

Tighter Bounds for Local Differentially Private Core Decomposition and Densest Subgraph

Monika Henzinger¹, A. R. Sricharan², and Leqi Zhu³

^{1,3}Institute of Science and Technology, Austria

²Faculty of Computer Science, UniVie Doctoral School of Computer Science DoCS, University of Vienna, Austria

Abstract

Computing the core decomposition of a graph is a fundamental problem that has recently been studied in the differentially private setting, motivated by practical applications in data mining. In particular, [DLR⁺22] gave the first mechanism for approximate core decomposition in the challenging and practically relevant setting of local differential privacy. One of the main open problems left by their work is whether the accuracy, i.e., the approximation ratio and additive error, of their mechanism can be improved. We show the first lower bounds on the additive error of approximate and exact core decomposition mechanisms in the centralized and local model of differential privacy, respectively. We also give mechanisms for exact and approximate core decomposition in the local model, with almost matching additive error bounds. Our mechanisms are based on a black-box application of continual counting. They also yield improved mechanisms for the approximate densest subgraph problem in the local model.

1 Introduction

Core decomposition and densest subgraph are fundamental problems that can be used to identify structure within graphs, with wide-ranging practical applications (see [DLR⁺22]). Recently, motivated by scenarios where the graph encodes confidential user information, there has been a growing interest in developing “privacy-preserving” algorithms for these problems. In particular, research has focused on algorithms that satisfy *differential privacy* [DMNS06], the de facto standard notion of privacy, which requires the distribution of the algorithm’s output to be nearly identical on graphs that differ by a single edge. Intuitively, this ensures that its output reveals little about the edges of the underlying graph.

Such algorithms have been designed for both the *centralized* and *local* models of differential privacy. In the centralized model, a *trusted* party runs the algorithm, with full access to all user data (here, the entire graph). The local model removes the trust assumption and is more distributed in nature. In this model, each user (here, a vertex in the graph) interactively discloses information about its data (here, its incident edges) in multiple *rounds* of communication with an *untrusted* server¹, which produces the output. All messages sent throughout the interaction, not just the output, must be differentially private.

In both models, a primary metric of interest is the *accuracy* of the output. This is fairly well understood in the centralized model, where there are fast and private (approximation) algorithms for core decomposition [DLL23] and densest subgraph [NV21, FHS22, DLR⁺22, DLL23, DKL23] with optimal and near-optimal additive error, respectively. However, there are still gaps in our understanding of this metric (and its trade-off with round complexity) in the local model. In a groundbreaking paper, [DLR⁺22] initiated the study of both problems in the local model by showing various trade-offs between approximation ratio, additive error, and round complexity. In particular, they gave (1) a private

¹The server may accidentally leak their messages, e.g., due to security breaches. However, it is not malicious. This is referred to as “honest but curious” in the literature.

$(2 + \eta)$ -approximate algorithm for core decomposition that has $O(\log^3 n)$ additive error and $O(\log n)$ round complexity, as well as (2) a private $(4 + \eta)$ -approximate algorithm for densest subgraph with the same trade-off. More recently, [DKLV23] gave private exact and $(2 + \eta)$ -approximate local algorithms for densest subgraph with improved additive error, roughly $O(\log^2 n)$, and with $O(n^2 \log n)$ and $O(\log n)$ round complexity, respectively. Concurrently with our work, [DLL23] gave exact and $(2 + \eta)$ -approximate local algorithms for core decomposition with $O(\log n)$ additive error and $O(n)$ and $O(\log n)$ round complexity, respectively. (See also Table 1.) All aforementioned error bounds hold *with high probability*, i.e., with probability at least $1 - n^{-c}$, for any fixed constant $c > 0$.

It is not clear, however, whether these are the best possible trade-offs. Thus, we seek to understand the following question:

What is the minimum additive error achievable, with high probability, for differentially private core decomposition (and densest subgraph) on n -vertex graphs, in the local model?

1.1 Our contributions

Lower bounds. We prove a lower bound on the additive error of approximate core decomposition in the centralized model, which carries over to the (less powerful) local model, regardless of round complexity. We also take the first step towards proving an accuracy and round complexity trade-off for exact core decomposition, by showing that 1-round algorithms must be inaccurate. To the best of our knowledge, these are the first lower bounds for core decomposition in either model. Specifically, we show the following results (see also Table 1). For simplicity, we consider mechanisms satisfying pure differential privacy, but similar bounds should hold for approximate differential privacy.

- *Lower bound for centralized and local model* (Theorem 6). For constant $\gamma \geq 1$, we show that any *centralized* algorithm for γ -approximate core decomposition has $\Omega(\gamma^{-1} \log n)$ additive error, with constant probability. Since centralized algorithms can simulate local ones, the bound also holds in the local model, regardless of the round complexity of the algorithm.
- *Lower bound for 1-round local model* (Theorem 7). We show that any local algorithm for exact core decomposition that uses a single round of communication has $\Omega(\sqrt{n})$ additive error, with constant probability, on a large family of graphs.

Upper bounds. To show that our first lower bound (which is already tight in the centralized model) is unlikely to be improved in the local model, we give local algorithms for exact and $(2 + \eta)$ -approximate core decomposition with improved, resp. nearly matching error, which exhibit new trade-offs with round complexity. In particular, we show that it is possible to achieve an additive error that does not have a linear dependence on the round complexity, answering a question of [DLR⁺22]. Concurrently with our work, [DLL23] showed that our first lower bound is, in fact, tight.

We obtain our results via a simple reduction from continual counting in the centralized model to core decomposition in the local model. In particular, in our algorithms, users employ continual counting mechanisms in a black-box manner. Thus, improving the accuracy of continual counting mechanisms (to match the known lower bound of [DNPR10]) would immediately improve the accuracy of our algorithms (to optimal).

Using a known (approximate) reduction from densest subgraph to core decomposition, our results for core decomposition immediately lead to local algorithms for 2-approximate and $(4 + \eta)$ -approximate densest subgraph with the same trade-offs.

Specifically, we show the following results (see also Table 1). We note that all our mechanisms satisfy pure differential privacy.

- *Exact core decomposition* (Theorem 3). We give a local algorithm for exact core decomposition that has $O(\log n \log \Delta)$ additive error, with high probability, and $O(n)$ round complexity, where Δ is the maximum degree of the graph.

| Problem | Apx. Factor | Additive Error | Rounds | Citation |
|--------------------|-------------|--------------------------------|-----------------------|-----------------------|
| Core Decomposition | 1 | $O(\log n)$ | $O(n)$ | [DLL23] |
| | 1 | $O(\log n \log \Delta)$ | $O(n)$ | [this work] |
| | 1 | $\Omega(\sqrt{n})$ | 1 | [this work] |
| | $2 + \eta$ | $O(\log n)$ | $O(\log n)$ | [DLL23] |
| | $2 + \eta$ | $O(\log^3 n)$ | $O(\log n)$ | [DLR ⁺ 22] |
| | $2 + \eta$ | $O(\log n \log \log n)$ | $O(\log^2 n)$ | [this work] |
| | γ | $\Omega(\gamma^{-1} \log n)$ | any | [this work] |
| Densest Subgraph | 1 | $O(\log(n\delta^{-1}) \log n)$ | $O(n^2 \log n)$ | [DKLV23] |
| | 1 | $\Omega(\sqrt{\log n})$ | any | [NV21] |
| | 2 | $O(\log n \log \Delta)$ | $O(n)$ | [this work] |
| | $2 + \eta$ | $O(\eta^{-1} \log^2 n)$ | $O(\eta^{-1} \log n)$ | [DKLV23] |
| | $4 + \eta$ | $O(\log^3 n)$ | $O(\log n)$ | [DLR ⁺ 22] |
| | $4 + \eta$ | $O(\log n \log \log n)$ | $O(\log^2 n)$ | [this work] |

Table 1: Summary of error bounds in the local model. Each upper bound is for a mechanism satisfying (ϵ, δ) -edge differential privacy and holds with high probability. If δ appears in the expression, then $0 < \delta < n^{-\Omega(1)}$; otherwise, $\delta = 0$. Some of our error bounds depend on the maximum degree, Δ , of the input graph. Lower bounds are for ϵ -differential privacy and hold with constant probability. To save space, the dependency on ϵ is omitted from each additive error upper bound.

- $(2 + \eta)$ -approximate core decomposition (Theorem 4). For constant $\eta > 0$, we give a local algorithm for $(2 + \eta)$ -approximate core decomposition that has nearly optimal additive error, $O(\log n \log \log n)$, with high probability, and $O(\log^2 n)$ round complexity.
- Approximate densest subgraph (Theorem 5). For constant $\eta > 0$, we give local algorithms for 2-approximate and $(4 + \eta)$ -approximate densest subgraph² with $O(\log n \log \Delta)$ and $O(\log n \log \log n)$ error, with high probability, and round complexity $O(n)$ and $O(\log^2 n)$, respectively.

1.2 Technical overview

Centralized lower bound. Our lower bound for γ -approximate core decomposition in the centralized model uses a standard “packing argument” [HT10]. The idea is to construct a collection of disjoint “bad” output classes for a certain input and show that the algorithm has a “large enough” probability of producing an output in each of the bad classes, due to differential privacy. The only subtlety is that we need to carefully construct the output classes to take advantage of the approximation guarantee.

Local lower bound. Our lower bound for 1-round, exact core decomposition in the local model is obtained by reduction to a problem in the *centralized* model, for which there is a strong lower bound. Specifically, it is known that, to privately answer $\Theta(n)$ random inner product queries on a secret $\{0, 1\}^n$ vector, most responses need to have $\Omega(\sqrt{n})$ additive error [De12]. Recently, [ELRS23] gave an elegant lower bound on the additive error of 1-round triangle counting algorithms in the local model using similar techniques.

In our case, we define a class of “query graphs”, one per inner product query, in which the coreness of a *fixed* vertex x (i.e., its assigned “score” in the core decomposition) is (roughly) the answer to the query. In addition to x , there are some *secret* vertices (and other vertices). For each secret vertex v , the existence of the edge $\{v, x\}$ is private information (which depends on the secret vector). Crucially, all neighborhoods of x and the secret vertices in *any possible query graph* appear in two specific query graphs, namely, ones corresponding to the all-ones and all-zeros query vectors, respectively. The neighborhoods of the remaining vertices do not depend on the secret vector.

Our approach to solving the inner product problem is now as follows. The centralized algorithm first simulates the 1-round local core decomposition algorithm on the two fixed graphs to determine

²Here, the algorithm returns a set of vertices, which induce a subgraph whose density is approximately maximal.

the messages that x and the secret vertices would send in *any* query graph and saves these messages. Subsequently, when answering a query, the centralized algorithm *reuses* the saved messages for x and the secret vertices (without further privacy loss) when it simulates the core decomposition algorithm on the corresponding query graph. This allows it to answer many queries correctly.

