Edge Differential Privacy for Algebraic Connectivity of Graphs

Bo Chen, Calvin Hawkins, Kasra Yazdani, Matthew Hale*

***Abstract*— Graphs are the dominant formalism for modeling multi-agent systems. The algebraic connectivity of a graph is particularly important because it provides the convergence rates of consensus algorithms that underlie many multi-agent control and optimization techniques. However, sharing the value of algebraic connectivity can inadvertently reveal sensitive information about the topology of a graph, such as connections in social networks. Therefore, in this work we present a method to release a graph's algebraic connectivity under a graph-theoretic form of differential privacy, called *edge differential privacy*. Edge differential privacy obfuscates differences among graphs' edge sets and thus conceals the absence or presence of sensitive connections therein. We provide privacy with bounded Laplace noise, which improves accuracy relative to conventional unbounded noise. The private algebraic connectivity values are analytically shown to provide accurate estimates of consensus convergence rates, as well as accurate bounds on the diameter of a graph and the mean distance between its nodes. Simulation results confirm the utility of private algebraic connectivity in these contexts.**

I. Introduction

Graphs are used to model a wide range of interconnected systems, including multi-agent control systems [1], social networks [2], and others [3]. Various properties of these graphs have been used to analyze controllers and dynamical processes over them, such as reaching a consensus [4], the spread of a virus [5], robustness to connection failures [6], and others. Graphs in these applications may contain sensitive information, e.g., one's close friendships in the case of a social network, and it is essential that these analyses do not inadvertently leak any such information.

Unfortunately, it is well-established that even graph-level analyses may inadvertently reveal sensitive information about individuals in them, such as the absence or presence of individual nodes in a graph [7] and the absence or presence of specific edges between them [8]. Similar privacy threats have received attention in the data science community, where graphs represent datasets and the goal is to enable data analysis while safeguarding the data of individuals in those datasets.

Differential privacy is one well-studied tool for doing so. Differential privacy is a statistical notion of privacy that has several desirable properties: (i) it is robust to side information, in that learning additional information about data-producing entities does not weaken privacy by much [9], and (ii) it is immune to post-processing, in that arbitrary post-hoc computations on private data do not weaken privacy [10].

*Department of Mechanical and Aerospace Engineering at the University of Florida, Gainesville, FL USA. Emails: {bo.chen,calvin.hawkins,kasra.yazdani,matthewhale}@ufl.edu. This work was supported in part by NSF under CAREER Grant #1943275 and by AFOSR under Grant #FA9550-19-1-0169.

There exist numerous differential privacy implementations for graph properties, including counts of subgraphs [8], degree distributions [11], and other frequent patterns in graphs [12]. These privacy mechanisms generally follow the pattern of computing the quantity of interest, adding carefully calibrated noise to it, and releasing its noisy form. Although simple, this approach strongly protects data with a suite of guarantees provided by differential privacy [10].

The need for privacy comes from the inferences that one can draw about a graph from these quantities, as detailed in [13]–[15]. Decades of research in algebraic graph theory have quantified connections between λ_2 and a myriad of other graph properties; see [16] for a summary. Accordingly, λ_2 implicates the same ability to draw inferences and hence gives rise to the same types of privacy concerns.

We therefore protect a graph's algebraic connectivity using edge differential privacy, which obfuscates the absence and/or presence of a pre-specified number of edges. A graph's algebraic connectivity (also called its Fiedler value) is equal to the second-smallest eigenvalue of its Laplacian. This value plays a central role in the study of multi-agent systems because it sets the convergence rates of consensus algorithms [17], which appear directly or in modified form in formation control [18], connectivity control [19], and many distributed optimization algorithms [20]. The private release of the algebraic connectivity of a graph would enable the computation of such convergence rates by a network analyst, while protecting sensitive properties of individuals in these graphs. This paper provides the means to do so.

Our implementation uses the recent bounded Laplace mechanism [21], which develops a mechanism that ensures that private scalars lie in a specified interval. The algebraic connectivity of a graph is bounded below by zero and above by the number of nodes in a graph, and we confine private outputs to this interval by applying the mechanism in [21] to the privatization of λ_2.

Contributions: We provide closed-form values for the sensitivity and other constants needed to define a privacy mechanism for algebraic connectivity, and this is the first contribution of this paper. The second contribution is bounding the error that privacy induces in the convergence rates of consensus. Differential privacy has made inroads in control applications ranging from LQ control [22], state estimation [23], [24], formation control [25], Markov decision processes [26], symbolic systems [27], multi-agent optimization [28], [29], reinforcement learning [30], location-based services [31], designing nonlinear observers [32], and others, due in part to the accurate analyses and high performance one can maintain even with privacy implemented. We show

that this is the case for the consensus setting as well. Our third contribution is the use of the private values of algebraic connectivity to bound other graph properties, namely the diameter of graphs and the mean distance between their nodes.

We note that [33] has developed a different approach to privacy for the eigendecomposition of a graph's adjacency matrix. Given our motivation by multi-agent systems, we focus on a graph's Laplacian, which commonly appears in multi-agent controllers, and we derive simpler forms for the distribution of noise required, as well as a privacy mechanism that does not require any post-processing.

