Estimating Power-Law Exponent with Edge Differential Privacy

Adam Tan
wat@sfu.ca
Simon Fraser University
Burnaby, BC, Canada

Mohamed Hefny
mohamed_hefny@sfu.ca
Simon Fraser University
Burnaby, BC, Canada

Keval Vora
keval@cs.sfu.ca
Simon Fraser University
Burnaby, BC, Canada

Abstract

Many real-world graphs have degree distributions that are well approximated by a power-law, and the corresponding scaling parameter α provides a compact summary of that structure which is useful for graph analysis and system optimization. When graphs contain sensitive relationship data, α must be estimated without revealing information about individual edges. This paper studies power-law exponent estimation under edge differential privacy. Instead of first releasing a noisy degree distribution and then fitting a power-law model, we propose privatizing only the low-dimensional sufficient statistics needed to estimate α, thereby avoiding the high distortion introduced by traditional approaches. Using these released statistics, we support both discrete approximation and likelihood-based numerical optimization for efficient parameter estimation. We develop edge-DP algorithms for both centralized and local DP models, compare degree release and log-statistic release in the local setting, and evaluate the resulting methods on various graph datasets across multiple privacy budgets and tail-cutoff settings.

ACM Reference Format:
Adam Tan, Mohamed Hefny, and Keval Vora. 2026. Estimating Power-Law Exponent with Edge Differential Privacy. In *Workshop on Secure and Private Data Management (SeQureDB '26), May 31-June 05, 2026, Bengaluru, India.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3807894.3810274

1 Introduction

Graph databases are useful in domains where relationships are central to the data, enabling efficient structure-based information retrieval. However, graphs often contain sensitive information, and there is a need to develop privacy-preserving graph analysis techniques that prevent sensitive information from being leaked.

Among the structural properties studied in graphs, degree distributions are especially important. Many real-world graphs exhibit scale-free behavior, in which a small number of nodes act as highly connected hubs while most nodes have relatively few connections. This pattern appears across domains such as the web, social networks, and online retail, among many others [31].

Such behavior is often modeled by a *power-law distribution* [8, 27, 31], typically expressed as $P(d) \propto d^{-\alpha}$, meaning that the probability $P(d)$ that a node has degree d decreases polynomially, with α being a constant parameter of the distribution known as the *scaling parameter*, also referred to as the power-law exponent. This means high-degree nodes are rare while low-degree nodes are much more common.

Estimating the scaling parameter α of a power-law distribution helps tailor graph algorithms and systems that rely on degree information [1, 7, 17, 22, 35, 37]. Practitioners typically estimate α by maximum likelihood, either through a closed-form discrete approximation or through numerical optimization [8]. For graphs containing sensitive relationship data, however, we must perform this estimation under privacy constraints so that the released parameter does not reveal sensitive information about individual edges or graph structure.

To estimate α while protecting sensitive graph information, we use differential privacy (DP) [11], which provides the formal framework for this goal by enabling data analysis while protecting sensitive graph information. Existing DP methods for graph analysis [10, 15, 29, 38, 41, 42] provide privacy guarantees for structural information in graphs. However, they do not study private estimation of the power-law scaling parameter α directly.

To address this gap, we develop algorithms for estimating the power-law scaling parameter α under *edge differential privacy* (edge-DP), which protects individual edges.

A common baseline for private α estimation, used for example by Hay et al. [13], is to first release a DP degree-distribution histogram, that is, counts of how many nodes have degree $0, 1, 2, \ldots$, and then fit a power-law model to the privatized histogram via MLE. However, when the goal is to estimate a single scalar parameter, this pipeline of releasing a histogram and then fitting a model is inefficient. The noise added to each degree count, together with smoothing, binning, and projection steps, can distort the tail of the degree distribution before fitting, which leads to inaccurate α estimates with high variance.

Instead, we privatize only the low-dimensional statistics needed for the α estimation. We start from the discrete approximation estimator for the power-law scaling parameter [8], decompose it into low-sensitivity sub-components, and apply the Laplace mechanism to each component before recombining them into a private estimate. We then show how the same privatized statistics can also support maximum likelihood estimation via numerical optimization, allowing us to obtain both discrete-approximation and numerical-optimization variants under edge-DP. Because the released quantities have low sensitivity, the Laplace mechanism adds relatively small noise.

We develop differentially private algorithms under both the centralized and local models. In the centralized model, a trusted curator has access to the entire graph and releases noisy estimates of the required low-dimensional statistics. In the local model, there is no trusted curator, and each node perturbs its own edge-related statistics before release. We study two local release strategies: degree release and log-statistic release.

SeQureDB '26, Bengaluru, India

ACM ISBN 979-8-4007-2219-6/2026/05
https://doi.org/10.1145/3807894.3810274

Table 1: Notations.

Symbols	Description
$G = (V, E)$	Graph with nodes V and edges E
d_v	Degree of node $v \in V$
$d_{\min}, d_{\max}$	Degree range of MLE fit $(d_{\min}, d_{\max})$
D'	Degrees in G between $d_{\min}$ and $d_{\max}$
T_{disc}, N	Statistics for α estimation
$\tilde{d}_v, \tilde{T}_{disc}, \tilde{N}$	DP estimates of d_v, T_{disc} and N
$\hat{\alpha}_{\text{CENTRAL}}, \hat{\alpha}_{\text{LOCAL}}$	Centralized DP and local DP α estimates

Together, these choices give two centralized edge-DP algorithms, one based on discrete approximation and another based on numerical optimization, and four local edge-DP variants obtained by combining degree release or log-statistic release with discrete approximation or numerical optimization.

We evaluate the accuracy of these methods on 6 publicly available graph datasets and 3 synthetic datasets. Our results show that directly privatizing the sufficient statistics needed to estimate α is more accurate and more stable than histogram-based fitting in the centralized model. Among the methods that privatize the sufficient statistics directly, numerical optimization is overall more accurate than discrete approximation.

2 Background and Setup

This section introduces the graph model, the power-law estimation setup, and the privacy definitions used throughout the paper. Table 1 summarizes the main notations.

2.1 Power-Law Degree Distribution

Graph. We consider a simple undirected graph $G = (V, E)$, where V denotes the set of nodes and E the set of edges. The degree of a node $v \in V$ is denoted by d_v.

Power-Law Degree Distribution. The degree distribution of a graph G is the probability distribution $P(d)$ that a randomly selected node $v \in V$ has degree $d_v = d$. Many real-world graphs are scale-free, with a small number of highly connected nodes and many low-degree nodes. A power-law degree distribution models scale-free graphs by $P(d) \propto d^{-\alpha}$ for $d \geq d_{\min}$, indicating the coexistence of a few highly connected nodes and many sparsely connected ones. Nodes with degrees $d_v \geq d_{\min}$ are called tail nodes.

The scaling parameter α is a single parameter that governs the heaviness of the distribution's right tail. It is typically below 3, with occasional exceptions [8]. Smaller values of α correspond to heavier tails and a higher likelihood of extreme high-degree nodes, whereas larger values imply a more rapid decay and consequently a more homogeneous connectivity structure within the network.

Discrete Power-Law Distribution. Since node degrees are integers, the degree distribution of tail nodes follows a discrete power-law distribution with parameter α. The node degrees in the tail are independent and identically distributed according to a truncated discrete power-law distribution with probability mass function $P(d \mid \alpha)$:

$$P(d \mid \alpha) = \frac{d^{-\alpha}}{Z(\alpha)} \qquad Z(\alpha) = \sum_{d=d_{\min}}^{d_{\max}} d^{-\alpha} \tag{1}$$

for $d \in \{d_{\min}, \ldots, d_{\max}\}$ where $d_{\max}$ represents a known upper bound on the node degrees and $Z(\alpha)$ is the normalizing constant based on Hurwitz-ζ-function [4]. The tail degrees with power-law distribution are captured in the multiset $D' = [d_v : v \in V \ \wedge \ d_{\min} \leq d_v \leq d_{\max}]$, and $N = |D'|$ is the number of degrees in the multiset D'.

