Randomized Quantization is All You Need for Differential Privacy in Federated Learning

Yeojoon Youn [1] **Zihao Hu** [1] **Juba Ziani** [1] **Jacob Abernethy** [1 2]

Abstract

Federated learning (FL) is a common and practical framework for learning a machine model in a decentralized fashion. A primary motivation behind this decentralized approach is data privacy, ensuring that the learner never sees the data of each local source itself. Federated learning then comes with two majors challenges: one is handling potentially complex model updates between a server and a large number of data sources; the other is that de-centralization may, in fact, be insufficient for privacy, as the local updates themselves can reveal information about the sources' data. To address these issues, we consider an approach to federated learning that combines quantization and differential privacy. Absent privacy, Federated Learning often relies on quantization to reduce communication complexity. We build upon this approach and develop a new algorithm called the **R**andomized **Q**uantization **M**echanism (RQM), which obtains privacy through a two-levels of randomization. More precisely, we randomly sub-sample feasible quantization levels, then employ a randomized rounding procedure using these sub-sampled discrete levels. We are able to establish that our results preserve "Renyi differential privacy" (Renyi DP). We empirically study the performance of our algorithm and demonstrate that compared to previous work it yields improved privacy-accuracy trade-offs for DP federated learning. To the best of our knowledge, this is the first study that solely relies on randomized quantization without incorporating explicit discrete noise to achieve Renyi DP guarantees in Federated Learning systems.

[*]Equal contribution [1]Georgia Institute of Technology, Atlanta, GA, USA [2]Google Research. Correspondence to: Yeojoon Youn <yjyoun92@gatech.edu>.

Workshop of Federated Learning and Analytics in Practice, colocated with 40^{th} International Conference on Machine Learning, Honolulu, Hawaii, USA.

1. Introduction

Federated Learning (FL) is an innovative approach to training on massive datasets, utilizing a multitude of devices like smartphones and IoT devices, each containing locally stored, privacy-sensitive data. At a basic level, privacy is maintained by storing local data on each end-user device without sharing it with the server. However, in some cases, device or local device data can be partially reconstructed from computed gradients (Zhu et al., 2019). This potential data leakage from gradients can be addressed through the use of privacy-preserving techniques. Equally important in the context of FL is communication efficiency. Given the extensive communication demands placed on many edge computing devices and the constraints of limited bandwidth, it is imperative to devise a training scheme that not only preserves privacy but also aligns with the requirements of efficient communication within FL systems. What is perhaps surprising, however, is that these two objectives are not necessarily in tension and *can even be aligned!* One way to improve communication overhead is to reduce bit complexity through stochastic rounding schemes, but we show that these randomization procedures, if designed carefully, provide additional benefits to data privacy.

Past studies, such as the work of Agarwal et al. (2018); Kairouz et al. (2021); Agarwal et al. (2021), have sought to tackle this issue by employing various forms of discrete additive DP noise in conjunction with quantization. This is because the application of continuous noise post-quantization would effectively render the noise update continuous, negating any benefits to communication efficiency. However, when these discrete additive noise methods are coupled with secure aggregation protocols (Bonawitz et al., 2017), aimed at preventing a server from inspecting individual local device updates, they encounter a challenge of biased estimation due to gradient clipping. To solve this, Chen et al. (2022) introduce the Poisson Binomial Mechanism (PBM), bypassing the use of additive noise and instead directly mapping continuous inputs to discrete values in an unbiased fashion.

Losses in accuracy compared to noise-free gradient updates that do not protect privacy in the strong sense afforded by differential privacy are essentially unavoidable. The ad-

dition of a privacy requirement inevitably constraints the learner's problem, and privacy must be traded-off with accuracy. Yet, the solution provided by Chen et al. (2022), while providing good performance, may still enjoy sub-optimal privacy-accuracy trade-offs. Can we develop new mechanisms with improved privacy-accuracy trade-off compared to the mechanism of Chen et al. (2022)?

Our starting point to address this question is to note that much research focusing on quantization in federated learning for the sake of communication complexity absent privacy (Alistarh et al., 2017; Reisizadeh et al., 2020; Haddadpour et al., 2021; Youn et al., 2022) reveal that performance degradation from quantization alone is somewhat minimal. Furthermore, quantization itself inherently reduces the amount of information encoded about the original input. While quantization in itself is insufficient for privacy, we posit that a two-stage approach, selecting a *randomized* quantization scheme followed by randomized rounding, can provide a viable approach to obtaining low communication complexity, formal *differential privacy* guarantees, while still enjoying good performance. Thus:

> *Can we harness randomization in quantization schemes to further improve privacy-accuracy trade-offs in differentially private federated learning?*

To address this question, we introduce what we call the *Randomized Quantization Mechanism*, or *RQM* for short. RQM achieves privacy entirely through randomly sub-sampling quantization levels followed by a (randomized) rounding procedure to a close-by quantization level.

Summary of Contributions. As mentioned above, our paper studies mechanisms for releasing gradients while satisfying Renyi differential privacy, and how our proposed mechanisms can be integrated in standard federated learning frameworks.

- In Section 2, we introduce our *Randomized Quantization Mechanism* that maps gradients to a randomized discrete grid in a way that preserves Renyi differential privacy.
- In Section 2.3, we provide theoretical evidence that our proposed Randomized Quantization Mechanism exhibits $\alpha-$Renyi differential privacy guarantees "locally", at the level of each single end-user device. Our theoretical guarantees hold for $\alpha \to +\infty$, implying in particular that they hold not just for Renyi but also for traditional $(\epsilon, 0)$-differential privacy.
- In Section 3, we provide experiments that highlight the performance of our mechanism. In particular, we show that RQM outperforms the state-of-the-art Poisson Binomial Mechanism (PBM) introduced by Chen et al. (2022) in two ways. First, we show that for any given α, RQM provides lower Renyi divergence hence better Renyi DP guarantees than the work of Chen et al. (2022). Second, we also show that when RQM is plugged in the standard differentially private federated learning framework, it leads to high model accuracy when compared to that demonstrated by Chen et al. (2022).

2. The Randomized Quantization Mechanism

In this section, we introduce our main building block for privacy in federated learning. This building block provides a mechanism for privately releasing a scalar aggregate statistic of a single user's data in the form of a new algorithm called the *Randomized Quantization Mechanism* (RQM). We remark that when dealing with f-dimensional vectors instead, we apply our Randomized Quantization Mechanism independently to each vector coordinate. The model explanation of a federated learning set-up with RQM can be found in Appx. D.

We first formally present our RQM mechanism, outlined in Algorithm 1. Since our mechanism relies on a discrete probability distribution to choose the quantization, we show how this probability distribution over the quantizations translates into a probability distribution over outcomes of our mechanism on any given input x; this distribution over outcomes is crucial to characterize the level of privacy obtained by our mechanism. Finally, we theoretically analyze the Renyi differential privacy guarantees of RQM by using this distribution over outcomes.

2.1. Randomized Quantization Mechanism

In this section, we assume that each user outputs a continuous scalar input x computed from their data; this can be viewed as the simplest case of local updates.

Our RQM algorithm is then comprised of three key components: (1) enlarging the output range beyond the input range and setting up evenly spaced quantization bins, (2) sub-sampling realized quantization levels, and (3) performing a randomized rounding procedure on the *sub-sampled* (and only those) discrete levels to map a value x to a quantization level. Each of these steps is crucial in ensuring the Renyi DP guarantees of the RQM, as we describe below. Formally, in each step, we perform the following operations:

1. We establish the output range of our mechanism by first augmenting the size of the input range. We do so by adding Δ to the upper bound c and subtracting Δ from the lower bound $-c$ on the input data.

This augmentation of the range is necessary for privacy: if we use the same range for the output, the quantization output for the maximum input ($x = c$) would subsequently always be c subsequently, leaking a lot of information about x. After this, we establish m initial, evenly spaced quantization levels $(B(0), B(1), \cdots, B(m-1))$ within this output range, which will be potential outputs of our mechanism.

2. Instead of using the entire set of quantization levels from step (1), we randomly sub-sample feasible quantization levels. We do so by including each discrete level for quantization with a carefully chosen probability q. The randomization of the quantization levels is necessary for privacy; otherwise, a value of x would always map to the fixed set of two quantization levels deterministically. This immediately breaks differential privacy.

3. We perform quantization on the sub-sampled discrete levels (and these sub-sampled levels only), achieving both robust privacy and unbiased estimation. We identify the quantization bin that houses the input x and perform randomized rounding on x within this interval. The specific probabilities employed for randomized rounding can be reviewed in Algorithm 1.

