

---

# Fast Optimal Locally Private Mean Estimation via Random Projections

---

**Hilal Asi**  
Apple Inc.  
hilal.asi94@gmail.com

**Vitaly Feldman**  
Apple Inc.  
vitaly.edu@gmail.com

**Jelani Nelson**  
UC Berkeley  
minilek@berkeley.edu

**Huy L. Nguyen**  
Northeastern University  
hu.nguyen@northeastern.edu

**Kunal Talwar**  
Apple Inc.  
kunal@kunaltalwar.org

## Abstract

We study the problem of locally private mean estimation of high-dimensional vectors in the Euclidean ball. Existing algorithms for this problem either incur sub-optimal error or have high communication and/or run-time complexity. We propose a new algorithmic framework, ProjUnit, for private mean estimation that yields algorithms that are computationally efficient, have low communication complexity, and incur optimal error up to a  $1 + o(1)$ -factor. Our framework is deceptively simple: each randomizer projects its input to a random low-dimensional subspace, normalizes the result, and then runs an optimal algorithm such as PrivUnitG in the lower-dimensional space. In addition, we show that, by appropriately correlating the random projection matrices across devices, we can achieve fast server run-time. We mathematically analyze the error of the algorithm in terms of properties of the random projections, and study two instantiations. Lastly, our experiments for private mean estimation and private federated learning demonstrate that our algorithms empirically obtain nearly the same utility as optimal ones while having significantly lower communication and computational cost.

## 1 Introduction

Distributed estimation of the mean, or equivalently the sum, of vectors  $v_1, \dots, v_n \in \mathbb{R}^d$  is a fundamental problem in distributed optimization and federated learning. For example, in the latter, each of  $n$  devices may compute some update to parameters of a machine learning model based on its local data, at which point a central server wishes to apply all updates to the model, i.e. add  $\sum_{i=1}^n v_i$  to the vector of parameters. The typical desire to keep local data private necessitates methods for computing this sum while preserving privacy of the local data on each individual device, so that the central server essentially only learns the noisy sum and (almost) nothing about each individual summand  $v_i$  [20, 11].

The gold standard for measuring privacy preservation is via the language of *differential privacy* [21]. In this work, we study this problem in the setting of *local differential privacy* (LDP). We consider (one-round) protocols for which there exists some randomized algorithm  $\mathcal{R} : \mathbb{R}^d \rightarrow \mathcal{M}$  (called the *local randomizer*), such that each device  $i$  sends  $\mathcal{R}(v_i)$  to the server. We say the protocol is  $\varepsilon$ -*differentially private* if for any  $v, v' \in \mathbb{R}^d$  and any event  $S \subseteq \mathcal{M}$ ,  $\Pr(\mathcal{R}(v) \in S) \leq e^\varepsilon \Pr(\mathcal{R}(v') \in S)$ . If  $\varepsilon = 0$  then the distribution of  $\mathcal{R}(v)$  is independent of  $v$ , and hence the output of  $\mathcal{R}(\cdot)$  reveals nothing about the local data (perfect privacy); meanwhile if  $\varepsilon = \infty$  then the distributions of  $\mathcal{R}(v)$  and  $\mathcal{R}(v')$  can be arbitrarily far, so that in fact one may simply set  $\mathcal{R}(x) = x$  and reveal local data in the clear (total lack of privacy). Thus,  $\varepsilon \geq 0$  is typically called the *privacy loss* of a protocol.

There has been much previous work on private protocols for estimating the mean  $\mu := \frac{1}{n} \sum_{i=1}^n v_i$  in the LDP setting. Henceforth we assume each  $v_i$  lives on the unit Euclidean sphere <sup>1</sup>  $\mathbb{S}_{d-1} \subset \mathbb{R}^d$ . Let  $\hat{\mu}$  be the estimate computed by the central server based on the randomized messages  $\mathcal{R}(v_1), \dots, \mathcal{R}(v_n)$  it receives. Duchi and Rogers [18] showed that the asymptotically optimal expected mean squared error  $\mathbb{E} \|\mu - \hat{\mu}\|_2^2$  achievable by any one-round protocol must be at least  $\Omega(\frac{d}{n \min(\varepsilon, \varepsilon^2)})$ , which is in fact achieved by several protocols [19, 9, 14, 24]. These protocols however achieved empirically different errors, with some having noticeably better constant factors than others.

Recent work of Asi et al. [7] sought to understand the optimal error achievable by any protocol. Let  $\mathcal{R}$  be any local randomizer satisfying  $\varepsilon$ -DP, and  $\mathcal{A}$  be an aggregation algorithm for the central server such that it computes its mean estimate as  $\hat{\mu} := \mathcal{A}(\mathcal{R}(v_1), \dots, \mathcal{R}(v_n))$ . Furthermore, suppose that the protocol is *unbiased*, so that  $\mathbb{E} \hat{\mu} = \mu$  for any inputs  $v_1, \dots, v_n$ . Lastly, let  $\mathcal{A}_{\text{PrivUnit}_\varepsilon}, \mathcal{R}_{\text{PrivUnit}_\varepsilon}$  denote the PrivUnit $_\varepsilon$  protocol of [9] (parameterized to satisfy  $\varepsilon$ -DP <sup>2</sup>). Let  $\text{Err}_{n,d}(\mathcal{A}, \mathcal{R})$  denote

$$\sup_{v_1, \dots, v_n \in \mathbb{S}^{d-1}} \|\mathcal{A}(\mathcal{R}(v_1), \dots, \mathcal{R}(v_n)) - \mu\|_2^2.$$

Asi et al. [7] proved the remarkable theorem that for any  $n, d$ :

$$\text{Err}_{n,d}(\mathcal{A}, \mathcal{R}) \geq \text{Err}_{n,d}(\mathcal{A}_{\text{PrivUnit}_\varepsilon}, \mathcal{R}_{\text{PrivUnit}_\varepsilon}).$$

Thus, PrivUnit is not only asymptotically optimal, but in fact actually optimal in a very strong sense (at least, amongst unbiased protocols).

While this work thus characterizes the optimal error achievable for  $\varepsilon$ -LDP mean estimation, there are other desiderata that are important in practice. The most important amongst them are the *device runtime* (the time to compute  $\mathcal{R}(v)$ ), the *server runtime* (the time to compute  $\mathcal{A}$  on  $(\mathcal{R}(v_1), \dots, \mathcal{R}(v_n))$ ), and the *communication cost* ( $\lceil \log_2 |\mathcal{M}| \rceil$  bits for each device to send its  $\mathcal{R}(v)$  to the server). The known error-optimal algorithms (PrivUnit [9] and PrivUnitG [7]) either require communicating  $d$  floats or have a slower device runtime of  $\Omega(e^\varepsilon d)$ . As mean estimation is often used as a subroutine in high-dimensional learning settings, this communication cost can be prohibitive and this has led to a large body of work on reducing communication cost [2, 27, 14, 26, 24, 12]. Server runtimes of these optimal algorithms are also slow, scaling as  $nd$ , whereas one could hope for nearly linear time  $\tilde{O}(n + d)$  (see Table 1).

Chen et al. [14] recently studied this tradeoff and proposed an algorithm called SQKR, which has an optimal communication cost of  $\varepsilon$  bits and device runtime only  $O(d \log^2 d)$ . However, this algorithm incurs error that is suboptimal by a constant factor, which can be detrimental in practice. Indeed our experiments in Section 4 demonstrate the significance of such constants as the utility of these algorithms does not match that of optimal algorithms even empirically, resulting for example in 10% degradation in accuracy for private federated learning over MNIST with  $\varepsilon = 10$ .

Feldman and Talwar [24] give a general approach to reducing communication via rejection sampling. When applied to PrivUnitG, it yields a natural algorithm that we call Compressed PrivUnitG. While it yields optimal error and near-optimal communication, it requires device run time that is  $O(e^\varepsilon d)$ . These algorithms are often used for large  $d$  (e.g. in the range  $10^5 - 10^7$ ) corresponding to large model sizes. The values of  $\varepsilon$  are often in the range 4-12 or more, which may be justifiable due to privacy being improved by aggregation or shuffling [10, 15, 22, 25]. For this range of values, the  $\Theta(e^\varepsilon d)$  device runtime is prohibitively large and natural approaches to reduce this in Feldman and Talwar [24] lead to increased error. To summarize, in the high-dimensional setting, communication-efficient local randomizers are forced to choose between high device runtime or suboptimal error (see Table 1).

Another related line of work is non-private communication efficient distributed mean estimation where numerous papers have recently studied the problem due to its importance in federated learning [37, 4, 29, 2, 26, 23, 32, 38, 39]. Similarly to our paper, multiple works in this line of work have used random rotations to design efficient algorithms [37–39]. However, the purpose of these works is to

<sup>1</sup>One often considers the problem for the vectors being of norm at most 1, rather than exactly 1. It is easy to show that vectors  $v$  in the unit ball in  $\mathbb{R}^d$  can be mapped to  $\mathbb{S}_d \subseteq \mathbb{R}^{d+1}$ , simply as  $(v, 1 - \|v\|_2)$ . Thus up to changing  $d$  to  $d + 1$ , the two problems are the same. Since we are interested in the case of large  $d$ , we choose the version that is easier to work with.

<sup>2</sup>There are multiple ways to set parameters of PrivUnit to achieve  $\varepsilon$ -DP; we assume the setting described by Asi et al. [7], which optimizes the parameters to minimize the expected mean squared error.

develop better quantization schemes for real-valued vectors to reduce communication to  $(1 + o(1)) \cdot d$  bits. This is different from our goal, which is to send  $k \ll d$  parameters while still obtaining the statistically optimal bounds up to a  $1 + o(1)$  factor. Moreover, in order to preserve privacy, our algorithms require new techniques to handle the norms of the projected vectors, and post-process them using different normalization schemes, in order to guarantee that the projection error after post-processing is negligible.

## 1.1 Contributions

Our main contribution is a new framework, ProjUnit, for private mean estimation which results in near-optimal, efficient, and low-communication algorithms. Our algorithm obtains the same optimal utility as PrivUnit and PrivUnitG but with a significantly lower polylogarithmic communication complexity, and device runtime that is  $O(d \log d)$  and server runtime  $\tilde{O}(n + d)$ . We also implement our algorithms and show that both the computational cost and communication cost are small empirically as well. Figure 1 plots the error as a function of  $\epsilon$  for several algorithms and demonstrates the superiority of our algorithms compared to existing low-communication algorithms (see more details in Section 4). Moreover, we show that the optimal error bound indeed translates to fast convergence in our private federated learning simulation.

At a high level, each local randomizer in our algorithm first projects the vector to a randomly chosen lower-dimensional subspace, and then runs an optimal local randomizer in this lower-dimensional space. At first glance, this is reminiscent of the use of random projections in the Johnson-Lindenstrauss (JL) transform or the use of various embeddings in prior work (such as [14]). However, unlike the JL transform and embeddings in prior work, in our application, each point uses its own projection matrix. The JL transform is designed to preserve *distances* between points, not the points themselves. In our application, a random projection is used to obtain a low-dimensional unbiased estimator for a point; however the variance of this estimator is quite large (of the order of  $d/k$ ). Our main observation is that while large, this variance is small compared to the variance added due to privacy when  $k$  is chosen appropriately. This fact allows us to use the projections as a pre-processing step. With some care needed to control the norm of the projected vector which is no longer fixed, we then run a local randomizer in the lower dimensional space. We analyze the expected squared error of the whole process and show that as long as the random projection ensemble satisfies certain specific properties, the expected squared error of our algorithm is within  $(1 + o(1))$  factors of the optimal; the  $o(1)$  term here falls with the projection dimension  $k$ . The required properties are easily shown to be satisfied by random orthogonal projections. We further show that more structured projection ensembles, that allow for faster projection algorithms, also satisfy the required properties, and this yields even faster device runtime.

Although these structured projections result in fast device runtime, they still incur an expensive computational cost for the server which needs to apply the inverse transformation for each client, resulting in runtime  $O(nd \log d)$ . Specifically, each device sends a privatized version  $\hat{v}_i$  of  $W_i v_i$ , and the server must then compute  $\sum_i W^\top \hat{v}_i$ . To address this, we use correlated transformations in order to reduce server runtime while maintaining optimal accuracy up to  $1 + o(1)$  factors. In our correlated ProjUnit protocol, the server pre-defines a random transformation  $W$ , which all devices then use to define  $W_i = S_i W$  where  $S_i$  is a sampling matrix. The advantage is then that  $\sum_i W_i^\top \hat{v}_i$  is replaced with  $\sum_i W^\top \hat{v}_i = W^\top (\sum_i S_i \hat{v}_i)$ , which can be computed more quickly as it requires only a single matrix-vector multiplication. The main challenge with correlated transformations is that the correlated client transformations result in increased variance. However, we show that the independence in choosing the sampling matrices  $S_i$  is sufficient to obtain optimal error.

Finally, we note without correlating projections each client using its own projection would imply that each projection needs to be communicated to the server. Doing this naively would require communicating  $kd$  real values completely defeating the benefits of our protocol. However the projection matrix does not depend on the input and therefore can be communicated cheaply using a seed to an appropriate pseudorandom generator.

## 2 A random projection framework for low-communication private algorithms

In this section we propose a new algorithm, namely ProjUnit, which has low communication complexity and obtains near-optimal error (namely, up to a  $1 + o(1)$  factor of optimum). The starting point

	Utility	Run-time (client)	Run-time (server)	Communication
Repeated PrivHS [19, 24]	$O(\text{OPT})$	$\varepsilon d$	$n \lceil \varepsilon \rceil$	$\varepsilon \cdot \text{poly}(\log d)$
PrivUnit [9]	OPT	$d$	$nd$	$d$
SQKR [14]	$O(\text{OPT})$	$d \log^2 d$	$n\varepsilon \log d + d \log^2 d$	$\varepsilon \log d$
CompPrivUnit [24]	$(1 + o(1)) \cdot \text{OPT}$	$e^\varepsilon d$	$nd \log d$	$\varepsilon \cdot \text{poly}(\log d)$
FastProjUnit (Section 2)	$(1 + o(1)) \cdot \text{OPT}$	$d \log d$	$nd \log d$	$\varepsilon \log^2 d$
FastProjUnit-corr (Section 3)	$(1 + o(1)) \cdot \text{OPT}$	$d \log d$	$n \log^3 d + n\varepsilon \log d + d \log d$	$\varepsilon \log^2 d$

Table 1: Comparison of Error-Run-time-Communication trade-offs for different algorithms for private mean estimation. The last two rows use our algorithms from Section 2 and Section 3 with a communication budget  $k \approx \varepsilon \log d$ . PrivUnitG is omitted as it has the same guarantees as PrivUnit except utility where it has  $(1 + o(1)) \cdot \text{OPT}$ . We omit constant factors from the run-time and communication complexities.

