
Private Vector Mean Estimation in the Shuffle Model: Optimal Rates Require Many Messages

Hilal Asi¹ Vitaly Feldman¹ Jelani Nelson² Huy L. Nguyen³ Kunal Talwar¹ Samson Zhou⁴

Abstract

We study the problem of private vector mean estimation in the shuffle model of privacy where n users each have a unit vector $v^{(i)} \in \mathbb{R}^d$. We propose a new multi-message protocol that achieves the optimal error using $\mathcal{O}(\min(n\epsilon^2, d))$ messages per user. Moreover, we show that any (unbiased) protocol that achieves optimal error requires each user to send $\Omega(\min(n\epsilon^2, d)/\log(n))$ messages, demonstrating the optimality of our message complexity up to logarithmic factors. Additionally, we study the single-message setting and design a protocol that achieves mean squared error $\mathcal{O}(dn^{d/(d+2)}\epsilon^{-4/(d+2)})$. Moreover, we show that *any* single-message protocol must incur mean squared error $\Omega(dn^{d/(d+2)})$, showing that our protocol is optimal in the standard setting where $\epsilon = \Theta(1)$. Finally, we study robustness to malicious users and show that malicious users can incur large additive error with a single shuffler.

1. Introduction

Vector mean estimation is a fundamental problem in federated learning, where a large number of distributed users can provide information to collaboratively train a machine learning model. Formally, there are n users that each have a real-valued vector $v^{(i)} \in \mathbb{R}^d$. In the vector mean estimation problem, the goal is to compute the average of the vectors $v = \frac{1}{n} \sum_{i=1}^n v^{(i)}$, whereas in the closely related vector aggregation problem, the goal is to compute the sum of the vectors $nv = \sum_{i=1}^n v^{(i)}$. As the privacy error scales with the norms of the vectors, we normalize and thus assume that $\|v^{(i)}\|_2 \leq 1$. The vectors could represent frequencies

¹Apple ²UC Berkeley ³Northeastern ⁴Texas A&M University.
Correspondence to: Hilal Asi <hilal.asi94@gmail.com>,
Vitaly Feldman <vitaly.edu@gmail.com>, Jelani Nelson <minilek@berkeley.edu>, Huy L. Nguyen <huylen-guyen@gmail.com>, Kunal Talwar <kunal@kunaltalwar.org>, Samson Zhou <samsonzhou@gmail.com>.

Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024. Copyright 2024 by the author(s).

of sequences of words in smartphone data for predictive text suggestions, shopping records for financial transactions or recommendation systems, various medical statistics for patients from different healthcare institutions, or gradient updates to be used to train a machine learning model. Thus, vector mean estimation and vector aggregation are used in a number of applications, such as deep learning through federated learning (Shokri & Shmatikov, 2015; Abadi et al., 2016; McMahan et al., 2017), frequent itemset mining (Sun et al., 2014), linear regression (Nguyễn et al., 2016), and stochastic optimization (Chaudhuri et al., 2011; Cheu et al., 2022).

Due to the sensitive nature of many of these data types, recent efforts have concentrated on facilitating federated analytics while preserving privacy. Differential privacy (DP) (Dwork et al., 2006) has emerged as a widely adopted rigorous mathematical definition that quantifies the amount of privacy leaked by a mechanism for any given individual user. In particular, local differential privacy (LDP) (Kasiviswanathan et al., 2011) demands that the distribution of the transcript of the communication protocol cannot be greatly affected by a change in a single distributed user's input. This approach enables the distributed collection of insightful statistics about a population, while protecting the private information of individual data subjects even with an untrusted curator who analyzes the collected statistics.

Unfortunately, in order to ensure privacy, the local model often requires a high amount of noise that results in poor accuracy of the resulting mechanisms. For example, in the simple case where $v^{(i)} \in \{0, 1\}$, i.e., binary summation, there exist private mechanisms with $\mathcal{O}(1)$ additive error in the central setting where the data curator is trusted (Dwork et al., 2016), but the additive error must be $\Omega(\sqrt{n})$ in the local model (Beimel et al., 2008; Chan et al., 2012). Consequently, the Encode, Shuffle, Analyze (ESA) model was proposed as an alternative distributed setting that could potentially result in a lower error (Bittau et al., 2017). The shuffle model of privacy is a special case of the ESA framework introduced by (Cheu et al., 2019), where a trusted shuffler receives and permutes a set of encoded messages from the distributed users, before passing them to an untrusted data curator. (Cheu et al., 2019) and (Erlingsson et al., 2019)

showed that for the important tasks of binary and real-valued summation, there are shuffle protocols that nearly match the accuracy of the optimal central DP mechanisms. Of note, (Balle et al., 2019; 2020; Ghazi et al., 2020; 2021) study the 1-dimensional real summation problem both under the lens of minimizing the error and the message complexity to achieve optimal error. In particular, (Ghazi et al., 2021) show that there is an optimal protocol that requires each user to send $1 + o(1)$ messages in expectation. However, the natural extension of their approach to d -dimensional mean estimation requires a number of messages that is exponential in d .

For d -dimensional mean estimation in the single-message setting, the most relevant works are that of (Scott et al., 2021; 2022), who study minimizing the mean-squared error of protocols that aim to compute the mean of vectors $u^{(1)}, \dots, u^{(n)} \in \mathbb{R}^d$, where each sampled vector $u^{(i)}$ consists of a number of coordinates sampled from the input vector $v^{(i)} \in \mathbb{R}^d$. (Scott et al., 2021; 2022) treat the sampled vectors $u^{(i)}$ as the true vectors and show a single-message shuffle protocol for estimating their mean. However, the mean-squared error of the overall protocol can be large, due to the large variance incurred by the procedure of sampling vectors $u^{(i)}$ from the true vectors.

Private vector mean estimation in the shuffle model is thus not well-understood, both under single-message and multi-message settings. In particular the following natural questions are open: first, in the multi-message setting, what is the total number of messages required in the shuffle model in order to obtain optimal rates for vector mean estimation. Secondly, what are the optimal algorithms for the single-messages setting.

Another desiderata in the design of distributed algorithms is that of robustness to malicious agents. In our context, we would like the system to be somewhat robust to one or a small number of clients that behave maliciously. For a problem like vector aggregation, a client can always misrepresent their input, and thus impact the sum; when vectors are restricted to having norm at most one, this can impact the true sum by at most two in the norm. The poisoning robustness of a protocol is defined to be ρ if the impact of an adversarial client on the computed sum is upper bounded, in Euclidean norm, by ρ . Thus a protocol that computes the exact sum has robustness 2. We would like to design protocols with robustness that is not much larger. We note that robustness of this kind has been previously been studied in other models of privacy (Cheu et al., 2021; Talwar, 2022). In the shuffle model with multiple messages, there are two different possible models from the robustness point of view.

In any implementation of a shuffle protocol that aims to achieve robustness, one must limit the number of contributions a single client can make: indeed if a single malicious

client can pretend to be a million different clients without being detected, one cannot hope to achieve any reasonable robustness. In typical implementations of a shuffler, such control can be achieved. For example, in a mix-net implementation of shuffling (Bittau et al., 2017), each client sends a non-anonymous but encrypted message to the first hop, where this first server can see who sent the message but not the contents of the message. This first hop can then validate that each sender sends at most one, or at most a predetermined number of messages to the server. When this bound is B and there are n clients, this server can implement this rate control using $\mathcal{O}(n)$ counters that can count up to B , for a total of $n \log_2(B + 1)$ bits of storage. We call this model, where each client can send a bounded number of messages to a single shuffler, the *multi-message shuffle* model.

This is distinct from a *multi-shuffler* model, where a client is allowed to send 1 message to each of B shufflers (or equivalently, B messages to a single shuffler with the constraint that there be at most 1 of each of B “types” of message). To ensure robustness, the shuffler would then need to rate-limit each type of message. When implemented in a mix-net setting as above, this multi-shuffler would require the first hop server to store $\mathcal{O}(nB)$ bits. It is easy to see that information-theoretically, a server cannot ensure nB separate rate limits using $o(nB)$ bits of state. For large B , this is significantly more than the $n \log_2(B + 1)$ bits that suffice for the multi-message shuffle.

Similarly in other implementations, e.g. those building on PrivacyPass or OHTTP tokens (Davidson et al., 2018; Thomson & Wood, 2023; Hendrickson et al., 2023), there is a server that implements the rate control at some step, and its cost scales as nB for multiple shufflers, compared to $n \log_2(B + 1)$ for a multi-message shuffle. Thus from an overhead point of view, these two models are significantly different. As a concrete example, when $n = 10^8$, and the vectors are $d = 10^6$ -dimensional, a d -message shuffle requires a few hundred megabytes of storage for the counters, whereas a multi-shuffler would require 12 terabytes of storage. It is thus much preferable to design algorithms that are robust in the multi-message shuffle model, rather than in the multi-shuffler model.

1.1. Our Contributions

In this work, we study the vector aggregation and vector mean estimation problems in the shuffle model of privacy, both in the single-message and multi-message settings, and from the viewpoint of robustness. We show the following results.

Multiple messages per user (Section 2). We consider the multi-message setting where users are allowed to send multiple messages. We propose a new protocol in the shuffle

model that obtains optimal mean squared error of $\tilde{\mathcal{O}}\left(\frac{d}{\varepsilon^2}\right)$ using $\tilde{\mathcal{O}}\left(\min(d, n\varepsilon^2)\right)$ messages per user, matching the performance of the central model of privacy (Bassily et al., 2014) up to logarithmic factors.

Theorem 1.1. *There exists an (ε, δ) -DP mechanism for vector aggregation that uses $\tilde{\mathcal{O}}\left(\min(d, n\varepsilon^2)\right)$ messages per user and achieves mean squared error $\tilde{\mathcal{O}}\left(\frac{d}{\varepsilon^2}\right)$.*

Moreover, we prove the following lower bound which shows that $\Omega(\min(n\varepsilon^2, d)/\log(n))$ messages are necessary in the shuffle model in order to obtain the optimal rate. The lower bound holds for any unbiased or summation protocol (as we define in Section 1.3).

Theorem 1.2. *For any (unbiased or summation) (ε, δ) -Shuffle DP protocol for vector aggregation that achieves the optimal mean squared error $\mathcal{O}\left(\frac{d}{\varepsilon^2}\right)$, must send $k = \Omega(\min(n\varepsilon^2, d)/\log(n))$ messages.*

Single message per user (Section 3). We also study the single-message setting where each user is allowed to send only a single message. We show that there exists a private protocol that can achieve mean squared error $\mathcal{O}\left(dn^{d/(d+2)}\varepsilon^{-4/(d+2)}\right)$.

Theorem 1.3. *For any $\varepsilon \in (0, 1)$, $\delta \in (0, 1)$, and $d, n \in \mathbb{N}$, there exists an (ε, δ) -DP protocol in the one-message shuffle model with mean squared error $\mathcal{O}_\delta\left(dn^{d/(d+2)}\varepsilon^{-4/(d+2)}\right)$.*

Though the mean squared error of Theorem 1.3 seems somewhat arbitrary, we show that it is tight for a single message per user shuffle.

Theorem 1.4. *Let \mathcal{P} be an (ε, δ) -DP protocol for vector aggregation on the unit ball \mathbb{B}_2^{d-1} in the one-message shuffle model with $\delta < \frac{1}{2}$. Then the mean squared error of \mathcal{P} satisfies $\text{MSE}(\mathcal{P}) = \Omega\left(dn^{d/(d+2)}\right)$.*

Robustness to malicious users (Section 4). We subsequently study the robustness of shuffle DP protocols to malicious users, who may distribute adversarial messages in an effort to induce the maximal possible mean squared error by a protocol.

We first show that for additive protocols in the multi-message shuffle model, each malicious user can induce additive mean squared error up to $\Omega\left(\frac{d}{\log^2(nd)}\right)$, for a total of $\Omega\left(\frac{kd}{\log^2(nd)}\right)$ additive mean squared error across k malicious users. More generally, we show the following result for the case of s shufflers.

Theorem 1.5. *Let $\varepsilon = \mathcal{O}(1)$ and $\delta < \frac{1}{nd}$. Then any (ε, δ) -DP mechanism for vector summation in which s shufflers take messages corresponding to a disjoint subset of the coordinates and returns the sum of the messages across*

n players with k malicious users has additive error mean squared error $\Omega\left(\frac{kd}{s\log^2(nd)}\right)$.

On the other hand, we show that our protocol is robust to malicious users when multiple shufflers exist: in this case, k malicious users can only induce error $\mathcal{O}(k)$, rather than $\Omega(kd)$. Since the input of each user is a vector with at most unit length, then our result essentially says that a malicious user can at most hide its input vector by generating the protocol for a different vector. By comparison, each malicious user in the context of Theorem 1.5 can be responsible for error $\Omega\left(\frac{d}{\log^2(nd)}\right)$, which can be significantly larger than the unit length of each input vector.

Thus our results show that a large class of accurate protocols in the multi-message shuffle model are inherently non-robust. While the multi-shuffler model can allow for better robustness, it comes at a significant additional cost. We remark that a trusted *aggregator* such as one built on top of PRIO (Corrigan-Gibbs & Boneh, 2017) can ensure high robustness as well as low overhead (c.f. (Rothblum et al., 2023)). While it is more complex to implement a trusted aggregator (compared to a shuffler), our results point to an important reason why a shuffler may not be sufficient when robustness is a concern.

1.2. Related Work

Mean estimation is a fundamental problem for data analytics and is the building block for many algorithms in stochastic optimization such as stochastic gradient descent. As a result, privacy-preserving frequency estimation has been extensively studied in applications of federated learning.

Real summation in the shuffle model. There has also been a line of work studying real summation, i.e., vector summation with $d = 1$, in the shuffle model. In the single-message shuffle model, (Balle et al., 2019) showed that the optimal additive error is $\tilde{\Theta}_\varepsilon(n^{1/6})$, whereas in the multi-message shuffle model, there exist protocols that achieve additive error $\mathcal{O}\left(\frac{1}{\varepsilon}\right)$ (Balle et al., 2020; Ghazi et al., 2020; 2021). In particular, (Balle et al., 2020; Ghazi et al., 2020) use the split-and-mix protocol of (Ishai et al., 2006) to achieve additive error $\mathcal{O}\left(\frac{1}{\varepsilon}\right)$, though at the cost of using at least 3 messages, each of length at least $\frac{\log \frac{1}{\delta}}{\log n}$. Subsequently, the protocol of (Ghazi et al., 2021) achieves near-optimal error, while using only $1 + o(1)$ messages per user in expectation.

Lower bounds for the multi-message shuffle model. For the problem of mean estimation, existing work does not have any lower bounds in the multi-shuffle model. However, for other problems such as private selection or parity learning, several recent papers have demonstrated new

lower bounds for the multi-message model (Cheu & Ullman, 2021; Chen et al., 2020a; Beimel et al., 2020). More precisely, (Cheu & Ullman, 2021) proved new lower bounds of $\Omega(\sqrt{D})$ on the sample complexity of selection from D candidates (and other learning problems) under the pan-privacy model, which implies lower bounds for the shuffle model. However, their results do not extend to our setting as high-dimensional mean estimation is not difficult in the pan-private model and thus do not translate to strong lower bounds for privacy in the shuffle model. Moreover, (Chen et al., 2020a) proved lower bounds of D/k for private selection for the multi-message model with k messages. Finally, (Beimel et al., 2020) consider the common element problem (which aims to identify an element that is common to all users) and prove non-trivial lower bounds for the multi-message model when k is small. However, these lower bounds are different from ours in two distinct ways: first, none of them hold for the problem of high-dimensional mean estimation, and secondly, they do not exhibit the same phase transition behavior that our lower bounds show, where an optimal rate is achieved only when $k \geq d$.