The rest of the paper is organized as follows. Section II provides background and problem statements. Section III develops the differential privacy mechanism for algebraic connectivity. Next, Section IV uses this mechanism to privately compute consensus convergence rates, and Section V applies it to bounding other graph properties. Then, Section VI provides simulation results and Section VII provides concluding remarks.

II. Background and Problem Formulation

In this section, we briefly review the necessary background on graph theory and differential privacy, followed by formal problem statements.

A. Background on Graph Theory

We consider an undirected, unweighted graph $G = (V, E)$ defined over a set of nodes $V = \{1, \dots, n\}$ with edge set $E \subset V \times V$. The pair (i, j) belongs to E if nodes i and j share an edge, and $(i, j) \notin E$ otherwise. Let $\mathcal{G}_n$ denote the set of all graphs on n nodes. We let $d_i = |\{j \in V \mid (i, j) \in E\}|$ denote the degree of node $i \in V$. The degree matrix $D(G) \in \mathbb{R}^{n \times n}$ is the diagonal matrix $D(G) = \operatorname{diag}\big(d_1, \dots, d_n\big)$. The adjacency matrix of G is

$$(H(G))_{ij} = \begin{cases} 1 & (i, j) \in E \\ 0 & \text{otherwise} \end{cases}.$$

We denote the Laplacian of the graph G by $L(G) = D(G) - H(G)$, which we simply refer to by L when the associated graph is clear from context.

Let the eigenvalues of L be ordered according to $\lambda_1(L) \leq \lambda_2(L) \leq \cdots \leq \lambda_n(L)$. The matrix L is symmetric and positive semidefinite, and thus $\lambda_i(L) \geq 0$ for all i. All graphs G have $\lambda_1(L) = 0$, and a seminal result shows that $\lambda_2(L) > 0$ if and only if G is connected [34]. Thus, λ_2 is often called the *algebraic connectivity* of a graph. Throughout this paper, we consider connected graphs.

The value of λ_2 encodes a great deal of information about G: its value is non-decreasing in the number of edges in G, and algebraic connectivity is closely related to graph diameter and various other algebraic properties of graphs [16]. The value of λ_2 also characterizes the performance of consensus algorithms. Specifically, worst-case disagreement in a consensus protocol decays proportionally to $e^{-\lambda_2 t}$ [35].

B. Background on Differential Privacy

Differential privacy is enforced by a *mechanism*, which is a randomized map. Given "similar" inputs, a differential privacy mechanism produces outputs that are approximately indistinguishable from each other. Formally, a mechanism must obfuscate differences between inputs that are *adjacent*[1].

Definition 1. Let $A \in \mathbb{N}$ be given, and fix a number of nodes $n \in \mathbb{N}$. Two graphs on n nodes, G and G', are adjacent if they differ by A edges. We express this mathematically via

$$\operatorname{Adj}_A(G, G') = \begin{cases} 1 & |E(G) \Delta E(G')| \leq A \\ 0 & \text{otherwise} \end{cases},$$

where $S_1 \Delta S_2 = (S_1 \backslash S_2) \cup (S_2 \backslash S_1)$ is the symmetric difference of two sets and $|\cdot|$ denotes cardinality. $\Diamond$

Thus, A is the number of edges whose absence or presence must be concealed by privacy. In other words, differential privacy for λ_2 must make any graph approximately indistinguishable from any graph within A edges of it when the private value of λ_2 is released.

Next, we briefly review differential privacy; see [10] for a complete exposition. A privacy mechanism $\mathcal{M}$ for a function f is obtained by first computing the function f on a given input x, and then adding noise to the output. The distribution of noise depends on the sensitivity of the function f to changes in its input, described below. It is the role of a mechanism to approximate functions of sensitive data with private responses, and we next state this formally.

Definition 2 (*Differential privacy; [10]*)**.** Let $\epsilon > 0$ and $\delta \in [0, 1)$ be given and fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Then a mechanism $\mathcal{M} : \Omega \times \mathcal{G}_n \to \mathbb{R}$ is (ϵ, δ)-differentially private if, for all adjacent graphs $G, G' \in \mathcal{G}_n$,

$$\mathbb{P}\big[\mathcal{M}(G) \in S\big] \leq \exp(\epsilon) \cdot \mathbb{P}\big[\mathcal{M}\left(G'\right) \in S\big] + \delta$$

for all sets S in the Borel σ-algebra over $\mathbb{R}$. $\Diamond$

The value of ϵ controls the amount of information shared, and typical values range from 0.1 to $\log 3$ [10]. The value of δ can be regarded as the probability that more information is shared than ϵ should allow, and typical values range from 0 to 0.05. Smaller values of both imply stronger privacy. Given ϵ and δ, a privacy mechanism must enforce Definition 2 for all graphs adjacent in the sense of Definition 1.

We next define the sensitivity of λ_2, which will be used later to calibrate the variance of privacy noise. With a slight abuse of notation, we treat λ_2 as a function $\lambda_2 : \mathcal{G}_n \to \mathbb{R}$.

