On the Inherent Anonymity of Gossiping

Rachid Guerraoui
Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland

Anne-Marie Kermarrec
Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland

Anastasiia Kucherenko
Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland

Rafael Pinot
Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland

Sasha Voitovych[1]
University of Toronto, Canada

Abstract

Detecting the *source of a gossip* is a critical issue, related to identifying *patient zero* in an epidemic, or the *origin of a rumor* in a social network. Although it is widely acknowledged that random and local gossip communications make source identification difficult, there exists no general quantification of the level of anonymity provided to the source. This paper presents a principled method based on *ε-differential privacy* to analyze the inherent source anonymity of gossiping for a large class of graphs. First, we quantify the fundamental limit of source anonymity any gossip protocol can guarantee in an arbitrary communication graph. In particular, our result indicates that when the graph has poor connectivity, no gossip protocol can guarantee any meaningful level of differential privacy. This prompted us to further analyze graphs with controlled connectivity. We prove on these graphs that a large class of gossip protocols, namely *cobra walks*, offers tangible differential privacy guarantees to the source. In doing so, we introduce an original proof technique based on the reduction of a gossip protocol to what we call a *random walk with probabilistic die out*. This proof technique is of independent interest to the gossip community and readily extends to other protocols inherited from the security community, such as the *Dandelion* protocol. Interestingly, our tight analysis precisely captures the *trade-off* between dissemination time of a gossip protocol and its source anonymity.

2012 ACM Subject Classification Security and privacy → Privacy-preserving protocols

Keywords and phrases Gossip protocol, Source anonymity, Differential privacy

Digital Object Identifier 10.4230/LIPIcs.DISC.2023.24

Related Version *Full Version*: `https://arxiv.org/abs/2308.02477` [28]

Acknowledgements The authors are thankful to Nirupam Gupta and Pierre-Louis Roman for fruitful discussions on the early version of the paper, and to the anonymous reviewers of DISC 2023 for their constructive comments.

1 Introduction

A gossip protocol (a.k.a., an epidemic protocol) is a distributed algorithm that disseminates information in a peer-to-peer system [47, 1, 34, 38, 19, 24]. Gossip protocols have been long used to model the propagation of infectious diseases [30, 37, 3], as well as rumors in social networks where users randomly exchange messages [17, 26]. It is commonly accepted that random and local communications between the users make source identification hard, and thus

[1] Part of the work was done when Sasha Voitovych was an intern at EPFL as part of the EPFL Excellence Research Internship Program.

provide *inherent* anonymity to the source of the gossip, i.e., anonymity that comes solely from the spreading dynamic without relying on any additional cryptographic primitives (as in [42]). Source anonymity in gossip protocols constitutes an active area of research. On the one hand, many works aim to establish *privacy guarantees* for the source of the gossip by concealing it against an adversary, e.g., hiding the whistleblower on social media [27, 25, 23, 26, 7, 22]. On the other hand, a large effort is put towards identifying *privacy limits* for the source of a gossip by designing adversarial strategies that accurately recover the source, e.g., "patient zero" identification in epidemics [33, 54, 46, 49, 9, 41].

Although a significant amount of research is dedicated to the investigation of source anonymity, existing approaches (as summarized in [33]) mainly focus on specific settings, such as locating the source of a gossip for a particular protocol, hiding it against a chosen adversarial strategy or examining the problem on a narrow family of graphs (trees, complete graphs, etc.). This prevents the results from being generalized, and it remains unclear how hard it is to recover the source of a gossip in general, naturally raising the following question.

What are the fundamental limits and guarantees on the inherent source anonymity of gossiping in a general setting?

We take an important step towards addressing this question by adapting the celebrated mathematical framework of ε-differential privacy (ε-DP) to our context [20, 21]. Although the concept is a gold standard to measure privacy leakage from queries on tabular databases, it can be also adapted to different privacy semantics and threat models [15]. In our context, we use ε-DP to measure the *inherent* source anonymity of gossiping in general graphs. We adopt a widely used threat model where the adversary aims to guess the source by monitoring the communications of a set of *curious* nodes in the graph [33, 46, 48, 54, 16, 23]. Using differential privacy enables us to overcome the limitations of previous work, as DP guarantees hold regardless of the exact strategy of the attacking adversary. Additionally, DP guarantees can be combined with any prior knowledge the adversary has on the location of the source, making our results generalizable. Our contributions can be summarized as follows.

1.1 Main results

We propose a mathematical framework that adapts the concept of differential privacy to quantify source anonymity in any graph (Section 3). In doing so, we highlight the importance of considering two types of adversaries: the *worst-case* and the *average-case*. For the worst-case adversary, we focus on privacy guarantees that hold *regardless* of the location of the curious nodes in the graph. In other words, these guarantees hold even if the adversary knows the communication graph in advance and chooses curious nodes strategically. For the average-case adversary, we focus on privacy guarantees that hold with high probability when curious nodes are chosen uniformly at random. Here, the adversary does not know the structure of the underlying communication graph in advance. Within our mathematical framework, we establish the following results for both adversarial cases.

Privacy limits. We first quantify a fundamental limit on the level of ε-DP any gossip protocol can provide on any graph topology (Section 4). This result indicates that no gossip protocol can ensure any level of differential privacy on poorly connected graphs. This motivates us to consider graphs with controlled connectivity, namely expander graphs. Expanders are an important family of strongly connected graphs that are commonly considered in the gossip protocols literature [8, 29, 11]. On this class, we get the following results.

Privacy guarantees. We prove that a large class of gossip protocols provides tangible differential privacy guarantees to the source (Section 5). We first consider the parameterized family of gossip protocols known as $(1+\rho)$-cobra walks [18, 11, 45, 6], which constitutes a natural generalization of a simple random walk. A cobra walk can be seen as an SIS (Susceptible-Infected-Susceptible) epidemic, a well-established model for analyzing the spread of epidemics and viruses in computer networks [30, 37]. In particular, a $(1+\rho)$-cobra walk is an instance of an SIS epidemic scheme where active nodes constitute the infectious set, the duration of the infectious phase is equal to one and every infected node can only infect one or two of its neighbors at a time. In order to establish differential privacy guarantees on this class of gossip protocols, we rely on the critical observation that the cobra walk has a quantifiable probability of mixing before hitting a curious node (see Section 1.2 for more details on this observation). This characteristic is not unique to cobra walks, as it is shared by several other types of gossip protocols. Accordingly, we also show how to generalize our privacy guarantees to the ρ-Dandelion protocol [7], first introduced as an anonymous communication scheme for Blockchains.

Dissemination time vs. privacy trade-off. As an important by-product of our analysis, we precisely capture the trade-off between dissemination time and privacy of a large class of gossip protocols operating on sufficiently dense graphs we call near-Ramanujan graphs (Section 7). The privacy-latency tension has been suggested several times in the literature [7, 5, 32]. However, our work presents the first formal proof of this long-standing empirical observation. Specifically, we show that our privacy guarantees are tight for both $(1+\rho)$-cobra walks [11] and ρ-Dandelion protocol [7]. Additionally, we give a tight analysis of the dissemination time as a function of parameter ρ. This analysis leads us to conclude that increasing parameter ρ results in a faster dissemination, but decreases privacy guarantees of the protocol, formally establishing the existence of a trade-off between privacy and dissemination time. As cobra walks are strongly related to SIS-epidemics, and Dandelion to anonymous protocols in peer-to-peer networks, our results are relevant for both epidemic and source anonymity communities.

1.2 Technical challenges & proof techniques

A major technical contribution of our paper is the privacy guarantee of $(1+\rho)$-cobra walks in non-complete graphs. The derivation of this result has been challenging to achieve for two reasons. Firstly, our objective is to establish differential privacy guarantees in general graphs, which is a more complex scenario than that of complete graphs (as seen in [5]), where any communication between pairs of nodes is equiprobable, and symmetry arguments can be utilized. Yet, this technique is no longer applicable to our work. The fact that no symmetry assumptions about graph structure can be made calls for new more sophisticated proof techniques. Second, cobra walks are challenging to analyze directly. State-of-the-art approaches analyzing the dissemination time of cobra walks circumvent this issue by analyzing a dual process instead, called BIPS [11, 12, 6]. There, the main idea is to leverage the duality of BIPS and cobra walks with respect to hitting times [11]. While hitting times provide sufficient information for analyzing the dissemination time of a cobra walk, they cannot be used to evaluate differential privacy, as they do not provide sufficient information about the probability distribution of the dissemination process. We overcome this difficulty through a two-step proof technique, described below.