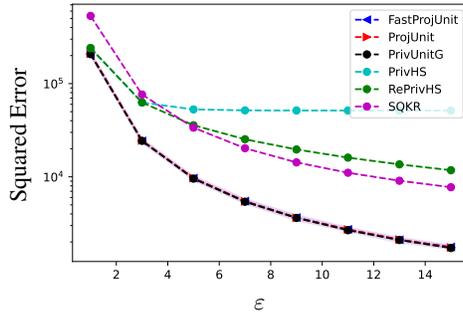


Figure 1: Squared error of different algorithms as a function of  $\varepsilon$  for  $d = 32768$  averaged over 50 runs with 90% confidence intervals. The lines for the top three algorithms are almost identical.

of our algorithms is a randomized projection map in  $\mathbb{R}^{k \times d}$  which we use to project the input vectors to a lower-dimensional space. The algorithm then normalizes the vector as a necessary pre-processing step. Finally, the local randomizer applies PrivUnitG [7] (see Algorithm 8 in Appendix) over the normalized projected vector and sends the response to the server. The server then applies the inverse transformation and aggregates the responses in order to estimate the mean. We present the full details of the client and server algorithms in Algorithm 1 and Algorithm 2.

To analyze this algorithm, we first present our general framework for an arbitrary distribution over projections. In the next sections we utilize different instances of the framework using different random projections such as random rotations and more structured transforms.

---

#### Algorithm 1 ProjUnit (client)

---

**Require:** Input vector  $v \in \mathbb{R}^d$ , Distribution over projections  $\mathcal{W}$ .

- 1: Randomly sample transform  $W \in \mathbb{R}^{k \times d}$  from  $\mathcal{W}$
  - 2: Project the input vector  $v_p = Wv$
  - 3: Normalize  $u = \frac{v_p}{\|v_p\|_2}$
  - 4: Let  $\hat{u} = \text{PrivUnitG}(u)$  (as in Algorithm 8)
  - 5: Send  $\hat{u}$  and (encoding of)  $W$  to server
- 

The following theorem states the privacy and utility guarantees of ProjUnit for a general distribution over transformation  $\mathcal{W}$  that satisfies certain properties. For ease of notation, let  $\mathcal{R}_{\text{PU}}$  denote the ProjUnit local randomizer of the client (Algorithm 1), and  $\mathcal{A}_{\text{PU}}$  denote the server aggregation of

---

**Algorithm 2** ProjUnit (server)

---

- 1: Receive  $\hat{u}_1, \dots, \hat{u}_1$  from clients with (encodings of) transforms  $W_1, \dots, W_n$
  - 2: Return the estimate  $\hat{\mu} = \frac{1}{n} \sum_{i=1}^n W_i^\top \hat{u}_i$
- 

ProjUnit (Algorithm 2). To simplify notation, we let

$$\text{Err}_{n,d}(\text{PrivUnitG}) = \text{Err}_{n,d}(\mathcal{A}_{\text{PrivUnitG}_\varepsilon}, \mathcal{R}_{\text{PrivUnitG}_\varepsilon}) = c_{d,\varepsilon} \frac{d}{n\varepsilon}$$

denote the error of the PrivUnitG  $\varepsilon$ -DP protocol where  $\mathcal{A}_{\text{PrivUnitG}_\varepsilon}, \mathcal{R}_{\text{PrivUnitG}_\varepsilon}$  denote the PrivUnitG protocol with optimized parameters (see Algorithm 8) and  $c_{d,\varepsilon} = O(1)$  is a constant [7]. We defer the proof to Appendix B.1.

**Theorem 1.** *Let  $k \leq d$  and assume that the transformations  $W_i \in \mathbb{R}^{k \times d}$  are independently chosen from a distribution  $\mathcal{W}$  that satisfies:*

1. *Bounded operator norm:*  $\mathbb{E} \left[ \|W_i^\top\|^2 \right] \leq d/k + \beta_{\mathcal{W}}$ .
2. *Bounded bias:*  $\left\| \mathbb{E} \left[ \frac{W_i^\top W_i v}{\|W_i v\|_2} \right] - v \right\|_2 \leq \sqrt{\alpha_{\mathcal{W}}}$  for all unit vectors  $v \in \mathbb{R}^d$ .

*Then for all unit vectors  $v_1, \dots, v_n \in \mathbb{R}^d$ , setting  $\hat{\mu} = \mathcal{A}_{\text{PU}}(\mathcal{R}_{\text{PU}}(v_1), \dots, \mathcal{R}_{\text{PU}}(v_n))$ , the local randomizers  $\mathcal{R}_{\text{PU}}$  are  $\varepsilon$ -DP and*

$$\mathbb{E} \left[ \left\| \hat{\mu} - \frac{1}{n} \sum_{i=1}^n v_i \right\|_2^2 \right] \leq \text{Err}_{n,d}(\text{PrivUnitG}) \cdot \left( 1 + \frac{k\beta_{\mathcal{W}}}{d} + O\left(\frac{\varepsilon + \log k}{k}\right) \right) + \alpha_{\mathcal{W}}.$$

## 2.1 ProjUnit using Random Rotations

Building on the randomized projection framework of the previous section, in this section we instantiate it with a random rotation matrix. In particular, we sample  $W \in \mathbb{R}^{k \times d}$  with the structure

$$W_H = \sqrt{\frac{d}{k}} S U, \tag{1}$$

where  $U \in \mathbb{R}^{d \times d}$  is a random rotation matrix such that  $U^\top U = I$ , and  $S \in \mathbb{R}^{k \times d}$  is a sampling matrix where each row has a single 1 in a uniformly random location (without repetitions). The following theorem states our guarantees for this distribution.

**Theorem 2.** *Let  $k \leq d$  and  $W \in \mathbb{R}^{k \times d}$  be a random rotation matrix sampled as described in (1). Then for all unit vectors  $v_1, \dots, v_n \in \mathbb{R}^d$ , setting  $\hat{\mu} = \mathcal{A}_{\text{PU}}(\mathcal{R}_{\text{PU}}(v_1), \dots, \mathcal{R}_{\text{PU}}(v_n))$ , the local randomizers  $\mathcal{R}_{\text{PU}}$  are  $\varepsilon$ -DP and*

$$\mathbb{E} \left[ \left\| \hat{\mu} - \frac{1}{n} \sum_{i=1}^n v_i \right\|_2^2 \right] \leq \text{Err}_{n,d}(\text{PrivUnitG}) \cdot \left( 1 + O\left(\frac{\varepsilon + \log k}{k}\right) \right) + O\left(\frac{1}{k^2}\right).$$

The proof follows directly from Theorem 1 and the following proposition which proves certain properties of random rotations. We defer the proof to Appendix B.2.

**Proposition 1.** *Let  $W \in \mathbb{R}^{k \times d}$  be a random rotation matrix sampled as described in (1). Then*

1. *Bounded operator norm:*  $\|W^\top\| \leq \sqrt{\frac{d}{k}}$ .
2. *Bounded bias:* for every unit vector  $v \in \mathbb{R}^d$ :  $\left\| \mathbb{E} \frac{W^\top W v}{\|W v\|} - v \right\| = O(1/k)$ .

We also have similar analysis for Gaussian transforms with an additional  $O(\sqrt{k/d})$  factor in the first term. We include the analysis in Appendix D.

## 2.2 Fast ProjUnit using the SRHT

While the random rotation based randomizer enjoys near-optimal error and low communication complexity, its runtime complexity is somewhat unsatisfactory as it requires calculating  $Wv$  for  $W \in \mathbb{R}^{k \times d}$ , taking time  $O(kd)$ . In this section, we propose a ProjUnit algorithm using the Subsampled Randomized Hadamard transform (SRHT), which is closely related to the fast JL transform [3]. We show that this algorithm has the same optimality and low-communication guarantees as the random rotations version, and additionally has an efficient implementation with  $O(d \log d)$  client runtime.

The SRHT ensemble contains matrices  $W_H \in \mathbb{R}^{k \times d}$  with the following structure:

$$W_H = \sqrt{\frac{d}{k}} SHD, \quad (2)$$

where  $S \in \mathbb{R}^{k \times d}$  is a sampling matrix where each row has a single 1 in a uniformly random location sampled without replacement,  $H \in \mathbb{R}^{d \times d}$  is the Hadamard matrix, and  $D \in \mathbb{R}^{d \times d}$  is a diagonal matrix where  $D_{ii}$  are independent samples from the Rademacher distribution, that is,  $D_{ii} \sim \text{Unif}\{-1, +1\}$ . The main advantage of the SRHT is that there exist efficient algorithms for matrix-vector multiplication with  $H$ .

The following theorem presents our main guarantees for the SRHT-based ProjUnit algorithm.

**Theorem 3.** *Let  $k \leq d$  and  $W$  be sampled from the SRHT ensemble as described in (2). Then for all unit vectors  $v_1, \dots, v_n \in \mathbb{R}^d$ , setting  $\hat{\mu} = \mathcal{A}_{\text{PU}}(\mathcal{R}_{\text{PU}}(v_1), \dots, \mathcal{R}_{\text{PU}}(v_n))$ , the local randomizers  $\mathcal{R}_{\text{PU}}$  are  $\varepsilon$ -DP and*

$$\mathbb{E} \left[ \left\| \hat{\mu} - \frac{1}{n} \sum_{i=1}^n v_i \right\|_2^2 \right] \leq \text{Err}_{n,d}(\text{PrivUnitG}) \cdot \left( 1 + O\left(\frac{\varepsilon + \log k}{k}\right) \right) + O\left(\frac{\log^2 d}{k}\right).$$

**Remark 1.** *The communication complexity of SRHT-based ProjUnit can be reduced to  $O(k \log d + \log^2 d)$ . To see this, note that  $\hat{u}$  is a  $k$ -dimensional vector. Moreover, the matrix  $W = \sqrt{d/k} \cdot SHD$  can be sent in  $O(k \log d)$  as follows:  $S$  has  $k$  rows, each with a single entry that contains 1, hence we can send the indices for each row using  $\log d$  bits for each row. Moreover,  $H$  is the Hadamard transform and need not be sent. Finally,  $D$  is a diagonal matrix with entries  $D_{ii} \sim \text{Unif}\{-1, +1\}$ . By standard techniques [36, 5], we only need the entries of  $D$  to be  $O(\log(d))$ -wise independent for Proposition 2 to hold. Thus  $O(\log^2 d)$  bits suffice to communicate a sampled  $D$ .*

The proof of the theorem builds directly on the following two properties of the SHRT.

**Proposition 2.** *Let  $W$  be sampled from the SRHT ensemble as described in (2). Then we have*

1. *Bounded operator norm:*  $\mathbb{E} [\|W^\top\|] = \mathbb{E} [\|W\|] \leq \sqrt{d/k}$ .
2. *Bounded bias:* for every unit vector  $v \in \mathbb{R}^d$ :  $\left\| \mathbb{E} \frac{W^\top W v}{\|W v\|} - v \right\| = O(\log(d)/\sqrt{k})$ .

Theorem 3 now follows directly from the bounds of Theorem 1 using  $\alpha_{\mathcal{W}} = O(\log^2(d)/k)$ . We defer the proof to Appendix B.3.

**Remark 2.** *While our randomizers in this section pay an additive term that does not decrease with  $n$  (e.g.  $\log^2(d)/k$  for SRHT), this term is negligible in most settings of interest. Indeed, using Theorem 3 and the fact that  $\text{Err}_{n,d}(\text{PrivUnitG}) = c_{d,\varepsilon} d/n\varepsilon$ , we get that the final error of our SRHT algorithm is roughly  $c_{d,\varepsilon} d/n\varepsilon(1 + o(1)) + O(\log^2(d)/k)$ . This implies that in the high-dimensional setting the bias term is negligible.*

However, to cover the whole spectrum of parameters, we develop a nearly unbiased versions of these algorithms in Appendix A. In particular, we show in Theorem 6 that our unbiased version has error

$$\text{Err}_{n,d}(\text{PrivUnitG}) \cdot \left( 1 + O\left(\frac{\varepsilon + \log k}{k} + \sqrt{\frac{\log(nd/k)}{k}}\right) \right).$$

## 3 Efficient Server Runtime via Correlated Sampling

One downside of the algorithms in the previous section is that server runtime can be costly: indeed, as each client uses an independent transformation, the server has to apply the inverse transformation

(matrix multiplication) for each client, resulting in runtime  $O(nd \log d)$ . In this section, we propose a new protocol that significantly reduces server runtime to  $O(n \log^3 d + d \log d + nk)$  while retaining similar error guarantees. The protocol uses correlated transformations between users which allows the server to apply an inverse transformation only a small number of times. However, clients cannot use the same transformation as this will result in large bias.

The protocol works as follows: the server samples  $U \in \mathbb{R}^{d \times d}$  from the Randomized Hadamard transform:  $U = HD$  where  $H \in \mathbb{R}^{d \times d}$  is the Hadamard transform, and  $D \in \mathbb{R}^{d \times d}$  is a diagonal matrix where each diagonal entry is independently sampled from the Rademacher distribution. Then, client  $i \in [n]$ , samples a random sampling matrix  $S_i \in \mathbb{R}^{k \times d}$ , and uses  $U$  to define the transform  $W_i \in \mathbb{R}^{k \times d}$ :

$$W_i = \sqrt{\frac{d}{k}} S_i U. \quad (3)$$

We describe the full details of the client and server algorithms for correlated ProjUnit in Algorithm 3 and Algorithm 4, and denote them  $\mathcal{R}_{\text{CPU}}$  and  $\mathcal{A}_{\text{CPU}}$ , respectively. We have the following guarantee. We defer the proof to Appendix C.

