Source Anonymity for Private Random Walk Decentralized Learning

Maximilian Egger, Svenja Lage, Rawad Bitar and Antonia Wachter-Zeh
*Technical University of Munich, Germany, {maximilian.egger, rawad.bitar, antonia.wachter-zeh}@tum.de
†Communications and Navigation, German Aerospace Center, Germany, {svenja.lage}@dlr.de

***Abstract*—This paper considers random walk-based decentralized learning, where at each iteration of the learning process, one user updates the model and sends it to a randomly chosen neighbor until a convergence criterion is met. Preserving data privacy is a central concern and open problem in decentralized learning. We propose a privacy-preserving algorithm based on public-key cryptography and anonymization. In this algorithm, the user updates the model and encrypts the result using a distant user's public key. The encrypted result is then transmitted through the network with the goal of reaching that specific user. The key idea is to hide the source's identity so that, when the destination user decrypts the result, it does not know who the source was. The challenge is to design a network-dependent probability distribution (at the source) over the potential destinations such that, from the receiver's perspective, all users have a similar likelihood of being the source. We introduce the problem and construct a scheme that provides anonymity with theoretical guarantees. We focus on random regular graphs to establish rigorous guarantees.**

I. Introduction

Machine learning models have the potential to provide significant benefits in a wide range of areas, including intelligent healthcare [1], [2], Internet of Things (IoT) [3] or Internet of Vehicles [4], [5]. However, the success of the models relies heavily on access to large and comprehensive datasets. Distributed learning in its various forms, e.g., federated and decentralized learning, emerged as a new paradigm for accessing massive amounts of personalized and private data generated by participating clients.

In federated learning [6], users maintain their data locally and only share locally updated models with the federator, who orchestrates the training process. Decentralized learning eliminates the need for a central authority. Instead, users take an active role in distributing model updates among themselves. The users can be modeled as vertices in a graph. Users who can communicate are connected with an edge. Two main types of algorithms are studied in the literature: *(i) gossip algorithms*, e.g., [7]–[13], in which at every iteration, all users update the model locally and share their update with all their neighbors; and *(ii) random walk-based algorithms*, e.g., [14]–[18], in which at every iteration, one designated user updates the model locally and shares the update with one of its neighbors chosen at random. In both cases, the algorithm proceeds until certain convergence criteria are met. We focus on random walk-based algorithms due to their low communication cost incurred per iteration. The name random walk-based algorithm stems from the machine learning model being passed sequentially among neighboring nodes, thus drawing a random walk (RW) on the graph.

Despite users' data being kept locally, privacy is not immediately preserved. For instance, users can infer updates to the model by comparing its state to the point where the RW was last observed. The current model might exhibit a bias towards the data of the user who recently updated it. By accumulating such observations, a user may potentially glean information about the other users' data, cf. [19], [20].

The main approach to conceal individual data updates is through the application of differential privacy [21], as done in [17], by injecting carefully designed random noise into the model updates. However, this comes at the cost of a trade-off between privacy and model precision [22]. The noise needed to ensure privacy grows with the number of updates that will be observed [23]. Therefore, without proper care, a large amount of noise may be needed, significantly reducing the algorithm's utility. While homomorphic encryption, as introduced in [24], presents a promising alternative as it allows for computation on encrypted data, the computational overhead of such algorithms often renders them impractical for handling large datasets or complex functions (see, for example [25]).

This paper introduces a novel privacy-preserving model that does not require differential privacy mechanisms and whose practicality surpasses homomorphic solutions. The core idea is to use public-key encryption and source anonymity as follows. The user updating the model encrypts it using the public key of the destination user. The model is transmitted through the graph in a way that, when reaching the destination, the identity of the transmitting user will be concealed. This model achieves two goals simultaneously: *(i) high utility*, the destination can decrypt the model and use it plainly; and *(ii) privacy*, by concealing the identity of the transmitter, the eavesdropping users would not be able to map the information they inferred to other users' data. Note that our privacy-preserving mechanism is compatible with differential privacy. Noise can be inserted into the model before encryption. The noise level needed here may be lower since concealing the identity of the user hinders the composition of multiple observations.

The main challenge in this model is to carefully design a choice of the destination node among all nodes to ensure that the identity of the transmitting user is concealed. We term this property "source anonymity", inspired by similar studies in other contexts, such as wireless sensor networks, e.g., [26], [27]. To the best of our knowledge, this work is the first to tackle anonymity for random walk-based learning. We investigate random regular graphs in which a rigorous analysis of our method can be carried out. We study the privacy leakage,

This project is funded by DFG (German Research Foundation) projects under Grant Agreement Nos. BI 2492/1-1 and WA 3907/7-1.

i.e., how well the source anonymity can be guaranteed under probabilistic guarantees. All proofs are deferred to an extended version and can be found in [28].

II. Model and Anonymity Notion

Consider a collaborative setup consisting of N users, also referred to as nodes. Each user i with $i \in [N] \triangleq \{1, \dots, N\}$, possesses a local dataset $\mathcal{D}_i$ and is capable of communicating with a subset of the other users $\mathcal{N}_i \subseteq [N]$, called its neighbors. We represent the users as vertices in a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where the set of vertices $\mathcal{V} = [N]$ corresponds to the users, and the edges $\mathcal{E}$ represent the communication links between them. The degree of a node $i \in \mathcal{V}$ is defined as the number of neighbors it has, i.e., $\deg(i) = |\{j \in \mathcal{V} : (i,j) \in \mathcal{E}\}| = |\mathcal{N}_i|$.

