Node-Differentially Private Estimation of the Number of Connected Components

Iden Kalemaj
Boston University, US
ikalemaj@bu.edu

Sofya Raskhodnikova
Boston University, US
sofya@bu.edu

Adam Smith
Boston University, US
ads22@bu.edu

Charalampos E. Tsourakakis
Boston University, US
tsourolampis@gmail.com

ABSTRACT

We design the first node-differentially private algorithm for approximating the number of connected components in a graph. Given a database representing an n-vertex graph G and a privacy parameter ε, our algorithm runs in polynomial time and, with probability $1 - o(1)$, has additive error $\widetilde{O}(\frac{\Delta^* \ln \ln n}{\varepsilon})$, where Δ^* is the smallest possible maximum degree of a spanning forest of G. Node-differentially private algorithms are known only for a small number of database analysis tasks. A major obstacle for designing such an algorithm for the number of connected components is that this graph statistic is not robust to adding one node with arbitrary connections (a change that node-differential privacy is designed to hide): *every* graph is a neighbor of a connected graph.

We overcome this by designing a family of efficiently computable Lipschitz extensions of the number of connected components or, equivalently, the size of a spanning forest. The construction of the extensions, which is at the core of our algorithm, is based on the forest polytope of G. We prove several combinatorial facts about spanning forests, in particular, that a graph with no induced Δ-stars has a spanning forest of degree at most Δ. With this fact, we show that our Lipschitz extensions for the number of connected components equal the true value of the function for the largest possible monotone families of graphs. More generally, on all monotone sets of graphs, the ℓ_∞ error of our Lipschitz extensions is nearly optimal.

1 INTRODUCTION

Counting the number of features in a graph, ranging from local structures such as edges, triangles, wedges and other small motifs [MSOI+02, XWC+22, CMMT22] to more global features such as the number of connected components [Fra78, KMAF10, KLM+14, KW20], is a family of foundational tasks in graph mining. Such counts are used both in feature-based approaches to anomaly detection in graphs [ATK15] and as quality criteria for generative models [XWC+22]. Specifically, the number of connected components is a basic statistic for understanding the structure of a network at the most fundamental level. Its uses range from determining the number of classes in a population [Goo49] to approximating such diverse statistics as the weight of the minimum spanning tree [CRT05] and the number of documented deaths in the Syrian war [CSS18].

When the underlying network data captures private relationships between people, publishing the number of connected components in the network can leak those individuals' sensitive information. *Differential privacy* [DMNS16] is a widely studied and deployed formal privacy guarantee for data analysis. The output distributions of a differentially private algorithm must look nearly indistinguishable for any two input databases that differ only in the data of a single individual. In this work, we focus on relational databases that store graphs where nodes represent individuals and edges capture relationships between them. There are two natural adaptations of differential privacy for graph databases: *edge differential privacy* and *node differential privacy* (or, more concisely, *edge-privacy* and *node-privacy*) [HLMJ09]. For edge-privacy, first investigated by Nissim et al. [NRS07], the indistinguishability requirement applies to any two graphs that differ in one edge. In contrast, for node-privacy (first studied by three concurrent works [BBDS13, KNRS13, CZ13]), it applies to any two graphs that differ in one node and all its adjacent edges. Node-privacy is more suitable for databases representing networks of people, since, in this context, it protects each individual (node) and all their connections (adjacent edges). However, node-privacy is much harder to attain than edge-privacy, since it hides much larger changes to the database and provides much stronger privacy protections.

We design the first node-private algorithm for approximating the number of connected components in a graph database. We show that our algorithm runs in polynomial time and give an upper bound on its additive error. The bound is low for real-world networks and typical graphs arising from random network models.

Connectivity of graph databases has not been previously investigated in the context of node-privacy. Let $f_{\mathrm{cc}}(G)$ denote the number of connected components in a graph G. Obtaining an edge-private algorithm for approximating f_{cc} is trivial. Without privacy guarantees, approximating f_{cc} with additive error has been studied in the literature on sublinear-time algorithms [CRT05, BKM14, KW20].

Instance-based accuracy. A major obstacle for designing a node-private algorithm for f_{cc} is that this graph statistic is sensitive to adding one node with arbitrary connections: *every* graph G can be changed to a connected graph G' by adding a new node adjacent to all nodes in G. By definition, the output distributions of a node-private algorithm $\mathcal{A}$ must be nearly indistinguishable on G and G'. If $f_{\mathrm{cc}}(G)$ is large, $\mathcal{A}$ must be

inaccurate on at least one of G and G'. Thus, meaningful worst case guarantees for approximating f_{cc} are impossible.

To overcome this lower bound, we provide instance-based accuracy guarantees. Our bound on the error is roughly Δ^*, the smallest possible maximum degree of a spanning forest of the input database G. In particular, Δ^* is always at most the maximum degree of the graph, though it can be much smaller. Previous works on node-private algorithms either do not analyze accuracy at all or analyze it through lenses that are too coarse (for example, by looking at the maximum degree) or too task-specific for our purposes; see the discussion in Section 1.2.

The core of our algorithm is a construction of a family of efficiently computable *Lipschitz extensions* of f_{cc} that provide good approximations to the true function values.

Background on Lipschitz extensions. The framework of Lipschitz extensions for designing node-private algorithms was proposed by three concurrent works [BBDS13, KNRS13, CZ13] and subsequently refined [RS16b, RS16a].

DEFINITION 1.1 (NODE-NEIGHBORING GRAPHS; NODE-DISTANCE; Δ-LIPSCHITZ FUNCTIONS). *Two graphs G and G' are* node-neighbors *if one can be obtained from the other by removing a vertex and all of its adjacent edges.*

The node-distance $d(G, G')$ *between two graphs G and G' is the smallest number of modifications needed to obtain G' from G, where each modification is either the removal of one node and all of its adjacent edges, or the insertion of one node with arbitrary edges incident on it.*

Let $\mathcal{G}$ denote the set of all graphs. A function $h \colon \mathcal{G} \to \mathbb{R}$ is Δ-Lipschitz on a set S of graphs if $|h(G) - h(G')| \le \Delta \cdot d(G, G')$ for all $G, G' \in S$. When $S = \mathcal{G}$, we refer to f as Δ-Lipschitz, without specifying S.

A Δ-Lipschitz function can be approximated node-privately with error $\Theta(\Delta)$ (with high probability and for constant privacy parameter) using the standard Laplace mechanism (see Theorem 2.2).

Given a function f and a parameter Δ, a Lipschitz extension f_Δ is defined by first identifying an *anchor set* of graphs on which f is Δ-Lipschitz. Then f_Δ is made identical to f on the anchor set and is extended outside the set so that it is Δ-Lipschitz on $\mathcal{G}$. Instead of approximating the sensitive function f, a node-private algorithm $\mathcal{A}$ can release f_Δ (for an appropriately chosen Δ) via the Laplace mechanism. The error of $\mathcal{A}$ on a graph G is $|f(G) - f_\Delta(G)|$ plus the noise added for privacy, which is $\Theta(\Delta)$. The Generalized Exponential Mechanism (GEM) [RS16b] can be used to select the value of Δ that approximately minimizes the expected error. The resulting algorithm has small error on graphs G for which there is a small Δ such that $f_\Delta(G) \approx f(G)$.

Lipschitz extensions can be defined for any metric space. A line of work in functional analysis seeks to understand general conditions under which, given a function f and a subset of its domain, a Lipschitz extension to the entire domain exists [BL98]. Such an extension always exists for real-valued f [McS34]. However, it is not known to be efficiently computable even when f is computable in polynomial time. The main challenge for utilizing the framework of Lipschitz extensions lies in designing efficiently computable extensions that match the desired function on a large subset of the domain.

Our Lipschitz extensions for f_{cc}. Our Lipschitz extensions for the number of connected components are computable in polynomial time. Given a parameter Δ and a graph G, the value of our Lipschitz extension $f_\Delta(G)$ is obtained by solving a linear program over the forest polytope of G. Let S_Δ be the anchor set of f_Δ, i.e., the set of graphs G with $f_\Delta(G) = f_{cc}(G)$. Then S_Δ contains all graphs that have a spanning forest of degree at most Δ. This feature of our construction allows us to bound the error of the algorithm by the smallest possible maximum degree of a spanning forest of the graph.

We show that our construction of Lipschitz extensions is nearly optimal. We consider two notions of optimality. First, we compare our sets S_Δ to the largest possible anchor sets. To make optimal anchor sets well defined, we require them to be *monotone*, i.e., if G is in the set, so are all of its induced subgraphs. As it turns out, the largest monotone anchor set is unique and characterized in terms of the *down-sensitivity* of f, a quantity introduced by the name of "empirical global sensitivity" by Chen and Zhou [CZ13]. The down-sensitivity of f at a graph G, denoted $DS_f(G)$, measures the maximum change in the value of f between any two node-neighboring induced subgraphs of G. Let $S_\Delta^* = \{G \mid DS_f(G) \le \Delta\}$. Then S_Δ^* is the largest monotone anchor set for the function f and parameter Δ. A Lipschitz extension for the set S_Δ^* exists for all functions f (it is defined and analyzed in Lemma A.3). In general, the evaluation of the extension takes exponential time even if f is efficiently computable. We show that $S_{\Delta-1}^* \subseteq S_\Delta$, i.e., the anchor set for our Lipschitz extension with parameter Δ contains the largest monotone anchor set for parameter $\Delta - 1$. When the Lipschitz extension is used in our algorithm, that difference of 1 barely changes the noise added for privacy. Moreover, our construction runs in polynomial time.

The second notion of optimality we consider is more general. It was introduced by Cummings and Durfee [CD20]. They study the ℓ_∞ error of the extension, measured locally over all induced subgraphs of any given graph. According to this measure, the error of our Lipschitz extension f_Δ on induced subgraphs of a graph G is at most twice the error of the best Lipschitz extension with parameter $\Delta - 1$ on the same set of graphs. The best extension does not have to be efficiently computable and is selected separately for each G.

1.1 Our Results

We consider databases that represent undirected, unweighted graphs G with vertex set $V(G)$ and edge set $E(G)$. We start by formally defining the privacy notion of our algorithm.

DEFINITION 1.2 (NODE-PRIVACY). *A randomized algorithm $\mathcal{A}$ is ε-node-private if for all node-neighboring graphs G, G' and all events S in the output space of $\mathcal{A}$,*

$$\Pr[\mathcal{A}(G) \in S] \le e^\varepsilon \Pr[\mathcal{A}(G') \in S].$$

Our first idea is to restate the problem of estimating f_{cc}, the number of connected components, in terms of estimating the size of a spanning forest. Let $f_{sf}(G)$ denote the number of

edges in a spanning forest of G. The reason f_{sf} is an easier function to work with is that it is monotone (nondecreasing) under the addition of nodes and edges, whereas f_{cc} can increase or decrease when we insert a new node with arbitrary edges. The relationship between the two functions is straightforward:

$$f_{\text{cc}}(G) = |V(G)| - f_{\text{sf}}(G). \tag{1}$$

The number of nodes in a graph can be easily estimated with an ε-node-private algorithm with additive error $O(1/\varepsilon)$ by using the standard Laplace mechanism. Henceforth, we state our results in terms of approximating f_{sf}, with the understanding that it is equivalent to approximating f_{cc} with additive error. Our main result is a polynomial-time, ε-node-private algorithm for approximating f_{sf}. Its error is upper bounded by $O(\Delta^* \ln \ln n)$ for constant ε.

