(will be inserted by the editor)

Sequential Stratified Regeneration: MCMC for Large State Spaces with an Application to Subgraph Count Estimation

Carlos H. C. Teixeira 1* · Mayank Kakodkar 2* · Vinícius Dias 1,3 · Wagner Meira Jr. 1 · Bruno Ribeiro 2

the date of receipt and acceptance should be inserted later

Abstract This work considers the general task of estimating the sum of a bounded function over the edges of a graph, given neighborhood query access and where access to the entire network is prohibitively expensive. To estimate this sum, prior work proposes Markov chain Monte Carlo (MCMC) methods that use random walks started at some seed vertex and whose equilibrium distribution is the uniform distribution over all edges, eliminating the need to iterate over all edges. Unfortunately, these existing estimators are not scalable to massive real-world graphs. In this paper, we introduce Ripple, an MCMC-based estimator that achieves unprecedented scalability by stratifying the Markov chain state space into ordered strata with a new technique that we denote sequential stratified regenerations. We show that the Ripple estimator is consistent, highly parallelizable, and scales well.

We empirically evaluate our method by applying Ripple to the task of estimating connected, induced subgraph counts given some input graph. Therein, we demonstrate that Ripple is accurate and can estimate counts of up to 12-node subgraphs, which is a task at a scale that has been considered unreachable, not only by prior MCMC-based methods but also by other sampling approaches. For instance, in this target application, we present results in which the Markov chain state space is as large as 10^{43} , for which Ripple computes estimates in less than 4 hours, on average.

Keywords Markov Chain Monte Carlo, Random Walk, Regenerative Sampling, Motif Analysis, Subgraph Counting, Graph Mining

⊠ Carlos H. C. Teixeira

E-mail: carlos@dcc.ufmg.br

¹ Universidade Federal de Minas Gerais, Belo Horizonte, Brazil

² Purdue University, West Lafayette, USA

³ Universidade Federal de Ouro Preto, Ouro Preto, Brazil

^{*} Equal contribution.

1 Introduction

This work considers the following general task: Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a simple graph, where \mathcal{V} is the set of vertices, \mathcal{E} is the set of edges, and \mathcal{E} contains at most a single edge between any pair of vertices and no self-loops. Our goal is to efficiently estimate the sum of a bounded function over all the edges of \mathcal{G} ,

$$\mu(\mathcal{E}) = \sum_{(u,v)\in\mathcal{E}} f(u,v) \tag{1}$$

where $f : \mathcal{E} \to \mathbb{R}$, $f(\cdot) < B$ is a bounded function for some constant $B \in \mathbb{R}$ under the following query model from Avrachenkov et al. (2016).

Assumption 1 (Query Model) Assume we are given arbitrary seed vertices and can query the neighborhood $\mathbf{N}(u) \triangleq \{v \in \mathcal{V} : (u, v) \in \mathcal{E}\}$ for any vertex $u \in \mathcal{V}$ such that accessing the entire graph \mathcal{G} is prohibitively expensive.

This setting arises naturally in the subgraph counting problem, which we study in Section 4 Simple Monte Carlo procedures are not useful because random vertex and edge queries are not directly available, and reservoir sampling would require iteration over all edges. Standard Markov chain Monte Carlo (MCMC) methods cannot estimate the quantity in Equation (1) and are limited to estimate $\mu(\mathcal{E})/|\mathcal{E}|$, because $|\mathcal{E}|$ in our task is unknown (Ribeiro and Towsley) 2012). Generally, under Assumption 1 Equation (1) is estimated using specialized MCMC estimators that use a random-walk-like Markov chain that has a uniform distribution over the edges \mathcal{E} as its equilibrium distribution. However, these estimators (Avrachenkov et al.) 2016) are impractical in large graphs because their running time is $O(|\mathcal{E}|)$.

Traditional MCMC methods are limited by their reliance on the Markov chain on \mathcal{G} reaching equilibrium or burning in. Because the rate of convergence to equilibrium depends on the spectral gap (Aldous and Fill), [2002), a significant number of Markov chain steps is needed to burn in in order to produce accurate estimates of Equation (1), particularly in large graphs. Parallel approaches that divide the state space into disjoint "chunks", which are to be processed in parallel (Wilkinson), [2006], [Neiswanger et al.], [2014), offer no respite because we cannot access the entire graph. In fact, \mathcal{G} may not even have disconnected components (i.e., disjoint chunks) that can be parallelized. Therefore, traditional MCMC on \mathcal{G} offers no meaningful parallelization opportunities and running times may be arbitrarily long.

Contributions. This work introduces sequential stratified regeneration (Ripple), a novel parallel MCMC technique that expands the application frontier of MCMC to large state-space graphs \mathcal{G} . Ripple stratifies the underlying Markov chain state space into ordered strata that need not be disjoint chunks, rather, they need to be connected. Markov chain regeneration (Nummelin 1978) is then used to compute estimates in each stratum sequentially, using a recursive method, which improves regeneration frequencies and reduces variance. Ripple offers an unprecedented level of efficiency and parallelism

for MCMC sampling on large state-space graphs while retaining the benefits of MCMC-based algorithms, such as low memory demand (polynomial w.r.t. output).

Surprisingly, the parallelism of Ripple comes from the regeneration rather than the stratification: the strata's job is to keep regeneration times short. We demonstrate that the estimates obtained by Ripple are consistent, among other theoretical guarantees. In addition, we empirically show the power of Ripple in a real-world application by specializing Equation (1) to subgraph counting in multi-million-node attributed graphs—to the best of our knowledge, a task at a scale that has been thought unreachable by any other MCMC method. Our specific contributions to the subgraph counting problem include streaming-based optimizations coupled with a parallel reservoir sampling algorithm, novel efficiency improvements to the random walk on the HON (Wang et al.) [2014) and a theoretical analysis of scalability in terms of running time and memory w.r.t. the subgraph size, verified empirically on large datasets.

2 Background and Prior Work

The MCMC random-walk-like Markov chain over the graph \mathcal{G} is defined as:

Definition 1 (Random Walk on \mathcal{G}) Given a simple graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, a simple random walk is a time-homogenous Markov chain $\boldsymbol{\Phi}$ with state space \mathcal{V} and transition probability $p_{\boldsymbol{\Phi}}(u,v) = 1/\mathbf{d}(u)$, when $(u,v) \in \mathcal{E}$ and 0 otherwise, where $\mathbf{d}(u) = |\mathbf{N}(u)|$ is the degree of u in \mathcal{G} and $\mathbf{N}(u) = \{v : (u,v) \in \mathcal{E}\}$ is the neighborhood of u.

It is easy to check that the above random walk can be sampled under Assumption 1 and that on a connected graph, this walk samples edges uniformly at random in a steady state (check Appendix B.1 for details). Our notation is summarized in Appendix A.

2.1 Regenerations in Discrete Markov Chains

The rate of convergence to stationarity of the random walk Φ from Definition depends on the spectral gap (Aldous and Fill) (2002). As such, practitioners are encouraged to run a single, long sample path, which prevents them from splitting the task among multiple cores. Usually, because the spectral gap is unknown or loosely bounded, practitioners use various diagnostics to eyeball if the chain has mixed (Rosenthal, 1995). The variance of an estimate computed from a stationary chain (Ribeiro and Towsley, 2012) also depends on the spectral gap.

A solution to the above problems is to *split* (Nummelin, 1978) the Markov chain using regenerations. Discrete Markov chains regenerate every time they

¹ The spectral gap is defined as $\delta = 1 - \max\{|\lambda_2|, |\lambda_{|\mathcal{V}|}|\}$, where λ_i denotes the *i*-th eigenvalue of the transition probability matrix of Φ .

enter a fixed state, which is referred to as a regeneration point. This naturally yields the definition of a $random\ walk\ tour\ (RWT)$.

Definition 2 (*RWT* over Φ) Given a time-homogenous Markov chain Φ over finite state space \mathcal{V} and a fixed point $x_0 \in \mathcal{V}$, an $RWT \mathbf{X} = (X_i)_{i=1}^{\xi}$ is a sequence of states visited by Φ between two consecutive visits to x_0 , that is, $X_1 = x_0$ and $\xi = \min\{i > 1 : X_{i+1} = x_0\}$ is the first return time to x_0 .

Because of the strong Markov property (Bremaud) [2001] Chap-2,Thm-7.1), RWTs started at x_0 are i.i.d. and can be used to estimate $\mu(\mathcal{E})$ from Equation (1) when $|\mathcal{E}|$ is unknown (Avrachenkov et al., [2016] [2018] Teixeira et al., [2018] Savarese et al., [2018] Cooper et al., [2016] Massoulié et al., [2006].

Lemma 1 (RWT Estimate) Given the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and the random walk Φ from Definition $\boxed{1}$ consider $f : \mathcal{E} \to \mathbb{R}$ bounded by B, and \mathcal{T} , a set of m RWTs started at $x_0 \in \mathcal{V}$ (Definition $\boxed{2}$) sampled in a parallel z core environment assuming each core samples an equal number of tours. Then,

$$\hat{\mu}_*(\mathcal{T}; f, \mathcal{G}) = \frac{\mathbf{d}(x_0)}{2m} \sum_{\mathbf{X} \in \mathcal{T}} \sum_{j=1}^{|\mathbf{X}|} f(X_j, X_{j+1}), \qquad (2)$$

is an unbiased and consistent estimator of $\mu(\mathcal{E}) = \sum_{(u,v) \in \mathcal{E}} f(u,v)$ if \mathcal{G} is connected, where each X_j refers to the jth state in the RWT $\mathbf{X} \in \mathcal{T}$.

The expected running time for sampling m tours is $O\left(m/z\frac{2|\mathcal{E}|}{\mathbf{d}(x_0)}\right)$, and when \mathcal{G} is non-bipartite, the variance of the estimate is bounded as

$$\operatorname{Var}\left(\hat{\mu}_{*}(\mathcal{T})\right) \leq \frac{3B^{2}}{m} \frac{\left|\mathcal{E}\right|^{2}}{\delta(\boldsymbol{\Phi})},\tag{3}$$

where $\delta(\Phi)$ is the spectral gap as defined in the beginning of this section.

The RWT Estimate can be considered a Las Vegas transformation of MCMC, which takes random time but yields unbiased estimates of objectives, such as Equation (1). The parallelism in the expected running time in Lemma 1 is directly due to the independence of RWTs. Moreover, confidence intervals for the RWT Estimate can be computed, because $\sqrt{m} \frac{\hat{\mu}_*(\mathcal{T}) - \mu(\mathcal{E})}{\hat{\sigma}(\mathcal{T})}$ approaches the standard normal distribution for sufficiently large m, where $\hat{\sigma}(\mathcal{T})^2$ is the empirical variance of $\hat{\mu}_*(\mathbf{X})$, the RWT Estimate computed using an individual tour $\mathbf{X} \in \mathcal{T}$.

2.2 Improving the Regeneration Frequency

From Lemma $\boxed{1}$ it is clear that increasing the degree of the regeneration point $\mathbf{d}(x_0)$ and spectral gap $\delta(\boldsymbol{\Phi})$ and decreasing $|\mathcal{E}|$ reduces the variance as well as the running time of the RWT Estimate. Avrachenkov et al. $\boxed{2016}$ showed that using the *supernode* in a contracted graph as a regeneration point achieves the above reductions.

Definition 3 (Contracted Graph) (Avrachenkov et al.) 2016) Given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ from Definition 1 and a set of vertices $I \subset \mathcal{V}$, a contracted graph is a multigraph \mathcal{G}_I formed by collapsing I into a single node ζ_I . The vertex set of \mathcal{G}_I is then given by $\mathcal{V}\setminus I \cup \{\zeta_I\}$, and its edge multiset is obtained by conditionally replacing each endpoint of each edge with ζ_I if it is a member of I and removing self-loops on ζ_I . We refer to the set I and the vertex ζ as the supernode.

Contractions benefit RWTs because the supernode degree $\mathbf{d}_{\mathcal{G}_I}(\zeta_I)$ in \mathcal{G}_I and the spectral gap $\delta(\boldsymbol{\Phi}_I)$ of the random walk on the contracted graph increase monotonically with |I| (Avrachenkov et al.) 2016). Moreover, RWTs can be sampled on \mathcal{G}_I without explicit construction, as we see next.

Remark 1 Let the multi-set $\mathbf{N}(\zeta_I) \triangleq \biguplus_{u \in I} \mathbf{N}_{\mathcal{G}}(u) \setminus I$ be the neighborhood of the supernode in \mathcal{G}_I from Definition $\boxed{3}$ Let Φ_I be the simple random walk on \mathcal{G}_I . An RWT $(X_i)_{i=1}^{\xi}$ on Φ_I from ζ_I is sampled by setting $X_1 = \zeta_I$, sampling X_2 u.a.r. from $\mathbf{N}(\zeta_I)$ and subsequently sampling transitions from Φ until the chain enters I, i.e., $\xi = \min\{i > 1 \colon X_{\xi+1} \in I\}$.

This construction naturally stratifies \mathcal{E} and decomposes $\mu(\mathcal{E})$ as $\mu(\mathcal{E}_*) + \mu(\mathcal{E} \setminus \mathcal{E}_*)$, where we can exactly compute the $\mu(\mathcal{E}_*)$ and compute an RWT Estimate of $\mu(\mathcal{E} \setminus \mathcal{E}_*)$ on the contracted graph. However, to compute the supernode degree, $\mathbf{d}_{\mathcal{G}_I}(\zeta_I)$; furthermore, to sample from $\mathbf{N}(\zeta_I)$, we need to enumerate the set of the edges incident on I in \mathcal{G} given by $\mathcal{E}_* \subset \mathcal{E}$. As such, a massive supernode I (which is crucial when $|\mathcal{E}|$ is large) makes enumerating \mathcal{E}_* prohibitively expensive. We overcome these issues and gain additional control over regenerations by further stratifying $|\mathcal{E}|$.

3 Sequential Stratified Regenerations

Ripple controls regeneration times through a sequential stratification of the vertices and edges of \mathcal{G} into ordered strata as illustrated in Figure $\boxed{1}$ which allows us to control the regeneration frequency and the RWT Estimate variance. For each stratum, we then construct a graph in which the supernode is created by collapsing all prior strata, from which RWTs can be sampled. We use the RWTs from the previous strata to estimate the degree of and sample transitions from the supernode. The core idea is described in two steps: Section $\boxed{3.1}$ details the stratification and conditions that it needs to satisfy and Section $\boxed{3.2}$ describes the recursion. Finally, we show that the estimator bias converges to zero asymptotically in the number of tours. Particularly for subgraph counting, we show that Ripple's time complexity is independent of the (higher-order) graph size ($|\mathcal{E}|$) and only depends polynomially on the diameter and maximum degree of the input graph and the subgraph size (Section $\boxed{4}$).

3.1 Sequential Stratification

Consider the following vertex and edge stratification procedure.

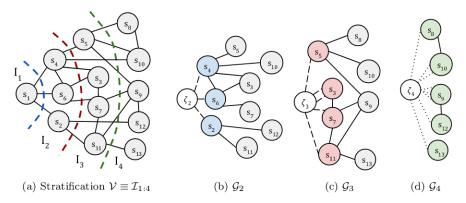


Fig. 1: Figure 1a shows a simple graph \mathcal{G} that is stratified into four strata $\{\mathcal{I}_1, \mathcal{I}_2, \mathcal{I}_3, \mathcal{I}_4\}$. Figures 1b to 1d show the second, third and fourth graph strata constructed by Definition 5 In the multi-graph \mathcal{G}_2 (Figure 1b), vertices in \mathcal{I}_1 are collapsed into ζ_2 and only edges incident on \mathcal{I}_2 are preserved. The edge set therefore contains \mathcal{J}_2 and the edges between ζ_2 and \mathcal{I}_2 . Consequently, self-loops on ζ_2 and edges between $\mathcal{I}_{3:4}$ are absent. Figures 1c and 1d follow suit. In each stratum \mathcal{G}_r , RWTs from ζ_r are started by sampling u.a.r. from the dotted edges and estimates are computed over the solid edges.

