vqSGD: Supplementary Material

A Experiments

We use our gradient quantization scheme to train a fully connected ReLU activated network with 1000 hidden nodes using the MNIST [24] and the Fashion MNIST [38] dataset (60000 data points with 10 classes for each). We use the *cross-entropy loss* function for the training the neural network with a total of $d = 795010$ parameters.

The dataset is divided equally among 100 workers. Each worker computes the local gradients and communicates the quantized gradient to the master which then aggregates and send the updated parameters. We plot the error at each iteration (Figure 1) and compare our results with QSGD quantization.

We use vqSGD with cross polytope scheme, $Q_{C_{cp}}$, along with the variance reduction technique with repetition parameter $s = 100$. Therefore, each local machine sends about $2060 = 100 \cdot \log(2d)$ bits per iteration whereas, QSGD requires 3825.05 bits for MNIST and 2266.79 bits for Fashion MNIST, of communication per iteration per machine (computed by averaging over the total bits of communication over 50 iterations) to communicate the quantized gradient. Our results indicate that vqSGD converges at a similar rate to QSGD while communicating much lesser bits.

We also our vqSGD with the cross polytope scheme, $Q_{C_{cp}}$, to train a ReLU network with 4000 hidden nodes using the CIFAR 10 dataset [23]. This dataset also has 10 classes, every other set up is same except now we have $d = 12332010$ parameters.

The dataset is again equally divided among 100 users. Using vqSGD, each machine send 2455 bits per iteration using the variance reduction scheme. On the other hand, for QSGD, the number of bits per machine per iteration is 4096.9 (computed by averaging over the total bits of communication over 50 iterations). As is evident from the plot in Figure 1, vqSGD communicates lesser number of bits to achieve similar performance.

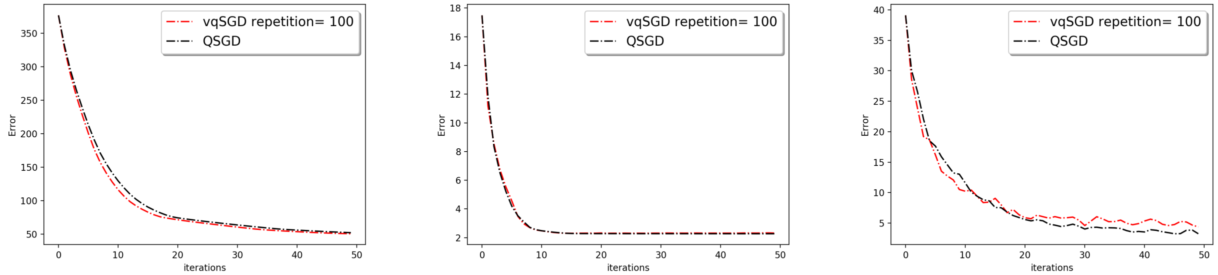

(a) MNIST (b) Fashion MNIST (c) CIFAR10

Figure 1: Convergence for fully connected ReLU network compared to QSGD

Further we experimentally show the performance of vqSGD using the cross polytope Q_{cp}, to solve the least squares problem and logistic regression for binary classification.

Least Squares: In the least square problem, we solve for $\boldsymbol{\theta^*} = \arg\min_{\boldsymbol{\theta}} \|A\boldsymbol{\theta} - \boldsymbol{b}\|_2^2$, where the matrix $A \in \mathbb{R}^{n\times d}$ and $\boldsymbol{\theta^*} \in \mathbb{R}^d$ are generated by sampling each entry from $\mathcal{N}(0, 1)$ and we set $\boldsymbol{b} = A\boldsymbol{\theta^*}$.

In order to show the performance of vqSGD, we simulate the iterations of distributed SGD with $n = 10000$ data samples distributed equally among $N = 500$ worker nodes. In every iteration of SGD, each worker node computes the local gradient on individual data batch and communicates the quantized version of the local gradient to the parameter server. The parameter server on receiving all the quantized gradients averages them and broadcasts the updated model to all the workers. The convergence of SGD is measured by the error term $\|\boldsymbol{\theta^*} - \boldsymbol{\theta_t}\|_2$, where $\boldsymbol{\theta_t}$ is the computed parameter at the end of t-th iteration of distributed SGD.

We compare the convergence of the least square problem for $d = 100, 200, 500$ against the state-of-the-art quantization schemes - DME [33] and QSGD [5]. The results are presented in Figure 2.

The results indicate that vqSGD achieves the same rate of convergence and accuracy as DME and QSGD while communicating only $\log(2d)$ bits and one real (l_2 norm of the vector form each server), whereas, DME (one bit stochastic quantization) and QSGD both require communication of about $\sqrt{d}$ bits and one real.

For the same problem setup, we also show the improvement in the performance of vqSGD using the repetition

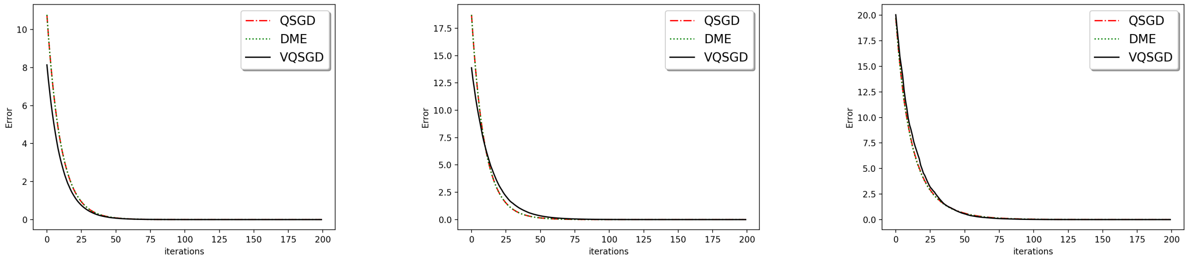

(a) d=100 (b) d=200 (c) d=500

Figure 2: Comparison of convergence for the least square problem with $d = 100, 200, 500$.

