

# Analyzing Graphs with Node Differential Privacy

Shiva Prasad Kasiviswanathan\*    Kobbi Nissim†    Sofya Raskhodnikova‡    Adam Smith‡

March 16, 2015

## Abstract

We develop algorithms for the private analysis of network data that provide accurate analysis of realistic networks while satisfying stronger privacy guarantees than those of previous work. We present several techniques for designing *node* differentially private algorithms, that is, algorithms whose output distribution does not change significantly when a node and all its adjacent edges are added to a graph. We also develop methodology for analyzing the accuracy of such algorithms on realistic networks.

The main idea behind our techniques is to “project” (in one of several senses) the input graph onto the set of graphs with maximum degree below a certain threshold. We design projection operators, tailored to specific statistics that have low sensitivity and preserve information about the original statistic. These operators can be viewed as giving a fractional (low-degree) graph that is a solution to an optimization problem described as a maximum flow instance, linear program, or convex program. In addition, we derive a generic, efficient reduction that allows us to apply any differentially private algorithm for bounded-degree graphs to an arbitrary graph. This reduction is based on analyzing the smooth sensitivity of the “naive” truncation that simply discards nodes of high degree.

## 1 Introduction

Data from social and communication networks have become a rich source of insights in the social and information sciences. Gathering, sharing and analyzing these data is challenging, however, in part because they are often highly sensitive (your Facebook friends or the set of people you email reveal a tremendous amount of information about you, as in, e.g., Jernigan and Mistree [11]). This paper develops algorithms for the private analysis of network data that provide accurate analysis of realistic networks while satisfying stronger privacy guarantees than those of previous work.

A recent line of work, starting from Dinur and Nissim [4], investigates rigorous definitions of privacy for statistical data analysis. Differential privacy (Dwork *et al.* [8, 5]), which emerged from this line of work, has been successfully used in the context of “tabular”, or “array” data. Roughly, differential privacy guarantees that changes to one person’s data will not significantly affect the output distribution of an analysis procedure.

For tabular data, it is clear which data “belong” to a particular individual. In the context of graph data, two interpretations of this definition have been proposed: *edge* and *node* differential privacy. Intuitively, edge differential privacy ensures that an algorithm’s output does not reveal the inclusion or removal of

---

\*General Electric Global Research, USA. [kasivisw@gmail.com](mailto:kasivisw@gmail.com). Part of this work was done while the author was a postdoc at Los Alamos National Laboratory and IBM T.J. Watson Research Center.

†Ben-Gurion University, Israel; [kobbi@cs.bgu.ac.il](mailto:kobbi@cs.bgu.ac.il).

‡Pennsylvania State University, USA; [{sofya, asmith}@cse.psu.edu](mailto:{sofya, asmith}@cse.psu.edu). S.R. is supported by NSF CAREER grant CCF-0845701 and NSF grant CDI-0941553 and A.S. is supported by NSF Awards CCF-0747294 and CDI-0941553 as well as Penn State Clinical & Translational Research Institute, NIH/NCRR Award UL1RR033184.

a particular edge in the graph, while node differential privacy hides the inclusion or removal of a node together with all its adjacent edges.

Node privacy is a strictly stronger guarantee, but until now there have been no node-private algorithms that can provide accurate analysis of the sparse networks that arise in practice. One challenge is that for many natural statistics, node privacy is impossible to achieve while getting accurate answers in the worst case. The problem, roughly, is that node-private algorithms must be robust to the insertion of a new node in the graph, but the properties of a sparse graph can be altered dramatically by the insertion of a well-connected node. For example, for common graph statistics – the number of edges, the frequency of a particular subgraph – the change can overwhelm the value of the statistic in sparse graphs.

In this paper we develop several techniques for designing differentially *node*-private algorithms, as well as a methodology for analyzing their accuracy on realistic networks. The main idea behind our techniques is to “project” (in one of several senses) the input graph onto the set of graphs with maximum degree below a certain threshold. The benefits of this approach are two-fold. First, node privacy is easier to achieve in bounded-degree graphs since the insertion of one node affects only a relatively small part of the graph. Technically, the *sensitivity* of a given query function may be much lower when the function is restricted to graphs of a given degree. Second, for realistic networks this transformation loses relatively little information when the degree threshold is chosen carefully.

The difficulty with this approach is that the projection itself may be very sensitive to a change of a single node in the original graph. We handle this difficulty via two different techniques. First, for a certain class of statistics, we design tailored projection operators that have low sensitivity and preserve information about a given statistic. These operators can be viewed as giving a fractional (low-degree) graph that is a solution to a convex optimization problem, typically given by a maximum flow instance or linear program. Using such projections we get algorithms for accurately releasing the number of edges in a graph, and counts of small subgraphs such as triangles,  $k$ -cycles, and  $k$ -stars (used as sufficient statistics for popular graph models) in a graph, and certain estimators for power law graphs (see Sections 4 and 5).

Our second technique is much more general: we analyze the “naive” projection that simply discards high-degree nodes in the graph. We give efficient algorithms for bounding the “local sensitivity” of this projection, which measures how sensitive it is to changes in a particular input graph. Using this, we derive a generic, efficient reduction that allows us to apply any differentially private algorithm for bounded-degree graphs to an arbitrary graph. The reduction’s loss in accuracy depends on how far the input graph is from having low degree. We use this to design algorithms for releasing the entire degree distribution of a graph.

Because worst-case accuracy guarantees are problematic for node-private algorithms, we analyze the accuracy of our algorithms under a mild assumption on the degree distribution of the input graph. The simplest guarantees are for the case where a bound  $D$  on the maximum degree of the graph is known, and the guarantees typically relate the algorithms’s accuracy to how quickly the query function can change when restricted to graphs of degree  $D$  (e.g., Corollary 6.5). However, real-world networks are not well-modeled by a graphs of a fixed degree, since they often exhibit influential, high-degree nodes. In our main results, we assume only that tail of the degree distribution decreases slightly more quickly than what trivially holds for all graphs. (If  $\bar{d}$  is the average degree in a graph, Markov’s inequality implies that the fraction of nodes with degree above  $t \cdot \bar{d}$  is at most  $1/t$ . We assume that this fraction goes down as  $1/t^\alpha$  for a constant  $\alpha > 1$  or  $\alpha > 2$ , depending on the result.) Our assumption is satisfied by all the well-studied social network models we know of, including so-called *scale-free* graphs [3].

## 1.1 Related Work

The initial statements of differential privacy [8, 5] considered databases that are arrays or sets – each individual’s information corresponds to an entry in the database, and this entry may be changed without affecting other entries. That paper also introduced the very basic technique for constructing differentially private function approximations, by the addition of Laplace noise calibrated to the *global sensitivity* of the function.<sup>1</sup> This notion naturally extends to the case of graph data, where each individual’s information corresponds to an *edge* in the graph (edge privacy). The basic technique of Dwork *et al.* [8] continues to give a good estimate, e.g., for counting the number of edges in a graph, but it ceases to provide good analyses even for some of the most basic functions of graphs (diameter, counting the number of occurrences of a small specified subgraph) as these functions exhibit high global sensitivity.

The first differentially private computations over graph data appeared in Nissim *et al.* [15] where it was shown how to estimate, with differential edge privacy, the cost of the minimum spanning tree and the number of triangles in a graph. These computations employed a different noise addition technique, where noise is calibrated to a more local variant of sensitivity, called *smooth sensitivity*. These techniques and results were further extended by Karwa *et al.* [12]. Hay *et al.* [10] showed that the approach of [8] can still be useful when combined with a post-processing technique for removing some of the noise. They use this technique for constructing a differentially edge-private algorithm for releasing the degree distribution of a graph. They also proposed the notion of differential node privacy and highlighted some of the difficulties in achieving it.

A different approach to graph data was suggested by Rastogi *et al.* [17], where the privacy is weakened to a notion concerning a Bayesian adversary whose prior distribution on the database comes from a specified family of distributions. Under this notion of privacy, and assuming that the adversary’s prior admits mainly negative correlations between edges, they give an algorithm for counting the occurrences of a specified subgraph. The notion they use, though, is weaker than differential edge privacy. We refer the reader to [12] for a discussion on how the assumptions about an attacker’s prior limit the applicability of the privacy definition.

The current work considers databases where nodes correspond to individuals, and edges correspond to relationships between these individuals. Edge privacy corresponds in this setting to a requirement that the properties of every relationship (such as its absence or presence) should be kept hidden, but the overall relationship pattern of an individual may be revealed. However, each individual’s information corresponds to *all* edges adjacent to her node and a more natural extension of differential privacy for this setting would be that this entire information should be kept hidden. This is what we call *node privacy* (in contrast with *edge privacy* guaranteed in prior work). A crucial deviation from edge privacy is that a change in the information of one individual can affect the information of all other individuals. We give methods that provide node privacy for a variety of types of graphs, including very sparse graphs.

Finally, motivated by examples from social networks Gehrke *et al.* [9] suggest a stronger notion than differential node privacy – called *zero-knowledge privacy* – and demonstrate that this stronger notion can be achieved for several tasks in extremely dense graphs. Zero-knowledge privacy, as they employ it, can be used to release quantities that can be computed from small, random induced subgraphs of a larger graph. Their techniques are not directly applicable to sparse graphs (since a random induced subgraph will contain very few edges, with high probability).

We note that while node privacy gives a very strong guarantee, it may not answer all privacy concerns in a social network. Kifer and Machanavajjhala [13] criticize differential privacy in the context of social

---

<sup>1</sup>Informally, global sensitivity of a function measures the largest change in the function outcome than can result from changing one of its inputs.

networks, noting that individuals can have a greater effect on a social network than just forming their own relationships (their criticism is directed at edge privacy, but it can also apply to node privacy).

**Concurrent Work.** In independent work, Blocki *et al.* [1] also consider node-level differential private algorithms for analyzing sparse graphs. Both our work and that of Blocki *et al.* are motivated by getting good accuracy on sparse graphs, and employ projections onto the set of low-degree graphs to do so. The two works differ substantially in the technical details. See Appendix A for a detailed comparison.

**Organization.** Section 2 defines the basic framework of node and edge privacy and gives background on sensitivity and noise addition that is needed in the remainder of the paper. Section 3 introduces a useful, basic class of queries that can be analyzed with node privacy, namely queries that are linear in the degree distribution. Section 4 gives our first projection technique based on maximum flow and applies it to privately estimate the number of edges in a graph (Section 4.2). Section 4.3 generalizes the flow technique to apply it to any concave function on degree. Section 5 provides a private (small) subgraph counting algorithm via linear programming. Finally, Section 6 describes our general reduction from privacy on all graphs to the design of algorithms that are private only on bounded-degree graphs, and applies it to privately release the (entire) degree distribution.

## 2 Preliminaries

**Notation.** We use  $[n]$  to denote the set  $\{1, \dots, n\}$ . For a graph,  $(V, E)$ ,  $\bar{d}(G) = 2|E|/|V|$  is the average degree of the graph  $G$  and  $\deg_v(G)$  denotes the degree of node  $v \in V$  in  $G$ . When the graph referenced is clear, we drop  $G$  in the notation. The asymptotic notation  $O_n(\cdot), o_n(\cdot)$  is defined with respect to growing  $n$ . Other parameters are assumed to be functions independent of  $n$  unless specified otherwise.

Let  $\mathcal{G}$  denote the set of unweighted, undirected finite *labeled* graphs, and let  $\mathcal{G}_n$  denote the set of graphs on at most  $n$  nodes and  $\mathcal{G}_{n,D}$  be the set of all graphs in  $\mathcal{G}_n$  with maximum degree  $D$ .

### 2.1 Graphs Metrics and Differential Privacy

We consider two metrics on the set of labeled graphs: node and edge distance. The *node distance*  $d_{\text{node}}(G, G')$  (also called *rewiring distance*) between graphs  $G$  and  $G'$  is the minimum number of nodes in  $G'$  that need to be changed (“rewired”) to obtain  $G$ . Rewiring allows one to add a new node (with an arbitrary set of edges to existing nodes), remove it entirely, or change its adjacency lists arbitrarily. In particular, a rewiring can affect the adjacency lists of all other nodes. Equivalently, let  $k$  is the number of nodes in the largest induced subgraph of  $G$  which equals the corresponding induced subgraph of  $G'$ . The node distance is  $d_{\text{node}}(G, G') = \max\{|V_G|, |V_{G'}|\} - k$ . Graphs  $G, G'$  are node neighbors if their node distance is 1.

The *edge distance*  $d_{\text{edge}}(G, G')$  is the minimum number of edges in  $G'$  that need to be changed (*i.e.*, added or deleted) to obtain  $G$ . We also count insertion or removal of an isolated node (to allow for graphs with different number of nodes). In this paper, distance between graphs refers to the node distance unless specified otherwise.

**Definition 2.1** ( $(\epsilon, \delta)$ -differential Privacy [8, 5, 6]). *A randomized algorithm  $\mathcal{A}$  is  $(\epsilon, \delta)$ -node-private (resp. edge-private) if for all events  $S$  in the output space of  $\mathcal{A}$ , and for all graphs  $G, G'$  at rewiring distance 1 (resp. edge-distance 1) we have:*

$$\Pr[\mathcal{A}(G) \in S] \leq \exp(\epsilon) \times \Pr[\mathcal{A}(G') \in S] + \delta.$$

When  $\delta = 0$ , the algorithm is  $\epsilon$ -differentially private. In this paper, if node or edge privacy is not specified, we mean node privacy by default.

In this paper, for simplicity of presentation, we assume that  $n = |V|$ , the number of nodes of the input graph  $G$ , is publicly known. This assumption is justified since, as we will see, one can get a very accurate estimate of  $|V|$  via a node-private query. Moreover, given a publicly known value  $n$ , one can force the input graph  $G = (V, E)$  to have  $n$  nodes without sacrificing differential node privacy: one either pads the graph with isolated nodes (if  $|V| < n$ ) or discards the  $|V| - n$  “excess” nodes with the largest labels (if  $|V| > n$ ) along with all their adjacent edges. Changing one node of  $G$  corresponds to a change of at most one node in the resulting  $n$ -node graph as long as the differentially private algorithms being run on the data do not depend on the labeling (i.e., they should be symmetric in the order of the labels).