Mean estimation in the LDP model. (Duchi & Rogers, 2019; Duchi et al., 2016) studied the vector mean estimation problem in the LDP model, showing how to achieve optimal error without accounting for any communication constraints. (Bhowmick et al., 2018) developed a new algorithm, PrivUnit, and proved it is optimal up to constants, and (Asi et al., 2022) show that PrivUnit with optimized parameters is the optimal mechanism. More recently, there have been several works that study LDP aggregation with low communication cost such as (Chen et al., 2020b; Feldman & Talwar, 2021; Asi et al., 2023). Another line of work considers improving the communication cost in the setting where the input vector of each user is k -sparse (Bassily & Smith, 2015; Fanti et al., 2016; Ye & Barg, 2018; Acharya & Sun, 2019; Zhou et al., 2022).

1.3. Preliminaries and problem setting

Notation. We let $\mathbb{S}^{d-1} = \{v \in \mathbb{R}^d : \|v\|_2 = 1\}$ denote the d -dimensional sphere. For a set $S \subseteq \mathbb{R}^d$ and a vector $v \in \mathbb{R}^d$, define $\text{dist}(v, S) = \inf_{u \in S} \|v - u\|_2^2$. Let $\mathbb{B}^{d-1} = \{v \in \mathbb{R}^d : \|v\|_2 \leq 1\}$ be the unit ball in d dimensions.

We recall the standard definition of differential privacy.

Definition 1.6 (Differential privacy). (Dwork et al., 2006) Given a privacy parameter $\varepsilon > 0$ and a failure parameter $\delta \in (0, 1)$, a randomized algorithm $\mathcal{A} : \mathcal{D} \rightarrow \mathcal{R}$ is (ε, δ) -differentially private if, for every neighboring datasets $D, D' \in \mathcal{D}$ and for all $U \subseteq \mathcal{R}$,

$$\Pr[\mathcal{A}(D) \in U] \leq e^\varepsilon \cdot \Pr[\mathcal{A}(D') \in U] + \delta.$$

We require the standard advanced composition of differen-

tial privacy.

Theorem 1.7 (Advanced composition of differential privacy (Dwork & Roth, 2014)). Let $\varepsilon, \delta \geq 0$ and $\delta' > 0$. The composition of k algorithms that are each (ε, δ) -differentially private is itself $(\tilde{\varepsilon}, \tilde{\delta})$ -differentially private, where

$$\tilde{\varepsilon} = \varepsilon \sqrt{2k \ln(1/\delta')} + k\varepsilon \left(\frac{e^\varepsilon - 1}{e^\varepsilon + 1} \right), \quad \tilde{\delta} = k\delta + \delta'.$$

Shuffle differential privacy. In the shuffle DP model, we have n users, each holding a vector $v_i \in \mathbb{S}^{d-1}$. A protocol in this model is a pair of procedures $(\mathcal{A}, \mathcal{R})$ where $\mathcal{R} : \mathbb{S}^{d-1} \rightarrow \mathcal{Z}^k$ is a local randomizer that each user applies to produce k messages in \mathcal{Z} . Then, a shuffler Π is applied to all messages output by the users, before applying an aggregation $\mathcal{A} : \mathcal{Z}^* \rightarrow \mathbb{R}^d$ over the shuffled messages to return an output

$$\hat{v} = \mathcal{A}(\Pi(\mathcal{R}(v_1), \dots, \mathcal{R}(v_n))).$$

We say that a protocol $(\mathcal{A}, \mathcal{R})$ is (ε, δ) -Shuffle DP if the algorithm that outputs $\Pi(\mathcal{R}(v_1), \dots, \mathcal{R}(v_n))$ is (ε, δ) -DP. Moreover, we define the mean squared error $\text{Err}(\mathcal{A}, \mathcal{R})$ of the protocol to be

$$\sup_{v_1, \dots, v_n \in \mathbb{S}^{d-1}} \mathbb{E} \left[\left\| \mathcal{A}(\Pi(\mathcal{R}(v_1), \dots, \mathcal{R}(v_n))) - \sum_{i=1}^n v_i \right\|_2^2 \right].$$

Throughout the paper, we use the notion of *unbiased* and *summation* protocols. More specifically, we say that a protocol $(\mathcal{A}, \mathcal{R})$ is unbiased if for all $v_1, \dots, v_n \in \mathbb{S}^{d-1}$

$$\mathbb{E}[\mathcal{A}(\Pi(\mathcal{R}(v_1), \dots, \mathcal{R}(v_n))] = \sum_{i=1}^n v_i.$$

We also let \mathcal{A}^+ denote the summation aggregation, that is,

$$\mathcal{A}^+(\Pi(\mathcal{R}(v_1), \dots, \mathcal{R}(v_n))) = \sum_{m \in \Pi(\mathcal{R}(v_1), \dots, \mathcal{R}(v_n))} m.$$

These notions will be useful for our lower and upper bounds.

1.3.1. KASHIN REPRESENTATION

We use Kashin's representation in our multi-message algorithms, which has the following property.

Lemma 1.8 (Kashin's representation). (Lyubarskii & Vershynin, 2010) Let $d \geq 1$. There exists a transformation $U_K \in \mathbb{R}^{2d \times d}$ and a constant C_K such that

$$(1) \quad U_K^T U_K = I_d$$

$$(2) \quad \text{For all } x \in \mathbb{S}^{d-1}, \|U_K x\|_\infty \leq \frac{C_K}{\sqrt{d}}.$$

Here we use the subscript K simply to denote Kashin's representation. We call the matrix U_K the Kashin transformation.

We remark that Kashin's representation was first used for locally-private mean estimation in (Feldman et al., 2021).

2. Multiple Messages

In this section, we study algorithms for vector aggregation in the shuffle model of privacy when each user is permitted to send multiple messages. In particular, we study the number of messages that each user should send so that the resulting protocol can achieve the same mean squared error as the optimal mechanism in the central setting of DP. We first show a lower bound for the number of messages that must be sent per user to achieve the best possible error while guaranteeing DP. We then give an algorithm with matching number of messages per user, while achieving the best possible error for DP protocols.

2.1. $\tilde{\Omega}(\min(n\epsilon^2, d))$ messages are necessary

In this section, we prove that any unbiased shuffle DP protocol that obtains optimal error must send at least $k \geq \Omega\left(\frac{\min(n\epsilon^2, d)}{\log n}\right)$ messages per user. We prove this lower bound for summation protocols in Section 2.1.1 and for any unbiased protocol in Section 2.1.2.

In our setting, we have n users with inputs $v_1, \dots, v_n \in \mathbb{R}^d$ where $\|v_i\|_2 \leq 1$. Each user applies a local randomizer $\mathcal{R}(v_i)$ which sends k messages, $\mathcal{R}(v_i) = (m_i^1, \dots, m_i^k)$, then an aggregation protocol \mathcal{A} is applied over the shuffled messages, producing an output $\hat{\mu} = \mathcal{A}(\Pi(\mathcal{R}(v_1), \mathcal{R}(v_2), \dots, \mathcal{R}(v_n)))$, where Π is the shuffling operation.

2.1.1. LOWER BOUND FOR SUMMATION PROTOCOLS

We begin by proving the lower bound for summation protocols where $\mathcal{A}^+(m_1, \dots, m_{nk}) = \sum_{i=1}^{nk} m_i$. Throughout this section, we assume that the aggregation protocol $\mathcal{A} = \mathcal{A}^+$ and that $\mathcal{R} : \mathbb{S}^{d-1} \rightarrow \mathcal{Z}^k$ where $\mathcal{Z} = \mathbb{R}^d$.

Theorem 2.1. *Let $\epsilon, \delta \leq 1$ and $\mathcal{R} : \mathbb{S}^{d-1} \rightarrow \mathcal{Z}^k$ be an (ϵ, δ) -Shuffle DP randomizer. If $\text{Err}(\mathcal{A}^+, \mathcal{R}) \leq \mathcal{O}(d/\epsilon^2)$ then $k \geq \Omega\left(\frac{\min(n\epsilon^2, d)}{\log n}\right)$.*

Towards proving this result, we first prove the following symmetry property that is satisfied by an optimal summation protocol. For a randomizer \mathcal{R} , let $\mathcal{R}^+(v_i)$ denote the summation of all messages in $\mathcal{R}(v_i)$, that is, $\mathcal{R}^+(v_i) = \sum_{j=1}^k \mathcal{R}(v_i)_j$. We defer the proof to Appendix A.1.1.

Lemma 2.2. *Let $\epsilon \leq 1$, $\mathcal{R} : \mathbb{S}^{d-1} \rightarrow \mathcal{Z}^k$ be (ϵ, δ) -Shuffle DP. There exists an (ϵ, δ) -Shuffle DP randomizer $\hat{\mathcal{R}} : \mathbb{S}^{d-1} \rightarrow \mathcal{Z}^k$ such that*

$$(1) \text{ Err}(\mathcal{A}^+, \hat{\mathcal{R}}) \leq \text{Err}(\mathcal{A}^+, \mathcal{R})$$

(2) (Symmetry) For all $u, v \in \mathbb{S}^{d-1}$,

$$\mathbb{E} \left[\left\| \hat{\mathcal{R}}^+(v) - v \right\|_2^2 \right] = \mathbb{E} \left[\left\| \hat{\mathcal{R}}^+(u) - u \right\|_2^2 \right]$$

Proof. (sketch) The new randomizer $\hat{\mathcal{R}}$ works as follows: first, it samples a rotation matrix $U \in \mathbb{R}^{d \times d}$ (known public randomness) such that $U^T U = I$, then sets

$$\hat{\mathcal{R}}(v) = U^T \mathcal{R}(Uv),$$

where $U^T \mathcal{R}(Uv)$ denotes multiplying each message in $\mathcal{R}(Uv)$ by U^T . The lemma then follows using standard algebraic manipulations (see Appendix A.1.1 for full proof). \square

The proof of the lower bound builds on the following reconstruction attack against summation protocols. The attack essentially iterates over all subsets of messages of size k and adds their sum to the output set. We argue that if the protocol has small error (less than n), then the input vector will be close to a vector in the output set.

Algorithm 1 Reconstruction attack against summation protocols

Input: Shuffled set of messages $W = \{m_i\}_{i \in [nk]} \in \mathbb{R}^d$
Output: A set $S \subseteq \mathbb{R}^d$

- 1: Initialize $S = \emptyset$
- 2: **for** $t = 1$ to $\binom{nk}{k}$ **do**
- 3: Pick a (new) set of k messages from W ; denote it by W_t
- 4: $S \leftarrow S \cup \{\sum_{m \in W_t} m\}$
- 5: **end for**
- 6: Return S

The following proposition states the guarantees of this reconstruction attack.

Proposition 2.3. *Let $v_1, \dots, v_n \in \mathbb{S}^{d-1}$, $\mathcal{R} : \mathbb{S}^{d-1} \rightarrow \mathcal{Z}^k$ be randomizer that satisfies the symmetry condition of Lemma 2.2. For an input set $W = \Pi(\mathcal{R}(v_1), \dots, \mathcal{R}(v_n))$, Algorithm 1 outputs a set $S \subset \mathbb{R}^d$ of size $\binom{nk}{k}$ such that*

$$\mathbb{E} [\text{dist}(v_1, S)] = \mathbb{E} \left[\min_{u \in S} \|v_1 - u\|_2^2 \right] \leq \frac{\text{Err}(\mathcal{A}^+, \mathcal{R})}{n},$$

where the expectation is over the randomness of the algorithm.

We can now provide the main idea for proving Theorem 2.1. We defer the full proof to Appendix A.1.2.

Proof. (sketch) Consider $d \leq n\epsilon^2/100$ and let $P = \{v_1, v_2, \dots, v_M\}$ be a ρ -packing of \mathbb{S}^{d-1} where $\rho = 1/10$.

We will prove the lower bounds by analyzing the algorithm over the following M datasets:

$$X_i = (v_i, v_1, \dots, v_1).$$

The main idea is to show that if an algorithm is accurate, then our reconstruction attack Algorithm 1 will return a set S of size $\binom{nk}{k} \approx 2^{k \log(n)}$ that contains v_i . If $k \ll d$, then the size of the reconstructed set S is much smaller than the size of the packing P which contradicts privacy.

More formally, let S_i be the output of the reconstruction attack (Algorithm 1) over the input $\Pi(\mathcal{R}(v_i), \mathcal{R}(v_1), \dots, \mathcal{R}(v_1))$, and let O_i be the projection of S_i to the packing P ; that is, $O_i = \{\text{Proj}_P(v) : v \in S_i\}$.

Proposition 2.3 states that $\mathbb{E}[\text{dist}(v_i, S_i)] \leq \frac{d}{n\varepsilon^2} \leq 1/100$, hence we get that

$$\Pr(v_i \in O_i) \geq \Pr(\text{dist}(v_i, S_i) < \rho) \geq 9/10,$$

where the first inequality follows as P is ρ -packing, and the second inequality follows from Markov's inequality.

On the other hand, note that

$$\begin{aligned} \sum_{i=1}^M \Pr(v_i \in O_1) &= \sum_{i=1}^M E[1\{v_i \in O_1\}] \\ &= E\left[\sum_{i=1}^M 1\{v_i \in O_1\}\right] \\ &\leq E[|O_1|] \leq \binom{nk}{k}. \end{aligned}$$

Hence there exists an $1 \leq i \leq M$ such that

$$\Pr(v_i \in O_1) \leq \frac{\binom{nk}{k}}{M}.$$

As the protocol is (ε, δ) -DP, we also have

$$\begin{aligned} \Pr(v_i \in O_1) &\geq \Pr(v_i \in O_i)e^{-\varepsilon} - \delta \\ &\geq \frac{9}{10e} \geq 1/6. \end{aligned}$$

Combining these together, and given that $M \geq 2^d$ for $\rho = 1/10$, we have that

$$2^d \leq 6 \binom{nk}{k} \leq 6(en)^k.$$

This implies that $k \geq \Omega(d/\log(n))$ whenever $d \leq n\varepsilon^2/100$.

Now consider $d \geq n\varepsilon^2/100$. The proof builds on a reduction (Proposition A.1) which converts an optimal protocol for d -dimensional inputs into an optimal protocol for d' -dimensional inputs where $d' = n\varepsilon^2/200$ with the same number of messages. The lower bound then follows immediately from the lower bound for small d .

We provide the full details of the proof and the missing proof for $d \geq n\varepsilon^2/100$ in Appendix A.1. \square

2.1.2. LOWER BOUND FOR UNBIASED PROTOCOLS

In this section, we prove the same lower bound for any aggregation strategy as long as it is unbiased. We assume that the aggregation protocol \mathcal{A} is unbiased; that is, for all $v_1, \dots, v_n \in \mathbb{S}^{d-1}$,

$$\mathbb{E}[\mathcal{A}(\Pi(\mathcal{R}(v_1), \mathcal{R}(v_2), \dots, \mathcal{R}(v_n)))] = \sum_{i=1}^n v_i.$$

The lower bound builds on the following reconstruction attack against unbiased protocols (Algorithm 2). The attack follows the same recipe as the attack against summation protocols (Algorithm 1) to iterate over all subsets of messages of size k . However, given k messages, now we apply a different reconstruction scheme that uses the aggregation \mathcal{A} with zero-mean dummy inputs, and finally taking expectations.