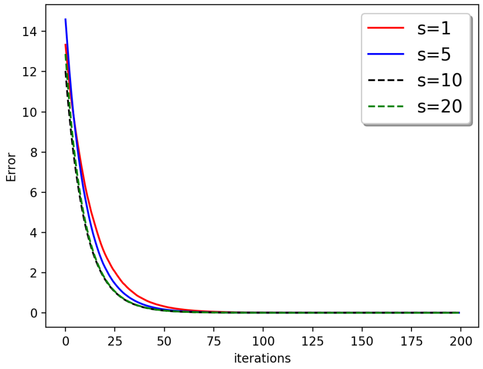

Figure 3: Convergence of $\boldsymbol{\theta_t}$ for $s = 1, 5, 10, 20$ for least square problem

technique for variance reduction. Recall that using repetition technique, each worker now sends s different indices instead of 1 which increases the communication to $s \log(2d)$ bits and 1 real. In Figure 3 we plot the convergence of the lease square problem with $d = 200$ with different values of $s = 1, 5, 10, 20$. We see the evident improvement in the convergence of vqSGD using this repetition scheme with increasing s.

Binary Classification: We compared the performance of vqSGD against DME and QSGD for the binary classification problem with logistic regression using various datasets from the UCI repository [10]. The logistic regression objective is defined as

$$\frac{1}{n}\sum_{i=1}^{n} \log(1 + \exp(-b_i \boldsymbol{a_i}^T \boldsymbol{\theta}) + \frac{1}{2n}\|\boldsymbol{\theta}\|_2^2, \tag{8}$$

where $\boldsymbol{\theta} \in \mathbb{R}^d$ is the parameter, $\boldsymbol{a_i} \in \mathbb{R}^d$ is the feature data and $b_i \in \{-1, +1\}$ is its corresponding label.

We partition the data into 20 equal-sized batches, each assigned to a different worker node. We calculate the classification error for different (test) datasets after training the parameter in the distributed settings (same as described in least square problem). Results of the experiments are presented in Table 4, where each entry is averaged over 20 different runs.

Method	DME	QSGD	vqSGD
a5a ($d = 122$)	0.238 ± 0.0003	0.238 ± 0.0002	0.2368 ± 0.0029
a9a ($d = 123$)	0.234 ± 0.0003	0.234 ± 0.00017	0.234 ± 0.0015
gisset-scale ($d = 5000$)	0.0947 ± 0.00384	0.10475 ± 0.006	0.1480 ± 0.0174
splice ($d = 60$)	0.467 ± 0.017	0.4505 ± 0.0352	0.16618 ± 0.0054

Table 4: Comparison in classification error (mean± standard deviation) for various UCI datasets

We note that for most datasets, with the exception of gisset-scale, vqSGD with $O(N \log d)$ bits of communication per iteration performs equally well or sometimes even better than QSGD and DME with $O(Nd)$ bits of communication per iteration.

B Missing proofs from Section 4

Proof of Lemma 1. $\mathbf{E}[Q_C(\boldsymbol{v})] = \sum_{i=1}^{|C|} a_i \cdot \boldsymbol{c_i} = pv.$ □

Proof of Lemma 2. From the definition of the quantization function,

$$\mathbf{E}[\|\boldsymbol{v} - Q_C(\boldsymbol{v})\|_2^2] = \mathbf{E}[\|Q_C(\boldsymbol{v})\|^2] - \|\boldsymbol{v}\|^2 \leq R^2.$$

This is true as C satisfies Condition (1) and therefore, each point $\boldsymbol{c_i} \in C$ has a bounded norm, $\|\boldsymbol{c_i}\| \leq R$. □

Proof of Theorem 3. Since $\hat{\boldsymbol{g}}$ is the average of N unbiased estimators, the fact that $\mathbf{E}[\hat{\boldsymbol{g}}] = \boldsymbol{g}$ follows from Lemma 1. For the variance computation, note that

$$\begin{aligned}\mathbf{E}[\|\boldsymbol{g} - \hat{\boldsymbol{g}}\|_2^2] &= \frac{1}{N^2}\left(\sum_{i=1}^{N} \mathbf{E}[\|\boldsymbol{g_i} - \hat{\boldsymbol{g}}_{\boldsymbol{i}}\|_2^2]\right) \qquad (\text{ since } \hat{\boldsymbol{g}}_{\boldsymbol{i}} \text{ is an unbiased estimator of } \boldsymbol{g}\ ) \\ &\leq \frac{R^2}{N^2}\sum_{i=1}^{N} \|g_i\|^2 \qquad (\text{from Lemma 2}).\end{aligned}$$

□

Proof of Proposition 4. The proof follows simply by linearity of expectations and Lemma 2.

$$\mathbf{E}\left[\|\boldsymbol{v} - \hat{\boldsymbol{v}}\|_2^2\right] = \mathbf{E}\left[\|\frac{1}{s}\sum_{i=1}^{s}\left(\boldsymbol{v} - Q_C(\boldsymbol{v})\right)\|^2\right] \leq \frac{1}{s}\cdot R^2.$$

□

C Missing proofs from Section 5

To prove Theorem 7, we first show the following lemma that allows us to union bound over the discrete set of points in an ε-net of a unit sphere. Consider an ε-net for the unit sphere $N(\varepsilon)$ for any $\varepsilon \leq 1/R$. We know that such a set exists with $|N(\varepsilon)| \leq \left(1 + \frac{2}{\varepsilon}\right)^d \leq \left(\frac{3}{\varepsilon}\right)^d$ [12].

Lemma 17. *Let C be a set of points in $\mathbb{R}^d$ such that $\|\boldsymbol{c}\|^2 \leq R^2$ for all $\boldsymbol{c} \in C$. If for each $\boldsymbol{y} \in N(\varepsilon)$, $\boldsymbol{y}^T\boldsymbol{c} \geq 2$ for some $\boldsymbol{c} \in C$, then it holds that for each $\boldsymbol{x} \in S^{d-1}$, there is a $\boldsymbol{c}' \in C$ such that $\boldsymbol{x}^T\boldsymbol{c}' \geq 1$.*

