On the Rao-Blackwellization and Its Application for Graph Sampling via Neighborhood Exploration

Chul-Ho Lee Xin Xu Do Young Eun

***Abstract*—We study how the so-called Rao-Blackwellization, which is a variance reduction technique via "conditioning" for Monte Carlo methods, can be judiciously applied for graph sampling through neighborhood exploration. Despite its popularity for Monte Carlo methods, it is little known for Markov chain Monte Carlo methods and has never been discussed for random walk-based graph sampling. We first propose two forms of Rao-Blackwellization that can be used as a swap-in replacement for virtually all (reversible) random-walk graph sampling methods, and prove that the 'Rao-Blackwellized' estimators reduce the (asymptotic) variances of their original estimators yet maintain their inherent unbiasedness. The variance reduction can translate into lowering the number of samples required to achieve a desired sampling accuracy. However, the sampling cost for neighborhood exploration, if required, may outweigh such improvement, even leading to higher total amortized cost. Considering this, we provide a generalization of Rao-Blackwellization, which allows one to choose a suitable extent of obtaining Rao-Blackwellized samples in order to achieve a right balance between sampling cost and accuracy. We finally provide simulation results via real-world datasets that confirm our theoretical findings.**

I. Introduction

Graph sampling via *random walk crawling* has been an effective means for sampling large complex networks, ranging from unstructured peer-to-peer networks to online social networks, to estimate network statistics such as degree distribution and clustering coefficient, and social-network properties, e.g., the fraction of users with some common property/interest and their relationships. In particular, the random walk-based graph sampling has been known as the sole vital solution for sampling online social networks, which can only be done via their public yet restrictive interfaces, only allowing local neighborhood queries [1], [2]. Such interface limitation clearly prevents from obtaining "independent" samples directly from networks or incurs very high cost using rejection sampling (to guess a user ID space) even if possible. In addition, it has been more preferable over the traditional graph-traversal algorithms for sampling, as it has many desirable properties with statistical guarantees [1]–[3] inherited from the popular Markov chain Monte Carlo (MCMC) techniques.

In an early stage of research in this area, the so-called Metropolis-Hastings random walk and reweighted random walk (or simple random walk with reweighting) methods [1]–[4], which now become off-the-shelf algorithms and used as baseline methods, were commonly used and extensively evaluated for their feasibility studies. Since then, the research has been diversified. One branch of studies is to "speed up" random walks, aiming at improving the estimation accuracy of random-walk samplers, or equivalently, lowering the number of samples needed to achieve a desired accuracy guarantee, by effectively overcoming the *slow-diffusion* problem of main random walks of the two popular methods. For instance, the work in [5] proposes weighted random walks having a (known) bias toward more relevant nodes, which is then corrected by a usual reweighting procedure. Our earlier work in [6] introduces the benefit of non-reversible Markov chains to create more 'directional' walks. On the other hand, while a target sampling function f is often up to designers or developers as a design choice and can be virtually anything, another branch of research is on how to carefully design such a function so as to uncover important network properties that are not just nodal and edge properties but more complicated topological properties. For instance, there have been recent studies that go beyond the nodal or edge properties and propose new estimators (with new f) to estimate subgraph patterns such as network motifs and graphlets [7], [8].

In this paper, we take an orthogonal yet uncharted angle. In contrast to the previous studies, we demonstrate that for a given sampling function f, it is *always* possible to find its related function that, in turn, improves the accuracy of the original estimator while maintaining its inherent unbiasedness. Our rationale behind the improvement is to exploit the *neighborhood information* of each node visited into the estimation, without modifying the trajectories of random-walk samplers. In a similar vein, there are a few recent studies, which take advantage of knowing the neighbors' information to improve the accuracy of their estimators [9], [10] or literally speed up random walks from the perspective of the hitting and cover times [11], [12]. However, the former studies require either "independent" samples [9] or a specific estimator running on top of simple random walk with reweighting [10], thereby limiting their applicability. The latter studies have no direct implication for unbiased graph sampling [11], [12].

Formally speaking, we study how the so-called Rao-Blackwellization [13] – a variance reduction technique for Monte Carlo methods (requiring independent samples), can be applied for random walk-based graph sampling. The underlying idea is, instead of using an empirical average of samples, to make use of averaging the samples' associated *conditional expected values*. The variance reduction then follows from the fact that "conditioning" reduces the variance while keeping the same expected value. While it is little known for MCMC methods [14], [15] and has never been discussed for graph

Chul-Ho Lee is with the Department of Electrical and Computer Engineering, Florida Institute of Technology, Melbourne, FL, and Xin Xu and Do Young Eun are with the Department of Electrical and Computer Engineering, North Carolina State University, Raleigh, NC. Email: clee@fit.edu; {xxu5, dyeun}@ncsu.edu. This work was supported in part by National Science Foundation under Grant Nos. CNS-1217341 and CNS-1423151.

sampling, we propose two forms of Rao-Blackwellization tailored for graph sampling for the first time in the literature, utilizing the transition probabilities of the random-walk samplers along with neighborhood information for conditioning. We also prove that our 'Rao-Blackwellized' estimators achieve *smaller* (asymptotic) variance while maintaining their inherent unbiasedness even with correlated samples returned from the underlying random-walk samplers. Our Rao-Blackwellized estimators are general enough and can be applied as a quick swap-in replacement for *virtually all* (reversible) random-walk graph sampling methods.

To translate our theoretical findings into usable algorithms in practice, we further consider a likely scenario where neighborhood exploration incurs higher cost for sampling than the baseline method, for which our Rao-Blackwellized estimators would result in fewer samples starting from the same total sampling budget. To address this issue, we generalize the two forms of Rao-Blackwellization so that random-walk samplers can make an *independent* decision on the use of Rao-Blackwellization at *each* node visited, and then prove their unbiasedness and provide variance analysis for the generalization. We show that, contrary to common belief, there exists a scenario in which partially utilizing the neighborhood information via Rao-Blackwellization can yield worse performance than not using the neighborhood information at all. We also explain how different sampling methods, providing varying qualities of samples yet with different costs, can be compared on the same ground, and introduce a *refined* form of the variance, re-scaled by a penalty for higher cost per sample (or a benefit for lower cost per sample). Our generalization and the refined form for the asymptotic variance allow one to seek for a more balanced tradeoff between better quality of Rao-Blackwellized samples and higher sampling cost. We finally provide simulation results obtained over real-world network datasets to support our analysis.

II. Preliminaries

We provide a mathematical background for random-walk graph sampling methods and review a variance reduction technique for Monte Carlo methods called *Rao-Blackwellization* to set the stage for the subsequent exposition.

A. Markov Chain Theory for Graph Sampling

Consider a connected, undirected and non-bipartite graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ to model a (social) network for sampling, with a set of nodes (users) $\mathcal{N} = \{1, 2, ..., n\}$ and a set of edges $\mathcal{E}$ (indicating neighboring relationships between users). The graph $\mathcal{G}$ is assumed to have no self-loops and no multiple edges. Let $N(i) = \{j \in \mathcal{N} : (i, j) \in \mathcal{E}\}$ be a set of neighbors of node i with degree $d(i) = |N(i)|$. The main operation of random-walk graph sampling methods is to launch a "random-walk" agent or multiple agents to crawl $\mathcal{G}$ in order to collect its representative samples and, in turn, unbiasedly estimate its nodal or network statistics.

