Shortest Paths and Distances with Differential Privacy

Adam Sealon
MIT
asealon@csail.mit.edu

ABSTRACT

We introduce a model for differentially private analysis of weighted graphs in which the graph topology $(\mathcal{V}, \mathcal{E})$ is assumed to be public and the private information consists only of the edge weights $w : \mathcal{E} \rightarrow \mathbb{R}^+$. This can express hiding congestion patterns in a known system of roads. Differential privacy requires that the output of an algorithm provides little advantage, measured by privacy parameters ϵ and δ , for distinguishing between neighboring inputs, which are thought of as inputs that differ on the contribution of one individual. In our model, two weight functions w, w' are considered to be neighboring if they have ℓ_1 distance at most one.

We study the problems of privately releasing a short path between a pair of vertices and of privately releasing approximate distances between all pairs of vertices. We are concerned with the *approximation error*, the difference between the length of the released path or released distance and the length of the shortest path or actual distance.

For the problem of privately releasing a short path between a pair of vertices, we prove a lower bound of $\Omega(|\mathcal{V}|)$ on the additive approximation error for fixed privacy parameters ϵ, δ . We provide a differentially private algorithm that matches this error bound up to a logarithmic factor and releases paths between all pairs of vertices, not just a single pair. The approximation error achieved by our algorithm can be bounded by the number of edges on the shortest path, so we achieve better accuracy than the worst-case bound for pairs of vertices that are connected by a low-weight path consisting of $o(|\mathcal{V}|)$ vertices.

For the problem of privately releasing all-pairs distances, we show that for trees we can release all-pairs distances with approximation error $O(\log^{2.5} |\mathcal{V}|)$ for fixed privacy parameters. For arbitrary bounded-weight graphs with edge weights in $[0, M]$ we can release all distances with approximation error $\tilde{O}(\sqrt{|\mathcal{V}|M})$.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PODS '16, June 26–July 1, 2016, San Francisco, CA, USA.

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4191-2/16/06...\$15.00

DOI: <http://dx.doi.org/10.1145/2902251.2902291>

1. INTRODUCTION

1.1 Differential privacy and our model

Privacy-preserving data analysis is concerned with releasing useful aggregate information from a database while protecting sensitive information about individuals in the database. *Differential privacy* [DMNS06] provides strong privacy guarantees while allowing many types of queries to be answered approximately. Differential privacy requires that for any pair of *neighboring* databases x and y , the output distributions of the mechanism on database x and database y are very close. In the traditional setting, databases are collections of data records and are considered to be neighboring if they are identical except for a single record, which is thought of as the information associated with a single individual. This definition guarantees that an adversary can learn very little about any individual in the database, no matter what auxiliary information the adversary may possess.

A new model We introduce a model of differential privacy for the setting in which databases are weighted graphs $(\mathcal{G}, w)$. In our setting, the graph topology $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is assumed to be public, and only the edge weights $w : \mathcal{E} \rightarrow \mathbb{R}^+$ must be kept private. Individuals are assumed to have only bounded influence on the edge weights. Consequently, two weight functions on the same unweighted graph are considered to be neighbors if they have ℓ_1 distance at most one.

The model is well suited to capture the setting of a navigation system which has access to current traffic data and uses it to direct drivers. The travel times along routes provided by the system should be as short as possible. However, the traffic data used by the system may consist of private information. For instance, navigation tools such as Google Maps and Waze estimate traffic conditions based on GPS locations of cars or based on traffic reports by users. On the other hand, the topology of the road map used by the system is non-private, since it is a static graph readily available to all users. We would like to provide information about paths and distances in the network without revealing sensitive information about the edge weights. In this paper we introduce a variant of differential privacy suitable for studying network routing, and provide both algorithms and lower bounds.

1.2 Our results

We will consider two classic problems in this model which are relevant for routing. First, we are interested in finding *short paths* between specified pairs of vertices. We cannot hope always to release the shortest path while preserving

privacy, but we would like to release a path that is not much longer than the shortest path. Here the *approximation error* is the difference in length between the released path and the shortest path. Second, we would like to release *distances* between pairs of vertices. The *approximation error* here is the absolute difference between the released distance and the actual distance. As we discuss below, releasing an accurate distance estimate for a single pair of vertices in our model is a straightforward application of the Laplace mechanism of [DMNS06]. We focus on the more difficult problem of releasing *all-pairs distances* privately with low error.

Interestingly, all of our error bounds are independent of the sum of the edge weights $\|w\|_1$, which corresponds most naturally to the size of the database under the usual setting for differential privacy. Instead, they depend only on the number of vertices $|\mathcal{V}|$ and edges $|\mathcal{E}|$ in the graph and the privacy parameters ϵ and δ . Our error bounds constitute additive error. Consequently, if the edge weights are large, the error will be small in comparison.

Approximate shortest paths We consider the problem of privately releasing an approximate *shortest path*. We provide a strong reconstruction based lower bound, showing that in general it is not possible under differential privacy to release a short path between a pair of vertices with additive error better than $\Omega(|\mathcal{V}|)$. Our lower bound is obtained by reducing the problem of reconstructing many of the rows of a database to the problem of finding a path with low error.

We also show that a simple ϵ -differentially private algorithm based on the Laplace mechanism of [DMNS06] comes close to meeting this bound. If the minimum-weight path between a pair of vertices s, t consists of at most k edges, then the weight of the path released by our algorithm is greater by at most $O(k \log |\mathcal{V}|)/\epsilon$. Since $k < |\mathcal{V}|$, the weight of the released path is $O(|\mathcal{V}| \log |\mathcal{V}|)/\epsilon$ greater than optimal. Note that when the edge weights are large, the length of a path can be much larger than $(|\mathcal{V}| \log |\mathcal{V}|)/\epsilon$, in which case our algorithm provides a good approximation. Moreover, for many networks we would expect that shortest paths should be compact and consist of few vertices, in which case the accuracy obtained will be much greater. The algorithm releases not only a single path but short paths between every pair of vertices without loss of additional privacy or accuracy.

Approximate all pairs distances For the problem of privately releasing all-pairs *distances*, standard techniques yield error $O(|\mathcal{V}| \log |\mathcal{V}|)/\epsilon$ for each query under ϵ -differential privacy. We obtain improved algorithms for two special classes of graphs, trees and arbitrary graphs with edges of bounded weight.

For trees we show that a simple recursive algorithm can release all-pairs distances with error $O(\log^{2.5} |\mathcal{V}|)/\epsilon$. Implicitly, the algorithm finds a collection $\mathcal{C}$ of paths with two properties. Every edge is contained in at most $\log |\mathcal{V}|$ paths, and the unique path between any pair of vertices can be formed from at most $4 \log |\mathcal{V}|$ paths in $\mathcal{C}$ using the operations of concatenation and subtraction. The first property implies that the query consisting of the lengths of all of the paths in $\mathcal{C}$ has global ℓ_1 sensitivity $\log |\mathcal{V}|$, so we can release estimates of the lengths of all paths in $\mathcal{C}$ with error roughly $\log |\mathcal{V}|$ using the standard Laplace mechanism [DMNS06]. The second property implies that we can use these estimates to compute the distance between any pair of vertices without too much increase in error. We first reduce

the approximate all-pairs distances problem to the problem of approximating all distances from a single fixed vertex. We then solve the single-source problem recursively, repeatedly decomposing the tree into subtrees of at most half the size.

In the special case of the path graph, releasing approximate all-pairs distances is equivalent to query release of threshold functions on the domain $X = \mathcal{E}$. The results of [DNPR10] yield the same error bound that we obtain for computing distances on the path graph. Consequently our algorithm for all-pairs distances on trees can be viewed as a generalization of a result for private query release of threshold functions.

For bounded-weight graphs we show that we can generate a set of vertices $\mathcal{Z} \subset \mathcal{V}$ such that distances between all pairs of vertices in $\mathcal{Z}$ suffice to estimate distances between all pairs of vertices in $\mathcal{V}$. The required property on $\mathcal{Z}$ is that every vertex $v \in \mathcal{V}$ is connected to a vertex $z_v \in \mathcal{Z}$ by a path consisting of few vertices. Since we can always find a set $\mathcal{Z}$ that is not too large, we can release distances between all pairs of vertices in $\mathcal{Z}$ with relatively small error. The distance between a pair of vertices $u, v \in \mathcal{V}$ can then be approximated by the distance between nearby vertices $z_u, z_v \in \mathcal{Z}$. In general, if edge weights are in $[0, M]$ for $\frac{1}{|\mathcal{V}|} < M\epsilon < |\mathcal{V}|$, then we can release all pairs shortest paths with an additive error of $\tilde{O}(\sqrt{|\mathcal{V}| M \epsilon^{-1} \log(1/\delta)})$ per path under (ϵ, δ) -differential privacy for any $\epsilon, \delta > 0$. Note that the distance between a pair of vertices can be as large as $|\mathcal{V}| \cdot M$, so this error bound is nontrivial. For specific graphs we can obtain better bounds by finding a smaller set $\mathcal{Z}$ satisfying the necessary requirements.

Additional problems We also consider two other problems in this model, the problem of releasing a nearly *minimal spanning tree* and the problem of releasing an almost *minimum weight perfect matching*. For these problems, the *approximation error* is the absolute difference in weight between the released spanning tree or matching and the minimum-weight spanning tree or matching.

Using a similar argument to the shortest path results of Section 5, we provide lower bounds and algorithms for these problems. Through a reduction from the problem of reconstructing many rows in a database, we show that it is not possible under differential privacy to release a spanning tree or matching with error better than $\Omega(|\mathcal{V}|)$. Using a simple algorithm based on the Laplace mechanism of [DMNS06], we can privately release a spanning tree or matching with error roughly $(|\mathcal{V}| \log |\mathcal{V}|)/\epsilon$. These results are presented in Appendix B.