Theorem 1.3 (Node-private algorithm for the size of the spanning forest). *There exists an ε-node-private algorithm $\mathcal{A}$ that, given an n-node graph G and a parameter $\varepsilon > 0$, runs in polynomial time. Let Δ^* be the smallest possible maximum degree of a spanning forest of G. If $E(G) \neq \emptyset$, then, with probability $1 - o(1)$, the output of the algorithm satisfies*

$$|\mathcal{A}(G) - f_{\text{sf}}(G)| \leq \Delta^* \cdot \widetilde{O}\Big(\frac{\ln \ln n}{\varepsilon}\Big). \tag{2}$$

Note that Δ^* is at most the maximum degree of G. Therefore, the error bound in (2) also holds when Δ^* is replaced with the maximum degree of G.

1.1.1 Our Lipschitz extensions for f_{sf}. Our main technical contribution and the key tool in our algorithm design is a construction of a family of Lipschitz extensions for the size of the spanning forest. Recall from Section 1 that the first challenge in designing a family of Lipschitz extension is identifying their anchor sets. The anchor set of our Lipschitz extension f_Δ contains the set of graphs that have a spanning forest of degree at most Δ (called a spanning Δ-forest). One can easily see that f_{sf} is Δ-Lipschitz on this set.

The construction of f_Δ is based on the forest polytope of the input graph. The vertices of the polytope, represented by vectors $x \in \{0, 1\}^E$, correspond to forest subgraphs of G, where the weight of an edge is an indicator for whether the edge is in the forest. The forest polytope contains all convex combinations of the forest subgraphs of G. We consider the degree-bounded forest polytope of a graph, parameterized by Δ. It introduces the additional requirement that the total weight of the edges incident on each vertex is at most Δ. Then $f_\Delta(G)$ equals the total maximum weight that can be assigned to the edges of G, while ensuring that the vector of weights is in the Δ-bounded forest polytope. We show that f_Δ is Δ-Lipschitz. It is also computable in polynomial time using an LP solver and an efficient linear separation oracle.

Our extensions and their analysis are described in Section 3.

1.1.2 Relation to Down-sensitivity. The *down-sensitivity* of a function, defined next, provides a useful lens for thinking about instance-dependent guarantees [CZ13, RS16a].

Definition 1.4 (Down-sensitivity [CZ13]). *For two graphs H, H', we write $H \preceq H'$ (or $H \prec H'$) if H is an induced (or proper*

induced) subgraph of H'. The down-sensitivity *of $f: \mathcal{G} \to \mathbb{R}$ at G, denoted by $DS_f(G)$, is defined as*

$$DS_f(G) = \max_{\substack{H \preceq H' \preceq G \\ H, H' \text{ neighbors}}} |f(H') - f(H)|.$$

The down-sensitivity of f_{sf} and of f_{cc} differ by at most 1 on all graphs (since their sum, $|V|$, changes by at most 1 between node-neighboring graphs).

Within the framework of Lipschitz extensions, the down-sensitivity characterizes the largest monotone anchor set that is possible for a Δ-Lipschitz extension. It plays an important role in the accuracy guarantees in previous work on node-private algorithms. We show that the error of our algorithm for the size of the spanning forest (and the number of connected components) can be bounded in terms of the down-sensitivity (Theorem 1.5 below). The theorem provides a comparable result to existing theorems on subgraph counts and degree distributions [CZ13, RS16a].

Theorem 1.5 (Down-sensitivity guarantee of algorithm for the size of the spanning forest). *If $E(G) \neq \emptyset$, with probability $1 - o(1)$, the algorithm in Theorem 1.3 satisfies*

$$|\mathcal{A}(G) - f_{\text{sf}}(G)| \leq DS_{f_{\text{sf}}}(G) \cdot \widetilde{O}\Big(\frac{\ln \ln n}{\varepsilon}\Big).$$

An analogous statement holds with f_{cc} replacing f_{sf} when we modify the algorithm according to (1).

To prove Theorem 1.5, we establish a connection between the down-sensitivity of f_{sf} at graph G and the existence of a degree-bounded spanning forest of G. Lemma 1.6 together with Theorem 1.3 immediately imply Theorem 1.5.

Lemma 1.6. *Let Δ^* denote the smallest possible maximum degree of a spanning forest of a graph G. Then $\Delta^* \leq DS_{f_{\text{sf}}}(G) + 1$.*

Lemma 1.6 is obtained from two combinatorial results about spanning forests of graphs. The first, Lemma 1.7, characterizes the down-sensitivity of f_{sf} via the size of induced stars of G. Given an integer $k \geq 1$ and vertices $v_0, \ldots, v_k \in V(G)$, we say that these vertices form an *induced k-star centered at v_0* in G if $(v_0, v_i) \in E(G)$ for all $i \in [k]$ and $(v_i, v_j) \notin E(G)$ for all $i, j \in [k]$.

Lemma 1.7. *For a graph G, let $s(G)$ denote the largest integer such that G has an induced $s(G)$-star. Then*

$$DS_{f_{\text{sf}}}(G) = s(G).$$

Note that $s(G)$ can be much lower than the maximum degree of G due to the requirement that the $s(G)$-star be induced. The second result, Lemma 1.8, states that a graph with no large induced stars has a low-degree spanning forest.

Lemma 1.8. *A graph with no induced Δ-stars has a spanning Δ-forest.*

Lemma 1.8 is the key link in the chain connecting the down-sensitivity of f_{sf} to our algorithm's accuracy. The proof is constructive: we give a procedure that adds one vertex at a time to the spanning forest, modifying it at each step to maintain the degree bound.

Down Sensitivity, Anchor Sets, and Our Extensions. As discussed in Section 1, one way to measure the quality of a Lipschitz extension is in terms of its anchor set. Given a function f and parameter Δ, the largest possible monotone anchor set of any Δ-Lipschitz extension of f is the set $S_\Delta^* = \{G \mid DS_f(G) \leq \Delta\}$. By "largest" we mean that every monotone subset of an anchor set is a subset of S_Δ^*. We give an explicit Lipschitz extension with anchor set S_Δ^* and a proof of the optimality of S_Δ^* in Section A. However, for general f, the construction need not be efficient, even if f is computable in polynomial time.

The Lipschitz extensions f_Δ (for f_{sf}) used in our main algorithm are efficiently computable and have anchor sets that nearly match those of the extension based on down-sensitivity: their anchor sets S_Δ contain all graphs with down-sensitivity at most $\Delta - 1$, which means they contain the largest possible monotone anchor set for Lipschitz parameter $\Delta - 1$.

LEMMA 1.9 (NEARLY OPTIMAL ANCHOR SETS). *Let G be a graph and $\Delta \geq DS_{f_{\mathrm{sf}}}(G) + 1$. Then $f_\Delta(G) = f_{\mathrm{sf}}(G)$. Consequently, for all $\Delta \geq 1$,*

$$S_{\Delta-1}^* \subseteq S_\Delta .$$

We prove Lemmas 1.6–1.9 in Section 4.

1.1.3 Optimality of our Lipschitz extensions in terms of ℓ_∞ error. Next, we analyze the optimality of our Lipschitz extension in the vein of results of Cummings and Durfee [CD20]. They design an algorithm for constructing a Lipschitz extension for general f. Their construction is 2-competitive with the optimal extension with the same Lipschitz parameter. To measure the error of an extension f_Δ, we define $\mathrm{Err}_G(f_\Delta, f)$ as $\max_{H \leq G} |f_\Delta(H) - f(H)|$. Cummings and Durfee compare $\mathrm{Err}_G(f_\Delta, f)$ to $\mathrm{Err}_G(f^*, f)$ for all functions f^* of bounded sensitivity. We give such a comparison for our Lipschitz extension for the size of the spanning forest.

DEFINITION 1.10 (FUNCTIONS OF BOUNDED SENSITIVITY). *Let $\mathcal{G}$ be the set of all undirected, unweighted graphs. For a Lipschitz parameter $\Delta > 0$, define*

$$\mathcal{F}_\Delta = \{f \colon \mathcal{G} \to \mathbb{R} \mid f \text{ is } \Delta\text{-Lipschitz}\}.$$

THEOREM 1.11 (OPTIMALITY OF OUR LIPSCHITZ EXTENSION). *Let $\Delta \geq 1$ and f_Δ be our Lipschitz extension with parameter Δ for the size of the spanning forest. If $\mathrm{Err}_G(f_\Delta, f_{\mathrm{sf}}) > 0$ then*

$$\mathrm{Err}_G(f_\Delta, f_{\mathrm{sf}}) \leq \left(2 \cdot \min_{f^* \in \mathcal{F}_{\Delta-1}} \mathrm{Err}_G(f^*, f_{\mathrm{sf}})\right) - 1. \quad (3)$$

Cummings and Durfee show that for all functions f, their Lipschitz extension, denoted by f_Δ', satisfies $\mathrm{Err}_G(f_\Delta', f) \leq 2 \min_{f^* \in \mathcal{F}_\Delta} \mathrm{Err}_G(f^*, f)$, Instead, in our result, the error of the extension is compared against the more restricted class of functions of sensitivity $\Delta - 1$. However, our extension is computable in polynomial time, as opposed to the exponential construction of [CD20]. The core of the proof is a lemma (Lemma 5.2) which, intuitively, explains the error of our Lipschitz extension f_Δ by attributing the error on a graph G to one of its induced subgraphs. To prove it, we use a combinatorial result of Win [Win89] on the decomposition of graphs with

no spanning Δ-forests. We prove Lemma 5.2 and Theorem 1.11 in Section 5.

Theorem 1.11 provides a strong type of optimality. In particular, it implies Lemma 1.9 on the optimality of the anchor sets S_Δ.

1.1.4 Performance of our algorithm on specific graph families. We show that our algorithm provides a good approximation to f_{cc} (the number of connected components) for two common graph models: the Erdős-Rényi model $G(n, p)$ for subconstant p, and the geometric graph model.

The $G(n, p)$ model generates a random n-node graph by independently adding an edge between each pair of nodes with probability p. Erdős and Rényi [ER60] showed that the connectivity of G changes with p. Consider the regime where $np = c$ for some constant $c > 0$. With probability $1 - o(1)$, the graph will have $f_{\mathrm{cc}} = \Omega(n)$ and maximum degree $O(\log n)$. By Theorem 1.3 and (1), our node-private algorithm's estimate of f_{cc} will then have additive error $\pm\widetilde{O}((\log n)/\varepsilon)$ and relative error $\tilde{O}((\log^2 n)/\varepsilon n)$.

A random geometric graph G is defined by a set of n vertices V in the unit square, i.e., $V \subseteq [0, 1] \times [0, 1]$, and a distance parameter $r \in (0, 1)$. An edge exists between two vertices in V if their Euclidean distance is at most r [Pen03]. Geometric graphs have low-degree spanning forests, namely, of degree at most 6 [BBK+11]. We provide an alternative proof of this fact, via Lemma 1.8: a geometric graph does not have induced 6-stars because it is impossible to fit 6 points in the unit disk so that the distance between all points is strictly greater than 1. By Lemma 1.8, a geometric graph has a spanning 6-forest. By Theorem 1.3 and (1), our algorithm gives a $\pm\tilde{O}((\ln \ln n)/\varepsilon)$ approximation to f_{cc} in geometric graphs. Random geometric graphs and their variants have been used to model certain aspects of real-world networks, including social graphs [BLM+15, FFV06] and mobile networks [PSSS13, Zin22].

