}_\pi(f^2) < \infty$ regardless of any initial distribution, where \xrightarrow{d} denotes convergence in distribution and $N(0, \sigma^2(f))$ is a Gaussian random variable with zero mean and variance $\sigma^2(f)$ given by (2). \square

Note that Theorems 1–2 (SLLN and CLT) do not require any assumption of aperiodicity [31]. However, for simplicity, we do not consider periodic Markov chains in our analysis

throughout the paper. In addition, we focus on *bounded* functions f (and thus $\mathbb{E}_\pi(f^2) < \infty$), which is typical in graph sampling applications.

The CLT says that the distribution of $\hat{\mu}_t(f)$ is asymptotically normal with variance $\sigma^2(f)$, and thus allows one to evaluate the asymptotic variance $\sigma^2(f)$ in order to decide approximately how many (correlated) samples are required to achieve a certain accuracy of the estimator $\hat{\mu}_t(f)$. Hence, the asymptotic variance $\sigma^2(f)$ has been an important criterion to rank the *efficiency* among competing Markov chains with the same π for the MCMC samplers [28, 26, 35, 18, 31, 27], although quantifying $\sigma^2(f)$ may not be easy. In particular, by noting that the asymptotic variance is independent of any initial distribution for which the CLT holds, the efficiency ordering over competing Markov chains with the same π (the smaller the asymptotic variance, the better the estimator performance) is still in effect even when the competing Markov chains are already in their stationary regimes (already ‘mixed’). Observe that from $X_0 \sim \pi$, $X_t \sim \pi$ for all t (the chain $\{X_t\}$ is in the stationary regime), and thus (2) becomes

$$\sigma^2(f) = \text{Var}(f(X_0)) + 2 \sum_{k=1}^{\infty} \text{Cov}(f(X_0), f(X_k)), \quad (3)$$

where $\text{Cov}(f(X_0), f(X_k)) = \mathbb{E}\{f(X_0)f(X_k)\} - \mathbb{E}_\pi^2(f)$ denotes the covariance between $f(X_0)$ and $f(X_k)$. As a special case, if the samples $\{X_t\}$ are *i.i.d.* and drawn directly from π , then $\sigma^2(f) = \text{Var}(f(X_0))$. Also, even if the competing Markov chains are already in their stationary regimes, the correlation structure over random samples given by each of these Markov chains can vary and significantly affect their asymptotic variances. Observe that reducing the temporal correlation over random samples can lead to smaller asymptotic variances. This intuition can be leveraged to improve the existing Markov chain samplers for unbiased graph sampling.

Motivated by the effectiveness of the asymptotic variance with its connection to the CLT, in this paper, we consider the asymptotic variance as a primary performance metric, and develop two random walk-based graph sampling methods, each of which guarantees the unbiased graph sampling with smaller asymptotic variance than its corresponding counterpart in the current networking literature. Before going into details, we next briefly review the existing two random walk-based graph sampling methods.

3. RANDOM WALK-BASED SAMPLING

3.1 Simple Random Walk with Re-weighting

We first review the SRW-rw, a.k.a., respondent-driven sampling [33], which has been recently used in [29, 12] for unbiased graph sampling. This method operates based upon a sequence of (correlated) random samples obtained by a SRW, together with a proper re-weighting process to ensure the unbiased sampling. It is essentially a special case of the importance sampling (a Monte Carlo method) applied for random samples generated by a Markov chain [6, 24, 13].

Consider a SRW on \mathcal{G} that moves from a node to one of its neighbors chosen uniformly at random (u.a.r.). Specifically, let $\{X_t\}$ be the Markov chain representing the sequence of visited nodes by the SRW, with its transition matrix $\mathbf{P} = \{P(i, j)\}_{i, j \in \mathcal{N}}$ given by $P(i, j) = 1/d(i)$ if $(i, j) \in \mathcal{E}$, and $P(i, j) = 0$ otherwise. It is well known that \mathbf{P} is irre-

ducible, and reversible with respect to a unique stationary distribution π for which $\pi(i) = d(i)/(2|\mathcal{E}|)$, $i \in \mathcal{N}$ [2].

Suppose that there are t random samples $\{X_s\}_{s=1}^t$ from the SRW. Then, for a function of interest f , choose a weight function $w : \mathcal{N} \rightarrow \mathbb{R}$ such that $w(i) = u(i)/\pi(i) = 2|\mathcal{E}|/(nd(i))$, $i \in \mathcal{N}$. Observe that from the SLLN in Theorem 1, as $t \rightarrow \infty$,

$$\hat{\mu}_t(wf) = \frac{1}{t} \sum_{s=1}^t w(X_s)f(X_s) \rightarrow \mathbb{E}_\pi(wf) = \mathbb{E}_\mathbf{u}(f) \text{ a.s.}$$

and thus the estimator $\hat{\mu}_t(wf)$ is unbiased for $\mathbb{E}_\mathbf{u}(f)$. However, this estimator itself is not practical, since n and $|\mathcal{E}|$ are typically unknown a priori. Instead, another estimator $\hat{\mu}_t(wf)/\hat{\mu}_t(w)$ is often used as an unbiased estimator for $\mathbb{E}_\mathbf{u}(f)$. Indeed, Theorem 1 asserts that $\hat{\mu}_t(wf)$ and $\hat{\mu}_t(w)$ converge to $\mathbb{E}_\mathbf{u}(f)$ and 1 almost surely, as $t \rightarrow \infty$, respectively. This yields

$$\frac{\hat{\mu}_t(wf)}{\hat{\mu}_t(w)} = \frac{\sum_{s=1}^t w(X_s)f(X_s)}{\sum_{s=1}^t w(X_s)} \rightarrow \mathbb{E}_\mathbf{u}(f) \text{ a.s.}$$

Hence, the estimator $\hat{\mu}_t(wf)/\hat{\mu}_t(w)$ can be made in such a way that we need to know $w(i)$ only up to a multiplicative constant. That is, if we set $w(i) = 1/d(i)$, $i \in \mathcal{N}$, then the estimator $\hat{\mu}_t(wf)/\hat{\mu}_t(w)$ remains intact, and is more practical as an unbiased estimator for $\mathbb{E}_\mathbf{u}(f)$. Throughout this paper, we refer to the estimator $\hat{\mu}_t(wf)/\hat{\mu}_t(w)$ with $w(i) = 1/d(i)$ ($i \in \mathcal{N}$) as the unbiased estimator for $\mathbb{E}_\mathbf{u}(f)$ in the SRW-rw [29, 12].

3.2 Metropolis-Hastings Algorithm

The MH algorithm [25, 16] was developed to construct a transition matrix \mathbf{P} of a time-reversible Markov chain $\{X_t\}$ with a given, desired stationary distribution π . Here, we only discuss the MH algorithm under the topological constraints of a graph \mathcal{G} in that transition from node i to $j \neq i$ is allowed only when they are neighbors of each other. The MH algorithm is defined as follows. At the current state i of X_t , the next state X_{t+1} is proposed with a *proposal* probability $Q(i, j)$, which is a state transition probability of an arbitrary irreducible Markov chain on the state space \mathcal{N} , where $Q(i, j) > 0$ if and only if $Q(j, i) > 0$, and $Q(i, j) = 0$ for all $(i, j) \notin \mathcal{E}$ ($i \neq j$). Let $\mathbf{Q} \triangleq \{Q(i, j)\}_{i, j \in \mathcal{N}}$ be a proposal (transition) matrix. The proposed state transition to $X_{t+1} = j$ is accepted with an *acceptance* probability

$$A(i, j) = \min \left\{ 1, \frac{\pi(j)Q(j, i)}{\pi(i)Q(i, j)} \right\}, \quad (4)$$

and rejected with probability $1 - A(i, j)$ in which case $X_{t+1} = i$. Thus, the transition probability $P(i, j)$ becomes, for $i \neq j$,

$$P(i, j) = Q(i, j)A(i, j) = \min \left\{ Q(i, j), Q(j, i) \frac{\pi(j)}{\pi(i)} \right\}, \quad (5)$$

with $P(i, i) = 1 - \sum_{j \neq i} P(i, j)$, which ensures that \mathbf{P} is reversible with respect to π . Note that the uniqueness of π is granted due to the irreducibility of \mathbf{Q} (so is \mathbf{P}) and the finite state space.

The MH algorithm, in addition to its popular applications for MCMC simulation, has been also widely used as a means for unbiased graph sampling [34, 29, 12, 15]. Specifically, the MH algorithm has been applied to construct a MHRW on \mathcal{G} achieving a uniform stationary distribution, i.e., $\pi = \mathbf{u}$. This is done with transition probabilities of a SRW as the

proposal probabilities, i.e., $Q(i, j) = 1/d(i)$ if $(i, j) \in \mathcal{E}$, and $Q(i, j) = 0$ otherwise. The resulting transition probability of the MHRW on \mathcal{G} becomes

$$P(i, j) = \begin{cases} \min \left\{ \frac{1}{d(i)}, \frac{1}{d(j)} \right\} & \text{if } (i, j) \in \mathcal{E}, \\ 0 & \text{if } (i, j) \notin \mathcal{E}, i \neq j, \end{cases} \quad (6)$$

and $P(i, i) = 1 - \sum_{j \neq i} P(i, j)$. Thus, \mathbf{P} is reversible with respect to $\pi = \mathbf{u}$, implying that for any function f , the estimator $\hat{\mu}_t(f)$ based upon random samples by the MHRW is unbiased for $\mathbb{E}_\mathbf{u}(f)$. This version of MH algorithm is summarized in Algorithm 1, where $X_t \in \mathcal{N}$ denotes the location of the MHRW at time t , and $d(X_t)$ denotes the degree of node X_t . Here, X_0 can be arbitrarily chosen.