Proof. Let $\boldsymbol{y} \in N(\varepsilon)$ be a net-point, and $\boldsymbol{c} \in C$ be such that $\boldsymbol{y}^T c \geq 2$. Note that all points $\boldsymbol{x} \in S^{d-1}$ in the ε-neighborhood of $\boldsymbol{y}$ can be written as $\boldsymbol{x} = \boldsymbol{y} + \tilde{\boldsymbol{y}}$, where $\tilde{\boldsymbol{y}} \in \mathbb{R}^d$ has norm at most ϵ. Therefore, $\boldsymbol{x}^T\boldsymbol{c} = \boldsymbol{y}^T\boldsymbol{c} + \tilde{\boldsymbol{y}}^T\boldsymbol{c} \geq 2 - \|\tilde{\boldsymbol{y}}\|\|\boldsymbol{c}\| \geq 1$. Since $N(\varepsilon)$ covers the entire unit sphere, it follows that for all points $\boldsymbol{x}$ on the unit sphere, there will be a $\boldsymbol{c} \in C$ such that $\boldsymbol{x}^T\boldsymbol{c} \geq 1$. □

Proof of Theorem 7. Let us choose the random set C of $t := e^{\frac{20d}{R^2} + 2\log d}$ points in $\mathbb{R}^d$ in the following way: Each coordinate of any $\boldsymbol{c} \in C$ is chosen independently from a zero-mean Gaussian distribution with variance $\sigma^2 := \frac{R^2}{9d}$. We say that a vector $\boldsymbol{x} \in S^{d-1}$ is a witness for C if $\boldsymbol{x}^T\boldsymbol{c} < 1$ for all $\boldsymbol{c} \in C$. We now show that with high probability, there is no witness for C in S^{d-1}. Using Lemma 17, it is sufficient to show that for any $\boldsymbol{x} \in N(\varepsilon)$, $\boldsymbol{x}^T\boldsymbol{c} \geq 2$ for some $\boldsymbol{c} \in C, \varepsilon \leq 1/R$.

Let us define E_1 to be the event that $\|\boldsymbol{c}\|^2 \leq R^2$ for all $\boldsymbol{c} \in C$. Since every entry of $\boldsymbol{c}$ is chosen from i.i.d. Gaussian, the norm $\|\boldsymbol{c}\|^2$ is distributed according to χ^2-distribution with variance $2d\sigma^4$. Since χ^2-distribution is subexponential [34][Eq. 2.18], for any $\boldsymbol{c} \in C$, we have, for any $l \geq 1$, $\Pr(\|\boldsymbol{c}\|^2 > d\sigma^2(l+1)) \leq e^{-dl/8}$. This implies,

$$\Pr(\|\boldsymbol{c}\|^2 > R^2) \leq e^{-\frac{1}{8}(R^2/\sigma^2 - d)} \leq e^{-d},$$

substituting the value of σ^2. Then by union bound over all $c \in C$.

$$\Pr[\bar{E}_1] \leq te^{-d} = e^{-d + 20d/R^2 + 2\log d} \leq e^{-\Omega(d)}, \tag{9}$$

for $R \geq 5$.

Let E_2 denote the event that for each $\boldsymbol{y} \in N(\varepsilon)$, there exists $\boldsymbol{c} \in C$ such that $\boldsymbol{y}^T\boldsymbol{c} \geq 2$. For any fixed $\boldsymbol{y} \in N(\varepsilon)$, and $\boldsymbol{c} \in C$, define $p_{\boldsymbol{y},\boldsymbol{c}}$ to be the probability that $\boldsymbol{y}^T\boldsymbol{c} \geq 2$. Note that since c has i.i.d. Gaussian entries, then for any fixed $\boldsymbol{y} \in N(\varepsilon)$, the inner product $\boldsymbol{y}^T\boldsymbol{c}$ is distributed according to $\mathcal{N}(0, \sigma^2)$. Using standard bounds for Gaussian distributions [8],

$$\begin{aligned} p_{\boldsymbol{y},\boldsymbol{c}} &:= \Pr[\boldsymbol{y}^T\boldsymbol{c} \geq 2] \geq \frac{2\sigma}{(\sigma^2+4)\sqrt{2\pi}} e^{-\frac{2}{\sigma^2}} \\ &\geq \frac{1}{\sqrt{2\pi}} \min(\sigma^{-1}, \sigma/4) e^{-\frac{2}{\sigma^2}} \geq \frac{R}{12\sqrt{2\pi d}} e^{-\frac{2}{\sigma^2}}, \end{aligned}$$

for any $R \leq 6\sqrt{d}$.

Since each $\boldsymbol{c}$ is chosen independently, the probability that $\boldsymbol{y}^T\boldsymbol{c} < 2$ for all $\boldsymbol{c} \in C$ is $(1 - p_{\boldsymbol{y},\boldsymbol{c}})^t \leq e^{-t \cdot p_{\boldsymbol{y},\boldsymbol{c}}}$. Now for $\varepsilon = 1/R$, by union bound, since $|N(\varepsilon)| \leq \left(\frac{3}{\varepsilon}\right)^d$,

$$\begin{aligned} \Pr[\bar{E_2}] &= \Pr[\exists\ \boldsymbol{y} \in N(\varepsilon) \text{ s.t. } \boldsymbol{y}^T\boldsymbol{c} < 2\ \forall\ \boldsymbol{c} \in C] \\ &\leq e^{-t \cdot p_{y,c} + d\log 3R} \\ &= e^{-t \cdot e^{-\frac{18d}{R^2} - \log(\frac{12\sqrt{2\pi d}}{R})} + d\log 3R} \\ &= e^{-e^{\frac{2d}{R^2} + 2\log d - \log(\frac{12\sqrt{2\pi d}}{R})} + d\log 3R} \\ &= e^{-d^2 e^{\frac{2d}{R^2} - \log(\frac{12\sqrt{2\pi d}}{R})} + d\log 3R} \\ &\leq e^{-\Omega(d)}. \end{aligned}$$

Therefore, $\Pr[\bar{E_1} \cup \bar{E_2}] \leq e^{-\Omega(d)}$. Then using Lemma 17 and Theorem 6, we obtain the statement of the theorem. □

Proof of Proposition 8. We prove this theorem by showing that the point set C_{RM} satisfies the characterization of Theorem 6. Since all points in C_{RM} have squared norm exactly d, from Lemma 1 and Lemma 2, the proof follows.

First note that the matrix with the points in C_{RM} as its rows, has the following structure:

$$H := \begin{bmatrix} H_p \\ -H_p \end{bmatrix}$$

where, H_p is the $2^p \times 2^p$ Hadamard matrix.