We sketch a potential learning algorithm that serves as motivation for our work, noting that our methods can be applied to any RW-based decentralized system that processes sensitive individual data. Starting with a random model $\mathbf{w}_0 \in \mathbb{R}^d$ and a designated node $i_0 \in \mathcal{V}$, the following procedure takes place at every iteration $g \geq 0$. Node i_g updates the model using (stochastic) gradient descent based on a predefined loss function $F(\mathbf{w}, \mathcal{D}_{i_g})$ on its local data $\mathcal{D}_{i_g}$, i.e., the model update reads $\mathbf{w}_{g+1} = \mathbf{w}_g - \eta \nabla F(\mathbf{w}_g, \mathcal{D}_{i_g})$. The updated model is then sent to a random neighbor $i_{g+1} \in \mathcal{N}_{i_g}$. The process repeats until certain convergence criteria are met.

The described learning approach implicitly leaks information about the nodes' private data through the shared model updates. The privacy problem becomes most pronounced when eavesdropping nodes can obtain information about the model update of the targeted neighbors. For instance, consider a situation where, at iteration g, a node i_g sends the model update to a neighbor $i_{g+1} \in \mathcal{N}_{i_g}$, who updates the model and then chooses $i_{g+2} = i_g \in \mathcal{N}_{i_{g+1}}$ as the next destination. In this case, node i_g can directly obtain the local model update (gradient) $\nabla F(\mathbf{w}_{g+1},\ \mathcal{D}_{i_{g+1}})$ of node i_{g+1}, and from this infer information about node i_{g+1}'s local data $\mathcal{D}_{i_{g+1}}$.

We propose a modification of this learning algorithm that does not allow model updates conducted by direct members and further hides the identity of the updating node. In particular, node i_{g+1} should not know the identity of the predecessor i_g, thereby providing a new notion of privacy through source anonymity. We focus on a single iteration in the following and hence drop the iteration index g.

We assume that each node generates a cryptographic key pair and publishes its public key. Upon updating the model, the current node $i \in \mathcal{V}$ selects a destination node $j \in \mathcal{V} \setminus \{i\}$ according to a predefined distribution p_{D_i}. This distribution may be both node-dependent and time-varying. The current node then encrypts the model using the public key of the destination node j. The RW continues independently moving on the graph, but no update is performed until the destination node is reached. Once the RW reaches node j, the model can be decrypted using node j's private key. After updating the model, the learning process proceeds. The time it takes for the RW to transition from node i to node j is known as the first hitting time, denoted by $T_{\text{FH}}(i,j)$. We define the first hitting time distribution as $p_{i \to j}(t)$, which represents the probability that the RW reaches node j for the first time at time t, starting from node i. The return time $T_{\text{FR}}(i)$ is the time it takes for an RW to return to a node i after leaving the node.

By encrypting the model for a designated destination node and choosing appropriate distributions p_{D_i}, we ensure that the destination node cannot determine the identity of the source node, thereby preserving the source node's anonymity. While the first hitting time distribution is assumed to be known to all parties, the time of the previous model update at node i is unknown to the destination and can only be statistically estimated to infer the source identity. This inherently provides a stronger anonymity than if the previous update times were known. Our formal privacy notion is defined as follows.

Definition 1. *Let H denote the entropy function. We say that an RW-learning model, as described above, ensures α-privacy if, for intervals $I_{i,j} \triangleq [E_{i,j} - \delta, E_{i,j} + \delta]$ and fixed $\delta > 0$,*

$$\min_{j \in \mathcal{V}} H\left(\left[\frac{\int_{I_{i,j}} p_{D_i}(j) p_{i \to j}(t) dt}{\sum_{i' \in \mathcal{V} \setminus \{j\}} \int_{I_{i',j}} p_{D_{i'}}(j) p_{i' \to j}(t) dt} \right]_{i \in \mathcal{V} \setminus \{j\}} \right) \geq \alpha, \tag{1}$$

where $\alpha > 0$ and $E_{i,j} \geq \delta$ for all $i, j \in \mathcal{V}$.

Imagine that node $j \in \mathcal{V}$ receives the RW and attempts to estimate the identity of the source node. To facilitate this, node j estimates the time interval for an RW to move from a source node to itself, denoted as $E_{i,j}$. This estimation can be accomplished through various means, including computing average values or incorporating supplementary side information. For every node $i \in \mathcal{V} \setminus \{j\}$, and an estimated path length around $E_{i,j}$, node j calculates the probability that node i was the sender. If the entropy over all possible source nodes is high, node j will struggle to distinguish the true source node from other nodes. In the optimal scenario, the possible source nodes appear uniformly distributed to the destination node, making it impossible to identify the true source. Finally, note that this condition must hold for every possible destination node $j \in \mathcal{V}$.

Concealing the source of an update incurs a cost in terms of increased runtime. Unlike the classical approach, where the model is updated at every time step, our method updates the model only when the RW reaches the designated destination node. On average, the first hitting time is given by

$$\mathbb{E}[T_{\text{FH}}] = \sum_{i,j \in \mathcal{V}, i \neq j} \mathbb{E}[T_{\text{FH}}(i,j)] P(i) p_{D_i}(j),$$

where $P(i)$ represents the likelihood of the RW being in node $i \in \mathcal{V}$. In practice, choosing suitable distributions p_{D_i} requires balancing two competing goals: achieving a high level of anonymity while also minimizing the runtime overhead.