Algorithm 1 MH algorithm for MHRW (at time t)

- 1: Choose node j u.a.r. from neighbors of X_t , i.e., $N(X_t)$
 - 2: Generate $p \sim U(0, 1)$
 - 3: **if** $p \leq \min \left\{ 1, \frac{d(X_t)}{d(j)} \right\}$ **then**
 - 4: $X_{t+1} \leftarrow j$
 - 5: **else**
 - 6: $X_{t+1} \leftarrow X_t$
 - 7: **end if**
-

REMARK 1. *The MH algorithm (Algorithm 1) does not need to know the self-transition probabilities $P(i, i)$ explicitly, nor does it require all the neighbors' degree information of the current node X_t at each time t . Instead, only the degree information of the randomly chosen neighbor j is enough for making decision whether or not to move to j .*

Recall that the above unbiased estimators are based on t random, consecutive samples obtained under SRW or MHRW, respectively. Observe that the SRW, currently at node i at time s can 'backtrack' to the previously visited node with probability $1/d(i)$, i.e., $X_{s+1} = X_{s-1}$, trapping the SRW temporarily in a local region. The situation can be worse for the regions in which nodes have small degrees (so higher chance of backtracking). Similarly, the MHRW at node i can also backtrack to the previously visited node after staying at node i for some random time. This slow 'diffusion' of SRW/MHRW over the space can, in turn, lead to highly duplicated random samples for a short to moderate time duration, thereby increasing the variance of the unbiased estimators. Recall that the asymptotic variance in (3) involves covariance terms $\text{Cov}(f(X_0), f(X_k))$. Thus, it would be beneficial for both SRW and MHRW (or precisely, their variants) to avoid backtracking to the previously visited node up to the extent possible in order to reduce the temporal correlation over random consecutive samples, while maintaining the same stationary distribution so that the aforementioned mathematical framework for the unbiased estimation remains intact. Thus motivated, for the rest of this paper, we investigate how to achieve this at *almost no additional cost*, and rigorously prove that our proposed sampling methods give smaller (no worse) asymptotic variance than the SRW (with re-weighting) and MHRW-based ones, respectively.

4. AVOID BACKTRACKING TO PREVIOUSLY VISITED NODE

In this section, we propose two random walk-based graph sampling methods – (i) **non-backtracking random walk**

with re-weighting and (ii) MH algorithm with delayed acceptance, each of which *theoretically guarantees* unbiased graph sampling with smaller asymptotic variance than the SRW-rw and the (original) MH algorithm, respectively. In particular, our proposed sampling methods require almost no additional cost, or more precisely, just remembering where the underlying random walk came from, when compared to the conventional methods. The reasoning behind the improvement of asymptotic variance is to modify each of SRW and MHRW, when making a transition from the current node to one of its neighbors, to reduce bias toward the previous state (one of the neighbors of the current node), while maintaining the same stationary distribution. Note that such directional bias breaks the time-reversibility of the SRW and MHRW. Thus, a common building block for our proposed sampling methods will be, for a given reversible Markov chain with its stationary distribution π , to construct a *non-reversible* Markov chain preserving the same π while avoiding (to the extent possible) transitions that backtrack to the state from which the chain just came. Our challenge here is to construct such a non-reversible chain with only one-step memory and theoretical guarantee for higher efficiency (smaller asymptotic variance). In what follows, we first explain a basic setup for this transformation and several relevant issues, and then present the details of our proposed methods.

4.1 From Reversible To Non-reversible Chains

Consider a generic random walk on \mathcal{G} , or a finite, irreducible, *time-reversible* Markov chain $\{X_t \in \mathcal{N}\}_{t \geq 0}$, with its transition matrix $\mathbf{P} = \{P(i, j)\}_{i, j \in \mathcal{N}}$ and stationary distribution $\pi = [\pi(i), i \in \mathcal{N}]$. Our goal here is to construct its related new random walk or a finite, irreducible, *non-reversible* Markov chain with the same π which avoids backtracking to the previously visited node, which in turn produces a smaller asymptotic variance than the original reversible chain. An important requirement is that this transformation should be done at no additional cost and in a distributed manner. It is worth noting that there have been other works [9, 10] showing that certain non-reversible Markov chains or lifted Markov chains mix substantially faster than their related reversible chains. While this concept has been also applied to design a fast and efficient average consensus algorithm [19, 23], it is still unknown how to construct such a non-reversible chain or lifted Markov chain in a fully distributed or decentralized fashion, not to mention how to do so for any arbitrarily given target stationary distribution π .

A general recipe for constructing a non-reversible Markov chain in an augmented state space: Let $X'_t \in \mathcal{N}$, $t = 0, 1, 2, \dots$, be the location of a new random walk at time t . At the current node X'_t , the next node X'_{t+1} is decided based upon not only the current node X'_t but also the previous node X'_{t-1} so as to avoid backtracking. Due to the dependency (memory) to the previous node, $\{X'_t\}_{t \geq 0}$ itself cannot be a Markov chain on the state space \mathcal{N} , regardless of the choice of transition matrix. This walk, however, can still be made Markovian on an augmented state space instead, defined by

$$\Omega \triangleq \{(i, j) : i, j \in \mathcal{N} \text{ s.t. } P(i, j) > 0\} \subseteq \mathcal{N} \times \mathcal{N} \quad (7)$$

with $|\Omega| < \infty$, and $Z'_t \triangleq (X'_{t-1}, X'_t) \in \Omega$ for $t \geq 1$. For notational simplicity, let e_{ij} denote state $(i, j) \in \Omega$. Note that

$e_{ij} \neq e_{ji}$. It is also possible that $e_{ii} \in \Omega$ for some i . A similar interpretation of a weighted random walk (or a reversible Markov chain) on the augmented state space can be also found in [2, Ch.3], although its purpose is not for the construction of a related non-reversible chain.

Let $\mathbf{P}' \triangleq \{P'(e_{ij}, e_{lk})\}_{e_{ij}, e_{lk} \in \Omega}$ be the transition matrix of an irreducible Markov chain $\{Z'_t \in \Omega\}_{t \geq 1}$ on the state space Ω . Here, by definition, $P'(e_{ij}, e_{lk}) = 0$ for all $j \neq l$. If the unique stationary distribution $\pi' \triangleq [\pi'(e_{ij}), e_{ij} \in \Omega]$ of the chain $\{Z'_t\}$ is given by

$$\pi'(e_{ij}) = \pi(i)P(i, j), \quad e_{ij} \in \Omega, \quad (8)$$

implying that $\pi'(e_{ij}) = \pi'(e_{ji})$ from the reversibility of the original chain $\{X_t\}$, then the probability of the new random walk $\{X'_t\}$ being at node j in the *steady-state* is the same as $\pi(j)$ for all j (the stationary distribution of the original reversible chain $\{X_t\}$). To see this, note that

$$\sum_{i \in \mathcal{N}: e_{ij} \in \Omega} \pi'(e_{ij}) = \sum_{i \in \mathcal{N}} \pi(i)P(i, j) = \pi(j), \quad \forall j \in \mathcal{N}, \quad (9)$$

where the first equality follows from $P(u, v) = 0, \forall (u, v) \notin \Omega$. In particular, for any original function of interest $f : \mathcal{N} \rightarrow \mathbb{R}$, choose another function $g : \Omega \rightarrow \mathbb{R}$ such that $g(e_{ij}) = f(j)$, and observe

$$\begin{aligned} \mathbb{E}_{\pi'}(g) &= \sum_{e_{ij} \in \Omega} g(e_{ij})\pi'(e_{ij}) = \sum_{j \in \mathcal{N}} \sum_{i \in \mathcal{N}} f(j)\pi(i)P(i, j) \\ &= \sum_{j \in \mathcal{N}} f(j)\pi(j) = \mathbb{E}_{\pi}(f). \end{aligned}$$

Then, the SLLN in Theorem 1 gives

$$\frac{1}{t} \sum_{s=1}^t g(Z'_s) = \frac{1}{t} \sum_{s=1}^t f(X'_s) \rightarrow \mathbb{E}_{\pi'}(g) = \mathbb{E}_{\pi}(f) \text{ a.s.}, \quad (10)$$

i.e., $\sum_{s=1}^t g(Z'_s)/t$ is a valid unbiased estimator for $\mathbb{E}_{\pi}(f)$. We thus define, for any given function $f : \mathcal{N} \rightarrow \mathbb{R}$,

$$\hat{\mu}'_t(f) \triangleq \frac{1}{t} \sum_{s=1}^t f(X'_s) \quad (11)$$

to be clearly distinguished from $\hat{\mu}_t(f)$ in (1) defined based on the original chain $\{X_t\}$, while $\hat{\mu}'_t(f)$ and $\hat{\mu}_t(f)$ are both unbiased estimators for $\mathbb{E}_{\pi}(f)$. In addition, the CLT in Theorem 2 implies

$$\begin{aligned} \sqrt{t} \cdot \left[\frac{1}{t} \sum_{s=1}^t g(Z'_s) - \mathbb{E}_{\pi'}(g) \right] &= \sqrt{t} \cdot [\hat{\mu}'_t(f) - \mathbb{E}_{\pi}(f)] \\ &\xrightarrow{d} N(0, \sigma'^2(f)), \quad (12) \end{aligned}$$

where $\sigma'^2(f)$ denotes the asymptotic variance of $\hat{\mu}'_t(f)$ (and also of $\sum_{s=1}^t g(Z'_s)/t$). Throughout the paper, we use the prime symbol ($'$) for any notation related to a newly defined process (e.g., $\{X'_t\}$) to differentiate it from its counterpart defined on the original process (e.g., $\{X_t\}$).