Definition 3. The sensitivity of λ_2 is the greatest difference between its values on Laplacians of adjacent graphs. Formally, given A,

$$\Delta\lambda_2 = \max_{\substack{G, G' \in \mathcal{G}_n \\ \operatorname{Adj}_A(G, G') = 1}} \big|\lambda_2(L) - \lambda_2(L')\big|,$$

[1]The word "adjacency" appears in two forms in this paper: for the adjacency matrix H above, and for the adjacency relation used by differential privacy. The adjacency matrix appears only in this section and only to defined the graph Laplacian. All uses of "adjacent" and "adjacency" in the rest of the paper pertain to differential privacy (not the adjacency matrix).

where L and L' are the Laplacians of G and G'. $\Diamond$

We next state the problems that we will solve.

Problem 1. *Given the adjacency relation in Definition 1, develop a mechanism to provide (ϵ, δ)-differentially privacy for the algebraic connectivity of a graph G.*

Problem 2. *Given a differentially private algebraic connectivity, quantify the accuracy of consensus protocol convergence rate estimates that use it.*

Problem 3. *Given a private algebraic connectivity, develop bounds on the expectation of other graph properties.*

Graph privacy threats have been observed in the data science community for other scalar-valued graph properties, such as counts of subgraphs and triangles [13], degree sequences [14], and numerous others [15]. These threats have been addressed by developing new mechanisms to provide differential privacy to the graph properties of interest. Given the possibility of privacy breaches associated with releasing λ_2 and the wide use of λ_2 in analyzing multi-agent systems, our solution to Problem 1 will develop techniques to protect λ_2 with differential privacy, and our solutions to Problems 2 and 3 will quantify privacy's impact where we expect λ_2 to be used.

III. Privacy mechanism for λ_2

In this section we develop the privacy mechanism that enforces edge differential privacy. We start by providing a bound on the sensitivity in Definition 3.

Lemma 1. Fix an adjacency parameter $A \in \mathbb{N}$. Then, with respect to the adjacency relation in Definition 1, the sensitivity of λ_2 satisfies $\Delta\lambda_2 \leq 2A$.

Proof: See [36, Appendix A]. ■

Noise is added by a mechanism, which is a randomized map used to implement differential privacy. The Laplace mechanism is widely used, and it adds noise from a Laplace distribution to sensitive data (or functions thereof). The standard Laplace mechanism has support on all of $\mathbb{R}$, though, for graphs on n nodes, λ_2 is known to lie in the interval $[0, n]$. One can add Laplace noise and then project the result onto $[0, n]$ (which is differentially private because the projection is merely post-processing), though similar approaches have been shown to produce highly inaccurate private data [37]. Instead, we use the bounded Laplace mechanism in [21]; though bounded Laplace noise appeared earlier in the privacy literature, to the best of our knowledge [21] is the first work to rigorously analyze its privacy properties. We state it in a form amenable to use with λ_2.

Definition 4. Let $b > 0$ and let $D = [0, n]$. Then the bounded Laplace mechanism $W_{\lambda_2} : \Omega \to D$, for each $\lambda_2 \in D$, is given by its probability density function $f_{W_{\lambda_2}}$ as

$$f_{W_{\lambda_2}}(x) = \begin{cases} 0 & \text{if } x \notin D \\ \frac{1}{C_{\lambda_2}(b)} \frac{1}{2b} e^{-\frac{|x-\lambda_2|}{b}} & \text{if } x \in D \end{cases},$$

where $C_{\lambda_2}(b) = \int_D \frac{1}{2b} e^{-\frac{|x-\lambda_2|}{b}} dx$ is a normalizing term. $\Diamond$

Next, we establish an algebraic relation for b which lets the bounded Laplace mechanism satisfy the theoretical guarantees of (ϵ, δ)-differential privacy in Definition 2.

Theorem 1 (*Privacy mechanism for λ_2; Solution to Problem 1*)**.** Let $\epsilon > 0$ and $\delta \in (0, 1)$ be given. Then for the bounded Laplace mechanism W_{λ_2} in Definition 4, choosing b according to

$$b \geq \frac{2A}{\epsilon - \log\left(\frac{1 - \frac{1}{2}e^{-\frac{2A}{b}}\left(1 + e^{-\frac{n}{b}-1}\right)}{1 - \frac{1}{2}\left(1 + e^{-\frac{n}{b}}\right)}\right) - \log(1-\delta)} \tag{1}$$

satisfies (ϵ, δ)-differentially privacy.

Proof: See [36, Appendix B]. ■

Note that b appears on both sides of (1). In [21], the authors provide an algorithm to solve for b using the bisection method, and we use this in the remainder of the paper to find b.

IV. Applications to Consensus

In this section, we solve Problem 2 and bound the error in consensus convergence rates when they are computed using private values of λ_2. As discussed in the introduction, consensus protocols underlie a number of multi-agent control and optimization algorithms, e.g., [17]–[20]. Consider a network of n agents running a consensus protocol with communication topology modeled by an undirected, unweighted graph G. To protect the connections in this graph, a differentially private version of λ_2 is used for analysis. This computation would be performed, for example, by a network analyst who only has access to the private value of λ_2 that has been shared by a data curator.

