Muffliato: Peer-to-Peer Privacy Amplification for Decentralized Optimization and Averaging

Edwige Cyffers [1,⋆] Mathieu Even[2,⋆], Aurélien Bellet [1], Laurent Massoulié [2]

⋆ Equal contribution

[1]Univ. Lille, Inria, CNRS, Centrale Lille, UMR 9189 - CRIStAL, F-59000 Lille

[2]Inria Paris - Département d'informatique de l'ENS, PSL Research University, Paris, France

Abstract

Decentralized optimization is increasingly popular in machine learning for its scalability and efficiency. Intuitively, it should also provide better privacy guarantees, as nodes only observe the messages sent by their neighbors in the network graph. But formalizing and quantifying this gain is challenging: existing results are typically limited to Local Differential Privacy (LDP) guarantees that overlook the advantages of decentralization. In this work, we introduce pairwise network differential privacy, a relaxation of LDP that captures the fact that the privacy leakage from a node u to a node v may depend on their relative position in the graph. We then analyze the combination of local noise injection with (simple or randomized) gossip averaging protocols on fixed and random communication graphs. We also derive a differentially private decentralized optimization algorithm that alternates between local gradient descent steps and gossip averaging. Our results show that our algorithms amplify privacy guarantees as a function of the distance between nodes in the graph, matching the privacy-utility trade-off of the trusted curator, up to factors that explicitly depend on the graph topology. Remarkably, these factors become constant for expander graphs. Finally, we illustrate our privacy gains with experiments on synthetic and real-world datasets.

1 Introduction

Training machine learning models traditionally requires centralizing data in a single server, raising issues of scalability and privacy. An alternative is to use Federated Learning (FL), where each user keeps her data on device [43, 35]. In *fully decentralized* FL, the common hypothesis of a central server is also removed, letting users, represented as nodes in a graph, train the model via peer-to-peer communications along edges. This approach improves scalability and robustness to central server failures, enabling lower latency, less power consumption and quicker deployment [42, 11, 53, 49, 1, 41, 38].

Another important dimension is privacy, as a wide range of applications deal with sensitive and personal data. The gold standard to quantify the privacy leakage of algorithms is Differential Privacy (DP) [20]. DP typically requires to randomly perturb data-dependent computations to prevent the final model from leaking too much information about any individual data point (e.g., through data memorization). However, decentralized algorithms do not only reveal the final model to the participating nodes, but also the results of some intermediate computations. A standard solution is to use Local Differential Privacy (LDP) [36, 19], where random perturbations are performed locally by each user, thus protecting against an attacker that would observe everything that users share. From the

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

algorithmic standpoint, local noise addition can be easily performed within decentralized algorithms, as done for instance in [33, 6, 15, 59, 57]. Unfortunately, LDP requires large amounts of noise, and thus provides poor utility.

In this work, we show that LDP guarantees give a very pessimistic view of the privacy offered by decentralized algorithms. Indeed, there is no central server receiving all messages, and the participating nodes can only observe the messages sent by their neighbors in the graph. So, a given node should intuitively leak less information about its private data to nodes that are far away. We formally quantify this privacy amplification for the fundamental brick of communication at the core of decentralized optimization: gossip algorithms. Calling *Muffliato* the combination of local noise injection with a gossip averaging protocol, we precisely track the resulting privacy leakage between each pair of nodes. Through gossiping, the private values and noise terms of various users add up, obfuscating their contribution well beyond baseline LDP guarantees: as their distance in the graph increases, the privacy loss decreases. We then show that the choice of graph is crucial to enforce a good privacy-utility trade-off while preserving the scalability of gossip algorithms.

Our results are particularly attractive in situations where nodes want stronger guarantees with respect to some (distant) peers. For instance, in social network graphs, users may have lower privacy expectations with respect to close relatives than regarding strangers. In healthcare, a patient might trust her family doctor more than she trusts other doctors, and in turn more than employees of a regional agency and so on, creating a hierarchical level of trust that our algorithms naturally match.

1.1 Contributions and outline of the paper

(i) Inspired by the recent definitions of Cyffers and Bellet [17], we introduce *pairwise network DP*, a relaxation of Local Differential Privacy which is able to quantify the privacy loss of a decentralized algorithm for each pair of distinct users in a graph.

(ii) We propose *Muffliato*[1], a privacy amplification mechanism composed of local Gaussian noise injection at the node level followed by gossiping for averaging the private values. It offers privacy amplification that increases as the distance between two nodes increases. Informally, the locally differentially private value shared by a node u is mixed with other contributions, to the point that the information that leaks to another node v can have a very small sensitivity to the initial value in comparison to the accumulated noise.

(iii) We analyze both synchronous gossip [18] and randomized gossip [11] under a unified privacy analysis with arbitrary time-varying gossip matrices. We show that the magnitude of the privacy amplification is significant: the average privacy loss over all the pairs in this setting reaches the optimal privacy-utility trade-off of a trusted aggregator, up to a factor $\frac{d}{\sqrt{\lambda_W}}$, where λ_W is the weighted graph eigengap and d the maximum degree of the graph. Remarkably, this factor can be of order 1 for expanders, yielding a sweet spot in the privacy-utility-scalability trade-off of gossip algorithms. Then, we study the case where the graph is itself random and private, and derive stronger privacy guarantees.

(iv) Finally, we develop and analyze differentially private decentralized Gradient Descent (GD) and Stochastic Gradient Descent (SGD) algorithms to minimize a sum of local objective functions. Building on *Muffliato*, our algorithms alternate between rounds of differentially private gossip communications and local gradient steps. We prove that they enjoy the same privacy amplification described above for averaging, up to factors that depend on the regularity of the global objective.

(v) We demonstrate the usefulness of our approach and analysis through experiments on synthetic and real-world datasets and network graphs. We illustrate how privacy is amplified between nodes in the graph as a function of their distance, and show how time-varying random graphs can be used to split the privacy loss more uniformly across nodes in decentralized optimization.

1.2 Related work

Gossip algorithms and decentralized optimization. Gossip algorithms [10, 18] were introduced to compute the global average of local vectors through peer-to-peer communication, and are at the core of many decentralized optimization algorithms. Classical decentralized optimization algorithms alternate

[1]The name is borrowed from the Harry Potter series: it designates a "spell that filled the ears of anyone nearby with an unidentifiable buzzing", thereby concealing messages from unintended listeners through noise injection.

between gossip communications and local gradient steps [47, 37, 38], or use dual formulations and formulate the consensus constraint using gossip matrices to obtain decentralized dual or primal-dual algorithms [53, 31, 24, 25, 39, 1]. We refer the reader to [48] for a broader survey on decentralized optimization. Our algorithms are based on the general analysis of decentralized SGD in [38].

LDP and privacy amplification mechanisms. Limitations of LDP for computing the average of the private values of n users have been studied, showing that for a fixed privacy budget, the expected squared error in LDP is n times larger than in central DP [12]. More generally, LDP is also known to significantly reduce utility for many learning problems [60, 56], which motivates the study of intermediate trust models. Cryptographic primitives, such as secure aggregation [21, 54, 9, 13, 34, 5, 52] and secure shuffling [16, 23, 3, 30, 29], as well as additional mechanisms such as amplification by subsampling [2] or amplification by iteration [27], can offer better utility for some applications, but cannot be easily applied in a fully decentralized setting, as they require coordination by a central server.

Privacy amplification through decentralization. The idea that decentralized communications can provide differential privacy guarantees was initiated by [7] in the context of rumor spreading. Closer to our work, [17] showed privacy amplification for random walk algorithms on complete graphs, where the model is transmitted from one node to another sequentially. While we build on their notion of Network DP, our work differs from [17] in several aspects: (i) our analysis holds for any graph and explicitly quantifies its effect, (ii) instead of worst-case privacy across all pairs of nodes, we prove pairwise guarantees that are stronger for nodes that are far away from each other, and (iii) unlike random walk approaches, gossip algorithms allow parallel computation and thus better scalability.

2 Setting and Pairwise Network Differential Privacy

We study a decentralized model where n nodes (users) hold private datasets and communicate through gossip protocols, that we describe in Section 2.1. In Section 2.2, we recall differential privacy notions and the two natural baselines for our work, central and local DP. Finally, we introduce in Section 2.3 the relaxation of local DP used throughout the paper: the *pairwise network DP*.

2.1 Gossip Algorithms

We consider a connected graph $G = (\mathcal{V}, \mathcal{E})$ on a set $\mathcal{V}$ of n users. An edge $\{u, v\} \in \mathcal{E}$ indicates that u and v can communicate (we say they are neighbors). Each user $v \in \mathcal{V}$ holds a local dataset $\mathcal{D}_v$ and we aim at computing averages of private values. This averaging step is a key building block for solving machine learning problems in a decentralized manner, as will be discussed in Section 4. From any graph, we can derive a gossip matrix.

Definition 1 (Gossip matrix)**.** *A gossip matrix over a graph G is a* symmetric *matrix $W \in \mathbb{R}^{n\times n}$ with non-negative entries, that satisfies $W\mathbb{1} = \mathbb{1}$* i.e. *W is stochastic ($\mathbb{1} \in \mathbb{R}^n$ is the vector with all entries equal to 1), and such that for any $u, v \in \mathcal{V}$, $W_{u,v} > 0$ implies that $\{u, v\} \in \mathcal{E}$ or $u = v$.*

The iterates of synchronous gossip [18] are generated through a recursion of the form $x^{t+1} = Wx^t$, and converge to the mean of initial values $x^0 \in \mathbb{R}^n$ at a linear rate $e^{-t\lambda_W}$, with λ_W defined below.

Definition 2 (Spectral gap)**.** *The spectral gap λ_W associated with a gossip matrix W is $\min_{\lambda\in \mathrm{Sp}(W)\setminus\{1\}}(1 - |\lambda|)$, where $\mathrm{Sp}(W)$ is the spectrum of W.*

The inverse of λ_W is the relaxation time of the random walk on G with transition probabilities W, and is closely related to the connectivity of the graph: adding edges improve mixing properties (λ_W increases), but can reduce scalability by increasing node degrees (and thus the per-iteration communication complexity). The rate of convergence can be accelerated to $e^{-t\sqrt{\lambda_W}}$ using re-scaled Chebyshev polynomials, leading to iterates of the form $x^t = P_t(W)x^0$ [8].

Definition 3 (Re-scaled Chebyshev polynomials)**.** *The re-scaled Chebyshev polynomials $(P_t)_{t\geqslant 0}$ with scale parameter $\gamma \in [1, 2]$ are defined by second-order linear recursion:*

$$P_0(X) = 1\,, \quad P_1(X) = X\,, \quad P_{t+1}(X) = \gamma X P_t(X) + (1-\gamma)P_{t-1}(X)\,,\ t \geqslant 2\,. \tag{1}$$

2.2 Rényi Differential Privacy

Differential Privacy (DP) quantifies how much information the output of an algorithm $\mathcal{A}$ leaks about the dataset taken as input [20]. DP requires to define an adjacency relation between datasets. In this

work, we adopt a user-level relation [44] which aims to protect the whole dataset $\mathcal{D}_v$ of a given user represented by a node $v \in \mathcal{V}$. Formally, $\mathcal{D} = \cup_{v\in\mathcal{V}}\mathcal{D}_v$ and $\mathcal{D}' = \cup_{v\in\mathcal{V}}\mathcal{D}'_v$ are adjacent datasets, denoted by $\mathcal{D} \sim \mathcal{D}'$, if there exists $v \in \mathcal{V}$ such that only $\mathcal{D}_v$ and $\mathcal{D}'_v$ differ. We use $\mathcal{D} \sim_v \mathcal{D}'$ to denote that $\mathcal{D}$ and $\mathcal{D}'$ differ only in the data of user v.

We use Rényi Differential Privacy (RDP) [45] to measure the privacy loss, which allows better and simpler composition than the classical (ϵ, δ)-DP. Note that any (α, ε)-RDP algorithm is also $(\varepsilon + \ln(1/\delta)/(\alpha - 1), \delta)$-DP for any $0 < \delta < 1$ [45].

Definition 4 (Rényi Differential Privacy)**.** *An algorithm $\mathcal{A}$ satisfies (α, ε)-Rényi Differential Privacy (RDP) for $\alpha > 1$ and $\varepsilon > 0$ if for all pairs of neighboring datasets $\mathcal{D} \sim \mathcal{D}'$:*

$$D_\alpha\left(\mathcal{A}(\mathcal{D})||\mathcal{A}(\mathcal{D}')\right) \leqslant \varepsilon\,, \tag{2}$$

where for two random variables X and Y, $D_\alpha\big(X \,||\, Y\big)$ is the Rényi divergence *between X and Y:*

$$D_\alpha\big(X \,||\, Y\big) = \tfrac{1}{\alpha-1}\ln\int\left(\tfrac{\mu_X(z)}{\mu_Y(z)}\right)^\alpha \mu_Y(z)dz\,.$$

with μ_X and μ_Y the respective densities of X and Y.

Without loss of generality, we consider gossip algorithms with a single real value per node (in that case, $\mathcal{D}_v = \{x_v\}$ for some $x_v \in \mathbb{R}$), and we aim at computing a private estimation of the mean $\bar{x} = (1/n)\sum_v x_v$. The generalization to vectors is straightforward, as done subsequently for optimization in Section 4. In general, the value of a (scalar) function g of the data can be privately released using the Gaussian mechanism [20, 45], which adds $\eta \sim \mathcal{N}(0, \sigma^2)$ to $g(\mathcal{D})$. It satisfies $(\alpha, \alpha\Delta_g^2/2\sigma^2)$-RDP for any $\alpha > 1$, where $\Delta_g = \sup_{\mathcal{D}\sim\mathcal{D}'} \|g(\mathcal{D}) - g(\mathcal{D}')\|$ is the sensitivity of g. We focus on the Gaussian mechanism for its simplicity (similar results could be derived for other DP mechanisms), and thus assume an upper bound on the L_2 inputs sensitivity.