1.2 Prior Work on Private Graph Analysis

Of the two natural adaptations of differential privacy to network data, edge-privacy is easier to achieve and (as a result) has been studied more extensively. Edge-private algorithms are known for a wide range of graph statistics and modeling tasks; see, for example, [NRS07, HLMJ09, BBDS12, GRU12, KS12, Upa13, WWW13, WWZX13, KRSY14, PGM14, LM14, ZCP+15, MCB15, NIR16, RRT19, ZN19, ALJ20, BGM22, DLR+22]. The number of connected components has not been studied explicitly under edge-privacy, but it is easy to release with additive error $\Theta(1/\varepsilon)$, since it can change by at most 1 with the insertion or removal of an edge.

Node privacy is a better fit for social network data, where one individual's data consists of all the relationships (edges) that the individual contributes to the graph. Its much stronger privacy protection makes it harder to attain. Existing work addresses subgraph counts [BBDS13, KNRS13, CZ13, DZBJ18, LML20]; degree and triangle distributions [RS16b, DLL16, LML20]; parameter estimation in stochastic block models [BCS15, BCSZ18, SU21]; training of graph neural networks [DMS+21]; and generating synthetic graphs [ZNF20]. See [RS16a, MUP+22] for

surveys on node-private algorithms and [XCK+21] for implementations of some of the algorithms.

Existing works on node-privacy that prove rigorous accuracy statements generally take one of two approaches. Some works assume that the input is generated by a distribution from a specific family and analyze how well their algorithm approximates the parameters of that distribution; that approach does not fit our setting of arbitrary fixed inputs. Other works seek to formulate instance-dependent guarantees based on specific features of a graph, notably the maximum degree [KNRS13, BBDS13, RS16b, DLL16]. Chen and Zhou [CZ13] give a finer-grained analysis, showing that the accuracy of their algorithms for subgraph counts (and implicitly those of [KNRS13]) relates to the down-sensitivity of the input; that approach was subsequently refined and strengthened [RS16b, RS16a].

Our work goes further. We identify a feature of the graph—Δ^*—that bounds the error of our algorithm and show that Δ^* can be bounded above in terms of the down-sensitivity of f_{sf} (and by extension that of f_{cc}). Implementing this approach requires new ideas and combinatorial results, since f_{sf} has a significantly different structure from that of subgraph counts: it is not a sum of quantities computable from the local view of each node, and it is not supermodular. Consequently, the connection to down-sensitivity is more subtle.

2 PRELIMINARIES

The most basic private mechanism for releasing a statistic f returns the value of f with additive noise scaled according to the global sensitivity of f. The noise follows a Laplace distribution. The Laplace distribution with mean 0 and standard deviation $\sqrt{2}b$, denoted by $\text{Lap}(b)$, has probability density $h(z) = \frac{e^{-|z|/b}}{2b}$.

DEFINITION 2.1 (GLOBAL SENSITIVITY [DMNS16]). *Given a function $f : \mathcal{G} \to \mathbb{R}$, its global sensitivity, GS_f, is defined as*

$$GS_f = \max_{neighbors\ G,\ G'} |f(G) - f(G')|.$$

Unless specified otherwise, we use GS_f w.r.t. node-neighbors.

THEOREM 2.2 (LAPLACE MECHANISM [DMNS16]). *The algorithm $\mathcal{A}$ that, given a graph G, outputs $\mathcal{A}(G) = f(G) + \text{Lap}(GS_f/\varepsilon)$ is ε-node-private.*

LEMMA 2.3 (TAIL OF LAPLACE RANDOM VARIABLE). *If $X \sim \text{Lap}(b)$, then $\Pr[|X| \geq t \cdot b] = e^{-t}$.*

Differential privacy is preserved under post-processing. Additionally, the outputs of multiple private algorithms can be combined to obtain an algorithm that has privacy protection linear in the number of composed algorithms.

LEMMA 2.4 (COMPOSITION AND POST-PROCESSING [DMNS16, DKM+06]). *If an algorithm $\mathcal{A}$ runs ε-node-private algorithms $\mathcal{A}_1, ..., \mathcal{A}_t$ and applies a randomized algorithm g to the outputs, then $\mathcal{A}(G) = g(\mathcal{A}_1(G), \ldots, \mathcal{A}_t(G))$ is $(t\varepsilon)$-node-private.*

3 A LIPSCHITZ EXTENSION FOR THE SIZE OF THE SPANNING FOREST

In this section, we prove Theorem 1.3 by giving a polynomial-time ε-node-private algorithm for approximating the size of the spanning forest of a graph G.

We start by defining a family of Lipschitz extensions for f_{sf} that are computable in polynomial time. Recall from Section 1.1.1 that our construction is based on the Δ-bounded forest polytope of the input graph.

DEFINITION 3.1 (LIPSCHITZ EXTENSION FOR THE SIZE OF THE SPANNING FOREST). *Given a vector $x \in \mathbb{R}^E$ and an edge $e \in E$, let $x(e)$ denote the value of the vector x at edge e. For a subset $S \subseteq V$, let $E[S]$ be the set of edges in the subgraph of G induced by S. Let $\delta(v)$ denote the set of edges incident to v. For a set of edges $F \subseteq E$, denote by $x(F)$ the value $\sum_{e \in F} x(e)$. Given $\Delta > 0$, the Δ-bounded forest polytope of G, denoted $\mathcal{P}_\Delta(G)$, consists of vectors $x \in \mathbb{R}^E$ that satisfy the following constraints:*

$$x(e) \geq 0 \qquad\qquad \forall\, e \in E; \qquad (4)$$
$$x(E[S]) \leq |S| - 1 \qquad \forall\, S \subseteq V,\ |S| \geq 2; \qquad (5)$$
$$x(\delta(v)) \leq \Delta \qquad\qquad \forall\, v \in V. \qquad (6)$$

The Lipschitz extension at G with parameter Δ is defined as $f_\Delta(G) = \max_{x \in \mathcal{P}_\Delta(G)} x(E)$.

Each f_Δ can be used to privately approximate f_{sf} by outputting a private approximation of f_Δ via the Laplace mechanism. The Laplace noise is scaled according to Δ, which is an upper bound on the global sensitivity of f_Δ. We use the Generalized Exponential Mechanism to privately select $\hat\Delta \in [1, n]$, such that $f_{\hat\Delta}(G)$ has the lowest expected error when used to approximate $f_{\text{sf}}(G)$ via the Laplace mechanism. Finally, we output the private approximation of $f_{\hat\Delta}$ via the Laplace mechanism. This procedure is described in Algorithm 1.

Algorithm 1 Node-Private Size of Spanning Forest

Input: Graph G, privacy parameter $\varepsilon > 0$, failure probability $\beta \in (0, 1)$.
1: Run Algorithm 4 (based on GEM [RS16b]) with parameters $\varepsilon/2$ and β and access to Algorithm 2, to obtain $\hat\Delta$.
2: $f_{\hat\Delta}(G) \leftarrow \texttt{EvalLipschitzExtension}(G, \hat\Delta)$.
3: Return $f_{\hat\Delta}(G) + Z$ where $Z \sim \text{Lap}(\frac{2\hat\Delta}{\varepsilon})$.

Algorithm 2 `EvalLipschitzExtension`

Input: Graph G, Lipschitz parameter $\Delta \in [1, n]$.
1: Solve the linear program from Definition 3.1 to obtain $f_\Delta(G) \leftarrow \max_{x \in \mathcal{P}_\Delta(G)} x(E)$.
2: Return $f_\Delta(G)$.

Lemma 3.3 summarizes the properties of our family of Lipschitz extensions that allow us to use the Generalized Exponential Mechanism with this family.

DEFINITION 3.2 (MONOTONE IN Δ, LIPSCHITZ UNDERESTI-MATES). *Let $h\colon \mathcal{G} \to \mathbb{R}$. The functions $\{h_\Delta\}_{\Delta \in [1, \Delta_{\max}]}$ are a family of monotone in Δ, Lipschitz underestimates for h if:*

(1) *(UNDERESTIMATION) $h_\Delta(G) \leq h(G)$ for all $\Delta \in [1, \Delta_{\max}]$ and all G.*

(2) *(MONOTONICITY) $h_{\Delta_1}(G) \leq h_{\Delta_2}(G)$ for all $\Delta_1 < \Delta_2$ and all G.*

(3) *(LIPSCHITZNESS) h_Δ is Δ-Lipschitz for all $\Delta \in [1, \Delta_{\max}]$.*

LEMMA 3.3 (PROPERTIES OF THE LIPSCHITZ EXTENSION). *The Lipschitz extensions $\{f_\Delta\}_{\Delta \in [1,n]}$ in Definition 3.1 are a family of monotone in Δ, Lipschitz underestimates for f_{sf}. Moreover, for all $\Delta \in [1, n]$ and all G, the following hold.*

(1) *If G has a spanning Δ-forest, then $f_\Delta(G) = f_{\mathrm{sf}}(G)$.*

(2) *f_Δ is computable in polynomial time.*

REMARK 3.4. *The Lipschitz constant Δ for f_Δ is tight. To see this, consider the graph G with Δ isolated vertices and the node-neighboring graph G' obtained from G by adding one vertex with edges to all vertices of G. Then $f_\Delta(G) = 0$ and $f_\Delta(G') = \Delta$.*

PROOF OF LEMMA 3.3. We first show that $f_\Delta(G) \leq f_{\mathrm{sf}}(G)$ for all Δ and G. Let $S_1, \ldots, S_k \subseteq V(G)$ be the vertex sets of the connected components of G. For all $x \in \mathcal{P}_\Delta(G)$ and for all $i \in [k]$, we have $x(E[S_i]) \leq |S_i| - 1$. Therefore, $x(E) \leq \sum_{i \in [k]} |S_i| - f_{\mathrm{cc}}(G) = n - f_{\mathrm{cc}}(G) = f_{\mathrm{sf}}(G)$.

Monotonicity in Δ follows from the fact that every vector $x \in \mathcal{P}_\Delta(G)$ also satisfies $x \in \mathcal{P}_{\Delta'}(G)$ for $\Delta' \geq \Delta$.

Next, we show that f_Δ is Δ-Lipschitz. Let G be obtained from $G' = (V', E')$ by removing a node v and all its adjacent edges. Then $f_\Delta(G) \leq f_\Delta(G')$, since for each $x \in \mathcal{P}_\Delta(G)$, there is a vector $x' \in \mathcal{P}_\Delta(G')$ with $x(E) = x'(E')$. Namely, x' has the same entries as x in E and has value 0 in the entries $E' \setminus E$.

Let $x' \in \mathcal{P}_\Delta(G')$ be such that $x'(E') = f_\Delta(G')$. Consider the vector x obtained from x' by omitting all the entries pertaining to edges in $\delta(v)$. Then $x'(E') - x(E) \leq \Delta$, since $x'(\delta(v)) \leq \Delta$. Also, $x \in \mathcal{P}_\Delta(G)$, and thus $x(E) \leq f_\Delta(G)$. We get

$$|f_\Delta(G') - f_\Delta(G)| = f_\Delta(G') - f_\Delta(G)$$
$$= x'(E') - x(E) + x(E) - f_\Delta(G) \leq \Delta + 0 = \Delta.$$

This concludes the proof that f_Δ is Δ-Lipschitz.

To prove Item 1, let F be the edges of a spanning Δ-forest of G, i.e., $|\delta(v) \cap F| \leq \Delta$ for all vertices v of G. Let x_F be the vector whose values are 1 for all $e \in F$ and 0 otherwise. Then $x_F \in \mathcal{P}_\Delta(G)$ and $x_F(E) = f_{\mathrm{sf}}(G)$. Thus $f_\Delta(G) \geq x_F(E) = f_{\mathrm{sf}}(G)$. By the underestimation property, we obtain $f_\Delta(G) = f_{\mathrm{sf}}(G)$.