Future directions In this work we describe differentially private algorithms and lower bounds for several natural graph problems related to network routing. We would like to explore how well our algorithms perform in practice on actual road networks and traffic data. We would also like to develop new algorithms for additional graph problems and to design improved all-pairs distance algorithms for additional classes of networks. In addition, lower bounds for accuracy of private all-pairs distances would be very interesting.

Scaling Our model requires that we preserve the indistinguishability of two edge weight functions which differ by at most 1 in ℓ_1 norm. The constant 1 here is arbitrary, and the error bounds in this paper scale according to it. For instance, if a single individual can only influence edge weights by $1/|\mathcal{V}|$

rather than 1 in ℓ_1 norm, then we can privately find a path between any pair of vertices whose length is only $O(\log |\mathcal{V}|)/\epsilon$ longer than optimal rather than $O(|\mathcal{V}| \log |\mathcal{V}|)/\epsilon$. The other results in this paper can be scaled similarly.

1.3 Previous work on graphs and differential privacy

Why a new model? The two main models which have been considered for applying differential privacy to network structured data are *edge* and *node* differential privacy [HLMJ09, BBDS13, KNRS13, KRSY11]. Under edge differential privacy, two graphs are considered to be neighbors if one graph can be obtained from the other by adding or removing a single edge. With node differential privacy, two graphs are neighbors if they differ only on the set of edges incident to a single vertex.

However, the notions of edge and node differential privacy are not well suited to shortest path and distance queries. Consider an algorithm that releases a short path between a specified pair of vertices. Any useful program solving this problem must usually at least output a path of edges which are in the graph. But doing so blatantly violates privacy under both the edge and node notions of differential privacy, since the released edges are private information. Indeed, an algorithm cannot release subgraphs of the input graph without violating both edge and node differential privacy because this distinguishes the input from a neighboring graph in which one of the released edges is removed.

What if we simply want to release the distance between a pair of vertices? In general it is not possible to release approximate distances with meaningful utility under edge or node differential privacy, since changing a single edge can greatly change the distances in the graph. Consider the unweighted path graph $\mathcal{P}$ and any pair of adjacent vertices x, y on the path. Removing edge (x, y) disconnects the graph, increasing the distance between x and y from 1 to ∞ . Even if the graph remains connected, the removal of an edge does not preserve approximate distances. Consider any pair of adjacent vertices x, y on the cycle graph $\mathcal{C}$. Here, removing edge (x, y) increases the distance between x and y from 1, the smallest possible distance, to $|\mathcal{V}| - 1$, the largest possible distance. Hence, the edge and node notions of differential privacy do not enable the release of approximate distances. This inadequacy motivates the new notion of differential privacy for graphs introduced in this work.

A histogram formulation A database consisting of elements from a data universe $\mathcal{U}$ can be described by a vector $D \in \mathbb{N}^{|\mathcal{U}|}$, where the i th component of the vector corresponds to the number of copies of the i th element of $\mathcal{U}$ in the database. More generally, we can allow the database to be a point in $\mathbb{R}^{|\mathcal{U}|}$, which corresponds to allowing a fractional number of occurrences of each element. Since an edge weight function w is an element of $\mathbb{R}^{|\mathcal{E}|}$ and the notions of sensitivity coincide, we can rephrase our model in the standard differential privacy framework.

Consequently, we can use existing tools for privately answering low sensitivity queries to yield incomparable results to those presented in this paper for the problem of releasing all-pairs distances with low error. If all edge weights are integers and the sum $\|w\|_1$ of the edge weights is known, we can release all-pairs distances with additive error

$$\tilde{O}(\sqrt{\|w\|_1 \cdot \log |\mathcal{V}|} \log^{1.5}(1/\delta)/\epsilon)$$

except with probability δ using the (ϵ, δ) -differentially private boosting mechanism of [DRV10]. The assumption that $\|w\|_1$ is known is not problematic, since we can privately release a good approximation. We can also extend the algorithm to noninteger edge weights for the queries in our setting, although we obtain a worse error bound. We do this by modifying the base synopsis generator of [DRV10] to produce a database in $(\tau\mathbb{N})^{|\mathcal{E}|}$ whose fractional counts are integer multiples of $\tau = \alpha/(2|\mathcal{V}|)$, where α will be the obtained additive approximation. The modified algorithm releases all-pairs distances with error $\tilde{O}(\sqrt{\|w\|_1 \cdot |\mathcal{V}|} \log^{4/3}(1/\delta)/\epsilon^{2/3})$.

This error bound has a better dependence on $|\mathcal{V}|$ than the algorithms for all-pairs distances in general and weighted graphs described in Section 4. However, it has a substantial dependence on $\|w\|_1$, while the error bound for all algorithms in the remainder of this paper are independent of $\|w\|_1$. In addition, the running time of the algorithm of [DRV10] is exponential in the database size, while all algorithms described below run in polynomial time.

Additional related work While the private edge weight model explored in this work is new, a few previous works have considered problems on graphs in related models. Nissim, Raskhodnikova and Smith [NRS07] consider the problem of computing the cost of the minimum spanning tree of a weighted graph. They introduce the notion of smooth sensitivity, which they use to compute the approximate MST cost. In their work, edge weights are bounded, and each weight corresponds to the data of a single individual. In contrast, we allow unbounded weights but assume a bound on the effect an individual can have on the weights.

Hsu et al. [HHR⁺14] consider the problem of privately finding high-weight matchings in bipartite graphs. In their model, which captures a private version of the allocation problem, two weightings of the complete bipartite graph are neighbors if they differ only in the weights assigned to the edges incident to a single left-vertex, which correspond to the preferences of a particular agent. They show that the problem is impossible to solve under the standard notion of differential privacy. They work instead under the relaxed notion of “joint differential privacy,” in which knowledge of the edges of the matching which correspond to some of the left-vertices cannot reveal the weights of edges incident to any of the remaining left-vertices.

A series of works has explored the problem of privately computing the size of all cuts $(S, \bar{S})$ in a graph. Gupta, Roth and Ullman [GRU12] show how to answer cut queries with $O(|\mathcal{V}|^{1.5})$ error. Blocki et al. [BBDS12] improve the error for small cuts. Relatedly, Gupta et al. [GLM⁺10] show that we can privately release a cut of close to minimal size with error $O(\log |\mathcal{V}|)/\epsilon$, and that this is optimal. Since the size of a cut is the number of edges crossing the cut, it can also be viewed as the sum of the weights of the edges crossing the cut in a $\{0, 1\}$ -weighting of the complete graph. Consequently the problem can be restated naturally in our model.

As we have discussed, the problem of approximating all threshold functions on a totally ordered set is equivalent to releasing approximate distances on the path graph. In addition to the work of Dwork et al. [DNPR10] described above, additional bounds and reductions for query release and learning of threshold functions are shown in Bun et al. [BNSV15].

2. THE PRIVACY MODEL

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ denote an undirected graph with vertex set $\mathcal{V}$ and edge set $\mathcal{E}$, and let $w : \mathcal{E} \rightarrow \mathbb{R}^+$ be a *weight function*. (The shortest path results in Section 5 also apply to directed graphs.) Let $V = |\mathcal{V}|$ and $E = |\mathcal{E}|$ be the number of vertices and edges, respectively.

Let $\mathcal{P}_{xy}$ denote the set of paths between a pair of vertices $x, y \in \mathcal{V}$. For any path $P \in \mathcal{P}_{xy}$, the weight $w(P)$ is the sum $\sum_{e \in P} w(e)$ of the weights of the edges of P . The distance $d_w(x, y)$ from x to y denotes the weighted distance $\min_{P \in \mathcal{P}_{xy}} w(P)$. We will denote the unweighted or hop distance from x to y by $h(x, y) = \min_{P \in \mathcal{P}_{xy}} \ell(P)$, where the hop length $\ell(P)$ of path $P = (v_0, \dots, v_\ell)$ is the number ℓ of edges on the path. Let the shortest path $\text{SP}_w(x, y)$ denote a path from x to y of minimum possible weight $w(\text{SP}_w(x, y)) = d_w(x, y)$.

We now formally define differential privacy in the private edge weight model.

DEFINITION 2.1. For any edge set $\mathcal{E}$, two weight functions $w, w' : \mathcal{E} \rightarrow \mathbb{R}^+$ are **neighboring**, denoted $w \sim w'$, if

$$\|w - w'\|_1 = \sum_{e \in \mathcal{E}} |w(e) - w'(e)| \leq 1.$$

DEFINITION 2.2. For any graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, let $\mathcal{A}$ be an algorithm that takes as input a weight function $w : \mathcal{E} \rightarrow \mathbb{R}^+$. $\mathcal{A}$ is (ϵ, δ) -**differentially private** on $\mathcal{G}$ if for all pairs of neighboring weight functions w, w' and for all sets of possible outputs S , we have that

$$\Pr[\mathcal{A}(w) \in S] \leq e^\epsilon \Pr[\mathcal{A}(w') \in S] + \delta.$$

If $\delta = 0$ we say that $\mathcal{A}$ is ϵ -differentially private on $\mathcal{G}$.

For a class $\mathcal{C}$ of graphs, we say that $\mathcal{A}$ is (ϵ, δ) -differentially private on $\mathcal{C}$ if $\mathcal{A}$ is (ϵ, δ) -differentially private on $\mathcal{G}$ for all $\mathcal{G} \in \mathcal{C}$.

We now define our accuracy criteria for the approximate shortest paths and distances problems.

DEFINITION 2.3. For the shortest path problem, the **error** of a path $P \in \mathcal{P}_{xy}$ between vertices x, y is the difference $w(P) - d_w(x, y)$ between the length of P and the length of the shortest path between x and y .

DEFINITION 2.4. For the approximate distances problem, the **error** is the absolute difference between the released distance between a pair of vertices x, y and the actual distance $d_w(x, y)$.