Step I: Reduction to a random walk with probabilistic die out. To establish ε-differential privacy, we essentially show that two executions of the same $(1+\rho)$-cobra walk that started from different sources are statistically indistinguishable to an adversary monitoring a set of curious nodes. In doing so, we design a novel proof technique that involves reducing the analysis of gossip dissemination in the presence of curious nodes, to a *random walk with probabilistic die out*. Such a protocol behaves as a simple random walk on the communication graph G, but it is killed at each step (i) if it hits a curious node, or otherwise (ii) with probability ρ. We show that disclosing the death site of such a random walk to the adversary results in a bigger privacy loss than all the observations reported by the curious nodes during the gossip dissemination. Then, we can reduce the privacy analysis of cobra walks to the study of such a random walk with probabilistic die out.

Step II: Analysis of a random walk with probabilistic die out. To study a random walk with probabilistic die out, we characterize the spectral properties of the (scaled) adjacency matrix $\boldsymbol{Q}$ corresponding to the subgraph of G induced by the non-curious nodes. In particular, we show that if curious nodes occupy a small part of every neighborhood in G, then the subgraph induced by non-curious nodes (i) is also an expander graph and (ii) has an almost-uniform first eigenvector. While (i) is a direct consequence of the Cauchy Interlacing Theorem, (ii) is more challenging to obtain. We need to bound $\boldsymbol{Q}$ from above and below by carefully designed matrices with an explicit first eigenvector. Combining (i) and (ii) allows us to precisely estimate the behavior of the random walk with probabilistic die out, which yields the desired differential privacy guarantees.

Generality of the proof. The reduction to a random walk with probabilistic die out is the most critical step of our proof. It is general and allows us to analyze several other protocols without having to modify the most technical part of the proof (Step II above). We demonstrate the generality of this technique by applying this reduction to the Dandelion protocol and obtain similar privacy guarantees to cobra walks.

1.3 Related work

Inherent anonymity of gossiping. To the best of our knowledge, only two previous works have attempted to quantify the inherent source anonymity of gossiping through differential privacy [5, 32]. The former work [5] is the first to analyze source anonymity using differential privacy. It measures the guarantees of a class of gossip protocols with a muting parameter (which we call "muting push" protocols) and contrasts these guarantees with the dissemination time of these protocols on a complete graph. Both the threat model and the nature of the technical results in [5] heavily depend on the completeness of the graph. In such a context, the analysis is considerably simplified for two reasons. Firstly, the presence of symmetry allows for the curious node locations to be ignored, rendering the average-case and the worst-case adversaries equivalent. Secondly, in contrast to what would happen in non-complete graphs, since any node can communicate with any other node in each round, a single round of communication is sufficient to hide the identity of the source. However, when considering the spread of epidemics or the propagation of information in social networks, communication graphs are seldom complete [43]. Our work highlights that non-completeness of the graph potentially challenges the differential privacy guarantees that gossip protocols can achieve and also makes it important to distinguish between average and worst-case threat models. Therefore, our results constitute a step toward a finer-grained analysis of the anonymity of gossiping in general graphs. Note that our work can be seen as a strict generalization of the

results of [5], since, in addition to cobra walks and Dandelion, we also show that our proof techniques described in Section 1.2 apply to "muting push" protocols (see Appendix E of the full version of the paper [28]).

The second approach [32] addresses a problem that appears to be similar to ours at first glance, as it aims to quantify source anonymity in non-complete graphs. However, the authors consider a different threat model, where an adversary can witness any communication with some probability instead of only those passing through the curious nodes. Furthermore, the paper only gives negative results and does not provide any differential privacy guarantees, which is the most technically challenging part of our paper.

Dissemination time vs. privacy trade-off. Several previous works [53, 4, 14, 51] have suggested the existence of a tension between source anonymity (i.e., privacy) and latency of message propagation. Under the threat model we consider in this work (with curious nodes), [7] conjectured that the Dandelion protocol would exhibit a trade-off between (their definition of) source anonymity and dissemination time. Later, works [5] and [32] provided more tangible evidence for the existence of a dissemination time vs. privacy trade-off when analyzing source anonymity through differential privacy. However, these works do not provide a tight analysis of the tension between dissemination time and privacy, hence making their observation incomplete. To the best of our knowledge, our work is the first to rigorously demonstrate the existence of a trade-off between the dissemination time of a gossip protocol and the privacy of its source thanks to the *tightness* of our analysis.

2 Preliminaries

For a vector $\boldsymbol{x} \in \mathbb{R}^m$, we denote by x_i its ith coordinate, i.e., $\boldsymbol{x} = (x_1, x_2, \ldots, x_m)^\top$. Similarly, for a matrix $\boldsymbol{M} \in \mathbb{R}^{m \times m'}$, we denote by M_{ij} its entry for the ith row and jth column. Furthermore, for any symmetric matrix $\boldsymbol{M} \in \mathbb{R}^{m \times m}$, we denote by $\lambda_1(\boldsymbol{M}) \geq \lambda_2(\boldsymbol{M}) \geq \ldots \geq \lambda_m(\boldsymbol{M})$ its eigenvalues. We use $\mathbf{1}_m \in \mathbb{R}^m$ to denote an all-one vector, $\boldsymbol{I}_m \in \mathbb{R}^{m \times m}$ to denote the identity matrix, $\boldsymbol{J}_m \in \mathbb{R}^{m \times m}$ to denote an all-one square matrix, and $\boldsymbol{O}_{m \times m'} \in \mathbb{R}^{m \times m'}$ to denote an all-zero matrix. Finally, for any $\boldsymbol{x} \in \mathbb{R}^m$, we denote by $||\boldsymbol{x}||_p \triangleq (\sum_{i=1}^m |x_i|^p)^{1/p}$ the ℓ_p norm of $\boldsymbol{x}$ for $p \in [1, \infty)$ and by $||\boldsymbol{x}||_\infty \triangleq \max_{i \in m} |x_i|$ the ℓ_∞ norm of $\boldsymbol{x}$.

Throughout the paper, we use the *maximum divergence* to measure similarities between probability distributions. We consider below a common measurable space (Ω, Σ) on which the probability measures are defined. Let μ, ν be two probability measures over Σ. The *max divergence* between μ and ν is defined as[2]

$$D_\infty(\mu \parallel \nu) \triangleq \sup_{\sigma \in \Sigma,\ \mu(\sigma) > 0} \ln \frac{\mu(\sigma)}{\nu(\sigma)}.$$

Furthermore, for two random variables X, Y with laws μ and ν respectively, we use the notation $D_\infty(X \parallel Y)$ to denote $D_\infty(\mu \parallel \nu)$.

[2] Note that we allow $\nu(\sigma) = 0$ in the definition. If $\nu(\sigma) = 0$ but $\mu(\sigma) > 0$ for some $\sigma \in \Sigma$, the max divergence is set to ∞ by convention.

2.1 Graph theoretical terminology

Consider an undirected connected graph $G = (V, E)$, where V is the set of nodes and E is the set of edges. G cannot have self-loops or multiple edges. For any $v \in V$, we denote by $N(v)$ the set containing the neighbours of v in G and by $\deg(v)$ the number of edges incident to v. Furthermore, G is said to be a *regular graph*, if there exists $d(G)$ such that $\deg(v) = d(G)$ for every $v \in V$; $d(G)$ is called the degree of the graph. Additionally, for a set $U \subseteq V$ and $v \in V$, we denote by $\deg_U(v)$ the number of neighbours of v contained in U, i.e., $\deg_U(v) = |N(v) \cap U|$. Below, we introduce some additional graph terminology.