In continuous time, a consensus protocol takes the form $\dot{x} = -L(G)x$, where $L(G)$ is the graph Laplacian. This protocol converges to the average of agents' initial state values with worst-case error bound at time t proportional to $e^{-\lambda_2 t}$ [35]. Let $r(t) = e^{-\lambda_2 t}$ be the true convergence rate for the network G. Let $\tilde{\lambda}_2$ be the output of the bounded Laplace mechanism with privacy parameters ϵ and δ. Let $\tilde{r}(t) = e^{-\tilde{\lambda}_2 t}$ be the corresponding convergence rate estimate. To compare the estimated convergence rate under privacy to the true convergence rate, we analyze $|\tilde{r}(t) - r(t)|$.

Note that as $t \to \infty$, both $\tilde{r}(t) \to 0$ and $r(t) \to 0$, which implies that $|\tilde{r}(t) - r(t)| \to 0$ as well. Although the error in the convergence rate estimate goes to 0 asymptotically, we are interested in analyzing the error at all values of t, which will give insight into the short-run utility of private convergence rate estimates. To accomplish this, we give a concentration bound that bounds the probability $\mathbb{P}\left(|\tilde{r}(t) - r(t)| \geq a\right)$ in terms of t, the true algebraic connectivity λ_2, and the level of the privacy encoded in b that is determined using ϵ and δ.

Theorem 2 (*Convergence rate concentration bound; Solution to Problem 2*)**.** Let $C_{\lambda_2}(b) = 1 - \frac{1}{2}\left(e^{-\frac{\lambda_2}{b}} + e^{-\frac{n-\lambda_2}{b}}\right)$.

Then, for $t > 0$ and a fixed λ_2 and b,

$$\mathbb{P}\left(|\tilde{r}(t) - r(t)| \geq a\right) \leq \frac{1}{C_{\lambda_2}(b)} \frac{1}{2a} \left(\rho_1(t) + \rho_2(t) - \rho_3(t)\right),$$

where

$$\rho_1(t) = \frac{e^{-\lambda_2(\frac{1}{b}+t)}\left(-bte^{\frac{\lambda_2}{b}} + bt + e^{\lambda_2 t} - 1\right)}{bt - 1},$$

$$\rho_2(t) = e^{-\lambda_2 t}\left(1 - e^{\frac{\lambda_2 - n}{b}}\right),\ \rho_3(t) = \frac{e^{-\lambda_2 t} - e^{\frac{\lambda_2 - n(bt+1)}{b}}}{bt+1}.$$

Proof: See [36, Appendix C]. ■

Taking limits of the bound in Theorem 2 shows that as $t \to \infty$, $\mathbb{P}\left(|\tilde{r}(t) - r(t)| \geq a\right) \to 0$ for all a, and thus this bound has the expected asymptotic behavior.

We can use Theorem 2 to further characterize the transient response of error in the estimated consensus convergence rate. Specifically, we can bound the time required for the error in the convergence rate estimate to be larger than some threshold a only with probability smaller than η. Formally, given a threshold $a > 0$ and probability $\eta > 0$, we bound the times t for which $\mathbb{P}\left(|\tilde{r}(t) - r(t)| \geq a\right) \leq \eta$.

Theorem 3. Fix $a > 0$ and $\eta \in (0,1)$. Let $\epsilon > 0$ and $\delta \in (0,1)$ be given and compute the scale parameter $b > 0$. Consider a graph on n nodes with algebraic connectivity λ_2. If $\lambda_2 \leq \frac{n}{2}$, then we have $\mathbb{P}\left(|\tilde{r}(t) - r(t)| \geq a\right) \leq \eta$ for

$$t \geq \frac{\left(e^{-\frac{\lambda_2}{b}} - e^{\frac{\lambda_2 - n}{b}}\right)\frac{b}{\lambda_2 e} + 2aC_{\lambda_2}(b)\eta + 1}{2aC_{\lambda_2}\eta b}.$$

If $\lambda_2 > \frac{n}{2}$, then the desired bound holds for $t \geq \frac{2aC_{\lambda_2}(b)\eta + 1}{2aC_{\lambda_2}(b)\eta b}$.

Proof: See [36, Appendix D]. ■

We note that the statistics of the differential privacy mechanism can be released without harming privacy. Therefore, the values of $C_{\lambda_2}(b)$ and b can be publicly released. A network analyst can thus compute these bounds for any choices of a and η of interest. Because the exact value of λ_2 is unknown, they can compute the maximum value of these two times to find a time after which the desired error bound always holds.

Notice that the two conditions on t in Theorem 3 only vary by a factor of $\left(e^{-\frac{\lambda_2}{b}} - e^{\frac{\lambda_2 - n}{b}}\right)\frac{b}{\lambda_2 e}$ in the numerator, and when λ_2 is large this term is negative. This means that if λ_2 is large, the required time for $\mathbb{P}\left(|\tilde{r}(t) - r(t)| \geq a\right) \leq \eta$ is smaller than if λ_2 was small. In Section VI, we provide simulation results and further commentary for Theorem 3.

Beyond the consensus protocol, the value of λ_2 is related to many other graph properties [16], and we next show how the private value of λ_2 can still be used to accurately bound two other properties of interest.

V. Bounding Other Graph Properties

There exist numerous inequalities relating λ_2 to other quantitative graph properties [16], [35], and one can therefore expect that the private λ_2 will be used to estimate other quantitative characteristics of graphs. To illustrate the utility of doing so, in this section we bound the graph diameter d and mean distance ρ in terms of the private value $\tilde{\lambda}_2$.

