

Quantization of distributed data for learning

Osama A. Hanna, Yahya H. Ezzeldin, Christina Fragouli and Suhas Diggavi

Abstract—We consider machine learning applications that train a model by leveraging data distributed over a trusted network, where communication constraints can create a performance bottleneck. A number of recent approaches propose to overcome this bottleneck through compression of gradient updates. However, as models become larger, so does the size of the gradient updates. In this paper, we propose an alternate approach to learn from distributed data that quantizes data instead of gradients, and can support learning over applications where the size of gradient updates is prohibitive. Our approach leverages the dependency of the computed gradient on data samples, which lie in a much smaller space in order to perform the quantization in the smaller dimension data space. At the cost of an extra gradient computation, the gradient estimate can be refined by conveying the difference between the gradient at the quantized data point and the original gradient using a small number of bits. Lastly, in order to save communication, our approach adds a layer that decides whether to transmit a quantized data sample or not based on its importance for learning. We analyze the convergence of the proposed approach for smooth convex and non-convex objective functions and show that we can achieve order optimal convergence rates with communication that mostly depends on the data rather than the model (gradient) dimension. We use our proposed algorithm to train ResNet models on the CIFAR-10 and ImageNet datasets, and show that we can achieve an order of magnitude savings over gradient compression methods. These communication savings come at the cost of increasing computation at the learning agent, and thus our approach is beneficial in scenarios where communication load is the main problem.

I. INTRODUCTION

Consider a machine learning application, where an agent wishes to train a model using data instances that are collected (and streamed potentially in real time) from a set of distributed terminal nodes. We assume that the data is generated/collected at distributed nodes which do not have enough data to locally build good learning models. The scenario we consider is for a trusted network, so there are no privacy constraints on the data. Application scenarios for this include wireless sensor networks and IoT networks that collect data which is to be used for learning, e.g., an autonomous vehicle collecting data from by-the-road sensors, and machine learning techniques used to overcome routing and tracking problems in wireless sensor networks [1], [2], [3]. Unlike work that aim for parallelizing the computation, we assume a setup where the learning is

O. A. Hanna, C. Fragouli and S. Diggavi are with the Electrical and Computer Engineering Department at the University of California, Los Angeles, CA 90095 USA (e-mail: {ohanna, christina.fragouli, suhasdiggavi}@ucla.edu). Yahya H. Ezzeldin was with the Electrical and Computer Engineering Department at the University of California, Los Angeles, CA 90095 USA. He is now with the Electrical and Computer Engineering Department at the University of Southern California, CA 90089 USA (e-mail: yessa@usc.edu).

This work was partially supported by NSF grants #2007714, #1955632 and by UC-NL grant LFR-18-548554 and by Army Research Laboratory under Cooperative Agreement W911NF-17-2-0196.




Fig. 1. Building intuition for data sampling: training on all samples vs. only samples inside the triangle leads to classifiers with approximately the same performance.

done centrally by an agent (server) from data that is naturally distributed. As a result, the key challenge that we address in this work, is that the terminal nodes are connected to the server through a weak communication fabric; for instance, through wireless, bandwidth constrained links. The question we ask is: how can we compress the data so as to reduce the number of bits communicated, without affecting the learning performance of the agent.

Our main observation is that, although the nodes may continuously be observing data, not all the data is equally useful - perhaps a subset of the data is all we need. This can be thought of as a generalization of the notion of support vectors in Support Vector Machines (SVM); effectively, we ask, which of the data instances are most helpful for learning. As a simple example, we experimentally verified that training a neural network classifier on the data instances inside the triangle in Figure 1 yields approximately the same validation accuracy as when we train the classifier on the full set of data instances. The challenge is that, we would like to decide which are the instances to communicate in an online manner, without having access to the whole dataset in advance.

How useful is an observed data instance depends on the corresponding stage of learning (eg., no data is useful if the model is already learned) and on the algorithm the agent uses for model training. In this paper we assume that the agent uses stochastic gradient descent, which is by far the most popular algorithm used for model training today.

