Privacy Amplification by Decentralization

Edwige Cyffers*[1], Aurélien Bellet[2]

[1]ENS Lyon [2]INRIA

Context: Central vs. Local Differential Privacy

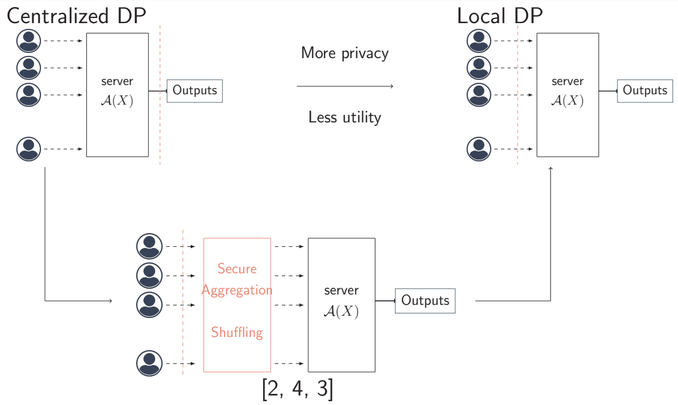

[2, 4, 3]

Summary of Contributions

- **New relaxation of LDP** that naturally arises in the decentralized setting
- **Formal privacy amplification results** for decentralized protocols based on random walks, see for instance [6]
- Gains confirmed by **numerical experiments**

Example of random walk

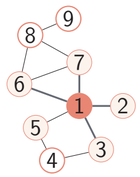

Time	vertex	value
1	1	39
2	**6**	**29**
3	8	27
4	7	23
5	1	22
6	**6**	**21**
7	7	22

Table: Local DP

Time	vertex	value
1	1	39
2	6	29
3	8	27
4	7	23
5	1	22
6	6	21
7	7	22

Table: Observation of 1

Open Questions

- Generalization to **arbitrary graphs**
- Dealing with **time-evolving topologies**
- Privacy amplification for other algorithms, e.g. **randomized gossip**

References

[1] B. Balle, G. Barthe, and M. Gaboardi. Privacy Amplification by Subsampling: Tight Analyses via Couplings and Divergences. In *NeurIPS*, 2018.

[2] B. Balle, J. Bell, A. Gascón, and K. Nissim. Differentially Private Summation with Multi-Message Shuffling. Technical report, arxiv:1906.09116, 2019.

[3] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth. Practical Secure Aggregation for Privacy-Preserving Machine Learning. In *CCS*, 2017.

[4] U. Erlingsson, V. Feldman, I. Mironov, A. Raghunathan, and K. Talwar. Amplification by Shuffling: From Local to Central Differential Privacy via Anonymity. In *SODA*, 2019.

[5] V. Feldman, I. Mironov, K. Talwar, and A. Thakurta. Privacy Amplification by Iteration. In *FOCS*, 2018.

[6] S. Ram, A. Nedić, and V. Veeravalli. Incremental stochastic subgradient algorithms for convex optimization. *SIAM Journal on Optimization*, 20(2):691–717, 2009.

Network Differential Privacy

Setting

- Graph of n nodes with edges for communication
- Each node u holds private database D_u

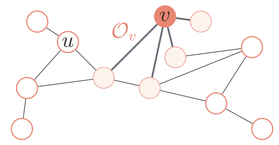

Definition 1 (Network Differential Privacy).

An algorithm $\mathcal{A}$ is (ε, δ)-network DP if for all pairs of distinct users $u, \mathbf{v} \in V$ and all pairs of datasets $D \sim_u D'$, we have

$$\mathbb{P}(\mathcal{O}_{\mathbf{v}}(\mathcal{A}(D))) \leq e^{\varepsilon}\mathbb{P}(\mathcal{O}_{\mathbf{v}}(\mathcal{A}(D'))) + \delta.$$

- **Intermediate point** between central and local DP
- View $\mathcal{O}_v$ of user v includes received messages and **depends of the algorithm**
- Variant with **collusion** (= merge of nodes)

Walk on a Ring

Algorithm 1: Real summation.

$\tau \leftarrow 0$; $a \leftarrow 0$;
for $k = 1$ *to* K **do**
  for $u = 1$ *to* n **do**
    if $a = 0$ **then**
      $\tau \leftarrow \tau + x_u^k + \text{Perturb}(\sigma)$;
      $a = n - 2$;
    else
      $\tau \leftarrow \tau + x_u^k$;
      $a \leftarrow a - 1$;
return τ

Aggregation

- Variant with noise equally split between nodes
- **Same privacy-utility** trade-off as a **trusted aggregator**
- Amplify privacy by $O(1/\sqrt{n})$ compared to local DP

Algorithm 2: Discrete histogram.