While there are infinitely many different transition matrices \mathbf{P}' leading to the unbiased estimator $\hat{\mu}'_t(f)$ for $\mathbb{E}_{\pi}(f)$, our primary goal is, at (almost) no additional cost and in a distributed manner, to find a transition matrix \mathbf{P}' that also guarantees smaller asymptotic variance. Under a rather restricted setting, R. Neal gave a partial answer to this in [27] saying that less backtracking (rendering the resulting

Markov chain $\{Z'_t\}$ non-reversible) can result in a smaller asymptotic variance. We restate his finding below.

THEOREM 3. [27, Theorem 2] *Suppose that $\{X_t\}$ is an irreducible, reversible Markov chain on the state space \mathcal{N} with transition matrix $\mathbf{P} = \{P(i, j)\}$ and stationary distribution π . Construct a Markov chain $\{Z'_t\}$ on the state space Ω with transition matrix $\mathbf{P}' = \{P'(e_{ij}, e_{lk})\}$ in which the transition probabilities $P'(e_{ij}, e_{lk})$ satisfy the following two conditions: for all $e_{ij}, e_{ji}, e_{jk}, e_{kj} \in \Omega$ with $i \neq k$,*

$$P(j, i)P'(e_{ij}, e_{jk}) = P(j, k)P'(e_{kj}, e_{ji}), \quad (13)$$

$$P'(e_{ij}, e_{jk}) \geq P(j, k). \quad (14)$$

Then, the Markov chain $\{Z'_t\}$ is irreducible and non-reversible with a unique stationary distribution π' in which $\pi'(e_{ij}) = \pi(i)P(i, j)$, $e_{ij} \in \Omega$. Also, for any function f , the asymptotic variance of $\hat{\mu}'_t(f)$ is no greater than that of $\hat{\mu}_t(f)$, i.e., $\sigma'^2(f) \leq \sigma^2(f)$. \square

REMARK 2. *The condition in (13) ensures that the resulting transition matrix \mathbf{P}' is stationary with respect to π' in (8) and, in turn, leads to the unbiased estimator $\hat{\mu}'_t(f)$ for $\mathbb{E}_{\pi}(f)$. Together with this condition, the condition in (14) – less backtracking to the previously visited node, brings out the improvement of asymptotic variance.*

Theorem 3 is quite versatile and provides a guideline on how to choose the transition matrix \mathbf{P}' of a Markov chain $\{Z'_t\}$ leading to smaller asymptotic variance, and thus will play an essential role in developing our graph sampling methods and subsequent analysis. Despite this large degree of freedom, it is still uncertain how to choose such a transition matrix \mathbf{P}' at no additional cost. While R. Neal suggested a procedure to find \mathbf{P}' , it generally poses significant cost, especially for improving the MH algorithm, as admitted in [27]. (See pp. 9–10 therein.) Recall that the MH algorithm (Algorithm 1) for MHRW only needs the degree information of a randomly chosen neighbor j of the current node X_t to decide whether or not to move j , as mentioned in Remark 1.[†] However, the procedure by R. Neal necessitates the *explicit knowledge* of all $P(i, i)$'s [27]. That is, the corresponding modified MHRW would require all the neighbors' degree information of the current node X'_t at each time t in order to choose the next node X'_{t+1} . Imagine such modified MHRW crawling over an OSN (say, Facebook) and located at a certain user's page. To simply decide where to go, the walk would have to visit all his/her friends' pages first and collect all their degree information (i.e., the number of friends) before making decision to move. This is clearly impractical for our graph sampling purpose. Therefore, in this paper, we set out to develop our own graph samplers with higher efficiency *without any such overhead*, by leveraging Theorem 3 as a building block.

4.2 Non-backtracking Random Walk with Re-weighting

We first introduce non-backtracking random walk with re-weighting (NBRW-rw) that ensures unbiased graph sampling, and then prove that NBRW-rw guarantees a smaller

[†]More generally, in the MH algorithm with any proposal matrix \mathbf{Q} , it is often unnecessary to know self-transition probabilities $P(i, i)$ explicitly, or does not require summing the probabilities of rejection for all possible proposals just to compute $P(i, i) = Q(i, i) + \sum_{j \neq i} Q(i, j)(1 - A(i, j))$.




Figure 1: Illustrating the transitions of an NBRW. The walker is currently located at node j (with $d(j)=4$) and just came from node i . From j , it will move to one of its neighbors except node i with equal probability.

asymptotic variance than SRW-rw. The non-backtracking random walk (NBRW) is a discrete-time random walk which ‘never’ backtracks (thus named non-backtracking) to the previous node (whenever possible) while preserving the same stationary distribution as that of a SRW. Thus, the proposed sampling method is to use an NBRW, instead of a SRW, to collect a sequence of samples by crawling over a target \mathcal{G} , and at the same time, to employ the same re-weighting process as is done for SRW in order to eliminate sampling bias induced from its non-uniform stationary distribution.

Consider an irreducible, reversible Markov chain $\{X_t\}_{t \geq 0}$ (a sequence of visited nodes) by the SRW, with its transition matrix \mathbf{P} and stationary distribution π . Then, the NBRW is defined as follows. A (discrete-time) random walk at the current node j with $d(j) \geq 2$ moves to the next node k , chosen u.a.r. from the neighbors of node j except the previous node i . If the current node j has only one neighbor ($d(j)=1$), the walk always returns to the previous node i . Figure 1 depicts this non-backtracking nature of the NBRW in its transitions over the nodes of \mathcal{G} . Here, an initial position of the NBRW can be arbitrarily chosen. The NBRW initially moves from the initial position to one of its neighbors with equal probability due to the absence of its ‘previous node’, and then proceeds as defined above thereafter.

Let $X'_t \in \mathcal{N}$, $t=0, 1, 2, \dots$, be the location of an NBRW. As before, we construct a Markov chain $\{Z'_t = (X'_{t-1}, X'_t)\}_{t \geq 1}$ with its transition matrix $\mathbf{P}' = \{P'(e_{ij}, e_{lk})\}_{e_{ij}, e_{lk} \in \Omega}$ given by, for all $e_{ij}, e_{jk} \in \Omega$ with $i \neq k$ ($d(j) \geq 2$),

$$P'(e_{ij}, e_{jk}) = \frac{1}{d(j)-1} > \frac{1}{d(j)} = P(j, k), \quad (15)$$

implying that $P'(e_{ij}, e_{ji}) = 0$. Also, $P'(e_{ij}, e_{ji}) = 1$ for any j with $d(j) = 1$. All other elements of \mathbf{P}' are zero. Clearly, \mathbf{P}' satisfies the conditions in (13)–(14). From Theorem 3, the Markov chain $\{Z'_t\}$ is irreducible and non-reversible with a unique stationary distribution

$$\pi'(e_{ij}) = \pi(i)P(i, j) = \frac{1}{2|\mathcal{E}|}, \quad e_{ij} \in \Omega. \quad (16)$$

That is, the probability of the NBRW being at node j in the steady-state is the same as $\pi(j)$. See (9). From (10)–(11) and Theorem 3, we also know that for any given function f of interest, $\hat{\mu}'_t(f)$ and $\hat{\mu}_t(f)$ are both unbiased estimators for $\mathbb{E}_{\pi}(f)$, and the asymptotic variance of $\hat{\mu}'_t(f)$ (based on the random samples by the NBRW) is no larger than that of $\hat{\mu}_t(f)$ (by the SRW), i.e., $\sigma'^2(f) \leq \sigma^2(f)$.

However, both unbiased estimators $\hat{\mu}'_t(f)$ and $\hat{\mu}_t(f)$ are for $\mathbb{E}_{\pi}(f)$, not $\mathbb{E}_{\mathbf{u}}(f)$. It is unclear whether such improvement for the asymptotic variance remains true even after a proper re-weighting to obtain *unbiased* samples. As explained in Section 3.1, the SRW-rw is to use the estimator $\hat{\mu}_t(wf)/\hat{\mu}_t(w)$ with $w(i) = 1/d(i)$ ($i \in \mathcal{N}$) in order to consis-

tently estimate $\mathbb{E}_{\mathbf{u}}(f)$. Since the stationary distribution of the NBRW remains the same as that of the SRW, we can also use the estimator $\hat{\mu}'_t(wf)/\hat{\mu}'_t(w)$ with the same weight function w , as a valid approximation of $\mathbb{E}_{\mathbf{u}}(f)$. Let $\sigma_{\mathbb{W}}^2(f)$ and $\sigma'_{\mathbb{W}}(f)$ denote the asymptotic variances of the estimators $\hat{\mu}_t(wf)/\hat{\mu}_t(w)$ and $\hat{\mu}'_t(wf)/\hat{\mu}'_t(w)$, respectively. To proceed, we need the following.

THEOREM 4 (SLUTSKY'S THEOREM). [4, pp.332]
Let $\{A_t\}$ and $\{B_t\}$ be the sequences of random variables. If $A_t \xrightarrow{d} A$, and B_t converges in probability to a non-zero constant b , then $A_t/B_t \xrightarrow{d} A/b$. \square

Now we state our main result.