For any fixed $\boldsymbol{x} \in S^{d-1}$, consider the sum $S(\boldsymbol{x}) := \sum_{\boldsymbol{c} \in C_{RM}} (\boldsymbol{x}^T\boldsymbol{c})^2$. We first show that $S(\boldsymbol{x}) \geq 2(d+1)$.

$$\begin{aligned} S(\boldsymbol{x}) &= \sum_{\boldsymbol{c} \in C_{RM}} (\boldsymbol{x}^T\boldsymbol{c})^2 = 2\sum_{\boldsymbol{h_i} \in H_p} (\boldsymbol{x}^T\boldsymbol{h_i})^2 \\ &= 2\|H_p\boldsymbol{x}\|^2 = 2(\boldsymbol{x}^T H_p^T)(H_p\boldsymbol{x}) \\ &\overset{(i)}{=} 2d \cdot \|x\|^2 = 2d. \end{aligned}$$

(i) follows from the fact that the columns of the Hadamard matrix are mutually orthogonal and therefore, $H_p^T H_p = d \cdot I_d$, where, I_d denotes the $d \times d$ identity matrix.

By an averaging argument, it then follows that there exists at least one $\boldsymbol{c} \in C_{RM}$ such that $|\boldsymbol{x}^T\boldsymbol{c}| \geq 1$. Since for every $\boldsymbol{c} \in C_{RM}$, there exists $-\boldsymbol{c} \in C_{RM}$, we get that $x^T\boldsymbol{c} \geq 1$ for some $\boldsymbol{c} \in C_{RM}$. □

D Missing proofs from Section 5.3

Proof of Proposition 9. The proof of Proposition 9 follows directly from Lemma 2 provided the point set C_{cp} satisfies Condition (1) with $R = \sqrt{d}$. We will now prove this fact.

Since each vertex is of the form $\pm\sqrt{d}\boldsymbol{e_i}$, it follows that all the vertices of $\textsc{conv}(C_{cp})$, and hence the entire convex hull lies inside a ball of radius $\sqrt{d}$, *i.e.*, , $\textsc{conv}(C_{cp}) \subset B_d(\mathbf{0}_d, \sqrt{d})$.

To prove that the unit ball is contained in the convex hull $\textsc{conv}(C_{cp})$, we pick any arbitrary point $v \in B_d(\mathbf{0}_d, 1)$ and show that it can written as a convex combination of points in C_{cp}. The fact follows from the solution to the system of linear equations (2) given in Equation (15). Note that the solution satisfies $a_i \geq 0$ and $\sum_i a_i = 1$ for any point $\boldsymbol{v} \in B_d(0, 1)$. $\square$

Proof of Lemma 10. Let $K := \textsc{conv}(N(\varepsilon))$ be the convex hull of the ε-net points of the unit sphere. Let $B_d(\mathbf{0}_d, r)$ be the inscribed ball in K for some $r < 1$. We show that $r \geq 1 - \varepsilon$.

Consider the face of K that is tangent to $B_d(\mathbf{0}_d, r)$ at point $\boldsymbol{z}$. We will show that $\|\boldsymbol{z}\|_2 \geq 1 - \varepsilon$. Extend the line joining $(\mathbf{0}_d, \boldsymbol{z})$ to meet $\mathcal{S}^{d-1}$ at point $\boldsymbol{x}$. Since $\boldsymbol{x} \in \mathcal{S}^{d-1}$, we know that there exists a net point $\boldsymbol{u}$ at a distance of at most ε from it. Therefore, the distance of $\boldsymbol{x}$ from K is upper bounded by ε, *i.e.*, $\text{dist}(\boldsymbol{x}, K) = \|\boldsymbol{x} - \boldsymbol{z}\| \leq \|\boldsymbol{x} - \boldsymbol{u}\| \leq \varepsilon$. Therefore $\|\boldsymbol{z}\| = 1 - \|\boldsymbol{x} - \boldsymbol{z}\| \geq 1 - \varepsilon$.

Therefore scaling all the points of $N(\varepsilon)$ by any $R \geq \frac{1}{1-\varepsilon}$ we see that $B_d(\mathbf{0}_d, 1) \subseteq \textsc{conv}(C)$. $\square$

E Missing proofs from Section 6

Proof of Proposition 12. First we show that the point set C_S satisfies Condition (1) with $R = 2d$. The fact that $\textsc{conv}(C_S) \subset B_d(\mathbf{0}_d, 2d)$ follows trivially from the observation that each point in $C_S \in B_d(\mathbf{0}_d, 2d)$.

To show that $B_d(\mathbf{0}_d, 1) \subset \textsc{conv}(C_S)$, consider any face of the convex hull, $F_{\boldsymbol{c}} := \textsc{conv}(C_S \setminus \{\boldsymbol{c}\})$, for some $\boldsymbol{c} \in C_S$. We show that $F_{\boldsymbol{c}}$ is at an ℓ_2 distance of at least 1 from $\mathbf{0}_d$. This in turn shows that any point outside the convex hull must be outside the unit ball as well.

First consider the case when $\boldsymbol{c} = -4\mathbf{1}_d$. We observe that the face $F_{\boldsymbol{c}}$ is contained in the hyperplane $H_{\boldsymbol{c}} := \{\boldsymbol{x} \in \mathbb{R}^d \mid \frac{1}{\sqrt{d}} \mathbf{1}_d^T \boldsymbol{x} = 2\sqrt{d}\}$, and therefore is at a distance of $O\left(\sqrt{d}\right)$ from the origin.

Now consider the case when $\boldsymbol{c} = 2d\ \boldsymbol{e_1}$. Let $\boldsymbol{w} = \frac{1}{\sqrt{\frac{9}{16} - \frac{1}{2d}}}(-\frac{3}{4} + \frac{1}{2d}, \frac{1}{2d}, \ldots, \frac{1}{2d})^T \in \mathbb{R}^d$ be a unit vector. We note that $F_{\boldsymbol{c}} \subset H_{\boldsymbol{c}}$, where $H_{\boldsymbol{c}} := \{\boldsymbol{x} \in \mathbb{R}^d \mid \boldsymbol{w}^T \boldsymbol{x} = \frac{1}{\sqrt{\frac{9}{16} - \frac{1}{2d}}}\}$ is the hyperplane defined by the unit normal vector $\boldsymbol{w}$ that is at a distance of at least 1 from $\mathbf{0}_d$.