The fact that we restrict our attention to learning through stochastic gradient descent, naturally raises the following question. Leveraging the literature on federated learning, we could perhaps ask the terminal nodes to calculate gradients on their observed data, and then compress the gradient updates, following a technique such as in [4], [5], [6], [7], [8], [9]. Indeed, this is a reasonable solution (and we compare against it in our evaluation section VI), but has a major drawback: the size of models (and gradient updates) gets larger and larger - while the dimension of data does not. Indeed, in the pursuit of improved accuracy, neural networks have evolved from thousands to millions to hundreds of millions of parameters [10], [11], [12]. However, the dimension (number of features) of

data-instances themselves, whether images, natural signals, or sensory data, has not and is not expected to change much. Given that, unlike the federated learning setup, our data is not private, and our agent does not require computational support, it makes sense to leverage the lower dimensionality of data to achieve lower communication cost. While there is work that aims to prune and reduce the size of machine learning models [13], [14], our methods are still useful in those cases whenever the data dimension is smaller than the model dimension.

In this paper we propose a scheme that is based on quantizing data instances. The proposed compression approach consists of two parts: (1) a selection scheme and (2) a quantization scheme. The selection scheme first decides which data instances are communicated depending on how important they are to learning; afterwards, these selected instances are quantized using our quantization scheme. The quantization scheme combines aspects of dataset quantization and gradient compensation (if nodes have enough computational capabilities), and aims to enable the learning agent to perform gradient updates with low (or no) performance loss. The optional gradient compensation step, in our scheme, enhances the learning performance by communicating only a few number of bits but comes at the cost of a gradient computation at a sensor node. We find that our scheme uses $\mathcal{O}(d \log_2(h))$ bits in the worst case to achieve the optimal convergence rate, where d and h are the dataset and model dimensions, respectively¹. For large models ($h \gg d$), this is much smaller than $\mathcal{O}(h)$, the information theoretic lower bound derived in [9] on the number of bits needed to convey quantized gradients to achieve optimal convergence rate. This is because we take advantage of the data dependency of gradients, which is ignored by the oracle used for the lower bounds in [9].

Our main contributions are:

- We exploit the fact that computed gradients are dependent on the datapoints to design DaQuSGD, a quantization approach that only communicates $\mathcal{O}(d \log_2 h)$ bits in the worst case per iteration, where d is the dimension of a single datapoint and h is the dimension of the model parameters.
- We additionally exploit the fact that only a small portion of datapoints could be sufficient for a reliable model update in order to design an online sample selection approach that acts as a preliminary sampling step to DaQuSGD. Our sample selection scheme provides savings in both communication and computation for the same performance.
- We analyze the convergence of our scheme for convex and non-convex objective functions, and prove that it achieves a similar convergence rate as unquantized SGD.
- We numerically show the benefits of using DaQuSGD for training ResNet models for CIFAR-10 and ImageNet datasets. We achieve accuracy comparable to unquantized stochastic gradient descent, while using an order of magnitude less bits compared to gradient quantization schemes.

Related work: Stochastic gradient descent has been very successful in training neural networks [4], [5], [15]. A number of works proposed quantized versions of stochastic gradient

descent, see for example [4], [5], [16], [17], [18], [6], [7], [8], [9]; these are based on gradient quantization and typically use $\mathcal{O}(h)$ bits/iteration to achieve the same convergence rate as the unquantized gradient descent, which is proved to be a fundamental limit in [9]² for memoryless gradient quantizers. Using memory with quantizers, sparsification methods were considered [19], [20], [21], [22], [23], that improve communication efficiency at the cost of increasing the gradient iterations until convergence³.