Maximum Likelihood Estimation for α. A common way to estimate α is through Maximum Likelihood Estimation (MLE). Clauset *et al.* [8] showed that a closed-form discrete approximation estimator for observed degrees $d_i \in D'$ can be defined as follows:

$$\hat{\alpha}_{\text{disc}} = 1 + \frac{N}{T_{\text{disc}}} \qquad T_{\text{disc}} = \sum_{i=1}^{N} \ln\left(\frac{d_i}{d_{\min} - 0.5}\right) \tag{2}$$

This discrete approximation provides a closed-form estimate which may generally be good enough for most practical purposes. Alternatively, $\hat{\alpha}_{\text{disc}}$ can be obtained by numerical optimization as described next.

The likelihood function for the discrete power-law model in Eq. 1 with $d_i \in D'$ is [4, 8]:

$$\ell(\alpha; D') = \sum_{i=1}^{N} \log P(d_i \mid \alpha) = -\alpha \sum_{i=1}^{N} \ln d_i - N \ln Z(\alpha)$$

With the aggregated statistic $T = \sum_{i=1}^{N} \ln d_i$, we can express the log-likelihood in terms of T as follows:

$$\ell(\alpha; T, N) = -\alpha T - N \ln Z(\alpha) \tag{3}$$

where the sufficient statistics for α are the pair (T, N), since $Z(\alpha)$ does not depend on observed d_i. Therefore, the maximum likelihood estimator is:

$$\hat{\alpha}_{\text{disc}} = \arg\max_{\alpha > 0} \ell(\alpha; T, N) = \arg\max_{\alpha > 0} \left[-\alpha T - N \ln Z(\alpha)\right] \tag{4}$$

2.2 Privacy Model

In graph data where sensitive information lies in connections between entities, edge differential privacy (edge-DP) [19] ensures analyses do not reveal individual edges. Similar to the original DP formulation by Dwork *et al.* [11], with edge-DP each edge is treated as an individual entry in a database (or graph) G.

Edge Differentially Private α Estimation. A randomized α estimation algorithm $\mathcal{A}(G)$ that takes input graph G and outputs some $\hat{\alpha}_{\text{disc}}$ value from output space $\mathcal{R}$ is *ε edge differentially private* (ε-edge-DP) if for all $R \subseteq \mathcal{R}$, and neighboring graphs G and G',

$$Pr(\mathcal{A}(G) \in R) \leq e^{\varepsilon} Pr(\mathcal{A}(G') \in R)$$

Here, neighboring graphs G and G' share the same set of nodes but differ in one edge (*i.e.*, the size of the symmetric difference of their edge sets is 1). The ε is referred to as the *privacy budget* as it governs the amount of random noise added to the $\hat{\alpha}_{\text{disc}}$ value.

Centralized and Local Edge-DP Models. Edge differential privacy can be realized under two models based on data visibility: the central model and the local model.

In the centralized model, a trusted curator maintains the entire graph and applies a randomized algorithm to ensure that the presence or absence of any single edge cannot be inferred. This is suitable for traditional database scenarios like a curator-managed social network, where the entire graph is safely accessible.

However, a trusted curator of data store having access to the entire graph can become impractical in modern systems that rely on decentralized or federated architectures. In local edge differential privacy (LEDP) [10, 12, 15, 29, 32], each node retains ground truth to their associated data, and aggregations on the graphs are constructed using ε-edge-DP queries to the nodes. Hence, the $\hat{\alpha}_{\text{disc}}$ computed under LEDP is based on the degree estimates that must be computed from individually perturbed edges.

2.3 Our Goal

Our goal is to design ε-edge-DP algorithms that estimate the power-law scaling parameter of a graph G over its fitted tail D' while preserving high utility under privacy constraints. In the centralized model, the algorithm accesses the entire graph and produces a private estimate $\hat{\alpha}_{\text{CENTRAL}}$. In the LEDP model, each node ensures ε-edge-DP locally on its own edge-related information, and the reports are aggregated to produce private estimate $\hat{\alpha}_{\text{LOCAL}}$.

3 Centralized Algorithms

We develop two DP algorithms in centralized model to estimate scaling parameter $\hat{\alpha}_{\text{CENTRAL}}$: one using discrete approximation and other via numerical optimization. Both algorithms first compute noisy statistics for T_{disc} and N as defined in Eq. 2 using Laplace mechanism. And then, we use these to compute $\hat{\alpha}_{\text{CENTRAL}}$ using two approaches (Section 3.2 and Section 3.3).

To compute the noisy statistics using Laplace mechanism, we first analyze their sensitivities as described next.

3.1 Sensitivity Analysis

Lemma 3.1. Global sensitivity of T_{disc} is bounded by $\mathrm{O}\big(\ln(\frac{d_{\min}+1}{d_{\min}})\big)$.

PROOF. Consider neighboring graphs G and G' that differ in exactly one edge (u, w). Only degrees of u and w change across these two graphs. Let d'_u and d'_w be respectively the degrees of nodes u and w in G'. Without loss of generality, $d_u - d'_u = 1$ and $d_w - d'_w = 1$. Define the per-node contribution $t_v(G)$ as:

$$t_v(G) = \begin{cases} \ln\left(\frac{d_v(G)}{d_{\min}-0.5}\right) & d_v(G) \geq d_{\min} \\ 0 & d_v(G) < d_{\min} \end{cases}$$

The change in a node's contribution will depend on whether its degree remains in the tail. If it does (*i.e.*, $d_v \geq d_{\min}$ and $d'_v \geq d_{\min}$), then the change in the contribution is:

$$\begin{aligned} t_v(G) - t_v(G') &= \ln\left(\frac{d_v}{d_{\min}-0.5}\right) - \ln\left(\frac{d'_v}{d_{\min}-0.5}\right) \\ &= \ln\left(\frac{d_v}{d'_v}\right) = \ln\left(\frac{d+1}{d}\right) \end{aligned}$$

where d is an arbitrary integer such that $d \geq d_{\min}$. Because this expression is strictly decreasing in d, the largest possible difference will occur when $d = d_{\min}$:

$$|t_v(G) - t_v(G')| \leq \ln\left(\frac{d_{\min}+1}{d_{\min}}\right)$$

When $d_{\min} = 1$, this difference becomes $\ln 2$. The difference decreases as $d_{\min}$ grows, because $\frac{d_{\min}+1}{d_{\min}}$ decreases as $d_{\min}$ increases.

Algorithm 1 DP $\hat{\alpha}_{\text{CENTRAL}}$ via Discrete Approximation

Input: Graph G, Minimum Degree $d_{\min}$, Maximum Degree $d_{\max}$, Privacy Budget $\varepsilon = \varepsilon_t + \varepsilon_n$
Output: DP Estimate of α

1: $D' = [d_v : \forall v \in G.V \wedge d_{\min} \leq d_v \leq d_{\max}]$
2: $T_{disc} \leftarrow \sum_{d \in D'} \ln\left(\frac{d}{d_{\min}-0.5}\right)$
3: $\tilde{T}_{disc} \leftarrow T_{disc} + \text{LAP}(\frac{2 \times ln(\frac{d_{\min}+1}{d_{\min}})}{\varepsilon_t})$
4: $\tilde{N} \leftarrow |D'| + \text{LAP}(\frac{2}{\varepsilon_n})$
5: **return** $1 + \frac{\tilde{N}}{\tilde{T}_{disc}}$

When a node's degree crosses from being below $d_{\min}$ to above $d_{\min}$ (or vice versa), then its contribution changes from 0 to $\ln\left(\frac{d_{\min}}{d_{\min}-0.5}\right)$. The absolute value of this quantity is also no larger than $\ln\left(\frac{d_{\min}+1}{d_{\min}}\right)$ when $d_{\min} \geq 1$.