Assumption 1. *There exists some constant $\Delta > 0$ such that for all $u \in \mathcal{V}$ and for any adjacent datasets $\mathcal{D} \sim_u \mathcal{D}'$, we have $\|x_u - x'_u\| \leqslant \Delta$.*

In central DP, a trusted aggregator can first compute the mean $\bar{x}$ (which has sensitivity Δ/n) and then reveal a noisy version with the Gaussian mechanism. On the contrary, in local DP where there is no trusted aggregator and everything that a given node reveals can be observed, each node must locally perturb its input (which has sensitivity Δ), deteriorating the privacy-utility trade-off. Formally, to achieve (α, ε)-DP, one cannot have better utility than:

$$\mathbb{E}\left[\left\|x^{\text{out}} - \bar{x}\right\|^2\right] \leqslant \frac{\alpha\Delta^2}{2n\varepsilon} \quad \text{for local DP}\,, \quad \text{and} \quad \mathbb{E}\left[\left\|x^{\text{out}} - \bar{x}\right\|^2\right] \leqslant \frac{\alpha\Delta^2}{2n^2\varepsilon} \quad \text{for central DP}\,,$$

where x^{out} is the output of the algorithm. This $1/n$ gap motivates the study of relaxations of local DP.

2.3 Pairwise Network Differential Privacy

We relax local DP to take into account privacy amplification between nodes that are distant from each other in the graph. We define a decentralized algorithm $\mathcal{A}$ as a randomized mapping that takes as input a dataset $\mathcal{D} = \cup_{v\in\mathcal{V}}(\mathcal{D}_v)$ and outputs the transcript of all messages exchanged between users in the network. A message between neighboring users $\{u, v\} \in \mathcal{E}$ at time t is characterized by the tuple $(u, m(t), v)$: user u sent a message with content $m(t)$ to user v, and $\mathcal{A}(\mathcal{D})$ is the set of all these messages. Each node v only has a partial knowledge of $\mathcal{A}(\mathcal{D})$, captured by its *view*:

$$\mathcal{O}_v\big(\mathcal{A}(\mathcal{D})\big) = \{(u, m(t), v) \in \mathcal{A}(\mathcal{D}) \quad \text{such that} \quad \{u, v\} \in \mathcal{E}\}\,.$$

This subset corresponds to direct interactions of v with its neighbors, which provide only an indirect information on computations in others parts of the graph. Thus, we seek to express privacy constraints that are personalized for each pair of nodes. This is captured by our notion of Pairwise Network DP.

Definition 5 (Pairwise Network DP)**.** *For $f : \mathcal{V} \times \mathcal{V} \to \mathbb{R}^+$, an algorithm $\mathcal{A}$ satisfies (α, f)-Pairwise Network DP (PNDP) if for all pairs of distinct users $u, v \in \mathcal{V}$ and neighboring datasets $D \sim_u D'$:*

$$D_\alpha\big(\mathcal{O}_v(\mathcal{A}(\mathcal{D}))\,||\,\mathcal{O}_v(\mathcal{A}(\mathcal{D}'))\big) \leqslant f(u, v)\,. \tag{3}$$

We note $\varepsilon_{u\to v} = f(u, v)$ the privacy leaked to v from u and say that u is $(\alpha, \varepsilon_{u\to v})$-PNDP with respect to v if only inequality (3) holds for $f(u, v) = \varepsilon_{u\to v}$.

By taking f constant in Definition 5, we recover the definition of Network DP [17]. Our pairwise variant refines Network DP by allowing the privacy guarantee to depend on u and v (typically, on their relative position in the graph). We assume that users are *honest but curious*: they truthfully follow the protocol, but may try to derive as much information as possible from what they observe. We refer to Appendix G for a natural adaptation of our definition and results to the presence of *colluding nodes* and to the protection of *groups* of users.

In addition to pairwise guarantees, we will use the *mean privacy loss* $\overline{\varepsilon}_v = \frac{1}{n}\sum_{u\in\mathcal{V}\setminus\{v\}} f(u,v)$ to compare with baselines LDP and trusted aggregator by enforcing $\overline{\varepsilon} = \max_{v\in\mathcal{V}} \overline{\varepsilon}_v \leqslant \varepsilon$. The value $\overline{\varepsilon}_v$ is the average of the privacy loss from all the nodes to v and thus does not correspond to a proper privacy guarantee, but it provides a convenient way to summarize our gains, noting that distant nodes — in ways that will be specified — will have better privacy guarantee than this average, while worst cases will remain bounded by the baseline LDP guarantee provided by local noise injection.

3 Private Gossip Averaging

In this section, we analyze a generic algorithm with arbitrary time-varying communication matrices for averaging. Then, we instantiate and discuss these results for synchronous communications with a fixed gossip matrix, communications using randomized gossip [11], and with Erdös-Rényi graphs.

3.1 General Privacy Analysis of Gossip Averaging

We consider gossip over time-varying graphs $(G_t)_{0\leqslant t\leqslant T}$, defined as $G_t = (\mathcal{V}, \mathcal{E}_t)$, with corresponding gossip matrices $(W_t)_{0\leqslant t\leqslant T}$. The *generic Muffliato* algorithm $\mathcal{A}^T$ for averaging $x = (x_v)_{v\in\mathcal{V}}$ corresponds to an initial noise addition followed by T gossip steps. Writing $W_{0:t} = W_{t-1}\dots W_0$, the iterates of $\mathcal{A}^T$ are thus defined by:

$$\forall v \in \mathcal{V}, x_v^0 = x_v + \eta_v \text{ with } \eta_v \sim \mathcal{N}(0,\sigma^2), \quad \text{and } x^{t+1} = W_t x^t = W_{0:t+1}(x+\eta)\,. \tag{4}$$

Note that the update rule at node $v \in \mathcal{V}$ writes as $x_v^{t+1} = \sum_{w\in\mathcal{N}_t(v)} (W_t)_{v,w} x_w^t$ where $\mathcal{N}_t(v)$ are the neighbors of v in G_t, so for the privacy analysis, the view of a node is:

$$\mathcal{O}_v\big(\mathcal{A}^T(\mathcal{D})\big) = \big\{\big(W_{0:t}(x+\eta)\big)_w \mid \quad \{v,w\}\in\mathcal{E}_t\,, \quad 0\leqslant t\leqslant T-1\big\} \cup \{x_v\}\,. \tag{5}$$

Theorem 1. *Let $T \geqslant 1$ and denote by $\mathcal{P}^T_{\{v,w\}} = \{s < T \,:\, \{v,w\}\in\mathcal{E}_s\}$ the set of time-steps with communication along edge $\{v,w\}$. Under Assumption 1, $\mathcal{A}^T$ is (α, f)-PNDP with:*

$$f(u,v) = \frac{\alpha\Delta^2}{2\sigma^2} \sum_{w\in\mathcal{V}} \sum_{t\in\mathcal{P}^T_{\{v,w\}}} \frac{(W_{0:t})^2_{u,w}}{\|(W_{0:t})_w\|^2}\,. \tag{6}$$

This theorem, proved in Appendix B, gives a tight computation of the privacy loss between every pair of nodes and can easily be computed numerically (see Section 5). Since distant nodes correspond to small entries in $W_{0:t}$, Equation 6 suggests that they reveal less to each other. We will characterize this precisely for the case of fixed communication graph in the next subsection.

Another way to interpret the result of Theorem 1 is to derive the corresponding mean privacy loss:

$$\overline{\varepsilon}_v = \frac{\alpha\Delta^2 T_v}{2n\sigma^2}\,, \tag{7}$$

where T_v is the total number of communications node v was involved with up to time T. Thus, in comparison with LDP, the mean privacy towards v is n/T_v times smaller. In other words, a node learns much less than in LDP as long as it communicates $o(n)$ times.

3.2 Private Synchronous *Muffliato*

We now consider *Muffliato* over a fixed graph (Algorithm 1). Note that we use gossip acceleration (see Definition 3). We start by analyzing the utility of *Muffliato*, which decomposes as an averaging error term vanishing exponentially fast, and a *bias* term due to the noise. General convergence rates are given in Appendix C, from which we extract the following result.

Algorithm 1: MUFFLIATO

Input: local values $(x_v)_{v\in\mathcal{V}}$ to average, gossip matrix W on a graph G, in T iterations, noise variance σ^2

$\gamma \leftarrow 2\frac{1-\sqrt{\lambda_W(1-\frac{\lambda_W}{4})}}{(1-\lambda_W/2)^2}$

for *all nodes v in parallel* **do**
  $x_v^0 \leftarrow x_v + \eta_v$ where $\eta_v \sim \mathcal{N}(0,\sigma^2)$
for $t = 0$ *to* $T-1$ **do**
  for *all nodes v in parallel* **do**
    for *all neighbors w defined by W* **do**
      Send x_v^t, receive x_w^t
    $x_v^{t+1} \leftarrow (1-\gamma)x_v^{t-1} + \gamma\sum_{w\in\mathcal{N}_v} W_{v,w}x_w^t$

Algorithm 2: RANDOMIZED MUFFLIATO

Input: local values $(x_v)_{v\in\mathcal{V}}$ to average, activation intensities $(p_{\{v,w\}})_{\{v,w\}\in\mathcal{E}}$, in T iterations, noise variance σ^2

for *all nodes v in parallel* **do**
  $x_v^0 \leftarrow x_v + \eta_v$ where $\eta_v \sim \mathcal{N}(0,\sigma^2)$
for $t = 0$ *to* $T-1$ **do**
  Sample $\{v_t, w_t\} \in \mathcal{E}$ with probability $p_{\{v_t,w_t\}}$
  v_t and w_t exchange $x_{v_t}^t$ and $x_{w_t}^t$
  Local averaging:
  $x_{v_t}^{t+1} = x_{w_t}^{t+1} = \frac{x_{v_t}^{t+1}+x_{w_t}^{t+1}}{2}$
  For $v \in \mathcal{V}\setminus\{v_t,w_t\}$, $x_v^{t+1} = x_v^t$

Table 1: Utility of *Muffliato* for several topologies under the constraint $\overline{\varepsilon} \leqslant \varepsilon$ for the classic gossip matrix where $W_{v,w} = \min(1/d_v, 1/d_w)$ and d_v is the degree of node v. $\tilde{\mathcal{O}}(\cdot)$ hides constant and logarithmic factors. Recall that utility is $\tilde{\mathcal{O}}(\alpha\Delta^2/n\varepsilon)$ for LDP and $\tilde{\mathcal{O}}(\alpha\Delta^2/n^2\varepsilon)$ for central DP.

Graph	Arbitrary	Expander	C-Torus	Complete	Ring
Algorithm 1	$\tilde{\mathcal{O}}\left(\frac{\alpha\Delta^2 d}{n^2\varepsilon\sqrt{\lambda_W}}\right)$	$\tilde{\mathcal{O}}\left(\frac{\alpha\Delta^2}{n^2\varepsilon}\right)$	$\tilde{\mathcal{O}}\left(\frac{\alpha\Delta^2 C}{n^{2-1/C}\varepsilon}\right)$	$\tilde{\mathcal{O}}\left(\frac{\alpha\Delta^2}{n\varepsilon}\right)$	$\tilde{\mathcal{O}}\left(\frac{\alpha\Delta^2}{n\varepsilon}\right)$
Algorithm 2	$\tilde{\mathcal{O}}\left(\frac{\alpha\Delta^2}{n^2\varepsilon\lambda_W}\right)$	$\tilde{\mathcal{O}}\left(\frac{\alpha\Delta^2}{n^2\varepsilon}\right)$	$\tilde{\mathcal{O}}\left(\frac{\alpha\Delta^2}{n^{2-2/C}\varepsilon}\right)$	$\tilde{\mathcal{O}}\left(\frac{\alpha\Delta^2}{n^2\varepsilon}\right)$	$\tilde{\mathcal{O}}\left(\frac{\alpha\Delta^2}{n\varepsilon}\right)$

Theorem 2 (Utility analysis)**.** *Let λ_W be the spectral gap of W. Muffliato (Algorithm 1) verifies, for any $t \geqslant T^{\text{stop}}$:*

$$\frac{1}{2n}\sum_{v\in\mathcal{V}}\mathbb{E}\left[\left\|x_v^t - \bar{x}\right\|^2\right] \leqslant \frac{3\sigma^2}{n}, \quad \textit{where} \quad T^{\text{stop}} = \frac{1}{\sqrt{\lambda_W}}\ln\left(\frac{n}{\sigma^2}\max\left(\sigma^2, \frac{1}{n}\sum_{v\in\mathcal{V}}\|x_v-\bar{x}\|^2\right)\right).$$

For the privacy guarantees, Theorem 1 still holds as accelerated gossip can be seen as a post-processing of the non-accelerated version. Thanks to the fixed graph, we can derive a more explicit formula.

Corollary 1. *Algorithm 1 satisfies $(\alpha, \varepsilon_{u\to v}^T(\alpha))$-PNDP for node u with respect to v, with:*

$$\varepsilon_{u\to v}^T(\alpha) \leqslant \frac{\alpha\Delta^2 n}{2\sigma^2}\max_{\{v,w\}\in\mathcal{E}} W_{v,w}^{-2}\sum_{t=1}^{T}\mathbb{P}\left(X^t = v | X^0 = u\right)^2,$$

where $(X^t)_t$ is the random walk on graph G, with transitions W.

This result allows us to directly relate the privacy loss from u to v to the probability that the random walk on G with transition probabilities given by the gossip matrix W goes from u to v in a certain number of steps. It thus captures a notion of distance between nodes in the graph. We also report the utility under fixed mean privacy loss $\overline{\varepsilon} = \max_{v\in\mathcal{V}}\overline{\varepsilon}_v \leqslant \varepsilon$ in Table 1 for various graphs, where one can see a utility-privacy trade-off improvement of $n\sqrt{\lambda_W}/d$, where d is the maximum degree, compared to LDP. Using expanders closes the gap with a trusted aggregator (i.e., central DP) up to constant and logarithmic terms. Remarkably, graph topologies that make gossip averaging efficient (i.e. with big $\sqrt{\lambda_W}/d$), such as exponential graphs or hypercubes [58], are also the ones that achieve optimal privacy amplification (up to logarithmic factors). In other words, *privacy, utility and scalability are compatible*.

3.3 Private Randomized *Muffliato*

Synchronous protocols require global coordination between nodes, which can be costly or even impossible in some settings. On the contrary, asynchronous protocols only require separated activation of edges: they are thus are more resilient to stragglers nodes and faster in practice. In asynchronous gossip, at a given time-step a single edge $\{u, v\}$ is activated independently from the past with probability

$p_{\{u,v\}}$, as described by Boyd et al. [11]. In our setting, randomized *Muffliato* (Algorithm 2) corresponds to instantiating our general analysis with $W^t = W_{\{v_t,w_t\}} = I_n - (e_{v_t} - e_{w_t})(e_{v_t} - e_{w_t})^\top/2$ if $\{v_t, w_t\}$ is sampled at time t. The utility analysis is similar to the synchronous case.