Definition 4 (Sequential Stratification) Given $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ from Definition 1 a function $\rho \colon \mathcal{V} \to \{1, \dots, R\}$ induces the stratification $(\mathcal{I}_r, \mathcal{J}_r)_{r=1}^R$ if $s \in \mathcal{I}_{\rho(s)}$, for each $s \in \mathcal{V}$, and $(u, v) \in \mathcal{J}_{\min(\rho(u), \rho(v))}$, for each $(u, v) \in \mathcal{E}$.

Note that these strata are pairwise disjoint and their union is the set of vertices and edges of the graph. Next, we describe the contracted graph over which RWTs are to be sampled in each stratum.

Definition 5 (r-th Graph Stratum) Let $\mathcal{A}_{i:j} \triangleq \bigcup_{x=i}^{j} \mathcal{A}_{x}$ be defined for any ordered tuple of sets. Let $(\mathcal{I}_r, \mathcal{J}_r)_{r=1}^R$ be the stratification induced by ρ from Definition $\boxed{4}$ on $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. The r-th graph stratum $\mathcal{G}_r = (\mathcal{V}_r, \mathcal{E}_r), r > 1$, is obtained by removing all edges not incident on \mathcal{I}_r and vertices that do not neighbor vertices in \mathcal{I}_r and subsequently contracting $\mathcal{I}_{1:r-1}$ into ζ_r according to Definition $\boxed{3}$ Further, let $\mathbf{\Phi}_r$ denote the simple random walk on \mathcal{G}_r .

It can be shown that the vertex set \mathcal{V}_r contains the r-th stratum \mathcal{I}_r , the r-th supernode ζ_r , obtained by collapsing $\mathcal{I}_{1:r-1}$, and vertices from subsequent strata neighboring \mathcal{I}_r , $\cup_{u \in \mathcal{I}_r} \mathbf{N}(u) \cap \mathcal{I}_{r+1:R}$. The edge multiset \mathcal{E}_r is the union of \mathcal{J}_r and edges that connect ζ_r to vertices in \mathcal{I}_r resulting from the graph contraction. A detailed example is shown in Figure 1 Note that when R = 2, Ripple reduces to the estimator from Avrachenkov et al. (2016).

Ergodicity-Preserving Stratification. Because the RWT Estimate is consistent only if the underlying graph is connected according to Lemma we have the following definition:

Definition 6 (Ergodicity-Preserving Stratification (*EPS*)) The stratification due to ρ from Definition 4 is an Ergodicity-Preserving Stratification if each graph stratum from Definition 5 is connected, i.e., Φ_r , r > 1, is irreducible.

We propose necessary and sufficient conditions on ρ that yield an EPS.

Proposition 1 ρ yields an EPS if the following three conditions are satisfied:

- (a) for at least one vertex in each connected component of \mathcal{G} , ρ evaluates to 1;
- (b) for each $u: \rho(u) = r$, there exists $v \in \mathbf{N}(u)$ such that $\rho(v) \leq r$; and
- (c) there exists $(u_0, v_0) \in \mathcal{E}$ such that $\rho(u_0) = r$ and $\rho(v_0) < r$.

Although the optimal stratification would depend on \mathcal{G} and the quantity being estimated, an ideal stratification would yield graph strata wherein the supernode degree and connectivity are maximized (Lemma 1) while minimizing the number of strata (because of the bias propagation described in Theorem 2). ρ needs to be efficient as well because we will see that it is evaluated at each step of the random walk and the Ripple estimators from Definitions 8 and 9 heavily depend on it. In Proposition 10 we show that return times to the supernode are inversely proportional to the fraction of vertices in \mathcal{I}_r connected to $\mathcal{I}_{1:r-1}$.

3.2 Recursive Regenerations

Assume for the moment that in each stratum, $r=2,\ldots,R$, we know the degree of the supernode $\mathbf{d}(\zeta_r)$ and can sample directly from $p_{\mathbf{\Phi}_r}(\zeta_r,\cdot)$, which is the transition probability out of ζ_r in the graph stratum \mathcal{G}_r . We could then sample RWTs \mathcal{T}_r and compute stratumwise RWT Estimates, which when combined as $\mu(\mathcal{J}_1) + \sum_{r=2}^R \hat{\mu}_r(\mathcal{T}_r)$ provide an unbiased estimate of $\mu(\mathcal{E})$ as a direct consequence of Lemma \mathbb{I} and the linearity of expectations. Unfortunately, the impracticality of this assumption, especially under Assumption \mathbb{I} (when R > 2), necessitates the following relaxation.

Definition 7 (Supernode Estimates, $\widehat{\mathbf{d}}(\zeta_r)$ and $\widehat{p}_{\boldsymbol{\Phi}_r}(\zeta_r, \cdot)$) Given an *EPS* of \mathcal{G} (Definition 6), the supernode estimates in the r-th graph stratum \mathcal{G}_r consist of the estimate of the degree $\widehat{\mathbf{d}}(\zeta_r)$ and a sample from some approximate transition probability out of the supernode $\widehat{p}_{\boldsymbol{\Phi}_r}(\zeta_r, \cdot)$. Let $\widehat{\boldsymbol{\Phi}}_r$ be the random walk on \mathcal{G}_r , where transitions are sampled according to $\boldsymbol{\Phi}_r$ everywhere except ζ_r , where they are sampled from $\widehat{p}_{\boldsymbol{\Phi}_r}(\zeta_r, \cdot)$.

Although $\widehat{\Phi}_r$ may not be reversible, RWTs on $\widehat{\Phi}_r$ retain pairwise independence and the benefits stated after Lemma [1] We leverage this fact in the following recursive solution that computes *supernode estimates* in the current stratum using supernode estimates and tours sampled in the previous strata.

Definition 8 (Ripple's Recurrence Relation) Given a graph \mathcal{G} stratified according to ρ (Definition 6) and some stratum r, $1 < r \le R$, assume access

to the result of previous recursive steps, i.e., the set of m_q RWTs $(\mathcal{T}_q^{\dagger})$, supernode degree estimates $\widehat{\mathbf{d}}(\zeta_q)$ and estimated transition probabilities out of the supernode $\widehat{p}_{\mathbf{\Phi}_q}(\zeta_q,\cdot)$ (Definition 7) for all $2 \leq q < r$. The estimate of the number of edges between \mathcal{I}_q and \mathcal{I}_r is given by

$$\widehat{\beta}_{q,r} = \frac{\widehat{\mathbf{d}}(\zeta_q)}{|\mathcal{T}_q^{\dagger}|} \sum_{\mathbf{X} \in \mathcal{T}_q^{\dagger}} \sum_{j=2}^{|\mathbf{X}|} \mathbf{1} \{ \rho(X_j) = r \},$$
(4)

where X_j is the j-th state visited in tour **X**, and by convention, $\widehat{\beta}_{1,r} = |\mathcal{E} \cap \mathcal{I}_1 \times \mathcal{I}_r|$ is exactly computed. The r-th supernode degree is then estimated as

$$\widehat{\mathbf{d}}(\zeta_r) = \sum_{q=1}^{r-1} \widehat{\beta}_{q,r} \,. \tag{5}$$

Transitions from $\widehat{p}_{\Phi_r}(\zeta_r,\cdot)$ are sampled by sampling $q \in \{1,\ldots,r-1\}$ with probability $\widehat{\beta}_{q,r}$ and then sampling u.a.r. from $\widehat{\mathbf{U}}_{q,r}$, which is defined as

$$\widehat{\mathbf{U}}_{q,r} = \biguplus_{\mathbf{X} \in \mathcal{T}_q^{\dagger}} \biguplus_{j=2}^{|\mathbf{X}|} \{ X_j \colon \rho(X_j) = r \} , \text{ when } q > 1 ,$$
 (6)

and as $\bigoplus_{u \in \mathcal{I}_1} \mathbf{N}(u) \cap \mathcal{I}_r$ by convention when q = 1, where \uplus is the multi-set union. $\widehat{\mathbf{U}}_{q,r}$, q > 1, is thus the multi-set of all states in \mathcal{I}_r visited by RWTs on $\widehat{\boldsymbol{\Phi}}_q$. An RWT so started stops when it reaches some state X', where $\rho(X') = r$.

Proposition $\overline{7}$ (Appendix \overline{C}) contains additional details for sampling RWTs on Φ_r . The above recursion therefore allows us to estimate supernode degrees and sample RWTs to compute an estimate of $\mu(\mathcal{E})$ from Equation $\overline{(1)}$ as follows:

Definition 9 (Ripple's μ **Estimator)** Given the supernode degree estimates $\widehat{\mathbf{d}}(\zeta_r)$ and RWTs \mathcal{T}_r^{\dagger} sampled in each graph stratum from Definition 8 and the edge strata \mathcal{J}_r , $2 \leq r \leq R$ based on an EPS of \mathcal{G} from Definition 6, the Ripple estimate is defined as

$$\hat{\mu}_{Ripple} = \mu(\mathcal{J}_1) + \sum_{r=2}^{R} \hat{\mu}\left(\mathcal{T}_{2:r}^{\dagger}; f\right), \tag{7}$$

where,
$$\hat{\mu}\left(\mathcal{T}_{2:r}^{\dagger}; f\right) = \frac{\widehat{\mathbf{d}}(\zeta_r)}{2|\mathcal{T}_r^{\dagger}|} \sum_{\mathbf{X} \in \mathcal{T}_r^{\dagger}} \sum_{j=2}^{|\mathbf{X}|-1} f(X_j, X_{j+1}),$$
 (8)

and X_j is the jth state visited by the $RWT \mathbf{X} \in \mathcal{T}_r^{\dagger}$. The dependence of \mathcal{T}_r^{\dagger} and $\hat{\mathbf{d}}(\zeta_r)$ on $\mathcal{T}_{2:r-1}^{\dagger}$ is suppressed for brevity.

This estimate of $\mu(\mathcal{E})$ is unbiased when the number of tours is infinite.

Theorem 1 The Ripple estimate from Definition 9 is a consistent estimator of $\mu(\mathcal{E})$ (asymptotically unbiased in the number of tours), that is,

$$\lim_{|\mathcal{T}_2^{\dagger}| \to \infty} \dots \lim_{|\mathcal{T}_R^{\dagger}| \to \infty} \mu(\mathcal{J}_1) + \sum_{r=2}^{R} \hat{\mu}\left(\mathcal{T}_{2:r}^{\dagger}; f\right) \stackrel{a.s}{=} \mu(\mathcal{E}).$$

In the finite regime, however, there exists a bias in each stratum that depends on the estimation bias in the previous strata, which we quantify as follows:

Theorem 2 Given the random walk Φ_r on the EPS-stratum \mathcal{G}_r from Definitions $\boxed{5}$ and $\boxed{6}$, the estimates of the degree and transition probability at the supernode $\widehat{\mathbf{d}}(\zeta_r)$ and $\widehat{p}_{\Phi_r}(\zeta_r,\cdot)$ from Definition $\boxed{7}$ and assuming aperiodic Φ_r , the bias of the Ripple estimate in the rth stratum from Equation $\boxed{8}$ is given by

$$\left| \mathbb{E} \left[\hat{\mu} \left(\mathcal{T}_{2:r}^{\dagger}; f \right) \left| \mathcal{T}_{2:r-1}^{\dagger} \right] - \mu(\mathcal{J}_r) \right| \leq \left(\lambda_r \nu_r + |1 - \lambda_r| \right) \frac{\sqrt{3} B |\mathcal{E}_r|}{\sqrt{\delta_r}} \,,$$

where δ_r is the spectral gap of Φ_r , B is the upper bound of f, $\nu_r = \|\widehat{p}_{\Phi_r}(\zeta_r, \cdot) - p_{\Phi_r}(\zeta_r, \cdot)\|_2$ is the L^2 distance between transition probabilities out of ζ_r (Aldous and Fill, 2002) (Definition 13) and $\lambda_r = \widehat{\mathbf{d}}(\zeta_r)/\mathbf{d}(\zeta_r)$.

Therefore, the bias in each stratum affects the bias in subsequent strata. Consequently, we control the empirical variance in each stratum by increasing the number of tours sampled (we detail this for subgraph counting in Section 4).

4 Applying Ripple to Count Subgraphs

We now focus on a concrete implementation of Ripple to count subgraphs on a given simple input graph G=(V,E,L) with vertices V, edges E, and attribute function L, which is assumed to be finite and undirected. In general, a subgraph induced by any $V'\subset V$ on G is given by $G(V')=(V',E\cap(V'\times V'),L)$. However, in this work, we are interested in subgraphs G(V') that are connected and where |V'|=k, referred to as a connected, induced subgraph (CIS) of size k or k-CIS. As such, the task is defined as

Definition 10 (Subgraph Count) Let $\mathcal{V}^{(k)}$ be the set of all k-CISs of graph G, let \sim denote the graph isomorphism equivalence relation (or any equivalence relation), and let \mathcal{H} be an arbitrary set of pairwise nonequivalent k-CISs. The subgraph count is defined as the $|\mathcal{H}|$ -dimensional vector $\mathcal{C}^{(k)} = (\mathcal{C}_H^{(k)})_{H \in \mathcal{H}}$, where $\mathcal{C}_H^{(k)} = \sum_{s \in \mathcal{V}^{(k)}} \mathbf{1}\{s \sim H\}$, and $\mathbf{1}\{\cdot\}$ is the indicator function.

Therefore, $\mathcal{C}^{(k)}$ contains the count of subgraphs in $\mathcal{V}^{(k)}$ equivalent to each subgraph in \mathcal{H} . We suppress the dependence of $\mathcal{C}^{(k)}$ on \mathcal{H} for simplicity.

Subgraph counting is challenging when k > 3 in real-world input graphs because $\mathcal{V}^{(k)}$ is not tractably enumerable and naively sampling k vertices to obtain CISs is challenging because $|\mathcal{V}^{(k)}|/|V|^k \to 0$ (as evidenced by Table 1). Next, we address this issue by reducing the subgraph counting problem to an edge sum (Equation (1)) over a higher-order graph that only provides neighborhood query access for large-real-world input graphs. We also propose a stratification strategy compatible with the access model and introduce novel solutions to improve speed and memory requirements. We defer the straightforward aspects to Appendix E wherein we summarize the entire algorithm (Algorithm 2).

4.1 MCMC on the Subgraph Space

Wang et al. (2014) proposed a network over subgraphs called the HON, which exposes neighborhood query access from Assumption 1 and is therefore amenable to MCMC solutions (which we optimize in Algorithm 1).

Definition 11 (Higher-Order Network (k-HON) (Wang et al., 2014)) The higher-order network or $HON\mathcal{G}^{(k)} = (\mathcal{V}^{(k)}, \mathcal{E}^{(k)})$ is a graph whose vertices are the set of all k-CIS contained in the input graph G, and (u, v) form an edge in $\mathcal{E}^{(k)}$ if they share all but k-1 vertices, that is, $|V(u) \cap V(v)| = k-1$.

In the k-1-HON, the subgraph induced by an edge $(u,v) \in \mathcal{E}^{(k-1)}$, i.e., $G(V(u) \cup V(v))$, is a k-CIS. Thus, the subgraph counts from Definition 10 can then be expressed as an edge sum over $\mathcal{G} \equiv \mathcal{G}^{(k-1)}$ as

$$\mathcal{C}^{(k)} = \mu(\mathcal{E}^{(k-1)}) = \sum_{(u,v)\in\mathcal{E}^{(k-1)}} \left(\frac{\mathbf{1}\{G(V(u)\cup V(v))\sim H\}}{\gamma(u,v)}\right)_{H\in\mathcal{H}}, \quad (9)$$

where $\gamma(u,v) = |\{(\ddot{u},\ddot{v}) \in \mathcal{E} \colon V(u) \cup V(v) \equiv V(\ddot{u}) \cup V(\ddot{v})\}|$ is the number of edges that represent the same subgraph as (u,v). The set of edges sampled by a random walk on $\mathcal{G}^{(k-1)}$ is called the *pairwise subgraph random walk* (PSRW). Having reduced the subgraph counting task to Equation (1), we proceed with implementing Ripple.