Finally, we prove Item 2. To compute $f_\Delta(G)$, we need to solve a linear program. Padberg and Wolsey [PW83] show the existence of a polynomial-time separation oracle for condition (5) of the linear program that involves an exponential number of constraints. Conditions (4) and (6) can be clearly checked in polynomial-time. Thus, there exists a polynomial time algorithm for solving the linear program in Definition 3.1 and computing $f_\Delta(G)$. □

3.1 From Extensions to the Main Algorithm

To formulate and analyze our main algorithm (and prove Theorem 1.3), we measure the "quality" of each extension f_Δ as

an approximation for $f_{\mathrm{sf}}(G)$ by its expected error when outputting $f_\Delta(G)$ privately via the Laplace mechanism, as in Algorithm 1. We would like to select $\hat{\Delta} \in [1, n]$, so that $f_{\hat{\Delta}}$ from our family of Lipschitz extensions minimizes the expected error. The Generalized Exponential Mechanism selects an approximation to the true minimizer $\hat{\Delta}$ in a private way. Consider any function $h\colon \mathcal{G} \to \mathbb{R}$ and a family $\{h_\Delta\}_{\Delta \in [1, \Delta_{\max}]}$ of monotone in Δ, Lipschitz underestimates for h. The quality of each function h_Δ is measured by its approximation error, defined as

$$err_h(\Delta, G) := |h_\Delta(G) - h(G)| + \Delta/\varepsilon. \qquad (7)$$

Note that $err_h(\Delta, G)$ is an upper bound for the expected error $\mathbb{E}[|h_\Delta(G) + \mathrm{Lap}(\Delta/\varepsilon) - h(G)|]$ by the triangle inequality and the fact that $\mathbb{E}[|\mathrm{Lap}(b)|] = b$. The Generalized Exponential Mechanism outputs $\hat{\Delta} \in [1, \Delta_{\max}]$ that approximately minimizes the quantity $err_h(\Delta, G)$.

THEOREM 3.5 (GEM [RS16B]). *Fix $\varepsilon > 0$ and $\beta \in (0, 1)$. Let $\{h_\Delta\}_{\Delta \in [1, \Delta_{\max}]}$ be a family of monotone in Δ, Lipschitz underestimates for $h\colon \mathcal{G} \to \mathbb{R}$. Then there exists an ε-node-private algorithm (Algorithm 4) obtained from the Generalized Exponential Mechanism, that outputs a value $\hat{\Delta}$ such that for all $\Delta \in [1, \Delta_{\max}]$ and all G, with probability at least $1 - \beta$, it holds*

$$err_h(\hat{\Delta}, G) \leq err_h(\Delta, G) \cdot O\Big(\ln \frac{\ln(\Delta_{\max})}{\beta} \Big).$$

Furthermore, the algorithm runs in polynomial time if all h_Δ are polynomial time computable.

We now we use Theorem 3.5 together with Lemma 3.3 to complete the proof of Theorem 1.3.

PROOF OF THEOREM 1.3. Let $\mathcal{A}$ be Algorithm 1 run with failure probability $\beta = \frac{1}{\ln \ln n}$. We first show that $\mathcal{A}$ is ε-node-private. Step 1 of algorithm $\mathcal{A}$ is $(\varepsilon/2)$-node-private by Theorem 3.5. Step 3 is also $(\varepsilon/2)$-node-private by Theorem 2.2. By composition (Lemma 2.4), algorithm $\mathcal{A}$ is ε-node-private.

We now bound the error of $\mathcal{A}$. With probability at least $1 - \beta/2$, we have

$$|\mathcal{A}(G) - f_{\mathrm{sf}}(G)| = |f_{\hat{\Delta}}(G) + Z - f_{\mathrm{sf}}(G)|$$
$$\leq |f_{\hat{\Delta}}(G) - f_{\mathrm{sf}}(G)| + |Z|$$
$$\leq |f_{\hat{\Delta}}(G) - f_{\mathrm{sf}}(G)| + \frac{2\hat{\Delta}}{\varepsilon} \ln\Big(\frac{2}{\beta}\Big).$$

The last inequality follows from Lemma 2.3. By the definition in (7), we have $|\mathcal{A}(G) - f_{\mathrm{sf}}(G)| \leq err_{f_{\mathrm{sf}}}(\hat{\Delta}, G) \cdot 2\ln(2/\beta)$ with probability at least $1 - \beta/2$.

Let Δ^* be the smallest value in $[1, n]$ such that G has a spanning Δ^*-forest. By Item 1 of Lemma 3.3, $|f_{\Delta^*}(G) - f_{\mathrm{sf}}(G)| = 0$. Therefore, $err_{f_{\mathrm{sf}}}(\Delta^*, G) = \frac{\Delta^*}{\varepsilon}$. By Lemma 3.3, the functions $\{f_\Delta\}_{\Delta \in [1,n]}$ are a family of monotone in Δ, Lipschitz underestimates for f_{sf}. Applying Theorem 3.5 with $\Delta_{\max} = n$, we have that with probability at least $1 - \beta$,

$$|\mathcal{A}(G) - f_{\mathrm{sf}}(G)| \leq \frac{\Delta^*}{\varepsilon} \cdot O\Big(\ln \Big(\frac{\ln n}{\beta}\Big) \cdot \ln \frac{1}{\beta}\Big).$$

Since $\beta = \frac{1}{\ln \ln n}$, we obtain the desired result. Since the Lipschitz extensions are computable in polynomial time (Item

Figure 1: Before and after the local repair at vertex v_1. Black solid edges are in the spanning forest. Dotted red edges are in the graph G, but not in the spanning forest.

2 of Lemma 3.3), algorithm $\mathcal{A}$ runs in polynomial time. This concludes the proof. □

4 DOWN-SENSITIVITY OF THE SIZE OF THE SPANNING FOREST

In this section, we prove Lemma 1.6 which establishes a connection between the down-sensitivity of f_{sf} and the existence of a bounded-degree spanning forest, and Lemma 1.9 which connects down-sensitivity to the anchor sets of our extension. Lemma 1.6 follows from two combinatorial results on spanning forests: Lemmas 1.7 and 1.8. We start by proving Lemma 1.8, which connects induced stars to the existence of bounded-degree spanning forests and is the key step in the proof of Lemma 1.6. The proofs of Lemmas 1.6 and 1.7 are deferred to Section 4.1. We prove Lemma 1.9 in Section 4.2

We first give an overview of the main ideas of the proof of Lemma 1.8. The proof is by induction on the number of vertices in G. Suppose G has no induced Δ-stars, i.e., $s(G) < \Delta$. Let v_0 be a vertex in G, which is not a cut vertex. (A vertex is a *cut vertex* if its removal from the graph increases the number of connected components.) Let G' be the subgraph of G induced by all vertices other than v_0. Then $s(G') \leq s(G) < \Delta$. By the inductive hypothesis, G' has a spanning forest F of degree at most Δ. Let v_1 be one of the neighbors of v_0 in G and add the edge (v_0, v_1) to F. Now F is a spanning forest of G, since v_0 is not a cut vertex. However, the degree of v_1 in F could now be $\Delta + 1$. We modify F to obtain a spanning forest of G of degree at most Δ by performing a sequence of "local repairs". If v_1 has degree $\Delta + 1$ in F, then v_1 has two neighbors a and b in F, such that a and b are adjacent in G (this is true since G has no induced Δ-stars). A local repair at v_1 replaces the edge (v_1, a) in F with the edge (a, b). See Fig. 1. We show that this operation does not introduce a cycle, and F is still a spanning forest of G. The degree of v_1 in F is now Δ, as desired, but the degree of a could have increased from Δ to $\Delta + 1$, thus calling for another local repair at a. At first glance, the local repair only pushes the problem around, rather than fixing it. However, with some care, we can show that the sequence of vertices $v_0, v_1, \ldots$ where we perform the local repair forms a path in a spanning forest of G. Thus, the sequence of local repairs will eventually conclude, yielding a spanning forest of G of degree at most Δ.

PROOF OF LEMMA 1.8. Let G be a graph on n vertices. The proof is by induction on n. If $n = 1$, then G has no induced Δ-stars for all $\Delta > 0$. Graph G also has a spanning Δ-forest for all $\Delta > 0$. Thus the lemma holds. By the same reasoning, the lemma holds if $E(G) = \emptyset$. Suppose $E[G] \neq \emptyset$ and that the lemma holds for every graph with $n - 1$ vertices. Then there exists some vertex $v_0 \in V(G)$ which is not isolated and not a cut vertex. E.g., consider a spanning forest of G and let v_0 be one of the leaves of the spanning forest. Let $G' = G \setminus \{v_0\}$. Then $s(G') \leq s(G) < \Delta$. By the inductive hypothesis, there exists a spanning forest F of G' of degree at most Δ.

Let vertex v_1 be a neighbor of v_0 in G (recall that vertex v_0 is not isolated). Let F_0 be the subgraph of G consisting of F and the edge (v_0, v_1). Since v_0 is not a cut vertex, then F_0 is a spanning forest of G. If all vertices in F_0 have degree at most Δ, then we have found the desired spanning forest. Otherwise, we perform a sequence of local repairs, as outlined in Algorithm 3. The output of Algorithm 3 is a sequence $\{(v_i, F_i)\}$ of vertices and subgraphs of G. The subgraph F_i is obtained after a local repair at vertex v_i (recall Fig. 1). In Claim 4.1, we show that all subgraphs F_i are spanning forests of G. Additionally, the vertices $v_0, v_1, \ldots$ are distinct and thus the sequence $\{(v_i, F_i)\}$ is finite (it must end once all vertices of G are output). By the description of Algorithm 3, the final F_i it outputs is a spanning forest of G of degree at most Δ, as desired. Applying Claim 4.1 thus concludes the proof. □

Algorithm 3 Repair a $(\Delta+1)$-forest to get a spanning Δ-forest

Input: Graph G, a vertex v_0, and a forest F_0, defined above
1: $i \leftarrow 1$
2: **while** max degree in F_{i-1} is greater than Δ **do**
3:     Let v_i be a vertex with degree at least $\Delta + 1$ in F_{i-1}.
4:     Let N be Δ neighbors of v_i in F_{i-1}, with $v_{i-1} \notin N$.
5:     Let $a_i, b_i \in N$ such that $(a_i, b_i) \in E(G)$.   ▷ Since $s(G) < \Delta$, such a_i, b_i exist.
6:     Let $F_i = (F_{i-1} \setminus \{(v_i, b_i)\}) \cup \{(a_i, b_i)\}$.
7:     Output (v_i, F_i).
8:     $i \leftarrow i + 1$

CLAIM 4.1. *For all $i \geq 0$ and pairs (v_i, F_i) output by Algorithm 3, the following hold:*

(a) *F_i is a spanning forest of G.*
(b) *All vertices in F_i have degree at most $\Delta + 1$, and at most one of them has degree $\Delta + 1$.*
(c) *(v_i, v_{i+1}) is an edge in both F_i and F_{i+1}, assuming F_i is not the final subgraph output by Algorithm 3.*
(d) *Vertices $v_0, \ldots, v_i$ are distinct, and they form a path in F_i.*

PROOF. We prove Item (a) by induction on i. The claim is true for $i = 0$, since adding (v_0, v_1) to F does not introduce a cycle, and the number of connected components of G and G' is the same since v_0 is not a cut vertex. Suppose the claim holds for F_{i-1}. We show that it also holds for F_i. First note that F_{i-1} and F_i have the same number of edges. Therefore, it suffices to show that F_i has no cycles. Suppose F_i has a cycle. Then this cycle must include the edge (a_i, b_i), whereas all other

edges of the cycle are in F_{i-1}. Suppose the vertex v_i is not included in the cycle. Then, by replacing the edge (a_i, b_i) with the edges (v_i, a_i) and (v_i, b_i), we would obtain a cycle in F_{i-1}, a contradiction. Now suppose v_i is included in the cycle. By replacing the edges (v_i, a_i) and (a_i, b_i) with the edge (v_i, b_i), we would again obtain a cycle in F_{i-1}, a contradiction. Thus, F_i is a spanning forest, which concludes the proof of Item (a).