Consider a discrete-time random walk (or Markov chain) $\{X_t \in \mathcal{N}\}_{t \geq 0}$ on $\mathcal{G}$, representing a sequence of nodes visited by the sampling agent according to transition probabilities $P(i, j) = \mathbb{P}[X_{t+1} = j | X_t = i]$, which form a transition matrix $\mathbf{P}$. We assume that $\{X_t\}$ is an ergodic (i.e., irreducible and aperiodic) chain with its unique stationary distribution $\boldsymbol{\pi} = [\pi(1), \pi(2), \ldots, \pi(n)]$. Similar to MCMC, the basic idea of graph sampling via random-walk crawling is approximating or estimating an expectation, for a *target function* $f : \mathcal{N} \to \mathbb{R}$,

$$\mathbb{E}_{\boldsymbol{u}}[f(X)] = \sum_{i \in \mathcal{N}} f(i) \frac{1}{n}$$

by an empirical average (or sample mean)

$$\hat{\mu}_t(f) = \frac{1}{t} \sum_{k=1}^{t} f(X_k) \tag{1}$$

over a realization (or sample path) of the random walk (sampling agent) $X_1, X_2, \ldots, X_t$ having $\boldsymbol{\pi} = \boldsymbol{u}$ as its stationary distribution, where $\boldsymbol{u} = [1/n, 1/n, \ldots, 1/n]$. The most famous method of this type is the Metropolis-Hastings random walk method [1], [2], [4], [6] which is an application of the Metropolis-Hastings algorithm [16], [17]. Regarding the target function f, if the goal of sampling is to estimate the average degree of nodes in $\mathcal{G}$, one can use a function $f(i) = d(i)$, leading to the target quantity $\mathbb{E}_{\boldsymbol{u}}[d(X)] = 2|\mathcal{E}|/n$. Also when estimating a fraction of nodes belonging to a certain 'membership' class A, an indicator function $f(i) = \mathbb{I}_A(i)$ can be used with noting that $\mathbb{E}_{\boldsymbol{u}}[\mathbb{I}_A(X)] = \mathbb{P}[i \in A]$, where $\mathbb{I}_A(i) = 1$ if $i \in A$ and $\mathbb{I}_A(i) = 0$ if $i \notin A$.

If the stationary distribution $\boldsymbol{\pi}$ is, however, not uniform, a weighted average (or a ratio estimator) is used instead of (1) as follows.*

$$\frac{\hat{\mu}_t(\psi)}{\hat{\mu}_t(w)} = \frac{\sum_{k=1}^{t} w(X_k) f(X_k)}{\sum_{k=1}^{t} w(X_k)}, \tag{2}$$

where 'importance' weights $w(i), i \in \mathcal{N}$, are generally given by $w(i) = 1/\pi(i)$ or $w(i) = c/\pi(i)$ with some constant c because they only need to be evaluated up to a multiplicative constant. For simplicity, we define a function $\psi : \mathcal{N} \to \mathbb{R}$ to denote the entrywise product of f and w, i.e., $\psi(i) = w(i) f(i), i \in \mathcal{N}$. Note that such proper weights are necessary to correct a *bias* introduced from the non-unform stationary distribution. The so-called reweighted random walk or simple random walk with reweighting [1], [2], [4], [6] is a popular example of this kind, with $w(i) = 1/d(i)$. In fact, any *unbiased* graph sampling methods adopting 'non-uniform' random walks belong to this class.

We below explain two important theorems for graph sampling, namely ergodic theorem and central limit theorem for Markov chains. The former has been a fundamental basis for the "asymptotic unbiasedness" of the estimator $\hat{\mu}_t(f)$ in (1) or $\hat{\mu}_t(\psi)/\hat{\mu}_t(w)$ in (2), while the latter has recently become popular in the graph sampling literature to mathematically (not just through numerical simulations) evaluate the quality or *efficiency* of the estimator [6], [8], [18], [21].

*It is sometimes referred to as the Hansen-Hurwitz estimator [1], [2], [5], [18], and is also called respondent-driven sampling in other disciplines [19], [20]. It is basically a form of MCMC importance sampling.

For an ergodic chain $\{X_t\}$ with its stationary distribution $\boldsymbol{\pi}$, the ergodic theorem [22] says that, for any bounded function $f : \mathcal{N} \to \mathbb{R}$,

$$\lim_{t\to\infty} \hat{\mu}_t(f) = \mathbb{E}_{\boldsymbol{\pi}}\left[f(X)\right] = \sum_{i\in\mathcal{N}} f(i)\pi(i) \quad \text{a.s.} \tag{3}$$

for any initial distribution on X_0. This implies that if $\boldsymbol{\pi} = \boldsymbol{u}$, $\hat{\mu}_t(f)$ is asymptotically *unbiased* for our target quantity $\mathbb{E}_{\boldsymbol{u}}\left[f(X)\right]$ (for unbiased graph sampling) as $t \to \infty$. Otherwise, the ratio estimator $\hat{\mu}_t(\psi)/\hat{\mu}_t(w)$ can be used instead and gives an unbiased estimate of $\mathbb{E}_{\boldsymbol{u}}\left[f(X)\right]$.

Moreover, the central limit theorem [23], [24] says that for any bounded function f and any initial distribution on X_0,

$$\sqrt{t}\left(\hat{\mu}_t(f) - \mathbb{E}_{\boldsymbol{\pi}}\left[f(X)\right]\right) \xrightarrow{d} N(0,\sigma^2) \quad \text{as } t \to \infty, \tag{4}$$

where $N(0,\sigma^2)$ is a normal (or Gaussian) random variable with zero mean and variance σ^2, and $\xrightarrow{d}$ denotes convergence in distribution. Here the *asymptotic variance* σ^2 is given by

$$\sigma^2 = \lim_{t\to\infty} t \cdot \mathrm{Var}[\hat{\mu}_t(f)] = \gamma_0 + 2\sum_{k=1}^{\infty}\gamma_k, \tag{5}$$

where $\gamma_0 = \mathrm{Var}_{\boldsymbol{\pi}}[f(X_0)]$ is the marginal variance of $f(X_0)$ with respect to $\boldsymbol{\pi}$ and $\gamma_k = \mathrm{Cov}_{\boldsymbol{\pi}}[f(X_0), f(X_k)]$ is the lag-k autocovariance of the *stationary* sequence $\{f(X_t)\}$. Note that the central limit theorem can also be similarly applied for the ratio estimator $\hat{\mu}_t(\psi)/\hat{\mu}_t(w)$ [6], [20], which shall be used in our analysis later. The asymptotic variance highly depends on the function f and the transition matrix $\mathbf{P}$ of the chain, so we use $\sigma^2(f,\mathbf{P})$ instead of σ^2 for the rest of this paper. Similarly we use $\gamma_k(f)$ instead of γ_k for all $k \geq 0$ to signify their dependence on f. On the contrary, for notational simplicity, we use $\mathbb{E}_{\boldsymbol{\pi}}(f)$ for $\mathbb{E}_{\boldsymbol{\pi}}\left[f(X)\right]$, unless otherwise necessary.

From the accuracy point of view, the smaller the variability (or variance) of an unbiased estimator, the better its accuracy (or efficiency). The variance reduction can have a large effect on the number of samples that need to be taken to obtain a desired accuracy guarantee. While asymptotic variance is the (rescaled) variance of an estimator in the limit, which allows one to mathematically evaluate the accuracy of the estimator, the asymptotic variance has also been a good indicator for the performance with *finite* samples. Thus, we consider the asymptotic variance as a main performance metric in this work, and focus on devising a very general method that can be used as a swap-in replacement for a large class of random-walk graph sampling methods in order to reduce their (asymptotic) variances while maintaining their inherent unbiasedness.

B. Rao-Blackwellization for Monte Carlo methods

Rao-Blackwellization is a generic approach to reducing the variance of a Monte Carlo estimator, which is to use "conditioning" [13], [25]. We first collect two simple facts regarding the conditioning. For two (possibly correlated) random variables X and Y, it is well known that $\mathbb{E}[\mathbb{E}[X|Y]] = \mathbb{E}[X]$, where $\mathbb{E}[X|Y]$ is the conditional expectation of X given Y, which is itself a random variable (a function of Y) [26]. In addition, the conditional variance formula [26] says

$$\mathrm{Var}[X] = \mathrm{Var}[\mathbb{E}[X|Y]] + \mathbb{E}[\mathrm{Var}[X|Y]]$$

implying that

$$\mathrm{Var}[X] \geq \mathrm{Var}[\mathbb{E}[X|Y]], \tag{6}$$

from $\mathbb{E}[\mathrm{Var}[X|Y]] \geq 0$, since a variance is always non-negative. In words, by conditioning we can always reduce the variance while keeping the same expected value.