THEOREM 5. For any function $f : \mathcal{N} \rightarrow \mathbb{R}$, the asymptotic variance of $\hat{\mu}'_t(wf)/\hat{\mu}'_t(w)$ is no larger than that of $\hat{\mu}_t(wf)/\hat{\mu}_t(w)$, i.e., $\sigma_{\mathbb{W}}'^2(f) \leq \sigma_{\mathbb{W}}^2(f)$, where the weight function w is given by $w(i) = 1/d(i)$, $i \in \mathcal{N}$. \square

PROOF. Since the estimator $\hat{\mu}_t(wf)/\hat{\mu}_t(w)$ remains invariant up to a constant multiple of w , without loss of generality, we can set $w(i) = u(i)/\pi(i) = 2|\mathcal{E}|/(nd(i))$. For any given f , observe that

$$\begin{aligned} \sqrt{t} \left[\frac{\hat{\mu}_t(wf)}{\hat{\mu}_t(w)} - \mathbb{E}_{\mathbf{u}}(f) \right] &= \sqrt{t} \left[\frac{\sum_{s=1}^t w(X_s) f(X_s)}{\sum_{s=1}^t w(X_s)} - \mathbb{E}_{\mathbf{u}}(f) \right] \\ &= \frac{t}{\sum_{s=1}^t w(X_s)} \sqrt{t} \left[\frac{\sum_{s=1}^t w(X_s) (f(X_s) - \mathbb{E}_{\mathbf{u}}(f))}{t} \right]. \end{aligned} \quad (17)$$

Define another function $h : \mathcal{N} \rightarrow \mathbb{R}$ such that

$$h(i) \triangleq w(i)(f(i) - \mathbb{E}_{\mathbf{u}}(f)), \quad i \in \mathcal{N},$$

implying $\mathbb{E}_{\pi}(h) = \sum_{i \in \mathcal{N}} h(i)\pi(i) = 0$. Then, from Theorems 1 and 2, we have, as $t \rightarrow \infty$,

$$\frac{1}{t} \sum_{s=1}^t w(X_s) \rightarrow 1 \text{ a.s.}, \text{ and } \sqrt{t} \left[\frac{1}{t} \sum_{s=1}^t h(X_s) \right] \xrightarrow{d} \mathbb{N}(0, \sigma^2(h)).$$

Since almost sure convergence implies convergence in probability [4], by Slutsky's theorem, from (17), we have

$$\sqrt{t} \left[\frac{\hat{\mu}_t(wf)}{\hat{\mu}_t(w)} - \mathbb{E}_{\mathbf{u}}(f) \right] \xrightarrow{d} \mathbb{N}(0, \sigma^2(h)), \text{ as } t \rightarrow \infty.$$

Together with (10) and (12), following the same lines above, we similarly have

$$\sqrt{t} \left[\frac{\hat{\mu}'_t(wf)}{\hat{\mu}'_t(w)} - \mathbb{E}_{\mathbf{u}}(f) \right] \xrightarrow{d} \mathbb{N}(0, \sigma'^2(h)), \text{ as } t \rightarrow \infty.$$

Hence, for a given f , the asymptotic variance of the estimator $\hat{\mu}_t(wf)/\hat{\mu}_t(w)$ is nothing but $\sigma_{\mathbb{W}}^2(f) = \sigma^2(h)$. Similarly, $\sigma_{\mathbb{W}}'^2(f) = \sigma'^2(h)$. Therefore, since Theorem 3 says that for any function f , $\sigma'^2(f) \leq \sigma^2(f)$, we also have $\sigma_{\mathbb{W}}'^2(f) \leq \sigma_{\mathbb{W}}^2(f)$. That is, the asymptotic variance of $\hat{\mu}'_t(wf)/\hat{\mu}'_t(w)$ is no larger than that of $\hat{\mu}_t(wf)/\hat{\mu}_t(w)$. \blacksquare

REMARK 3. In [3], the NBRW was originally considered for regular graphs with $d(i) = d > 3$, $\forall i \in \mathcal{N}$, and shown to lead to faster mixing rate (i.e., faster rate of convergence to its stationary distribution) than that of the SRW. In contrast, for any general (connected, undirected, not necessarily regular) graph \mathcal{G} , we show that the NBRW-rw ensures not only the unbiased graph sampling but also smaller asymptotic variance than the SRW-rw.

4.3 Metropolis-Hastings Algorithm with Delayed Acceptance

We turn our attention to improving the MH algorithm. For any given, desired stationary distribution π , we propose Metropolis-Hastings algorithm with delayed acceptance (MHDA), which theoretically guarantees smaller asymptotic variance than the (generic) MH algorithm with proposal matrix \mathbf{Q} that constructs a reversible Markov chain with arbitrary π . In particular, we demonstrate that MHDA can be applied, as a special case, to construct a (non-Markovian) random walk on a graph \mathcal{G} which not only achieves a uniform stationary distribution $\pi = \mathbf{u}$ for unbiased graph sampling, but leads to higher efficiency than MHRW by the MH algorithm (Algorithm 1). We emphasize that the only additional overhead here is remembering the previously visited node (one of the neighbors of the current node) from which the random walk came.

Interpreting a reversible MH Markov chain as a semi-Markov chain: Consider an irreducible, reversible Markov chain $\{X_t \in \mathcal{N}\}_{t \geq 0}$ by the MH algorithm with its transition matrix $\mathbf{P} = \{P(i, j)\}_{i, j \in \mathcal{N}}$ given by (5), and any arbitrarily given target stationary distribution π . Recall that the MH algorithm is nothing but a repetition of proposing a state transition with proposal probability $Q(i, j)$ that is then accepted with an acceptance probability $A(i, j)$ in (4) or rejected otherwise. Observe that the process $\{X_t\}$, after entering into state (node) i , stays at state i for a geometrically distributed time duration with mean $1/(1 - P(i, i))$, and then moves to another state $j \in N(i)$. Formally, define a Markov chain $\{\tilde{X}_m \in \mathcal{N}\}_{m \geq 0}$ with its transition matrix $\tilde{\mathbf{P}} \triangleq \{\tilde{P}(i, j)\}_{i, j \in \mathcal{N}}$ given by, for $j \neq i$,

$$\begin{aligned} \tilde{P}(i, j) &= \frac{P(i, j)}{1 - P(i, i)} = \frac{Q(i, j)A(i, j)}{\sum_{j \neq i} Q(i, j)A(i, j)} \\ &= \frac{\min\{Q(i, j), Q(j, i)\pi(j)/\pi(i)\}}{\sum_{j \neq i} \min\{Q(i, j), Q(j, i)\pi(j)/\pi(i)\}}, \end{aligned} \quad (18)$$

with $\tilde{P}(i, i) = 0$. It is not difficult to see that the chain $\{\tilde{X}_m\}$ is irreducible, and reversible with respect to a unique stationary distribution $\tilde{\pi} \triangleq [\tilde{\pi}(i), i \in \mathcal{N}]$, given by

$$\tilde{\pi}(i) \propto \pi(i)(1 - P(i, i)), \quad i \in \mathcal{N}.$$

Also, we define a function $\gamma : \mathcal{N} \rightarrow \mathbb{R}$ such that, for $i \in \mathcal{N}$,

$$\gamma(i) \triangleq 1 - P(i, i) = \sum_{j \neq i} \min\{Q(i, j), Q(j, i)\pi(j)/\pi(i)\}, \quad (19)$$

and define a sequence $\{\xi_m\}_{m \geq 0}$ for which ξ_m depends solely on $\{\tilde{X}_m\}_{m \geq 0}$ and is geometrically distributed with parameter $\gamma(\tilde{X}_m)$. It thus follows that $\mathbb{E}\{\xi_m | \tilde{X}_m = i\} = 1/\gamma(i)$, $i \in \mathcal{N}$. The process $\{X_t\}$ can now be interpreted as a semi-Markov chain with embedded Markov chain $\{\tilde{X}_m\}$ and respective sojourn times $\{\xi_m\}$. Suppose that the random walk by the MH algorithm (or the process $\{X_t\}$) enters node j of a graph \mathcal{G} , depicted in Figure 2, at time $t=1$ ($X_1 = j$). If we consider a sample path $(X_1, X_2, \dots, X_7) = (j, j, j, i, i, j, k)$, then we have corresponding sequences $(\tilde{X}_1, \tilde{X}_2, \tilde{X}_3, \tilde{X}_4) = (j, i, j, k)$ and $(\xi_1, \xi_2, \xi_3) = (3, 2, 1)$. Note that the standard definition of a semi-Markov process allows the sojourn time ξ_m to depend on both \tilde{X}_m and \tilde{X}_{m+1} (and so we are dealing with a special case). From the theory of semi-Markov processes (e.g., [32]), one can easily recover the stationary




Figure 2: An example graph \mathcal{G}

distribution π as

$$\pi(i) \propto \tilde{\pi}(i)/\gamma(i), \quad i \in \mathcal{N}. \quad (20)$$

The above interpretation has been similarly given in the MCMC literature [24, 11]. In particular, it is known that, for any given function $f : \mathcal{N} \rightarrow \mathbb{R}$,

$$\hat{\mu}_{m,\text{MH}}(f) \triangleq \frac{\sum_{l=1}^m \xi_l f(\tilde{X}_l)}{\sum_{l=1}^m \xi_l} \quad (21)$$

converges almost surely to $\mathbb{E}_{\pi}(f)$, as $m \rightarrow \infty$, and thus $\hat{\mu}_{m,\text{MH}}(f)$ is also an unbiased estimator for $\mathbb{E}_{\pi}(f)$ [24]. This definition of $\hat{\mu}_{m,\text{MH}}(f)$ enables more tractable analysis on its asymptotic variance, denoted as $\sigma_{\text{MH}}^2(f)$, by connecting it to its counterpart in the importance sampling for Markov chains [24, 11]. Note that for sufficiently large t (also m), the (original) unbiased estimator $\hat{\mu}_t(f) = \sum_{s=1}^t f(X_s)/t$ can be written as $\hat{\mu}_{m,\text{MH}}(f)$ plus some negligible term (after setting the same initial point $\tilde{X}_1 = X_1$), because it is always possible to find m such that $\sum_{l=1}^m \xi_l \leq t < \sum_{l=1}^{m+1} \xi_l$. Also, in the limit $t, m \rightarrow \infty$, $\hat{\mu}_t(f)$ and $\hat{\mu}_{m,\text{MH}}(f)$ are the same. We thus focus on estimators in the form of $\hat{\mu}_{m,\text{MH}}(f)$ in our subsequent analysis.