Originally, we had a more complex, direct argument. We briefly mention it here to give an idea of what is going on “under the hood” of the reduction. Specifically, we showed that, on a class of random graphs (similar to the query graphs instantiated in our reduction), most transcripts of a 1-round core decomposition algorithm have the property that, conditioned on seeing the transcript, the coreness of some (fixed) vertex still has high variance, $\Omega(n)$. This implies that the standard deviation of the additive error is $\Omega(\sqrt{n})$. We prefer the reduction argument, as it is simpler and gives a stronger result.

Local upper bounds. On the algorithmic side, our starting point is the $(2 + \eta)$ -approximate core decomposition algorithm of [DLR⁺22], which essentially implements an approximate version of the classic, *exact* peeling algorithm for core decomposition. To remove the approximation factor, it is natural to consider implementing the original algorithm, which is what we do to obtain our local algorithm for exact core decomposition. At a high level, the main technical challenge here is for each user (vertex) to privately disclose its degree as the server continually deletes subsets of users from the graph (based on the disclosed degrees). In [DLR⁺22], each user simply adds fresh Laplace noise to its actual degree and discloses the resulting sum. In this case, the standard deviation of the noise (which dominates the additive error) needs to be *linear in the number of disclosures*, i.e., deletions, to ensure privacy. This is acceptable for the approximate peeling algorithm, where $O(\log^2 n)$ deletions occur, but not for the exact algorithm, where $\Omega(n)$ deletions might occur.

A *continual counting* mechanism [DNPR10] privately discloses all (noisy) prefix sums of an input sequence of integers. There are *adaptive* variants, where the elements of the input sequence arrive online, one at a time, and the mechanism discloses the sum after each arrival. Our new approach is for each user to use an adaptive continual counting mechanism to track the *sum of its degree changes* due to deletions. It can then disclose its degree as its (noisy) initial degree plus the last sum.

There are several points that make implementing this non-trivial. (a) First, note that each user runs its own continual counting mechanism whose input (apart from its neighborhood list) is given by the server. Thus, if the inputs to the mechanisms (on two graphs that differ by a single edge) can differ for each user, even slightly, then the total privacy loss can be $\Omega(n)$, which would require each mechanism to add prohibitively large noise to maintain overall privacy. Through a careful analysis, we show that the total privacy loss is constant. (b) Next, since the outputs of the counting mechanisms have noise, the users do not report their actual degrees, but rather *noisy* ones, so the error could amplify over time. We prove that the exact peeling algorithm is *robust*, in the sense that the error of its output is bounded by the maximum error of the noisy degrees (which is fairly small, with high probability). (c) Finally, in the standard formalism of the local model via local randomizers, users do not have persistent memory. It is possible to simulate *any* user/server algorithm (with persistent memory) in the local randomizer model (with computationally unbounded users).

Our local algorithm for $(2 + \eta)$ -approximate core decomposition is obtained by applying the same thought process to the $(2 + \eta)$ -approximate algorithm of [DLR⁺22]. The difference here is that, since the number of deletions is $O(\log^2 n)$, the counting mechanisms have to disclose far fewer prefix sums and, hence, achieve better accuracy.

Finally, [DLL23] showed that their exact core decomposition can be post-processed to obtain a 2-approximate densest subgraph with roughly the same additive error (in the centralized model). Similarly, [DLR⁺22] implicitly proved this for their $(2 + \eta)$ -approximate core decomposition, obtaining a $(4 + \eta)$ -approximate densest subgraph (in the local model). We generalize their arguments to any approximation ratio, filling in details of the proof of [DLL23] (see the paragraph before Lemma 9). Specifically, we give a sufficient condition that allows a γ -approximate core decomposition to yield a 2γ -approximate densest subgraph with the same (asymptotic) additive error (in either model). Then we show that the condition is satisfied by our 1-approximate and $(2 + \eta)$ -approximate algorithms, hence obtaining 2-approximate

$(4 + \eta)$ -approximate densest subgraph algorithms, respectively, with the same trade-offs.

Related work. Concurrently with our work, there is a local algorithm for exact core decomposition with additive error $O(\log n)$ [DLL23]. Similar to our local algorithm, this algorithm privately implements the classic, exact peeling algorithm. However, they use different techniques (a multi-dimensional variant of the sparse vector technique). Continual counting has been used in the centralized model, along with other techniques, to obtain a $(2 + \eta)$ -approximate private densest subgraph algorithm with additive error $O(\log^{2.5} n)$, with high probability, which runs in near-linear time [FHS22]. There, the privacy analysis has a similar flavor as ours, but is more complex, while the accuracy analysis is quite different. Finally, there are insertions-only and deletions-only graph algorithms under continual observation that have used the idea of tracking the difference sequence of a desired quantity (e.g. degree of a vertex) via continual counting in the centralized model [FHO21].

2 Preliminaries

Core decomposition and densest subgraph In this paper, a *graph* is always simple (i.e., no self-loops or parallel edges), undirected, and unweighted. Let $G = (V, E)$ be a graph. We use the notation $V(G)$ and $E(G)$ to denote the set of vertices and edges of G , respectively. A graph H is a *subgraph* of G if $V(H) \subseteq V(G)$. Given a subset U of vertices of G , the subgraph of G *induced* by U , denoted $G[U]$, is the *largest* subgraph H of G such that $V(H) = U$, i.e., $|E(H)|$ is maximized. A vertex u of G is *adjacent* to (or, a *neighbor* of) a vertex $v \neq u$ of G if and only if $\{u, v\} \in E(G)$. The *degree* of a vertex v in G , denoted $\deg_G(v)$, is the number of vertices adjacent to v in G . Two graphs G and G' (on the same set of vertices) *differ* on a single edge (say e) if the symmetric difference of their edge sets, $(E(G) \setminus E(G')) \cup (E(G') \setminus E(G))$, consists of a single element (namely, e).

The *coreness* of a vertex v in a graph G , denoted $k_G(v)$, is the largest integer k such that there exists a subgraph H of G that contains v such that, for every vertex $u \in V(H)$, the (induced) degree of u in H is at least k , i.e., $\deg_H(u) \geq k$. We omit the subscript when it is clear from the context. For integer $k \geq 0$, the k -*core* of the graph G is the subgraph of G induced by the vertices of G with coreness at least k in G , i.e., $G[U]$, where $U = \{v \in V \mid k_G(v) \geq k\}$ ³. In the *core decomposition problem*, the input is a graph $G = (V, E)$ and the goal is to return a vector $(k_G(v) : v \in V)$ containing the coreness of each vertex of G (assuming a fixed ordering of the vertices of G).

The *density* of a graph G , denoted $\rho(G)$, is the quantity $|E(G)|/|V(G)|$. A *densest subgraph* of G is a subgraph H of G such that $\rho(H)$ is maximized, among all subgraphs of G . Notice that, for any subset U of vertices of G , if H is a subgraph of G such that $V(H) = U$, then $\rho(H) \leq \rho(G[U])$. In the *densest subgraph problem*, the input is a graph $G = (V, E)$ and the goal is to return a subset of vertices $U \subseteq V(G)$ that induce a densest subgraph of G .

Differential privacy For any $\epsilon, \delta > 0$, a randomized algorithm \mathcal{A} mapping (inputs in) \mathcal{X} to (outputs in) \mathcal{Y} is (ϵ, δ) -*differentially private* [DMNS06] if for all *neighboring* inputs $x, x' \in \mathcal{X}$ and for all subsets $Y \subseteq \mathcal{Y}$ of outputs, $\Pr(\mathcal{A}(x) \in Y) \leq e^\epsilon \Pr(\mathcal{A}(x') \in Y) + \delta$. Here, the definition of neighboring depends on the context and the problem. If $\delta = 0$, we simply write ϵ -*differentially private*.

Fact 1. Let \mathcal{A} be an ϵ -differentially private algorithm mapping \mathcal{X} to \mathcal{Y} .

1. (*Composition*) Suppose that \mathcal{B} is an ϵ' -differentially private algorithm mapping $\mathcal{X} \times \mathcal{Y}$ to \mathcal{Z} . Then the *composition* of \mathcal{A} and \mathcal{B} , i.e., the randomized algorithm mapping \mathcal{X} to \mathcal{Z} via $x \mapsto \mathcal{B}(x, \mathcal{A}(x))$, is $(\epsilon + \epsilon')$ -differentially private.
2. (*Post-processing*) Suppose that \mathcal{B} is a randomized algorithm mapping \mathcal{Y} to \mathcal{Z} . Then the *post-processing* of \mathcal{A} by \mathcal{B} , i.e., the randomized algorithm mapping \mathcal{X} to \mathcal{Z} via $x \mapsto \mathcal{B}(\mathcal{A}(x))$, is ϵ -differentially private.

³An alternative, equivalent definition is that the k -core of G is the (unique) largest subgraph of G in which every vertex has degree at least k .

The ℓ_1 -sensitivity of a function $f : \mathcal{X} \rightarrow \mathbb{R}^d$, denoted $\Delta_1(f)$, is the supremum of the quantity $\|f(x) - f(x')\|_1 = \sum_{i=1}^d |f(x)_i - f(x')_i|$, over all neighboring $x, x' \in \mathcal{X}$. The Laplace distribution (centered at 0) with scale $b > 0$ has probability density function $f(x) = \frac{1}{2b} \exp(-|x|/b)$. We write $X \sim \text{Lap}(b)$ or just $\text{Lap}(b)$ to denote a random variable X that is distributed according to the Laplace distribution with scale b .

Fact 2. If random variable $X \sim \text{Lap}(b)$, then, for any $\beta > 0$, $\Pr[|X| > b \log(1/\beta)] \leq \beta$.

Fact 3 (Theorem 3.6 in [DR14]). Let $f : \mathcal{X} \rightarrow \mathbb{R}^d$ be any non-constant function and let $\epsilon > 0$. Suppose that, for each $i \in [d]$, random variable $X_i \sim \text{Lap}(\Delta_1(f)/\epsilon)$. Then the Laplace mechanism $\mathcal{A}(x) = f(x) + (X_1, \dots, X_d)$ is ϵ -differentially private.

Local differential privacy In the *local* model of differential privacy, there are n users, each with *private data*, who communicate in synchronous *rounds* with a *server* (or *curator*). In each round, each user can send a message to the server, based on its private data, its local memory, the messages (from the server) that it has received so far, and possibly some *fresh* local randomness. Upon receiving all user messages, the server can either broadcast a message to all users, or end the interaction and produce an output. If it broadcasts a message, then the users will receive it at the end of the round. Subsequently, they may write information to their local memories, which persists into future rounds, before continuing to the next round.