▶ **Definition 1** (Vertex cut & connectivity). *A* vertex cut *of G is a subset of vertices $K \subseteq V$ whose removal disconnects G or leaves just one vertex. A minimum vertex cut of G is a vertex cut of the smallest size. The size of a minimum vertex cut for G, denoted $\kappa(G)$, is called the vertex connectivity of G.*

Consider an undirected connected graph $G = (V, E)$ of size n where V is an ordered set of nodes. We denote by $\boldsymbol{A}$ the adjacency matrix of G, i.e., $A_{vu} = 1$ if $\{v, u\} \in E$ and $A_{vu} = 0$ otherwise. We also denote by $\hat{\boldsymbol{A}} = \boldsymbol{D}^{-1/2}\boldsymbol{A}\boldsymbol{D}^{-1/2}$ the normalized adjacency matrix of G, where $\boldsymbol{D}$ is the diagonal degree matrix, i.e., $D_{vu} = \deg(v)$ if $v = u$ and 0 otherwise. Since $\hat{\boldsymbol{A}}$ is a symmetric and normalized matrix, the eigenvalues of $\hat{\boldsymbol{A}}$ are real valued and $\lambda_1(\hat{\boldsymbol{A}}) = 1$. Using this terminology, the *spectral expansion* of G is defined as

$$\lambda(G) \triangleq \max\{|\lambda_2(\hat{\boldsymbol{A}})|, |\lambda_n(\hat{\boldsymbol{A}})|\}. \tag{1}$$

▶ **Definition 2** (Expander graph). *Consider an undirected regular graph G. If $d(G) = d$ and $\lambda(G) \leq \lambda$, then G is said to be a (d, λ)-expander graph.*

2.2 Gossip protocols

Consider an undirected connected communication graph $G = (V, E)$ where two nodes $u, v \in V$ can directly communicate if and only if $\{u, v\} \in E$. One node $s \in V$, called the *source*, holds a unique gossip g to be propagated throughout the graph. In this context, *a gossip protocol* is a predefined set of rules that orchestrates the behavior of the nodes with regard to the propagation of g. Essentially, the goal of a protocol is that with probability 1 every node in G eventually receives g. We assume discrete time steps and synchronous communication, i.e., the executions proceed in rounds of one time step.[3] While every node in G has access to the global clock, we assume that the execution of the protocol starts at a time $t_\star \in \mathbb{Z}$, which is *only* known to the source s.

Execution of a gossip protocol. At any point of the execution of the protocol, a node $u \in V$ can either be active or non-active. Only active nodes are allowed to send messages during the round. A gossip protocol always starts with the source s being the only active node, and at every given round $t + 1$ active nodes are the nodes that received the gossip at round t. We will use $X_t \subseteq V$ to denote the set of active nodes at the beginning of round $t \geq t_\star$ and set $X_{t_\star} = \{s\}$ by convention. Denoting by $(u \to v)$ a communication between nodes u and v, we define $\mathcal{C}$ to be the set of all possible communications in G, i.e., $\mathcal{C} = \{(u \to v) : \{u, v\} \in E\} \cup \{(u \to u) : u \in V\}$. Note that we allow an active node

[3] Although, for clarity, we focus on a synchronous communication, our analysis of privacy guarantees in Section 5 readily extends to an asynchronous setting.

u to send a fictitious message to itself to stay active in the next communication round. Then, the t^{th} round of an execution for a given protocol $\mathcal{P}$ can be described by a pair (X_t, C_t), where $X_t \subseteq V$ is a set of active nodes, and C_t is the (multi)set of communications of $\mathcal{C}$ which happened at round t. We denote by S the random variable characterizing the *execution* of the protocol. Naturally, an *execution* is described by a sequence of rounds, i.e., $S = \{(X_t, C_t)\}_{t \geq t_\star}$. We define *expected dissemination time* of the protocol as the expected number of rounds for all nodes to receive the gossip during an execution. Finally, we denote $\mathcal{E}$ the set of all possible executions.

Cobra and random walk. Coalescing-branching random walk protocol (a.k.a., cobra walk) [18, 11, 45, 6] is a natural generalization of a simple random walk that is notably useful to model and understand Susceptible-Infected-Susceptible (SIS) epidemic scheme [30, 37]. We consider a $(1+\rho)$-cobra walk as studied in [11] with $\rho \in [0,1]$[4]. This is a gossip protocol where at every round $t \geq t_\star$, each node $u \in X_t$ samples a token from a Bernoulli distribution with parameter ρ. If the token equals zero, u samples uniformly at random a node v from its neighbors $N(u)$ and communicates the gossip to it, i.e., $(u \to v)$ is added to C_t. If the token equals one, the protocol *branches*. Specifically, u independently samples two nodes v_1 and v_2 at random (with replacement) from its neighbors and communicates the gossip to both of them, i.e., $(u \to v_1)$, and $(u \to v_2)$ are added to C_t. At the end of the round, each node $u \in X_t$ deactivates. Note that, when $\rho = 0$, this protocol degenerates into a simple random walk on the graph; hence it has a natural connection with this random process.

Dandelion protocol. Dandelion is a gossip protocol designed to enhance source anonymity in the Bitcoin peer-to-peer network. Since it was introduced in [7], it has received a lot of attention from the cryptocurrency community. Dandelion consists of two phases: (i) the anonymity phase, and (ii) the spreading phase. The protocol is parameterized by $\rho \in [0,1)$, the probability of transitioning from the anonymity phase to the spreading phase. Specifically, the phase of the protocol is characterized by a token $anonPhase \in \{0,1\}$ held by a global oracle and initially equal to 0. At the beginning of each round of the Dandelion execution, if $anonPhase = 1$ the global oracle sets $anonPhase = 0$ with probability ρ and keeps $anonPhase = 1$ with probability $1-\rho$. Once $anonPhase = 0$, the global oracle stops updating the token. Based on this global token, at each round, active nodes behave as follows. If the $anonPhase = 1$, the execution is in the anonymity phase and an active node u samples a node v uniformly at random from its neighborhood $N(u)$ and communicates the gossip to it, i.e., $(u \to v)$ is added to C_t. Afterwards, node u deactivates, i.e., in the anonymity phase only one node is active in each round. If the $anonPhase = 0$, the execution is in the spreading phase. Then the gossip is broadcast, i.e., each node $u \in X_t$ communicates the gossip to all of its neighbors and for $\forall v \in N(u)$, $(u \to v)$ is added to C_t.

[4] Some prior works also study k-cobra walks with branching parameter $k \geq 3$ [18]. We do not consider this class, since our negative result for a 2-cobra walk (Theorem 34 in the full version of the paper [28]) implies that a k-cobra walk for any $k \geq 3$ does not satisfy a reasonable level of differential privacy.

3 Mathematical framework for source anonymity in general graphs

Given a source and a gossip protocol, we fix the probability space $(\mathcal{E}, \Sigma, \mathbb{P})$, where Σ is the standard cylindrical σ-algebra on $\mathcal{E}$ (as defined in Appendix A.1 of [55]) and $\mathbb{P}$ is a probability measure characterizing the executions of the protocol. In the remaining, to avoid measurability issues, we only refer to subsets of $\mathcal{E}$ from Σ.

3.1 Measuring source anonymity with differential privacy

We now describe the mathematical framework we use to quantify source anonymity of gossiping. We consider a threat model where an external adversary has access to a subset $F \subset V$ of size $f < n-1$ of *curious* nodes. Curious nodes in F execute the protocol correctly, but report their communications to the adversary. The adversary aims to identify the source of the gossip using this information. We distinguish two types of adversaries, namely worst-case and average-case, depending on the auxiliary information they have on the graph.

Threat models: worst-case and average-case adversaries. On the one hand, a *worst-case* adversary is aware of the structure of the graph G and may choose the set of curious nodes to its benefit. On the other hand, the *average-case* adversary is not aware of the topology of G before the start of the dissemination, hence the set of curious nodes is chosen uniformly at random among all subsets of V of size f. We assume that the messages shared in the network are unsigned and are passed unencrypted. Also, the contents of transmitted messages (containing the gossip) do not help to identify the source of the gossip. In other words, adversaries can only use the information they have on the dissemination of the gossip through the graph to locate the source. We also assume that the adversary does not know the exact starting time $t_\star \in \mathbb{Z}$ of the dissemination. To formalize the observation received by the external adversary given a set of curious nodes F, we introduce a function $\Psi^{(F)}$ that takes as input communications C from a single round and outputs only the communications of C visible to the adversary. Note that a communication $(v \to u)$ is visible to the adversary if and only if either v or u belongs to F. Consider an execution $S = \{(X_t, C_t)\}_{t \geq t_\star}$ of a gossip protocol, and denote by t_{ADV} the first round in which one of the curious nodes received the gossip. Then we denote by $S_{\text{ADV}} = \{\Psi^{(F)}(C_t)\}_{t \geq t_{\text{ADV}}}$ the random variable characterizing the observation of the adversary for the whole execution. Note that the adversary does not know $t_\star$, hence it cannot estimate how much time passed between $t_\star$ and t_{ADV}.

