

Number of Connected Components in a Graph: Estimation via Counting Patterns

Ashish Khetan, Harshay Shah, and Sewoong Oh*
UIUC

Abstract

Due to the limited resources and the scale of the graphs in modern datasets, we often get to observe a sampled subgraph of a larger original graph of interest, whether it is the worldwide web that has been crawled or social connections that have been surveyed. Inferring a global property of the original graph from such a sampled subgraph is of a fundamental interest. In this work, we focus on estimating the number of connected components. It is a challenging problem and, for general graphs, little is known about the connection between the observed subgraph and the number of connected components of the original graph. In order to make this connection, we propose a highly redundant and large-dimensional representation of the subgraph, which at first glance seems counter-intuitive. A subgraph is represented by the counts of patterns, known as network motifs. This representation is crucial in introducing a novel estimator for the number of connected components for general graphs, under the knowledge of the spectral gap of the original graph. The connection is made precise via the Schatten k -norms of the graph Laplacian and the spectral representation of the number of connected components. We provide a guarantee on the resulting mean squared error that characterizes the bias variance tradeoff. Experiments on synthetic and real-world graphs suggest that we improve upon competing algorithms for graphs with spectral gaps bounded away from zero.

1 Introduction

With the increasing size of modern datasets, a common network analysis task involves sampling a graph, due to restrictions on memory, communication, and computation resources. From such a subgraph with sampled nodes and their interconnections, we want to infer some global properties of the original graph that are relevant to the application in hand. This paper focuses on the task of inferring the number of connected components. It is a fundamental graph property of interest in various applications such as estimating the weight of the minimum spanning trees [5, 2], estimating the number of classes in a population [12], and visualizing large networks [19].

In the sampled subgraph, the count of connected components in general can be smaller as well as larger than that of the original graph. Some connected components might not be sampled at all, whereas the connected nodes in the original graph is not guaranteed to be connected in the subgraph. It is not clear how the true number of components is related to the complex structure of the sampled graph. For general graphs, it is unknown how to unravel the complex relationship between the sampled subgraph and the global property of interest. In this paper, we propose

*Author emails are: ashish.khetan09@gmail.com, hrshah4@illinois.edu, swoh@illinois.edu

encoding the sampled subgraph by counting patterns in the subgraph, and show that it makes its connection to the number of connected components transparent.

We represent a graph by a vector of counts of all possible patterns, also known as network motifs. For example, the first and second entries in this count vector encodes the number of nodes and (twice) the number of edges, respectively. Later entries encode the count of increasingly complex patterns: the number of times a pattern is repeated in the graph. This vector is clearly a redundant over-representation whose dimension scales super exponentially in the graph size. Perhaps surprisingly, for the purpose of approximately inferring a global property, it suffices to have the first few hundred dimensions of this vector, corresponding to the counts of small patterns. For counting those patterns, we introduce novel algorithms, and give a precise characterization of how the complexity (the size of the patterns included in the estimation) trades off with accuracy (the mean squared error).

Problem statement and our proposed approach. We want to estimate the number of connected components in a simple graph $G = (V, E)$ from a sampled subset of its nodes and the corresponding subgraph. Let N be the number of vertices and $\text{cc}(G)$ the number of connected components in G . We consider the subgraph sampling model, that is, a subset of vertices is sampled at random and the induced subgraph is observed. We consider a Bernoulli sampling model, where each vertex is sampled independently with a probability p . Let Ω be the set of randomly observed vertices, and G_Ω be the corresponding induced subgraph, i.e. $G_\Omega = (\Omega, E_\Omega)$ where $(i, j) \in E_\Omega$ if $i, j \in \Omega$ and $(i, j) \in E$. We want to estimate $\text{cc}(G)$ from G_Ω . We propose a novel spectral approach, which makes transparent the relation between the counts of patterns and the number of connected components.

We propose characterizing the number of connected components as the count of zero eigenvalues of its Laplacian matrix $L \in \mathbb{R}^{n \times n}$ given by

$$L \equiv D - A, \tag{1}$$

where $D = \text{diag}(A\mathbf{1})$ is the diagonal matrix of the degrees, and A is the adjacency matrix of the graph G . The rank of L reveals $\text{cc}(G)$ as

$$\begin{aligned} \text{cc}(G) &= N - \text{rank}(L) \\ &= N - \sum_{i \in [N]} \mathbb{I}[\sigma_i(L) > 0], \end{aligned} \tag{2}$$

where the $\sigma_i(L)$'s are the singular values of the graph Laplacian L . Using this relation directly for estimation is an overkill as estimating the singular values is more challenging than estimating $\text{cc}(G)$. Instead, we use a few steps of functional approximations to relate to the pattern counts. By Gershgorin's circle theorem, we have $\sigma_i(L) \leq 2d_{\max}$, where d_{\max} is the maximum degree in G . We therefore normalize L by $1/\beta$ for some $\beta \geq 2d_{\max}$ to ensure all eigenvalues lie in the unit interval $[0, 1]$ and denote it by $\tilde{L} = (1/\beta)L$. For any constant $0 < \alpha < 1$ that separates the zero and non-zero eigenvalues such that $\alpha < \min_i \{\sigma_i(\tilde{L}) : \sigma_i(\tilde{L}) \neq 0\}$, we consider the following approximation of the rank function. We approximate the step function in (2) by a continuous piecewise linear function

$H_\alpha : [0, 1] \rightarrow [0, 1]$ illustrated in Figure 4:

$$\begin{aligned} H_\alpha(x) &= \begin{cases} 1 & \text{if } x \in [\alpha, 1], \\ \frac{x}{\alpha} & \text{if } x \in [0, \alpha]. \end{cases}, \quad \text{and} \\ \text{cc}(G) &= N - \sum_{i \in [N]} H_\alpha(\sigma_i(\tilde{L})), \end{aligned} \quad (3)$$

where we used the fact that the approximation is exact under our assumption that the spectral gap is lower bounded by α . To connect it to the pattern counts, we propose a further approximation using a polynomial function $f_\alpha : \mathbb{R} \rightarrow \mathbb{R}$ of a finite degree m . Precisely, for $f_\alpha(x) = a_1x + \dots + a_mx^m$ (e.g. Figure 4), we immediately have the following relation:

$$\sum_{i=1}^N f_\alpha(\sigma_i(\tilde{L})) = \sum_{k=1}^m \frac{a_k}{\beta^k} \|L\|_k^k, \quad (4)$$

where $\|L\|_k^k$ is the Schatten- k norm of L which is defined as sum of k -th power of its singular values: $\|L\|_k^k \equiv \sum_{i=1}^N \sigma_i(L)^k$. As we choose $f_\alpha(x)$ to be a close approximation of the desired $H_\alpha(x)$, we have the following approximate relation: $\text{cc}(G) \approx N - \sum_{k=1}^m (a_k/\beta^k) \|L\|_k^k$, which can be made arbitrarily close by choosing a larger degree m .

Finally, we propose using the fact that $\|L\|_k^k = \text{Tr}(L^k)$ is a sum of the weights of all length k closed walks. Once we compute the (weighted) count of those walks for each pattern, this gives a direct formula to approximate the number of connected components from the counts. This approximation can be made as accurate as we want, by choosing the right order m in the polynomial approximation. Unlike the singular values, the (weighted) counts can be directly estimated from the sampled subgraph in a statistically efficient manner. We introduce a novel unbiased estimator $\hat{\Theta}_k(G_\Omega)$ for Schatten- k norms of L in Section 2 that uses the counts of patterns in the sampled subgraph, and appropriately aggregates the estimated counts of the original graph. Together with a polynomial approximation $f_\alpha(x)$, this gives a novel estimator:

$$\widehat{\text{cc}}(G_\Omega, \alpha, \beta, m) \equiv N - \sum_{k=1}^m \frac{a_k}{\beta^k} \hat{\Theta}_k(G_\Omega), \quad (5)$$

where $\hat{\Theta}_k(G_\Omega)$ is an unbiased estimate of Schatten- k norm of L defined in (12) and a_k 's are the coefficients in the polynomial approximation $f_\alpha(x) = a_1x + \dots + a_mx^m$ as defined as in (22).

Related work. The connection between the number of connected components in the original graph and the counts of various patterns in the sampled graph has been explored in [9, 10] for limited classes of graphs with particular structures. These estimators are customized for two simple extreme cases of *forests* and *unions of disjoint cliques*, and rely only on the counts of a few extremely simple patterns.

For a *forest* $G = (V, E)$, the estimator introduced in [9] exploits the simple relation that the number of connected components is $\text{cc}(G) = |V| - |E|$. Hence, we only need to estimate the number of edges. This is a straightforward procedure that uses the counts of k -stars in the sampled subgraph for $k \in \{0, 1, \dots\}$. A k -star is a graph with one central node with k adjacent nodes, mutually disjoint.

For a *union of disjoint cliques* $G = (V, E)$, the estimator introduced in [9] exploits the simple relation that the number of connected components is $\text{cc}(G) = \sum_{k=1}^{|V|} \{ \# \text{ of cliques of size } k \}$. We only need to estimate the number of cliques of each size k in the original graph. This is straightforward as the observed size of the cliques follow a multinomial distribution. This requires only the counts of k -cliques in the sampled subgraph for $k \in \{1, 2, \dots\}$ [12, 9]. A k -clique is a fully connected graph with k nodes. These approaches have recently been extended in [15] to include chordal graphs, which introduces a novel idea of smoothing to achieve a strong performance guarantees. However, none of these methods can be applied to our setting where we consider the original graph to be a general graph.

Contributions. We pose the problem of estimating the number of connected components as a spectral estimation problem of estimating the rank of the graph Laplacian. This is further split into two tasks of first estimating the Schatten k -norms of the Laplacian and then applying a functional approximation.

We propose an unbiased estimator of the Schatten k -norm $\|L\|_k^k$ based on the counts of patterns in the subsampled graph, known as k -cyclic pseudographs. The main challenge is in estimating the diagonal entries of L (which is the degree of each node), that is critical in computing the weighted counts of the k -cyclic pseudographs. To overcome this challenge, in Section 3 we introduce an estimator that uses a novel idea of partitioning the subsampled graph and stitching the estimated degrees in each partition together.

Combining the estimated Schatten norms with polynomial approximation of $H_\alpha(x)$ in (3), we introduce a novel estimator of the number of connected components. To the best of our knowledge, this is the first estimator with theoretical guarantees for general graphs. We provide a sharp characterization of the bias-variance tradeoff of our estimator in Section 5. Numerical experiments on synthetic and real-world graphs show that the proposed method improves upon the competing baseline, with a comparable run time.

2 Unbiased estimator of Schatten- k norms of a graph Laplacian

In this section, we focus on the unnormalized L as Schatten norms are homogeneous and the normalization can be applied afterwards. We first provide an alternative method for computing $\|L\|_k$, and show how it leads to a novel estimator of the Schatten norm from a sampled subgraph. We use an alternative expression of the Schatten k -norm of a positive semidefinite L as the trace of the k -th power:

$$(\|L\|_k)^k = \text{Tr}(L^k). \quad (6)$$

Such a sum of the diagonal entries is the sum of weights of all closed walks of length k , where the weight of a walk is defined as follows. A length- k closed walk in $G = (V, E)$ is a sequence of vertices $w = (w_1, w_2, \dots, w_k, w_{k+1})$ with $w_1 = w_{k+1}$ and either $(w_i, w_{i+1}) \in E$ or $w_i = w_{i+1}$ for all $i \in [k]$. Note that we allow repeated nodes and repeated edges. Essentially, these are walks in a graph G augmented by self-loops at each of the nodes. We define the *weight* of a walk w in G to be

$$\mu_G(w) \equiv \prod_{i=1}^k L_{w_i w_{i+1}}, \quad (7)$$

which is the product of the weights along the walk and $L = D - A$ is the graph Laplacian. It follows from (6) that

$$\|L\|_k^k = \sum_{w: \text{all length } k \text{ closed walks}} \mu_G(w). \quad (8)$$

Even though this formula holds for any general matrix L , it simplifies significantly for graph Laplacians, as its all non-zero off-diagonal entries are -1 (and its diagonal entries are the degrees of the nodes). Consider a length-3 walk $w = (u, v, v, u)$ whose pattern is shown in the subgraph A_2 in Figure 1. This walk has weight $\mu_G(w) = (-1)^2 d_v$, where d_v is the degree of node v . Similarly, a walk (u, u, u, u) of pattern A_1 in Figure 1 has weight $\mu_G(w) = d_u^3$, and a walk $w = (u, v, x, u)$ of pattern A_3 has weight $\mu_G(w) = (-1)^3$.

In general, for a node u in a walk w of length k , let s_u denote the number of self-loops traversed in the walk on node u . Then, it follows that

$$\mu_G(w) = (-1)^{(k - \sum_{u \in w} s_u)} \prod_{u \in w} d_u^{s_u}, \quad (9)$$

where d_u is the degree of node u in G . The weight of a walk is ± 1 if there are no self loops in the walk. Otherwise, its absolute value is the product of the degrees of the vertices corresponding to the self loops, and its sign is determined by how many non-self loop edges there are.

The first critical step in our approach is to partition the summation in Eq. (8) according to the pattern of the respective walk, which will make (i) counting those walks of the same pattern more efficient; and (ii) also de-biasing straight forward (see Equation (12)) under ransom sampling. We refer to component-wise scaling w.r.t. the inverse of the probability of being sampled as de-biasing, which is a critical step in our approach and will be explained in detail later in this section. Following the notations from enumeration of small cycles in [1] and [14], we use the family of patterns called *k-cyclic pseudographs*:

$$\|L\|_k^k = \sum_{H \in \mathbb{H}_k} \sum_{w: \mathcal{H}(w)=H} \mu_G(w), \quad (10)$$

where \mathbb{H}_k is the set of *patterns* that have k edges, and $\{w : \mathcal{H}(w) = H\}$ is the set of walks on G that have the same pattern H . We give formal definitions below. *k-cyclic pseudographs* expand the standard notion of simple *k-cyclic graphs*, and include multi-edges and loops, which explains the name *pseudograph*.