Theorem 3 (Utility analysis)**.** *Let $\lambda(p)$ be the spectral gap of graph G with weights $(p_{\{v,w\}})_{\{v,w\}\in\mathcal{E}}$. Randomized* Muffliato *(Algorithm 2) verifies, for all $t \geqslant T^{\mathrm{stop}}$:*

$$\frac{1}{2n}\sum_{v\in\mathcal{V}} \mathbb{E}\left[\left\|x_v^t - \bar{x}\right\|^2\right] \leqslant \frac{2\sigma^2}{n}\,, \quad \textit{where } T^{\mathrm{stop}} = \frac{1}{\lambda(p)}\ln\left(\frac{n}{\sigma^2}\max\left(\sigma^2, \frac{1}{n}\sum_{v\in\mathcal{V}}\left\|x_v^0 - \bar{x}\right\|^2\right)\right).$$

To compare with synchronous gossip (Algorithm 1), we note that activation probabilities can be derived from a gossip matrix W by taking $p_{\{u,v\}} = 2W_{\{u,v\}}/n$ implying that $\lambda(p) = 2\lambda_W/n$, thus requiring n times more iterations to reach the same utility as the synchronous applications of matrix W. However, for a given time-horizon T and node v, the number of communications v can be bounded with high probability by a T/n multiplied by a constant whereas Algorithm 1 requires $d_v T$ communications. Consequently, as reported in Table 1, for a fixed privacy mean $\bar{\varepsilon}_v$, Algorithm 2 has the same utility as Algorithm 1, up to two differences: the degree factor d_v is removed, while $\sqrt{\lambda_W}$ degrades to λ_W as we do not accelerate randomized gossip (see Remark 1 below). Randomized gossip can thus achieve an optimal privacy-utility trade-off with large-degree graphs, as long as the spectral gap is small enough.

Remark 1 (Accelerating Randomized *Muffliato*)**.** *For simplicity, Randomized* Muffliato *(Algorithm 2) is not accelerated, while* Muffliato *(Algorithm 1) uses Chebychev acceleration to obtain a dependency on $\sqrt{\lambda_W}$ rather than λ_W. Thus, and as illustrated by Table 1, Algorithm 2 does not improve over Algorithm 1 for all values of d (maximum degree), n and λ_W. However, Algorithm 2 can be accelerated using a continuized version of Nesterov acceleration [24, 25], thus replacing $\lambda(p)$ in the expression of T^{stop} in Theorem 3, by $\sqrt{\lambda(p)/(dn)}$. Doing so, using randomized communications improve privacy guarantees over Algorithm 1 for all graphs considered in Table 1.*

3.4 Erdös-Rényi Graphs

So far the graph was considered to be public and the amplification only relied on the secrecy of the messages. In practice, the graph may be sampled randomly and the nodes need only to know their direct neighbors. We show that we can leverage this through the weak convexity of Rényi DP to amplify privacy between non-neighboring nodes. We focus on Erdös-Rényi graphs, which can be built without central coordination by picking each edge independently with the same probability q. For $q = c\ln(n)/n$ where $c > 1$, Erdös-Rényi graphs are good expanders with node degrees $d_v = \mathcal{O}(\log n)$ and λ_W concentrating around 1 [32]. We obtain the following privacy guarantees.

Theorem 4 (*Muffliato* on a random Erdös-Rényi graph)**.** *Let $\alpha > 1$, $T \geqslant 0$, $\sigma^2 \geqslant \frac{\Delta^2\alpha(\alpha-1)}{2}$ and $q = c\frac{\ln(n)}{n}$ for $c > 1$. Let $u, v \in \mathcal{V}$ be distinct nodes. After running Algorithm 1 with these parameters, node u is $(\alpha, \varepsilon^T_{u\to v}(\alpha))$-PNDP with respect to v, with:*

$$\varepsilon^T_{u\to v}(\alpha) \leqslant \begin{cases} \dfrac{\alpha\Delta^2}{2\sigma^2} & \textit{with probability } q\,, \\ \dfrac{\alpha\Delta^2}{\sigma^2}\dfrac{Td_v}{n-d_v} & \textit{with probability } 1-q\,. \end{cases}$$

This results shows that with probability q, u and v are neighbors and there is no amplification compared to LDP. The rest of the time, with probability $1-q$, the privacy matches that of a trusted aggregator up to a degree factor $d_v = \mathcal{O}(\log n)$ and $T = \tilde{\mathcal{O}}(1/\sqrt{\lambda_W}) = \tilde{\mathcal{O}}(1)$ [32]. In particular, if several rounds of gossip averaging are needed, as in the next section for SGD, changing the graph mitigates the privacy loss of the rounds where two nodes are neighbors thanks to the rounds where they are not.

4 Private Decentralized Optimization

We now build upon *Muffliato* to design decentralized optimization algorithms. Each node $v \in \mathcal{V}$ possesses a data-dependent function $\phi_v : \mathbb{R}^D \to \mathbb{R}$ and we wish to *privately* minimize the function

$$\phi(\theta) = \frac{1}{n}\sum_{v\in\mathcal{V}}\phi_v(\theta)\,, \quad \text{with } \phi_v(\theta) = \frac{1}{|\mathcal{D}_v|}\sum_{x_v\in\mathcal{D}_v}\ell_v(\theta, x_v)\,, \quad \theta\in\mathbb{R}^D\,, \tag{8}$$

Algorithm 3: MUFFLIATO-SGD and MUFFLIATO-GD

Input: initial points $\theta_v^0 \in \mathbb{R}^D$, number of iterations T, step sizes $\nu > 0$, noise variance σ^2, gossip matrices $(W_t)_{t\geqslant 0}$, local functions ϕ_v, number of communication rounds K

for $t = 0$ *to* $T-1$ **do**

  for *all nodes* v *in parallel* **do**

    Compute $\hat{\theta}_v^t = \theta_v^t - \nu \nabla_\theta \ell_v(\theta_v^t, x_v^t)$ where $x_v^t \sim \mathcal{L}_v$

  $\theta_v^{t+1} = \text{MUFFLIATO}\big((\hat{\theta}_v^t)_{v\in\mathcal{V}}, W_t, K, \nu^2\sigma^2\big)$

where $\mathcal{D}_v$ is the (finite) dataset corresponding to user v for data lying in a space $\mathcal{X}_v$, and $\ell_v : \mathbb{R}^D \times \mathcal{X}_v \to \mathbb{R}$ a loss function. We assume that ϕ is μ-strongly convex, and each ϕ_v is L-smooth, and denote $\kappa = L/\mu$. We note that our results can be extended to the general convex and smooth setting. Denoting by $\theta^\star$ the minimizer of ϕ, for some non-negative $(\zeta_v^2)_{v\in\mathcal{V}}$, $(\rho_v^2)_{v\in\mathcal{V}}$ and all $v \in \mathcal{V}$, we assume:

$$\|\nabla\phi_v(\theta^\star) - \nabla\phi(\theta^\star)\|^2 \leqslant \zeta_v^2 \quad , \quad \mathbb{E}\left[\|\nabla\ell_v(\theta^\star, x_v) - \nabla\phi(\theta^\star)\|^2\right] \leqslant \rho_v^2 \,, \quad x_v \sim \mathcal{L}_v \,,$$

where $\mathcal{L}_v$ is the uniform distribution over $\mathcal{D}_v$. We write $\bar{\rho}^2 = \frac{1}{n}\sum_{v\in\mathcal{V}} \rho_v^2$ and $\bar{\zeta}^2 = \frac{1}{n}\sum_{v\in\mathcal{V}} \zeta_v^2$.

We introduce Algorithm 3, a private version of the classical decentralized SGD algorithm studied in [38]. Inspired by the optimal algorithm MSDA of Scaman et al. [53] that alternates between K Chebychev-accelerated gossip communications and expensive dual gradient computations, our Algorithm 3 alternates between K Chebychev-accelerated gossip communications and cheap local stochastic gradient steps. This alternation reduces the total number of gradients leaked, a crucial point for achieving good privacy. Note that in Algorithm 3, each communication round uses a potentially different gossip matrix W_t. In the results stated below, we fix $W_t = W$ for all t and defer the more general case to Appendix F, where different independent Erdös-Rényi graphs with same parameters are used at each communication round.

Remark 2. *Our setting encompasses both GD and SGD.* Muffliato-*GD is obtained by removing the stochasticity,* i.e., *setting* $\ell_v(\cdot) = \phi_v(\cdot)$. *In that case,* $\bar{\rho}^2 = 0$.

Theorem 5 (Utility analysis of Algorithm 3)**.** *For suitable step-size parameters, for a total number of* T^{stop} *computations and* $T^{\text{stop}}K$ *communications, with:*

$$T^{\text{stop}} = \tilde{\mathcal{O}}(\kappa) \,, \quad \textit{and} \quad K = \left\lceil \sqrt{\lambda_W}^{-1} \ln\left(\max\left(n, \frac{\bar{\zeta}^2}{D\sigma^2 + \bar{\rho}^2}\right)\right)\right\rceil \,,$$

the iterates $(\theta^t)_{t\geqslant 0}$ *generated by Algorithm 3 verify* $\mathbb{E}\left[\phi(\tilde{\theta}^{\text{out}}) - \phi(\theta^\star)\right] = \tilde{\mathcal{O}}(\frac{D\sigma^2+\bar{\rho}^2}{\mu n T^{\text{stop}}})$ *where* $\tilde{\theta}^{\text{out}} \in \mathbb{R}^D$ *is a weighted average of the* $\bar{\theta}^t = \frac{1}{n}\sum_{v\in\mathcal{V}} \theta_v^t$ *until* T^{stop}.

For the following privacy analysis, we need a bound on the sensitivity of gradients with respect to the data. To this end, we assume that for all v and x_v, $\ell_v(\cdot, x_v)$ is $\Delta_\phi/2$ Lipschitz[2].

Theorem 6 (Privacy analysis of Algorithm 3)**.** *Let* u *and* v *be two distinct nodes in* $\mathcal{V}$. *After* T *iterations of Algorithm 3 with* $K \geqslant 1$, *node* u *is* $(\alpha, \varepsilon_{u\to v}^T(\alpha))$*-PNDP with respect to* v, *with:*

$$\varepsilon_{u\to v}^T(\alpha) \leqslant \frac{T\Delta_\phi^2\alpha}{2\sigma^2} \sum_{k=0}^{K-1} \sum_{w:\{v,w\}\in\mathcal{E}} \frac{(W^k)_{u,w}^2}{\|(W^k)_w\|^2} \,. \tag{9}$$

Thus, for any $\varepsilon > 0$, *Algorithm 3 with* $T^{\text{stop}}(\kappa, \sigma^2, n)$ *steps and for* K *as in Theorem 5, there exists* f *such that the algorithm is* (α, f)*-pairwise network DP, with:*

$$\forall v \in \mathcal{V} \,, \quad \bar{\varepsilon}_v \leqslant \varepsilon \quad \textit{and} \quad \mathbb{E}\left[\phi(\tilde{\theta}^{\text{out}}) - \phi(\theta^\star)\right] \leqslant \tilde{\mathcal{O}}\left(\frac{\alpha D\Delta_\phi^2 d}{\mu n^2 \varepsilon\sqrt{\lambda_W}} + \frac{\bar{\rho}^2}{nL}\right) \,,$$

where $d = \max_{v\in\mathcal{V}} d_v$.

[2]This assumption can be replaced by the more general Assumption 2 given in Appendix F.

The term $\frac{\bar{\rho}^2}{nL}$ above (which is equal to zero for *Muffliato*-GD, see Remark 2) is privacy independent. It is typically dominated by the first term, which corresponds to the utility loss due to privacy. Comparing Theorem 6 with the results for *Muffliato* (Table 1 in Section 3.2), the only difference lies in the factor $D\Delta_\phi^2/\mu$. Note that Δ_ϕ^2 plays the role of the sensitivity Δ^2, and D appears naturally due to considering D-dimensional parameters. On the other hand, μ is directly related to the complexity of the optimization problem through the condition number κ: the easier the problem, the better the privacy-utility trade-off of our algorithm. Regarding the influence of the graph, the same discussion as after Corollary 1 applies here. In particular, for expander graphs like the exponential graph of [58], the factor $d/\sqrt{\lambda_W}$ is constant. In this case, converting to standard (ϵ, δ)-DP gives $\tilde{\mathcal{O}}\big(\frac{D\Delta_\phi^2}{\mu n^2 \epsilon^2}\big)$, recovering the optimal privacy-utility trade-off of central DP [4, 55] up to logarithmic factors. Remarkably, we achieve this optimal rate under a near-linear gradient complexity of $T^{\mathrm{stop}}(\kappa, \sigma^2, n)n = \tilde{\mathcal{O}}(\kappa n)$ and near-linear total number of messages $T^{\mathrm{stop}}(\kappa, \sigma^2, n)Kn = \tilde{\mathcal{O}}(\kappa n)$.

Remark 3 (Time-varying graphs)**.** *The analysis of* Muffliato-*GD/SGD presented in this section (Theorems 5 and 6) assumes constant gossip matrices $W_t = W$. A more general version of these results is presented in Appendix F to handle time-varying matrices and graphs. This can be used to model randomized communications (as previously described for gossip averaging in Section 3.3) as well as user dropout (see experiments in Appendix H). Time-varying graphs can also be used to split the privacy loss more uniformly across the different nodes by avoiding that nodes have the same neighbors across multiple gossip computations. We illustrate this experimentally for decentralized optimization in Section 5, where we randomize the graph after each gradient step of* Muffliato-*GD.*

5 Experiments

In this section, we show that pairwise network DP provides significant privacy gains in practice even for moderate size graphs. We use synthetic graphs and real-world graphs for gossip averaging. For decentralized optimization, we solve a logistic regression problem on real-world data with time-varying Erdos-Renyi graphs, showing in each case clear gains in privacy compared to LDP. The code used to obtain these results is available at `https://github.com/totilas/muffliato`.

Averaging on synthetic graphs. We generate synthetic graphs with $n = 2048$ nodes and define the corresponding gossip matrix according to the Hamilton scheme. Note that the privacy guarantees of *Muffliato* are deterministic for a fixed W, and defined by Equation 4. For each graph, we run *Muffliato* for the theoretical number of steps required for convergence, and report in Figure 1(a) the pairwise privacy guarantees aggregated by shortest path lengths between nodes, along with the LDP baseline for comparison. *Exponential graph* (generalized hypercube): this has shown to be an efficient topology for decentralized learning [58]. Consistently with our theoretical result, privacy is significantly amplified. The shortest path completely defines the privacy loss, so there is no variance. *Erdos-Renyi graph* with $q = c \log n/n$ $(c \geqslant 1)$ [22], averaged over 5 runs: this has nearly the same utility-privacy trade-off as the exponential graph but with significant variance, which motivates the time-evolving version mentioned in Remark 3. *Grid:* given its larger mixing time, it is less desirable than the two previous graphs, emphasizing the need for careful design of the communication graph. *Geometric random graph:* two nodes are connected if and only if their distance is below a given threshold, which models for instance Bluetooth communications (effective only in a certain radius). We sample nodes uniformly at random in the square unit and choose a radius ensuring full connectivity. While the shortest path is a noisy approximation of the privacy loss, the Euclidean distance is a very good estimator as shown in Appendix H.