Algorithm 1 Randomized Quantization Mechanism

Input: $c > 0, x \in [-c, c]$, extend the upper bound and lower bound by Δ, the maximum number of quantization levels m, include a certain quantization level with probability q
Set X^{max}: $X^{\text{max}} = c + \Delta$, max and min value of quantization levels is respectively X^{max}, $-X^{\text{max}}$.
Quantization bins: $i = 0, 1, \cdots, m-1 \rightarrow B(i) = -X^{\text{max}} + \frac{2iX^{\text{max}}}{m-1}$.
sub-sample feasible quantization levels:
    Always include $B(0), B(m-1)$ & $i = 1, 2, \cdots, m-2 \rightarrow$ include $B(i)$ with probability q.
    sub-sampled indices of quantization levels $\rightarrow i_1(=0), i_2, \cdots, i_l(=m-1)$.
Quantization step:
    Find $i_{j^*}(i_1 \leq i_{j^*} \leq i_l)$ that satisfies $x \in [B(i_{j^*}), B(i_{j^*+1})]$.
    Do randomized rounding on x in this interval.

$$z = \begin{cases} i_{j^*+1}, & \text{with probability } \frac{x - B(i_{j^*})}{B(i_{j^*+1}) - B(i_{j^*})} \\ i_{j^*}, & \text{o/w} \end{cases}$$

return z

2.1.1. Flexibility of RQM hyperparameters

The hyperparameters within our RQM algorithm offer enhanced flexibility, allowing for a more nuanced hyperparameter optimization when compared to PBM. RQM has in fact three hyperparameters Δ, q, m, while PBM has two hyperparameters θ, m (See Algorithm 2 in Chen et al. (2022)). At a fixed number of discrete levels m, i.e. at a fixed level of communication complexity, this allows us to search over a bigger space of output distributions of quantization levels than Chen et al. (2022) through the choice of (q, Δ). In Section 3.2, we show that this leads to RQM achieving better privacy-accuracy trade-offs than PBM.

2.2. Resulting Discrete Distribution of Outcomes

Given an input x and parameters m, q, Δ, we can compute the discrete probability distribution of outputs $Q(x)$ of RQM over the set of potential quantization levels $B(0), B(1), \cdots B(m-1)$. This discrete probability distribution given in Lemma E.1 (See equation (3)) and the full proof of Lemma E.1 are provided in Appx. E.1. Equation (3) exhibits different cases for $\Pr(Q(x) = i)$. These cases depend on i) how the i-th quantization level compares to j, where j is defined to be such that $x \in [B(j), B(j+1))$ and ii) on the two special cases $i = 0$ or $i = m-1$. The probabilities corresponding to these two extreme values of i differ from the rest in that we always incorporate the 0-th and $(m-1)$-th discrete level, which influences our probability calculations.

Figures 3a and 3b in Appx. C provide some insights into how to derive the distribution of outputs and gives some intuition for Lemma E.1. In Figure 3a, the solid lines corresponding to the $(0, 3, 4, 7, 9, 10, 14, 15)$-th discrete levels have been selected for quantization, while the dotted lines $(1, 2, 5, 6, 8, 11, 12, 13)$-th discrete levels have been thrown away. To have $Q(x) = 10$, the 10-th discrete level must always be chosen while the 11-th discrete level must not be chosen by the sub-sampling step of our algorithm; the probability of this happening is $q(1-q)$. The probability of the next quantization level bigger than x being the 14-th level, as shown in Figure 3a, is similarly given by $q(1-q)^2$ (levels 12 and 13 must not be sub-sampled, but 14 must be). Then, the likelihood of x transitioning to the 10-th discrete level due to randomized rounding between the 10-th and 14-th levels is $\frac{B(14)-x}{B(14)-B(10)}$. I.e., the situation described in Figure 3a happens with probability $q^2(1-q)^3 \frac{B(14)-x}{B(14)-B(10)}$. For a complete analysis, we must also account for randomized rounding intervals $[B(10), B(12)], [B(10), B(13)], [B(10), B(15)]$ and aggregate all these probabilities.

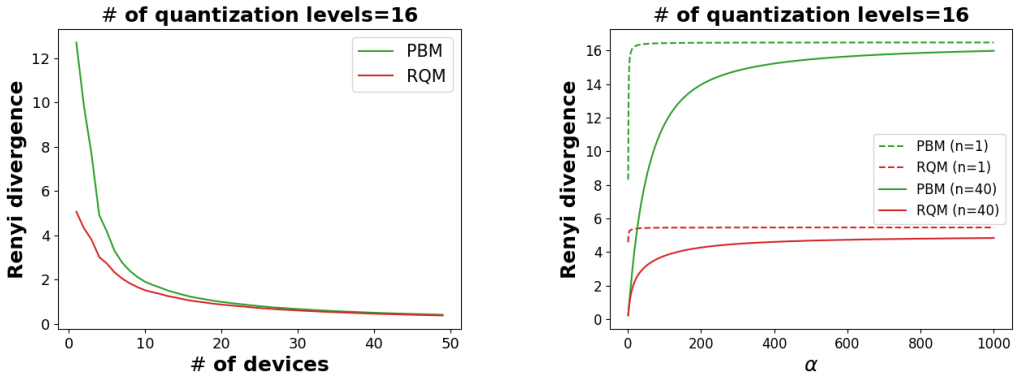

Figure 1. The left figure illustrates the inverse relationship between the number of devices n and the upper bound of the Renyi divergence. The right figure indicates how the Renyi divergence increases as α increases.

2.3. Theoretical Analysis of RQM's Privacy Guarantees

We now provide a theoretical analysis of the level of Renyi-differential privacy achieved by our single-dimensional RQM mechanism. The full proof of Theorem 2.1 is provided in Appx. E.2.

Theorem 2.1. *Let $c, \Delta > 0$, $m \in \mathbb{N}$, and $q \in (0,1)$ be parameters of Algorithm 1. Consider two scalars x and x' in $[-c, c]$, $P_{Q(x)}$ the distribution of outputs of the quantization mechanism ran on scalar x, and $P_{Q(x')}$ the distribution of outputs of the quantization mechanism ran on scalar x'. We have:*

$$D_\alpha(P_{Q(x)}||P_{Q(x')}) \leq D_\infty(P_{Q(x)}||P_{Q(x')}) \\ \leq \log\left(2(1-q)^2\left(1+\frac{c}{\Delta}\right)\right) + m\log\frac{1}{1-q}. \quad (1)$$

Our bound focuses on the case where $\alpha \to +\infty$, in which case (α, ϵ)-Renyi differential privacy is in fact the same as $(\epsilon, 0)$-differential privacy as per the traditional definition of DP. There, we note that the privacy level $\epsilon = \log\left(2(1-q)^2\left(1+\frac{c}{\Delta}\right)\right) + m\log\frac{1}{1-q}$ that we obtain increases linearly on m, the number of quantizations level. This makes sense as a large number of quantization levels allows one to encode more information about the initial scalar x, in turn leading to less privacy and higher ϵ's. We also note that as Δ increases, ϵ decreases, and we obtain more privacy; once again, this follows the intuition from Section 2.1 that when we increase the output range, we better protect the privacy of extreme values of x that are close to c or $-c$. As expected, when $\Delta = 0$, $\epsilon \to +\infty$ and our privacy guarantees are trivial, highlighting the fact that augmenting the range of output values beyond $[-c, c]$ is an unavoidable step to obtain reasonable privacy guarantees.

3. Experiments

In this section, we conduct experiments designed to complement our theoretical results and illustrate how RQM performs compared to PBM in terms of the privacy-accuracy trade-off. We first establish that RQM provides superior Renyi DP guarantees for the multiple-device scenario by numerically calculating the upper bound of Renyi divergence. Then, we employ the same parameters that yielded improved Renyi DP to demonstrate that RQM also excels in accuracy in our federated learning experiments.

3.1. Numerical Renyi Privacy Guarantees

In Section 2.3, we characterized the privacy guarantee of our Randomized Quantization Mechanism in the special case in which $\alpha \to +\infty$. Our privacy guarantees hold at the *local* level, in that it protects against a strong adversary that can see the output $Q(x_i)$ of each device i (but not the input data x_i). A weaker but potentially interesting adversary can only see the output of the trusted, secure aggregator SecAgg[1]. In this case, assuming we have n devices that provide scalar inputs $x_1, \ldots, x_n$—where x_i is the input of device i—to the quantization mechanism RQM[2], we are interested in an adversary that only sees the aggregate quantity $\sum_{i=1}^n Q(x_i)$. Given two vectors of inputs x and x' that only differ in the input x_i of a single device i, we obtain (α, ϵ)-Renyi differential privacy with

$$\epsilon \triangleq D_\alpha\left(P_{\sum_{i=1}^n Q(x_i)}||P_{\sum_{i=1}^n Q(x_i')}\right),$$

as studied by Chen et al. (2022).

We use Equation (3) to numerically compute and plot this

[1]The learner only sees the output of SecAgg and is such an adversary.

[2]We can think of these inputs as summary statistics computed on databases $D_1, \ldots, D_n$ before adding privacy; e.g., these could be a single coordinate of a clipped gradient.