Since all other faces are symmetric, the proof for the case $\boldsymbol{c} = 2d\ \boldsymbol{e_i}, i \in [d]$ follows similarly.

Privacy: We now show that the quantization scheme is ϵ-differentially private for any $\epsilon > \log 7$. From the definition of ϵ-DP, it is sufficient to show that for any $\boldsymbol{x}, \boldsymbol{y} \in B_d(\mathbf{0}_d, 1)$, and every $\boldsymbol{c} \in C_S$,

$$\frac{\Pr[Q_{C_S}(\boldsymbol{x}) = \boldsymbol{c}]}{\Pr[Q_{C_S}(\boldsymbol{y}) = \boldsymbol{c}]} \leq 7.$$

Since $\boldsymbol{x}, \boldsymbol{y} \in \textsc{conv}(C_S)$, we can express them as the convex combination of points in C_S. Let $\boldsymbol{x} = \sum_{\boldsymbol{c} \in C_S} a_c^{(\boldsymbol{x})} \boldsymbol{c}$. Similarly, let $\boldsymbol{y} = \sum_{\boldsymbol{c} \in C_S} a_c^{(\boldsymbol{y})} \boldsymbol{c}$. Then, from the construction of the quantization function Q_{C_S}, we know that

$$\frac{\Pr[Q_{C_S}(\boldsymbol{x}) = \boldsymbol{c}]}{\Pr[Q_{C_S}(\boldsymbol{y}) = \boldsymbol{c}]} = \frac{a_c^{(\boldsymbol{x})}}{a_c^{(\boldsymbol{y})}}.$$

We now show that the ratio $\frac{a_c^{(\boldsymbol{x})}}{a_c^{(\boldsymbol{y})}}$ is at most 7 for *any* pair $\boldsymbol{x}, \boldsymbol{y} \in B_d(\mathbf{0}_d, 1)$ and any $\boldsymbol{c} \in C_S$. The privacy bound follows from this observation.

First, consider the case $\boldsymbol{c} = -4\mathbf{1}_d$. From the closed form solution for any $\boldsymbol{x} \in \textsc{conv}(C_S)$ described in Equation (6), we know that $a_c^{(\boldsymbol{x})} = \frac{1}{3} - \frac{\sum_{i=1}^d x_i}{6d}$. For any $\boldsymbol{x} \in B_d(\mathbf{0}_d, 1)$, $\sum_{i=1}^d x_i \in [-\|\boldsymbol{x}\|_1, \|\boldsymbol{x}\|_1] \subseteq \left[-\sqrt{d}, \sqrt{d}\right]$. Therefore, $a_c^{(\boldsymbol{x})} \in \left[\frac{1}{3} - \frac{1}{6\sqrt{d}}, \frac{1}{3} + \frac{1}{6\sqrt{d}}\right]$. It then follows that for any $\boldsymbol{x}, \boldsymbol{y} \in B_d(\mathbf{0}_d, 1)$ and $\boldsymbol{c} = -4\mathbf{1}_d$,

$$\frac{a_c^{(\boldsymbol{x})}}{a_c^{(\boldsymbol{y})}} \leq \frac{\frac{1}{3} + \frac{1}{6\sqrt{d}}}{\frac{1}{3} - \frac{1}{6\sqrt{d}}} = 1 + \frac{2}{2\sqrt{d} - 1} \leq 3$$

Now we consider the case when $\boldsymbol{c} = 2d\ \boldsymbol{e_1}$. Then from the closed from solution in Equation (6), we get that for any $\boldsymbol{x} \in \textsc{conv}(C_S)$ the coefficient $a_c^{(\boldsymbol{x})} = \frac{x_1}{2d}\left(1 - \frac{2}{3d}\right) - \frac{\sum_{i=2}^d x_i}{3d^2} + \frac{2}{3d}$. Note that this quantity is maximized for

$\boldsymbol{x} = \boldsymbol{e_1}$ and minimized for $\boldsymbol{x} = -\boldsymbol{e_1}$. Therefore the ratio for any $\boldsymbol{x}, \boldsymbol{y} \in B_d(\mathbf{0}_d, 1)$ and $\boldsymbol{c} = 2d\ \boldsymbol{e_1}$ is at most

$$\frac{a_c^{(\boldsymbol{x})}}{a_c^{(\boldsymbol{y})}} \leq \frac{7d-2}{d+2} \leq 7$$

The ratio for all other vertices can be computed in a similar fashion and is bounded by the same quantity. □

Proof of Proposition 13. First, we show that C_H satisfies Condition (1) with $R = 2d$. The fact that $\textsc{conv}(C_H) \subset B_d(\mathbf{0}_d, 2d)$ is trivial and follows since every point in C_H is contained in $B_d(\mathbf{0}_d, 2d)$.

To show that $B_d(\mathbf{0}_d, 1) \subset C_H$, consider any $\boldsymbol{x} \in B_d(\mathbf{0}_d, 1)$, and the closed form solution for the coefficients a_i given by Equation (7). We now show that these coefficients indeed give a convex combination. Note that $a_i := \frac{1}{d+1}\left(1 + \frac{\boldsymbol{c}^T\boldsymbol{x}}{4d}\right) \geq 0$. This holds since $\boldsymbol{c}^T\boldsymbol{x} \geq \|\boldsymbol{c}\|\|\boldsymbol{x}\| \geq -2d$. Moreover, from the property of Hadamard matrices,

$$\sum_{i=1}^{d+1} a_i = \frac{1}{d+1}\begin{bmatrix}1 & \dots & 1\end{bmatrix} H_p^T \begin{bmatrix}1 \\ \frac{\boldsymbol{x}}{2\sqrt{d}}\end{bmatrix} = 1.$$

The last equality follows from the following property of the Hadamard matrices that can be proved using induction.

$$\begin{bmatrix}1 & \dots & 1\end{bmatrix} H_p^T = \begin{bmatrix}2^p & 0 & \dots & 0\end{bmatrix}.$$

Therefore, any $\boldsymbol{x} \in B_d(\mathbf{0}_d, 1)$ can be expressed as a convex combination of the points in C_H, *i.e.*, , $\boldsymbol{x} = \sum_{i=1}^{d+1} a_i \boldsymbol{c_i}$, for $\boldsymbol{c_i} \in C_H$.