A *local mechanism* \mathcal{M} specifies (randomized) algorithms for the server and for each user, which dictate the messages they send in each round and what they write to their local memories. The *input* of \mathcal{M} is the private data of each user. An *execution* of \mathcal{M} on an input generates a *transcript*, which consists of the messages sent by the server and all users in every round, if they run the algorithms specified by \mathcal{M} , with each user having the corresponding private input data. The *round complexity* of a mechanism is the maximum number of rounds in any transcript. A local mechanism is ϵ -differentially private if it is ϵ -differentially private, when viewed as a randomized algorithm mapping private user data to transcripts (provided the notion of neighboring user data is defined).

When the local model is applied to graphs, there is an underlying *input graph*, G , whose vertices are the users (and public knowledge). The private data of each user (vertex) is its set of neighbors in G . Hence, only users u and v know whether $\{u, v\}$ is an edge in G . Two input graphs are *neighboring* if they differ in at most one edge. In other words, a local mechanism in this setting is ϵ -(*edge*) *differentially private* if its transcript is ϵ -differentially private on neighboring input graphs.

The prevailing formalization of local differential privacy used for graph data [DLR⁺22, ELRS23, DKLV23] was adapted from the general (non-graph) setting [JMNR19]. In this formalism, *local randomizers* are used to describe the computations carried out by the users to generate the (differentially private) messages they send in each round. Specifically, an ϵ -*local randomizer* is a randomized algorithm R that takes as input a subset of vertices of the input graph and returns an output such that, for all subsets of vertices N, N' (corresponding to possible neighborhoods of a user) that differ in at most one element and any subset Y of outputs, $\Pr[R(N) \in Y] \leq e^\epsilon \Pr[R(N') \in Y]$. Paraphrasing [DLR⁺22], a *local mechanism* is specified by a potentially infinite set of local randomizers, \mathcal{R} , together with a function, \mathcal{A} , describing how the interaction proceeds. A *0-round transcript* is the empty sequence. Given a t -round transcript π , $\mathcal{A}(\pi)$ either returns \perp , indicating the interaction ends, or a pair $(S_U^{t+1}, S_R^{t+1}, S_\epsilon^{t+1})$ encoding the set S_U^{t+1} of users who participate in round $t + 1$, the set $S_R^{t+1} \subseteq \mathcal{R}$ containing the local randomizer assigned to each participating user, and the corresponding privacy parameters S_ϵ^{t+1} . Suppose that S_O^{t+1} contains the outputs generated by the local randomizers after each participating user runs its assigned local randomizer on its private data (i.e., neighborhood in the input graph) in round $t + 1$. Then the concatenation $\pi \odot (S_U^{t+1}, S_R^{t+1}, S_\epsilon^{t+1}, S_O^{t+1})$ is a $(t + 1)$ -round transcript. In Appendix A, we show that local mechanisms in our user/server model can be simulated in the local randomizer formalism. Given this, we prefer to describe our local mechanisms in terms of server and user algorithms, instead of local randomizers.

Continual counting In the *continual counting* problem, the input is a sequence of integers (a *stream*) up to some fixed, finite length, T . The goal is to output all prefix sums of the stream. A *continual counting* mechanism allows the elements of the input stream to be *inserted*, one at a time. After each insertion, the mechanism outputs the *count* (prefix sum) so far. Two streams (of the same length) are *neighboring* if they differ in at most one element, by at most one. The definition of differential privacy for neighboring input streams captures non-adaptive counting mechanisms. We refer to [JRSS23] for a formal definition of *adaptive* counting mechanisms, where the elements of the stream can depend on the outputs of the mechanism.

The *binary-tree* mechanism is a differentially private adaptive continual counting mechanism that was introduced by [DNPR10] and [CSS10]. Using the sparse-vector technique, [DNRR15] gave an improved version of the binary tree mechanism, which we call the *sparse-vector* counting mechanism, where the error at a time step is a function of the true count at that time step.

Theorem 1. [DNPR10, CSS10] *Let $\epsilon > 0$ be a constant. There is an ϵ -differentially private adaptive counting mechanism such that, with probability at least $1 - \beta$, at any single time step $t \in [T]$, the error is bounded by*

$$\text{err}_{\text{BT}}(t, \beta, \epsilon) = O\left(\epsilon^{-1} \cdot \log t \sqrt{\log 1/\beta} \cdot \max\left\{\sqrt{\log t}, \sqrt{\log 1/\beta}\right\}\right)$$

Theorem 2. [DNRR15] *Let $\epsilon > 0$ be a constant. There is an ϵ -differentially private adaptive counting mechanism such that, with probability at least $1 - \beta$, at any single time step $t \in [T]$, the error is bounded by*

$$O\left(\text{err}_{\text{BT}}(\min\{n_t, t\}, \beta, \epsilon) + \epsilon^{-1} \cdot \log(t/\beta)\right)$$

where n_t is the true continual count at time t and err_{BT} is the error of the binary tree mechanism.

3 Core decomposition via continual counting

Let $\epsilon > 0$ be any positive constant. In this section, we describe an ϵ -edge differentially private mechanism for core decomposition on n -vertex graphs, in the local model. The mechanism runs for at most n rounds in any execution and has $O(\epsilon^{-1} \log^2 n)$ additive error with high probability. It can also be made memoryless. In the following, we refer to the coreness of a vertex that is returned by the mechanism as the *estimated* coreness of that vertex.

Theorem 3. *For any $\epsilon > 0$, there is a memoryless local ϵ -edge differentially private mechanism \mathcal{M} that returns an estimate of the coreness of each vertex in an n -vertex graph $G = (V, E)$ such that:*

- \mathcal{M} runs for at most n rounds in any execution, and
- for each vertex $v \in V$ with actual coreness $k(v)$, the estimated coreness $\tilde{k}(v)$ of v satisfies, with high probability:

$$k(v) - \alpha \leq \tilde{k}(v) \leq k(v) + \alpha$$

where $\alpha = O(\epsilon^{-1} \log n \log \Delta)$ and Δ is the maximum degree of G .

At a high level, we implement a private version of the classic *peeling algorithm* of Matula and Beck [MB83]. This algorithm deletes vertices in increasing order of their coreness by repeatedly performing the following steps, until the graph is empty:

1. compute the minimum degree, d , among all vertices currently in the graph,
2. iteratively delete vertices with degree at most d from the graph, until none are left, and
3. report the coreness of each vertex deleted in the preceding step as d .

To do so, a natural approach is for each vertex to send its *noisy* degree to the server in each round. That is, it sends its degree, plus some random noise to ensure privacy. Given these noisy degrees, the server can determine the set of vertices to be deleted in that round and broadcast it to every vertex. Upon receiving the server’s message, each (surviving) vertex can determine which of its neighbors have been deleted and update its degree accordingly.

The main difficulty with this approach is generating the noisy degrees. In particular, a vertex may participate in many rounds (e.g., $\Omega(n)$ on a line graph) and, hence, send many noisy degrees. Naïvely using fresh noise each time requires too much noise to account for the privacy loss.

Our observation is that substantially less noise is needed if each vertex v uses a private *continual counting* mechanism, \mathcal{C}_v , instead. Specifically, v first generates a noisy degree by adding Laplace noise to its initial degree and sends this in the first round. If v is not deleted in a round, then it inserts the (absolute) *change* in its degree, as a result of zero or more of its neighbors being deleted in the round, into \mathcal{C}_v . Its initial noisy degree, minus the count that it receives from \mathcal{C}_v , is the noisy degree that it sends in the next round.

Algorithms 1 and 2 contain the pseudocode for the server and vertices, respectively.

Algorithm 1: Server algorithm for coreness estimation on a graph $G = (V, E)$, in the local model.

```

1  $d \leftarrow 0$ 
2 foreach round  $t = 1, 2, \dots, n$  do
3   receive noisy degree  $\tilde{d}_t(v)$  from each vertex  $v \in A_t := V \setminus (S_1 \cup \dots \cup S_{t-1})$ 
4    $d \leftarrow \max\{d, \min\{\tilde{d}_t(v) \mid v \in A_t\}\}$ 
5    $S_t \leftarrow \{v \in A_t \mid \tilde{d}_t(v) \leq d\}$ 
6   broadcast  $S_t$  to all vertices in  $A_t$  and set estimate  $\tilde{k}(v) = d$  for each vertex  $v \in S_t$ 
7   if  $V \setminus (S_1 \cup \dots \cup S_t) = \emptyset$  then return estimate vector  $(\tilde{k}(v) : v \in V)$ 

```

Algorithm 2: User algorithm for coreness estimation on a graph $G = (V, E)$ in the local model; code for each vertex $v \in V$. We assume that v is given its set of neighbors, N_v , in G , as input and that v maintains an $\epsilon/2$ -differentially private continual counting mechanism, \mathcal{C}_v , supporting up to n adaptive insertions.

```

1  $\tilde{d}_1(v) \leftarrow |N_v| + \text{Lap}(4/\epsilon)$ 
2 foreach round  $t = 1, 2, \dots, n$  do
3   send  $\tilde{d}_t(v)$  to server and receive message  $S_t$  from server
4   if  $v \in S_t$  then terminate
5    $\tilde{d}_{t+1}(v) \leftarrow \tilde{d}_1(v) - \mathcal{C}_v.\text{INSERT}(|N_v \cap S_t|)$ 

```

3.1 Privacy

The key point is as follows: if the transcript (containing the sequence of outputs of each continual counting mechanism) is the same for the first $t - 1$ rounds in two executions of the mechanism on neighboring graphs, then the sequence of inputs inserted into each continual counting mechanism is nearly the same in both executions.

Observation 1. Suppose the mechanism is executed on two graphs that differ by an edge e . If the transcript is the same in the first $t - 1$ rounds in both executions, then the following holds.

1. The set of vertices deleted (i.e., the message broadcast by the server) in each round $r < t$ is the same in both executions.
2. For each vertex that is not an endpoint of e , its initial degree and subsequent changes in degree are the same in each round $r < t$ in both executions.

3. For each endpoint v of e , its initial degree differs by one in the two executions. It has a change in degree that differs (by one) in a round $r < t$ only if the other endpoint was deleted in round $r - 1$, and v itself was not deleted in round $r - 1$ (or earlier).

Since the initial noisy degree is differentially private, the transcript of the first round is differentially private. By the preceding observation, to ensure that a transcript of subsequent rounds has roughly equal probability of occurring when the mechanism is executed on two neighboring graphs, it suffices for the counting mechanisms to have roughly equal probability of producing the same outputs, when executed on streams that *collectively* differ in at most one input (namely, at most one change in degree of an endpoint of the differing edge). A subtle point is that each counting mechanism needs to support *adaptive* insertions, as each input to the mechanism depends on its previous outputs. In Appendix B, we formally prove the following lemma.

Lemma 1. *If the continual counting mechanism of each vertex is $\epsilon/2$ -differentially private and supports up to n adaptive insertions, then the transcript is ϵ -edge differentially private.*

3.2 Accuracy

The noisy degrees sent to the server may differ significantly from the actual degrees. Hence, in each round, the server may incorrectly delete some vertices or fail to delete others. We show that if there is a bound on the error of the noisy degrees (of all vertices and in all rounds), then the same bound holds on the error of the coreness estimates. The proof is in Appendix B.