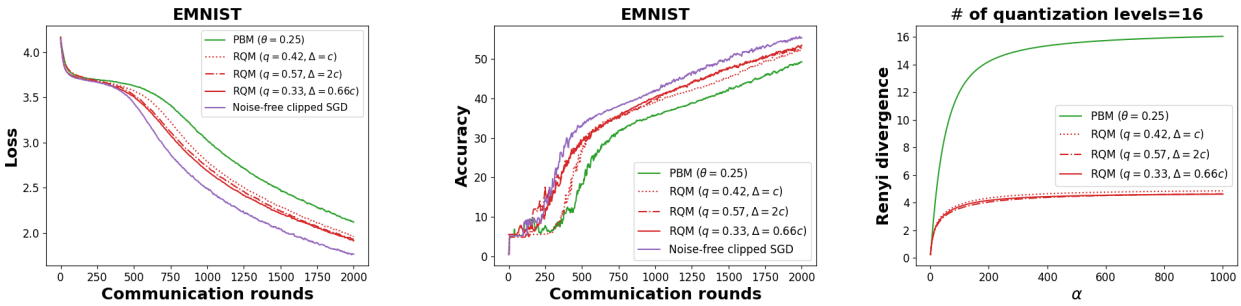

Figure 2. Comparing RQM with PBM and noise-free clipped SGD on EMNIST. All three RQMs with different hyperparameters outperform PBM in both a loss plot (Left) and an accuracy plot (Middle). These RQMs also show better Renyi DP guarantees than PBM (Right).

Renyi divergence for finite α and for $n \geq 1$ in Figure 1. We compare it to the Renyi divergence of the Poisson Binomial Mechanism of Chen et al. (2022); we note that we do not compare to the upper bound provided by Chen et al. (2022) that may not be tight, but instead to the actual Renyi divergence computed numerically and *exactly*. In both cases, we plot the nearly *worst-case* (over x, x') Renyi divergence, which is approximately maximized when all x_i's are either $-c$ or c[3]; to do so, we generate x and x' by taking $x_1 = c, x_1' = -c$ and randomly assigning a value of either c or $-c$ to $x_2, \cdots, x_n$. The hyperparameter choice details for the DP experiment are provided in Appx. F.1.

The left picture of Figure 1 fixes $\alpha = 2$ and compares the Renyi divergences of PBM and RQM when n increases. At the same time, the Renyi divergences seem to converge to similar values for $n \to +\infty$, and we note that our framework performs better for finite values of n, with a noticeably increasing performance gap as the number of user devices n becomes smaller. The right picture of Figure 1 fixes $n = 1$, then $n = 40$, and compares the Renyi divergence of PBM and RQM for a large range of $\alpha \in [0, 1000]$; we see significant disparities in the levels of Renyi privacy guaranteed by PBM and RQM, with RQM vastly outperforming (i.e., guaranteeing a lower Renyi divergence hence a better privacy guarantee than) PBM, with the gap in privacy guarantees increasing as $\alpha \to +\infty$.

3.2. Federated Learning Experiments

In this section, we expand our experiments to provide insights beyond the privacy guarantees of the Randomized Quantization Mechanism itself, and to take into account how RQM integrates with the rest of the federated learning framework described in Appx. D. We implement multi-dimensional RQM within the federated DP-SGD algorithm described in Algorithm 2.

[3]More details on this are discussed in Appx. F.1.1.

We evaluate the privacy-accuracy trade-off of our algorithms against the current leading approach, the Poisson Binomial Mechanism (Chen et al., 2022), and an ideal noise-free clipped SGD benchmark that does not provide any differential privacy guarantee. The classification task for our federated learning experiment is performed on the EMNIST dataset (Cohen et al., 2017). The experimental setup details are provided in Appx F.2.

Experimental Results. The left and middle plots of Figure 2 clearly show that all three RQMs with different hyperparameter pairs show improved performance (in terms of loss and accuracy) on the EMNIST dataset than PBM. Among three RQMs, $(\Delta, q) = (0.66c, 0.33)$ achieves the best accuracy. The performance of our three RQMs are still worse than noise-free clipped SGD: this is unavoidable because noise-free clipped SGD only focus on accuracy without providing any privacy guarantees, and is an ideal, impossible-to-achieve benchmark with privacy.

The right plot in Figure 2 replicates experiment of Section 3.1 that were aimed at showcasing the privacy level achieved by RQM compared to PBM. The figure shows that the improved accuracy of the three RQMs compared to PBM in the left and middle figures does not come at the cost of privacy. In fact, the three plots together demonstrate that all three instantiations of RQM provide both better performance and better Renyi DP guarantees than PBM. I.e., in our experiments, RQM improves the *privacy-accuracy trade-off* of federated differentially private stochastic gradient descent compared to the current state of the art.

Acknowledgements

We are grateful to Peter Kairouz, Wei-Ning Chen, Ayfer Özgür, and all other anonymous reviewers for their invaluable comments. This work was made possible through the support of the AI4OPT Institute under the NSF Award 2112533, and the support of the NSF Award IIS-1910077.

References

Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL https://www.tensorflow.org/. Software available from tensorflow.org.

Agarwal, N., Suresh, A. T., Yu, F. X. X., Kumar, S., and McMahan, B. cpsgd: Communication-efficient and differentially-private distributed sgd. *Advances in Neural Information Processing Systems*, 31, 2018.

Agarwal, N., Kairouz, P., and Liu, Z. The skellam mechanism for differentially private federated learning. *Advances in Neural Information Processing Systems*, 34: 5052–5064, 2021.

Alistarh, D., Grubic, D., Li, J., Tomioka, R., and Vojnovic, M. Qsgd: Communication-efficient sgd via gradient quantization and encoding. *Advances in neural information processing systems*, 30, 2017.

Bell, J. H., Bonawitz, K. A., Gascón, A., Lepoint, T., and Raykova, M. Secure single-server aggregation with (poly) logarithmic overhead. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1253–1269, 2020.

Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., and Seth, K. Practical secure aggregation for privacy-preserving machine learning. In *proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1175–1191, 2017.

Chaudhuri, K., Guo, C., and Rabbat, M. Privacy-aware compression for federated data analysis. In *Uncertainty in Artificial Intelligence*, pp. 296–306. PMLR, 2022.

Chen, W.-N., Ozgur, A., and Kairouz, P. The poisson binomial mechanism for unbiased federated learning with secure aggregation. In *International Conference on Machine Learning*, pp. 3490–3506. PMLR, 2022.

Cohen, G., Afshar, S., Tapson, J., and Van Schaik, A. Emnist: Extending mnist to handwritten letters. In *2017 international joint conference on neural networks (IJCNN)*, pp. 2921–2926. IEEE, 2017.

Dwork, C., Kenthapadi, K., McSherry, F., Mironov, I., and Naor, M. Our data, ourselves: Privacy via distributed noise generation. In *Advances in Cryptology-EUROCRYPT 2006: 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28-June 1, 2006. Proceedings 25*, pp. 486–503. Springer, 2006.

Dwork, C., Roth, A., et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.

Erlingsson, Ú., Pihur, V., and Korolova, A. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, pp. 1054–1067, 2014.

Gandikota, V., Kane, D., Maity, R. K., and Mazumdar, A. vqsgd: Vector quantized stochastic gradient descent. In *International Conference on Artificial Intelligence and Statistics*, pp. 2197–2205. PMLR, 2021.

Geyer, R. C., Klein, T., and Nabi, M. Differentially private federated learning: A client level perspective. *arXiv preprint arXiv:1712.07557*, 2017.

Guo, C., Chaudhuri, K., Stock, P., and Rabbat, M. Privacy-aware compression for federated learning through numerical mechanism design. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J. (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 11888–11904. PMLR, 23–29 Jul 2023. URL https://proceedings.mlr.press/v202/guo23a.html.

Haddadpour, F., Kamani, M. M., Mokhtari, A., and Mahdavi, M. Federated learning with compression: Unified analysis and sharp guarantees. In *International Conference on Artificial Intelligence and Statistics*, pp. 2350–2358. PMLR, 2021.

Kairouz, P., Liu, Z., and Steinke, T. The distributed discrete gaussian mechanism for federated learning with secure aggregation. In *International Conference on Machine Learning*, pp. 5201–5212. PMLR, 2021.

Levy, D., Sun, Z., Amin, K., Kale, S., Kulesza, A., Mohri, M., and Suresh, A. T. Learning with user-level privacy. *Advances in Neural Information Processing Systems*, 34, 2021.

Li, T., Liu, Z., Sekar, V., and Smith, V. Privacy for free: Communication-efficient learning with differential privacy using sketches. *arXiv preprint arXiv:1911.00972*, 2019.

McMahan, B., Moore, E., Ramage, D., Hampson, S., and y Arcas, B. A. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pp. 1273–1282. PMLR, 2017a.

McMahan, H. B., Ramage, D., Talwar, K., and Zhang, L. Learning differentially private recurrent language models. *arXiv preprint arXiv:1710.06963*, 2017b.

Mironov, I. Rényi differential privacy. In *2017 IEEE 30th computer security foundations symposium (CSF)*, pp. 263–275. IEEE, 2017.