Privacy: We now show that the quantization scheme is ϵ-differentially private for any $\epsilon > 0.4$. From the definition of ϵ-DP, it is sufficient to show that for any $\boldsymbol{x}, \boldsymbol{y} \in B_d(\mathbf{0}_d, 1)$, and any $\boldsymbol{c} \in C_H$,

$$\frac{\Pr[Q_{C_H}(\boldsymbol{x}) = \boldsymbol{c}]}{\Pr[Q_{C_S}(\boldsymbol{y}) = \boldsymbol{c}]} \leq 1 + \sqrt{2}$$

Since $\boldsymbol{x}, \boldsymbol{y} \in \textsc{conv}(C_S)$, we can express them as the convex combination of points in C_H. Let $\boldsymbol{x} = \sum_{\boldsymbol{c} \in C_H} a_c^{(\boldsymbol{x})}$. Similarly, let $\boldsymbol{y} = \sum_{\boldsymbol{c} \in C_H} a_c^{(\boldsymbol{y})}$. Then, from the construction of the quantization function Q_{C_H}, we know that

$$\frac{\Pr[Q_{C_H}(\boldsymbol{x}) = \boldsymbol{c}]}{\Pr[Q_{C_H}(\boldsymbol{y}) = \boldsymbol{c}]} = \frac{a_c^{(\boldsymbol{x})}}{a_c^{(\boldsymbol{y})}}. \tag{10}$$

From the closed form solution in Equation (7), we know that for any $\boldsymbol{x} \in \textsc{conv}(C_H)$, the coefficient of $\boldsymbol{c}$ in the convex combination of $\boldsymbol{x}$ is given by $a_c^{(\boldsymbol{x})} = \frac{1}{d+1}\left(1 + \frac{\boldsymbol{c}^T\boldsymbol{x}}{4d}\right)$. Plugging this in Equation (10), we get

$$\frac{\Pr[Q_{C_H}(\boldsymbol{x}) = \boldsymbol{c}]}{\Pr[Q_{C_H}(\boldsymbol{y}) = \boldsymbol{c}]} = \frac{a_c^{(\boldsymbol{x})}}{a_c^{(\boldsymbol{y})}} = \frac{1 + \frac{\boldsymbol{c}^T\boldsymbol{x}}{4d}}{1 + \frac{\boldsymbol{c}^T\boldsymbol{y}}{4d}} = 1 + \frac{\frac{\boldsymbol{c}^T(\boldsymbol{x}-\boldsymbol{y})}{4d}}{1 + \frac{\boldsymbol{c}^T\boldsymbol{y}}{4d}} \tag{11}$$

$$\leq 1 + \frac{\|\boldsymbol{c}\|_2\|\boldsymbol{x}-\boldsymbol{y}\|_2}{4d - \|\boldsymbol{c}\|_2} \qquad \text{for } \boldsymbol{y} = -\frac{\boldsymbol{c}}{\|\boldsymbol{c}\|_2} \tag{12}$$

$$\leq 1 + \frac{2\sqrt{2}d}{4d - 2d} \tag{13}$$

(since $\|\boldsymbol{x}-\boldsymbol{y}\|_2 \leq \sqrt{2}$ and $\|\boldsymbol{c}\|_2 = 2d$.)

$$= 1 + \sqrt{2} \tag{14}$$

This concludes the proof of Proposition 13. □

Proof of Proposition 14. The fact that $Q_{\tilde{C}_{cp}}$ satisfies Condition (1) with $R = 2\sqrt{d}$ follows from the proof of Proposition 9. For any $v \in \mathbb{R}^d$, we can compute the convex combinations as

$$a_i = \begin{cases} \frac{v_i}{2\sqrt{d}} + \frac{\gamma}{2d} & \text{if } v_i > 0 \text{ and } i \leq d \\ -\frac{v_i}{2\sqrt{d}} + \frac{\gamma}{2d} & \text{if } v_i \leq 0 \text{ and } i > d \\ \frac{\gamma}{2d} & \text{otherwise} \end{cases} \tag{15}$$

where, $\gamma := 1 - \frac{\|\boldsymbol{v}\|_1}{2\sqrt{d}}$, is a non-negative quantity for every $\boldsymbol{v} \in B_d(\mathbf{0}_d, 1)$.
To prove the privacy guarantees of this scheme, we first state a few observations:

- Since $\|\boldsymbol{v}\|_1 \in [-\sqrt{d}, \sqrt{d}]$, the quantity $\gamma \in [1/2, 3/2]$.
- For any coordinate $i \in [d]$, if $x_i > 0$, then the coefficients $a_i = \frac{|x_i|}{2\sqrt{d}} + \frac{\gamma}{2d}$, and $a_{d+i} = \frac{\gamma}{2d}$.
- Similarly, if $x_i \leq 0$, then the coefficients $a_i = \frac{\gamma}{2d}$, and $a_{d+i} = \frac{|x_i|}{2\sqrt{d}} + \frac{\gamma}{2d}$.
- For any $\boldsymbol{x} \in B_d(\mathbf{0}_d, 1)$, $x_i \in [-1, 1]$

Let $\boldsymbol{x}, \boldsymbol{y} \in B_d(\mathbf{0}_d, 1)$ and for any $\boldsymbol{c} \in \tilde{C}_{cp}$, we need to upper bound the following quantity to prove the privacy guarantees of the scheme:

$$p_c := \frac{\Pr[Q_{\tilde{C}_{cp}}(\boldsymbol{x}) = \boldsymbol{c}]}{\Pr[Q_{\tilde{C}_{cp}}(\boldsymbol{y}) = \boldsymbol{c}]}$$

Note that it is sufficient to consider only one of the points $\boldsymbol{c} = 2\sqrt{d}\boldsymbol{e_j}$ in the following four scenarios:

1. $x_i > 0, y_i > 0$, then $p_c = \frac{\frac{|x_i|}{2\sqrt{d}} + \frac{\gamma x}{2d}}{\frac{|y_i|}{2\sqrt{d}} + \frac{\gamma y}{2d}} \leq \frac{\frac{1}{2\sqrt{d}} + \frac{3}{4d}}{\frac{1}{4d}} \leq O(\sqrt{d})$.
2. $x_i > 0, y_i \leq 0$, then $p_c = \frac{\frac{|x_i|}{2\sqrt{d}} + \frac{\gamma x}{2d}}{\frac{\gamma y}{2d}} \leq \frac{\frac{1}{2\sqrt{d}} + \frac{3}{4d}}{\frac{1}{4d}} \leq O(\sqrt{d})$.
3. $x_i \leq 0, y_i > 0$, then $p_c = \frac{\frac{\gamma x}{2d}}{\frac{|y_i|}{2\sqrt{d}} + \frac{\gamma y}{2d}} = \frac{\frac{3}{4d}}{\frac{1}{4d}} \leq 3$
4. $x_i \leq 0, y_i \leq 0$, then $p_c = \frac{\frac{\gamma x}{2d}}{\frac{\gamma y}{2d}} = \frac{\frac{3}{4d}}{\frac{1}{4d}} \leq 3$.