Therefore, the sensitivity of $t_v(G)$ is no larger than $\ln\left(\frac{d_{\min}+1}{d_{\min}}\right)$ for any node, and since a single edge affects at most two nodes, the total sensitivity of $T_{\text{disc}} = \sum_v t_v(G)$ is bounded by:

$$\Delta T_{\text{disc}} = \max_{\text{neighbors } G,G'} |T_{\text{disc}}(G) - T_{\text{disc}}(G')| \leq 2 \ln\left(\frac{d_{\min}+1}{d_{\min}}\right)$$

□

The above $d_{\min}$ dependent bound on the global sensitivity of T_{disc} is important because it remains small for all relevant choices of $d_{\min}$. Even at the smallest value, $d_{\min} = 1$, the sensitivity is only $2 \ln 2 \approx 1.386$. As $d_{\min}$ increases, this bound decreases, so for a fixed privacy budget the released value $\tilde{T}_{\text{disc}}$ stays closer to the true statistic. The overall effect of $d_{\min}$ on final estimation accuracy, however, depends on additional factors and is evaluated in Section 5.

Lemma 3.2. Global sensitivity of N is at most 2.

PROOF. Consider neighboring graphs G and G' that differ in exactly one edge. The only nodes that can have their tail membership changed are the endpoints of the edge. Each endpoint can either enter the tail (if previously out of the tail) or exit the tail (if previously in the tail). Since only two nodes are affected, the total change in the number of nodes in the tail is:

$$\Delta N = \max_{\text{neighbors } G,G'} |N(G) - N(G')| \leq 2$$

□

3.2 $\hat{\alpha}_{\text{CENTRAL}}$ via Discrete Approximation

Using global sensitivities ΔT_{disc} and ΔN, the DP estimate $\hat{\alpha}_{\text{CENTRAL}}$ is computed using the Laplace mechanism [11].

Algorithm 1 computes $\hat{\alpha}_{\text{CENTRAL}}$ using the discrete approximation estimator from Eq. 2. Lines 1-2 compute D' and T_{disc} using node degrees. Lines 3-4 add Laplace noise proportional to the global sensitivities to compute noisy $\tilde{T}_{\text{disc}}$ and $\tilde{N}$. The ε budget is split into ε_t and ε_n while adding Laplace noise for $\tilde{T}_{\text{disc}}$ and $\tilde{N}$ respectively. Finally, line 5 inserts the noisy estimates into Eq. 2 to obtain the DP estimate $\hat{\alpha}_{\text{CENTRAL}}$; by post-processing, this step has no privacy loss.

Lemma 3.3. Using the Laplace Mechanism and Sequential Composition [11], $\hat{\alpha}_{\text{CENTRAL}}$ computed by Algorithm 1 is $(\varepsilon_t + \varepsilon_n)$-edge differentially private.

Algorithm 2 DP $\hat{\alpha}_{\text{CENTRAL}}$ via Numerical Optimization

Input: Graph G, Minimum Degree $d_{\min}$, Maximum Degree $d_{\max}$, Privacy Budget $\varepsilon = \varepsilon_t + \varepsilon_n$
Output: DP Estimate of α
1: $D' = [d_v : \forall v \in G.V \wedge d_{\min} \le d_v \le d_{\max}]$
2: $T_{disc} \leftarrow \sum_{d \in D'} \ln\left(\frac{d}{d_{\min}-0.5}\right)$
3: $\tilde{T}_{disc} \leftarrow T_{disc} + \text{LAP}(\frac{2 \times ln(\frac{d_{\min}+1}{d_{\min}})}{\varepsilon_t})$
4: $\tilde{N} \leftarrow |D'| + \text{LAP}(\frac{2}{\varepsilon_n})$
5: **return** $\underset{\alpha>0}{\text{argmax}}\ \text{LOGLIKELIHOOD}(\tilde{T}_{disc}, \tilde{N}, d_{\min}, d_{\max}, \alpha)$ ▷ Algo. 3

Algorithm 3 Log-likelihood score

Input: T_{disc}, N, Minimum Degree $d_{\min}$, Maximum Degree $d_{\max}$, Alpha α
Output: Log-Likelihood score for α
1: $Z \leftarrow \sum_{d=d_{\min}}^{d_{\max}} d^{-\alpha}$
2: **if** $Z \le 0$ **then**
3: **return** $-\infty$
4: **end if**
5: $S \leftarrow T_{\text{disc}} + N \times ln(d_{\min} - 0.5)$
6: **return** $-\alpha \times S - N \times ln(Z)$

3.3 $\hat{\alpha}_{\text{CENTRAL}}$ via Numerical Optimization

Instead of the above closed-form estimation, we can estimate $\hat{\alpha}_{\text{CENTRAL}}$ by numerically optimizing the discrete log-likelihood. The key idea is to reuse the same noisy $\tilde{T}_{\text{disc}}$ and $\tilde{N}$ estimates, as described next.

From the definition of T_{disc} in Eq. 2:

$$T_{\text{disc}} = \sum_{i=1}^{N} \ln\left(\frac{d_i}{d_{\min}-0.5}\right) = \sum_{i=1}^{N} \ln d_i - N \ln(d_{\min}-0.5)$$

Therefore,

$$\sum_{i=1}^{N} \ln d_i = T_{\text{disc}} + N \ln(d_{\min}-0.5)$$

Hence, we can compute the DP estimate of this sum using the noisy statistics $\tilde{T}_{\text{disc}}$ and $\tilde{N}$:

$$\widetilde{\sum \ln d_i} = \tilde{T}_{\text{disc}} + \tilde{N} \ln(d_{\min}-0.5) \tag{5}$$

Hence, based on Eq. 4, our central DP MLE is:

$$\hat{\alpha}_{\text{CENTRAL}} = \underset{\alpha>0}{\arg\max}\left[-\alpha \widetilde{\sum \ln d_i} - \tilde{N} \ln Z(\alpha)\right]$$

Algorithm 2 computes $\hat{\alpha}_{\text{CENTRAL}}$ by maximizing this DP log-likelihood objective. The computation of noisy statistics $\tilde{T}_{\text{disc}}$ and $\tilde{N}$ in lines 1-4 is same as in that in the previous algorithm, using Laplace noise with split budgets. In our experiments, we set $\varepsilon_t = \varepsilon_n = \varepsilon/2$ for simplicity. Using these DP estimates, MLE is numerically computed as post-processing step on line 5, with the log-likelihood computation shown in Algorithm 3.

Lemma 3.4. Using the Laplace Mechanism and Sequential Composition [11], $\hat{\alpha}_{\text{CENTRAL}}$ computed by Algorithm 2 is $(\varepsilon_t + \varepsilon_n)$-edge differentially private.

Discussion. The normalization constant $Z(\alpha)$ needs to know $d_{\max}$. In practice, if the maximum degree of the graph is known or can be assumed to be public knowledge (*e.g.*, maximum degree in social graphs is often visible or reported), it can directly be used without adding noise. On the other hand, $d_{\max}$ is just a single scalar number, so its private estimation requires much less ε budget than that for the entire degree distribution. Furthermore for power-law distributions, setting $d_{\max}$ conservatively high enough (*e.g.*, $d_{\max} = |V|$) has very little impact on the MLE with no privacy cost.

4 Local Algorithms

In the local model, each node perturbs its own information before releasing. Hence, the aggregator would never see the raw node degrees, and instead operate on noisy per-node statistics produced by local DP mechanisms.