Both d and ρ measure graph size and provide insight into how easily information can be transferred across a network [38]. We estimate each one in terms of the private λ_2 and bound the error induced in these estimates by privacy. Similar bounds can be simply derived, e.g., on minimal/maximal degree, edge connectivity, etc., because their bounds are proportional to λ_2 [39].

We first recall bounds from the literature.

Lemma 2 (*Diameter and Mean Distance Bounds [40]*)**.** For an undirected, unweighted graph G of order n, define

$$\overline{d}(\lambda_2, \alpha) = \left(2\sqrt{\frac{\lambda_n}{\lambda_2}}\sqrt{\frac{\alpha^2 - 1}{4\alpha}} + 2\right)\left(\log_\alpha \frac{n}{2}\right)$$

$$\overline{\rho}(\lambda_2, \alpha) = \left(\sqrt{\frac{\lambda_n}{\lambda_2}}\sqrt{\frac{\alpha^2 - 1}{4\alpha}} + 1\right)\left(\frac{n}{n-1}\right)\left(\frac{1}{2} + \log_\alpha \frac{n}{2}\right).$$

Then for any fixed $\lambda_2 > 0$ and any $\alpha > 1$, the diameter d and mean distance ρ of the graph G are bounded via

$$\underline{d}(\lambda_2) = \frac{4}{n\lambda_2} \leq d \leq \overline{d}(\lambda_2, \alpha)$$

$$\underline{\rho}(\lambda_2) = \frac{2}{(n-1)\lambda_2} + \frac{n-2}{2(n-1)} \leq \rho \leq \overline{\rho}(\lambda_2, \alpha).$$

The least upper bounds can be derived by using the values of α_d and α_ρ that minimize $\overline{d}(\lambda_2, \alpha)$ and $\overline{\rho}(\lambda_2, \alpha)$, respectively. ■

A list of α_d and α_ρ can be found in Table 1 in [40]. To quantify the impacts of using the private λ_2 in these bounds, we next bound the expectations of d and ρ when computed with $\tilde{\lambda}_2$. These bounds use the upper incomplete gamma function $\Gamma(\cdot,\cdot)$ and the imaginary error function $\operatorname{erfi}(\cdot)$, defined as

$$\Gamma(s,x) = \int_x^\infty t^{s-1}e^{-t}dt \quad \text{and} \quad \operatorname{erfi}(x) = \frac{2}{\sqrt{\pi}}\int_0^x e^{t^2}dt.$$

We use $\tilde{d}$ and $\tilde{\rho}$, respectively, to denote the diameter and mean distance when computed with $\tilde{\lambda}_2$. Their expected values can be bounded as follows.

Theorem 4 (*Expectation bounds for d and ρ; Solution to Problem 3*)**.** For any $\lambda_2 > 0$, denote its private value by $\tilde{\lambda}_2$. Then, when bounded using $\tilde{\lambda}_2$, the expectations of the diameter, $\mathbb{E}[\tilde{d}]$, and mean distance, $\mathbb{E}[\tilde{\rho}]$, obey

$$\frac{4}{n\mathbb{E}[\tilde{\lambda}_2]} \leq \mathbb{E}[\tilde{d}] \leq \mathbb{E}[\overline{d}(\tilde{\lambda}_2, \alpha_d)] \qquad \text{and}$$

$$\frac{2}{(n-1)\mathbb{E}[\tilde{\lambda}_2]} + \frac{n-2}{2(n-1)} \leq \mathbb{E}[\tilde{\rho}] \leq \mathbb{E}[\overline{\rho}(\tilde{\lambda}_2, \alpha_\rho)],$$

where

$$\mathbb{E}[\bar{d}(\tilde{\lambda}_2, \alpha_d)] = \left[2\sqrt{\frac{\lambda_n(\alpha_d^2-1)}{4\alpha_d}}\mathbb{E}\left[\sqrt{\frac{1}{\tilde{\lambda}_2}}\right]+2\right]\left[\log_{\alpha_d}\frac{n}{2}\right],$$

$$\mathbb{E}[\bar{\rho}(\tilde{\lambda}_2, \alpha_\rho)] = \left[\sqrt{\frac{\lambda_n(\alpha_\rho^2-1)}{4\alpha_\rho}}\mathbb{E}\left[\frac{1}{\sqrt{\tilde{\lambda}_2}}\right]+1\right] \cdot \left[\frac{n}{n-1}\right]\cdot\left[\frac{1}{2}+\log_{\alpha_\rho}\frac{n}{2}\right].$$

We can compute the expectation terms with $\tilde{\lambda}_2$ via

$$\mathbb{E}\left[\frac{1}{\sqrt{\tilde{\lambda}_2}}\right] = \frac{1}{C_{\lambda_2}(b)}\frac{1}{2b}\left(\sqrt{\pi}\sqrt{b}e^{-\frac{\lambda_2}{b}}\left(\operatorname{erfi}\left(\sqrt{\frac{\lambda_2}{b}}\right)\right)\right.$$
$$\left.+\sqrt{b}e^{\frac{\lambda_2}{b}}\left(\Gamma\left(\frac{1}{2},\frac{n}{b}\right)-\Gamma\left(\frac{1}{2},\frac{\lambda_2}{b}\right)\right)\right),$$

$$\mathbb{E}[\tilde{\lambda}_2] = \frac{1}{2C_{\lambda_2}(b)}\left(2\lambda_2 + be^{-\frac{\lambda_2}{b}} - be^{-\frac{n-\lambda_2}{b}} - ne^{-\frac{n-\lambda_2}{b}}\right).$$

Proof: See [36, Appendix E]. ■

Remark 1. A larger ϵ indicates weaker privacy, and it results in a smaller value of b and a distribution of privacy noise that is more tightly concentrated about its mean. Thus, a larger ϵ implies that the expected value $E[\tilde{\lambda}_2]$ is closer to the exact, non-private λ_2, which leads to smaller disagreements in the bounds on the exact and expected values of d and ρ.