Definition 1. Let $C_k = (V_k, E_k)$ denote the undirected simple cycle with k nodes. An unlabelled and undirected pseudograph $H = (V_H, E_H)$ is called a *k-cyclic pseudograph* for $k \geq 3$ if there exists an onto node-mapping from $C_k = (V_k, E_k)$, i.e. $f : V_k \rightarrow V_H$, and a one-to-one edge-mapping $g : E_k \rightarrow E_H$ such that $g(e) = (f(u_e), f(v_e))$ for all $e = (u_e, v_e) \in E_k$. We use \mathbb{H}_k to denote the set of all *k-cyclic pseudographs*. We use $c(H)$ to the number of different node mappings f from C_k to a *k-cyclic pseudograph* H . Each closed walk w of length k is associated with one of the graphs in \mathbb{H}_k , as there is a unique H that the walk is an Eulerian cycle of under a one-to-one mapping of the nodes. We denote this graph by $\mathcal{H}(w) \in \mathbb{H}_k$.

Figure 1 shows examples of all 3-cyclic pseudographs. $\mathbb{H}_3 = \{A_1, A_2, A_3\}$ and each one is a distinct pattern that can be mapped from a triangle graph C_3 . In the case of A_1 , there is only one

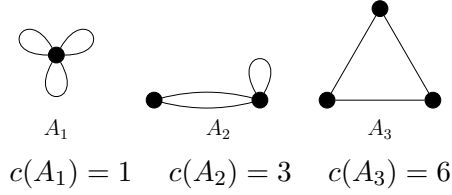


Figure 1: The 3-cyclic pseudographs $\mathbb{H}_3 = \{A_1, A_2, A_3\}$.

mapping from C_3 to A_1 and corresponding multiplicity is $c(A_1) = 1$. Also, a walk $w = (u, u, u, u)$ on the graph G has pattern A_1 , which we denote by $\mathcal{H}(w) = A_1$. In the case of A_2 , any of the three nodes can be mapped to the left-node of A_2 , which gives $c(A_2) = 3$. In the case of A_3 , each permutation of the three nodes are distinct, which gives $c(A_3) = 6$. We show more examples of length 4 in Figure 2. k -cyclic pseudographs for larger k can be enumerated as well (e.g. [14]).




Figure 2: The 4-cyclic pseudographs \mathbb{H}_4 .

For a pattern H , let S_H denote the set of self-loops in H , and s_u denote the number of self-loops at node u in the walk w . Then the summation of walks can be partitioned according to their patterns as:

$$\|L\|_k^k = \sum_{H \in \mathbb{H}_k} (-1)^{k-|S_H|} \left\{ \sum_{w: \mathcal{H}(w)=H} \prod_{u \in w} d_u^{s_u} \right\}, \quad (11)$$

which follows from substituting (9) in (10). This expression does not require the (computation of) singular values and leads to a natural unbiased estimator given a sampled subgraph. As the probability of a walk being sampled depends only on the pattern, we introduce a novel estimator $\hat{\Theta}_k(G_\Omega)$ of $\|L\|_k^k$ that de-biases each pattern separately:

$$\hat{\Theta}_k(G_\Omega) = \sum_{H \in \mathbb{H}_k} \frac{(-1)^{k-|S_H|}}{p^{|V_H|}} \left\{ \sum_{w: \mathcal{H}(w)=H} \theta_w(G_\Omega) \mathbb{I}(w \subseteq G_\Omega) \right\}, \quad (12)$$

where $|V_H|$ is the number of nodes in H , $p^{|V_H|}$ is the probability that walk with pattern H is sampled (i.e. all edges involves in the walk are present in the sampled subgraph G_Ω), and $\mathbb{I}(w \subseteq G_\Omega)$ denotes the indicator that all nodes in the walk w are sampled. $\theta_w(G_\Omega)$ is defined below.

As the degrees of the nodes in the original graph are unknown, it is challenging to estimate the polynomial of the degrees $\prod_{u \in w} d_u^{s_u}$ in Eq. (11), from the sampled graph. To this end, we introduce

a novel estimator $\theta_w(G_\Omega)$ in Section 3, which is unbiased; it satisfies

$$\mathbb{E}[\theta_w(G_\Omega) | \mathbb{I}(w \subseteq G_\Omega)] = \prod_{u \in w} d_u^{s_u}.$$

It immediately follows by taking the expectation of (12), that $\widehat{\Theta}_k(G_\Omega)$ is unbiased, i.e.

$$\mathbb{E}_\Omega[\widehat{\Theta}_k(G_\Omega)] = \|L\|_k^k. \quad (13)$$

3 An unbiased estimator of the polynomial of the degrees

Our strategy to get an unbiased estimator of $\prod_{u \in w} d_u^{s_u}$ is to first partitioning the nodes in the original graph G to get a more insightful factorization of $\prod_{u \in w} d_u^{s_u}$ in Eq. (16) (see Figure 3) that removes dependences between the summands, and next by estimating each term independently in the factorization.




Figure 3: We are partitioning the original graph G with respect to a length- $(k = 6)$ closed walk $w = (1, 2, 2, 2, 3, 3, 1)$. Its corresponding k -cyclic pseudograph $\mathcal{H}(w) \in \mathbb{H}_6$ is shown on the top. Red nodes are not connected to either 2 or 3, which are the nodes of interest in w as they have self loops. Blue nodes are only connected to 2, purple to only 3, and blue to both 2 and 3.

Consider a concrete task of estimating $\prod_{u \in w} d_u^{s_u} = (d_2)^2(d_3)^1 = 6^2 \times 6$, for a walk $w = (1, 2, 2, 2, 3, 3, 1)$ in the observed subgraph G_Ω . Note that we only see G_Ω , whose degrees are very different from the original graph. For instance, node 2 now has degree 3 and node 3 has degree 3 in the sampled graph. Further, these random variables (the observed degrees) are correlated, making estimation challenging. To make such correlations apparent, we first give a novel partitioning of the nodes V in the following.

3.1 Partitioning V

Our strategy is first to partition the nodes V in the original graph, with respect to a walk $w = (w_1, \dots, w_{k+1})$ of interest. For a closed walk w , let $U = \{u_1, \dots, u_\ell\}$ denote the set of nodes in w

that have at least one self-loop, let $\ell = |U|$ denote its cardinality, and let $\{s_1, \dots, s_\ell\}$ denote the number of self-loops at each node. In the running example, we have $U = \{u_1 = 2, u_2 = 3\}$, $\ell = 2$, $s_1 = 2$, and $s_2 = 1$. As our goal is to estimate $(d_2)^2(d_3)$, we partition the nodes with respect to how they relate to the nodes in $U = \{2, 3\}$. Concretely, there are four partitions: nodes that are not connected to either 2 or 3 (shown in red in Figure 3), nodes that are only connected to 2 (shown in green), nodes that are only connected to 3 (shown in purple), and nodes that are connected to both 2 and 3 (shown in blue). Nodes in each partition contribute in different ways to the target quantity $(d_2)^2(d_3)$, which will be precisely captured in the *factorization* in Eq. (16). In general, we need to consider all such variations in the partitioning, which gives

$$V = \bigcup_{T \subseteq U} V_{T, U \setminus T}, \quad (14)$$

where $V_{T, T'} = \{\bigcap_{v \in T} \partial v\} \cap \{\bigcap_{v \in T'} \partial v^c\}$ is the set of nodes that are adjacent to all nodes in T but are not adjacent to any nodes in T' , and ∂v denotes the neighborhood of node v and ∂v^c denotes the complement of ∂v . We let $V_{\emptyset, U} = \bigcap_{v \in U} \partial v^c$ and $V_{U, \emptyset} = \bigcap_{v \in U} \partial v$. Essentially, we are labelling each node according to which nodes in U it is adjacent to, and grouping those nodes with the same label. In the running example, $V = V_{U, \emptyset} \cup V_{\{3\}, \{2\}} \cup V_{\{2\}, \{3\}} \cup V_{\emptyset, U}$, where the partitions are subset of nodes in blue, purple, green, and red, respectively.

Let $d_{T, U \setminus T} = |V_{T, U \setminus T}|$ denote the size of a partition such that

$$d_u = \sum_{T \in \mathcal{T}_u} d_{T, U \setminus T}, \quad (15)$$

for any $u \in U$ where $\mathcal{T}_u = \{T \subseteq U | u \in T\}$ is the set of subsets of U containing u . For example, $d_2 = 6$ which is the sum of blue and green nodes, and $d_3 = 6$ which is the sum of blue and purple nodes.

We are partitioning the neighborhood of u such that each term can be separately estimated. This ensures we handle the correlations among the degrees of different nodes in w correctly. The quantity of interest is

$$\begin{aligned} \prod_{i \in [\ell]} d_{u_i}^{s_i} &= \prod_{i \in [\ell]} \left(\sum_{T \in \mathcal{T}_{u_i}} d_{T, U \setminus T} \right)^{s_i} \\ &= \sum_{(T_1^{(1)}, \dots, T_1^{(s_1)}, \dots, T_\ell^{(1)}, \dots, T_\ell^{(s_\ell)}) \in (\mathcal{T}_{u_1})^{s_1} \times \dots \times (\mathcal{T}_{u_\ell})^{s_\ell}} \left\{ \prod_{j=1}^{\ell} \prod_{i=1}^{s_j} d_{T_j^{(i)}, U \setminus T_j^{(i)}} \right\}, \end{aligned} \quad (16)$$

where $T_j^{(i)}$ is a i -th choice of a set in \mathcal{T}_{u_j} that contains the node u_j for $i \in [s_j]$, and $[\ell] = \{1, \dots, \ell\}$ denotes the set of positive integers up to ℓ . The second equation follows directly from exchanging the product and the summation. This alternative expression is crucial in designing an unbiased estimator, since each term in the summation can now be estimated separately as follows.

Consider a task of estimating a single term in (16), and we merge those $T_j^{(i)}$'s that happen to be identical:

$$\prod_{j=1}^{\ell} \prod_{i=1}^{s_j} d_{T_j^{(i)}, U \setminus T_j^{(i)}} = \prod_{T \in \mathbb{T}} (d_{T, U \setminus T})^{t_T}, \quad (17)$$

where $\mathbb{T} = \{T_1^{(1)}, \dots, T_1^{(s_1)}, \dots, T_\ell^{(1)}, \dots, T_\ell^{(s_\ell)}\}$ is the current set of partitions allowing for multiple entries of the same set, and t_T is the multiplicity, i.e. how many times a set T appears in the set $\mathbb{T} = (T_1^{(1)}, \dots, T_1^{(s_1)}, \dots, T_\ell^{(1)}, \dots, T_\ell^{(s_\ell)})$. Each term in the right-hand side can be now separately estimated, as (a) $V_{T,U \setminus T}$'s are disjoint and (b) we know for the sampled subgraph the membership of each sampled node. This follows from the fact that, conditioned on the event that $\{w \subseteq \Omega\}$, we know how the sampled nodes in Ω are connected to any node in $\{w_i\}_{i=1}^{k+1}$ and in particular those with self-loops denoted by U . Hence, for any node in Ω the membership (or the color in the Figure 3) is trivially revealed. Therefore, we can handle (the degrees $d_{T,U \setminus T}$ in) each partition separately, and estimate each monomial in $\prod_{T \in \mathbb{T}} (d_{T,U \setminus T})^{t_T}$. The problem is reduced to the task of estimating $d_{T,U \setminus T}^s$ for some integer s and some partition $V_{T,U \setminus T}$.

3.2 Unbiased estimator of $d_{T,U \setminus T}^s$

From the original graph $G = (V, E)$ (where the size of each partition is denoted by $d_{T,U \setminus T}$), we observe a sampled subgraph $G_\Omega = (\Omega, E_\Omega)$ (where the size of each partition in G_Ω is denoted by $d_{T,U \setminus T}(\Omega)$), and we let $d_{T,U \setminus T}(w)$ denote the size of the partition intersecting the walk $w = (w_1, \dots, w_{k+1})$. Precisely, $d_{T,U \setminus T}(\Omega) \equiv |V_{T,U \setminus T} \cap \Omega|$, and $d_{T,U \setminus T}(w) \equiv |V_{T,U \setminus T} \cap \{w_i\}_{i=1}^{k+1}|$. We do not allow multiple counts when computing the size, such that $d_{\{2\},\{3\}}(\Omega) = 2$ and $d_{\{2\},\{3\}}(w) = 1$, in the example.