We explore two approaches in this model, both with Laplace mechanism. The first approach computes LEDP degrees and uses them for $\hat{\alpha}_{\text{LOCAL}}$ estimation. The second approach is consistent with the central model; here, each node releases the noisy log-function statistic required to compute $\tilde{T}_{\text{disc}}$. These approaches result in four algorithms depending on the use of closed form discrete approximation versus numerical optimization using noisy statistics.

Approach 1: Release Degree Statistic. LEDP degree is computed with each node releasing its noisy degree [32]. Hence, the contribution from each node v is simply its DP degree estimate $\tilde{d}_v$. Hence, $\tilde{T}_{\text{disc}}$ and $\tilde{N}$ are defined as:

$$\tilde{T}_{\text{disc}} = \sum_{\tilde{d}_v \ge d_{\min}} \ln\left(\frac{\tilde{d}_v}{d_{\min}-0.5}\right) \qquad \tilde{N} = \sum_{\tilde{d}_v \ge d_{\min}} 1 \tag{6}$$

Approach 2: Release Log Statistic. Here, the contribution from each node v is modeled as $\tilde{c_v}$:

$$\tilde{c_v} = \begin{cases} \widetilde{\ln}\left(\frac{d_v}{d_{\min}-0.5}\right) & \tilde{d}_v \ge d_{\min} \\ 0 & \tilde{d}_v < d_{\min} \end{cases} \tag{7}$$

where $\tilde{d}_v$ is the DP estimate of d_v, and $\widetilde{\ln}\left(\frac{d_v}{d_{\min}-0.5}\right)$ denotes the DP estimate of $\ln\left(\frac{d_v}{d_{\min}-0.5}\right)$. Hence, $\tilde{T}_{\text{disc}}$ and $\tilde{N}$ are defined as:

$$\tilde{T}_{\text{disc}} = \sum_{v} \tilde{c_v} \qquad \tilde{N} = \sum_{\tilde{c_v} \ge \ln\left(\frac{d_{\min}}{d_{\min}-0.5}\right)} 1 \tag{8}$$

4.1 Sensitivity Analysis

We analyze the global sensitivity of the degree statistic and log statistic to guide the Laplace noise addition.

Lemma 4.1. Global sensitivity of node degree is 1.

PROOF. Adding or removing one edge can only change the degree of the endpoints of that edge by 1. □

Lemma 4.2. Global sensitivity of log statistic from Eq. 7 is at most $\ln(\frac{d_{\min}+1}{d_{\min}})$.

PROOF. The proof follows a similar argument to that for Lemma 3.1 proof. Consider neighboring graphs G and G' that differ in exactly one edge (u, w). Only degrees of u and w change across these two graphs. Let d'_u and d'_w be respectively the degrees of nodes u and w in G'. Without loss of generality, $d_u - d'_u = 1$ and $d_w - d'_w = 1$.

Algorithm 4 DP $\hat{\alpha}_{\text{LOCAL}}$ via Degree Release

Input: Graph G, Minimum Degree $d_{\min}$, Maximum Degree $d_{\max}$, Privacy Budget ε
Output: DP Estimate of α

/* Step 1: Release local degree statistic */
1: **for each** $v \in G.V$ **do**
2: $\tilde{d_v} \leftarrow d_v + Lap(\frac{2}{\varepsilon})$ ▷ $\frac{\varepsilon}{2}$ budget split
3: RELEASE $\tilde{d_v}$
4: **end for**

/* Step 2: Aggregate local statistics */
5: $\tilde{T}_{\text{disc}} \leftarrow 0$; $\tilde{N} \leftarrow 0$
6: **for** $v \in G.V$ **do**
7: $\tilde{d} \leftarrow$ DEGREESTATISTIC(v)
8: **if** $\tilde{d} \geq d_{\min}$ **then**
9: $\tilde{T}_{\text{disc}} \leftarrow \tilde{T}_{\text{disc}} + \ln(\frac{\tilde{d}}{d_{\min}-0.5})$
10: $\tilde{N} \leftarrow \tilde{N} + 1$
11: **end if**
12: **end for**

/* Step 3: Estimate α */
/* Option A: Discrete Approximation */
13: $\hat{\alpha}_{\text{LOCAL}} \leftarrow 1 + \frac{\tilde{N}}{\tilde{T}_{\text{disc}}}$
/* Option B: Numerical Optimization */
14: $\hat{\alpha}_{\text{LOCAL}} \leftarrow \underset{\alpha>0}{\text{argmax}}\, \text{LOGLIKELIHOOD}(\tilde{T}_{\text{disc}}, \tilde{N}, d_{\min}, d_{\max}, \alpha)$
15: **return** $\hat{\alpha}_{\text{LOCAL}}$

The change in a node's contribution c_v will depend on whether its degree remains in the tail. If it does (*i.e.*, $d_v \geq d_{\min}$ and $d'_v \geq d_{\min}$), then the change in the contribution is:

$$c_v(G) - c_v(G') = \ln\left(\frac{d_v}{d_{\min}-0.5}\right) - \ln\left(\frac{d'_v}{d_{\min}-0.5}\right) = \ln\left(\frac{d_v}{d'_v}\right) = \ln\left(\frac{d+1}{d}\right)$$

where d is an arbitrary integer such that $d \geq d_{\min}$. Because this expression is strictly decreasing in d, the largest possible difference will occur when $d = d_{\min}$:

$$|c_v(G) - c_v(G')| \leq \ln\left(\frac{d_{\min}+1}{d_{\min}}\right)$$

When $d_{\min} = 1$, this difference becomes $\ln 2$. The difference decreases as $d_{\min}$ grows, because the logarithm function is monotonically increasing.

When a node's degree crosses from being below $d_{\min}$ to above $d_{\min}$ (or vice versa), then its contribution changes from 0 to $\ln\left(\frac{d_{\min}}{d_{\min}-0.5}\right)$. The absolute value of this quantity is also no larger than $\ln\left(\frac{d_{\min}+1}{d_{\min}}\right)$ when $d_{\min} \geq 1$.

Thus in all cases the sensitivity of c_v is at most $\ln\left(\frac{d_{\min}+1}{d_{\min}}\right)$. □

The global sensitivity of c_v benefits from the similar $d_{\min}$-dependent bound as in the central model.

4.2 $\hat{\alpha}_{\text{LOCAL}}$ via Degree Release

Algorithm 4 shows the LEDP computation for $\hat{\alpha}_{\text{LOCAL}}$ using Laplace mechanism for degree release. In the first step (lines 1-4), each node releases its DP degree estimate $\tilde{d_v}$ computed using Laplace noise proportional to sensitivity 1. The ε privacy budget is divided by 2 as each edge is used twice to compute the degree estimates of its two endpoints. The second step (lines 5-12) aggregates the local releases to compute noisy statistics $\tilde{T}_{disc}$ and $\tilde{N}$ as defined in Eq. 6. This aggregation is post-processing using the DP degree estimates and has no privacy loss. Finally, these noisy statistics are used to estimate $\hat{\alpha}_{\text{LOCAL}}$ in step 3. This results in the following two options.

Option A:* $\hat{\alpha}_{\text{LOCAL}}$ *via Discrete Approximation. As shown on line 13, the noisy estimates are plugged into Eq. 2 for discrete approximation of the DP estimate $\hat{\alpha}_{\text{LOCAL}}$.

Option B:* $\hat{\alpha}_{\text{LOCAL}}$ *via Numerical Optimization. Numerical optimization is performed using noisy $\tilde{T}_{disc}$ and $\tilde{N}$ based on the same analysis for Eq. 5. Our local DP MLE is:

$$\hat{\alpha}_{\text{LOCAL}} = \underset{\alpha>0}{\arg\max}\left[-\alpha\widetilde{\sum \ln d_i} - \tilde{N}\ln Z(\alpha)\right] \quad (9)$$

which is shown on line 14 in Algorithm 4.