Lemma 2. *Let $G = (V, E)$ be a graph. If each noisy degree differs from its corresponding actual degree by at most α , then, for every vertex $v \in V$ with actual coreness $k(v)$, the estimated coreness $\tilde{k}(v)$ of v satisfies $k(v) - \alpha \leq \tilde{k}(v) \leq k(v) + \alpha$. Furthermore, every vertex in $G[U]$ has (induced) degree at least $\tilde{k}(v) - \alpha$, where U is the set of all vertices u with estimated coreness $\tilde{k}(u) \geq \tilde{k}(v)$.*

Note that the second statement of Lemma 2 is only needed for the accuracy proof of the approximate densest subgraph algorithm in Appendix C.2. By Laplace tail bounds with $\beta = 1/n^{\Omega(1)}$ (Fact 2) and a union bound over all vertices, with high probability, $\deg_G(v) - O(\epsilon^{-1} \log n) \leq \tilde{d}_1(v) \leq \deg_G(v) + O(\epsilon^{-1} \log n)$ for all $v \in V$. Hence, the additive error of $\tilde{d}_{t+1}(v)$ is $O(\epsilon^{-1} \log n)$, plus the error of the continual counting mechanism, with high probability. Note that the length of any input sequence is bounded by $T = n$ and its prefix sums are bounded by $n_T = \Delta$. Hence, taking $\beta = n^{-\Omega(1)}$, the sparse-vector counting mechanism has additive error $O(\epsilon^{-1} \log n \log \Delta)$ with high probability (Theorem 2). Combining this with Lemma 2 immediately yields the following.

Lemma 3. *With high probability, for every vertex $v \in V$, the estimated coreness of v differs from the actual coreness of v by at most $O(\epsilon^{-1} \log n \log \Delta)$ when using the sparse-vector counting mechanism.*

3.3 Further applications

[DLR⁺22] gave a memoryless local ϵ -edge differentially private mechanism by building on [LSY⁺22] for approximate core decomposition. We show in Appendix C.1 that our framework can be used in this setting to obtain the following guarantee.

Theorem 4. *For any $\epsilon, \eta > 0$, there is a memoryless local ϵ -edge differentially private mechanism \mathcal{M} that returns an estimate of the coreness of each vertex in an n -vertex graph $G = (V, E)$ such that:*

- \mathcal{M} runs for $O(\log_{1+\eta} n \log n)$ rounds in any execution, and
- for each vertex $v \in V$ with actual coreness $k(v)$, the estimated coreness $\tilde{k}(v)$ of v satisfies:

$$k(v) - \alpha \leq \tilde{k}(v) \leq (2 + \eta)k(v) + \alpha,$$

where $\alpha = O(\epsilon^{-1} \log n \log \log_{1+\eta} n)$ with probability $1 - n^{-\Omega(1)}$.

We also show in Appendix C.2 how to obtain the following guarantees for densest subgraph using our results for coreness approximation.

Theorem 5. *For any $\epsilon, \eta > 0$, there are memoryless local ϵ -edge differentially private mechanisms \mathcal{M} and \mathcal{M}' that return subsets of vertices U and U' from a given n -vertex graph G , respectively, such that:*

- *\mathcal{M} and \mathcal{M}' run for at most n and $O(\log_{1+\eta} n \log \log n)$ rounds, respectively, in any execution.*
- *With probability $1 - n^{-\Omega(1)}$, the density of $G[U]$ and $G[U']$ is at least $\rho^*/2 - \alpha$ and $\rho^*/(4 + \eta) - \alpha'$, respectively, where ρ^* is the maximum density of any subgraph of G , $\alpha = O(\epsilon^{-1} \log n \log \Delta)$, $\alpha' = O(\epsilon^{-1} \log_{1+\eta} n \log \log n)$, and Δ is the maximum degree of G .*

4 Lower bounds for core decomposition

Let $\epsilon > 0$ be any positive constant. We first give a lower bound in the centralized model (that carries over to the local model) and then give a lower bound for the local model.

Theorem 6. *Let $\gamma \geq 1$ be a constant and let V be a set of n vertices. Suppose that \mathcal{M} is an ϵ -edge differentially private mechanism in the centralized model that estimates the coreness of every vertex in a given graph G on V such that, for all vertices $v \in V$ with actual coreness $k(v)$, the estimated coreness $\tilde{k}(v)$ of v satisfies*

$$\gamma^{-1}k(v) - \alpha \leq \tilde{k}(v) \leq \gamma k(v) + \alpha \text{ for all } v \in V \text{ simultaneously with probability at least } p.$$

Then $\alpha = \Omega(\gamma^{-1} \log(np))$.

Proof. Let $d = \lceil (2\alpha + 1)\gamma \rceil$ and let G be any $(d + 1)$ -regular graph on V . For each vertex $v \in V$, let G_v be the same as G , except that all $d + 1$ edges incident to v are removed. Observe that, in G_v , v has coreness 0, while all other vertices $u \neq v$ have coreness at least d . Let \mathcal{C}_v be the set of all coreness estimate vectors \tilde{k} such that $\tilde{k}(v) \leq \alpha$ and $\tilde{k}(u) > \alpha$ for all $u \neq v$. Observe that \mathcal{C}_v and \mathcal{C}_u are disjoint for $u \neq v$. Furthermore, since $\gamma^{-1}d - \alpha = \alpha + 1$, $\mathcal{M}(G_v)$ must return $\tilde{k}(v) \leq \alpha$ and for all $u \neq v$, $\tilde{k}(u) \geq \alpha + 1$, i.e., an estimate in \mathcal{C}_v with probability at least p .

Since G and G_v differ by $d + 1$ edges and \mathcal{M} satisfies ϵ -edge differential privacy, we have that $\Pr(\mathcal{M}(G) \in \mathcal{C}_v) \geq e^{-\epsilon(d+1)} \Pr(\mathcal{M}(G_v) \in \mathcal{C}_v) \geq e^{-\epsilon(d+1)}p$. Since \mathcal{C}_u and \mathcal{C}_v are disjoint for $u \neq v$, $1 \geq \Pr(\bigcup_{v \in V} (\mathcal{M}(G) \in \mathcal{C}_v)) = \sum_{v \in V} \Pr(\mathcal{M}(G) \in \mathcal{C}_v) \geq ne^{-\epsilon(d+1)}p$. By rearranging, $\lceil (2\alpha + 1)\gamma \rceil = d \geq \ln(np) - 1$. Therefore, $\alpha = \Omega(\gamma^{-1} \log(np))$. \square

Let V be a set of $2n + 1$ vertices and let $x \in V$ be a fixed vertex. Consider any ϵ -edge differentially private mechanism \mathcal{M} that *non-interactively* (i.e., in a single round) estimates the coreness of x in a given graph on V , in the local model. We show that there is a large family of graphs on which \mathcal{M} has constant probability of returning an estimate with $\Omega(\sqrt{n})$ error for x .

Theorem 7. *For any constant $\epsilon > 0$, there exists a constant $0 < \eta < \frac{1}{2}$ such that the following holds. Suppose that \mathcal{M} is a non-interactive ϵ -edge differentially private local mechanism that estimates the coreness of a fixed vertex, x , in an arbitrary n -vertex graph such that, if $k(x)$ is the actual coreness of x , then the estimated coreness $\tilde{k}(x)$ of x satisfies:*

$$k(x) - \alpha \leq \tilde{k}(x) \leq k(x) + \alpha \text{ with probability at least } \frac{1}{2} + \eta.$$

Then there is a family of n -vertex graphs of size $2^{\Omega(n)}$ on which $\alpha = \Omega(\sqrt{n})$.

Our approach is to reduce to a known lower bound on the error of privately answering a linear number of random *inner product queries* on a secret dataset $X \in \{0, 1\}^n$. Here, two datasets X and X' are *neighboring* if they differ in at most one coordinate, a query is specified by a vector $Q \in \mathbb{R}^n$, and

the *error* of a response r to query Q is $|r - \langle Q, X \rangle|$. [ELRS23] were the first to use this approach to prove lower bounds in the local model.

Roughly speaking, the lower bound says that no ϵ -differentially private mechanism (with a trusted curator) can answer $O(n)$ random inner product queries in $\{-1, 1\}^n$ so that, with constant probability, a large fraction of the answers have $o(\sqrt{n})$ error. (Otherwise, an “attacker” can use such queries to “reconstruct” X with high accuracy, violating privacy.) In Appendix D, we formally state the lower bound and use it to prove the following modified variant, adapted to inner product queries in $\{0, 1\}^n$. The idea is to convert a mechanism answering inner product queries in $\{0, 1\}^n$ into a mechanism answering inner product queries in $\{-1, 1\}^n$ with roughly the same error and privacy loss.

Theorem 8. *For any constants $\epsilon > 0$ and $\frac{1}{20} > \delta > 0$, there is a constant $0 < \eta < \frac{1}{2}$ such that no (ϵ, δ) -differentially private mechanism can answer $m = O(n)$ random inner product queries in $\{0, 1\}^n$ on a secret dataset $X \in \{0, 1\}^n$ such that, with probability at least $\Omega(\sqrt{\delta})$, a $(\frac{1}{2} + \eta)$ -fraction of its answers have $o(\sqrt{n})$ error.*

To apply Theorem 8, we construct a 2ϵ -differentially private mechanism \mathcal{N} that answers m random inner product queries in $\{0, 1\}^n$. For each query, \mathcal{N} simulates \mathcal{M} on a *query graph* in which the coreness of x is roughly the intended answer to the query. We show that, with constant probability, a large fraction of \mathcal{N} ’s responses will have the same error as \mathcal{M} . Therefore, \mathcal{M} has $\Omega(\sqrt{n})$ error with constant probability. All proofs are in Appendix E.

Query graphs Fix a partition (A, B) of $V \setminus \{x\}$ with $|A| = |B|$ and an enumeration, a_1, \dots, a_n , of the vertices in A . Let $Q \in \{0, 1\}^n$ be an arbitrary inner product query. The *query graph* for Q on dataset X is the graph $G_X(Q)$ on V defined as follows.

1. For all $i \in [n]$, $X_i \in \{0, 1\}$ indicates whether a_i is adjacent to x and $Q_i \in \{0, 1\}$ indicates whether a_i is adjacent to (0) no vertex in B or (1) every vertex in B .
2. The vertices of B form a clique, which x is not adjacent to.

The idea is that Q is used to bound the coreness of vertices in A . Specifically, if $Q_i = 1$, then a_i is adjacent to every vertex in B and, hence, has coreness at least $n - 1$. Otherwise, a_i can only be adjacent to x , so it has coreness at most 1. Thus, the coreness of x is roughly the number of neighbors $a_i \in A$ such that $Q_i = 1$, i.e., $\langle Q, X \rangle$. (If Q is the all-zero vector, then the coreness of x may still be 1.)