III. Optimal Source Anonymity

To elucidate the implications of the privacy notation outlined in the previous section, we examine the specific instance of Random Regular Graphs (RRG). The uniform structure of RRGs enables a thorough and precise analysis, leveraging analytical expressions for the first hitting time and return time of the RWs. This characteristic makes RRGs a more accessible and illustrative choice for demonstrating our model, compared to other, more complex random graphs. We show in this section how to select the distributions p_{D_i} such that, in the absence of additional side information, our model yields an optimal

privacy guarantee; i.e., that the potential source nodes appear uniformly distributed from the perspective of the destination node. We therefore give a concise introduction to RRGs and their associated first hitting time distribution. We then address the topic of source node anonymity within this context.

A. On the First Hitting Time of RRGs

RRGs are characterized by a degree distribution, that, for all nodes, is a degenerate probability density function, such that $P(k) = \mathbb{1}_{\{c\}}(k)$, where c is the degree of the RRG, and $\mathbb{1}_{\{c\}}(k) = 1$ iff $c = k$ and 0 else. Hence, every node within the graph exhibits a uniform degree. We focus on RRGs in which the degree parameter $c \geq 3$, for which the graph consists of a single, connected component if N is large enough [29].

For RRGs, the authors of [30] presented approximate expressions for the first hitting time distribution. Notably, they considered two distinct cases for the RW between two nodes. In the shortest path (SP) scenario, the RW follows the direct path between node i and node j. This includes scenarios in which the RW may backtrack some of its steps or even recede. Formally, a path belongs to this case if the subnetwork consisting of the nodes and edges along the trajectory forms a tree network, and the distance between node i and node j in this subnetwork is the same as in the entire network. All other paths belong to the opposite case, denoted as $\neg$SP.

We are interested in the first hitting time between nodes i and j, which, in an RRG, only depends on their distance ℓ (i.e., the length of the shortest path). According to [30], we have $P(T_{\text{FH}} = t|\ell) = P(T_{\text{FH}} = t|\ell, \text{SP})P(\text{SP}|\ell) + P(T_{\text{FH}} = t|\ell, \neg\text{SP})P(\neg\text{SP}|\ell)$, where the hitting time distributions conditioned on the two distinct cases SP and $\neg$SP are given by $P(T_{\text{FH}} = t|\ell, \text{SP}) = \frac{\ell}{t}\binom{t}{\frac{t+\ell}{2}}(1-c^{-1})^{\frac{t+\ell}{2}}c^{\frac{\ell-t}{2}}$ for $(t-\ell)$ even and

$$P(T_{\text{FH}} = t|\ell, \neg\text{SP}) = \left(e^{\frac{c'}{N}} - 1\right)e^{-c'\frac{t-\ell}{N}}, \tag{2}$$

for $t > \ell$, where c' is defined as $c' \triangleq (c-2)/(c-1)$. The probabilities for each case are given by

$$P(\text{SP}|\ell) = \left(\frac{1}{c-1}\right)^{\ell} + \frac{1}{N} \text{ and } P(\neg\text{SP}|\ell) = 1 - P(\text{SP}|\ell).$$

Additionally, we can express the expected value of the first hitting time distribution, conditioned on the distance between two nodes, as $\mathbb{E}[T_{\text{FH}}|\ell] = \mathbb{E}[T_{\text{FH}}|\ell, \text{SP}]P(\text{SP}|\ell) + \mathbb{E}[T_{\text{FH}}|\ell, \neg\text{SP}]P(\neg\text{SP}|\ell)$ with $\mathbb{E}[T_{\text{FH}}|\ell, \text{SP}] = \frac{c}{c-2}\ell$ and

$$\mathbb{E}[T_{\text{FH}}|\ell, \neg\text{SP}] = \ell + \frac{1}{1 - e^{-\frac{c'}{N}}}. \tag{3}$$

In the non-shortest path scenario, we applied the more accurate result [30, Eq. 54] for $P(T_{\text{FH}} > t|\ell, \neg\text{SP})$ to calculate the expectation, yielding a result that differs marginally from the original expression in [30, Eq. 70].

B. Source Anonymity in Light of RRGs

We demonstrate how to select the distributions p_{D_i} to achieve optimal α-privacy, where optimal refers to a value of α equal to the entropy of a uniform distribution. For simplicity, we assume that all distributions p_{D_i} are time-invariant. Furthermore, due to the regular structure of the graph, we posit that for an arbitrary node i, the probability distribution $p_{D_i}(j)$ depends solely on the distance between nodes i and j. Let $A(\ell)$ be the number of nodes situated at distance $\ell \geq 1$, which we assume is constant for every node i. We can then equivalently express the distribution p_{D_i} as $p_{D_i}(j) = \frac{p(\ell)}{A(\ell)}$, where p denotes a probability distribution over node distances. To achieve high-probability control in arbitrary cases, we constrain the support of p by excluding direct neighbors. Denote the resulting support of p as $[\ell_1, \ell_2]$, where ℓ_2 is bounded from above by the diameter of the graph.