Averaging on real-world graphs. We consider the graphs of the Facebook ego dataset [40], where nodes are the friends of a given user (this central user is not present is the graph) and edges encode the friendship relation between these nodes. Ego graphs typically induce several clusters corresponding to distinct communities: same high school, same university, same hobbies... For each graph, we extract the giant connected component, choose a user at random and report its privacy loss with respect to other nodes. The privacy loss given by LDP is only relevant within the cluster of direct neighbors: privacy guarantees with respect to users in other communities are significantly better, as seen in Figure 1(b). We observe this consistently across other ego graphs (see Appendix H). This is in line with one of our initial motivation: our pairwise guarantees are well suited to situations where nodes want stronger privacy with respect to distant nodes.

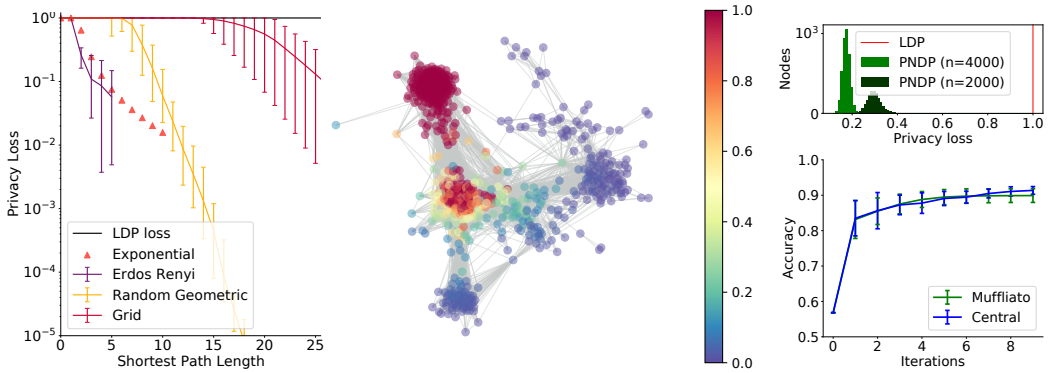

Figure 1: (a) Left: Privacy loss of *Muffliato* in pairwise NDP on synthetic graphs (best, worst and average in error bars over nodes at a given distance), confirming a significant privacy amplification as the distance increases. (b) Middle: Privacy loss of *Muffliato* from a node chosen at random on a Facebook ego graph, showing that leakage is very limited outside the node's own community. (c) Right: Privacy loss and utility of *Muffliato*-GD when using different Erdös-Rényi graphs after each gradient step, compared to a baseline based on a trusted aggregator.

Logistic regression on real-world data. Logistic regression corresponds to minimizing Equation 8 with loss function $\ell(\theta; x, y) = \ln(1 + \exp(-y\theta^\top x))$ where $x \in \mathbb{R}^d$ and $y \in \{-1, 1\}$. We use a binarized version of UCI Housing dataset.[3] We standardize the features and normalize each data point x to have unit L_2 norm so that the logistic loss is 1-Lipschitz for any (x, y). We split the dataset uniformly at random into a training set (80%) and a test set and further split the training set across users. After each gradient step of *Muffliato*-GD, we draw at random an Erdös-Rényi graph of same parameter q to perform the gossiping step and run the theoretical number of steps required for convergence. For each node, we keep track of the privacy loss towards the first node (note that all nodes play the same role). We report the pairwise privacy loss for this node with respect to all others for $n = 2000$ and $n = 4000$ in Figure 1(c) (top). We see that, as discussed in Remark 3, time-varying graphs are effective at splitting the privacy loss more uniformly across nodes: the privacy gains over LDP are clear with respect to all nodes. As captured by our theory, these gains increase with the number of nodes n in the system, and they also concentrate better around the mean. We compare the utility of *Muffliato*-GD to a federated learning alternative which uses the same parameters but aggregates noisy model updates using a *trusted* central server rather than by gossiping. As seen in Figure 1(c) (bottom), both approaches behave similarly in terms of accuracy across iterations.

6 Conclusion

We showed that gossip protocols amplify the LDP guarantees provided by local noise injection as values propagate in the graph. Despite the redundancy of gossip that, at first sight, could be seen as an obstacle to privacy, the privacy amplification turns out to be significant: it can nearly match the optimal privacy-utility trade-off of the trusted curator. From the fundamental building block — noise injection followed by gossip — that we analyzed under the name *Muffliato*, one can easily extend the analysis to other decentralized algorithms, such as the dual approach proposed in [53]. Our results are motivated by the typical relation between proximity in the communication graph and lower privacy expectations. Other promising directions are to assume that closer people are more similar, which leads to smaller individual privacy accounting [26], to design new notions of similarity between nodes in graphs that match the privacy loss variations, and to study privacy attacks [51].

Acknowledgments This work was supported by grants ANR-20-CE23-0015 (Project PRIDE) ANR-20-THIA-0014 program "AI_PhD@Lille", ANR-19-P3IA-0001(PRAIRIE Institute), and from the MSR-INRIA joint centre. We thank the organizers of NeurIPS@Paris 2021 for the in-person event that allowed us to meet by chance and to start this piece of work.

[3] https://www.openml.org/d/823

References

[1] Sulaiman A. Alghunaim and Ali H. Sayed. Linear convergence of primal-dual gradient methods and their performance in distributed optimization, 2019. URL https://arxiv.org/abs/1904.01196.

[2] Borja Balle, Gilles Barthe, and Marco Gaboardi. Privacy Amplification by Subsampling: Tight Analyses via Couplings and Divergences. In *NeurIPS*, 2018.

[3] Borja Balle, James Bell, Adrià Gascón, and Kobbi Nissim. Differentially Private Summation with Multi-Message Shuffling. Technical report, arxiv:1906.09116, 2019.

[4] Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private Empirical Risk Minimization: Efficient Algorithms and Tight Error Bounds. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 464–473, Philadelphia, PA, USA, October 2014. IEEE.

[5] James Henry Bell, Kallista A. Bonawitz, Adrià Gascón, Tancrède Lepoint, and Mariana Raykova. Secure single-server aggregation with (poly)logarithmic overheads. In *CCS*, 2020.

[6] Aurélien Bellet, Rachid Guerraoui, Mahsa Taziki, and Marc Tommasi. Personalized and Private Peer-to-Peer Machine Learning. In *AISTATS*, 2018.

[7] Aurélien Bellet, Rachid Guerraoui, and Hadrien Hendrikx. Who started this rumor? Quantifying the natural differential privacy guarantees of gossip protocols. In *DISC*, 2020.

[8] Raphaël Berthier, Francis Bach, and Pierre Gaillard. Accelerated gossip in networks of given dimension using jacobi polynomial iterations. *SIAM Journal on Mathematics of Data Science*, 2(1): 24–47, 2020. doi: 10.1137/19M1244822. URL https://doi.org/10.1137/19M1244822.

[9] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H. Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical Secure Aggregation for Privacy-Preserving Machine Learning. In *CCS*, 2017.

[10] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52(6):2508–2530, 2006. doi: 10.1109/TIT.2006.874516.

[11] Stephen Boyd, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah. Randomized gossip algorithms. *IEEE transactions on information theory*, 52(6):2508–2530, 2006.

[12] T.-H. Hubert Chan, Elaine Shi, and Dawn Song. Optimal Lower Bound for Differentially Private Multi-party Aggregation. In *ESA*, 2012.

[13] T.-H. Hubert Chan, Elaine Shi, and Dawn Song. Privacy-preserving stream aggregation with fault tolerance. In *Financial Cryptography*, 2012.

[14] Konstantinos Chatzikokolakis, Miguel E. Andrés, Nicolás E. Bordenabe, and Catuscia Palamidessi. Broadening the Scope of Differential Privacy Using Metrics. In De Cristofaro, Emiliano, Wright, and Matthew, editors, *The 13th Privacy Enhancing Technologies Symposium*, volume 7981 of *Lecture Notes in Computer Science*, pages 82–102, Bloomington, Indiana, United States, July 2013. Springer. doi: 10.1007/978-3-642-39077-7. URL https://hal.inria.fr/hal-00767210.

[15] Hsin-Pai Cheng, Patrick Yu, Haojing Hu, Syed Zawad, Feng Yan, Shiyu Li, Hai Helen Li, and Yiran Chen. Towards Decentralized Deep Learning with Differential Privacy. In *CLOUD*, 2019.

[16] Albert Cheu, Adam D. Smith, Jonathan Ullman, David Zeber, and Maxim Zhilyaev. Distributed Differential Privacy via Shuffling. In *EUROCRYPT*, 2019.

[17] Edwige Cyffers and Aurélien Bellet. Privacy amplification by decentralization, 2020. URL https://arxiv.org/abs/2012.05326.

[18] A. G. Dimakis, S. Kar, J. M. F. Moura, M. G. Rabbat, and A. Scaglione. Gossip algorithms for distributed signal processing. *Proceedings of the IEEE*, 98(11):1847–1864, 2010.

[19] John C Duchi, Michael I Jordan, and Martin J Wainwright. Local privacy and statistical minimax rates. In *FOCS*, 2013.

[20] Cynthia Dwork and Aaron Roth. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3-4):211–407, 2014.

[21] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our Data, Ourselves: Privacy Via Distributed Noise Generation. In *EUROCRYPT*, 2006.

[22] P. Erdös and A. Rényi. On random graphs i. *Publicationes Mathematicae Debrecen*, 6:290, 1959.

[23] Ulfar Erlingsson, Vitaly Feldman, Ilya Mironov, Ananth Raghunathan, and Kunal Talwar. Amplification by Shuffling: From Local to Central Differential Privacy via Anonymity. In *SODA*, 2019.

[24] Mathieu Even, Hadrien Hendrikx, and Laurent Massoulié. Asynchrony and acceleration in gossip algorithms. Technical report, arXiv:2011.02379, 2020.

[25] Mathieu Even, Raphaël Berthier, Francis Bach, Nicolas Flammarion, Hadrien Hendrikx, Pierre Gaillard, Laurent Massoulié, and Adrien Taylor. Continuized accelerations of deterministic and stochastic gradient descents, and of gossip algorithms. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 28054–28066. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper/2021/file/ec26fc2eb2b75aece19c70392dc744c2-Paper.pdf.

[26] Vitaly Feldman and Tijana Zrnic. Individual Privacy Accounting via a Rényi Filter. In *NeurIPS*, 2021.

[27] Vitaly Feldman, Ilya Mironov, Kunal Talwar, and Abhradeep Thakurta. Privacy Amplification by Iteration. In *FOCS*, 2018.

[28] Vitaly Feldman, Ilya Mironov, Kunal Talwar, and Abhradeep Thakurta. Privacy amplification by iteration. *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, Oct 2018. doi: 10.1109/focs.2018.00056. URL http://dx.doi.org/10.1109/FOCS.2018.00056.

[29] Vitaly Feldman, Audra McMillan, and Kunal Talwar. Hiding Among the Clones: A Simple and Nearly Optimal Analysis of Privacy Amplification by Shuffling. Technical report, arXiv:2012.12803, 2020.

[30] Badih Ghazi, Noah Golowich, Ravi Kumar, Pasin Manurangsi, Rasmus Pagh, and Ameya Velingker. Pure Differentially Private Summation from Anonymous Messages. Technical report, arXiv:2002.01919, 2020.

[31] Hadrien Hendrikx, Francis Bach, and Laurent Massoulié. An accelerated decentralized stochastic proximal algorithm for finite sums. In *Advances in Neural Information Processing Systems*, 2019.

[32] Christopher Hoffman, Matthew Kahle, and Elliot Paquette. Spectral gaps of random graphs and applications. *International Mathematics Research Notices*, 2021(11):8353–8404, May 2019. ISSN 1687-0247. doi: 10.1093/imrn/rnz077. URL http://dx.doi.org/10.1093/imrn/rnz077.

[33] Zhenqi Huang, Sayan Mitra, and Nitin Vaidya. Differentially Private Distributed Optimization. In *ICDCN*, 2015.

[34] Bargav Jayaraman, Lingxiao Wang, David Evans, and Quanquan Gu. Distributed learning without distress: Privacy-preserving empirical risk minimization. In *NeurIPS*, 2018.

[35] Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, Rafael G. L. D'Oliveira, Hubert Eichner, Salim El Rouayheb, David Evans, Josh Gardner, Zachary Garrett, Adrià Gascón, Badih Ghazi, Phillip B. Gibbons, Marco Gruteser, Zaid Harchaoui, Chaoyang He, Lie He, Zhouyuan Huo, Ben Hutchinson, Justin Hsu, Martin Jaggi, Tara Javidi, Gauri Joshi, Mikhail Khodak, Jakub Konecný, Aleksandra Korolova, Farinaz Koushanfar, Sanmi Koyejo, Tancrède Lepoint, Yang Liu, Prateek Mittal, Mehryar Mohri, Richard Nock, Ayfer Özgür, Rasmus Pagh, Hang Qi, Daniel Ramage, Ramesh Raskar, Mariana Raykova, Dawn Song, Weikang Song, Sebastian U. Stich, Ziteng Sun, Ananda Theertha Suresh, Florian Tramèr, Praneeth Vepakomma, Jianyu Wang, Li Xiong, Zheng Xu, Qiang Yang, Felix X. Yu, Han Yu, and Sen Zhao. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210, 2021.

[36] Shiva Prasad Kasiviswanathan, Homin K. Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam D. Smith. What Can We Learn Privately? In *FOCS*, 2008.

[37] Anastasia Koloskova, Sebastian Stich, and Martin Jaggi. Decentralized stochastic optimization and gossip algorithms with compressed communication. In *International Conference on Machine Learning*, volume 97, pages 3478–3487. PMLR, 2019.

[38] Anastasia Koloskova, Nicolas Loizou, Sadra Boreiri, Martin Jaggi, and Sebastian Stich. A unified theory of decentralized SGD with changing topology and local updates. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 5381–5393. PMLR, 13–18 Jul 2020. URL https://proceedings.mlr.press/v119/koloskova20a.html.