Suppose now that our target distribution is a joint distribution $\boldsymbol{\pi}$ of two variables (X,Y) and we are interested in estimating the expectation of a function f of one variable X, i.e., $\mathbb{E}_{\boldsymbol{\pi}}[f(X)]$. If $i.i.d.$ samples $X_1, X_2, \ldots, X_t$ are drawn directly from the marginal distribution of X, then a straightforward estimator is

$$\bar{\mu}_t^1(f) = \frac{1}{t}\sum_{k=1}^{t} f(X_k).$$

Now consider the following Rao-Blackwellized estimator with $i.i.d.$ samples $Y_1, Y_2, \ldots, Y_t$ drawn from the marginal distribution of Y:

$$\bar{\mu}_t^2(f) = \frac{1}{t}\sum_{k=1}^{t} \mathbb{E}[f(X)|Y_k].$$

Clearly, both are *unbiased*, because $\mathbb{E}[\bar{\mu}_t^1(f)] = \mathbb{E}_{\boldsymbol{\pi}}[f(X)]$ and

$$\mathbb{E}[\bar{\mu}_t^2(f)] = \mathbb{E}_{\boldsymbol{\pi}}[\mathbb{E}[f(X)|Y]] = \mathbb{E}_{\boldsymbol{\pi}}[f(X)].$$

However, the variance of $\bar{\mu}_t^2(f)$ can only be lower, i.e.,

$$\mathrm{Var}[\bar{\mu}_t^1(f)] = \frac{\mathrm{Var}_{\boldsymbol{\pi}}[f(X)]}{t} \geq \frac{\mathrm{Var}_{\boldsymbol{\pi}}[\mathbb{E}[f(X)|Y]]}{t} = \mathrm{Var}[\bar{\mu}_t^2(f)],$$

since we assume (for now) that $\{X_k\}$ are $i.i.d.$ and the inequality is from (6). This procedure is often referred to as Rao-Blackwellization [13], [25]. As long as the conditional expectation $\mathbb{E}[X|Y]$ is available analytically and Y is easy to sample, one can enjoy the Rao-Blackwellization, which reduces the variance of an unbiased estimator while keeping its unbiasedness.

III. Applying Rao-Blackwellization for Graph Sampling

The Rao-Blackwellization technique has widely been used as an effective variance-reduction technique for Monte Carlo methods (e.g., [13] and references thereafter) and even applied for sampling of search engine quality metrics [27], which *all* require "independent" samples. It is, however, little known for MCMC methods, or effectively for "correlated" samples, and in this regard, a few studies have only been done in Statistics, along which it was shown that the transition probabilities of Markov chains (so correlated samples) can be used for the conditioning of Rao-Blackwellization [14], [15]. In order to apply their finding to graph sampling, however, the stationary distribution of random walk samplers needs be *uniform*, clearly limiting its applicability.

Thus, we below show how the Rao-Blackwellization can be generally applied for graph sampling even when the stationary distribution of random walk samplers is *non-uniform*,

which is much more popular in the existing graph sampling methods [1], [2], [5]–[7], [28]. In particular, we prove the asymptotic unbiasedness and *smaller* asymptotic variance of the Rao-Blackwellized estimators. This improvement can still be achieved in practice, since the transition probabilities of an underlying random walk (or Markov chain) are used for conditioning and such knowledge is obtainable via neighborhood exploration [11], [12] or when neighborhood information is available [9], [10].

A. Uniform Stationary Distribution

Consider a sampling agent that moves over $\mathcal{G}$ according to the transition matrix $\mathbf{P}$ of an ergodic, reversible[†] Markov chain $\{X_t\}_{t\geq 0}$, where X_t is the location of the agent at time t. We first consider the case of a *uniform* stationary distribution $\boldsymbol{\pi}=\boldsymbol{u}$. For any bounded target function f, we define another function $\mathbf{P}f:\mathcal{N}\to\mathbb{R}$ such that

$$\mathbf{P}f(i)=\mathbb{E}[f(X_{k+1})|X_k=i]=\sum_{j\in\mathcal{N}}P(i,j)f(j),\ i\in\mathcal{N}. \quad (7)$$

Here we use the notation $\mathbf{P}f$ in order to indicate that it can be interpreted as an n-dimensional (column) vector which is the product of an $n\times n$ matrix $\mathbf{P}$ and an n-dimensional vector f. From the definition of $\mathbf{P}f$, we can then see that

$$\mathbb{E}_{\boldsymbol{\pi}}[\mathbf{P}f(X_k)]=\mathbb{E}_{\boldsymbol{\pi}}\left[\mathbb{E}[f(X_{k+1})|X_k]\right]=\mathbb{E}_{\boldsymbol{\pi}}[f(X)], \quad (8)$$

since both X_k and X_{k+1} have the same distribution $\boldsymbol{\pi}$ in the stationary regime.

We consider the following Rao-Blackwellized estimator of $\mathbb{E}_{\boldsymbol{\pi}}[f(X)]$, which was first introduced in [14], [15].

$$\hat{\mu}_t(\mathbf{P}f)=\frac{1}{t}\sum_{k=1}^{t}\mathbf{P}f(X_k). \quad (9)$$

It follows from the ergodic theorem for Markov chains in (3) with a function f replaced by $\mathbf{P}f$ that $\hat{\mu}_t(\mathbf{P}f)$ is asymptotically *unbiased*, i.e.,

$$\lim_{t\to\infty}\hat{\mu}_t(\mathbf{P}f)=\mathbb{E}_{\boldsymbol{\pi}}[\mathbf{P}f(X_k)]=\mathbb{E}_{\boldsymbol{\pi}}[f(X)]\ \text{ a.s.}, \quad (10)$$

where the second equality is from (8).

We next discuss how efficient the Rao-Blackwellized estimator $\hat{\mu}_t(\mathbf{P}f)$ would be by comparing its variance to that of the original estimator $\hat{\mu}_t(f)$. We first look at the *marginal* variance or 'quality' of *each* of correlated samples used in the estimator, assuming that samples are marginally drawn from $\boldsymbol{\pi}$ while possibly correlated, and show that the marginal variance under the Rao-Blackwellized estimator can only be smaller, as was the case of $i.i.d.$ samples in Section II-B.

To proceed, we collect notations and simple facts in the theory of ergodic, *reversible* Markov chains [22]. For functions $f,g:\mathcal{N}\to\mathbb{R}$, define their scalar product with respect to $\boldsymbol{\pi}$ as $\langle f,g\rangle_{\boldsymbol{\pi}}=\sum_{i\in\mathcal{N}}f(i)g(i)\pi(i)$. For a reversible chain $\{X_t\}$ and functions f,g, we have the so-called 'self-adjoint' property

$$\langle f,\mathbf{P}g\rangle_{\boldsymbol{\pi}}=\sum_{i\in\mathcal{N}}f(i)(\mathbf{P}g)(i)\pi(i)=\sum_{i\in\mathcal{N}}f(i)\sum_{j\in\mathcal{N}}P(i,j)g(j)\pi(i)$$
$$=\sum_{j\in\mathcal{N}}\sum_{i\in\mathcal{N}}P(j,i)f(i)g(j)\pi(j)=\langle\mathbf{P}f,g\rangle_{\boldsymbol{\pi}} \quad (11)$$

where the second last equality is from the reversibility. Iteratively applying the self-adjoint property yields

$$\langle\mathbf{P}^{m_1}f,\mathbf{P}^{m_2}g\rangle_{\boldsymbol{\pi}}=\langle\mathbf{P}^{m_1+m_2}f,g\rangle_{\boldsymbol{\pi}}=\langle f,\mathbf{P}^{m_1+m_2}g\rangle_{\boldsymbol{\pi}} \quad (12)$$

for any positive integers m_1,m_2, where $\mathbf{P}^m$ is the m-step transition matrix (or the m-th power of the transition matrix). We then have the following. Due to space constraint, we refer to our technical report [29] for all the proofs in this paper.