We also prove Item (b) by induction. The claim is true for F_0, since all vertices in F_0 have the same degree as in F, except for v_1, whose degree increases by 1 between F and F_0. Since $\deg_F(v_1) \leq \Delta$, then $\deg_{F_0}(v_1) \leq \Delta + 1$. Additionally, v_0 has degree exactly 1, which is at most Δ.

Suppose the claim is true for $F_0, \ldots, F_{i-1}$. If all vertices in F_{i-1} have degree at most Δ, this concludes the proof. Otherwise, there exists a unique vertex v_i with $\deg_{F_{i-1}}(v_i) = \Delta + 1$. Let a_i, b_i be the neighbors of v_i defined in Step 5. Note that the only vertices whose degrees change between F_i and F_{i-1} are a_i and v_i. The degree of v_i decreases by 1, and thus $\deg_{F_i}(v_i) = \Delta$. The degree of a_i increases by 1. By the inductive hypothesis, $\deg_{F_{i-1}}(a_i) \leq \Delta$, and thus $\deg_{F_i}(a_i) \leq \Delta + 1$. All other vertices in F_i have the same degree as in F_{i-1}, thus their degree is at most Δ. This concludes the proof of Item (b).

We prove Item (c). Fix iteration i and suppose F_i is not the final subgraph output by Algorithm 3. By Item (b), vertex v_i has degree $\Delta+1$, whereas all other vertices in F_{i-1} have degree at most Δ. Let a_i, b_i be the neighbors of v_i defined in Step 5. By the proof of Item (b), all vertices in F_i, except a_i, have degree at most Δ. Since F_i is not the final forest, it must hold that $\deg_{F_i}(a_i) = \Delta + 1$. Therefore, $v_{i+1} = a_i$. Since (v_i, a_i) is an edge in F_i, we obtain that (v_i, v_{i+1}) is an edge in F_i. Note that in iteration $i + 1$, we pick neighbors a_{i+1} and b_{i+1} of v_{i+1} such that $a_{i+1} \neq v_i$ and $b_{i+1} \neq v_i$. Therefore, the edge (v_i, v_{i+1}) remains unchanged in F_{i+1}. This concludes the proof of Item (c).

We prove Item (d) by induction. It clearly holds for $i = 0$. Suppose it is true for iteration i. We prove it for iteration $i + 1$. By Item (c), the edge (v_i, v_{i+1}) is in F_i. Note that $v_{i+1} \neq v_{i-1}$, since $v_{i+1} = a_i$, and we specifically choose $a_i \neq v_{i-1}$. Suppose $v_{i+1} = v_j$ for some $j < i - 1$. By the inductive hypothesis, the path from v_j to v_i is in F_i. Together with the edge (v_i, v_j), it forms a cycle in F_i, a contradiction. Therefore, v_{i+1} is distinct from all vertices $v_0, \ldots, v_i$. We now show that the vertices $v_0, \ldots, v_{i+1}$ form a path in F_{i+1}. By a similar argument as for $v_{i+1} = a_i$, we can show that b_i is distinct from all vertices $v_0, \ldots, v_i$. Therefore, all edges on the path from v_0 to v_i are also in the forest F_{i+1}. Since (v_i, v_{i+1}) is an edge in F_{i+1}, we obtain that the vertices $v_0, \ldots, v_{i+1}$ form a path in F_{i+1}. □

4.1 Completing the Proof of Lemma 1.6

In this section, we complete the proof of Lemma 1.6. To that end, we first prove Lemma 1.7 which characterizes the down-sensitivity of f_{sf} via the size of induced stars of the graph. Lemma 1.6 follows easily from Lemmas 1.7 and 1.8.

PROOF OF LEMMA 1.7. We first show that $DS_{f_{\text{sf}}}(G) \geq s(G)$. Let H' be a maximum induced star of G, i.e., H' is an induced $s(G)$-star, and let v be its central vertex. Let $H = H' \setminus \{v\}$, i.e., H consists of $s(G)$ isolated vertices. It follows that $f_{\text{sf}}(H) = 0$

and $f_{\text{sf}}(H') = s(G)$. Since $H \leq H' \leq G$ and H, H' are node-neighboring, then $D_{f_{\text{sf}}}(G) \geq f_{\text{sf}}(H') - f_{\text{sf}}(H) = s(G)$.

We now show that $DS_{f_{\text{sf}}}(G) \leq s(G)$. Let $H \prec H'$ be two node-neighboring induced subgraphs of G that differ at node v and satisfy $DS_{f_{\text{sf}}}(G) = f_{\text{sf}}(H') - f_{\text{sf}}(H)$. If $\deg_{H'}(v) = 0$ then $f_{\text{sf}}(H') = f_{\text{sf}}(H)$, and the lemma holds. Now assume $\deg_{H'}(v) \geq 1$. Then $f_{\text{sf}}(H') \geq f_{\text{sf}}(H)$, and in particular, $f_{\text{sf}}(H') - f_{\text{sf}}(H) = |V(H')| - f_{\text{cc}}(H') - |V(H)| + f_{\text{cc}}(H) = f_{\text{cc}}(H) - f_{\text{cc}}(H') + 1$. Let S_v be the vertices of the connected component of H' containing v, and let $S_1, \ldots, S_k \subseteq S_v$ be the sets of vertices of the connected components of the subgraph induced by $S_v - \{v\}$. Note that $k \geq 1$, since $\deg_{H'}(v) \geq 1$. Then $f_{\text{cc}}(H) - f_{\text{cc}}(H') = k - 1$, and thus $DS_{f_{\text{sf}}}(G) = f_{\text{sf}}(H') - f_{\text{sf}}(H) = k$. It remains to prove that $k \leq s(G)$. For each $i \in [k]$, select some vertex $v_i \in S_i$ that is adjacent to v in the graph H'. Then the set of vertices $\{v_i \mid i \in [k]\}$ is an independent set in G because the vertices v_i are in different connected components of H. As a result, the vertices $\{v_i \mid i \in [k]\} \cup \{v\}$ induce a star centered at v in graph G. Therefore, $k \leq s(G)$, which concludes the proof. □

PROOF OF LEMMA 1.6. We show the following equivalent statement: If $DS_{f_{\text{sf}}}(G) \leq \Delta - 1$, then G has a spanning Δ-forest. Consider G such that $DS_{f_{\text{sf}}}(G) \leq \Delta - 1$. By Lemma 1.7, $s(G) = DS_{f_{\text{sf}}}(G) \leq \Delta - 1$. Thus, G has no induced Δ-stars, and by Lemma 1.8 we obtain that G has a spanning Δ-forest, as desired. □

4.2 Anchor Sets of our Extensions

In this section, we prove Lemma 1.9, which connects down-sensitivity to the anchor sets of our extensions.

PROOF OF LEMMA 1.9. Fix $\Delta > 0$ and a graph G such that $DS_{f_{\text{sf}}}(G) \leq \Delta - 1$. By Lemma 1.6, G has a spanning Δ-forest. Item 1 of Lemma 3.3 gives that $f_{\Delta}(G) = f_{\text{sf}}(G)$. Therefore, if $G \in S_{\Delta-1}^*$, then G has a Δ-forest, and as a result $G \in S_{\Delta}$. Thus $S_{\Delta-1}^* \subseteq S_{\Delta}$. □

5 OPTIMALITY OF OUR LIPSCHITZ EXTENSIONS

In this section, we prove Theorem 1.11, which says that our Lipschitz extension for f_{sf}, the size of the spanning forest, is close to optimal for this function. Our proof relies on the following combinatorial fact proved by Win [Win89] about graphs with no spanning Δ-forest. For a graph G and vertex set $X \subseteq V(G)$, we denote by $G \setminus X$ the graph obtained from G by removing all vertices in X and their adjacent edges.

LEMMA 5.1 (LEMMA 1, THEOREM 1 IN [WIN89]). *Let $\Delta \geq 2$. If a graph G has no spanning Δ-forest, then there exists an induced subgraph $S \prec G$ and a vertex set $X \subset V(S)$, such that:*

(1) S has a spanning Δ-tree.
(2) G has no edges between vertices in $G \setminus V(S)$ and vertices in $S \setminus X$.
(3) $f_{\text{cc}}(S \setminus X) \geq |X|(\Delta - 2) + 2$.

The key ingredient in our proof of Theorem 1.11 is the following lemma which, intuitively, explains the error of our Lipschitz extension f_{Δ} by attributing the error on a graph G

to one of its induced subgraphs. Recall from Lemma 3.3 that $f_\Delta(G) \le f_{\mathrm{sf}}(G)$ for all graphs G and that f_Δ can err only on graphs with no spanning Δ-forest. Lemma 5.2 shows that every such graph G has an induced subgraph H such that $f_\Delta(G)$ is nearly as large as possible for any Δ-Lipschitz lower bound on $f_{\mathrm{sf}}(G)$ because $f_\Delta(G)$ cannot deviate too far from $f_{\mathrm{sf}}(H)$.

LEMMA 5.2. *Let $\Delta \ge 1$. If a graph G has no spanning Δ-forest, then there exists a proper induced subgraph $H \prec G$ such that*

$$f_\Delta(G) \ge f_{\mathrm{sf}}(H) + (\Delta - 1)d(G, H) + 1. \tag{8}$$

PROOF. Let $G = (V, E)$ be a graph with no spanning Δ-forest.

To see that (8) holds for $\Delta = 1$, let H be a subgraph consisting of a single vertex of G. Then $f_{\mathrm{sf}}(H) = 0$, and the right-hand size of (8) evaluates to 1. Since G has no spanning Δ-forest, G has at least one edge. Thus, $f_\Delta(G) \ge 1$, and (8) holds for $\Delta = 1$.

Now fix some $\Delta \ge 2$. The proof is by induction on $|V| + |E|$, the total number of vertices and edges of G. Since G has no spanning Δ-forest, it must have a vertex of degree at least $\Delta + 1$. The smallest such graph (in terms of $|V| + |E|$) is a $(\Delta + 1)$-star.

For the base case, let G be a $(\Delta + 1)$-star. Denote its central vertex by v. Let $H = G \setminus \{v\}$. Then $f_\Delta(G) = \Delta$, $f_{\mathrm{sf}}(H) = 0$, and $(\Delta - 1)d(G, H) + 1 = \Delta$, since $d(G, H) = 1$. Thus, the claim holds for G (and is tight).

For the inductive step, consider a graph G that contains a $(\Delta + 1)$-star as a proper subgraph. See Fig. 2 for an illustration of the subgraphs of G considered in the rest of the proof.