Data quantization was considered in [24] for generalized linear models learning. However, in a generalized linear model, the gradient dimension is the same as the data dimension and the gradient in this case is only a scaled version of the data instance. In this paper, we consider data quantization for a general learning model. Our work is the first, as far as we know, that leverages dataset-based quantization to achieve the optimal convergence rate at a communication cost of $\mathcal{O}(d \log_2(h))$ that mostly depends on the dataset and not the model dimension. To further reduce communication cost, we propose a sample selection algorithm. There is a rich body of work on sample selection in centralized settings [25], [26], [27], [28], however, these are either computationally expensive, cannot be applied to deep neural networks, require the knowledge of full dataset, or cannot be applied in a distributed setting. For example, if every node has only one sample, applying the CRAIG algorithm proposed in [27] locally on each node’s data would result in transmitting all the samples. Schemes that are based on influence function [29], [30], [31] are also related. Influence functions approximate the effect of training samples on the model predictions during testing. This requires the computation of the gradient and Hessian of the model, which can be computationally expensive in deep networks. Moreover, influence functions are not well understood for deep models and non-convex functions [32] and are shown to be fragile in deep learning [33]. We present and analyze a simple sample selection scheme that requires only the computation of the forward path of the model. Importance sampling, such as [28], studies weighted sampling methods with the goal of reducing the stochastic gradient variance. However, in such work, a point is going to be sampled in every iteration (but with a non-uniform optimized distribution). As a result, offering no communication cost saving. Furthermore, [28] also requires knowledge of non-homogeneous upper bounds on the per sample gradient. Such knowledge is not currently feasible for deep models; this is also the reason why their approach [28] is only evaluated using an SVM model. Subsampling has also been utilized in [34] in order to reduce the size of required dataset augmentation. The proposed schemes in [34] require full knowledge of the dataset and targets to improve the model accuracy, in contrast to our goal of reducing the communication load. Our sample selection scheme can be applied for deep networks in distributed settings. This approach offers reduction in both communication and computation costs without sacrificing performance.

²The reason why the lower bound in [9] can be broken is discussed in Section III.

³Effective implementations of these also use $\mathcal{O}(h)$ bits per gradient iteration.

¹Our theoretical analysis relies on a smoothness assumption of the gradient as a function of the data point; that is, if the true and quantized data points are close, then their gradients are also close.

There is a large body of research on active learning that aims to perform learning with few labeled instances, if unlabeled data are available but labels are expensive to obtain, by selecting few unlabeled instances to be labeled by an oracle [35], [36], [37], [38]. In contrast, all our data are labeled and we aim to decide what samples to not transmit to save communication.

Our work also differs from traditional signal compression [39]: our quantization does not aim to maintain the ability to reconstruct data, but instead the ability to perform gradient updates with low (or no) performance loss.

Notations: In the rest of the paper, we use the following notation convention. We denote with $[a : b]$ the set of integers from a to $b \geq a$; we also use the shorthand $[b]$ to denote integers from 1 to b . We use lowercase letters for scalars, uppercase letters for constants, and lowercase boldface letters for vectors. Calligraphic upper letters are used to denote sets of vectors or scalars.

II. SETUP

A learning agent wants to use data from a space $\mathcal{Z} \subseteq \mathbb{R}^d$ to generate a model in a space $\mathcal{W} \subseteq \mathbb{R}^h$. The data samples are collected at distributed nodes that can communicate with the agent using at most r bits per transmission. The overall dataset $\mathcal{D} = \{\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(N)}\}$ consists of N samples from the space \mathcal{Z} which is assumed to be bounded, i.e., $\|\mathbf{z}\|_2 \leq B$; we refer to such \mathcal{Z} as being B -bounded. We also assume that \mathcal{W} is bounded such that $\|\mathbf{w} - \mathbf{w}'\|_2^2 \leq D^2, \forall \mathbf{w}, \mathbf{w}' \in \mathcal{W}$; we refer to such \mathcal{W} as being D -relatively bounded. The objective of the learning agent is to minimize the empirical risk of the output model. For a given model $\mathbf{w} \in \mathcal{W}$, this risk is given by

$$L(\mathbf{w}) = \frac{1}{N} \sum_{i=1}^N \ell(\mathbf{w}, \mathbf{z}^{(i)}), \quad (1)$$

where $\ell : \mathcal{W} \times \mathcal{Z} \rightarrow \mathbb{R}^+$ is a loss function. The loss function ℓ is known at all nodes in the system.