Lemma 4.3. Using the Laplace Mechanism and Sequential Composition [11], $\hat{\alpha}_{\text{LOCAL}}$ computed by Algorithm 4 is ε-edge differentially private.

Algorithm 5 DP $\hat{\alpha}_{\text{LOCAL}}$ via Log Statistic Release

Input: Graph G, Minimum Degree $d_{\min}$, Maximum Degree $d_{\max}$, Privacy Budget ε
Output: DP Estimate of α

/* Step 1: Release local log statistic */
1: **for each** $v \in G.V$ **do**
2: $\tilde{c_v} \leftarrow \ln(\frac{d_v}{d_{\min}-0.5}) + Lap(\frac{2\times\ln(\frac{d_{\min}+1}{d_{\min}})}{\varepsilon})$ ▷ $\frac{\varepsilon}{2}$ budget split
3: RELEASE $\tilde{c_v}$
4: **end for**

/* Step 2: Aggregate local statistics */
5: $\tilde{T}_{\text{disc}} \leftarrow 0$; $\tilde{N} \leftarrow 0$
6: **for** $v \in G.V$ **do**
7: $\tilde{c} \leftarrow$ LOGSTATISTIC(v)
8: **if** $\tilde{c} \geq \ln(\frac{d_{\min}}{d_{\min}-0.5})$ **then**
9: $\tilde{T}_{\text{disc}} \leftarrow \tilde{T}_{\text{disc}} + \tilde{c}$
10: $\tilde{N} \leftarrow \tilde{N} + 1$
11: **end if**
12: **end for**

/* Step 3: Estimate α */
/* Option A: Discrete Approximation */
13: $\hat{\alpha}_{\text{LOCAL}} \leftarrow 1 + \frac{\tilde{N}}{\tilde{T}_{\text{disc}}}$
/* Option B: Numerical Optimization */
14: $\hat{\alpha}_{\text{LOCAL}} \leftarrow \underset{\alpha>0}{\text{argmax}}\, \text{LOGLIKELIHOOD}(\tilde{T}_{\text{disc}}, \tilde{N}, d_{\min}, d_{\max}, \alpha)$
15: **return** $\hat{\alpha}_{\text{LOCAL}}$

4.3 $\hat{\alpha}_{\text{LOCAL}}$ via Log Statistic Release

Algorithm 5 shows the LEDP computation for $\hat{\alpha}_{\text{LOCAL}}$ using Laplace mechanism for log statistic release. In the first step (lines 1-4), each node releases its DP estimate of the log statistic using Laplace noise proportional to global sensitivity $\ln(\frac{d_{\min}+1}{d_{\min}})$ with the privacy budget split between two edge endpoints. The second step (lines 5-12) aggregates local contributions as post-processing to compute noisy statistics $\tilde{T}_{disc}$ and $\tilde{N}$ as defined in Eq. 8. These noisy statistics are used to estimate $\hat{\alpha}_{\text{LOCAL}}$ in step 3, resulting in following two options.

Table 2: Graph datasets & their power-law scaling parameter α.

Graph	Nodes	Edges	Power-law α	
			$d_{\min} = 1$	$d_{\min} = 3$
wiki	7,115	414,750	1.176	1.474
enron	36,692	367,661	1.494	1.918
brightkite	58,228	428,156	1.551	1.982
ego-twitter	81,306	4,841,488	1.187	1.372
gplus	107,614	60,989,732	1.126	1.222
stanford	281,903	2,312,497	1.459	2.218
syn-power-0	100,000	1,477,208	2.000	2.000
syn-power-1	100,000	4,010,327	2.500	2.500
syn-power-2	100,000	997,299	3.000	3.000

Option A: $\hat{\alpha}_{\text{LOCAL}}$ via Discrete Approximation. Eq. 2 is used for discrete approximation of the DP estimate $\hat{\alpha}_{\text{LOCAL}}$ using the noisy estimates (line 13 in Algorithm 5).

Option B: $\hat{\alpha}_{\text{LOCAL}}$ via Numerical Optimization. Numerical optimization is performed (line 14 in Algorithm 5) using noisy $\tilde{T}_{disc}$ and $\tilde{N}$ for local DP MLE defined in Eq. 9.

Lemma 4.4. Using the Laplace Mechanism and Sequential Composition [11], $\hat{\alpha}_{\text{LOCAL}}$ computed by Algorithm 5 is ε-edge differentially private.

5 Experimental Evaluation

In this section, we evaluate the accuracy of our ε-edge-DP α estimation algorithms and answer the following research questions:

RQ1. Does adding noise directly to sub-components of α estimator provide better estimates compared to the degree distribution based power-law fitting?

RQ2. How does the accuracy compare for numerical optimization using noisy estimates instead of directly using the closed form discrete approximation?

RQ3. Does degree release based approach in local model provide higher accuracy compared to solutions based on local log statistic release?

RQ4. How does the choice of $d_{\min}$ affect the accuracy and stability of private α estimation?

Algorithms. With different combinations for estimation methods (discrete approximation versus numerical optimization) and local statistic release (degree versus log statistic), we evaluate the centralized edge DP and LEDP algorithms listed in Table 3.

We compare against the ε-edge-DP degree distribution based power-law fitting approach from Hay et al. [13] which is developed for central model. This is called **BASE**.

Datasets. We test our algorithms on 9 graph datasets: 6 publicly-available datasets from repository [28] based on SNAP [21] (treated to be undirected) and 3 synthetic datasets. Table 2 summarizes the datasets. The power-law scaling parameter α values are mostly below 3 for $d_{\min}$ between 1 and 3; this is consistent with previous observations [8] where α is typically below 3. The synthetic datasets are generated using INC-POWERLAW generator [2] that produces a simple random graph conforming with a degree sequence corresponding to the given scaling parameter α.

Table 3: Algorithm labels used in the evaluation.

Label	Model	Release	Estimator	Reference
DA	Centralized	–	Discrete Approx.	Algorithm 1
NO	Centralized	–	Numerical Opt.	Algorithm 2
DA/DR	Local	Degree	Discrete Approx.	Algorithm 4 (A)
DA/LR	Local	Log-Statistic	Discrete Approx.	Algorithm 5 (A)
NO/DR	Local	Degree	Numerical Opt.	Algorithm 4 (B)
NO/LR	Local	Log-Statistic	Numerical Opt.	Algorithm 5 (B)

Table 4: Summary for centralized model. Mean l_1 is averaged over all runs and datasets, max l_1 is the worst case, and std. range gives the per-dataset standard deviation range over 20 runs. Bold marks the lowest mean and max for each $d_{\min}$.

Method	Mean l_1 (%)	Max l_1 (%)	Std. range
	$d_{\min} = 1$		
BASE	15.88	76.70	5.249–16.443
DA	9.57	17.69	0.00063–0.01459
NO	**0.0049**	**0.0989**	0.00075–0.02512
	$d_{\min} = 3$		
BASE	21.79	89.80	6.962–16.850
DA	5.28	9.47	0.00111–0.02877
NO	**0.0066**	**0.1115**	0.00079–0.03200

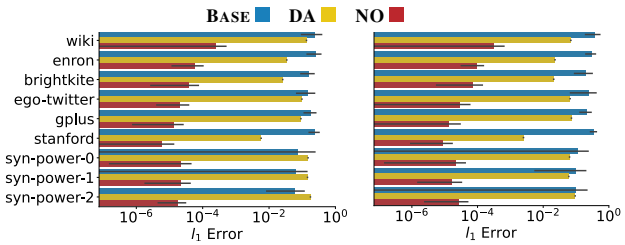

(a) $d_{\min} = 1$ **(b)** $d_{\min} = 3$

Figure 1: Performance of centralized DP algorithms. l_1 errors of NO, DA, and BASE.