Lemma 1: $\mathrm{Var}_{\boldsymbol{\pi}}[\mathbf{P}f(X_k)]\leq\mathrm{Var}_{\boldsymbol{\pi}}[f(X)]$ for any f. □

While the marginal variance of each sample in the Rao-Blackwellized estimator is lower, its smaller asymptotic variance is not guaranteed due to the (often complicated) covariance structure of $\{f(X_t)\}$, which differs significantly from the case of $i.i.d.$ samples. Thanks to the result in [15], however, the form of Rao-Blackwellization in (9) indeed brings an improvement in the asymptotic variance. Let $\sigma^2(f,\mathbf{P})$ and $\sigma^2(\mathbf{P}f,\mathbf{P})$ be the asymptotic variances of the estimators $\hat{\mu}_t(f)$ and $\hat{\mu}_t(\mathbf{P}f)$, respectively. We below restate the result in [15].

Theorem 1: For any function f, $\sigma^2(\mathbf{P}f,\mathbf{P})\leq\sigma^2(f,\mathbf{P})$, and the amount of reduction is given by

$$\sigma^2(f,\mathbf{P})-\sigma^2(\mathbf{P}f,\mathbf{P})=\mathrm{Var}_{\boldsymbol{\pi}}[f(X_0)+\mathbf{P}f(X_0)]$$
$$=\gamma_0(f)+2\gamma_1(f)+\gamma_2(f)\geq 0. \quad \square$$

As an example, consider the Metropolis-Hastings random walk (MHRW) method, having $\boldsymbol{\pi}=\boldsymbol{u}$. For a MHRW sampler, assuming the current node i, one of its neighbors is proposed as the next node to move to, with probability $Q(i,j)=1/d(i)$. This proposed transition is accepted with probability $A(i,j)=\min\{1,d(i)/d(j)\}$, and rejected with $1-A(i,j)$ in which case the agent stays at node i. Thus, the transition probabilities are given by

$$P(i,j)=\begin{cases}\min\{1/d(i),1/d(j)\}, & \text{if } j\in N(i)\\ 0, & \text{if } j\notin N(i) \text{ and } j\neq i,\end{cases}$$

with $P(i,i)=1-\sum_{j\neq i}P(i,j)$. The resulting chain is reversible with respect to $\boldsymbol{\pi}=\boldsymbol{u}$. Therefore, the Rao-Blackwellization can be done with the estimator $\hat{\mu}_t(\mathbf{P}f)$, as long as neighborhood information such as the degrees of neighbors and the values of f at the neighbors is available when visiting node i.

The theoretical guarantees provided by Theorem 1 clearly allow one to benefit from 'Rao-Blackwellizing' the MHRW estimator.[‡] It has, however, been shown in many studies [1], [2], [4], [6] that the estimates obtained by the so-called reweighted random walk (RRW) method are *consistently* more accurate than the MHRW method, so it has been recommended over the MHRW method for practical purposes [2], [4]. Then a natural follow-up question is "How can we apply such Rao-Blackwellization to the RRW estimator, or more generally, any (reversible) random-walk estimators requiring the correction

[†]A Markov chain is said to be *reversible* with respect to $\boldsymbol{\pi}$, if it satisfies the reversibility condition, i.e., $\pi(i)P(i,j)=\pi(j)P(j,i)$ for all $i,j\in\mathcal{N}$.

[‡]It it worth noting that Theorem 1 holds for *any* stationary distribution $\boldsymbol{\pi}$ [15]. With the estimator $\hat{\mu}_t(\mathbf{P}f)$ in (9), its (direct) applicability is, however, only limited to the case of $\boldsymbol{\pi}=\boldsymbol{u}$ for *unbiased* graph sampling.

of a bias from $\boldsymbol{\pi} \neq \boldsymbol{u}$?". We below tackle this challenging question. As an affirmative answer, we first propose a Rao-Blackwellized *ratio* estimator as a modification of the *pure* ratio estimator in (2) for the first time in the literature, and prove that our proposed estimator indeed guarantees asymptotic unbiasedness yet smaller asymptotic variance than the one in (2).

B. Non-Uniform Stationary Distribution

We turn our attention to the boarder yet important case of $\boldsymbol{\pi} \neq \boldsymbol{u}$, for which it is *unknown* how to obtain a Rao-Blackwellized estimate of $\mathbb{E}_{\boldsymbol{u}}(f)$, not $\mathbb{E}_{\boldsymbol{\pi}}(f)$. To this end, for any bounded target function f, we propose the following Rao-Blackwellized ratio estimator of $\mathbb{E}_{\boldsymbol{u}}(f)$:

$$\frac{\hat{\mu}_t(\mathbf{P}\psi)}{\hat{\mu}_t(\mathbf{P}w)} = \frac{\sum_{k=1}^{t} \mathbf{P}\psi(X_k)}{\sum_{k=1}^{t} \mathbf{P}w(X_k)} = \frac{\sum_{k=1}^{t} \sum_{j\in N(X_k)} P(X_k,j)w(j)f(j)}{\sum_{k=1}^{t} \sum_{j\in N(X_k)} P(X_k,j)w(j)}, \quad (13)$$

with weights $w(i) = c/\pi(i)$ for some constant $c > 0$. Recall that $\psi(i) = w(i)f(i)$. Let $\sigma^2_{\boldsymbol{w}}(\mathbf{P}f, \mathbf{P})$ and $\sigma^2_{\boldsymbol{w}}(f, \mathbf{P})$ be the asymptotic variances of $\hat{\mu}_t(\mathbf{P}\psi)/\hat{\mu}_t(\mathbf{P}w)$ and $\hat{\mu}_t(\psi)/\hat{\mu}_t(w)$, respectively. We below state our main result on the Rao-Blackwellized ratio estimator of $\mathbb{E}_{\boldsymbol{u}}(f)$.

Theorem 2: For any function f and weights $w(i) = c/\pi(i)$ with a constant $c > 0$, $\hat{\mu}_t(\mathbf{P}\psi)/\hat{\mu}_t(\mathbf{P}w)$ is asymptotically unbiased for $\mathbb{E}_{\boldsymbol{u}}(f)$ and $\sigma^2_{\boldsymbol{w}}(\mathbf{P}f, \mathbf{P}) \leq \sigma^2_{\boldsymbol{w}}(f, \mathbf{P})$. Furthermore, the amount of variance reduction is given by

$$\sigma^2_{\boldsymbol{w}}(f, \mathbf{P}) - \sigma^2_{\boldsymbol{w}}(\mathbf{P}f, \mathbf{P}) = \mathrm{Var}_{\boldsymbol{\pi}}[h(X_0) + \mathbf{P}h(X_0)] \geq 0, \quad (14)$$

where the function $h : \mathcal{N} \to \mathbb{R}$ is defined as

$$h(i) = [\psi(i) - w(i)\mathbb{E}_{\boldsymbol{u}}(f)]/(cn), \quad i \in \mathcal{N}. \quad (15)$$

As an example of the 'non-uniform' random-walk sampling methods, consider the reweighted random walk (RRW) method that uses a simple random walk (SRW) for moving over $\mathcal{G}$, followed by a post-hoc reweighting procedure [1], [2], [4], [6]. For a SRW sampler, at the current node i, the next node is chosen *uniformly at random* among i's neighbors. The resulting Markov chain is reversible with stationary probabilities $\pi(i) = d(i)/(2|\mathcal{E}|)$. To correct the bias toward high-degree nodes, the RRW method is essentially using the ratio estimator in (2) with weights $w(i) = 1/d(i)$. Therefore, our proposed Rao-Blackwellized estimator for the RRW method becomes

$$\frac{\hat{\mu}_t(\mathbf{P}\psi)}{\hat{\mu}_t(\mathbf{P}w)} = \frac{\sum_{k=1}^{t} \sum_{j\in N(X_k)} [f(j)/d(j)]/d(X_k)}{\sum_{k=1}^{t} \sum_{j\in N(X_k)} [1/d(j)]/d(X_k)}, \quad (16)$$

which is guaranteed to be asymptotically unbiased and have smaller asymptotic variance than the original one due to Theorem 2. This improvement again can be achieved by exploiting neighborhood information, as was the case for 'Rao-Blackwellizing' the MHRW method.