[39] Dmitry Kovalev, Elnur Gasanov, Alexander Gasnikov, and Peter Richtarik. Lower bounds and optimal algorithms for smooth and strongly convex decentralized optimization over time-varying networks. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 22325–22335. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper/2021/file/bc37e109d92bdc1ea71da6c919d54907-Paper.pdf.

[40] Jure Leskovec and Julian Mcauley. Learning to discover social circles in ego networks. In F. Pereira, C.J. Burges, L. Bottou, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012. URL https://proceedings.neurips.cc/paper/2012/file/7a614fd06c325499f1680b9896beedeb-Paper.pdf.

[41] Xiangru Lian, Ce Zhang, Huan Zhang, Cho-Jui Hsieh, Wei Zhang, and Ji Liu. Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent. In *NIPS*, 2017.

[42] Cassio G. Lopes and Ali H. Sayed. Incremental adaptive strategies over distributed networks. *IEEE Transactions on Signal Processing*, 55(8):4064–4077, 2007. doi: 10.1109/TSP.2007.896034.

[43] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-Efficient Learning of Deep Networks from Decentralized Data. In Aarti Singh and Jerry Zhu, editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 1273–1282. PMLR, 20–22 Apr 2017.

[44] H. Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. Learning Differentially Private Recurrent Language Models. In *ICLR*, 2018.

[45] Ilya Mironov. Renyi differential privacy. *CoRR*, abs/1702.07476, 2017. URL http://arxiv.org/abs/1702.07476.

[46] Mohar, Y Alavi, G Chartrand, and OR Oellermann. The laplacian spectrum of graphs. *Graph theory, combinatorics, and applications*, 1991.

[47] Angelia Nedic and Asuman Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, 54(1):48–61, 2009. doi: 10.1109/TAC.2008.2009515.

[48] Angelia Nedic, Alex Olshevsky, and Michael G. Rabbat. Network topology and communication-computation tradeoffs in decentralized optimization. *Proceedings of the IEEE*, 106(5):953–976, May 2018.

[49] Giovanni Neglia, Gianmarco Calbi, Don Towsley, and Gayane Vardoyan. The role of network topology for distributed machine learning. In *INFOCOM*, 2019.

[50] Helen Nissenbaum. Privacy as contextual integrity. *Washington Law Review*, 79(1):119–157, February 2004. ISSN 0043-0617.

[51] Dario Pasquini, Mathilde Raynal, and Carmela Troncoso. On the privacy of decentralized machine learning. Technical report, arXiv:2205.08443, 2022.

[52] César Sabater, Aurélien Bellet, and Jan Ramon. Distributed Differentially Private Averaging with Improved Utility and Robustness to Malicious Parties. Technical report, arXiv:2006.07218, 2020.

[53] Kevin Scaman, Francis Bach, Sébastien Bubeck, Yin Tat Lee, and Laurent Massoulié. Optimal algorithms for smooth and strongly convex distributed optimization in networks. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 3027–3036. PMLR, 06–11 Aug 2017. URL https://proceedings.mlr.press/v70/scaman17a.html.

[54] Elaine Shi, T.-H. Hubert Chan, Eleanor G. Rieffel, Richard Chow, and Dawn Song. Privacy-Preserving Aggregation of Time-Series Data. In *NDSS*, 2011.

[55] Di Wang, Minwei Ye, and Jinhui Xu. Differentially private empirical risk minimization revisited: Faster and more general. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *NeurIPS*, volume 30. Curran Associates, Inc., 2017.

[56] Di Wang, Marco Gaboardi, and Jinhui Xu. Empirical Risk Minimization in Non-interactive Local Differential Privacy Revisited. In *NeurIPS*, 2018.

[57] Jie Xu, Wei Zhang, and Fei Wang. A(dp)^{2}sgd: Asynchronous decentralized parallel stochastic gradient descent with differential privacy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.

[58] Bicheng Ying, Kun Yuan, Yiming Chen, Hanbin Hu, Pan Pan, and Wotao Yin. Exponential graph is provably efficient for decentralized deep training. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021.

[59] Xueru Zhang, Mohammad Mahdi Khalili, and Mingyan Liu. Improving the Privacy and Accuracy of ADMM-Based Distributed Algorithms. In *ICML*, 2018.

[60] Kai Zheng, Wenlong Mou, and Liwei Wang. Collect at Once, Use Effectively: Making Non-interactive Locally Private Learning Possible. In *ICML*, 2017.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] We define our relaxation of Local DP formally in Section 2, the gossip averaging is analyzed in Section 3 and show substantial amplification, supported by the experiments in Section 5. We define and prove guarantees for decentralized optimization in Section 4.
   (b) Did you describe the limitations of your work? [Yes] We discuss each theorem after in its subsection and show the effective magnitude of the privacy amplification in the experiments (Section 5).
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] We have included a broader impact statement at the end of the supplementary.

(d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]

2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] All our theorems and corollaries have their complete set of assumptions.
   (b) Did you include complete proofs of all theoretical results? [Yes] , except when proving similar results, we do not repeat the full proof but only explain the main differences.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] We provide the code needed to reproduce the results in the supplementary material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Most of our experiments have no hyperparameters to tune, and we provide the hyperparameters in Annex for Figure 1(c).
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] see Figures 1(c) and 1(a).
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No] All of the simulations ran in a few minutes on a regular laptop.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] We use Houses Dataset and Facebook Ego dataset and cite them.
   (b) Did you mention the license of the assets? [Yes] The dataset Houses is in the public domain, as indicated on the link provided and the Facebook ego dataset in under BSD license.
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] We have included our code in the supplementary.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Preliminary Lemmas and Notations

We conduct our analysis with the Gaussian mechanism, introduced in Section 2, that we recall here for readability.

Lemma 1 (Gaussian mechanism)**.** *For $\alpha > 1$, noise variance σ^2, sensitivity $\Delta > 0$ and $x, y \in \mathbb{R}$ such that $|x - y| \leqslant \Delta$, we have:*

$$D_\alpha\big(\mathcal{N}(x, \sigma^2) \,||\, \mathcal{N}(y, \sigma^2)\big) \leqslant \frac{\alpha \Delta^2}{2\sigma^2} \,.$$

For the privacy analysis when the graph is private and randomly sampled (Appendix E), we use the following result on the weak convexity of the Rényi divergence [28].

Lemma 2 (Quasi-convexity of Rényi divergence [28])**.** *Let $(\mu_i)_{i\in\mathcal{I}}$ and $(\nu_i)_{i\in\mathcal{I}}$ be probability distributions over shared space, such that for all $i \in \mathcal{I}$, we have $D_\alpha(\mu_i||\nu_i) \leqslant c/(\alpha - 1)$ for some $c \in (0, 1]$. Let ρ be a distribution over $\mathcal{I}$ and μ_ρ and ν_ρ be obtained by sampling i from ρ, and outputing a sample from μ_i and ν_i. Then, we have:*

$$D_\alpha(\mu_\rho||\nu_\rho) \leqslant (1 + c)\mathbb{E}\left[D_\alpha(\mu_i||\nu_i) \mid i \sim \rho\right] \,.$$

In the following, we will use the notation $u \sim v$ to denote that two nodes u and v are neighbors.

B General Privacy Analysis (Theorem 1)

Proof. We need to bound the privacy loss that occurs from the following view:

$$\mathcal{O}_v\big(\mathcal{A}^T(\mathcal{D})\big) = \big\{\big(W_{0:t}(x + \eta)\big)_w \mid \quad \{v, w\} \in \mathcal{E}_t \,, \quad 0 \leqslant t \leqslant T - 1\big\} \cup \{x_v\} \,.$$

We have:

$$D_\alpha\big(\mathcal{O}_v(\mathcal{A}^T(\mathcal{D})) \,||\, \mathcal{O}_v(\mathcal{A}^T(\mathcal{D}'))\big) \leqslant \sum_{t=0}^{T-1} \sum_{w\in\mathcal{N}_t(v)} D_\alpha\big(\big(W_{0:t}(x + \eta)\big)_w \,||\, \big(W_{0:t}(x' + \eta)\big)_w\big) \,.$$

We have $(W_{0:t}(x'+\eta))_w - (W_{0:t}(x+\eta))_w \sim \mathcal{N}((W_{0:t}(x'+\eta))_w - (W_{0:t}(x+\eta))_w, \sigma^2\|(W_{0:t})_w\|^2)$ with a sensitivity verifying $|(W_{0:t}(x'+\eta))_w - (W_{0:t}(x+\eta))_w|^2 \leqslant \Delta^2 (W_{0:t})^2_{u,w}$ under Assumption 1 and $\mathcal{D} \sim_u \mathcal{D}'$. Thus, using Lemma 1, we have:

$$D_\alpha\big(\big(W_{0:t}(x + \eta)\big)_w \,||\, \big(W_{0:t}(x' + \eta)\big)_w\big) \leqslant \frac{\alpha\Delta^2}{2\sigma^2} \frac{(W_{0:t})^2_{u,w}}{\|(W_{0:t})_w\|^2} \,,$$

leading to the desired $f(u, v)$. The mean privacy loss is then obtained by summing the above inequality for $u \neq v$ and $t < T$:

$$\begin{aligned}
\bar{\varepsilon}_v = \frac{1}{n} \sum_{u\neq v} f(u, v) &\leqslant \frac{1}{n} \sum_{u\in\mathcal{V}} \frac{\alpha\Delta^2}{2\sigma^2} \sum_{w\in\mathcal{V}} \sum_{t\in\mathcal{P}^T_{\{v,w\}}} \frac{(W_{0:t})^2_{u,w}}{\|(W_{0:t})_w\|^2} \\
&= \frac{1}{n} \frac{\alpha\Delta^2}{2\sigma^2} \sum_{t\in\mathcal{P}^T_{\{v,w\}}} \sum_{w\in\mathcal{V}} \sum_{u\in\mathcal{V}} \frac{(W_{0:t})^2_{u,w}}{\|(W_{0:t})_w\|^2} \\
&= \frac{1}{n} \frac{\alpha\Delta^2}{2\sigma^2} \sum_{t\in\mathcal{P}^T_{\{v,w\}}} \sum_{w\in\mathcal{V}} 1 \\
&= \frac{\alpha\Delta^2 T_v}{2n\sigma^2} \,,
\end{aligned}$$

where $T_v = \sum_{w\in\mathcal{V}} |\mathcal{P}^T_{\{v,w\}}|$ is exactly the number of communications node v is involved in, up to time T. □

C Synchronous *Muffliato*

C.1 Utility Analysis (Theorem 2)

While the main text presents the result for the 1-dimensional case for simplicity, we prove here the convergence for the general case where each node holds a vector of dimension D. Theorem 2 is then a direct consequence of Equation (10) for $D = 1$.

Theorem 7 (Utility analysis)**.** *For any $T \geqslant 0$, the iterates $(x^T)_{T\geqslant 0}$ of* Muffliato *(Algorithm 1) verify, for λ_W defined in Definition 2 and $\bar{x} = \frac{1}{n}\sum_{v\in\mathcal{V}} x_v$:*

$$\frac{1}{2n}\sum_{v\in\mathcal{V}} \mathbb{E}\left[\left\|x_v^T - \bar{x}\right\|^2\right] \leqslant \left(\frac{1}{n}\sum_{v\in\mathcal{V}} \|x_v - \bar{x}\|^2 + D\sigma^2\right) e^{-T\sqrt{\lambda_W}} + \frac{D\sigma^2}{n}\,. \tag{10}$$

Proof. For $t \geqslant 0$ and $y \in \mathbb{R}^{\mathcal{V}\times D}$, using results from Berthier et al. [8], we have, for a vector $y \in \mathbb{R}^{\mathcal{V}\times D}$ such that $\sum_{v\in\mathcal{V}} y_v = 0$, $\|P_t(W)y\|^2 \leqslant 2(1-\sqrt{\lambda_W})^2\|y\|^2$. In particular:

$$\left\|P_t(W)(y - \bar{y}\mathbb{1}^\top)\right\| \leqslant 2(1-\sqrt{\lambda_W})^t\left\|y - \bar{y}\mathbb{1}^\top\right\|^2,$$

where $\mathbb{1}$ with the vector with all entries equal to 1. Since

$$x^t = P_t(W)\big(x + \mathcal{N}(0, \sigma^2 I_{\mathcal{V}\times D})\big)\,, \quad t \geqslant 0\,,$$

we obtain that, for $\eta \sim \mathcal{N}(0, \sigma^2 I_{\mathcal{V}\times D})$ and $\bar{\eta} = \frac{1}{n}\sum_{v\in\mathcal{V}} \eta_v \mathbb{1}^\top \in \mathbb{R}^{\mathcal{V}\times D}$, using bias-variance decomposition twice:

$$\begin{aligned}
\frac{1}{2}\mathbb{E}\left[\left\|x^t - \bar{x}\right\|^2\right] &= \frac{1}{2}\mathbb{E}\left[\left\|P_t(W)(x^{(0)} - \bar{x})\right\|^2\right] \\
&= \frac{1}{2}\|P_t(W)(x + \eta - \bar{x} - \bar{\eta})\|^2 + \frac{1}{2}\mathbb{E}\left[\|P_t(W)\bar{\eta}\|^2\right] \\
&\leqslant (1-\sqrt{\lambda_W})^t\mathbb{E}\left[\|x + \eta - \bar{x} - \bar{\eta}\|^2\right] + \frac{D\sigma^2}{2n} \\
&\leqslant (1-\sqrt{\lambda_W})^t\left(\mathbb{E}\left[\|x - \bar{x}\|^2\right] + nD\sigma^2\right) + \frac{D\sigma^2}{2n}\,.
\end{aligned}$$

□

The precision $\frac{3D\sigma^2}{n}$ is thus reached for

$$T^{\text{stop}}\big(W, (x_v)_{v\in\mathcal{V}}, D\sigma^2\big) \leqslant \sqrt{\lambda_W}^{-1}\ln\left(\frac{n}{D\sigma^2}\max\left(D\sigma^2, \frac{1}{n}\sum_{v\in\mathcal{V}}\|x_v - \bar{x}\|^2\right)\right)\,.$$

C.2 Privacy Analysis (Corollary 1)

Proof of Corollary 1. For a fixed gossip matrix W, we have $W_{0:t} = W^t$, so that Theorem 1 reads:

$$f(u,v) = \frac{\alpha\Delta^2}{2\sigma^2}\sum_{t<T}\sum_{w:\{v,w\}\in\mathcal{E}} \frac{(W^t)^2_{u,w}}{\|(W^t)_w\|^2}\,.$$

Since W is bi-stochastic, $\frac{(W^t)^2_{u,w}}{\|(W^t)_w\|^2} \leqslant n \times (W^t)^2_{u,w} = n\mathbb{P}\left(X^t = u | X^0 = w\right)^2$, using Cauchy-Schwarz inequality.