Algorithm 2 Reconstruction attack against unbiased protocols

Input: Shuffled set of messages $W = \{m_i\}_{i \in [nk]}$
Output: A set $S \subseteq \mathbb{R}^d$

- 1: Initialize $S = \emptyset$
- 2: **for** $t = 1$ to $\binom{nk}{k}$ **do**
- 3: Pick a (new) set of k messages from W ; denote it by W_t
- 4: Calculate u_t to be $\mathbb{E}_{\tilde{v}_2, \dots, \tilde{v}_n \sim \text{Unif}(\mathbb{S}^{d-1})} [\mathcal{A}(\Pi(W_t, \mathcal{R}(\tilde{v}_2), \dots, \mathcal{R}(\tilde{v}_n)))]$
- 5: $S \leftarrow S \cup \{u_t\}$
- 6: **end for**
- 7: Return S

The following proposition states the guarantees of this reconstruction attack against unbiased protocols. We defer the proof to Appendix A.2.1.

Proposition 2.4. *Let $v_1, \dots, v_n \in \mathbb{S}^{d-1}$ where $v_1 \sim \text{Unif}(\mathbb{S}^{d-1})$, $\mathcal{R} : \mathbb{S}^{d-1} \rightarrow \mathcal{Z}^k$ and \mathcal{A} be an unbiased protocol. For an input set $W = \Pi(\mathcal{R}(v_1), \dots, \mathcal{R}(v_n))$, Algorithm 2 outputs a set $S \subset \mathbb{R}^d$ of size $\binom{nk}{k}$ such that*

$$\mathbb{E}[\text{dist}(v_1, S)] = \mathbb{E}\left[\min_{u \in S} \|v_1 - u\|_2^2\right] \leq \frac{\text{Err}(\mathcal{A}, \mathcal{R})}{n},$$

where the expectation is over the randomness of v_1 and the algorithm.

We can now prove our main lower bound for unbiased protocols. The proof is similar to the proof of Theorem 2.1 using the new construction attack. We defer it to Appendix A.2.2.

Theorem 2.5. *Let $\varepsilon \leq 1$, $\mathcal{R} : \mathbb{S}^{d-1} \rightarrow \mathcal{Z}^k$ be (ε, δ) -shuffle DP, and \mathcal{A} be an unbiased protocol. If $\text{Err}(\mathcal{A}, \mathcal{R}) \leq \mathcal{O}(d/\varepsilon^2)$ then $k \geq \Omega\left(\frac{\min(n\varepsilon^2, d)}{\log n}\right)$.*

2.2. Optimal multi-message protocol

In this section, we briefly overview a private protocol that achieves the optimal mean squared error for vector aggregation. The protocol requires that each user sends $\mathcal{O}(d)$ messages in expectation.

We adapt the 1-dimensional mechanism of (Ghazi et al., 2021) to vector aggregation by requiring that each user separately performs the scalar aggregation on each coordinate and padding the resulting messages to vectors in the natural manner, before sending the messages. Due to standard composition theorems, the privacy parameter for each coordinate must have a smaller privacy budget, so that the overall privacy loss across the d coordinates is still ε . We describe the local randomizer in Algorithm 3 and the aggregation in Algorithm 4. Our algorithms use the optimal 1-dimensional algorithm of (Ghazi et al., 2021): we let $\mathcal{R}_{\text{GKMPs}}^{(\varepsilon, \delta)}$ denote their local randomizer with parameters (ε, δ)

Algorithm 3 Local randomizer for vector aggregation

Input: $v^{(i)} \in \mathbb{S}^{d-1}$, privacy parameters (ε, δ)
Output: $S^{(i)} \subset \mathbb{R}^d$

- 1: Let $S^{(i)} = \emptyset$ and $U_K \in \mathbb{R}^{2d \times d}$ be the Kashin transformation with constant C_K
- 2: Set $u^{(i)} = \frac{\sqrt{d}}{C_K} U_K v^{(i)}$
- 3: **for** $j = 1$ to $2d$ **do**
- 4: Let $S_j = \mathcal{R}_{\text{GKMPs}}^{(\varepsilon_0, \delta_0)}(u_j^{(i)})$ where $\varepsilon_0 = \frac{\varepsilon}{2\sqrt{2d \log(2/\delta)}}$
 and $\delta_0 = \frac{\delta}{2d}$
- 5: Update $S^{(i)} = S^{(i)} \cup \{m \cdot e_j : m \in S_j\}$
- 6: **end for**
- 7: Output $S^{(i)}$

Algorithm 4 Aggregation for vector aggregation

Input: Shuffled messages $M \subset \mathbb{R}^{2d}$
Output: $\hat{v} \in \mathbb{R}^d$

- 1: Let $U_K \in \mathbb{R}^{2d \times d}$ be the Kashin transformation with constant C_K
- 2: Calculate $\hat{u} = \sum_{m \in M} m$
- 3: Output $\hat{v} = \frac{C_K}{\sqrt{d}} U_K^T \hat{u}$

We have the following result for our protocol. The proof is standard and we defer it to Appendix A.3.

Theorem 2.6. *Let $\mathcal{R} : \mathbb{S}^{d-1} \rightarrow \mathbb{R}^{2d}$ be the local randomizer in Algorithm 3 and $\mathcal{A} : (\mathbb{R}^{2d})^* \rightarrow \mathbb{R}^d$ be the aggregation in Algorithm 4. Then, \mathcal{R} is (ε, δ) -Shuffle DP randomizer, each user sends $d \cdot \left(1 + \tilde{\mathcal{O}}_\varepsilon\left(\frac{\log(1/\delta)}{\sqrt{n}}\right)\right)$ messages in expectation, and the protocol has error*

$$\text{Err}(\mathcal{A}, \mathcal{R}) \leq \mathcal{O}\left(\frac{d \log(1/\delta)}{\varepsilon^2}\right).$$

Finally, we note that it is possible to achieve this rate with $\mathcal{O}(n\varepsilon^2)$ messages using the protocols in (Chen et al., 2023): their protocols work in the shuffle model and send $\mathcal{O}(n\varepsilon^2)$ bits per users in T rounds. However, their approach can also work in a single round if the coordinates are independent (which is the case if the Kashin representation is applied). Overall, we conclude that there is a protocol for the shuffle model that requires $\mathcal{O}(\min(n\varepsilon^2, d))$ messages.

3. Single Message per User

In this section, we study private vector summation when each user is only allowed to send a single message. We first give an algorithm for this setting in Section 3.1 and then show that the algorithm is near-optimal in Section 3.2.

3.1. A Single-Message Protocol

In this section, we describe a simple protocol for private vector summation in the shuffle model that achieves near-optimal error when each user can only send a single message and ε is constant. Indeed, both the protocol and the corresponding analysis can be viewed as a generalization of (Balle et al., 2019) from the aggregation of real numbers to real-valued vectors.

The protocol first picks a granularity r so that all messages will only correspond to vectors whose coordinates are multiples of r . Each user i then randomly rounds each coordinate of their input $v^{(i)}$ to one of the two neighboring multiples of r to form a vector $\tilde{v}^{(i)}$. Each user then performs randomized response to determine whether the message $w^{(i)}$ they send is their randomly rounded input $\tilde{v}^{(i)}$ or a message selected uniform at random from the set $[r]^d$ of all possible rounded messages. The local randomizer appears in Algorithm 5.

The analyzer takes the set $\{w^{(i)}\}_{i \in [n]}$ of messages and computes their vector sum $z = \sum_{i=1}^n w^{(i)}$. It then adjusts each coordinate $j \in [d]$ of z to account for the expected noise from randomized response, so that the expectation of the corrected z_j is precisely the sum of the inputs $\sum_{i=1}^n v_j^{(i)}$. We provide the full details in Algorithm 6.

We first note that since each vector in $[r]^d$ can be encoded as a integer in $[r^d]$, then the privacy guarantees of (Balle et al., 2019) for the local randomizer holds as follows:

Lemma 3.1 (Theorem 3.1 in (Balle et al., 2019)). *The mechanism in Algorithm 6 is (ε, δ) -DP for the number of buckets $k = (r + 1)^d$ and $\gamma \geq \min\left(1, \max\left(\frac{14k}{(n-1)\varepsilon^2} \log \frac{2}{\delta}, \frac{27k}{(n-1)\varepsilon}\right)\right)$.*

We now upper bound the mean squared error of Algorithm 6.

Theorem 1.3. *For any $\varepsilon \in (0, 1)$, $\delta \in (0, 1)$, and $d, n \in \mathbb{N}$, there exists an (ε, δ) -DP protocol in the one-message shuffle*

Algorithm 5 Local randomizer for single message per user

Input: $v^{(i)} \in \mathbb{S}^{d-1}$, parameters r, c, d, n
Output: $w^{(i)} \in \{0, 1, \dots, r\}^d$

```

1:  $\gamma \leftarrow \frac{c(r+1)}{n}$ 
2: for  $j = 1$  to  $j = d$  do
3:    $\tilde{v}^{(i)}_j \sim \left\lfloor rv_j^{(i)} \right\rfloor + \text{Ber} \left( rv_j^{(i)} - \left\lfloor rv_j^{(i)} \right\rfloor \right)$ 
4: end for
5: Sample  $b \sim \text{Ber}(\gamma)$ 
6: if  $b = 0$  then
7:    $w^{(i)} \leftarrow \tilde{v}^{(i)}$ 
8: else
9:    $w^{(i)} \sim \text{Unif}([r]^d)$ 
10: end if

```

Algorithm 6 Aggregation for bucket-based randomized response

Input: $w^{(i)} \in \{0, 1, \dots, r\}^d$ for $i \in [n]$ and c from Algorithm 5

Output: $\tilde{v} \in [0, n]^d$

```

1:  $z \leftarrow \frac{1}{r} \sum_{i=1}^n w^{(i)}$ 
2: for  $j = 1$  to  $j = d$  do
3:    $\tilde{v}_j \leftarrow \frac{2z_j - c(r+1)}{2 - 2\gamma} = \left( z_j - \frac{c(r+1)}{2} \right) / (1 - \gamma)$ 
4: end for
5: Return  $\tilde{v}$ 

```

model with mean squared error $\mathcal{O}_\delta(dn^{d/(d+2)}\varepsilon^{-4/(d+2)})$.

Proof. Consider Algorithm 6. The mechanism is (ε, δ) -private by the choice of

$$\gamma \geq \min \left(1, \max \left(\frac{14k}{(n-1)\varepsilon^2} \log \frac{2}{\delta}, \frac{27k}{(n-1)\varepsilon} \right) \right)$$

and Lemma 3.1.

The mean squared error is at most

$$\sup_{\{v^{(i)}\}} \mathbb{E} \left[\sum_{j=1}^d (\tilde{v}_j - v_j)^2 \right] = \sup_{\{v^{(i)}\}} \mathbb{E} \left[\sum_{j=1}^d \left(\tilde{v}_j - \sum_{i=1}^n v_j^{(i)} \right)^2 \right]$$

For a real number x , let $F(x) = \frac{x - c(r+1)/2}{1 - c(r+1)/n}$, so that F is the debiasing function applied coordinate-wise to z .

$$\begin{aligned} & \sup_{\{v^{(i)}\}} \mathbb{E} \left[\sum_{j=1}^d (\tilde{v}_j - v_j)^2 \right] \\ &= \sup_{\{v^{(i)}\}} \mathbb{E} \left[\sum_{j=1}^d \left(F(z_j) - \sum_{i=1}^n v_j^{(i)} \right)^2 \right] \\ &= \sup_{\{v^{(i)}\}} \mathbb{E} \left[\sum_{j=1}^d \left(F \left(\frac{1}{r} \sum_{i=1}^n w_j^{(i)} \right) - \sum_{i=1}^n v_j^{(i)} \right)^2 \right]. \end{aligned}$$

Note that by construction $\mathbb{E}[F(z_j)] = \sum_{i=1}^n v_j^{(i)}$, for all $j \in [d]$. Thus the cross terms cancel, so that we further have

$$\begin{aligned} & \sup_{\{v^{(i)}\}} \mathbb{E} \left[\sum_{j=1}^d (\tilde{v}_j - v_j)^2 \right] \\ &= \sup_{\{v^{(i)}\}} \mathbb{E} \left[\sum_{j=1}^d \sum_{i=1}^n \left(F \left(\frac{w_j^{(i)}}{r} \right) - \sum_{i=1}^n v_j^{(i)} \right)^2 \right] \\ &= \sup_{\{v^{(i)}\}} \sum_{j=1}^d \sum_{i=1}^n \mathbb{V} \left[F \left(\frac{w_j^{(i)}}{r} \right) \right], \end{aligned}$$

where we use \mathbb{V} to denote the variance. Note that after debiasing, the γ fraction of the coordinates that were randomly generated from the uniform distribution, due to $b \sim \text{Ber}(\gamma)$, do not contribute variance. Hence the mean-squared error is at most

$$\begin{aligned} \sup_{\{v^{(i)}\}} \mathbb{E} \left[\sum_{j=1}^d (\tilde{v}_j - v_j)^2 \right] &= \frac{nd}{(1-\gamma)^2} \sup_{x_1^{(1)}} \mathbb{V} \left[\frac{w_1^{(1)}}{r} \right] \\ &\leq \frac{nd}{(1-\gamma)^2} \left(\frac{1-\gamma}{4r^2} + \frac{\gamma}{2} \right). \end{aligned}$$

Recall that we set $\gamma = \frac{c(k+1)}{n}$ for a parameter c , which we require to guarantee

$$\gamma \geq \min \left(1, \max \left(\frac{14k}{(n-1)\varepsilon^2} \log \frac{2}{\delta}, \frac{27k}{(n-1)\varepsilon} \right) \right),$$

to satisfy privacy. Then we have

$$\begin{aligned} \sup_{\{v^{(i)}\}} \mathbb{E} \left[\sum_{j=1}^d (\tilde{v}_j - v_j)^2 \right] &\leq \frac{nd}{(1-\gamma)^2} \left(\frac{1}{4r^2} + \frac{c(k+1)}{2n} \right) \\ &\leq \frac{nd}{(1-\gamma)^2} \left(\frac{1}{4r^2} + \frac{c((r+1)^d + 1)}{2n} \right). \end{aligned}$$

By setting $c(r+1)^{d+2} = \mathcal{O}(n)$ for $c = \mathcal{O}(\frac{1}{\varepsilon^2} \log \frac{1}{\delta})$, we have that the above quantity is minimized at $r = (\mathcal{O}_\delta(\frac{n}{c}))^{1/(d+2)}$. Thus since $c = \mathcal{O}(\frac{1}{\varepsilon^2} \log \frac{1}{\delta})$, then the mean squared error is at most $\mathcal{O}_\delta(dn^{d/(d+2)}\varepsilon^{-4/(d+2)})$. \square

3.2. Lower Bound

In this section, we show that our protocol in Section 3.1 is near-optimal by proving that for any $\varepsilon = \mathcal{O}(1)$, the mean squared error of any protocol that gives ε -DP in the shuffle model in which each user sends a single message is $\Omega(dn^{d/(d+2)})$. The main intuition is that we can partition the space into blocks of size length $\frac{1}{r}$, so that there are r^d hypercubes in total. Although r is a parameter that can be chosen at the protocol's discretion, there are two sources

of error for any private protocol that result in two opposing tensions on the value of r .

The first source of error is that due to the privacy guarantees, the output distribution for an input $v^{(i)}$ to a player i may overlap with the output distribution for any input in $[r]^d$. In this case, the message may be decoded to some other vector with large distance from $v^{(i)}$, resulting in large mean squared error. In particular, larger values of r force the output of the local randomizer to have less signal about the true block containing the input $v^{(i)}$, since the output distribution must intersect with that of more possible inputs. This is formalized in [Lemma B.3](#).

The second source of error is that any vector inside a block may incur error from the message representing the block, due to the partition of the space. In particular, the message may be decoded correctly for the block, but the set of all vectors within the block has large diameter, and so the resulting mean squared error is large. Specifically, smaller values of r result in blocks with larger diameter, which again force the output of the local randomizer to have less signal about the true input $v^{(i)}$ within each block. This is formalized in [Lemma B.4](#). The resulting lower bound then follows from optimizing r with respect to the two possible sources of error, resulting in the following theorem.