4.2 Ergodicity-Preserving Stratification for Subgraph Counting

Toward using Ripple, we propose an Ergodicity-Preserving Stratification of \mathcal{G} via the stratification function ρ .

Proposition 2 (*EPS* for subgraphs) Consider the set of n_1 seed subgraphs \mathcal{I}_1 whose vertex sets in G are pairwise non-intersecting. Let $V(\mathcal{I}_1) \triangleq \bigcup_{\ddot{s} \in \mathcal{I}_1} V(\ddot{s})$ be the set of all vertices in G forming subgraphs in \mathcal{I}_1 . Let DIST(u) be the shortest path distance from $u \in V$ to any vertex in $V(\mathcal{I}_1)$. Define ρ as

$$\rho(s) = 1 + \sum_{u \in V(s)} (\text{dist}(u) + \mathbf{1}\{u \in V(\mathcal{I}_1) \backslash V^*)\}) ,$$

where V^* is the largest connected subset of V(s) such that $V^* \subseteq V(\ddot{s})$ for some seed vertex $\ddot{s} \in \mathcal{I}_1$ with ties broken arbitrarily. If \mathcal{I}_1 contains a subgraph from each connected component of G, the stratification from Definition 4 generated using ρ is an Ergodicity-Preserving Stratification (Definition 6).

DIST can be precomputed for all $u \in V$ using a single BFS in O(|V| + |E|), and ρ can be computed in O(k). Although R is unknown a priori, it is upper bounded as $(k-1) \cdot D_G$, where D_G is the diameter of G and the Ripple estimator simply ignores *empty* strata, i.e., strata in which the estimated degree of the supernode $\mathbf{d}(\zeta_r) = 0$. To control bias, we aim to reduce $\max_{u \in V} \mathrm{DIST}(u)$ by recruiting seed subgraphs in \mathcal{I}_1 , which are far apart in G.

4.3 Miscellaneous Optimizations

Controlling Memory through Streaming. In each pair of strata r < t, Definition 8 uses tours \mathcal{T}_r^{\dagger} to compute $\hat{\mu}(\mathcal{T}_{2:r}^{\dagger};f)$, $\hat{\beta}_{r,t}$ and $\hat{\mathbf{U}}_{r,t}$, which are, respectively, the estimates of $\mu(\mathcal{J}_r)$ and the size of and sample from the set of vertices in \mathcal{I}_t connected to \mathcal{I}_r . Although $\hat{\mu}(\mathcal{T}_{2:r}^{\dagger};f)$ and $\hat{\beta}_{r,t}$ can be computed as running sums, storing $\hat{\mathbf{U}}_{r,t}$ requires memory on the order of the sum of all tour lengths, which is random. Our solution is to use Algorithm R (Vitter, 1985), to sample a fixed-size (M) sample without replacement from all the tours in \mathcal{T}_r^{\dagger} (See Appendix E.1). We note that although the hyperparameter M controls memory, it may introduce bias when the number of tours $|\mathcal{T}_r^{\dagger}| > M$ due to (possible) oversampling, which we observe in Figure 6 (Appendix F).

Speeding up Subgraph Random Walks. To sample a random walk in the HON, naively sampling u.a.r. from the neighborhood of a k-1-CIS requires $O(k^4\Delta_G)$ operations, where Δ_G is the maximum degree in the input graph (see Appendix E.2). In Algorithm 1 we propose a rejection sampling algorithm that does so efficiently using articulation points (Hopcroft and Tarjan, 1973).

```
Algorithm 1: Efficient Neighborhood Sampling in \mathcal{G}^{(k-1)}
```

```
Input: k - 1-CIS s, Graph G
    Output: x \sim \text{UNIF}(\mathbf{N}_{\mathcal{G}^{(k-1)}}(s))
1 Let \deg_s = \sum_{u \in V(s)} \mathbf{d}(u) and \mathcal{A}_s be the articulating points of s
   while True do
         Sample u from V(s) w.p. \propto \deg_s -\mathbf{d}(u);
                                                                      // u is the vertex to remove
         Sample a from V(s)\setminus\{u\} w.p. \propto \mathbf{d}(a)
5
         Sample v \sim \text{UNIF}(\mathbf{N}(a));
                                                                           //\ v is the vertex to add
6
         BIAS = |N(v) \cap V(s) \setminus \{u\}|;
                                                                                  // v's sampling bias
         if UNIF(0,1) < 1/BIAS then
              x = G(V(s) \cup \{v\} \setminus \{u\})
9
              if u \neq v and (u \notin A_s \text{ or } x \text{ is connected}) then
                                                                                // Connectivity Check
10
                   return x:
```

Proposition 3 Given a subgraph $s \in \mathcal{V}^{(k-1)}$, Algorithm 1 samples u.a.r. from $\mathbf{N}_{\mathcal{G}^{(k-1)}}(s)$ in $O(k^2 \frac{\Delta_s + k|A_s|}{k - |A_s|})$ expected time, where $\Delta_s \triangleq \max_{u \in V(s)} \mathbf{d}_G(u)$ is the maximum degree of vertices in s, and A_s contains articulation points of s.

Therefore, the running time of Algorithm 1 is $\in O(k\Delta_s + k^2)$ when s is dense $(|A_s| \approx 0)$ and increases to $O(k^2\Delta_s + k^4)$ for sparse subgraphs, which is faster than the naive algorithm.

From Error Bounds to Tour Counts. Ripple auto-decides the number of RWTs required in each stratum based on an approximate error bound ϵ provided as input such that the number of tours $\to \infty$ as $\epsilon \to 0$, and the Ripple estimate

converges to the ground truth (Theorem $\boxed{1}$). Specifically, RWTs are sampled until we satisfy

$$\hat{\sigma}(\mathcal{T}_r^{\dagger}; f_1) / \sqrt{|\mathcal{T}_r^{\dagger}|} \le \epsilon \,\hat{\mu}(\mathcal{T}_r^{\dagger}; f_1) \,,$$
 (10)

where $\hat{\mu}(\mathcal{T}_r^{\dagger}; f_1)$ is the Ripple estimate from Equation (8) of the number of edges in the r-th graph stratum \mathcal{G}_r (i.e., $f_1(\cdot) = 1$), and $\hat{\sigma}^2(\mathcal{T}_r^{\dagger}; f_1) = \widehat{\text{Var}}_{\mathbf{X} \sim \mathcal{T}_r^{\dagger}}(\hat{\mu}(\mathbf{X}; f_1))$ is the former's empirical variance over tours.

Performance Guarantees. Ignoring the complexity of loading the input graph into memory, we show that for subgraph counting, the memory and time requirements of Ripple are a polynomial in k. In Appendix $\boxed{\mathbb{E}}$, we state and prove a detailed version in which the complexity also depends polynomially on the diameter and maximum degree of G and is invariant to |V| and |E|.

Proposition 4 Assuming a constant m RWTs sampled per stratum and ignoring graph loading, the Ripple estimator for k-CIS counts detailed in Appendix E-Algorithm 2 has total memory and time complexity in $\widehat{O}(k^3 + |\mathcal{H}|)$ and $\widehat{O}(k^7 + |\mathcal{H}|)$, respectively, when all factors other than k and $|\mathcal{H}|$ are ignored. More details for subgraph counting with Ripple are provided in Appendix E

5 Experiments and Results

We now evaluate the Ripple estimator for k-node subgraph (k-CIS) counts on large-real-world networks. We show that Ripple outperforms the state-of-the-art method in terms of time and space and that Ripple converges to the ground truth for various pattern sizes as hyperparameters are varied. Additional experiments that evaluate the parallelism, etc., are deferred to Appendix F Our code is available at https://github.com/dccspeed/ripple.

- Execution environment. Our experiments were performed on a dual Intel Xeon Gold 6254 CPU with 72 virtual cores (total) at 3.10 GHz and 392 GB of RAM. In addition, this machine is equipped with a fast SSD NVMe PCIex4 with 800 GB of free space available.
- Baselines. We use Motivo (Bressan et al.) 2019), a fast and parallel C++ system for subgraph counting, as the baseline because it is the only method capable of counting large patterns (k>6), to the best of our knowledge. Additionally, notice that existing MCMC methods for subgraph counting, such as IMPRG (Chen and Lui) 2018) and RGPM (Teixeira et al.) 2018), cannot count beyond k=5 in practice.
- Datasets. We use large networks from SNAP (Leskovec and Krevl) [2014), representing diverse domains, which have been used to evaluate many subgraph counting algorithms (Bressan et al., [2018] [2019). Table [1] presents the basic features of these datasets, including the order of magnitude of the Ripple estimates of the subgraph counts $|\mathcal{V}^{(k)}|$, k = 6, 8, 10, 12.
- Hyper-parameters \mathcal{I}_1 , M and ϵ . Finally, we evaluate the trade-off between accuracy and resource consumption by varying the aforementioned hyperparameters, detailed in Sections [4.2] and [4.3] (M is evaluated in Appendix F)

Graph	$ \mathbf{V} $	$ \mathbf{E} $	D_{G}	$\Delta_{ m G}$	$ Magn \mathcal{V}^{(6)} $	$rac{\mathbf{itude\ of}}{ \mathcal{V}^{(8)} }$	Est. # $ \mathcal{V}^{(10)} $	$ \mathcal{V}^{(12)} $
Amazon	334,863	925,872	44	549	10^{11}	10^{15}	10^{19}	10^{22}
DBLP	317,080	1,049,866	21	343	10^{12}	10^{16}	10^{19}	10^{23}
Cit-Pat.	3,774,768	16,518,948	22	793	10^{14}	10^{18}	10^{22}	10^{26}
Pokec	1,632,803	30,622,564	11	14,854	10^{18}	10^{25}	10^{32}	10^{38}
LiveJ.	3,997,962	34,681,189	17	14,815	10^{19}	10^{25}	10^{32}	10^{38}
Orkut	3,072,441	117, 185, 083	9	33,313	10^{21}	10^{28}	10^{35}	10^{43}

Table 1: The graphs that we used along with their diameter D_G , maximum degree Δ_G and the estimated orders of magnitude of k-CIS counts, $|\mathcal{V}^{(k)}|$.

5.1 Scalability Assessment

We start by assessing the scalability of the methods when estimating k-CIS counts for $k \geq 6$. To the best of our knowledge, Motivo is the only existing method capable of estimating these patterns. Motivo has two phases: a build-up phase, which constructs an index table in the disk, and a sampling phase that queries this table. We only measure the time taken by the build-up phase and the out-of-core (disk) usage because this is a bottleneck for Motivo. As such, we report the best-case scenario for Motivo, and the reported values are lower bounds for the actual time and space requirement. For Ripple, we report the total time and the RAM usage as the space cost because our method works purely in memory. Both methods were executed using all threads available.

In Tables 2 and 3 we compare the running time and space usage of Ripple and Motivo. We also report their rate of increase in terms of the subgraph size k in columns $^{\text{Time}^{(k)}}/_{\text{Time}^{(k-2)}}$ and $^{\text{Space}^{(k)}}/_{\text{Space}^{(k-2)}}$. We fix $\epsilon = 0.003$, $|\mathcal{I}_1| = 10^4$ and $M = 10^7$ based on the analysis in Section 5.2 and Appendix F For Motivo, we follow the authors' suggestions. In Appendix F Table 5 we report the dispersion $\frac{\max - \min}{\max}$ of the Ripple estimates generated in the measured runs to ensure that the results are not arbitrary.

Running time Scalability (Table 2). Although Motivo outperforms Ripple for k=6,8, it does not scale well for k=10,12, where the execution terminates because of insufficient storage space. Particularly, for DBLP, Motivo required approximately 10 minutes to process 10-CIS but almost 9 hours for 12-CIS, a growth rate of $58\times$. On the other hand, Ripple not only succeeded in all configurations in less than 4 hours on average but also exhibited a smoother growth in running time, with the largest increase ratio being $2.7\times$, observed for DBLP and LiveJournal when k went from 8 to 10. Furthermore, $\frac{\text{Time}^{(k)}}{\text{Time}^{(k-2)}} < (\frac{k}{(k-2)})^7$ in all cases according to Proposition 10

Space Scalability (Table 3). The trends in space usage mirror those of the running time, where we see an almost exponential increase w.r.t. k for Motivo compared to a near constant increase for Ripple despite its polynomial complexity (Proposition 10). For example, in Amazon, Motivo's space demand increases by $7.5 \times$ when k goes from 6 to 8 and increases to $12 \times$ from 10 to

		Motivo Build-up only		Ripple (ϵ =	Ripple	
Dataset	k	Time (hrs)	$\frac{\operatorname{Time}^{(k)}}{\operatorname{Time}^{(k-2)}}$	Time (hrs)	$\frac{\operatorname{Time}^{(k)}}{\operatorname{Time}^{(k-2)}}$	gain (hrs)
Amazon	6	$\boldsymbol{0.002 \pm 0.000}$	_	0.020 ± 0.000	_	-0.018
	8	0.006 ± 0.000	$3 \times$	0.029 ± 0.000	$1.4 \times$	-0.023
Amazon	10	0.082 ± 0.000	$13.7 \times$	0.056 ± 0.000	$1.9 \times$	+0.026
	12	3.630 ± 0.002	$44.3 \times$	$\boldsymbol{0.095 \pm 0.002}$	$1.7 \times$	+3.535
	6	0.002 ± 0.000	_	0.013 ± 0.000	_	-0.011
DBLP	8	$\boldsymbol{0.007 \pm 0.000}$	$3.5 \times$	0.030 ± 0.000	$2.3 \times$	-0.023
DBLP	10	0.156 ± 0.000	$22.3 \times$	$\boldsymbol{0.082 \pm 0.000}$	$2.7 \times$	+0.074
	12	9.099 ± 0.002	$58.3 \times$	$\boldsymbol{0.105 \pm 0.002}$	$1.3 \times$	+8.994
	6	$\boldsymbol{0.022 \pm 0.000}$	_	0.033 ± 0.000	_	-0.011
ъ.	8	0.098 ± 0.000	$4.5 \times$	0.051 ± 0.000	$1.5 \times$	+0.047
Patents	10	> 1.1 hrs, crashed	_	$\boldsymbol{0.090 \pm 0.001}$	$1.8 \times$	_
	12	> 0.5 hrs, crashed	_	0.117 ± 0.003	$1.3 \times$	_
	6	0.012 ± 0.000	_	0.459 ± 0.142	_	-0.447
Pokec	8	$\boldsymbol{0.128 \pm 0.000}$	$10.7 \times$	0.759 ± 0.282	$1.7 \times$	-0.631
1 OKeC	10	5.965 ± 0.000	$46.6 \times$	$\boldsymbol{1.400 \pm 0.592}$	$1.8 \times$	+4.565
	12	> 1.5 hrs, crashed	_	$\boldsymbol{1.469 \pm 0.334}$	$1 \times$	_
	6	0.024 ± 0.000	_	0.351 ± 0.009	_	-0.327
LiveJ.	8	$\boldsymbol{0.205 \pm 0.000}$	$8.5 \times$	0.642 ± 0.074	$1.8 \times$	-0.437
LiveJ.	10	> 2.3 hrs, crashed	_	$\boldsymbol{1.76 \pm 1.550}$	$2.7 \times$	_
	12	> 0.7 hrs, crashed	_	$\boldsymbol{2.189 \pm 1.350}$	$1.2 \times$	_
	6	0.032 ± 0.000	_	0.669 ± 0.026	_	-0.637
Orkut	8	$\boldsymbol{0.585 \pm 0.006}$	$18.3 \times$	1.744 ± 0.983	$2.6 \times$	-1.159
Orkut	10	> 8.9 hrs, crashed	_	$\boldsymbol{2.633 \pm 1.065}$	$1.5 \times$	_
	12	> 1.8 hrs, crashed	_	$\boldsymbol{3.967 \pm 3.162}$	$1.5 \times$	_

Table 2: Running time comparison between Ripple and Motivo. The last column shows that for large k, Ripple provides gains of up to 9 hours when Motivo can run to completion. Motivo crashes for large k on large graphs.

12. Ripple's largest rate of increase is $1.4 \times$ when k goes from 6 to 8 for DBLP, and it saves up to 600 GB of space when Motivo does not crash.