Reisizadeh, A., Mokhtari, A., Hassani, H., Jadbabaie, A., and Pedarsani, R. Fedpaq: A communication-efficient federated learning method with periodic averaging and quantization. In *International Conference on Artificial Intelligence and Statistics*, pp. 2021–2031. PMLR, 2020.

Rényi, A. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, volume 4, pp. 547–562. University of California Press, 1961.

Youn, Y., Kumar, B., and Abernethy, J. Accelerated federated optimization with quantization. In *Workshop on Federated Learning: Recent Advances and New Challenges (in Conjunction with NeurIPS 2022)*, 2022.

Zhu, L., Liu, Z., and Han, S. Deep leakage from gradients. *Advances in neural information processing systems*, 32, 2019.

A. Related Work

Both communication complexity and privacy concerns have been driving forces behind the development of Federated Learning. Federated optimization often uses two types of privacy-preserving techniques hand-in-hand. One is secure multi-party computation, which protects the communication between local devices and the learner, preventing an attacker from intercepting messages sent between them (Bell et al., 2020; Bonawitz et al., 2017). One is information-theoretic privacy guarantees such as differential privacy (Dwork et al., 2014) that prevent inference of any given single local device's data from observed summary outputs (such as local gradient updates or the learner's model itself). For example, McMahan et al. (2017b) and Geyer et al. (2017) add a calibrated amount of Gaussian noise to the average of clipped local device updates based on the FedAvg (McMahan et al., 2017a) algorithm.

In this paper, we focus on providing robust Renyi differential privacy guarantees in federated optimization while maintaining high communication efficiency and good accuracy. Previous methods have often used an approach based on quantization followed by the addition of discrete noise to achieve both differential privacy guarantees and low communication efficiency. Agarwal et al. (2018) introduces the first communication-efficient federated optimization algorithm with differential privacy by incorporating quantization with the binomial mechanism. Kairouz et al. (2021) and Agarwal et al. (2021) employ discrete Gaussian and Skellam mechanisms, respectively, in conjunction with quantization and secure aggregation for enhanced privacy. However, the above methods lead to biased estimation due to the necessity of modular clipping. To address this issue, Chen et al. (2022) and Chaudhuri et al. (2022) propose unbiased mechanisms with improved privacy-accuracy trade-offs. Chen et al. (2022) encodes local devices' gradients into a parameter of the binomial distribution, allowing their mechanism to generate a sample from this distribution without the need for additive discrete noise. In contrast, rather than using known privacy mechanisms, Chaudhuri et al. (2022) introduces the *Minimum Variance Unbiased mechanism* (MVU) to enhance the privacy-utility trade-off by solving an optimization problem designed to minimize the output variance of the mechanism while adhering to local differential privacy and unbiasedness constraints. Enhancing this model, Guo et al. (2023) propose a more scalable MVU mechanism with better privacy-utility trade-off, achieved through a new interpolation procedure in the numerical design process. Despite their progress in improving the privacy-utility trade-off, these methods do not fully exploit the privacy advantages offered by randomized quantization itself.

Our research is not the first to leverage compression techniques to achieve both communication efficiency and provable privacy benefits without incorporating additive discrete noise (Li et al., 2019; Gandikota et al., 2021). (Li et al., 2019) assume a Gaussian input vector distribution for their sketching algorithms to ensure differential privacy guarantees, which might not be strictly necessary. Gandikota et al. (2021) provide multiple vector quantization schemes which require a lower bound on the achievable level of privacy ϵ. Then, they add differential private mechanisms such as randomized response or Rappor (Erlingsson et al., 2014) to the vector quantization schemes to achieve DP for any $\epsilon > 0$. However, and to the best of our knowledge, our Randomized Quantization Mechanism is the first investigation that exclusively utilizes randomization of the quantization itself to attain improved Renyi DP guarantees within Federated Learning frameworks.

B. Preliminaries

B.1. Differential Privacy

Given that we employ the concept of differential privacy to demonstrate the privacy guarantees of our Randomized Quantization Mechanism, we first present the original definition of differential privacy.

Definition B.1. ((Approximate) Differential Privacy (Dwork et al., 2006)) For $\epsilon, \delta \geq 0$, a randomized mechanism $\mathcal{M} : \mathcal{D} \to \mathcal{R}$ satisfies (ϵ, δ)-differential privacy if for any neighbor dataset $D, D' \in \mathcal{D}$ differing by the addition or removal of a single user's records, it holds that

$$\Pr(\mathcal{M}(D) \in E) \leq e^{\epsilon} \cdot \Pr(\mathcal{M}(D') \in E) + \delta$$

for all events $E \subset \mathcal{R}$.

In this paper, we also consider a variant of standard differential privacy called *Renyi Differential Privacy* (or *Renyi DP*), introduced in the seminal work of Mironov (2017). We develop mechanisms that guarantee Renyi DP and by extension traditional DP. The use of Renyi DP allows for tight privacy accounting throughout the training iterations. Renyi differential privacy relies on first understanding the notion of *Renyi divergence*:

Definition B.2. (Renyi Divergence (Rényi, 1961)) Let P and Q be probability distributions defined over $\mathcal{R}$. The Renyi

divergence of order $\alpha > 1$ is defined as

$$D_\alpha(P||Q) := \frac{1}{\alpha - 1} \log \mathbb{E}_{x\sim Q}\Big[\Big(\frac{P(x)}{Q(x)}\Big)^\alpha\Big].$$

Then, Renyi differential privacy is defined as follows:

Definition B.3. (Renyi Differential Privacy (Mironov, 2017)) A randomized mechanism $\mathcal{M} : \mathcal{D} \to \mathcal{R}$ satisfies (α, ϵ)-Renyi differential privacy if for any neighbor dataset $D, D' \in \mathcal{D}$ it holds that

$$D_\alpha(P_{\mathcal{M}(D)}||P_{\mathcal{M}(D')}) \leq \epsilon. \tag{2}$$

When $\alpha \to \infty$, (α, ϵ)-Renyi DP in fact recovers standard $(\epsilon, 0)$-DP. However, Renyi DP provides a finer-grained definition of privacy in that its guarantees can be tailored to the specific value of α and corresponding Renyi divergence that one considers. We now state a major property of the Renyi divergence that is useful to our theoretical analysis.

Lemma B.4. *(Monotonicity) D_α is nondecreasing in α. I.e., $D_\alpha(P||Q) \leq D_{\alpha'}(P||Q)$ for all $1 \leq \alpha \leq \alpha' \leq \infty$.*

B.2. User-Level Privacy

In the context of federated learning, we employ differential privacy to mask the contribution of any individual local device, making it challenging for a potential adversary to discern whether a local device's dataset was utilized in the training process. As such, we need to extend the traditional item-level definition of differential privacy (Definition B.1) by redefining what we mean by neighboring datasets. In this context, two datasets are considered neighboring if one dataset can be created by changing any subset of data points of a single user from the other dataset. This user-level perspective is relatively standard and is the same as the one studied by McMahan et al. (2017b) and Levy et al. (2021).

Definition B.5. (User-level DP (Levy et al., 2021)) For $\epsilon, \delta \geq 0$, a randomized mechanism $\mathcal{M} : \mathcal{D} \to \mathcal{R}$ satisfies (ϵ, δ)-user level DP if for any neighbor dataset $D, D' \in \mathcal{D}$ satisfying $\mathrm{d}_{\text{user}}(D, D') \leq 1$, it holds that

$$\Pr(\mathcal{M}(D) \in E) \leq e^\epsilon \cdot \Pr(\mathcal{M}(D') \in E) + \delta$$

for all events $E \subset \mathcal{R}$, where d_{user} is defined with n users as

$$D = (D_1, \cdots, D_n), \text{ where } D_i = \{z_{i,1}, \cdots, z_{i,m_i}\} \to \mathrm{d}_{\text{user}}(D, D') := \sum_{i=1}^{n} 1\{D_i \neq D_i'\}$$

C. An Example of RQM

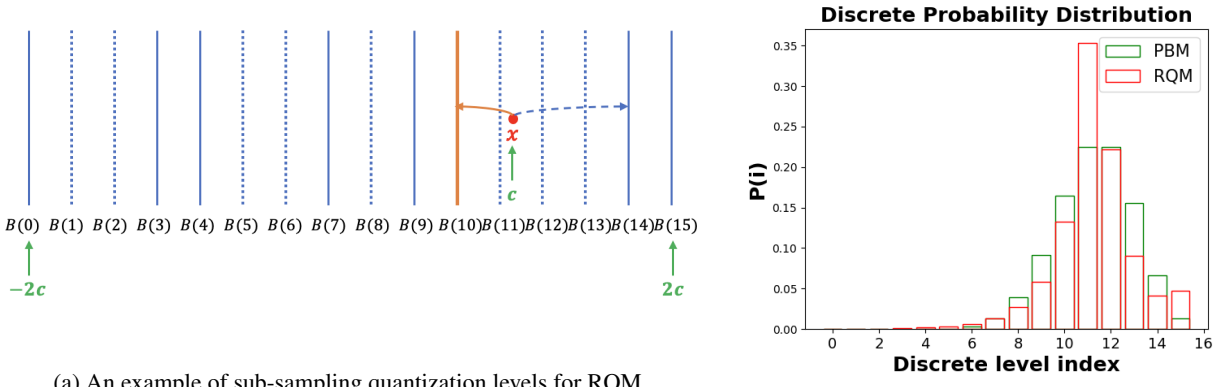

(a) An example of sub-sampling quantization levels for RQM.