Consider a sequence of pairs (\tilde{X}_m, ξ_m) . From the success in the NBRW-rw, one may ask what if the reversible embedded Markov chain $\{\tilde{X}_m\}_{m \geq 0}$ is replaced by a related stochastic process $\{\tilde{X}'_m \in \mathcal{N}\}_{m \geq 0}$, or more precisely, a non-reversible Markov chain $\{(\tilde{X}'_{m-1}, \tilde{X}'_m)\}_{m \geq 1}$ on the augmented state space Ω , which avoids backtracking transitions to the extent possible, while preserving the same stationary distribution $\tilde{\pi}$. Another question can be whether this transformation guarantees that the estimator in (21) based on (\tilde{X}'_m, ξ_m) remains *unbiased* for $\mathbb{E}_{\pi}(f)$ and also have *higher efficiency* than the original one. Our answer is in the affirmative, and this is the reasoning behind the improvement of our proposed MHDA over the standard MH algorithm. We stress here that, in contrast to the NBRW, backtracking transitions in the process $\{\tilde{X}'_m\}$ should be *avoided only up to the extent possible*[‡] so as to maintain the arbitrarily given original stationary distribution $\tilde{\pi}$. Thus, the extension from the MH algorithm to our proposed MHDA becomes necessarily more involved than the case of NBRW.

Description of MHDA: Let $X'_t \in \mathcal{N}$, $t = 0, 1, 2, \dots$, be the position of a random walk (or the state of a stochastic process). We also define the augmented state space Ω in (7) based on the reversible embedded Markov chain $\{\tilde{X}_m\}$, where $\tilde{P}(i, i) = 0$ and so $e_{ii} \notin \Omega$ for all i .

MHDA is described as follows. Suppose that node i is the previous node from which the walk came. MHDA first operates just like the MH algorithm. At the current node (state) $X'_t = j \neq i$, the next node $X'_{t+1} = k \in N(j)$ is proposed with probability $Q(j, k)$ ($j \neq k$). Then, the proposed transition to k is accepted with probability $A(j, k)$ in (4), and

[‡]If $\{\tilde{X}'_m\}$ is made purely non-backtracking just like we did for NBRW, then we lose unbiasedness for the resulting MH-based estimator in general.




Figure 3: Illustrating a difference between MH algorithm and MHDA: a walker moves from node j to node k with probability $P(j, k)$ in the MH algorithm, but with larger probability $P(j, k) + P(j, i)Q'(e_{ij}, e_{jk})A'(e_{ij}, e_{jk})$ in the MHDA.

rejected with probability $1 - A(j, k)$ in which case $X'_{t+1} = j$. Here, in contrast to the MH algorithm, MHDA renders the accepted transition to $X'_{t+1} = k$ temporarily *pending*, and applies another procedure to proceed with the actual transition to k .[§]

Specifically, for the accepted transition to k , if $k \neq i$, then the ‘actual’ transition takes place, i.e., $X'_{t+1} = k$, with probability $P(j, k) = Q(j, k)A(j, k)$ as in the MH algorithm. On the other hand, if $k = i$, then the transition to node i is *delayed* (thus named ‘delayed acceptance’). The next node X'_{t+1} is again proposed with another proposal probability $Q'(e_{ij}, e_{jk})$, which is a transition probability of an arbitrary Markov chain on the state space Ω , where $Q'(e_{ij}, e_{lk}) = 0$ for all $j \neq l$, and $Q'(e_{ij}, e_{jk}) > 0$ if and only if $Q'(e_{kj}, e_{ji}) > 0$. The (second) proposed transition to $X'_{t+1} = k$ is accepted with another acceptance probability $A'(e_{ij}, e_{jk})$, and rejected with probability $1 - A'(e_{ij}, e_{jk})$ in which case $X'_{t+1} = i$ (backtracking occurs). That is, transition probability $P(j, i) = Q(j, i)A(j, i)$ in the MH algorithm is leveraged to create another transition opportunity from j to $k \neq i$ in the MHDA. So, the transition from j to $k \neq i$ occurs with larger probability $P(j, k) + P(j, i)Q'(e_{ij}, e_{jk})A'(e_{ij}, e_{jk})$ than the MH algorithm (w.p. $P(j, k)$). This is also illustrated in Figure 3 where the thickness of arrows represents the corresponding transition probabilities from node j to other node (including self-transition). The new acceptance probability $A'(e_{ij}, e_{jk})$ will be specified shortly.

In summary, under MHDA, the walker stays at each node for the same random amount of time as it would be under the MH algorithm, while reducing the bias toward the previous node when making transitions to one of its neighbors.

Analysis of MHDA: Let \tilde{X}'_m , $m \geq 0$, be the sequence of nodes visited by the walk, which moves over \mathcal{G} according to the MHDA. The process $\{\tilde{X}'_m\}$ is clearly different from the reversible, embedded Markov chain $\{\tilde{X}_m\}$ for the MH algorithm. Also, let ξ'_m , $m \geq 0$, be the respective sojourn time at node \tilde{X}'_m . Note that the MHDA behaves differently from the MH algorithm (performs the additional procedure) only when a proposed transition from node j to node $k \neq j$ (occurring with probability $Q(j, k)$) is accepted with probability $A(j, k)$ in the MH algorithm. Thus, ξ'_m is also geometrically distributed with parameter $\gamma(\tilde{X}'_m)$. See (19) for $\gamma(\cdot)$. That is, given that $\tilde{X}'_{m_1} = \tilde{X}'_{m_2} = i$, the sojourn times ξ'_{m_1} and ξ'_{m_2} have identical distributions. There-

[§]If the transition to $k = j$ was accepted after a proposal with $Q(j, j) > 0$, then the MHDA accepts the transition as in the MH algorithm without any further action.

fore, the MHDA, similar to the MH algorithm, can be also characterized by a sequence of the pairs (\tilde{X}'_m, ξ'_m) . As an example, if the random walk by the MHDA (or the process $\{\tilde{X}'_t\}$) enters node i in Figure 2 at time $t = 1$ ($X'_1 = i$) and $(X'_1, X'_2, \dots, X'_7) = (i, i, l, j, j, j, k)$, then we consequently have $(\tilde{X}'_1, \tilde{X}'_2, \tilde{X}'_3, \tilde{X}'_4) = (i, l, j, k)$ and $(\xi'_1, \xi'_2, \xi'_3) = (2, 1, 3)$. We define, for any given $f : \mathcal{N} \rightarrow \mathbb{R}$,

$$\hat{\mu}'_{m, \text{MHDA}}(f) \triangleq \frac{\sum_{i=1}^m \xi'_i f(\tilde{X}'_i)}{\sum_{i=1}^m \xi'_i}. \quad (22)$$

We prove below that $\hat{\mu}'_{m, \text{MHDA}}(f)$ converges almost surely to $\mathbb{E}_\pi(f)$, implying that $\hat{\mu}'_{m, \text{MHDA}}(f)$ is an unbiased estimator for $\mathbb{E}_\pi(f)$. We also prove that, after showing the CLT holds for $\hat{\mu}'_{m, \text{MHDA}}(f)$, the asymptotic variance of $\hat{\mu}'_m(f)$, denoted as $\sigma'^2_{\text{MHDA}}(f)$, is smaller than its counterpart $\sigma^2_{\text{MH}}(f)$ for the MH algorithm.

To this end, we first explain how to properly choose the new acceptance $A'(e_{ij}, e_{jk})$ so that the process $\{\tilde{X}'_m\}$ has the same stationary distribution as that of the reversible embedded chain $\{\tilde{X}_m\}$, while, at the same time, the process $\{\tilde{X}'_m\}$ reduces backtracking transitions. Instead of the process $\{\tilde{X}'_m\}$, we deal with its related non-reversible Markov chain defined on the augmented state space Ω by consulting the general recipe for this purpose in Section 4.1. Recall the state space Ω in (7) obtained from the transition matrix $\tilde{\mathbf{P}} = \{\tilde{P}(i, j)\}$ of the chain $\{\tilde{X}_m\}$. We define $\tilde{Z}'_m \triangleq (\tilde{X}'_{m-1}, \tilde{X}'_m) \in \Omega$ for $m \geq 1$, and $\tilde{\mathbf{P}}' \triangleq \{\tilde{P}'(e_{ij}, e_{lk})\}_{e_{ij}, e_{lk} \in \Omega}$ to be the transition matrix of a Markov chain $\{\tilde{Z}'_m\}_{m \geq 1}$. For instance, consider a sample path $(\tilde{X}'_1, \tilde{X}'_2, \tilde{X}'_3, \tilde{X}'_4) = (i, l, j, k)$ in the above example. We have $(\tilde{Z}'_2, \tilde{Z}'_3, \tilde{Z}'_4) = ((i, l), (l, j), (j, k))$. If the chain $\{\tilde{Z}'_m\}$ has a unique stationary distribution $\tilde{\pi}' \triangleq [\tilde{\pi}'(e_{ij}), e_{ij} \in \Omega]$ given by

$$\tilde{\pi}'(e_{ij}) = \tilde{\pi}(i)\tilde{P}(i, j), \quad e_{ij} \in \Omega, \quad (23)$$

implying that $\tilde{\pi}'(e_{ij}) = \tilde{\pi}'(e_{ji})$ from the reversibility of the embedded chain $\{\tilde{X}_m\}$, then the *steady-state* probability of the process $\{\tilde{X}'_m\}$ being at node j is the same as $\tilde{\pi}(j)$ for all j . From the description of MHDA, observe that, for all $e_{ij}, e_{jk} \in \Omega$ with $i \neq k$ ($d(j) \geq 2$),

$$\tilde{P}'(e_{ij}, e_{jk}) = \tilde{P}(j, k) + \tilde{P}(j, i)Q'(e_{ij}, e_{jk})A'(e_{ij}, e_{jk}), \quad (24)$$

while $\tilde{P}'(e_{ij}, e_{ji}) = 1 - \sum_{k \neq i} \tilde{P}'(e_{ij}, e_{jk})$, as is also shown in Figure 3(b). Note that $\tilde{P}'(e_{ij}, e_{jk})$ specifies the next node of the random walk by MHDA, *given that* the walk has to move from the current node to one of its neighbors (its sojourn time is over). Thus, $\tilde{P}(j, k)$ and $\tilde{P}(j, i)$ are used here instead of $P(j, k)$ and $P(j, i)$, respectively. In addition, for any j with $d(j) = 1$, we have $\tilde{P}'(e_{ij}, e_{ji}) = \tilde{P}(j, i) = 1$, $(i, j) \in \mathcal{E}$, since $Q'(e_{ij}, e_{ji}) = 1$ (due to the stochastic matrix $\{Q'(e_{ij}, e_{lk})\}$ and $A'(e_{ij}, e_{ji}) = 1$ which is shown below).