It is interesting to note that the estimator in [10] happens to be the same as the one in (16) with $f(i) = \mathbb{I}_A(i)$ for some label A. While their estimator is fundamentally different and uses different reported values for sampling, after the reweighting procedure, its final form coincides with (16).[§] Thus, the estimator in [10] can now be formally interpreted under our Rao-Blackwellization framework. In particular, as a byproduct of our Theorem 2, the estimator in [10] is also guaranteed to have smaller (asymptotic) variance, which was not shown in [10].

We also emphasize that our proposed Rao-Blackwellization procedure and its resulting theoretical guarantees (Theorem 2) are not limited to only the RRW method, but can directly be applied to virtually all 'non-uniform' random-walk-based methods (e.g., [5], [7], [28]), each of which employs a reversible random walk (Markov chain) along with a reweighting procedure, correcting a bias from its non-uniform stationary distribution. Even more importantly, our results can equally work for the so-called respondent-driven sampling (or more broadly MCMC importance sampling) and its applications in different disciplines (e.g., [19], [20]), although our focus here is limited to unbiased graph sampling.

We have discussed two forms of Rao-Blackwellization (for uniform and non-uniform $\boldsymbol{\pi}$), which are applicable for a very large class of random-walk graph sampling methods, with theoretical guarantees on their first-/second-order behaviors. An essential prerequisite here is that neighborhood information can be readily available at each node [9], [10] or the cost of neighborhood exploration can be ignored [11], [12] for Rao-Blackwellization, in which case the resulting variance reduction translates into lowering the number of samples to achieve a desired sampling accuracy. However, if such neighborhood information is no longer freely available and/or if one has to manually explore all the neighbors at every node visited for Rao-Blackwellization, then it can incur high computational and/or query cost to obtain one sample. Such an increase in cost per sample may offset the reduction in the number of samples (or in the asymptotic variance), even leading to higher total amortized cost. Therefore, Rao-Blackwellization should be utilized *judiciously*.

The problem can boil down to choosing whether to keep performing Rao-Blackwellization or not *throughout* for sampling, considering the sampling costs. However, since such a hard selection between two options greatly limits our design choices, a natural extension would be to generalize the two forms of Rao-Blackwellization so that the sampler can make an *independent* decision on the use of Rao-Blackwellization at *each* node visited. By doing so, one can control the level of neighborhood exploration or the amount of Rao-Blackwellized samples to manage the total expenditure of sampling costs with the hope to still enjoy the benefit of neighboring information.

[§]To see this a bit further, consider the numerator of each estimator. As mentioned in [10], their estimator originally follows from the one in [9] for independent samples, which has noting to do with Rao-Blackwellization, and uses a sampled value at node i given by $\sum_{j\in N(i)} f(j)/d(j)$, followed by reweighting with $w(i) = 1/d(i)$. In contrast, our sampled value at node i is simply $\mathbf{P}\psi(i)$, which is an *average* of its neighbors' (already re-weighted) values $\sum_{j\in N(i)} w(j)f(j)$.

IV. Neighborhood Exploration: From Theory To Practice

We introduce a generalization of Rao-Blackwellization and show its mathematical properties, including the asymptotic unbiasedness and non-trivial variance performance. We also discuss how different instances of this generalization, generating varying qualities of samples with different costs, can be compared on the same ground by taking into account the sampling costs into the variance computation, so that a 'right' instance of the generalization can be selected under more practical cost constraint.

A. Mathematical Model

As before, consider that a random-walk sampler crawls over $\mathcal{G}$ according to the transition matrix $\mathbf{P}$ of an ergodic, reversible chain $\{X_t\}$ with its stationary distribution $\boldsymbol{\pi}$. As a generalization of Rao-Blackwellization, when arriving at node i, the random-walk sampler can perform the 'Rao-Blackwellized' sampling, i.e., obtaining $\mathbf{P}f(i)$, with probability $q(i)$, and the original sampling, i.e., obtaining $f(i)$, with probability $1-q(i)$. While the trajectory of the sampler remains the same, the sampler can now make an independent decision of randomly choosing a node sample from either $\mathbf{P}f(i)$ or $f(i)$ at each node i visited. This generalization greatly extends our design choices, with different location-dependent probabilities $q(i)$.

For any bounded target function f, define a related function $\tilde{f}:\mathcal{N}\times\{0,1\}\to\mathbb{R}$ such that

$$\tilde{f}(i,0)=f(i),\ \text{and}\ \tilde{f}(i,1)=\mathbf{P}f(i),\ i\in\mathcal{N}.$$

To capture the location-dependent sampling, given the current position $X_t=i$ of the sampler, we define by Y_t an independent Bernoulli random variable, taking the value 1 with probability $q(i)$ and the value 0 with $1-q(i)$. Then, letting $\boldsymbol{q}=[q(1),q(2),\ldots,q(n)]$, we consider the following estimator with a joint sequence $\{(X_k,Y_k)\}_{t\geq 0}$.

$$\hat{\mu}_t(f,\boldsymbol{q})=\frac{1}{t}\sum_{k=1}^{t}\tilde{f}(X_k,Y_k), \tag{17}$$

where we use $\hat{\mu}_t(f,\boldsymbol{q})$ to indicate its dependence on the choice of $\boldsymbol{q}$. Note that the definition of the estimator here is for $\mathbb{E}_{\boldsymbol{\pi}}(f)$ with general $\boldsymbol{\pi}$. This will remain the same for the uniform $\boldsymbol{\pi}=\boldsymbol{u}$, while its related ratio estimator will be used for any non-uniform $\boldsymbol{\pi}$ as shall be shown later. Before going into the details, we have the following.

Theorem 3: In order for $\hat{\mu}_t(f,\boldsymbol{q})$ to be an asymptotically unbiased estimator of $\mathbb{E}_{\boldsymbol{\pi}}(f)$ for *any* function f, $\boldsymbol{q}$ has to be constant, i.e., $q(i)=q$ for all i. □

This result is somewhat counter-intuitive. At first glance, one might expect the asymptotic unbiasedness for any $\boldsymbol{q}$, as the (marginal) expectation of $\mathbf{P}f$ with respect to $\boldsymbol{\pi}$ is the same as that of f as seen from (8), and what we are doing would seem to be just a random mixture of two marginally unbiased values. The subtle point to note here is that the unbiasedness for $\mathbf{P}f$ or f is *only after* taking weighted sum with $\boldsymbol{\pi}$, and unless $\boldsymbol{q}$ is constant, the resulting expected value of $\hat{\mu}_t(f,\boldsymbol{q})$, in the limit, with general $\boldsymbol{q}$ is no longer the same as $\mathbb{E}_{\boldsymbol{\pi}}(f)$ and quickly becomes out of control. For this reason, since we are mainly interested in (asymptotically) unbiased estimators, we hereafter focus on the constant case $q(i)=q$ for all i.

B. Asymptotic Unbiasedness and Variance Performance

Since the estimator form is different depending on whether the stationary distribution $\boldsymbol{\pi}$ is uniform or not, we separate the cases for ease of exposition. We first consider the case of a uniform stationary distribution $\boldsymbol{\pi}=\boldsymbol{u}$, which is the one for the MHRW method. As a special case, if $q=1$, then the estimator in (17) becomes $\hat{\mu}_t(f,1)=\hat{\mu}_t(\mathbf{P}f)$, which is the case performing Rao-Blackwellization throughout at every node visited. Similarly for $q=0$, the estimator becomes $\hat{\mu}_t(f,0)=\hat{\mu}_t(f)$, which is the original estimator before applying Rao-Blackwellization. Let $\sigma^2(f,\mathbf{P},q)$ be the asymptotic variance of $\hat{\mu}_t(f,q)$. We then have the following.