(b) Distribution of outputs $Q(x)$ under PBM and RQM.

Figure 3. An example of RQM with input $x = c$ and parameters $\Delta = c,\ m = 16$. The right figure (b) provides some insights on what the distribution induced by our quantization mechanism looks like, and seems to evidence that its shape differs from that of PBM.

D. Model

We consider a federated learning set-up comprised of three types of entities: there are n end-user devices, one secure aggregator called *SecAgg*, and one learner. The learner's goal is to learn a machine learning model, parameterized by a f-dimensional vector $w \in \mathbb{R}^f$, using the data on the devices through Stochastic Gradient Descent (SGD). However, the learner does not access the data from the devices directly, both for communication efficiency and privacy reasons. Rather, at each time step t:

1. Each end-user device i computes a clipped gradient $g_t^i \in [-c, c]^f$ locally using the data on that device only. This gradient is then encoded into an integer z_t^i. This integer can be seen as the index of a discrete level in a discretization of the space of potential gradients.

2. The secure aggregator receives one message z_t^i from each device i, which encodes information about the gradient g_t^i computed by i. The aggregator aggregates them into a single message $z_t = \sum_i z_t^i$.

3. The server decodes z_t, computes the corresponding gradient $\hat{g}_t$, and takes a gradient step $w_{t+1} \leftarrow w_t - \eta \hat{g}_t$.

The traditional approach to federated learning releases gradients exactly; this approach is, however, i) inefficient from a communication complexity perspective and ii) vulnerable when it comes to privacy. We address i) by discretizing (or "quantizing") the space of possible values of the gradients to a grid of size m per coordinate of the gradient; in turn, we require only $f \times \log m$ bits to represent a single update by a single device. Regarding ii), we note that it is well-understood that releasing exact gradients can lead to privacy violations in that the secure aggregator and the learner can recover information about device i's dataset D_i through the gradient itself. To address this issue, instead of releasing the gradient g_t^i directly, device i releases a noisy quantization level $z_t^i = RQM(g_t^i)$, where RQM is a Randomized Quantization Mechanism that must satisfy Renyi differential privacy. The entire setup is described formally in Algorithm 2.

What algorithm 2 describes is essentially the well-known, generic *Differentially Private Stochastic Gradient Descent* approach to federated learning (McMahan et al., 2017b). The focus and novelty of our work, however, come from the design of the private mechanism RQM itself, which is proposed in Section 2.

Algorithm 2 Distributed DP-SGD with RQM

Input: N local devices, each local device dataset $D_i \in \mathcal{D}$ $(i = 1, \cdots, N)$, clipping threshold c, RQM parameters (Δ, m, q), server learning rate η, initial vector w_0, loss function $f(w, D)$
for $t = 0, \cdots, T-1$ **do**
  Server broadcasts w_t to n sampled local devices from total N local devices;
  for each local device i in parallel **do**
    $g_t^i \leftarrow \text{Clip}(\nabla f(w_t, D_i))$;
    $z_t^i \leftarrow \text{RQM}(g_t^i)$;
    send z_t^i to the secure aggregator SecAgg.
  end for
  SecAgg outputs $z_t = \sum_{i=1}^n z_t^i$;
  server decodes $\hat{g}_t \leftarrow -(c + \Delta) + \frac{2z_t(c+\Delta)}{n(m-1)}$;
  server finds $w_{t+1} \leftarrow w_t - \eta \hat{g}_t$.
end for

E. Missing Proofs

E.1. Proof of Lemma E.1

Lemma E.1. *Let $m \in \mathbb{N}$, and $q \in (0, 1)$ be parameters of Randomized Quantization Mechanism Q. Define evenly spaced m quantization levels $B(0), \cdots, B(m-1)$ as in Algorithm 1. Let j be the unique integer such that $x \in [B(j), B(j+1))$.*

The probability distribution of outcomes of the Randomized Quantization Mechanism is given by:

$$Pr\Big(Q(x)=i\Big)=\begin{cases}(1-q)^{j-i}\Big((1-q)^{m-j-2}\frac{B(m-1)-x}{B(m-1)-B(i)}+\sum_{k=j+1}^{m-2}q(1-q)^{k-j-1}\frac{B(k)-x}{B(k)-B(i)}\Big), & i=0,\\ q(1-q)^{j-i}\Big((1-q)^{m-j-2}\frac{B(m-1)-x}{B(m-1)-B(i)}+\sum_{k=j+1}^{m-2}q(1-q)^{k-j-1}\frac{B(k)-x}{B(k)-B(i)}\Big), & 0<i\le j,\\ q(1-q)^{i-j-1}\Big((1-q)^{j}\frac{x-B(0)}{B(i)-B(0)}+\sum_{k=1}^{j}q(1-q)^{j-k}\frac{x-B(k)}{B(i)-B(k)}\Big), & j+1\le i<m-1,\\ (1-q)^{i-j-1}\Big((1-q)^{j}\frac{x-B(0)}{B(i)-B(0)}+\sum_{k=1}^{j}q(1-q)^{j-k}\frac{x-B(k)}{B(i)-B(k)}\Big), & i=m-1.\end{cases} \tag{3}$$

Proof of Lemma E.1 We divide the range of i into four cases-$0<i\le j$, $i=0$, $j+1\le i<m-1$, $i=m-1$- and compute the discrete probability $\Pr(Q(x)=i)$ for each case. The core proof idea of Lemma E.1 is centered on evaluating the probability of each potential interval that can be used for randomized rounding for x. Subsequently, the probability that $Q(x)=i$ arises due to randomized rounding within a given interval is computed. Thus, when $i\le j$ and $k\ge j+1$, we define the event E_i and F_k as below to use this notation for calculating the probability of each potential interval that can be used for the randomized rounding step in Algorithm 1.

$$\begin{aligned}E_i&:\ \text{the event of } i\text{-th discrete level being used for randomized rounding}\\ F_k&:\ \text{the event of } k\text{-th discrete level being used for randomized rounding}\end{aligned} \tag{4}$$

From the above definition of two events, $E_i\cap F_k$ indicates an event of the interval $[B(i),B(k)]$ being used for randomized rounding. In this event, this also means $i_{j^*}=i$ and $i_{j^*+1}=k$ in Algorithm 1. Now, let's deep dive into how we can exactly calculate $\Pr(Q(x)=i)$ for each case of four ranges.

(I) $0<i\le j$:

First, Let us consider the case when $0<i\le j$. Similar to the logic in Section 2.2, to have $Q(x)=i$, the i-th discrete level must always be chosen while the $(i+1)$-th, $\cdots$, j-th discrete levels must not be chosen by the sub-sampling step of our algorithm. The probability of this happening is $q(1-q)^{j-i}$. Thus, we can use the definition of the event E_i in (4) for this case.

$$\Pr(E_i)=\Pr(i:\ \text{chosen},\ (i+1,\cdots,j):\ \text{not chosen})=q(1-q)^{j-i} \tag{5}$$

Let us denote k as an index of the next quantization level bigger than x. The possible ks are $j+1,\cdots,m-1$. When $k\in[j+1,m-2]$, the probability of the next quantization level bigger than x being the k-th level is similarly given by $q(1-q)^{k-j-1}$. Thus, we can use the definition of the event F_k in (4) for this case.

$$\Pr(F_k)=\Pr(k:\ \text{chosen},\ (j+1,\cdots,k-1):\ \text{not chosen})=q(1-q)^{k-j-1} \tag{6}$$

Then, the likelihood of x transitioning to the i-th discrete level due to the randomized rounding between i-th and k-th levels is $\frac{B(k)-x}{B(k)-B(i)}$. This means

$$\Pr(Q(x)=i|E_i\cap F_k)=\frac{B(k)-x}{B(k)-B(i)} \tag{7}$$

Therefore, for $k\in[j+1,m-2]$, by combining (5), (6), (7), we get

$$\begin{aligned}&\Pr((Q(x)=i)\cap E_i\cap F_k)\\ =&\Pr(E_i\cap F_k)\cdot\Pr(Q(x)=i|E_i\cap F_k)\\ =&\Pr(E_i)\cdot\Pr(F_k)\cdot\Pr(Q(x)=i|E_i\cap F_k)\ (\because\ E_i,F_k:\ \text{independent})\\ =&q(1-q)^{j-i}\cdot q(1-q)^{k-j-1}\cdot\frac{B(k)-x}{B(k)-B(i)}\end{aligned} \tag{8}$$