Suppose (8) holds for all proper subgraphs of G. Let $S \le G$ and $X \subset V(S)$ be the subgraph and the vertex set given by Lemma 5.1. Then

$$f_\Delta(G) \ge f_\Delta(S) + f_\Delta(G \setminus V(S)) = f_{\mathrm{sf}}(S) + f_\Delta(G \setminus V(S)), \tag{9}$$

where the equality holds by Item 1 of Lemma 3.3, since S has a spanning Δ-tree. We apply Item 1 of Lemma 5.1 and (1):

$$
\begin{aligned}
& f_{\mathrm{sf}}(S) - f_{\mathrm{sf}}(S \setminus X) \\
&\quad = |V(S)| - 1 - |V(S \setminus X)| + f_{\mathrm{cc}}(S \setminus X) \\
&\quad \ge |X| - 1 + |X|(\Delta - 2) + 2 = |X|(\Delta - 1) + 1, \tag{10}
\end{aligned}
$$

where the inequality holds by Item 3 of Lemma 5.1. Combining (9) and the lower bound on $f_{\mathrm{sf}}(S)$ implied by (10) yields

$$f_\Delta(G) \ge f_{\mathrm{sf}}(S \setminus X) + |X|(\Delta - 1) + 1 + f_\Delta(G \setminus V(S)). \tag{11}$$

First, consider the case when $G \setminus V(S)$ has a spanning Δ-forest. By Item 1 of Lemma 3.3, $f_\Delta(G \setminus V(S)) = f_{\mathrm{sf}}(G \setminus V(S))$. Let $H = G \setminus X$. By Lemma 5.1, there are no edges between $G \setminus V(S)$ and $S \setminus X$. Thus, $f_{\mathrm{sf}}(S \setminus X) + f_{\mathrm{sf}}(G \setminus V(S)) = f_{\mathrm{sf}}(G \setminus X) = f_{\mathrm{sf}}(H)$. Also, $|X| = d(G, H)$. Substituting these two equalities into (11), we obtain that (8) holds in this case.

Finally, we consider the case when $G \setminus V(S)$ has no spanning Δ-forest. Let $G_1 = G \setminus V(S)$. By Item 3 of Lemma 5.1, $V(S) \ne \emptyset$, so G_1 is a proper subgraph of G. By the inductive hypothesis, there exists an induced subgraph $H_1 \prec G_1$ such that

$$f_\Delta(G \setminus V(S)) = f_\Delta(G_1) \ge f_{\mathrm{sf}}(H_1) + d(G_1, H_1)(\Delta - 1) + 1.$$

Figure 2: The subgraphs in Lemma 5.2. Graph $G[X]$ is the subgraph of G induced by X. There are no edges between $S \setminus X$ and $G_1 = G \setminus V(S)$. Graph H_1 is an induced subgraph of G_1. Finally, $H = G \setminus X$ and $H' = H_1 \cup (S \setminus X)$.

Plugging this bound on $f_\Delta(G \setminus V(S))$ into (11), we get

$$
\begin{aligned}
f_\Delta(G) \ge & f_{\mathrm{sf}}(S \setminus X) + f_{\mathrm{sf}}(H_1) \\
& + (|X| + d(G_1, H_1))(\Delta - 1) + 2.
\end{aligned}
$$

Let $H' = (S \setminus X) \cup H_1$. By Item 2 of Lemma 5.1, there are no edges between G_1 and $S \setminus X$. Hence, $f_{\mathrm{sf}}(S \setminus X) + f_{\mathrm{sf}}(H_1) = f_{\mathrm{sf}}(H')$. Also note that $|X| + d(G_1, H_1) = d(G, H')$. We obtain

$$f_\Delta(G) \ge f_{\mathrm{sf}}(H') + (\Delta - 1)d(G, H') + 2.$$

Thus, (8) holds in this case, too, concluding the proof. □

We now prove Theorem 1.11 using Lemma 5.2. Except for Lemma 5.2, which is at the core of our argument, the rest of the proof builds on an argument of Cummings and Durfee [CD20].

PROOF OF THEOREM 1.11. Let $G = (V, E)$ be a graph on which our Lipschitz extension errs, i.e., $\mathrm{Err}_G(f_\Delta, f_{\mathrm{sf}}) > 0$. The proof proceeds by induction on $|V| + |E|$, the total number of vertices and edges of G. By Item 1 of Lemma 3.3, since $\mathrm{Err}_G(f_\Delta, f_{\mathrm{sf}}) > 0$, graph G has no spanning Δ-forest. Therefore, we can consider the same base case as in the proof of Lemma 5.2.

For the base case, let G be a $(\Delta + 1)$-star. Let v be its central vertex. For this graph, the left-hand side of (3) equals 1. Let $f^* \in \mathcal{F}_{\Delta-1}$. By definition, $|f^*(G) - f^*(G \setminus \{v\})| \le \Delta - 1$. In contrast, $f_{\mathrm{sf}}(G) = \Delta + 1$ and $f_{\mathrm{sf}}(G \setminus \{v\}) = 0$. The quantity $\max\{|f^*(G) - f_{\mathrm{sf}}(G)|, |f^*(G \setminus \{v\}) - f_{\mathrm{sf}}(G \setminus \{v\})|\}$ is minimized when $f^*(G) = \Delta$ and $f^*(G \setminus \{v\}) = 1$ (under the Lipschitz requirement on f^*). Thus, the right-hand side of (3) equals 1, and the theorem holds when G is a $(\Delta + 1)$-star.

For the inductive step, consider a graph G that contains a $(\Delta + 1)$-star as a proper subgraph and $\mathrm{Err}_G(f_\Delta, f_{\mathrm{sf}}) > 0$. Suppose the theorem holds for all proper subgraphs of G. We show that it also holds for G. Let $\widetilde{H} \le G$ be the induced subgraph of G that maximizes the quantity $|f_\Delta(H) - f_{\mathrm{sf}}(H)|$ over all induced subgraphs $H \le G$. Then $\mathrm{Err}_{\widetilde{H}}(f_\Delta, f_{\mathrm{sf}}) > 0$.

First, consider the case when $\widetilde{H} \neq G$, i.e., $\widetilde{H} \prec G$. Then

$$\mathrm{Err}_G(f_\Delta, f_{\mathrm{sf}}) = \max_{H \leq G} |f_\Delta(H) - f_{\mathrm{sf}}(H)|$$

$$= \max_{H \leq \widetilde{H}} |f_\Delta(H) - f_{\mathrm{sf}}(H)| \tag{12}$$

$$\leq 2 \min_{f^* \in \mathcal{F}_{\Delta-1}(\mathcal{G})} \max_{H \leq \widetilde{H}} |f^*(H) - f_{\mathrm{sf}}(H)| - 1 \tag{13}$$

$$\leq 2 \min_{f^* \in \mathcal{F}_{\Delta-1}(\mathcal{G})} \max_{H \leq G} |f^*(H) - f_{\mathrm{sf}}(H)| - 1, \tag{14}$$

where (12) follows by our choice of $\widetilde{H}$, the inequality in (13) holds by the inductive hypothesis, and (14) holds since all induced subgraphs of $\widetilde{H}$ are also induced subgraphs of G. Thus, the theorem holds for G.

Finally, suppose $\widetilde{H} = G$, i.e., $\mathrm{Err}_G(f_\Delta, f_{\mathrm{sf}}) = |f_\Delta(G) - f_{\mathrm{sf}}(G)|$. Then $\mathrm{Err}_G(f_\Delta, f_{\mathrm{sf}}) = f_{\mathrm{sf}}(G) - f_\Delta(G)$ by the under-estimation property of our extension (Lemma 3.3). Moreover, since $f_\Delta(G) \neq f_{\mathrm{sf}}(G)$, Item 1 of Lemma 3.3 implies that G has no spanning Δ-forest. Thus, we can apply Lemma 5.2. It gives us that G has an induced subgraph $H \leq G$ such that $f_\Delta(G) \geq f_{\mathrm{sf}}(H) + (\Delta - 1)d(G, H) + 1$. For all $f^* \in \mathcal{F}_{\Delta-1}$, we have that $|f^*(G) - f^*(H)| \leq (\Delta - 1) \cdot d(G, H)$. Thus,

$$f_\Delta(G) \geq f_{\mathrm{sf}}(H) + |f^*(G) - f^*(H)| + 1$$

$$\geq f_{\mathrm{sf}}(H) + f^*(G) - f^*(H) + 1.$$

We use this inequality to get a bound on the error:

$$\mathrm{Err}_G(f_\Delta, f_{\mathrm{sf}}) = f_{\mathrm{sf}}(G) - f_\Delta(G)$$

$$\leq f_{\mathrm{sf}}(G) - f_{\mathrm{sf}}(H) - f^*(G) + f^*(H) - 1$$

$$\leq |f_{\mathrm{sf}}(G) - f^*(G)| + |f^*(H) - f_{\mathrm{sf}}(H)| - 1$$

$$\leq 2 \cdot \mathrm{Err}_G(f^*, f_{\mathrm{sf}}) - 1.$$

Since this holds for all $f^* \in \mathcal{F}_{\Delta-1}$, we finally obtain that $\mathrm{Err}_G(f_\Delta, f_{\mathrm{sf}}) \leq 2 \min_{f^* \in \mathcal{F}_{\Delta-1}} \mathrm{Err}_G(f^*, f_{\mathrm{sf}}) - 1$. $\square$

6 CONCLUSION

In this work, we presented the first node-differentially private algorithm for releasing the number of connected components of an undirected graph, a foundational query in graph databases. Our algorithm provides instance-based accuracy guarantees that are parameterized by the smallest possible maximum degree of a spanning forest of the input graph and relies on an efficiently computable family of Lipschitz extensions. The functions in this family can be evaluated in polynomial time by solving an LP with an exponentially large number of constraints using the ellipsoid algorithm with an efficient separation oracle. An interesting open direction is designing a faster algorithm with the same privacy and utility guarantees.

REFERENCES

[ALJ20] Faraz Ahmed, Alex X. Liu, and Rong Jin. Publishing social network graph eigenspectrum with privacy guarantees. *IEEE Transactions on Network Science and Engineering*, 7(2):892–906, 2020.

[ATK15] Leman Akoglu, Hanghang Tong, and Danai Koutra. Graph based anomaly detection and description: a survey. *Data Mining and Knowledge Discovery*, 29(3):626–688, 2015.

[BBDS12] Jeremiah Blocki, Avrim Blum, Anupam Datta, and Or Sheffet. The Johnson-Lindenstrauss transform itself preserves differential privacy. In *Proceedings, IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 410–419, 2012.

[BBDS13] Jeremiah Blocki, Avrim Blum, Anupam Datta, and Or Sheffet. Differentially private data analysis of social networks via restricted

sensitivity. In *Proceedings, Innovations in Theoretical Computer Science (ITCS)*, pages 87–96. ACM, 2013.

[BBK+11] József Balogh, Béla Bollobás, Michael Krivelevich, Tobias Müller, and Mark Walters. Hamilton cycles in random geometric graphs. *The Annals of Applied Probability*, 21(3):1053–1072, 2011.

[BCS15] Christian Borgs, Jennifer T. Chayes, and Adam D. Smith. Private graphon estimation for sparse graphs. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1369–1377, 2015.

[BCSZ18] Christian Borgs, Jennifer T. Chayes, Adam D. Smith, and Ilias Zadik. Revealing network structure, confidentially: Improved rates for node-private graphon estimation. In *Proceedings, IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 533–543, 2018.

[BGM22] Jeremiah Blocki, Elena Grigorescu, and Tamalika Mukherjee. Privately estimating graph parameters in sublinear time. In *Proceedings, International Colloquium on Automata, Languages and Programming (ICALP)*, volume 229, pages 26:1–26:19, 2022.

[BKM14] Petra Berenbrink, Bruce Krayenhoff, and Frederik Mallmann-Trenn. Estimating the number of connected components in sublinear time. *Information Processing Letters*, 114(11):639–642, 2014.

[BL98] Yoav Benyamini and Joram Lindenstrauss. *Geometric nonlinear functional analysis*. American Mathematical Society, 1998.