Theorem 1.4. *Let \mathcal{P} be an (ε, δ) -DP protocol for vector aggregation on the unit ball \mathbb{B}_2^{d-1} in the one-message shuffle model with $\delta < \frac{1}{2}$. Then the mean squared error of \mathcal{P} satisfies $\text{MSE}(\mathcal{P}) = \Omega(dn^{d/(d+2)})$.*

The proof of this result is technical and we defer it to [Appendix B](#).

4. Robustness to Malicious Users

We first observe that our multi-message protocol is not robust against malicious users in the single-shuffle setting, in the sense that a single malicious user can additively incur much larger than constant mean squared error, even though their input vector has at most unit length. In fact, each user can incur up to $\Omega\left(\frac{k}{\log^2(nd)}\right)$ additive mean squared error.

Theorem 4.1. *Let $\varepsilon = \mathcal{O}(1)$ and $\delta < \frac{1}{nd}$. Then any (ε, δ) -shuffle DP mechanism for vector summation that takes the sum of the messages across n players with k malicious users has additive error $\Omega\left(\frac{kd}{\log^2(nd)}\right)$.*

[Theorem 1.5](#) then follows from a simple power mean inequality. On the other hand, we observe that [Algorithm 4](#) is robust against malicious users in the setting where a separate shuffler is responsible for the messages corresponding to each coordinate of a user.

Lemma 4.2. *Suppose that a separate shuffler handles the messages for each coordinate from all users in [Algorithm 4](#).*

Then the mean squared error induced by k malicious users is at most $\mathcal{O}(k)$.

Impact Statement

This paper presents work whose goal is to advance the field of Private Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

Abadi, M., Chu, A., Goodfellow, I. J., McMahan, H. B., Mironov, I., Talwar, K., and Zhang, L. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pp. 308–318, 2016. [1](#)

Acharya, J. and Sun, Z. Communication complexity in locally private distribution estimation and heavy hitters. In *Proceedings of the 36th International Conference on Machine Learning, ICML*, pp. 51–60, 2019. [4](#)

Asi, H., Feldman, V., and Talwar, K. Optimal algorithms for mean estimation under local differential privacy. In *International Conference on Machine Learning, ICML, USA*, pp. 1046–1056, 2022. [4, 17](#)

Asi, H., Feldman, V., Nelson, J., Nguyen, H., and Talwar, K. Fast optimal locally private mean estimation via random projections. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. [4](#)

Balle, B., Bell, J., Gascón, A., and Nissim, K. The privacy blanket of the shuffle model. In *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Proceedings, Part II*, pp. 638–667, 2019. [2, 3, 7, 23, 24, 25](#)

Balle, B., Bell, J., Gascón, A., and Nissim, K. Private summation in the multi-message shuffle model. In *CCS '20: 2020 ACM SIGSAC Conference on Computer and Communications Security*, pp. 657–676, 2020. [2, 3](#)

Bassily, R. and Smith, A. D. Local, private, efficient protocols for succinct histograms. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC*, pp. 127–135, 2015. [4](#)

Bassily, R., Smith, A. D., and Thakurta, A. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pp. 464–473, 2014. [3](#)

Beimel, A., Nissim, K., and Omri, E. Distributed private data analysis: Simultaneously solving how and what. In *Advances in Cryptology - CRYPTO 2008, 28th Annual*

International Cryptology Conference. Proceedings, pp. 451–468, 2008. 1

Beimel, A., Haitner, I., Nissim, K., and Stemmer, U. On the round complexity of the shuffle model. In *Theory of Cryptography Conference*, pp. 683–712. Springer, 2020. 4

Bhowmick, A., Duchi, J. C., Freudiger, J., Kapoor, G., and Rogers, R. Protection against reconstruction and its applications in private federated learning. *CoRR*, abs/1812.00984, 2018. 4

Bittau, A., Erlingsson, Ú., Maniatis, P., Mironov, I., Raghunathan, A., Lie, D., Rudominer, M., Kode, U., Tinnés, J., and Seefeld, B. Prochlo: Strong privacy for analytics in the crowd. In *Proceedings of the 26th Symposium on Operating Systems Principles*, pp. 441–459, 2017. 1, 2

Chan, T. H., Shi, E., and Song, D. Optimal lower bound for differentially private multi-party aggregation. In *Algorithms - ESA 2012 - 20th Annual European Symposium. Proceedings*, pp. 277–288, 2012. 1

Chaudhuri, K., Monteleoni, C., and Sarwate, A. D. Differentially private empirical risk minimization. *J. Mach. Learn. Res.*, 12:1069–1109, 2011. 1

Chen, L., Ghazi, B., Kumar, R., and Manurangsi, P. On distributed differential privacy and counting distinct elements. *arXiv:2009.09604 [cs.CR]*, 2020a. 4

Chen, W., Kairouz, P., and Özgür, A. Breaking the communication-privacy-accuracy trilemma. In *Proceedings of the 33rd Annual Conference on Advances in Neural Information Processing Systems (NeurIPS)*, 2020b. 4

Chen, W.-N., Song, D., Ozgur, A., and Kairouz, P. Privacy amplification via compression: Achieving the optimal privacy-accuracy-communication trade-off in distributed mean estimation. *arXiv:2304.01541 [stat.ML]*, 2023. 7

Cheu, A. and Ullman, J. The limits of pan privacy and shuffle privacy for learning and estimation. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pp. 1081–1094, 2021. 4

Cheu, A., Smith, A. D., Ullman, J. R., Zeber, D., and Zhilyaev, M. Distributed differential privacy via shuffling. In *Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Proceedings, Part I*, pp. 375–403, 2019. 1

Cheu, A., Smith, A., and Ullman, J. Manipulation attacks in local differential privacy. *Journal of Privacy and Confidentiality*, 11(1), Feb. 2021. doi: 10.29012/jpc.754. URL <https://journalprivacyconfidentiality.org/index.php/jpc/article/view/754>. 2

Cheu, A., Joseph, M., Mao, J., and Peng, B. Shuffle private stochastic convex optimization. In *The Tenth International Conference on Learning Representations, ICLR*, 2022. 1

Corrigan-Gibbs, H. and Boneh, D. Prio: Private, robust, and scalable computation of aggregate statistics. In Akella, A. and Howell, J. (eds.), *14th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2017, Boston, MA, USA, March 27-29, 2017*, pp. 259–282. USENIX Association, 2017. 3

Davidson, A., Goldberg, I., Sullivan, N., Tankersley, G., and Valsorda, F. Privacy pass: Bypassing internet challenges anonymously. *Proc. Priv. Enhancing Technol.*, 2018(3): 164–180, 2018. 2

Duchi, J. C. Introductory lectures on stochastic convex optimization. In *The Mathematics of Data*, IAS/Park City Mathematics Series. American Mathematical Society, 2018. 14

Duchi, J. C. and Rogers, R. Lower bounds for locally private estimation via communication complexity. In *Conference on Learning Theory, COLT*, pp. 1161–1191, 2019. 4

Duchi, J. C., Wainwright, M. J., and Jordan, M. I. Minimax optimal procedures for locally private estimation. *CoRR*, abs/1604.02390, 2016. 4

Dwork, C. and Roth, A. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4):211–407, 2014. 4, 22

Dwork, C., McSherry, F., Nissim, K., and Smith, A. D. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography, Third Theory of Cryptography Conference, TCC, Proceedings*, pp. 265–284, 2006. 1, 4

Dwork, C., McSherry, F., Nissim, K., and Smith, A. D. Calibrating noise to sensitivity in private data analysis. *J. Priv. Confidentiality*, 7(3):17–51, 2016. 1

Erlingsson, U., Feldman, V., Mironov, I., Raghunathan, A., Talwar, K., and Thakurta, A. Amplification by shuffling: From local to central differential privacy via anonymity. In *Proceedings of the Thirtieth ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2019. 1

Fanti, G., Pihur, V., and Erlingsson, Ú. Building a RAP-POR with the unknown: Privacy-preserving learning of associations and data dictionaries. *Proc. Priv. Enhancing Technol.*, 2016(3):41–61, 2016. 4

Feldman, V. and Talwar, K. Lossless compression of efficient private local randomizers. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, pp. 3208–3219. PMLR, 2021. 4

Feldman, V., Guzmán, C., and Vempala, S. S. Statistical query algorithms for mean vector estimation and stochastic convex optimization. *Math. Oper. Res.*, 46(3):912–945, 2021. 5

Ghazi, B., Manurangsi, P., Pagh, R., and Velingker, A. Private aggregation from fewer anonymous messages. In *Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Proceedings, Part II*, pp. 798–827, 2020. 2, 3

Ghazi, B., Kumar, R., Manurangsi, P., Pagh, R., and Sinha, A. Differentially private aggregation in the shuffle model: Almost central accuracy in almost a single message. In *Proceedings of the 38th International Conference on Machine Learning, ICML*, pp. 3692–3701, 2021. 2, 3, 7, 22

Hendrickson, S., Iyengar, J., Pauly, T., Valdez, S., and Wood, C. A. Rate-Limited Token Issuance Protocol. Internet-Draft draft-ietf-privacypass-rate-limit-tokens-01, Internet Engineering Task Force, March 2023. URL <https://datatracker.ietf.org/doc/draft-ietf-privacypass-rate-limit-tokens/01/>. Work in Progress. 2

Ishai, Y., Kushilevitz, E., Ostrovsky, R., and Sahai, A. Cryptography from anonymity. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), USA, Proceedings*, pp. 239–248. IEEE Computer Society, 2006. 3

Kasiviswanathan, S. P., Lee, H. K., Nissim, K., Raskhodnikova, S., and Smith, A. D. What can we learn privately? *SIAM J. Comput.*, 40(3):793–826, 2011. 1

Lyubarskii, Y. and Vershynin, R. Uncertainty principles and vector quantization. *IEEE Trans. Inf. Theory*, 56(7): 3491–3501, 2010. 4

McMahan, B., Moore, E., Ramage, D., Hampson, S., and y Arcas, B. A. Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS*, pp. 1273–1282, 2017. 1

Nguyễn, T. T., Xiao, X., Yang, Y., Hui, S. C., Shin, H., and Shin, J. Collecting and analyzing data from smart device users with local differential privacy. *CoRR*, abs/1606.05053, 2016. 1

Rothblum, G. N., Omri, E., Chen, J., and Talwar, K. Pine: Efficient norm-bound verification for secret-shared vectors, 2023. 3

Scott, M., Cormode, G., and Maple, C. Applying the shuffle model of differential privacy to vector aggregation. In Pirk, H. and Heinis, T. (eds.), *Proceedings of the The British International Conference on Databases*, volume 3163 of *CEUR Workshop Proceedings*, pp. 50–59, 2021. 2

Scott, M., Cormode, G., and Maple, C. Aggregation and transformation of vector-valued messages in the shuffle model of differential privacy. *IEEE Trans. Inf. Forensics Secur.*, 17:612–627, 2022. 2

Shokri, R. and Shmatikov, V. Privacy-preserving deep learning. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pp. 1310–1321. ACM, 2015. 1

Sun, C., Fu, Y., Zhou, J., Gao, H., et al. Personalized privacy-preserving frequent itemset mining using randomized response. *The Scientific World Journal*, 2014, 2014. 1

Talwar, K. Differential secrecy for distributed data and applications to robust differentially secure vector summation. In Celis, L. E. (ed.), *3rd Symposium on Foundations of Responsible Computing, FORC 2022, June 6–8, 2022, Cambridge, MA, USA*, volume 218 of *LIPICS*, pp. 7:1–7:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. 2

Thomson, M. and Wood, C. A. Oblivious HTTP. Internet-Draft draft-ietf-ohai-ohttp-08, Internet Engineering Task Force, March 2023. URL <https://ietf-wg-ohai.github.io/oblivious-http/draft-ietf-ohai-ohttp.html>. Work in Progress. 2

Ye, M. and Barg, A. Optimal schemes for discrete distribution estimation under locally differential privacy. *IEEE Trans. Inf. Theory*, 64(8):5662–5676, 2018. 4

Zhou, M., Wang, T., Chan, T. H., Fanti, G., and Shi, E. Locally differentially private sparse vector aggregation. In *43rd IEEE Symposium on Security and Privacy, SP*, pp. 422–439, 2022. 4

A. Missing Proofs from Section 2

A.1. Missing Proofs from Section 2.1.1

In this section, we provide the missing proof for the lower bound for the setting of summation protocols (Section 2.1.1).

A.1.1. PROOF OF LEMMA 2.2

Lemma 2.2. *Let $\varepsilon \leq 1$, $\mathcal{R} : \mathbb{S}^{d-1} \rightarrow \mathcal{Z}^k$ be (ε, δ) -Shuffle DP. There exists an (ε, δ) -Shuffle DP randomizer $\hat{\mathcal{R}} : \mathbb{S}^{d-1} \rightarrow \mathcal{Z}^k$ such that*

$$(1) \quad \text{Err}(\mathcal{A}^+, \hat{\mathcal{R}}) \leq \text{Err}(\mathcal{A}^+, \mathcal{R})$$

$$(2) \quad (\text{Symmetry}) \text{ For all } u, v \in \mathbb{S}^{d-1},$$

$$\mathbb{E} \left[\left\| \hat{\mathcal{R}}^+(v) - v \right\|_2^2 \right] = \mathbb{E} \left[\left\| \hat{\mathcal{R}}^+(u) - u \right\|_2^2 \right]$$

Proof. The new randomizer $\hat{\mathcal{R}}$ works as follows: first, it samples a rotation matrix $U \in \mathbb{R}^{d \times d}$ (known public randomness) such that $U^T U = I$, then sets

$$\hat{\mathcal{R}}(v) = U^T \mathcal{R}(Uv),$$

where $U^T \mathcal{R}(Uv)$ denotes multiplying each message in $\mathcal{R}(Uv)$ by U^T .

To prove privacy, we have to prove that $\Pi(U^T \mathcal{R}(Uv_1), U^T \mathcal{R}(Uv_2), \dots, U^T \mathcal{R}(Uv_n))$ is (ε, δ) -DP. As U is known, it is sufficient to prove that $\Pi(\mathcal{R}(Uv_1), \mathcal{R}(Uv_2), \dots, \mathcal{R}(Uv_n))$ is (ε, δ) -DP. This follows directly from the fact that $\Pi(\mathcal{R}(v_1), \mathcal{R}(v_2), \dots, \mathcal{R}(v_n))$ is (ε, δ) -DP, and that the hamming distance between $X = (v_1, \dots, v_n)$ and $X' = (v'_1, \dots, v'_n)$ is the same as the hamming distance between $X_U = (Uv_1, \dots, Uv_n)$ and $X'_U = (Uv'_1, \dots, Uv'_n)$.

For utility, we have

$$\begin{aligned} \text{Err}(\mathcal{A}^+, \hat{\mathcal{R}}) &= \sup_{v_1, \dots, v_n} \mathbb{E} \left[\left\| \mathcal{A}^+(\Pi(\hat{\mathcal{R}}(v_1), \hat{\mathcal{R}}(v_2), \dots, \hat{\mathcal{R}}(v_n))) - \sum_{i=1}^n v_i \right\|_2^2 \right] \\ &= \sup_{v_1, \dots, v_n} \mathbb{E} \left[\left\| \sum_{i=1}^n \hat{\mathcal{R}}^+(v_i) - v_i \right\|_2^2 \right] \\ &= \sup_{v_1, \dots, v_n} \mathbb{E} \left[\left\| \sum_{i=1}^n U^T \mathcal{R}(Uv_i) - v_i \right\|_2^2 \right] \\ &= \sup_{v_1, \dots, v_n} \mathbb{E} \left[\left\| U^T \sum_{i=1}^n (\mathcal{R}(Uv_i) - Uv_i) \right\|_2^2 \right] \\ &= \sup_{v_1, \dots, v_n} \mathbb{E} \left[\left\| \sum_{i=1}^n (\mathcal{R}(Uv_i) - Uv_i) \right\|_2^2 \right] \\ &= \text{Err}(\mathcal{A}, \mathcal{R}). \end{aligned}$$

For the third claim, note that $\hat{\mathcal{R}}(-v) = U^T \mathcal{R}(-Uv)$. As U and $-U$ has the same distribution, we can also write $\hat{\mathcal{R}}(-v) = -U^T \mathcal{R}(Uv)$ which is the same as the distribution of $-\hat{\mathcal{R}}(v)$.