Methodology. We set ε privacy budget to 1.0 for our experiments. We also conducted experiments where the ε value is varied between 0.1 and 5 to study performance across different privacy budgets. We report results for $d_{\min}$ values of 1 and 3; while we also considered with $d_{\min}$ values 5 and 10, the non-private MLE of α was outside of the $[0, \infty]$ range which is semantically invalid. For accuracy metric, we measure the l_1 error compared to the non-private α parameter of each dataset. Each experiment was repeated 20 times and we report the mean and standard deviation.

Detailed results for each dataset and ε are available in the extended version [34].

5.1 Centralized Algorithms

To answer RQ1 for $\hat{\alpha}_{\text{CENTRAL}}$ estimates, we compare DA and NO with BASE. Figure 1 and Table 4 summarize the results.

NO is the strongest central method. Its mean l_1 error is roughly three orders of magnitude lower than DA at both $d_{\min} = 1$ and $d_{\min} = 3$, and its worst-case error is about two orders of magnitude lower. This advantage holds on every individual dataset. On average, DA is more accurate than BASE, at both $d_{\min}$ values, though BASE outperforms DA on the three synthetic power-law datasets at $d_{\min} = 1$. The gap between NO and DA is because DA uses a closed-form approximation of α, while NO optimizes the exact discrete log-likelihood.

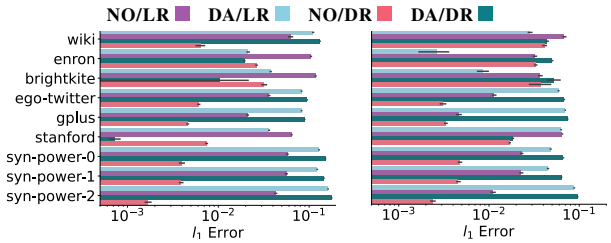

(a) $d_{\min} = 1$ **(b)** $d_{\min} = 3$

Figure 2: Performance of local DP algorithms. l_1 errors across different combinations for MLE (discrete approximation versus numerical optimization) and local statistic release (degree versus log statistic).

Increasing $d_{\min}$ worsens performance for BASE on every dataset and for NO on most datasets; DA, in contrast, improves at $d_{\min} = 3$ on every dataset.

5.2 Local Algorithms

To answer RQ2 and RQ3 for $\hat{\alpha}_{\text{LOCAL}}$ estimates, we compare NO/LR, DA/LR, NO/DR and DA/DR. Figure 2 and Table 5 summarize the results.

NO/DR is the strongest local variant overall. Mean l_1 error follows NO/DR < NO/LR < DA/LR < DA/DR at both $d_{\min} = 1$ and $d_{\min} = 3$ (Table 5), and NO/DR has the lowest per-dataset error on 6 of 9 datasets at each $d_{\min}$. Release mode interacts with the estimator: NO prefers degree release, while DA prefers log-statistic release, on most datasets at both $d_{\min}$ values. All four local variants have lower worst-case error than BASE. Increasing $d_{\min}$ helps all local variants except NO/DR; at $d_{\min} = 3$, DA/LR improves on 8 of 9 datasets, NO/LR on 7 of 9, and DA/DR on 6 of 9, while NO/DR worsens on 7 of 9.

5.3 Sensitivity to Privacy Budget

We analyze how l_1 error changes with privacy budget ε at fixed $d_{\min} = 1$. Figure 3, Figure 4, and Figure 5 show results for three datasets.

Syn-power-1. Figure 3a (central, $d_{\min} = 1$) shows a nearly flat trend for mean l_1 error across varying ε, with both central variants changing very little as privacy budget increases. In this dataset, NO has far lower error than DA at every ε. DA stays nearly same throughout while NO drops as ε increases.

Figure 3b (local, $d_{\min} = 1$) shows all four local curves decreasing as ε increases. Within the log-statistic release family, DA/LR is better at low privacy budgets ($\varepsilon \leq 0.5$), but NO/LR becomes better for $\varepsilon \geq 1$. Within the degree-release family, DA/DR is slightly better at $\varepsilon = 0.1$, while NO/DR has lower error for $\varepsilon \geq 0.3$. Overall, degree-release variants (NO/DR, DA/DR) remain below log-statistic-release variants (NO/LR, DA/LR), and the best local accuracy is obtained by NO/DR.

Comparing across the two models, central variants are less sensitive to ε, while local variants benefit more from larger ε. In absolute error, central NO is below all local variants across the plotted range.

Brightkite. Figure 4a (central, $d_{\min} = 1$) shows a nearly flat trend as ε increases from 0.1 to 5, with small variation across the full range. In this dataset, NO has far lower error than DA at every ε.

Table 5: Summary for local model. Mean l_1 is averaged over all runs and datasets, max l_1 is the worst case, and std. range gives the per-dataset standard deviation range over 20 runs. Bold marks the lowest mean and max for each $d_{\min}$.

Method	Mean l_1 (%)	Max l_1 (%)	Std. range
	$d_{\min} = 1$		
DA/LR	8.72	16.27	0.01644–0.10478
DA/DR	9.14	17.76	0.00295–1.07710
NO/LR	6.26	11.93	0.03661–0.30503
NO/DR	**1.03**	**3.31**	0.00817–0.18210
	$d_{\min} = 3$		
DA/LR	4.61	8.88	0.01388–0.12399
DA/DR	5.94	9.62	0.00528–0.85213
NO/LR	3.06	**7.21**	0.02496–0.22257
NO/DR	**1.63**	7.91	0.00927–0.99561

Figure 4b (local, $d_{\min} = 1$) shows that local error decreases overall as ε increases. The direct estimation helps the degree-release family across the full ε range, while in the log-statistic family the two estimators are close and trade places across ε, overall, DA/DR is the lowest error local method at every ε value shown. At smaller ε, both degree-release variants (NO/DR, DA/DR) sit below the log-statistic release variants (NO/LR, DA/LR), this advantage narrows with larger ε and NO/DR rises above the log-statistic curves at $\varepsilon = 5$. DA/DR, however, remains the best local option throughout.

Finally, comparing across the two models, the same pattern as Figure 3 holds: central curves are less sensitive to ε, while local curves improve as ε increases. In absolute error, central variants remain below all local variants across the full range.

Ego-twitter. Figure 5a (central, $d_{\min} = 1$) shows a similar trend as previous datasets. Figure 5b (local, $d_{\min} = 1$) shows a stronger dependence on ε, with all local curves decreasing as privacy budget increases. Within the degree-release family, NO/DR has lower error than DA/DR at every ε. Within the log-statistic release family, NO/LR has no valid estimate at $\varepsilon = 0.1$ because all runs hit the $\hat{\alpha} < 0$ clamp, but for $\varepsilon \geq 0.3$ it is consistently lower error than DA/LR. Overall, NO/DR is the lowest-error local method at every ε value shown, and by $\varepsilon \geq 1$ both numerical-optimization variants outperform their discrete-approximation counterparts.

Finally, comparing across two models, the same pattern as previous datasets holds: central curves are less sensitive to ε, while local curves improve overall as ε increases. In absolute error, central NO remains below all reported local estimates across the entire range.

5.4 Summary of Findings

Across all datasets, privacy budgets, and both privacy models, we highlight four overall observations.

Observation 1: Direct Sub-Component Privatization Outperforms Degree-Distribution Fitting. To answer RQ1, our results show that adding noise directly to the α-estimator sub-components is consistently better than degree-distribution based fitting. In the centralized setting, NO achieves lower error and better stability than BASE on every dataset, and DA does so in aggregate mean and on all real-world networks. The main reason is that our method perturbs only low-dimensional sufficient statistics (T_{disc}, N), while degree-distribution based fitting injects noise into many histogram bins and can distort the tail used to estimate α.

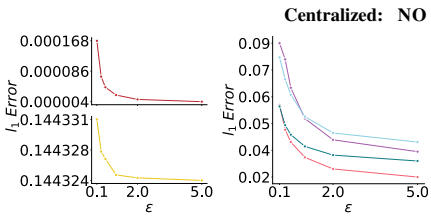

(a) Centralized. (b) Local.