Therefore, the privacy guarantees hold for any $\epsilon > O(\log d)$. □

Proof of Theorem 15. First we show that $\hat{\boldsymbol{v}} = PQ_{C,\epsilon}(\boldsymbol{v}) = \frac{1}{p-q}\sum_{i=1}^{|C|}(\mathbf{1}_{\{\boldsymbol{y}=\boldsymbol{c_i}\}} - q)\boldsymbol{c_i}$ is an unbiased estimator of v. From linearity of expectations, we have

$$\mathbf{E}[\hat{\boldsymbol{v}}] = \frac{1}{p-q}\sum_{i=1}^{|C|}(\Pr[\boldsymbol{y} = \boldsymbol{c_i}] - q)\boldsymbol{c_i}, \tag{16}$$

where, the expectation is taken over the randomness of both the quantization and RR scheme. Recall that

$$\boldsymbol{y} := \text{RR }_p(Q_C(\boldsymbol{v}), C) \in C,$$

where $p = \frac{e^\epsilon}{e^\epsilon + |C| - 1}$. Therefore,

$$\begin{aligned}\Pr(\boldsymbol{y} = \boldsymbol{c_i}) &= \sum_{j=1}^{|C|}\Pr[\boldsymbol{y} = \boldsymbol{c_i}|Q_C(\boldsymbol{v}) = \boldsymbol{c_j}] \cdot \Pr[Q_C(\boldsymbol{v}) = \boldsymbol{c_j}] \\ &= (p-q)a_i + q.\end{aligned}$$

Therefore $\mathbf{E}[\hat{\boldsymbol{v}}] = \frac{1}{p-q}\sum_{i=1}^{|C|}(p-q)a_i\boldsymbol{c_i} = \boldsymbol{v}$.

Now we bound the variance of the estimator

$$\begin{aligned}
\mathbf{E}[\|\boldsymbol{v}-\hat{\boldsymbol{v}}\|^2] &= \mathbf{E}\left[\|\sum_{i=1}^{|C|}\left(\frac{1}{p-q}(\mathbf{1}_{\{\boldsymbol{y}=\boldsymbol{c_i}\}}-q)-a_i\right)\boldsymbol{c_i}\|^2\right] \\
&\le \sum_{i=1}^{|C|}\mathbf{E}\left[\left(\frac{1}{p-q}(\mathbf{1}_{\{\boldsymbol{y}=\boldsymbol{c_i}\}}-q)-a_i\right)^2\|\boldsymbol{c_i}\|^2\right] \\
&= \sum_{i=1}^{|C|}\textsc{Var}\left[\left(\frac{1}{p-q}(\mathbf{1}_{\{\boldsymbol{y}=\boldsymbol{c_i}\}}-q)\right)\|\boldsymbol{c_i}\|^2\right] \\
&= \left(\frac{1}{p-q}\right)^2\sum_{i=1}^{|C|}\textsc{Var}(\mathbf{1}_{\{\boldsymbol{y}=\boldsymbol{c_i}\}})\|\boldsymbol{c_i}\|^2 \\
&= O(|C|R^2),
\end{aligned}$$

since $\|c_i\|^2 \le R^2$ and $\textsc{Var}(\mathbf{1}_{\{\boldsymbol{y}=\boldsymbol{c_i}\}}) \le 1/4$.

Privacy Now we show that our scheme is ϵ differentially private where ϵ is the input parameter to the RR algorithm. For any two points $\boldsymbol{v},\boldsymbol{w} \in B_d(\mathbf{0}_d,1)$,

$$\frac{PQ_{C,\epsilon}(\boldsymbol{v})=y}{PQ_{C,\epsilon}(\boldsymbol{w})=y} = \frac{\sum_{i=1}^{|C|}\Pr(y|Q_C(\boldsymbol{v})=\boldsymbol{c_i})\Pr(Q_C(\boldsymbol{v})=\boldsymbol{c_i})}{\sum_{j=1}^{|C|}\Pr(y|Q_C(\boldsymbol{w})=\boldsymbol{c_j})\Pr(Q_C(\boldsymbol{w})=\boldsymbol{c_j})} \tag{17}$$

$$\le \frac{\max_i\Pr(y|Q_C(\boldsymbol{v})=\boldsymbol{c_i})\sum_{i=1}^{|C|}\Pr(Q_C(\boldsymbol{v})=\boldsymbol{c_i})}{\min_j\Pr(y|Q_C(\boldsymbol{w})=\boldsymbol{c_j})\sum_{i=1}^{|C|}\Pr(Q_C(\boldsymbol{w})=\boldsymbol{c_j})} \tag{18}$$

$$= \frac{\max_i\Pr(y|Q_C(\boldsymbol{v})=\boldsymbol{c_i})}{\min_j\Pr(y|Q_C(\boldsymbol{w})=\boldsymbol{c_j})} \le e^{\epsilon} \tag{19}$$

we are using the following privacy property of Randomized Rounding [36] mechanism in Equation (19)

$$\sup_{i,j}\frac{\Pr(y|Q_C(\boldsymbol{v})=\boldsymbol{c_i})}{\Pr(y|Q_C(\boldsymbol{w})=\boldsymbol{c_j})} \le e^{\epsilon} \quad \forall \boldsymbol{v},\boldsymbol{w}$$