Let us focus on a particular walk w on a graph G , its corresponding U and a fixed $T \subseteq U$, such that $V_{T,U \setminus T}$ and $d_{T,U \setminus T}$ are fixed. Now $d_{T,U \setminus T}(\Omega)$ is a random variable representing how many nodes in the partition $V_{T,U \setminus T}$ are sampled. Conditioned on the fact that w is sampled, and hence a $d_{T,U \setminus T}(w)$ sampled nodes are already observed, the remaining $(d_{T,U \setminus T} - d_{T,U \setminus T}(w))$ nodes are sampled i.i.d. with probability p . Hence, conditioned on $\{w \subseteq \Omega\}$, the size of the sampled partition is distributed as

$$d_{T,U \setminus T}(\Omega) \sim \text{Binom}(d_{T,U \setminus T} - d_{T,U \setminus T}(w), p) + d_{T,U \setminus T}(w). \quad (18)$$

This leads to a natural unbiased estimator of the monomial $d_{T,U \setminus T}^s$ as

$$\tilde{d}_{T,U \setminus T}^{(s)} = \langle (A^{-1})_{s+1}, \bar{d} \rangle, \quad (19)$$

where $\bar{d} = [1, d_{T,U \setminus T}(\Omega), d_{T,U \setminus T}(\Omega)^2, \dots, d_{T,U \setminus T}(\Omega)^s]^\top$ is a column vector in \mathbb{R}^{s+1} of the monomials of the observed size of the partition, A is the unique matrix satisfying

$$\mathbb{E}[\bar{d}] = A [1, d_{T,U \setminus T}, \dots, d_{T,U \setminus T}^s]^\top, \quad (20)$$

and $(A^{-1})_{s+1}$ is the $(s+1)$ -th row of A^{-1} . One can check immediately that $\mathbb{E}[\tilde{d}_{T,U \setminus T}^{(s)}] = \langle (A^{-1})_{s+1}, \mathbb{E}[\bar{d}] \rangle = d_{T,U \setminus T}^s$, hence giving the desired unbiased estimator. The matrix A is a lower triangular matrix which depends only on s , p and the structure of the walk via $d_{T,U \setminus T}(w)$. In terms of these three parameters, the required vector $(A^{-1})_{s+1}$ has a closed form expression, and hence the estimator can be computed in a straight forward manner. It uses the moments of a binomial distribution, which can be computed immediately.

An example of A for $s = 3$ is given in (51), where one should plug-in $\ell = d_{T,U \setminus T}(w) + 1$, $\omega = d_{T,U \setminus T} + 1$ and $\tilde{\tau} = d_{T,U \setminus T}(\Omega)$. This leads to an unbiased estimator of $\prod_{i \in [\ell]} d_{u_i}^{s_i}$ by replacing

(19) into (17) and (16):

$$\theta(w, G_\Omega) = \sum_{(T_1^{(1)}, \dots, T_1^{(s_1)}, \dots, T_\ell^{(1)}, \dots, T_\ell^{(s_\ell)}) \in (\mathcal{T}_{u_1})^{s_1} \times \dots \times (\mathcal{T}_{u_\ell})^{s_\ell}} \left\{ \prod_{T \in \mathbb{T}} \tilde{d}_{T, U \setminus T}^{(t_T)} \right\}. \quad (21)$$

By construction, it is immediate that the estimator is unbiased: $\mathbb{E}[\theta(w, G_\Omega) | I(w \subseteq \Omega)] = \prod_{u \in w} d_u^{s_u}$.

4 Polynomial approximation

The remaining goal in our approach is to design a polynomial approximation of the target function $H_\alpha : [0, 1] \rightarrow [0, 1]$ defined in (3) for a fixed scalar $\alpha \in (0, 1)$. Concretely, for a given integer m , we want a degree- m polynomial approximation $f(x)$ of $H_\alpha(x)$ such that (i) $f(0) = 0$; (ii) the approximation error (as measured by the ℓ_∞ norm) is small in the interval $[\alpha, 1]$; and (iii) we can provide an upper bound on the approximation error: $\max_{x \in [\alpha, 1]} |H_\alpha(x) - f(x)|$. The first condition can be met by any function with proper scaling and shifting, and strictly enforcing it ensures that we make fair comparisons. The second condition ensures we have a good approximation, as the non-zero singular values only lie in the interval $[\alpha, 1]$. In particular, the approximation error outside of this interval is irrelevant. The last condition ensures we get the desired performance guarantees for the estimation error of the number of connected components. The (upper bound on the) approximation error of the polynomial function directly translates into the end-to-end error on the estimation.

A first attempt might be to use a Chebyshev approximation [6, 16] directly on $H_\alpha(x)$. This is optimal in terms of achieving a target ℓ_∞ error with the smallest degree in all regimes of $[0, 1]$. However, we only care about the ℓ_∞ error in $[\alpha, 1]$. As shown in magenta curve in Figure 4, the Chebyshev approximation $C(x)$ unnecessarily fits the curve in $(0, \alpha]$, resulting in larger error in $[\alpha, 1]$.

A natural fix is to use filter design techniques, e.g. Parks-McClellan algorithm [17], where Chebyshev polynomials have been applied to design high pass filters with similar constraints as ours. This will give a polynomial approximation with small approximation error in the desired pass band of $[\alpha, 1]$. However, these techniques do not come with the desired approximation guarantee that we seek.

One approach proposed in [23] does come with a provable error bound. This approximation $B(x)$ composes a Chebyshev approximation of a constant degree q with the CDF of a beta distribution of degree $(m/q - 1)/2$. The beta distribution boosts the approximation of the function in the interval $[\alpha, 1]$, thus providing an error bound of $O((c_\alpha)^m)$, where c_α is a constant that depends on α . Figure 4 shows that $B(x)$ (in green) still unnecessarily fits the curve in $(0, \alpha]$, as it starts with a (lower-degree) Chebyshev approximation of $H_\alpha(x)$.

Our goal is to design a new polynomial approximation that ignores the region $(0, \alpha]$ completely, such that it achieves improved performance in $[\alpha, 1]$, and also comes with a provable error bound. We propose using a parametric family that ensures $f_b(0) = 0$:

$$f_b(x) = 1 - \prod_{i=1}^m (1 - b_i x), \quad (22)$$

for a vector $b = [b_1, \dots, b_m] \in \mathbb{R}^m$. We provide an upper bound on the approximation error achieved by the optimal b^* , provide a choice of \tilde{b} in a closed form that achieves the same error bound, and provide a heuristic for locally searching for the optimal b^* to improve upon the closed-form \tilde{b} .

$H_\alpha(x)$ and its approximations

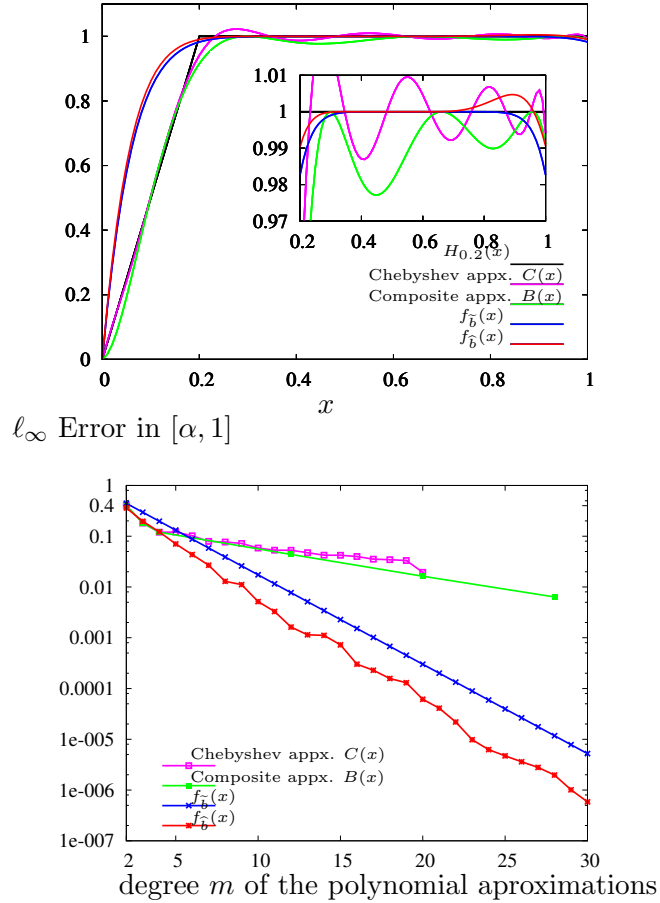


Figure 4: Top: four polynomial approximations to $H_{0.2}(x)$ of degree $k = 10$, $f_{\tilde{b}}(x)$ with \tilde{b} chosen according to Algorithm 1, $f_{\tilde{b}}(x)$ with $\tilde{b} = (2/1.2)\mathbf{1}$ as prescribed above, Chebyshev approximation $C(x)$, and the composite approximation $B(x)$ from [23]. Bottom: approximation error achieved by the proposed $f_{\tilde{b}}(x)$ improves upon other polynomial functions.

Proposition 2. For any $\alpha \in (0, 1)$ and $m \geq 2$, the optimal parameter $b^* \in \arg \min_{b \in \mathbb{R}^m} \max_{x \in [\alpha, 1]} |H_\alpha(x) - f_b(x)|$ achieves error bounded by

$$\max_{x \in [\alpha, 1]} |H_\alpha(x) - f_{b^*}(x)| \leq \left(\frac{1 - \alpha}{1 + \alpha} \right)^m. \quad (23)$$

A proof is provided in Section A.9. As the optimal b^* is challenging to find, one option is to simplify the optimization by searching over a smaller space. By constraining all b_i 's to be the same,

solving for minimum ℓ_∞ error gives a closed form solution $\tilde{b} \equiv (2/(1 + \alpha))[1, \dots, 1]$ that achieves the bound in (23) with equality, i.e. $\max_{x \in [\alpha, 1]} |H_\alpha(x) - f_{\tilde{b}}(x)| = ((1 - \alpha)/(1 + \alpha))^m$.

For practical use, we prescribe a slightly better approximation using a local search algorithm in Algorithm 1. The approximation guarantee is compared for $\alpha = 0.2$ and varying m in Figure 4 against the analytical choice $f_{\tilde{b}}(x)$, the standard Chebyshev approximation $C(x)$ of the first kind, and the approximation $B(x)$ from [23]. The proposed $f_{\hat{b}}(x)$ significantly improves upon both, achieving a faster convergence. The key idea is to exploit the fact that we care about approximating only in the regime of $[\alpha, 1]$. There might be other techniques to design better polynomial approximation than ours, e.g. [17], but might not come with a performance guarantee.

The inset in the top panel of Figure 4 illustrates how the proposed \hat{b} in red admits more fluctuations to achieve smaller ℓ_∞ error, compared to the uniform choice of \tilde{b} . In Algorithm 1, starting from a moderate perturbation around \tilde{b} , we iteratively identify the point x' achieving the maximum error and update b such that the error at x' is decreased. This approximation can be done offline for many random initializations for the desired α and m ; the one with minimum error can be stored for later use.

Algorithm 1 Local search for a polynomial approximation

Require: degree m , α , number of iterations T , step size $\delta > 0$

Ensure: $\hat{b} \in \mathbb{R}^m$

$b_i \leftarrow (2/(1 + \alpha)) + \text{U}[-1, 1]$ for all $i \in [m]$

for $t = 1$ **to** T **do**

$x' \leftarrow \arg \max_{x \in [\alpha, 1]} |\prod_{i \in [m]} (1 - b_i x)|$

$b \leftarrow b + \text{sign}(1 - f_b(x')) \times \delta \times \nabla_b f_b(x')$

end for

5 Main results

The polynomial approximation $f_{\hat{b}}(x)$ of the form (22) can easily be translated into the standard polynomial with coefficients $a = (a_1, \dots, a_m)$ such that $f_{\hat{b}}(x) = a_1 x + \dots + a_m x^m$. Together with the Schatten norm estimator $\hat{\Theta}_k(G_\Omega)$ in (12), this gives the proposed estimate $\widehat{\text{cc}}(G_\Omega, \alpha, \beta, m)$ in (5). We first give an upper bound on the multiplicative error for a special case of union of cliques, and give a general bound in Theorem 4. The overall procedure achieves the following, for a special case of *union of cliques*, which are also called *transitive graphs*:

Theorem 3. *If the underlying graph G is a disjoint union of cliques with clique sizes ω_i , for each connected component $1 \leq i \leq \text{cc}(G)$, $\omega_{\max} \equiv \max_i \{\omega_i\}$ and $\omega_{\min} \equiv \min_i \{\omega_i\}$, then for any choice of $\beta \geq \omega_{\max}$ and $\alpha \leq \omega_{\min}/\beta$, and any integer $m \geq 1$, there exist a function $g(m) = O(m!)$ and a constant $C > 0$ such that for $\omega_{\min} > C$,*

$$\frac{\mathbb{E}[(\widehat{\text{cc}}(G_\Omega, \alpha, \beta, m) - \text{cc}(G))^2]}{\text{cc}(G)^2} \leq \frac{g(m)(1 - p^m)}{\text{cc}(G)^2 p \beta^2} \sum_{i=1}^{\text{cc}(G)} \left(\omega_i^4 (1 + (\omega_i p)^{1-2m}) \right) + \gamma^{2m} \frac{N^2}{\text{cc}(G)^2}, \quad (24)$$

where $\gamma = (1 - \alpha)/(1 + \alpha)$. Moreover, if there exist some positive constants c_i 's such that $\omega_i^3 p \geq c_i m!$ or $\omega_i p^3 \leq c_i/m!$ for all i , then (24) holds with $g(m) = O(c^m)$ for some constant $c > 0$.

A proof of Theorem 3 is provided in a longer version of this paper. This clearly shows the tradeoff between the variance (the first term in the RHS) and the bias (the second term in the RHS). If we choose larger m , our functional approximation becomes more accurate resulting in a smaller bias. However, this will require counting larger patterns in the estimate of $\widehat{\Theta}_k(G_\Omega)$, leading to a larger variance.

In general, the complexity of our estimator for union of cliques is $O(m \times \text{cc}(G))$, as all relevant quantities to compute $\widehat{\Theta}_k(G_\Omega)$ can be pre-computed and stored in a table for all combinations of k and the size of the observed cliques. At execution time, we only need to look up one number for each clique we observe and for each $\widehat{\Theta}_k(G_\Omega)$ we are estimating. Hence, the above guarantee also characterizes the trade-off between the computational complexity and the accuracy. For example, when spectral gap α is small, we need large m with longer run-time to get bias as small as we need. We emphasize here that our estimator is generic and does not assume the true graph is union of cliques. The same generic estimator happens to be more efficient, when the *observed* subgraph is a union of cliques.

Consider the bias term, which captures how the error increases for graphs with smaller spectral gap in L . The normalized spectral gap for union of cliques is $\omega_{\min}/\omega_{\max}$, and balanced components result in a small spectral gap and a more accurate estimation.