---

**Algorithm 3** Correlated ProjUnit (client)

---

**Require:** Input vector  $v \in \mathbb{R}^d$ .

- 1: Randomly sample diagonal  $D$  from the Rademacher distribution based on predefined seed
  - 2: Sample  $S \in \mathbb{R}^{k \times d}$  where each row is chosen uniformly at random without replacement from standard basis vectors  $\{e_1, \dots, e_d\}$
  - 3: Project the input vector  $v_p = SHDv$  where  $H \in \mathbb{R}^{d \times d}$  is the Hadamard matrix
  - 4: Normalize  $u = \frac{v_p}{\|v_p\|_2}$
  - 5: Let  $\hat{u} = \text{PrivUnitG}(u)$  (as in Algorithm 8)
  - 6: Send  $\hat{u}$ , and (an encoding of)  $S$  to server
- 

---

**Algorithm 4** Correlated ProjUnit (server)

---

- 1: Receive  $\hat{u}_1, \dots, \hat{u}_1$  from clients with (encodings of) transforms  $S_1, \dots, S_n$
  - 2: Sample diagonal matrices  $D$  from Rademacher distribution based on predefined seed
  - 3: Let  $U = HD$  where  $H \in \mathbb{R}^{d \times d}$  is the Hadamard matrix
  - 4: Return the estimate  $\hat{\mu} = \frac{1}{n} U^\top \sum_{i=1}^n S_i^\top \hat{u}_i$
- 

**Theorem 4.** *Let  $k \leq d$ . Then for all unit vectors  $v_1, \dots, v_n \in \mathbb{R}^d$ , setting  $\hat{\mu} = \mathcal{A}_{\text{CPU}}(\mathcal{R}_{\text{CPU}}(v_1), \dots, \mathcal{R}_{\text{CPU}}(v_n))$ , the local randomizers  $\mathcal{R}_{\text{CPU}}$  are  $\varepsilon$ -DP and*

$$\mathbb{E} \left[ \left\| \hat{\mu} - \frac{1}{n} \sum_{i=1}^n v_i \right\|_2^2 \right] \leq \text{Err}_{n,d}(\text{PrivUnitG}) \left( 1 + O\left(\frac{\varepsilon + \log k}{k}\right) \right) + O\left(\frac{\log^2 d}{k}\right).$$

Moreover, server runtime is  $O(n \log(d) \log^2(nd) + d \log d + nk)$ .

## 4 Experiments

We conclude the paper with several experiments that demonstrate the performance of our proposed algorithms, comparing them to existing algorithms in the literature. We conduct our experiments in two different settings: the first is synthetic data, where we aim to test our algorithms and understand their performance for our basic task of private mean estimation, comparing them to other algorithms. In our second setting, we seek to evaluate the performance of our algorithms for *private federated learning* which requires private mean estimation as a subroutine for DP-SGD.

### 4.1 Private mean estimation

In our synthetic-data experiment, we study the basic private mean estimation problem, aspiring to investigate the following aspects:

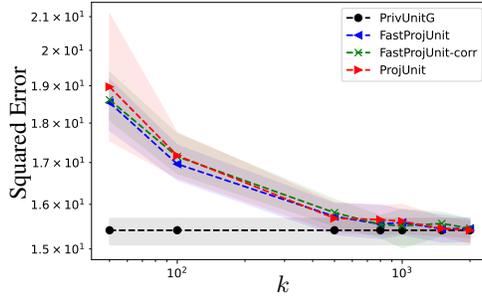


Figure 2: Performance of ProjUnit, FastProjUnit and their correlated versions with 90% confidence intervals as a function of  $k$  for  $d = 32768$ ,  $\text{NumRep} = 30$ ,  $n = 50$ , and  $\varepsilon = 10$

1. Utility of ProjUnit algorithms as a function of the communication budget
2. Utility of our low-communication algorithms compared to the optimal utility and other existing low-communication algorithms
3. Run-time complexity of our algorithms compared to existing algorithms

Our experiments<sup>3</sup> measure the error of different algorithms for estimating the mean of a dataset. To this end, we sample unit vectors  $v_1, \dots, v_n \in \mathcal{R}^d$  by normalizing samples from the normal distribution  $\mathcal{N}(\mu, 1/d)$  (where  $\mu \in \mathbb{R}^d$  is a random unit vector), and apply a certain privacy protocol  $\text{NumRep}$  times to estimate the mean  $\sum_{i=1}^n v_i/n$ , producing mean squared errors  $e_1, \dots, e_{\text{NumRep}}$ . Our final error estimate is then the mean  $\frac{1}{\text{NumRep}} \sum_{i=1}^{\text{NumRep}} e_i$ . We test the performance of several algorithms: ProjUnit (Subsection 2.1), FastProjUnit (Subsection 2.2), FastProjUnit-corr (Section 3), PrivUnitG [7], CompPrivUnitG [7, 24], PrivHS [19], RePrivHS [19, 24], SQKR [14].

In Figure 2, we plot the error for our ProjUnit algorithms as a function of the communication budget  $k$ . We consider a high-dimensional regime where  $d = 2^{15}$  with a small number of users  $n = 50$  and a bounded communication budget  $k \in [1, 2000]$ . Our plot shows that our ProjUnit algorithms obtain the same near-optimal error as PrivUnitG for  $k$  as small as 1000. Moreover, the plots show that the correlated versions of our ProjUnit algorithms obtain nearly the same error.

To translate this into concrete numbers, the transform  $W$  can be communicated using a small seed ( $\sim 128$  bits) in practice, or using  $k \log d + \log^2 d \sim 20400$  bits or less than 3kB for  $d = 10^6$ . Sending the  $k$ -dimensional vector of 32-bit floats would need an additional 4kB. Thus the total communication cost is between 4 and 8kB. This can be further reduced by using a compressed version of PrivUnitG in the projected space, which requires the client to send a 128-bit seed. In this setting, the communication cost is a total of 256 bits. Thus in the sequel, we primarily focus on the  $k = 1000$  version of our algorithms.

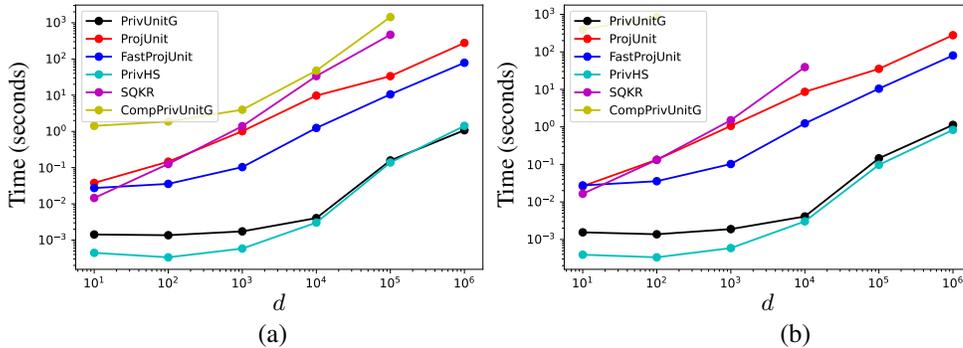


Figure 3: Run-time (in seconds) as a function of the dimension for (a)  $\varepsilon = 10$  and (b)  $\varepsilon = 16$ . The plots for some algorithms are not complete as they did not finish within the cut-off time.

<sup>3</sup>The code is also available online on <https://github.com/apple/ml-projunit>

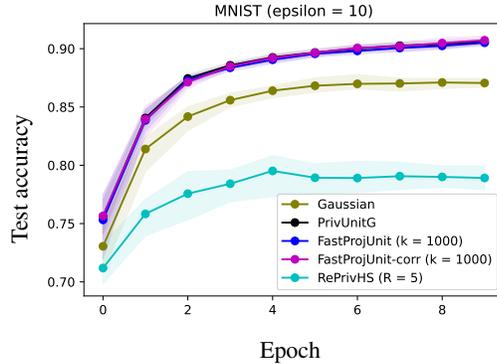


Figure 4: Test accuracy on the MNIST dataset with 90% confidence intervals as a function of epoch for  $\epsilon = 10.0$ .

Next, we compare the performance of our low-communication algorithms against existing low-communication algorithms: PrivHS and SQKR. In Figure 1, we plot the error as a function of the privacy parameter for each algorithm using the best choice of  $k$  (bound on communication) for each algorithm. In particular, we choose  $k = \epsilon$  for SQKR, num repetitions  $R = \epsilon/2$  for repeated PrivHS,  $k = 1000$  for ProjUnit and FastProjUnit. Moreover, in this experiment we set  $n = 1$  and NumRep = 50 to estimate the variance of each method. The figure shows that PrivHS and SQKR, while having low-communication complexity, suffer a significantly worse utility than (near-optimal) PrivUnitG. On the other side, both our ProjUnit algorithms obtain nearly the same error as PrivUnitG with a bounded communication budget of  $k = 1000$ .

In our third experiment in Figure 3, we plot the runtime of each algorithm as a function of the privacy parameter. Here, we use  $n = 1$ , NumRep = 10 and measure the run-time of each method for different values of the dimension  $d$  and privacy parameter  $\epsilon$ , allowing each method to run for 1 hour before interrupting. As expected from our theoretical analysis, the runtime of ProjUnit using random rotations is noticeably slower than the (high communication cost) PrivUnitG. However, our SRHT-based FastProjUnit is substantially faster and has a comparable run-time to PrivUnitG. Moreover, for large  $\epsilon$  and  $d$ , the run-time of compressed PrivUnitG becomes too costly compared to our algorithms due to the  $e^\epsilon d$  time complexity.

## 4.2 Private federated learning

Having demonstrated the effectiveness of our methods for private mean estimation, in this section we illustrate the improvements offered by our algorithms for private federated learning. Similarly to the experimental setup in [13], we consider the MNIST [31] dataset and train a convolutional network (see Table 2) using DP-SGD [1] with 10 epochs and different sub-routines for privately estimating the mean of gradients at each batch. In order to bound the sensitivity, we clip the gradients to have  $\ell_2$ -norm 1, and run DP-SGD with batch size of 600, step-size equal to 0.1, and momentum of 0.5.

Figure 4 shows our results for this experiment, where we plot the test accuracy as a function of the epoch for each method. The plots demonstrate that our ProjUnit algorithms obtain similar performance to PrivUnitG, and better performance than the Gaussian mechanism or PrivHS. For the Gaussian mechanism, we set  $\delta = 10^{-5}$  and add noise to satisfy  $(\epsilon, \delta)$ -DP using the analysis in [8]. We did not run SQKR in this experiment as it is not sufficiently computationally efficient for this experiment and has substantially worse performance in the experiments of the previous section. We also did not run the MVU mechanism [13] as their experiments show that it is worse than the Gaussian mechanism which has worse performance than our methods.

Finally, our private algorithms obtain accuracy roughly 91%, whereas the same model trained without privacy obtains around 98%. This degradation in accuracy is mostly due to the choice of the optimization algorithm (DP-SGD with clipping); indeed, even without adding any noise, DP-SGD with clipping achieves around 91% accuracy, suggesting that other private optimization algorithms with different clipping strategies (e.g. [34, 6]) may tighten this gap further. As this is not the main focus of our work, we leave this investigation to future work.

## Acknowledgments

HLN is supported by NSF CAREER grant 1750716 and NSF grant 2311649.

## References

- [1] Martin Abadi, Andy Chu, Ian Goodfellow, Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 23rd Annual ACM Conference on Computer and Communications Security (CCS)*, pages 308–318, 2016.
- [2] Naman Agarwal, Ananda Theertha Suresh, Felix Xinnan X Yu, Sanjiv Kumar, and Brendan McMahan. cpsgd: Communication-efficient and differentially-private distributed sgd. In *Proceedings of the 31st Annual Conference on Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [3] Nir Ailon and Bernard Chazelle. The fast Johnson-Lindenstrauss transform and approximate nearest neighbors. *SIAM Journal on Computing*, 39(1):302–322, 2009.
- [4] Dan Alistarh, Demjan Grubic, Jerry Li, Ryota Tomioka, and Milan Vojnovic. Qsgd: Communication-efficient sgd via gradient quantization and encoding. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30, pages 1709–1720. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/6c340f25839e6acdc73414517203f5f0-Paper.pdf>.
- [5] N. Alon and J. H. Spencer. *The Probabilistic Method*. Wiley-Interscience, second edition, 2000.
- [6] Hilal Asi, John Duchi, Alireza Fallah, Omid Javidi, and Kunal Talwar. Private adaptive gradient methods for convex optimization. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, pages 383–392, 2021.
- [7] Hilal Asi, Vitaly Feldman, and Kunal Talwar. Optimal algorithms for mean estimation under local differential privacy. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, 2022.
- [8] Borja Balle and Yu-Xiang Wang. Improving the gaussian mechanism for differential privacy: Analytical calibration and optimal denoising. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pages 394–403. PMLR, 2018.
- [9] Abhishek Bhowmick, John Duchi, Julien Freuder, Gaurav Kapoor, and Ryan Rogers. Protection against reconstruction and its applications in private federated learning. *arXiv:1812.00984 [stat.ML]*, 2018.
- [10] Andrea Bittau, Úlfar Erlingsson, Petros Maniatis, Ilya Mironov, Ananth Raghunathan, David Lie, Mitch Rudominer, Ushasree Kode, Julien Tinnes, and Bernhard Seefeld. Prochlo: Strong privacy for analytics in the crowd. In *Proceedings of the 26th Symposium on Operating Systems Principles, SOSP '17*, page 441–459. New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450350853. doi: 10.1145/3132747.3132769. URL <https://doi.org/10.1145/3132747.3132769>.
- [11] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H. Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the Annual ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 1175–1191, 2017.
- [12] Kamalika Chaudhuri, Claire Monteleoni, and Anand D. Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12:1069–1109, 2011.
- [13] Kamalika Chaudhuri, Chuan Guo, and Mike Rabbat. Privacy-aware compression for federated data analysis. In *Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence*, volume 180, pages 296–306. PMLR, 2022. URL <https://proceedings.mlr.press/v180/chaudhuri22a.html>.