Learning algorithm. We are interested in gradient-based learning algorithms. Let $\nabla L(\mathbf{w})$ be the gradient of $L(\mathbf{w})$ and $g_{\mathbf{z}}(\mathbf{w}) = \nabla \ell(\mathbf{w}, \mathbf{z})$ be the gradient of the loss function (w.r.t. \mathbf{w}). It is well known that if $\mathbf{z} \in \mathcal{S}$ is sampled uniformly at random, $g_{\mathbf{z}}(\mathbf{w})$ is an unbiased stochastic gradient of $L(\mathbf{w})$, i.e., satisfies $\mathbb{E}[g_{\mathbf{z}}(\mathbf{w})] = \nabla L(\mathbf{w})$. The gradient $g_{\mathbf{z}}(\mathbf{w})$ can be calculated at the node sampling \mathbf{z} from the dataset. Let $\widehat{g}_{\mathbf{z}}(\mathbf{w})$ be a general stochastic estimate of the gradient $\nabla L(\mathbf{w})$, calculated by the learning agent. This can be the true stochastic gradient $g_{\mathbf{z}}(\mathbf{w})$ estimated from a datapoint \mathbf{z} , or a low-precision mapping of $g_{\mathbf{z}}(\mathbf{w})$ [7], [40], [41], [6], [9], or a function of both $g_{\mathbf{z}}(\mathbf{w})$ and \mathbf{z} . Using this estimate, the model is updated at iteration $j + 1$ as

$$\mathbf{w}^{(j+1)} = \mathbf{w}^{(j)} - \eta \widehat{g}_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)}), \quad (2)$$

where η is the learning rate.

System operational properties. We assume that the learning agent (sometimes denoted as the server) and the distributed nodes make up a trusted network with the server being the strong capable entity at its center. As a result, we assume no privacy requirements and no computational constraints at the server. Additionally, the distributed nodes are connected to the

server through a weak wireless communication fabric. Thus, our focus, in this work, is on how to quantize data such that the learning algorithm converges with the same unquantized rate while reducing the communication overhead⁴.

Quantization. At the j -th iteration, a node s_j has access to a datapoint $\mathbf{z}^{(j)}$. The learning agent generates the stochastic gradient $\widehat{g}_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j-1)})$ by receiving information from the s_j -th node through an r -bit quantizer Q_j , where we assume that Q_j has access to both the current datapoint $\mathbf{z}^{(j)}$, and the latest model $\mathbf{w}^{(j-1)}$. An r -bit quantizer consists of mappings (Q_j^e, Q_j^d) , with an encoder mapping $Q_j^e : \mathcal{Z} \times \mathcal{W} \rightarrow \{0, 1\}^{\bar{r}}$ and a decoder mapping $Q_j^d : \{0, 1\}^{\bar{r}} \rightarrow \mathbb{R}^h$, where $\bar{r} \in \{0, r\}$. Note that in our setup we allow the encoder to transmit zero bits (by picking $\bar{r} = 0$) for some samples, which corresponds to not sending the sample. The overall quantizer $Q_j = Q_j^d \circ Q_j^e$ captures the combined effect of the encoder and decoder. Note that in the quantization schemes that directly quantize gradients, the encoder mapping Q_j^e first computes the stochastic gradient $g_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j-1)})$ from the unquantized datapoint $\mathbf{z}^{(j)}$ and then maps $g_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j-1)})$ to r bits. In this paper, we explore an alternative approach where Q_j^e can make use of the datapoint, $\mathbf{z}^{(j)}$, directly. We highlight that we do not assume any extra information at the distributed nodes, i.e., they still only have access to the latest model and a full-precision datapoint.

Assumptions on the loss function. In the following sections, we assume that the function $\ell(\mathbf{w}, \mathbf{z})$ is C_w -smooth in \mathbf{w} for all $\mathbf{z} \in \mathcal{Z}$, i.e., the gradient $g_{\mathbf{z}}(\mathbf{w}) = \partial \ell(\mathbf{w}, \mathbf{z}) / \partial \mathbf{w}$ is C_w -Lipschitz continuous in \mathbf{w} for all $\mathbf{z} \in \mathcal{Z}$. In other words

$$\|g_{\mathbf{z}}(\mathbf{w}) - g_{\mathbf{z}}(\mathbf{w}')\|_2 \leq C_w \|\mathbf{w} - \mathbf{w}'\|_2, \quad \forall \mathbf{z} \in \mathcal{Z}. \quad (3)$$