▶ Remark 3. For Dandelion, the adversary actually also has access to the value of $anonPhase$ in round t, i.e., we have $S_{\text{ADV}} = \{\Psi^{(F)}(C_t), anonPhase_t\}_{t \geq t_{\text{ADV}}}$. We omit this detail from the main part of the paper for simplicity of presentation, but it does not challenge our results on privacy guarantees. See Appendix C.4 in the full version of the paper [28] for more details.

Measuring source anonymity. We formalize source anonymity below by adapting the well-established definition of differential privacy. In the remaining of the paper, for a random variable A, we will write $A^{(s)}$ to denote this random variable conditioned on the node $s \in V \setminus F$ being the source. In our setting, we say that a gossip protocol satisfies differential privacy if for any $u, v \in V$ the random sequences $S^{(v)}_{\text{ADV}}$ and $S^{(u)}_{\text{ADV}}$ are statistically indistinguishable. More formally, we define differential privacy as follows.

▶ **Definition 4** (Differential privacy). *Consider an undirected graph $G = (V, E)$ and a set of curious nodes $F \subset V$. Then, a gossip protocol satisfies* ε-differential privacy *(ε-DP) for the set F if, for any two nodes $v, u \in V \setminus F$, the following holds true*

$$D_\infty \left(S^{(v)}_{ADV} \parallel S^{(u)}_{ADV} \right) \leq \varepsilon.$$

When establishing differential privacy guarantees against a *worst-case adversary*, we aim to find a value ε which only depends on the number of curious nodes f, and is *independent* of the identity of the nodes in F. Accordingly, we say that a gossip protocol satisfies *ε-DP against a worst-case adversary* if it satisfies ε-DP for any set $F \subset V$ such that $|F| = f$.

When establishing differential privacy against an *average-case adversary*, we aim to find a value of ε for which the protocol satisfies ε-DP *with high probability*[5] when choosing the f curious nodes uniformly at random from V. Formally, let $\mathcal{U}_f(V)$ be the uniform distribution over all subsets of V of size f, a gossip protocol satisfies *ε-DP against an average-case adversary* if

$$\mathbb{P}_{F\sim\mathcal{U}_f(V)}\left[\max_{v,u\in V\setminus F} D_\infty\left(S^{(v)}_{\text{ADV}} \parallel S^{(u)}_{\text{ADV}}\right) \leq \varepsilon\right] \geq 1 - \frac{1}{n}. \tag{2}$$

3.2 Semantic of source anonymity

Differential privacy is considered the gold standard definition of privacy, since ε-DP guarantees hold *regardless* of the strategy of the adversary and any prior knowledge it may have on the location of the source. Yet, the values of ε are notoriously hard to interpret [39, 31]. To better understand the semantic of our definition of differential privacy, we consider below two simple examples of adversarial strategies: maximum a posteriori and maximum likelihood estimations. For these strategies, we derive bounds on the probability of an adversary successfully guessing the source in an effort to give a reader an intuition on the meaning of the parameter ε. The proofs are given in Appendix F of the full version of the paper [28].

Maximum a posteriori strategy. Maximum a posteriori (MAP) strategy can be described as follows. Suppose an adversary has an a priori distribution p that assigns to every node in $V\setminus F$ a probability of being the source of the gossip. Intuitively, p corresponds to the set of beliefs the adversary has on the origin of the gossip before observing the dissemination. This prior might reflect information acquired from any auxiliary authority or some expert knowledge on the nature of the protocol. Suppose the adversary observes an event σ. Then, a MAP-based adversary "guesses" which node is the most likely to be the source, assuming event σ occurred and assuming the source has been sampled from the prior distribution p. Such guess is given by

$$\hat{s}_{MAP} = \operatorname*{argmax}_{v\in V\setminus F} \mathbb{P}_{s\sim p}\left[v = s \mid S^{(s)}_{\text{ADV}} \in \sigma\right] = \operatorname*{argmax}_{v\in V\setminus F} \mathbb{P}\left[S^{(v)}_{\text{ADV}} \in \sigma\right] p(v). \tag{3}$$

Using ε-DP, we can upper bound the success probability of such a guess. Suppose the protocol satisfies ε-DP, then the probability of correctly identifying a source $s \sim p$ conditioned on σ happening is upper bounded as follows

$$\mathbb{P}_{s\sim p}\left[\hat{s}_{MAP} = s \mid S^{(s)}_{\text{ADV}} \in \sigma\right] \leq \exp(\varepsilon) p\left(\hat{s}_{MAP}\right). \tag{4}$$

Such an upper bound has a simple interpretation. Note that $p(\hat{s}_{MAP})$ characterizes the maximum probability of a successfully guessing $\hat{s}_{MAP}$ based solely on adversary's prior knowledge. Then, the upper bound above states that the probability of a successful guess after observing the dissemination is amplified by a factor of at most $\exp(\varepsilon)$ compared to success probability of a guess based on a priori knowledge only.

[5] An event is said to hold with high probability on graph G of size n, if it holds with probability $\geq 1 - 1/n$.

Maximum likelihood strategy. Maximum likelihood estimation (MLE) occupies a prominent place [23, 49, 50, 46] in the literature, both for designing source location attacks, and for defending against adversaries that follow an MLE strategy. This method is a special instance of MAP estimator in (3) with a uniform prior distribution $p = \mathcal{U}(V \setminus F)$ on the source. We can show that, if the protocol satisfies ε-DP, such guess has a bounded success probability.

$$\mathbb{P}_{s\sim\mathcal{U}(V\setminus F)}\left[\hat{s}_{MLE} = s \mid S^{(s)}_{\text{ADV}} \in \sigma\right] \leq \frac{\exp(\varepsilon)}{n-f}. \tag{5}$$

4 Fundamental limits of source anonymity: lower bound on ε

We start by studying the fundamental limits of differential privacy in general graphs. Specifically, we aim to show that vertex connectivity constitutes a hard threshold on the level of source anonymity gossiping can provide. First, we present a warm-up example indicating that in a poorly connected graph, no gossip protocol can achieve any meaningful level of differential privacy against a worst-case adversary. We then validate this intuition by devising a universal lower bound on ε that applies for any gossip protocol and any undirected connected graph. Complete proofs related to this section can be found in Appendix B of the full version of the paper [28].

4.1 Warm-up

Consider a non-complete graph $G = (V, E)$ and $K \subset V$, a vertex cut of G. Then, by definition, deleting K from G partitions the graph into two disconnected subgraphs. When $f \geq |K|$, a worst-case adversary can take F such that $K \subseteq F$. Then, the curious nodes can witness all the communications that pass from one subgraph to the other. Intuitively, this means that any two nodes that are not in the same subgraph are easily distinguishable by the adversary. Hence, differential privacy cannot be satisfied. This indicates that the level of differential privacy any gossip protocol can provide in a general graph fundamentally depends on the connectivity of this graph. To validate this first observation and determine the fundamental limits of gossiping in terms of source anonymity, we now determine a lower bound on ε.

4.2 Universal lower bound on ε

We present, in Theorem 5, a universal lower bound on ε which holds for any gossip protocol, on any connected graph and for both the worst-case and the average-case adversaries.

▶ **Theorem 5.** *Consider an undirected connected graph $G = (V, E)$ of size n, a number of curious nodes $f > 1$, and an arbitrary gossip protocol $\mathcal{P}$. If $\mathcal{P}$ satisfies ε-DP against an average-case or a worst-case adversary, then*

$$\varepsilon \geq \ln(f - 1).$$

Moreover, if $\kappa(G) \leq f$, then $\mathcal{P}$ cannot satisfy ε-DP with $\varepsilon < \infty$ against a worst-case adversary.

Proof sketch. To establish the above lower bound, we assume that the adversary simply predicts that the first non-curious node to contact the curious set is the source of the gossip. As the definition of differential privacy does not assume a priori knowledge of the adversarial strategy, computing the probability of success for this attack provides a lower bound on ε.