Summing over the neighbors $w \sim v$, we obtain, for $t < T$:

$$\begin{aligned}
\sum_{w\sim v} \frac{\alpha}{2\sigma^2} \sum_{w:\{v,w\}\in\mathcal{E}} \frac{(W^t)^2_{u,w}}{\|(W^t)_w\|^2} &\leqslant \frac{\alpha n}{2\sigma^2} \sum_{w\sim v} \mathbb{P}\left(X^t = u | X^0 = w\right)^2 \\
&\leqslant \frac{\alpha n}{2\sigma^2} \left(\sum_{w\sim v} \mathbb{P}\left(X^t = u | X^0 = w\right) \right)^2 \\
&\leqslant \frac{\alpha n}{2\sigma^2} \frac{1}{\min_{w\sim v} W^2_{v,w}} \left(\sum_{w\sim v} W_{v,w} \mathbb{P}\left(X^t = u | X^0 = w\right) \right)^2 \\
&\leqslant \frac{\alpha n}{2\sigma^2} \frac{1}{\min_{w\sim v} W^2_{v,w}} \mathbb{P}\left(X^{t+1} = u | X^0 = v\right)^2 ,
\end{aligned}$$

where the last line is obtained by observing that:

$$\sum_{w\sim v} W_{v,w} \mathbb{P}\left(X^t = u | X^0 = w\right) = \mathbb{P}\left(X^{t+1} = u | X^0 = v\right) ,$$

by conditioning on the first step of the random walk. This leads to Corollary 1. □

C.3 First Line of Table 1

The results in the first line of Table 1 are obtained by observing that v is involved in $d_v T$ communications up to time T, leading to $\bar{\varepsilon}_v = \frac{\alpha\Delta^2 d_v T}{2\sigma^2 n^2}$. Using Theorem 2, we have a utility of $3\sigma^2/n$ for $T^{\text{stop}} = \lambda_W^{-1/2} \ln\left(\frac{n}{\sigma^2} \max\left(\sigma^2, \frac{1}{n}\sum_{v\in\mathcal{V}} \left\|x_v^0 - \bar{x}\right\|^2\right)\right)$ steps. Thus, imposing $\bar{\varepsilon}_v \leqslant \varepsilon$ for a fixed $\varepsilon > 0$ gives us $\sigma^2 = \frac{\alpha\Delta^2 d T^{\text{stop}}}{2\sigma^2 n^2}$, leading to a utility of

$$\tilde{\mathcal{O}}\Big(\frac{\alpha\Delta^2 d}{2\sigma^2\sqrt{\lambda_W}}\Big) .$$

We then instantiate this formula on graphs with known spectral gaps, as described for instance in Mohar et al. [46].

D Randomized *Muffliato*

D.1 Utility Analysis (Theorem 3)

As in the synchronous case, we prove a more general convergence result that holds for D-dimensional inputs. Then, Theorem 3 follows directly from (11).

Theorem 8 (Utility analysis)**.** *For any $T \geqslant 0$, the iterates $(x^T)_{T\geqslant 0}$ of randomized* Muffliato *(Algorithm 2) verify:*

$$\frac{1}{2n}\sum_{v\in\mathcal{V}} \mathbb{E}\left[\left\|x_v^T - \bar{x}\right\|^2\right] \leqslant \left(\frac{1}{n}\sum_{v\in\mathcal{V}} \left\|x_v^0 - \bar{x}\right\|^2 + D\sigma^2\right) e^{-T\lambda_2(p)} + \frac{D\sigma^2}{n} . \tag{11}$$

Proof of Theorem 3. For $t \geqslant 0$ and $y \in \mathbb{R}^{\mathcal{V}\times D}$, using results from Boyd et al. [11], we have:

$$\mathbb{E}\left[\left\|W(t)(y - \bar{y}\mathbb{1}^\top)\right\|^2\right] \leqslant (1-\lambda(p))^t \left\|y - \bar{y}\mathbb{1}^\top\right\|^2 ,$$

where $\mathbb{1}$ with the vector with all entries equal to 1 and $\bar{y} = \frac{1}{n}\sum_{v\in\mathcal{V}} y_v$. The rest of the proof follows as in the proof of Theorem 2 with two bias-variance decompositions. □

The precision $\frac{2D\sigma^2}{n}$ is thus reached for

$$T^{\text{stop}}\big(W, (x_v)_{v\in\mathcal{V}}, \sigma^2\big) \leqslant \lambda(p)^{-1} \ln\left(\frac{n}{D\sigma^2} \max\left(D\sigma^2, \frac{1}{n}\sum_{v\in\mathcal{V}} \|x_v - \bar{x}\|^2\right)\right) . \tag{12}$$

D.2 Privacy Analysis

In terms of privacy, randomized *Muffliato* satisfies the following guarantees, obtained by applying Theorem 1.

Corollary 2. *After T iterations of randomized* Muffliato*, and conditionally on the edges sampled, node $u \in \mathcal{V}$ is $(\alpha, \varepsilon^T_{u\to v}(\alpha))$-PNDP with respect to v, with:*

$$\varepsilon^T_{u\to v}(\alpha) \leqslant \frac{\alpha\Delta^2}{2\sigma^2} \sum_{w\sim v} \sum_{t\in\mathcal{P}^T_{\{v,w\}}} \frac{(W_{0:t})^2_{uw}}{\|(W_{0:t})_w\|^2} \,.$$

Taking the mean over $u \neq v$ yields:

$$\bar{\varepsilon}_v = \frac{1}{n} \sum_{u\in\mathcal{V}\setminus\{v\}} \varepsilon^T_{u\to v}(\alpha) \leqslant \frac{\alpha\Delta^2}{2n\sigma^2} T_v \,,$$

where $T_v = \sum_{t<T} \sum_{w\sim v} \mathbb{1}_{\{\{v,w\}=\{v_t,w_t\}\}}$ the number of communications node v is involved in up to time T.

Note that T_v is a Binomial random variable of parameters (T, π_v) where $\pi_v = \sum_{w\sim v} p_{\{v,w\}}$.

We now explain how we obtain the second line of Table 1.

Note that a choice $p_{\{v,w\}} = 2W_{v,w}/n$ for some given gossip matrix W yields probability activations. For the sake of comparison with *Muffliato* with a fixed matrix, we place ourselves in this case. This leads to $\pi_v = 2/n$, so that

$$\mathbb{E}\left[\bar{\varepsilon}_v\right] = \frac{\alpha\Delta^2 T}{2n^2\sigma^2} \,,$$

and for any $C > 0$,

$$\mathbb{P}\left(T_v - \mathbb{E}T_v \geqslant C\right) \leqslant \exp\big(-\frac{C^2}{T}\big) \,,$$

using Hoeffding's inequality. We take as time-horizon $T = T^{\text{stop}} \geqslant 1/\lambda(p)$ defined in Theorem 3, leading to

$$\mathbb{P}\left(T_v - \mathbb{E}T_v \geqslant C\right) \leqslant \exp\big(-\frac{C^2\lambda_W}{n}\big) \,, \quad \mathbb{E}\left[\bar{\varepsilon}_v\right] = \tilde{\mathcal{O}}\big(\frac{\alpha\Delta^2}{2n\sigma^2\lambda_W}\big) \,,$$

since $\lambda(p) = \frac{2\lambda_W}{n}$ in our case.

The same methodology as in the synchronous case (imposing $\bar{\varepsilon}_v \leqslant \varepsilon$ for the time horizon T^{stop}, deriving σ^2 from this and thus the resulting utility) leads to the second line of Table 1.

E *Muffliato* on Private Random Graphs (Theorem 4)

We fix all nodes, and in particular u the attacked node, and v the observer. We assume that G is drawn randomly. Edges $\{v, w\}$ are drawn independently from one another. The result we prove below is just slightly more general than Theorem 4 that is recovered for $\mathbb{P}\left(\{u, v\} \in \mathcal{E}\right) = q$ (Erdös-Rényi random graph).

We make the following assumption: node v is only aware of its direct neighbors in the topology of graph G, and conditionally on $\{v\} \cup \mathcal{N}(v)$, the law of the graph is invariant under any permutation over the set $\mathcal{V} \setminus (\{v\} \cup \mathcal{N}(v))$.

Theorem 9 (*Muffliato* with a random graph)**.** *Let $\alpha > 1$, $T \geqslant 0$, $\sigma^2 \geqslant \frac{\Delta^2\alpha(\alpha-1)}{2}$ and let the above assumptions on G be satisfied. After running Algorithm 1 with these parameters, node u is $(\alpha, \varepsilon^T_{u\to v}(\alpha))$-PNDP with respect to v, with:*

$$\varepsilon^T_{u\to v}(\alpha) \leqslant \begin{cases} \dfrac{\alpha\Delta^2}{2\sigma^2} & \textit{with probability } \mathbb{P}\left(\{u, v\} \in \mathcal{E}\right) \,, \\ \dfrac{\alpha\Delta^2}{\sigma^2}\dfrac{Td_v}{n-d_v} & \textit{with probability } 1 - \mathbb{P}\left(\{u, v\} \in \mathcal{E}\right) \,. \end{cases}$$

Proof. If $\{u,v\} \in \mathcal{E}$, we cannot do better than $\varepsilon^T_{u\to v} \leqslant \frac{\alpha\Delta^2}{2\sigma^2}$: v only sees $x_u^{(0)} + \mathcal{N}(0,\sigma^2)$ and then next messages can be seen as post-processing of this initial message and thus do not induce further loss. This happens with probability $\mathbb{P}\left(\{u,v\} \in \mathcal{E}\right)$.

Now, we reason conditionally on $\{v\} \cup \mathcal{N}(v)$ and $u \notin \mathcal{N}(v)$. Using the averaged formula (6), we have:

$$\frac{1}{n}\left(\sum_{w\in\mathcal{N}(v)} \frac{\alpha\Delta^2}{2\sigma^2} + \sum_{w\in\mathcal{V}\setminus(\mathcal{N}(v)\cup\{v\})} \varepsilon^T_{w\to v}(\alpha)\right) \leqslant \frac{\alpha\Delta^2 d_v T}{2n\sigma^2}\,.$$

Here we adapted the proof of the formula: to obtain the right-handside, a value $\varepsilon^T_{w\to v}$ bigger than $\frac{\alpha\Delta^2}{2\sigma^2}$ was taken, so that the formula above is also true. Then, using the fact that node v only sees its neighbors, we can use Lemma 2 that allows us to take the mean conditionally on $v \cup \mathcal{N}(v)$ (for $\sigma^2 \geqslant \frac{\Delta^2\alpha(\alpha-1)}{2}$), leading to

$$\frac{1}{n}\left(\sum_{w\in\mathcal{N}(v)} \frac{\alpha\Delta^2}{2\sigma^2} + \sum_{w\in\mathcal{V}\setminus(\mathcal{N}(v)\cup\{v\})} \mathbb{E}\left[\varepsilon^T_{w\to v}(\alpha) \mid v\cup\mathcal{N}(v)\right]\right) \leqslant \frac{\alpha\Delta^2 d_v T}{n\sigma^2}\,.$$

In fact, we write it with the expected value, but all nodes are equal since node v only sees its neighbors. Using the invariance under permutation of $\mathbb{E}\left[\varepsilon^T_{w\to v}(\alpha) v\cup\mathcal{N}(v)\right]$ over $w \in \mathcal{V}\setminus(\mathcal{N}(v)\cup\{v\})$, we have that:

$$\frac{1}{n}\left(\sum_{w\in\mathcal{N}(v)} \frac{\alpha\Delta^2}{2\sigma^2} + (n-d_v)\varepsilon^T_{u\to v}\right) \leqslant \frac{\alpha\Delta^2 d_v T}{n\sigma^2}\,.$$

Rearranging this inequality gives the result. □

F Differentially Private Decentralized Optimization

We consider Algorithm 3 with general time-varying matrices W_t. We assume that for all $t \geqslant 0$, $\lambda_{W_t} \geqslant \lambda$ for some fixed $\lambda > 0$. An instance of this setting is to sample different Erdös-Rényi random graphs at each communication round and adapt W_t accordingly. Such graphs have a spectral gap that concentrates around 1, so that for $\lambda = 1/2$, with high probability $\lambda_{W_t} \geqslant \lambda$ will be verified [32].

F.1 Proof of Theorem 5 (Utility Analysis)

As before, we have a more general convergence result.

Theorem 10 (Utility analysis of Algorithm 3)**.** *Let* $K \geqslant \left\lceil \sqrt{\lambda}^{-1} \ln\left(\max\left(n, \frac{\bar{\zeta}^2}{D\sigma^2+\bar{\rho}^2}\right)\right)\right\rceil$. *For a suitable choice of step size parameters, the iterates* $(\theta^t)_{t\geqslant 0}$ *generated by Algorithm 3 verify:*

$$\mathbb{E}\left[\phi(\tilde{\theta}^T) - \phi(x^\star)\right] \leqslant \tilde{\mathcal{O}}\left(\frac{\bar{\rho}^2 + D\sigma^2}{n\mu T} + L\left\|\theta^0 - \theta^\star\right\|^2 e^{-\frac{T}{2\kappa}}\right)\,,$$

where $\tilde{\theta}^T = \frac{\sum_{t<T}\omega^t\bar{\theta}^t}{\sum_{t<T}\omega^t}$ *is a weighted average along the trajectory of the means* $\bar{\theta}^t = \frac{1}{n}\sum_{v\in\mathcal{V}}\theta^t_v$.

The proof of Theorem 10 is a direct consequence of Theorem 2 in Koloskova et al. [38], and especially the formula in their Appendix A.4. We apply their result with $\bar{\rho}^2 + D\sigma^2$ instead of their $\bar{\sigma}^2$, $\tau = 1$ (no varying topology), and p such that $1 - p = 2(1-\sqrt{\lambda})^K \leqslant 2\min(\frac{1}{n}, \frac{D\sigma^2+\bar{\rho}^2}{\bar{\zeta}^2})$.

F.2 Proof of Theorem 6 (Privacy Analysis)

The function Lipschitzness can in fact be replaced by a more general assumption.