Moreover, the gradient $g_{\mathbf{z}}(\mathbf{w})$ is assumed to be C_z -Lipschitz continuous in \mathbf{z} for all $\mathbf{w} \in \mathcal{W}$, i.e.,

$$\|g_{\mathbf{z}}(\mathbf{w}) - g_{\mathbf{z}'}(\mathbf{w})\|_2 \leq C_z \|\mathbf{z} - \mathbf{z}'\|_2, \quad \forall \mathbf{w} \in \mathcal{W}. \quad (4)$$

While (3) is a standard assumption [7], [23], [9], the assumption in (4) is particular to our setup as we relate gradient updates to the datapoints. For this, we need that if a quantized point \mathbf{z}_Q is close to its original point \mathbf{z} , then the gradients $g_{\mathbf{z}}(\mathbf{w}), g_{\mathbf{z}_Q}(\mathbf{w})$ are also close. We employ the assumption of Lipschitz continuity of the gradient in \mathbf{z} in (4) to formalize this notion. Note that this assumption is implied if $\|\frac{\partial g_{\mathbf{z}}(\mathbf{w})}{\partial \mathbf{z}}\|_2 \leq C_z$ [42], which is the case in many loss functions (for instance, in the loss functions in the examples in Section III below) since standard theoretical analysis in learning theory assume a bounded space for \mathbf{z} and \mathbf{w} [43]. We also highlight that in many cases (as in the loss function in the Example 2 in Section III below, where $C_z = \mathcal{O}(h)$), C_z might depend on h or d . However, we will show that the effect of C_z can be removed from the convergence rate with an increase in the number of bits that is $\mathcal{O}(d \log_2(C_z))$. We observed in our numerical evaluation, on the CIFAR-10 dataset with ResNet-18, that our DaQuSGD approach gives a ratio $\frac{\|g_{\mathbf{z}}(\mathbf{w}) - g_{\mathbf{z}'}(\mathbf{w})\|_2}{\|\mathbf{z} - \mathbf{z}'\|_2}$ that is upper bounded by 1 when calculated using the sampled \mathbf{z} and taking \mathbf{z}' as its quantized version \mathbf{z}_Q .

⁴Please note a general fact in wireless that uplink is much more costly than downlink as the latter is a single broadcast transmission while uplink requires multiple transmissions depending on the number of nodes in the system.

III. PRELIMINARIES AND MOTIVATION

A. Convergence of stochastic gradient descent.

In our setup, the learning agent aims to learn a hypothesis of dimension h , using stochastic estimates of the gradient of the empirical risk function L , where these estimates are conveyed with a very low precision representation $\{0, 1\}^R$, that is, $R \ll h$. In this work, we are interested in providing convergence guarantees when learning from low precision estimates.

It is well known that for smooth convex risk functions, if stochastic gradient descent is performed with *unbiased* stochastic gradients that have *bounded variance* σ^2 , then it converges with $\mathcal{O}\left(\frac{\sigma}{\sqrt{n}}\right)$ [42]. The exact theorem is provided in Appendix A for completeness. It was also shown in [44] that under the mentioned assumptions, $\mathcal{O}\left(\frac{\sigma}{\sqrt{n}}\right)$ is a lower bound on the convergence rate of stochastic gradient descent. Hence, throughout the paper we refer to $\mathcal{O}\left(\frac{1}{\sqrt{n}}\right)$ as the order-optimal convergence rate; we take away that to achieve it, it suffices to construct a quantized stochastic gradient $\hat{g}_z(\mathbf{w})$ that satisfies the unbiasedness and bounded variance, where the variance bound should be $\mathcal{O}(\sigma^2)$ (cannot grow with other system parameters such as h, d). We follow this well known technique to prove the order-optimal convergence of DaQuSGD in Section IV. Similar to the smooth convex risk functions case, for smooth non-convex risk functions, the stochastic gradient descent converges to a local optimal point provided that the stochastic gradients are unbiased and have bounded variance. The main theorem statement [15], is reiterated in Appendix A for completeness.