We can perform a similar computation for $k = m-1$. However, the probability of event F_{m-1} is different from that of Equation (6) because the $(m-1)$-th level is always chosen by Algorithm 1. Thus, we have

$$\Pr(F_{m-1}) = \Pr(m-1: \text{ chosen}, \ (j+1, \cdots, m-2): \text{ not chosen}) = (1-q)^{m-j-2} \tag{9}$$

Therefore, for $k = m-1$, by combining Equations (5), (9), and (7), we obtain

$$\begin{aligned}
&\Pr((Q(x)=i) \cap E_i \cap F_{m-1}) \\
=&\Pr(E_i \cap F_{m-1}) \cdot \Pr(Q(x)=i | E_i \cap F_{m-1}) \\
=&\Pr(E_i) \cdot \Pr(F_{m-1}) \cdot \Pr(Q(x)=i | E_i \cap F_{m-1}) \\
=&q(1-q)^{j-i} \cdot (1-q)^{m-j-2} \cdot \frac{B(m-1)-x}{B(m-1)-B(i)}
\end{aligned} \tag{10}$$

Finally, by combining Equations (8) and (10), we get

$$\begin{aligned}
&\Pr(Q(x)=i) \\
=&\sum_{k=j+1}^{m-1} \Pr((Q(x)=i) \cap E_i \cap F_k) \\
=&q(1-q)^{j-i}\Big((1-q)^{m-j-2}\frac{B(m-1)-x}{B(m-1)-B(i)} + \sum_{k=j+1}^{m-2} q(1-q)^{k-j-1}\frac{B(k)-x}{B(k)-B(i)}\Big)
\end{aligned} \tag{11}$$

(II) $i = 0$:

We can compute $\Pr(Q(x)=0)$ in a similar way as in case (I). However, the probability of event E_0 is different from that of Equation (5) because the 0-th level is always chosen by Algorithm 1. Thus, we have

$$\Pr(E_0) = \Pr(0: \text{ chosen}, \ (1, \cdots, j): \text{ not chosen}) = (1-q)^j \tag{12}$$

Therefore, in (11), by substituting E_i into E_0, we get

$$\begin{aligned}
&\Pr(Q(x)=0) \\
=&\sum_{k=j+1}^{m-1} \Pr((Q(x)=0) \cap E_0 \cap F_k) \\
=&(1-q)^j\Big((1-q)^{m-j-2}\frac{B(m-1)-x}{B(m-1)-B(0)} + \sum_{k=j+1}^{m-2} q(1-q)^{k-j-1}\frac{B(k)-x}{B(k)-B(0)}\Big)
\end{aligned} \tag{13}$$

(III) $j+1 \le i < m-1$:

For i within this range, we can similarly compute $\Pr(Q(x)=i)$ as in (I). To obtain $Q(x)=i$, the i-th discrete level must always be chosen while the $(j+1)$-th, $\cdots$, $(i-1)$-th discrete levels must not be chosen by the sub-sampling step of our algorithm. Thus, since $i \ge j+1$, the probability of i-th discrete level being used for randomized rounding can be expressed by using F_i (refer to (6)).

$$\Pr(F_i) = \Pr(i: \text{ chosen}, \ (j+1, \cdots, i-1): \text{ not chosen}) = q(1-q)^{i-j-1} \tag{14}$$

Let us denote k as an index of the just previous level less than x. The possible k's are $0, \cdots, j$. Then, for $k \in [1, j]$, the probability of the k-th discrete level being used for randomized rounding can be represented by utilizing E_k (refer to (5)).

$$\Pr(E_k) = \Pr(k: \text{ chosen}, \ (k+1, \cdots, j): \text{ not chosen}) = q(1-q)^{j-k} \tag{15}$$

Then, the likelihood of x transitioning to the i-th discrete level due to the randomized rounding between k-th and i-th levels is $\frac{x-B(k)}{B(i)-B(k)}$. This means

$$\Pr(Q(x)=i | E_k \cap F_i) = \frac{x-B(k)}{B(i)-B(k)} \tag{16}$$

Therefore, for $k \in [1, j]$, by combining (14), (15), (16), we get

$$\begin{aligned}&\Pr((Q(x)=i)\cap E_k \cap F_i)\\ =&\Pr(F_i)\cdot\Pr(E_k)\cdot\Pr(Q(x)=i|E_k\cap F_i)\\ =&q(1-q)^{i-j-1}\cdot q(1-q)^{j-k}\cdot\frac{x-B(k)}{B(i)-B(k)}\end{aligned} \tag{17}$$

We can similarly calculate for $k=0$ by using (12).

$$\begin{aligned}&\Pr((Q(x)=i)\cap E_0 \cap F_i)\\ =&\Pr(F_i)\cdot\Pr(E_0)\cdot\Pr(Q(x)=i|E_0\cap F_i)\\ =&q(1-q)^{i-j-1}\cdot(1-q)^{j}\cdot\frac{x-B(0)}{B(i)-B(0)}\end{aligned} \tag{18}$$

Finally, by combining (17) and (18)

$$\begin{aligned}&\Pr(Q(x)=i)\\ =&\sum_{k=0}^{j}\Pr((Q(x)=i)\cap E_k\cap F_i)\\ =&q(1-q)^{i-j-1}\Big((1-q)^j\frac{x-B(0)}{B(i)-B(0)}+\sum_{k=1}^{j}q(1-q)^{j-k}\frac{x-B(k)}{B(i)-B(k)}\Big)\end{aligned} \tag{19}$$

(IV) $i=m-1$:

We can calculate $\Pr(Q(x)=m-1)$ in a similar way compared to case (III). However, the $(m-1)$-th level should be always chosen by Algorithm 1, we rely on Equation (9) rather than Equation (14). We obtain:

$$\begin{aligned}&\Pr(Q(x)=m-1)\\ =&\sum_{k=0}^{j}\Pr((Q(x)=m-1)\cap E_k\cap F_{m-1})\\ =&(1-q)^{m-j-2}\Big((1-q)^j\frac{x-B(0)}{B(m-1)-B(0)}+\sum_{k=1}^{j}q(1-q)^{j-k}\frac{x-B(k)}{B(m-1)-B(k)}\Big)\end{aligned} \tag{20}$$

Therefore, we finally get Equation (3) of Lemma E.1 from combining cases (I), (II), (III), and (IV).

E.2. Proof of Theorem 2.1

We use Lemma B.4 to find an upper bound on $D_\alpha(P_{Q(x)}||P_{Q(x')})$.

$$\begin{aligned}D_\alpha(P_{Q(x)}||P_{Q(x')})&\le D_\infty(P_{Q(x)}||P_{Q(x')})\\ &=\sup_{i\in\{0,1,\cdots,m-1\}}\log\Big(\frac{\Pr(Q(x)=i)}{\Pr(Q(x')=i)}\Big)\\ &=\max\Big(\sup_{i\in\{0,m-1\}}\log\Big(\frac{\Pr(Q(x)=i)}{\Pr(Q(x')=i)}\Big),\sup_{i\in\{1,\cdots,m-2\}}\log\Big(\frac{\Pr(Q(x)=i)}{\Pr(Q(x')=i)}\Big)\\ &\le\max\Big(\log\Big(\frac{1}{\min_{i\in\{0,m-1\}}\Pr(Q(x')=i)}\Big),\log\Big(\frac{q}{\min_{i\in\{1,\cdots,m-2\}}\Pr(Q(x')=i)}\Big)\end{aligned}$$

The second inequality comes from $\Pr(Q(x)=i)\le 1$ for any i and $\Pr(Q(x)=i)\le q$ for $i\in\{1,2,\cdots,m-2\}$. $\Pr(Q(x)=i)$ is less than or equal to q for $i\in\{1,2,\cdots,m-2\}$ because

$$\begin{aligned}\Pr(Q(x)=i)&=\Pr(Q(x)=i|i:\text{chosen})\Pr(i:\text{chosen})\\ &=\Pr(Q(x)=i|i:\text{chosen})\times q\\ &\le q(\Pr(Q(x)=i|i:\text{chosen})+\Pr(Q(x)\neq i|i:\text{chosen}))=q\end{aligned}$$

We establish a value for j that makes it so that $-c$ falls within the range of values between $B(j)$ and $B(j+1)$. Since $\min_{i\in\{0,m-1\}} \Pr(Q(x')=i) \geq \Pr(Q(-c)=m-1)$ and $\min_{i\in\{1,\cdots,m-2\}} \Pr(Q(x')=i) \geq \Pr(Q(-c)=m-2)$, we obtain

$$
\begin{aligned}
&D_\alpha(P_{Q(x)}||P_{Q(x')}) \\
&\quad\leq \max\left(\log\left(\frac{1}{\min_{i\in\{0,m-1\}}\Pr(Q(x')=i)}\right), \log\left(\frac{q}{\min_{i\in\{1,\cdots,m-2\}}\Pr(Q(x')=i)}\right)\right) \\
&\quad\leq \max\left(\log\left(\frac{1}{\Pr(Q(-c)=m-1)}\right), \log\left(\frac{q}{\Pr(Q(-c)=m-2)}\right)\right) \\
&\quad= \max\left(\log\left(\frac{1}{(1-q)^{m-2}\cdot\frac{-c-B(0)}{B(m-1)-B(0)} + \sum_{k=1}^{j} q(1-q)^{m-2-k}\cdot\frac{-c-B(k)}{B(m-1)-B(k)}}\right)\right. \\
&\left.\quad, \log\left(\frac{q}{q\left((1-q)^{m-3}\frac{-c-B(0)}{B(m-2)-B(0)} + \sum_{k=1}^{j} q(1-q)^{m-3-k}\frac{-c-B(k)}{B(m-2)-B(k)}\right)}\right)\right) \\
&\quad\leq \log\left(\frac{1}{(1-q)^{m-2}\cdot\frac{-c-B(0)}{B(m-1)-B(0)}}\right) \\
&\quad= \log\frac{1}{(1-q)^{m-2}\cdot\frac{\Delta}{2c+2\Delta}} \\
&\quad= \log\left(\frac{2(1-q)^2(c+\Delta)}{\Delta}\right) + m\log\frac{1}{1-q}
\end{aligned}
$$

To go from the third to the fourth and fifth line, we used Lemma E.1.