3. PRELIMINARIES

We will now introduce a few basic tools which will be used throughout the remainder of this paper. A number of differential privacy techniques incorporate noise sampled according to the Laplace distribution. We define the distribution and state a concentration bound for sums of Laplace random variables.

DEFINITION 3.1. The **Laplace distribution** with scale b , denoted $\text{Lap}(b)$, is the distribution with probability density function given by

$$p(x) = \frac{1}{2b} \exp(-|x|/b).$$

If Y is distributed according to $\text{Lap}(b)$, then for any $t > 0$ we have that $\Pr[|Y| > t \cdot b] = e^{-t}$.

LEMMA 3.1. (CONCENTRATION OF LAPLACE RANDOM VARIABLES [CSS10]) Let $X_1, \dots, X_t$ be independent random variables distributed according to $\text{Lap}(b)$, and let $X = \sum_i X_i$. Then for all $\gamma \in (0, 1)$ we have that with probability at least $1 - \gamma$,

$$|X| < 4b\sqrt{t} \ln(2/\gamma) = O(b\sqrt{t} \log(1/\gamma)).$$

One of the first and most versatile differentially private algorithms is the Laplace mechanism, which releases a noisy answer with error sampled from the Laplace distribution with scale proportional to the sensitivity of the function being computed.

DEFINITION 3.2. For any function $f : \mathcal{X} \rightarrow \mathbb{R}^k$, the **sensitivity**

$$\Delta f = \max_{\substack{w, w' \in \mathcal{X} \\ w \sim w'}} \|f(w) - f(w')\|_1$$

is the largest amount f can differ in ℓ_1 norm between neighboring inputs.

In our setting we have $\mathcal{X} = (\mathbb{R}^+)^{\mathcal{E}}$.

LEMMA 3.2. (LAPLACE MECHANISM [DMNS06]) Given any function $f : \mathcal{X} \rightarrow \mathbb{R}^k$, the Laplace mechanism on input $w \in \mathcal{X}$ independently samples $Y_1, \dots, Y_k$ according to $\text{Lap}(\Delta f/\epsilon)$ and outputs

$$\mathcal{M}_{f, \epsilon}(w) = f(w) + (Y_1, \dots, Y_k).$$

The Laplace mechanism is ϵ -differentially private.

Finally, we will need the following results on the composition of differentially private algorithms.

LEMMA 3.3. (BASIC COMPOSITION, E.G. [DKM⁺06]) For any $\epsilon, \delta \geq 0$, the adaptive composition of k (ϵ, δ) -differentially private mechanisms is $(k\epsilon, k\delta)$ -differentially private.

LEMMA 3.4. (COMPOSITION [DRV10, DR13]) For any $\epsilon, \delta, \delta' \geq 0$, the adaptive composition of k (ϵ, δ) -differentially private mechanisms is $(\epsilon', k\delta + \delta')$ -differentially private for

$$\epsilon' = \sqrt{2k \ln(1/\delta')} \cdot \epsilon + k\epsilon(e^\epsilon - 1)$$

which is $O(\sqrt{k \ln(1/\delta')} \cdot \epsilon)$ provided $k \leq 1/\epsilon^2$. In particular, if $\epsilon' \in (0, 1)$, $\delta, \delta' \geq 0$, the composition of k (ϵ, δ) -differentially private mechanisms is $(\epsilon', k\delta + \delta')$ -differentially private for $\epsilon = O(\epsilon' / \sqrt{k \ln(1/\delta')})$.

4. COMPUTING DISTANCES

In this section we consider the problem of releasing distance oracle queries in the private edge weights model. Since neighboring weight functions differ by at most 1 in ℓ_1 norm, the weight of any path also changes by at most 1. Consequently, a single distance oracle query is sensitivity-1, and so we can use the Laplace mechanism (Lemma 3.2) to answer it privately after adding noise proportional to $1/\epsilon$. However, what if we would like to learn all-pairs distances privately?

There are V^2 pairs (s, t) of vertices, so we can achieve ϵ -differential privacy by adding to each query Laplace noise proportional to V^2/ϵ . We can do better using approximate differential privacy ($\delta > 0$) and Lemma 3.4 (Composition). Adding Laplace noise proportional to $1/\epsilon'$ to

Figure 1: The partition used in Algorithm 1 to separate a tree into subtrees of size at most $V/2$.

each query results in a mechanism that is $(V\epsilon'\sqrt{2\ln(1/\delta)} + V^2\epsilon'(e^{\epsilon'} - 1), \delta)$ -differentially private for any $\delta > 0$. Taking $\epsilon' = \epsilon/O(V\sqrt{\ln 1/\delta})$ for $\epsilon < 1$, we obtain a mechanism that is (ϵ, δ) -differentially private by adding $O(V\sqrt{\ln 1/\delta})/\epsilon$ noise to each query.

The other natural approach is to release an ϵ -differentially private version of the graph by adding $\text{Lap}(1/\epsilon)$ noise to each edge. This will be the basis for our approach in Algorithm 3 for computing approximate shortest paths. With probability $1 - \gamma$, all E Laplace random variables will have magnitude $(1/\epsilon)\log(E/\gamma)$, so the length of every path in the released synthetic graph is within $(V/\epsilon)\log(E/\gamma)$ of the length of the corresponding path in the original graph. Therefore with probability $1 - \gamma$ we have that all pairs distances in the released synthetic graph will be within $(V/\epsilon)\log(E/\gamma)$ of the corresponding distances in the original graph.

Both of these approaches result in privately releasing all pairs distances with additive error roughly V/ϵ . It is natural to ask whether this linear dependence on V is the best possible. In the remainder of this section we present algorithms which substantially improve on this bound for two special classes of graphs, trees and arbitrary graphs with edges of bounded weight. For trees we can obtain all-pairs distances with error $\text{polylog}(V)/\epsilon$, and for graphs with edges in $[0, M]$ we can obtain all-pairs distances with error roughly $\sqrt{VM}/\epsilon$.

4.1 Distances in trees

For trees it turns out to be possible to release all-pairs distances with far less error. We will first show a single-source version of this result for rooted trees, which we will then use to obtain the full result. The idea is to split the tree into subtrees of at most half the size of the original tree. As long as we can release the distance from the root to each subtree with small error, we can then recurse on the subtrees.

The problem of privately releasing all-pairs distances for the path graph can be restated as the problem of privately releasing threshold functions for a totally ordered data universe of size $E = V - 1$. Dwork et al. [DNPR10] showed that we can privately maintain a running counter for T timesteps, with probability $1 - \gamma$ achieving additive accu-

racy of $O(\log(1/\gamma)\log^{2.5} T/\epsilon)$ at each timestep. This algorithm essentially computes all threshold queries for a totally ordered universe of size T , which is equivalent to computing the distance from an endpoint to each other vertex of the path graph. We match the accuracy of the [DNPR10] algorithm and generalize the result to arbitrary trees.

In Appendix A we provide an alternate algorithm for the path graph which achieves the same asymptotic bounds. This result can be viewed as a restatement of the [DNPR10] algorithm.

THEOREM 4.1. (SINGLE-SOURCE DISTANCES ON ROOTED TREES) *Let $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ be a tree with root vertex v_0 , and let $\epsilon > 0$. Then there is an algorithm that is ϵ -differentially private on $\mathcal{T}$ that on input edge weights $w : \mathcal{E} \rightarrow \mathbb{R}^+$ releases approximate distances from v_0 to each other vertex, where with probability $1 - \gamma$ the approximation error on each released distance is*

$$O(\log^{1.5} V \cdot \log(1/\gamma))/\epsilon$$

for any $\gamma \in (0, 1)$.

PROOF. Given the tree $\mathcal{T}$ and root v_0 , we will partition $\mathcal{V}$ into subtrees of size at most $V/2$ as shown in Figure 1 and recursively obtain distances in each subtree. There exists some vertex v^* such that the subtree rooted at v^* contains more than $V/2$ vertices while the subtree rooted at each child of v^* has at most $V/2$ vertices. The topology of the graph is public, so we can release vertex v^* . Let $v_1, \dots, v_t$ be the children of v^* , and let $\mathcal{T}_i = (\mathcal{V}_i, \mathcal{E}_i)$ be the subtree rooted at v_i for each $i \in [t]$. Let $\mathcal{T}_0 = (\mathcal{V}_0, \mathcal{E}_0)$ be the subtree rooted at v_0 consisting of the remaining vertices $\mathcal{V} \setminus (\mathcal{V}_1 \cup \dots \cup \mathcal{V}_t)$.

We can compute and release distances between the following pairs of vertices, adding $\text{Lap}(\log V/\epsilon)$ noise to each distance:

- (v_0, v^*)
- (v^*, v_i) for each $i \in [t]$

Since v^* is the parent of $v_1, \dots, v_t$ in the tree rooted at v_0 , the path from v_0 to v^* in T contains none of the edges (v^*, v_i) for $i \in [t]$, so the function which releases these $t + 1$ distances is sensitivity-1.

Algorithm 1: Rooted tree distances

Inputs: Tree $\mathcal{T} = (\mathcal{V}, \mathcal{E})$, root $v_0 \in \mathcal{V}$, edge weights $w : \mathcal{E} \rightarrow \mathbb{R}^+$, and number of vertices n of original tree.