Differential privacy “composes” well, in the sense that privacy is preserved (albeit with slowly degrading parameters) even when the adversary gets to see the outcome of multiple differentially private algorithms run on the same data set.

**Lemma 2.2** (Composition, Post-processing [14, 7]). *If an algorithm  $\mathcal{A}$  runs  $t$  randomized algorithms  $\mathcal{A}_1, \dots, \mathcal{A}_t$ , each of which is  $(\epsilon, \delta)$ -differentially private, and applies an arbitrary (randomized) algorithm  $g$  to their results, i.e.,  $\mathcal{A}(G) = g(\mathcal{A}_1(G), \dots, \mathcal{A}_t(G))$ , then  $\mathcal{A}$  is  $(t\epsilon, t\delta)$ -differentially private.*

## 2.2 Calibrating Noise to Sensitivity

**Output Perturbation.** One common method for obtaining efficient differentially private algorithms for approximating real-valued functions is based on adding a small amount of random noise to the true answer. In this paper, we use two families of random distributions to add noise: Laplace and Cauchy. A *Laplace* random variable with mean 0 and standard deviation  $\sqrt{2}\lambda$  has density  $h(z) = (1/(2\lambda))e^{-|z|/\lambda}$ . We denote it by  $\text{Lap}(\lambda)$ . A *Cauchy* random variable with median 0 and median absolute deviation  $\lambda$  has density  $h(z) = 1/(\lambda\pi(1 + (z/\lambda)^2))$ . We denote it by  $\text{Cauchy}(\lambda)$ .

**Global Sensitivity.** The most basic framework for achieving differential privacy, Laplace noise is scaled according to the *global sensitivity* of the desired statistic  $f$ . This technique extends directly to graphs as long as we measure sensitivity with respect to the same metric as differential privacy. Below, we define these (standard) notions in terms of node distance and node privacy. Recall that  $\mathcal{G}_n$  is the set of all  $n$ -node graphs.

**Definition 2.3** (Global Sensitivity [8]). *The  $\ell_1$ -global node sensitivity of a function  $f : \mathcal{G}_n \rightarrow \mathbb{R}^p$  is:*

$$\Delta f = \max_{G, G' \text{ node neighbors}} \|f(G) - f(G')\|_1.$$

For example, the number of edges in a graph has node sensitivity  $n$  (when we restrict our attention to  $n$ -node graphs), since rewiring a node can add or remove at most  $n$  nodes. In contrast, the number of nodes in a graph has node sensitivity 1, even when we consider graphs of all sizes (not just a fixed size  $n$ ).

**Theorem 2.4** (Laplace Mechanism [8]). *The algorithm  $\mathcal{A}(G) = f(G) + \text{Lap}(\Delta f/\epsilon)^p$  (i.e., adds i.i.d. noise  $\text{Lap}(\Delta f/\epsilon)$  to each entry of  $f$ ), is  $\epsilon$ -node-private.*

Thus, we can release the number of nodes  $|V|$  in a graph with noise of expected magnitude  $1/\epsilon$  while satisfying node differential privacy. Given a public bound  $n$  on the number of nodes, we can release the number of edges  $|E|$  with additive noise of expected magnitude  $(n - 1)/\epsilon$  (the global sensitivity for releasing edge count is  $n - 1$ ).

**Local Sensitivity.** The magnitude of noise added by the Laplace mechanism depends on  $\Delta f$  and the privacy parameter  $\epsilon$ , but not on the database  $G$ . For many functions, this approach yields high noise, not reflecting the function’s typical insensitivity to individual inputs. Nissim *et al.* [15] proposed a local measure of sensitivity, defined next.

**Definition 2.5** (Local Sensitivity [15]). *For a function  $f : \mathcal{G}_n \rightarrow \mathbb{R}^p$  and a graph  $G \in \mathcal{G}_n$ , the local sensitivity of  $f$  at  $G$  is  $LS_f(G) = \max_{G'} \|f(G) - f(G')\|_1$ , where the maximum is taken over all node neighbors  $G'$  of  $G$ .*

Note that, by Definitions 2.3 and 2.5, the global sensitivity  $\Delta f = \max_G LS_f(G)$ . One may think of the local sensitivity as a discrete analogue of the magnitude of the gradient of  $f$ .

A straightforward argument shows that every differentially private algorithm must add distortion at least as large as the local sensitivity on many inputs. However, finding algorithms whose error matches the local sensitivity is *not* straightforward: an algorithm that releases  $f$  with noise magnitude proportional to  $LS_f(G)$  on input  $G$  is not, in general, differentially private [15], since the noise magnitude itself can leak information.

**Smooth Bounds on  $LS$ .** Nissim *et al.* [15] propose the following approach: instead of using the local sensitivity, select noise magnitude according to a *smooth* upper bound on the local sensitivity, namely, a function  $S$  that is an upper bound on  $LS_f$  at all points and such that  $\ln(S(\cdot))$  has low global sensitivity. The level of smoothness is parameterized by a number  $\beta$  (where smaller numbers lead to a smoother bound) which depends on  $\epsilon$ .

**Definition 2.6** (Smooth Bounds [15]). *For  $\beta > 0$ , a function  $S : \mathcal{G}_n \rightarrow \mathbb{R}$  is a  $\beta$ -smooth upper bound on the local sensitivity of  $f$  if it satisfies the following requirements:*

$$\begin{aligned} \text{for all } G \in \mathcal{G}_n : & \quad S(G) \geq LS_f(G); \\ \text{for all neighbors } G, G' \in \mathcal{G}_n : & \quad S(G) \leq e^\beta S(G'). \end{aligned}$$

One can add noise proportional to smooth bounds on the local sensitivity using a variety of distributions. We state here the version based on the Cauchy distribution.

**Theorem 2.7** (Calibrating Noise to Smooth Bounds [15]). *Let  $f : \mathcal{G}_n \rightarrow \mathbb{R}^p$  be a real-valued function and let  $S$  be a  $\beta$ -smooth bound on  $LS_f$ . If  $\beta \leq \epsilon/(\sqrt{2}p)$ , the algorithm  $\mathcal{A}(G) = f(G) + \text{Cauchy}(\sqrt{2}S(G)/\epsilon)^p$  (adding i.i.d.  $\text{Cauchy}(\sqrt{2}S(G)/\epsilon)$  to each coordinate of  $f$ ) is  $\epsilon$ -differentially private.*

From the properties of Cauchy distribution, the algorithm of the previous theorem has median absolute error  $(\sqrt{2}S(G))/\epsilon$  (the median absolute error is the median of the random variable  $|\mathcal{A}(G) - f(G)|$ , where  $\mathcal{A}(G)$  is the released value and  $f(G)$  is the query answer). Note that the *expected* error of Cauchy noise is not defined. One can get a similar result with an upper bound on any finite moment of the error using different heavy-tailed probability distributions [15]. We use Cauchy noise here for simplicity.

To compute smooth bounds efficiently, it is convenient to break the expression defining it down into tractable components. For every distance  $t$ , consider the largest local sensitivity attained on graphs at distance at most  $t$  from  $G$ . The *local sensitivity of  $f$  at distance  $t$*  is:

$$LS^{(t)}(G) = \max_{G' \in \mathcal{G}_n : d_{\text{node}}(G, G') \leq t} LS_f(G').$$

Now the smooth sensitivity is:  $S_{f, \beta}^*(G) = \max_{t=0, \dots, n} e^{-t\beta} LS^{(t)}(G)$ . Many smooth bounds on the local sensitivity have a similar form, with  $LS^{(t)}$  being replaced by some other function  $C^{(t)}(G)$  with the property that  $C^{(t)}(G) \leq C^{(t+1)}(G')$  for all pairs of neighbors  $G, G'$ . For example, our bounds on the sensitivity of naive truncation have this form (Proposition 6.4, Section 6).

### 2.3 Sensitivity and Privacy on Bounded-degree Graphs

A graph is  $D$ -bounded if it has maximum degree at most  $D$ . The degree bound  $D$  can be a function of the number of nodes in the graph. We can define a variant of differential privacy that constrains an algorithm only on these bounded-degree graphs.

**Definition 2.8** (Bounded-degree  $(\epsilon, \delta)$ -differential Privacy). *A randomized algorithm  $\mathcal{A}$  is  $(\epsilon, \delta)_D$ -node-private (resp.  $(\epsilon, \delta)_D$ -edge-private) if for all pairs of  $D$ -bounded graphs  $G_1, G_2 \in \mathcal{G}_{n,D}$  that differ in one node (resp. edge), we have  $\Pr[\mathcal{A}(G) \in \mathcal{S}] \leq e^\epsilon \Pr[\mathcal{A}(G') \in \mathcal{S}] + \delta$ .*

In bounded-degree graphs, the difference between edge privacy and node privacy is relatively small. For example, an  $(\epsilon, 0)_D$ -edge-private algorithm is also  $(\epsilon D, 0)_D$ -node-private (and a similar statement can be made about  $(\epsilon, \delta)$  privacy, with a messier growth in  $\delta$ ).

The notion of global sensitivity defined above (from previous work) can also be refined to consider only how the function may change within  $\mathcal{G}_{n,D}$ , and we can adjust the Laplace mechanism correspondingly to add less noise while satisfying  $(\epsilon, 0)_D$ -differential privacy.

**Definition 2.9** (Global Sensitivity on Bounded Graphs). *The  $\ell_1$ -global node sensitivity on  $D$ -bounded graphs of a function  $f : \mathcal{G}_n \rightarrow \mathbb{R}^p$  is:*

$$\Delta_D f = \max_{G, G' \in \mathcal{G}_{n,D}: d_{\text{node}}(G, G')=1} \|f(G) - f(G')\|_1.$$

**Observation 2.10** (Laplace Mechanism on Bounded Graphs). *The algorithm  $\mathcal{A}(G) = f(G) + \text{Lap}(\Delta_D f / \epsilon)^p$  is  $(\epsilon, 0)_D$ -node-private.*

### 2.4 Assumptions on Graph Structure

Let  $p_G$  denote the degree distribution of the graph  $G$ , i.e.,  $p_G(k) = |\{v : \deg_v(G) = k\}| / |V|$ . Similarly,  $P_G$  denotes the cumulative degree distribution, i.e.,  $P_G(k) = |\{v : \deg_v(G) \geq k\}| / |V|$ . Recall that  $\bar{d}(G) = 2|E| / |V|$  is the average degree of  $G$ .

**Assumption 2.11** ( $\alpha$ -decay). *Fix  $\alpha \geq 1$ . A graph  $G$  satisfies  $\alpha$ -decay if for all<sup>2</sup> real numbers  $t > 1$ ,  $P_G(t \cdot \bar{d}) \leq t^{-\alpha}$ .*

Note that *all* graphs satisfy 1-decay (by Markov's inequality). The assumption is nontrivial for  $\alpha > 1$ , but it is nevertheless satisfied by almost all widely studied classes of graphs. So-called “scale-free” networks (those that exhibit a heavy-tailed degree distribution) typically satisfy  $\alpha$ -decay for  $\alpha \in (1, 2)$ . Random graphs satisfy  $\alpha$ -decay for essentially arbitrarily large  $\alpha$  since their degree distributions have tails that decay exponentially (more precisely, for any  $\alpha$  we can find a constant  $c_\alpha$  such that, with high probability,  $\alpha$ -decay holds when  $t > c_\alpha$ ). Regular graphs satisfy the assumption with  $\alpha = \infty$ . Next we consider an implication of  $\alpha$ -decay.

**Lemma 2.12.** *Consider a graph  $G$  on  $n$  nodes that satisfies  $\alpha$ -decay for  $\alpha > 1$ , and let  $D > \bar{d}$ . Then the number of edges in  $G$  adjacent to nodes of degree at least  $D$  is  $O(\bar{d}^\alpha n / D^{\alpha-1})$ .*

<sup>2</sup>Our results hold even when this condition is satisfied only for sufficiently large  $t$ . For simplicity, we use a stronger assumption in our presentation.

*Proof.* The ratio of the number of edges touching nodes of degree at least  $D$  over  $n$  is at most

$$\begin{aligned} \sum_{i \geq D} p_G(i) \cdot i &= D \cdot P_G(D) + \sum_{j > D} P_G(j) \leq D \left( \frac{\bar{d}}{D} \right)^\alpha + \sum_{j > D} \left( \frac{\bar{d}}{j} \right)^\alpha \\ &\leq \bar{d}^\alpha (D^{1-\alpha} + \sum_{j > D}^\infty j^{-\alpha}) = O(\bar{d}^\alpha / D^{\alpha-1}). \end{aligned}$$

The first inequality uses the  $\alpha$ -decay assumption and the fact that  $D > \bar{d}$ . Therefore, it follows that the number of edges touching nodes of degree at least  $D$  is at most  $n \cdot O(\bar{d}^\alpha / D^{\alpha-1})$ .  $\square$

### 3 Linear Queries in the Degree Distribution

The first, and simplest, queries we consider are functions linear in the degree distribution. In many cases, these can be released directly with node privacy, though they also highlight why bounding the degree leads to such a drastic reduction in sensitivity. Suppose we are given a function  $h : \mathbb{N} \rightarrow \mathbb{R}^{\geq 0}$  that takes nonnegative real values. We can extend it to a function on graphs as follows:

$$F_h(G) \stackrel{\text{def}}{=} \sum_{v \in G} h(\deg_v(G)),$$

where  $\deg_v$  is the degree of the node  $v$  in  $G$ . We will drop the superscript in  $F_h$  when  $h$  is clear from the context. The query  $F_h$  can also be viewed as the inner product of  $\vec{h} = (h(0), \dots, h(n-1))$  with the degree distribution  $p_G$ , scaled up by  $n$ , i.e.,  $F_h(G) = n \langle \vec{h}, p_G \rangle$ .