Theorem 4: For any function f and any $q\in[0,1]$, $\hat{\mu}_t(f,q)$ is an asymptotically unbiased estimator of $\mathbb{E}_{\boldsymbol{\pi}}(f)$. Its asymptotic variance is given by

$$\begin{aligned}\sigma^2(f,\mathbf{P},q)=\sigma^2(f,\mathbf{P})(1-q)+\sigma^2(\mathbf{P}f,\mathbf{P})q\\ -2\left[\gamma_1(f)-\gamma_2(f)\right]q(1-q),\end{aligned} \tag{18}$$

where $\gamma_k(f)=\mathrm{Cov}_{\boldsymbol{\pi}}[f(X_0),f(X_k)]$. □

We next consider the case of any non-uniform $\boldsymbol{\pi}$, covering a much broader class of random-walk graph sampling methods, and obtain a similar result to Theorem 4. Only difference here is the need of a ratio estimator to correct a bias from the non-uniform $\boldsymbol{\pi}$, although it is non-trivial to prove its mathematical properties. Similar to the function $\tilde{f}$ above, for any bounded f, we define a weight function $\tilde{w}:\mathcal{N}\times\{0,1\}\to\mathbb{R}$ such that

$$\tilde{w}(i,0)=w(i),\ \text{and}\ \tilde{w}(i,1)=\mathbf{P}w(i),\ i\in\mathcal{N},$$

and also another function $\tilde{\psi}:\mathcal{N}\times\{0,1\}\to\mathbb{R}$ such that

$$\begin{aligned}&\tilde{\psi}(i,0)=\psi(i)=w(i)f(i),\ \text{and}\\ &\tilde{\psi}(i,1)=\mathbf{P}\psi(i)=\sum_{j\in\mathcal{N}}P(i,j)w(j)f(j),\ i\in\mathcal{N},\end{aligned}$$

with weights $w(i)=c/\pi(i)$ for some constant $c>0$. We then consider the following ratio estimator of $\mathbb{E}_{\boldsymbol{u}}(f)$:

$$\frac{\hat{\mu}_t(\psi,q)}{\hat{\mu}_t(w,q)}=\frac{\sum_{k=1}^{t}\tilde{\psi}(X_k,Y_k)}{\sum_{k=1}^{t}\tilde{w}(X_k,Y_k)}. \tag{19}$$

Here we can recover the degenerate cases by setting $q=0$ or $q=1$. Letting $\sigma^2_{\boldsymbol{w}}(f,\mathbf{P},q)$ be its asymptotic variance, we have the following.

Theorem 5: For any function f and any q, $\hat{\mu}_t(\psi,q)/\hat{\mu}_t(w,q)$ is asymptotically unbiased for $\mathbb{E}_{\boldsymbol{u}}(f)$. Its asymptotic variance is also given by

$$\begin{aligned}\sigma^2_{\boldsymbol{w}}(f,\mathbf{P},q)=\sigma^2(h,\mathbf{P})(1-q)+\sigma^2(\mathbf{P}h,\mathbf{P})q\\ -2\left[\gamma_1(h)-\gamma_2(h)\right]q(1-q),\end{aligned} \tag{20}$$

where h is in (15) and $\gamma_k(h)=\mathrm{Cov}_{\boldsymbol{\pi}}[h(X_0),h(X_k)]$. □

Theorems 4 and 5 show that the estimators in (17) and (19) both produce asymptotically unbiased samples with their asymptotic variances expressed as a quadratic function of q. As shown in Lemma 1, Rao-Blackwellization always makes the quality (or marginal variance) of every sample better than that of the usual sample. Even with the non-trivial autocovariance structures of $\{f(X_t)\}$ and $\{\mathbf{P}f(X_t)\}$, we have shown in Theorems 1 and 2 that the same ordering relationship also holds for the whole estimator with correlated samples. Thus, it is quite natural to expect that the more we incorporate Rao-Blackwellized samples with higher values of q, the resulting asymptotic variance would be always decreasing. Our next result, however, shows this is not always true.

Corollary 1: In the setting of Theorem 4, if

$$\max\{\gamma_1(f),\ \gamma_0(f)+4\gamma_1(f)\} \geq \gamma_2(f) \tag{21}$$

then $\sigma^2(f,\mathbf{P},q)$ is monotonically decreasing in $q \in [0,1]$ from $\sigma^2(f,\mathbf{P},0) = \sigma^2(f,\mathbf{P})$ down to $\sigma^2(f,\mathbf{P},1) = \sigma^2(\mathbf{P}f,\mathbf{P})$. Otherwise, there exists $q^* \in (0,1)$ such that $\sigma^2(f,\mathbf{P},q^*) > \sigma^2(f,\mathbf{P})$. The same goes for Theorem 5 with f replaced by h in (15). □

The condition in (21) is satisfied if $\gamma_k(f)$ (lag-k autocovariance of $f(X_t)$) is monotonically decreasing in k. This will be the case for any f if all the eigenvalues of $\mathbf{P}$ are non-negative [13], as typically seen in the so-called *lazy* random walk [30] or likely in the MHRW methods with self-transitions. In this case, as expected, the more we incorporate the neighborhood information into the sampler via the Rao-Blackwellization, the smaller the asymptotic variance of the obtained samples, albeit with potentially higher costs.

When the condition in (21) is not satisfied, however, Corollary 1 says the Rao-Blackwellization via neighborhood exploration, even if there is no additional cost for this operation, could be harmful and even produce worse samples than the pure random-walk sampling method *without* using any neighborhood information at all in the first place. This will be more likely the situation when there are negative eigenvalues for $\mathbf{P}$, inducing negative $\gamma_1(f)$ and $\gamma_0(f)+4\gamma_1(f) < \gamma_2(f)$. This issue is further compounded by the fact that it often costs more to obtain Rao-blackwellized samples, as we show next.

C. Comparing Sampling Methods From Practical Standpoint

We here illustrate how to *fairly* compare different sampling methods, producing varying qualities of samples but with different costs, on the same ground. Consider our original estimator $\hat{\mu}_t(f)$ for estimating $\mathbb{E}_{\boldsymbol{\pi}}(f)$. Observe that the "mean square error" of the estimator can be written as

$$\mathbb{E}\left[(\hat{\mu}_t(f)-\mathbb{E}_{\boldsymbol{\pi}}(f))^2\right] = \mathrm{Var}[\hat{\mu}_t(f)] + (\mathbb{E}[\hat{\mu}_t(f)]-\mathbb{E}_{\boldsymbol{\pi}}(f))^2,$$

where $\mathrm{Var}[\hat{\mu}_t(f)] = \mathbb{E}\left[(\hat{\mu}_t(f)-\mathbb{E}[\hat{\mu}_t(f)])^2\right]$. On one hand, it is known that the squared bias $(\mathbb{E}[\hat{\mu}_t(f)]-\mathbb{E}_{\boldsymbol{\pi}}(f))^2$ is of size $O(t^{-2})$ [31]. This becomes negligible when the chain is close to its stationary distribution, which is often the case for large t (more precisely, after the 'mixing time' of the chain [30]). It is obviously zero from the beginning for *stationary* chains. On the other hand, from the central limit theorem for Markov chains in (4), we know that $\mathrm{Var}[\hat{\mu}_t(f)]$ is of size $O(t^{-1})$ and the asymptotic variance is $\sigma^2(f,\mathbf{P}) = \lim_{t\to\infty} t \cdot \mathrm{Var}[\hat{\mu}_t(f)]$. Thus, for sufficiently large t, we have [31]

$$\mathrm{Var}[\hat{\mu}_t(f)] \approx \sigma^2(f,\mathbf{P})/t, \tag{22}$$

thereby implying that for sufficiently large t, the mean square error of the estimator is crucially influenced by the variance $\mathrm{Var}[\hat{\mu}_t(f)]$ of the estimator, which is again approximated by $\sigma^2(f,\mathbf{P})/t$. From this perspective, smaller asymptotic variance indicates that the estimator requires fewer samples to reach a desired accuracy even for the case of *finite* samples.

Now we are ready to discuss how to incorporate the sampling costs into the evaluation of estimation accuracy. Consider the original estimator $\hat{\mu}_t(f)$ and its Rao-Blackwellized version $\hat{\mu}_t(\mathbf{P}f)$ both obtained by a MHRW sampler for an illustration purpose. From (22) and Theorem 1, we have

$$\mathrm{Var}[\hat{\mu}_t(f)] \approx \sigma^2(f,\mathbf{P})/t \geq \sigma^2(\mathbf{P}f,\mathbf{P})/t \approx \mathrm{Var}[\hat{\mu}_t(\mathbf{P}f)].$$

Despite this variance reduction, the Rao-Blackwellized estimator may use *more costly* samples, as the sampler may have to manually explore all the neighbors at every node visited for Rao-Blackwellization.