□

Proof of Theorem 16. First we show that $\hat{\boldsymbol{v}} = \frac{1}{(1-2p)}\sum_{j=1}^{|C|}(y_j-p)\,\boldsymbol{c_j}$ is an unbiased estimator of v. From linearity of expectations, we have

$$\mathbf{E}[\hat{\boldsymbol{v}}] = \frac{1}{(1-2p)}\sum_{j=1}^{|C|}(\mathbf{E}[y_j]-p)\,\boldsymbol{c_j}, \tag{20}$$

where, the expectation is taken over the randomness of both the quantization and RAPPOR scheme.
Recall that

$$\boldsymbol{y} := \textsc{Rappor}\ _p(\text{1-}\textsc{hot}\ (Q_C(\boldsymbol{v}),C)) \in \{0,1\}^{|C|}.$$

Each entry of the vector $\boldsymbol{y}$ is an independent binary random variable and

$$\begin{aligned}
\mathbf{E}[y_j] = \Pr(y_j=1) &= \sum_{i=1}^{|C|}\Pr(y_j,Q_C(\boldsymbol{v})=\boldsymbol{c_i}) \\
&= \sum_{i=1}^{|C|}Pr(y_j|Q_C(\boldsymbol{v})=\boldsymbol{c_i})Pr(Q_C(\boldsymbol{v})=\boldsymbol{c_i}) \\
&= Pr(y_j|Q_C(\boldsymbol{v})=\boldsymbol{c_j})Pr(Q_C(\boldsymbol{v})=\boldsymbol{c_j}) \\
&+ \sum_{i\neq j}^{|C|}Pr(y_j|Q_C(\boldsymbol{v})=\boldsymbol{c_i})Pr(Q_C(\boldsymbol{v})=\boldsymbol{c_i}) \\
&= (1-p)a_j + p(1-a_j) = p + (1-2p)a_j.
\end{aligned} \tag{21}$$

Plugging Equation (21) in Equation (20) , we get

$$\mathbf{E}(\hat{\boldsymbol{v}}) = \frac{1}{(1-2p)}\sum_{j=1}^{|C|}(p+(1-2p)a_j-p)\,\boldsymbol{c_j} = \sum_{j=1}^{|C|} a_j\boldsymbol{c_j} = \boldsymbol{v} \tag{22}$$

Now we show a bound on the variance of the estimate

$$\begin{aligned}
\mathbf{E}\left[\|\boldsymbol{v}-\hat{\boldsymbol{v}}\|_2^2\right] &= \mathbf{E}\left[\|\sum_{j=1}^{|C|} a_j\boldsymbol{c_j} - \frac{1}{(1-2p)}\sum_{j=1}^{|C|}(y_j-p)\,\boldsymbol{c_j}\|_2^2\right] && (23)\\
&= \sum_{j=1}^{|C|}\mathbf{E}\left(a_j - \frac{(y_j-p)}{(1-2p)}\right)^2 |\boldsymbol{c_j}|^2 && (24)
\end{aligned}$$

(all the cross terms are 0 as they are mutually independent and $\mathbf{E}\left(a_j - \frac{y_j-p}{1-2p}\right) = 0$)

$$\begin{aligned}
&= \sum_{j=1}^{|C|}\mathrm{var}\left(\frac{y_j-p}{1-2p}\right)|\boldsymbol{c_j}|^2 && (25)\\
&= \left(\frac{1}{1-2p}\right)^2\sum_{j=1}^{|C|}\mathrm{var}(y_j)\,|\boldsymbol{c_j}|^2 = O(|C|R^2) && (26)
\end{aligned}$$

Equation (26) comes form the fact that y_j is a binary random variable and $Var(y_j) = Pr(y_j)(1-Pr(y_j)) \leq \frac{1}{4}$ and $|\boldsymbol{c_j}|^2 \leq R^2$.

Privacy Now we show that our scheme is ϵ differentially private where ϵ is the input parameter to the RAPPOR algorithm. For any two points $\boldsymbol{v}, \boldsymbol{w} \in B_d(\mathbf{0}_d, 1)$,

$$\begin{aligned}
\frac{PQ_{C,\epsilon}(\boldsymbol{v}) = y}{PQ_{C,\epsilon}(\boldsymbol{w}) = y} &= \frac{\sum_{i=1}^{|C|}\Pr(y|Q_C(\boldsymbol{v}) = \boldsymbol{c_i})\Pr(Q_C(\boldsymbol{v}) = \boldsymbol{c_i})}{\sum_{j=1}^{|C|}\Pr(y|Q_C(\boldsymbol{w}) = \boldsymbol{c_j})\Pr(Q_C(\boldsymbol{w}) = \boldsymbol{c_j})} && (27)\\
&\leq \frac{\max_i \Pr(y|Q_C(\boldsymbol{v}) = \boldsymbol{c_i})\sum_{i=1}^{|C|}\Pr(Q_C(\boldsymbol{v}) = \boldsymbol{c_i})}{\min_j \Pr(y|Q_C(\boldsymbol{w}) = \boldsymbol{c_j})\sum_{i=1}^{|C|}\Pr(Q_C(\boldsymbol{w}) = \boldsymbol{c_j})} && (28)\\
&= \frac{\max_i \Pr(y|Q_C(\boldsymbol{v}) = \boldsymbol{c_i})}{\min_j \Pr(y|Q_C(\boldsymbol{w}) = \boldsymbol{c_j})} \leq e^{\epsilon} && (29)
\end{aligned}$$

By the privacy property of RAPPOR [15] mechanism , we are using the following fact in equation (29)

$$\sup_{i,j}\frac{\Pr(y|Q_C(\boldsymbol{v}) = \boldsymbol{c_i})}{\Pr(y|Q_C(\boldsymbol{w}) = \boldsymbol{c_j})} \leq e^{\epsilon} \quad \forall \boldsymbol{v}, \boldsymbol{w}$$

Communication : Now we show that for the RAPPOR based scheme the expected communication is linear in $|C|$. Say y is the output when RAPPOR is applied to one hot encoded binary string. Without loss of generality say the the bit string is $\boldsymbol{e_i}$. The output y is generated as follows

$$\Pr(y_j = 1) = \begin{cases} p & \text{if } j \neq i \\ (1-p) & \text{if } j = i \end{cases}$$

So the expected sparsity (l_0 norm) of the output is

$$\begin{aligned}
\mathbf{E}[\|y\|_0] &= \sum_{i}^{|C|} y_i = (|C|-1)p + (1-p)\\
&= |C|p + (1-2p) = O(|C|)
\end{aligned}$$

□