Several natural quantities can be expressed as linear queries. The number of edges in the graph, for example, corresponds to half the identity function, that is,  $h(i) = i/2$  (since the sum of the degrees is twice the number of edges). The number of nodes in the graph is obtained by choosing the constant function  $h(i) = 1$ . The number of nodes with degrees in a certain range – say above a threshold  $D$  – also falls into this category. Less obviously, certain subgraph counting queries, namely, the number of  $k$ -stars for a given  $k$ , can be obtained by taking  $h(i) = \binom{i}{k}$  for  $i \geq k$  (and  $h(i) = 0$  for  $i < k$ ).

The sensitivity of these linear queries depends on the maximum value that  $h$  can take as well as the largest jump in  $h$  over the interval  $\{0, \dots, n-1\}$ . Let

$$\|h'\|_\infty \stackrel{\text{def}}{=} \max_{0 \leq i < n-1} |h(i+1) - h(i)|.$$

We refer to  $\|h'\|_\infty$  as the *maximum slope* of  $h$ . This quantity depends on  $n$ , though we leave  $n$  out of the notation for clarity. Let

$$\|h\|_\infty \stackrel{\text{def}}{=} \left( \max_{0 \leq i \leq n-1} |h(i)| \right).$$

**Lemma 3.1.** *The sensitivity of  $F_h$  on  $\mathcal{G}_n$  is at most  $\Delta F_h \leq \|h\|_\infty + (n-1) \cdot \|h'\|_\infty$ . If there is a value  $j \in \{0, \dots, n-1\}$  such that  $h(j) = 0$ , then  $\Delta F_h \leq 2(n-1) \|h'\|_\infty$ .*

*Proof.* Let  $G_1, G_2 \in \mathcal{G}_n$  be two graphs with  $d_{\text{node}}(G_1, G_2) = 1$ . Let  $j$  be the vertex that was rewired in  $G_1$  to create  $G_2$ . The summand in  $F_h(G) = \sum_{v \in G} h(\deg_v(G))$  corresponding to  $j$  changes by at most  $\max_{0 \leq i \leq n-1} |h(i)|$ . The degrees of all other nodes  $i \neq j$  can go up or down by at most 1, and so the corresponding summands each change by at most  $\|h'\|_\infty$ . The total change in  $F_h$  is thus at most



$(\max_{0 \leq i \leq n-1} |f(i)|) + (n-1) \cdot \|f'\|_\infty$ , as desired. If there is some value at which  $f(i) = 0$ , then the maximum value of  $f$  can be at most  $n-1$  times the maximum slope, yielding a total sensitivity bound of  $2(n-1)\|f'\|_\infty$ .  $\square$

This simple rule immediately gives us tight bounds on the sensitivity of several natural functions, such as the number of nodes, number of edges and the number of  $k$ -stars for a given  $k$ ).

- (1) **Number of Nodes:** If  $h(i) = 1$  for all  $i$ , then  $F_h(G)$  counts the number of nodes in  $G$  and the sensitivity of  $F_h$  is at most 1 (which is tight).
- (2) **Number of Edges:** If  $h(i) = i/2$ , then  $F_h(G)$  counts the twice the number of edges in  $G$ . It has maximum slope  $1/2$  and hence sensitivity the sensitivity of  $F_h$  is at most  $n-1$  (which is tight).
- (3) **Number of  $k$ -stars:** If  $h(i) = \binom{i}{k}$ , then  $F_h(G)$  counts the number of  $k$ -stars in  $G$ . The maximum slope of  $h$  is  $\binom{n-1}{k} - \binom{n-2}{k} = \binom{n-2}{k-1}$ . We thus estimate the sensitivity of  $F_h$  to be at most  $\binom{n-1}{k} + (n-1)\binom{n-2}{k-2}$ , which is tight. The second, simpler bound in the lemma yields  $2(n-1)\binom{n-2}{k-2}$ , which is slightly loose though both bounds are  $\Theta(n^{k-1})$  for fixed  $k$ .

A similar calculation shows that the sensitivity of  $k$ th moment of the degree distribution is  $\Theta(n^k)$ .

- (4) **Common Estimators for Power Law Coefficients:** Many real-world networks exhibit heavy-tailed degree distributions, and a common goal of analysts is to identify the coefficient of a power law that best fits the data (we note that power laws are not the only heavy-tailed distributions, but they are very popular). One well-studied approach to identifying the power law coefficient is to treat the degrees as  $n$  independent samples from a power law distribution (Clauset *et al.* [3]). In that case, the maximum likelihood estimator for the exponent is  $1 + n/M(G)$  where  $M(G) = \sum_{v \in V} \ln(\deg_v)$ . Note that  $M$  is a linear function of the degree distribution (as  $M(G) = F_h(G)$  with  $h(i) = \ln(i)$  for  $i \geq 1$  and  $h(0) = 0$ ) with maximum slope  $\ln(2) - \ln(1) = \ln(2)$  and maximum value  $\ln(n-1)$ . The sensitivity of  $M$  is  $\Theta(n)$ . Therefore, applying the Laplace mechanism directly is problematic, since the noise (of magnitude  $O(n/\epsilon)$ ) will swamp the value of the query. In Section 4.3, we propose a different approach (based on convex programming) for privately releasing these estimators.
- (5) **Counting Nodes in an Interval:** If  $f = \chi_{[a,b]}$  where  $\chi_{[a,b]}(i) = 1$  if  $a \leq i \leq b$ , and 0 otherwise, then  $F_f$  counts the number of nodes of degree between thresholds  $a$  and  $b$ . However, the sensitivity  $\Delta F_f = \Theta(n)$ , making the answer to this query useless once Laplace noise has been added.

We can reduce the sensitivity of this query by *tapering* the characteristic function of the interval. Given an interval  $[a, b]$ , consider the tapered step function  $f_{t,a,b}(i) = \max\{0, 1 - t \cdot \text{dist}(i, [a, b])\}$ , where  $\text{dist}(i, [a, b])$  denotes the distance from  $i$  to the nearest point in the interval  $[a, b]$ . The maximum slope of  $f_{t,a,b}$  is  $t$ , so  $\Delta F_{f_{t,a,b}} = 2tn$ . Answers to this query may be meaningful for any  $t = o(1)$  (since then the sensitivity will be  $o(n)$ ). We will find this sort of “smoothed” counting query to be useful when estimating how many nodes of high degree there are in a graph (see Proposition 6.4, Section 6).

- (6) **Low-order Fourier Coefficients:** Consider the function  $f_k(x) = \exp(i2\pi kx/n)$  (where  $i = \sqrt{-1}$ ). Now  $F_{f_k}$  gives the  $k$ th-order Fourier coefficient of the degree distribution. A simple calculation shows that the sensitivity  $\Delta F_{f_k}$  is  $O(k)$ .

The linear queries already give us a toolkit for analyzing graphs with node privacy, much as linear queries (over the data points) give a powerful basic toolkit for the differentially private analysis of conventional data sets (as in the SuLQ framework of Blum *et al.* [2]). The difference, of course, is that we need to consider slowly varying functions in order to keep the sensitivity low.

**Graphs of Bounded Degree** Notice that the techniques mentioned above for bounding the sensitivity of a linear query work better in bounded-degree graphs. Specifically, the sensitivity of  $F_h$  on  $D$ -bounded graphs is at most

$$\Delta F_h \leq \|h\|_\infty + D\|h'\|_\infty. \quad (1)$$

This motivates the approaches in the remainder of the paper, which seek to first bound the degree via a projection step.

## 4 Flow-based Lipschitz Extensions

We now present our flow-based technique. In Section 4.1, we define a flow function and show that it has low global node sensitivity and, on bounded-degree graphs, it correctly computes the number of edges in the graph. In Section 4.2, we design a node-private algorithm for releasing the number of edges in a graph based on this flow function.

### 4.1 Flow Graph

**Definition 4.1** (Flow graph). *Given an (undirected) graph  $G = (V, E)$ , let  $V_\ell = \{v_\ell \mid v \in V\}$  and  $V_r = \{v_r \mid v \in V\}$  be two copies of  $V$ , called the left and the right copies, respectively. Let  $D$  be a natural number less than  $n$ . The flow graph of  $G$  with parameter  $D$ , a source  $s$  and a sink  $t$  is a directed graph on nodes  $V_\ell \cup V_r \cup \{s, t\}$  with the following capacitated edges: edges of capacity  $D$  from the source  $s$  to all nodes in  $V_\ell$  and from all nodes in  $V_r$  to the sink  $t$ , and unit-capacity edges  $(u_\ell, v_r)$  for all edges  $\{u, v\}$  of  $G$ . Let  $v_{\text{fl}}(G)$  denote the value of the maximum flow in the flow graph of  $G$ .*

**Lemma 4.2.** *The global node sensitivity  $\Delta v_{\text{fl}} \leq 2D$ .*

*Proof.* Consider adding a node  $v$  with some adjacent edges to  $G$ . This corresponds to adding to the flow graph nodes  $v_\ell$  and  $v_r$ , edges  $(s, v_\ell)$  and  $(v_r, t)$  of capacity  $D$ , and some edges from  $v_\ell$  to  $V_r$  and from  $V_\ell$  to  $v_r$ . Since we only add edges, the old maximum flow is a valid flow in the new flow graph. So the value of the maximum flow cannot decrease.

To see that it can increase by at most  $2D$ , consider a minimum  $(s, t)$ -cut  $E_{\text{cut}}$  in the old graph. That cut  $E_{\text{cut}} \cup \{s, v_r\} \cup \{v_\ell, t\}$  in the new graph has capacity only  $2D$  larger than that of the cut  $E_{\text{cut}}$ . Thus, the capacity of minimum cut, and hence the maximum flow, increase by at most  $2D$ . By symmetry, removing a node from  $G$  cannot increase the maximum flow, and can decrease it by at most  $2D$ . Therefore, rewiring can change the maximum flow by at most  $2D$ .  $\square$

**Lemma 4.3.** *For all graphs  $G$ , the value  $v_{\text{fl}}(G) \leq 2f_e(G)$ . Moreover, if  $G$  is  $D$ -bounded then  $v_{\text{fl}}(G) = 2f_e(G)$ .*

*Proof.* The first statement holds because the cut that crosses all unit-capacity edges has capacity  $2 \cdot f_e(G)$ .

To prove the second statement, consider the following flow in the flow graph of  $G$ : all unit-capacity edges are carrying one unit of flow, and all edges of the form  $(s, v_\ell)$  and  $(v_r, t)$  carry flow  $\deg_v$ . Clearly, this flow is valid and has value  $\sum_{v \in V(G)} \deg_v = 2 \cdot f_e(G)$ , i.e.,  $v_{\text{fl}}(G) \geq 2 \cdot f_e(G)$ .

To see that this flow is maximum, consider the cut  $\{s\} \cup V_\ell$ . Its value is equal to the number of unit-capacity edges. Recall that for each edge  $\{u, v\}$  of  $G$ , we add edges  $(u_\ell, v_r)$  and  $(v_\ell, u_r)$  to the flow graph. Thus, the value of this cut is  $2 \cdot f_e(G)$ . By max-flow min-cut theorem,  $v_{\text{fl}}(G) \leq 2 \cdot f_e(G)$ .  $\square$

## 4.2 Algorithm for Releasing the Number of Edges

In this section, we design a node-private algorithm for releasing the number of edges. The main challenge in applying the methodology from the previous section is that we need to select a good threshold  $D$  that balances two conflicting goals: keeping the sensitivity low and retaining as large a fraction of the graph as possible.

Given a graph  $G$ , let  $f_e(G)$  be the number of edges in  $G$ . Observe that the global node sensitivity of the edge count,  $\Delta f_e$ , is at most  $n$  because rewiring (or adding/removing) a node can change this count by at most  $n$ . So releasing  $f_e$  with Laplace noise of the magnitude  $n/\epsilon$  is  $\epsilon$ -node-private. The resulting approximate count is accurate if the number of edges in the input graph  $G$  is large. The following algorithm allows us to release an accurate count even when this number is low, provided that  $G$  satisfies  $\alpha$ -decay, a natural assumption discussed in Section 2.4.

---

### Algorithm 1 $\epsilon$ -Node-Private Algorithm for Releasing $f_e(G)$

---

Input: parameters  $\epsilon, D, n$ , and graph  $G$  on  $n$  nodes.

- 1: Let  $\hat{e}_1 = f_e(G) + \text{Lap}(\frac{2n}{\epsilon})$  and threshold  $\tau = \frac{n \ln n}{\epsilon}$ .
  - 2: If  $\hat{e}_1 \geq 3\tau$ , **return**  $\hat{e}_1$ .
  - 3: Else compute the flow value  $v_{\text{fl}}(G)$  given in Definition 4.1 with  $D$ .
  - 4: **return**  $\hat{e}_2 = v_{\text{fl}}(G)/2 + \text{Lap}(\frac{2D}{\epsilon})$ .
- 

**Lemma 4.4.** *Algorithm 1 is an  $\epsilon$ -node-private algorithm that takes a graph  $G$  and parameters  $\epsilon, n, D$ , and outputs an approximate count for  $f_e(G)$  (number of edges in  $G$ ).*

1. If  $f_e(G) \geq (5n \ln n)/\epsilon$ , then with probability at least  $1 - 1/\ln n$ , Algorithm 1 outputs  $\hat{e}_1$  with

$$|\hat{e}_1 - f_e(G)| \leq (2n \ln \ln n)/\epsilon.$$

2. If  $G$  satisfies  $\alpha$ -decay for  $\alpha > 1$ ,  $D > \bar{d}$ , and  $f_e(G) < (n \ln n)/\epsilon$ , then with probability at least  $1 - 2/\ln n$ , Algorithm 1 outputs  $\hat{e}_2$  and

$$|\hat{e}_2 - f_e(G)| = O\left(\frac{2D \ln \ln n}{\epsilon} + \frac{\bar{d}^\alpha}{D^{\alpha-1}}\right).$$

The algorithm runs in  $O(nf_e(G))$  time.