For the final claim, note that

$$\begin{aligned}\mathbb{E} \left[\left\| \hat{\mathcal{R}}^+(v) - v \right\|_2^2 \right] &= \mathbb{E} \left[\left\| U^T \mathcal{R}(Uv) - v \right\|_2^2 \right] \\ &= \mathbb{E} \left[\left\| U^T (\mathcal{R}(Uv) - Uv) \right\|_2^2 \right] \\ &= \mathbb{E} \left[\left\| (\mathcal{R}(Uv) - Uv) \right\|_2^2 \right].\end{aligned}$$

The claim follows as Uv_1 and Uv_2 have the same distribution for any v_1 and v_2 in the unit ball. \square

Proposition 2.3. *Let $v_1, \dots, v_n \in \mathbb{S}^{d-1}$, $\mathcal{R} : \mathbb{S}^{d-1} \rightarrow \mathcal{Z}^k$ be randomizer that satisfies the symmetry condition of Lemma 2.2. For an input set $W = \Pi(\mathcal{R}(v_1), \dots, \mathcal{R}(v_n))$, Algorithm 1 outputs a set $S \subset \mathbb{R}^d$ of size $\binom{nk}{k}$ such that*

$$\mathbb{E} [\text{dist}(v_1, S)] = \mathbb{E} \left[\min_{u \in S} \|v_1 - u\|_2^2 \right] \leq \frac{\text{Err}(\mathcal{A}^+, \mathcal{R})}{n},$$

where the expectation is over the randomness of the algorithm.

Proof. Let $\Delta(v) = E[\mathcal{R}^+(v) - v]$ be the bias of \mathcal{R}^+ over v . Note that the error of the protocol over dataset (u_1, \dots, u_n) is

$$\begin{aligned}E \left[\left\| \mathcal{A}^+(\Pi(\mathcal{R}(u_1), \mathcal{R}(u_2), \dots, \mathcal{R}(u_n))) - \sum_{i=1}^n u_i \right\|_2^2 \right] \\ = E \left[\left\| \sum_{i=1}^n \mathcal{R}^+(u_i) - u_i \right\|_2^2 \right] \\ = \sum_{i=1}^n E \left[\left\| \mathcal{R}^+(u_i) - u_i \right\|_2^2 \right] + \sum_{i \neq j \in [n]} E \left[\mathcal{R}^+(u_i) - u_i \right]^T E \left[\mathcal{R}^+(u_j) - u_j \right] \\ = \sum_{i=1}^n E \left[\left\| \mathcal{R}^+(u_i) - u_i \right\|_2^2 \right] + \sum_{i \neq j \in [n]} \Delta(u_i)^T \Delta(u_j).\end{aligned}$$

For input dataset $X = (u, u, \dots, u)$, this implies

$$\begin{aligned}E \left[\left\| \mathcal{A}^+(\Pi(\mathcal{R}(u), \mathcal{R}(u), \dots, \mathcal{R}(u))) - nu \right\|_2^2 \right] \\ = nE \left[\left\| \mathcal{R}^+(u) - u \right\|_2^2 \right] + \binom{n}{2} \|\Delta(u)\|_2^2 \\ \geq nE \left[\left\| \mathcal{R}^+(u) - u \right\|_2^2 \right].\end{aligned}$$

As \mathcal{R} satisfies the symmetry assumption that $E[\|\mathcal{R}^+(v) - v\|_2^2] = E[\|\mathcal{R}^+(u) - u\|_2^2]$ for all $u, v \in \mathbb{S}^{d-1}$, and since the error is bounded by d/ε^2 , we have that

$$E \left[\left\| \mathcal{R}^+(v_1) - v_1 \right\|_2^2 \right] \leq \frac{d}{n\varepsilon^2}.$$

Finally, note that $\mathcal{R}^+(v_1) \in S$ as the attack of Algorithm 1 iterates over all possible subsets of size k and adds their sum to S . Hence, there exists t such that $W_t = \mathcal{R}(v_1)$, in which case the algorithm will add $\mathcal{R}^+(v_1)$ to S . \square

Given the previous attack, we are now ready to prove our lower bound.

A.1.2. PROOF OF THEOREM 2.1

Theorem 2.1. Let $\varepsilon, \delta \leq 1$ and $\mathcal{R} : \mathbb{S}^{d-1} \rightarrow \mathcal{Z}^k$ be an (ε, δ) -Shuffle DP randomizer. If $\text{Err}(\mathcal{A}^+, \mathcal{R}) \leq \mathcal{O}(d/\varepsilon^2)$ then $k \geq \Omega\left(\frac{\min(n\varepsilon^2, d)}{\log n}\right)$.

Proof. Let $\text{Err}(\mathcal{A}, \mathcal{R}) \leq C \cdot d/\varepsilon^2$ for some universal constant $1 \leq C < \infty$. Based on Lemma 2.2, we can assume that the randomizer R satisfies the symmetry property:

$$\mathbb{E} \left[\left\| \mathcal{R}^+(v) - v \right\|_2^2 \right] = \mathbb{E} \left[\left\| R^+(u) - u \right\|_2^2 \right], \quad \text{for all } u, v \in \mathbb{S}^{d-1}.$$

First, we prove the lower bounds for $d \leq n\varepsilon^2/100C$. Let $P = \{v_1, v_2, \dots, v_M\}$ be a ρ -packing of the unit ball such that $M = 2^{d \log(1/\rho)}$ (the existence of such packing is standard in the literature (Duchi, 2018)). We will prove the lower bounds by analyzing the algorithm over the following M datasets:

$$X_i = (v_i, v_1, \dots, v_1).$$

Let S_i be the output of the reconstruction attack (Algorithm 1) over the input $\Pi(\mathcal{R}(v_i), \mathcal{R}(v_1), \dots, \mathcal{R}(v_1))$, and let O_i be the projection of S_i to the packing P ; that is, $O_i = \{\text{Proj}_P(v) : v \in S_i\}$.

Proposition 2.3 states that $\mathbb{E}[\text{dist}(v_i, S_i)] \leq \frac{Cd}{n\varepsilon^2} \leq 1/100$, hence we get that

$$\mathbf{Pr}[v_i \in O_i] \geq \mathbf{Pr}[\text{dist}(v_i, S_i) < \rho] \geq 9/10,$$

where the first inequality follows as P is ρ -packing, and the second inequality follows from markov inequality.

On the other hand, note that

$$\begin{aligned} \sum_{i=1}^M \mathbf{Pr}[v_i \in O_1] &= \sum_{i=1}^M \mathbb{E}[1\{v_i \in O_1\}] \\ &= \mathbb{E} \left[\sum_{i=1}^M 1\{v_i \in O_1\} \right] \\ &\leq \mathbb{E}[|O_1|] \leq \binom{nk}{k}. \end{aligned}$$

Hence there exists an $1 \leq i \leq M$ such that

$$\mathbf{Pr}[v_i \in O_1] \leq \frac{\binom{nk}{k}}{M}.$$

As the protocol is (ε, δ) -DP, we also have

$$\begin{aligned} \mathbf{Pr}[v_i \in O_1] &\geq \mathbf{Pr}[v_i \in O_i] e^{-\varepsilon} - \delta \\ &\geq \frac{9}{10e} \geq 1/6. \end{aligned}$$

Combining these together, and given that $M \geq 2^d$ for $\rho = 1/10$, we have that

$$2^d \leq 6 \binom{nk}{k} \leq 6(en)^k.$$

This implies that $k \geq \Omega(d/\log(n))$ whenever $d \leq n\varepsilon^2/100C$.

Now we prove the lower bound for $d \geq n\varepsilon^2/100$. The proof builds on the following proposition which states that we can convert an optimal protocol for d -dimensional inputs into an optimal protocol for d' -dimensional inputs where $d' = n\varepsilon^2/200$ with the same number of messages. We defer the proof to Appendix A.1.3.

Proposition A.1. Let $d' = n\varepsilon^2/200C \geq 1$ and $d \geq 2d'$. Let $\mathcal{R} : \mathbb{S}^{d-1} \rightarrow \mathcal{Z}^k$ be an (ε, δ) -shuffle DP protocol with error $\text{Err}(\mathcal{A}^+, \mathcal{R}) \leq \mathcal{O}(d/\varepsilon^2)$. There exists $\mathcal{R}' : \mathbb{B}^{d'-1} \rightarrow \mathcal{Z}^k$ that is (ε, δ) -shuffle DP such that $\text{Err}(\mathcal{A}^+, \mathcal{R}') \leq \mathcal{O}(d'/\varepsilon^2)$.

Now, let \mathcal{A}^+ and $\mathcal{R} : \mathbb{B}^{d-1} \rightarrow \mathcal{Z}^k$ be a protocol that obtains error $\text{Err}(\mathcal{A}^+, \mathcal{R}) \leq \mathcal{O}(d/\varepsilon^2)$ using k messages. Proposition A.1 implies that there is a randomizer $\mathcal{R}' : \mathbb{B}^{d'-1} \rightarrow \mathcal{Z}^k$ such that $\text{Err}(\mathcal{A}, \mathcal{R}') \leq \mathcal{O}(d'/\varepsilon^2)$ for $d' = n\varepsilon^2/200C$. As $d' \leq n\varepsilon^2/100C$, this shows that $k \geq \Omega(d'/\log(n)) = \Omega(n\varepsilon^2/\log(n))$. \square

A.1.3. PROOF OF PROPOSITION A.1

To prove Proposition A.1, we need the following lemma which shows that we can convert any summation protocol into another one where the error is split evenly across coordinates.

We use the following notation. For a permutation $\pi : [d] \rightarrow [d]$ and a vector $v \in \mathbb{R}^d$, we let $\hat{v} = v(\pi)$ denote the shuffling of the coordinates of v based on π , that is $\hat{v}_j = v_{\pi(j)}$.

Lemma A.2. If $\mathcal{R} : \mathbb{S}^{d-1} \rightarrow \mathcal{Z}^k$ is (ε, δ) -shuffle DP then there exists $\hat{\mathcal{R}} : \{\frac{-1}{\sqrt{d}}, \frac{1}{\sqrt{d}}\}^d \rightarrow \mathcal{Z}^k$ that is (ε, δ) -shuffle DP and for $j \in [d]$ and $v_1, \dots, v_n \in \{\frac{-1}{\sqrt{d}}, \frac{1}{\sqrt{d}}\}^d$,

$$\mathbb{E} \left[\left| \left(\sum_{i=1}^n \hat{\mathcal{R}}^+(v_i) - \sum_{i=1}^n v_i \right)_j \right|^2 \right] \leq \frac{\text{Err}(\mathcal{A}^+, \mathcal{R})}{d}.$$

Proof. $\hat{\mathcal{R}}$ will use shared public randomness to shuffle the coordinates of each vector and flip the signs of each coordinate. This will ensure that all coordinates will have the same marginal distribution for their error.

More precisely, let $\pi : [d] \rightarrow [d]$ be a random permutation of the coordinates picked uniformly at random, and let $s_1, \dots, s_d \sim \text{Ber}(1/2)$. Our new randomizer $\hat{\mathcal{R}}$ over input v will first transform the input vector v into \hat{v} where

$$\hat{v} = s \cdot v(\pi),$$

where the multiplication is element-wise. Then, we run $\mathcal{R}(\hat{v})$ to get messages $\hat{m}_1, \dots, \hat{m}_k$. For each of these messages, we apply the inverse transformation, and output m_1, \dots, m_k where

$$m_i = s \cdot \hat{m}_i(\pi^{-1}).$$

Note that

$$\begin{aligned} \left| \left(\sum_{i=1}^n \hat{\mathcal{R}}^+(v_i) - \sum_{i=1}^n v_i \right)_j \right|^2 &= \left| \sum_{i=1}^n m_{i,j} - v_j \right|^2 \\ &= \left| s_j \sum_{i=1}^n \hat{m}_{i,\pi^{-1}(j)} - \hat{v}_{\pi^{-1}(j)} \right|^2 \\ &= \left| \sum_{i=1}^n \hat{m}_{i,\pi^{-1}(j)} - \hat{v}_{\pi^{-1}(j)} \right|^2 \\ &= \left(\sum_{i=1}^n \mathcal{R}^+(\hat{v}_i) - \sum_{i=1}^n \hat{v}_i \right)_{\pi^{-1}(j)}. \end{aligned}$$

As $\hat{v}_1, \dots, \hat{v}_n$ are uniformly random vectors from $\{-1, +1\}^d/\sqrt{d}$, we get that $\left(\sum_{i=1}^n \hat{\mathcal{R}}^+(v_i) - \sum_{i=1}^n v_i \right)_j^2$ have the same

distribution for all $j \in [d]$. The claim now follows since

$$\begin{aligned}
 \mathbb{E} \left[\sum_{j=1}^d \left(\sum_{i=1}^n \hat{\mathcal{R}}^+(v_i) - \sum_{i=1}^n v_i \right)_j^2 \right] &= \mathbb{E} \left[\sum_{j=1}^d \left(\sum_{i=1}^n \mathcal{R}^+(\hat{v}_i) - \sum_{i=1}^n \hat{v}_i \right)_{\pi^{-1}(j)}^2 \right] \\
 &= \mathbb{E} \left[\left\| \sum_{i=1}^n \mathcal{R}^+(\hat{v}_i) - \sum_{i=1}^n \hat{v}_i \right\|_2^2 \right] \\
 &\leq \text{Err}(\mathcal{A}^+, \mathcal{R}).
 \end{aligned}$$

□

We are now ready to prove Proposition A.1.

Proof. (of Proposition A.1) \mathcal{R}' will work as follows for a d' -dimensional input v' : first, apply Kashin representation $U \in \mathbb{R}^{2d' \times d'}$ to get $w = Uv' \in \mathbb{R}^{2d'}$ such that $\|w\|_\infty \leq 2/\sqrt{d'}$. Then, we convert w into a binary vector u by setting for all $i \in [2d']$

$$u_i = \begin{cases} 2\text{sign}(w_i)/\sqrt{d'} & \text{with probability } \frac{w_i\sqrt{d'}+2}{4} \\ -2\text{sign}(w_i)/\sqrt{d'} & \text{with probability } \frac{-w_i\sqrt{d'}+2}{4} \end{cases}$$

Note that $E[u_i] = w_i$ and that $E[(w_i - u_i)^2] \leq 4/d'$ since $|u_i| \leq 2/\sqrt{d'}$.

Now, let $\hat{\mathcal{R}}$ be the randomizer guaranteed from Lemma A.2 for the randomizer \mathcal{R} . Our local randomizer \mathcal{R}' will do the following: it constructs $v \in \mathbb{R}^d$ by setting $v = (u, 0, \dots, 0)$ then applies $\hat{\mathcal{R}}$ over v to generate k messages m_1, \dots, m_k . Finally, it truncates the messages to the first $2d'$ coordinates and applies the inverse Kashin transformation to the messages, that is, sends $U^T m_1[1 : 2d'], \dots, U^T m_k[1 : 2d']$

Privacy of \mathcal{R}' follows immediately from privacy of \mathcal{R} . It remains to prove an upper bound on the error for \mathcal{R}' .