5.2 Accuracy and Convergence Assessment

Next, we evaluate the accuracy and convergence of Ripple on small and large subgraph patterns, where the former refers to subgraph sizes in which the number of isomorphic subgraphs can be exactly computed using ESCAPE (Pinar et al., [2017]), i.e., $k \leq 5$.

Accuracy on Small k. For $k \in \{3,5\}$, we evaluate the L2-norm between the Ripple estimate and the exact value of the count vector $\mathcal{C}^{(k)}$ (Equation [9]) of all non-isomorphic subgraph patterns. Figure 2 shows results for k=5 (where the number of patterns of interest $|\mathcal{H}|=21$) for different settings of the parameters ϵ and $|\mathcal{I}_1|$. In all datasets, we note that the L2-norm decreases as ϵ decreases from 0.3 to 0.003 and as $|\mathcal{I}_1|$ increases from 100 to 10^4 . Between the worst setting, $(\epsilon, |\mathcal{I}_1|) = (0.3, 100)$, and the best $(\epsilon, |\mathcal{I}_1|) = (0.003, 10^4)$, we see an error reduction close to an order of magnitude. This is due to Theorem 2 and Lemma 1 because reducing ϵ increases the number of tours, lowers the error

		Motivo Build-up only		Ripple (ϵ)	Ripple	
Dataset	k	Space (GB)	$\frac{\operatorname{Space}^{(k)}}{\operatorname{Space}^{(k-2)}}$	Space (GB)	$\frac{\operatorname{Space}^{(k)}}{\operatorname{Space}^{(k-2)}}$	gain (GB)
Amazon	6	$\boldsymbol{0.53 \pm 0.00}$	_	4.69 ± 0.06	_	-4.16
	8	4.00 ± 0.00	$7.5 \times$	5.73 ± 0.12	$1.2 \times$	-1.73
Amazon	10	48.00 ± 0.00	$12 \times$	$\boldsymbol{7.38 \pm 0.36}$	$1.3 \times$	+40.62
	12	559 ± 0.00	$11.6 \times$	$\boldsymbol{9.09 \pm 1.02}$	$1.2 \times$	+549.91
	6	0.50 ± 0.00	_	4.58 ± 0.02	_	-4.08
DBLP	8	4.00 ± 0.00	8×	6.31 ± 0.00	$1.4 \times$	-2.31
DDLL	10	50.00 ± 0.00	$12.5 \times$	$\boldsymbol{7.99 \pm 0.01}$	$1.3 \times$	+42.01
	12	611.00 ± 0.00	$12.2 \times$	10.45 ± 0.02	$1.3 \times$	+600.55
	6	7.00 ± 0.00	_	11.50 ± 0.05	_	-4.5
Patents	8	66.00 ± 0.00	$9.4 \times$	13.80 ± 0.03	$1.2 \times$	+52.2
ratents	10	> 800, crashed	_	15.85 ± 0.08	$1.1 \times$	> 800
	12	> 800, crashed	_	18.12 ± 0.10	$1.1 \times$	> 800
	6	3.7 ± 0.00	_	13.69 ± 0.06	_	-9.99
Pokec	8	36.00 ± 0.00	$9.7 \times$	17.17 ± 0.03	$1.3 \times$	18.83
Рокес	10	407.00 ± 0.00	$11.3 \times$	20.31 ± 0.01	$1.2 \times$	+386.69
	12	$> 800, \ crashed$	_	22.82 ± 0.03	$1.1 \times$	> 800
	6	7.70 ± 0.00	_	18.26 ± 0.02	_	-10.56
LiveJ.	8	73.00 ± 0.00	$9.5 \times$	21.26 ± 0.00	$1.2 \times$	+51.74
LiveJ.	10	> 800, crashed	_	24.43 ± 0.72	$1.1 \times$	> 800
	12	$> 800, \ crashed$	_	27.75 ± 0.00	$1.1 \times$	> 800
	6	$\textbf{7.90} \pm \textbf{0.00}$	_	40.38 ± 0.00	_	-32.48
Orkut	8	78.00 ± 0.000	$9.9 \times$	43.49 ± 0.00	$1.1 \times$	+34.51
Orkut	10	$> 800, \ crashed$	_	46.63 ± 0.00	$1.1 \times$	> 800
	12	$> 800,\ crashed$	_	49.73 ± 0.00	$1.1 \times$	> 800

Table 3: Space usage comparison between Ripple and Motivo.Motivo runs out of disk space for larger datasets in which $k \geq 10$, while Ripple scales almost linearly. Ripple saves up to 600 GB of space when Motivo can run.

and therefore leads to reduced error propagation. Increasing \mathcal{I}_1 also reduces the number of strata and therefore error propagation. Results for k=5 using the L- ∞ norm are deferred to Appendix F-Figure 5.

Convergence for Large k. When k > 5, subgraph counts for real-world graphs are computationally intractable. Therefore, we show that Ripple converges in these cases as we increase the computing effort. Consider the hypothesis that sparse patterns are frequent in power-law networks as k increases. To glean empirical evidence for this, we choose an appropriate pattern set \mathcal{H} and equivalence relationship in Definition $\boxed{10}$ and we use Ripple to compute the total number of k-CIS s and the number of sparse subgraphs and stars. A subgraph is defined as sparse if its density lies between 0 and 0.25, according to $\boxed{\text{Liu and Wong}}$ ($\boxed{2008}$). In Figure $\boxed{3}$ we show that Ripple converges for all datasets, and as expected, most patterns are sparse, with close to half of the patterns in many of the studied networks being stars. This proportion is attenuated in DBLP and Patents, where dense substructures naturally emerge from collaboration/citation among the authors that these graphs represent.

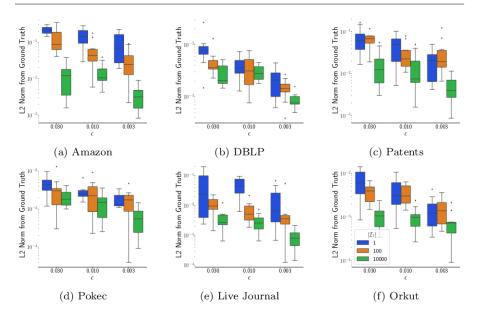


Fig. 2: Accuracy and convergence analysis for 5-CIS s. We plot the L2-norm between the Ripple estimate and the exact value of the count vector $\mathcal{C}^{(5)}$ (Equation 9) of all non-isomorphic subgraph patterns against various configurations of the parameters ϵ and $|\mathcal{I}_1|$. As expected, the accuracy improves as the error bound ϵ decreases and the number of seed subgraphs $|\mathcal{I}_1|$ increases. Each box and whisker represents 10 runs.

6 Related Work

For better presentation, we split this section into two parts: (1) parallel *MCMC* techniques and (2) methods for subgraph counting.

Parallel MCMC through Splitting. Since Nummelin (1978); Athreya and Ney (1978), multiple techniques have been proposed to circumvent the burn-in period by splitting the chain into i.i.d. sample paths. This approach allows practitioners to compute unbiased estimates in parallel and determine confidence intervals. Perfect sampling methods based on coupling (Propp and Wilson, 1996) require the transitions to be monotonic w.r.t. some ordering over the state space, and annealing/tempering (Neal, 2001) methods require some notion of temperature, which are absent in general graph random walks. Methods such as (Mykland et al., 1995; Jacob et al., 2020; Glynn and Rhee, 2014) require a minorization condition to hold, albeit implicitly.

Regeneration point-based methods on finite state chains (Cooper et al., 2016) Massoulié et al., 2006) Avrachenkov et al., 2016, 2018 Savarese et al., 2018 Teixeira et al., 2018) are more general because they only rely on standard ergodicity conditions. Although Cooper et al., (2016); Massoulié et al.

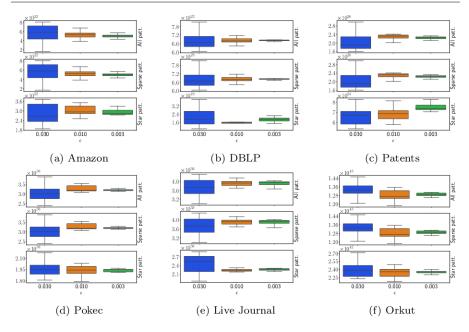


Fig. 3: Convergence of Ripple estimates of 12-CIS pattern counts. We estimate the total number of subgraphs $|\mathcal{V}^{(12)}|$ and the number of sparse patterns and stars. Estimates over 10 runs are presented as box and whiskers plots, which exhibit a reduction in variance as ϵ increases. Indeed, almost all patterns are sparse, and the most frequent substructure is a star.

(2006) used tours to estimate graph properties, Avrachenkov et al. (2016, 2018) proposed supernodes to reduce running times. The studies in Savarese et al. (2018); Teixeira et al. (2018) further used supernode-based tours to estimate gradients in *RBM*s and to count subgraphs. To the best of our knowledge, no existing regeneration point method controls running times through stratification.

Subgraph Counting through Sampling. Many random walk algorithms have been proposed to sample subgraphs, with some methods only capable of estimating subgraph pattern distributions, which is much easier than estimating counts. The studies of GUISE (Bhuiyan et al., 2012) and RSS (Matsuno and Gionis 2020) use a Metropolis-Hastings (Hastings 1970) walk, and the latter improves the mixing time of the underlying Markov chain using canonical paths (Sinclair, 1992). Waddling (Han and Sethu, 2016) and IMPRG Chen and Lui (2018) perform a simple random walk over the input graph and use specialized estimators to sample 5-node patterns. Although PSRW (Wang et al., 2014) first proposed the HON-based random walk and RGPM (Teixeira et al., 2018) used tours on it to estimate subgraph counts, both are limited to $k \leq 5$ due to the size of the HON.

Multiple attempts to Monte Carlo sample subgraphs have been proposed whose scaling is limited because of the complexity of computing either the importance weights, rejection rate or variance (Kashtan et al.) 2004 Wernicke 2006; Iyer et al.) 2018; Yang et al. 2018; Wang et al. 2018). Efficient methods that sample dense regions/subgraphs are unfortunately not extensible to sparse patterns Jain and Seshadhri (2017) 2020). Motivo (Bressan et al.) 2018 2019) is an example of color-coding methods in which an index table is built using a deterministic dynamic programming algorithm, which is then exploited to sample subgraphs uniformly and independently. However, CC methods suffer from the exponential time and space complexities associated with building and accessing the index table. Motivo proposed succinct index tables and efficient out-of-core I/O mechanisms to ameliorate this issue and extended the applicability of CC methods to larger subgraphs. Please, check Ribeiro et al. (2019) for an extensive survey on subgraph counting methods.

7 Conclusions

In this paper, we propose the Ripple estimator that uses sequentially stratified regenerations to control the running time of a random walk tour-based MCMC. We prove that the estimator is consistent (w.r.t. the number of random walk tours) and that the time and memory complexity of our implementation for the subgraph counting problem is linear in the number of patterns of interest and polynomial in the subgraph size. We empirically verify our claims on multiple graph datasets and show that Ripple can accurately estimate subgraph counts with a smaller memory footprint compared to that of the state-of-theart Motivo (Bressan et al., 2019). Ripple is currently the only subgraph pattern count estimator that can estimate k = 10, 12 node patterns in million-node graphs. Beyond our specific application, Ripple provides a promising way to expand the sphere of influence of regenerative simulation in discrete reversible MCMC.

References

Aldous D, Fill JA (2002) Reversible markov chains and random walks on graphs

Athreya KB, Ney P (1978) A new approach to the limit theory of recurrent markov chains. Trans of the American Mathematical Society 245:493–501

Avrachenkov K, Ribeiro B, Sreedharan JK (2016) Inference in osns via lightweight partial crawls. In: ACM SIGMETRICS, pp 165–177

Avrachenkov K, Borkar VS, Kadavankandy A, Sreedharan JK (2018) Revisiting random walk based sampling in networks: evasion of burn-in period and frequent regenerations. Computational social networks 5(1):1-19

Bhuiyan MA, Rahman M, Rahman M, Al Hasan M (2012) Guise: Uniform sampling of graphlets for large graph analysis. In: 2012 IEEE 12th Int. Conf. on Data Mining, IEEE, pp 91–100

- Bremaud P (2001) Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues. Texts in Applied Mathematics, Springer New York
- Bressan M, Chierichetti F, Kumar R, Leucci S, Panconesi A (2018) Motif counting beyond five nodes. ACM TKDD 12(4)
- Bressan M, Leucci S, Panconesi A (2019) Motivo: Fast motif counting via succinct color coding and adaptive sampling. Proc VLDB Endow
- Chen X, Lui JC (2018) Mining graphlet counts in online social networks. ACM TKDD 12(4):1-38
- Cooper C, Radzik T, Siantos Y (2016) Fast low-cost estimation of network properties using random walks. Internet Mathematics
- Diaconis P, Stroock D (1991) Geometric bounds for eigenvalues of markov chains. The Annals of Applied Probability pp 36–61
- Geman S, Geman D (1984) Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. IEEE Trans on pattern analysis and machine intelligence
- Geyer CJ (1992) Practical markov chain monte carlo. Statistical science
- Glynn PW, Rhee Ch (2014) Exact estimation for markov chain equilibrium expectations. Journal of Applied Probability 51(A):377–389
- Han G, Sethu H (2016) Waddling random walk: Fast and accurate mining of motif statistics in large graphs. In: ICDM, IEEE, pp 181–190
- Hastings WK (1970) Monte carlo sampling methods using markov chains and their applications. Biometrika
- Hopcroft J, Tarjan R (1973) Algorithm 447: efficient algorithms for graph manipulation. Communications of the ACM 16(6):372–378
- Iyer AP, Liu Z, Jin X, Venkataraman S, Braverman V, Stoica I (2018) {ASAP}: Fast, approximate graph pattern mining at scale. In: OSDI, pp 745–761
- Jacob PE, O'Leary J, Atchadé YF (2020) Unbiased markov chain monte carlo methods with couplings. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 82(3):543–600
- Jain S, Seshadhri C (2017) A fast and provable method for estimating clique counts using turán's theorem. In: WWW, WWW '17, pp 441–449
- Jain S, Seshadhri C (2020) Provably and efficiently approximating near-cliques using the turán shadow: Peanuts. In: WWW 2020, pp 1966-1976
- Kashtan N, Itzkovitz S, Milo R, Alon U (2004) Efficient sampling algorithm for estimating subgraph concentrations and detecting network motifs. Bioinformatics 20(11):1746–1758
- Leskovec J, Krevl A (2014) SNAP Datasets: Stanford large network dataset collection. http://snap.stanford.edu/data
- Liu G, Wong L (2008) Effective pruning techniques for mining quasi-cliques. In: ECML-PKDD
- Massoulié L, Le Merrer E, Kermarrec AM, Ganesh A (2006) Peer counting and sampling in overlay networks: random walk methods. In: PODC
- Matsuno R, Gionis A (2020) Improved mixing time for k-subgraph sampling. In: Proc. of the 2020 SIAM Int. Conf. on Data Mining, SIAM, pp 568–576
- Mykland P, Tierney L, Yu B (1995) Regeneration in markov chain samplers. Journal of the American Statistical Association 90(429):233–241

Neal RM (2001) Annealed importance sampling. Statistics and computing Neiswanger W, Wang C, Xing E (2014) Asymptotically exact, embarrassingly parallel mcmc. UAI