We first demonstrate the result for the average-case adversary. Assume that F is sampled uniformly at random from V. We can show that there exists $v \in V$ such that the attack implemented by the adversary succeeds with large enough probability when v is the source of the gossip. This fact essentially means that this v is easily distinguishable from any other node in the graph, which yields the lower bound $\varepsilon \geq \ln(f-1)$ in the average case. We now consider the worst-case adversary. Assume that F can be chosen by the adversary. As the lower bound $\varepsilon \geq \ln(f-1)$ holds with positive probability when F is chosen at random, there exists at least one set F for which it holds. Choosing this set of curious nodes establishes the claim for the worst-case adversary. Furthermore, when $\kappa(G) \leq f$, we follow the intuition from Section 4.1 to build a set F that disconnects the graph. Using this set, we prove that ε cannot be finite. ◀

Theorem 5 shows that the connectivity of the graph is an essential bottleneck for differential privacy in a non-complete graph. This stipulates us to study graphs with controlled connectivity, namely (d, λ)-expander graphs. Note that in a (d, λ)-expander, the vertex connectivity does not exceed d. Hence, Theorem 5 implies that no gossip protocol can satisfy any meaningful level of differential privacy against a worst-case adversary on a (d, λ)-expander if $f \geq d$. Considering this constraint, while studying a gossip against a worst-case adversary, we only focus on cases where the communication graph G has a large enough degree d.

5 Privacy guarantees: upper bound on ε

We now present a general upper bound on ε that both holds for $(1+\rho)$-cobra walks and ρ-Dandelion on d-regular graphs with fixed expansion, i.e., (d, λ)-expander graphs. Complete proofs related to this section can be found in Appendix C of the full version of the paper [28]. Our privacy guarantees are quite technical, which is justified by the intricacies of the non-completeness of the graph. Recall that, in the case of complete topologies analyzed in [5], after one round of dissemination all information on the source is lost unless a curious node has been contacted. However, in a general expander graph, this property does not hold anymore. Indeed, even after multiple rounds of propagation, the active set of the protocol can include nodes that are close to the location of the source s. Thus, differential privacy may be compromised.

5.1 Adversarial density

The attainable level of source anonymity for a given protocol is largely influenced by the location of curious nodes. However, accounting for all possible placements of curious nodes is a very challenging and intricate task. To overcome this issue and state our main result, we first introduce the notion of *adversarial density* that measures the maximal fraction of curious nodes that any non-curious node may have in its neighborhood. Upper bounding the adversarial density of a graph is a key element to quantifying the differential privacy guarantees of a gossip protocol. Formally, this notion is defined as follows.

▶ **Definition 6.** *Consider an undirected connected d-regular graph $G = (V, E)$, and an arbitrary set of curious nodes $F \subseteq V$. The* adversarial density *of F in G, denoted α_F, is the maximal fraction of curious nodes that any node $v \in V \setminus F$ has in its neighborhood. Specifically,*

$$\alpha_F \triangleq \max_{v \in V \setminus F} \frac{\deg_F(v)}{d}.$$

For any set of curious nodes F, we have $\alpha_F \leq f/d$. Hence, even when F is chosen by a worst-case adversary, the adversarial density is always upper bounded by f/d. However, for the average-case adversary we can obtain a much tighter bound, stated in Lemma 7 below.

▶ **Lemma 7.** *Consider an undirected connected d-regular graph $G = (V, E)$ of size n and a set of curious nodes $F \sim \mathcal{U}_f(V)$, with adversarial density α_F. We denote $\beta = f/n$ and $\gamma = \ln(n)/(ed)$, where e is Euler's constant. Then, with probability at least $1 - 1/n$, $\alpha_F \leq \alpha$ with*

$$\alpha \leq 4e\frac{\max\{\gamma, \beta\}}{1 + \max\{\ln(\gamma) - \ln(\beta), 0\}}.$$

Furthermore, if there exist $\delta > 0, c > 0$ such that $f/n > c$ and $d > \ln(n)/(c^2\delta^2)$ then a similar statement holds with $\alpha \leq (1 + \delta)\beta$.

We deliberately state this first lemma in a very general form. This allows us to precisely quantify how the upper bound on the adversarial density improves as f decreases. To make this dependency clearer, we provide special cases in which the bound on α_F is easily interpreted. First, assume that $d \in \omega_n(\log(n))$ and $f/n \in \Omega_n(1)$. Then, α_F is highly concentrated around f/n, up to a negligible multiplicative constant, when n is large enough. On the other hand, when the ratio f/n becomes subconstant, the concentration becomes looser. In particular, if $d \in \omega_n(\log(n))$ and $f/n \in o_n(1)$, then $\alpha_F \in o_n(1)$ with high probability. Finally, if f/n drops even lower (e.g., when $f/n \in n^{-\Omega_n(1)}$), we get $\alpha_F \in O_n(1/d)$ or $\alpha_F \in n^{-\Omega_n(1)}$ with high probability for any d.

5.2 General upper bound on ε

Thanks to Lemma 7 bounding adversarial density, we can now state our main theorem providing a general upper bound on ε for $(1 + \rho)$-cobra walks and ρ-Dandelion.

▶ **Theorem 8.** *Consider an undirected connected (d, λ)-expander graph $G = (V, E)$ of size n, let f be the number of curious nodes, and let $\mathcal{P}$ be a $(1 + \rho)$-cobra walk with $\rho < 1$. Set $\alpha = f/d$ (resp. set α as in Lemma 7). If $\lambda < 1 - \alpha$, then $\mathcal{P}$ satisfies ε-DP against a worst-case adversary (resp. an average-case adversary) with*

$$\varepsilon = \ln(\rho(n - f) + f) - 2\tilde{T}\ln(1 - \alpha) - \tilde{T}\ln(1 - \rho) - \ln(1 - \lambda) + \ln(24),$$

and $\tilde{T} = \left\lceil \log_{\frac{\lambda}{1-\alpha}}\left(\frac{1-\alpha}{4(n-f)}\right)\right\rceil \left(\log_{\frac{\lambda}{1-\alpha}}(1 - \alpha) + 2\right) + 2.$

The above statement also holds if $\mathcal{P}$ is a ρ-Dandelion protocol with $\rho < 1$.

Note that the upper bound on ε in Theorem 8 improves as the number of curious nodes f decreases (since α decreases with f) or when the expansion improves (as λ decreases, $\tilde{T}$ also decreases). Yet, there is a complex interplay between the parameters n, f, d, and λ above. Additionally, we point out that for a worst-case adversary the privacy guarantees can be established only if $f/d < 1$. For the average-case, this assumption can be dropped, and we are able to establish positive results for f as high as $\Theta_n(n)$.

6 Proof sketch for Theorem 8

Although results for worst-case and average-case adversaries have their own technical specificity, they both share the same general idea. Specifically, we introduce a random process that helps bounding from above the value of ε. This random process resembles a random walk

that at each step reveals its position to the adversary with some probability that depends on ρ and on the state of the process. We call this process a *random walk with probabilistic die out*. Then, we show that such random walk mixes sufficiently well before its position is revealed, which provides indistinguishability between any two possible sources.

The first half of our proof (step I) relies on the reduction of a gossip protocol to a random walk with probabilistic die out. This part is slightly different for different protocols, but for simplicity we only present step I for the cobra walk. In the second half (step II), we only analyze a random walk with probabilistic die out. It is hence universal and applies to both cobra walks and Dandelion protocols. Complete proofs for both protocols can be found in Appendix C of the full version of the paper [28].

6.1 Step I: reduction to a random walk with probabilistic die out

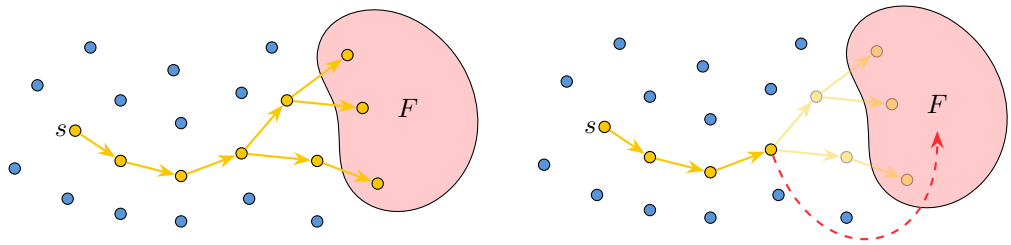

(a) Dissemination of a cobra walk.

(b) Dissemination of a random walk with die out.

Figure 1 Illustration of the reduction from a cobra walk (Fig. 1a) to a random walk with probabilistic die out (Fig. 1b). In Fig. 1a, the dissemination continues after the walk branches and hits the curious set F in several places. In the random walk with die out, instead of letting the dissemination branch, we stop the dissemination as soon as the cobra walk branches and report the position of the branching node.