- [14] Wei-Ning Chen, Peter Kairouz, and Ayfer Özgür. Breaking the communication-privacy-accuracy trilemma. In *Proceedings of the 33rd Annual Conference on Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [15] Albert Cheu, Adam Smith, Jonathan Ullman, David Zeber, and Maxim Zhilyaev. Distributed differential privacy via shuffling. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2019*, pages 375–403, Cham, 2019. Springer International Publishing. ISBN 978-3-030-17653-2.
- [16] Michael B. Cohen, Jelani Nelson, and David P. Woodruff. Optimal approximate matrix product in terms of stable rank. In *Proceedings of the 41st International Colloquium on Automata, Languages and Programming (ICALP)*, pages 11:1–11:14, 2016. Full version at <https://arxiv.org/abs/1507.02268v3>.
- [17] K.R. Davidson and Stanislaw Szarek. Local operator theory, random matrices and banach spaces. *Handbook on the Geometry of Banach spaces, Vol. 1*, pages 317–366, 01 2003.
- [18] John Duchi and Ryan Rogers. Lower bounds for locally private estimation via communication complexity. In *Proceedings of the 32nd Annual Conference on Learning Theory (COLT)*, pages 1161–1191, 2019.
- [19] John C. Duchi, Michael I. Jordan, and Martin J. Wainwright. Minimax optimal procedures for locally private estimation. *Journal of the American Statistical Association*, 113(521):182–215, 2018.
- [20] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Advances in Cryptology (EUROCRYPT 2006)*, 2006.
- [21] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the Third Theory of Cryptography Conference*, pages 265–284, 2006.
- [22] Ulfar Erlingsson, Vitaly Feldman, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Abhradeep Thakurta. Amplification by shuffling: From local to central differential privacy via anonymity. In *Proceedings of the Thirtieth ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2019.
- [23] Fartash Faghri, Iman Tabrizian, Ilia Markov, Dan Alistarh, Daniel M. Roy, and Ali Ramezani-Kebrya. Adaptive gradient quantization for data-parallel sgd. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- [24] Vitaly Feldman and Kunal Talwar. Lossless compression of efficient private local randomizers. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, pages 3208–3219. PMLR, 2021.
- [25] Vitaly Feldman, Audra McMillan, and Kunal Talwar. Hiding among the clones: A simple and nearly optimal analysis of privacy amplification by shuffling. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 954–964, 2022. doi: 10.1109/FOCS52979.2021.00096. arXiv:2012.12803 [cs.LG].
- [26] Venkata Gandikota, Daniel Kane, Raj Kumar Maity, and Arya Mazumdar. vqsgd: Vector quantized stochastic gradient descent. *arXiv preprint arXiv:1911.07971*, 2019.
- [27] Antonious M. Girgis, Deepesh Data, Suhas Diggavi, Peter Kairouz, and Ananda Theertha Suresh. Shuffled model of federated learning: Privacy, communication and accuracy trade-offs, 2020.
- [28] Parikshit Gopalan, Daniel Kane, and Raghu Meka. Pseudorandomness via the discrete fourier transform. pages 903–922, 10 2015. doi: 10.1109/FOCS.2015.60.
- [29] Jakub Konečný and Peter Richtárik. Randomized distributed mean estimation: Accuracy vs. communication. *Frontiers in Applied Mathematics and Statistics*, 4:62, 2018.

- [30] Pravesh K. Kothari and Raghu Meka. Almost optimal pseudorandom generators for spherical caps: Extended abstract. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, STOC '15, page 247–256, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450335362. doi: 10.1145/2746539.2746611. URL <https://doi.org/10.1145/2746539.2746611>.
- [31] Yann LeCun, Corinna Cortes, and CJ Burges. MNIST handwritten digit database, 1998. URL <http://yann.lecun.com/exdb/mnist>. ATT Labs [Online].
- [32] P. Mayekar and H. Tyagi. Limits on gradient compression for stochastic optimization. In *2020 IEEE International Symposium on Information Theory (ISIT)*, pages 2658–2663, 2020. doi: 10.1109/ISIT44484.2020.9174075.
- [33] Noam Nisan. Pseudorandom generators for space-bounded computation. *Combinatorica*, 12: 449–461, 1992.
- [34] Venkatadheeraj Pichapati, Ananda Theertha Suresh, Felix X. Yu, Sashank J. Reddi, and Sanjiv Kumar. AdaClip: Adaptive clipping for private SGD. *arXiv:1908.07643 [cs.LG]*, 2020.
- [35] Tamás Sarlós. Improved approximation algorithms for large matrices via random projections. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 143–152, 2006.
- [36] Jeanette P. Schmidt, Alan Siegel, and Aravind Srinivasan. Chernoff–hoeffding bounds for applications with limited independence. *SIAM Journal on Discrete Mathematics*, 8(2):223–250, 1995. doi: 10.1137/S089548019223872X.
- [37] Ananda Theertha Suresh, Felix X. Yu, Sanjiv Kumar, and H. Brendan McMahan. Distributed mean estimation with limited communication. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017.
- [38] Shay Vargaftik, Ran Ben-Basat, Amit Portnoy, Gal Mendelson, Yaniv Ben-Itzhak, and Michael Mitzenmacher. Drive: One-bit distributed mean estimation. *Proceedings of the 34th Annual Conference on Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [39] Shay Vargaftik, Ran Ben-Basat, Amit Portnoy, Gal Mendelson, Yaniv Ben Itzhak, and Michael Mitzenmacher. Eden: Communication-efficient and robust distributed mean estimation for federated learning. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, 2022.

## A Nearly unbiased ProjUnit randomizers

Our randomizers in Section 2 and Section 3 may have  $O(\log(d)/k)$  bias which can become relatively large when  $n \gg d$ . In this section, we propose a different normalization technique which allows to provide a sufficiently small bound on the bias, while still enjoying the same guarantees as the fast ProjUnit algorithm. We develop versions of this algorithm for both random rotations (Appendix A.1) and the SRHT transform (Appendix A.2).

### A.1 Unbiased variant of ProjUnit using random rotations

In this section, we describe the modification for random rotation matrices. These transformations are not as efficient as SRHT hence we only present the simple non-correlated version; in the next section we present our unbiased and correlated sampling procedure for SRHT transforms.

For rotationally symmetric distributions of matrices, we slightly modify the algorithm by scaling the output of PrivUnitG by a fixed factor  $c$  so that it leads to an unbiased estimate of  $v$  i.e.  $\mathbb{E}[W^\top \hat{u}] = v$ .

---

**Algorithm 5** Unbiased version of ProjUnit using random rotations (client)

---

**Require:** Input vector  $v \in \mathbb{R}^d$ .

- 1: Randomly sample a rotation matrix  $W \in \mathbb{R}^{k \times d}$  as described in (1)
  - 2: Project the input vector  $v_p = Wv$
  - 3: Normalize  $u = \frac{v_p}{\|v_p\|_2}$
  - 4: Let  $\hat{u} = c \cdot \text{PrivUnitG}(u)$  where  $c = \sqrt{\frac{k}{d} \frac{\Gamma((d+1)/2)\Gamma(k/2)}{\Gamma((k+1)/2)\Gamma(d/2)}}$
  - 5: Send  $\hat{u}$  and (encoding of)  $W$  to server
- 

We present the details of this modification in Algorithm 5 and state its guarantees in the following theorem. We let  $\mathcal{R}_{\text{PU}}$  denote the local randomizer described in Algorithm 5.

**Theorem 5.** *Let  $k \leq d$ . For all unit vectors  $v_1, \dots, v_n \in \mathbb{R}^d$ , setting  $\hat{\mu} = \mathcal{A}_{\text{PU}}(\mathcal{R}_{\text{PU}}(v_1), \dots, \mathcal{R}_{\text{PU}}(v_n))$ , the local randomizers  $\mathcal{R}_{\text{PU}}$  are  $\varepsilon$ -DP and*

$$\mathbb{E} \left[ \left\| \hat{\mu} - \frac{1}{n} \sum_{i=1}^n v_i \right\|_2^2 \right] \leq \text{Err}_{n,d}(\text{PrivUnitG}) \cdot \left( 1 + O\left(\frac{\varepsilon + \log k}{k}\right) \right).$$

*Proof.* The proof proceeds in the same way as Theorem 1. We break the error down into two terms:

$$\begin{aligned} \mathbb{E} \left[ \left\| \hat{\mu} - \frac{1}{n} \sum_{i=1}^n v_i \right\|_2^2 \right] &= \mathbb{E} \left[ \left\| \frac{1}{n} \sum_{i=1}^n W_i^\top \hat{u}_i - v_i \right\|_2^2 \right] \\ &= \mathbb{E} \left[ \left\| \frac{1}{n} \sum_{i=1}^n W_i^\top \hat{u}_i - cW_i^\top u_i + cW_i^\top u_i - v_i \right\|_2^2 \right] \\ &\stackrel{(i)}{=} \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \left[ \|W_i^\top \hat{u}_i - cW_i^\top u_i\|_2^2 \right] + \frac{1}{n^2} \mathbb{E} \left[ \left\| \sum_{i=1}^n cW_i^\top u_i - v_i \right\|_2^2 \right] \\ &\leq \frac{c^2}{n} \max_{i \in [n]} \mathbb{E} \left[ \|W_i^\top\|_2^2 \right] \cdot \text{Err}_{1,k}(\text{PrivUnitG}) + \frac{1}{n^2} \mathbb{E} \left[ \left\| \sum_{i=1}^n cW_i^\top u_i - v_i \right\|_2^2 \right]. \end{aligned}$$

where (i) follows from the fact that PrivUnitG is unbiased and  $\mathbb{E}[\hat{u}_i] = cu_i$ . The first term is bounded in the same way as before noting that  $c = 1 + O(1/k)$  (see Lemma F.1). To analyze the second term, we first show that  $\mathbb{E}[cW_i^\top u_i] = v_i$  using a change of variables. Let  $W'_i = W_i P_i^\top$  where  $P_i$  is the rotation matrix such that  $P_i v_i = e_1$ , the first standard basis vector. Due to the rotational symmetry of the uniform distribution over rotation matrices,  $W'_i$  is also a random rotation matrix. Note that  $W_i = W'_i P_i$ , hence

$$\begin{aligned}
\mathbb{E}_{W_i}[cW_i^\top u_i] &= \mathbb{E}_{W'_i} \left[ \frac{c}{\|W'_i P_i v_i\|_2} P_i^\top W'^\top_i W'_i P_i v_i \right] \\
&= cP_i^\top \mathbb{E} \left[ \underbrace{\frac{1}{\|W'_i e_1\|_2} W'^\top_i W'_i e_1}_z \right]
\end{aligned}$$

Notice that  $z_j = \frac{1}{\|W'_i e_1\|_2} e_j^\top W'^\top_i W'_i e_1 = \langle W'_i e_j, \frac{1}{\|W'_i e_1\|_2} W'_i e_1 \rangle$ . Because  $W'_i$  is a random rotation matrix (re-scaled by  $\sqrt{d/k}$ ), Lemma F.1 implies that  $\mathbb{E}[z_1] = \sqrt{\frac{d}{k} \frac{\Gamma((k+1)/2)\Gamma(d/2)}{\Gamma((d+1)/2)\Gamma(k/2)}} = \frac{1}{c} = 1 + O(1/k)$  and  $\mathbb{E}[z_j] = 0$  for all  $j > 1$ . Thus,  $\mathbb{E}[z] = \frac{1}{c} e_1$  and  $\mathbb{E}[cW_i^\top u_i] = P_i^\top e_1 = P_i^\top P_i v_i = v_i$ . Because  $cW_i^\top u_i$  is an unbiased estimator of  $v_i$ , we have

$$\begin{aligned}
\mathbb{E} \left[ \left\| \sum_{i=1}^n cW_i^\top u_i - v_i \right\|_2^2 \right] &= \sum_{i=1}^n \mathbb{E} \left[ \|cW_i^\top u_i - v_i\|_2^2 \right] \\
&= \sum_{i=1}^n \mathbb{E} \left[ \|cW_i^\top u_i\|_2^2 + \|v_i\|_2^2 - 2c v_i^\top W_i^\top u_i \right] \\
&= \sum_{i=1}^n \mathbb{E} \left[ \|cW_i^\top u_i\|_2^2 + \|v_i\|_2^2 - 2c \|W_i^\top v_i\|_2 \right] \\
&\leq n + \sum_{i=1}^n c^2 \mathbb{E} \left[ \|W_i^\top\|_2^2 \right] \\
&\leq O(nd/k).
\end{aligned}$$

Combining all of these together, the claim follows by noting that  $\text{Err}_{1,d}(\text{PrivUnitG})/n = \text{Err}_{n,d}(\text{PrivUnitG}) = \Theta(d/n\varepsilon)$ .  $\square$

## A.2 Nearly unbiased SRHT-based randomizers

While rescaling by a constant was sufficient to debias the random rotation based randomizer, it is not clear whether such rescaling can debias the SRHT ProjUnit randomizer as it is not rotationally symmetric. To address this, we propose a different normalization technique for the SRHT randomizer which allows to provide tighter upper bounds on the bias. We provide the details for our new client and server protocols in Algorithm 6 and Algorithm 7, respectively.

Let  $\mathcal{R}_{\text{UPU}}$  denote the unbiased ProjUnit local randomizer of the client (Algorithm 6), and  $\mathcal{A}_{\text{UPU}}$  denote the server aggregation of unbiased ProjUnit (Algorithm 7). We have the following guarantees for this procedure.