Among many possible choices for the acceptance probability $A'(e_{ij}, e_{jk})$ in the MHDA, we have the following.

PROPOSITION 1. *For any given $\{Q'(e_{ij}, e_{lk})\}$, suppose that the acceptance probability $A'(e_{ij}, e_{jk})$ is given by*

$$A'(e_{ij}, e_{jk}) = \min \left\{ 1, \frac{P^2(j, k)Q'(e_{kj}, e_{ji})}{P^2(j, i)Q'(e_{ij}, e_{jk})} \right\}. \quad (25)$$

Then, the resulting transition matrix $\tilde{\mathbf{P}}'$, and $\tilde{\mathbf{P}}$ satisfy conditions (13)–(14). \square

PROOF. See our technical report [21]. \blacksquare

From Theorem 3 and Proposition 1, the Markov chain $\{\tilde{Z}'_m\}$ with its transition matrix $\tilde{\mathbf{P}}'$ as in (24) and (25), is irreducible and non-reversible with a unique stationary distribution $\tilde{\pi}'$ in (23). This also implies that the process $\{\tilde{X}'_m\}$ has the same stationary distribution $\tilde{\pi}$, as explained before. We now present our main result.

THEOREM 6. *Consider a given, desired stationary distribution $\pi = [\pi(i), i \in \mathcal{N}]$. Under the MHDA with any given $\{Q'(e_{ij}, e_{lk})\}$ and its corresponding $A'(e_{ij}, e_{jk})$ in (25), for any given function $f : \mathcal{N} \rightarrow \mathbb{R}$, as $m \rightarrow \infty$, $\hat{\mu}'_{m, \text{MHDA}}(f)$ converges almost surely to $\mathbb{E}_\pi(f)$, and also the asymptotic variance of $\hat{\mu}'_{m, \text{MHDA}}(f)$ is no larger than that of $\hat{\mu}_{m, \text{MH}}(f)$, i.e., $\sigma'^2_{\text{MHDA}}(f) \leq \sigma^2_{\text{MH}}(f)$.* \square

PROOF. See our technical report [21]. \blacksquare

An application of MHDA for unbiased graph sampling:

We explain how MHDA can be applied for unbiased graph sampling applications. In particular, we present how to construct a (discrete-time) random walk by MHDA, named Metropolis-Hastings Random walk with Delayed Acceptance (MHRW-DA), on \mathcal{G} that achieves the uniform stationary distribution, i.e., $\pi = \mathbf{u}$. The MHRW-DA here operates as an extension of Algorithm 1 with the following choice of $\{Q'(e_{ij}, e_{lk})\}$: for all $e_{ij}, e_{jk} \in \Omega$ with $i \neq k$ ($d(j) \geq 2$),

$$Q'(e_{ij}, e_{jk}) = 1/(d(j) - 1), \quad (26)$$

implying that $Q'(e_{ij}, e_{ji}) = 0$. Also, $Q'(e_{ij}, e_{ji}) = 1$ for any j with $d(j) = 1$. All other elements are zero. While $\{Q'(e_{ij}, e_{lk})\}$ is the same as the transition matrix of NBRW, a ‘Metropolizing’ step, which is done with $A'(e_{ij}, e_{jk})$ in (25), must follow in order to ensure that the stationary distribution is uniform and the resulting estimator is unbiased. In other words, $A'(e_{ij}, e_{jk})$ in (25) becomes

$$A'(e_{ij}, e_{jk}) = \min \left\{ 1, \min \left\{ \frac{1}{d(j)^2}, \frac{1}{d(k)^2} \right\} / \min \left\{ \frac{1}{d(j)^2}, \frac{1}{d(i)^2} \right\} \right\}.$$

This version of the MHDA is summarized in Algorithm 2, where $X'_t \in \mathcal{N}$ is the location of MHRW-DA at time t and $Y_t \in \mathcal{N}$ indicates the previous node from which the MHRW-DA came ($Y_t \neq X'_t$). Here, X'_0 can be chosen arbitrarily. Since there is no notion of ‘previous node’ Y_0 at time $t = 0$, MHRW-DA initially behaves the same as MHRW until it moves from the initial position to one of its neighbors, and then proceeds as described in Algorithm 2 thereafter.

Theorem 6 states that the MHDA works for any given stationary distribution π , while allowing us to freely choose the new proposal probabilities $\{Q'(e_{ij}, e_{lk})\}$ as desired. Thus, Algorithm 2 for MHRW-DA is nothing but a ‘special case’ of the MHDA. Theorem 6 asserts that MHRW-DA produces unbiased samples with higher efficiency than the corresponding MHRW (Algorithm 1). Again, we emphasize that the only additional overhead for MHRW-DA, compared to the MHRW, is *remembering where it came from*, Y_t . Note that the degree of the previous node Y_t is already known and can easily be retrieved, while the degree information of another randomly chosen neighbor is also necessary anyway even in the MH algorithm (to decide whether or not to move there).

5. SIMULATION RESULTS

In this section, we present simulation results to support our theoretical findings. To this end, we use the following

Algorithm 2 MHDA for MHRW-DA (at time t)

```

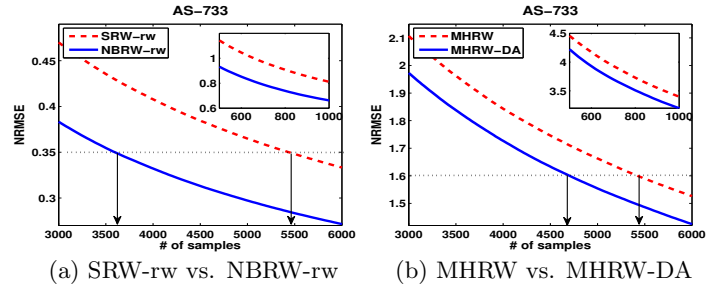
1: Choose node  $i$  u.a.r. from neighbors of  $X'_t$ , i.e.,  $N(X'_t)$ 
2: Generate  $p \sim U(0, 1)$ 
3: if  $p \leq \min\left\{1, \frac{d(X'_t)}{d(i)}\right\}$  then
4:   if  $Y_t = i$  and  $d(X'_t) > 1$  then
5:     Choose node  $k$  u.a.r. from  $N(X'_t) \setminus \{i\}$ 
6:     Generate  $q \sim U(0, 1)$ 
7:     if  $q \leq \min\left\{1, \min\left\{1, \left(\frac{d(X'_t)}{d(k)}\right)^2\right\} \max\left\{1, \left(\frac{d(i)}{d(X'_t)}\right)^2\right\}\right\}$ 
8:       then
9:          $X'_{t+1} \leftarrow k$  and  $Y_{t+1} \leftarrow X'_t$ 
10:      else
11:         $X'_{t+1} \leftarrow i$  and  $Y_{t+1} \leftarrow X'_t$ 
12:      end if
13:    else
14:       $X'_{t+1} \leftarrow i$  and  $Y_{t+1} \leftarrow X'_t$ 
15:    end if
16:  else
17:     $X'_{t+1} \leftarrow X'_t$  and  $Y_{t+1} \leftarrow Y_t$ 
18:  end if

```

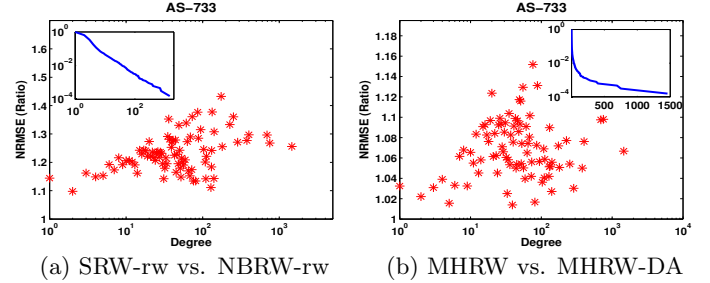
real-world network datasets [1]: **AS-733** – an undirected graph of autonomous systems (ASs) composed of 6474 nodes and 13233 edges, where nodes represent ASs and edges exist according to AS-AS peering relationships; **Road-PA** – a road network of Pennsylvania, forming an undirected graph with 1088092 nodes and 3083796 edges, where nodes represent intersections and endpoints and edges represent the roads connecting them; **Web-Google** – a directed web graph with 875713 nodes and 5105039 edges, where nodes represent web pages and directed edges represent hyperlinks between them. For our simulation, we use an undirected version of this web graph. To ensure graph connectivity, we also use the largest connected component (LCC) of each graph, where the LCC sizes of AS-733, Road-PA, and Web-Google graphs are 6474, 1087562, and 855802, respectively. Here, the average degrees of AS-733, Road-PA, and Web-Google graphs are 4.09, 2.83, and 10.03, while their maximum degrees are 1459, 9, and 6332, respectively.