Consider a $(1+\rho)$-cobra walk started at s and denote $W^{(s)}$ the random variable indicating the last position of the cobra walk before it either branches or hits a curious node. More formally, if the round at which the cobra walk branches or contacts a curious node for the first time is τ, then the active set at this round would be $X_\tau^{(s)} = \{W^{(s)}\}$, with $W^{(s)} \in V \setminus F$. We first show that disclosing $W^{(s)}$ to the adversary reveals more information about the source than $S_{\text{ADV}}^{(s)}$. Intuitively, this follows from the Markov property of the active set $\left\{X_t^{(s)}\right\}_{t \geq t_\star}$ of the cobra walk. In fact, by definition of τ, we have $\tau \leq t_{\text{ADV}}$. Hence, the sequence of adversarial observations $S_{\text{ADV}}^{(v)}$ can be obtained from $X_\tau^{(s)} = \left\{W^{(s)}\right\}$ via a randomized mapping independent of the initial source s. Then, using the data processing inequality Theorem 14 of [40]) we show that for any two possible sources $u, v \in V \setminus F$, we have

$$D_\infty \left(S_{\text{ADV}}^{(v)} \parallel S_{\text{ADV}}^{(u)} \right) \leq D_\infty \left(W^{(v)} \parallel W^{(u)} \right). \tag{6}$$

This means that it suffices to obtain an upper bound on $D_\infty \left(W^{(v)} \parallel W^{(u)} \right)$ for any $u, v \in V \setminus F$ to obtain an appropriate value for ε. Then, we note that $W^{(s)}$ can be described as the death site of a process we refer to as *random walk with probabilistic die out*, which was started at s. Such a process constitutes a random walk which is killed at each step either (i) if it hits a curious node, or otherwise (ii) with probability ρ. We illustrate this process in Figure 1 and how it relates to the cobra walk.

6.2 Step II: upper bounding the max divergence between death sites

The rest of the proof is dedicated to analyzing the probability distribution of the death site of such a process. Let $\boldsymbol{Q} = \hat{\boldsymbol{A}}[V \setminus F]$ be the principled submatrix of $\hat{\boldsymbol{A}}$ induced by the rows and columns of $V \setminus F$ and let $\boldsymbol{R}$ be a diagonal matrix of size $(n-f) \times (n-f)$ such that $R_{ww} = \deg_F(w)/d$ for every $w \in V \setminus F$. Then, $W^{(s)}$ can be described as an absorbing Markov chain. More precisely, let nodes from $V \setminus F$ be transient states, and equip every node $w \in V \setminus F$ with an absorbing state $\mathrm{sink}(w)$ which corresponds to the event of dying at w. The transition matrix of our absorbing Markov chain can be written in a block form as

$$\boldsymbol{P} = \begin{bmatrix} (1-\rho)\boldsymbol{Q} & \boldsymbol{O}_{(n-f)\times(n-f)} \\ \rho\boldsymbol{I}_{n-f} + (1-\rho)\boldsymbol{R} & \boldsymbol{I}_{n-f} \end{bmatrix}. \tag{7}$$

In the above, $\boldsymbol{P}_{xy}$ denotes the transition probability from a state y to a state x. The first $n-f$ columns correspond to transition probabilities from transient states $w \in V \setminus F$ and the last $n-f$ ones correspond to transition probabilities from absorbing states $\mathrm{sink}(w)$ for $w \in V \setminus F$. The probability of transitioning between two transient states $v, u \in V \setminus F$ (top-left block of $\boldsymbol{P}$) is defined similarly to a simple random walk on G, multiplied by the probability of not branching $(1-\rho)$. The transition probability between w and $\mathrm{sink}(w)$ (bottom-left block of $\boldsymbol{P}$) is naturally defined as the probability of branching plus the probability of contacting a curious node at the current step without branching.

According to the above, being absorbed in $\mathrm{sink}(w)$ corresponds to the event $W^{(s)} = w$. Hence, using $\boldsymbol{Q}$ and $\boldsymbol{R}$ to compute a closed form expression for absorbing probabilities of the above Markov chain, we can rewrite $D_\infty\left(W^{(v)} \parallel W^{(u)}\right)$ as follows

$$D_\infty\left(W^{(v)} \parallel W^{(u)}\right) = \max_{w \in V \setminus F} \ln \frac{(\boldsymbol{I}_{n-f} - (1-\rho)\boldsymbol{Q})^{-1}_{vw}}{(\boldsymbol{I}_{n-f} - (1-\rho)\boldsymbol{Q})^{-1}_{uw}}. \tag{8}$$

To conclude the proof, we now need to upper bound the right-hand side (8). To do so, we first note that, as per Theorem 3.2.1 in [35], we can use the following series decomposition,

$$(\boldsymbol{I}_{n-f} - (1-\rho)\boldsymbol{Q})^{-1} = \sum_{t=0}^{\infty} (1-\rho)^t \boldsymbol{Q}^t. \tag{9}$$

This means that we can reduce the computation of $D_\infty\left(W^{(v)} \parallel W^{(u)}\right)$ to analyzing the powers of the matrix $\boldsymbol{Q}^t$. Furthermore, for large values of t, we can approximate $\boldsymbol{Q}^t$ by a one-rank matrix using the first eigenvalue and the first eigenvector of $\boldsymbol{Q}$. This motivates us to study the spectral properties of $\boldsymbol{Q}$. We begin by showing that $\boldsymbol{Q}$ is dominated by its first eigenvalue. To further estimate the coordinates of the first eigenvector of $\boldsymbol{Q}$, we need to introduce subsidiary matrices $\overline{\boldsymbol{Q}}$ and $\underline{\boldsymbol{Q}}$. We carefully design these matrices to have an explicit first eigenvector and so that their entries bound from above and below respectively those of $\boldsymbol{Q}$. Using these two properties, we obtain a measure of how far the first eigenvector of $\boldsymbol{Q}$ is from the uniform vector $\mathbf{1}_{n-f}/\sqrt{n-f}$. By controlling spectral properties of $\boldsymbol{Q}$, we establish efficient one-rank approximations of high powers of $\boldsymbol{Q}$. Applying this to (8), we obtain an upper bound on the max divergence between $W^{(v)}$ and $W^{(u)}$, for any $u, v \in V \setminus F$. Specifically, assuming that the adversarial density $\alpha_F < 1 - \lambda$, we get

$$D_\infty\left(W^{(v)} \parallel W^{(u)}\right) \leq \ln(\rho(n-f)+f) - 2\tilde{T}\ln(1-\alpha_F) - \tilde{T}\ln(1-\rho) - \ln(1-\lambda) + \ln(24),$$

where $\tilde{T} = \left\lceil \log_{\frac{\lambda}{1-\alpha_F}}\left(\frac{1-\alpha_F}{4(n-f)}\right)\right\rceil \left(\log_{\frac{\lambda}{1-\alpha_F}}(1-\alpha_F) + 2\right) + 2$. Finally, substituting (6) in the above, and upper bounding α_F as per Section 5.1 we get the expected result.

7 Trade-off: Dissemination time vs. privacy

Note that when the gossip protocol parameter ρ decreases, the privacy guarantees in Theorem 8 improve. Yet, this worsens the dissemination time, which suggests the existence of a *trade-off* between the dissemination time and the source anonymity of the protocol. In this section, we formalize this observation by showing the tightness of Theorem 8 on a family of strong expanders called *near-Ramanujan graphs*. Intuitively, for dense enough graph topologies, most terms in Theorem 8 vanish, hence considerably simplifying the analysis of the result. Near-Ramanujan graphs can be defined as follows.

▶ **Definition 9** (Near-Ramanujan family of graphs). *Let $\mathcal{G}$ be an infinite family of regular graphs. $\mathcal{G}$ is called near-Ramanujan if there exists a constant $c > 0$ such that $\lambda(G) \leq cd(G)^{-1/2}$ for any graph $G \in \mathcal{G}$ of large enough size.*

This choice of graph family is motivated by the fact that near-Ramanujan graphs naturally arise in the study of dense random regular graphs. In fact, for any large enough n and any $3 \leq d \leq n/2$ (with dn even) a random d-regular graph on n nodes is near-Ramanujan with high probability as shown in [10, 52]. That means that almost every d-regular graph is near-Ramanujan. Besides using near-Ramanujan graphs, we assume the topologies to be dense enough, i.e., $d \in n^{\Omega_n(1)}$. Refining the statement of Theorem 8 to this family of graphs, we obtain the following corollary.