- Nummelin E (1978) A splitting technique for harris recurrent markov chains. magazine for "u r probability theory and related areas 43(4):309-318
- Pinar A, Seshadhri C, Vishal V (2017) Escape: Efficiently counting all 5-vertex subgraphs. In: WWW, pp 1431–1440
- Propp JG, Wilson DB (1996) Exact sampling with coupled markov chains and applications to statistical mechanics. Random Structures & Algorithms
- Ribeiro B, Towsley D (2012) On the estimation accuracy of degree distributions from graph sampling. In: CDC
- Ribeiro P, Paredes P, Silva ME, Aparicio D, Silva F (2019) A survey on subgraph counting: Concepts, algorithms and applications to network motifs and graphlets. arXiv preprint arXiv:191013011
- Robert C, Casella G (2013) Monte Carlo statistical methods. Springer Science & Business Media
- Rosenthal JS (1995) Minorization conditions and convergence rates for markov chain monte carlo. Journal of the American Statistical Association
- Savarese PH, Kakodkar M, Ribeiro B (2018) From monte carlo to las vegas: Improving restricted boltzmann machine training through stopping sets. In: AAAI
- Sinclair A (1992) Improved bounds for mixing rates of markov chains and multicommodity flow. Combinatorics, probability and Computing 1(4)
- Teixeira CH, Cotta L, Ribeiro B, Meira W (2018) Graph pattern mining and learning through user-defined relations. In: ICDM, IEEE, pp 1266–1271
- Vitter JS (1985) Random sampling with a reservoir. ACM Trans on Mathematical Software (TOMS)
- Wang P, Lui JCS, Ribeiro B, Towsley D, Zhao J, Guan X (2014) Efficiently estimating motif statistics of large networks. ACM TKDD 9(2)
- Wang P, Zhao J, Zhang X, Li Z, Cheng J, Lui JCS, Towsley D, Tao J, Guan X (2018) Moss-5: A fast method of approximating counts of 5-node graphlets in large graphs. IEEE TKDE
- Wernicke S (2006) Efficient detection of network motifs. IEEE/ACM Trans Comput Biol Bioinformatics
- Wilkinson DJ (2006) Parallel bayesian computation. Statistics Textbooks and Monographs 184:477
- Yang C, Lyu M, Li Y, Zhao Q, Xu Y (2018) Ssrw: a scalable algorithm for estimating graphlet statistics based on random walk. In: Int. Conf. on Database Systems for Advanced Applications, Springer, pp 272–288
- Yao P, Zheng L, Zeng Z, Huang Y, Gui C, Liao X, Jin H, Xue J (2020) A locality-aware energy-efficient accelerator for graph mining applications. In: MICRO

A Notation

The most important notations from the paper are summarized in Table 4

Symbol	Explanation
G = (V, E)	The graph where we have neighbor query access and whose edge sum
	is being computed.
$N(u), \mathbf{d}(u)$	Neighborhood and degree of a vertex in \mathcal{G} if no subscript is specified.
$\mu(\mathcal{E})$	The sum over edges in \mathcal{E} (or some subset) of some function f .
Φ , $p_{\Phi}(u,v)$, $\pi_{\Phi}(u)$	The random walker on \mathcal{G} , its transition probability and stationary
	distribution.
$\mathbf{X}, \boldsymbol{\xi}, \boldsymbol{\mathcal{T}}$	an RWT (tour), its length and a set of m RWT s.
$\hat{\mu}_*(\mathcal{T};f,\mathcal{G})$	an RWT Estimate of $\mu(\mathcal{E})$.
	The spectral gap of the transition probability matrix of a chain.
$\hat{\sigma}(\cdot)^2$	Empirical variance of an RWT Estimate.
ζ_I,\mathcal{G}_I	Collapsed state and graph obtained by collapsing $I \subset \mathcal{V}$.
q < r < t $\rho \colon \mathcal{V} \to \{1, \dots, R\}$	Strata ids always used in the same order $1 \le q < r < t \le R$.
$ \begin{array}{c} \rho \colon V \to \{1, \dots, n\} \\ \mathcal{I}_r, \mathcal{J}_r \end{array} $	Stratification function. r-th vertex and edge stratum.
$\mathcal{G}_r = (\mathcal{V}_r, \mathcal{E}_r)$	r-th graph stratum.
ζ_r, \mathcal{T}_r	Supernode in each stratum and a set of m_r perfectly sampled tours
\$1,71	from ζ_r .
$\mathbf{d}(\zeta_r), p_{\mathbf{\Phi}_r}(\zeta_r, \cdot)$	Degree and transition probability out of the supernode.
$\widehat{\mathbf{d}}(\zeta_r), \widehat{p}_{\mathbf{d}_r}(\zeta_r, \cdot)$	Estimated degree and transition probability out of the supernode.
$\mathcal{T}_{a}^{\dagger}$	m_q RWTs samples using supernode estimates.
$\widehat{\widehat{\beta}}_{a,n}$	The estimate of the number of edges between \mathcal{I}_q and \mathcal{I}_r .
$\widehat{\mathbf{d}}(\zeta_r), \widehat{p}_{\mathbf{\Phi}_r}(\zeta_r, \cdot)$ \mathcal{T}_q^{\dagger} $\widehat{\beta}_{q,r}$ $\widehat{\mathbf{U}}_{q,r}$	Multiset of states visited by \mathcal{T}_q^{\dagger} that lie in \mathcal{I}_r .
$\hat{\mu}_{Ripple},\hat{\mu}\left(\mathcal{T}_{2:r}^{\dagger};f ight)$	Overall and per-stratum Ripple estimate.
$\delta_r, \nu_r, \lambda_r$	Spectral gap and the errors in the supernode estimates in the r -th
-,,-,,-,	stratum.
G = (V, E, L)	The labelled input graph in which we want to count subgraphs.
G(V')	Subgraph induced by V' in G .
\mathcal{H}	Nonequivalent (non-isomorphic) patterns of interest.
$\mathcal{C}^{(k)} = (\mathcal{C}_H^{(k)})_{H \in \mathcal{H}}$ $\mathcal{G}^{(k)} = (\mathcal{V}^{(k)}, \mathcal{E}^{(k)})$	The subgraph pattern count vector.
$\mathcal{G}^{(k)} = (\mathcal{V}^{(k)}, \mathcal{E}^{(k)})$	The k-HON that provides neighborhood query access and is used to
,	count subgraphs.
$\gamma(u,v)$	Number of edges in $\mathcal{E}^{(k-1)}$ that represent the same subgraph as (u, v) .
DIST(u)	Shortest path distance from $u \in V$ to any seed vertex in $V(\mathcal{I}_1)$.
V^*	The largest connected subset of $V(s)$ that constitutes an intersection
	between s and $\ddot{s} \in \mathcal{I}_1$.
M	Reservoir size.
Δ_G, D_G	Maximum degree in and diameter of G .
\mathcal{A}_s	Articulation points in s.
ϵ	Per-stratum error bound used to control tour count.

Table 4: Table of Notations

B Proofs for Section 2

$B.1\ MCMC$ Estimates

Given a graph \mathcal{G} , when the $|\mathcal{E}|$ is unknown, the MCMC estimate of $\mu(\mathcal{E})/|\mathcal{E}|$ is given by:

Proposition 5 (MCMC Estimate (Geyer) 1992; Geman and Geman 1984; Hastings, 1970)) When $\mathcal G$ from Definition 1 is connected, the random walk Φ is reversible and

positive recurrent with stationary distribution $\pi_{\Phi}(u) = \mathbf{d}(u)/2|\mathcal{E}|$. Then, the MCMC estimate

$$\hat{\mu}_0\left((X_i)_{i=1}^t\right) = \frac{1}{t-1} \sum_{i=1}^{t-1} f(X_i, X_{i+1}),$$

computed using an arbitrarily started sample path $(X_i)_{i=1}^t$ from Φ is an asymptotically unbiased estimate of $\mu(\mathcal{E})/|\mathcal{E}|$. When \mathcal{G} is non-bipartite, i.e., Φ is aperiodic, and t is large, $\hat{\mu}_0$ converges to $\mu(\mathcal{E})/|\mathcal{E}|$ as

$$\left| \mathbb{E}[\hat{\mu}_0 \left((X_i)_{i=1}^t \right)] - \mu(\mathcal{E}) / |\mathcal{E}| \right| \le B \frac{C}{t \delta(\boldsymbol{\Phi})},$$

where $\delta(\Phi)$ is the spectral gap of Φ and $C \triangleq \sqrt{\frac{1-\pi_{\Phi}(X_1)}{\pi_{\Phi}(X_1)}}$ such that $f(\cdot) \leq B$.

Proof (Asymptotic unbiasedness) Because \mathcal{G} is undirected, finite and connected, Φ is a finite state space, irreducible, time-homogeneous Markov chain and is therefore positive recurrent (Bremaud, [2001] 3-Thm.3.3). The reversibility and stationary distribution holds from the detailed balance test (Bremaud, [2001] 2-Cor.6.1) because

$$\pi_{\mathbf{\Phi}}(u) \, p_{\mathbf{\Phi}}(u,v) = \pi_{\mathbf{\Phi}}(v) \, p_{\mathbf{\Phi}}(v,u) = \frac{\mathbf{1}\{(u,v) \in \mathcal{E}\}}{2|\mathcal{E}|} \, .$$

The ergodic theorem (Bremaud) 2001, 3-Cor.4.1) then applies because f is bounded and we have

$$\lim_{t \to \infty} \frac{1}{t-1} \sum_{i=1}^{t-1} f(X_i, X_{i+1}) = \sum_{(u,v) \in \mathcal{V} \times \mathcal{V}} \pi_{\mathbf{\Phi}}(u) \, p_{\mathbf{\Phi}}(u,v) f(u,v) = \frac{\mu(\mathcal{E})}{|\mathcal{E}|} \,.$$

Proof (Bias) Let the *i*-step transition probability of Φ be given by $p_{\Phi}^{i}(u,v)$. The bias at the *i*-th step is given by

$$\begin{aligned} \operatorname{BIAS}_{i} &= \left| \mathbb{E}\left[f(X_{i}, X_{i+1}) \right] - \sum_{(u,v) \in \mathcal{V} \times \mathcal{V}} \pi_{\mathbf{\Phi}}(u) \, p_{\mathbf{\Phi}}(u,v) f(u,v) \right| \\ &= \left| \sum_{(u,v) \in \mathcal{V} \times \mathcal{V}} p_{\mathbf{\Phi}}^{i}(X_{1}, u) \, p_{\mathbf{\Phi}}(u,v) f(u,v) - \sum_{(u,v) \in \mathcal{V} \times \mathcal{V}} \pi_{\mathbf{\Phi}}(u) \, p_{\mathbf{\Phi}}(u,v) f(u,v) \right| \\ &\leq B \left| \sum_{u \in \mathcal{V}} p_{\mathbf{\Phi}}^{i}(X_{1}, u) \sum_{v \in \mathcal{V}} p_{\mathbf{\Phi}}(u,v) - \sum_{u \in \mathcal{V}} \pi_{\mathbf{\Phi}}(u) \sum_{v \in \mathcal{V}} p_{\mathbf{\Phi}}(u,v) \right| \\ &\leq B \left| \sum_{u \in \mathcal{V}} p_{\mathbf{\Phi}}^{i}(X_{1}, u) - \sum_{u \in \mathcal{V}} \pi_{\mathbf{\Phi}}(u) \right| \leq B \sum_{u \in \mathcal{V}} \left| p_{\mathbf{\Phi}}^{i}(X_{1}, u) - \pi_{\mathbf{\Phi}}(u) \right|, \end{aligned}$$

where $f(\cdot) \leq B$, and the final inequality is due to Jensen's inequality. From Diaconis and Stroock [1991] Prop-3),

$$\operatorname{bias}_{i} \leq B \sqrt{\frac{1 - \pi_{\boldsymbol{\Phi}}(X_{1})}{\pi_{\boldsymbol{\Phi}}(X_{1})}} \beta_{*}^{i},$$

where $\beta_* = 1 - \delta(\Phi)$ is the SLEM of Φ . Because of Jensen's inequality and by summing a GP.

$$\left|\mathbb{E}[\hat{\mu}_0\left((X_i)_{i=1}^t\right)] - \frac{\mu(\mathcal{E})}{|\mathcal{E}|}\right| \leq \frac{1}{t-1}\sum_{i=1}^{t-1}\mathrm{bias}_i \leq \frac{B}{t-1}\sqrt{\frac{1-\pi_{\mathbf{\Phi}}(X_1)}{\pi_{\mathbf{\Phi}}(X_1)}}\frac{1-\beta_*^t}{1-\beta_*}\,.$$

Assuming that $\beta_*^t \approx 0$ and $t-1 \approx t$ when t is sufficiently large completes the proof.

Lemma 2 (Avrachenkov et al. (2016)) Let Φ be a finite state space, irreducible, timehomogeneous Markov chain, and let \xi denote the return time of RWT started from some $x_0 \in \mathcal{S}$ as defined in Definition 2. If Φ is reversible, then

$$\mathbb{E}\left[\xi^2\right] \le \frac{3}{\pi_{\mathbf{\Phi}}(x_0)^2 \delta(\mathbf{\Phi})},\tag{11}$$

where $\pi_{\Phi}(x_0)$ is the stationary distribution of x_0 , and $\delta(\Phi)$ is the spectral gap of Φ . When Φ is not reversible, the second moment of return times is given by Equation (12).

Proof Using (Aldous and Fill, 2002, Eq 2.21), we have

$$\mathbb{E}\left[\xi^{2}\right] = \frac{1 + 2\mathbb{E}_{\pi_{\Phi}}(T_{x_{0}})}{\pi_{\Phi}(x_{0})},\tag{12}$$

where $\mathbb{E}_{\pi_{\Phi}}(T_{x_0})$ is the expected hitting time of x_0 from the steady state. Combining (Aldous and Fill) [2002] Lemma 2.11 & Eq 3.41) and accounting for continuization yields

$$\mathbb{E}_{\pi_{\mathbf{\Phi}}}(T_{x_0}) \leq \frac{1}{\pi_{\mathbf{\Phi}}(x_0)\delta(\mathbf{\Phi})} \text{ and therefore, } \mathbb{E}\left[\xi^2\right] \leq \frac{1 + \frac{2}{\pi_{\mathbf{\Phi}}(x_0)\delta(\mathbf{\Phi})}}{\pi_{\mathbf{\Phi}}(x_0)} < \frac{3}{\pi_{\mathbf{\Phi}}(x_0)^2\delta(\mathbf{\Phi})},$$

because $\pi_{\mathbf{\Phi}}(x_0)$ and $\delta(\mathbf{\Phi})$ lie in the interval (0,1).

Proposition 6 Given a positive recurrent Markov chain Φ over state space ${\mathcal S}$ and a set of m RWTs \mathcal{T} and assuming an arbitrary ordering over \mathcal{T} , where $\mathbf{X}^{(i)}$ is the ith RWT in \mathcal{T} , $\mathbf{X}^{(i)}$ and $|\mathbf{X}^{(i)}|$ are i.i.d. processes such that $\mathbb{E}[|\mathbf{X}^{(i)}|] < \infty$, and when the tours are stitched together as defined next, the sample path is governed by Φ . For $t \geq 1$, define $\Phi_t = X_{t-R_{N_t}}^{N_t}$, where $R_i = \sum_{i'=1}^{i-1} |\mathbf{X}^i|$ when i > 1 and $R_1 = 0$ and $N_t = \max\{i : R_i < t\}$.

Proof R_i is a sequence of stopping times. Therefore, the strong Markov property (Bremaud, [2001], 2-Thm.7.1) states that sample paths before and after R_i are independent and are governed by Φ . Because Φ is positive recurrent and x_0 is visited i.o., the regenerative cycle theorem (Bremaud) 2001, 2-Thm.7.4) states that these trajectories are identically distributed and are equivalent to the tours \mathcal{T} sampled according to Definition 2 $\mathbb{E}[|\mathbf{X}^{(i)}|] < \infty$ due to positive recurrence.