As a test case, we consider the estimation of the degree distribution of each graph – $\mathbb{P}\{D_G = d\}$ (pdf) and $\mathbb{P}\{D_G > d\}$ (ccdf), to evaluate and compare our proposed NBRW-rw and MHRW-DA (MHDA in Algorithm 2) against SRW-rw and MHRW (MH algorithm in Algorithm 1), respectively. As mentioned before, to estimate $\mathbb{P}\{D_G = d\}$, we just need to choose a function $f(i) = \mathbf{1}_{\{d(i)=d\}}$, $i \in \mathcal{N}$, for the corresponding estimators. Similarly, we choose $f(i) = \mathbf{1}_{\{d(i)>d\}}$ for $\mathbb{P}\{D_G > d\}$. To measure the estimation accuracy, we use the following normalized root mean square error (NRMSE) [5, 30, 20], $\sqrt{\mathbb{E}\{(\hat{x}(t) - x)^2\}}/x$, where $\hat{x}(t)$ is the estimated value out of t samples and x is the (ground-truth) real value. ($x = \lim_{t \rightarrow \infty} \hat{x}(t)$ from unbiasedness.) In all simulations, an initial position of each random walk is drawn from its stationary distribution as similarly used in [5], unless otherwise specified. In practical implementations, one can employ a ‘burn-in’ period to drive the random walk close to its steady-state [12]. Each data point reported here for AS-733 graph is obtained from 10^4 independent simulations, while, for Road-PA and Web-Google graphs, the data points are based on 10^5 and $5 \cdot 10^5$ simulations, respectively.

We first present the simulation results for AS-733 graph whose ‘actual’ degree distribution is almost a ‘power-law’ as depicted in Figure 5 (insets). Figure 4 shows that NBRW-rw



(a) SRW-rw vs. NBRW-rw (b) MHRW vs. MHRW-DA
Figure 4: AS-733 graph. NRMSE (averaged over all possible degrees d) of the estimator of $\mathbb{P}\{D_G = d\}$, when we vary the number of samples; the insets are for smaller samples.



(a) SRW-rw vs. NBRW-rw (b) MHRW vs. MHRW-DA

Figure 5: AS-733 graph. NRMSE ratio (per degree d) when estimating $\mathbb{P}\{D_G = d\}$ with 10^4 samples; the insets represent the ‘actual’ degree distribution (ccdf) in (a) log-log scale, (b) semi-log scale.

(resp. MHRW-DA) outperforms SRW-rw (resp. MHRW) in terms of the required number of samples (cost) to achieve the same level of estimation error when estimating $\mathbb{P}\{D_G = d\}$, as expected from our theoretical results. Here, the NBRW-rw (resp. MHRW-DA) brings out about 35% (resp. 14%) cost saving on average, when compared to the SRW-rw (resp. MHRW). In addition, we plot, in Figure 5, the NRMSE ratio of SRW-rw (resp. MHRW) to the case of NBRW-rw (resp. MHRW-DA) for every degree d when estimating $\mathbb{P}\{D_G = d\}$ with 10^4 samples. It clearly shows the improvement of our proposed methods for each degree d (all data points are above one). We also provide the NRMSE curve (with its ratio), in Figure 6, for the comparison between NBRW-rw (resp. MHRW-DA) and SRW-rw (resp. MHRW) when estimating $\mathbb{P}\{D_G > d\}$ with 10^4 samples, which is again clearly consistent with our theoretical findings. In addition, we conduct another simulation to see the impact of non-stationary start for each random walk on the sampling accuracy, for which an initial position of each SRW and NBRW is drawn from a uniform distribution, while the initial position for MHRW and MHRW-DA is chosen with a probability proportional to node degree. Under this setting, we measure NRMSE of the estimator of $\mathbb{P}\{D_G = d\}$, and observe that NBRW-rw and MHRW-DA still outperform SRW-rw and MHRW, respectively, as shown in Figure 7. Note that there is not much difference between the stationary start and non-stationary start cases. (See Figures 4 and 7.)

We next present the simulation results for Road-PA graph in which every node has small degree, ranging from 1 to 9, and the actual degree distribution (pdf) is given in Figure 9 (inset). As seen from Figure 8, SRW-rw (resp. MHRW) requires *more than twice larger samples* than the case of NBRW-rw (resp. MHRW-DA) to attain the same level of accuracy for the estimation of $\mathbb{P}\{D_G = d\}$. Specifically, the

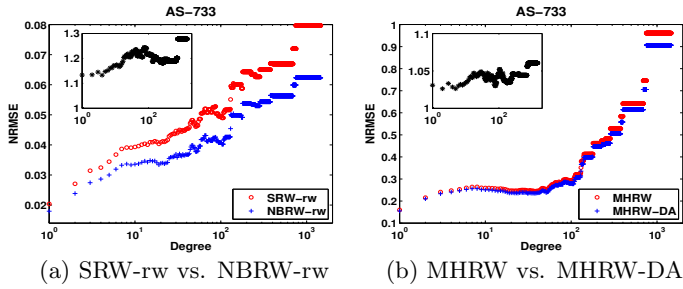


Figure 6: AS-733 graph. NRMSE (per degree d) when estimating $\mathbb{P}\{D_G > d\}$ with 10^4 samples; the insets show NRMSE ratio.




Figure 7: AS-733 graph. NRMSE (averaged over all possible d) of the estimator of $\mathbb{P}\{D_G = d\}$, when each random walk does not start in the stationary regime.

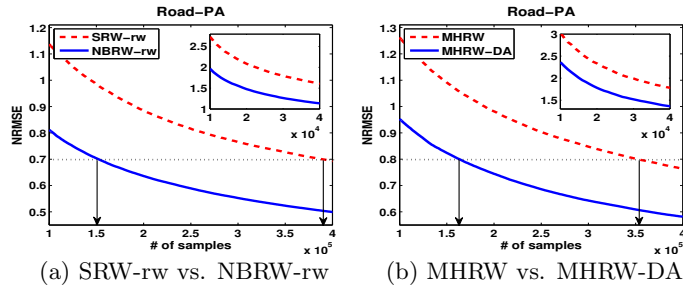


Figure 8: Road-PA graph. NRMSE (averaged over all possible d) of the estimator of $\mathbb{P}\{D_G = d\}$, while varying the number of samples; the insets are for smaller samples.

NBRW-rw and MHRW-DA leads to about 60% and 54% cost saving on average. Also, as before, Figure 9 shows the NRMSE ratio of SRW-rw (resp. MHRW) to the case of NBRW-rw (resp. MHRW-DA) for every degree d when estimating $\mathbb{P}\{D_G = d\}$ with $5 \cdot 10^5$ samples, and Figure 10 depicts the NRMSE curve (with its ratio) for the estimation of $\mathbb{P}\{D_G > d\}$ with $5 \cdot 10^5$ samples, which are all in good agreement with our theoretical findings. We observe that the NBRW-rw and MHRW-DA are remarkably effective for Road-PA graph, as the graph structure with small node degrees makes the less-backtracking feature more favorable.

We finally provide the simulation results for Web-Google graph whose actual degree distribution is more like a power-law as shown in Figure 12 (insets). Figure 11 demonstrates that NBRW-rw (resp. MHRW-DA) surpasses SRW-rw (resp. MHRW) overall for the estimation of $\mathbb{P}\{D_G = d\}$, although their improvements are not as large as before. Again, Figure 12 shows the NRMSE ratio of SRW-rw (resp. MHRW) to the case of NBRW-rw (resp. MHRW-DA) for every degree d in estimating $\mathbb{P}\{D_G = d\}$ with $5 \cdot 10^5$ samples, and Figure 13 depicts the NRMSE curve (with its ratio) for the estimation of $\mathbb{P}\{D_G > d\}$ with $5 \cdot 10^5$ samples. Clearly, NBRW-rw per-

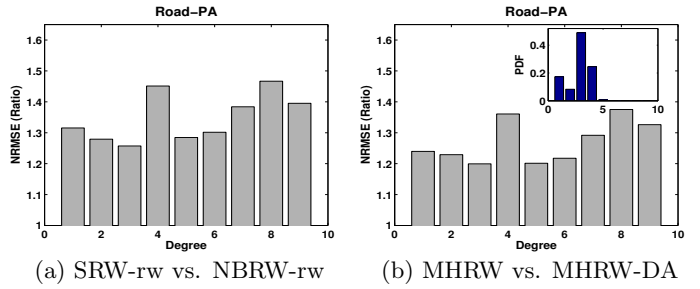


Figure 9: Road-PA graph. NRMSE ratio (per degree d) when estimating $\mathbb{P}\{D_G = d\}$ with $5 \cdot 10^5$ samples; the inset represents the ‘actual’ degree distribution (pdf).

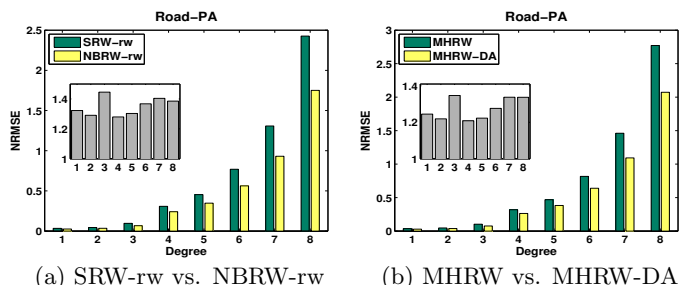


Figure 10: Road-PA graph. NRMSE (per degree d) for the estimation of $\mathbb{P}\{D_G > d\}$ with $5 \cdot 10^5$ samples; the insets show NRMSE ratio.

forms better than SRW-rw for each degree d (all data points are above one), as expected from our theoretical results. We also observe similar results when comparing MHRW-DA and MHRW. There is, however, just one data point below one (in the ratio) for the estimation of $\mathbb{P}\{D_G = d\}$. We admit that such an ‘outlier’ may be possible, since our theoretical results hold in the asymptotic sense. Nonetheless, MHRW-DA leads to an overall performance improvement (over all possible d). Due to space constraints, we refer to our technical report [21] for more simulation results including the impact of non-stationary start for each random walk on the estimation accuracy under Road-PA and Web-Google graphs.