Assumption 2. *We assume that, for some* $\Delta^2_\phi > 0$, *for all* v *in* $\mathcal{V}$, *and for all adjacent datasets* $\mathcal{D} \sim_v \mathcal{D}'$ *on* v, *we have:*

$$\sup_{\theta\in\mathbb{R}^D}\ \sup_{(x_v,x'_v)\in\mathcal{D}_v\times\mathcal{D}'_v} \left\|\nabla_x\ell(\theta,x_v) - \nabla_x\ell(\theta,x'_v)\right\|^2 \leqslant \Delta^2_\phi\,.$$

Theorem 11 (Privacy analysis of Algorithm 3)**.** *Let $(W_t)_{0 \leqslant t < T}$ be a sequence of gossip matrices of spectral gap larger than λ. Let u and v be two distinct nodes in $\mathcal{V}$. After T iterations of Algorithm 3 with $K \geqslant 1$, node u is $(\alpha, \varepsilon^T_{u \to v}(\alpha))$-PNDP with respect to v, with:*

$$\varepsilon^T_{u \to v}(\alpha) \leqslant \frac{\Delta_\phi^2 \alpha}{2\sigma^2} \sum_{t=0}^{T-1} \sum_{k=0}^{K-1} \sum_{w : \{v,w\} \in \mathcal{E}_t} \frac{(W_t^k)^2_{u,w}}{\left\| (W_t^k)_w \right\|^2} \,. \tag{13}$$

Thus, for any $\varepsilon > 0$, Algorithm 3 with $T^{\text{stop}}(\kappa, \sigma^2, n)$ steps and for K as in Theorem 5, there exists f such that the algorithm is (α, f)-pairwise network DP, with:

$$\forall v \in \mathcal{V}\,, \quad \overline{\varepsilon}_v \leqslant \varepsilon \quad \textit{and} \quad \mathbb{E}\left[\phi(\tilde{\theta}^{\text{out}}) - \phi(\theta^\star)\right] \leqslant \tilde{\mathcal{O}}\left(\frac{\alpha D \Delta_\phi^2 \bar{d}}{n^2 \mu \varepsilon \sqrt{\lambda}} + \frac{\bar{\rho}^2}{nL}\right) ,$$

where $\bar{d} = \sup_{v \in \mathcal{V}} \bar{d}_v$ for $\bar{d}_v = \frac{1}{T} \sum_{t<T} |\{w \in \mathcal{V} : \{v,w\} \in \mathcal{E}_t\}|$ the mean degree of node v throughout time.

Proof of Theorem 11. The information leaked by u to v up to iteration T of Algorithm 3 consists in the T (stochastic) gradients locally computed at node u and gossiped through the graph, using the *Muffliato* algorithm. Using Theorem 1 (with Assumption 1 satisfied using Assumption 2) and a post processing inequality, round of communication t leads to a privacy leak of:

$$\frac{\Delta_\phi^2 \alpha}{2\sigma^2} \sum_{k=0}^{K-1} \sum_{w : \{v,w\} \in \mathcal{E}_t} \frac{(W_t^k)^2_{u,w}}{\left\| (W_t^k)_w \right\|^2} \,,$$

where $\mathcal{E}_t$ are the edges of the graph drawn at time t. We obtain the first inequality of Theorem 6 by summing this over $t < T$.

For the second inequality, we have, by summing:

$$\overline{\varepsilon}_v = \frac{1}{n} \sum_{u \neq v} \varepsilon^T_{u \to v}(\alpha) \leqslant \frac{KT\bar{d}_v \Delta_\phi^2 \alpha}{2n\sigma^2} \leqslant \frac{KT\bar{d} \Delta_\phi^2 \alpha}{2n\sigma^2} \,.$$

In order to reach a precision $\frac{D\sigma^2 + \bar{\rho}^2}{n}$, are required $T = \mathcal{O}(\kappa \ln(D\sigma^2/n))$ iterations. Using $K = \tilde{\mathcal{O}}(1/\sqrt{\lambda})$, we have:

$$\overline{\varepsilon}_v = \mathcal{O}\left(\frac{\bar{d} \Delta_\phi^2 \alpha}{2n\sigma^2} \kappa \sqrt{\lambda}^{-1} \ln(D\sigma^2/n)\right) .$$

Taking σ^2 such that $\frac{\bar{d} \Delta_\phi^2 \alpha}{2n\sigma^2} \kappa \sqrt{\lambda}^{-1} \ln(\sigma^2/n) = \varepsilon$ yields the desired result. □

G Extensions to Collusion and Group Privacy

In this section, we discuss natural extensions of our privacy definitions to the case of colluding nodes, and to group privacy.

G.1 Presence of Colluding Nodes

G.1.1 Definitions

The notions of pairwise network DP we introduced in Section 2.3 can naturally be extended to account for potential collusions. For $V \subset \mathcal{V}$ a set of colluding nodes, we define the *view* of the colluders as:

$$\mathcal{O}_V\big(\mathcal{A}(\mathcal{D})\big) = \bigcup_{v \in V} \mathcal{O}_V\big(\mathcal{A}(\mathcal{D})\big) \,,$$

or equivalently as:

$$\mathcal{O}_V\big(\mathcal{A}(\mathcal{D})\big) = \{(u, m(t), v) \in \mathcal{A}(\mathcal{D}) \quad \text{such that} \quad \{u,v\} \in \mathcal{E}\,, v \in V\} \,.$$

Below, $\mathcal{P}(\mathcal{V})$ denotes the powerset of $\mathcal{V}$.

Definition 6 (Pairwise Network DP with colluders)**.** *For $f : \mathcal{V} \times \mathcal{P}(\mathcal{V}) \to \mathbb{R}^+$, an algorithm $\mathcal{A}$ satisfies (α, f)-pairwise network DP if for all users $u \in \mathcal{V}$, pairs of neighboring datasets $D \sim_u D'$, and any potential set of colluders $V \in \mathcal{V}$ such that $u \notin V$, we have:*

$$D_\alpha\big(\mathcal{O}_V(\mathcal{A}(\mathcal{D})) \,||\, \mathcal{O}_V(\mathcal{A}(\mathcal{D}'))\big) \leqslant f(u, V)\,. \tag{14}$$

We note $f(u, V) = \varepsilon_{u\to V}$ the privacy leaked to the colluding nodes V from u and say that u is $(\alpha, \varepsilon_{u\to V})$-PNDP with respect to V if only inequality (14) holds for $f(u, V)$. Finally, if for a function $f : \mathcal{V} \times \mathcal{P}(\mathcal{V}) \to \mathbb{R}$, inequality (14) holds for all $(u, V) \in \mathcal{V} \times \mathcal{P}(\mathcal{V})$ such that $u \notin V$, we say that $\mathcal{A}$ is (α, f)-pairwise NDP.

This definition quantifies the privacy loss of a node u with respect to the collusion of any possible subset V of nodes, and thus generalizes the definition of the main text (which corresponds to restricting V such that $|V| = 1$).

The proofs of this section are actually direct consequences of the proof techniques of our results without colluders, by replacing $\mathcal{O}_v$ (the view of a colluder) by $\mathcal{O}_V$ (the view of the colluding set). Roughly speaking, V can be seen as a unique abstract node, resulting from the fusion of all its nodes.

G.1.2 Adapting Theorem 1 and the Resulting Corollaries

For $w \in \mathcal{V}$ and $V \in \mathcal{P}(\mathcal{V})$, let $\mathcal{P}^t_{\{V,w\}} = \{s < t \,:\, \exists v \in V\,, \{v, w\} \in \mathcal{E}_s\}$ the times (up to time t) at which an edge $\{v, w\}$ for any $v \in V$ is activated *i.e.* the times at which there is a communication between w and a colluder. For $t \geqslant 0$ and $V \subset \mathcal{V}$, let $\mathcal{N}_t(V)$ be the neighbors in G_t of the colluders set V, defined as:

$$\mathcal{N}_t(V) = \{w \in \mathcal{V} \setminus V \,|\, \exists v \in V\,, \{v, w\} \in \mathcal{E}_t\}\,.$$

For $T \geqslant 1$, $\sum_{t<T} |\mathcal{N}_t(V)|$ is thus the total number of communications in which colluders are involved with.

Theorem 12. *Assume that Assumption 1 holds. Let $T \geqslant 1$, $u \in \mathcal{V}$ and $V \subset \mathcal{V}$ such that $u \notin V$, and $\alpha > 1$. Then the algorithm $\mathcal{A}^T$ is (α, f)-PNDP with::*

$$f(u, V) \leqslant \frac{\alpha\Delta^2}{2\sigma^2} \sum_{w\in\mathcal{V}} \sum_{t\in\mathcal{P}^T_{\{V,w\}}} \frac{(W_{0:t})^2_{u,w}}{\|(W_{0:t})_w\|^2}\,. \tag{15}$$

Consequently, there exists f such that $\mathcal{A}^T$ is (α, f)-PNDP with:

$$\bar{\varepsilon}_V = \frac{1}{n} \sum_{u\in\mathcal{V}\setminus V} f(u, V) \leqslant \frac{\alpha\Delta^2}{2n\sigma^2} \sum_{t<T} |\mathcal{N}_t(V)|\,, \tag{16}$$

where $\sum_{t<T} |\mathcal{N}_t(V)|$ is the total number of communications a colluding set V is involved with, up to time T.

We now consider the synchronous *Muffliato* algorithm (Algorithm 1) with colluders.

Corollary 3. *Let $u \in \mathcal{V}$ and $V \in \mathcal{P}(\mathcal{V})$ such that $u \notin V$, $\alpha > 0$. After T iterations of Algorithm 1, node u is $(\alpha, \varepsilon^T_{u\to V}(\alpha))$-PNDP with respect to V, with:*

$$\varepsilon^T_{u\to V}(\alpha) \leqslant \frac{\alpha n}{2\sigma^2} \max_{w\sim V} W^{-2}_{v,w} \sum_{t=1}^{T} \mathbb{P}\left(X^t \in V | X^0 = u\right)^2\,,$$

where $(X_t)_t$ is the random walk on graph G, with transitions W, and $w \sim V$ if $w \notin V$ and if there exists $v \in \mathcal{V}$ such that $\{v, w\} \in \mathcal{E}$.

Corollary 4. *There exists $f : \mathcal{V} \times \mathcal{P}(\mathcal{V}) \to \mathbb{R}^+$ such that Algorithm 1 after T steps is (α, f)-PNDP with the following privacy-utility guarantees for any $V \subset \mathcal{V}$:*

$$\bar{\varepsilon}_V = \frac{1}{n} \sum_{u\in\mathcal{V}\setminus V} f(u, V) \leqslant \varepsilon \quad , \qquad \frac{1}{2n} \sum_{v\in\mathcal{V}} \mathbb{E}\left[\left\|x^{\text{out}}_v - \bar{x}\right\|^2\right] \leqslant \tilde{\mathcal{O}}\left(\alpha \frac{d_V \Delta^2}{\varepsilon n^2 \sqrt{\lambda_W}}\right)\,,$$

where x^{out} is the output of Algorithm 1 after $T^{\text{stop}}(x, W, \sigma^2)$ steps for $\sigma^2 = \frac{d_V \Delta^2}{2\alpha\varepsilon}$, and $\tilde{\mathcal{O}}$ hides logarithmic factors in n and ε. d_V is the degree of set V, defined as the number of $w \in \mathcal{V} \setminus V$ such that there exists $v \in V$, $\{v, w\} \in \mathcal{E}$.

Corollary 5 (*Muffliato* on a random graph with collusions)**.** *Let $\alpha > 1$, $T \geqslant 0$, $\sigma^2 \geqslant \frac{\Delta^2\alpha(\alpha-1)}{2}$ and $q = c\frac{\ln(n)}{n}$ for $c > 1$. Let $u \in \mathcal{V}$ and $V \in \mathcal{P}(\mathcal{V})$ such that $u \notin V$. After running Algorithm 1 on an Erdos Rényi random graph of parameters (n, q) and under the assumptions of Theorem 4, node u is $(\alpha, \varepsilon^T_{u\to v}(\alpha))$-PNDP with respect to colluders V, with:*

$$\varepsilon^T_{u\to V}(\alpha) \leqslant \begin{cases} \dfrac{\alpha}{2\sigma^2} & \text{with probability } 1-(1-q)^{|V|} \\ \dfrac{\alpha}{\sigma^2}\dfrac{Td_V}{n-d_V} & \text{with probability } (1-q)^{|V|} \end{cases} .$$

G.1.3 Compensating for Collusions with Time-Varying Graph Sampling

We now consider the decentralized optimization problem of Section 4 in the presence of colluders, and analyze its privacy with time-varying graph sampling as in Appendix F.2.

The motivation for this is that, if the graph is fixed, node u will suffer from poor privacy guarantees (the same as in LDP) with respect to the colluding set V as soon as u is in $\mathcal{N}(V) = \mathcal{N}_0(V)$ (i.e., one of colluders in V is a neighbor of u). Even if the graph is sampled randomly, this will happen with probability that increases with $|V|$. In contrast, for time-varying random graphs sampled independently at each communication round and for sufficiently many communication rounds (i.e., large enough condition number κ), it becomes unlikely that u is in $\mathcal{N}_t(V)$ for many rounds t, and therefore the privacy guarantees with respect to V can improve.

Below, we consider Algorithm 3 with time-varying graphs (and associated gossip matrices $(W_t)_{t\geqslant 0}$) sampled in an *i.i.d.* fashion at each communication round as Erdös-Rényi graphs of parameters $n, q = \frac{c\ln(n)}{n}$ for some $c > 1$, such that they verify $\lambda_{W_t} \geqslant \lambda > 0$ for all t (as noted before, this happens with high probability for λ of order 1 [32]). In this context, we have the following result.