**Theorem 6.** *Let  $k \leq d$  and  $\delta = k/n^2d$ . Assume  $k \geq \max\{\varepsilon + \log k, \log^2(nd)\}$ . Then for all unit vectors  $v_1, \dots, v_n \in \mathbb{R}^d$ , setting  $\hat{\mu} = \mathcal{A}_{\text{UPU}}(\mathcal{R}_{\text{UPU}}(v_1), \dots, \mathcal{R}_{\text{UPU}}(v_n))$ , the local randomizers  $\mathcal{R}_{\text{UPU}}$  are  $\varepsilon$ -DP and*

$$\mathbb{E} \left[ \left\| \hat{\mu} - \frac{1}{n} \sum_{i=1}^n v_i \right\|_2^2 \right] \leq \text{Err}_{n,d}(\text{PrivUnitG}) \cdot \left( 1 + O \left( \frac{\varepsilon + \log k}{k} + \sqrt{\frac{\log^2(nd)}{k}} \right) \right).$$

---

**Algorithm 6** Nearly Unbiased ProjUnit (client)

---

**Require:** Input vector  $v \in \mathbb{R}^d$ , Bias bound probability  $\delta$ .

- 1: Randomly sample diagonal  $D$  from the Rademacher distribution based on predefined seed
- 2: Sample  $S \in \mathbb{R}^{k \times d}$  where each row is chosen uniformly at random without replacement from standard basis vectors  $\{e_1, \dots, e_d\}$
- 3: Let  $W = \sqrt{d/k}SHD$
- 4: Set  $C = 1 + \Theta(\sqrt{\log^2(k/\delta)/k})$
- 5: Project the input vector  $v_p = Wv$
- 6: Complete the norm then normalize:

$$u = \begin{cases} \frac{1}{\sqrt{C}} \left( v_p, \sqrt{C - \|v_p\|_2^2} \right), & \text{if } \|v_p\|_2^2 \leq C \\ \left( \frac{v_p}{\|v_p\|_2}, 0 \right), & \text{otherwise} \end{cases}$$

- 7: Let  $\hat{u} = \text{PrivUnitG}(u)$
  - 8: Send  $C, \hat{u}$  and (encoding of)  $S$  to server
- 

---

**Algorithm 7** Nearly Unbiased ProjUnit (server)

---

- 1: Receive  $C, \hat{u}_1, \dots, \hat{u}_n$ , from clients with (encodings of) transforms  $S_1, \dots, S_n$
- 2: Sample the diagonal matrices  $D$  from the Rademacher distribution based on predefined seed
- 3: Let  $U = HD$  for where  $H \in \mathbb{R}^{d \times d}$  is the Hadamard matrix
- 4: Return the estimate

$$\hat{\mu} = \frac{\sqrt{C}}{n} U^\top \sum_{i=1}^n S_i^\top \hat{u}_i[1:k]$$


---

*Proof.* Note that  $\hat{\mu} = \frac{\sqrt{C}}{n} \sum_{i=1}^n W_i^\top \hat{u}_i[1:k]$ . Thus we get

$$\mathbb{E} \left[ \left\| \hat{\mu} - \frac{1}{n} \sum_{i=1}^n v_i \right\|_2^2 \right] \tag{4}$$

$$\begin{aligned} &= \mathbb{E} \left[ \left\| \frac{\sqrt{C}}{n} \sum_{i=1}^n W_i^\top \hat{u}_i[1:k] - v_i \right\|_2^2 \right] \\ &= \frac{1}{n} \max_{i \in [n]} \mathbb{E} \left[ \left\| \sqrt{C} W_i^\top \hat{u}_i[1:k] - v_i \right\|_2^2 \right] + \frac{1}{n^2} \sum_{i \neq j} \mathbb{E} \langle \sqrt{C} W_i^\top \hat{u}_i[1:k] - v_i, \sqrt{C} W_j^\top \hat{u}_j[1:k] - v_j \rangle \end{aligned} \tag{5}$$

Now we upper bound both terms in (5) separately. Note that Corollary F.1 implies that with probability at least  $1 - n\delta$  we have  $\|W_i v_i\|_2^2 \leq (1 + C_1 \sqrt{\log^2(k/\delta)/k}) \|v_i\|_2^2$  for all  $i \in [n]$ . We let  $E$  denote the event that this event holds. Note that  $P(E) \geq 1 - \bar{\delta}$  where  $\bar{\delta} = n\delta$ .

We begin with the second term in (5). Let  $C = 1 + C_1 \sqrt{\log^2(k/\delta)/k}$  for appropriate constant  $C_1 > 0$ . Note that  $W_i = \sqrt{d/k} S_i H D$ , taking expectations over the randomness of the local randomizer, and noticing that PrivUnitG is unbiased, we have that for  $i \neq j$

$$\mathbb{E} \langle \sqrt{C} W_i^\top \hat{u}_i[1:k] - v_i, \sqrt{C} W_j^\top \hat{u}_j[1:k] - v_j \rangle = \mathbb{E} \langle \sqrt{C} W_i^\top u_i[1:k] - v_i, \sqrt{C} W_j^\top u_j[1:k] - v_j \rangle \tag{6}$$

Since  $W_i = \sqrt{d/k}S_iHD$  and  $W_j = \sqrt{d/k}S_jHD$ , we have

$$\begin{aligned}
& \mathbb{E} \left[ \langle \sqrt{C}W_i^\top u_i[1:k] - v_i, \sqrt{C}W_j^\top u_j[1:k] - v_j \rangle \right] \\
&= \mathbb{E} \left[ \langle W_i^\top W_i v_i - v_i, W_j^\top W_j v_j - v_j \rangle \mid E \right] P(E) \\
&+ \mathbb{E} \left[ \langle \sqrt{C}W_i^\top W_i v_i / \|W_i v_i\|_2 - v_i, \sqrt{C}W_j^\top W_j v_j / \|W_j v_j\|_2 - v_j \rangle \mid E^c \right] P(E^c) \\
&\leq v_i^\top \mathbb{E}[(W_i^\top W_i - I)(W_j^\top W_j - I) \mid E] v_j + \bar{\delta}(2Cd/k + 2) \\
&\leq v_i^\top \mathbb{E}[DH(d/kS_i^\top S_i - I)(d/kS_j^\top S_j - I)HD \mid E] v_j + \bar{\delta}(2Cd/k + 2) \\
&\stackrel{(i)}{\leq} \|\mathbb{E}[d/kS_i^\top S_i - I \mid E]\|_2 \|\mathbb{E}[d/kS_j^\top S_j - I \mid E]\|_2 + \bar{\delta}(2Cd/k + 2) \\
&\leq O((\bar{\delta}d/k)^2) + \bar{\delta}(2Cd/k + 2) \leq O(\bar{\delta}d/k),
\end{aligned}$$

where inequality (i) follows since  $\|\mathbb{E}[d/kS_j^\top S_j - I \mid E]\|_2 \leq O(\bar{\delta}/k)$  since

$$I = \mathbb{E}[(d/k)S_i^\top S_i] = \mathbb{E}[(d/k)S_i^\top S_i \mid E]P(E) + (1 - P(E))\mathbb{E}[(d/k)S_i^\top S_i \mid E^c]$$

which implies that

$$\mathbb{E}[(d/k)S_i^\top S_i \mid E] = \frac{I - (1 - P(E))\mathbb{E}[(d/k)S_i^\top S_i \mid E^c]}{P(E)}.$$

Since  $P(E) \geq 1 - \bar{\delta}$  and  $\|(d/k)S_i^\top S_i\|_2 \leq (d/k)$ , this shows that  $\|(\mathbb{E}[(d/k)S_i^\top S_i \mid E] - I)v\|_2 = O(\bar{\delta}d/k)$ .

Now we proceed to analyze the first term in (5). Note that for any  $i \in [n]$

$$\begin{aligned}
\mathbb{E} \left[ \left\| \sqrt{C}W_i^\top \hat{u}_i[1:k] - v_i \right\|_2^2 \right] &= \mathbb{E} \left[ \left\| \sqrt{C}W_i^\top \hat{u}_i[1:k] - \sqrt{C}W_i^\top u_i[1:k] + \sqrt{C}W_i^\top u_i[1:k] - v_i \right\|_2^2 \right] \\
&\stackrel{(i)}{=} C \mathbb{E} \left[ \left\| W_i^\top \hat{u}_i[1:k] - W_i^\top u_i[1:k] \right\|_2^2 \right] + \mathbb{E} \left[ \left\| \sqrt{C}W_i^\top u_i[1:k] - v_i \right\|_2^2 \right] \\
&\stackrel{(ii)}{\leq} C \mathbb{E} \left[ \left\| W_i^\top \right\|_2^2 \right] \cdot C \cdot \text{Err}_{1,k+1}(\text{PrivUnitG}) + \mathbb{E} \left[ \left\| \sqrt{C}W_i^\top u_i[1:k] - v_i \right\|_2^2 \right].
\end{aligned}$$

where (i) follows since  $\mathbb{E}[\hat{u}_i] = u_i$  as PrivUnitG is unbiased, and (ii) since PrivUnitG is applied for  $k+1$  dimensional vectors of squared norm  $C$ , hence its error is  $C \cdot \text{Err}_{1,k+1}(\text{PrivUnitG})$ . For the first term, as  $\|W_i\|_2^2 \leq d/k$  and  $C = 1 + C_1 \sqrt{\log^2(k/\delta)/k}$ , we have:

$$\begin{aligned}
C^2 \left\| W_i^\top \right\|_2^2 \cdot \text{Err}_{1,k+1}(\text{PrivUnitG}) &\leq C^2 \frac{d}{k} c_{k+1,\varepsilon} \frac{k+1}{\varepsilon} \\
&= C^2 \frac{d}{\varepsilon} c_{d,\varepsilon} \frac{c_{k+1,\varepsilon}}{c_{d,\varepsilon}} (1 + 1/k) \\
&= C^2 \frac{d}{\varepsilon} c_{d,\varepsilon} \cdot \left( 1 + O\left(\frac{\varepsilon + \log k}{k}\right) \right) \\
&= \text{Err}_{1,d}(\text{PrivUnitG}) \cdot \left( 1 + O\left(\frac{\varepsilon + \log k}{k} + \sqrt{\frac{\log^2(k/\delta)}{k}}\right) \right),
\end{aligned}$$

where the third step follows from Proposition 5. For the second term, we have

$$\begin{aligned}
\mathbb{E} \left[ \left\| \sqrt{C}W_i^\top u_i[1:k] - v_i \right\|_2^2 \right] &= \mathbb{E} \left[ \left\| \sqrt{C}W_i^\top u_i[1:k] - \sqrt{C}W_i^\top W_i v_i + \sqrt{C}W_i^\top W_i v_i - v_i \right\|_2^2 \right] \\
&\leq 2C \mathbb{E} \left[ \left\| W_i^\top u_i[1:k] - W_i^\top W_i v_i \right\|_2^2 \right] + 2 \mathbb{E} \left[ \left\| \sqrt{C}W_i^\top W_i v_i - v_i \right\|_2^2 \right] \\
&\leq 2C \mathbb{E} \left[ \left\| W_i^\top u_i[1:k] - W_i^\top W_i v_i \right\|_2^2 \right] + 2 \mathbb{E} \left[ \left\| \sqrt{C}W_i^\top W_i v_i - W_i^\top W_i v_i \right\|_2^2 \right] \\
&\quad + 2 \mathbb{E} \left[ \left\| W_i^\top W_i v_i - v_i \right\|_2^2 \right] \\
&\leq 2C \mathbb{E} \left[ \left\| W_i^\top u_i[1:k] - W_i^\top W_i v_i \right\|_2^2 \right] + 2(\sqrt{C} - 1)^2 d/k + 2(d/k - 1),
\end{aligned}$$

where the second inequality follows since  $\mathbb{E}[W_i^\top W_i] = I$ . Now we have

$$\begin{aligned} \mathbb{E} \left[ \left\| W_i^\top u_i[1:k] - W_i^\top W_i v_i \right\|_2^2 \right] &= \mathbb{E} \left[ \left\| W_i^\top W_i v_i / \sqrt{C} - W_i^\top W_i v_i \right\|_2^2 \mid E \right] \mathbb{P}(E) \\ &\quad + \mathbb{E} \left[ \left\| W_i^\top W_i v_i / \|W_i v_i\|_2 - W_i^\top W_i v_i \right\|_2^2 \mid E^c \right] \mathbb{P}(E^c) \\ &\leq (1/\sqrt{C} - 1)^2 d/k + 2\bar{\delta}(d/k + (d/k)^2) \\ &\leq O \left( \frac{d\sqrt{\log^2(k/\delta)}}{k^{3/2}} + \frac{\bar{\delta}d^2}{k^2} \right). \end{aligned}$$

Overall, putting these back in Inequality (5), we get

$$\begin{aligned} \mathbb{E} \left[ \left\| \hat{\mu} - \frac{1}{n} \sum_{i=1}^n v_i \right\|_2^2 \right] &= \frac{1}{n} \text{Err}_{1,d}(\text{PrivUnitG}) \cdot \left( 1 + O \left( \frac{\varepsilon + \log k}{k} + \sqrt{\frac{\log^2(k/\delta)}{k}} \right) \right) \\ &\quad + O \left( \frac{1}{n} \left( \frac{d\sqrt{\log^2(k/\delta)}}{k^{3/2}} + \frac{n\delta d^2}{k^2} + \frac{d}{k} \right) + (n\delta d/k)^2 \right). \end{aligned}$$

Noting that  $\text{Err}_{1,d}(\text{PrivUnitG})/n = \text{Err}_{n,d}(\text{PrivUnitG}) = c_{d,\varepsilon} \cdot \frac{d}{n\varepsilon}$  for some constant  $c_{d,\varepsilon}$ , this implies the theorem given that  $\delta = k/n^2 d$ .  $\square$

## B Missing proofs for Section 2

### B.1 Proof of Theorem 1

First, note that the claim about privacy follows directly from the privacy guarantees of PrivUnitG [7] as our algorithm applies PrivUnitG over a certain input vector with unit norm.