B.2 Proof of Lemma 1

Proof (Unbiasedness and Consistency) Because \mathcal{G} is connected, Φ is positive recurrent with steady state $\pi_{\Phi}(u) \propto \mathbf{d}(u)$ due to Proposition 5. Consider the reward process $F^{(i)} = \mathbf{d}(u)$ $\sum_{j=1}^{|\mathbf{X}^{(i)}|} f(X_j^{(i)}, X_{j+1}^{(i)}), i \geq 1$. From Proposition 6, $F^{(i)}$ and $|\mathbf{X}^{(i)}|$ are i.i.d. sequences with finite first moments, because $F^{(i)} \leq B|\mathbf{X}^{(i)}|$. Let N_t and R_i be as defined in Proposition 6. Therefore, from the renewal reward theorem (Bremaud) [2001] 3-Thm.4.2), we have

$$\frac{\mathbb{E}[F^{(i)}]}{\mathbb{E}[|\mathbf{X}^{(i)}|]} = \lim_{t \to \infty} \frac{\sum_{i=1}^{N_t} F^{(i)}}{t} = \lim_{t \to \infty} \frac{\sum_{i=1}^{N_t} F^{(i)}}{R_{N_t}} \cdot \frac{R_{N_t}}{t} = \frac{\sum_{i=1}^{N_t} F^{(i)}}{R_{N_t}} \,,$$

where the final equality holds because $\lim_{t\to\infty}\frac{R_{N_t}}{t}=1-\lim_{t\to\infty}\frac{t-R_{N_t}}{t}$, and $\lim_{t\to\infty}\frac{t-R_{N_t}}{t}$ converges to 0 as $t\to\infty$ because $|\mathbf{X}^{(\cdot)}|<\infty$ w.p. 1 because $\boldsymbol{\varPhi}$ is positive recurrent. From Proposition 6 and the definition of $F^{(i)}$, $\sum_{i=1}^{N_t}F^{(i)}=\sum_{j=1}^{R_{N_t}}f(\varPhi_j,\varPhi_{j+1})$, and because f and $\pi_{\boldsymbol{\varPhi}}$ are bounded, we have from the ergodic theorem (Bremaud) 2001) 3-Cor.4.1),

$$\frac{\mathbb{E}[F^{(i)}]}{\mathbb{E}[|\mathbf{X}^{(i)}|]} = \lim_{t \to \infty} \frac{\sum_{j=1}^{R_{N_t}} f(\Phi_j, \Phi_{j+1})}{R_{N_t}} \stackrel{a.s}{=} \sum_{(u,v) \in \mathcal{V} \times \mathcal{V}} \pi_{\mathbf{\Phi}}(u) p_{\mathbf{\Phi}}(u,v) g(u,v) = \frac{2\mu(\mathcal{E})}{2|\mathcal{E}|} \ .$$

From Kac's formula Aldous and Fill 2002 Cor.2.24), $1/\mathbb{E}[|\mathbf{X}^{(i)}|] = \pi_{\mathbf{\Phi}}(x_0) = \frac{\mathbf{d}(x_0)}{2|\mathcal{E}|}$, and

$$\mathbb{E}\left[\frac{\mathbf{d}(x_0)}{2}F^{(i)}\right] \stackrel{a.s}{=} \mu(\mathcal{E}).$$

 $\hat{\mu}_*(\mathcal{T}; f, \mathcal{G})$ is unbiased by linearity of expectations on the summation over \mathcal{T} , and consistency is a consequence of Kolmogorov's SLLN (Bremaud, 2001, 1-Thm.8.3).

Proof (Running Time) From Kac's formula (Aldous and Fill) [2002] Cor.2.24), $\mathbb{E}[|\mathbf{X}^{(i)}|] = \frac{2|\mathcal{E}|}{\mathbf{d}(x_0)}$. From Proposition 6 tours can be sampled independently and thus parallelly. All cores will sample an equal number of tours in expectation, yielding the running time bound.

Proof (Variance) Because $f(\cdot) < B$, and tours are i.i.d., the variance is given by

$$\operatorname{Var}(\hat{\mu}_*(\mathcal{T})) = \operatorname{Var}\left(\frac{\mathbf{d}(x_0)}{2m} \sum_{\mathbf{X} \in \mathcal{T}} \sum_{j=1}^{|\mathbf{X}|} f(X_j, X_{j+1})\right) \le \frac{\mathbf{d}(x_0)^2 B^2}{4m} \operatorname{Var}(|\mathbf{X}|).$$

From Lemma 2 and Kac's formula (Aldous and Fill) 2002 Cor.2.24), Var ($|\mathbf{X}|$) is given by

$$\operatorname{Var}(|\mathbf{X}|) \leq \frac{3}{\pi_{\mathbf{\Phi}}(x_0)^2 \delta(\mathbf{\Phi})} - \frac{1}{\pi_{\mathbf{\Phi}}(x_0)^2} \leq \frac{3}{\pi_{\mathbf{\Phi}}(x_0)^2 \delta(\mathbf{\Phi})} = \frac{12|\mathcal{E}|^2}{\mathbf{d}(x_0)^2 \delta(\mathbf{\Phi})}.$$

C Proofs for Section 3

Assumption 2 For each \mathcal{G}_r , $1 < r \le R$ from Definition 5 assume $\mathbf{d}(\zeta_r)$ is known and that $p_{\Phi_r}(\zeta_r, \cdot)$ can be sampled from.

Proposition 7 (RWTs in Φ_r) Under Assumption 2 given access only to the original chain Φ and stratifying function ρ , let Φ_r be the random walk in the graph stratum \mathcal{G}_r from Definition 5. To sample an RWT $(X_i)_{i=1}^{\xi}$ over Φ_r from the supernode ζ_r , we set $X_1 = \zeta_r$, sample $X_2 \sim p_{\Phi_r}(\zeta_r, \cdot)$, and then, until $\rho(X_{\xi+1}) < r$, we sample

$$X_{i+1} \sim \text{Unif}\left(\mathbf{N}_{G_r}\left(X_i\right)\right) \equiv \begin{cases} \text{Unif}(\mathbf{N}_{\mathcal{G}}(X_i)) & \text{if } \rho(X_i) = r \\ \text{Unif}(\mathbf{N}_{\mathcal{G}}(X_i) \cap \mathcal{I}_r) & \text{if } \rho(X_i) > r \end{cases}.$$

Proof The proof is a direct consequence of Definition 5 and Definition 1

Proposition 8 (Perfectly Stratified Estimate) Under Assumption 2 given the EPS (Definition 6) stratum \mathcal{G}_r (Definition 5), bounded $f: \mathcal{E} \to \mathbb{R}$ and a set of m RWTs \mathcal{T}_r over Φ_r from ζ_r from Proposition 7 the per stratum estimate is given by

$$\hat{\mu}(\mathcal{T}_r; f, \mathcal{G}_r) = \frac{\mathbf{d}(\zeta_r)}{2m} \sum_{\mathbf{X} \in \mathcal{T}_r} \sum_{j=2}^{|\mathbf{X}|-1} f(X_j, X_{j+1}), \qquad (13)$$

where X_j is the jth state visited in the RWT $\mathbf{X} \in \mathcal{T}_r$. For all r > 1, $\hat{\mu}(\mathcal{T}_r; f, \mathcal{G}_r)$ is an unbiased and consistent estimator of $\mu(\mathcal{J}_r) = \sum_{(u,v) \in \mathcal{J}_r} f(u,v)$, where \mathcal{J}_r is the r-th edge stratum defined in Definition [4].

Proof Define $f' \colon \mathcal{E}_r \to \mathbb{R}$ as $f'(u,v) \triangleq \mathbf{1}\{u,v \neq \zeta_r\}f(u,v)$. By Definition 2 in each RWT $\mathbf{X} \in \mathcal{T}_r$, $f'(X_1,X_2) = f'(X_{|\mathbf{X}|},X_{|\mathbf{X}|+1}) = 0$, and therefore, $\hat{\mu}(\mathcal{T}_r;f,\mathcal{G}_r) = \hat{\mu}_*(\mathcal{T};f',\mathcal{G}_r)$, where $\hat{\mu}_*$ is the RWT Estimate from Lemma 1 Moreover, because \mathcal{G}_r is connected,

$$\mathbb{E}\left[\hat{\mu}(\mathcal{T}_r;f,\mathcal{G}_r)\right] = \mathbb{E}\left[\hat{\mu}_*(\mathcal{T};f',\mathcal{G}_r)\right] = \sum_{(u,v)\in\mathcal{E}_r} f'(u,v) = \sum_{(u,v)\in\mathcal{J}_r} f(u,v)\,,$$

where the final equality holds because \mathcal{E}_r is the union of \mathcal{J}_r and edges incident on the supernode. Consistency is also due to Lemma $\boxed{1}$

C.1 Proof of Proposition 1

Proof Proposition 1 (a) is necessary because when Proposition 1 (a) does not hold, there exists a component such that the minimum value of ρ in that component is $\ddot{r} > 0$ such that in $\mathcal{G}_{\ddot{r}}$ (Definition 5), and the supernode $\zeta_{\ddot{r}}$ will be disconnected from all vertices. If Proposition 1 (b) is violated, a vertex \ddot{u} exists that is disconnected in $\mathcal{G}_{\rho(\ddot{u})}$, and if Proposition 1 (c) is violated, the supernode is disconnected. Finally, it is easily seen that these conditions sufficiently guarantee that each stratum is connected, and the stratification is an EPS.

C.2 Proof of Theorem 1

We begin by defining the multi-set containing the end points of edges between vertex strata.

Definition 12 Given \mathcal{G} stratified into R strata, $\forall 1 \leq q < t \leq R$ define border multi-sets as

$$\mathcal{B}_{q,t} \triangleq \{v \, \forall (u,v) \in \mathcal{E} : u \in \mathcal{I}_q \text{ and } v \in \mathcal{I}_t\}.$$

The degree of the supernode in \mathcal{G}_r (Definition 5) is then given by $\mathbf{d}(\zeta_r) = \sum_{q=1}^{r-1} |\mathcal{B}_{q,r}|$, and transitions out of ζ_r can be sampled by sampling $q \in \{1, \ldots, r-1\}$ w.p. $\propto |\mathcal{B}_{q,r}|$ and then by uniformly sampling from $\mathcal{B}_{q,r}$.

Proposition 9 Given the setting in Definitions 8 and 9 for all $1 \le r < t \le R$,

$$\lim_{|\mathcal{T}_{2}^{\dagger}| \to \infty} \dots \lim_{|\mathcal{T}_{r}^{\dagger}| \to \infty} \widehat{\beta}_{r,t} \stackrel{a.s}{=} |\mathcal{B}_{r,t}|, \qquad (14)$$

$$\lim_{|\mathcal{T}_{2}^{\dagger}| \to \infty} \dots \lim_{|\mathcal{T}_{r}^{\dagger}| \to \infty} \widehat{\mathbf{U}}_{r,t} \sim \text{UNIF}(\mathcal{B}_{r,t}),$$
(15)

$$\lim_{|\mathcal{T}_{2}^{\dagger}| \to \infty} \dots \lim_{|\mathcal{T}_{r}^{\dagger}| \to \infty} p_{\widehat{\boldsymbol{\Phi}}_{r}}(\mathbf{X}) = p_{\boldsymbol{\Phi}_{r}}(\mathbf{X}), \quad \forall \mathbf{X} \in \mathcal{T}_{r}^{\dagger}.$$
(16)

i.e., each tour in \mathcal{T}_r^{\dagger} is perfectly sampled from Φ_r .

Proof (By Strong Induction) The base case for r = 1 holds by the base case in Definition 8. Now assume that Proposition 9 holds for all strata up to and including r - 1. Because of the inductive claim and by Definition 12.

$$\lim_{|\mathcal{T}_{2}^{\dagger}| \to \infty} \dots \lim_{|\mathcal{T}_{r-1}^{\dagger}| \to \infty} \widehat{\mathbf{d}}(\zeta_{r}) = \sum_{q=1}^{r-1} \widehat{\beta}_{q,r} \stackrel{a.s}{=} \sum_{q=1}^{r-1} |\mathcal{B}_{q,r}| = \mathbf{d}(\zeta_{r}),$$
and similarly,
$$\lim_{|\mathcal{T}_{2}^{\dagger}| \to \infty} \dots \lim_{|\mathcal{T}_{r-1}^{\dagger}| \to \infty} \widehat{p}_{\boldsymbol{\Phi}_{r}}(\zeta_{r}, \cdot) \equiv p_{\boldsymbol{\Phi}_{r}}(\zeta_{r}, \cdot)$$

because the inductive claim makes the procedure of sampling transitions out of ζ_r in Definition 8 equivalent to Definition 12 Equation (16) holds because transition probabilities at all states other than ζ_r are equivalent in Φ_r and $\widehat{\Phi}_r$ according to Definition 7 Now recall that

$$\widehat{\beta}_{r,t} = \frac{\widehat{\mathbf{d}}(\zeta_r)}{|\mathcal{T}_r^{\dagger}|} \sum_{\mathbf{X} \in \mathcal{T}_r^{\dagger}} \sum_{j=2}^{|\mathbf{X}|} \mathbf{1} \{ \rho(X_j) = t \}.$$

Because $\widehat{\mathbf{d}}(\zeta_r) = \mathbf{d}(\zeta_r)$ and the tours are sampled perfectly,

$$\lim_{|\mathcal{T}_2^{\dagger}| \to \infty} \dots \lim_{|\mathcal{T}_{r-1}^{\dagger}| \to \infty} \widehat{\beta}_{r,t} = \widehat{\mu}_* \left(\mathcal{T}_r^{\dagger}; f' \right) ,$$

where $f'(u,v) = \mathbf{1}\{\rho(v) = t\}$ and $\hat{\mu}_*$ is from Lemma $\boxed{1}$ from which we also use the consistency guarantee to show that under an *EPS*, Equation $\boxed{14}$ holds as

$$\lim_{|\mathcal{T}_2^{\dagger}| \to \infty} \dots \lim_{|\mathcal{T}_r^{\dagger}| \to \infty} \widehat{\beta}_{r,t} \stackrel{a.s}{=} \sum_{(u,v) \in \mathcal{E}_r} f'(u,v) = |\mathcal{B}_{r,t}|.$$

Because of Proposition 6 concatenating tours $\mathbf{X} \in \mathcal{T}_q^{\dagger}$ yields a sample path from Φ_r , and these samples are distributed according to π_{Φ_r} as $|\mathcal{T}_{r'}^{\dagger}| \to \infty$, $r' \le r$. Therefore,

$$\lim_{|\mathcal{T}_2^\dagger| \to \infty} \dots \lim_{|\mathcal{T}_r^\dagger| \to \infty} \uplus_{\mathbf{X} \in \mathcal{T}_q^\dagger} \uplus_{j=2}^{|\mathbf{X}|} \left\{ X_j \colon \rho(X_j) = t \right\} \sim \pi_{\mathbf{\Phi}_r}' \,,$$

where $\pi'_{\Phi_r}(u) \propto \mathbf{1}\{\rho(u) = t\}\mathbf{d}_{\mathcal{G}_r}(u)$, which is equivalent to $\text{UNIF}(\mathcal{B}_{r,t})$ by Definitions and 12 thus proving Equation (15).

Proof (Main Theorem) Combining Proposition 9 and Proposition 8 proves Theorem 1

C.3 Proof of Theorem 2

Definition 13 (L^2 **Distance between** $\widehat{\pi}$ and π (Aldous and Fill, 2002)) The L^2 distance between discrete probability distribution $\widehat{\pi}$ and reference distribution π with sample space Ω is given by $\|\widehat{\pi} - \pi\|_2 = \sum_{i \in \Omega} \frac{(\widehat{\pi}(i) - \pi(i))^2}{\pi(i)}$.

Definition 14 (**Distorted chain**) Given a Markov chain Φ over finite state space S and an arbitrary $x_0 \in S$, let $\widehat{\Phi}$ be the distorted chain such that $\forall u \neq x_0, \ p_{\widehat{\Phi}}(u, \cdot) = p_{\Phi}(u, \cdot)$, and $p_{\widehat{\Phi}}(x_0, \cdot)$ is an arbitrary distribution with support $\sup(p_{\widehat{\Phi}}(x_0, \cdot)) \subseteq \sup(p_{\Phi}(x_0, \cdot))$. The distortion is given by $\|p_{\widehat{\Phi}}(x_0, \cdot) - p_{\Phi}(x_0, \cdot)\|$ as defined in Definition 13.