Let $v'_1, \dots, v'_n \in \mathbb{B}^{d'-1}$ and let u_1, \dots, u_n be their corresponding binary vectors from the above procedure. Let $v_i = (u_i, 0, \dots, 0) \in \mathbb{R}^d$. Lemma A.2 guarantees that for all $j \in [d]$ we have

$$\mathbb{E} \left[\left\| \left(\sum_{i=1}^n \hat{\mathcal{R}}^+(v_i) - v_i \right)_j \right\|^2 \right] \leq \frac{\text{Err}(\mathcal{A}, \mathcal{R})}{d}.$$

Thus, when truncating to the first $2d'$ coordinates of $\hat{\mathcal{R}}$, we have

$$\mathbb{E} \left[\left\| \sum_{i=1}^n \hat{\mathcal{R}}^+(v_i)[1 : 2d'] - v_i[1 : 2d'] \right\|^2 \right] \leq \frac{d'}{d} \text{Err}(\mathcal{A}, \mathcal{R}).$$

Now, let us analyze the error of \mathcal{R}' . Note that

$$\begin{aligned}
 \mathbb{E} \left[\left\| \sum_{i=1}^n \mathcal{R}'^+(v'_i) - v'_i \right\|^2 \right] &= \mathbb{E} \left[\left\| \sum_{i=1}^n U^T \hat{\mathcal{R}}^+(v_i)[1 : 2d'] - v'_i \right\|^2 \right] \\
 &= \mathbb{E} \left[\left\| \sum_{i=1}^n \hat{\mathcal{R}}^+(v_i)[1 : 2d'] - U v'_i \right\|^2 \right] \\
 &\leq \mathbb{E} \left[\left\| \sum_{i=1}^n \hat{\mathcal{R}}^+(v_i)[1 : 2d'] - u_i + u_i - w_i \right\|^2 \right] \\
 &\leq 2\mathbb{E} \left[\left\| \sum_{i=1}^n \hat{\mathcal{R}}^+(v_i)[1 : 2d'] - u_i \right\|^2 \right] + 2\mathbb{E} \left[\left\| \sum_{i=1}^n u_i - w_i \right\|^2 \right] \\
 &= 2\mathbb{E} \left[\left\| \sum_{i=1}^n \hat{\mathcal{R}}^+(v_i)[1 : 2d'] - v_i[1 : 2d'] \right\|^2 \right] + 2\mathbb{E} \left[\left\| \sum_{i=1}^n u_i - w_i \right\|^2 \right] \\
 &\leq 2 \frac{d'}{d} \text{Err}(\mathcal{A}, \mathcal{R}) + \frac{8n}{d'} \\
 &\leq \mathcal{O}(d'/\varepsilon^2),
 \end{aligned}$$

where the last inequality follows since $\text{Err}(\mathcal{A}, \mathcal{R}) \leq \mathcal{O}(d/\varepsilon^2)$ and $d' \geq n\varepsilon^2/200$.

□

A.2. Missing Proofs from Section 2.1.2

A.2.1. PROOF OF PROPOSITION 2.4

Proposition 2.4. *Let $v_1, \dots, v_n \in \mathbb{S}^{d-1}$ where $v_1 \sim \text{Unif}(\mathbb{S}^{d-1})$, $\mathcal{R} : \mathbb{S}^{d-1} \rightarrow \mathcal{Z}^k$ and \mathcal{A} be an unbiased protocol. For an input set $W = \Pi(R(v_1), \dots, R(v_n))$, Algorithm 2 outputs a set $S \subset \mathbb{R}^d$ of size $\binom{n^k}{k}$ such that*

$$\mathbb{E} [\text{dist}(v_1, S)] = \mathbb{E} \left[\min_{u \in S} \|v_1 - u\|_2^2 \right] \leq \frac{\text{Err}(\mathcal{A}, \mathcal{R})}{n},$$

where the expectation is over the randomness of v_1 and the algorithm.

Proof. The proof builds on the arguments of Lemma 3.1 in (Asi et al., 2022) used in the local privacy model. Let P denote the uniform distribution over the sphere \mathbb{S}^{d-1} . First, note that as Algorithm 2 iterates over all possible subsets of messages of size k , we have that $W_t = \mathcal{R}(v_1)$ for some t , hence the set S has the point

$$u_t = \mathbb{E}_{\tilde{v}_2, \dots, \tilde{v}_n \sim P} [\mathcal{A}(\Pi(\mathcal{R}(v_1), \mathcal{R}(\tilde{v}_2), \dots, \mathcal{R}(\tilde{v}_n)))] \in S$$

We define $\hat{\mathcal{R}}_i$ to be

$$\hat{\mathcal{R}}_i(v_i) = \mathbb{E}_{v_j \sim P, j \neq i} [\mathcal{A}(\Pi(\mathcal{R}(v_1), \dots, \mathcal{R}(v_n)))].$$

Note that $\hat{\mathcal{R}}_1(v_1) \in S$ and that $\mathbb{E}[\hat{\mathcal{R}}_i(v)] = v$ for all $v \in \mathbb{S}^{d-1}$. We define

$$\hat{\mathcal{R}}_{\leq i}(v_1, \dots, v_i) = \mathbb{E}_{v_j \sim P, j > i} \left[\mathcal{A}(\Pi(\mathcal{R}(v_1), \dots, \mathcal{R}(v_n))) - \sum_{j=1}^i v_j \mid v_{1:i} \right],$$

and $\hat{\mathcal{R}}_0 = 0$. We now have

$$\begin{aligned}
 & \mathbb{E}_{v_1, \dots, v_n \sim P} \left[\left\| \mathcal{A}(\Pi(\mathcal{R}(v_1), \dots, \mathcal{R}(v_n))) - \sum_{i=1}^n v_i \right\|_2^2 \right] \\
 &= \mathbb{E}_{v_1, \dots, v_n \sim P} \left[\left\| \hat{\mathcal{R}}_{\leq n}(v_1, \dots, v_n) \right\|_2^2 \right] \\
 &= \mathbb{E}_{v_1, \dots, v_n \sim P} \left[\left\| \hat{\mathcal{R}}_{\leq n}(v_1, \dots, v_n) - \hat{\mathcal{R}}_{\leq n-1}(v_1, \dots, v_{n-1}) + \hat{\mathcal{R}}_{\leq n-1}(v_1, \dots, v_{n-1}) \right\|_2^2 \right] \\
 &\stackrel{(i)}{=} \mathbb{E}_{v_1, \dots, v_n \sim P} \left[\left\| \hat{\mathcal{R}}_{\leq n}(v_1, \dots, v_n) - \hat{\mathcal{R}}_{\leq n-1}(v_1, \dots, v_{n-1}) \right\|_2^2 \right] + \mathbb{E}_{v_1, \dots, v_{n-1} \sim P} \left[\left\| \hat{\mathcal{R}}_{\leq n-1}(v_1, \dots, v_{n-1}) \right\|_2^2 \right] \\
 &\stackrel{(ii)}{=} \sum_{i=1}^n \mathbb{E}_{v_1, \dots, v_i \sim P} \left[\left\| \hat{\mathcal{R}}_{\leq i}(v_1, \dots, v_i) - \hat{\mathcal{R}}_{\leq i-1}(v_1, \dots, v_{i-1}) \right\|_2^2 \right] \\
 &\stackrel{(iii)}{\geq} \sum_{i=1}^n \mathbb{E}_{v_i \sim P} \left[\left\| E_{v_1, \dots, v_{i-1} \sim P} [\hat{\mathcal{R}}_{\leq i}(v_1, \dots, v_i) - \hat{\mathcal{R}}_{\leq i-1}(v_1, \dots, v_{i-1})] \right\|_2^2 \right] \\
 &\stackrel{(iv)}{=} \sum_{i=1}^n \mathbb{E}_{v_i \sim P} \left[\left\| \hat{\mathcal{R}}_i(v_i) - v_i \right\|_2^2 \right]
 \end{aligned}$$

where (i) follows since $\mathbb{E}_{v_n \sim P} [\hat{\mathcal{R}}_{\leq n}(v_1, \dots, v_n)] = \hat{\mathcal{R}}_{\leq n-1}(v_1, \dots, v_{n-1})$, (ii) follows by induction, (iii) follows from Jensen's inequality, and (iv) follows since $E_{v_1, \dots, v_{i-1} \sim P} [\hat{\mathcal{R}}_{\leq i}(v_1, \dots, v_i)] = \hat{\mathcal{R}}_i(v_i) - v_i$ and $E_{v_1, \dots, v_{i-1} \sim P} [\hat{\mathcal{R}}_{\leq i-1}(v_1, \dots, v_{i-1})] = 0$.

Now, as $\hat{\mathcal{R}}_i$ has the same distribution for all i because of the shuffling operator, we get that

$$\mathbb{E}_{v_i \sim P} \left[\left\| \hat{\mathcal{R}}_1(v_1) - v_1 \right\|_2^2 \right] \leq \text{Err}(\mathcal{A}, \mathcal{R})/n.$$

Thus, as $\text{dist}(v_1, S) \leq \left\| \hat{\mathcal{R}}_1(v_1) - v_1 \right\|_2^2$, the claim follows. \square

A.2.2. PROOF OF THEOREM 2.5

Theorem 2.5. *Let $\varepsilon \leq 1$, $\mathcal{R} : \mathbb{S}^{d-1} \rightarrow \mathcal{Z}^k$ be (ε, δ) -shuffle DP, and \mathcal{A} be an unbiased protocol. If $\text{Err}(\mathcal{A}, \mathcal{R}) \leq \mathcal{O}(d/\varepsilon^2)$ then $k \geq \Omega\left(\frac{\min(n\varepsilon^2, d)}{\log n}\right)$.*

Proof. The proof will follow the proof of Theorem 2.1 using the new reconstruction attack of Algorithm 2. Let $\text{Err}(\mathcal{A}, \mathcal{R}) \leq C \cdot d/\varepsilon^2$ for some universal constant $1 \leq C < \infty$.

First, we prove the lower bounds for $d \leq n\varepsilon^2/100C$. Note that Proposition 2.4 and Markov inequality imply that there is a set $A \subset \mathbb{S}^{d-1}$ such that $\Pr_{v \sim \text{Unif}(\mathbb{S}^{d-1})}[A] \geq 1/2$ and for all $v \in A$ and $v_2, \dots, v_n \in \mathbb{S}^{d-1}$, letting S_v be the output of Algorithm 2 over the input $\Pi(R(v), R(v_2), \dots, R(v_n))$, Markov inequality implies

$$\Pr [\text{dist}(v, S_v) \leq 4d/n\varepsilon^2] \geq 1/2.$$

As $\Pr_{v \sim \text{Unif}(\mathbb{S}^{d-1})}[A] \geq 1/2$, this implies that there is a ρ -packing of the unit ball $P = \{v_1, v_2, \dots, v_M\} \subset A$ such that $M = 2^{d \log(1/\rho)-1}$ and $\Pr(\text{dist}(v_i, S_{v_i}) \leq 4Cd/n\varepsilon^2) \geq 1/2$.

We will prove the lower bounds by analyzing the algorithm over the following M datasets:

$$X_i = (v_i, v_1, \dots, v_1),$$

for $i \in [M]$.

Let S_i be the output of the reconstruction attack (Algorithm 2) over the shuffled messages $\Pi(\mathcal{R}(v_i), \mathcal{R}(v_1), \dots, \mathcal{R}(v_1))$. We define the projection set of S_i to the packing P to be $O_i = \{\text{Proj}_P(v) : v \in S_i\}$. Proposition 2.4 now implies that for all $i \in [M]$, $\text{dist}(v_i, S_i) \leq Cd/n\varepsilon^2 \leq \rho$ with probability $1/2$, hence as P is ρ -packing we have that

$$\Pr[v_i \in O_i] \geq \Pr[\text{dist}(v_i, S_i) \leq \rho] \geq 9/10.$$

On the other hand, note that for O_1

$$\begin{aligned} \sum_{i=1}^M \Pr[v_i \in O_1] &= \sum_{i=1}^M \mathbb{E}[1\{v_i \in O_1\}] \\ &= \mathbb{E}\left[\sum_{i=1}^M 1\{v_i \in O_1\}\right] \\ &\leq \mathbb{E}[|O_1|] \leq \binom{nk}{k}. \end{aligned}$$

Hence there exists an $1 \leq i \leq M$ such that

$$\Pr[v_i \in O_1] \leq \frac{\binom{nk}{k}}{M}.$$

As the protocol is (ε, δ) -DP, we also have

$$\begin{aligned} \Pr[v_i \in O_1] &\geq \Pr[v_i \in O_i] e^{-\varepsilon} - \delta \\ &\geq \frac{9}{10e} - \delta \geq 1/6 \end{aligned}$$

Combining these together, and given that $M \geq 2^d/2$ for $\rho = 1/10$, we have that

$$2^d \leq 12 \binom{nk}{k} \leq 6(en)^k.$$

This implies that $k \geq \Omega(d/\log(n))$ whenever $d \leq n\varepsilon^2/100C$.

Now we prove the lower bound for $d \geq n\varepsilon^2/100C$. The proof builds on the following proposition which states that we can convert an optimal protocol for d -dimensional inputs into an optimal protocol for d' -dimensional inputs where $d' = n\varepsilon^2/200C$ with the same number of messages. We defer the proof to Appendix A.2.3.

Proposition A.3. *Let $d' = n\varepsilon^2/200C \geq 1$ and $d \geq 2d'$. Let $\mathcal{R} : \mathbb{S}^{d-1} \rightarrow \mathcal{Z}^k$ be an (ε, δ) -Shuffle DP randomizer with aggregation \mathcal{A} that is unbiased such that $\text{Err}(\mathcal{A}, \mathcal{R}) \leq \mathcal{O}(d/\varepsilon^2)$. There exists $\mathcal{R}' : \mathbb{S}^{d'-1} \rightarrow \mathcal{Z}^k$ and aggregation \mathcal{A}' that is unbiased and (ε, δ) -Shuffle DP such that $\text{Err}(\mathcal{A}', \mathcal{R}') \leq \mathcal{O}(d'/\varepsilon^2)$.*

Let \mathcal{A} and $\mathcal{R} : \mathbb{S}^{d-1} \rightarrow \mathcal{Z}^k$ be an unbiased (ε, δ) -Shuffle DP protocol that obtains error $\text{Err}(\mathcal{A}, \mathcal{R}) \leq \mathcal{O}(d/\varepsilon^2)$ using k messages. Proposition A.3 implies that there is a randomizer $\mathcal{R}' : \mathbb{S}^{d'-1} \rightarrow \mathcal{Z}^k$ and aggregation \mathcal{A}' that is (ε, δ) -Shuffle DP and unbiased such that $\text{Err}(\mathcal{A}', \mathcal{R}') \leq \mathcal{O}(d'/\varepsilon^2)$ for $d' = n\varepsilon^2/200C$. As $d' \leq n\varepsilon^2/100C$, the lower bound we proved above shows that $k \geq \Omega(d'/\log(n)) = \Omega(n\varepsilon^2/\log(n))$. \square

A.2.3. PROOF OF PROPOSITION A.3

The proof will follow the proof of Appendix A.1.3 with general aggregation \mathcal{A} . To this end, in the next lemma we show that we can convert any unbiased protocol into another unbiased one where the error is split evenly across coordinates.