1. Let v^* be the unique vertex such that the subtree rooted at v^* has more than $V/2$ vertices while the subtree rooted at each child of v^* has at most $V/2$ vertices.
2. Let $v_1, \dots, v_t$ be the children of v^* .
3. Let $\mathcal{T}_i$ be the subtree rooted at v_i for $i \in [t]$, and let $\mathcal{T}_0 = \mathcal{T} \setminus \{\mathcal{T}_1, \dots, \mathcal{T}_t\}$.
4. Sample $X_{v^*, \mathcal{T}} \leftarrow \text{Lap}(\log V/\epsilon)$ and let $d(v^*, \mathcal{T}) = d_w(v_0, v^*) + X_{v^*, \mathcal{T}}$.
5. Let $d(v_0, \mathcal{T}) = 0$.
6. For each $i \in [t]$, sample $X_{v_i, \mathcal{T}} \leftarrow \text{Lap}(\log V/\epsilon)$ and let $d(v_i, \mathcal{T}) = d(v^*, \mathcal{T}) + w((v^*, v_i)) + X_{v_i, \mathcal{T}}$.
7. Recursively compute distances in each subtree $\mathcal{T}_0, \dots, \mathcal{T}_t$.
8. For each vertex $v \in \mathcal{V}$, if $v \in \mathcal{T}_i$, then let $d(v, \mathcal{T}) = d(v_i, \mathcal{T}) + d(v, \mathcal{T}_i)$.
9. Release $d(v, \mathcal{T})$ for all $v \in \mathcal{V}$.

We then recursively repeat the procedure on each subtree $\mathcal{T}_0, \dots, \mathcal{T}_t$ until we reach trees containing only a single vertex, adding $\text{Lap}(\log V/\epsilon)$ noise to each released value. Each subtree has size at most $V/2$, so the depth of recursion is bounded by $\log V$. The subtrees $\mathcal{T}_0, \dots, \mathcal{T}_t$ are also disjoint. Consequently the function which releases all of the distances at depth d in the recursion has sensitivity 1 for any d . Therefore the function which releases all distances queried in this recursive procedure has sensitivity at most $\log V$. Since we add $\text{Lap}(\log V/\epsilon)$ noise to each coordinate of this function, the algorithm outlined above is an instantiation of the Laplace mechanism (Lemma 3.2) and is ϵ -differentially private.

We now show how these queries suffice to compute the distance from root vertex v_0 to each other vertex with small error. The algorithm samples at most $2V$ Laplace random variables distributed according to $\text{Lap}(\log V/\epsilon)$, so by a union bound, with probability $1 - \gamma$ all of these have magnitude $O(\log V \log(V/\gamma))/\epsilon$. Consequently to obtain an error bound of roughly $O(\log^3 V)/\epsilon$ it suffices to show that any distance in $\mathcal{T}$ can be represented as a sum of $O(\log V)$ of the noisy distances released in the algorithm. We will use Lemma 3.1 to obtain a slightly better bound on the error terms.

Let set Q consist of the pairs of vertices corresponding to distance queries made by the algorithm above. We prove the following statement by induction. For any vertex $u \in \mathcal{V}$, there is a path from v_0 to u in the graph $(\mathcal{V}, Q)$ consisting of at most $2 \log V$ edges. The base case $V = 1$ is vacuous. For $V > 1$, note that since subtrees $\mathcal{T}_0, \dots, \mathcal{T}_t$ partition the vertex set, u must lie in one of them. If $u \in \mathcal{T}_i$, then by the inductive assumption on $\mathcal{T}_i$, there is a path from v_i to u in $(\mathcal{V}, Q)$ consisting of at most $2 \log(V/2) = 2 \log V - 2$ edges. If $i = 0$ this already suffices. Otherwise, if $i > 0$, note that $(v_0, v^*) \in Q$ and $(v^*, v_i) \in Q$. Consequently there is a path in $(\mathcal{V}, Q)$ from v_0 to u consisting of at most $2 \log V$ edges.

This means that there is a set of at most $2 \log V$ distances queried such that the distance from v_0 to u consists of the

sum of these distances. Consequently by adding these distances we obtain an estimate for the distance from v_0 to u whose error is the sum of at most $2 \log V$ independent random variables each distributed according to $\text{Lap}(\log V/\epsilon)$. By Lemma 3.1, we have that with probability at least $1 - \gamma$, we compute an estimate of $d_w(v_0, u)$ with error $O(\log^{1.5} V \cdot \log(1/\gamma))/\epsilon$ for any $\gamma \in (0, 1)$. Since differentially private mechanisms are preserved under post-processing, this algorithm satisfies ϵ -differential privacy and computes distances from v^* to each other vertex in T , where with probability at least $1 - \gamma$ each distance has error at most $O(\log^{1.5} V \cdot \log(1/\gamma))/\epsilon$. $\square$

We now extend this result to obtain all-pairs distances for trees.

THEOREM 4.2. (ALL-PAIRS DISTANCES ON TREES) *For any tree $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ and $\epsilon > 0$ there is an algorithm that is ϵ -differentially private on $\mathcal{T}$ and on input edge weights $w : \mathcal{E} \rightarrow \mathbb{R}^+$ releases all-pairs distances such that with probability $1 - \gamma$, all released distances have approximation error*

$$O(\log^{2.5} V \cdot \log(1/\gamma))/\epsilon$$

for any $\gamma \in (0, 1)$. For each released distance, with probability $1 - \gamma$ the approximation error is $O(\log^{1.5} V \cdot \log(1/\gamma))/\epsilon$.

PROOF. Arbitrarily choose some root vertex v_0 . Use the ϵ -differentially private algorithm of Theorem 4.1 to obtain approximate distances from v_0 to each other vertex of $\mathcal{T}$. We now show that this suffices to obtain all-pairs distances.

Consider any pair of vertices x, y , and let z be their lowest common ancestor in the tree rooted at v_0 . Then

$$\begin{aligned} d_w(x, y) &= d_w(z, x) + d_w(z, y) \\ &= d_w(v_0, x) + d_w(v_0, y) - 2d_w(v_0, z). \end{aligned}$$

Since with probability $1 - 3\gamma$ we can compute each of these three distances with error $O(\log^{1.5} V \cdot \log(1/\gamma)/\epsilon)$, it follows that we can compute the distance between x and y

Algorithm 2: Bounded-weight distances

Inputs: $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, $w : \mathcal{E} \rightarrow \mathbb{R}^+$, $k, M, \gamma, \epsilon' > 0$, k -covering $\mathcal{Z}$.

1. For all $y, z \in \mathcal{Z}$, sample $X_{y,z} \leftarrow \text{Lap}(|\mathcal{Z}|/\epsilon')$ and output $a_{y,z} := d_w(y, z) + X_{y,z}$.
2. For $v \in V$, let $z(v) \in \mathcal{Z}$ denote a vertex in $\mathcal{Z}$ with $h(v, z(v)) \leq k$.
3. The approximate distance between vertices $u, v \in V$ is given by $a_{z(u), z(v)}$.

with error at most four times this, which is still $O(\log^{1.5} V \cdot \log(1/\gamma))/\epsilon$. By a union bound, for any $\gamma \in (0, 1)$, with probability at least $1 - \gamma$ each error among the $V(V-1)/2$ all-pairs distances released is at most $O(\log^{1.5} V \cdot \log(V/\gamma))/\epsilon = O(\log^{2.5} V \cdot \log(1/\gamma))/\epsilon$. $\square$

4.2 Distances in bounded-weight graphs

THEOREM 4.3. *For all $\mathcal{G}, \delta, \gamma, M$, and $\epsilon \in (0, 1)$, if $1/V < M\epsilon < V$ then there is an algorithm $\mathcal{A}$ that is (ϵ, δ) -differentially private on $\mathcal{G}$ such that for all $w : \mathcal{E} \rightarrow [0, M]$, with probability $1 - \gamma$, $\mathcal{A}(w)$ outputs approximate all-pairs distances with additive error*

$$O\left(\sqrt{VM\epsilon^{-1} \cdot \log(1/\delta)} \cdot \log(VM\epsilon/\gamma)\right)$$

per distance. For any $\epsilon > 0$, $\mathcal{G}, \gamma, M$, if $1/V < M\epsilon < V^2$ then there is an algorithm $\mathcal{B}$ that is ϵ -differentially private on $\mathcal{G}$ such that for all $w : \mathcal{E} \rightarrow [0, M]$, with probability $1 - \gamma$, $\mathcal{B}(w)$ outputs approximate all-pairs distances with additive error

$$O\left((VM)^{2/3} \epsilon^{-1/3} \log(VM\epsilon/\gamma)\right)$$

per distance.

To achieve this result, we will find a small subset $\mathcal{Z}$ of $\mathcal{V}$ such that every vertex $v \in \mathcal{V}$ is near some vertex $z \in \mathcal{Z}$. We will need the following definition, introduced in [MM75].

DEFINITION 4.1. *A subset $\mathcal{Z} \subset \mathcal{V}$ of vertices is a k -covering if for every $v \in \mathcal{V}$ there is some $z \in \mathcal{Z}$ such that $h(v, z) \leq k$, where h is the hop-distance.*

A k -covering is sometimes called a k -dominating set (e.g. in [CN82]). The following lemma shows that we can find a sufficiently small k -covering for any graph $\mathcal{G}$.

LEMMA 4.4. [MM75] *If $V \geq k + 1$, then $\mathcal{G}$ has a k -covering of size at most $\lceil V/(k + 1) \rceil$.*

PROOF. Consider any spanning tree $\mathcal{T}$ of $\mathcal{G}$ and any vertex $x \in \mathcal{V}$ that is an endpoint of one of the longest paths in $\mathcal{T}$. For $0 \leq i \leq k$, let $\mathcal{Z}_i$ be the subset of $\mathcal{V}$ consisting of vertices whose distance from x in $\mathcal{T}$ is congruent to i modulo $k + 1$. It can be shown that each $\mathcal{Z}_i$ is a k -covering of $\mathcal{T}$ and therefore of $\mathcal{G}$. But the $k + 1$ sets $\mathcal{Z}_i$ form a partition of $\mathcal{V}$. Therefore some $\mathcal{Z}_i$ has $|\mathcal{Z}_i| \leq \lfloor V/(k + 1) \rfloor$, as desired. $\square$

THEOREM 4.5. *For any $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, $k, \delta > 0$, and $\gamma, \epsilon \in (0, 1)$, let $\mathcal{Z}$ be a k -covering of $\mathcal{G}$, and let $Z = |\mathcal{Z}|$. Then there is an algorithm $\mathcal{A}$ which is (ϵ, δ) -differentially private*

on $\mathcal{G}$ such that for any edge weight function $w : \mathcal{E} \rightarrow [0, M]$, with probability $1 - \gamma$, $\mathcal{A}(w)$ releases all-pairs distances with approximation error

$$O\left(kM + Z\epsilon^{-1} \log(Z/\gamma) \sqrt{\log(1/\delta)}\right)$$

per distance.