*Proof.* First, we show that Algorithm 1 is  $\epsilon$ -node-private. By Theorem 2.4, an algorithm that releases estimate  $\hat{e}_1$  is  $\epsilon/2$ -node-private because the node sensitivity of  $f_e(G)$  is  $n$ . Similarly, an algorithm that releases estimate  $\hat{e}_2$  is  $\epsilon/2$ -node-private because, by Lemma 4.2, the node sensitivity of  $v_{\text{fl}}(G)/2$  is  $D$ . By the composition lemma (Lemma 2.2), the algorithm that releases both estimates is  $\epsilon$ -node-private, and Algorithm 1 can be simulated by an algorithm that knows only  $\hat{e}_1$  and  $\hat{e}_2$ , but has no access to  $G$ .

Second, we analyze the accuracy of the output. Setting  $\delta = 1/\ln n$  in Lemma B.1, we get that with probability at least  $1 - 1/\ln n$ , the estimate  $\hat{e}_1$  satisfies:

$$|\hat{e}_1 - f_e(G)| \leq \frac{2n \ln \ln n}{\epsilon} = 2\tau. \tag{2}$$

We say the estimate  $\hat{e}_1$  is *good* if it satisfies the condition in Equation (2). Notice that if  $f_e(G) \geq 5\tau$  and  $\hat{e}_1$  is good then  $\hat{e}_1 \geq 3\tau$ . That is, if  $f_e(G) \geq 5\tau$ , i.e., as in Case 1 of the lemma statement, then

with probability at least  $1 - 1/\ln n$  (which is the probability that  $\hat{e}_1$  is good), Algorithm 1 outputs  $\hat{e}_1$ . In other words, if  $f_e(G) \geq 5\tau$ , then with probability at least  $1 - 1/\ln n$ , Algorithm 1 outputs  $\hat{e}_1$  with  $|\hat{e}_1 - f_e(G)| \leq (2n \ln \ln n)/\epsilon$ .

Next we look at Case 2 of the lemma statement. If  $f_e(G) < \tau$ , then with probability at least  $1 - 1/\ln n$  (the same probability that  $\hat{e}_1$  is good),  $\hat{e}_1 < 3\tau$  and the algorithm outputs  $\hat{e}_2$ . Let  $e_h$  denote the number of edges adjacent to nodes of degree at least  $D$ . Lemma 2.12 implies that  $e_h = O(\bar{d}^\alpha/D^{\alpha-1})$ . By Lemma 4.3, the estimate  $v_{\text{fl}}(G)/2$  of the number of edges satisfies:

$$f_e(G) - \frac{v_{\text{fl}}(G)}{2} \leq e_h,$$

as the edges incident on the nodes of degree greater than  $D$  are the only edges which are not accounted by  $v_{\text{fl}}(G)/2$ .

Setting  $\delta = 1/\ln n$  in Lemma B.1, we get that with probability at least  $1 - 1/\ln n$ , the estimate  $\hat{e}_2$

$$\left| \hat{e}_2 - \frac{v_{\text{fl}}(G)}{2} \right| \leq \frac{2D \ln \ln n}{\epsilon}.$$

Combining the above statements we get that with probability at least  $1 - 1/\ln n$ ,

$$\begin{aligned} |\hat{e}_2 - f_e(G)| &\leq \left| \hat{e}_2 - \frac{v_{\text{fl}}(G)}{2} \right| + \left| \frac{v_{\text{fl}}(G)}{2} - f_e(G) \right| \leq \frac{2D \ln \ln n}{\epsilon} + e_h \\ &\leq \frac{2D \ln \ln n}{\epsilon} + O\left(\frac{\bar{d}^\alpha}{D^{\alpha-1}}\right). \end{aligned}$$

Combining above statements, we get that under the conditions of Case 2, with probability at least  $1 - (1/\ln n + 1/\ln n)$ , the algorithm outputs  $\hat{e}_2$  and

$$|\hat{e}_2 - f_e(G)| = O\left(\frac{2D \ln \ln n}{\epsilon} + \frac{\bar{d}^\alpha}{D^{\alpha-1}}\right).$$

This completes the accuracy analysis. The running time of the algorithm is dominated by the cost of computing max-flow in  $G$ , which can be done in  $O(nf_e(G))$  time [16].  $\square$

Using this lemma, and setting  $D = n^{1/\alpha}$ , we get the following theorem about privately releasing edge counts.

**Theorem 4.5** (Releasing Edge Counts Privately). *There is a node differentially private algorithm which, given constants  $\alpha > 1$ ,  $\epsilon > 0$ , and a graph  $G$  on  $n$  nodes, computes with probability at least  $1 - 2/(\ln n)$  an  $(1 \pm o_n(1))$ -approximation to  $f_e(G)$  (the number of edges in  $G$ ) if either of the following holds:*

1. If  $f_e(G) \geq (5n \ln n)/\epsilon$ .
2. If  $G$  satisfies  $\alpha$ -decay and  $f_e(G) = \omega(n^{1/\alpha}(\ln n)^{\alpha+1})$ .

*Proof.* **Case 1.** First assume that  $f_e(G) \geq 5n \ln n/\epsilon = 5\tau$  (where  $\tau = n \ln n/\epsilon$ ), then with probability at least  $1 - 1/\ln n$ , Algorithm 1 outputs  $\hat{e}_1$  and

$$|\hat{e}_1 - f_e(G)| \leq \frac{2n \ln \ln n}{\epsilon} \equiv \frac{|\hat{e}_1 - f_e(G)|}{f_e(G)} = o_n(1) \text{ as } f_e(G) \geq 5\tau.$$

Hence,  $\hat{e}_1$  is an  $(1 \pm o_n(1))$ -approximation to  $f_e(G)$ .

**Case 2.** Let  $G$  satisfy  $\alpha$ -decay and  $f_e(G) = \omega(n^{1/\alpha}(\ln n)^{\alpha+1})$ . Let us assume that  $\tau > f_e(G) = \omega(n^{1/\alpha}(\ln n)^{\alpha+1})$ . In this case, with probability at least  $1 - 2/\ln n$ , Algorithm 1 outputs  $\hat{e}_2$  and

$$|\hat{e}_2 - f_e(G)| = O\left(\frac{2D \ln \ln n}{\epsilon} + \frac{\bar{d}^\alpha}{D^{\alpha-1}}\right).$$

Set  $D = n^{1/\alpha}$ . Since average degree  $\bar{d} \leq 6\epsilon \ln n = O(\ln n)$ , therefore  $D > \bar{d}$ . Therefore, with probability at least  $1 - 2/\ln n$ , Algorithm 1 outputs  $\hat{e}_2$  and

$$\frac{|\hat{e}_2 - f_e(G)|}{f_e(G)} = \frac{O\left(\frac{2D \ln \ln n}{\epsilon} + \frac{\bar{d}^\alpha}{D^{\alpha-1}}\right)}{f_e(G)} = \frac{O\left(\frac{2n^{1/\alpha} \ln \ln n}{\epsilon} + \frac{n(\ln n)^\alpha}{n^{(\alpha-1)/\alpha}}\right)}{f_e(G)} = o_n(1).$$

as  $f_e(G) = \omega(n^{1/\alpha}(\ln n)^{\alpha+1})$ . Now if  $f_e(G) \geq 5\tau$ , then analysis in Case 1 implies the result. If  $f_e(G)$  lies between  $\tau$  and  $5\tau$ , then Algorithm 1 returns either  $\hat{e}_1$  and  $\hat{e}_2$ , and in both cases arguments similar as above can be used to show that we get the desired approximation ratio. This completes the proof of the theorem.  $\square$

### 4.3 Extension to Concave Query Functions

The flow-based technique of the previous section can be generalized considerably. In this section, we look at linear queries in the degree distribution in which the function  $h$  specifying the query is itself concave, meaning that its increments  $h(i+1) - h(i)$  are non-increasing as  $i$  goes from 0 to  $n-2$ . The number of edges in the graph is an example of such a query, since the increments of  $h(i) = i/2$  are constant.<sup>3</sup>

For mathematical convenience, we assume that the function  $h$  is in fact defined on the real interval  $[0, n-1]$  and is increasing and concave on that set (meaning that for all  $x, y \in [0, n-1]$ , we have  $h((x+y)/2) \leq (h(x) + h(y))/2$ ). It is always possible to extend a (discrete) function on  $\{0, \dots, n-1\}$  with nonincreasing increments to a concave function on  $[0, n-1]$  by interpolating linearly between each adjacent pair of values  $h(i), h(i+1)$ . Note that the maximum of  $h$  is preserved by this transformation, and the largest increment  $|h(i+1) - h(i)|$  equals the Lipschitz constant of the new function (defined as  $\sup_{x,y \in [0, n-1]} \frac{|h(x) - h(y)|}{|x - y|}$ ).

Given a graph  $G$  on at most  $n$  nodes, a concave function  $h$  on  $[0, n-1]$  and a threshold  $D$ , we define an optimization problem as follows: construct the flow graph (Definition 4.1) as before, but make the objective to maximize  $obj_h(Fl) = \sum_{v \in V} h(Fl(v))$ , where  $Fl(v)$  is the units of flow passing from  $s$  to  $v_\ell$  in the flow  $Fl$ . Let  $opt_h(G)$  denote the maximum value of the objective function over all feasible flows. The constraints of this optimization problem are all linear.

This new optimization problem is no longer a maximum flow problem (nor even a linear program), but the concavity of  $h$  ensures that it still a convex optimization problem and can be solved in polynomial time using convex programming techniques. Note that we need  $h$  be to concave only for computational efficiency purposes, and one could define the above flow graph and optimization problem for all  $h$ .

**Proposition 4.6.** *For every increasing function  $h : [0, n-1] \rightarrow \mathbb{R}^{\geq 0}$ ,*

1. *If  $G$  is  $D$ -bounded, then  $opt_h = F_h(G)$  (that is, the value of the optimization problem equals the correct value of the query).*

<sup>3</sup>There is some possible confusion here: any query of the form  $F_h$  described in Section 3 is linear in the degree distribution of the graph. Our additional requirement here is that the ‘‘little’’ function  $h$  be concave in the degree argument  $i$ .

2. The optimum  $opt_h$  has global sensitivity at most  $\|f\|_\infty + D\|f'\|_\infty$  on  $\mathcal{G}_n$ , where  $\|f\|_\infty = \max_{0 \leq x \leq D} h(x)$  and  $\|f'\|_\infty$  is the Lipschitz coefficient of  $h$  on  $[0, D]$  (that is, the global sensitivity of the optimization problem's value is at most the sensitivity of  $F_h$  on  $D$ -bounded graphs).

3. If  $h$  is concave then  $opt_h(G)$  can be computed to arbitrary accuracy in polynomial (in  $n$ ) time.

*Proof.* 1. If  $G$  is  $D$ -bounded, then there is a flow that routes  $\deg_v$  units of flow through nodes  $v_\ell$  and  $v_r$  for every node  $v$  of  $G$ . This is the maximum flow that can pass through  $v_\ell$  since it has  $\deg_v$  outgoing capacity-1 edges; similarly for  $v_r$ . This flow thus maximizes  $obj_h$ , and has value  $F_h(G)$ .

2. Consider graphs  $G', G$  (of arbitrary degree) where  $G'$  is obtained by inserting a node  $u$  in  $G$ . First, note that  $opt_h(G') \geq opt_h(G)$ , since the maximum flow for  $G$  is a feasible flow for  $G'$ . Second, note that given any feasible flow  $Fl$  for  $G'$ , we can obtain a feasible flow for  $G$  by removing the flow passing through  $u_\ell$  (which equals  $Fl(u_\ell)$ ) and the flow passing through  $u_r$  (which equals  $Fl(u_r)$ ). Removing the flow through  $u_\ell$  causes the objective function to drop by at most  $\|h\|_\infty = \max_{0 \leq x \leq D} h(x)$ . Removing the flow through  $u_r$  causes the total flow through left nodes other than  $u_\ell$  to drop by at most  $D$  units. This induces a change of at most  $\|h'\|_\infty D$ . So  $opt_h(G) \geq opt_h(G') - \|h\|_\infty + D\|h'\|_\infty$ , as desired.

3. If  $h$  is concave, then  $opt_h(G)$  can be computed using a convex program which we solve using an ellipsoid or interior point method.  $\square$

Thus, as with the number of edges, we can ask a query which matches  $F_h$  on  $D$ -bounded graphs but whose global sensitivity on the whole space is bounded by its sensitivity of the set of  $D$ -bounded graphs.

The MLE for power laws described in Section 3 is an interesting example where Proposition 4.6 could be used. There is a natural concave extension for the power law MLE: set  $f(x) = x$  for  $0 \leq x < 1$  and  $f(x) = 1 + \ln(x)$  for  $x \geq 1$ . The sensitivity of  $F_f$  on  $D$ -bounded graphs is  $\Delta_D f \leq 1 + \ln(D) + D$  (this follows from Equation (1)). In graphs with few high-degree nodes of degree greater than  $D$ , this leads to a much better private approximation to the power-law MLE in low-degree graphs than suggested in Section 3.

## 5 LP-based Lipschitz Extensions

In this section, we show how to privately release the number of (not necessarily induced copies) of a specified small template graph  $H$  in the input graph  $G$ . For example,  $H$  can be a triangle, a  $k$ -cycle, a length- $k$  path, a  $k$ -star ( $k$  nodes connected to a single common neighbor), or a  $k$ -triangle ( $k$  nodes connected to a pair of common neighbors that share an edge). Let  $f_H(G)$  denote the number of (not necessarily induced) copies of  $H$  in  $G$ , where  $H$  is a connected graph on  $k$  nodes.

### 5.1 LP-based Function

**Definition 5.1** (Function  $v_{LP}(G)$ ). *Given an (undirected) graph  $G = ([n], E)$  and a number  $D \in [n]$ , consider the following LP. The LP has a variable  $x_C$  for every copy of the template graph  $H$  in  $G$ . Let  $\Delta_D f$  denote the global node sensitivity of function  $f$  in  $D$ -bounded graphs. Then the LP corresponding to  $G$  is*

$$\text{maximize } \sum_{\text{copies } C \text{ of } H} x_C \text{ subject to:}$$

specified as follows:

$$0 \leq x_C \leq 1 \text{ for all variables } x_C$$

$$S_v \leq \Delta_D f_H \text{ for all nodes } v \in [n], \quad \text{where } S_v = \sum_{C:v \in V(C)} x_C.$$

We denote the value that maximizes this linear program by  $v_{\text{LP}}(G)$ .