Figure 3: Syn-power-1: l_1 vs ε.

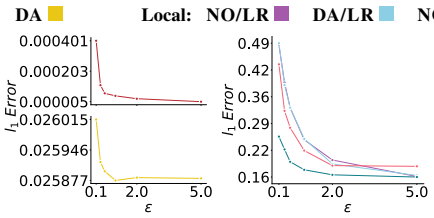

(a) Centralized. (b) Local.

Figure 4: Brightkite: l_1 vs ε.

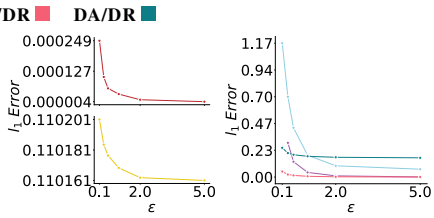

(a) Centralized. (b) Local.

Figure 5: Ego-twitter: l_1 vs ε.

Observation 2: Numerical Optimization is More Accurate Than Direct Approximation Under DP Noise. To answer RQ2, NO has lower l_1 error than DA in both models. In the centralized model, NO's mean l_1 error is more than two orders of magnitude lower than that of DA on every dataset, and nearly three orders of magnitude lower on aggregate. In the local model, NO/DR is far lower than DA/DR, and NO/LR is lower than DA/LR on average (Table 5). The difference in accuracy is because DA uses a closed-form approximation of α, while NO optimizes the exact discrete log-likelihood.

Observation 3: In Local DP, Degree Release Helps Numerical Optimization but Log-Statistic Release Helps Direct Approximation. To answer RQ3, the effect of release mode depends on the estimator family. NO/DR has lower mean l_1 error than NO/LR on most datasets, while DA/LR has lower mean l_1 error than DA/DR on most datasets.

Observation 4: Effect of $d_{\min}$ is Method-Dependent. To answer RQ4, the effect of increasing $d_{\min}$ from 1 to 3 is method-dependent. In the centralized setting, DA improves on every dataset, while NO worsens on most. In the local setting, DA/LR, DA/DR, and NO/LR improve on most datasets, while NO/DR worsens on most. The preferred $d_{\min}$ therefore depends on which estimator is used.

6 Related Work

Differentially Private Degree Distribution Release. A common approach for privately estimating α is to fit a power-law model to the privatized degree distribution histogram. Hay et al. [13] propose an efficient edge-DP method for releasing the degree distribution and show the fitting power-law estimate. Similarly, works like [14, 33, 39, 40] develop edge-DP degree distribution techniques. DP degree distribution release has also been studied under node-DP [9, 14, 24, 25, 36]. Compared to DP degree distribution use for α estimation, we approach the problem by adding noise only to the few statistics used for α estimation, hence avoiding inaccuracies from binning/smoothing/projection for noisy degree histograms.

Differentially Private Graph Algorithms. Beyond private degree distribution, several DP graph algorithms have been developed. They broadly fall into two styles covering both centralized and local models. The first style involves publishing private version of the graph (or a synthetic graph) [23, 32, 40, 42–44], whereas the second style is releasing DP estimations for specific graph queries [6, 10, 14–16, 19, 20, 29]. Our work relates follows the second style where we aim to release only the private estimate of a single parameter α.

Differentially Private Likelihood Optimization. Many likelihood-based estimators can be written as an optimization problem, *e.g.*, $\hat{\theta} = \arg\max_\theta \ell(\theta; D)$. There are ways to make such estimators private in the central DP model. One approach is to privatize the optimization itself, for example by perturbing the objective or by adding noise to the optimizer before release [5]. Another way is the McSherry-Talwar selection mechanism E_q^ε [26] which is used to select one output from candidates with quality scores (*e.g.*, log-likelihood). E_q^ε samples probabilistically, favoring higher-quality candidates, and providing differential privacy while selecting near-optimal outputs. More generally, DP M-Estimators can be computed by noisy iterative methods (*e.g.*, adding noise to gradients or Newton steps) that approximately solve the same estimation problem [3]. These DP ideas are relevant here because α is defined through a likelihood maximization problem, and our approach is to privately release only the few summary numbers for the estimation.

7 Conclusion and Future Work

We proposed methods to estimate the power-law exponent α of a graph's degree distribution under edge differential privacy. Our approach privatizes only the small set of sufficient statistics needed for the estimation, aiming to reduce tail distortion and error. We developed both centralized and local edge-DP algorithms with discrete-approximation and numerical-optimization variants, and our experiments showed our direct sufficient-statistic privatization approach is more accurate and stable than histogram-based fitting.

Interesting directions for future work remain, as discussed next.

Randomized-Response Tail Counts. Randomized response [18] is a local DP method for estimating counts of binary attributes. In our setting, it could be used to estimate the tail size $N = \sum_v \mathbf{1}\{d_v \geq d_{\min}\}$ by having each node privatize and send only the 1-bit tail-membership indicator, rather than sending noisy degrees and estimating N by thresholding the noisy values.

Friendship-Paradox Sampling under LDP. Although our current estimators assume one privatized report per node, an open direction is to study settings with partial participation or limited communication. Friendship-paradox sampling [30] provides an alternative way to obtain degree observations that over-represent high-degree nodes, integrating this with LDP would require designing a protocol that collects the necessary neighbor-sampled degree information privately and understanding the privacy/utility tradeoff.

Acknowledgements

This work is supported by the National Cybersecurity Consortium and the Natural Sciences and Engineering Research Council of Canada.

References

[1] Takuya Akiba, Yoichi Iwata, and Yuichi Yoshida. Fast Exact Shortest-path Distance Queries on Large Networks by Pruned Landmark Labeling. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, SIGMOD '13, page 349–360, 2013.

[2] Daniel Allendorf, Ulrich Meyer, Manuel Penschuck, Hung Tran, and Nick Wormald. Engineering Uniform Sampling of Graphs with a Prescribed Power-law Degree Sequence. In *2022 Proceedings of the Symposium on Algorithm Engineering and Experiments (ALENEX)*, pages 27–40. SIAM, 2022.

[3] Marco Avella-Medina, Casey Bradshaw, and Po-Ling Loh. Differentially Private Inference via Noisy Optimization. *The Annals of Statistics*, 51(5):2067 – 2092, 2023.

[4] Heiko Bauke. Parameter Estimation for Power-law Distributions by Maximum Likelihood Methods. *The European Physical Journal B*, 58(2):167–173, July 2007.

[5] Kamalika Chaudhuri, Claire Monteleoni, and Anand D. Sarwate. Differentially Private Empirical Risk Minimization. *Journal of Machine Learning Research*, 12(29):1069–1109, 2011.

[6] Shixi Chen and Shuigeng Zhou. Recursive Mechanism: Towards Node Differential Privacy and Unrestricted Joins. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, SIGMOD '13, page 653–664, 2013.

[7] Fan Chung and Linyuan Lu. Connected Components in Random Graphs with Given Expected Degree Sequences. *Annals of combinatorics*, 6(2):125–145, 2002.

[8] Aaron Clauset, Cosma Rohilla Shalizi, and Mark EJ Newman. Power-Law Distributions in Empirical Data. *SIAM Review*, 51(4):661–703, November 2009.

[9] Wei-Yen Day, Ninghui Li, and Min Lyu. Publishing Graph Degree Distribution with Node Differential Privacy. In *Proceedings of the International Conference on Management of Data*, SIGMOD '16, page 123–138, 2016.

[10] Laxman Dhulipala, Quanquan C. Liu, Sofya Raskhodnikova, Jessica Shi, Julian Shun, and Shangdi Yu. Differential Privacy from Locally Adjustable Graph Algorithms: k-Core Decomposition, Low Out-Degree Ordering, and Densest Subgraphs. In *IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 754–765, 2022.