Lemma 3 Given a finite state, positive recurrent Markov chain Φ over state space S, let $\widehat{\Phi}$ be the chain distorted at some $x_0 \in S$ from Definition 14. Let

$$\mathcal{X} = \{(X_1, \dots, X_{\xi}) \colon X_1 = x_0, \, \xi = \min\{t > 0 \colon X_{t+1} = x_0\}, \, p_{\Phi}(X_1, \dots, X_{\xi}) > 0\},\,$$

denote the set of all possible arbitrary lengths RWTs that begin and end at x_0 from Definition 2. Given a tour $\mathbf{Y} \in \mathcal{X}$ sampled from Φ and a bounded function $F \colon \mathcal{X} \to \mathbb{R}$,

$$\mathbb{E}_{\boldsymbol{\Phi}} \left[\frac{p_{\widehat{\boldsymbol{\Phi}}}(Y_1, Y_2)}{p_{\boldsymbol{\Phi}}(Y_1, Y_2)} F(\mathbf{Y}) \right] = \mathbb{E}_{\widehat{\boldsymbol{\Phi}}} \left[F(\mathbf{Y}) \right], \tag{17}$$

where \mathbb{E}_{Φ} and $\mathbb{E}_{\widehat{\Phi}}$ are expectations under the distribution of tours sampled from Φ and $\widehat{\Phi}$.

Proof All tours in \mathcal{X} are of finite length because of the positive recurrence of $\boldsymbol{\Phi}$. The ratio of the probability of sampling the tour $\mathbf{Y} = (Y_1, \dots, Y_{\mathcal{E}'})$ from the chain $\widehat{\boldsymbol{\Phi}}$ to $\boldsymbol{\Phi}$ is given by

$$\frac{p_{\widehat{\boldsymbol{\Phi}}}(\mathbf{Y})}{p_{\mathbf{\Phi}}(\mathbf{Y})} = \frac{\prod_{j=1}^{\xi'} p_{\widehat{\boldsymbol{\Phi}}}(Y_j, Y_{j+1})}{\prod_{j=1}^{\xi'} p_{\mathbf{\Phi}}(Y_j, Y_{j+1})} = \frac{p_{\widehat{\boldsymbol{\Phi}}}(Y_1, Y_2)}{p_{\mathbf{\Phi}}(Y_1, Y_2)},\tag{18}$$

because $p_{\Phi}(Y_j,\cdot) = p_{\widehat{\Phi}}(Y_j,\cdot)$, $\forall 1 < j \leq \xi'$ because $Y_j \neq x_0$ by the definitions of \mathcal{X} and $\widehat{\Phi}$. Because $\sup(p_{\widehat{\Phi}}(x_0,\cdot)) \subseteq \sup(p_{\Phi}(x_0,\cdot))$, $\sup(p_{\widehat{\Phi}}(\mathbf{Y})) \subseteq \sup(p_{\Phi}(\mathbf{Y}))$. The theorem statement therefore directly draws from the definition of importance sampling (Robert and Casella, 2013) Def 3.9) with the importance weights derived in Equation (18).

П

Lemma 4 Given a simple random walk Φ on the connected non-bipartite graph $\mathcal G$ from Definition 1 let $\widehat{\Phi}$ be the chain distorted at some $x_0 \in \mathcal S$ from with distortion ν Definition 14 Let $\lambda = \widehat{\mathbf d}(x_0)/\mathbf d(x_0)$. Let $f \colon \mathcal E \to \mathbb R$ bounded by B, and $F(\mathbf X) = \sum_{j=1}^{|\mathbf X|} f(X_j, X_{j+1})$, where $\mathbf X$ is an RWT as defined in Appendix B.2. The bias of an RWT Estimate (Equation 2) computed using tours sampled over $\widehat{\Phi}$ and using $\widehat{\mathbf d}(x_0)$ as the degree is given by

$$\text{bias} = \left| \mathbb{E}_{\widehat{\boldsymbol{\Phi}}} \left[\frac{\widehat{\mathbf{d}}(x_0)}{2} F(\mathbf{X}) \right] - \mu(\mathcal{E}) \right| \leq (\lambda \nu + |1 - \lambda|) \frac{\sqrt{3}B|\mathcal{E}|}{\sqrt{\delta}},$$

where δ is the spectral gap of Φ , and B is the upper bound of f.

Proof From Lemma 3 and Lemma 1 we have, respectively,

$$\mathbb{E}_{\widehat{\boldsymbol{\Phi}}}\left[\frac{\widehat{\mathbf{d}}(x_0)}{2}F(\mathbf{X})\right] = \mathbb{E}_{\boldsymbol{\Phi}}\left[\frac{\widehat{\mathbf{d}}(x_0)}{2}\frac{p_{\widehat{\boldsymbol{\Phi}}}(X_1, X_2)}{p_{\boldsymbol{\Phi}}(X_1, X_2)}F(\mathbf{X})\right],$$
$$\mu(\mathcal{E}) = \mathbb{E}_{\boldsymbol{\Phi}}\left[\frac{\mathbf{d}(x_0)}{2}F(\mathbf{X})\right].$$

Subtracting the two, squaring both sides and using the Cauchy-Schwarz inequality decomposes the squared bias into

$$\begin{aligned} \text{BIAS} &= \left| \mathbb{E}_{\mathbf{\Phi}} \left[\left(\frac{\widehat{\mathbf{d}}(x_0)}{\mathbf{d}(x_0)} \frac{p_{\widehat{\mathbf{\Phi}}}(X_1, X_2)}{p_{\mathbf{\Phi}}(X_1, X_2)} - 1 \right) \frac{\mathbf{d}(x_0)}{2} F(\mathbf{X}) \right] \right|. \\ \text{BIAS}^2 &\leq \underbrace{\mathbb{E} \left[\left(\frac{\widehat{\mathbf{d}}(x_0)}{\mathbf{d}(x_0)} \frac{p_{\widehat{\mathbf{\Phi}}}(x_0, X_2)}{p_{\widehat{\mathbf{\Phi}}}(x_0, X_2)} - 1 \right)^2 \right]}_{\text{BIAS}_{\text{dist}}} \underbrace{\mathbb{E} \left[\left(\frac{\mathbf{d}(x_0)}{2} F(\mathbf{X}) \right)^2 \right]}_{\text{BIAS}_{\text{spectral}}}, \end{aligned}$$

where the expectation is under Φ . Using definitions from the theorem statement,

$$\begin{aligned} \text{BIAS}_{\text{dist}} &= \frac{\widehat{\mathbf{d}}(x_0)^2}{\mathbf{d}(x_0)^2} \mathbb{E}\left[\left(\frac{p_{\widehat{\Phi}}(x_0, X_2)}{p_{\Phi}(x_0, X_2)} \right)^2 \right] + 1 - 2 \frac{\widetilde{\mathbf{d}}(x_0)}{\mathbf{d}(x_0)} \mathbb{E}\left[\frac{p_{\widehat{\Phi}}(x_0, X_2)}{p_{\Phi}(x_0, X_2)} \right] \\ &= \lambda^2 (1 + \nu^2) + 1 - 2\lambda = \lambda^2 + \lambda^2 \nu^2 + 1 - 2\lambda \\ &= \lambda^2 \nu^2 + (1 - \lambda)^2 \le (\lambda \nu + |1 - \lambda|)^2 \ . \end{aligned}$$

Because $F(\mathbf{X}) \leq B\xi$, the tour length, from Lemma 2, we see that

$$\mathrm{bias}_{\mathrm{spectral}} \leq \frac{\mathbf{d}(x_0)^2 B^2}{4} \frac{3}{\pi_{\mathbf{\Phi}}(x_0)^2 \delta} = \frac{3B^2 |\mathcal{E}|^2}{\delta} \,,$$

and combining both biases completes the proof for BIAS.

Proof (Main Theorem) Note that by linearity of expectations

$$\mathbb{E}\left[\hat{\mu}\left(\mathcal{T}_{2:r}^{\dagger};f\right)|\mathcal{T}_{2:r-1}^{\dagger}\right] = \mathbb{E}\left[\frac{\widehat{\mathbf{d}}(\zeta_r)}{2|\mathcal{T}_r^{\dagger}|} \sum_{\mathbf{X}\in\mathcal{T}_r^{\dagger}} \sum_{j=2}^{|\mathbf{X}|-1} f(X_j, X_{j+1})\right],$$

$$= \mathbb{E}_{\mathbf{X}\sim\widehat{\boldsymbol{\Phi}}_r}\left[\frac{\widehat{\mathbf{d}}(\zeta_r)}{2} \sum_{j=1}^{|\mathbf{X}|} f'(X_j, X_{j+1})\right],$$

where **X** is an RWT on $\widehat{\boldsymbol{\varPhi}}_r$ that depends on $\mathcal{T}_{2:r-1}^\dagger$ and $f'(u,v) \triangleq \mathbf{1}\{u,v \neq \zeta_r\}f(u,v)$. Applying Lemma 4 completes the proof because $\widehat{\boldsymbol{\varPhi}}_r$ is a distorted chain by Definition 14

D Proofs for Section 4

D.1 Proof of Proposition 2

Proof From Wang et al. (2014) Thm-3.1), we know that each disconnected component of G leads to a disconnected component in $\mathcal{G}^{(k-1)}$, and if \mathcal{I}_1 contains a subgraph in each connected component, Proposition 1 (a) is satisfied. We now prove that $\forall s \in \mathcal{V}^{(k-1)}$, if $\rho(s) = r > 1$, $\exists s' \in \mathbf{N}(s) \colon \rho(s') < r$ which simultaneously satisfies Proposition 1 (b) and Proposition 1 (c)

W.l.o.g. let the vertex with the smallest distance from the seed vertices be denoted by $\hat{u} = \operatorname{argmin}_{u \in V(s)} \operatorname{dist}(u)$. When $\operatorname{dist}(\hat{u}) > 0$, there exists $v \in \mathbf{N}_G(\hat{u})$ such that $\operatorname{dist}(v) < \operatorname{dist}(\hat{u})$ by the definition of dist. More concretely, v would be the penultimate vertex in the shortest path from the seed vertices to \hat{u} . Let $v' \neq \hat{u}$ be a nonarticulating vertex of s, which is possible because any connected graph has at least 2 nonarticulating vertices. Let $s_1 = G(V(s) \setminus \{v'\} \cup \{v\}) \in \mathcal{V}^{(k-1)}$. Now, $\rho(s_1) < \rho(s)$ because v' has been replaced with a vertex at necessarily a smaller distance and because the indicator in the definition of v ho will always be 0 in this case. Moreover, $s_1 \circ s = G(V(s) \cup \{v\}) \in \mathcal{V}^{(k)}$, and hence an edge exists between the two.

When $\operatorname{dist}(\hat{u}) = 0$, there exists $v \in \mathbf{N}_G(\hat{u})$ such that $\operatorname{dist}(v) = 0$. There exists a nonarticulating $v' \in V(s) \backslash V^*$ because otherwise V^* would have been disconnected. Observing that $\operatorname{dist}(v') + \mathbf{1}\{v' \in V(\mathcal{I}_1) \backslash V^*\} > 0$ completes the proof of ergodicity. \square

D.2 Proof of Proposition 3

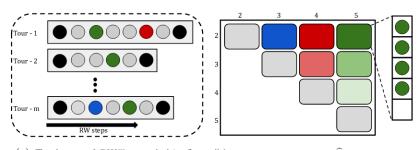
Proof (Sampling Probability) Consider the lines Lines $\boxed{3}$ to $\boxed{5}$. The probability of sampling the pair (u, v) from $V(s) \times \mathbf{N}_G(V(s))$ is given by

$$\begin{split} P(u,v) &= \sum_{a \in V(s) \backslash \{u\}} P(v|a,u) P(a|u) P(u) \\ &= \sum_{a \in V(s) \backslash \{u\}} \frac{\mathbf{1}\{v \in \mathbf{N}(a)\}}{\mathbf{d}(a)} \frac{\mathbf{d}(a)}{\deg_s - \mathbf{d}(u)} \frac{\deg_s - \mathbf{d}(u)}{(k-1-1)\deg_s} \\ &\propto \sum_{a \in V(s) \backslash \{u\}} \mathbf{1}\{v \in \mathbf{N}(a)\} = |N(v) \cap V(s) \backslash \{u\}| = \text{bias}\,, \end{split}$$

where bias is defined in Line 6 and corrected for in Line 7. After the rejection, therefore, $(u,v) \sim \text{UNIF}(V(s) \times \mathbf{N}_G(V(s)))$.

Line $\boxed{9}$ constitutes an importance sampling with unit weight for pairs (u,v), where removing u from and adding v to V(s) produces a k-1-CIS and zero otherwise. In Line $\boxed{9}$ because removing a nonarticulating vertex and adding another vertex to s cannot lead to a disconnected subgraph, we can avoid a DFS when $u \notin \mathcal{A}_s$. This completes the proof.

Proof (Time Complexity) Assuming access to a precomputed vector of degrees, the part up to Line $1 ext{is } O(k-1^2)$. In each proposal, Lines $2 ext{and } 4 ext{are } O(k-1)$, and Line $5 ext{is } O(\Delta_s)$. Line $6 ext{is } O(k-1)$, and the expected complexity of Line $9 ext{is } O(k-1^2 |A_s|/k-1)$ because in expectation only $|A_s|/k-1$ graph traversals will be required. The acceptance probability is $\geq 1/k-1$ is Line $7 ext{and } \geq \frac{k-1-|A_s|}{k-1}$. The expected number of proposals is therefore $\leq \frac{k-1^2}{k-1-|A_s|}$. As such, the expected time complexity is $O(k-1^2(1+\frac{\Delta_s+k-1|A_s|}{k-1-|A_s|}))$.



(a) \mathcal{T}_2 , the set of RWTs sampled in \mathcal{G}_2 . (b) The reservoir matrix, $\widehat{\mathbf{U}}_{r,t}$ for $2 \leq r < t \leq R$.

Fig. 4: Parallel RWTs and Reservoirs: Figure 4a shows the set of m RWTs sampled on \mathcal{G}_2 in parallel, where the supernode ζ_2 is colored black. The gray, blue, red and green colors represent states in stratum 2–5, respectively. Figure 4b shows the upper triangular reservoir matrix in which the cell in the r-th row and t-th column contains samples from $\widehat{\mathbf{U}}_{r,t}$.

E Additional Implementation Details

E.1 Parallel Sampling with a Reservoir Matrix.

Given a reasonably large M and the number of strata R, we initialize an upper triangular matrix of empty reservoirs $[\widehat{\mathbf{U}}_{r,t}]_{2 \leq r < t \leq R}$ and a matrix of atomic counters $[\widehat{\mathbf{M}}_{q,r}]_{2 \leq r < t \leq R}$ initialized to 0. In each stratum r, while being sampled in parallel whenever a tour enters the t-th stratum, $\widehat{\mathbf{M}}_{r,t}$ is incremented, and with a probability $\min(1, {}^{\mathbb{M}}/\widehat{\mathbf{M}}_{r,t})$, the state is inserted into a random position in the reservoir $\widehat{\mathbf{U}}_{r,t}$ and rejected otherwise. The only contention between threads in this scheme is at the atomic counter and in the rare case where two threads choose the same location to overwrite, wherein ties are broken based on the value of the atomic counter at the insertion time, guaranteeing thread safety. The space complexity of a reservoir matrix is therefore $O(R^2 \mathrm{M})$.

A toy example of this matrix is presented in Figure 4 where R=5, and the RWTs are being sampled on the graph stratum \mathcal{G}_2 . Whenever (non-gray) states in $\mathcal{I}_{3:5}$ are visited, they are inserted into the corresponding reservoirs— $\widehat{\mathbf{U}}_{2,5}$ is depicted in detail.

E.2 PSRW Neighborhood

The neighborhood of a k-CIS s in $\mathcal{G}^{(k)}$ is the set of all vertices $u, v \in V$ such that replacing u with v in s yields a k-CIS. Formally,

$$\mathbf{N}_{\mathcal{G}^{(k)}}(s) \equiv \left\{ (u, v) \in V(s) \times \mathbf{N}_G(V(s)) : G(V(s) \cup \{v\} \setminus \{u\}) \in \mathcal{V}^{(k)} \right\}, \tag{19}$$

where $\mathbf{N}_G(V(s)) = \cup_{x \in V(s)} \mathbf{N}_G(x)$ is the union of the neighborhood of each vertex in s. The size of the neighborhood is then $\mathrm{O}(k\,\mathbf{N}_G(V(s))) \in \mathrm{O}(k^2\Delta_G)$ because $\mathbf{N}_G(V(s)) \in \mathrm{O}(k\Delta_G)$, where Δ_G is the maximum degree in G. Each potential neighbor further requires a connectivity check in the form of a BFS or DFS, which implies that the naive neighborhood sampling algorithm requires $\mathrm{O}(k^4\Delta_G)$ time.