Consider the variance term, and as an extreme example, consider the case when all cliques are of the same size $\omega = N/\text{cc}(G)$. It immediately follows that for $\beta = \omega_{\max} = \omega_{\min}$ and $\alpha = 1$, there is no bias and $\gamma = 0$. Further assuming $\omega p > 1$, we can choose some small $m = O(1)$ to minimize the variance which scales as $O(N^2/(\text{cc}(G)^3 p))$. Hence, to achieve arbitrarily small error, it is sufficient to have sample size Np scale as $(N/\text{cc}(G))^3$. This implies that finite multiplicative error is guaranteed only for $\text{cc}(G) = \Omega(N^{2/3})$.

Such a condition on $\text{cc}(G)$ increasing with respect to N seems to be unavoidable in general. Consider a case when $\text{cc}(G) = cN$ for some constant c . Then, we need $m = (1/2) \log_\gamma(\delta c^2/2)$ to make the bias as small as we want, say $\delta/2$. Suppose the connected components are balanced such that $\omega_{\max} = O(1)$, then the variance term will be at most $\delta/2$, if $p = \Omega(m^\varepsilon/N)$, where ε depends on γ and c .

Note that the best known guarantees for estimators tailored for union of cliques still require $\text{cc}(G) = \Omega(N^{1-\varepsilon})$ for small but strictly positive p , where the ε can be made arbitrarily small with a small sampling probability p (e.g. [9, 15]).

We run synthetic experiments on a graph of $N = 3775$ nodes and union of 50 cliques, each of size $\{51, 52, \dots, 100\}$. Figure 5 shows that we improve upon three competing estimators for a broad range of p . $\widehat{\text{cc}}_{\text{chordal}}$ is the best known estimator for chordal graphs from [15], and $\widehat{\text{cc}}_{\text{clique}}$ is a smoothed version of $\widehat{\text{cc}}_{\text{chordal}}$ explicitly using the knowledge that the underlying graph is a union of cliques. These are tailored for chordal graphs and cliques, respectively, and cannot be applied to general graphs. Our generic algorithm, with an appropriate choices of α , β , and m , outperforms these approaches for unions of cliques. In particular, when p is small, variance dominates and choosing small m helps, whereas when p is large, bias dominates and choosing large m helps.

Theorem 4. *For any graph G with size of connected components ω_i , for each component $1 \leq i \leq \text{cc}(G)$, and degree of each node $d_j^{(i)}$, for $1 \leq j \leq \omega_i$ with $d_i \equiv \max_j \{d_j^{(i)}\}$, and $d_{\max} \equiv \max_{i,j} \{d_j^{(i)}\}$, $d_{\min} \equiv \min_{i,j} \{d_j^{(i)}\}$, for any choice of $\beta \geq 2d_{\max}$ and $\alpha \leq \sigma_{\min}(L)/\beta$, and any integer $m \geq 1$, there*

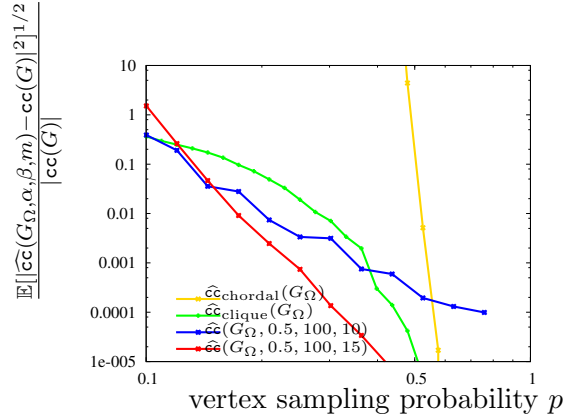


Figure 5: The proposed estimator for two choices of the degree m of the polynomial. With the right choice of m , we improve upon competing estimators when the original graph is a union of cliques.

exist a function $g(m) = O(2^{m^2})$ such that,

$$\frac{\mathbb{E}[(\widehat{\text{cc}}(G_\Omega, \alpha, \beta, m) - \text{cc}(G))^2]}{\text{cc}(G)^2} \leq \frac{g(m)(1-p^m)}{\text{cc}(G)^2 p \beta^2} \sum_{i=1}^{\text{cc}(G)} \left(\omega_i^2 d_i^2 \left(1 + (d_i p)^{1-2m} \right) \right) + \gamma^{2m} \frac{N^2}{\text{cc}(G)^2}, \quad (25)$$

where $\gamma = (1 - \alpha)/(1 + \alpha)$. Moreover, if there exist some positive constants c_i 's such that $d_i^3 p \geq c_i 2^{m^2}$ or $d_i p^3 \leq c_i / 2^{m^2}$ for all i , then (25) holds with $g(m) = O(c^m)$ for some constant $c > 0$.

A proof of Theorem 4 is provided in a longer version of this paper. This guarantee shows a similar bias-variance tradeoff, with similar dependence on m , which controls the computational complexity and α which is the normalized spectral gap of the original graph Laplacian. The main difference in this generic setting is how computational complexity depends on m . Since we need to estimate $\theta_w(G_\Omega)$ which is an unbiased estimate of $\prod_{u \in w} d_u^{s_u}$, we need to compute it separately for each observed walk w of length k that involves at least one self loop. For the other walks, we exploit a recent algorithm in counting patterns from [14] inspired by a celebrated result from [1], and compute their weighted counts. This can be made as a look-up table, and overall the complexity scales as $O(\text{cc}(G) \times \omega_{\max}^3)$ for $m \leq 7$ and for larger m scales as $O(\text{cc}(G) \times \omega_{\max}^{m/2} \times 2^{m/2})$. If one has faster algorithms for counting patterns those can be seamlessly included in the procedure, for example using recent advances in recursive methods for counting structures from [8]. Our code is publicly available at [url-anonymized](https://github.com/epi-anonymized).

We run experiments in Figure 6 on a graph of size $N = 5000$ with 50 components each drawn from Erdős-Rényi graph with probability half $G_{100,0.5}$. A moderate $m = 5$ is sufficient to achieve multiplicative error as small as 0.002, which implies that we make a small mistake in one out of ten instances. Note that $\widehat{\text{cc}}_{\text{chordal}}$ and $\widehat{\text{cc}}_{\text{clique}}$ cannot be applied as the observed subgraph is neither cliques nor chordal. A heuristic is proposed in [15], which is explained in Section 5.1 such that $\widehat{\text{cc}}_{\text{chordal}}$ can be applied. As the bias does not depend on p , this experiment implies that with only

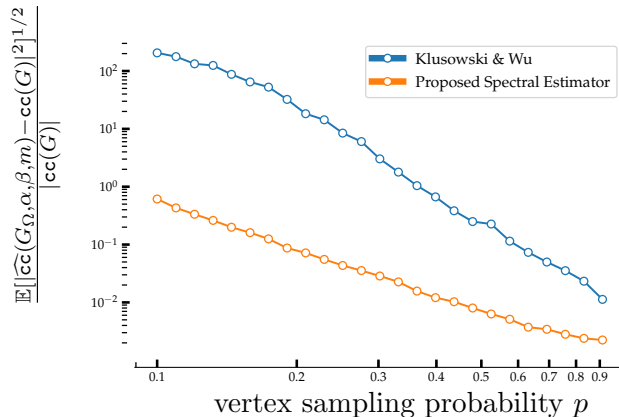


Figure 6: The proposed estimator $\widehat{cc}(G_{\Omega}, 0.5, 80, 5)$ improves upon a competing estimator from [15].

$m = 5$ the bias is already smaller than 0.002 and the variance is dominating. This is due to the fact that union of Erdős-Rényi graphs exhibit large spectral gaps. The variance decreases linearly in this log-log scale, with respect to the sampling probability p .

For the example of union of Erdős-Rényi graphs $G_{n,q}$ of the same size with $cc(G) = N/n$ connected components, the (normalized) spectral gap is $\Omega(1)$ for a large enough q . This exhibits the desired spectral gap, as long as q is sufficiently large, e.g. $q = \Omega(\log n/n)$. The ideal case is when $q = 1$, which recovers the union of cliques. The normalized spectral gap is one, which is the maximum possible value. On the other hand, the spectral gap can be also made arbitrarily small. Consider a union of n -cycles, where each component is a cycle of length n . In this case, the normalized spectral gap scales as $O(1/n^2)$, which can be quite small. For general graphs, the difficulty (both in computational complexity and sample complexity) depends on the spectral gap of the original graph. If spectral gap is small, then we need higher degree polynomial approximation functions to make the bias small, which in turn requires larger patterns to be counted. This increases the computational complexity and also the variance in the estimate. More samples are required to account for this increased variance.

5.1 Real-world graphs

We run our estimator on two real-world graphs. **ACL** [18] is an academic citation network that consists of 115,311 citations between 18,664 papers published in ACL (Association of Computational Linguistics) conferences, journals and workshops between years 1965 and 2016. **Hep-PH** [11] is an academic citation network comprising 412,533 citations between 34,546 Hep-PH (high energy physics phenomenology) arXiv manuscripts between years 1992 and 2002.

Estimating connected components on real graphs is challenging as they have large condition numbers and long cycles. In [15], this is dealt with by triangulating the real graph to turn it into a chordal graph before sampling. Alternatively, in this section, to make the estimation problem tractable, we apply the following two modifications to real-world graphs. First, we add random edges between nodes that belong to the same connected component to improve connectivity. If a component has m edges, we add extra qm edges at random. We choose $q = 0.4$ for **ACL** and $q = 0.6$ for **Hep-PH**. Secondly, we trim the degree of extremely high-degree nodes in the network. If a node

has degree larger than the 95 percentile of the degree distribution, we randomly remove its edges until it reaches the degree of the 95 percentile. For our experiments, we use the first 5000 nodes in each dataset, for computational efficiency. After modifying the graphs, the number of connected components in ACL and Hep-PH are 118 and 133.

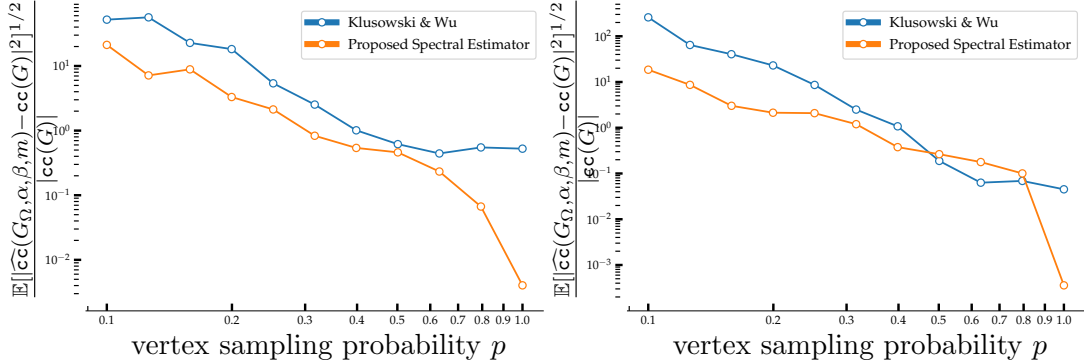


Figure 7: On (modified) real-world graphs, ACL (left) and Hep-PH (right), the proposed approach improves upon the competing baseline approach for most regime of the sampling probability p .

We compare the performance of the proposed algorithm on two real-world datasets to the algorithm proposed in [15]. As this estimator is customized for chordal graphs, we first triangulate the subsampled ACL and Hep-PH via a minimal chord completion algorithm based on maximum cardinality search [4] and then apply their smoothed estimator on subsampled graphs of the real-world networks, as proposed in [15]. Each data point in Figure 7 is averaged over 50 random instances of subsampling. Due to the triangulation, the approach of [15] is biased even when the sampling probability is close to one. For most of the values of p , the proposed estimator outperforms the baseline approach. We set $m = 5$ for the proposed spectral approach.

6 Conclusion

We address the problem of estimating the number of connected components in an undirected simple graph, when only a subgraph is observed, where the nodes are chosen uniformly at random. Existing methods relied on special structures of the graphs, such as union of disjoint cliques, union of disjoint trees, and chordal graphs. Applying a key insight of viewing the number of connected components as a spectral property that depends on the singular values of the graph Laplacian, we propose a novel spectral approach to this problem. Based on the fact that the number of connected components are the number of zero-valued singular values of the graph Laplacian, we make several innovations. First, we propose weighted count of small patterns (which are called network motifs) to estimate the k -th Schatten norm of the graph Laplacian. Next, to get an estimate of the (monomials of the) degrees, we propose a novel partitioning scheme that gives an unbiased estimate of the desired quantity to be used in the estimate of the k -th Schatten norms. We propose a polynomial approximation of the linearly interpolated step function, and prove an upper bound on the approximation guarantee. Putting these together, we introduce the first estimator with provable performance guarantees, that works for graphs with positive spectral gaps.

We next discuss several challenges in applying this framework to real world graphs.

Counting patterns. When the underlying graph G is a disjoint union of cliques, computational complexity of our estimator is $O(m \times \text{cc}(G))$. In this case, for any clique of size k , count of all possible patterns in it and the estimates of $\theta_w(G_\Omega)$ for any walk w , characterized by the degrees of the self loops it involves, can be pre-computed and stored in a table for look-up at the time of execution. $\theta_w(G_\Omega)$ is an unbiased estimate of the polynomial of the node degrees $\prod_{u \in w} d_u^{s_u}$. For general graphs, to compute $\widehat{\Theta}_k(G_\Omega)$, we need to compute $\theta_w(G_\Omega)$ separately for each observed walk w that has at least one self loop. For the walks that do not involve any self-loop, we can use matrix multiplication based pattern counting algorithms proposed in [14] for $m \leq 7$. For $m > 7$, one can use homomorphism based a recent recursive algorithm from [8]. Therefore, for general graphs the major computational complexity arises in computing $\theta_w(G_\Omega)$ for walks w that has at least one self-loop. In a different sampling scenario, if we have the additional information of the degree of each node that we observe then computing $\theta_w(G_\Omega)$ can be made fast for all the walks w . Another option is to apply recent advances in sampling-based methods for counting patterns, including wedge sampling [21], the 3-path sampling [13], Moss [22], GRAFT [20], and using Hamiltonian paths for debiasing [7]. However, it is not immediate how to include the estimation of the monomials of the degrees into these existing fast methods.