[11] Cynthia Dwork and Aaron Roth. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4):211–407, August 2014.

[12] Talya Eden, Quanquan C Liu, Sofya Raskhodnikova, and Adam Smith. Triangle Counting with Local Edge Differential Privacy. *Random Structures & Algorithms*, 66(4):e70002, 2025.

[13] Michael Hay, Chao Li, Gerome Miklau, and David Jensen. Accurate Estimation of the Degree Distribution of Private Networks. In *Ninth IEEE International Conference on Data Mining*, pages 169–178, 2009.

[14] Yihua Hu, Hao Ding, and Wei Dong. N2E: A General Framework to Reduce Node-Differential Privacy to Edge-Differential Privacy for Graph Analytics. *Proceedings of the ACM on Management of Data*, 3(6), December 2025.

[15] Jacob Imola, Takao Murakami, and Kamalika Chaudhuri. Locally Differentially Private Analysis of Graph Statistics. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 983–1000. USENIX Association, August 2021.

[16] Jacob Imola, Takao Murakami, and Kamalika Chaudhuri. Communication-Efficient Triangle Counting under Local Differential Privacy. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 537–554, Boston, MA, August 2022. USENIX Association.

[17] Minhao Jiang, Ada Wai-Chee Fu, Raymond Chi-Wing Wong, and Yanyan Xu. Hop Doubling Label Indexing for Point-to-Point Distance Querying on Scale-Free Networks. *Proceedings of the VLDB Endowment*, 7(12), 2014.

[18] Peter Kairouz, Sewoong Oh, and Pramod Viswanath. Extremal Mechanisms for Local Differential Privacy. *Advances in Neural Information Processing Systems*, 27, 2014.

[19] Vishesh Karwa, Sofya Raskhodnikova, Adam Smith, and Grigory Yaroslavtsev. Private Analysis of Graph Structure. *Proceedings of the VLDB Endowment*, 4(11):1146–1157, August 2011.

[20] Shiva Prasad Kasiviswanathan, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. Analyzing Graphs with Node Differential Privacy. In *Theory of Cryptography*, pages 457–476, Berlin, Heidelberg, 2013.

[21] Jure Leskovec and Andrej Krevl. SNAP Datasets: Stanford Large Network Dataset Collection. http://snap.stanford.edu/data, June 2014.

[22] Ye Li, Leong Hou U, Man Lung Yiu, and Ngai Meng Kou. An Experimental Study on Hub Labeling Based Shortest Path Algorithms. *Proceedings of the VLDB Endowment*, 11(4):445–457, 2017.

[23] Zhetao Li, Yong Xiao, Haolin Liu, Xiaofei Liao, Ye Yuan, and Junzhao Du. Dynamic Graph Publication With Differential Privacy Guarantees for Decentralized Applications. *IEEE Transactions on Computers*, 74(5):1771–1785, 2025.

[24] Ganghong Liu, Xuebin Ma, and Wuyungerile Li. Publishing Node Strength Distribution With Node Differential Privacy. *IEEE Access*, 8:217642–217650, 2020.

[25] Shang Liu, Yang Cao, Takao Murakami, and Masatoshi Yoshikawa. A Crypto-Assisted Approach for Publishing Graph Statistics with Node Local Differential Privacy. In *IEEE International Conference on Big Data (Big Data), Osaka, Japan, December 17–20, 2022*, pages 5765–5774, 2022.

[26] Frank McSherry and Kunal Talwar. Mechanism Design via Differential Privacy. In *Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. IEEE, October 2007.

[27] Michael Mitzenmacher. A Brief History of Generative Models for Power Law and Lognormal Distributions. *Internet mathematics*, 1(2):226–251, 2004.

[28] Pranay Mundra and Quanquan C. Liu. mundrapranay/DistributedLEDPGraphAlgos: VLDB Artifact Release, Zenodo , June 2025.

[29] Pranay Mundra, Charalampos Papamanthou, Julian Shun, and Quanquan C Liu. Practical and Accurate Local Edge Differentially Private Graph Algorithms. *Proceedings of the VLDB Endowment*, 18(11):4199–4213, 2025.

[30] Buddhika Nettasinghe and Vikram Krishnamurthy. Maximum Likelihood Estimation of Power-law Degree Distributions via Friendship Paradox-based Sampling. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 15(6), May 2021.

[31] M. E. J. Newman. Power Laws, Pareto Distributions and Zipf's Law. *Contemporary Physics*, 46(5):323–351, 2005.

[32] Zhan Qin, Ting Yu, Yin Yang, Issa Khalil, Xiaokui Xiao, and Kui Ren. Generating Synthetic Decentralized Social Graphs with Local Differential Privacy. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, CCS '17, page 425–438, 2017.

[33] Jenni Reuben. Towards a Differential Privacy Theory for Edge-Labeled Directed Graphs. In *SICHERHEIT 2018*, pages 273–278. Gesellschaft für Informatik e.V., Bonn, 2018.

[34] Adam Tan, Mohamed Hefny, and Keval Vora. Estimating power-law exponent with edge differential privacy. *arXiv:2604.20274*, 2026.

[35] Wai Teng Tang, Ruizhe Zhao, Mian Lu, Yun Liang, Huynh Phung Huyng, Xibai Li, and Rick Siow Mong Goh. Optimizing and Auto-tuning Scale-free Sparse Matrix-vector Multiplication on Intel Xeon Phi. In *IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*, pages 136–145. IEEE, 2015.

[36] Jonathan Ullman and Adam Sealfon. Efficiently Estimating Erdos-Renyi Graphs with Node Differential Privacy. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

[37] Keval Vora. Lumos: Dependency-Driven Disk-based Graph Processing. In *USENIX Annual Technical Conference (USENIX ATC 19)*, pages 429–442, 2019.

[38] Songlei Wang, Yifeng Zheng, Xiaohua Jia, and Haibo Hu. PrivAGM: Secure Construction of Differentially Private Directed Attributed Graph Models on Decentralized Social Graphs. *Proceedings of the VLDB Endowment*, 18(11):4682–4694, July 2025.

[39] Yue Wang and Xintao Wu. Preserving Differential Privacy in Degree-Correlation based Graph Generation. *Transactions on Data Privacy*, 6(2):127–145, August 2013.

[40] Chengkun Wei, Shouling Ji, Changchang Liu, Wenzhi Chen, and Ting Wang. AsgLDP: Collecting and Generating Decentralized Attributed Graphs With Local Differential Privacy. *IEEE Transactions on Information Forensics and Security*, 15:3239–3254, 2020.

[41] Minze Xu, Zhentai Xie, Zhibin Wang, Guangzhan Wang, Longbin Lai, Yuan Zhang, Chen Tian, and Sheng Zhong. Sectric: Towards Accurate, Privacy-Preserving and Efficient Triangle Counting. *Proceedings of the VLDB Endowment*, 18(10):3382–3395, June 2025.

[42] Quan Yuan, Zhikun Zhang, Linkang Du, Min Chen, Peng Cheng, and Mingyang Sun. PrivGraph: Differentially Private Graph Data Publication by Exploiting Community Information. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 3241–3258, 2023.

[43] Sen Zhang, Haibo Hu, Qingqing Ye, and Jianliang Xu. PrivDPR: Synthetic Graph Publishing with Deep PageRank under Differential Privacy. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.1*, KDD '25, page 1936–1947, 2025.

[44] Yuxuan Zhang, Jianghong Wei, Xiaojian Zhang, Xuexian Hu, and Wenfen Liu. A Two-Phase Algorithm for Generating Synthetic Graph Under Local Differential Privacy. In *Proceedings of the 8th International Conference on Communication and Network Security*, ICCNS '18, page 84–89, 2018.