Initially, we consider a scenario in which the node lacks any supplementary side information beyond the graph's structure and the distributions $(p_{D_i})_{i\in\mathcal{V}}$. Under these conditions, a natural choice for $E_{i,j}$ in (1) is the expected first hitting time $E_{i,j} = \mathbb{E}[T_{\text{FH}}(i,j)]$ from node i to node j. Notably, in an RRG, the expected first hitting time $E_{i,j}$ depends solely on the nodes' distance ℓ. The integral in the numerator of (1) simplifies to

$$\int_{\mathbb{E}[T_{\text{FH}}(i,j)]-\delta}^{\mathbb{E}[T_{\text{FH}}(i,j)]+\delta} p_{D_i}(j)p_{i\to j}(t)dt = \frac{p(\ell)}{A(\ell)} \sum_{t=\mathbb{E}[T_{\text{FH}}|\ell]-\delta}^{\mathbb{E}[T_{\text{FH}}|\ell]+\delta} P\left(T_{\text{FH}} = t|\ell\right). \tag{4}$$

When computing the sum in (4), we focus on the hitting time probability within an interval centered around the expected time. However, as a consequence of (3), it follows that $\mathbb{E}[T_{\text{FH}}|\ell] \geq \mathbb{E}[T_{\text{FH}}|\ell, \neg\text{SP}]P(\neg\text{SP}|\ell) \sim \left(l + \frac{N}{c'}\right)P(\neg\text{SP}|\ell)$ and hence, $\mathbb{E}[T_{\text{FH}}|\ell] \gg l$. For values of t in this regime, we have $P(T_{\text{FH}} = t|\ell, \text{SP}) \in O(t^{-\frac{3}{2}}e^{-\gamma t})$ with $\gamma = \log(c) - \frac{1}{2}\log(c-1) - \log(2) > 0$, which, for all c becomes asymptotically negligible compared to the contribution from the non-shortest path. This observation justifies the following simplification.

Assumption 1. *The first hitting time distribution within the interval $I_\ell \triangleq [\mathbb{E}[T_{FH}|\ell] - \delta, \mathbb{E}[T_{FH}|\ell] + \delta]$ is dominated by the non-shortest path scenario for all $\ell \in [\ell_1, \ell_2]$, i.e., we assume that $\Pr(T_{FH} = t|\ell) \triangleq \Pr(T_{FH} = t|\ell, \neg SP)\Pr(\neg SP|\ell)$ for every $t \in I_\ell, \ell \in [\ell_1, \ell_2]$.*

Under this assumption, the integral presented in (4) admits a closed-form solution. To simplify the notation, let

$$K_\delta(\ell) \triangleq \frac{P(\neg\text{SP}|\ell)}{A(\ell)}\left(e^{\frac{c'}{N}} - 1\right)e^{c'\frac{\delta+\ell}{N}}\frac{e^{-c'\frac{2\delta+1}{N}} - 1}{e^{-\frac{c'}{N}} - 1}.$$

Lemma 1. *Let $i, j \in \mathcal{V}$ be two nodes within distance $\ell \in [\ell_1, \ell_2]$ and let $\delta > 0$ be fixed. Under Assumption 1, the integral in (4) is given by*

$$\int_{\mathbb{E}[T_{FH}|\ell]-\delta}^{\mathbb{E}[T_{FH}|\ell]+\delta} p_{D_i}(j)p_{i\to j}(t)dt = p(\ell)E_\delta(\ell)$$

where $E_\delta(\ell) = K_\delta(\ell)e^{-c'\frac{\mathbb{E}[T_{FH}|\ell]}{N}}$.

We find that the privacy notion now only depends on the distance ℓ between two nodes. Consequently, in the context of RRGs, our privacy notion requires that the distribution

$$W_\delta(\ell) \triangleq \frac{p(\ell)E_\delta(\ell)}{\sum_{\ell'} p(\ell')E_\delta(\ell')}$$

on $\ell \in [\ell_1, \ell_2]$ remains sufficiently close to a uniform distribution over all possible values of ℓ. Notably, this security notion is particularly satisfied when $p(\ell)E_\delta(\ell)$ remains constant across

all values of $\ell \in [\ell_1, \ell_2]$. The following choice of $p(\ell)$ achieves optimal privacy in the sense of (1) for RRGs.

Lemma 2. *Let $\delta > 0$. For every $\ell \in [\ell_1, \ell_2]$, choose $p(\ell)$ as*

$$p^\star_\delta(\ell) \triangleq \frac{1}{E_\delta(\ell) \sum_{\ell'=\ell_1}^{\ell_2} \frac{1}{E_\delta(\ell')}}$$

and $p^\star_\delta(\ell) = 0$ elsewhere. Then W_δ is uniform on the support $[\ell_1, \ell_2]$ and consequently $H(W_\delta) = \log(\ell_2 - \ell_1 + 1)$.

This choice of $p(\ell)$ guarantees the destination node cannot identify the source node with better accuracy than by simply choosing uniformly at random from all possible nodes. This demonstrates how to achieve optimal α-privacy with the maximal value of α, if no further side information is available.

Remark 1. *For RRGs, we analyze the distributions over distances. Equivalently, the entropy can be formulated expanding each distance ℓ with all source nodes in distance ℓ.*

IV. Source Anonymity under Side Information

In practice, the situation is more complex than initially described. A node can gather side information about the first hitting time by recalling the most recent visit of the RW. This additional information can, in turn, affect the node's ability to accurately infer the identity of the source node. This consideration becomes particularly crucial when the return time between two visits of the RW is short, as it effectively eliminates certain nodes as potential sources. If the model was re-encrypted between two returns to node j, and node j is the designated destination, the return time represents an upper bound on the first hitting time. If the distributions p_{D_i} were chosen as before, the destination node could indeed make a more informed guess about the source node, surpassing the uniform case. To alleviate this problem, users can choose p_{D_i} such that even with this additional side information, the deviation from the uniform distribution remains bounded with high probability. To optimize the model for such cases, the user selects a design parameter κ as an upper bound for the first hitting time. Following this, $E_{i,j}$ in Definition 1 can be chosen as $E_{i,j} = \mathbb{E}[T_{\text{FH}}|\ell, T_{\text{FH}} \leq \kappa]$. We will first demonstrate how to choose p_{D_i} in this setting and then analyze the probabilistic guarantees when a different value κ' is observed during a random walk, where κ' refers to the return of the random walk used as side information by the destination node as described above. In line with Assumption 1, we assume the following.