B. Direct gradient quantization is not always efficient.

The work in [9] shows that if gradients $g_z(\mathbf{w})$ of the risk function $L(\mathbf{w})$ lie in a unit ℓ_2 ball and the quantizer has no memory, then the minimum number of bits required for quantizing gradient to guarantee optimal convergence rate (same as unquantized gradient descent) is $\Omega(h)$ bits per iteration; this lower bound is derived using an oracle with access only to $g_z(\mathbf{w})$ and not \mathbf{z} . For memory-based quantizers, practical implementations typically require $\Omega(h)$ bits per iteration, even when sparsification methods [5], [45], [19] are used with quantization [23].

This subsection is dedicated to examples where the $\Omega(h)$ lower bound does not hold. In these examples, we exploit the structure of the gradient to develop efficient custom quantization strategies that break the $\Omega(h)$ lower bound. These examples do not suffer from the $\Omega(h)$ lower bound since the possible gradient values do not span the whole unit ℓ_2 ball; the gradients live in a low dimensional space through their implicit dependence on \mathbf{z} . Thus, this subsection motivates why in the following Section IV, we develop an approach that relies on data quantization instead of gradient quantization, while still aiming to achieve the optimal convergence rate.

Example 1: Consider the simple case where $d = 1, h \gg d$, and the loss function is $\ell(\mathbf{w}, \mathbf{z}) = \frac{\mathbf{z}}{\sqrt{h}} \mathbf{1}^T \mathbf{w}$, where $\mathcal{W} = \{\mathbf{w} \in \mathbb{R}^h \mid \|\mathbf{w}\|_2 \leq 1\}$, and $\mathcal{Z} = \{\mathbf{z} \in \mathbb{R} \mid |\mathbf{z}| \leq 1\}$. In this case, the best bound we can get is $\|g_z(\mathbf{w})\|_2 \leq 1$. Consider a quantizer

that only sends 1 bit per data sample to the learning agent and operates as follows. At each time j , Q_j^e, Q_j^d are chosen to be

$$Q_j^e = Q^e(x) = \begin{cases} 1 & \text{if with probability } \frac{x+1}{2} \\ 0 & \text{if with probability } \frac{1-x}{2} \end{cases},$$

$$Q_j^d(x) = Q^d(x) = 2x - 1. \quad (5)$$

The stochastic gradient is chosen to be $\frac{\mathbf{z}_Q^{(j)}}{\sqrt{h}} \mathbf{1}$, where $\mathbf{z}_Q^{(j)}$ is the output of decoder Q_j^d . It is not difficult to see that the constructed gradient is an unbiased estimate of $\nabla L(\mathbf{w})$ with variance that is bounded by 1. Hence, using this unbiased estimate, we can achieve the optimal convergence rate. In this example, the gradient takes values only on the line segment $\left\{\frac{\mathbf{z}}{\sqrt{h}} \mathbf{1} \mid |\mathbf{z}| \leq 1\right\}$; it does not span the whole unit ℓ_2 ball.

Example 2 Let us consider the loss function $\ell(\mathbf{w}, \mathbf{z}) = \frac{1}{\sqrt{h}} f(\mathbf{w}^T \mathbf{v}(\mathbf{z}))$, where

$$\mathbf{v}(\mathbf{z}) = [1, \mathbf{z}, \mathbf{z}^2, \dots, \mathbf{z}^{h-1}],$$

$\mathcal{W} = \{\mathbf{w} \in \mathbb{R}^h \mid \|\mathbf{w}\|_2 \leq 1\}$, $\mathcal{Z} = \{\mathbf{z} \in \mathbb{R} \mid |\mathbf{z}| \leq 1\}$, and f is a general 1-smooth function. For instance, f could be the logistic regression loss, i.e.,

$$f(\mathbf{w}^T \mathbf{v}(\mathbf{z})) = f_y(\mathbf{z}^T \mathbf{v}(\mathbf{z})) = \log(1 + \exp(-y\mathbf{w}^T \mathbf{v}(\mathbf{z}))).$$

We here assume that $d = 1$ for simplicity, and include the $d > 1$ generalization at the end of this part. In this case, the best bound we can get on the norm of the stochastic gradient is $\|g_z(\mathbf{w})\|_2 \leq 1$ as \mathbf{z} is allowed to take the value $\mathbf{z} = 1$.