VI. Simulations

In this section, we present consensus simulation results and numerical results for the bounds on graph measurements when using the private λ_2 in graph analysis.

Consider a network of $n = 10$ agents with $\lambda_2 = 1$ and a true convergence rate of $r(t) = e^{-\lambda_2 t} = e^{-t}$. The network operator wishes to privatize λ_2 using the bounded Laplace mechanism with $\epsilon = 0.4$, $\delta = 0.05$, and $A = 1$. Solving for b with the algorithm in [21] yields $b \geq 7.39$, and selecting $b = 7.39$ ensures $(0.4, 0.05)-$differential privacy. Let $\tilde{\lambda}_2$ be the private output of the bounded Laplace mechanism. Then, for a recipient of $\tilde{\lambda}_2$, the estimated consensus convergence rate is $\tilde{r}(t) = e^{-\tilde{\lambda}_2 t}$. Let $\mathbb{P}\big(|\tilde{r}(t) - r(t)| \leq a\big)$ be the the probability of the error of the convergence rate estimate being less than a at time t. Intuitively, $\mathbb{P}\big(|\tilde{r}(t) - r(t)| \leq a\big)$ should be close to 1 as t grows.

We can lower bound $\mathbb{P}\big(|\tilde{r}(t) - r(t)| \leq a\big)$ by noting that

$$\mathbb{P}\big(|\tilde{r}(t) - r(t)| \leq a\big) = 1 - \mathbb{P}\big(|\tilde{r}(t) - r(t)| \geq a\big), \quad (2)$$

which we can use Theorem 2 to bound. To that end, Figure 1 shows how $\mathbb{P}\big(|\tilde{r}(t) - r(t)| \leq a\big)$ changes with time for $a = 0.2$ and shows 500 sample convergence rate estimates for values of λ_2 privatized with the parameters from above.

These simulations show that for a sufficiently large t, $\mathbb{P}\big(|\tilde{r}(t) - r(t)| \leq a\big)$ is close to 1 and that the times t for which this occurs are often small. This occurs because the bounded Laplace mechanism outputs $\tilde{\lambda}_2 \in [0, n]$, and thus $\tilde{r}(t) \to 0$ as $t \to \infty$ for any $\tilde{\lambda}_2$, while the true convergence rate $r(t)$ also converges to 0. These results

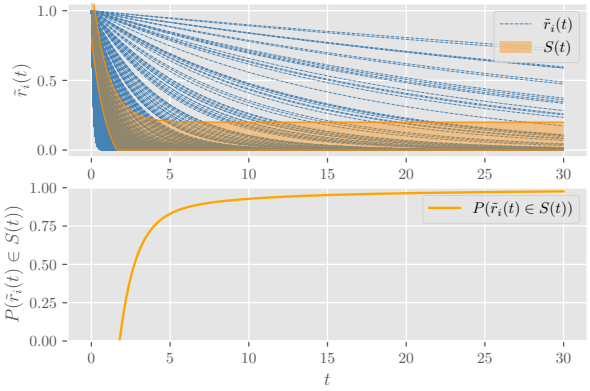

Fig. 1. The top plot shows 500 sample convergence rate estimates $\tilde{r}_i(t) = e^{-\tilde{\lambda}_2^i t}$, where $\tilde{\lambda}_2^i$ is the output of the i^{th} trial of the bounded Laplace mechanism with $\epsilon = 0.4$, $\delta = 0.05$, and $A = 1$ for a network of $n = 10$ agents with $\lambda_2 = 1$. The set $S(t)$ shown on these plots is defined as $S(t) = \{\tilde{r}_i(t) : \tilde{r}_i(t) = e^{-\tilde{\lambda}_2^i t}$ and $|\tilde{r}_i(t) - e^{-\lambda_2 t}| \leq a\}$ with $a = 0.2$. The bottom plot shows the lower bound on $\mathbb{P}(\tilde{r}_i(t) \in S)$ obtained by using Theorem 1 in (2). This lower bound approaches 1 relatively quickly and is consistent with the sample convergence rates shown in the top plot.