Assumption 2. *$\mathbb{E}[T_{FH}|\ell, T_{FH} \leq \kappa']$ is dominated by the non-shortest path scenario, that is $\mathbb{E}[T_{FH}|\ell, T_{FH} \leq \kappa'] \approx \mathbb{E}[T_{FH}|\ell, \neg SP, T_{FH} \leq \kappa'] \Pr(\neg SP|\ell, T_{FH} \leq \kappa')$.*

The expected value obtained under additional side information can be calculated as follows.

Proposition 1. *For every $\kappa' > \ell$, we have*

$$\mathbb{E}[T_{FH}|\ell, \neg SP, T_{FH} \leq \kappa'] = \frac{\ell}{1 - e^{-c'/N(\kappa' - \ell)}} + \frac{1}{1 - e^{-c'/N}}.$$

The following generalization of Lemma 1 gives a closed-form expression for the corresponding modification of (4) with additional side information.

Lemma 3. *Let $i, j \in \mathcal{V}$ be two nodes within distance $\ell \in [\ell_1, \ell_2]$. Let $\delta > 0$ be fixed and assume that it is known that the first hitting time is restricted by $\kappa' \in (\ell, \infty)$. Under Assumption 1, we have*

$$\int_{\mathbb{E}[T_{FH}|\ell, T_{FH} \leq \kappa'] - \delta}^{\mathbb{E}[T_{FH}|\ell, T_{FH} \leq \kappa'] + \delta} p_{D_i}(j) p_{i \to j}(t) dt = p(\ell) E_{\delta, \kappa'}(\ell),$$

where $E_{\delta,\kappa'}(\ell) = K_\delta(\ell) e^{-c' \frac{\mathbb{E}[T_{FH}|\ell, T_{FH} \leq \kappa']}{N}}$.

We refine the solution given in Lemma 2 for the optimal distribution of $p^\star_\delta(\ell)$ to the case where certain side information $\kappa' = \kappa$ is assumed. Afterwards, we provide guarantees on privacy when the actual side information κ' differs from the design parameter κ. Given κ, we choose the destination node distributions as follows.

Lemma 4. *Let $\kappa, \delta > 0$. For every $\ell \in [\ell_1, \ell_2]$, choose $p(\ell)$ as*

$$p^\star_{\kappa,\delta}(\ell) = \frac{s}{E_{\delta,\kappa}(\ell)} = \frac{1}{E_{\delta,\kappa}(\ell) \sum_{\ell'=\ell_1}^{\ell_2} \frac{1}{E_{\delta,\kappa}(\ell')}},$$

and $p^\star_{\kappa,\delta}(\ell) = 0$ elsewhere. With $W_{\delta,\kappa'}$ defined as

$$W_{\delta,\kappa'}(\ell) \triangleq \frac{p^\star_{\kappa,\delta}(\ell) E_{\delta,\kappa'}(\ell)}{\sum_{\ell'} p^\star_{\kappa,\delta}(\ell') E_{\delta,\kappa'}(\ell')}, \tag{5}$$

we have that $W_{\delta,\kappa}$, i.e., $W_{\delta,\kappa'}(\ell)$ for $\kappa' = \kappa$, is uniform on the support $[\ell_1, \ell_2]$ and consequently $\mathrm{H}(W_{\delta,\kappa}) = \log(\ell_2 - \ell_1 + 1)$.

Let $p^\star_{\kappa,\delta}(\ell)$ be the optimal distribution for the case $\kappa' = \kappa$. To analyze the privacy for $\kappa' \neq \kappa$, we study the distribution $W_{\delta,\kappa'}(\ell)$, which inherits the support $[\ell_1, \ell_2]$ from $p^\star_{\kappa,\delta}$. Our objective is to show that, with high probability, the deviation between $W_{\delta,\kappa'}$ and the uniform distribution of $W_{\delta,\kappa}$ is bounded, facilitating a bound for the α-privacy guarantee. Using Proposition 1, we can bound the total variation and the entropy as follows.

Theorem 1. *Let $\kappa, \kappa', \delta > 0$. For $W_{\delta,\kappa'}$ as in (5), we have*

$$d_{TV}(W_{\delta,\kappa'} \| W_{\delta,\kappa}) \leq \frac{1}{\ell_2 - \ell_1} \sum_{\ell=\ell_1}^{\ell_2} (e^{\frac{c'}{N} \varphi \varepsilon_\ell(\kappa, \kappa')} - 1),$$

where $\varepsilon_\ell(\kappa, \kappa') = \ell \left| \frac{1}{1 - e^{-c'/N(\kappa' - \ell)}} - \frac{1}{1 - e^{-c'/N(\kappa - \ell)}} \right|$, and $\varphi \triangleq \Pr(\neg SP|\ell_2, T_{FH} \leq \kappa')$.