▶ **Corollary 10.** *Let $\mathcal{P}$ be a $(1+\rho)$-cobra walk and let $\mathcal{G}$ be a family of d-regular near-Ramanujan graphs with n nodes and $d \in n^{\Omega_n(1)}$. Suppose $f/d \in 1 - \Omega_n(1)$ (resp. $f/n \in 1 - \Omega_n(1)$). Then, for any $G \in \mathcal{G}$ of large enough size n and any $\rho \in 1 - \Omega_n(1)$, $\mathcal{P}$ satisfies ε-DP against a worst-case adversary (resp. an average-case adversary) for some*

$$\varepsilon \in \ln\left(\rho(n-f)+f\right) + O_n(1).$$

The above statement also holds if $\mathcal{P}$ is a ρ-Dandelion protocol with $\rho < 1$.

From Corollary 10, when $\rho = 0$, we obtain a level of differential privacy that matches, up to an additive constant, the universal lower bound $\varepsilon \geq \ln(f-1)$. Accordingly, $\rho = 0$ leads to an *optimal* differential privacy guarantee. However, in this case, both the cobra walk and the Dandelion protocol degenerate into simple random walks with dissemination time in $\Omega_n(n \log(n))$ [2]. Increasing ρ parameter makes the dissemination faster, but potentially worsens the privacy guarantees.

Studying Dandelion and cobra walks, we show that the result in Corollary 10 is tight up to an additive constant. Then, we formally validate our intuition that decreasing ρ increases the dissemination time by providing corresponding tight guarantees on dissemination time. Finally, to put our results in perspective, we compare them to a random walk (optimal privacy but high dissemination time), and to a 2-cobra walk (optimal dissemination time with bad, completely vacuous, privacy guarantees). We summarize our findings for both worst-case and average-case adversaries in Table 1 and the detailed analysis can be found in Appendix D of the full version of the paper [28].

8 Summary & future directions

This paper presents an important step towards quantifying the inherent level of source anonymity that gossip protocols provide on general graphs. We formulate our results through the lens of differential privacy. First, we present a universal lower bound on the level of

Table 1 Summary of the tension between differential privacy of a $(1+\rho)$-cobra walk and Dandelion gossip and their dissemination time on dense near-Ramanujan graphs. Graphs have diameter D and consist of n nodes, f of which are curious. Note that the upper bounds on ε hold under assumptions in Corollary 10. Lower bounds on ε hold assuming $f/n \in 1 - \Omega_n(1)$, and for cobra walk we also assume $f \in n^{\Omega_n(1)}$. Dissemination time bounds for cobra walk and Dandelion hold for $\rho \in \omega_n\left(\sqrt{\log(n)/n}\right)$ and $\rho \in \Omega_n(1/n)$ respectively.

Protocol	Privacy (ε)	Dissemination time	References
Random walk	$\ln(f) + \Theta_n(1)$	$\Theta_n(n \log(n))$	Corollary 10, Theorem 5, [2]
ρ-Dandelion	$\ln(\rho(n-f)+f) + \Theta_n(1)$	$\Theta_n\left(\frac{1}{\rho} + D\right)$	Corollary 10, Theorem 45 [28], Theorem 49 [28]
$(1+\rho)$-Cobra walk	$\ln(\rho(n-f)+f) + \Theta_n(1)$	$O_n\left(\frac{\log(n)}{\rho^3}\right)$, $\Omega_n\left(\frac{\log(n)}{\rho}\right)$	Corollary 10, Theorem 32 [28], Theorem 44 [28]
2-Cobra walk	$\ln(n) + \Omega_n(1)$	$\Theta_n(\log(n))$	Theorem 32 [28], Theorem 44 [28]

differential privacy an arbitrary gossip protocol can satisfy. Then, we devise an in-depth analysis of the privacy guarantees of $(1+\rho)$-cobra walk and ρ-Dandelion protocols on expander graphs. When $\rho = 0$, the protocols spread the gossip via a random walk, which achieves optimal privacy, but has poor dissemination time. On the other hand, we show that increasing ρ improves the dissemination time while the privacy deteriorates. In short, our tight analysis allows to formally establish the trade-off between dissemination time and the level of source anonymity these protocols provide. An interesting open research question would be to establish whether this "privacy vs dissemination time" trade-off is fundamental or if there exists a class of gossip protocols that could circumvent this trade-off.

We consider differential privacy, because, unlike other weaker notions of privacy (e.g., MLE-based bounds), it can be applied against an *arbitrary* strategy of the adversary, factoring in *any* prior beliefs an adversary may have about the location of the source and the nature of the gossip protocol. This makes differential privacy strong and resilient. However, differential privacy is often criticized for being too stringent in some settings. Consequently, a number of possible interesting relaxations have been proposed in the literature such as Pufferfish [36] and Renyi differential privacy [44]. Adapting our analysis to these definitions constitutes an interesting open direction as it would enable consideration of less stringent graphs structures and probability metrics.

Finally, we believe that our results could be applied to solve privacy related problems in other settings. For example, it was recently observed in [13] that sharing sensitive information via a randomized gossip can amplify the privacy guarantees of some learning algorithms, in the context of privacy-preserving decentralized machine learning. However, this work only considers the cases when the communication topology is a clique or a ring. We believe that the techniques we develop in this paper can be useful to amplify privacy of decentralized machine learning on general topologies. This constitutes an interesting open problem.

References

1 Huseyin Acan, Andrea Collevecchio, Abbas Mehrabian, and Nick Wormald. On the push&pull protocol for rumor spreading. *SIAM Journal on Discrete Mathematics*, 31(2):647–668, 2017. doi:10.1145/2767386.2767416.

2 David J. Aldous. Lower bounds for covering times for reversible markov chains and random walks on graphs. *Journal of Theoretical Probability*, 2:91–100, 1989.

3 Yeganeh Alimohammadi, Christian Borgs, and Amin Saberi. Algorithms using local graph features to predict epidemics. In *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2022)*, 2022. doi:10.1137/1.9781611977073.136.

4 Amos Beimel and Shlomi Dolev. Buses for anonymous message delivery. *Journal of Cryptology*, 16(1), 2003. doi:10.1007/s00145-002-0128-6.

5 Aurélien Bellet, Rachid Guerraoui, and Hadrien Hendrikx. Who started this rumor? Quantifying the natural differential privacy of gossip protocols. In *International Symposium on Distributed Computing (DISC 2020)*, 2020. doi:10.4230/LIPIcs.DISC.2020.8.

6 Petra Berenbrin, George Giakkoupis, and Peter Kling. Tight bounds for coalescing-branching random walks on regular graphs. In *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2018)*, 2018.

7 Shaileshh Bojja Venkatakrishnan, Giulia Fanti, and Pramod Viswanath. Dandelion: Redesigning the bitcoin network for anonymity. In *Proceedings of the ACM on Measurement and Analysis of Computing Systems (SIGMETRICS 2017)*, 2017. doi:10.1145/3078505.3078528.

8 Stephen Boyd, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52(6):2508–2530, 2006. doi:10.1109/TIT.2006.874516.

9 Yun Chai, Youguo Wang, and Liang Zhu. Information sources estimation in time-varying networks. *IEEE Transactions on Information Forensics and Security*, 16:2621–2636, 2021. doi:10.1109/TIFS.2021.3050604.

10 Nicholas A. Cook, Larry Goldstein, and Tobias Johnson. Size biased couplings and the spectral gap for random regular graphs. *Annals of Probability*, 46:72–125, 2018. doi:10.1214/17-AOP1180.

11 Colin Cooper, Tomasz Radzik, and Nicolas Rivera. The coalescing-branching random walk on expanders and the dual epidemic process. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing (PODC 2016)*, 2016. doi:10.1145/2933057.2933119.

12 Colin Cooper, Tomasz Radzik, and Nicolás Rivera. Improved cover time bounds for the coalescing-branching random walk on graphs. In *Proceedings of the 29th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA 2017)*, 2017. doi:10.1145/3087556.3087564.

13 Edwige Cyffers and Aurélien Bellet. Privacy amplification by decentralization. In *International Conference on Artificial Intelligence and Statistics (AIStat 2020)*, 2020.

14 Debajyoti Das, Sebastian Meiser, Esfandiar Mohammadi, and Aniket Kate. Anonymity trilemma: Strong anonymity, low bandwidth overhead, low latency-choose two. In *IEEE Symposium on Security and Privacy (SP)*, pages 108–126, 2018. doi:10.1109/SP.2018.00011.

15 Damien Desfontaines and Balazs Pejo. Sok: Differential privacies. In *Proceedings on Privacy Enhancing Technologies Symposium (PETS 2020)*, 2020.