Lemma 4. *The coreness of x in $G_X(Q)$ is either $\langle Q, X \rangle$ or $\langle Q, X \rangle + 1$.*

Answering random queries Naïvely answering queries by simulating \mathcal{M} on each of the corresponding query graphs leads to prohibitively large privacy loss. Our key observation is that the vertices whose neighborhoods in the query graphs depend on the secret dataset X , namely $\{x\} \cup A$, only have a few possible neighborhoods among all query graphs on X . This suggests that, to reduce the privacy loss, we can generate only a few messages for these vertices and *reuse* them in all simulations.

Observation 2. The following holds for any secret dataset $X \in \{0, 1\}^n$.

1. The neighbors of x are the same in all query graphs on X .
2. Each vertex in A only has two possible neighborhoods among all query graphs on X (namely, $\{x\}$ and $B \cup \{x\}$, if $X_i = 1$, and \emptyset and B otherwise).
3. For all datasets X , the neighborhood of each vertex $b \in B$ in the query graph $G_X(Q)$ is a function of Q , i.e., $(B \setminus \{b\}) \cup \{a_i \in A \mid Q_i = 1\}$.

Specifically, our mechanism \mathcal{N} answers m random inner product queries $Q^{(1)}, \dots, Q^{(m)} \in \{0, 1\}^n$ on a secret dataset $X \in \{0, 1\}^n$ as follows.

1. Let $\mathbf{0} \in \{0, 1\}^n$ and $\mathbf{1} \in \{0, 1\}^n$ be the all-zeros and all-ones vectors, respectively.
 - (a) Run the user algorithm of x (specified by \mathcal{M}) on its neighborhood in $G_X(\mathbf{0})$ to obtain the message π_x .
 - (b) For each $i \in [n]$: run the user algorithm of a_i (specified by \mathcal{M}) on its neighborhood in $G_X(\mathbf{0})$ and $G_X(\mathbf{1})$ to obtain the messages $\pi_{a_i}(0)$ and $\pi_{a_i}(1)$, respectively.
2. For each query $Q^{(j)}$, *simulate* a run of \mathcal{M} on $G_X(Q^{(j)})$ as follows:
 - (a) Run the user algorithm of each vertex $b \in B$ (specified by \mathcal{M}) on its neighborhood in the query graph $G_X(Q^{(j)})$ to obtain the message $\pi_b^{(j)}$.
 - (b) Run the server's algorithm (specified by \mathcal{M}) on the transcript
$$\{\pi_x\} \cup \{\pi_{a_i}(Q_i^{(j)}) \mid i \in [n]\} \cup \{\pi_b^{(j)} \mid b \in B\}$$
to obtain an estimate $\tilde{k}_j(x)$ of the coreness of vertex x in the query graph $G_X(Q^{(j)})$.
3. Return $(\tilde{k}_j(x) : j \in [m])$.

Analysis Since we run the user algorithm of each vertex in $\{x\} \cup A$ on its respective neighborhood in at most two graphs and \mathcal{M} is ϵ -edge differentially private, the messages generated in the first step are collectively 2ϵ -edge differentially private. Since the pairs of graphs $G_X(\mathbf{0}), G_{X'}(\mathbf{0})$ and $G_X(\mathbf{1}), G_{X'}(\mathbf{1})$ each differ in at most one edge when datasets X and X' are neighboring, it follows that the first step is 2ϵ -differentially private. The other steps can be viewed as simply post-processing the outputs (messages) of the first step. Therefore, the entire procedure is 2ϵ -differentially private.

Lemma 5. \mathcal{N} is 2ϵ -differentially private.

Lemma 4 implies that \mathcal{N} produces a response to a query with additive error exceeding $\alpha + 1$ only if the simulation of \mathcal{M} on the corresponding query graph produces an estimate of the coreness of x that has additive error exceeding α . We note that, since the messages of vertices in $\{x\} \cup A$ are reused, the simulations are *not* independent. Nonetheless, Markov's inequality can be used to bound the probability that a large number of simulations have error exceeding α .

Lemma 6. Suppose that \mathcal{M} has error exceeding α with probability at most β . Then, with probability at least $1 - 1/\gamma$, at least a $(1 - \gamma\beta)$ -fraction of the responses of \mathcal{N} have error at most $\alpha + 1$.

To summarize, we have shown that, if \mathcal{M} estimates the coreness of a fixed vertex with error exceeding α with probability at most $\beta = \frac{1}{2} - \eta$, then with constant probability (say $1 - 1/1.001 \approx 0.001$), a large fraction ($1 - 1.001\beta \approx \frac{1}{2} + \eta$) of the responses of \mathcal{M} have error at most $\alpha + 1$. Since \mathcal{N} is 2ϵ -differentially private, it is $(2\epsilon, \delta)$ -differentially private for any $\delta > 0$. Therefore, by the contrapositive of Theorem 8, $\alpha = \Omega(\sqrt{n})$.

Acknowledgements We would like to thank an anonymous reviewer for suggesting the strategy to simulate the user/server version of the local model (in which users have local memory) in the local randomizer model.

Funding *M. Henzinger and A. R. Sricharan:* This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 101019564) and the Austrian Science Fund (FWF) project Z 422-N, project I 5982-N, and project P 33775-N, with additional funding from the *netidee SCIENCE Stiftung*, 2020–2024.



References (with cited pages)

- CSS10** T. H. Hubert Chan, Elaine Shi, and Dawn Song.
[Private and continual release of statistics.](#)
IACR Cryptol. ePrint Arch. 2010, p. 76. (7)
- De12** Anindya De.
[Lower bounds in differential privacy.](#)
In *Theory of Cryptography Conference, TCC 2012*. Springer. (3 and 19)
- DKLV23** Michael Dinitz, Satyen Kale, Silvio Lattanzi, and Sergei Vassilvitskii.
[Improved differentially private densest subgraph: Local and purely additive.](#)
CoRR 2023. (1, 2, 3, and 6)
- DLL23** Laxman Dhulipala, George Z. Li, and Quanquan C. Liu.
[Near-optimal differentially private k-core decomposition.](#)
CoRR 2023. (1, 2, 3, 4, 5, and 18)
- 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 *Symposium on Foundations of Computer Science, FOCS 2022*. IEEE. (1, 2, 3, 4, 6, 9, 14, 17, and 18)
- DMNS06** Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith.
[Calibrating noise to sensitivity in private data analysis.](#)
In *Theory of Cryptography Conference, TCC 2006*. Springer. (1 and 5)
- DNPR10** Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N. Rothblum.
[Differential privacy under continual observation.](#)
In *Symposium on Theory of Computing, STOC 2010*. ACM. (2, 4, and 7)
- DNRR15** Cynthia Dwork, Moni Naor, Omer Reingold, and Guy N. Rothblum.
[Pure differential privacy for rectangle queries via private partitions.](#)
In *Conference on the Theory and Application of Cryptology and Information Security, ASIACRYPT 2015*. Springer. (7)
- DR14** Cynthia Dwork and Aaron Roth.
[The algorithmic foundations of differential privacy.](#)
Found. Trends Theor. Comput. Sci. 2014, 9(3-4):211–407. (6)
- ELRS23** Talya Eden, Quanquan C. Liu, Sofya Raskhodnikova, and Adam D. Smith.
[Triangle counting with local edge differential privacy.](#)
In *International Colloquium on Automata, Languages, and Programming, ICALP 2023*. (3, 6, and 11)
- FHO21** Hendrik Fichtenberger, Monika Henzinger, and Wolfgang Ost.
[Differentially private algorithms for graphs under continual observation.](#)
In *European Symposium on Algorithms, ESA 2021*. LIPIcs. (5)
- FHS22** Alireza Farhadi, MohammadTaghi Hajiaghayi, and Elaine Shi.
[Differentially private densest subgraph.](#)
In *International Conference on Artificial Intelligence and Statistics, AISTATS 2022*. PMLR. (1 and 5)
- HT10** Moritz Hardt and Kunal Talwar.
[On the geometry of differential privacy.](#)
In *Symposium on Theory of Computing, STOC 2010*. ACM. (3)

- JMNR19** Matthew Joseph, Jieming Mao, Seth Neel, and Aaron Roth.
The role of interactivity in local differential privacy.
 In *Symposium on Foundations of Computer Science, FOCS 2019*. IEEE. (6)
- JRSS23** Palak Jain, Sofya Raskhodnikova, Satchit Sivakumar, and Adam D. Smith.
The price of differential privacy under continual observation.
 In *International Conference on Machine Learning, ICML 2023*. PMLR. (7)
- LSY⁺22** Quanquan C. Liu, Jessica Shi, Shangdi Yu, Laxman Dhulipala, and Julian Shun.
Parallel batch-dynamic algorithms for k-core decomposition and related graph problems.
 In *Symposium on Parallelism in Algorithms and Architectures, SPAA 2022*. ACM. (9 and 17)
- MB83** David W. Matula and Leland L. Beck.
Smallest-last ordering and clustering and graph coloring algorithms.
J. ACM 1983, 30(3):417–427. (7)
- NV21** Dung Nguyen and Anil Vullikanti.
Differentially private densest subgraph detection.
 In *International Conference on Machine Learning, ICML 2021*. PMLR. (1 and 3)

A Formal definition of differential privacy in the local model

We formally define our model in this section.

A *local mechanism* \mathcal{M} is defined by a pair of functions $f : \mathcal{T} \times \mathcal{L} \times \{0, 1\}^* \rightarrow \Sigma^* \times \mathcal{L}$ and $g : \mathcal{T} \times \{0, 1\}^* \rightarrow \Sigma^*$, where \mathcal{T} is the set of all possible \mathcal{M} -transcripts⁴, Σ^* is the set of all possible messages, \mathcal{L} is the set of all possible local memory states, and $\{0, 1\}^*$ is the set of all binary strings. An *execution* of a local mechanism $\mathcal{M} = \{f, g\}$ on an input graph $G = (V, E)$ generates a \mathcal{M} -*transcript* $\hat{\tau}$ as follows. Initially, the local memory of each user v is $\ell_v^0 := (v, N_v)$, where N_v is the set of neighbors of v in G , and the (0 -round) \mathcal{M} -transcript $\hat{\tau}$ is empty. At the start of the $(t + 1)$ -th *round*, each user v simultaneously generates fresh random bits $\gamma_v^{t+1} \in \{0, 1\}^*$ and evaluates $(\sigma_v^{t+1}, \ell_v^{t+1}) := f(\hat{\tau}, \ell_v^t, \gamma_v^{t+1})$. Then the local memory of each user v is set to ℓ_v^{t+1} and $\hat{\tau}$ is updated to $\hat{\tau} \odot (\sigma_v^{t+1} : v \in V)$, where \odot is the string concatenation operation. Subsequently, the server generates fresh random bits $\gamma_s^{t+1} \in \{0, 1\}^*$ and $\hat{\tau}$ is updated to $\hat{\tau} \odot g(\hat{\tau}, \gamma_s^{t+1})$. (At this point, we say $\hat{\tau}$ is a $(t + 1)$ \mathcal{M} -round transcript.) If $g(\hat{\tau}, \gamma_s^{t+1})$ is not the empty string, then the next round begins. Otherwise, the execution ends and the resulting $\hat{\tau}$ is a *complete* \mathcal{M} -transcript (on input G).