We show next that the optimal convergence rate can be achieved with only $\log_2(h)$ bits per sample, instead of the expected $\Omega(h)$ bits. Let $Q^e : [-1, 1] \rightarrow \{0, 1\}$, $Q^d : \{0, 1\} \rightarrow [-1, 1]$ be defined as in (5). The algorithm runs as follows. Each node transmits $1 + \lceil \log_2(h) \rceil$ bits representing the set of values $\{Q^e(f'(\mathbf{w}^T \mathbf{v}(\mathbf{z})))\} \cup \{Q^e(\mathbf{z}^{2^j})\}_{j=1}^{\lceil \log_2(h) \rceil}$, where f' is the derivative of the function f . Let $b_j(i)$ be the j -th least significant bit in the binary representation of the integer i . The learning agent constructs $Q(\mathbf{z}^i) = \prod_{j=1}^{\lceil \log_2(h) \rceil} Q^d \circ Q^e(\mathbf{z}^{2^j})^{b_j(i)} \forall i \in [0 : k]$ and the quantized stochastic gradient $\hat{g}_z(\mathbf{w}) = \frac{1}{\sqrt{h}} Q^d \circ Q^e(f'(\mathbf{w}^T \mathbf{v}(\mathbf{z}))) [1, Q(\mathbf{z}^1), \dots, Q(\mathbf{z}^k)]^T$. In Appendix B, we show that this algorithm achieves the optimal convergence rate and that in this example $\Omega(\log(h))$ is a lower bound on the communication cost for any SGD algorithm that achieves the optimal convergence rate.

Extension for $d > 1$: If $d > 1$, the algorithm can be extended by sending a quantized version of $f'(\mathbf{w}^T \mathbf{v}(\mathbf{z}))$ together with a quantized version of $[\mathbf{z}_1^{2^j}, \dots, \mathbf{z}_d^{2^j}]$ for each $j \in [1 : \lceil \log_2(k) \rceil]$ with d bits using the DataQ with stochastic quantization that is described in Section IV below. This would require $1 + d \lceil \log_2(k) \rceil$ bits in total. Note that when $d > 1$, each element of $\mathbf{v}(\mathbf{z})$ corresponds to a monomial $\prod_{i=1}^d \mathbf{z}_i^{x_i}$, $\sum_{i=1}^d x_i \leq k$, hence, the size of the model is $h = \mathcal{O}(k^d)$.

IV. DATASET-QUANTIZED STOCHASTIC GRADIENT DESCENT (DAQUSGD)

This section presents DaQuSGD, our proposed approach for dataset-based quantized stochastic gradient descent, which is summarized below and in the pseudocode in Algorithm 1.

Algorithm 1 DaQuSGD

Initialize $\mathbf{w}^{(0)}$, hyperparameter: m number of levels in data quantizer; Suppose η_j follows a learning rate schedule.

for $j = 1$ **to** n **do**

Node selected in the j -th iteration:

 Sample datapoint $\mathbf{z}^{(j)}$

$\mathbf{z}_Q^{(j)}, \mathbf{b}_z^{(j)} \leftarrow \text{DataQ}(\mathbf{z}^{(j)}, m)$;

$\triangleright \mathbf{b}_z^{(j)}$ is the lossless binary compression of $\mathbf{z}_Q^{(j)}$

$\Delta \leftarrow g_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j-1)}) - g_{\mathbf{z}_Q^{(j)}}(\mathbf{w}^{(j-1)})$

$\mathbf{b}_g^{(j)} \leftarrow \text{GradCorrQ}(\Delta, m)$;

$\triangleright \mathbf{b}_g^{(j)}$ is binary compression of estimate $\widehat{\Delta}$ of Δ

 Send $\mathbf{b}_z^{(j)}$ and $\mathbf{b}_g^{(j)}$ to the learning agent

Learning Agent:

 Reconstruct $\mathbf{z}_Q^{(j)}$ and $\widehat{\Delta}$ from $\mathbf{b}_z^{(j)}$ and $\mathbf{b}_g^{(j)}$

$\widehat{g} \leftarrow g_{\mathbf{z}_Q^{(j)}}(\mathbf{w}^{(j-1)}) + \widehat{\Delta}$

$\mathbf{w}^{(j)} \leftarrow \mathbf{w}^{(j-1)} - \eta_j \widehat{g}$

 Broadcast $\mathbf{w}^{(j)}$ to all nodes

Define $\text{DataQ}(\mathbf{z}, m)$:

// The DataQ subroutine

 Hyperparameters: Bound B on the value of $\|\mathbf{z}\|_2$.