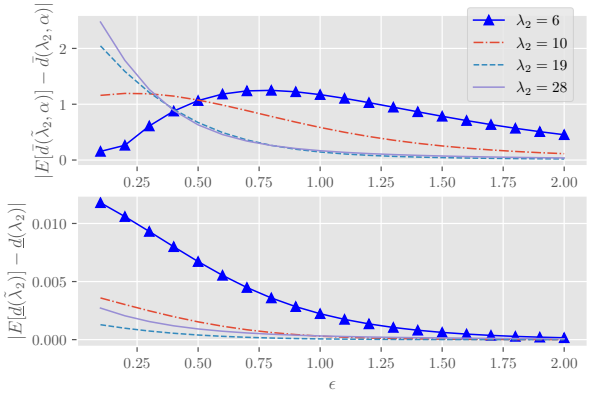

Fig. 2. The top plot shows the distance between the exact and expected upper bounds for d. The bottom plot shows the distance between the corresponding lower bounds.

also show that Theorem 1 is consistent with intuition as highlighted in Figure 1, namely that as t grows the error in any estimated convergence rate using the output of the bounded Laplace mechanism converges to 0 eventually.

We next present simulation results for using the private value of λ_2 to estimate d and ρ. We consider networks of $n = 30$ agents with different edge sets and hence different values of λ_2. We let $\lambda_n = n$ and therefore the upper bounds on d and ρ in Theorem 4 can reach their worst-case values. We apply the bounded Laplace mechanism with $\delta = 0.05$ and a range of $\epsilon \in [0.1, 2]$. To illustrate the effects of privacy in bounding diameter, we compute the distance between the exact (non-private) upper bound on diameter in Lemma 2 and the expected (private) upper bound on diameter in Theorem 4. This distance is shown in the upper plot in Figure 2, and the lower plot shows the analogous distance for the diameter lower bounds. Figure 3 shows the corresponding upper- and lower-bound distances for ρ.

In all plots, there is a general decrease in the distance between the exact and private bounds as ϵ grows. Recalling that larger ϵ implies weaker privacy, these simulations confirm

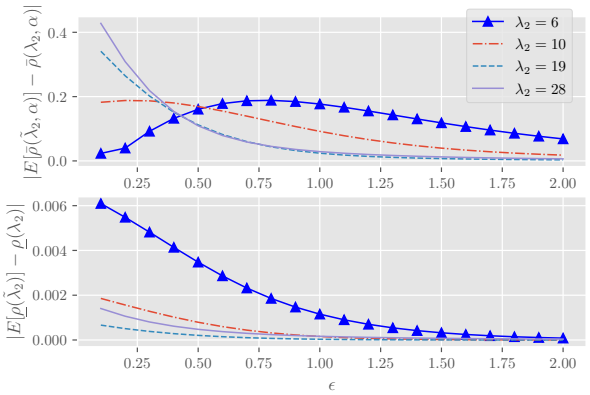

Fig. 3. The top plot shows the distance between the exact and expected upper bounds for ρ. The bottom plot shows the distance between the corresponding lower bounds.

that weaker privacy guarantees result in smaller differences between the exact and expected bounds for d and ρ, as predicted in Remark 1.

VII. Conclusions

This paper presented a differential privacy mechanism for the algebraic connectivity of undirected, unweighted graphs. Bounded noise was used to provide private values that are still accurate, and the private values of algebraic connectivity were shown to give accurate estimates of consensus protocol convergence rates, and the diameter and mean distance of a graph. Future work includes the development of new privacy mechanisms for other algebraic graph properties.

References

[1] Wei Ren, R. W. Beard, and E. M. Atkins, "A survey of consensus problems in multi-agent coordination," in *Proceedings of the 2005, American Control Conference, 2005.*, 2005.

[2] J. Scott, "Social network analysis," *Sociology*, vol. 22, no. 1, pp. 109–127, 1988.

[3] M. D. Shirley and S. P. Rushton, "The impacts of network topology on disease spread," *Eco. Complexity*, vol. 2, no. 3, pp. 287–299, 2005.

[4] Y. Zheng, L. Wang, and Y. Zhu, "Consensus of heterogeneous multi-agent systems," vol. 5, no. 16, pp. 1881–1888.

[5] P. Van Mieghem, J. Omic, and R. Kooij, "Virus spread in networks," *IEEE/ACM Transactions on Networking*, vol. 17, no. 1, pp. 1–14, 2009.

[6] S. Freitas and D. H. Chau, "Evaluating graph vulnerability and robustness using tiger," 2020.

[7] S. P. Kasiviswanathan, K. Nissim, S. Raskhodnikova, and A. Smith, "Analyzing graphs with node differential privacy," in *Proceedings of the 10th Theory of Cryptography Conference on Theory of Cryptography*. Springer-Verlag, 2013, p. 457–476.

[8] V. Karwa, S. Raskhodnikova, A. Smith, and G. Yaroslavtsev, "Private analysis of graph structure," *ACM Trans. Database Syst.*, vol. 39, no. 3, 2014.

[9] S. P. Kasiviswanathan and A. Smith, "On the 'semantics' of differential privacy: A bayesian formulation," *Journal of Privacy and Confidentiality*, vol. 6, no. 1, Jun. 2014.

[10] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," vol. 9, no. 3, pp. 211–407.

[11] W.-Y. Day, N. Li, and M. Lyu, "Publishing graph degree distribution with node differential privacy," in *Proceedings of the 2016 International Conference on Management of Data*, 2016, p. 123–138.

[12] E. Shen and T. Yu, "Mining frequent graph patterns with differential privacy," in *Proceedings of the 19th ACM International Conference on Knowledge Discovery and Data Mining*, 2013, pp. 545–553.