Other sampling techniques. In practical settings, sampling nodes uniformly at random might be unrealistic. Our estimator generalizes naturally to a broader class of sampling schemes, which we call *graph sampling*. Consider a scenario where you first sample an *unlabelled* mother graph H_0 of the same size as G (the graph of interest) from any distribution (in particular we do not require any independence on the sampled edges). Then, we apply a permutation drawn uniformly at random to assign node labels to the unlabelled graph H_0 . Let H denote this labelled graph, which we use to sample the original graph of interest G . Specifically, for all edges in H , we observe whether the corresponding edge is present or not in G . Namely, we observe the adjacency matrix of G , but masked by the adjacency matrix of H . The random permutation ensures that the sampling probability for a pattern only depends on the shape of the pattern and not the specific labels of the nodes involved, making our algorithm extendible up to properly applying the debiasing as per the new sampling model. The model studied in this paper is a special case of graph sampling where H_0 is a clique of random size, drawn according to a binomial distribution.

On the other hand, more practical sampling scenarios are adaptive to the topology of the graph, creating selection biases. Examples include crawling a connected path from a starting node, sampling higher degree nodes, or sampling via random walks. These create dependencies among the topology and the sampling, which we believe is outside the scope of this paper, but nevertheless poses an interesting new research direction.

7 Proofs

We provide the proof sketch of the main results. Recall $\widehat{\text{cc}}(G_\Omega, \alpha, \beta, m) = N - \sum_{k=1}^m \frac{a_k}{\beta^k} \widehat{\Theta}_k(G_\Omega)$, where $\widehat{\Theta}_k(G_\Omega)$ is an unbiased estimate of Schatten- k norm of L defined in (12) and a_k 's are coefficients of polynomial $f_b(x)$ defined in (22). We show in (28) that for the proposed estimator, bias is bounded as

$$\begin{aligned} |\text{cc}(G) - \mathbb{E}[\widehat{\text{cc}}(G_\Omega, \alpha, \beta, m)]| &\leq (N - \text{cc}(G))((1 - \alpha)/(1 + \alpha))^m \\ &= (N - \text{cc}(G))\gamma^m, \end{aligned} \tag{26}$$

where $\gamma \equiv (1 - \alpha)/(1 + \alpha)$. For the choice of $\tilde{b} \equiv (2/(1 + \alpha))[1, \dots, 1]$, the coefficient of x^k in $f_b(x)$ is bounded as $|a_k| \leq \binom{m}{k} 2^k \leq 2^m 2^k \leq 4^m$. Therefore the mean square error is bounded by

$$\begin{aligned} \mathbb{E}[(\widehat{\text{cc}}(G_\Omega, \alpha, \beta, m) - \text{cc}(G))^2] &= \\ 2^{4m} \text{Var} \left(\sum_{k=1}^m \frac{1}{\beta^k} \widehat{\Theta}_k(G_\Omega) \right) &+ \left((N - \text{cc}(G)) \gamma^m \right)^2. \end{aligned} \quad (27)$$

Since the Schatten norm estimator is unbiased, $\mathbb{E}[\widehat{\Theta}_k(G_\Omega)] = \|L\|_k^k$, we have $\mathbb{E}[\widehat{\text{cc}}(G_\Omega, \alpha, \beta, m)] = N - \sum_{i=1}^N f_\alpha(\sigma_i(\tilde{L}))$, where f_α is defined in Equation (4). Note that β is chosen such that the non-zero eigenvalues of $\tilde{L} = L/\beta$ are bounded between α and 1. With the proposed choice of polynomial function f_α , we have $f_\alpha = f_b$. Using, Equation (3), $\text{cc}(G) = N - \sum_{i=1}^N H_\alpha(\sigma_i(\tilde{L}))$, along with $\max_{x \in [\alpha, 1]} |H_\alpha(x) - f_b(x)| = ((1 - \alpha)/(1 + \alpha))^m$, where $\tilde{b} \equiv (2/(1 + \alpha))[1, \dots, 1]$, we have,

$$\begin{aligned} \left| \text{cc}(G) - \mathbb{E}[\widehat{\text{cc}}(G_\Omega, \alpha, \beta, m)] \right| &= \sum_{i=1}^N \left(H_\alpha(\sigma_i(\tilde{L})) - f_b(x) \right) \\ &= \sum_{i: \sigma_i(\tilde{L})=0} \left(H_\alpha(\sigma_i(\tilde{L})) - f_b(x) \right) + \sum_{i: \sigma_i(\tilde{L}) \neq 0} \left(H_\alpha(\sigma_i(\tilde{L})) - f_b(x) \right) \\ &\leq \text{cc}(G) \left(H_\alpha(0) - f_b(0) \right) + (N - \text{cc}(G)) \max_{x \in [\alpha, 1]} |H_\alpha(x) - f_b(x)| \\ &\leq (N - \text{cc}(G)) ((1 - \alpha)/(1 + \alpha))^m = (N - \text{cc}(G)) \gamma^m. \end{aligned} \quad (28)$$

For the two cases: (a) when the underlying graph G is disjoint union of cliques, and (b) a general graph G with maximum degree d_{\max} , we provide bounds on the variance of the Schatten k -norm estimator that leads to the bounds on mean square error using Equation (27).

Denote each connected component of G by $G^{(i)}$ for $1 \leq i \leq \text{cc}(G)$, and let $G_\Omega^{(i)}$ denote the randomly observed subgraph of the connected component $G^{(i)}$. Then, we have,

$$\text{Var} \left(\sum_{k=1}^m \frac{1}{\beta^k} \widehat{\Theta}_k(G_\Omega) \right) = \sum_{i=1}^{\text{cc}(G)} \text{Var} \left(\sum_{k=1}^m \frac{1}{\beta^k} \widehat{\Theta}_k(G_\Omega^{(i)}) \right). \quad (29)$$

Note that, our estimator of Schatten k -norm naturally decomposes, and can be computed separately for each connected component $G^{(i)}$ and then added together to get the estimate for the graph G .

7.1 Proof of Theorem 3

The following lemma provides bound on the variance of Schatten k -norm estimator for a clique graph. We give a proof in Section A.1.

Lemma 1. *For a clique graph G on ω vertices, there exists a universal positive constant C such that for $\omega \geq C$, variance of Schatten k -norm estimator $\widehat{\Theta}_k(G_\Omega)$ is bounded by*

$$\text{Var}(\widehat{\Theta}_k(G_\Omega)) \leq g(k) \frac{\omega^{2k+2}}{p} \left(1 + \frac{1}{(\omega p)^{2k-1}} \right), \quad (30)$$

where $g(k) = O(k!)$. Moreover, if there exists a positive constant c such that $\omega^3 p \geq ck!$ or $\omega p^3 \leq c/k!$ then (30) holds with $g(k) = \text{poly}(k)$.

Using Equations (27) and (29) along with Lemma 1, and the fact that $\beta \geq \omega$, Theorem 3 follows immediately. For a proof of this lemma, we refer to the longer version of this paper.

References

- [1] N. Alon, R. Yuster, and U. Zwick. Finding and counting given length cycles. *Algorithmica*, 17(3):209–223, 1997.
- [2] P. Berenbrink, B. Krayenhoff, and F. Mallmann-Trenn. Estimating the number of connected components in sublinear time. *Information Processing Letters*, 114(11):639–642, 2014.
- [3] Daniel Berend and Tamir Tassa. Improved bounds on bell numbers and on moments of sums of random variables. *Probability and Mathematical Statistics*, 30(2):185–205, 2010.
- [4] Anne Berry, Jean RS Blair, Pinar Heggernes, and Barry W Peyton. Maximum cardinality search for computing minimal triangulations of graphs. *Algorithmica*, 39(4):287–298, 2004.
- [5] B. Chazelle, R. Rubinfeld, and L. Trevisan. Approximating the minimum spanning tree weight in sublinear time. *SIAM Journal on computing*, 34(6):1370–1379, 2005.
- [6] Pafnuti Lvovich Chebyshev. *Théorie of the mechanisms known as parallel e logrammes*. Critical Academy of Science Printing, 1853.
- [7] Xiaowei Chen and John CS Lui. Mining graphlet counts in online social networks. In *Data Mining (ICDM), 2016 IEEE 16th International Conference on*, pages 71–80. IEEE, 2016.
- [8] Radu Curticapean, Holger Dell, and Dániel Marx. Homomorphisms are a good basis for counting small subgraphs. *arXiv preprint arXiv:1705.01595*, 2017.
- [9] O. Frank. Estimation of the number of connected components in a graph by using a sampled subgraph. *Scandinavian Journal of Statistics*, pages 177–188, 1978.
- [10] O. Frank and F. Harary. Cluster inference by using transitivity indices in empirical graphs. *Journal of the American Statistical Association*, 77(380):835–840, 1982.
- [11] Johannes Gehrke, Paul Ginsparg, and Jon Kleinberg. Overview of the 2003 kdd cup. *ACM SIGKDD Explorations Newsletter*, 5(2):149–151, 2003.
- [12] L. A. Goodman. On the estimation of the number of classes in a population. *The Annals of Mathematical Statistics*, pages 572–579, 1949.
- [13] Madhav Jha, C Seshadhri, and Ali Pinar. Path sampling: A fast and provable method for estimating 4-vertex subgraph counts. In *Proceedings of the 24th International Conference on World Wide Web*, pages 495–505. International World Wide Web Conferences Steering Committee, 2015.
- [14] A. Khetan and S. Oh. Spectrum estimation from a few entries. *arXiv preprint arXiv:1703.06327*, 2017.

- [15] J. M. Klusowski and Y. Wu. Estimating the number of connected components in a graph via subgraph sampling. *Technical report*, 2017. available at <http://www.stat.yale.edu/~yw562/preprints/cc.pdf>.
- [16] John C Mason and David C Handscomb. *Chebyshev polynomials*. CRC Press, 2002.
- [17] T Parks and J McClellan. Chebyshev approximation for nonrecursive digital filters with linear phase. *IEEE Transactions on Circuit Theory*, 19(2):189–194, 1972.
- [18] Dragomir R. Radev, Pradeep Muthukrishnan, Vahed Qazvinian, and Amjad Abu-Jbara. The acl anthology network corpus. *Language Resources and Evaluation*, pages 1–26, 2013.
- [19] D. Rafiei. Effectively visualizing large networks through sampling. In *Visualization, 2005. VIS 05. IEEE*, pages 375–382. IEEE, 2005.
- [20] Mahmudur Rahman, Mansurul Alam Bhuiyan, and Mohammad Al Hasan. Graft: An efficient graphlet counting method for large graph analysis. *IEEE Transactions on Knowledge and Data Engineering*, 26(10):2466–2478, 2014.
- [21] Comandur Seshadhri, Ali Pinar, and Tamara G Kolda. Triadic measures on graphs: The power of wedge sampling. In *Proceedings of the 2013 SIAM International Conference on Data Mining*, pages 10–18. SIAM, 2013.
- [22] Pinghui Wang, Jing Tao, Junzhou Zhao, and Xiaohong Guan. Moss: A scalable tool for efficiently sampling and counting 4-and 5-node graphlets. *arXiv preprint arXiv:1509.08089*, 2015.
- [23] Y. Zhang, M. J. Wainwright, and M. I. Jordan. Distributed estimation of generalized matrix rank: Efficient algorithms and lower bounds. *arXiv preprint arXiv:1502.01403*, 2015.

Appendix

A Proofs

A.1 Proof of Lemma 1




Figure 8: For a set of cliques G , we get a sampled G_Ω . An example of a length- $(k = 6)$ closed walk $w = (1, 2, 2, 2, 3, 3, 1)$ and its corresponding k -cyclic pseudograph $\mathcal{H}(w) \in \mathbb{H}_6$.