When the variable  $x_C$  takes values 1 or 0, it signifies the presence or absence of the corresponding copy of  $H$  in  $G$ . The first type of constraints restricts these variables to  $[0, 1]$ . The second type of constraints says that every node can participate in at most  $\Delta_D f_H$  copies of  $H$ . This is the largest number of copies of  $H$  in which a node can participate in a  $D$ -bounded graph.

**Observation 5.2.**  $\Delta_D f_H \leq k \cdot D \cdot (D - 1)^{k-2}$ , where  $k$  is the number of nodes in  $H$ .

*Proof.* Let  $G$  be a  $D$ -bounded graph and  $v$  be a node of  $G$ . The global node sensitivity of  $f_H$  in  $D$ -bounded graphs is the number of copies of  $H$  in which  $v$  can participate. Recall that  $H$  is connected and assume w.l.o.g. that its node set is  $[k]$ . There are  $k$  nodes  $u$  of  $H$  to which node  $v$  can get mapped to,  $D$  choices of neighbors of  $v$  in  $G$  to which the smallest neighbor of  $u$  can get mapped to, and at most  $D - 1$  choices of unused neighbors of previous nodes to which each remaining node of  $H$  can get mapped to.  $\square$

**Lemma 5.3.** The global node sensitivity  $\Delta_{v_{\text{LP}}} \leq \Delta_D f_H \leq k \cdot D \cdot (D - 1)^{k-2}$ .

*Proof.* The second inequality follows from Observation 5.2. It remains to prove the first inequality.

Consider adding a node  $i$  with some adjacent edges to  $G$ . This corresponds to adding to the LP of  $G$  the variables  $x_C$ , where  $C$  are the copies of  $H$  that involve the new node  $i$ , as well as the constraints that restrict these variables to be in  $[0, 1]$  and the constraint  $S_i \leq \Delta_D f_H$ . In addition, the sums  $S_v$  for variables  $v$  other than  $i$  can have new variables added to them.

Every solution, where the old variables are set as in a feasible solution to the old LP and the new variables are set to 0, is feasible for the new LP. Therefore, the value of the LP could not have decreased. Now suppose the value of the LP increased from  $opt$  to  $opt'$ . Consider a solution that maximizes the new objective. The setting of the old variables in this solution gives a feasible solution to the old LP. So the contribution of the old variables to the objective  $\sum x_C$  is at most  $opt$ . Since  $S_i \leq \Delta_D f_H$ , the contribution of the new variables to the objective is at most  $\Delta_D f_H$ . Thus,  $opt' \leq opt + \Delta_D f_H$ , giving the first inequality.  $\square$

**Lemma 5.4.** For all graphs  $G$ , the value  $v_{\text{LP}}(G) \leq f_H(G)$ . Moreover, if  $G$  is  $D$ -bounded then  $v_{\text{LP}}(G) = f_H(G)$ .

*Proof.* The first statement follows because the objective function of the LP has  $f_H(G)$  variables. Now suppose  $G$  is a  $D$ -bounded graph. Then the solution, where  $x_C = 1$  for all variables, is feasible and maximizes the objective function. Since the number of variables is equal to the number of copies of the template  $H$  in the graph  $G$ , the value of the LP is  $f_H(G)$ .  $\square$

## 5.2 Releasing Counts of Small Subgraphs

The LP-based function from the previous section can be used to privately release small subgraph counts. If  $f_H(G)$  is relatively large then the Laplace mechanism will give an accurate estimate. Using the LP-based

function, we can release  $f_H(G)$  accurately when  $f_H(G)$  is much smaller, provided that  $G$  satisfies  $\alpha$ -decay. In this section, we work out the details of the algorithm for the special case when  $H$  has 3 nodes, i.e., is the triangle or the 2-star, but the underlying ideas apply even when  $H$  is some other small subgraph.

---

**Algorithm 2**  $\epsilon$ -Node-Private Algorithm for Releasing Subgraph Count  $f_H(G)$

---

Input: parameters  $\epsilon, D, n$ , template graph  $H$  on 3 nodes, and graph  $G$  on  $n$  nodes.

- 1: Let  $\hat{f}_1 = f_H(G) + \text{Lap}(\frac{6n^2}{\epsilon})$  and threshold  $\zeta = \frac{n^2 \ln n}{\epsilon}$ .
  - 2: If  $\hat{f}_1 \geq 7\zeta$ , **return**  $\hat{f}_1$ .
  - 3: Compute the value  $v_{\text{LP}}(G)$  given in Definition 5.1 using  $D$ .
  - 4: **return**  $\hat{f}_2 = v_{\text{LP}}(G) + \text{Lap}(\frac{6D^2}{\epsilon})$ .
- 

**Lemma 5.5.** *Algorithm 2 is an  $\epsilon$ -node-private polynomial time algorithm that takes a graph  $G$ , parameters  $\epsilon, D, n$ , and a connected template graph  $H$  on 3 nodes, and outputs an approximate count for  $f_H(G)$  (the number of copies of  $H$  in  $G$ ).*

1. If  $f_H(G) \geq (13n^2 \ln n)/\epsilon$ , then with probability at least  $1 - 1/\ln n$ , Algorithm 2 outputs  $\hat{f}_1$  and

$$|\hat{f}_1 - f_H(G)| \leq (6n^2 \ln \ln n)/\epsilon.$$

2. If  $G$  satisfies  $\alpha$ -decay for  $\alpha > 1$ ,  $D > \bar{d}$ , and  $f_H(G) < (n^2 \ln n)/\epsilon$ , then with probability at least  $1 - 2/\ln n$ , Algorithm 2 outputs  $\hat{f}_2$  and

$$|\hat{f}_2 - f_H(G)| \leq \frac{6D^2 \ln \ln n}{\epsilon} + t_h,$$

$$\text{where } t_h = \begin{cases} O(\bar{d}^\alpha n \cdot D^{2-\alpha}) & \text{if } \alpha > 2, \\ O(\bar{d}^\alpha n \cdot \ln n) & \text{if } \alpha = 2, \\ O(\bar{d}^\alpha n \cdot n^{2-\alpha}) & \text{if } 1 < \alpha < 2. \end{cases}$$

*Proof.* First, we show that Algorithm 2 is  $\epsilon$ -node-private. By Theorem 2.4, an algorithm that releases estimate  $\hat{f}_1$  is  $\epsilon/2$ -node-private because the node sensitivity of  $f_H$  by Observation 5.2 is at most  $3n^2$ . Similarly, an algorithm that releases estimate  $\hat{f}_2$  is  $\epsilon/2$ -node-private because, by Lemma 5.3, the node sensitivity of  $v_{\text{LP}}(G)$  is at most  $3D^2$ . Thus, by the composition lemma (Lemma 2.2), the algorithm that releases both estimates is  $\epsilon$ -node-private, and Algorithm 2 can be simulated by an algorithm that knows only  $\hat{f}_1$  and  $\hat{f}_2$ , but has no access to  $G$ .

Next we analyze the accuracy of the algorithm. Setting  $\delta = 1/\ln n$  in Lemma B.1, we get that with probability at least  $1 - 1/\ln n$ ,

$$|\hat{f}_1 - f_H(G)| \leq \frac{6n^2 \ln \ln n}{\epsilon} = 6\zeta. \quad (3)$$

We say the estimate  $\hat{f}_1$  is *good* if it satisfies the condition in Equation (3). If  $f_H(G) \geq (13n^2 \ln n)/\epsilon$  and  $\hat{f}_1$  is good, then  $\hat{f}_1 \geq 7\zeta$ . That is, if  $f_H(G) \geq 13\zeta$ , i.e., as in Case 1 in the lemma statement, then with probability at least  $1 - 1/\ln n$  (which is the probability that  $\hat{f}_1$  is good), Algorithm 2 outputs  $\hat{f}_1$ . Therefore, in Case 1, with probability at least  $1 - 1/\ln n$ , Algorithm 2 outputs  $\hat{f}_1$  and  $|\hat{f}_1 - f_H(G)| \leq (6n^2 \ln \ln n)/\epsilon$ .

Now we analyze Case 2 of the lemma. If  $f_H(G) < \zeta$ , then with probability at least  $1 - 1/\ln n$  (the same probability that  $\hat{f}_1$  is good),  $\hat{f}_1 < 7\zeta$  and the algorithm outputs  $\hat{f}_2$ .



With probability at least  $1 - 1/\ln n$ ,

$$|\hat{f}_2 - f_H(G)| \leq |\hat{f}_2 - v_{\text{LP}}(G)| + |v_{\text{LP}}(G) - f_H(G)| \leq \frac{6D^2 \ln \ln n}{\epsilon} + t_h,$$

where  $t_h$  is defined in Lemma 5.6. Therefore, in Case 2, with probability at least  $1 - (1/\ln n + 1/\ln n)$ , the algorithm outputs  $\hat{f}_2$  and

$$|\hat{f}_2 - f_H(G)| \leq \frac{6D^2 \ln \ln n}{\epsilon} + t_h.$$

This completes the accuracy analysis.  $\square$

**Lemma 5.6.** *If  $H$  has 3 nodes and  $G$  satisfies  $\alpha$ -decay for  $\alpha > 1$  and  $D \geq \bar{d}$  then  $v_{\text{LP}}(G) \geq f_H(G) - t_h$ , where  $t_h = O(\bar{d}^\alpha n D^{2-\alpha})$  if  $\alpha > 2$ ,  $t_h = O(\bar{d}^2 n \ln n)$  if  $\alpha = 2$ , and  $t_h = O(\bar{d}^\alpha n n^{2-\alpha})$  if  $1 < \alpha < 2$ .*

*Proof.* Let  $t_h$  denote the number of copies of  $H$  that include a node of degree greater than  $D$ . Lemma 5.4 implies that  $v_{\text{LP}}(G) \geq f_H(G) - t_h$ . It remains to bound  $t_h$ .

$$\begin{aligned} \frac{t_h}{n} &\leq \frac{1}{n} \sum_{v: \deg_v > D} \binom{\deg_v}{2} = \sum_{i \geq D} p_G(i) \binom{i}{2} = \sum_{i \geq D} p_G(i) (1 + 2 + \dots + (i-1)) \\ &= \binom{D}{2} \sum_{i \geq D} p_G(i) + \sum_{j > D} \left( \sum_{i \geq j} (i-1) p_G(i) \right) = \binom{D}{2} P_G(D) + \sum_{j > D} O\left(\frac{\bar{d}^\alpha}{j^{\alpha-1}}\right) \\ &\leq \binom{D}{2} \cdot \left(\frac{\bar{d}}{D}\right)^\alpha + O(\bar{d}^\alpha) \cdot \sum_{j > D} \frac{1}{j^{\alpha-1}} \end{aligned} \quad (4)$$

If  $\alpha > 2$ , then

$$\sum_{j > D} \frac{1}{j^{\alpha-1}} \leq \int_D^{n-1} \frac{1}{j^{\alpha-1}} dj = \frac{D^{2-\alpha} - (n-1)^{2-\alpha}}{\alpha - 2} = O(D^{2-\alpha}).$$

If  $\alpha = 2$ , then

$$\sum_{j > D} \frac{1}{j} \leq \int_D^{n-1} \frac{1}{j} dj = (\ln(n-1) - \ln D) = O(\ln n).$$

If  $1 < \alpha < 2$ , then

$$\sum_{j > D} \frac{1}{j^{\alpha-1}} \leq \int_D^{n-1} \frac{1}{j^{\alpha-1}} dj = \frac{(n-1)^{2-\alpha} - D^{2-\alpha}}{2 - \alpha} = O(n^{2-\alpha}).$$

Substituting the various values for  $\sum_{j > D} \frac{1}{j^{\alpha-1}}$  in Equation (4) completes the proof.  $\square$

Using Lemmas 5.5 and 5.6 with a carefully chosen threshold degree  $D$ , we get the following theorem about privately releasing counts of subgraphs on 3 nodes. A private value of  $\bar{d}$  can be obtained using Theorem 4.5.

**Theorem 5.7** (Releasing Subgraph Counts Privately). *There is a node differentially private algorithm which, given constants  $\alpha > 1$ ,  $\epsilon > 0$ , a connected template graph  $H$  on 3 nodes, and a graph  $G$  on  $n$  nodes, computes with probability at least  $1 - 2/(\ln n)$  an  $(1 \pm o_n(1))$ -approximation to  $f_H(G)$  (the number of copies of  $H$  in  $G$ ) if either of the following holds:*

1. If  $f_H(G) \geq (13n^2 \ln n)/\epsilon$ .
2. If  $G$  satisfies  $\alpha$ -decay, has average degree at most  $\bar{d} > 1$ , and either of the following holds: (a)  $f_H(G) = \omega(\bar{d}^2 n^{2/\alpha} \ln n)$  if  $\alpha > 2$ , (b)  $f_H(G) = \omega(\bar{d} n \ln^2 n)$  if  $\alpha = 2$ , or (c)  $f_H(G) = \omega(\bar{d}^\alpha n^{3-\alpha} \ln n)$  if  $1 < \alpha < 2$ .

*Proof.* **Case 1.** If  $f_H(G) \geq (13n^2 \ln n)/\epsilon = 13\zeta$  (where  $\zeta = (n^2 \ln n)/\epsilon$ ), then with probability at least  $1 - 1/\ln n$ , Algorithm 2 output  $\hat{f}_1$  and

$$\left| \hat{f}_1 - f_H(G) \right| \leq \frac{6n^2 \ln \ln n}{\epsilon} \equiv \frac{|\hat{f}_1 - f_H(G)|}{f_H(G)} = o_n(1) \text{ as } f_H(G) \geq 13\tau.$$

Thus, the algorithm produces an  $(1 \pm o_n(1))$ -approximation to  $f_H(G)$ .