For accuracy, note that  $\hat{\mu} = \frac{1}{n} \sum_{i=1}^n W_i^\top \hat{u}_i$ , therefore

$$\begin{aligned} \mathbb{E} \left[ \left\| \hat{\mu} - \frac{1}{n} \sum_{i=1}^n v_i \right\|_2^2 \right] &= \mathbb{E} \left[ \left\| \frac{1}{n} \sum_{i=1}^n W_i^\top \hat{u}_i - v_i \right\|_2^2 \right] \\ &= \mathbb{E} \left[ \left\| \frac{1}{n} \sum_{i=1}^n W_i^\top \hat{u}_i - W_i^\top u_i + W_i^\top u_i - v_i \right\|_2^2 \right] \\ &\stackrel{(i)}{=} \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \left[ \left\| W_i^\top \hat{u}_i - W_i^\top u_i \right\|_2^2 \right] + \frac{1}{n^2} \mathbb{E} \left[ \left\| \sum_{i=1}^n W_i^\top u_i - v_i \right\|_2^2 \right] \\ &\leq \frac{1}{n} \max_{i \in [n]} \mathbb{E} \left[ \left\| W_i^\top \right\|_2^2 \right] \cdot \text{Err}_{1,k}(\text{PrivUnitG}) + \frac{1}{n^2} \mathbb{E} \left[ \left\| \sum_{i=1}^n W_i^\top u_i - v_i \right\|_2^2 \right]. \end{aligned}$$

where (i) follows since  $\mathbb{E}[\hat{u} \mid W_i = w_i] = u$  as PrivUnitG is unbiased. Now we analyze each of these two terms separately. For the first term, as  $\mathbb{E}[\|W_i\|_2^2] \leq d/k + \beta_{\mathcal{W}}$  for all  $i \in [n]$  we have that

is bounded by

$$\begin{aligned}
\max_{i \in [n]} \mathbb{E} \left[ \|W^\top\|_2^2 \right] \cdot \text{Err}_{1,k}(\text{PrivUnitG}) &\leq \left( \frac{d}{k} + \beta_{\mathcal{W}} \right) c_{k,\varepsilon} \frac{k}{\varepsilon} \\
&= \left( \frac{d}{\varepsilon} + \frac{\beta_{\mathcal{W}} k}{\varepsilon} \right) c_{d,\varepsilon} \frac{c_{k,\varepsilon}}{c_{d,\varepsilon}} \\
&= \left( \frac{d}{\varepsilon} + \frac{\beta_{\mathcal{W}} k}{\varepsilon} \right) c_{d,\varepsilon} \cdot \left( 1 + O\left( \frac{\varepsilon + \log k}{k} \right) \right) \\
&= \text{Err}_{1,d}(\text{PrivUnitG}) \cdot \left( 1 + \frac{\beta_{\mathcal{W}} k}{d} + O\left( \frac{\varepsilon + \log k}{k} \right) \right),
\end{aligned}$$

where the third step follows from Proposition 5. For the second term,

$$\begin{aligned}
\mathbb{E} \left[ \left\| \sum_{i=1}^n W_i^\top u_i - v_i \right\|_2^2 \right] &= \sum_{i=1}^n \sum_{j \neq i} \mathbb{E} [\langle W_i^\top u_i - v_i, W_j^\top u_j - v_j \rangle] + \sum_{i=1}^n \mathbb{E} [\|W_i^\top u_i - v_i\|_2^2] \\
&\leq \sum_{i=1}^n \sum_{j \neq i} \|\mathbb{E} W_i^\top u_i - v_i\|_2 \cdot \|\mathbb{E} W_j^\top u_j - v_j\|_2 + \sum_{i=1}^n \mathbb{E} [\|W_i^\top u_i - v_i\|_2^2] \\
&\leq n(n-1)\alpha_{\mathcal{W}} + \sum_{i=1}^n \mathbb{E} [\|W_i^\top u_i\|_2^2 + \|v_i\|_2^2 - 2v_i^\top W_i^\top u_i] \\
&= n(n-1)\alpha_{\mathcal{W}} + \sum_{i=1}^n \mathbb{E} [\|W_i^\top u_i\|_2^2 + 1 - 2\|W_i v_i\|_2] \\
&\leq n(n-1)\alpha_{\mathcal{W}} + n \max_{i \in [n]} \mathbb{E} [\|W_i^\top\|_2^2] + n.
\end{aligned}$$

Overall, this shows that

$$\begin{aligned}
\mathbb{E} \left[ \left\| \hat{\mu} - \frac{1}{n} \sum_{i=1}^n v_i \right\|_2^2 \right] &\leq \text{Err}(\text{PrivUnitG}_d, n) \cdot \left( 1 + O\left( \frac{\varepsilon + \log k}{k} \right) \right) \\
&\quad + O\left( \frac{d}{nk} \right) + \frac{1}{n} + \frac{(n-1)\alpha_{\mathcal{W}}}{n}.
\end{aligned}$$

Noticing that  $\text{Err}_{n,d}(\text{PrivUnitG}) = c_{d,\varepsilon} \cdot \frac{d}{n\varepsilon}$  for some constant  $c_{d,\varepsilon}$  (see [7]), this implies that

$$\mathbb{E} \left[ \left\| \hat{\mu} - \frac{1}{n} \sum_{i=1}^n v_i \right\|_2^2 \right] \leq \text{Err}_{n,d}(\text{PrivUnitG}) \cdot \left( 1 + O\left( \frac{\varepsilon + \log k}{k} \right) \right) + \alpha_{\mathcal{W}}.$$

This completes the proof.

## B.2 Proof of Proposition 1

The first item follows immediately as  $U \in \mathbb{R}^{d \times d}$  is a random rotation matrix where  $U^\top U = I$ , hence  $\|U\| \leq 1$ .

For the second item, we use a change of variables. Let  $W' = WP^\top$  where  $P$  is the rotation matrix such that  $Pv = e_1$ , the first standard basis vector. Recall that the rotation matrix  $P$  is orthogonal i.e.  $P^\top = P^{-1}$ . Due to the rotational symmetry of rotation matrices,  $W'$  is a random also a random rotation matrix. Note that  $W = W'P$ .

$$\begin{aligned}
\mathbb{E}_W \left[ \frac{W^\top W v}{\|W v\|_2} \right] &= \mathbb{E}_{W'} \left[ \frac{1}{\|W' P v\|_2} P^\top W'^\top W' P v \right] \\
&= P^\top \mathbb{E} \left[ \underbrace{\frac{1}{\|W' e_1\|_2} W'^\top W' e_1}_z \right]
\end{aligned}$$

Notice that  $z_j = \frac{1}{\|W'e_1\|_2} e_j^\top W'^\top W'e_1 = \langle W'e_j, \frac{1}{\|W'e_1\|_2} W'e_1 \rangle$ . First, note that  $\mathbb{E}[z_j] = 0$  for all  $j > 1$ . Moreover,  $z_1 = \frac{1}{\|W'e_1\|_2} \langle W'e_1, W'e_1 \rangle = 1 + O(1/k)$ . We let  $c = \mathbb{E}[z_1]$ . Thus,  $\mathbb{E}[z] = ce_1$  and  $\mathbb{E}[W^\top Wv / \|Wv\|_2] = cP^\top e_1 = cP^\top Pv = cv$ . Therefore,  $\left\| \mathbb{E}\left[\frac{W^\top Wv}{\|Wv\|_2} - v\right] \right\|_2 = |c - 1| \|v\|_2 = O(1/k)$ .

### B.3 Proof of Proposition 2

The bound on the operator norm is straightforward and follows from the fact that the Hadamard transform has operator norm bounded by 1.

Next we bound the bias. Let  $\delta = \min(1/d^2, k/2d)$  and let  $E_1$  denote the event where  $\|Wv\| \in 1 \pm O(\ln(k/\delta)/\sqrt{k})$ . By Corollary F.1,  $E_1$  happens with probability  $1 - \delta$ . Note that  $W^\top W$  is PSD and  $\mathbb{E}[W^\top W] = I$ . Thus,

$$\begin{aligned} \left\| \mathbb{E}\left[\frac{W^\top Wv}{\|Wv\|} - W^\top Wv\right] \right\| &\leq \left\| \mathbb{E}\left[\frac{W^\top Wv}{\|Wv\|} - W^\top Wv \mid E_1\right] \right\| + \left\| \mathbb{E}\left[\frac{W^\top Wv}{\|Wv\|} - W^\top Wv \mid \overline{E_1}\right] \right\| \mathbb{P}(\overline{E_1}) \\ &\leq \left\| \mathbb{E}\left[\left(\frac{1}{\|Wv\|} - 1\right) W^\top Wv \mid E_1\right] \right\| + (\|W^\top\| + 1)\mathbb{P}(\overline{E_1}) \\ &\leq \left\| \mathbb{E}\left[\left(\frac{1}{\|Wv\|} - 1\right) W^\top W \mid E_1\right] \right\| + (\|W^\top\| + 1)\mathbb{P}(\overline{E_1}) \\ &\leq \left\| \mathbb{E}\left[\left|\frac{1}{\|Wv\|} - 1\right| W^\top W \mid E_1\right] \right\| + (\|W^\top\| + 1)\mathbb{P}(\overline{E_1}) \\ &\leq O(\ln(k/\delta)/\sqrt{k}) + (\sqrt{d/k} + 1)\delta \end{aligned}$$

Substituting in the value of  $\delta$  gives the desired bound on the bias, by noticing that  $\left\| \mathbb{E}[W^\top W \mid E_1] \right\|_2 \leq 2$  since for any unit vector  $x$ ,

$$\begin{aligned} x^\top Ix &= \mathbb{E}[x^\top W^\top Wx] \\ &= \mathbb{E}[x^\top W^\top Wx \mid E_1]\mathbb{P}(E_1) + \mathbb{E}[x^\top W^\top Wx \mid \overline{E_1}]\mathbb{P}(\overline{E_1}) \\ &\geq \mathbb{E}[x^\top W^\top Wx \mid E_1](1 - \delta). \end{aligned}$$

In other words,  $\mathbb{E}[x^\top W^\top Wx \mid E_1] \leq \frac{1}{1-\delta}$  for all  $x \in \mathbb{R}^d$  with unit norm.

### C Proof of Theorem 4

The proof of this result follows from the next proposition.

**Proposition 3.** *Let  $k \leq d$ ,  $G \geq 1$  be an integer,  $U_1, \dots, U_G$  and  $W_1, \dots, W_n$  be sampled as described in (3). Moreover,  $U_j$  for  $j \in [G]$  and  $W_i$  for  $i \in [n]$  satisfy:*

1. *Bounded operator norm:*  $\|U_j^\top\| \leq 1$ .
2. *Bounded bias:*  $\left\| \mathbb{E}\left[\frac{W_i^\top W_i v}{\|W_i v\|_2} - v\right] \right\|_2 \leq \sqrt{\alpha_W}$  for all unit vectors  $v \in \mathbb{R}^d$ .

Then for all unit vectors  $v_1, \dots, v_n \in \mathbb{R}^d$ , setting  $\hat{\mu} = \mathcal{A}_{\text{CPU}}(\mathcal{R}_{\text{CPU}}(v_1), \dots, \mathcal{R}_{\text{CPU}}(v_n))$ ,

$$\mathbb{E}\left[\left\| \hat{\mu} - \frac{1}{n} \sum_{i=1}^n v_i \right\|_2^2\right] \leq \text{Err}_{n,d}(\text{PrivUnitG}) \cdot \left(1 + O\left(\frac{\varepsilon + \log k}{k} + \frac{n\varepsilon \log^2(nd)}{Gdk}\right)\right) + \alpha_W.$$

Before proving the proposition, we can now prove Theorem 4.

*Proof.* The proof follows from Proposition 3 by noting that the server returns  $\sum_{i=1}^n W_i^\top \hat{u}_i / n$ . The first property holds immediately from the definition of  $U_j$ . Moreover, for the second property,

Proposition 2 implies that  $\alpha_{\mathcal{W}} = O(\log^2(d)/k)$ . Since  $\text{Err}_{n,d}(\text{PrivUnitG}) = \Theta(d/n\varepsilon)$ , the claim about utility follows.

Now we prove the part regarding runtime. First, note that calculating the matrix  $D$  can be done efficiently using standard techniques [36, 5]. The server has to calculate the quantity

$$U^\top \sum_{i=1}^n S_i^\top \hat{u}_i.$$

Note that the summation has vectors which are  $k$ -sparse, therefore can be done in time  $O(nk)$ . Then, we have a multiplication step by Hadamard transform, which can be done in  $O(d \log d)$ .  $\square$

We now prove Proposition 3.

*Proof.* Note that  $\hat{\mu} = \frac{1}{n} \sum_{i=1}^n W_i^\top \hat{u}_i$ . Therefore we have

$$\begin{aligned} \mathbb{E} \left[ \left\| \hat{\mu} - \frac{1}{n} \sum_{i=1}^n v_i \right\|_2^2 \right] &= \mathbb{E} \left[ \left\| \frac{1}{n} \sum_{i=1}^n W_i^\top \hat{u}_i - v_i \right\|_2^2 \right] \\ &= \mathbb{E} \left[ \left\| \frac{1}{n} \sum_{i=1}^n W_i^\top \hat{u}_i - W_i^\top u_i + W_i^\top u_i - v_i \right\|_2^2 \right] \\ &\stackrel{(i)}{=} \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \left[ \|W_i^\top \hat{u}_i - W_i^\top u_i\|_2^2 \right] + \frac{1}{n^2} \mathbb{E} \left[ \left\| \sum_{i=1}^n W_i^\top u_i - v_i \right\|_2^2 \right] \\ &\leq \frac{1}{n} \max_{i \in [n]} \mathbb{E} \left[ \|W_i^\top\|_2^2 \right] \cdot \text{Err}_{1,k}(\text{PrivUnitG}) + \frac{1}{n^2} \mathbb{E} \left[ \left\| \sum_{i=1}^n W_i^\top u_i - v_i \right\|_2^2 \right]. \end{aligned}$$

where (i) follows since  $\mathbb{E}[\hat{u}] = u$  as PrivUnitG is unbiased. Now we analyze each of these two terms separately. For the first term, as  $\mathbb{E}[\|W_i\|^2] \leq d/k$  for all  $i \in [n]$  we have that it is bounded by

$$\begin{aligned} \max_{i \in [n]} \mathbb{E} \left[ \|W_i^\top\|_2^2 \right] \cdot \text{Err}_{1,k}(\text{PrivUnitG}) &\leq \frac{d}{k} c_{k,\varepsilon} \frac{k}{\varepsilon} = \frac{d}{\varepsilon} c_{d,\varepsilon} \frac{c_{k,\varepsilon}}{c_{d,\varepsilon}} \\ &= \frac{d}{\varepsilon} c_{d,\varepsilon} \cdot \left( 1 + O\left(\frac{\varepsilon + \log k}{k}\right) \right) \\ &= \text{Err}_{1,d}(\text{PrivUnitG}) \cdot \left( 1 + O\left(\frac{\varepsilon + \log k}{k}\right) \right), \end{aligned}$$

where the third step follows from Proposition 5. For the second term,

$$\mathbb{E} \left[ \left\| \sum_{i=1}^n W_i^\top u_i - v_i \right\|_2^2 \right] = \sum_{i=1}^n \sum_{j \neq i} \mathbb{E} [\langle W_i^\top u_i - v_i, W_j^\top u_j - v_j \rangle] + \sum_{i=1}^n \mathbb{E} \left[ \|W_i^\top u_i - v_i\|_2^2 \right].$$