16 Claudia Díaz, Stefaan Seys, Joris Claessens, and Bart Preneel. Towards measuring anonymity. In *Privacy Enhancing Technologies*, pages 54–68, 2003.

17 Benjamin Doerr, Mahmoud Fouz, and Tobias Friedrich. Social networks spread rumors in sublogarithmic time. In *Proceedings of the forty-third annual ACM symposium on Theory of computing (STOC 2011)*, 2011.

18 Chinmoy Dutta, Gopal Pandurangan, Rajmohan Rajaraman, and Scott Roche. Coalescing-branching random walks on graphs. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA 2013)*, 2013. doi:10.1145/2817830.

19 Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. Consensus in the presence of partial synchrony. *Journal of the ACM*, 35(2):288–323, 1988. doi:10.1145/42282.42283.

20 Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3*, pages 265–284, 2006.

21 Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2013. doi:10.1561/0400000042.

22 Giulia Fanti, Peter Kairouz, Sewoong Oh, Kannan Ramchandran, and Pramod Viswanath. Rumor source obfuscation on irregular trees. In *International Conference on Measurement and Modeling of Computer Systems, (SIGMETRICS 2016)*, 2016. doi:10.1145/2896377.2901471.

23 Giulia Fanti, Peter Kairouz, Sewoong Oh, Kannan Ramchandran, and Pramod Viswanath. Hiding the rumor source. *IEEE Transactions on Information Theory*, 63(10):6679–6713, 2017. doi:10.1109/TIT.2017.2696960.

24 Chryssis Georgiou, Seth Gilbert, Rachid Guerraoui, and Dariusz R Kowalski. Asynchronous gossip. *Journal of the ACM*, 60(2):1–42, 2013. doi:10.1145/2450142.2450147.

25 Chryssis Georgiou, Seth Gilbert, and Dariusz R. Kowalski. Confidential gossip. In *International Conference on Distributed Computing Systems (DISC 2011)*, 2011. doi:10.1109/ICDCS.2011.71.

26 George Giakkoupis, Rachid Guerraoui, Arnaud Jégou, Anne-Marie Kermarrec, and Nupur Mittal. Privacy-conscious information diffusion in social networks. In *International Symposium on Distributed Computing (DISC 2015)*, 2015.

27 Karol Gotfryd, Marek Klonowski, and Dominik Pająk. On location hiding in distributed systems. In *Structural Information and Communication Complexity*, pages 174–192, 2017.

28 Rachid Guerraoui, Anne-Marie Kermarrec, Anastasiia Kucherenko, Rafael Pinot, and Sasha Voitovych. On the Inherent Anonymity of Gossiping (Full Version), 2023. arXiv:2308.02477.

29 Zeyu Guo and He Sun. Gossip vs. markov chains, and randomness-efficient rumor spreading. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2014)*, 2014.

30 Herbert W Hethcote. The mathematics of infectious diseases. *SIAM review*, 42(4):599–653, 2000. doi:10.1137/S0036144500371907.

31 Justin Hsu, Marco Gaboardi, Andreas Haeberlen, Sanjeev Khanna, Arjun Narayan, Benjamin C Pierce, and Aaron Roth. Differential privacy: An economic method for choosing epsilon. In *IEEE 27th Computer Security Foundations Symposium*, pages 398–410, 2014.

32 Yufan Huang, Richeng Jin, and Huaiyu Dai. Differential privacy and prediction uncertainty of gossip protocols in general networks. In *IEEE Global Communications Conference (GLOBECOM 2020)*, 2020. doi:10.1109/GLOBECOM42002.2020.9322558.

33 Jiaojiao Jiang, Sheng Wen, Shui Yu, Yang Xiang, and Wanlei Zhou. Identifying propagation sources in networks: State-of-the-art and comparative studies. *IEEE Communications Surveys & Tutorials*, 19(1):465–481, 2017. doi:10.1109/COMST.2016.2615098.

34 Richard Karp, Christian Schindelhauer, Scott Shenker, and Berthold Vocking. Randomized rumor spreading. In *41st Annual Symposium on Foundations of Computer Science (FOCS 2000)*, 2000. doi:10.1109/SFCS.2000.892324.

35 John G Kemeny and J Laurie Snell. *Finite markov chains*. Springer New York, NY, 1960.

36 Daniel Kifer and Ashwin Machanavajjhala. Pufferfish: A framework for mathematical privacy definitions. *ACM Transactions on Database Systems (TODS)*, 39(1):1–36, 2014. doi:10.1145/2514689.

37 István Z Kiss, Joel C Miller, Péter L Simon, et al. Mathematics of epidemics on networks. *Cham: Springer*, 598:31, 2017.

38 Dariusz R Kowalski and Christopher Thraves Caro. Estimating time complexity of rumor spreading in ad-hoc networks. In *International Conference on Ad-Hoc Networks and Wireless (ADHOC-NOW 2013)*, 2013.

39 Jaewoo Lee and Chris Clifton. How much is enough? Choosing ε for differential privacy. In *Information Security: 14th International Conference*, pages 325–340, 2011.

40 Friedrich Liese and Igor Vajda. On divergences and informations in statistics and information theory. *IEEE Transactions on Information Theory*, 52(10):4394–4412, 2006. doi:10.1109/TIT.2006.881731.

41 Xuecheng Liu, Luoyi Fu, Bo Jiang, Xiaojun Lin, and Xinbing Wang. Information source detection with limited time knowledge. In *Proceedings of the Twentieth ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 2019)*, pages 389–390, 2019. doi:10.1145/3323679.3326626.

42 Yang Liu, Junfeng Wu, Ian R. Manchester, and Guodong Shi. Gossip algorithms that preserve privacy for distributed computation part I: The algorithms and convergence conditions. In *IEEE Conference on Decision and Control (CDC 2018)*, 2018. doi:10.1109/CDC.2018.8619783.

43 Guy Melancon. Just how dense are dense graphs in the real world? A methodological note. In *Proceedings of the AVI Workshop on BEyond Time and Errors: Novel Evaluation Methods for Information Visualization (BELIV 2006)*, 2006. doi:10.1145/1168149.1168167.

44 Ilya Mironov. Rényi differential privacy. In *2017 IEEE 30th computer security foundations symposium (CSF)*, pages 263–275. IEEE, 2017. doi:10.1109/CSF.2017.11.

45 Michael Mitzenmacher, Rajmohan Rajaraman, and Scott Roche. Better bounds for coalescing-branching random walks. *ACM Transactions on Parallel Computing*, 5(1):1–23, 2018. doi:10.1145/3209688.

46 Pedro C. Pinto, Patrick Thiran, and Martin Vetterli. Locating the source of diffusion in large-scale networks. *Physical Review Letters*, 109(6), 2012. doi:10.1103/PhysRevLett.109.068702.

47 Boris Pittel. On spreading a rumor. *SIAM Journal on Applied Mathematics*, 47(1):213–223, 1987. doi:10.1137/0147013.

48 Michael K Reiter and Aviel D Rubin. Crowds: Anonymity for web transactions. *ACM transactions on information and system security (TISSEC)*, 1(1):66–92, 1998. doi:10.1145/290163.290168.

49 D. Shah and T. Zaman. Rumors in a network: Who's the culprit? *IEEE Transactions on Information Theory*, 57(8):5163–5181, 2011. doi:10.1109/TIT.2011.2158885.

50 Devavrat Shah and Tauhid Zaman. Detecting sources of computer viruses in networks: Theory and experiment. In *Proceedings of ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 2010)*, 2010.

51 Robin Snader and Nikita Borisov. A tune-up for tor: Improving security and performance in the tor network. In *Network and Distributed System Security Symposium (NDSS 2008)*, 2008.

52 Konstantin E. Tikhomirov and Pierre Youssef. The spectral gap of dense random regular graphs. *The Annals of Probability*, 2019.

53 Parv Venkitasubramaniam and Venkat Anantharam. Anonymity under light traffic conditions using a network of mixes. In *46th Annual Allerton Conference on Communication, Control, and Computing (ALLERTON 2008)*, 2008. doi:10.1109/ALLERTON.2008.4797721.

54 Matthew K. Wright, Micah Adler, Brian Neil Levine, and Clay Shields. An analysis of the degradation of anonymous protocols. In *Network and Distributed System Security Symposium (NDSS 2002)*, 2002.

55 Yi Yu, Tengyao Wang, and Richard J. Samworth. A useful variant of the Davis–Kahan theorem for statisticians. *Biometrika*, 102:315–323, 2014.