It is also worth noting that a direct comparison between SRW-rw (or NBRW-rw) and MH algorithm (or MHDA) may not be appropriate. Recently, [12] numerically shows a counter-example that MHRW can be more efficient, although SRW-rw has been shown to be better than the MHRW over several numerical simulations [29, 12]. In addition, [6] proved, through several examples, that there is no clear winner between the importance sampling for reversible Markov chains (whose special case is the SRW-rw) and the MH algorithm. The MH algorithm also is valuable because it can be used to construct a reversible chain with any given stationary distribution. We thus have focused on improving each of the SRW-rw and the MH algorithm separately.

6. CONCLUDING REMARKS

We demonstrated, in theory and simulation, that our proposed NBRW-rw and MHDA guarantee unbiased graph sampling, while also achieving higher sampling efficiency than SRW-rw and MH algorithm, respectively. While the focus of this paper was on the unbiased graph sampling, we cannot stress enough the versatile applicability of the MHDA for *any* non-uniform node sampling (e.g., intentionally creating a known bias toward preferable nodes), not to mention

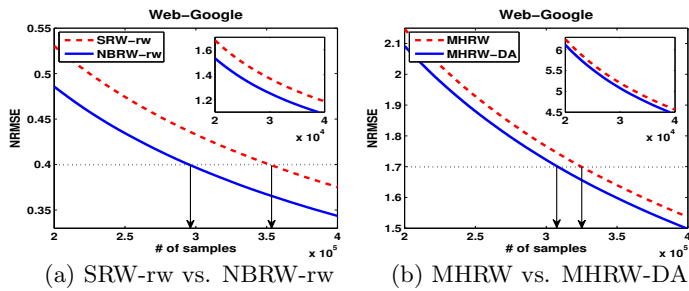


Figure 11: Web-Google graph. NRMSE (averaged over all possible d) of the estimator of $\mathbb{P}\{D_G = d\}$ with different number of samples; the insets are for smaller samples.

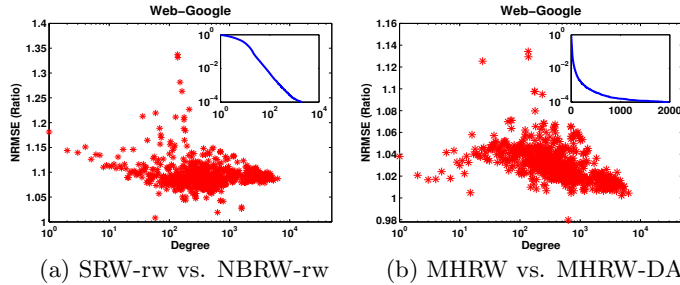


Figure 12: Web-Google graph. NRMSE ratio (per degree d) when estimating $\mathbb{P}\{D_G = d\}$ with $5 \cdot 10^5$ samples; the insets represent the ‘actual’ degree distribution (ccdf) in (a) log-log scale, (b) semi-log scale.

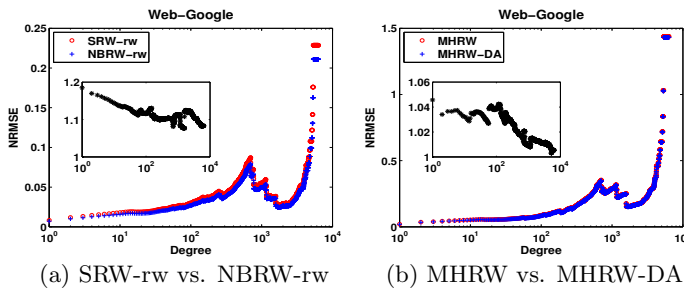


Figure 13: Web-Google graph. NRMSE (per degree d) when estimating $\mathbb{P}\{D_G > d\}$ with $5 \cdot 10^5$ samples; the insets show NRMSE ratio.

its improvement over the famous MH algorithm in sampling efficiency. We expect that the MHTDA can be applied to many other problems beyond the unbiased graph sampling.

Acknowledgement

The authors thank the anonymous reviewers and especially our shepherd, Don Towsley, for many constructive comments and suggestions that greatly improved the quality of this paper. This work was supported in part by NSF CNS-0545893.

7. REFERENCES

- [1] Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data/>.
- [2] D. Aldous and J. Fill. *Reversible Markov Chains and Random Walks on Graphs*. monograph in preparation.
- [3] N. Alon, I. Benjamini, E. Lubetzky, and S. Sodin. Non-backtracking random walks mix faster. *Communications in Contemporary Mathematics*, 9(4):585–603, 2007.
- [4] R. B. Ash and C. A. Doléans-Dade. *Probability and measure theory*. Academic Press, second edition, 2000.
- [5] K. Avrachenkov, B. Ribeiro, and D. Towsley. Improving random walk estimation accuracy with uniform restarts. In *WAW*, 2010.
- [6] F. Bassetti and P. Diaconis. Examples comparing importance

- sampling and the Metropolis algorithm. *Illinois Journal of Mathematics*, 50(1):67–91, 2006.
- [7] P. Berenbrink, C. Cooper, T. R. R. Elsässer, and T. Sauerwald. Speeding up random walks with neighborhood exploration. In *ACM SODA*, 2010.
- [8] S. Boyd, P. Diaconis, and L. Xiao. Fastest mixing markov chain on a graph. *SIAM Review*, 46(4):667–689, 2004.
- [9] F. Chen, L. Lovász, and I. Pak. Lifting markov chains to speed up mixing. In *ACM STOC*, 1999.
- [10] P. Diaconis, S. Holmes, and R. M. Neal. Analysis of a nonreversible markov chain sampler. *Annals of Applied Probability*, 10(3):726–752, 2000.
- [11] R. Douc and C. P. Robert. A vanilla Rao-Blackwellization of Metropolis-Hastings algorithms. *Annals of Statistics*, 39(1):261–277, 2011.
- [12] M. Gjoka, M. Kurant, C. T. Butts, and A. Markopoulou. Practical recommendations on crawling online social networks. *IEEE JSAC*, 2011.
- [13] S. Goel and M. J. Salganik. Respondent-driven sampling as Markov chain Monte Carlo. *Statistics in Medicine*, 28(17):2202–2229, 2009.
- [14] P. J. Green and A. Mira. Delayed rejection in reversible jump metropolis-hastings. *Biometrika*, 88(4):1035–1053, 2001.
- [15] M. A. Hasan and M. J. Zaki. Output space sampling for graph patterns. In *VLDB*, 2009.
- [16] W. K. Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.
- [17] S. Ikeda, I. Kubo, and M. Yamashita. The hitting and cover times of random walks on finite graphs using local degree information. *Theoretical Computer Science*, 410(1):94–100, January 2009.
- [18] G. L. Jones. On the Markov chain central limit theorem. *Probability Surveys*, 1:299–320, 2004.
- [19] K. Jung and D. Shah. Fast gossip via nonreversible random walk. In *IEEE ITW*, 2006.
- [20] M. Kurant, M. Gjoka, C. T. Butts, and A. Markopoulou. Walking on a graph with a magnifying glass: stratified sampling via weighted random walks. In *ACM SIGMETRICS*, 2011.
- [21] C.-H. Lee, X. Xu, and D. Y. Eun. Beyond random walk and Metropolis-Hastings samplers: Why you should not backtrack for unbiased graph sampling. Technical report, Dept. of ECE, North Carolina State University, April 2012.
- [22] D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov chains and mixing times*. American Mathematical Society, 2009.
- [23] W. Li and H. Dai. Accelerating distributed consensus via lifting markov chains. In *IEEE ISIT*, 2007.
- [24] S. Malefaki and G. Iliopoulos. On convergence of properly weighted samples to the target distribution. *Journal of Statistical Planning and Inference*, 138(4):1210–1225, 2008.
- [25] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21(6):1087–1092, 1953.
- [26] A. Mira. Ordering and improving the performance of monte carlo markov chains. *Statistical Science*, 16(4):340–350, 2001.
- [27] R. M. Neal. Improving asymptotic variance of MCMC estimators: non-reversible chains are better. Technical report, No. 0406, Dept. of Statistics, University of Toronto, July 2004.
- [28] P. H. Peskun. Optimum monte-carlo sampling using markov chains. *Biometrika*, 60:607–612, 1973.
- [29] A. H. Rasti, M. Torbjazi, R. Rejaie, N. Duffield, W. Willinger, and D. Stutzbach. Respondent-driven sampling for characterizing unstructured overlays. In *INFOCOM*, 2009.
- [30] B. Ribeiro and D. Towsley. Estimating and sampling graphs with multidimensional random walks. In *IMC*, 2010.
- [31] G. O. Roberts and J. S. Rosenthal. General state space Markov chains and MCMC algorithms. *Probability Surveys*, 1:20–71, 2004.
- [32] S. M. Ross. *Stochastic processes*. John Wiley & Son, second edition, 1996.
- [33] M. J. Salganik and D. D. Heckathorn. Sampling and estimation in hidden populations using respondent-driven sampling. *Sociological Methodology*, 34:193–239, 2004.
- [34] D. Stutzbach, R. Rejaie, N. Duffield, S. Sen, and W. Willinger. On unbiased sampling for unstructured peer-to-peer networks. *IEEE/ACM Transactions on Networking*, 17(2):377–390, 2009.
- [35] S. S. Wu and M. T. Wells. An extension of the metropolis algorithm. *Communications in Statistics – Theory and Methods*, 34(3):585–596, 2005.