For the second term note that

$$\begin{aligned} \sum_{i=1}^n \mathbb{E} \left[ \|W_i^\top u_i - v_i\|_2^2 \right] &= \sum_{i=1}^n \mathbb{E} \left[ \|W_i^\top u_i\|_2^2 + \|v_i\|_2^2 - 2v_i^\top W_i^\top u_i \right] \\ &= \sum_{i=1}^n \mathbb{E} \left[ \|W_i^\top u_i\|_2^2 + 1 - 2\|W_i v_i\|_2 \right] \\ &\leq n \max_{i \in [n]} \mathbb{E} \left[ \|W_i^\top\|_2^2 \right] + n \leq n(d/k + 1). \end{aligned}$$

For the first term, we have

$$\begin{aligned} &\mathbb{E} [\langle W_i^\top u_i - v_i, W_j^\top u_j - v_j \rangle] \\ &= \mathbb{E} [\langle W_i^\top (u_i - W v_i) + W_i^\top W_i v_i - v_i, W_j^\top (u_j - W_j v_j) + W_j^\top W_j v_j - v_j \rangle] \\ &= \mathbb{E} [\langle W_i^\top (u_i - W v_i), W_j^\top (u_j - W_j v_j) \rangle] + \mathbb{E} [\langle W_i^\top W_i v_i - v_i, W_j^\top (u_j - W_j v_j) \rangle] \\ &\quad + \mathbb{E} [\langle W_i^\top (u_i - W v_i), W_j^\top W_j v_j - v_j \rangle] + \mathbb{E} [\langle W_i^\top W_i v_i - v_i, W_j^\top W_j v_j - v_j \rangle] \end{aligned}$$

Because  $\mathbb{E}_{S_i} [\frac{d}{k} S_i^\top S_i] = I$ ,  $H^\top H = I$ , and  $D^\top D = I$ , we can evaluate the second term:

$$\mathbb{E}_{S_i} [\langle W_i^\top W_i v_i - v_i, W_j^\top (u_j - W_j v_j) \rangle] = \mathbb{E}_{S_i} \left[ \left\langle \frac{d}{k} D^\top H^\top S_i^\top S_i H D v_i - v_i, W_j^\top (u_j - W_j v_j) \right\rangle \right] = 0$$

Similarly, we can evaluate the fourth term:

$$\mathbb{E} [\langle W_i^\top W_i v_i - v_i, W_j^\top W_j v_j - v_j \rangle] = \mathbb{E} \left[ \left\langle \frac{d}{k} D^\top H^\top S_i^\top S_i H D v_i - v_i, W_j^\top W_j v_j - v_j \right\rangle \right] = 0$$

The third term is similar. Thus, we only need to bound the first term. First we give an upper bound that holds with probability 1.

$$\begin{aligned} \langle W_i^\top (u_i - W_i v_i), W_j^\top (u_j - W_j v_j) \rangle &\leq \|W_i^\top\| \|W_j^\top\| (\|W_i\| + 1) (\|W_j\| + 1) \\ &\leq O(d^2/k^2) \end{aligned}$$

Let  $E_1$  be the event that  $\|W_i v_i\|, \|W_j v_j\| \in 1 \pm O(\ln(k/\delta)/\sqrt{k})$  where  $\delta$  is a parameter to be chosen later. We will split the expectation depending on the event  $E_1$ .

$$\begin{aligned} &\langle W_i^\top (u_i - W_i v_i), W_j^\top (u_j - W_j v_j) \rangle \\ &= \left( \frac{1}{\|W_i v_i\|} - 1 \right) \left( \frac{1}{\|W_j v_j\|} - 1 \right) v_i^\top W_i^\top W_i W_j^\top W_j v_j \\ &= \frac{d^2}{k^2} \left( \frac{1}{\|W_i v_i\|} - 1 \right) \left( \frac{1}{\|W_j v_j\|} - 1 \right) v_i^\top U^\top S_i^\top S_i \underbrace{U U^\top}_I S_j^\top S_j U v_j \end{aligned}$$

Note that  $M := S_i^\top S_i S_j^\top S_j$  is PSD (both  $S_i^\top S_i$  and  $S_j^\top S_j$  are diagonal matrices and so is the product). Furthermore,  $\mathbb{E}[\frac{d}{k} S_i^\top S_i] = I$ . Thus

$$\begin{aligned} &\mathbb{E} \left[ \frac{d^2}{k^2} \left( \frac{1}{\|W_i v_i\|} - 1 \right) \left( \frac{1}{\|W_j v_j\|} - 1 \right) v_i^\top U^\top M U v_j, E_1 \right] \cdot \Pr[E_1] \\ &\leq \mathbb{E} \left[ \frac{d^2}{4k^2} \left( \frac{1}{\|W_i v_i\|} - 1 \right) \left( \frac{1}{\|W_j v_j\|} - 1 \right) (v_i^\top + v_j^\top) U^\top M U (v_i + v_j), E_1 \right] \cdot \Pr[E_1] \\ &\leq \mathbb{E} [O(\ln^2(k/\delta) d^2/k^3) (v_i^\top + v_j^\top) U^\top M U (v_i + v_j)] \\ &= O(\ln^2(k/\delta)/k) \end{aligned}$$

Therefore,

$$\mathbb{E} [\langle W_i^\top (u_i - W_i v_i), W_j^\top (u_j - W_j v_j) \rangle] \leq O(\ln^2(k/\delta)/k) + O(d^2/k^2) \cdot \Pr[E_1]$$

We now complete the proof of the claim. Combining the analysis, we get

$$\begin{aligned} \mathbb{E} \left[ \left\| \hat{\mu} - \frac{1}{n} \sum_{i=1}^n v_i \right\|_2^2 \right] &\leq \text{Err}_{n,d}(\text{PrivUnitG}) \cdot \left( 1 + O\left(\frac{\varepsilon + \log k}{k}\right) \right) \\ &\quad + O\left(\frac{d}{nk} + \frac{d^2 \delta}{Gk^2} + \frac{\ln^2(k/\delta)}{Gk} + \alpha_{\mathcal{W}}\right). \end{aligned}$$

Noticing that  $\text{Err}_{1,d}(\text{PrivUnitG})/n = \text{Err}_{n,d}(\text{PrivUnitG}) = c_{d,\varepsilon} \cdot \frac{d}{n\varepsilon}$  for some constant  $c_{d,\varepsilon}$  (see [7]), and  $\delta \leq k/(nd)$ , this implies that

$$\mathbb{E} \left[ \left\| \hat{\mu} - \frac{1}{n} \sum_{i=1}^n v_i \right\|_2^2 \right] \leq \text{Err}_{n,d}(\text{PrivUnitG}) \cdot \left( 1 + O\left(\frac{\varepsilon + \log k}{k} + \frac{n\varepsilon \ln^2(k/\delta)}{Gdk}\right) \right) + \alpha_{\mathcal{W}}.$$

This proves the claim.  $\square$

## D ProjUnit using Gaussian transforms

Building on the randomized projection framework of the previous section, in this section we instantiate it with the Gaussian transform. In particular, we sample  $W \in \mathcal{R}^{k \times d}$  from the Gaussian distribution where  $W$  has i.i.d.  $\mathcal{N}(0, 1/k)$  entries. The following theorem states our guarantees for this distribution.

**Theorem 7.** *Let  $k \leq d$  and  $W \in \mathcal{R}^{k \times d}$  be sampled from the Gaussian distribution where  $W$  has i.i.d.  $\mathcal{N}(0, 1/k)$  entries. Then for all unit vectors  $v_1, \dots, v_n \in \mathbb{R}^d$ , setting  $\hat{\mu} = \mathcal{A}_{\text{PU}}(\mathcal{R}_{\text{PU}}(v_1), \dots, \mathcal{R}_{\text{PU}}(v_n))$ ,*

$$\mathbb{E} \left[ \left\| \hat{\mu} - \frac{1}{n} \sum_{i=1}^n v_i \right\|_2^2 \right] \leq \text{Err}(\text{PrivUnit}G_d, n) \left( 1 + O \left( \sqrt{\frac{k}{d}} + \frac{\varepsilon + \log k}{k} \right) \right) + O \left( \frac{1}{k^2} \right).$$

The proof follows directly from Theorem 1 and the following proposition which proves certain properties of the Gaussian transform.

**Proposition 4.** *Consider  $W \in \mathbb{R}^{k \times d}$  with i.i.d.  $\mathcal{N}(0, 1/k)$  entries and a fixed  $v \in \mathbb{R}^d$ . Then*

1. *Bounded operator norm:*

$$\mathbb{E} \|W^\top\|^2 \leq \frac{d}{k} \left( 1 + O \left( \sqrt{\frac{k}{d}} \right) \right).$$

2. *Bounded bias: for every unit vector  $v \in \mathbb{R}^d$*

$$\left\| \mathbb{E} \frac{W^\top W v}{\|W v\|} - v \right\| = O(1/k).$$

*Proof.* For the first item, we rely on standard results in random matrix theory. If we let  $Z$  denote the top singular value of  $\sqrt{k}W^\top$ , then (17, Theorem 2.13) shows that for any  $t$ ,  $\Pr(Z > \sqrt{d} + \sqrt{k} + t) < \exp(-t^2/2)$ . This implies that  $\text{median}(Z) \leq \sqrt{d} + \sqrt{k} + 2$ . Further, by the isoperimetric inequality,  $Z$  is concentrated around its median with subGaussian tails, so that the second moment of  $Z - \text{median}(Z)$  is at most  $O(1)$ . Thus the second moment of  $Z$  is at most  $\text{median}(Z)^2 + O(1) \leq (\sqrt{d} + \sqrt{k} + 2)^2 + O(1)$ . Scaling this by  $k$ , we conclude that  $\mathbb{E}[\|W^\top\|_{\text{op}}^2] \leq \frac{d}{k}(1 + 2\sqrt{\frac{k}{d}} + \frac{O(1)}{k})$ .

For the second item, we use a change of variables. Let  $W' = WP^\top$  where  $P$  is the rotation matrix such that  $Pv = e_1$ , the first standard basis vector. Recall that the rotation matrix  $P$  is orthogonal i.e.  $P^\top = P^{-1}$ . Due to the rotational symmetry of the normal distribution,  $W'$  is a random matrix with i.i.d.  $\mathcal{N}(0, 1/k)$  entries. Note that  $W = W'P$ .

$$\begin{aligned} \mathbb{E}_W [W^\top u] &= \mathbb{E}_{W'} \left[ \frac{1}{\|W'Pv\|_2} P^\top W'^\top W'Pv \right] \\ &= P^\top \mathbb{E} \left[ \underbrace{\frac{1}{\|W'e_1\|_2} W'^\top W'e_1}_z \right] \end{aligned}$$

Notice that  $z_j = \frac{1}{\|W'e_1\|_2} e_j^\top W'^\top W'e_1 = \langle W'e_j, \frac{1}{\|W'e_1\|_2} W'e_1 \rangle$ . Because  $W'$  has i.i.d.  $\mathcal{N}(0, 1/k)$  entries,  $z_1 = \frac{1}{\|W'e_1\|_2}$  is  $1/\sqrt{k}$  times a  $\chi$  random variable with  $k$  degrees of freedom. We let  $\frac{1}{c} = \mathbb{E}[z_1] = \frac{1}{\sqrt{k}} \cdot \frac{\sqrt{2}\Gamma((k+1)/2)}{\Gamma(k/2)} = 1 - O(1/k)$  and  $\mathbb{E}[z_j] = 0 \forall j > 1$ . Thus,  $\mathbb{E}[z] = \frac{1}{c}e_1$  and  $\mathbb{E}[W_i^\top u_i] = \frac{1}{c}P_i^\top e_1 = \frac{1}{c}P^\top Pv = \frac{1}{c}v$ . Therefore,  $\left\| \mathbb{E} \left[ \frac{W^\top W v}{\|W v\|_2} - v \right] \right\|_2 = \|1/c - 1\|v\|_2 = O(1/k)$ .  $\square$

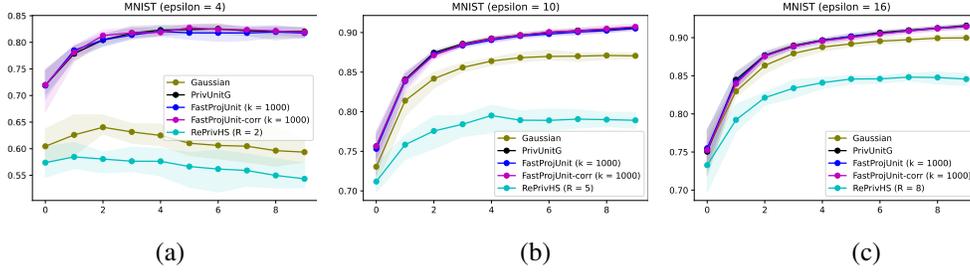


Figure 5: Test accuracy on the MNIST dataset as a function of epoch for (a)  $\epsilon = 4.0$ , (b)  $\epsilon = 10.0$  and (c)  $\epsilon = 16.0$ .

## E Additional plots for the MNIST experiment

We present additional details for the MNIST experiment including the description of the neural network (Table 2) and additional plots with different values of the privacy parameters for the MNIST experiment. In Figure 5, we present more plots for the MNIST experiment where we train models with several privacy parameters  $\epsilon \in \{4, 10, 16\}$ .

Layer	Parameters
Convolution + tanh	16 filters of $8 \times 8$ , stride 2, padding 2
Average pooling	$2 \times 2$ , stride 1
Convolution + tanh	32 filters of $4 \times 4$ , stride 2, padding 0
Average pooling	$2 \times 2$ , stride 1
Fully connected + tanh	32 units
Fully connected + tanh	10 units

Table 2: Architecture for convolutional network model.