**Case 2.** Let  $G$  satisfy  $\alpha$ -decay. Let us first assume that  $f_H(G) < \zeta$ . Set  $D$  as follows:

- i. If  $\alpha > 2$ , then set  $D = n^{1/\alpha} \bar{d}$ ,
- ii. If  $\alpha = 2$ , then set  $D = \bar{d} \sqrt{n \ln n}$ ,
- iii. If  $1 < \alpha < 2$ , then set  $D = \sqrt{\bar{d}^\alpha n^{3-\alpha}}$ .

Note that given  $\alpha$ , we can get a node-private  $(1 \pm o_n(1))$ -approximation to  $\bar{d}$  using Theorem 4.5, and use the approximation instead of the true value of  $\bar{d}$  here without changing the asymptotic accuracy.<sup>4</sup> Therefore, for simplicity, we assume that exact value of  $\bar{d}$  is known in the rest of the analysis. If  $\alpha > 2$ , then  $D > \bar{d}$  as  $n^{1/\alpha} > 1$ . If  $\alpha = 2$ , then  $D > \bar{d}$  as  $\sqrt{n \ln n} > 1$ . If  $1 < \alpha < 2$ , then  $D > \bar{d}$  as  $n^{(3-\alpha)/2} > \sqrt{n} > \bar{d}^{\alpha/2}$ . Therefore, in all the three settings  $D > \bar{d}$ . Under these conditions, Algorithm 2 with probability at least  $1 - 2/\ln n$  outputs  $\hat{f}_2$  and

$$|\hat{f}_2 - f_H(G)| \leq \frac{6D^2 \ln \ln n}{\epsilon} + t_h.$$

Substituting for  $t_h$  and  $D$  in the above equation gives

$$|\hat{f}_2 - f_H(G)| \leq \begin{cases} O(\bar{d}^2 n^{2/\alpha} \ln \ln n) & \text{if } \alpha > 2, \\ O(\bar{d} n \ln n \ln \ln n) & \text{if } \alpha = 2, \\ O(\bar{d}^\alpha n^{3-\alpha} \ln \ln n) & \text{if } 1 < \alpha < 2. \end{cases}$$

Notice that from the assumption of the theorem statement, in all the three  $\alpha$  settings,  $|\hat{f}_2 - f_H(G)| = o_n(f_H(G))$ .

If  $f_H(G) \geq 7\zeta$ , then the analysis in Case 1 implies the result, and if  $f_H(G)$  lies between  $\zeta$  and  $7\zeta$ , then Algorithm 2 returns either  $\hat{f}_1$  and  $\hat{f}_2$ , and in both cases arguments similar as above can be used to show that we get the desired approximation ratio. Thus, the algorithm produces an  $(1 \pm o_n(1))$ -approximation.  $\square$

## 6 Generic Reduction to Node Privacy in Bounded-Degree Graphs

We now turn to another, more general approach to getting more the accurate queries by looking at bounded degree graphs. Recall that if we had a promise that all degrees were at most  $D$ , then for many natural queries we could add less noise and still satisfy differential privacy. The question is, how can we enforce such a

<sup>4</sup>Since we assume that the average degree of  $G$  is greater than 1, the Case 2 of Theorem 4.5 always applies here and we can obtain a private  $(1 \pm o_n(1))$ -approximation to  $\bar{d}$ .

promise? Given an input graph  $G$ , possibly of large maximum degree, it is tempting to simply answer all queries with respect to a “truncated” version  $T(G)$ , in which nodes of very large degree have been removed. This is delicate, however, since the truncated graph  $T(G)$  may change a lot when a single node of  $G$  is changed. That is, it could be that the *local sensitivity* of the “truncation” operator (viewed as a map from  $\mathcal{G}_n$  to  $\mathcal{G}_{n,D}$ ) is very high, making queries on the truncated graph also high-sensitivity.

More generally, consider a projection operator  $T : \mathcal{G}_n \rightarrow \mathcal{G}_{n,D}$  which takes an arbitrary graph and outputs a  $D$ -bounded graph. We may define the (local, global, smooth) sensitivity of  $T$  in terms of the node distance  $d_{\text{node}}(T(G_1), T(G_2))$  where  $G_1$  and  $G_2$  differ in one node.

Given a query  $f$  defined on  $D$ -bounded graphs, it is easy to see that the local sensitivity of a *composed* query  $f \circ T$  is bounded by the product  $LS_T(G) \cdot \Delta_D f$  (one can see this as a discrete analogue of the chain rule from calculus). Our main lemma is that we can bound the smooth sensitivity similarly. We use the definition of  $\beta$ -smooth upper bound on local sensitivity from 2.6.

**Lemma 6.1** (Smooth Bounds on Composed Functions). *Let  $T : \mathcal{G}_n \rightarrow \mathcal{G}_{n,D}$ . If  $S_T(G)$  is a  $\beta$ -smooth upper bound on the local sensitivity of  $T$  (measured w.r.t. node distance), then  $S_{f \circ T}(G) = S_T(G) \cdot \Delta_D f$  is a  $\beta$ -smooth bound on the local sensitivity of  $f \circ T$ .*

Given a smooth upper bound on the local sensitivity of  $F_f \circ T$ , we can use Theorem 2.7 to obtain a private algorithm for releasing  $F_f$  on all graphs in  $\mathcal{G}_n$ .

Instead of using smooth sensitivity, we can also use a differentially private upper bound on the local sensitivity, inspired by Dwork and Lei [7] and Karwa *et al.* [12]. This give a general technique to transform any algorithm that is private on  $D$ -bounded graphs to one which is private for all graphs.

**Lemma 6.2** (Generic Reduction [12]). *Let  $T : \mathcal{G}_n \rightarrow \mathcal{G}_{n,D}$ . Suppose  $L_\epsilon$  is an  $(\epsilon, \delta_1)$ -differentially private algorithm (on all graphs in  $\mathcal{G}_n$ ) that outputs a real value such that  $\Pr[L_\epsilon(G) > LS_T(G)] \geq 1 - \delta_2$  (where  $LS_T$  is measured w.r.t. node distance).*

*Suppose that  $A$  is a  $(\epsilon, 0)_D$ -differentially private algorithm. Then the following algorithm is  $(2\epsilon, e^\epsilon \delta_2 + \delta_1)$ -differentially private: compute  $\hat{L} = L_\epsilon(G)$ , then run  $A$  on input  $T(G)$  with privacy parameter  $\epsilon' = \epsilon / \hat{L}$  and finally output the pair  $(\hat{L}, A(T(G)))$ .*

**Naive Truncation.** This is the simplest truncation operator. Consider the operator  $T_{\text{naive}}$  that deletes all nodes of degree greater than  $D$  in  $G = (V, E)$ . This may have high local sensitivity (for example, rewiring one node may change the degrees of many nodes from  $D$  to  $D + 1$ , resulting in a drastic increase in the number of nodes deleted by  $T_{\text{naive}}$ . This projector is computable in  $O(n + m)$  time, where  $n = |V|$  and  $m = |E|$ . The following simple lemma analyzes the sensitivity of this truncation operation.

**Lemma 6.3.** *Given a threshold  $D$ , the local sensitivity of naive truncation (w.r.t. node distance) is 1 plus the number of nodes with degree either  $D$  or  $D + 1$ .*

*Proof.* Consider two graphs  $G, G'$  with  $d_{\text{node}}(G, G') = 1$ . Let  $v$  be the nodes that was rewired to go from  $G$  to  $G'$ . The local sensitivity of naive truncation is the absolute difference between number of nodes with degree greater than  $D$  in  $G$  and the number of nodes with degree greater than  $D$  in  $G'$ . Since the degrees of all nodes other than  $v$  change by at most 1, the absolute difference is at most 1 (for node  $v$ ) plus the number of nodes with degree either  $D$  or  $D + 1$  in  $G$ .  $\square$

The following proposition bounds the local and smooth sensitivity of naive truncation. The last two parts of this proposition allow us to employ Lemmas 6.1 and 6.2, respectively.

**Proposition 6.4** (Bounding the Sensitivity of Naive Truncation). *Given a graph  $G$ , let  $N_k(G)$  denote the number of nodes in  $G$  with degrees in the range  $[D - k, D + k + 1]$ . Let  $C_k(G) = 1 + k + N_k(G)$ . Then*

1.  $C_0(G)$  is the local sensitivity of naive truncation at  $G$ .
2. For any graph  $G'$  within rewiring distance  $k + 1$  of  $G$ , the local sensitivity of naive truncation between  $G$  and  $G'$  is at most  $C_k(G)$ .
3.  $S_{T_{\text{naive}}}(G) = \max_{k \geq 0} e^{-\beta k} C_k(G)$  is a  $\beta$ -smooth upper bound on the local sensitivity of naive truncation. Moreover, if  $N_{\frac{1}{\beta} \ln n}(G) \leq \ell$  (that is, if there are  $\ell$  nodes in  $G$  with degrees in the range  $D \pm \ln n / \beta$ ), then

$$S_{T_{\text{naive}}}(G) \leq \ell + 1/\beta + 1.$$

4. Consider the tapered interval query given by the function  $f_{t,D,D+1}$  (defined in Section 3, Item (5)) for some  $t \in (\frac{1}{n}, 1]$ . The algorithm that returns

$$L(G) = 1 + F_{f_{t,D,D+1}}(G) + \frac{2tn \log(1/\delta)}{\epsilon} + \text{Lap}\left(\frac{2tn}{\epsilon}\right)$$

is  $(\epsilon, 0)$ -node-private and returns a value larger than  $LS_{T_{\text{naive}}}(G)$  with probability at least  $1 - \delta$ .

*Proof.* The proof of Part 1 follows from Lemma 6.3.

Part 2: By rewiring a set  $k + 1$  nodes, the degree of the remaining nodes could be changed by at most  $k + 1$ . Therefore, other than the  $k + 1$  nodes getting rewired, the only other nodes that can contribute towards the sensitivity of truncation are the nodes that have degrees between  $[D - k, D + k + 1]$  in  $G$ . The proof of the Part 2 follows.

Part 3: First, note that the statement of Part 2 is not sufficient for this part: we need to show that the bound that is smooth in the sense of Definition 2.6. The first requirement is trivial:  $S_{T_{\text{naive}}}$  is an upper bound on  $LS_{T_{\text{naive}}}$  since  $C_0(G)$  is one of the terms in the maximum. For smoothness, consider two neighboring graphs  $G$  and  $G'$ . For every  $k$ , we have  $N_k(G) \leq N_{k+1}(G') + 1$  (the rewired node contributes at most 1, and the remaining nodes change degrees by at most 1) and hence  $C_k(G) \leq C_{k+1}(G')$ . Let  $k^*$  be the value of  $k$  that maximizes the expression defining  $S_{T_{\text{naive}}}(G)$ . Then

$$S_{T_{\text{naive}}}(G) = e^{-\beta k^*} C_{k^*}(G) \leq e^\beta \cdot e^{-\beta(k^*+1)} C_{k^*+1}(G') \leq e^\beta S_{T_{\text{naive}}}(G').$$

Thus  $S_{T_{\text{naive}}}$  is a smooth upper bound on  $LS_{T_{\text{naive}}}$ .

Next suppose that  $N_{\frac{1}{\beta} \ln n}(G) \leq \ell$ . Note that

$$S_{T_{\text{naive}}}(G) \leq \left(\max_k e^{-\beta k} k\right) + \left(\max_k e^{-\beta k} C_k(G) + 1\right).$$

The first term is at most  $e^{-1}/\beta < 1/\beta$  (since the maximum of the continuous function  $g(x) = xe^{-\beta x}$  is  $e^{-1}/\beta$ ). The second term can be broken into two ranges: first if  $k > \frac{\ln n}{\beta}$ , then  $e^{-\beta k} < \frac{1}{en}$ , so  $e^{-\beta k} C_k(G) \leq \frac{n+1}{en} \leq 1$ . Second, if  $k \leq \ln n / \beta$  then  $e^{-\beta k} (N_k(G) + 1) \leq N_k(G) + 1 \leq \ell + 1$ . Putting the bounds together proves the claim.

Part 4: The linear query  $F_{f_{t,D,D+1}}(G)$  returns an upper bound on the number of nodes of degree  $D$  or  $D + 1$ . Thus,  $1 + F_{f_{t,D,D+1}}(G)$  is an upper bound on the local node sensitivity of  $T_{\text{naive}}$ , and releasing

$$1 + F_{f_{t,D,D+1}}(G) + \text{Lap}\left(\frac{2tn}{\epsilon}\right),$$

is  $(\epsilon, 0)$ -differentially private as the global sensitivity of  $1 + F_{f_t, D, D+1}(G)$  can be bounded as  $2tn$  (Section 3, Item (5)) Finally, because we add the (constant) offset  $\frac{2tn \log(1/\delta)}{\epsilon}$  to obtain  $L(G)$ , the tail bound for Laplace random variables (Lemma B.1) shows that  $L(G)$  exceeds  $1 + F_{f_t, D, D+1}(G)$  with probability at least  $1 - \delta$ .  $\square$

## 6.1 Using Naive Truncation: Deterministic and Randomized Cutoffs

The smooth sensitivity bound of Proposition 6.4 depends on the number of nodes immediately around the cutoff  $D$ . Thus, even if a graph  $G$  is  $D$ -bounded, truncating exactly at  $D$  may lead to a large smooth sensitivity bound. We get a much better bound on the noise by truncating slightly above the maximum degree. The following corollary follows by adding Cauchy noise as per Theorem 2.7.

**Corollary 6.5.** *For every  $\epsilon > 0$ , every threshold  $D > \sqrt{2}(\ln n)/\epsilon$  and every real-valued function  $f : \mathcal{G}_{n,D} \rightarrow \mathbb{R}$ , there is a  $\epsilon$ -node-private algorithm that outputs  $f(G)$  with median error  $O(\Delta_{\hat{D}} f / \epsilon^2)$ , where  $\hat{D} = D + 2 \ln(n)/\epsilon \leq 2D$ .*

**Randomizing the Degree Threshold** One obvious problem with the truncation technique is that we may not know the maximum degree in the graph, or the maximum degree may be very large. Indeed, as have seen in the algorithms for counting subgraphs, it often makes sense to project to a degree threshold well below the maximum degree in a graph. In that case, the smooth sensitivity bound of Proposition 6.4 could be large.