For a k -cyclic pseudograph $H = (V_H, E_H)$, let S_H denote the set of self-loops in H . Recall that in (12) Schatten k -norm estimator of Laplacian L is given by

$$\hat{\Theta}_k(G_\Omega) = \sum_{H \in \mathbb{H}_k} \frac{(-1)^{k-|S_H|}}{p^{|V_H|}} \left\{ \sum_{w: \mathcal{H}(w)=H} \theta(w, G_\Omega) \mathbb{I}(w \subseteq G_\Omega) \right\}.$$

For a clique graph G , the analysis of the above estimator simplifies significantly. We illustrate this with an example in Figure 8. Consider a length $k = 6$ walk $w = (1, 2, 2, 2, 3, 3, 1)$ with a corresponding k -cyclic pseudograph $\mathcal{H}(w)$. In general, the degree estimator $\theta(w, G_\Omega)$ is chosen such that $\mathbb{E}[\theta(w, G_\Omega)] = d_2^2 d_3$ where d_2 and d_3 are the degrees of nodes 2 and 3 in G , respectively. This simplifies significantly for a clique graph due to the fact that the degree of those nodes in a closed walk w are the same. Note that our estimator is general, and does not use this information or the fact that the underlying graph component is a clique. It is only the analysis that simplifies. Therefore, for a clique graph G , the degree estimator $\theta(w, G_\Omega)$ satisfies $\mathbb{E}[\theta(w, G_\Omega) | w \subseteq G_\Omega] = (\omega - 1)^{|S_H|}$, where ω is the size of the clique G . In the example, we have $\omega = 7$ and $|S_H| = 3$, therefore $\mathbb{E}[\theta(w, G_\Omega)] = 6^3$. Hence, it is best to further partition \mathbb{H}_k according to the number of nodes $\ell = |V_H|$ and the number of self-loops $s = |S_H|$. Precisely, we define

$$\mathbb{H}_{k,\ell,s} \equiv \{H(V_H, E_H) \in \mathbb{H}_k : |V_H| = \ell \text{ and } |S_H| = s\},$$

for $\ell = 1, s = k$ and $2 \leq \ell \leq k, 0 \leq s \leq k - \ell$. There are total $|\{w \in W : \mathcal{H}(w) \in \mathbb{H}_{k,\ell,s}\}| \leq \ell^{2k-s} \omega^\ell$ corresponding walks in this set. Here, W denotes the collection of all length k closed walks on a complete graph of ω vertices. We slightly overload the notion of complete graph to refer to an

undirected graph with not only all the $\omega(\omega - 1)/2$ simple edges but also with ω self loops as well. when G is a clique graph, the estimator (12) can be re-written as

$$\widehat{\Theta}_k(G_\Omega) = \sum_{\ell=1}^k \sum_{s=0}^{k-\ell+1} \sum_{H \in \mathbb{H}_{k,\ell,s}} \left\{ \frac{(-1)^{k-s}}{p^\ell} \times \sum_{w \in W: \mathcal{H}(w)=H} \theta(w, G_\Omega) \mathbb{I}(w \subseteq G_\Omega) \right\}. \quad (31)$$

Given this unbiased estimator, we provide an upper bound on the variance of each of the partitions. For any two walks $w, w' \in W$, let $|w \cap w'|$ denote the number of overlapping unique vertices of walks w and w' . We have,

$$\begin{aligned} \text{Var}(\widehat{\Theta}_k(G_\Omega)) &= \sum_{\ell=1}^k \sum_{s=0}^{k-\ell+1} \sum_{H \in \mathbb{H}_{k,\ell,s}} \left\{ \frac{1}{p^{2\ell}} \right. \\ &\quad \times \left. \sum_{w \in W: \mathcal{H}(w)=H} \text{Var}(\theta(w, G_\Omega) \mathbb{I}(w \subseteq G_\Omega)) \right\} \\ &+ 2 \sum_{\tilde{\ell}=0}^k \sum_{\substack{w, w' \in W \\ |w \cap w'| = \tilde{\ell}}} \text{Cov}(\theta(w, G_\Omega) \mathbb{I}(w \subseteq G_\Omega), \theta(w', G_\Omega) \mathbb{I}(w' \subseteq G_\Omega)) \\ &\quad \times \frac{(-1)^{|S_{\mathcal{H}(w)}| + |S_{\mathcal{H}(w')}|}}{p^{|V_{\mathcal{H}(w)}| + |V_{\mathcal{H}(w')}|}}. \end{aligned} \quad (32)$$

The following technical lemma provides upper bounds on the variance and covariance terms. We provide a proof in Section A.2.

Lemma 2. *Under the hypothesis of Lemma 1, for a length- k walk w over ℓ distinct nodes with $s \geq 1$ self-loops, the following holds:*

$$\text{Var}(\theta(w, G_\Omega) \mathbb{I}(w \subseteq G_\Omega)) \leq f(\ell, s) \left(p^{\ell-1} \omega^{2s-1} + \omega p^{\ell+1-2s} \right) + p^\ell \omega^{2s}, \quad (33)$$

and when $\ell = 1, s = k$, we have,

$$\text{Var}(\theta(w, G_\Omega) \mathbb{I}(w \subseteq G_\Omega)) \leq f(k) \omega^{2k-1} + g(k) \omega p^{2-2k} + p \omega^{2k}, \quad (34)$$

and for any length- k walks w_1, w_2 over ℓ_1, ℓ_2 distinct nodes with $\tilde{\ell}$ unique overlapping nodes, $|w_1 \cap w_2| = \tilde{\ell}$, $s_1, s_2 \geq 1$ self-loops respectively, the covariance term can be upper bounded by:

$$\begin{aligned} &\text{Cov}(\theta(w_1, G_\Omega) \mathbb{I}(w_1 \subseteq G_\Omega), \theta(w_2, G_\Omega) \mathbb{I}(w_2 \subseteq G_\Omega)) \\ &\leq f(\ell', s') p^{((\ell_1 + \ell_2 - \tilde{\ell}) - (s_1 + s_2))} \left((\omega p)^{(s_1 + s_2)} + \omega p \right), \end{aligned} \quad (35)$$

for some function $f(\ell, s) = O(k!)$, $g(k) = \text{poly}(k)$, where p is the vertex sampling probability. $\ell' \equiv \max\{\ell_1, \ell_2\}$ and $s' \equiv \max\{s_1, s_2\}$.

We use this lemma to get bound on $\text{Var}(\widehat{\Theta}_k(G_\Omega))$. First, we get a bound on the total variance term. For a walk $w \in W$ with $\mathcal{H}(w) \in \mathbb{H}_{k,\ell,s}$ with $1 \leq s \leq k-2$, using (33), we have,

$$\begin{aligned} & \frac{\omega^\ell}{p^{2\ell}} \text{Var}\left(\theta(w, G_\Omega) \mathbb{I}(w \subseteq G_\Omega)\right) \\ & \leq f(\ell, s) \left(\frac{\omega^{\ell+2s-1}}{p^{\ell+1}} + \frac{\omega^{\ell+1}}{p^{\ell+2s-1}} \right) + \frac{\omega^{2s+\ell}}{p^\ell} \\ & \leq f(k) \left(\frac{\omega^{2k-3}}{p^3} + \frac{\omega^3}{p^{2k-3}} \right) + \frac{\omega^{2k-2}}{p^2}. \end{aligned} \quad (36)$$

For a walk $w \in W$ with $\mathcal{H}(w) \in \mathbb{H}_{k,\ell,s}$ with $\ell = 1, s = k$, using (34), we have,

$$\begin{aligned} & \frac{\omega^\ell}{p^{2\ell}} \text{Var}\left(\theta(w, G_\Omega) \mathbb{I}(w \subseteq G_\Omega)\right) \\ & \leq f(k) \omega^{2k} p^{-2} + g(k) \omega^2 p^{-2k} + w^{2k+1} p^{-1}. \end{aligned} \quad (37)$$

For a walk w with $s = 0$, $\theta(w, G_\Omega) = 1$, and, we have,

$$\frac{\omega^\ell}{p^{2\ell}} \text{Var}\left(\theta(w, G_\Omega) \mathbb{I}(w \subseteq G_\Omega)\right) \leq \frac{\omega^\ell}{p^\ell}. \quad (38)$$

Combining, Equations (36), (37), and (38), and using $|\{w \in W : \mathcal{H}(w) \in \mathbb{H}_{k,\ell,s}\}| \leq \ell^{2k-s} \omega^\ell$, we have

$$\begin{aligned} & \sum_{\ell=1}^k \sum_{s=0}^{k-\ell+1} \sum_{H \in \mathbb{H}_{k,\ell,s}} \left\{ \frac{1}{p^{2\ell}} \sum_{w \in W: \mathcal{H}(w)=H} \mathbb{I}(w \subseteq G) \times \right. \\ & \quad \left. \text{Var}\left(\theta(w, G_\Omega) \mathbb{I}(w \subseteq G_\Omega)\right) \right\} \\ & \leq f(k) \omega^2 p^{-2k} + \omega^{2k+1} p^{-1}, \end{aligned} \quad (39)$$

and if $\omega p^3 \leq 1/f(k)$, then the above quantity is bounded by $g(k) \omega^2 p^{-2k} (1 + o(1))$. If $\omega^3 p \geq f(k)$, then the above quantity is bounded by $\omega^{2k+1} p^{-1} (1 + o(1))$.

Consider covariance term of two length- k walks w_1, w_2 over ℓ_1, ℓ_2 distinct nodes with $\tilde{\ell}$ unique overlapping nodes, $|w_1 \cap w_2| = \tilde{\ell}$, and $s_1, s_2 \geq 1$ self-loops with $s_1 + s_2 < 2k$. Since there are a total of $f(k) \omega^{\ell_1 + \ell_2 - \tilde{\ell}}$ such walks, using (35), we have

$$\begin{aligned} & \omega^{\ell_1 + \ell_2 - \tilde{\ell}} \text{Cov}\left(\theta(w_1, G_\Omega) \mathbb{I}(w_1 \subseteq G_\Omega), \theta(w_2, G_\Omega) \mathbb{I}(w_2 \subseteq G_\Omega)\right) \\ & \times \frac{(-1)^{|S_{\mathcal{H}(w_1)}| + |S_{\mathcal{H}(w_2)}|}}{p^{|V_{\mathcal{H}(w_1)}| + |V_{\mathcal{H}(w_2)}|}} \end{aligned} \quad (40)$$

$$\leq \frac{\omega^{\ell_1 + \ell_2 - \tilde{\ell}}}{p^{\tilde{\ell} + s_1 + s_2}} \left((\omega p)^{s_1 + s_2} + \omega p \right) \quad (41)$$

$$\leq \frac{\omega^{\ell_1 + \ell_2 + s_1 + s_2}}{(\omega p)^{\tilde{\ell}}} + \omega p \frac{\omega^{\ell_1 + \ell_2 - \tilde{\ell}}}{p^{s_1 + s_2 + \tilde{\ell}}} \quad (42)$$

$$\leq \omega^{2k+1} + \omega p \frac{\omega^2}{p^{2k-1}} \quad (43)$$

where in (42), for the first term the maximum is achieved at $\ell_1 = 1, s = k, \ell_2 = 2, s_2 = k - 2, \tilde{\ell} = 0$, for the second term the maximum is achieved at same ℓ_1, s_1, ℓ_2, s_2 with $\tilde{\ell} = 1$. Note that in the expression (41), the first term is dominating when $\omega p \geq 1$, and the second term is dominating when $\omega p < 1$.

When $s_1 + s_2 = 2k$ the two walks are self loop walks on single node with $s_1 = s_2 = k$. Since the graph G is a clique graph, $\theta(w, G_\Omega)$ for self loop walks depends only upon observed size of the clique, and $\theta(w_1, G_\Omega) = \theta(w_2, G_\Omega)$. Therefore, the expression in (40) can be bounded as:

$$\begin{aligned} & \omega^{\ell_1 + \ell_2 - \tilde{\ell}} \text{Cov}\left(\theta(w_1, G_\Omega)\mathbb{I}(w_1 \subseteq G_\Omega), \theta(w_2, G_\Omega)\mathbb{I}(w_2 \subseteq G_\Omega)\right) \\ & \quad \times \frac{(-1)^{|S_{\mathcal{H}(w_1)}| + |S_{\mathcal{H}(w_2)}|}}{p^{|V_{\mathcal{H}(w_1)}| + |V_{\mathcal{H}(w_2)}|}} \\ & \leq \frac{\omega^{2\ell}}{p^{2\ell}} \text{Var}\left(\theta(w, G_\Omega)\mathbb{I}(w \subseteq G_\Omega)\right) \\ & \leq f(k)\omega^{2k+1}p^{-2} + g(k)\omega^3p^{-2k} + w^{2k+2}p^{-1}, \end{aligned} \quad (44)$$

where (44) follows from (37).

Lemma 1 follows immediately by combining (32) with (66), (43) and (44).

A.2 Proof of Lemma 2

We use the following technical lemma to get bounds on conditional variance and covariance of the estimator θ . We provide a proof in Section A.3.

Lemma 3. *Under the hypothesis of Lemma 1, for length- k walks w_1, w_2 over ℓ_1, ℓ_2 distinct nodes with $s_1, s_2 \geq 1$ self-loops respectively, the conditional variance of estimator $\theta(w_1, G_\Omega)$, defined in (21), given that all the nodes in the walk are sampled can be upper bounded by*

$$\begin{aligned} & \text{Var}\left(\theta(w_1, G_\Omega) \middle| w_1 \subseteq G_\Omega\right) \\ & \leq f(\ell_1, s_1) \left(p^{-1}(\omega - \ell_1)^{2s_1 - 1} + (\omega - \ell_1)p^{1 - 2s_1} \right), \text{ and} \end{aligned} \quad (45)$$

$$\begin{aligned} & \mathbb{E}\left[\theta(w_1, G_\Omega)\theta(w_2, G_\Omega) \middle| \mathbb{I}(w_1 \subseteq G_\Omega)\mathbb{I}(w_2 \subseteq G_\Omega)\right] \\ & \leq f(\ell', s')p^{-(s_1 + s_2)} \left((\omega p)^{s_1 + s_2} + \omega p \right), \end{aligned} \quad (46)$$

for some function $f(\ell, s) = O(k!)$, where p is the vertex sampling probability. $\ell' \equiv \max\{\ell_1, \ell_2\}$ and $s' \equiv \max\{s_1, s_2\}$. Moreover, for a length k walk w with $\ell = 1$, and $s = k$,

$$\text{Var}\left(\theta(w, G_\Omega) \middle| w \subseteq G_\Omega\right) \leq \frac{f(k)(\omega - 1)^{2k - 1}}{p} + \frac{g(k)(\omega - 1)}{p^{2k - 1}}, \quad (47)$$

for some function $g(k) = \text{poly}(k)$.