[13] X. Ding, X. Zhang, Z. Bao, and H. Jin, "Privacy-preserving triangle counting in large graphs," in *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. Association for Computing Machinery, 2018, p. 1283–1292.

[14] M. Hay, C. Li, G. Miklau, and D. Jensen, "Accurate estimation of the degree distribution of private networks," in *2009 Ninth IEEE International Conference on Data Mining*, 2009, pp. 169–178.

[15] C. Task and C. Clifton, "A guide to differential privacy theory in social network analysis," in *International Conference on Advances in Social Networks Analysis and Mining*, 2012, pp. 411–417.

[16] N. M. M. de Abreu, "Old and new results on algebraic connectivity of graphs," *Linear Algebra and its Applications*, vol. 423, no. 1, pp. 53–73, 2007.

[17] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Transactions on Automatic Control*, vol. 49, no. 9, pp. 1520–1533, 2004.

[18] W. Ren and E. Atkins, "Distributed multi-vehicle coordinated control via local information exchange," *International Journal of Robust and Nonlinear Control*, vol. 17, pp. 1002–1033, 2007.

[19] M. C. De Gennaro and A. Jadbabaie, "Decentralized control of connectivity for multi-agent systems," in *Proceedings of the 45th IEEE Conference on Decision and Control*, 2006, pp. 3628–3633.

[20] A. Nedić, A. Olshevsky, and W. Shi, *Decentralized Consensus Optimization and Resource Allocation*, 2018, pp. 247–287.

[21] N. Holohan, S. Antonatos, S. Braghin, and P. Mac Aonghusa, "The bounded laplace mechanism in differential privacy," *arXiv preprint arXiv:1808.10410*, 2018.

[22] K. Yazdani, A. Jones, K. Leahy, and M. Hale, "Differentially private lq control," *arXiv preprint arXiv:1807.05082*, 2018.

[23] K. Yazdani and M. Hale, "Error bounds and guidelines for privacy calibration in differentially private kalman filtering," in *2020 American Control Conference (ACC)*, 2020, pp. 4423–4428.

[24] J. Le Ny and G. J. Pappas, "Differentially private filtering," *IEEE Transactions on Automatic Control*, vol. 59, no. 2, pp. 341–354, 2013.

[25] C. Hawkins and M. Hale, "Differentially private formation control," in *2020 59th IEEE Conference on Decision and Control (CDC)*, 2020.

[26] P. Gohari, M. Hale, and U. Topcu, "Privacy-preserving policy synthesis in markov decision processes," in *2020 59th IEEE Conference on Decision and Control (CDC)*, 2020.

[27] A. Jones, K. Leahy, and M. Hale, "Towards differential privacy for symbolic systems," in *2019 American Control Conference (ACC)*, 2019, pp. 372–377.

[28] M. Hale and M. Egerstedty, "Differentially private cloud-based multi-agent optimization with constraints," in *2015 American Control Conference (ACC)*, 2015, pp. 1235–1240.

[29] D. Han, K. Liu, H. Sandberg, S. Chai, and Y. Xia, "Privacy-preserving dual averaging with arbitrary initial conditions for distributed optimization," *IEEE Transactions on Automatic Control*, pp. 1–1, 2021.

[30] P. Gohari, B. Chen, B. Wu, M. Hale, and U. Topcu, "Privacy-preserving kickstarting deep reinforcement learning with privacy-aware learners," 2021.

[31] B. Niu, Q. Li, X. Zhu, G. Cao, and H. Li, "Achieving k-anonymity in privacy-aware location-based services," in *IEEE INFOCOM 2014 - IEEE Conference on Computer Communications*, 2014, pp. 754–762.

[32] J. Le Ny, "Privacy-preserving nonlinear observer design using contraction analysis," in *2015 54th IEEE Conference on Decision and Control (CDC)*, 2015, pp. 4499–4504.

[33] Y. Wang, X. Wu, and L. Wu, "Differential privacy preserving spectral graph analysis," in *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 2013, pp. 329–340.

[34] M. Fiedler, "A property of eigenvectors of nonnegative symmetric matrices and its application to graph theory," *Czechoslovak Mathematical Journal*, vol. 25, no. 4, pp. 619–633, 1975.

[35] M. Mesbahi and M. Egerstedt, *Graph Theoretic Methods in Multiagent Networks*, 2010.

[36] B. Chen, C. Hawkins, K. Yazdani, and M. Hale, "Edge differential privacy for algebraic connectivity of graphs," *arXiv preprint arXiv:2104.00654*, 2021.

[37] P. Gohari, B. Wu, C. Hawkins, M. Hale, and U. Topcu, "Differential privacy on the unit simplex via the dirichlet mechanism," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 2326–2340, 2021.

[38] M. J. Paldino, W. Zhang, Z. D. Chu, and F. Golriz, "Metrics of brain network architecture capture the impact of disease in children with epilepsy," *NeuroImage: Clinical*, vol. 13, pp. 201–208, 2017.

[39] M. Fiedler, "Algebraic connectivity of graphs," vol. 23.

[40] B. Mohar, "Eigenvalues, diameter, and mean distance in graphs," *Graph. Comb.*, 1991.