E.2.1 Articulation Points

Apart from the rejection sampling algorithm from Algorithm 1 we use articulation points to efficiently compute the subgraph bias γ from Equation 9. Specifically, given the k-1-CIS, s, $\gamma(s) = {\kappa - \mathcal{A}_s \choose 2}$, \mathcal{A}_s is the set of articulation points of s. This draws directly from

(Wang et al.) 2014) Sec-3.3) and the definition of articulation points. Hopcroft and Tarjan (1973) showed that for any simple graph s the set of articulation points can be computed in O(|V(s)| + |E(s)|) time.

```
Algorithm 2: Ripple for Subgraph Counting
      Input: Input graph G, Order k, Set of subgraph patterns \mathcal{H} of interest
      Input: Initial vertex stratum \mathcal{I}_1, Reservoir Size M and Error Bound \epsilon
      Output: \hat{\mu}, an asymptotically unbiased estimate of \mathcal{C}^{(k)}
      /* Initialization
 1 \hat{\mu} = 0, \hat{\beta}_{q,t} = 0, \hat{\mathbf{U}}_{q,t} = \emptyset, \forall 1 \leq q < t \leq R;
2 Run BFS for stratification \rho \colon \mathcal{V}^{(k-1)} \to \{1, \dots, R\}, with \mathcal{I}_1 (Proposition 2)
      /* Exact computation in the first stratum
 3 for each u \in \mathcal{I}_1, v \in \mathbf{N}_{\mathcal{G}^{(k-1)}}(u) do
4 | Update \widehat{\beta}_{1,\rho(v)} += 1, \widehat{\mathbf{U}}_{1,\rho(v)} \cup = v
5 | Update \widehat{\mu} += \left(\frac{\mathbf{1}\{u \circ v \sim H\}}{\gamma(u \circ v)}\right)_{H \in \mathcal{H}};
                                                                                                                           // Equation
      /* Estimate remaining strata
  6 for r \in 2, ..., R do
             Initialize \hat{\mu}_r = 0, m_r = 0
             parallel while Equation (10) is not satisfied do
  8
  9
                    Sample q from \{1, \ldots, \overline{r-1}\} w.p. \widehat{\beta}_{q,r}
                    Sample u from \widehat{\mathbf{U}}_{q,r};
                                                                                                                          // Equation (6)
10
                    Sample v \sim \text{UNIF}(\mathbf{N}_{\mathcal{G}^{(k-1)}}(u));
                                                                                                                            // Algorithm 1
11
                     while \rho(v) \ge r do
12
                            Update \hat{\mu}_r += \left(\frac{\mathbf{1}\{u \circ v \sim H\}}{\gamma(u \circ v)}\right)_{H \in \mathcal{H}};
                                                                                                                          // Equation (8)
13
                            if \rho(v) > r then
 14
                                  Update \widehat{\beta}_{r,\rho(v)} += 1;
                                                                                                                          // Equation (4)
 15
                                  Update \widehat{\mathbf{U}}_{r,\rho(v)} \cup = v;
                                                                                                                          // Equation (6)
 16
 17
                            /* Proposition 7 and Algorithm 1
                            if \rho(u) = r then
18
 19
                                 Sample v \sim \text{UNIF}(\mathbf{N}_{\mathcal{G}^{(k-1)}}(u))
20
                            else
 21
                                   while \rho(v) \neq r do
                                    Sample v \sim \text{UNIF}(\mathbf{N}_{\mathcal{G}^{(k-1)}}(u))
 22
23
             Compute \widehat{\deg}_r = \sum_{q=1}^{r-1} \widehat{\beta}_{q,r};
24
                                                                                                                          // Equation (5)
            \begin{split} \hat{\mu} +&= \frac{\widehat{\deg}_r}{2\,m_r} \hat{\mu}_r \ ; \\ \text{Update } \widehat{\beta}_{r,t} *&= \frac{\widehat{\deg}_r}{m_r}, \, \forall t > r \ ; \end{split}
                                                                                                       // Equations (7) and (8)
// Equation (4)
26
27 return \hat{\mu}
```

E.3 Proof of Proposition 4

Proposition 10 (Extended Version of Proposition $\boxed{4}$) We assume a constant number of tours m in each stratum and ignore graph loading. The Ripple estimator of k-CIS counts

described in Algorithm 2 has space complexity in

$$O(k^3 D_G^2 M + |\mathcal{H}|) \equiv \widehat{O}(k^3 + |\mathcal{H}|),$$

where \widehat{O} ignores all factors other than k and $|\mathcal{H}|$, M is the size of the reservoir from Sec-

tion 4.3 D_G is the diameter of G, and $|\mathcal{H}|$ is the number of patterns of interest. The total number of random walk steps is given by $O(k^3mD_G\Delta_GC_{REJ})$, where C_{REJ} is the number of rejections in Line 21 of Algorithm 2 Δ_G is the largest degree in G, and the total time complexity is $\widehat{O}(k^7 + |\mathcal{H}|)$.

Remark 2 In practice, we adapt the proposals in Algorithm 1 to minimize C_{REJ} using heuristics over the values of DIST (\cdot) from Proposition 2.

Lemma 5 Given a graph stratum \mathcal{G}_r from Definition $\overline{\mathcal{G}}_r$ for some r>1, define $\alpha_r=1$ $|\{u \in \mathcal{I}_r : \mathbf{N}(u) \cap \mathcal{I}_{1:r-1} \neq \emptyset\}|/|\mathcal{I}_r|$ as the fraction of vertices in the r-th vertex stratum that share an edge with a previous stratum. The return time ξ_r of the chain $\pmb{\Phi}_r$ to the supernode $\zeta_r \in \mathcal{V}_r$ follows $\mathbb{E}_{\Phi_r}[\xi_r] \leq \frac{2\bar{\mathbf{d}}_r}{\alpha_r}$, where $\bar{\mathbf{d}}_r$ is the average degree in \mathcal{G} of all vertices in \mathcal{I}_r .

Proof Because $\alpha_r \mathcal{I}_r$ vertices have at least one edge incident on ζ_r , $\mathbf{d}_{\mathcal{G}_r}(\zeta_r) \geq \alpha_r \mathcal{I}_r$. From Definition 5 because all edges not incident on \mathcal{I}_r are removed from \mathcal{G}_r , $\operatorname{Vol}(\mathcal{G}_r) \leq$ $2\sum_{u\in\mathcal{I}_r}\mathbf{d}_{\mathcal{G}}(u)$. Therefore, from Lemma 1.

$$\mathbb{E}_{\boldsymbol{\varPhi}_r}[\boldsymbol{\xi}_r] = \frac{\operatorname{Vol}(\mathcal{G}_r)}{\mathbf{d}(\boldsymbol{\zeta}_r)} \leq \frac{2\sum_{u \in \mathcal{I}_r} \mathbf{d}_{\mathcal{G}}(u)}{\alpha_r \mathcal{I}_r} = \frac{2\bar{\mathbf{d}}_r}{\alpha_r} \;.$$

Proposition 11 The Ergodicity-Preserving Stratification from Proposition 2 is such that $\alpha_r = 1$ for all r > 1 as defined in Lemma 5, and consequently, the diameter of each graph stratum is ≤ 4 . The total number of strata $R \in O(k D_G)$, where D_G is the diameter of G.

Proof We show in Appendix D.1 that for each vertex $s \in \mathcal{V}^{(k-1)}$, if $\rho(s) = r > 1$, there exists $s' \in \mathbf{N}(s)$ such that $\rho(s') < r$. This implies that $\alpha_r = 1$. In \mathcal{G}_r , therefore, from ζ_r , all vertices in \mathcal{I}_r are at unit distance from ζ_r , and vertices in $\mathbf{N}(\mathcal{I}_r) \setminus \mathcal{I}_r$ are at a distance of 2 from ζ_r . Because no other vertices are present in \mathcal{G}_r , this completes the proof of the first part. Trivially, $R \leq (k-1) \cdot \max_{u \in V} \operatorname{dist}(u) \in \operatorname{O}(k \cdot D_G)$.

Proof (Memory Complexity) From Algorithm 2 we compute a single count estimate per stratum and maintain reservoirs and inter-partition edge count estimates for each $2 \le q < 1$ $t \leq R$. Because a reservoir $\hat{\mathbf{U}}_{q,t}$ needs $O(k_{\rm M})$ space (Appendix E.1), the total memory requirement is $O(R^2 k_{\rm M})$, where R is the number of strata. From Proposition 11 plugging $R \in O(kD_G)$, and because storing the output $\hat{\mu}$ requires $O(|\mathcal{H}|)$ memory the proof is completed.

Proof (Time Complexity) The stratification requires a single BFS $\in O(|V| + |E|)$ from Section 4.2. In Line 3, the estimation phase starts by iterating over the entire higher-order neighborhood of each subgraphs in \mathcal{I}_1 . Based on Appendix E.2.1, Line 5 is in $O(k^2)$. Because the size of the higher-order neighborhood of each subgraph is $O(k^2\Delta_G)$ from Appendix E.2. the initial estimation phase will require $O(|\mathcal{I}_1| k^4 \Delta_G)$ time.

In all other strata r = 2, ..., R, we assume that m tours are sampled in Line 8. Starting each tour (Lines $\boxed{9}$ to $\boxed{11}$) requires order of magnitude R time, leading to a total time of $O(mR^2) \in O(mk^2D_G^2)$ because $R \in O(kD_G)$ from Proposition $\boxed{11}$ The total time for these ancilliary procedures is $O(mk^2D_G^2 + |\mathcal{I}_1|k^4\Delta_G)$

Therefore, the time complexity of bookkeeping and setup is $O(mk^2D_G^2 + |\mathcal{I}_1|k^4\Delta_G +$ |V|+|E|) $\in \widehat{O}(k^4)$. The time complexity at each random walk step is $O(k-1^2\Delta_G+k-1^4)$ $\widehat{\mathcal{O}}(k^4)$ from Appendix D.2 and Appendix E.2.1 We assume that the expected number of rejections in Line 21 is given by C_{REJ} . The total number of random walk steps is given by $O(R m C_{REJ})$ times the expected tour length. By Lemma 5 and proposition 11 the expected

Graph	Rel. dispersion of estimates					
Отарп	6	8	10	12		
Amazon	0.203	0.241	0.285	0.268		
DBLP	0.023	0.023	0.041	0.054		
Patents	0.037	0.083	0.093	0.123		
Pokec	0.065	0.044	0.037	0.046		
LiveJ.	0.050	0.060	0.033	0.066		
Orkut	0.021	0.761	0.053	0.031		

Table 5: Dispersion, $\frac{\max - \min}{\max}$, of Ripple's estimates of $|\mathcal{V}^{(k)}|$ computed using $\epsilon = 0.003$, $|\mathcal{I}_1| = 10^4$ and $M = 10^7$ for k = 6, 8, 10, 12 used in the analysis in Section 5.1 The selected hyper-parameters provide reasonably similar estimates over 10 independent runs. Orkut for k = 8 exhibits the largest dispersion because of the presence of a single outlier.

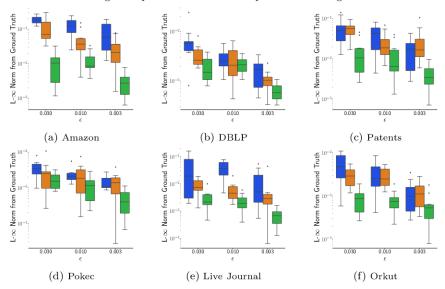


Fig. 5: Accuracy and convergence analysis for 5-CIS s using L- ∞ norm. The y-axes show the L- ∞ norm between the Ripple estimate and the ground truth vector $\mathcal{C}^{(5)}$, containing counts of all possible non-isomorphic subgraph patterns for various settings of ϵ and \mathcal{I}_1 . As expected, the accuracy improves as ϵ decreases and $|\mathcal{I}_1|$ increases. Each box and whisker represents 10 runs.

tour length is $O(\Delta_{\mathcal{G}^{(k-1)}}) \equiv O(k^2 \Delta_G)$. Therefore, the total number of random walk steps is $O(k^3 m D_G \Delta_G C_{\text{REJ}})$.

 $O(|\mathcal{H}|)$ time is to print the output $\hat{\mu}$. We assume that updating $\hat{\mu}$ is amortized in constant order if we use a hashmap to store elements of the vector, and because updating a single key in said hashmap is by Equation (9) increments, the proof is completed.

F Additional Results

We now present the results of additional experiments performed on Ripple. Table 5 shows the dispersion, $\frac{\max - \min}{\max}$, of the estimates that were used to measure the running time and space utilization of Ripple in Section 5.1 Figure 5 shows the L- ∞ norm from the ground truth for k=5 with $M=10^7$ while ϵ and \mathcal{I}_1 vary.

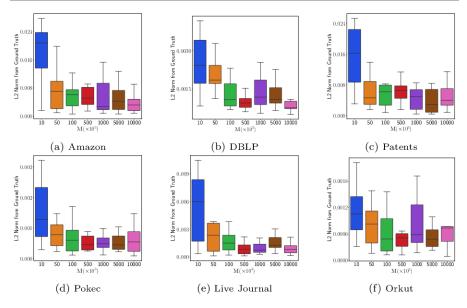


Fig. 6: Sensitivity of Ripple to the reservoir capacity ${\tt M}$ for k=5. We verify that a larger reservoir improves the accuracy of Ripple estimates in all graphs because it reduces oversampling bias. Each box and whisker plot represents 10 runs.

Trade-off between Convergence and Reservoir size. Next, we measure the effect of the reservoir capacity M on accuracy, as discussed in Section [4.3] We vary M from 50000 to 10^7 while keeping the other parameters fixed as $\epsilon = 0.003$ and $|\mathcal{I}_1| = 10^4$ and measure the L2-norm between the Ripple estimate and the exact value of the count vector $\mathcal{C}^{(5)}$, such as in Section [5.2] We see that larger reservoirs reduce oversampling bias and improve the convergence and accuracy in all datasets.

Scalability on Number of Threads (Figure 7). In this experiment, because of Equation 10, we set $\epsilon = 0.001$ to force a larger number of tours, thereby increasing the load per core and ensuring a sufficient workload. Further, we fix k = 5, set $|\mathcal{I}_1| = 10^4$, $|\mathcal{I}_1| = 10^4$, and compute running times over 10 executions while excluding the graph read time, which is not parallel. We observe that our implementation does not scale linearly: as we double the number of cores, the running time decreases by $|\mathcal{I}_1| = 10^4$, and $|\mathcal{I}_2| = 10^4$. Local profiling using hardware performance counters (Linux's perf) suggests that this overhead is an outcome of increased random-access patterns of in-memory graph data, which limits the overall use of the underlying processing pipeline. Indeed, sub-optimal access patterns of graph data are a known issue that is currently handled by dedicated accelerator hardware deploying optimized and specific caching mechanisms and memory access policies for workloads dominated by subgraph enumeration (Yao et al.) 2020).

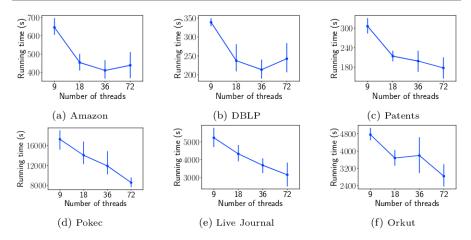


Fig. 7: Scalability w.r.t. the Number of Threads for 5-CIS. Despite a noticeable reduction in running time, the scalability is not linear. In fact, as we double the number of cores, the running time decreases by approximately a quarter instead of half, which may be due to the memory bandwidth limit coupled with the lack of memory locality, which is a ubiquitous problem in graph mining algorithms.