PROOF. There are Z^2 pairwise distances between vertices in $\mathcal{Z}$. We can compute and release noisy versions of each of these distances, adding $\text{Lap}(Z/\epsilon')$ noise to each. For any $\gamma \in (0, 1)$, with probability $1 - \gamma$ we have that each of these Z^2 noise variables has magnitude at most $(Z/\epsilon') \log(Z^2/\gamma)$. By Lemma 3.4, for any $\delta > 0$ releasing this is (ϵ, δ) -differentially private for $\epsilon' = O(\epsilon/\sqrt{\ln 1/\delta})$.

But this information allows the recipient to compute approximate distances between any pair of vertices $x, y \in \mathcal{V}$, as follows. Since $\mathcal{Z}$ is a k -covering of $\mathcal{G}$, we can find $z_x, z_y \in \mathcal{Z}$ which are at most k vertices from x and y . Since the maximum weight is M , the weight of the shortest path between x and z_x is at most kM , and similarly for y and z_y . Consequently

$$|d_w(x, y) - d_w(z_x, z_y)| \leq 2kM.$$

But we have released an estimate of $d_w(z_x, z_y)$ with noise distributed according to $\text{Lap}(Z/\epsilon)$. Consequently with probability $1 - \gamma$ each of these estimates differs from $d_w(z_x, z_y)$ by at most $(Z/\epsilon) \log(Z^2/\gamma)$. $\square$

We obtain a slightly weaker result under pure differential privacy.

THEOREM 4.6. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. For any $k > 0$ and $\gamma \in (0, 1)$, if $\mathcal{Z}$ is a k -covering of $\mathcal{G}$ of size Z , then there is an algorithm $\mathcal{A}$ that is ϵ -differentially private on $\mathcal{G}$ such that for any $w : \mathcal{E} \rightarrow [0, M]$, with probability $1 - \gamma$, $\mathcal{A}(w)$ releases all-pairs distances with approximation error*

$$O(kM + Z^2\epsilon^{-1} \log(Z/\gamma))$$

per distance.

PROOF. There are at most Z^2 pairwise distances between vertices in $\mathcal{Z}$, so we can release approximations of each distance, adding $\text{Lap}(Z^2/\epsilon)$ noise to each distance. With probability $1 - \gamma$ each of these Z^2 noise variables has magnitude at most $(Z^2/\epsilon) \log(Z^2/\gamma)$. By Lemma 3.3, releasing these distances is ϵ -differentially private. As above, since $\mathcal{Z}$ is a k -covering of $\mathcal{G}$, we can find $z_x, z_y \in \mathcal{Z}$ which are at most k vertices from x and y . Consequently $d_w(x, z_x) \leq kM$ and $d_w(y, z_y) \leq kM$, so

$$|d_w(x, y) - d_w(z_x, z_y)| \leq 2kM.$$

Figure 2: The graph used for the lower bound of Lemma 5.2.

But the released estimate for $d_w(z_x, z_y)$ has noise distributed according to $\text{Lap}(kM)$, so with probability $1 - \gamma$ each of these estimates differs from $d_w(x, y)$ by at most $O(Z^2 \epsilon^{-1} \log(Z^2/\gamma))$, as desired. $\square$

We now conclude the proof of Theorem 4.3.

PROOF OF THEOREM 4.3. Combine Lemma 4.4 with Theorem 4.5 for $k = \lfloor \sqrt{V/(M\epsilon)} \rfloor$, and combine Lemma 4.4 with Theorem 4.6 for $k = \lfloor V^{2/3}/(M\epsilon)^{1/3} \rfloor$. $\square$

Note that if we are only interested in the $V - 1$ distances from a single source, then directly releasing noisy distances and applying Lemma 3.4 yields (ϵ, δ) -differential privacy with error distributed according to $\text{Lap}(b)$ for $b = O(\sqrt{V \log 1/\delta})/\epsilon$, which has the same dependence on V as the bound provided by the theorem for releasing all pairs distances.

For some graphs we may be able to find a smaller k -covering than that guaranteed by Lemma 4.4. Then we can use Theorems 4.5 and 4.6 to obtain all-pairs distances with lower error. For instance, we have the following.

THEOREM 4.7. Let $\mathcal{G}$ be the $\sqrt{V} \times \sqrt{V}$ grid. Then for any $\epsilon, \gamma \in (0, 1)$ and $\delta > 0$, we can release with probability $1 - \gamma$ all-pairs distances with additive approximation error

$$V^{1/3} \cdot O\left(M + \epsilon^{-1} \log(V/\gamma) \sqrt{\log 1/\delta}\right)$$

while satisfying (ϵ, δ) -differential privacy.

PROOF. Let $\mathcal{V} = [\sqrt{V}] \times [\sqrt{V}]$, and let $\mathcal{Z} \subset \mathcal{V}$ consist of vertices $(i, j) \in \mathcal{V}$ with i, j both one less than a multiple of $V^{1/3}$. Then $\mathcal{Z}$ is a $2V^{1/3}$ -covering of $\mathcal{G}$, and also $|\mathcal{Z}| \leq V^{1/3}$. Consequently Theorem 4.5 implies the desired conclusion. $\square$

5. FINDING SHORTEST PATHS

5.1 Lower bound

In this section we present a lower bound on the additive error with which we can privately release a short path between a pair of vertices. The argument is based on a reduction from the problem of reconstructing a large fraction of the entries of a database. We show that an adversary can use an algorithm which outputs a short path in a graph to produce a vector with small Hamming distance to an input vector, which is impossible under differential privacy. To that end, we exhibit a “hard” graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and a family of weight functions, and provide a correspondence between inputs $x \in \{0, 1\}^n$ and weight functions $w : \mathcal{E} \rightarrow \{0, 1\}$.

THEOREM 5.1. There exists a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and vertices $s, t \in \mathcal{V}$ such that for any algorithm $\mathcal{A}$ that is (ϵ, δ) -differentially private on $\mathcal{G}$, there exist edge weights $w : \mathcal{E} \rightarrow$

$\{0, 1\}$ for which the expected approximation error of the path $\mathcal{A}(w)$ from s to t is at least $\alpha = (V - 1) \cdot \left(\frac{1 - (1 + e^\epsilon)\delta}{1 + e^{2\epsilon}}\right)$. In particular, for sufficiently small ϵ and δ , $\alpha \geq 0.49(V - 1)$.

Consequently, any differentially private algorithm for approximate shortest paths must on some inputs have additive error proportional to the number of vertices. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be the $(n + 1)$ -vertex graph with vertex set $\mathcal{V} = \{0, \dots, n\}$ and two parallel edges $e_i^{(0)}$ and $e_i^{(1)}$ between each pair of consecutive vertices $i - 1, i$, as shown in Figure 2. (For simplicity we have defined this is a multigraph, but it can be transformed into a simple graph with the addition of n extra vertices, changing the bound obtained by a factor of 2.)

LEMMA 5.2. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be the graph defined in the previous paragraph. For any α , let $\mathcal{A}$ be an algorithm that is (ϵ, δ) -differentially private on $\mathcal{G}$ that on input edge weights $w : \mathcal{E} \rightarrow \{0, 1\}$ produces a path from vertex $s = 0$ to vertex $t = n$ with expected approximation error at most α . Then there exists a $(2\epsilon, (1 + e^\epsilon)\delta)$ -differentially private algorithm $\mathcal{B}$ which on input $x \in \{0, 1\}^n$ produces $y \in \{0, 1\}^n$ such that the expected Hamming distance $d_H(x, y)$ is at most α .

PROOF. Given an input $x \in \{0, 1\}^n$, the corresponding edge weight function w_x is given by $w_x(e_i^{(x_i)}) = 0$ and $w_x(e_i^{(1-x_i)}) = 1$. That is, for each pair of consecutive vertices $i - 1, i$, one of the edges between them will have weight 0 and the other will have weight 1 as determined by the i th bit of the input x .

The algorithm $\mathcal{B}$ is as follows. On input $x \in \{0, 1\}^n$, apply $\mathcal{A}$ to $(\mathcal{G}, w_x)$, and let $\mathcal{P}$ be the path produced. Define $y \in \{0, 1\}^n$ as follows. Let $y_i = 0$ if $e_i^{(0)} \in \mathcal{P}$ and $y_i = 1$ otherwise. Output y .

We first show that this procedure is differentially private. Given neighboring inputs $x, x' \in \{0, 1\}^n$ which differ only on a single coordinate $x_i \neq x'_i$, we have that the associated weight functions w_x and $w_{x'}$ have ℓ_1 distance 2, since they disagree only edges $e_i^{(0)}$ and $e_i^{(1)}$. Consequently, since $\mathcal{A}$ is (ϵ, δ) -differentially private, we have that for any set of values S in the range of $\mathcal{A}$,

$$\begin{aligned} \Pr[\mathcal{A}(w_x) \in S] &\leq e^\epsilon (e^\epsilon \Pr[\mathcal{A}(w_{x'}) \in S] + \delta) + \delta \\ &= e^{2\epsilon} \Pr[\mathcal{A}(w_{x'}) \in S] + (1 + e^\epsilon)\delta. \end{aligned}$$

But algorithm $\mathcal{B}$ only accesses the database x through $\mathcal{A}$. Consequently by the robustness of differential privacy to post-processing, we have that $\mathcal{B}$ is $(2\epsilon, (1 + e^\epsilon)\delta)$ -differentially private.