Theorem 2. *Let $\rho \triangleq d_{TV}(W_{\delta,\kappa'} \| W_{\delta,\kappa})$. For $p^\star_{\kappa,\delta}(\ell)$ as in Eq. (5), the entropy of $W_{\delta,\kappa'}$ is lower bounded by*

$$\mathrm{H}(W_{\delta,\kappa'}) \geq (1 - \rho) \log(\ell_2 - \ell_1 + 1) + \rho \log(\rho) - \rho.$$

The distribution of κ' is captured by the first return time of the RW. The deliberate choice of the destination node does not affect the stochasticity of the RW, and hence the return time is independent of our method. Let $\mathcal{K} \subset \mathbb{N}$ be such that $\Pr(\kappa' \in \mathcal{K}) \geq 1 - \delta'$ Our goal is that for a given δ, a desired probability $1 - \delta'$ and a chosen $\mathcal{K} \subset \mathbb{N}$, we optimize

$$\min_{\kappa} \max_{\kappa' \in \mathcal{K}} |\mathrm{H}(W_{\delta,\kappa}) - \mathrm{H}(W_{\delta,\kappa'})|.$$

We make the following assumption on the first return time of an RW on an RRG, which can be separated into retroceding

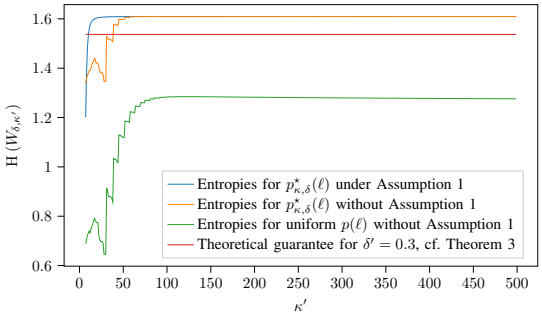

Fig. 1. RW-model for $N = 300$ nodes and degree $c = 3$ optimized for $\kappa = 634$ and $\delta = 5$. We compare the entropies over varying side information κ', with and without Assumption 1, compared with no countermeasure and the theoretical guarantee from Theorem 3 using $\delta' = 0.3$. Uniform distribution supported on $[\ell_1, \ell_2] = [2, 6]$ has entropy $\log(\ell_2 - \ell_1 + 1) \approx 1.61$.

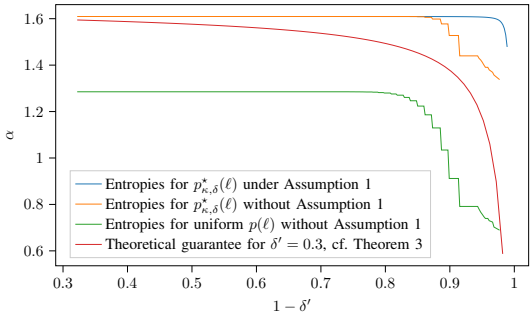

Fig. 2. α-privacy over probabilistic guarantee for $N = 300$, $c = 3$ and $\delta = 5$ with and without Assumption 1, compared with no countermeasure and the theoretical guarantee from Theorem 3.

(RETRO) and non-retroceding trajectories (¬RETRO). The distributions are known from [31].

Assumption 3. *The first return time on RRGs is dominated by non-retroceding scenarios, i.e.,* $\Pr(T_{FR} = t) = \Pr(T_{FR} = t|\neg RETRO)$*, which is described by* $Pr(T_{FR} > t|\neg RETRO) = e^{-c'\frac{t-2}{N-2}}$ *for* $t \geq 3$ *and* $\Pr(T_{FR} > t|\neg RETRO) = 1$ *otherwise.*

We further have $\mathbb{E}[T_{\text{FR}}|\neg\text{RETRO}] = 2 + (1 - e^{-c'\frac{1}{N-2}})^{-1}$. By the choice of $p(\ell)$ supported on $[\ell_1, \ell_2]$, each node selects the destination node in a distance of at least ℓ_1. Hence, at any destination node, the minimal observed return time is $T_{\text{FR}} \geq 2\ell_1$ and hence the probability that a destination node observes a return time of $T_{\text{FR}} \geq t \geq 2\ell_1$ is given by

$$\Pr(T_{\text{FR}} \geq t|\neg\text{RETRO}, T_{\text{FR}} \geq 2\ell_1) = e^{-c'\frac{t-2\ell_1}{N-2}}. \tag{6}$$

We justify Assumption 3 by the fact that we are interested in first return times of at least $2\ell_1$. In this case, the probability $\Pr(\neg\text{RETRO}) = 1/(c-1)$ diminishes further, making $\Pr(\text{RETRO}|T_{\text{FR}} \geq 2\ell_1)$ the dominating component. With this at hand, we have the following main result of our paper that quantifies α-privacy under a probabilistic guarantee.

Theorem 3. *Let* $\delta' > 0$ *and* $d \triangleq \ell_2 - \ell_1$*. Then, there exists a value for* κ *such that with probability at least* $1 - \delta'$ *for a certain* $\mathcal{K}$ *s.t.* $\Pr(\kappa' \in \mathcal{K}) \geq 1 - \delta'$*, the entropy observed by every destination node* j *is bounded by*

$$\max_{\kappa} \min_{\kappa' \in \mathcal{K}} \text{H}\left(W_{\delta,\kappa'}\right) \geq \underbrace{(1-\rho)\log(d+1) + \rho\log(\rho) - \rho}_{\alpha},$$

where, for $\varphi \triangleq \Pr(\neg SP|\ell_2, T_{FH} \leq \kappa')$*,*

$$\rho = O\left(\frac{1}{d}\frac{e^{\frac{c'}{2N}(\ell_2+1)\varphi\varepsilon} - e^{\frac{c'}{2N}\ell_1\varphi\varepsilon}}{e^{\frac{c'}{2N}\varphi\varepsilon} - 1} - 1\right),$$

and the asymptotic behavior of ε *is given by*
$\varepsilon \triangleq \frac{\ell_2}{2} O\left(\max\left\{(1-\delta')e^{-(2\ell_1-\ell_2)\frac{c'}{N}} - 1, \frac{-2-\log(1-\delta')}{2\log(1-\delta')}\right\}\right)$.