We say that \mathcal{M} is ϵ -*edge differentially private* if the output (transcript) of \mathcal{M} is ϵ -differentially private on graphs that differ by at most one edge. We say that \mathcal{M} is *memoryless* if f is the identity function on its second argument, i.e., for all $\hat{\tau}, \ell, \gamma$, there exists σ such that $f(\hat{\tau}, \ell, \gamma) = (\sigma, \ell)$. In particular, in a memoryless local mechanism, the local memory of each user v is always set to (v, N_v) .

Note that our definition implicitly assumes that the server includes each message it has received (i.e., the transcript so far) in its next “message” to the users (which is $\hat{\tau} \odot g(\hat{\tau}, \gamma_s^{t+1})$). This is why each user’s “message” is a function of the transcript so far, even if each user is memoryless. Similarly, for each user’s “message” to the server. Hence, since the server does not have any private data, it does not need local memory. In particular, its message can include the random string γ_s^{t+1} it sampled.

An ϵ -differentially private memoryless local mechanism $\mathcal{M} = \{f, g\}$ in our model can be modeled by a protocol \mathcal{A} , as defined in [DLR⁺22] and presented in Section 2. The main difference is that there is no server in the latter model and \mathcal{A} needs to map a partial \mathcal{A} -transcript to an assignment of local randomizers to the participating users of the next round, one per user. To handle this, we hard-code the server’s message in \mathcal{M} that is sent at the end of the previous round into the local randomizer assigned to each user in the current round of \mathcal{A} . To simulate the random bits of the server, one fixed user u sends

⁴To distinguish between transcripts arising from local mechanisms and protocols using local randomizers, we will use the terminology \mathcal{M} -transcript and \mathcal{A} -transcript, respectively, in this section.

additional random bits along with their message. This allows the users to agree on the random string generated by the server in the previous round.

More formally, let $\mathcal{M} = \{f, g\}$ be any ϵ -edge differentially private memoryless local mechanism. We define a protocol \mathcal{A} in the local randomizer model so that each t -round \mathcal{A} -transcript *corresponds* to a t -round \mathcal{M} -transcript. Indeed, the (empty) 0-round \mathcal{A} -transcript *corresponds* to the (empty) 0-round \mathcal{M} -transcript. Now consider any t -round \mathcal{A} -transcript τ that corresponds to a t -round \mathcal{M} -transcript $\hat{\tau}$. If $\hat{\tau}$ is non-empty and the last message of the server in $\hat{\tau}$ is the empty string, then we define $\mathcal{A}(\tau) = \perp$. Otherwise, we define $\mathcal{A}(\tau) = (S_U^{t+1}, S_R^{t+1}, S_\epsilon^{t+1})$, where

- the set S_U^{t+1} of participating users is V ,
- the local randomizer $S_R^{t+1}(v)$ assigned to user v generates fresh random bits $\gamma_v \in \{0, 1\}^*$ (and $\gamma_s \in \{0, 1\}^*$, if $v = u$) and then outputs $S_O^{t+1}(v) = \sigma_v$ for $v \neq u$ and $S_O^{t+1}(u) = (\sigma_u, \gamma_s)$, where $(\sigma_v, (v, N_v)) = f(\hat{\tau}, (v, N_v), \gamma_v)$, i.e., σ_v is the first parameter of $f(\hat{\tau}, (v, N_v), \gamma_v)$ ⁵, and
- the privacy parameter of $S_R^{t+1}(v)$ is $S_\epsilon^{t+1}(v) = \epsilon$.

In this case, the resulting $(t + 1)$ -round \mathcal{A} -transcript

$$\tau' = \tau \odot (S_U^{t+1}, S_R^{t+1}, S_\epsilon^{t+1}, S_O^{t+1})$$

corresponds to the $(t + 1)$ -round \mathcal{M} -transcript

$$\hat{\tau}' = \hat{\tau} \odot (\sigma_v : v \in V) \odot g(\hat{\tau} \odot (\sigma_v : v \in V), \gamma_s),$$

We note that the set \mathcal{R} of local randomizers used by \mathcal{A} essentially consists of all functions f with the first parameter hard-coded to be any value of \mathcal{T} , which is used to encode the corresponding \mathcal{M} -transcript and, in particular, the messages sent by the server to the users, as given by g . Furthermore, we emphasize that, in the specification of $\mathcal{A}(\tau')$, the corresponding \mathcal{M} -transcript $\hat{\tau}'$ is used to determine if $\mathcal{A}(\tau') = \perp$ and, if $\mathcal{A}(\tau') \neq \perp$, to determine the output of each local randomizer. Since (the output of) each local randomizer is chosen as a function of the corresponding \mathcal{M} -transcript, it follows that \mathcal{A} is ϵ -edge differentially private.

B Missing proofs for Algorithm 1 and Algorithm 2

Lemma 1. *If the continual counting mechanism of each vertex is $\epsilon/2$ -differentially private and supports up to n adaptive insertions, then the transcript is ϵ -edge differentially private.*

Proof. We view a transcript as a vector $\pi = (\pi_1, \dots, \pi_n)$ of functions $\pi_t : V \rightarrow \mathbb{R} \cup \{\perp\}$, for $t \in [n]$, where $\pi_t(v)$ is the message sent by vertex $v \in V$ in round t , or \perp if it sends no message. In particular, since the message, S_t , broadcast by the server in round t is a deterministic function of (π_1, \dots, π_t) , we omit it from the transcript.

Consider two graphs G and G' that differ in an edge e . Let Π and Π' denote the (random) transcript of the mechanism when executed on G and G' , respectively. We claim that for any set of (valid) transcripts \mathcal{A} , $\Pr(\Pi \in \mathcal{A}) \leq e^\epsilon \Pr(\Pi' \in \mathcal{A})$. Indeed, fix a transcript $\pi \in \mathcal{A}$. We will show that $\Pr(\Pi = \pi) \leq e^\epsilon \Pr(\Pi' = \pi)$. We split the analysis into two steps, each consisting of one mechanism that we analyze separately: (1) The mechanism \mathcal{M}_1 that discloses $\tilde{d}_1(v)$ for each vertex v , and (2) the mechanism \mathcal{M}_2 that discloses $\tilde{d}_t(v)$ for each vertex v and all $t > 1$.

(1) Consider a vertex $v \in V$. If $v \notin e$, then the degree of v is the same in both graphs, while for the two endpoints of e , the difference is exactly 1. Consider the vector containing the degrees of all vertices, in some fixed order. Then the ℓ_1 -norm of the difference of the vector for G and the vector for G' is 2. Thus, adding Laplacian noise with parameter $\epsilon/4$ to each entry in the vector guarantees that \mathcal{M}_1 is $\epsilon/2$ -edge differentially private.

⁵Note the use of the \mathcal{M} -transcript $\hat{\tau}$.

(2) For each vertex $v \in V$, let

$$p_v := \prod_{t=2}^n \Pr(\Pi_t(v) = \pi_t(v) \mid \wedge_{r=1}^{t-1} \Pi_r = \pi_r) \text{ and } p'_v := \prod_{t=2}^n \Pr(\Pi'_t(v) = \pi_t(v) \mid \wedge_{r=1}^{t-1} \Pi'_r = \pi_r).$$

By definition, $\Pi_1(v)$ is the value returned by \mathcal{M}_1 . For $t > 1$, conditioned on the transcript of the first $t - 1$ rounds being $(\pi_1, \dots, \pi_{t-1})$ in both executions, the sequence of messages, S_1, \dots, S_{t-1} , broadcast by the server in the first $t - 1$ rounds is the same in both executions. If v participates in round t , i.e., $v \notin S_1 \cup \dots \cup S_{t-1}$, then $\Pi_t(v)$ is the count returned by \mathcal{C}_v immediately after v inserts the input $c_t(v) := -|N_v \cap S_{t-1}|$ (its $(t - 1)$ th change in degree in the execution on G). Otherwise, $\Pi_t(v) = \perp$.

It follows that p_v is the probability that \mathcal{C}_v returns the sequence of outputs $(\pi_1(v), \pi_2(v), \dots)$ on the (adaptive) sequence of inputs $(c_1(v), c_2(v), \dots)$. Defining $c'_1(v), c'_2(v), \dots$ similarly for the execution on G' , we have that p'_v is the probability that \mathcal{C}_v returns the sequence of outputs $(\pi_1(v), \pi_2(v), \dots)$ on the sequence of inputs $(c'_1(v), c'_2(v), \dots)$.

By Observation 1, the inputs to \mathcal{C}_v are the same in both executions conditioned on the previous outputs being the same, i.e., $(c_1(v), c_2(v), \dots) = (c'_1(v), c'_2(v), \dots)$ for each vertex $v \notin e$. By coupling the random bits of each vertex $v \notin e$ in the two executions, this implies that each output of \mathcal{C}_v has the same probability of occurrence under both executions (conditioned on the previous outputs being equal). It follows that $p_v = p'_v$ for $v \notin e$ and, hence, $\prod_{v \notin e} p_v = \prod_{v \notin e} p'_v$.

On the other hand, for each endpoint $v \in e$, and $t > 1$, $c_t(v) \neq c'_t(v)$ only if S_{t-1} contains the other endpoint $u \in e$, but not v . In this case, $(c_2(v), c_3(v), \dots)$ and $(c'_2(v), c'_3(v), \dots)$ differ by one element, namely $|c_t(v) - c'_t(v)| = 1$ and $(c_2(u), c_3(u), \dots)$ equals $(c'_2(u), c'_3(u), \dots)$. If this is case does not happen, i.e., if u and v belong to the same set S_{t-1} , then $(c_2(v), c_3(v), \dots)$ and $(c'_2(v), c'_3(v), \dots)$ do not differ. Since \mathcal{C}_u and \mathcal{C}_v are $\epsilon/2$ -differentially private, it follows that $p_u = p'_u$ and $p_v \leq e^{\epsilon/2} p'_v$. Hence, $\prod_{v \in e} p_v \leq e^{\epsilon/2} \prod_{v \in e} p'_v$.