We now show that the expected number of coordinates in which y disagrees with x is at most α . The shortest path from s to t in $\mathcal{G}$ has length 0, so the expected length of the path $\mathcal{P}$ produced by $\mathcal{A}$ is at most α . Consequently in expectation $\mathcal{P}$ consists of at most α edges $e_i^{(1-x_i)}$. But

Algorithm 3: Private shortest paths

Inputs: $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, $w : \mathcal{E} \rightarrow \mathbb{R}^+$, $\gamma > 0$. For each edge $e \in \mathcal{E}$, do the following:

1. Sample $X_e \leftarrow \text{Lap}(1/\epsilon)$
2. Let $w'(e) = w(e) + X_e + (1/\epsilon) \log(E/\gamma)$, where $E = |\mathcal{E}|$

Output w' .

The approximate shortest path between a pair of vertices $x, y \in \mathcal{V}$ is taken to be the shortest path $\text{SP}_{w'}(x, y)$ in the weighted graph $(\mathcal{G}, w')$.

$y_i \neq x_i$ only if $e_i^{(1-x_i)} \in \mathcal{P}$, so it follows that the expected Hamming distance $d_H(x, y) \leq \alpha$. $\square$

We will now prove two simple and standard lemmas concerning the limits of identifying rows of the input of a differentially private algorithm. In the first lemma, we show that a differentially private algorithm cannot release a particular row of its input with probability greater than $\frac{e^\epsilon + \delta}{1 + e^\epsilon}$. In essence, for $\delta = 0$ this can be interpreted as a statement about the optimality of the technique of randomized response [War65].

LEMMA 5.3. *If algorithm $\mathcal{B} : \{0, 1\}^n \rightarrow \{0, 1\}$ is (ϵ, δ) -differentially private, then if we uniformly sample a random input $X \leftarrow U_n$, we have that for all i ,*

$$\Pr[\mathcal{B}(X) \neq X_i] \geq \frac{1 - \delta}{1 + e^\epsilon}$$

PROOF. We have that $\Pr[\mathcal{B}(X) = X_i]$ is given by

$$\begin{aligned} & \frac{1}{2} \Pr[\mathcal{B}(X_{-i}, 0) = 0] + \frac{1}{2} \Pr[\mathcal{B}(X_{-i}, 1) = 1] \\ & \leq \frac{e^\epsilon}{2} \cdot (\Pr[\mathcal{B}(X_{-i}, 1) = 0] + \Pr[\mathcal{B}(X_{-i}, 0) = 1]) + \delta \\ & = e^\epsilon \Pr[\mathcal{B}(X) \neq X_i] + \delta \end{aligned}$$

so since $\Pr[\mathcal{B}(X) = X_i] = 1 - \Pr[\mathcal{B}(X) \neq X_i]$,

$$\Pr[\mathcal{B}(X) \neq X_i] \geq \frac{1 - \delta}{1 + e^\epsilon}.$$

$\square$

This immediately implies the following result.

LEMMA 5.4. *If algorithm $\mathcal{B} : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is (ϵ, δ) -differentially private, then for some $x \in \{0, 1\}^n$ we have that the expected Hamming distance $d_H(\mathcal{B}(x), x)$ is at least $\frac{n(1-\delta)}{1+e^\epsilon}$.*

PROOF. By Lemma 5.3, projecting onto any coordinate i , the probability that $\Pr[\mathcal{B}(X)_i \neq X_i]$ for uniformly random $X \leftarrow U_n$ is at least $\frac{1-\delta}{1+e^\epsilon}$. Consequently the expected number of coordinates on which $\mathcal{B}(X)$ differs from X is $\mathbf{E}(d_H(\mathcal{B}(X), X)) \geq \frac{n(1-\delta)}{1+e^\epsilon}$. Since this holds for X uniformly random, in particular we have that there exists some $x \in \{0, 1\}^n$ such that

$$\mathbf{E}(d_H(\mathcal{B}(x), x)) \geq \frac{n(1-\delta)}{1+e^\epsilon}.$$

$\square$

We now conclude the proof of Theorem 5.1.

PROOF OF THEOREM 5.1. Assume that there is some algorithm which is (ϵ, δ) -differentially private on $\mathcal{G}$ and always produces a path of error at most α between s and t . Then by Lemma 5.2 we have that there is a $(2\epsilon, (1+e^\epsilon)\delta)$ -differentially private algorithm $\mathcal{B}$ which for all $x \in \{0, 1\}^n$ produces $y \in \{0, 1\}^n$ such that the expected Hamming distance $d_H(x, y)$ is less than α . But by Lemma 5.4, for some x the expected Hamming distance $\mathbf{E}(d_H(x, y)) \geq \frac{n(1-(1+e^\epsilon)\delta)}{1+e^{2\epsilon}} = \alpha$, yielding a contradiction. $\square$

5.2 Upper bound

In this section we show that an extremely simple algorithm matches the lower bound of the previous section up to a logarithmic factor, for fixed ϵ, δ . Consider a direct application of the Laplace mechanism (Lemma 3.2), adding $\text{Lap}(1/\epsilon)$ noise to each edge weight and releasing the resulting values. With high probability all of these $E < V^2$ noise variables will be small, providing a bound on the difference in the weight of any path between the released graph and the original graph. Consequently we can show that if we take the shortest path in the released graph, with 99% probability the length of the same path in the original graph is $O(V \log V)/\epsilon$ longer than optimal.

This straightforward application of the Laplace mechanism almost matches the lower bound of the previous section. Surprisingly, with the same error bound it releases not just a short path between a single pair of vertices but short paths between all pairs of vertices.

One drawback of this argument is that the error depends on the size of the entire graph. In practice we may expect that the shortest path between most pairs of vertices consist of relatively few edges. We would like the error to depend on the number of hops on the shortest path rather than scaling with the number of vertices. We achieve this with a post-processing step that increases the weight of all edges, introducing a preference for few-hop paths. We show that if there is a short path with only k hops, then our algorithm reports a path whose length is at most $O(k \log V/\epsilon)$ longer.

THEOREM 5.5. *For all graphs $\mathcal{G}$ and $\gamma \in (0, 1)$, Algorithm 3 is ϵ -differentially private on $\mathcal{G}$ and computes paths between all pairs of vertices such that with probability $1 - \gamma$, for all pairs of vertices $s, t \in \mathcal{V}$, if there exists a k -hop path of weight W in $(\mathcal{G}, w)$, the path released has weight at most $W + (2k/\epsilon) \log(E/\gamma)$.*

In particular, if the shortest path in $\mathcal{G}$ has k hops, then Algorithm 3 releases a path only $(2k/\epsilon) \log(E/\gamma)$ longer than optimal. This error term is proportional to the number of hops on the shortest path. Noting that the shortest path between any pair of vertices consists of fewer than V hops, we obtain the following corollary.

COROLLARY 5.6. *For any $\gamma \in (0, 1)$, with probability $1 - \gamma$ Algorithm 3 computes paths between every pair of vertices with approximation error at most $(2V/\epsilon) \log(E/\gamma)$.*

PROOF OF THEOREM 5.5. Each random variable X_e is distributed according to $\text{Lap}(1/\epsilon)$, so with probability $1 - \gamma$ we have that $|X_e| \leq (1/\epsilon) \log(1/\gamma)$ for any $\gamma \in (0, 1)$. By a union bound, with probability $1 - \gamma$ all $E < V^2$ of these random variables have magnitude at most $(1/\epsilon) \log(E/\gamma)$. Conditioning on this event, for each edge $e \in \mathcal{E}$, the modified weight computed by the algorithm satisfies

$$w(e) \leq w'(e) \leq w(e) + (2/\epsilon) \log(E/\gamma).$$

Therefore, for any k -hop path $\mathcal{P}$ we have that

$$w(\mathcal{P}) \leq w'(\mathcal{P}) \leq w(\mathcal{P}) + (2k/\epsilon) \log(E/\gamma).$$

For any $s, t \in \mathcal{V}$, if $\mathcal{Q}$ is the path from s to t produced by the algorithm and $\mathcal{Q}'$ is any path from s to t , then we have that $\mathcal{Q}$ is a shortest path under w' , so

$$w(\mathcal{Q}) \leq w'(\mathcal{Q}) \leq w'(\mathcal{Q}') \leq w(\mathcal{Q}') + (2\ell(\mathcal{Q}')/\epsilon) \log(E/\gamma).$$

□

Acknowledgments I am grateful to Salil Vadhan and Shafi Goldwasser for their guidance and encouragement throughout this project. Thanks to Jon Ullman, Tommy MacWilliam and anonymous referees for helpful suggestions. This work was supported by a DOE CSGF Fellowship, NSF Frontier CNS-1413920, DARPA W911NF-15-C-0236, and the Simons Foundation (agreement dated 6-5-12).