The result above is based on selecting an optimal value for κ that minimizes the upper bound on the entropy. When optimal uniformity of $W_{\delta,\kappa'}$ for $\kappa' = \kappa$ on $[\ell_1, \ell_2]$ should be achieved for the average return time, and hence $\kappa = \mathbb{E}[T_{\text{FR}}|\neg\text{RETRO}]$, we have the following result. The result also holds for $\kappa \to \infty$, which corresponds to the case studied in Section III.

Corollary 1. *When choosing* $p^{\star}_{\kappa,\delta}(\ell)$ *to match the average case, i.e., for* $\kappa = \mathbb{E}[T_{FR}|\neg RETRO]$*, and for* $\kappa \to \infty$*, Theorem 3 holds with* ε *replaced by* 2ε*.*

V. Numerical Evaluations

We evaluate our methods and theoretical guarantees on an RRG with $N = 300$ nodes and degree $c = 3$. Therefore, we consider the parameters $\delta = 5$, $\delta' = 0.3$, along with the choice of $\kappa = 634$ and $p^{\star}_{\kappa,\delta}(\ell)$ resulting from Theorem 3 as well as $\ell_1 = 2$ and $\ell_2 = 6$. We approximate $\Pr(\neg\text{SP}|\ell_2, T_{\text{FH}} \leq \kappa') \approx \Pr(\neg\text{SP}|\ell_2)$ for all κ'. In Fig. 1, we compare the entropy in Definition 1 for various side information κ' under Assumption 1 to the case where the optimal solution is found based on Lemma 4 using the exact properties of RRGs, i.e., without the relaxation in Assumption 1. As a baseline, we plot the source node anonymity without any countermeasure, i.e., for choosing $p_{\delta}(\ell)$ uniformly on $[\ell_1, \ell_2]$, and we also show our worst-case bound provided by Theorem 3. For the same choice of κ, we plot in Fig. 2 the minimum entropy resulting from Theorem 3 as a function of the probability $1 - \delta'$, illustrating how anonymity degrades as this probability varies. With approximately 90% probability, our method still achieves near-optimal anonymity. Lastly, we analyze source node anonymity for $\kappa' = \kappa$ over the mean iteration time $T_{\ell_1,\ell_2} = \sum_{\ell=\ell_1}^{\ell_2} p^{\star}_{\kappa,\delta}(\ell)\mathbb{E}[T_{\text{FH}}|\ell]$ determined by the choice of $p^{\star}_{\kappa,\delta}(\ell)$ and the average first hitting times $\mathbb{E}[T_{\text{FH}}|\ell]$. We observe an almost linear increase of α in T_{ℓ_1,ℓ_2} (cf. [18]).

VI. Conclusion

We considered the problem of privacy in decentralized random walk-based learning algorithms. Instead of resorting to only applying differential privacy guarantees, we formulated a new privacy notion based on revealing the model update, but concealing the identity of the owner of the revealed update. To that end, public key cryptography is used by the sender to encrypt the update with the public key of a designated destination, ensuring that no intermediate node can decrypt the model update. The choice of the destination is the key component. We designed a probability distribution over the choice of the destination that ensures that with high probability, the destination will not be able to guess the identity of the source.

REFERENCES

[1] G. J. Simon and C. Aliferis., Eds., *Artificial Intelligence and Machine Learning in Health Care and Medical Sciences. Best Practices and Pitfalls*. Springer Cham, 2024.

[2] Z. Lian, Q. Yang, W. Wang, Q. Zeng, M. Alazab, H. Zhao, and C. Su, "DEEP-FEL: Decentralized, Efficient and Privacy-Enhanced Federated Edge Learning for Healthcare Cyber Physical System," *IEEE Transactions on Network Science and Engineering*, vol. 9, no. 5, pp. 3558–3569, 2022.

[3] S. Messaoud, A. Bradai, S. H. R. Bukhari, P. T. A. Quang, O. B. Ahmed, and M. Atri, "A Survey on Machine Learning in Internet of Things: Algorithms, Strategies, and Applications," *Internet of Things*, vol. 12, p. 100314, 2020.

[4] E. S. Ali, M. B. Hassan, and R. A. Saeed, "Machine Learning Technologies in Internet of Vehicles," in *Intelligent Technologies for Internet of Vehicles*, N. Magaia, G. Mastorakis, C. Mavromoustakis, E. Pallis, and E. K. Markakis, Eds., 2021, pp. 225–252.

[5] L. Barbieri, S. Savazzi, M. Brambilla, and M. Nicoli, "Decentralized Federated Learning for Extended Sensing in 6G Connected Vehicles," *Vehicular Communications*, vol. 33, p. 100396, 2022.

[6] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-Efficient Learning of Deep Networks from Decentralized Data," in *International Conference on Artificial Intelligence and Statistics*, 2016.