Init. $\tau \in \mathbb{N}^L$ with γn uniformly random elements;
for $k = 1$ *to* K **do**
  for $u = 1$ *to* n **do**
    $y_u^k \leftarrow RR_\gamma(x_u^k)$;
    $\tau[y_u^k] \leftarrow \tau[y_u^k] + 1$;
for $i = 0$ *to* $L - 1$ **do**
  $\tau[i] \leftarrow \frac{\tau[i] - \gamma/L}{1-\gamma}$;
return τ

Histogram computation

- Based on randomized response
- Messages are multisets of answers so far $\rightarrow$ amplification by **shuffling** [4]
- Amplify privacy by $\mathcal{O}_\delta(1/\sqrt{n})$ compared to local DP

Theorem 2.

Let $\varepsilon, \delta > 0$. Algorithm 1 outputs an unbiased estimate of the sum $\bar{x}$ with standard deviation $\sqrt{\lfloor Kn/(n-1) \rfloor}\sigma$ and satisfies $(\sqrt{2K\ln(1/\delta')}\varepsilon + K\varepsilon(e^\varepsilon - 1), K\delta + \delta')$-network DP for any $\delta' > 0$.

Theorem 3.

Let $\epsilon, \delta \in (0, \frac{1}{100})$, and $n > 1000$. Let $\gamma = L/(\exp(12\varepsilon_0\sqrt{\frac{\log(1/\delta)}{n}}) + L - 1)$. Algorithm 2 outputs an unbiased estimate of the histogram with an expected number of random responses of $\gamma n(K+1)$ and satisfies $(\sqrt{2K\ln(1/\delta')}\varepsilon + K\varepsilon(e^\varepsilon - 1), K\delta + \delta')$-network DP for any $\delta' > 0$.

Walk on a Complete Graph

Theorem 4.

Let $\varepsilon, \delta > 0$. The random walk, run for $T > 0$ steps, satisfies $(\varepsilon', N_v\delta + \delta' + \hat{\delta})$-network DP for all $\delta', \hat{\delta} > 0$ with

$$\varepsilon' = \sqrt{2N_v\log(1/\delta)}\frac{\sqrt{2}\varepsilon}{n^{1/4}} + 2\sqrt{2N_v\gamma_n\log(1/\delta)}\varepsilon + N_v\frac{\sqrt{2}\varepsilon}{n^{1/4}}\left(e^{\frac{\sqrt{2}\varepsilon}{n^{1/4}}} - 1\right) + 4N_v\gamma_n\varepsilon(e^\varepsilon - 1),$$

where $N_v = \frac{T}{n} + \sqrt{\frac{3}{2}T\log(1/\hat{\delta})}$ and $\gamma_n = 1 - (1 - \frac{1}{n})^{\frac{\sqrt{n}}{2}}$.

- Theorem for **real aggregation**, elements of proof:
  - All vertices play a similar role
  - Law of return in a specific node $\mathcal{G}(\frac{1}{n})$
  - Bounding information obtained between two passages, function of the length + use amplification by subsampling [1]
- **Extends to histogram and gradient descent** via amp. by shuffling [4] and iteration [5]
- Experiments with simulated random walks show **important empirical gains** in all tasks

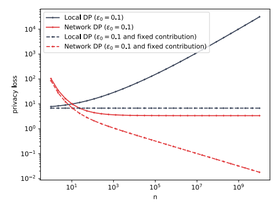

Figure: Real summation (theoretical)

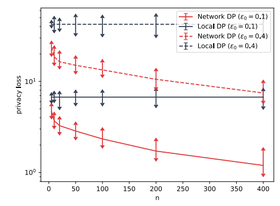

Figure: Discrete histograms

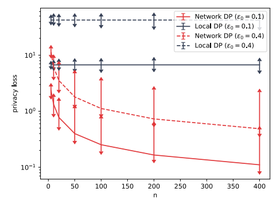

Figure: Real summation (empirical)

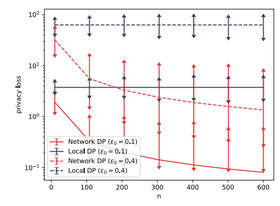

Figure: Gradient descent

* I am looking for a Master's internship, contact me ! edwige.cyffers@ens-lyon.fr