F. More Details about Experiments in the Main Paper

F.1. DP Experiment

Hyperparameter Choice. We fix the number of discrete levels m as 16 for both RQM and PBM to compare privacy guarantees between the two algorithms *at equal communication complexity*. We set the value c to be 1.5[4]. Figure 1 is then obtained using the following parameters: $\theta = 0.25$ for the hyperparameter of PBM[5]—we provide additional experiments comparing our results to PBM with different values of $\theta \in [0, 0.4]$ in Appx. G—and $\Delta = c$ with a corresponding fine-tuned $q = 0.42$ for RQM.

F.1.1. Nearly Worst-Case Renyi Divergence

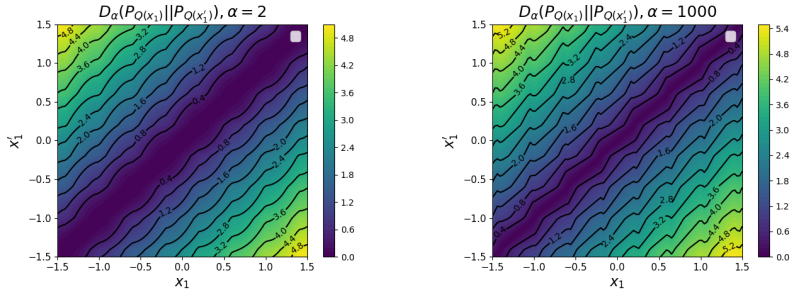

Figure 4. Both the left ($\alpha = 2$) and the right ($\alpha = 1000$) 2d plot illustrate how the Renyi divergence $D_\alpha(P_{Q(x_1)}||P_{Q(x_1')})$ changes with respect to the value of x_1 and x_1' for the single-device scenario. Here, we follow the hyperparameter choice right above.

[4]Our mechanism is in fact scale-invariant for DP guarantees, and the choice of c itself does not matter at a given constant ratio between Δ and c.

[5]See Algorithm 2 in Chen et al. (2022) for a description of the parameters.

Under a single-device case, the peak Renyi divergence $D_\alpha(P_{Q(x_1)}||P_{Q(x_1')})$ occurs predominantly around $(x_1, x_1') = (c, -c)$ and $(-c, c)$ (See Figure 4). Further, when we retain x_1' at $-c$, as per Figure 5, it's discernible that Renyi divergence $D_\alpha(P_{Q(x_1)}||P_{Q(-c)})$ increases as x_1 transitions from $-c$ to c. In instances of larger α, minor fluctuations at quantization levels are observed, followed by a swift incline in the Renyi divergence. However, considering these fluctuations as negligible, we deduce that the distance between distributions $P_{Q(x_1)}$ and $P_{Q(x_1')}$ rises almost monotonically as x_1 distances itself from x_1'. Thus, in a multiple-device situation, we can judiciously choose $x_1 = c$ and $x_1' = -c$ to represent the scenario of worst-case Renyi divergence. Looking at a basic case involving two devices, the divergence $D_\alpha(P_{Q(c)+Q(x_2)}||P_{Q(-c)+Q(x_2)})$ peaks around when $x_2 = c$, as seen in the left plot of Figure 6. This suggests that for a three-device scenario to maximize Renyi divergence, we could consider $x_1 = c, x_1' = -c, x_2 = c$. Here, $P_{Q(c)+Q(c)}$ and $P_{Q(-c)+Q(c)}$ represent the farthest separated distributions for a given x_1, x_1', x_2. Then, for the maximum Renyi divergence of a three-device scenario, we find that the divergence $D_\alpha(P_{Q(c)+Q(c)+Q(x_3)}||P_{Q(-c)+Q(c)+Q(x_3)})$ is maximized around when $x_3 = c$, as shown in the middle plot of Figure 6. Following this logic for a four-device case, the divergence $D_\alpha(P_{Q(c)+Q(c)+Q(c)+Q(x_4)}||P_{Q(-c)+Q(c)+Q(c)+Q(x_4)})$ peaks around $x_4 = c$ (depicted in the right plot of Figure 6). Thus, in the context of multiple devices, the above logic suggests that by choosing all x_is and x_i's from either $-c$ or c, we can do the privacy analysis of nearly worst-case Renyi divergence as in Section 3.1.

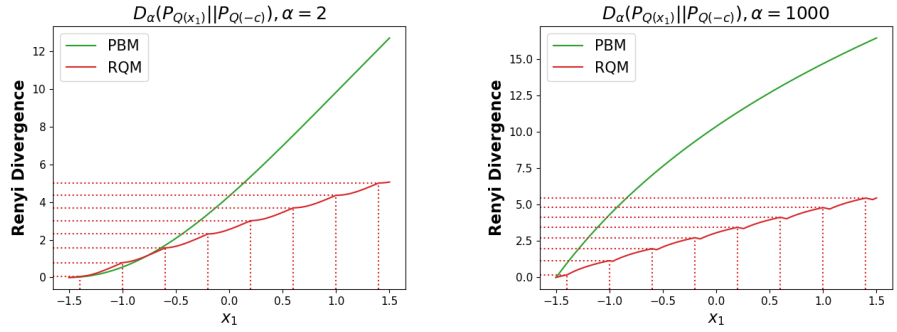

Figure 5. Both the left ($\alpha = 2$) and the right ($\alpha = 1000$) plot illustrate how the Renyi divergence $D_\alpha(P_{Q(x_1)}||P_{Q(-c)})$ changes as x_1 increases from $-c$ to c for the single-device scenario. Here, we fix $x_1' = -c$.

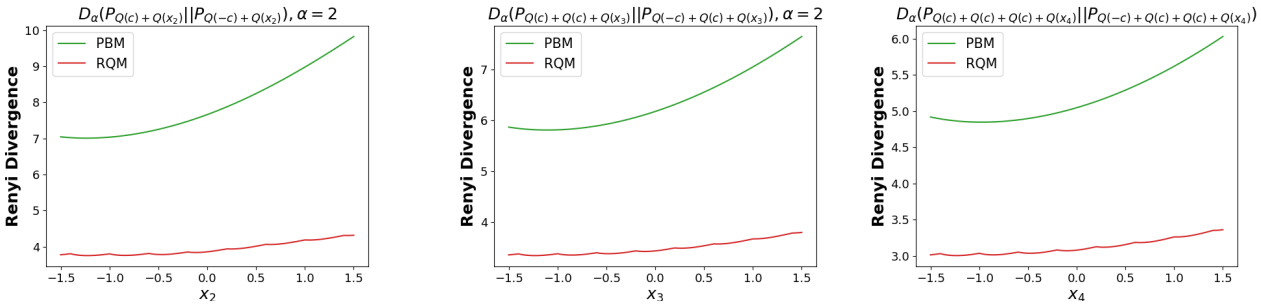

Figure 6. Each plot illustrates how the Renyi divergence changes as x_2 (left), x_3 (middle), and x_4 (right) increases from $-c$ to c for the multiple-devices scenario.

F.2. FL Experiment

Implementation environment. We adopt the implementation setup as outlined in Chen et al. (2022). To implement our algorithm, we utilize TensorFlow (Abadi et al., 2015) and the TensorFlow Federated (TFF) library. Our computation resources include 2 NVIDIA RTX A5000 GPUs. We simulate a federated learning scenario involving a total of $3,400$ local devices, with $n = 40$ local devices participating in each round. The total number of communication rounds is set at $2,000$.

Dataset & Training model. We perform image classification on the EMNIST dataset, which is comprised of 62 classes. We employ a Convolutional Neural Network (CNN) as the learning model for our training purposes.