Lemma A.4. *If $\mathcal{R} : \mathbb{S}^{d-1} \rightarrow \mathcal{Z}^k$ is (ε, δ) -shuffle DP randomizer and \mathcal{A} is unbiased, then there exists $\mathcal{R}' : \left\{\frac{-1}{\sqrt{d}}, \frac{1}{\sqrt{d}}\right\}^d \rightarrow \mathcal{Z}^k$ and \mathcal{A}' that is (ε, δ) -shuffle DP and unbiased such that for $j \in [d]$ and $v_1, \dots, v_n \in \left\{\frac{-1}{\sqrt{d}}, \frac{1}{\sqrt{d}}\right\}^d$,*

$$\mathbb{E}\left[\left|\left(\mathcal{A}'(\Pi(\mathcal{R}'(v_1), \dots, \mathcal{R}'(v_n))) - \sum_{i=1}^n v_i\right)_j\right|^2\right] \leq \frac{\text{Err}(\mathcal{A}, \mathcal{R})}{d}.$$

Proof. \mathcal{R}' will use shared public randomness to shuffle the coordinates of each vector and flip the signs of each coordinate. This will ensure that all coordinates will have the same marginal distribution for their error.

More precisely, let $\pi : [d] \rightarrow [d]$ be a random permutation of the coordinates picked uniformly at random, and let $s_1, \dots, s_d \sim \text{Ber}(1/2)$. Our new randomizer \mathcal{R}' over input v has

$$\mathcal{R}'(v) = \mathcal{R}(s \cdot v(\pi)),$$

where $(s \cdot v(\pi))_j = s_j v_{\pi(j)}$ is element-wise product.

Moreover, we define \mathcal{A}' given kn messages $m_i \in \mathcal{Z}$

$$\mathcal{A}'(m_1, \dots, m_{kn}) = s \cdot \mathcal{A}(m_1, \dots, m_{kn})(\pi^{-1}).$$

First, note that the privacy of \mathcal{R}' follows immediately from the privacy of \mathcal{R} . Moreover, \mathcal{A}' is unbiased as \mathcal{A} is unbiased:

$$\begin{aligned} \mathbb{E}[\mathcal{A}'(\Pi(\mathcal{R}'(v_1), \dots, \mathcal{R}'(v_n)))] &= s \cdot \mathbb{E}[\mathcal{A}(\Pi(\mathcal{R}(s \cdot v_1(\pi)), \dots, \mathcal{R}(s \cdot v_n(\pi))))(\pi^{-1})] \\ &= s \cdot \sum_{i=1}^n s \cdot v_i(\pi)(\pi^{-1}) \\ &= \sum_{i=1}^n v_i, \end{aligned}$$

where the last equality follows since $s \cdot s = 1^d$ and $v_i(\pi)(\pi^{-1}) = v_i$.

Now it remains to prove the claim about the error of \mathcal{R}' and \mathcal{A}' . Letting $\hat{v}_i = s \cdot v_i(\pi)$, note that $v = s \cdot \hat{v}(\pi^{-1})$, thus we get

$$\begin{aligned} \left| \left(\mathcal{A}'(\Pi(\mathcal{R}'(v_1), \dots, \mathcal{R}'(v_n))) - \sum_{i=1}^n v_i \right)_j \right|^2 &= \left| \left(s \cdot \mathcal{A}(\Pi(\mathcal{R}(s \cdot v_1(\pi)), \dots, \mathcal{R}(s \cdot v_n(\pi))))(\pi^{-1}) - \sum_{i=1}^n v_i \right)_j \right|^2 \\ &= \left| \left(s \cdot \mathcal{A}(\Pi(\mathcal{R}(\hat{v}_1), \dots, \mathcal{R}(\hat{v}_n)))(\pi^{-1}) - \sum_{i=1}^n s \cdot \hat{v}_i(\pi^{-1}) \right)_j \right|^2 \\ &= \left| \left(\mathcal{A}(\Pi(\mathcal{R}(\hat{v}_1), \dots, \mathcal{R}(\hat{v}_n))) - \sum_{i=1}^n \hat{v}_i \right)_{\pi^{-1}(j)} \right|^2 \end{aligned} \tag{1}$$

Summing over all coordinates,

$$\begin{aligned} \mathbb{E} \left[\sum_{j=1}^d \left(\mathcal{A}'(\Pi(\mathcal{R}'(v_1), \dots, \mathcal{R}'(v_n))) - \sum_{i=1}^n v_i \right)_j^2 \right] &= \mathbb{E} \left[\sum_{j=1}^d \left(\sum_{i=1}^n \mathcal{A}(\Pi(\mathcal{R}(\hat{v}_1), \dots, \mathcal{R}(\hat{v}_n))) - \sum_{i=1}^n \hat{v}_i \right)_{\pi^{-1}(j)}^2 \right] \\ &= \mathbb{E} \left[\left\| \mathcal{A}(\Pi(\mathcal{R}(\hat{v}_1), \dots, \mathcal{R}(\hat{v}_n))) - \sum_{i=1}^n \hat{v}_i \right\|_2^2 \right] \\ &\leq \text{Err}(\mathcal{A}, \mathcal{R}). \end{aligned}$$

Finally, the claim now follows since for all $j \in [d]$, $|(\mathcal{A}'(\Pi(\mathcal{R}'(v_1), \dots, \mathcal{R}'(v_n))) - \sum_{i=1}^n v_i)_j|^2$ have the same distribution and hence the same expectation: indeed, let

$$A = |(\mathcal{A}(\Pi(\mathcal{R}(\hat{v}_1), \dots, \mathcal{R}(\hat{v}_n))) - \sum_{i=1}^n \hat{v}_i)|^2 \quad \text{and } t = \pi^{-1}(j).$$

Equation (1) shows that $|(\mathcal{A}'(\Pi(\mathcal{R}'(v_1), \dots, \mathcal{R}'(v_n))) - \sum_{i=1}^n v_i)_j|^2 = A_t$. Now note that $\hat{v}_1, \dots, \hat{v}_n$ are uniformly random vectors from $\{-1, +1\}^d / \sqrt{d}$ and $\pi^{-1}(j)$ is random coordinate from $[d]$, hence the distribution of A_t is the same for all j . \square

We are now ready to prove Proposition A.3 which will follow the proof of Proposition A.1.

Proof. (of Proposition A.3) We will construct \mathcal{R}' and \mathcal{A}' as follows: for a d' -dimensional input $v' \in \mathbb{S}^{d'-1}$, \mathcal{R}' will first apply the Kashin representation $U \in \mathbb{R}^{2d' \times d'}$ to get $w = Uv' \in \mathbb{R}^{2d'}$ such that $\|w\|_\infty \leq 2/\sqrt{d'}$. Then, it converts w into a binary vector u by setting for all $i \in [2d']$

$$u_i = \begin{cases} 2\text{sign}(w_i)/\sqrt{d'} & \text{with probability } \frac{w_i\sqrt{d'}+2}{4} \\ -2\text{sign}(w_i)/\sqrt{d'} & \text{with probability } \frac{-w_i\sqrt{d'}+2}{4} \end{cases}$$

Note that $E[u_i] = w_i$ and that $E[(w_i - u_i)^2] \leq 4/d'$ since $|u_i| \leq 2/\sqrt{d'}$.

Now, let $\hat{\mathcal{R}}$ and $\hat{\mathcal{A}}$ be the randomizer and aggregation guaranteed from Lemma A.2 for the randomizer \mathcal{R} and aggregation \mathcal{A} . Our \mathcal{R}' will construct $v \in \mathbb{R}^d$ by setting $v = (u, 0, \dots, 0)$ then

$$\mathcal{R}'(v') = \hat{\mathcal{R}}(v).$$

Moreover, we define $\mathcal{A}' : \mathcal{Z}^{nk} \rightarrow \mathbb{R}^{d'}$ to be

$$\mathcal{A}'(m_1, \dots, m_{nk}) = U^T \hat{\mathcal{A}}(m_1, \dots, m_{nk})[1 : 2d']$$

We need to argue that \mathcal{R}' is (ε, δ) -Shuffle DP, that \mathcal{A}' is unbiased, and to prove the claim about utility.

Privacy of \mathcal{R}' follows immediately from privacy of $\hat{\mathcal{R}}$. As for unbiasedness, let $v'_1, \dots, v'_n \in \mathbb{S}^{d'-1}$ and let u_1, \dots, u_n and w_1, \dots, w_n be their corresponding vectors from the above procedure. Let $v_i = (u_i, 0, \dots, 0) \in \mathbb{R}^d$. Note that

$$\begin{aligned} \mathbb{E}[\mathcal{A}'(\Pi(\mathcal{R}'(v'_1), \dots, \mathcal{R}'(v'_n)))] &= \mathbb{E}\left[U^T \hat{\mathcal{A}}(\Pi(\hat{\mathcal{R}}(v_1), \dots, \hat{\mathcal{R}}(v_n)))[1 : 2d']\right] \\ &= U^T \sum_{i=1}^n \mathbb{E}[v_i[1 : 2d']] \\ &= U^T \sum_{i=1}^n \mathbb{E}[u_i] \\ &= U^T \sum_{i=1}^n w_i \\ &= U^T \sum_{i=1}^n Uv'_i \\ &= \sum_{i=1}^n v'_i. \end{aligned}$$

It remains to prove an upper bound on the error of \mathcal{R}' and \mathcal{A}' . First, note that Lemma A.2 guarantees that for all $j \in [d]$ we have

$$E[\left\| \hat{\mathcal{A}}(\Pi(\hat{\mathcal{R}}(v_1), \dots, \hat{\mathcal{R}}(v_n))) - \sum_{i=1}^n v_i \right\|_j^2] \leq \frac{\text{Err}(\mathcal{A}, \mathcal{R})}{d}.$$

Thus, when truncating to the first $2d'$ coordinates of $\hat{\mathcal{R}}$, we have

$$E[\left\| \hat{\mathcal{A}}(\Pi(\hat{\mathcal{R}}(v_1), \dots, \hat{\mathcal{R}}(v_n)))[1 : 2d'] - \sum_{i=1}^n v_i[1 : 2d'] \right\|_j^2] \leq \frac{2d'}{d} \text{Err}(\mathcal{A}, \mathcal{R}).$$

Now, let us analyze the error of \mathcal{R}' . Note that

$$\begin{aligned}
 & \mathbb{E} \left[\left\| \mathcal{A}'(\Pi(\mathcal{R}'(v'_1), \dots, \mathcal{R}'(v'_n))) - \sum_{i=1}^n v'_i \right\|^2 \right] \\
 &= \mathbb{E} \left[\left\| U^T \hat{\mathcal{A}}(\Pi(\hat{\mathcal{R}}(v_1), \dots, \hat{\mathcal{R}}(v_n)))[1 : 2d'] - \sum_{i=1}^n v'_i \right\|^2 \right] \\
 &= \mathbb{E} \left[\left\| \hat{\mathcal{A}}(\Pi(\hat{\mathcal{R}}(v_1), \dots, \hat{\mathcal{R}}(v_n)))[1 : 2d'] - \sum_{i=1}^n U v'_i \right\|^2 \right] \\
 &\leq \mathbb{E} \left[\left\| \hat{\mathcal{A}}(\Pi(\hat{\mathcal{R}}(v_1), \dots, \hat{\mathcal{R}}(v_n)))[1 : 2d'] - \sum_{i=1}^n u_i - u_i + w_i \right\|^2 \right] \\
 &\leq 2\mathbb{E} \left[\left\| \hat{\mathcal{A}}(\Pi(\hat{\mathcal{R}}(v_1), \dots, \hat{\mathcal{R}}(v_n)))[1 : 2d'] - \sum_{i=1}^n u_i \right\|^2 \right] + 2\mathbb{E} \left[\left\| \sum_{i=1}^n u_i - w_i \right\|^2 \right] \\
 &= 2\mathbb{E} \left[\left\| \hat{\mathcal{A}}(\Pi(\hat{\mathcal{R}}(v_1), \dots, \hat{\mathcal{R}}(v_n)))[1 : 2d'] - \sum_{i=1}^n v_i[1 : 2d'] \right\|^2 \right] + 2\mathbb{E} \left[\left\| \sum_{i=1}^n u_i - w_i \right\|^2 \right] \\
 &\leq 4 \frac{d'}{d} \text{Err}(\mathcal{A}, \mathcal{R}) + \frac{8n}{d'} \\
 &\leq \mathcal{O}(d'/\varepsilon^2),
 \end{aligned}$$

where the last inequality follows since $\text{Err}(\mathcal{A}, \mathcal{R}) \leq \mathcal{O}(d/\varepsilon^2)$ and $d' \geq n\varepsilon^2/200$. □

A.3. Missing Proofs for Section 2.2

Our d -dimensional algorithm builds on the 1-dimensional algorithm by (Ghazi et al., 2021). We let $\mathcal{R}_{\text{GKMPs}}^{(\varepsilon, \delta)}$ denote the local randomizer with parameters (ε, δ) and \mathcal{A}^+ is their aggregation (which is summation over messages). Their protocol has the following guarantees for 1-dimensional summation.

Lemma A.5. (Ghazi et al., 2021) *There is a local randomizer $\mathcal{R}_{\text{GKMPs}}^{(\varepsilon, \delta)} : [0, 1] \rightarrow \mathbb{R}^*$ that is (ε, δ) -Shuffle DP such that each user sends $1 + \tilde{\mathcal{O}}_\varepsilon \left(\frac{\log(1/\delta)}{\sqrt{n}} \right)$ in expectation and has error*

$$\text{Err}(\mathcal{R}_{\text{GKMPs}}^{(\varepsilon, \delta)}, \mathcal{A}^+) \leq \mathcal{O}(1/\varepsilon^2).$$

We also used advanced composition in our privacy proof.

Lemma A.6 (Advanced composition (Dwork & Roth, 2014)). *If A_1, \dots, A_k are randomized algorithms that each is (ε, δ) -DP, then their composition $(A_1(D), \dots, A_k(D))$ is $(\sqrt{2k \log(1/\delta)}\varepsilon + k\varepsilon(e^\varepsilon - 1), \delta' + k\delta)$ -DP where D is the input dataset.*

Now we present the guarantees of our protocol.

Theorem 2.6. *Let $\mathcal{R} : \mathbb{S}^{d-1} \rightarrow \mathbb{R}^{2d}$ be the local randomizer in Algorithm 3 and $\mathcal{A} : (\mathbb{R}^{2d})^* \rightarrow \mathbb{R}^d$ be the aggregation in Algorithm 4. Then, \mathcal{R} is (ε, δ) -Shuffle DP randomizer, each users sends $d \cdot \left(1 + \tilde{\mathcal{O}}_\varepsilon \left(\frac{\log(1/\delta)}{\sqrt{n}} \right) \right)$ messages in expectation, and the protocol has error*

$$\text{Err}(\mathcal{A}, \mathcal{R}) \leq \mathcal{O} \left(\frac{d \log(1/\delta)}{\varepsilon^2} \right).$$

Proof. First, note that the guarantees of Kashin representation imply that each $|u_j^{(i)}| \leq 1$, hence we can use $\mathcal{R}_{\text{GKMPs}}$. The claim from privacy follows from the fact the $\mathcal{R}_{\text{GKMPs}}$ is $(\varepsilon_0, \delta_0)$ -Shuffle DP and advanced composition of d such mechanisms.