One can get a substantially better bound by *randomizing* the cutoff. Given a target threshold  $D$ , consider an algorithm that picks a random threshold in a range of bounded by a constant multiple of  $D$  (say, between  $2D$  and  $3D$ ). We show that the smooth sensitivity of naive truncation is (likely to be) close to the *average* number of nodes of a random degree in the range, saving a factor of roughly  $D$  in the introduced noise.

**Lemma 6.6** (Randomized Cutoff Lemma). *Fix  $\beta > 0$ , a graph  $G$  on  $n$  nodes, and an integer  $D > 0$ . Let  $P_G(D)$  be the fraction of nodes in  $G$  of degree greater than  $D$ , and let  $\hat{D}$  be uniformly random in the range  $\{2D + 1, \dots, 3D\}$ . If  $T_{naive}$  is the naive truncation at degree  $\hat{D}$ , then*

$$\mathbb{E}_{\hat{D}}[S_{T_{naive}}(G)] \leq 2 \frac{n P_G(D)}{D} \cdot \frac{\ln n}{\beta} + \frac{1}{\beta} + 1.$$

*Proof.* From Proposition 6.4, we know that  $S_{T_{naive}}(G)$  is at most  $\ell + 1 + 1/\beta$  if there are  $\ell$  nodes in the range  $\hat{D} \pm \frac{\ln n}{\beta}$ . Every node can contribute to  $\ell$  for at most  $\frac{2 \ln n}{\beta}$  choices of  $\hat{D}$  (they have to lie in the range  $\hat{D} \pm \frac{\ln n}{\beta}$  to contribute). Thus the expectation of  $\ell$  is at most  $n P_G(D) \cdot \frac{2 \ln n}{\beta D}$ , and hence the expectation of  $S_{T_{naive}}(G)$  satisfies the claimed bound.  $\square$

## 6.2 Application of Naive Truncation for Releasing Degree Distribution

For concreteness, we work out one application of the naive truncation idea to releasing an approximation to the entire degree distribution (rather than releasing specific functions of that distribution). Our goal is to output a vector  $\hat{p}$  that minimizes the  $\ell_1$ -error  $\|\hat{p} - p_G\|_1$ , where  $p_G$  is the (true) degree distribution of the graph. If the error is  $o(1)$ , then  $\hat{p}$  provides an estimate with vanishing error for all of the entries of degree distribution.

We use Lemma 6.1 to get a smooth bound on local sensitivity. The global sensitivity  $\Delta_{\hat{D}} \|\hat{p} - p_G\|_1 \leq 2\hat{D}$ .

---

**Algorithm 3**  $\epsilon$ -Node-Private Algorithm for Releasing Degree Distributions
 

---

Input: parameters  $\epsilon, D, n$ , and a graph  $G$  on  $n$  nodes.

- 1: Pick  $\hat{D} \in_R \{2D + 1, \dots, 3D\}$ .
  - 2: Compute the naive truncation  $T_{\text{naive}}(G)$  with threshold  $\hat{D}$  and the  $\beta$ -smooth bound  $S_{T_{\text{naive}}}(G)$  with  $\beta = \epsilon/(\sqrt{2}(\hat{D} + 1))$  (as in Theorem 2.7).
  - 3: Output  $\hat{p} = p_{T_{\text{naive}}(G)} + \text{Cauchy}\left(\frac{2\sqrt{2}\hat{D}}{\epsilon} S_{T_{\text{naive}}}(G)\right)^{\hat{D}+1}$  (that is, add i.i.d. Cauchy noise with median absolute deviation  $\frac{2\sqrt{2}\hat{D}}{\epsilon} S_{T_{\text{naive}}}(G)$  to the entries of the degree distribution of  $T_{\text{naive}}(G)$ ).
- 

**Theorem 6.7.** *Algorithm 3 is an  $\epsilon$ -node-private algorithm that takes a graph  $G$  and parameters  $n, D, \epsilon$ , and outputs a vector  $\hat{p}$  such that, if  $G$  satisfies  $\alpha$ -decay for  $\alpha > 1$  and  $D > \frac{4}{\epsilon} \ln n$  and  $D > \bar{d}$  where  $\bar{d} = \bar{d}(G)$  is the average degree in  $G$ , then with probability at least  $1/2$  we have*

$$\|\hat{p} - p_G\|_1 = O\left(\frac{\bar{d}^\alpha \ln n \ln(D)}{\epsilon^2 D^{\alpha-2}} + \frac{D^3 \ln(D)}{n \epsilon^2}\right) = \tilde{O}\left(\frac{1}{\epsilon^2} \left(\frac{\bar{d}^\alpha}{D^{\alpha-2}} + \frac{D^3}{n}\right)\right),$$

and the  $\tilde{O}$  notation hides constants depending on  $\alpha$  and polylogarithmic factors in  $n$ .

*Proof.* The  $\epsilon$ -differential privacy property follows from the smooth sensitivity framework (Theorem 2.7) and the fact that the global sensitivity of  $n \cdot p_G$  on  $\hat{D}$ -bounded graphs is  $2\hat{D}$ .

We now turn to the accuracy claim of the algorithm. It is more convenient to bound  $n \cdot \|\hat{p} - p_G\|_1$ . There are two sources of error: (i) the nodes that are lost due to truncation and (ii) the noise that is introduced for privacy. We bound each of these sources separately.

The truncation has two effects on the degree distribution: first, nodes with degrees above  $\hat{D}$  are removed; second, nodes adjacent to the removed nodes have their degrees decreased. We can bound the second term by bounding the number of edges adjacent to nodes of degree over  $\hat{D}$ . Following the notation of Lemma 2.12, and applying the bounds there, the contribution to  $n \cdot \|\hat{p} - p_G\|_1$  due to truncation is at most

$$\underbrace{n P_G(\hat{D})}_{\text{removed nodes}} + \underbrace{\sum_{i > \hat{D}} i p_G(i)}_{\text{edges adjacent to removed nodes}} \leq n \frac{\bar{d}^\alpha}{\hat{D}^\alpha} + O\left(\frac{n \bar{d}^\alpha}{\hat{D}^{\alpha-1}}\right) = O\left(\frac{n \bar{d}^\alpha}{D^{\alpha-1}}\right). \quad (5)$$

Here we used the fact that  $\hat{D} = O(D)$ .

Now we analyze the contribution to  $n \cdot \|\hat{p} - p_G\|_1$  due to Cauchy noise addition. Let

$$\lambda = (\sqrt{2}(2\hat{D})/\epsilon) S_{T_{\text{naive}}}(G).$$

We need to bound the deviation of  $\sum_{i=1}^{\hat{D}+1} |\text{Cauchy}(\lambda)|$ . Let  $X_1, \dots, X_{\hat{D}+1}$  be  $\hat{D} + 1$  i.i.d. Cauchy(1) random variables. Firstly note that for any positive  $\lambda$ , the probability distribution of  $\sum_{i=1}^{\hat{D}+1} |\text{Cauchy}(\lambda)|$  is same as the probability distribution of  $\lambda \sum_{i=1}^{\hat{D}+1} |X_i|$  (since if  $X \sim \text{Cauchy}(1)$  then, for  $\lambda > 0$ ,  $\lambda X \sim \text{Cauchy}(\lambda)$ ). Here,  $\lambda$  itself is a random variable. By the randomized cutoff lemma (Lemma 6.6), the expected value of  $\lambda$  is

$$\begin{aligned} \mathbb{E}[\lambda] &= \frac{\sqrt{2}(2\hat{D})}{\epsilon} \mathbb{E}[S_{T_{\text{naive}}}(G)] \\ &\leq \frac{\sqrt{2}(2\hat{D})}{\epsilon} \left(2n \cdot P_G(D) \cdot \frac{\ln n}{\beta D} + \frac{1}{\beta} + 1\right) = O\left(n \cdot \frac{\bar{d}^\alpha}{D^{\alpha-1}} \cdot \frac{\ln n}{\epsilon^2} + \frac{D^2}{\epsilon^2}\right), \end{aligned}$$

where we are using the facts that  $\hat{D} = O(D)$  and  $\beta = \epsilon/(\sqrt{2}(\hat{D} + 1))$  to simplify the expression.

Now by Markov inequality (where the probability is only over the choice of  $\hat{D}$ ),  $\Pr_{\hat{D}}(\lambda \geq 4\mathbb{E}[\lambda]) \leq 1/4$ . Let  $E_1$  denote the event that  $\lambda \leq 4\mathbb{E}[\lambda]$ .

It remains to bound the random variable  $\sum_{i=1}^{\hat{D}+1} |X_i|$ . This variable does not have a finite expectation but does satisfy some weak concentration bound: By Lemma B.2 (with  $\delta = 1/4$ ), we get that for a sufficiently large  $\hat{D}$ ,

$$\Pr_{X_1, \dots, X_n} \left[ \sum_{i=1}^{\hat{D}+1} |X_i| \geq 8(\hat{D} + 1) \ln(4(\hat{D} + 1)) \right] \leq \frac{1}{4}$$

Let  $E_2$  denote the event that  $\sum_{i=1}^{\hat{D}+1} |X_i| \leq 8(\hat{D} + 1) \ln(4(\hat{D} + 1))$ . Conditioned on event  $E_1$  and  $E_2$  the  $\ell_1$ -error added for privacy is at most

$$\begin{aligned} \lambda \cdot \sum_{i=1}^{\hat{D}+1} |X_i| &= O \left( n \cdot \frac{\bar{d}^\alpha}{D^{\alpha-1}} \cdot \frac{\ln n}{\epsilon^2} + \frac{D^2}{\epsilon^2} \right) \cdot 8(\hat{D} + 1) \ln(4(\hat{D} + 1)) \\ &= O \left( \frac{n \bar{d}^\alpha \ln n \ln(D)}{\epsilon^2 D^{\alpha-2}} + \frac{D^3 \ln(D)}{\epsilon^2} \right). \end{aligned}$$

Removing the conditioning we get with probability at least  $1/2$ , the contribution to  $n \cdot \|\hat{p} - p_G\|_1$  from the noise added for privacy is at most

$$O \left( \frac{n \bar{d}^\alpha \ln n \ln(D)}{\epsilon^2 D^{\alpha-2}} + \frac{D^3 \ln(D)}{\epsilon^2} \right).$$

Combining the  $\ell_1$ -error for truncation (Equation (5)) and the above bound on the  $\ell_1$ -error for privacy completes the proof of the theorem.  $\square$

We note that one can get slightly better bounds on the error by considering an algorithm that uses different noise distributions other than Cauchy. We stick to Cauchy noise here for simplicity. For the following corollary, we set  $D = \bar{d}^{\frac{\alpha}{\alpha+1}} n^{\frac{1}{\alpha+1}}$  in the previous theorem.

**Corollary 6.8** (Releasing Degree Distribution Privately). *There is a node differentially private algorithm running in  $O(|E|)$  time which, given  $\alpha > 1$ ,  $\epsilon > 0$ , and a graph  $G = (V, E)$  on  $n$  nodes, computes an approximate degree distribution with  $\ell_1$  error (with probability at least  $1/2$ )*

$$\|\hat{p} - p_G\|_1 = \tilde{O} \left( \bar{d}^{\frac{3\alpha}{\alpha+1}} / \left( \epsilon^2 n^{\frac{\alpha-2}{\alpha+1}} \right) \right)$$

if  $G$  satisfies  $\alpha$ -decay and has average degree at most  $\bar{d} > 1$ . In particular, this error goes to 0 for any constant  $\alpha > 2$  when  $\bar{d}$  is polylogarithmic in  $n$ .

*Proof.* If we know the average degree  $\bar{d}$ , we can set  $D = \bar{d}^{\frac{\alpha}{\alpha+1}} n^{\frac{1}{\alpha+1}}$  in the Theorem 6.7 to get the desired accuracy bound. Note that  $D > \bar{d}$  (since  $n > \bar{d}$ ).

The algorithm does not receive  $\bar{d}$  as input. However, given  $\alpha$ , we can get a node-private  $(1 \pm o_n(1))$ -approximation to  $\bar{d}$  using Theorem 4.5, and use the approximation instead of the true value of  $\bar{d}$  here without changing the asymptotic accuracy.  $\square$

## Acknowledgments

We thank Madhav Jha for pointing out an error in an earlier version of the Randomized Cutoff Lemma.