Hyperparameter Choice. We adhere to the same hyperparameters for our FL experiments as those of Section 3.1: $m = 16$, $\theta = 0.25$ for PBM, $\Delta = c, q = 0.42$ for RQM. To highlight the flexibility of the choice of hyperparameters for RQM (Section 2.1.1), we also plot results of two more pairs $(\Delta, q) = (2c, 0.57)$ and $(\Delta, q) = (0.66c, 0.33)$. For clipping threshold c, we choose 2.9731×10^{-5}.

G. More Experimental Results

In Section 3, we opted for $\theta = 0.25$ as a hyperparameter for PBM and identified the pairs (Δ, q) - hyperparameters for RQM - that yielded a superior privacy-accuracy trade-off. Here, in this section, we extend our investigation by conducting additional experiments with a range of θ values within the [0, 0.4] interval, seeking to underscore the overarching superiority of RQM in achieving a better privacy-accuracy trade-off compared to PBM across varied θ values. For these subsequent trials, we choose $\theta = 0.15$ and $\theta = 0.35$.

G.1. More DP Experiments

In line with the experimental protocol detailed in Section 3.1, we conduct additional trials using differing θ values: $\theta = 0.15$ and $\theta = 0.35$. Consistent with the results reported in 3.1, these extended trials again demonstrate RQM's superior performance, reflected in its lower Renyi divergence, and thus enhanced privacy guarantee when compared with PBM.

Additional Experiment with $\theta = 0.15$. We choose $(\Delta, q) = (2.33c, 0.42)$ for RQM to compare with PBM with $\theta = 0.15$. The results of this experiment can be found in Figure 7.

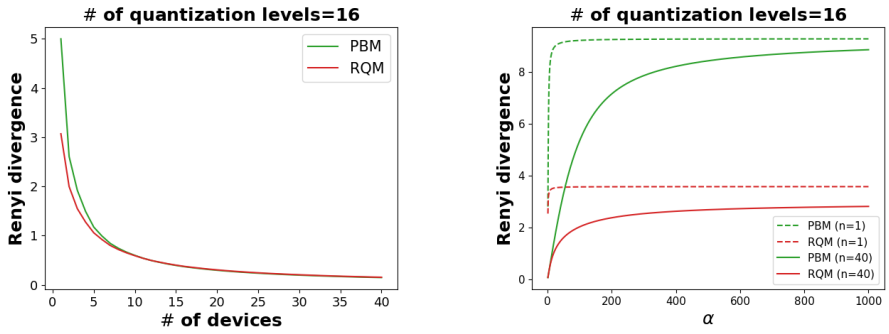

Figure 7. The results of an additional experiment about Numerical Renyi privacy guarantees with $\theta = 0.15$.

Additional Experiment with $\theta = 0.35$. We choose $(\Delta, q) = (0.429c, 0.49)$ for RQM to compare with PBM with $\theta = 0.35$. The results of this experiment can be found in Figure 8.

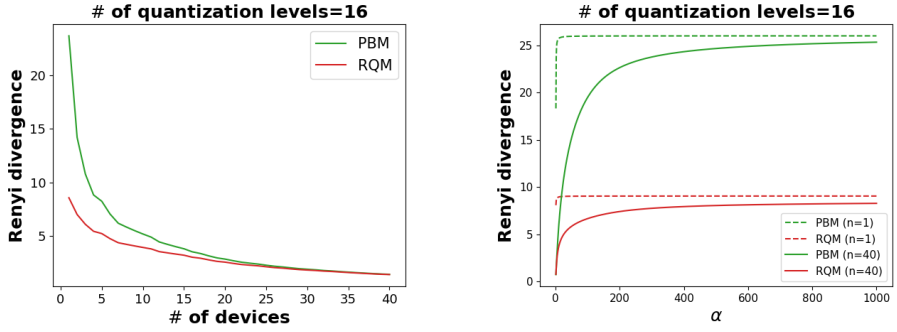

Figure 8. The results of an additional experiment about Numerical Renyi privacy guarantees with $\theta = 0.35$.

G.2. More FL Experiments

We follow the same FL experimental setup as in Section 3.2. For hyperparameter choice, we adhere to the same hyperparameters for the additional FL experiments as those of Appx. G.1. We also plot results of two more pairs (Δ, q) for each $\theta = 0.15$ and $\theta = 0.35$. Echoing the findings of Section 3.2, our additional analysis shows that, for each value of θ, all three instantiations of RQM consistently outperform PBM in terms of both performance and Renyi DP guarantees.

Additional Experiment with $\theta = 0.15$. Other than $(\Delta, q) = (2.33c, 0.42)$, we also plot results of two more pairs $(\Delta, q) = (4c, 0.5)$ and $(\Delta, q) = (c, 0.23)$. All three RQMs achieve similar accuracy, which is higher than one achieved by PBM with $\theta = 0.15$. The results of this experiment can be found in Figure 9.

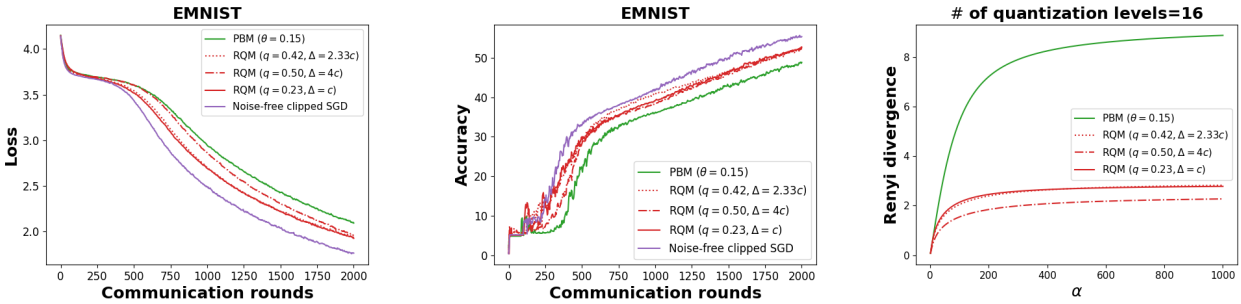

Figure 9. Comparing RQM with PBM and noise-free clipped SGD on EMNIST (Additional FL experiment with $\theta = 0.15$).

Additional Experiment with $\theta = 0.35$. The total number of communication rounds for this additional experiment is 1700. For the hyperparameters of RQM, other than $(\Delta, q) = (0.429c, 0.49)$, we also plot results of two more pairs $(\Delta, q) = (c, 0.65)$ and $(\Delta, q) = (0.25c, 0.37)$. All three RQMs achieve higher accuracy than one achieved by PBM with $\theta = 0.35$. Among three RQMs, $(\Delta, q) = (c, 0.65)$ achieves the best accuracy. The results of this experiment can be found in Figure 10.

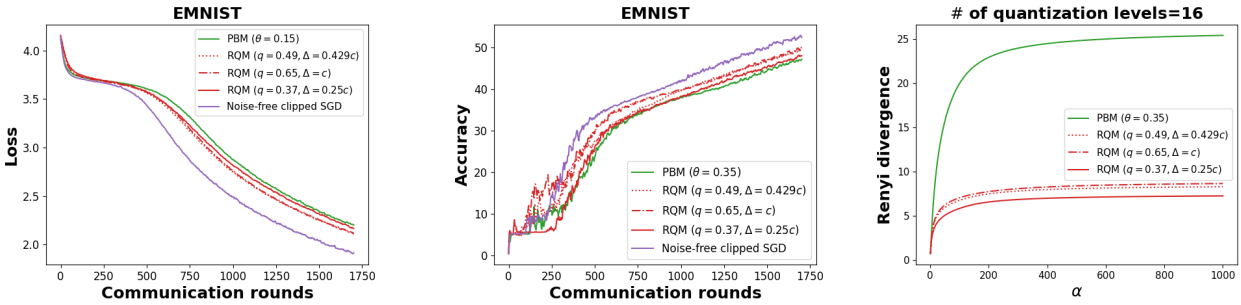

Figure 10. Comparing RQM with PBM and noise-free clipped SGD on EMNIST (Additional FL experiment with $\theta = 0.35$).

H. Discussion

In conclusion, this paper introduces a novel algorithm, the Randomized Quantization Mechanism (RQM). The RQM achieves privacy through a two-tiered process of randomization, which includes (1) the random subsampling of viable quantization levels, and (2) the application of a randomized rounding process with these subsampled discrete levels. We have theoretically demonstrated the Renyi differential privacy guarantees of RQM for a single end-user device and provided empirical evidence of its superior performance in the *privacy-accuracy trade-off* compared to the state-of-the-art Poisson Binomial Mechanism (PBM). In the future, it would be worthwhile to further examine the Renyi DP guarantees of RQM for multiple-device scenarios and the case of multi-dimensional RQM from a theoretical standpoint. Furthermore, increasing the flexibility of RQM hyperparameters by assigning unique probability values q_i to each ith discrete level presents an intriguing avenue for further enhancing the privacy-accuracy trade-off.