To see its impact on the actual variance performance, we suppose that $\hat{\mu}_t(f)$ uses a unit cost per sample, while $\hat{\mu}_t(\mathbf{P}f)$ requires κ times more cost per sample on average. For a fair comparison, the total expenditure for sampling should remain the same for both cases. Thus, if $\hat{\mu}_t(f)$ uses t samples, its corresponding Rao-Blackwellized version should be based on t/κ samples, i.e., $\hat{\mu}_{t/\kappa}(\mathbf{P}f)$. Therefore, their variances become

$$\begin{aligned}\mathrm{Var}[\hat{\mu}_t(f)] &\approx \sigma^2(f,\mathbf{P})/t,\\ \mathrm{Var}[\hat{\mu}_{t/\kappa}(\mathbf{P}f)] &\approx \sigma^2(\mathbf{P}f,\mathbf{P})/(t/\kappa) = \kappa\cdot\sigma^2(\mathbf{P}f,\mathbf{P})/t\\ &\approx \kappa\cdot\mathrm{Var}[\hat{\mu}_t(\mathbf{P}f)],\end{aligned}$$

respectively. We can see that depending on the actual amount of cost κ per sample, the amount of variance reduction by Rao-Blackwellization can be completely offset, even leading to higher variance in the end. This demonstrates that higher sampling cost per sample is penalized and correctly reflected in the variance computation.

Repeating the same argument as above, we can deduce that the following *refined* variance of our generalized, Rao-Blackwellized estimator can be used for comparison with different values of q:

$$\kappa(q)\cdot\mathrm{Var}[\hat{\mu}_t(f,q)] \approx \kappa(q)\cdot\sigma^2(f,\mathbf{P},q)/t. \tag{23}$$

A similar form of the refined variance can also be made for the generalized, Rao-Blackwellized *ratio* estimator (for any non-uniform $\boldsymbol{\pi}$), or even for different sampling methods that produce varying qualities of samples yet with different costs. It is worth noting that the refined-variance form allows us to keep the same number t of samples used, which simplifies the comparison, as $\kappa(q)$ is explicitly introduced and can be used as a penalty for higher cost per sample. Here we use $\kappa(q)$ instead

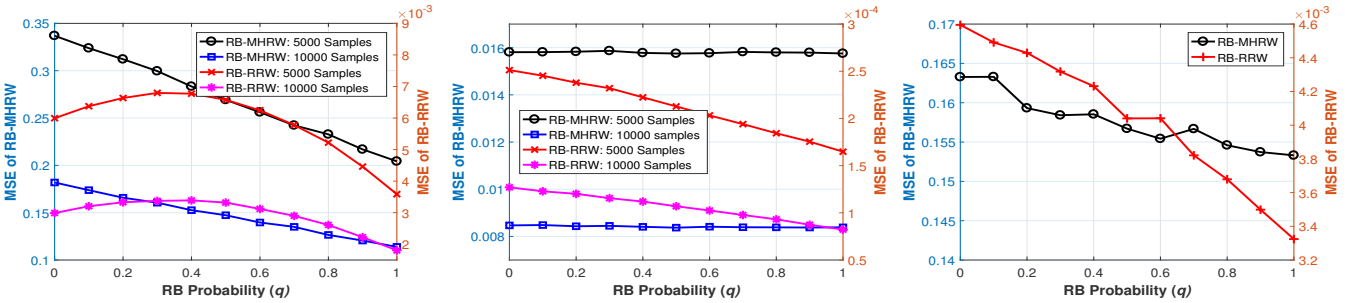

(a) Average degree of AS graph (b) Average clustering coefficient of AS graph (c) Average degree of Digg graph

Fig. 1. The MSE results of RB-MHRW(q) and RB-RRW(q) when changing the probability q.

of κ to reflect that the sampling cost $\kappa(q)$ would be a function of q and most likely be increasing in q, as the amount of neighborhood exploration, if required, increases with larger q. Therefore, in conjunction with our findings from Corollary 1, Rao-Blackwellization via neighborhood exploration surely offers a great improvement in the (asymptotic) variance if its associated cost is negligible; however, it should be used with great care if otherwise. ¶

V. Simulation Results

In this section, we present simulation results to support our theoretical findings. To this end, we consider two real-world network datasets. One is AS graph, an undirected graph of autonomous systems composed of 6474 nodes and 13233 edges, and the other is Digg graph, which is a social graph of the social news aggregator Digg's users, with 270,535 nodes and 1,731,658 edges. For our simulation, we use an undirected version of this graph. To ensure graph connectivity, we use the largest connected component of each graph.

As test cases, we consider the estimation of the average degree of each graph with a target sampling function $f(i)=d(i)$, and also the average clustering coefficient [1], [2], [21] with a target function $f(i)$ set to be a fraction of the number of connections among neighbors of i to the maximum possible connections among them, whose formal definition can be found in [21]. To measure the estimation accuracy, we use the "mean square error (MSE)" of each estimator, which is given by $\mathbb{E}[(\hat{x}-x)^2]$, where $\hat{x}$ represents the estimated value and x is the ground-truth value. As explained above, the MSE of an estimator can be approximated by its asymptotic variance divided by the number of samples used, as long as the number of samples is not that small. We can thus verify our theoretical findings through the evaluation of the MSE performance. In every simulation, an initial position of each random walk is drawn from its stationary distribution. Each data point reported here is obtained from 10^5 independent simulations. Due to space constraint, we here only provide representative simulation results.

In Figure 1, we compare the performance of our generalized, Rao-Blackwellized estimators for the MHRW and RRW methods under the *same* number of samples, which corresponds to the case that neighborhood information is available [9], [10] or neighborhood exploration can easily be done [11], [12]. Since our generalization is parameterized by probability q in performing the Rao-Blackwellized sampling at each node visited, we denote, by RB-MHRW(q) and RB-RRW(q), our Rao-Blackwellized estimators for the MHRW and RRW methods, respectively. Note that RB-MHRW(0) is the pure MHRW method and RB-MHRW(1) is the full use of Rao-Blackwellization. Similarly for RB-RRW(q).

In Figure 1(a), we report the MSE results of RB-MHRW(q) and RB-RRW(q) in estimating the average degree of AS graph with 5000 and 10000 samples, when varying the value of q. The left y-axis is for MSE of RB-MHRW(q) and the right y-axis is for that of RB-RRW(q). We can see that the MSE performance of RB-MHRW(q) exhibits a decreasing behavior in q, culminating at $q=1$, for both 5000 and 10000 samples. Here the larger the number of samples, the lower the MSE, which is well expected due to the (asymptotic) unbiasedness. However, for RB-RRW(q), there is no such a decreasing behavior, although RB-RRW(1) is always better than RB-RRW(0) as expected. This is the cautionary case as pointed out by Corollary 1.

Similarly, in Figure 1(b) and Figure 1(c), we report the MSE results of RB-MHRW(q) and RB-RRW(q) for estimating the average clustering coefficient of AS graph with 5000 samples and the average degree of Digg graph with 10^6 samples, respectively. The MSE of RB-RRW(q) is generally decreasing in q, with the best performance at $q=1$. On the contrary, the Rao-Blackwellization only provides a little improvement for the MHRW method. This trend is opposite to that of Figure 1(a). Therefore, we can see that the amount of improvement by Rao-Blackwellization and the benefit of its partial use can highly depend on the target function f and the underlying chain $\mathbf{P}$ (or the graph topology), among others.

In Figure 2, we present the MSE results of RB-MHRW(q) and RB-RRW(q) for estimating the average degree of AS graph under the consideration of sampling costs. While a precise cost model is not yet available in the literature, we use the following cost model. With a given sampling budget (in the number of samples), we deduct a unit cost (simply one sample) per sample for a usual sampling operation, while deducting a non-unit cost, defined as $\max\{1, \alpha d(i)\}$ with a control variable α, per sample for a Rao-Blackwellized

¶There always exists optimal $q^\star \in [0,1]$ to minimize the refined asymptotic variance, achieving the right balance between sampling cost and accuracy. The exact value of $q^\star$ clearly depends on the cost model for κ, the target function f and the underlying chain $\mathbf{P}$ (or the graph topology), among others. The development of a precise cost model is beyond the scope of this paper.