[7] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Gossip Algorithms: Design, Analysis and Applications," in *Proceedings IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies.*, vol. 3, 2005, pp. 1653–1664.

[8] D. Shah, "Gossip Algorithms," *Foundations and Trends in Networking*, vol. 3, no. 1, pp. 1–125, 2009.

[9] J. Lu, C. Y. Tang, P. R. Regier, and T. D. Bow, "Gossip Algorithms for Convex Consensus Optimization Over Networks," *IEEE Transactions on Automatic Control*, vol. 56, pp. 2917–2923, 2010.

[10] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized Gossip Algorithms," *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2508–2530, 2006.

[11] A. Nedic and A. Ozdaglar, "Distributed Subgradient Methods for Multi-Agent Optimization," *IEEE Transactions on Automatic Control*, vol. 54, pp. 48–61, 2009.

[12] A. Koloskova, S. Stich, and M. Jaggi, "Decentralized Stochastic Optimization and Gossip Algorithms with Compressed Communication," in *International Conference on Machine Learning*, 2019, pp. 3478–3487.

[13] J. C. Duchi, A. Agarwal, and M. J. Wainwright, "Dual Averaging for Distributed Optimization: Convergence Analysis and Network Scaling," *IEEE Transactions on Automatic control*, vol. 57, no. 3, pp. 592–606, 2011.

[14] B. Johansson, M. Rabi, and M. Johansson, "A Randomized Incremental Subgradient Method for Distributed Optimization in Networked Systems," *SIAM J. Optim.*, vol. 20, pp. 1157–1170, 2009.

[15] G. Ayache, V. Dassari, and S. El Rouayheb, "Walk for Learning: A Random Walk Approach for Federated Learning from Heterogeneous Data," *IEEE Journal on Selected Areas in Communications*, vol. 41, no. 4, pp. 929–940, 2023.

[16] J. C. Duchi, A. Agarwal, M. Johansson, and M. I. Jordan, "Ergodic Mirror Descent," *SIAM Journal on Optimization*, vol. 22, no. 4, pp. 1549–1578, 2012.

[17] G. Ayache and S. E. Rouayheb, "Private Weighted Random Walk Stochastic Gradient Descent," *IEEE Journal on Selected Areas in Information Theory*, vol. 2, no. 1, pp. 452–463, 2021.

[18] M. Egger, G. Ayache, R. Bitar, A. Wachter-Zeh, and S. El Rouayheb, "Self-Duplicating Random Walks for Resilient Decentralized Learning on Graphs," in *IEEE Global Communications Conference*, 2024, pp. 2960–2965.

[19] H. Yin, A. Mallya, A. Vahdat, J. M. Alvarez, J. Kautz, and P. Molchanov, "See through Gradients: Image Batch Recovery via GradInversion," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2021, pp. 16 332–16 341.

[20] H. Hu, Z. Salcic, L. Sun, G. Dobbie, P. S. Yu, and X. Zhang, "Membership Inference Attacks on Machine Learning: A Survey," *ACM Computing Surveys*, vol. 54, no. 11s, Sep. 2022.

[21] C. Dwork and A. Roth, "The Algorithmic Foundations of Differential Privacy," *Foundations and Trends in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.

[22] M. Li, Y. Tian, J. Zhang, D. Fan, and D. Zhao, "The Trade-Off Between Privacy and Utility in Local Differential Privacy," in *International Conference on Networking and Network Applications (NaNA)*, 2021, pp. 373–378.

[23] P. Kairouz, S. Oh, and P. Viswanath, "The Composition Theorem for Differential Privacy," ser. International Conference on Machine Learning 2015, 2015, p. 1376–1385.

[24] R. L. Rivest, L. Adleman, and M. L. Dertouzos, "On Data Banks and Privacy Homomorphisms," *Foundations of Secure Computation, Academia Press*, pp. 169–179, 1978.

[25] J.-W. Lee, H. Kang, Y. Lee, W. Choi, J. Eom, M. Deryabin, E. Lee, J. Lee, D. Yoo, Y.-S. Kim, and J.-S. No, "Privacy-Preserving Machine Learning With Fully Homomorphic Encryption for Deep Neural Network," *IEEE Access*, vol. 10, pp. 30 039–30 054, 2022.

[26] B. Alomair, A. Clark, J. Cuellar, and R. Poovendran, "Statistical Framework for Source Anonymity in Sensor Networks," in *IEEE Global Telecommunications Conference GLOBECOM 2010*, 2010, pp. 1–6.

[27] Y. Yang, M. Shao, S. Zhu, and G. Cao, "Towards Statistically Strong Source Anonymity for Sensor Networks," *ACM Transactions on Sensor Networks (TOSN)*, vol. 9, no. 3, pp. 1–23, 2013.

[28] M. Egger, S. Lage, R. Bitar, and A. Wachter-Zeh, "Source anonymity for private random walk decentralized learning," *arXiv preprint arXiv:2505.07011*, 2025.

[29] B. Bollobás, *Random Graphs*, 2nd ed., ser. Cambridge Studies in Advanced Mathematics, 2001.

[30] I. Tishby, O. Biham, and E. Katzav, "Analytical Results for the Distribution of First-Passage Times of Random Walks on Random Regular Graphs," *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2022, no. 11, p. 113403, Nov 2022.

[31] ——, "Analytical Results for the Distribution of First Return Times of Random Walks on Random Regular Graphs," *Journal of Physics A: Mathematical and Theoretical*, vol. 54, no. 32, p. 325001, Jul. 2021.