Recall for a clique graph, we have $\mathbb{E}[\theta(w, G_\Omega) | w \subseteq G_\Omega] = (\omega - 1)^s$. Therefore, we have,

$$\begin{aligned}
& \text{Var}\left(\theta(w, G_\Omega)\mathbb{I}(w \subseteq G_\Omega)\right) \\
&= \mathbb{E}\left[\text{Var}\left(\theta(w, G_\Omega)\mathbb{I}(w \subseteq G_\Omega) \middle| \mathbb{I}(w \subseteq G_\Omega)\right)\right] \\
&\quad + \text{Var}\left(\mathbb{E}\left[\theta(w, G_\Omega)\mathbb{I}(w \subseteq G_\Omega) \middle| \mathbb{I}(w \subseteq G_\Omega)\right]\right) \\
&= p^\ell \text{Var}\left(\theta(w, G_\Omega) \middle| w \subseteq G_\Omega\right) + p^\ell(1 - p^\ell)(\omega - 1)^{2s} \\
&\leq f(\ell, s)\left(p^{\ell-1}(\omega - \ell)^{2s-1} + (\omega - \ell)p^{\ell+1-2s}\right) \\
&\quad + p^\ell(1 - p^\ell)(\omega - 1)^{2s} \\
&\leq f(\ell, s)\left(p^{\ell-1}\omega^{2s-1} + \omega p^{\ell+1-2s}\right) + p^\ell\omega^{2s}, \tag{48}
\end{aligned}$$

where the inequality follows from Equation (45). Similarly, for a walk w with $\ell = 1$, and $s = k$, we have,

$$\text{Var}\left(\theta(w, G_\Omega)\mathbb{I}(w \subseteq G_\Omega)\right) \leq f(k)\omega^{2k-1} + g(k)\omega p^{2-2k} + p\omega^{2k},$$

where we used the inequality in (47).

For covariance term, we have,

$$\begin{aligned}
& \text{Cov}\left(\theta(w_1, G_\Omega)\mathbb{I}(w_1 \subseteq G_\Omega), \theta(w_2, G_\Omega)\mathbb{I}(w_2 \subseteq G_\Omega)\right) \\
&\leq \mathbb{E}\left[\theta(w_1, G_\Omega)\theta(w_2, G_\Omega) \middle| \mathbb{I}(w_1 \subseteq G_\Omega)\mathbb{I}(w_2 \subseteq G_\Omega)\right] p^{(\ell_1+\ell_2-\bar{\ell})} \\
&\leq f(\ell', s')p^{((\ell_1+\ell_2-\bar{\ell})-(s_1+s_2))}\left((\omega p)^{(s_1+s_2)} + \omega p\right), \tag{49}
\end{aligned}$$

where the inequality follows from Equation (46).

A.3 Proof of Lemma 3

When G is a union of disjoint cliques, the estimator $\theta(w, G_\Omega)$ defined in (21) has a compact representation. This follows from the fact that for any two nodes i and j that are connected in G_Ω , the neighborhoods of i and j in G_Ω exactly coincide. If this happens, then the estimator $\theta(w, G_\Omega)$ simplifies as follows. Consider a walk w with s self-loops, k edges (including self loops), and ℓ distinct nodes. Define a random integer $\tilde{\tau}$ as the degree of a node in the clique that w belongs to in the sampled graph G_Ω , conditioned on the fact that all nodes in w are sampled (if a walk w is sampled then it must belong to the same clique.). The randomness comes from the sampling of Ω . It is straightforward that $\tilde{\tau} \sim \text{Binom}(\omega - \ell, p) + (\ell - 1)$ as there are already $(\ell - 1)$ neighbors from the walk w and the rest of $((\omega - 1) - (\ell - 1))$ nodes are sampled in Ω with probability p , where ω is the size of the clique in the original graph G . For notational simplification, define a random integer $\tau \equiv \tilde{\tau} - (\ell - 1)$ that is distributed as $\tau \sim \text{Binom}(\omega - \ell, p)$. In the example in Figure 8, for the walk $w = (1, 2, 2, 2, 3, 3, 1)$, we have $\ell = 3$, $s = 3$, $\omega = 7$ and a random instance of $\tilde{\tau} = 4$ that is $\tau = 2$. We claim that the estimator $\theta(w, G_\Omega)$ given in (21) is a function of only τ , ℓ , s and p , and can be simplified as follows, and give a proof in Section A.4.

Lemma 4. *When the underlying G is a union of disjoint cliques and G_Ω is a subgraph obtained via vertex sampling with a probability p , for a length k walk w in G_Ω with ℓ distinct nodes and s self-loops, we have*

$$\theta(w, G_\Omega) = \langle A_{s+1}^{-1}, \bar{\tau} \rangle, \quad (50)$$

where $\tau \equiv \tilde{\tau} - (\ell - 1)$ and $\tilde{\tau}$ is the degree of any node in the clique that w belongs to in the sampled graph G_Ω , $\bar{\tau} = [1, \tau, \tau^2, \dots, \tau^s]$ is a column vector of monomials of τ up to degree s , and A_{s+1}^{-1} is the $(s+1)$ -th row of the inverse of the matrix $A \in \mathbb{R}^{(s+1) \times (s+1)}$ satisfying $A\bar{\omega} = \mathbb{E}[\bar{\tau}]$, for $\bar{\omega} = [1, (\omega - 1), (\omega - 1)^2, \dots, (\omega - 1)^s]$, a column vector of monomials of $(\omega - 1)$. Further, A is a lower-triangular matrix that depends only on ℓ and p , such that $\max_{i \in [s+1]} |(A^{-1})_{s+1, i}| = O(p^{-s})$.

In the running example, $s = 3$, $\tilde{\tau} = 4$, and $\ell = 3$, and therefore we have

$$A = \begin{bmatrix} 1 & 0 & 0 & 0 \\ p - \ell p & p & 0 & 0 \\ \ell^2 p^2 - \ell p^2 - \ell p + p & -2\ell p^2 + p^2 + p & p^2 & 0 \\ A_{41} & A_{42} & -3\ell p^3 + 3p^3 & p^3 \end{bmatrix}, \quad (51)$$

where $A_{41} = -\ell^3 p^3 + 3\ell^2 p^2 + \ell p^3 - 3\ell p^2 - \ell p + p$ and $A_{42} = 3\ell^2 p^3 - 6\ell p^2 - p^3 + 3p^2 + p$. For any s , corresponding A can be computed immediately from the moments of a Binomial distribution up to degree s . Since $\bar{\omega} = \mathbb{E}[A^{-1}\bar{\tau}]$, this representation immediately reveals that $\mathbb{E}[\theta(w, G_\Omega) | w \subseteq G_\Omega] = \mathbb{E}[\langle A_{s+1}^{-1}, \bar{\tau} \rangle] = (\omega - 1)^s$. Note that τ is conditioned on the event that all the nodes in w are sampled.

With this definition, the variance of $\theta(w, G_\Omega)$ can be upper bounded as follows. We will let $f(\ell, s)$ denote a function over ℓ and s that captures the dependence in ℓ and s that may change from line to line, and only track the dependence in ω and p .

$$\begin{aligned} \text{Var}\left(\theta(w, G_\Omega) \middle| w \subseteq G_\Omega\right) &\leq f(\ell, s) \max_{i \in [s+1]} \text{Var}((A^{-1})_{s+1, i} \tau^{i-1}) \\ &\leq f(\ell, s) p^{-2s} \text{Var}(\tau^s) \\ &= f(\ell, s) p^{-2s} \left(\mathbb{E}[\tau^{2s}] - \mathbb{E}[\tau^s]^2 \right) \\ &\leq f(\ell, s) p^{-2s} \left((\omega - \ell)^{2s} p^{2s} + f(s) (\omega - \ell)^{2s-1} p^{2s-1} \right. \\ &\quad \left. + (\omega - \ell)p - (\omega - \ell)^{2s} p^{2s} \right) \\ &\leq f(\ell, s) \left(p^{-1} (\omega - \ell)^{2s-1} + (\omega - \ell) p^{1-2s} \right), \end{aligned} \quad (52)$$

where the first inequality follows from the fact that $\bar{\tau}$ is an $s+1$ dimensional vector, the second inequality follows from that fact that $\max_i |(A^{-1})_{s+1, i}| = O(p^{-s})$ from Lemma 4 and $\max_i = \text{Var}(\tau^{i-1}) = \text{Var}(\tau^s)$, and in the third inequality we used the fact that $\tau \sim \text{Binom}(\omega - \ell, p)$ and a result from [3] that $\mathbb{E}[(\text{Binom}(d, p))^s] \leq \sum_{j=1}^s S(s, j) (dp)^j$ where $S(s, j)$ is the Sterling number of second kind. $S(s, s) = 1$, $S(s, 1) = 1$ and $S(s, j) \leq f(s)$, for $2 \leq j \leq s-1$. We also used Jensen's inequality $\mathbb{E}[(\text{Binom}(d, p))^s] \geq \mathbb{E}[(\text{Binom}(d, p))]^s$. This proves the desired bound in (33).

To prove the bound in (47), observe that when $\ell = 1$, $\tau \sim \text{Binom}(\omega - 1, p)$, and we can tighten the above set of inequalities

$$\begin{aligned}
\text{Var}\left(\theta(w, G_\Omega) \middle| w \subseteq G_\Omega\right) &\leq g(k) \max_{i \in [k+1]} \text{Var}((A^{-1})_{k+1, i} \tau^{i-1}) \\
&\leq g(k) p^{-2k} \text{Var}(\tau^k) \\
&= g(k) p^{-2k} \left(\mathbb{E}[\tau^{2k}] - \mathbb{E}[\tau^k]^2 \right) \\
&\leq g(k) p^{-2k} \left((\omega - 1)^{2k} p^{2k} + f(k) (\omega - 1)^{2k-1} p^{2k-1} \right. \\
&\quad \left. + (\omega - 1)p - (\omega - 1)^{2k} p^{2k} \right) \\
&\leq f(k) p^{-1} (\omega - 1)^{2k-1} + g(k) (\omega - 1) p^{1-2k}. \tag{53}
\end{aligned}$$

For the covariance term, conditioned on the event that both the walks w_1 and w_2 are observed, distribution of the random degree integer of each walk is $\tilde{\tau}_1 = \tilde{\tau}_2 = \tilde{\tau}_{12}$, where $\tilde{\tau}_{12} \sim \text{Binom}(\omega - (\ell_1 + \ell_2 - \tilde{\ell}), p) + (\ell_1 + \ell_2 - \tilde{\ell} - 1)$. Therefore, the shifted Binomial random variable of each walk $\tau_1 = \tilde{\tau}_1 - (\ell_1 - 1) = \tilde{\tau}_{12} - (\ell_1 - 1)$, and $\tau_2 = \tilde{\tau}_2 - (\ell_2 - 1) = \tilde{\tau}_{12} - (\ell_2 - 1)$. For the walk w_1 , with number of nodes ℓ_1 , self loops s_1 , lets denote the matrix A in (50) by A_1 , and similarly for walk w_2 , denote it by A_2 . Then we have,

$$\begin{aligned}
&\mathbb{E}\left[\theta(w_1, G_\Omega)\theta(w_2, G_\Omega) \middle| \mathbb{I}(w_1 \subseteq G_\Omega)\mathbb{I}(w_2 \subseteq G_\Omega)\right] \\
&\leq f(\ell', s') \max_{i_1 \in [s_1+1], i_2 \in [s_2+1]} \mathbb{E}\left[(A_1^{-1})_{s_1+1, i_1} \tau_1^{i_1-1} (A_2^{-1})_{s_2+1, i_2} \tau_2^{i_2-1}\right] \tag{54}
\end{aligned}$$

$$\leq f(\ell', s') p^{-(s_1+s_2)} \mathbb{E}\left[\tau_1^{s_1} \tau_1^{s_2}\right] \tag{55}$$

$$\leq f(\ell', s') p^{-(s_1+s_2)} \left((\omega p)^{s_1+s_2} + \omega p \right), \tag{56}$$

where the first inequality follows from the fact that $\bar{\tau}_1$ is an $s_1 + 1$ dimensional vector, and $\bar{\tau}_2$ is an $s_2 + 1$ dimensional vector, the second inequality follows from that fact that $\max_i |(A_1^{-1})_{s_1+1, i_1}| = O(p^{-s_1})$ from Lemma 4 and $\max_{i_1, i_2} \mathbb{E}[\tau_1^{i_1-1} \tau_2^{i_2-1}] = \mathbb{E}[\tau_1^{s_1} \tau_2^{s_2}]$, and in the third inequality we used a result from [3] that $\mathbb{E}[(\text{Binom}(d, p))^s] \leq \sum_{j=1}^s S(s, j) (dp)^j$ where $S(s, j)$ is the Sterling number of second kind. $S(s, j) \leq f(s)$, for $1 \leq j \leq s$. This proves the desired bound in (46).

A.4 Proof of Lemma 4

We are left to prove that the estimator $\theta(w, G_\Omega)$ simplifies as in (50) when the original graph is a clique graph(or a union of disjoint cliques).