Finally, since each vertex independently generates its message in each round t :

$$\begin{aligned} \Pr(\Pi = \pi) &= \prod_{v \in V} \prod_{t=1}^n \Pr(\Pi_t(v) = \pi_t(v) \mid \wedge_{r=1}^{t-1} \Pi_r = \pi_r) = \prod_{v \notin e} p_v \times \prod_{v \in e} p_v \leq \prod_{v \notin e} p'_v \times e^{\epsilon/2} \prod_{v \in e} p'_v \\ &= e^{\epsilon/2} \prod_{v \in V} \prod_{t=1}^n \Pr(\Pi'_t(v) = \pi_t(v) \mid \wedge_{r=1}^{t-1} \Pi'_r = \pi_r) = e^{\epsilon/2} \Pr(\Pi' = \pi). \end{aligned}$$

This shows that \mathcal{M}_2 is $\epsilon/2$ -edge differentially private. By simple composition, the complete mechanism, which consists of the output of \mathcal{M}_1 and \mathcal{M}_2 , is ϵ -edge differentially private. \square

Lemma 2. *Let $G = (V, E)$ be a graph. If each noisy degree differs from its corresponding actual degree by at most α , then, for every vertex $v \in V$ with actual coreness $k(v)$, the estimated coreness $\tilde{k}(v)$ of v satisfies $k(v) - \alpha \leq \tilde{k}(v) \leq k(v) + \alpha$. Furthermore, every vertex in $G[U]$ has (induced) degree at least $\tilde{k}(v) - \alpha$, where U is the set of all vertices u with estimated coreness $\tilde{k}(u) \geq \tilde{k}(v)$.*

Proof. Fix any vertex $v \in V$. Let t be the round in which v is deleted, let d_t (resp. $\tilde{d}_t(v)$) be the value of the variable d (resp. message from vertex v) stored by (resp. received by) the server at the end of round t , and let r be the first round in which the server sets d to d_t . By definition, $\tilde{k}(v) = d_t$, $r \leq t$, and d_t is the minimum noisy degree sent in round r , i.e., $\tilde{d}_r(u) \geq d_t$ for each $u \in U$, where U is the set of non-deleted vertices at the start of round r . Since r is the first round where d is set to d_t and d is non-decreasing, it follows that U is the set of all vertices u with estimated coreness $\tilde{k}(u) \geq d_t$. By definition, $v \in U$ and every vertex in $G[U]$ has (actual) induced degree at least $d_t - \alpha$ since $\tilde{d}_r(u) \geq d_t$ and the noisy degree differs in every round at most α from its corresponding degree. By definition, G contains a subgraph H containing v in which every vertex has induced degree at least $k(v)$. Note that $k(v)$ is the largest integer for which this holds, so $k(v) \geq d_t - \alpha$ and, hence, $\tilde{k}(v) = d_t \leq k(v) + \alpha$. This shows the second inequality of the lemma.

Consider the first vertex, u , of H to be deleted. Let t' be the round in which u is deleted and let $d_{t'}$ be the value of d stored at the end of round t' . By definition, no vertex in H (in particular, v) has been deleted at the start of round t' . Hence, $t' \leq t$ and u sends a noisy degree $\tilde{d}_{t'}(u)$ that is at least $k(v) - \alpha$. On the other hand, since u is deleted in round t' , the noisy degree that it sends in that round is at most $d_{t'}$. Since d is non-decreasing, it follows that $k(v) - \alpha \leq d_{t'} \leq d_t = \tilde{k}(v)$. This shows the first inequality of the lemma. \square

C Further applications of continual counting

C.1 Approximate core decomposition

Let $\epsilon, \eta > 0$ be any positive constants. [DLR⁺22] gave a memoryless local ϵ -edge differentially private mechanism by building on [LSY⁺22] for approximate core decomposition on n -vertex graphs $G = (V, E)$ with the following one-sided *multiplicative* error guarantee: for every vertex $v \in V$ with actual coreness $k(v)$, the *estimated* coreness $\tilde{k}(v)$ of v (i.e., as reported by the mechanism) satisfies:

$$k(v) - \alpha \leq \tilde{k}(v) \leq (2 + \eta)k(v) + \alpha,$$

where $\alpha = O(\epsilon^{-1} \log_{1+\eta} n \log^2 n)$ with probability $1 - n^{-\Omega(1)}$.

Their mechanism has $T = O(\log_{1+\eta} n \log n)$ rounds, divided into $O(\log n)$ disjoint *phases*, each consisting of $O(\log_{1+\eta} n)$ consecutive rounds. In each round, the non-deleted vertices send their current noisy degrees to the server. Each noisy degree is generated by adding fresh “discrete Laplace” noise to the corresponding actual degree. If the round occurs during phase ϕ , then the server replies with the set of all vertices with noisy degree at most $(2 + \eta)^\phi$ (in that round), which are to be deleted, and estimates the coreness of each such vertex as $(2 + \eta)^{\phi-1}$. The surviving vertices update their degrees and continue to the next round.

It can be shown that, if the standard deviation of each noisy degree is $O(\epsilon^{-1}T)$, then the transcript of each round is ϵT^{-1} -edge differentially private. Therefore, by composition, the entire transcript is ϵ -edge differentially private. The accuracy of the estimates follows from the following lemma, using the fact that each noisy degree differs from its corresponding actual degree by $O(\epsilon^{-1}T \log n)$ with probability $1 - n^{-\Omega(1)}$. This lemma is implicitly proved in [DLR⁺22], so we omit it.

Lemma 7. *Suppose that every noisy degree differs from its corresponding actual degree by at most α . Then, for every vertex $v \in V$ with actual coreness $k(v)$:*

- *the estimated coreness $\tilde{k}(v)$ of v satisfies $k(v) - O(\alpha) \leq \tilde{k}(v) \leq (2 + \eta)k(v) + O(\alpha)$ and*
- *every vertex in $G[U]$ has (induced) degree at least $\tilde{k}(v)/(2 + \eta) - O(\alpha)$, where U is the set of all vertices u with estimated coreness $\tilde{k}(u) \geq \tilde{k}(v)$.*

Similar to Section 3, we observe that it is possible to obtain smaller additive error if each vertex v generates a noisy initial degree using Laplace noise, and then generates its later noisy degrees using a private continual counting mechanism, \mathcal{C}_v , which tracks the absolute change in its degree in each round. In particular, since Observation 1 also holds in this setting, to ensure that the transcript is ϵ -edge differentially private, it suffices to use $\text{Lap}(\epsilon/4)$ noise and for each \mathcal{C}_v to be $\epsilon/2$ -differentially private. The privacy proof is nearly verbatim the same as Lemma 1, so we omit it.

Each vertex makes $T = O(\log_{1+\eta} n \log n)$ insertions. Taking $\beta = (T \cdot n^{\Omega(1)})^{-1}$, the binary tree mechanism guarantees an error bound of $O(\epsilon^{-1} \log n \log \log_{1+\eta} n)$ (Theorem 1). Combining this with Lemma 7 immediately yields the following.

Theorem 4. *For any $\epsilon, \eta > 0$, there is a memoryless local ϵ -edge differentially private mechanism \mathcal{M} that returns an estimate of the coreness of each vertex in an n -vertex graph $G = (V, E)$ such that:*

- *\mathcal{M} runs for $O(\log_{1+\eta} n \log n)$ rounds in any execution, and*

- for each vertex $v \in V$ with actual coreness $k(v)$, the estimated coreness $\tilde{k}(v)$ of v satisfies:

$$k(v) - \alpha \leq \tilde{k}(v) \leq (2 + \eta)k(v) + \alpha,$$

where $\alpha = O(\epsilon^{-1} \log n \log \log_{1+\eta} n)$ with probability $1 - n^{-\Omega(1)}$.

C.2 Approximate densest subgraph

It can be shown that if $G[U^*]$ is a densest subgraph of G , then the (induced) degree of every vertex in $G[U^*]$ is at least its density, ρ^* . (In particular, if some vertex has degree less than ρ^* , then removing it yields a subgraph with higher density, which is impossible.) Hence, $G[U^*]$ is contained in the ρ^* -core of G and ρ^* is at most the maximum coreness, k^* , of any vertex in G .

Lemma 8 (Folklore). *Let G be any graph. If ρ^* is maximum density of any subgraph of G and k^* is the maximum coreness of any vertex of G , then $\rho^* \leq k^*$.*

Let H be the k^* -core of G . Since every vertex in H has (induced) degree at least k^* , there are at least $k^*|V(H)|/2$ edges in H . Therefore, the density of H is $|E(H)|/|V(H)| \geq k^*/2 \geq \rho^*/2$ and returning H yields a 2-approximate solution to the densest subgraph problem on G .

[DLL23] and [DLR⁺22] leveraged the preceding observation to develop differentially private 2-approximate and $(4 + \eta)$ -approximate solutions to the densest subgraph problem in the centralized and local model, respectively. Specifically, their mechanisms return a set of vertices U and U' such that the density of $G[U]$ and $G[U']$ is at least $\rho^*/2$ and $\rho^*/(4 + \eta)$, respectively. Their procedures may be summarized as follows.

- (1) Compute an estimate $\tilde{k}(v)$ of the coreness of every vertex v in G .
- (2) Return $\tilde{U}^* = \{v \in V : \tilde{k}(v) = \tilde{k}^*\}$, where $\tilde{k}^* = \max\{\tilde{k}(v) : v \in V\}$.

Since step (2) is simply post-processing, the entire procedure has the same differential privacy guarantees as the coreness estimation procedure employed in step (1). The next lemma gives sufficient conditions for the accuracy of the estimation to scale with the accuracy of the underlying coreness estimation procedure. Roughly speaking, [DLL23] assumes that their coreness estimation procedure satisfies the second property, without proof, while [DLR⁺22] implicitly proves it.

Lemma 9. *Let $\gamma \geq 1$ be a constant. Suppose that for each vertex $v \in V$ with actual coreness $k(v)$:*

- the estimated coreness $\tilde{k}(v)$ of v satisfies $k(v) - \alpha \leq \tilde{k}(v) \leq \gamma k(v) + \alpha$ and
- every vertex in $G[U]$ has (induced) degree at least $\tilde{k}(v)/\gamma - \alpha$, where U is the set of all vertices u with estimated coreness $\tilde{k}(u) \geq \tilde{k}(v)$.

Then the density of $G[\tilde{U}^*]$ is at least $\rho^*/2\gamma - O(\alpha)$.

Proof. By the first assumption, $\tilde{k}^* = \max\{\tilde{k}(v) \mid v \in V\} \geq \max\{k(v) - \alpha \mid v \in V\} = k^* - \alpha$. By the second assumption, every vertex in $G[\tilde{U}^*]$ has induced degree at least $\tilde{k}^*/\gamma - \alpha \geq k^*/\gamma - (1 + 1/\gamma)\alpha$. It follows that $|E(G[\tilde{U}^*])| \geq |V(G[\tilde{U}^*])|(k^*/\gamma - (1 + 1/\gamma)\alpha)/2$ and the density of $G[\tilde{U}^*]$ is $|E(G[\tilde{U}^*])|/|V(G[\tilde{U}^*])| \geq (k^*/\gamma - (1 + 1/\gamma)\alpha)/2$. By Lemma 8, this is at least $\rho^*/2\gamma - O(\alpha)$. \square

By employing our coreness estimation procedures from Sections 3 and C.1, which both satisfy the conditions of Lemma 9 (see Lemmas 3 and 7), we immediately obtain the following.