6. REFERENCES

- [BBDS12] Jeremiah Blocki, Avrim Blum, Anupam Datta, and Or Sheffet. The johnson-lindenstrauss transform itself preserves differential privacy. In *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, pages 410–419. IEEE, 2012.
- [BBDS13] Jeremiah Blocki, Avrim Blum, Anupam Datta, and Or Sheffet. Differentially private data analysis of social networks via restricted sensitivity. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pages 87–96. ACM, 2013.
- [BNSV15] Mark Bun, Kobbi Nissim, Uri Stemmer, and Salil Vadhan. Differentially private release and learning of threshold functions. *arXiv preprint arXiv:1504.07553*, 2015.
- [CN82] GJ Chang and GL Nemhauser. The k -domination and k -stability problem on graphs. *Techn. Report*, 540, 1982.
- [CSS10] TH Hubert Chan, Elaine Shi, and Dawn Song. Private and continual release of statistics. In *Automata, Languages and Programming*, pages 405–417. Springer, 2010.
- [DKM⁺06] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Advances in Cryptology-EUROCRYPT 2006*, pages 486–503. Springer, 2006.
- [DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography*, pages 265–284. Springer, 2006.
- [DNPR10] Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N Rothblum. Differential privacy under continual observation. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 715–724. ACM, 2010.
- [DR13] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Theoretical Computer Science*, 9(3-4):211–407, 2013.
- [DRV10] Cynthia Dwork, Guy N Rothblum, and Salil Vadhan. Boosting and differential privacy. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 51–60. IEEE, 2010.
- [GLM⁺10] Anupam Gupta, Katrina Ligett, Frank McSherry, Aaron Roth, and Kunal Talwar. Differentially private combinatorial optimization. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1106–1125. Society for Industrial and Applied Mathematics, 2010.
- [GRU12] Anupam Gupta, Aaron Roth, and Jonathan Ullman. Iterative constructions and private data release. In *Theory of Cryptography*, pages 339–356. Springer, 2012.
- [HHR⁺14] Justin Hsu, Zhiyi Huang, Aaron Roth, Tim Roughgarden, and Zhiwei Steven Wu. Private matchings and allocations. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 21–30. ACM, 2014.
- [HLMJ09] Michael Hay, Chao Li, Gerome Miklau, and David Jensen. Accurate estimation of the degree distribution of private networks. In *Data Mining, 2009. ICDM'09. Ninth IEEE International Conference on*, pages 169–178. IEEE, 2009.
- [KNRS13] Shiva Prasad Kasiviswanathan, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. Analyzing graphs with node differential privacy. In *Theory of Cryptography*, pages 457–476. Springer, 2013.
- [KRSY11] Vishesh Karwa, Sofya Raskhodnikova, Adam Smith, and Grigory Yaroslavtsev. Private analysis of graph structure. *Proceedings of the VLDB Endowment*, 4(11):1146–1157, 2011.
- [MM75] A Meir and JW Moon. Relations between packing and covering numbers of a tree. *Pacific J. Math*, 61(1):225–233, 1975.
- [NRS07] Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. Smooth sensitivity and sampling in private data analysis. In *Proceedings of the*

thirty-ninth annual ACM symposium on Theory of computing, pages 75–84. ACM, 2007.

[War65] Stanley L Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.

APPENDIX

A. DISTANCES IN THE PATH GRAPH

In this section we give an explicit description of the private all-pairs distance algorithm for the path graph $\mathcal{P} = (\mathcal{V}, \mathcal{E})$ with vertex set $\mathcal{V} = [V]$ and edge set $\mathcal{E} = \{(i, i+1) : i \in [V-1]\}$. This result is a restatement of a result of [DNPR10], and is generalized to trees in Section 4.1. This alternate argument for the special case of the path graph is included for illustration.

The idea behind the construction is as follows. We designate a small set of hubs and store more accurate distances between consecutive hubs. As long as we can accurately estimate the distance from any vertex to the nearest hub and the distance between any pair of hubs, we can use these distances to obtain an estimate of the distance between any pair of vertices. Given any pair of vertices $x, y \in \mathcal{V}$, in order to estimate the distance $\text{dist}(x, y)$, we first find the hubs h_x and h_y nearest to x and y . We estimate the distance $\text{dist}(x, y)$ by adding our estimates for the distances $\text{dist}(x, h_x)$, $\text{dist}(h_x, h_y)$, and $\text{dist}(h_y, y)$.

Instead of simply using a single set of hubs, we will use a hierarchical tree of hubs of different levels. There will be many hubs of low level and only a small number of hubs of high level. Each hub will store an estimate of the distance to the next hub at the same level and every lower level. Hubs higher in the hierarchy will allow us to skip directly to distant vertices instead of accruing error by adding together linearly many noisy estimates. In order to estimate the distance between a particular pair of vertices $x, y \in \mathcal{V}$, we will only consider a small number of hubs on each level of the hierarchy. Since the total number of levels is not too large, this will result in a much more accurate differentially private estimate of the distance between x and y .

THEOREM A.1. *Let $\mathcal{P} = (\mathcal{V}, \mathcal{E})$ be the path graph on V vertices. For any $\epsilon > 0$, there is algorithm $\mathcal{A}$ that is ϵ -differentially private on $\mathcal{P}$ that on input $w : \mathcal{E} \rightarrow \mathbb{R}^+$ releases approximate all-pairs distances such that for each released distance, with probability $1 - \gamma$ the approximation error is $O(\log^{1.5} V \log(1/\gamma))/\epsilon$ for any $\gamma \in (0, 1)$.*

PROOF. For fixed k , define nested subsets $\mathcal{V} = S_0 \supset S_1 \supset \dots \supset S_{k-1}$ of the vertex set $\mathcal{V} = [V]$ as follows. Let

$$S_i = \{x \in [V] : V^{i/k} \mid x\}.$$

That is, S_0 consists of all the vertices, and in general S_i consists of one out of every $V^{1/k}$ vertices on the path. Then $|S_i| = V^{(k-i)/k}$. Let $s_{i,1}, s_{i,2}, \dots, s_{i,|S_i|}$ denote the elements of S_i in increasing order, for each i . Using the Laplace mechanism (Lemma 3.2), release noisy distances between each pair of consecutive vertices $s_{i,j}, s_{i,j+1}$ of each set, adding noise proportional to $\text{Lap}(k/\epsilon)$. Note that since the vertices in each S_i are in increasing order, each edge of $\mathcal{P}$ is only considered for a single released difference from each set. Consequently the total sensitivity to release all of these distances

is k , so releasing these noisy distances is ϵ -differentially private. Using post-processing and these special distances, we will compute approximate all-pairs distances with small error.

For any pair of vertices x, y , consider the path $\mathcal{P}[x, y]$ in $\mathcal{P}$ between x and y , and let i be the largest index such that S_i contains multiple vertices of $\mathcal{P}[x, y]$. We must have that $S_i \cap \mathcal{P}[x, y] < 2V^{1/k}$, since otherwise S_{i+1} would contain at least two vertices on $\mathcal{P}[x, y]$. Let x_i, y_i denote the first and last vertices in $S_i \cap \mathcal{P}[x, y]$. For $j < i$, let x_j denote the first vertex in $S_j \cap \mathcal{P}[x, x_i]$ and let y_j denote the last vertex in $S_j \cap \mathcal{P}[y_i, y]$. There are at most $1 + V^{1/k}$ vertices of S_j in $\mathcal{P}[x_j, x_{j+1}]$, since otherwise this interval would contain another vertex of S_{j+1} . Similarly, there are at most $1 + V^{1/k}$ vertices of S_j in $\mathcal{P}[y_{j+1}, y_j]$. Therefore we can express the distance from x_j to x_{j+1} as the sum of at most $V^{1/k}$ distances which were estimated in S_j , and similarly for the distance from y_{j+1} to y_j .

Putting this all together, we can estimate the distance from $x = x_0$ to $y = y_0$ as the sum of at most $2(i+1)V^{1/k} \leq 2kV^{1/k}$ approximate distances which were released. But each of these distances is released with noise distributed according to $\text{Lap}(k/\epsilon)$. Consequently the total error on the estimated distance from x to y is the sum of at most $2kV^{1/k}$ random variables distributed according to $\text{Lap}(k/\epsilon)$. Taking $k = \log V$, the error is the sum of at most $4 \log V$ variables distributed according to $\text{Lap}(\log V/\epsilon)$. By Lemma 3.1, with probability at least $1 - \gamma$ the sum of these $4 \log V$ variables is bounded by $O(\log^{1.5} V \log(1/\gamma))/\epsilon$ for any $\gamma \in (0, 1)$. Consequently this is an ϵ -differentially private algorithm which releases all-pairs distances in the path graph $\mathcal{P}$ such that for any $\gamma \in (0, 1)$, with probability at least $1 - \gamma$ the error in each released distance is at most $O(\log^{1.5} V \log(1/\gamma))/\epsilon$. $\square$

B. OTHER GRAPH PROBLEMS

In this section we consider some additional queries on graphs in the private edge weights model.

B.1 Almost-minimum spanning tree

We first consider the problem of releasing a low-cost spanning tree. The work of [NRS07] showed how to privately approximate the cost of the minimum spanning tree in a somewhat related privacy setting. We seek to release an actual tree of close to minimal cost under our model. Using techniques similar to the lower bound for shortest paths from Section 5.1, we obtain a lower bound of $\Omega(V)$ for the error of releasing a low-cost spanning tree, and show that the Laplace mechanism yields a spanning tree of cost $O(V \log V)$

Figure 3: (Left) The graph used in the reduction of Lemma B.2. (Right) A single gadget in the graph used in the reduction of Lemma B.5.

more than optimal. Note that in this section edge weights are permitted to be negative.

THEOREM B.1. *There exists a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ such that for any spanning tree algorithm $\mathcal{A}$ that is (ϵ, δ) -differentially private on $\mathcal{G}$, there exist edge weights $w : \mathcal{E} \rightarrow \{0, 1\}$ such that the expected weight of the spanning tree $\mathcal{A}(w)$ is at least*

$$\alpha = (V - 1) \cdot \left(\frac{1 - (1 + e^\epsilon)\delta}{1 + e^{2\epsilon}} \right)$$

longer than the weight of the minimum spanning tree. In particular, for sufficiently small ϵ and δ , $\alpha \geq 0.49(V - 1)$.

We first prove a lemma reducing the problem of reidentifying rows in a database to privately finding an approximate minimum spanning tree. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be the $(n+1)$ -vertex graph with vertex set $\mathcal{V} = \{0, \dots, n\}$ and a pair of edges $e_i^{(0)}$ and $e_i^{(1)}$ between vertex 0 and each vertex $i > 0$, as shown in Figure 3. (As in Lemma 5.2, this is a multigraph, but we can transform it into a simple graph by adding n extra vertices, changing the bound obtained by a factor of 2.)