## F Helper Lemmas

### F.1 Helper lemmas for random rotations

**Lemma F.1.** *Let  $x$  be a random unit vector on the unit ball of  $\mathbb{R}^d$  and  $z$  be the projection of  $x$  on to the last  $k$  coordinates. We have*

$$\left| \mathbb{E}[\|z\|] - \sqrt{k/d} \right| = O\left(\frac{1}{\sqrt{kd}}\right)$$

*Proof.* We represent  $d$  dimensional vector  $x$  using spherical coordinates as follows.

$$\begin{aligned} x_1 &= \cos(\phi_1) \\ x_2 &= \sin(\phi_1) \cos(\phi_2) \\ &\dots \\ x_{d-1} &= \sin(\phi_1) \cdots \sin(\phi_{d-2}) \cos(\phi_{d-1}) \\ x_d &= \sin(\phi_1) \cdots \sin(\phi_{d-2}) \sin(\phi_{d-1}) \end{aligned}$$

The squared length of the projection is

$$\sum_{i=d-k+1}^d x_i^2 = \sin^2(\phi_1) \cdots \sin^2(\phi_{d-k})$$

Recall the surface area element is  $\sin^{d-2}(\phi_1) \sin^{d-3}(\phi_2) \cdots \sin(\phi_{d-2}) d\phi_1 \cdots d\phi_{d-1}$ .

For  $k \geq 2$ , the expected length is

$$\frac{\int_0^\pi \cdots \int_0^\pi \int_0^{2\pi} (\sin(\phi_1) \cdots \sin(\phi_{d-k})) \sin^{d-2}(\phi_1) \sin^{d-3}(\phi_2) \cdots \sin(\phi_{d-2}) d\phi_1 \cdots d\phi_{d-1}}{S_{d-1}}$$

where  $S_{d-1}$  is the surface area of the unit ball, which is  $S_{d-1} = \frac{2\pi^{d/2}}{\Gamma(d/2)}$ .

We first evaluate the integral for each sine power.

**Claim F.1.** For integer  $n \geq 1$  we have

$$\int_0^\pi \sin^n x dx = \frac{\Gamma((n+1)/2)}{\Gamma(1+n/2)} \sqrt{\pi}$$

*Proof.* For  $n \geq 2$ , we have

$$\begin{aligned} \int \sin^n x dx &= - \int \sin^{n-1} x d(\cos x) \\ &= - \sin^{n-1} x \cos x + (n-1) \int \sin^{n-2} x \cos^2 x dx \\ &= - \sin^{n-1} x \cos x + (n-1) \int \sin^{n-2} x (1 - \sin^2 x) dx \end{aligned}$$

Thus,

$$\begin{aligned} \int_0^\pi \sin^n x dx &= \frac{n-1}{n} \int_0^\pi \sin^{n-2} x dx - \frac{\sin^{n-1} x \cos x \Big|_0^\pi}{n} \\ &= \frac{(n-1)/2}{n/2} \int_0^\pi \sin^{n-2} x dx \end{aligned}$$

The claim then follows using induction with base cases  $\int_0^\pi \sin x dx = 2$  and  $\int_0^\pi dx = \pi$ .  $\square$

$$\begin{aligned} &\int_0^\pi \cdots \int_0^\pi \int_0^{2\pi} (\sin(\phi_1) \cdots \sin(\phi_{d-k})) \sin^{d-2}(\phi_1) \sin^{d-3}(\phi_2) \cdots \sin(\phi_{d-2}) d\phi_1 \cdots d\phi_{d-1} \\ &= 2\pi \int_0^\pi \cdots \int_0^\pi \sin^{d-1}(\phi_1) \cdots \sin^k(\phi_{d-k}) \sin^{k-2}(\phi_{d-k+1}) \cdots \sin(\phi_{d-2}) d\phi_1 \cdots d\phi_{d-2} \\ &= 2\pi \frac{\Gamma(d/2)}{\Gamma((d+1)/2)} \sqrt{\pi} \cdots \frac{\Gamma((k+1)/2)}{\Gamma((k+2)/2)} \sqrt{\pi} \cdot \frac{\Gamma((k-1)/2)}{\Gamma(k/2)} \sqrt{\pi} \cdots \frac{\Gamma(1)}{\Gamma(3/2)} \sqrt{\pi} \\ &= 2\pi^{d/2} \frac{\Gamma((k+1)/2)}{\Gamma((d+1)/2) \Gamma(k/2)} \end{aligned}$$

The expected length is

$$\begin{aligned} \mathbb{E}[\|z\|] &= \frac{\Gamma((k+1)/2) \Gamma(d/2)}{\Gamma((d+1)/2) \Gamma(k/2)} \\ &= \frac{\Gamma(k) 2^d \Gamma(d/2)^2}{2^k \Gamma(k/2)^2 \Gamma(d)} \\ &= \frac{\sqrt{k} \left(1 - \frac{1}{4k} + O(1/k^2)\right)}{\sqrt{d} \left(1 - \frac{1}{4d} + O(1/d^2)\right)} \end{aligned}$$

In the second line, we use the Legendre duplication formula  $\Gamma(k/2) \Gamma((k+1)/2) = 2^{1-k} \sqrt{\pi} \Gamma(k)$ .

In the third line, we use the Stirling's approximation  $\Gamma(z) = \sqrt{2\pi/z} (z/e)^z (1 + 1/(12z) + O(1/z^2))$ .  $\square$

## F.2 SRHT Analysis

The Subsampled Randomized Hadamard Transform (SRHT) is the random matrix ensemble defined as  $W = \sqrt{\frac{d}{k}}SHD$ . Here  $D \in \mathbb{R}^{d \times d}$  is a diagonal matrix with independent uniform  $\pm 1$  values on its diagonal, and  $H \in \mathbb{R}^{d \times d}$  is the normalized Hadamard transform ( $H_{i,j} = (-1)^{\langle v(i), v(j) \rangle} / \sqrt{d}$ , where  $v(i)$  is the  $(\log_2 d)$ -dimensional vector obtained by writing  $i$  in binary). The matrix  $S \in \mathbb{R}^{k \times d}$  is a sampling matrix. The fact that the SRHT preserves the Euclidean norm of any fixed vector with large probability has been known for some time [3, 16, 35], though different works have analyzed slightly different variants of the SRHT, all having to do with how  $S$  is defined.

In this work, we make use of the SRHT in which  $S$  samples without replacement: that is, each row of  $S$  has a 1 in a uniformly random entry and zeroes elsewhere, and no two rows of  $S$  are equal. The tightest known analysis of the SRHT [16] analyzes the SRHT with a different sampling matrix:  $S_\eta = \text{diag}(\eta)$ , where  $\eta_1, \dots, \eta_d$  are independent Bernoulli random variables each with expectation  $k/d$  (so that we sample a *random* number of rows from  $HD$ , which is equal to  $k$  only in expectation).

The following is a special case of Theorem 9 in the full version of [16]

**Theorem 8** ([16]). *Suppose  $W = \sqrt{\frac{d}{k}}S_\eta HD$  for  $S = \text{diag}(\eta)$ , where  $\eta_1, \dots, \eta_d$  is a sequence of independent, uniform Bernoulli random variables each with expectation  $k/d$ . Then for some constant  $C > 0$ , for any fixed  $u \in \mathbb{R}^d$  of unit Euclidean norm and  $\delta \in (0, 1)$ ,*

$$\Pr_{\eta, D}(\|Wu\|_2^2 - 1 > C\sqrt{\log(1/\delta)\log(k/\delta)/k}) < \delta$$

An analysis of the SRHT using sampling without replacement then follows as a corollary.

**Corollary F.1.** *Suppose  $W = \sqrt{\frac{d}{k}}SHD$  is obtained with  $S$  being a  $k \times d$  sampling matrix without replacement. Then for some constant  $C > 0$ , for any fixed  $u \in \mathbb{R}^d$  of unit Euclidean norm and  $\delta \in (0, 1)$ ,*

$$\Pr_{\eta, D}(\|Wu\|_2^2 - 1 > C\sqrt{\log^2(k/\delta)/k}) < \delta$$

*Proof.* Consider  $W' = \sqrt{\frac{d}{k}}S_\eta HD$  with Bernoulli parameter  $k/d$ , as in Theorem 8. Then for any  $\delta' \in (0, 1)$  and fixed unit vector  $u \in \mathbb{R}^d$ ,  $\Pr_{\eta, D}(E) < \delta'$ , where  $E$  is the event that  $\|W'u\|_2^2 - 1 > C\sqrt{\log(1/\delta')\log(k/\delta')/k}$ . But we also have

$$\begin{aligned} \Pr(E) &\geq \Pr(E \cap (\|\eta\|_1 = k)) \\ &= \Pr(E \mid \|\eta\|_1 = k) \cdot \Pr(\|\eta\|_1 = k) \\ &= \Pr(E \mid \|\eta\|_1 = k) \cdot \Theta(1/\sqrt{k}). \end{aligned}$$

Note  $\Pr(E \mid \|\eta\|_1 = k)$  is exactly  $\Pr(\|Wu\|_2^2 - 1 > C\sqrt{\log(1/\delta')\log(k/\delta')/k})$ , where  $W$  is defined by sampling without replacement. Thus we have

$$\Pr(\|Wu\|_2^2 - 1 > C\sqrt{\log(1/\delta')\log(k/\delta')/k}) < C\delta'\sqrt{k}.$$

The claim then follows by applying the above with  $\delta' = \delta/(C\sqrt{k})$ .  $\square$

## G Details of PrivUnitG

For completeness, in this section we provide the full details of PrivUnitG which was proposed by Asi et al. [7]. Roughly, this algorithm uses the normal distribution to approximate the uniform distribution over the sphere for large dimensions. We refer the reader to [7] for more details about PrivUnitG.

In the algorithm,  $\Phi$  and  $\phi$  denote the Cumulative distribution function and probability density function for a Gaussian random variable  $\mathcal{N}(0, I_d)$ . There are multiple ways to set the parameters of PrivUnitG to achieve  $\varepsilon$ -DP; in our paper, we use the optimized parameters as described by Asi et al. [7], which allow to minimize the expected mean squared error (see Proposition 4 in [7]).

We note that Algorithm 8 describes the clients' algorithm (local randomizers) in the PrivUnitG protocol. The server aggregation simply adds all messages received from clients. Thus, we let  $\mathcal{R}_{\text{PrivUnitG}_\varepsilon}$  denote the local randomizer in Algorithm 8 (with optimized parameters to satisfy  $\varepsilon$ -DP) and let  $\mathcal{A}_{\text{PrivUnitG}_\varepsilon}$  denote the additive server aggregation.

---

**Algorithm 8** PrivUnitG( $p, q$ )

---

**Require:**  $v \in \mathbb{S}^{d-1}, q \in [0, 1], p \in [0, 1]$

- 1: Draw  $z \sim \text{Ber}(p)$
- 2: Let  $U = \text{N}(0, \sigma^2)$  where  $\sigma^2 = 1/d$
- 3: Set  $\gamma = \Phi_{\sigma^2}^{-1}(q) = \sigma \cdot \Phi^{-1}(q)$
- 4: **if**  $z = 1$  **then**
- 5:   Draw  $\alpha \sim U \mid U \geq \gamma$
- 6: **else**
- 7:   Draw  $\alpha \sim U \mid U < \gamma$
- 8: Draw  $V^\perp \sim \text{N}(0, \sigma^2(I - vv^T))$
- 9: Set  $V = \alpha v + V^\perp$
- 10: Calculate

$$m = \sigma \phi(\gamma/\sigma) \left( \frac{p}{1-q} - \frac{1-p}{q} \right)$$

- 11: Return  $\frac{1}{m} \cdot V$
- 

We also use the following useful result on the error of PrivUnitG for different dimensions. Recall that  $\text{Err}_{n,d}(\text{PrivUnitG}) = c_{d,\varepsilon} \frac{d}{n\varepsilon}$ . Then we have the following.

**Proposition 5** (Proposition 5, [7]). *Fix  $\varepsilon > 0$ . For any  $1 \leq k \leq d$ ,*

$$1 - O\left(\frac{\varepsilon + \log k}{k} + \frac{\varepsilon}{k}\right) \leq \frac{c_{k,\varepsilon}}{c_{d,\varepsilon}} \leq 1 + O\left(\frac{\varepsilon + \log k}{k} + \frac{\varepsilon}{k}\right).$$

## H Compressed PrivUnitG

Compressing the PrivUnit (resp. PrivUnitG) algorithm, using the technique of Feldman and Talwar [24], requires a pseudorandom generator that generates samples from a unit ball (resp. Gaussian) and fools spherical caps. As observed in [24], such PRGs with small seed length are known [30, 28]. However, the constructions in those works are optimized for seed length, and the computational cost of expanding a seed to a vector is a large polynomial. In this section, we argue that for inputs  $x$  having  $b$  bits of precision, we can compress PrivUnit/PrivUnitG to small seed length with a relatively efficient algorithm for seed expansion.

We will rely on Nisan's generator [33] which says that any space  $S$  computation that consumes  $N$  bits of randomness can be  $\delta$ -fooled using a random seed of length  $O(\log N(S + \log N/\delta))$ . Moreover, the computational cost of generating a pseudorandom string from a random seed is  $O(N \log N)$ . In our set up, the test that privacy of PrivUnit/PrivUnitG depends on is the  $[g \cdot x \geq \gamma]$ , when  $g$  is chosen from the Gaussian distribution. This probability that this test passes for the Gaussian distribution is  $e^{-c\varepsilon}$  for some constant  $c$ , and thus it suffices to set  $\delta$  to be  $e^{-c\varepsilon} \beta$  to ensure that mechanism satisfies  $(\varepsilon + 2\beta)$ -DP. For the rest of this discussion, we will set  $\beta = \varepsilon\tau/2$  which leads to  $\varepsilon' < \varepsilon(1 + \tau)$ . We can set  $\tau$  to be inverse polynomial as the dependence of the parameters on  $\tau$  will be logarithmic.

The test of interest for us can be implemented in  $S = O(\log d + b)$  space, and requires  $N = db$  bits of randomness. Plugging in these values and  $\delta = \varepsilon\tau e^{-c\varepsilon}/2$ , we get seed length  $O(\log db(b + c\varepsilon + \log db/\varepsilon\tau))$  and each expansion from seed to value requires run time  $O(db \log db)$ . Each run of PrivUnitG requires  $O(e^{c\varepsilon})$  expected random strings, leading to a run time of  $O(e^{c\varepsilon} bd \log db)$ .

For our algorithm, we run this on a  $k$ -dimensional vector instead of a  $d$ -dimensional one, with  $b = \log d$ . This gives us seed length  $O(\log(k \log d) \cdot (\log d + \varepsilon + \log(k \log d/\varepsilon)))$ . Given the projected vector, the run time is  $O(e^{c\varepsilon} k \log^2 d)$ .