## References

- [1] BLOCKI, J., BLUM, A., DATTA, A., AND SHEFFET, O. Differentially Private Data Analysis of Social Networks via Restricted Sensitivity. In *ITCS (To appear)* (2013).
- [2] BLUM, A., DWORK, C., MCSHERRY, F., AND NISSIM, K. Practical Privacy: The SuLQ Framework. In *PODS (2005)*, ACM, pp. 128–138.
- [3] CLAUSET, A., SHALIZI, C. R., AND NEWMAN, M. E. J. Power-Law Distributions in Empirical Data. *SIAM Review* 51, 4 (2009), 661–703.
- [4] DINUR, I., AND NISSIM, K. Revealing Information While Preserving Privacy. In *PODS (2003)*, ACM, pp. 202–210.
- [5] DWORK, C. Differential Privacy. In *ICALP (2006)*, LNCS, Springer, pp. 1–12.
- [6] DWORK, C., KENTHAPADI, K., MCSHERRY, F., MIRONOV, I., AND NAOR, M. Our Data, Ourselves: Privacy Via Distributed Noise Generation. In *EUROCRYPT (2006)*, LNCS, Springer, pp. 486–503.
- [7] DWORK, C., AND LEI, J. Differential Privacy and Robust Statistics. In *STOC (2009)*, pp. 371–380.
- [8] DWORK, C., MCSHERRY, F., NISSIM, K., AND SMITH, A. Calibrating noise to sensitivity in private data analysis. In *TCC (2006)*, LNCS, Springer, pp. 265–284.
- [9] GEHRKE, J., LUI, E., AND PASS, R. Towards Privacy for Social Networks: A Zero-Knowledge Based Definition of Privacy. In *TCC (2011)*, pp. 432–449.
- [10] HAY, M., LI, C., MIKLAU, G., AND JENSEN, D. Accurate Estimation of the Degree Distribution of Private Networks. In *ICDM (2009)*, pp. 169–178.
- [11] JERNIGAN, C., AND MISTREE, B. F. T. Gaydar: Facebook Friendships Expose Sexual Orientation. *First Monday* 14, 10 (2009).
- [12] KARWA, V., RASKHODNIKOVA, S., SMITH, A., AND YAROSLAVTSEV, G. Private analysis of graph structure. *PVLDB* 4, 11 (2011), 1146–1157.
- [13] KIFER, D., AND MACHANAVAJHALA, A. No Free Lunch in Data Privacy. In *SIGMOD (2011)*, pp. 193–204.
- [14] MCSHERRY, F., AND MIRONOV, I. Differentially Private Recommender Systems: Building Privacy into the Net. In *KDD (2009)*, ACM New York, NY, USA, pp. 627–636.
- [15] NISSIM, K., RASKHODNIKOVA, S., AND SMITH, A. Smooth sensitivity and sampling in private data analysis. In *Symp. Theory of Computing (STOC) (2007)*, ACM, pp. 75–84. Full paper: <http://www.cse.psu.edu/~asmith/pubs/NRS07>.
- [16] ORLIN, J. B. Max flows in  $O(nm)$  time, or better, 2012. [http://jorlin.scripts.mit.edu/docs/papersfolder/O\(nm\)MaxFlow.pdf](http://jorlin.scripts.mit.edu/docs/papersfolder/O(nm)MaxFlow.pdf).
- [17] RASTOGI, V., HAY, M., MIKLAU, G., AND SUCIU, D. Relationship Privacy: Output Perturbation for Queries with Joins. In *PODS (2009)*, pp. 107–116.



## A Comparison to Concurrent Work

Blocki *et al.* [1] provide algorithm for analyzing graph data with node-level differential privacy. They proceed from a similar intuition to ours, developing low-sensitivity projections onto the set of graphs of a given maximum degree. However, the results of the two papers are not directly comparable. This section discusses the differences between the two works.

Specifically, Blocki *et al.* have two main results on node privacy, both of which are incomparable to our corresponding results.

- First, Blocki *et al.* show that for every function  $f : \mathcal{G}_{n,D} \rightarrow \mathbb{R}$ , there exists an extension  $g : \mathcal{G}_n \rightarrow \mathbb{R}$  that agrees with  $f$  on  $\mathcal{G}_{n,D}$  and that has global sensitivity  $\Delta g = \Delta_D f$ . The resulting function need not be computable efficiently.

In contrast, we give explicit, *efficient* constructions of such extensions for several families of functions (the number of edges, linear functions of the degree distribution defined by concave queries, and subgraph counting queries).

- Second, Blocki *et al.* give a specific projection from arbitrary graphs to graphs of a particular degree  $\mu : \mathcal{G}_n \rightarrow \mathcal{G}_{n,D}$ , along with a smooth upper bound on its local sensitivity. They propose to use this for answering queries which have low node sensitivity on  $\mathcal{G}_{n,D}$ .

We give a similar result for a different projection (naive truncation). As in their work, we propose to compose this projection with queries that have low sensitivity when restricted to graphs of bounded degree (Lemma 6.1), though we also observe that more general types of composition are also possible (Lemma 6.2).

The results for these different projections are similar in that both techniques have low smooth sensitivity (depending only on  $\epsilon$ ) when the input graph has degree less than the input threshold  $D$ .

To the best of our understanding, the accuracy results are nevertheless incomparable. The Blocki *et al.* projection has a bicriteria approximation guarantee: on input  $D$  and  $G$ , their projection function is guaranteed to output a graph of degree at most  $D$  such that the distance  $d_{\text{node}}(G, \mu(G)) \leq 4d_{\text{node}}(G, \mathcal{G}_{n,D/2})$ . (No such guarantee is possible for naive truncation, which may be arbitrarily worse than the optimal projection even onto graphs of degree smaller than  $D$ .) Nonetheless, the sensitivity bound for  $\mu$  can be quite a bit higher than the one we present for naive truncation, resulting in lower noise added for privacy (similarly, there are graphs for which the other projection is less sensitive).

Our approach has a considerable efficiency advantage: the naive truncation procedure we propose runs in  $O(n + m)$  time for a graph with  $n$  vertices and  $m$  edges, whereas the projection of Blocki *et al.* seems to require solving a linear program with  $n + \binom{n}{2}$  variables and  $\Theta(n^2)$  constraints.

The final accuracy guarantees for our algorithms are stated for graphs that satisfy a mild tail bound on the degree distribution, called  $\alpha$ -decay. In contrast, Blocki *et al.* only give accuracy guarantees for graphs with bounded degree.

Finally, Blocki *et al.* also consider *edge* privacy, and give a simple, elegant projection operator that has constant edge sensitivity. There is no analogue of that result in this paper, which focuses on node privacy.

## B Useful Tail Bounds

We make use of the following tail inequalities. The first is standard:

**Lemma B.1** (Laplace Tail Bounds). *For all  $\delta \in (0, 1)$ , with probability at least  $1 - \delta$ , the magnitude of the Laplace noise  $|\text{Lap}(\lambda)| \leq \lambda \cdot \ln(1/\delta)$ .*

The second bound is straightforward, but we include a proof for completeness.

**Lemma B.2** (Sums of Absolute Values of Cauchy Random Variables). *There exists an absolute constant  $c$  such that: If  $X_1, \dots, X_n$  are i.i.d. Cauchy(1) random variables, then for all  $\delta < \frac{1}{2}$  and for  $n > c$ ,*

$$\Pr \left[ \sum_{i=1}^n |X_i| \geq \frac{n}{\delta} \left( 3 + \ln \left( \frac{n}{\delta} \right) \right) \right] \leq \delta.$$

The bound of  $n(3 + \ln(n/\delta))/\delta$  can be simplified to  $2n \ln(n/\delta)/\delta$  by increasing  $c$ .

*Proof.* Let  $A$  be the event that all of the  $|X_i|$  are at most  $N \stackrel{\text{def}}{=} \frac{4n}{3\delta}$ .

**Claim B.3.** *For sufficiently large  $n$ , we have  $\Pr[A] \geq 1 - \delta/2$ .*

*Proof.* The c.d.f. of the Cauchy distribution is  $\frac{1}{2} + \frac{\arctan(x)}{\pi}$ , so the probability that any given  $|X_i|$  exceeds a threshold  $x$  is  $\frac{1}{2} - \frac{\arctan(x)}{\pi}$ . When  $x = \frac{4n}{3\delta}$ , we can take the limit as  $n/\delta$  grows to see that the tail probability scales as  $\frac{3\delta}{2\pi n} + O(\delta^2/n^2)$ , which is larger than  $\delta/2n$  for large enough  $n$  (and  $\delta < 1/2$ ). By a union bound, the probability that all of the  $|X_i|$  lie below  $N$  is at least  $1 - \delta/2$ , as desired.  $\square$

**Claim B.4.** *Conditioned on  $A$ ,  $\mathbb{E}[|X_i|] \leq \frac{2+\ln N}{2}$  for each  $i = 1, \dots, n$ .*

*Proof.* Conditioned on  $A$ , each  $|X_i|$  has finite expectation

$$\begin{aligned} \mathbb{E}(|X_i| | A) &\leq \frac{1}{\Pr[A]} \int_0^N x \cdot \frac{2}{\pi(1+x^2)} dx \\ &= \frac{1}{\Pr[A]} \left( \int_0^1 \frac{2x}{\pi(1+x^2)} dx + \int_1^N \frac{2x}{\pi(1+x^2)} dx \right) \\ &\leq \frac{1}{\Pr[A]} \left( \int_0^1 \frac{2}{\pi} dx + \int_1^N \frac{2x}{\pi(2x^2)} dx \right) = \frac{1}{\Pr[A]} \left( \frac{\ln(N)}{\pi} + \frac{2}{\pi} \right) \leq \frac{\ln N + 2}{2}, \end{aligned}$$

where the last inequality uses the fact that  $\delta < 1/2$  (so that  $\pi(1 - \delta/2) \geq 2$ ).  $\square$

We can combine these two statements to get the desired result: by Markov's inequality, the probability that  $\sum_i |X_i|$  exceeds  $n(\ln(N) + 2)/\delta$  is at most  $\delta/2$  conditioned on  $A$ . Thus, the total probability that the sum exceeds this bound is  $\Pr[A] + (1 - \Pr[A])\delta/2 \leq \delta$ . Plugging in the value of  $N$  gives the desired result.  $\square$

## C Other Truncation Operators

We consider several other natural truncation operators  $T$ :

- *Minimum Node-Deletion:* Let  $T_{\min\text{-node}}(G)$  be the largest induced subgraph of  $G$  which has maximum degree at most  $D$  (there may be several such subgraphs of the same size; we define  $T_{\min\text{-node}}$  to output an arbitrary one).  $T_{\min\text{-node}}(G)$  is a *projection* of  $G$  onto  $\mathcal{G}_{n,D}$  in the node metric, that is, it outputs a graph that minimizes the node distance  $d_{\text{node}}(G, T_{\min\text{-node}}(G))$ . In general, this operator never deletes more nodes than naive truncation, but the minimum node-deletion projector may delete fewer nodes than naive truncation, since a first deletion may reduce the degrees of other nodes, obviating the need to delete them.

In contrast with naive truncation, computing the minimum node-deletion is NP-hard in general (for example, projecting onto the set of degree-0 graphs corresponds to finding a minimum vertex cover). Moreover, the projection may have large local sensitivity (see Proposition C.1). The projector still has nice properties: The *number* of nodes deleted by this method has low sensitivity: it can go up or down by at most 1 when we rewire 1 node. And, as with naive truncation, we can come up with useful bounds on the local (and smooth) sensitivity of this operator.

- *Minimum Edge-Deletion:* Another natural approach would be to remove the minimum number of edges that leave a graph with maximum degree  $D$ . Perhaps surprisingly, this truncation operator can be computed in polynomial time, via a generalization of Edmonds' algorithm for matching in general graphs. However, the set of edges it removes may vary a lot with small changes in the graph, and we do not know of a good analysis of its sensitivity.

Below, we give explicit upper bounds on the local and smooth sensitivity of minimum node-deletion. These bounds allow us to use them in the reductions of Lemmas 6.1 and 6.2. These bounds are easy to compute given the projection  $T_{\min\text{-node}}(G)$  (though, as noted earlier, the projection is NP-hard to compute in general). One drawback of this approach is that the bound on the local sensitivity equals (twice) the distance between  $G$  and  $T_{\min\text{-node}}(G)$ , which may be large in absolute value. In contrast, the sensitivity of naive truncation depends only on the number of nodes with degrees near the threshold.

**Proposition C.1.** *Let  $T_{\min\text{-node}}$  be the minimum node-deletion projector, and let  $d_{\text{node}}(G, \mathcal{G}_{n,D})$  denote the node distance between  $G$  and  $T_{\min\text{-node}}(G)$ . Then we have:*

1. *The local sensitivity of minimum node-deletion projector is at most  $2(d_{\text{node}}(G, \mathcal{G}_{n,D}) + 1)$ .*
2. *For any graph  $G'$  within rewiring distance  $k + 1$  of  $G$ , the local sensitivity of minimum node-deletion is at most  $2(d_{\text{node}}(G, \mathcal{G}_{n,D}) + k + 1)$ .*
3.  *$S_{T_{\min\text{-node}}}(G) = 2 \max_{k \geq 0} e^{-\beta k} (d_{\text{node}}(G, \mathcal{G}_{n,D}) + 1 + k)$  is a  $\beta$ -smooth upper bound on the local sensitivity of  $T_{\min\text{-node}}$ . Moreover, if  $d_{\text{node}}(G, \mathcal{G}_{n,D}) + 1 > 1/\beta$ , then  $S_{T_{\min\text{-node}}}(G) = 2(d_{\text{node}}(G, \mathcal{G}_{n,D}) + 1)$ .*
4. *The algorithm,*

$$L(G) = 2(d_{\text{node}}(G, \mathcal{G}_{n,D}) + 1) + \text{Lap}\left(\frac{2}{\epsilon}\right) + \frac{2 \log(1/\delta)}{\epsilon}$$

*is  $(\epsilon, 0)$ -differentially private and returns a value that is greater than the local sensitivity of  $T_{\min\text{-node}}$  on  $G$  with probability at least  $1 - \delta$ .*

*Proof.* Part 1: As noted earlier, the distance  $d_{\text{node}}(G, \mathcal{G}_{n,D})$  has node sensitivity 1 (as does distance to any particular set of inputs [8, Sec. 3.2]). If  $G$  and  $G'$  differ in a single node, then by the triangle inequality

$$\begin{aligned} & d_{\text{node}}(T_{\text{min-node}}(G), T_{\text{min-node}}(G')) \\ & \leq d_{\text{node}}(T_{\text{min-node}}(G), G) + d_{\text{node}}(G, G') + d_{\text{node}}(G', T_{\text{min-node}}(G')) \\ & \leq d_{\text{node}}(G, \mathcal{G}_{n,D}) + 1 + (d_{\text{node}}(G, \mathcal{G}_{n,D}) + 1) = 2d_{\text{node}}(G, \mathcal{G}_{n,D}) + 2. \end{aligned}$$

Given Part 1, the proofs of Parts 2, 3 and 4 are similar to the analogous statements for the naive truncation (Proposition 6.4). For Part 4, note that the global sensitivity of  $2d_{\text{node}}(G, \mathcal{G}_{n,D}) + 2$  is just 2, therefore, releasing

$$2d_{\text{node}}(G, \mathcal{G}_{n,D}) + 2 + \text{Lap}\left(\frac{2}{\epsilon}\right),$$

is  $(\epsilon, 0)$ -differentially private. □

The last two parts of this proposition allow us to employ Lemmas 6.1 and 6.2, respectively.