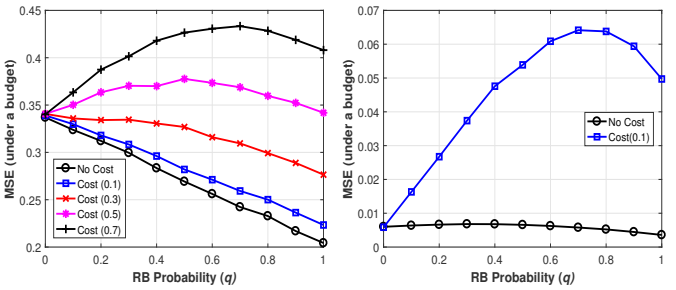

(a) RB-MHRW(q) (b) RB-RRW(q)

Fig. 2. The impact of sampling costs on the MSEs of RB-MHRW(q) and RB-RRW(q) for estimating the average degree of AS graph, with different choices of q.

operation at node i. This model reflects that obtaining Rao-Blackwellized samples from higher-degree nodes is more penalized. The MSE results are obtained with a budget of 5000 samples for different choices of α and varying q. Each result labeled with 'Cost(α)' indicates the one under a given α. 'No Cost' indicates the MSE result when the unit cost is deducted from the budget regardless of whether to perform the Rao-Blackwellized sampling. As expected from our discussion in IV-C, when we increase the sampling cost for each neighborhood exploration with increasing α, the amount of variance reduction by Rao-Blackwellization becomes completely offset, eventually leading to higher variance. We can also see that such an impact is intensified for RB-RRW(q), because the underlying random walk is SRW having a more bias toward high-degree nodes in the stationary regime. Since the decision for the Rao-Blackwellized sampling with q can only be done *independently* over i due to Theorem 3, the bias toward high-degree nodes increases the chance of the Rao-Blackwellized sampling at the high-degree nodes, incurring more sampling costs.

VI. CONCLUSION

We have studied, under our proposed Rao-Blackwellization framework, how much benefit one can achieve by knowing neighborhood information or performing neighbor exploration for random walk-based graph sampling. We have first shown two general forms of Rao-Blackwellization, which can be applied for a very large class of sampling methods, and proven that they always bring an improvement in the variance while maintaining the unbiasedness. We have further provided its non-trivial generalization and related mathematical proprieties. In particular, our results have pointed out possible danger of exploiting neighborhood information if done without care, which is in stark contrast to common belief. However, we also advocate the use of neighborhood information via Rao-Blackwellization, if its associated cost is negligible or such information is freely available.

REFERENCES

[1] M. Gjoka, M. Kurant, C. T. Butts, and A. Markopoulou, "Walking in facebook: A case study of unbiased sampling of OSNs," in *IEEE INFOCOM*, Mar. 2010.

[2] ——, "Practical recommendations on crawling online social networks," *IEEE JSAC*, vol. 29, no. 9, pp. 1872–1892, 2011.

[3] D. Stutzbach, R. Rejaie, N. Duffield, S. Sen, and W. Willinger, "On unbiased sampling for unstructured peer-to-peer networks," *IEEE/ACM Trans. on Networking*, vol. 17, no. 2, pp. 377–390, 2009.

[4] A. H. Rasti, M. Torkjazi, R. Rejaie, N. Duffield, W. Willinger, and D. Stutzbach, "Respondent-driven sampling for characterizing unstructured overlays," in *IEEE INFOCOM*, Apr. 2009.

[5] 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*, Jun. 2011.

[6] 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," in *ACM SIGMETRICS*, Jun. 2012.

[7] P. Wang, J. C. S. Lui, B. Ribeiro, D. Towsley, J. Zhao, and X. Guan, "Efficiently estimating motif statistics of large networks," *ACM Trans. on KDD*, vol. 9, no. 2, pp. 8:1–8:27, Nov. 2014.

[8] X. Chen, Y. Li, P. Wang, and J. C. S. Lui, "A general framework for estimating graphlet statistics via random walk," *CoRR*, vol. abs/1603.07504, 2016.

[9] A. Dasgupta, R. Kumar, and D. Sivakumar, "Social sampling," in *ACM SIGKDD*, Aug. 2012.

[10] P. Wang, B. Ribeiro, J. Zhao, J. C. S. Lui, D. Towsley, and X. Guan, "Practical characterization of large networks using neighborhood information," *CoRR*, vol. abs/1311.3037, 2013.

[11] P. Berenbrink, C. Cooper, R. Elsässer, T. Radzik, and T. Sauerwald, "Speeding up random walks with neighborhood exploration," in *ACM-SIAM SODA*, Jan. 2010.

[12] 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*, vol. 410, no. 1, pp. 94–100, Jan. 2009.

[13] J. S. Liu, *Monte Carlo strategies in scientific computing*. Springer-Verlag, 2004.

[14] W. Wefelmeyer, "Judging MCMC estimators by their asymptotic variance," in *Prague Stochastics*, Aug. 1998.

[15] I. W. McKeague and W. Wefelmeyer, "Markov chain Monte Carlo and Rao-Blackwellization," *Journal of Statistical Planning and Inference*, vol. 85, no. 1–2, pp. 171–182, Apr. 2000.

[16] 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*, vol. 21, no. 6, pp. 1087–1092, 1953.

[17] W. K. Hastings, "Monte carlo sampling methods using markov chains and their applications," *Biometrika*, vol. 57, no. 1, pp. 97–109, 1970.

[18] X. Wang, R. T. B. Ma, Y. Xu, and Z. Li, "Sampling online social networks via heterogeneous statistics," in *IEEE INFOCOM*, Apr. 2015.

[19] M. J. Salganik and D. D. Heckathorn, "Sampling and estimation in hidden populations using respondent-driven sampling," *Sociological Methodology*, vol. 34, pp. 193–239, 2004.

[20] S. Goel and M. J. Salganik, "Respondent-driven sampling as Markov chain Monte Carlo," *Statistics in Medicine*, vol. 28, no. 17, pp. 2202–2229, 2009.

[21] X. Xu, C.-H. Lee, and D. Y. Eun, "A general framework of hybrid graph sampling for complex network analysis," in *IEEE INFOCOM*, Apr. 2014.

[22] P. Brémaud, *Markov chains: Gibbs fields, Monte Carlo simulation, and queues*. Springer-Verlag, 1999.

[23] G. O. Roberts and J. S. Rosenthal, "General state space Markov chains and MCMC algorithms," *Probability Surveys*, vol. 1, pp. 20–71, 2004.

[24] O. Haggstrom and J. Rosenthal, "On variance conditions for Markov chain CLTs," *Electronic Communications in Probability*, vol. 12, pp. 454–464, Dec. 2007.

[25] G. Casella and C. P. Robert, "Rao-Blackwellisation of sampling schemes," *Biometrika*, vol. 83, no. 1, pp. 81–94, Mar. 1996.

[26] S. M. Ross, *Stochastic processes*, 2nd ed. John Wiley & Son, 1996.

[27] Z. Bar-Yossef and M. Gurevich, "Efficient search engine measurements," in *WWW*, May 2007.

[28] K. Avrachenkov, B. Ribeiro, and D. Towsley, "Improving random walk estimation accuracy with uniform restarts," in *WAW*, Dec. 2010.

[29] C.-H. Lee, X. Xu, and D. Y. Eun, "On the Rao-Blackwellization and its application for graph sampling via neighborhood exploration," *Technical Report*, Jul. 2016, https://www.dropbox.com/sh/2xwapqlv42rsr6k/AACEV1B82HdX0FBFZcK2kXtya/RBtechrep.pdf?raw=1.

[30] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov chains and mixing times*. American Mathematical Society, 2009.

[31] P. H. Peskun, "Optimum monte-carlo sampling using markov chains," *Biometrika*, vol. 60, pp. 607–612, 1973.