[BLM+15] Anthony Bonato, Marc Lozier, Dieter Mitsche, Xavier Pérez-Giménez, and Paweł Prałat. The domination number of on-line social networks and random geometric graphs. In *International Conference on Theory and Applications of Models of Computation*, pages 150–163, 2015.

[CD20] Rachel Cummings and David Durfee. Individual sensitivity preprocessing for data privacy. In *Proceedings, ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 528–547, 2020.

[CMMT22] Tianyi Chen, Brian Matejek, Michael Mitzenmacher, and Charalampos E Tsourakakis. Algorithmic tools for understanding the motif structure of networks. *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD*, pages 3–19, 2022.

[CRT05] Bernard Chazelle, Ronitt Rubinfeld, and Luca Trevisan. Approximating the minimum spanning tree weight in sublinear time. *SIAM Journal on Computing*, 34(6):1370–1379, 2005.

[CSS18] Beidi Chen, Anshumali Shrivastava, and Rebecca C Steorts. Unique entity estimation with application to the syrian conflict. *The Annals of Applied Statistics*, 12(2):1039–1067, 2018.

[CZ13] Shixi Chen and Shuigeng Zhou. Recursive mechanism: towards node differential privacy and unrestricted joins. In *Proceedings, ACM International Conference on Management of Data (SIGMOD)*, pages 653–664, 2013.

[DKM+06] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, volume 4004, pages 486–503, 2006.

[DLL16] Wei-Yen Day, Ninghui Li, and Min Lyu. Publishing graph degree distribution with node differential privacy. In *Proceedings, ACM International Conference on Management of Data (SIGMOD)*, pages 123–138, 2016.

[DLR+22] Laxman Dhulipala, Quanquan C. Liu, Sofya Raskhodnikova, Jessica Shi, Julian Shun, and Shangdi Yu. Differential privacy from locally adjustable graph algorithms: k-core decomposition, low out-degree ordering, and densest subgraphs. In *Proceedings, IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 754–765, 2022.

[DMNS16] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. *Journal of Privacy and Confidentiality*, 7(3):17–51, 2016.

[DMS+21] Ameya Daigavane, Gagan Madan, Aditya Sinha, Abhradeep Guha Thakurta, Gaurav Aggarwal, and Prateek Jain. Node-level differentially private graph neural networks, 2021.

[DZBJ18] Xiaofeng Ding, Xiaodong Zhang, Zhifeng Bao, and Hai Jin. Privacy-preserving triangle counting in large graphs. In *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM)*, pages 1283–1292. ACM, 2018.

[ER60] Paul Erdős and Alfréd Réyni. On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, page 17–61, 1960.

[FFV06] Abraham D Flaxman, Alan M Frieze, and Juan Vera. A geometric preferential attachment model of networks. *Internet Mathematics*, 3(2):187–205, 2006.

[Fra78] Ove Frank. Estimation of the number of connected components in a graph by using a sampled subgraph. *Scandinavian Journal of Statistics*, 5(4):177–188, 1978.

[Goo49] Leo A Goodman. On the estimation of the number of classes in a population. *The Annals of Mathematical Statistics*, 20(4):572–579,

1949.

[GRU12] Anupam Gupta, Aaron Roth, and Jonathan Ullman. Iterative constructions and private data release. In *Proceedings, Theory of Cryptography Conference (TCC)*, volume 7194, pages 339–356, 2012.

[HLMJ09] Michael Hay, Chao Li, Gerome Miklau, and David D. Jensen. Accurate estimation of the degree distribution of private networks. In *Proceedings, SIAM International Conference on Data Mining (ICDM)*, pages 169–178, 2009.

[KLM+14] Raimondas Kiveris, Silvio Lattanzi, Vahab Mirrokni, Vibhor Rastogi, and Sergei Vassilvitskii. Connected components in MapReduce and beyond. In *Proceedings of the ACM Symposium on Cloud Computing*, pages 18:1–18:13, 2014.

[KMAF10] U Kang, Mary McGlohon, Leman Akoglu, and Christos Faloutsos. Patterns on the connected components of terabyte-scale graphs. In *Proceedings, SIAM International Conference on Data Mining (ICDM)*, pages 875–880, 2010.

[KNRS13] Shiva Prasad Kasiviswanathan, Kobbi Nissim, Sofya Raskhodnikova, and Adam D. Smith. Analyzing graphs with node differential privacy. In *Proceedings, Theory of Cryptography Conference (TCC)*, pages 457–476, 2013.

[KRSY14] Vishesh Karwa, Sofya Raskhodnikova, Adam D. Smith, and Grigory Yaroslavtsev. Private analysis of graph structure. *ACM Transactions on Database Systems*, 39(3):22:1–22:33, 2014.

[KS12] Vishesh Karwa and Aleksandra B. Slavkovic. Differentially private graphical degree sequences and synthetic graphs. In *Privacy in Statistical Databases*, volume 7556, pages 273–285. Springer, 2012.

[KW20] Jason M Klusowski and Yihong Wu. Estimating the number of connected components in a graph via subgraph sampling. *Bernoulli*, 26(3):1635–1664, 2020.

[LM14] Wentian Lu and Gerome Miklau. Exponential random graph estimation under differential privacy. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 921–930. ACM, 2014.

[LML20] Ganghong Liu, Xuebin Ma, and Wuyungerile Li. Publishing node strength distribution with node differential privacy. *IEEE Access*, 8:217642–217650, 2020.

[MCB15] Yvonne Mülle, Chris Clifton, and Klemens Böhm. Privacy-integrated graph clustering through differential privacy. In *Proceedings of the Workshops of the EDBT/ICDT 2015 Joint Conference*, volume 1330, pages 247–254, 2015.

[McS34] Edward James McShane. Extension of range of functions. *Bulletin of the American Mathematical Society*, 40:837–842, 1934.

[MSOI+02] Ron Milo, Shai Shen-Orr, Shalev Itzkovitz, Nadav Kashtan, Dmitri Chklovskii, and Uri Alon. Network motifs: simple building blocks of complex networks. *Science*, 298(5594):824–827, 2002.

[MT07] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *Proceedings, IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 94–103, 2007.

[MUP+22] Tamara T. Mueller, Dmitrii Usynin, Johannes C. Paetzold, Daniel Ruecker, and Georgios Kaissis. SoK: Differential privacy on graph-structured data, 2022.

[NIR16] Hiep H. Nguyen, Abdessamad Imine, and Michaël Rusinowitch. Detecting communities under differential privacy. In *Proceedings of the ACM Workshop on Privacy in the Electronic Society (WPES@CCS)*, pages 83–93, 2016.

[NRS07] Kobbi Nissim, Sofya Raskhodnikova, and Adam D. Smith. Smooth sensitivity and sampling in private data analysis. In *Proceedings, ACM Symposium on Theory of Computing (STOC)*, pages 75–84, 2007.

[Pen03] Mathew D. Penrose. *Random Geometric Graphs*, volume 5. Oxford University Press, 2003.

[PGM14] Davide Proserpio, Sharon Goldberg, and Frank McSherry. Calibrating data to sensitivity in private data analysis. *Proceedings of the VLDB Endowment*, 7(8):637–648, 2014.

[PSSS13] Yuval Peres, Alistair Sinclair, Perla Sousi, and Alexandre Stauffer. Mobile geometric graphs: Detection, coverage and percolation. *Probability Theory and Related Fields*, 156(1):273–305, 2013.

[PW83] Manfred W. Padberg and Laurence A. Wolsey. Trees and cuts. In *Combinatorial Mathematics*, volume 75 of *North-Holland Mathematics Studies*, pages 511–517. North-Holland, 1983.

[RRT19] Leyla Roohi, Benjamin I. P. Rubinstein, and Vanessa Teague. Differentially-private two-party egocentric betweenness centrality. In *IEEE Conference on Computer Communications (INFOCOM)*, pages 2233–2241, 2019.

[RS16a] Sofya Raskhodnikova and Adam D. Smith. Differentially private analysis of graphs. In *Encyclopedia of Algorithms*, pages 543–547. 2016.

[RS16b] Sofya Raskhodnikova and Adam D. Smith. Lipschitz extensions for node-private graph statistics and the generalized exponential

mechanism. In *Proceedings, IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 495–504, 2016.

[SU21] Adam Sealfon and Jonathan R. Ullman. Efficiently estimating Erdos-Renyi graphs with node differential privacy. *J. Priv. Confidentiality*, 11(1), 2021.

[Upa13] Jalaj Upadhyay. Random projections, graph sparsification, and differential privacy. In *ASIACRYPT (1)*, volume 8269, pages 276–295, 2013.

[Win89] Sein Win. On a connection between the existence of k-trees and the toughness of a graph. *Graphs Comb.*, 5(1):201–205, 1989.

[WWW13] Yue Wang, Xintao Wu, and Leting Wu. Differential privacy preserving spectral graph analysis. In *Advances in Knowledge Discovery and Data Mining Pacific-Asia Conference (PAKDD)*, volume 7819, pages 329–340, 2013.

[WWZX13] Yue Wang, Xintao Wu, Jun Zhu, and Yang Xiang. On learning cluster coefficient of private networks. *Social Network Analysis and Mining*, 3(4):925–938, 2013.

[XCK+21] Siyuan Xia, Beizhen Chang, Karl Knopf, Yihan He, Yuchao Tao, and Xi He. DPGraph: A benchmark platform for differentially private graph analysis. In *Proceedings, ACM International Conference on Management of Data (SIGMOD)*, pages 2808–2812. ACM, 2021.

[XWC+22] Sheng Xiang, Dong Wen, Dawei Cheng, Ying Zhang, Lu Qin, Zhengping Qian, and Xuemin Lin. General graph generators: experiments, analyses, and improvements. *The VLDB Journal*, 31(5):897–925, 2022.

[ZCP+15] Jun Zhang, Graham Cormode, Cecilia M. Procopiuc, Divesh Srivastava, and Xiaokui Xiao. Private release of graph statistics using ladder functions. In *Proceedings, ACM International Conference on Management of Data (SIGMOD)*, pages 731–745, 2015.

[Zin22] Anatoliy Zinovyev. Space-stretch tradeoff in routing revisited. In *International Symposium on Distributed Computing (DISC)*, pages 37:1–37:16, 2022.

[ZN19] Sen Zhang and Weiwei Ni. Graph embedding matrix sharing with differential privacy. *IEEE Access*, 7:89390–89399, 2019.

[ZNF20] Sen Zhang, Weiwei Ni, and Nan Fu. Community preserved social graph publishing with node differential privacy. In *Proceedings, SIAM International Conference on Data Mining (ICDM)*, pages 1400–1405, 2020.

A DOWN-SENSITIVITY ERROR GUARANTEES FOR APPROXIMATING GENERAL FUNCTIONS

In this section, we prove two results related to down-sensitivity. The first, Theorem A.2, is a result of Raskhodnikova and Smith [RS16a] which states that for every monotone nondecreasing function f, there exists a node-private algorithm for approximating f whose error is bounded by the down-sensitivity of f. A function $f: \mathcal{G} \to \mathbb{R}$ is monotone nondecreasing if $f(H) \leq f(H')$ for all $H \leq H'$. Then, in Lemma A.3, we show that the largest monotone anchor set for a function f and Lipschitz constant Δ is characterised by the down-sensitivity of f.

We start by proving Lemma A.1, which is used in the proof of Theorem A.2.

Lemma A.1 (Lipschitz extension based on down-sensitivity). *Let $f: \mathcal{G} \to \mathbb{R}$ be a monotone nondecreasing function. Given $\Delta > 0$, define the family of functions*

$$\widehat{f_\Delta}(G) = \min_{\substack{H \leq G \\ DS_f(H) \leq \Delta}} \Big(f(H) + \Delta \cdot d(H, G) \Big).$$

The functions $\widehat{f_\Delta}(G)$ are a family of monotone in Δ, Lipschitz underestimates for f. Moreover, if $DS_f(G) \leq \Delta$, then $\widehat{f_\Delta}(G) = f(G)$.