We use the notations introduced in Section 3 and Section A.3. Consider a closed walk w of length k on ℓ distinct nodes with $U = \{u_1, \dots, u_{\tilde{\ell}}\}$ set of nodes in it that have at least one self-loop, $|U| = \tilde{\ell}$, and a total of s self loops. If the underlying graph is a clique graph the partition of V defined in (14), for any $T \subseteq U$, is as follows:

$$V_{T, U \setminus T} = \begin{cases} \emptyset & \text{if } |T| < |U| - 1 \\ d_{U, \emptyset} & \text{if } T = U \\ v & \text{if } T = U \setminus v, \text{ for any } v \in U. \end{cases} \tag{57}$$

Recall that $d_{U,\emptyset} \equiv \bigcap_{v \in U} \partial v$. Therefore, we have,

$$d_{T,U \setminus T}(w) = \begin{cases} 0 & \text{if } |T| < |U| - 1 \\ \ell - \tilde{\ell} & \text{if } T = U \\ 1 & \text{if } T = U \setminus v, \text{ for any } v \in U. \end{cases} \quad (58)$$

If the underlying clique graph is of size ω then $|d_{U,\emptyset}| = \omega - \tilde{\ell}$. Using the fact $d_{T,U \setminus T}(\Omega) \sim \text{Binom}(d_{T,U \setminus T} - d_{T,U \setminus T}(w), p) + d_{T,U \setminus T}(w)$, as explained in Section 3, we have,

$$d_{T,U \setminus T}(\Omega) \sim \begin{cases} 0 & \text{if } |T| < |U| - 1 \\ \text{Binom}(\omega - \ell, p) + (\ell - \tilde{\ell}) & \text{if } T = U \\ 1 & \text{if } T = U \setminus v, \text{ for any } v \in U. \end{cases}$$

Using Equation (57) it is immediate that that degree of any node $u \in U$ is $d_u = d_{U,\emptyset} + (\tilde{\ell} - 1)$, and hence, $\mathbb{E}[\theta(w, G_\Omega) | w \in G_\Omega] = (d_{U,\emptyset} + (\tilde{\ell} - 1))^s = (\omega - 1)^s$. Therefore, an alternative characterization of the estimator defined in (21) is as following: $\theta(w, G_\Omega)$, conditioned on the event that all the nodes in the walk w are sampled, is a random variable dependent only upon $d_{U,\emptyset}(\Omega) \sim \text{Binom}(\omega - \ell, p) + (\ell - \tilde{\ell})$ such that its conditional expectation is $\mathbb{E}[\theta(w, G_\Omega) | w \in G_\Omega] = (\omega - 1)^s$. With change of notations it is immediate that the estimator defined in (21) is same as the estimator in (50) when the underlying graph is a clique graph (or a disjoint union of cliques).

By the definition of A , it follows that the diagonal entries are exactly $\text{diag}([1, p, \dots, p^s])$ and the bottom-left off-diagonal entries are all $\Theta(1)$ with respect to p , and the top-right off-diagonal entries are all zeros. Applying the inverse to this lower triangular matrix, it follows that A^{-1} is also a lower triangular matrix with diagonal entries $\text{diag}([1, p^{-1}, \dots, p^{-s}])$ and the bottom-left off-diagonal entries are all $\Theta(1)$. It follows that $\max_{i \in [s+1]} |(A^{-1})_{s+1,i}| = O(p^{-s})$.

A.5 Proof of Theorem 4

The following lemma provides bound on variance of Schatten k -norm estimator for a connected general graph with maximum degree d_{\max} . We provide a proof in Section A.6

Lemma 5. *For a connected graph G on ω vertices with maximum degree d_{\max} , variance of Schatten k -norm estimator $\widehat{\Theta}_k(G_\Omega)$ is bounded by*

$$\text{Var}(\widehat{\Theta}_k(G_\Omega)) \leq h(k) \frac{\omega^2 d_{\max}^{2k}}{p} \left(1 + \frac{1}{(d_{\max} p)^{2k-1}} \right), \quad (59)$$

where $h(k) = O(2^{k^2})$. Moreover, if there exists a positive constant c such that $d_{\max}^3 p \geq c 2^{k^2}$ or $d_{\max} p^3 \leq c/2^{k^2}$ then (59) holds with $h(k) = \text{poly}(k)$.

Using Equations (27) and (29) along with Lemma 5, Theorem 4 follows immediately.

A.6 Proof of Lemma 5

We use the notations introduced in Section 3 and Section A.1. Denote the size of the connected component by ω and let d_{\max} be the maximum degree of any node in the connected component.

The following technical lemma provides upper bounds on the variance and covariance terms. We provide a proof in Section A.7.

Lemma 6. *Under the hypothesis of Lemma 5, for a length- k walk w over ℓ distinct nodes with $s \geq 1$ self-loops, the following holds:*

$$\text{Var}\left(\theta(w, G_\Omega)\mathbb{I}(w \subseteq G_\Omega)\right) \leq h(k)\left(p^\ell d_{\max}^{2s} + d_{\max}p^{\ell+1-2s}\right), \quad (60)$$

and when $\ell = 1$, $s = k$, we have,

$$\begin{aligned} & \text{Var}\left(\theta(w, G_\Omega)\mathbb{I}(w \subseteq G_\Omega)\right) \\ & \leq f(k)d_{\max}^{2k-1} + g(k)d_{\max}p^{2-2k} + pd_{\max}^{2k}, \end{aligned} \quad (61)$$

and for any length- k walks w_1, w_2 over ℓ_1, ℓ_2 distinct nodes with $\tilde{\ell}$ unique overlapping nodes, $|w_1 \cap w_2| = \tilde{\ell}$, $s_1, s_2 \geq 1$ self-loops respectively, the covariance term can be upper bounded by:

$$\begin{aligned} & \text{Cov}\left(\theta(w_1, G_\Omega)\mathbb{I}(w_1 \subseteq G_\Omega), \theta(w_2, G_\Omega)\mathbb{I}(w_2 \subseteq G_\Omega)\right) \\ & \leq h(k)p^{((\ell_1+\ell_2-\tilde{\ell})-(s_1+s_2))}\left((d_{\max}p)^{(s_1+s_2)} + d_{\max}p\right), \end{aligned} \quad (62)$$

for some function $h(k) = O(2^{k^2})$, $f(k) = O(k!)$, and $g(k) = \text{poly}(k)$, where p is the vertex sampling probability.

The total count of length k closed cycles on ℓ distinct nodes in a general graph on ω nodes graph with maximum degree d_{\max} is bounded by $f(k)\omega d_{\max}^{\ell-1}$. It follows from the observation that fixing a node in the cycle, there are at most $d_{\max}^{\ell-1}$ paths to $\ell - 1$ -hop neighbors. That is $|w \in W : \mathcal{H}(w) \in \mathbb{H}_{k,\ell,s}| \leq f(k)\omega d_{\max}^{\ell-1}$ for any $1 \leq s \leq k$.

We use these inequalities to get bound on variance and covariance terms in (32).

For a walk $w \in W$ with $\mathcal{H}(w) \in \mathbb{H}_{k,\ell,s}$ with $1 \leq s \leq k - 2$, using (60), we have,

$$\begin{aligned} & \frac{\omega d_{\max}^{\ell-1}}{p^{2\ell}} \text{Var}\left(\theta(w, G_\Omega)\mathbb{I}(w \subseteq G_\Omega)\right) \\ & \leq h(k)\left(\frac{d_{\max}^{\ell+2s-1}}{p^\ell} + \frac{d_{\max}^\ell}{p^{\ell+2s-1}}\right) \\ & \leq h(k)\omega\left(\frac{d_{\max}^{2k-3}}{p^2} + \frac{d_{\max}^2}{p^{2k-3}}\right). \end{aligned} \quad (63)$$

For a walk $w \in W$ with $\mathcal{H}(w) \in \mathbb{H}_{k,\ell,s}$ with $\ell = 1$, $s = k$, using (61), we have,

$$\begin{aligned} & \frac{\omega}{p^{2\ell}} \text{Var}\left(\theta(w, G_\Omega)\mathbb{I}(w \subseteq G_\Omega)\right) \leq \\ & f(k)\omega d_{\max}^{2k-1}p^{-2} + g(k)\omega d_{\max}p^{-2k} + \omega d_{\max}^{2k}p^{-1}. \end{aligned} \quad (64)$$

For a walk w with $s = 0$, $\theta(w, G_\Omega) = 1$, and, we have,

$$\frac{\omega d_{\max}^{\ell-1}}{p^{2\ell}} \text{Var}\left(\theta(w, G_\Omega)\mathbb{I}(w \subseteq G_\Omega)\right) \leq \frac{\omega d_{\max}^{\ell-1}}{p^\ell}. \quad (65)$$

Combining, Equations (63), (64), and (65), and using $|w \in W : \mathcal{H}(w) \in \mathbb{H}_{k,\ell,s}| \leq f(k)\omega d_{\max}^{\ell-1}$, we have

$$\begin{aligned} & \sum_{\ell=1}^k \sum_{s=0}^{k-\ell+1} \sum_{H \in \mathbb{H}_{k,\ell,s}} \left\{ \frac{1}{p^{2\ell}} \sum_{w \in W: \mathcal{H}(w)=H} \mathbb{I}(w \subseteq G) \times \right. \\ & \quad \left. \text{Var} \left(\theta(w, G_\Omega) \mathbb{I}(w \subseteq G_\Omega) \right) \right\} \\ & \leq h(k)\omega d_{\max} p^{-2k} + \omega d_{\max}^{2k} p^{-1}, \end{aligned} \quad (66)$$

and if $d_{\max} p^3 \leq 1/h(k)$, then the above quantity is bounded by $g(k)\omega d_{\max} p^{-2k}(1 + o(1))$. If $d_{\max}^3 p \geq h(k)$, then the above quantity is bounded by $\omega d_{\max}^{2k} p^{-1}(1 + o(1))$.

Analysis of covariance terms in (32) follows along the similar lines as that of the clique graph case and the result in Lemma 5 follows immediately.

A.7 Proof of Lemma 6

We give a lemma similar to Lemma 3 for the case of a general graph that provides a bound on conditional variance and conditional covariance terms. We give a proof in Section A.8.

Lemma 7. *Under the hypothesis of Lemma 5, for length- k walks w_1, w_2 over ℓ_1, ℓ_2 distinct nodes with $s_1, s_2 \geq 1$ self-loops respectively, the conditional variance of estimator $\theta(w_1, G_\Omega)$, defined in (21), given that all the nodes in the walk are sampled can be upper bounded by*

$$\begin{aligned} \text{Var} \left(\theta(w_1, G_\Omega) \middle| w_1 \subseteq G_\Omega \right) & \leq \\ & h(k) \left(d_{\max}^{2s_1} + d_{\max} p^{1-2s_1} \right), \text{ and} \end{aligned} \quad (67)$$

$$\begin{aligned} \mathbb{E} \left[\theta(w_1, G_\Omega) \theta(w_2, G_\Omega) \middle| \mathbb{I}(w_1 \subseteq G_\Omega) \mathbb{I}(w_2 \subseteq G_\Omega) \right] & \leq \\ & h(k) p^{-(s_1+s_2)} \left((d_{\max} p)^{s_1+s_2} + d_{\max} p \right), \end{aligned} \quad (68)$$

for some function $h(k) = O(2^{k^2})$, where p is the vertex sampling probability. Moreover, for a length k walk w with $\ell = 1$, and $s = k$,

$$\text{Var} \left(\theta(w, G_\Omega) \middle| w \subseteq G_\Omega \right) \leq f(k) p^{-1} d_{\max}^{2k-1} + g(k) d_{\max} p^{1-2k}, \quad (69)$$

for some function $f(k) = O(k!)$, and $g(k) = \text{poly}(k)$.

Using the above lemma, proof of Lemma 6 follows along the lines of the proof of Lemma 2.

A.8 Proof of Lemma 7

Recall that for a general graph $\theta(w, G_\Omega)$ is an unbiased estimator of $\prod_{u \in w} d_u^{s_u}$ and is given in (21). It is easy to see that for any given walk w on ℓ distinct nodes and with s self-loops,

$$\text{Var} \left(\theta(w, G_\Omega) \middle| w \subseteq G_\Omega \right) \leq h(k) \max_{\mathbb{T}} \left\{ \text{Var} \left(\left\{ \prod_{T \in \mathbb{T}} \tilde{d}_{T, U \setminus T}^{(t_T)} \right\} \right) \right\},$$

where $h(k) = O(2^{k^2})$. It follows from the fact that there are at most $k/2$ distinct nodes with self loops and hence at most $2^{k/2-1}$ partitions in (14) which leads to at most $2^{k^2/4}$ summation terms in (16). Further $\prod_{T \in \mathbb{T}} \hat{d}_{T, U \setminus T}^{(t_T)}$ is the product of independent random variables. Observe that using Lemma 3, we have

$$\text{Var}\left(\hat{d}_{T, U \setminus T}^{(t_T)}\right) \leq f(k) \left(p^{-1} d_{\max}^{2t_T-1} + d_{\max} p^{1-2t_T} \right) \quad (70)$$

and $\mathbb{E}[\hat{d}_{T, U \setminus T}^{(t_T)}] \leq d_{\max}^{t_T}$. Using the fact that for independent random variables X_1, X_2, \dots, X_n ,

$$\text{Var}(X_1 X_2 \cdots X_n) = \prod_{i=1}^n \left(\text{Var}(X_i) + (\mathbb{E}[X_i])^2 \right) - \prod_{i=1}^n (\mathbb{E}[X_i])^2, \quad (71)$$

we have,

$$\begin{aligned} & \max_{\mathbb{T}} \left\{ \text{Var} \left(\left\{ \prod_{T \in \mathbb{T}} \hat{d}_{T, U \setminus T}^{(t_T)} \right\} \right) \right\} \\ & \leq \prod_{T \in \mathbb{T}} f(k) \left(p^{-1} d_{\max}^{2t_T-1} + d_{\max} p^{1-2t_T} + d_{\max}^{2t_T} \right) \\ & \leq f(k) \left(d_{\max}^{2s} + d_{\max} p^{1-2s} \right) \end{aligned} \quad (72)$$

where in the last inequality we used that $\sum_{T \in \mathbb{T}} t_T = s$. (67) follows from collecting the above inequalities. (68) follows from the definition of $\theta(w, G_\Omega)$ given in (21) and the proof of (46) of Lemma 3. (69) follows directly from (47) of Lemma 3.

A.9 Proof of Proposition 2

$$\max_{x \in [\alpha, 1]} |H_\alpha(x) - f_{b^*}(x)| \leq \max_{x \in [\alpha, 1]} |H_\alpha(x) - f_{\tilde{b}}(x)|, \quad (73)$$

$$= \max \left\{ |H_\alpha(\alpha) - f_{\tilde{b}}(\alpha)|, |H_\alpha(1) - f_{\tilde{b}}(1)| \right\} \quad (74)$$

$$= \left(\frac{1-\alpha}{1+\alpha} \right)^m \quad (75)$$

where $\tilde{b} \equiv (2/(1+\alpha))[1, \dots, 1]$.