LEMMA B.2. *Let $\mathcal{G}$ be the graph defined in the previous paragraph, and let $\epsilon, \delta \geq 0$. For any α , let $\mathcal{A}$ be an algorithm that is (ϵ, δ) -differentially private on $\mathcal{G}$ that on input $w : \mathcal{E} \rightarrow \{0, 1\}$ produces a spanning tree whose weight in expectation is at most α greater than optimal. Then there exists a $(2\epsilon, (1 + e^\epsilon)\delta)$ -differentially private algorithm $\mathcal{B}$ which on input $x \in \{0, 1\}^n$ produces $y \in \{0, 1\}^n$ such that the expected Hamming distance $d_H(x, y)$ is at most α .*

PROOF. The outline of the proof is the same as that of Lemma 5.2. For input $x \in \{0, 1\}^n$, the corresponding edge weight function w_x is given by $w_x(e_i^{(x_i)}) = 0$ and $w_x(e_i^{(1-x_i)}) = 1$, where x_i is the i th bit of x .

We define algorithm $\mathcal{B}$ as follows. On input $x \in \{0, 1\}^n$, apply $\mathcal{A}$ to $(\mathcal{G}, w_x)$, and let $\mathcal{T}$ be the tree produced. Define $y \in \{0, 1\}^n$ by setting $y_i = 0$ if $e_i^{(0)} \in \mathcal{T}$ and $y_i = 1$ otherwise. Output y .

It is straightforward to verify that $\mathcal{B}$ is $(2\epsilon, (1 + e^\epsilon)\delta)$ -differentially private. We now bound the expected Hamming distance of x and y . The minimum spanning tree in $\mathcal{G}$ has weight 0, so the expected weight of $\mathcal{T}$ is at most α and $\mathcal{T}$ must consist of at most α edges $e_i^{(1-x_i)}$. But $y_i \neq x_i$ only if $e_i^{(1-x_i)} \in \mathcal{T}$, so in expectation $d_H(x, y) \leq w(\mathcal{T}) \leq \alpha$. $\square$

We now complete the proof of Theorem B.1.

PROOF OF THEOREM B.1. Assume that there is some (ϵ, δ) -differentially private algorithm which on all inputs produces a spanning tree with expected weight less than α more than optimal. By Lemma B.2, there is a $(2\epsilon, (1 + e^\epsilon)\delta)$ -differentially private algorithm which for all $x \in \{0, 1\}^n$ produces $y \in \{0, 1\}^n$ with expected Hamming distance less than α . But then Lemma 5.4, for some x the expected Hamming distance $\mathbf{E}(d_H(x, y)) \geq \frac{n(1 - (1 + e^\epsilon)\delta)}{1 + e^{2\epsilon}} = \alpha$, yielding a contradiction. $\square$

We now show that the Laplace mechanism (Lemma 3.2) almost matches this lower bound.

THEOREM B.3. *For any $\epsilon, \gamma > 0$ and $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, there is an algorithm $\mathcal{A}$ that is ϵ -differentially private on $\mathcal{G}$ that on input $w : \mathcal{E} \rightarrow \mathbb{R}$ releases with probability $1 - \gamma$ a spanning tree of weight at most $((V - 1)/\epsilon) \log(E/\gamma)$ larger than optimal.*

PROOF. Consider the algorithm that adds noise X_e distributed according to $\text{Lap}(1/\epsilon)$ for each edge $e \in \mathcal{E}$ and releases the minimum spanning tree on the resulting graph $(\mathcal{G}, w')$. This is ϵ -differentially private, since it is post-processing of the Laplace mechanism. We now show that the resulting error is small. By a union bound, with probability $1 - \gamma$ we have that $|X_e| \leq (1/\epsilon) \log(E/\gamma)$ for every $e \in \mathcal{E}$. Consequently, conditioning on this event, if $\mathcal{T}$ is the spanning tree released by the algorithm and $\mathcal{T}^*$ is the minimum spanning tree, then we have that

$$\begin{aligned} w(\mathcal{T}) &\leq w'(\mathcal{T}) + \frac{V-1}{\epsilon} \cdot \log(E/\gamma) \\ &\leq w'(\mathcal{T}^*) + \frac{V-1}{\epsilon} \cdot \log(E/\gamma) \\ &\leq w(\mathcal{T}^*) + \frac{2(V-1)}{\epsilon} \cdot \log(E/\gamma). \end{aligned}$$

$\square$

B.2 Low-weight matching

We now consider the problem of releasing a minimum weight matching in a graph in our model. As for the minimum spanning tree problem, a minor modification of the lower bound for shortest paths from Section 5.1 yields a similar result. For comparison, [HHR⁺14] use similar reconstruction techniques to obtain a lower bound for a matching problem on bipartite graphs in a somewhat different model in which all edge weights are in $[0, 1]$ and neighboring graphs can differ on the weights of the edges incident to a single left vertex. We show a lower bound of $\Omega(V)$ for the error of releasing a low-weight matching tree, and show that the matching released by the Laplace mechanism has weight $O(V \log V)$ greater than optimal.

The theorems in this section are stated for the problem of finding a minimum weight perfect matching. We can also obtain identical results for the problem of finding a minimum weight matching which is not required to be perfect, and for the corresponding maximum weight matching problems. Our results apply to both bipartite matching and general matching. Note that in this section edge weights are permitted to be negative.

THEOREM B.4. *There exists a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ such that for any perfect matching algorithm $\mathcal{A}$ which is (ϵ, δ) -differentially private on $\mathcal{G}$, there exist edge weights $w : \mathcal{E} \rightarrow \{0, 1\}$ such that the expected weight of the matching $\mathcal{A}(w)$ is at least*

$$\alpha = \left(\frac{V}{4} \right) \cdot \left(\frac{1 - (1 + e^\epsilon)\delta}{1 + e^{2\epsilon}} \right)$$

larger than the weight of the min-cost perfect matching. In particular, for sufficiently small ϵ, δ , $\alpha \geq 0.12 \cdot V$.

The following lemma reduces the problem of reidentification in a database to finding a low-cost matching. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be the $4n$ -vertex graph with vertex set $\mathcal{V} = \{(b_1, b_2, c) : b_1, b_2 \in \{0, 1\}, c \in [n]\}$ and edges from $(0, b, c)$ to $(1, b', c)$ for every $b, b' \in \{0, 1\}$, $c \in [n]$. That is, $\mathcal{G}$ consists of n disconnected hourglass-shaped gadgets as shown in Figure 3.

LEMMA B.5. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be the graph defined in the previous paragraph. For any α , let $\mathcal{A}$ be an algorithm that is (ϵ, δ) -differentially private on $\mathcal{G}$ that on input $w : \mathcal{E} \rightarrow \{0, 1\}$ produces a perfect matching of expected weight at most α*

greater than optimal. Then there exists a $(2\epsilon, (1 + e^\epsilon)\delta)$ -differentially private algorithm $\mathcal{B}$ which on input $x \in \{0, 1\}^n$ produces $y \in \{0, 1\}^n$ with expected Hamming distance to x at most α .

PROOF. For any input $x \in \{0, 1\}^n$, the corresponding weight function w_x assigns weight 1 to the edge connecting vertex $(0, 1, i)$ to $(1, 1 - x_i, i)$ for each $i \in [n]$, and assigns weight 0 to the other $3n$ edges. The algorithm $\mathcal{B}$ is as follows. On input $x \in \{0, 1\}^n$, apply $\mathcal{A}$ to $(\mathcal{G}, w_x)$, and let $\mathcal{M}$ be the matching produced. Define $y \in \{0, 1\}^n$ as follows. Let $y_i = 0$ if the edge from $(0, 1, i)$ to $(1, 0, i)$ is in the matching, and $y_i = 1$ otherwise. Output y .

Algorithm $\mathcal{B}$ is clearly $(2\epsilon, (1 + e^\epsilon)\delta)$ -differentially private. We will have that $y_i \neq x_i$ only if the edge from $(0, 1, i)$ to $(1, 1 - x_i, i)$ is in the matching, so the expected Hamming distance is at most the expected size of the matching produced by $\mathcal{A}$, which is at most α . $\square$

We now conclude the proof of Theorem B.4.

PROOF OF THEOREM B.4. The result follows by combining Lemma B.5 with Lemma 5.4. $\square$

Using the Laplace mechanism (Lemma 3.2), we obtain a nearly-matching upper bound.

THEOREM B.6. For any $\epsilon, \gamma > 0$ and $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ containing a perfect matching, there is an algorithm $\mathcal{A}$ that is ϵ -differentially private on $\mathcal{G}$ that on input $w : \mathcal{E} \rightarrow \mathbb{R}$ releases with probability $1 - \gamma$ a perfect matching of weight at most $(V/\epsilon) \log(E/\gamma)$ larger than optimal.

PROOF. Consider the algorithm that adds noise X_e distributed according to $\text{Lap}(1/\epsilon)$ for each edge $e \in \mathcal{E}$ and releases the minimum-weight perfect matching on the resulting graph $(\mathcal{G}, w')$. This is ϵ -differentially private, since it is post-processing of the Laplace mechanism. We now show that the resulting error is small. With probability $1 - \gamma$ we have that $|X_e| \leq (1/\epsilon) \log(E/\gamma)$ for every $e \in \mathcal{E}$. Consequently, conditioning on this event, if $\mathcal{M}$ is the matching released by the algorithm and $\mathcal{M}^*$ is the minimum-weight matching, then we have that

$$\begin{aligned} w(\mathcal{M}) &\leq w'(\mathcal{M}) + \frac{V}{2\epsilon} \cdot \log(E/\gamma) \\ &\leq w'(\mathcal{M}^*) + \frac{V}{2\epsilon} \cdot \log(E/\gamma) \\ &\leq w(\mathcal{M}^*) + \frac{V}{\epsilon} \cdot \log(E/\gamma). \end{aligned}$$

$\square$