Now we analyze the error of the protocol. We have

$$\begin{aligned}
 \mathbb{E} \left[\left\| \hat{v} - \sum_{i=1}^n v^{(i)} \right\|_2^2 \right] &= \mathbb{E} \left[\left\| \frac{C_K}{\sqrt{d}} U_K^T \hat{u} - \frac{C_K}{\sqrt{d}} U_K^T \sum_{i=1}^n u^{(i)} \right\|_2^2 \right] \\
 &= \frac{C_K^2}{d} \mathbb{E} \left[\left\| \hat{u} - \sum_{i=1}^n u^{(i)} \right\|_2^2 \right] \\
 &= \frac{C_K^2}{d} \mathbb{E} \left[\left\| \sum_{m \in M} m - u^{(i)} \right\|_2^2 \right] \\
 &\leq \frac{C_K^2}{d} \frac{d}{\varepsilon_0^2} \\
 &\leq \mathcal{O} \left(\frac{d \log(1/\delta)}{\varepsilon^2} \right),
 \end{aligned}$$

where the last inequality follows from the guarantees of the $\mathcal{R}_{\text{GKMPs}}$ protocol which has error $1/\varepsilon_0^2$ in each coordinate. The claim follows. \square

B. Missing Proofs from Section 3

To prove the lower bound, we first note that it suffices to assume that the local randomizer has bounded outputs and that the analyzer simply adds up all of the messages sent by the users, as shown by the next lemma.

Lemma B.1. *Let $\mathcal{P} = (\mathcal{R}, \mathcal{A})$ be an n -party protocol for vector aggregation in the single-message shuffle model. Let V be a random variable on $\left[-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}}\right]^d$ and suppose that users sample their inputs from the distribution V^n . Then there exists a protocol $\mathcal{P}' = (\mathcal{R}', \mathcal{A}')$ with user outputs $u_1, \dots, u_n \in \mathbb{R}^d$ such that:*

$$(1) \quad \mathcal{A}'(u_1, \dots, u_n) = \sum_{i=1}^n u_i \text{ and } \mathcal{R}' \text{ maps to } \left[-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}}\right]^d.$$

$$(2) \quad \text{MSE}(\mathcal{P}', V) \leq \text{MSE}(\mathcal{P}, V)$$

$$(3) \quad \text{If } \mathcal{S} \circ \mathcal{R}^n \text{ is } (\varepsilon, \delta)\text{-DP, then } \mathcal{S} \circ (\mathcal{R}')^n \text{ is } (\varepsilon, \delta)\text{-DP.}$$

Proof. The proof is similar to Lemma 4.1 in (Balle et al., 2019), generalizing from scalars to vectors. Let $\mathcal{R}' = f \circ \mathcal{R}$ be the post-processing local randomizer that uses the posterior mean estimator $f(u) = \mathbb{E}[V \mid \mathcal{V} = u]$ is the minimum MSE estimator. Then \mathcal{R}' maps to $\left[-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}}\right]^d$ as claimed.

Observe that for any estimator h of $Z := V_1 + \dots + V_n$ given the input $U = \{u_1, \dots, u_n\}$, we have

$$\begin{aligned}
 \text{MSE}(h, U) &= \mathbb{E} [(h(u) - Z)^2 \mid U] \\
 &= \mathbb{E} [Z^2 \mid U] - 2h(u) \cdot \mathbb{E}[Z \mid U] + (h(u))^2.
 \end{aligned}$$

This quantity is minimized over the choice of h at $h(u) = \mathbb{E}[Z \mid U]$.

Finally, since f is a post-processing local randomizer, then $\mathcal{S} \circ (\mathcal{R}')^n$ is (ε, δ) -DP by the post-processing property of DP. \square

Lemma B.2. Let $\mathcal{P} = (\mathcal{R}, \mathcal{A})$ be an n -party protocol for vector aggregation in the single-message shuffle model such that $\mathcal{R} : \left[-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}}\right]^d \rightarrow \left[-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}}\right]^d$ and \mathcal{A} is vector summation. Suppose V^n are n copies of a random variable V . Then

$$\text{MSE}(\mathcal{P}, V^n) \geq n\mathbb{E} [\|\mathcal{R}(V) - V\|_2^2].$$

Proof. The proof generalizes Lemma 4.2 in (Balle et al., 2019) from scalar inputs to vector inputs. Note that we can decompose the mean-squared error as follows.

$$\begin{aligned} \text{MSE}(\mathcal{P}, V^n) &= \mathbb{E} \left[\left\| \sum_{i=1}^n \mathcal{R}(V_i) - V_i \right\|_2^2 \right] \\ &= \sum_i \mathbb{E} [\|\mathcal{R}(V_i) - V_i\|_2^2] + \sum_{i \neq j} \mathbb{E} [\langle \mathcal{R}(V_i) - V_i, \mathcal{R}(V_j) - V_j \rangle] \\ &= \sum_i \mathbb{E} [\|\mathcal{R}(V_i) - V_i\|_2^2] + \sum_{i \neq j} \langle \mathbb{E} [\mathcal{R}(V_i) - V_i], \mathbb{E} [\mathcal{R}(V_i) - V_i] \rangle \\ &\geq n\mathbb{E} [\|\mathcal{R}(V) - V\|_2^2]. \end{aligned}$$

□

Consider the partition P of the hypercube $[0, 1]^d$ into r^d disjoint hypercubes with side length $\frac{1}{r}$. Let $I = \left\{ \frac{m}{r} - \frac{1}{2r} \mid m \in [r] \right\}$ and $J = I^d$. For each $a \in J$, we use $J(a)$ to denote the hypercube of P that contains J . For any $b \in J$, we use the notation $p_{a,b}$ to denote the probability that the randomizer maps a to $J(b)$.

Lemma B.3. Let $r \geq 32$. For any $b \in J$, we have

$$\frac{1}{r^d} \sum_{a \in J \setminus b} \left(\min \left(\left\| a - b \right\|_2 - \frac{\sqrt{d}}{2r}, 0 \right) \right)^2 \geq \frac{d}{2048}.$$

Proof. Let B be a hypercube with length $\frac{1}{8}$ centered at b . Note that we have $\mathbf{Pr}[a \in J \setminus B] \geq \frac{1}{2}$. For $a \in J \setminus B$, we have $\|a - b\|_2 \geq \frac{\sqrt{d}}{16}$. Then for $r \geq 64$, we have $\left(\|a - b\|_2 - \frac{1}{2r}\right)^2 \geq \frac{d}{32^2}$. Hence we have

$$\frac{1}{r^d} \sum_{a \in J \setminus b} \left(\left\| a - b \right\|_2 - \frac{1}{2r} \right)^2 \geq \frac{1}{2} \cdot \frac{d}{32^2} = \frac{d}{2048}.$$

□

Lemma B.4. The mean-squared error of the randomizer \mathcal{R} on the random variable V is at least:

$$\mathbb{E} [\|\mathcal{R}(V) - V\|_2^2] \geq \sum_{b \in J} \min \left(\frac{d(1 - p_{b,b})}{4r^{2+d}}, \min_{a \in J} p_{a,b} \cdot \frac{d}{2048} \right).$$

Proof. For the cases where the randomizer maps V to a value outside of its hypercube, we have:

$$\begin{aligned} \mathbb{E} [\|\mathcal{R}(V) - V\|_2^2] &= \sum_{b \in J} \mathbb{E} [\|\mathcal{R}(b) - b\|_2^2] \cdot \mathbf{Pr}[V = b] \\ &= \frac{1}{r^d} \sum_{b \in J} \mathbb{E} [\|\mathcal{R}(b) - b\|_2^2] \\ &\geq \frac{1}{r^d} \sum_{b \in J} (1 - p_{b,b}) \cdot \frac{d}{4r^2} \\ &= \sum_{b \in J} \frac{d(1 - p_{b,b})}{4r^{2+d}}. \end{aligned}$$

We also have

$$\begin{aligned}
 \mathbb{E} [\|\mathcal{R}(V) - V\|_2^2] &= \frac{1}{r^d} \sum_{b \in J} \mathbb{E} [\|\mathcal{R}(b) - b\|_2^2] \\
 &\geq \frac{1}{r^d} \sum_{b \in J} \sum_{a \in J \setminus b} p_{a,b} \left(\min \left(\|a - b\|_2 - \frac{\sqrt{d}}{2r}, 0 \right) \right)^2 \\
 &\geq \frac{1}{r^d} \sum_{b \in J} \min_{a \in J} p_{a,b} \sum_{a \in J \setminus b} \left(\min \left(\|a - b\|_2 - \frac{\sqrt{d}}{2r}, 0 \right) \right)^2 \\
 &\geq \sum_{b \in J} \min_{a \in J} p_{a,b} \cdot \frac{d}{2048},
 \end{aligned}$$

where the last inequality is from [Lemma B.3](#). Hence, we have

$$\mathbb{E} [\|\mathcal{R}(V) - V\|_2^2] \geq \sum_{b \in J} \min \left(\frac{d(1 - p_{b,b})}{4r^{2+d}}, \min_{a \in J} p_{a,b} \cdot \frac{d}{2048} \right).$$

□

Lemma B.5 (Lemma 4.5 in [\(Balle et al., 2019\)](#)). *Let $\mathcal{R} : [0, 1]^d \rightarrow [0, 1]^d$ be a local randomizer such that the shuffled protocol $\mathcal{M} = \mathcal{S} \circ \mathcal{R}^n$ is (ε, δ) -DP with $\delta < \frac{1}{2}$. Then for any $a, b \in J$ with $a \neq b$, we have either $p_{b,b} < 1 - \frac{e^{-\varepsilon}}{2}$ or $p_{a,b} \geq \frac{1}{n} \cdot \left(\frac{1}{2} - \delta\right)$.*

We are now ready to prove the lower bound.

Proof. By [Lemma B.2](#), we have $\text{MSE}(\mathcal{P}, V^n) \geq n \mathbb{E} [\|\mathcal{R}(V) - V\|_2^2]$. By [Lemma B.4](#), we have $\mathbb{E} [\|\mathcal{R}(V) - V\|_2^2] \geq \sum_{b \in J} \min \left(\frac{d(1 - p_{b,b})}{4r^{2+d}}, \min_{a \in J} p_{a,b} \cdot \frac{d}{2048} \right)$. Therefore by [Lemma B.5](#),

$$\begin{aligned}
 \text{MSE}(\mathcal{P}, V) &\geq n \sum_{b \in J} \min \left(\frac{d(1 - p_{b,b})}{4r^{2+d}}, \min_{a \in J} p_{a,b} \cdot \frac{d}{2048} \right) \\
 &\geq n \sum_{b \in J} \min \left(\frac{de^{-\varepsilon}}{4r^{2+d}}, \frac{1}{n} \cdot \left(\frac{1}{2} - \delta\right) \cdot \frac{d}{2048} \right) \\
 &\geq nr^d \min \left(\frac{de^{-\varepsilon}}{4r^{2+d}}, \frac{1}{n} \cdot \left(\frac{1}{2} - \delta\right) \cdot \frac{d}{2048} \right).
 \end{aligned}$$

The quantity is maximized for $r = \mathcal{O}(n^{1/(d+2)})$ with value $\Omega(dn^{d/(d+2)})$. □

C. Missing Proofs from Section 4

Theorem 4.1. *Let $\varepsilon = \mathcal{O}(1)$ and $\delta < \frac{1}{nd}$. Then any (ε, δ) -shuffle DP mechanism for vector summation that takes the sum of the messages across n players with k malicious users has additive error $\Omega\left(\frac{kd}{\log^2(nd)}\right)$.*

Proof. Suppose \mathcal{A} is a protocol in which 1) each user receives an input v and outputs d messages from a randomizer $\mathcal{R}(v)$, 2) after shuffling, the protocol collects the messages and outputs the sum of the messages. We consider casework on the distribution of the output of the randomizer \mathcal{R} .

Firstly, suppose that for an input vector v , $\Pr \left[\max_{m \in \mathcal{R}(v)} \|m\|_2 \geq \frac{1}{\sqrt{d}\alpha} \right] \geq \frac{1}{dn^2}$, for some parameter $\alpha > 1$ to be fixed. Note that a single malicious user can then run the randomizer \mathcal{R} on inputs $v^{(1)}$ and $v^{(2)}$ a total of $\mathcal{O}(dn^2)$ times and with probability 0.99, find a message m such that $\|m\|_2 \geq \frac{1}{\sqrt{d}\alpha}$. Note that the malicious user can send the message m a total of d times, which contributes L_2 norm $\frac{\sqrt{d}}{\alpha}$. Since each malicious user previously had a unit vector, then the mean squared error induced by each malicious user is at least $\frac{d}{\alpha^2}$. Therefore, k malicious users can induce mean squared error $\frac{kd}{\alpha^2}$.

Secondly, suppose that for an input vector v , we have $\sup \left\langle \frac{m}{\|m\|_2}, v \right\rangle > \frac{100\sqrt{\log nd}}{\sqrt{d}}$. We claim this would violate privacy. Note that for a random vector u , we have by the rotational invariance of Gaussians,

$$\Pr \left[\left\langle \frac{m}{\|m\|_2}, u \right\rangle > \frac{100\sqrt{\log nd}}{\sqrt{d}} \right] < \frac{1}{10n^2d^2}.$$

With probability at least $\frac{1}{10nd}$, none of the nd messages has correlation at least $\frac{100\sqrt{\log nd}}{d}$ with u . Thus we would be able to distinguish between the cases where the inputs are the neighboring datasets (v, v, \dots, v) and (u, v, \dots, v) , which contradicts (ε, δ) -differential privacy for $\varepsilon = \mathcal{O}(1)$ and $\delta < \frac{1}{nd}$,

It remains to consider the case where $\max_{m \in \mathcal{R}(v) \cup \mathcal{R}(u)} \|m\|_2 < \frac{1}{\sqrt{d}\alpha}$ and $\sup \left\langle \frac{m}{\|m\|_2}, v \right\rangle \leq \frac{100\sqrt{\log nd}}{\sqrt{d}}$. Note that in this case, we have

$$\begin{aligned} \left\langle \sum_{i \in [d]} m_i, v \right\rangle &= \sum_{i \in [d]} \|m_i\|_2 \cdot \left\langle \frac{m_i}{\|m_i\|_2}, v \right\rangle \\ &\leq \sup_{i \in [d]} \|m_i\|_2 \cdot d \cdot \sup_{i \in [d]} \left\langle \frac{m_i}{\|m_i\|_2}, v \right\rangle \\ &\leq \frac{100\sqrt{\log nd}}{\sqrt{d}} \cdot d \cdot \frac{1}{\sqrt{d}\alpha} = \frac{100\sqrt{\log nd}}{\alpha}. \end{aligned}$$

Thus for $\alpha > 200\sqrt{\log nd}$ and an elementary vector v , we have that

$$\left\langle \sum_{i \in [d]} m_i, v \right\rangle \leq \frac{1}{2},$$

and thus the mean squared error for the input (v, v, \dots, v) would be at least $\frac{n}{2}$.

Hence for $n > kd$, the mean squared error induced by k malicious users is at least $\Omega\left(\frac{kd}{\log^2(nd)}\right)$. \square

Theorem 1.5. *Let $\varepsilon = \mathcal{O}(1)$ and $\delta < \frac{1}{nd}$. Then any (ε, δ) -DP mechanism for vector summation in which s shufflers take messages corresponding to a disjoint subset of the coordinates and returns the sum of the messages across n players with k malicious users has additive error mean squared error $\Omega\left(\frac{kd}{s \log^2(nd)}\right)$.*

Proof. For $i \in [s]$, let d_i be the number of coordinates for which the i -th shuffler is responsible. Then we have $d_1 + \dots + d_s$. By Theorem 4.1, there exists a set of messages for which k malicious users can induce mean squared error $\Omega\left(\frac{kd_i^2}{\log^2(nd)}\right)$ through sum of the messages in the i -th shuffler. Now, we have that the mean squared error is $\sum_{j \in [n]} \|x_j\|_2^2 C\left(\sum_{i \in [s]} \frac{kd_i^2}{\log^2(nd)}\right)$, which is minimized at $\Omega\left(\frac{kd^2}{s \log^2(nd)}\right)$ for $d_1 = \dots = d_s = \frac{d}{s}$ by a standard power mean inequality. \square