Theorem 5. *For any $\epsilon, \eta > 0$, there are memoryless local ϵ -edge differentially private mechanisms \mathcal{M} and \mathcal{M}' that return subsets of vertices U and U' from a given n -vertex graph G , respectively, such that:*

- \mathcal{M} and \mathcal{M}' run for at most n and $O(\log_{1+\eta} n \log \log n)$ rounds, respectively, in any execution.
- With probability $1 - n^{-\Omega(1)}$, the density of $G[U]$ and $G[U']$ is at least $\rho^*/2 - \alpha$ and $\rho^*/(4 + \eta) - \alpha'$, respectively, where ρ^* is the maximum density of any subgraph of G , $\alpha = O(\epsilon^{-1} \log n \log \Delta)$, $\alpha' = O(\epsilon^{-1} \log_{1+\eta} n \log \log n)$, and Δ is the maximum degree of G .

D Inner product queries

[De12] proved a lower bound on the additive error of any differentially private mechanism that answers $m = O(n)$ random inner product queries in $\{-1, 1\}^n$ on a secret dataset in $\{0, 1\}^n$. Specifically, the lower bound says that if the error is $O(\sqrt{n})$ on a $(\frac{1}{2} + \eta)$ -fraction of the responses with probability $\Omega(\sqrt{\delta})$, then the mechanism is not (ϵ, δ) -differentially private.

Theorem 9 (Theorem 4.1 of [De12]). *For any $n \in \mathbb{N}$, $\epsilon > 0$, and $1/20 > \delta > 0$, there exists positive constants α, γ , and $\eta < 1/2$ such that any mechanism \mathcal{M} that answers $m = \alpha n$ random inner product queries $Q^{(1)}, \dots, Q^{(m)} \in \{-1, 1\}^n$ on a secret dataset $X \in \{0, 1\}^n$ satisfying*

$$\Pr_{\mathcal{M}, Q^{(1)}, \dots, Q^{(m)}} \left[\Pr_{i \in [m]} [|\mathcal{M}(X)_i - \langle X, Q^{(i)} \rangle| \leq \gamma \sqrt{n}] \geq \frac{1}{2} + \eta \right] \geq 3\sqrt{\delta}$$

where $\mathcal{M}(X)_i$ denotes the response of \mathcal{M} on query $Q^{(i)}$, is not (ϵ, δ) -differentially private.

To obtain the lower bound on inner product queries in $\{0, 1\}^n$ that we desire, we show that any ϵ -differentially private mechanism $\tilde{\mathcal{M}}$ that answers m random inner product queries in $\{0, 1\}^n$ can be converted into a 2ϵ -differentially private mechanism \mathcal{M} that answers m random inner product queries in $\{-1, 1\}^n$. The idea is to observe that, for any $Q \in \{-1, 1\}^n$, we may write $\langle Q, X \rangle = 2\langle \tilde{Q}, X \rangle - \langle \mathbf{1}, X \rangle$, where $\tilde{Q} \in \{0, 1\}^n$ is such that $\tilde{Q}_i = \frac{1}{2}(Q_i + 1)$ and $\mathbf{1} \in \{1\}^n$ is the all-ones vector. Using this, we can simulate $\tilde{\mathcal{M}}$ on X , scale each response by a factor of 2, and then subtract a noisy version of $\langle \mathbf{1}, X \rangle$ from each response. More precisely, to answer m random inner product queries in $\{-1, 1\}^n$ on $X \in \{0, 1\}^n$, \mathcal{M} does the following:

1. Generate $y \sim \text{Lap}(1/\epsilon)$ and release $\tilde{x} = \langle \mathbf{1}, X \rangle + y$.
2. Run $\tilde{\mathcal{M}}$ on X to obtain responses (r_1, r_2, \dots, r_m) .
3. Return $(2r_1 - \tilde{x}, 2r_2 - \tilde{x}, \dots, 2r_m - \tilde{x})$.

The first step is ϵ -differentially private by Lemma 3. The second step is ϵ -differentially private by assumption. Therefore, by composition, the first two steps are 2ϵ -differentially private. The last step is simply post-processing.

Notice that since \mathcal{M} answers m random inner product queries $\tilde{Q}^{(1)}, \dots, \tilde{Q}^{(m)} \in \{0, 1\}^n$ on X , the vectors $Q^{(1)}, \dots, Q^{(m)} \in \{-1, 1\}^n$, where $Q_i^{(j)} = 2\tilde{Q}_i^{(j)} - 1$ are random. Furthermore, the error of the j th answer is $|(2r_j - \tilde{x}) - \langle Q^{(j)}, X \rangle| \leq |2r_j - 2\langle \tilde{Q}^{(j)}, X \rangle| + |2\langle \tilde{Q}^{(j)}, X \rangle - \langle \mathbf{1}, X \rangle - \langle Q^{(j)}, X \rangle| + |y| = 2|r_j - \langle \tilde{Q}^{(j)}, X \rangle| + |y|$. By Laplace tail bounds (Lemma 2), $|y| \leq O(\epsilon^{-1} \log n)$ with high probability. Therefore, if the error of j th answer of $\tilde{\mathcal{M}}$ is $o(\sqrt{n})$, then so is the error of the j th answer of \mathcal{M} , assuming ϵ is a constant. In conclusion, we obtain the following theorem.

Theorem 8. *For any constants $\epsilon > 0$ and $\frac{1}{20} > \delta > 0$, there is a constant $0 < \eta < \frac{1}{2}$ such that no (ϵ, δ) -differentially private mechanism can answer $m = O(n)$ random inner product queries in $\{0, 1\}^n$ on a secret dataset $X \in \{0, 1\}^n$ such that, with probability at least $\Omega(\sqrt{\delta})$, a $(\frac{1}{2} + \eta)$ -fraction of its answers have $o(\sqrt{n})$ error.*

E Missing proofs from Section 4

Lemma 4. *The coreness of x in $G_X(Q)$ is either $\langle Q, X \rangle$ or $\langle Q, X \rangle + 1$.*

Proof. Let $A' = \{a_i \in A \mid Q_i = X_i = 1\}$. Notice that $|A'| = \langle Q, X \rangle$. If $|A'| = 0$, then every neighbor (if any) of x has degree 1. Hence, the coreness of x is at most $1 = \langle Q, X \rangle + 1$.

Now suppose $|A'| \neq 0$. We will show that the coreness of x equals $|A'| = \langle Q, X \rangle$. By construction, every vertex in A' is adjacent to x and every vertex in B . It follows that $A' \cup B \cup \{x\}$ induces a subgraph

of $G_X(Q)$ in which every vertex has (induced) degree at least $|A'| \geq 1$. Hence, the coreness of x is at least $|A'|$.

Next consider any set of vertices S containing x . If S contains a vertex $a_i \in A \setminus A'$ that is adjacent to x , then $Q_i = 0$ and the degree of a_i in the subgraph of $G_X(Q)$ induced by S is $1 \leq |A'|$. Otherwise, every neighbor of x in the subgraph of $G_X(Q)$ induced by S is in A' and, hence, the induced degree of x is at most $|A'|$. Hence, in either case the coreness of x is at most $|A'|$. \square

Lemma 5. \mathcal{N} is 2ϵ -differentially private.

Proof. Let $X, X' \in \{0, 1\}^n$ be two datasets that differ only at the i th coordinate. Consider any valid combination of messages, $\{\pi_x\} \cup \{\pi_{a_j}(0) \mid j \in [n]\} \cup \{\pi_{a_j}(1) \mid j \in [n]\}$, generated in the first step. Since X and X' differ only at the i th coordinate, the neighborhoods of x in $G_X(\mathbf{0}), G_{X'}(\mathbf{0})$ differ in only one edge, namely $\{x, a_i\}$. Similarly, for the neighborhoods of a_i in the pairs of graphs $G_X(\mathbf{0}), G_{X'}(\mathbf{0})$ and $G_X(\mathbf{1}), G_{X'}(\mathbf{1})$. The neighborhoods of the other vertices in A are the same in each pair of graphs $G_X(\mathbf{0}), G_{X'}(\mathbf{0})$ and $G_X(\mathbf{1}), G_{X'}(\mathbf{1})$. It follows that the ratio of the probabilities of seeing the (partial) transcript $\{\pi_x\} \cup \{\pi_{a_j}(0) \mid j \in [n]\}$ on the graphs $G_X(\mathbf{0}), G_{X'}(\mathbf{0})$ is at most e^ϵ , since the transcript is ϵ -edge differentially private. Similarly, for the ratio of the probabilities of seeing the (partial) transcript $\{\pi_{a_j}(1) \mid j \in [n]\}$ on the graphs $G_X(\mathbf{1}), G_{X'}(\mathbf{1})$. Thus, the ratio of the probabilities of seeing the entire combination of messages in the first step on datasets X and X' is at most $e^\epsilon \cdot e^\epsilon = e^{2\epsilon}$ and the first step is 2ϵ -differentially private. The remaining steps do not require knowledge of X and, hence, only serve as post-processing. Therefore, \mathcal{N} is 2ϵ -differentially private. \square

Lemma 6. Suppose that \mathcal{M} has error exceeding α with probability at most β . Then, with probability at least $1 - 1/\gamma$, at least a $(1 - \gamma\beta)$ -fraction of the responses of \mathcal{N} have error at most $\alpha + 1$.

Proof. For $j \in \{1, \dots, m\}$, let $X_j \in \{0, 1\}$ be the indicator random variable for whether the response to the j th query has error exceeding $\alpha + 1$, i.e., $|\tilde{k}_j(x) - \langle Q^{(j)}, X \rangle| > \alpha + 1$. By Lemma 4, if $k_j(x)$ is the actual coreness of x in $G_X(Q^{(j)})$, then $|k_j(x) - \langle Q^{(j)}, X \rangle| \leq 1$. Hence, by triangle inequality, $|\tilde{k}_j(x) - \langle Q^{(j)}, X \rangle| \leq |\tilde{k}_j(x) - k_j(x)| + |k_j(x) - \langle Q^{(j)}, X \rangle| \leq |k_j(x) - k_j(x)| + 1$. Since \mathcal{M} has error exceeding α with probability at most β , it follows that $E[X_j] \leq \beta$. Hence, by linearity of expectation, $E[\sum_{j=1}^m X_j] = \sum_{j=1}^m E[X_j] \leq \beta m$ and, by Markov's inequality, $\Pr(\sum_{i=1}^m X_i \geq \gamma\beta m) \leq 1/\gamma$. Therefore, with probability $1 - 1/\gamma$, at least $(1 - \gamma\beta)m$ responses of \mathcal{N} have error at most $\alpha + 1$. \square