Proof. To show that the functions underestimate f, note that $\widehat{f_\Delta}(G) \leq f(G) + \Delta \cdot d(G, G) = f(G)$. To prove monotonicity in Δ, fix Δ_1, Δ_2 and let $H \leq G$ such that $DS_f(H) \leq \Delta_2$.

Consider the empty graph $\emptyset$. Then $f(H) + \Delta_2 \cdot d(H, G) \geq f(\emptyset) + \Delta_1 d(\emptyset, G) \geq \widehat{f_{\Delta_1}}(G)$, since $\emptyset \leq H$ and f is monotone nondecreasing. Since the inequality holds for all H with $DS_f(H) \leq \Delta_2$, we obtain that $\widehat{f_{\Delta_2}}(G) \geq \widehat{f_{\Delta_1}}(G)$.

We now show Δ-Lipschitzness. Let G, G' be two neighboring graphs so that $V(G') = V(G) \cup \{v\}$, for $v \notin V$. We show that $\widehat{f_\Delta}(G) \leq \widehat{f_\Delta}(G') \leq \widehat{f_\Delta}(G) + \Delta$, which implies $|\widehat{f_\Delta}(G') - \widehat{f_\Delta}(G)| \leq \Delta$.

We start by showing that $\widehat{f_\Delta}(G) \leq \widehat{f_\Delta}(G')$. Let $H' \leq G'$ be the induced subgraph of G' such that $DS_f(H') \leq \Delta$ and $\widehat{f_\Delta}(G') = f(H') + \Delta \cdot d(H', G')$. Suppose that $v \notin V(H')$. Then, H is an induced subgraph of G and $d(H', G) = d(H', G') - 1$. It follows that $\widehat{f_\Delta}(G) \leq f(H') + \Delta \cdot d(H', G) = f(H') + \Delta \cdot (d(H', G') - 1) = \widehat{f_\Delta}(G') - \Delta$. Suppose now that $v \in V(H')$. Let $H = H' \setminus \{v\}$. Then H is an induced subgraph of G and $d(H', G') = d(H, G)$. Additionally, $DS_f(H) \leq DS_f(H') \leq \Delta$. Since f is monotone nondecreasing, it follows that $\widehat{f_\Delta}(G) \leq f(H) + \Delta \cdot d(H, G) \leq f(H') + \Delta \cdot d(H', G') = \widehat{f_\Delta}(G')$. In either case, we conclude that $\widehat{f_\Delta}(G) \leq \widehat{f_\Delta}(G')$.

We now show that $\widehat{f_\Delta}(G') \leq \widehat{f_\Delta}(G) + \Delta$. Let $H \leq G$ so that $DS_f(H) \leq \Delta$ and $\widehat{f_\Delta}(G) = f(H) + \Delta \cdot d(H, G)$. Since $H \leq G'$, then $\widehat{f_\Delta}(G') \leq f(H) + \Delta \cdot d(H, G') = f(H) + \Delta \cdot (d(H, G) + 1) = \widehat{f_\Delta}(G) + \Delta$. This concludes the proof of Lipschitzness.

Finally, we show that if $DS_f(G) \leq \Delta$ then $\widehat{f_\Delta}(G) = f(G)$. Suppose $DS_f(G) \leq \Delta$. Then f is Δ-Lipschitz for the set of induced subgraphs of G. Together with the fact that f is monotone nondecreasing, we obtain $f(G) - f(H) \leq \Delta \cdot d(H, G)$ for all induced subgraphs H of G. Therefore $\widehat{f_\Delta}(G) = f(G)$. □

We now use Lemma A.1 to prove Theorem A.2.

THEOREM A.2 (THEOREM 2 OF [RS16A]). *Let $f: \mathcal{G} \to \mathbb{R}$ be a monotone nondecreasing function. Let $\mathcal{G}^{(n)}$ denote the set of all n-node graphs and suppose that $\max_{G' \in \mathcal{G}^{(n)}} DS_f(G) \to \infty$ and $n \to \infty$. There exists an ε-node-private algorithm $\mathcal{A}_f$ that given an n-node graph G and privacy parameter $\varepsilon > 0$, with probability $1 - o(1)$ satisfies,*

$$|\mathcal{A}_f(G) - f(G)| \leq \frac{DS_f(G) + 1}{\varepsilon} \cdot \widetilde{O}\Big(\ln\ln \max_{G' \in \mathcal{G}^{(n)}} DS_f(G')\Big),$$

PROOF. Consider the family of functions defined in Lemma A.1, which are monotone in Δ, Lipschitz underestimates of f. Let $\beta = \frac{1}{\ln\ln \max_{G'} DS_f(G')}$ denote the failure probability. Note that $\beta = o(1)$. Let $\mathcal{A}_f$ be the algorithm that runs the algorithm of Theorem 3.5 with this family of Lipschitz extensions and parameters $\varepsilon/2, \beta/2$ to select a threshold $\hat{\Delta} \in [1, \max_{G'} DS_f(G')]$. It then outputs $f_{\hat{\Delta}}(G) + \mathrm{Lap}(\hat{\Delta}/(2\varepsilon))$. By composition (Lemma 2.4), $\mathcal{A}_f$ is ε-node-private.

We analyze the error of $\mathcal{A}_f$. By a similar argument as in the proof of Theorem 1.3, we have that with probability at least $1 - \beta/2$, it holds

$$|\mathcal{A}_f(G) - f(G)| \leq err_{f_{sf}}(\hat{\Delta}, G) \cdot 2\ln(2/\beta).$$

Let $\widetilde{\Delta} = DS_f(G) + 1$ (note that $\widetilde{\Delta} \geq 1$). By Lemma A.1, $err(\widetilde{\Delta}) = \frac{DS_f(G) + 1}{\varepsilon}$. Applying Theorem 3.5 with our choice of β concludes the proof. □

Finally, we prove Lemma A.3.

LEMMA A.3. *Given a function $f: \mathcal{G} \to \mathbb{R}$, a constant $\Delta > 0$, and a Δ-Lipschitz function $\hat{f}: \mathcal{G} \to \mathbb{R}$, let $S_{\hat{f}} = \{G \mid \hat{f}(G) = f(G)\}$. Let $S'_{\hat{f}}$ denote the largest monotone subset of $S_{\hat{f}}$ Then $S'_{\hat{f}} \subseteq \{G \mid DS_f(G) \leq \Delta\}$.*

PROOF. Suppose $G \in S'_{\hat{f}}$. Since $S'_{\hat{f}}$ is monotone, all induced subgraphs of G are also in $S'_{\hat{f}}$. Then,

$$\Delta \geq \max_{\substack{H \leq H' \leq G \\ H, H' \text{neighbors}}} |\hat{f}(H) - \hat{f}(H')| \tag{15}$$

$$= \max_{\substack{H \leq H' \leq G \\ H, H' \text{neighbors}}} |f(H) - f(H')| = DS_f(G). \tag{16}$$

The inequality in (15) holds since $\hat{f}$ is Δ-Lipschitz. The first equality in (16) holds since $H, H' \in S'_{\hat{f}} \subseteq S_{\hat{f}}$ for all induced subgraphs of G. The second equality in (16) is the definition of down-sensitivity. Therefore, $DS_f(G) \leq \Delta$ for all $G \in S'_{\hat{f}}$. □

B GEM

In this section, we state Algorithm 4 which applies the Generalized Exponential Mechanism of Raskhodnikova and Smith [RS16b] to the task of threshold selection for a family of Lipschitz extensions. This completes the description of Algorithm 1. Given a (possibly infinite) family of Lipschitz extensions, Algorithm 4 first obtains finitely many score functions $q_i, i \in I$ of bounded sensitivity, defined in terms of the Lipschitz extensions[1]. It then uses GEM to select the optimal score function.

GEM is a generalization of, and also includes as a subroutine, the Exponential Mechanism of McSherry and Talwar [MT07]. In the node-privacy setting, both algorithms take as input a graph G and score functions $q_i: \mathcal{G} \to \mathbb{R}, i \in I$ of bounded sensitivity. The goal of both algorithms is to output an index $\hat{i} \in I$ such that $q_{\hat{i}}$ approximately minimizes the value on G amongst all score functions, i.e., $q_{\hat{i}}(G) \approx \min_{i \in I} q_i(G)$. The Exponential Mechanism assumes a common upper bound Δ on the sensitivities of the score functions. GEM uses the fact that the score functions might have different sensitivities, and provides a better utility guarantee in the case when the optimal score function has a much lower sensitivity than the maximum of all sensitivities.

---

[1]The functions q_i defined in Step 4 of Algorithm 4 can actually have high sensitivity, since h can have high sensitivity. However, we only consider Lipschitz extensions that underestimate the true function h (see Lemma 3.3 and Theorem 3.5). In this case, minimizing q_i is equivalent to minimizing $h(G) - h_i(G) + \frac{i}{\varepsilon}$, and since the minimization is over i, we can treat $h(G)$ as a constant. So we can equivalently define the functions q_i as $-h_i(G) + \frac{i}{\varepsilon}$, which indeed have low sensitivity.

Algorithm 4 Threshold Selection for Lipschitz Extensions via GEM, adapted for node-privacy

---

Input: Function $h: \mathcal{G} \rightarrow \mathbb{R}$, access to subroutine `EvalLipschitzExtension` for evaluating the family of Lipschitz extensions $\{h_\Delta\}_{\Delta \in [1, \Delta_{\max}]}$, privacy parameter $\varepsilon > 0$, failure probability $\beta \in (0, 1)$

1: Let $k = \lfloor \log_2(\Delta_{\max}) \rfloor$, $I = \{2^0, \ldots, 2^k\}$, and $t = \frac{2\log(k/\beta)}{\varepsilon}$.

2: **for** $i \in I$ **do**:

3:     $h_i(G) \leftarrow \texttt{EvalLipschitzExtension}(G, i)$.

4:     $q_i(G) \leftarrow |h_i(G) - h(G)| + \frac{i}{\varepsilon}$.

5: **for** $i \in I$ **do**:

6:     $s_i(G) \leftarrow \max_{j \in I} \dfrac{(q_i(G) + ti) - (q_j(G) + tj)}{i + j}$.

7: To obtain $\hat{\imath}$, run the Exponential Mechanism [MT07] with graph G, score functions $\{s_i : i \in I\}$, and privacy parameter ε.

8: Return $\hat{\imath}$.

In Theorem B.1, we define the Exponential Mechanism and state its privacy guarantee. Algorithm 4 describes GEM tailored to the specific task of selecting an optimal Δ from a (possibly infinite) family of Lipschitz extensions $\{h_\Delta\}_{\Delta \in [1, \Delta_{\max}]}$. The utility, privacy, and running time guarantees of Algorithm 4 are stated in Theorem 3.5.

THEOREM B.1 (EXPONENTIAL MECHANISM [MT07]). *Given a graph G, finitely many score functions $q_i: \mathcal{G} \rightarrow \mathbb{R}, i \in I$ with global sensitivity at most Δ, and a privacy parameter $\varepsilon > 0$, the Exponential Mechanism $\mathcal{A}$ outputs an index $\hat{\imath} \in I$ drawn from the distribution $\Pr[\mathcal{A}(G) = i] \propto \exp\left(\dfrac{\varepsilon \cdot q_i(G)}{2\Delta}\right)$. Algorithm $\mathcal{A}$ is ε-node-private.*