Proposition 1. *Let $\alpha > 1$, $T \geqslant 0$, $\sigma^2 \geqslant \frac{\Delta^2\alpha(\alpha-1)}{2}$. Let $u \in \mathcal{V}$ and $V \in \mathcal{P}(\mathcal{V})$ such that $u \notin V$. After running Algorithm 3 under the assumptions described above and the function assumptions of Theorem 6, node u is $(\alpha, \varepsilon^T_{u\to v}(\alpha))$-PNDP with respect to colluders V, with:*

$$\varepsilon^T_{u\to V}(\alpha) \leqslant \frac{\Delta_\phi^2\alpha}{\sigma^2}\sum_{t=0}^{T^{\text{stop}}}\beta_t + (1-\beta_t)\frac{K|\mathcal{N}_t(V)|}{n-|\mathcal{N}_t(V)|} ,$$

where $T^{\text{stop}} = \tilde{\mathcal{O}}(\kappa)$ and $K = \tilde{\mathcal{O}}(1/\sqrt{\lambda})$ (see Theorem 5), $(\beta_t)_t$ are i.i.d. *Bernoulli random variables of parameter $\mathbb{P}(\exists v \in \mathcal{V}, \{u, v\} \in \mathcal{E}_t) = 1-(1-q)^{|V|}$, and $|\mathcal{N}_t(V)|$ is the number of neighbors of V in the graph sampled at iteration t, of order $\frac{c|V|\ln(n)}{n}$.*

G.1.4 Discussion

Generally speaking, our bounds degrade in presence of colluding nodes. This is a fundamental limitation of our approach that considers only privacy amplification due to decentralization. By definition, our privacy guarantees can only provide amplification as long as the view of the attackers is smaller than the one of the omniscient attacker considered in local differential privacy, i.e $\mathcal{O}_V(\mathcal{A}^T) \subsetneq \mathcal{A}^T$. A condition for having equality corresponds to observing all messages that are transmitted. In the case of a fixed graph, this can be characterized by the fact that V contains a dominating set for the graph. For Erdos Rényi graphs or exponential graphs, there exists dominating sets of size $\mathcal{O}(\log n)$, thus it is meaningless to expect guarantees for all possible sets of colluding nodes of that size. However, if some/most colluding nodes are actually far from the target node u in the graph, then good privacy amplification can still be achieved. This can be precisely measured by Equation 15.

G.2 Group Privacy

Symmetrically to the problem of collusions, where there are several attackers, one can study group privacy, where privacy guarantees are computed towards a group rather than a single individual. This is useful when some users have correlated data (e.g., close family members). The usual notion of group privacy [see e.g., 20] would provide a guarantee against all groups of users of a given size. However, this standard definition would provide pessimistic guarantees in some cases as it does not take advantage of the fact that groups may correspond to specific subgraphs. Hence, we propose the following definition.

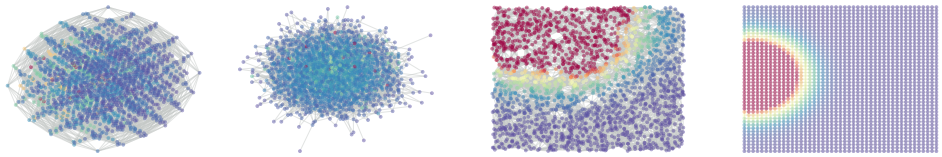

Figure 2: Level of the privacy loss for each node (color) with respect to a fixed node in the graph. These graphs corresponds to the graphs used in Figure 1(a): from left to right, exponential graph, Erdos-Renyi graph, geometric random graph and grid.

Definition 7 (Group Pairwise Network DP). *For $f : \mathcal{V} \times \mathcal{P}(\mathcal{V}) \to \mathbb{R}^+$, an algorithm $\mathcal{A}$ satisfies (α, f)-group pairwise network DP if for all set of users $U \subset \mathcal{V}$, pairs of neighboring datasets $D \sim_U D'$, and any vertex $v \in \mathcal{V}$, we have:*

$$D_\alpha\big(\mathcal{O}_v(\mathcal{A}(\mathcal{D})) \,||\, \mathcal{O}_v(\mathcal{A}(\mathcal{D}'))\big) \leqslant f(U, v)\,. \tag{17}$$

The modifications of the theorems are similar to the case of collusion, summing over the nodes in U (instead of summing over the nodes in V for the case of collusion). The following theorem gives the general case corresponding to Theorem 1. The other results can be adapted in the same way.

Theorem 13. *Let $T \geqslant 1$ and denote by $\mathcal{P}^T_{\{v,w\}} = \{s < T \,:\, \{v, w\} \in \mathcal{E}_s\}$ the set of time-steps with communication along edge $\{v, w\}$. Under Assumption 1, $\mathcal{A}^T$ is (α, f)-group PNDP for with:*

$$f(U, v) = \frac{\alpha\Delta^2}{2\sigma^2} \sum_{w\in\mathcal{V}} \sum_{t\in\mathcal{P}^T_{\{v,w\}}} \frac{\sum_{u\in U}(W_{0:t})^2_{u,w}}{\|(W_{0:t})_w\|^2}\,. \tag{18}$$

Note that in some configurations, this bound bound is clearly sub-optimal, as summing does not take into account cases where the information gathered by some of the nodes in the group can be seen as a post-processing of information received by other nodes in the group. In such cases, analyzing group privacy by considering a modified graph where the nodes in the group are merged into a single node, with edge/weights adjusted accordingly, would yield better results.

H Additional Numerical Experiments

H.1 Extra Synthetic graphs

Figure 1(a) summarizes the result of *Muffliato* according to the shortest path length. However, other characteristics of the topology can play a role in the privacy leakage. Thus, we show the graph representation for each of the synthetic graphs we considered in Figure 2.

We also report in Figure 3 how privacy guarantees improve when n increases for the exponential graph. We see that privacy guarantees improve with n: distance between nodes increases, but also the number of nodes with whom the contribution of a specific node is mixed. This is especially significant for pairs of nodes that are not direct neighbors but at short distance of each other.

H.2 Proof of Fixed Privacy Loss for Exponential Graphs

For an exponential graph, the pairwise privacy loss is fully determined by the shortest length path, i.e $f(u, v) = g(d(u, v))$ where $d : \mathcal{V} \to \mathbb{N}$ is the function that returns the length of the shortest path between u and v.

This result is a consequence of the invariance per vertex permutation in the hypercube. For the hypercube with 2^m vertices, each vertex can be represented by a m-tuple in $\{0, 1\}$, where there is an edge if and only if two vertices share all values of their tuple but one. Let us now fix two pairs of vertices (u, v) and (u', v') with the same distance between them. To prove that their privacy loss is the same, it is sufficient to exhibit an graph isomorphism Φ that sends (u, v) on (u', v').

By construction, $d(u, v)$ corresponds to the number of coordinates that differ between their tuple, and the same holds for (u', v'). The set of equal coordinates $Fix(u, v)$ is thus of the same size

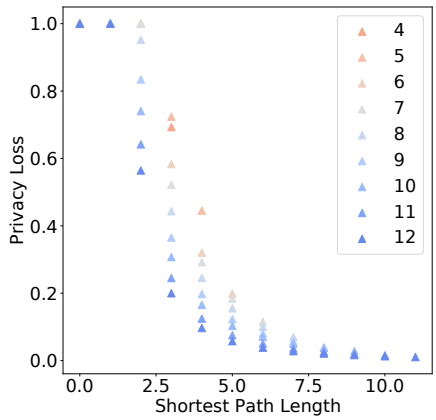

Figure 3: Privacy loss for the exponential graph with respect to the number of nodes n (following powers of 2).

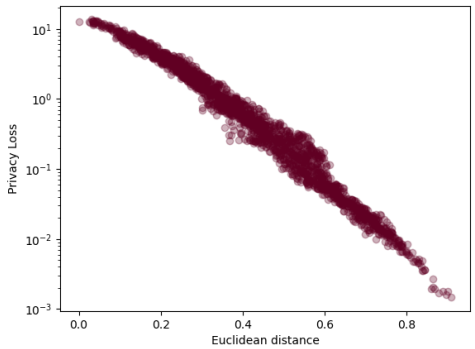

Figure 4: Privacy towards all the nodes as function of the Euclidean distance in a random geometric graph. We see a high level of correlation between the Euclidean distance and the privacy loss.

$m - d(u, v)$ than $Fix(u', v')$. Hence we can construct a bijective function b of the coordinates that is stable for the set of fixed coordinates

$$b(Fix(u, v)) = Fix(u', v') \quad b(\{1, 2, \ldots, m\} \setminus Fix(u, v)) = \{1, 2, \ldots, m\} \setminus Fix(u', v')$$

Finally, noting that 0 and 1 play the same role, we match each coordinate accordingly to the value defined by our couple. We thus define our isomorphism per coordinate $\Phi(w) = (\Phi_1(w), \ldots, \Phi_m(w))$ with $\Phi_i(w) = s(w_{b^{-1}(i)})$ where s is the identity function if $u_{b^{-1}(i)} = u_i$ and the swap function otherwise. This function is clearly an isomorphism: by construction it is a bijection, and the edges still exist if and only if the vertices differ on only one coordinate. We have $\Phi(u) = u'$ and $\Phi(v) = v'$ and thus the privacy loss is equal between the two pairs of vertices.

H.3 Random Geometric Graphs

Geometric graphs are examples of possible use cases of Pairwise Network Differential Privacy. Constructing edges when nodes are at a distance below a given threshold naturally models short-range wireless communications such as Bluetooth. In this situation, the Euclidean distance between nodes is a convenient indicator for setting the privacy loss. Indeed, it is a parameter that we can measure, and it can match the users expectations in terms of privacy loss. For instance, if direct neighbors in the graph correspond to people within 5 meters around the sender, some information are bound to be available to them independently of what may be revealed by the communication itself: sensitive attributes such a gender, age, or overall physical fitness are leaked simply from physical proximity. However, the user might have stronger privacy expectations for people far away. Hence, having privacy guarantees as function of the Euclidean distance can be particularly interesting.

Figure 5: Privacy loss on the 9 other Facebook Ego graphs, following the same methodology as in Figure 1(b).

Our experiments show that the privacy loss is extremely well correlated to the Euclidean distance, as represented in Figure 4. It is thus possible to design algorithms where one could impose Pairwise Network DP for a function $f(u, v) = g(\|z_u - z_v\|)$ where g is a non-increasing function and z_u and z_v are the geolocation of nodes u and v.

H.4 Facebook Ego Graphs

We report figure on the other nine graphs of the Facebook Ego dataset, following the same methodology and scale. Across these graphs, we can see that privacy losses depending on visible communities is consistent through datasets, and become more consistent as the number of nodes increase.

H.5 Logistic Regression on Houses Dataset

We report in Table 2 the parameters used in the experiments of Figure 1(c).

Table 2: Parameter for the logistic regression

Parameters	Value
# of trials	10
Step-size	0.7
# of nodes	2000 or 4000
probability of edges q	$\log(n)/n$
score	Mean accuracy

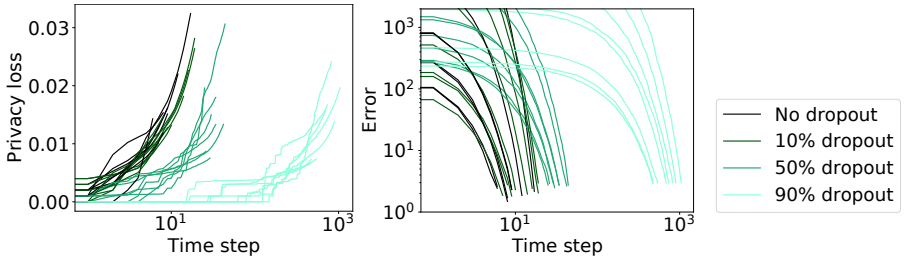

Figure 6: Simulation (10 runs) of a dropout scenario (with different levels of dropout) with $n = 1000$, where at each gossip step an Erdös-Renyi graph with parameter $p = 0.002$ over the set of available users. (a) Left: Privacy loss of a node across iterations ;(b) Right: Convergence of the same runs to the mean.

H.6 Modeling User Dropout using Time-Varying Graphs

In Theorem 1, we analyzed the very generic case where the gossip matrix is arbitrary at each time-step. Then, for deriving the convergence rate and obtaining closed-form privacy-utility trade-offs and using acceleration, we focused on fixed gossip matrices until the convergence of each gossip averaging step. In this subsection, we show empirically that we can still reach a good privacy-utility trade-offs when the gossip matrix change at each communication.

In particular, our experiment focuses on modelling user dropouts. Specifically, at each time step, the availability of each node is modeled by an independent Bernouilli random variable and we draw a new Erdos Renyi graph over the set of available nodes. We assume that at each step, each node has the same given probability to be active so as to ensure convergence to the mean of the values. One could design more sophisticated dropout models, as long as the contributions of nodes remain balanced.

We vary the expected level of available nodes at each step from 10 to 90%. Note that the number of iterations needed for convergence becomes stochastic: we thus set an arbitrary number of iterations experimentally, and average several runs. For simplicity, we do not perform gossip acceleration. We report several runs at each dropout level in Figure 6(a), and report the privacy loss and its standard deviation at each dropout level in Table 3. As expected, the convergence time increases with the proportion of inactive users, but the achievable privacy-utility trade-off is not significantly impacted. Therefore, our approach scales gracefully difficulty to the situations where there is dropout.

I Broader Impact Assessment

Our work promotes increased privacy in federated ML. The potential longer-term benefits of our work in this respect are a wider adoption of privacy-preserving and fully decentralized ML solutions by service providers thanks to the improved privacy-utility guarantees, as well as better confidence of users and the general public in the ability of decentralized ML systems to avoid catastrophic data leaks. In particular, our work shows that the advantages of decentralized methods over centralized ones have been overlooked, due to the lack of privacy analysis able to capture the benefits of decentralized algorithms.

Conversely, there are potential risks of accidental or deliberate misuse of our work in the sense that it could give a false sense of privacy to users if weak privacy parameters are used in deployment. This

Table 3: Total Privacy loss in function of dropout

Dropout	Privacy loss
No dropout	$(1.7 \pm 0.6) \times 10^{-2}$
10% dropout	$(1.8 \pm 0.6) \times 10^{-2}$
50% dropout	$(1.5 \pm 0.7) \times 10^{-2}$
90% dropout	$(1.4 \pm 0.6) \times 10^{-2}$

applies to all work on differential privacy. More applied research is needed towards developing a methodology to choose appropriate privacy parameters in a data-driven manner and to reliably assess the provided protection in practical use-cases.

Our work specifically proposes a varying privacy budget that depends on the relation between users, which might be misused for giving smaller privacy guarantees than the ones that would be designed otherwise. Modularity in privacy guarantees has however already been studied, for instance in [14] where the privacy budget is a function of a metric on the input space. Informally, defining what is an acceptable privacy budget based on the context in which some information is revealed is in line with the idea of contextual integrity [50]. According to Helen Nissenbaum's theory, the privacy expectations should take into account five elements: the sender, the receiver, the message, the medium of transmission and the purpose. Mathematically, adapting the privacy guarantee to the receiver and promoting peer-to-peer communications for building a global model thus naturally fits this view. In particular, contextual integrity emphasizes that the privacy budget should not only depend on the information being transmitted, but also on who receives it.