$\mathbf{z}^+ \leftarrow \max(\mathbf{z}, 0)$, $\mathbf{z}^- \leftarrow \max(-\mathbf{z}, 0)$

$\mathbf{a} \leftarrow \lfloor \frac{(m-1)\mathbf{z}^+}{B} \rfloor$, $\mathbf{b} \leftarrow \lfloor \frac{(m-1)\mathbf{z}^-}{B} \rfloor$; $\mathbf{a}, \mathbf{b} \in \mathbb{N}^d$

$\mathcal{S} = \{(\mathbf{a}, \mathbf{b}) \in \mathbb{N}^d \times \mathbb{N}^d \mid \|\mathbf{a}\|_1 + \|\mathbf{b}\|_1 \leq (m-1)^2\}$.

$\mathbf{b}_z^{(\mathbf{a}, \mathbf{b})} \leftarrow \text{Index}_{\mathcal{S}}((\mathbf{a}, \mathbf{b}))$ // Index of (\mathbf{a}, \mathbf{b}) in set \mathcal{S} by enumeration

Reconstruction. $\mathbf{z}_Q \leftarrow [\mathbf{a} - \mathbf{b}] \frac{B}{m-1}$;

Return $\mathbf{z}_Q, \mathbf{b}_z^{(\mathbf{a}, \mathbf{b})}$

Comment: The reconstruction step is also performed at the learning agent.

The encoder Q^e consists of two components: a datapoint quantization step, followed by a gradient correction step. Given a datapoint \mathbf{z} , a node applies a data quantizer $\text{DataQ}(\cdot, m)$ parameterized by an integer m , to create $\mathbf{z}_Q = \text{DataQ}(\mathbf{z}, m)$, the quantized version of \mathbf{z} . DataQ uniformly quantizes each feature of $|\mathbf{z}|$ into m quantization levels as well as implicitly communicates the sign of each feature. At a high level, the quantizer DataQ can produce \mathbf{z}_Q such that $\|g_{\mathbf{z}_Q}(\mathbf{w}) - g_{\mathbf{z}}(\mathbf{w})\|_2$ is reasonably small. However, the quantized estimate $g_{\mathbf{z}_Q}(\mathbf{w})$ is no longer an unbiased estimate of $\nabla L(\mathbf{w})$. In this case, we can only guarantee the convergence of $\nabla L(\mathbf{w}^{(n)})$ to a neighborhood of zero (i.e., to a neighborhood of a local optimal \mathbf{w}).

If we want to guarantee convergence to a local optimal point, then in next step a very low-precision amendment of $g_{\mathbf{z}_Q}(\mathbf{w})$ is sent to ensure that the central agent has access to an unbiased estimate of $\nabla L(\mathbf{w})$ with bounded variance. Although the error $\Delta = g_{\mathbf{z}}(\mathbf{w}) - g_{\mathbf{z}_Q}(\mathbf{w})$ exists in a large dimensional space \mathbb{R}^h , the fact that $\|g_{\mathbf{z}_Q}(\mathbf{w}) - g_{\mathbf{z}}(\mathbf{w})\|_2$ is small enables us to quantize Δ with a small number of bits (either 1 bit if the terminal nodes and the algorithm have shared randomness or $\lceil \log_2(h) \rceil + 1$ otherwise). Thus the main quantization load

is done in the data space of dimension $d \ll h$.

In the following subsections, we describe and analyze the two components of DaQuSGD: Datapoint Quantization DataQ and Gradient Correction GradCorrQ. For the different versions of DaQuSGD (with and without gradient correction), the order-wise convergence rates and their associated communication costs are summarized in Table I and compared with the state-of-the-art memoryless gradient quantization schemes.

A. Datapoint quantization [DataQ]

$\text{DataQ}(\cdot, m)$ is parametrized by a tuning parameter m capturing the number of levels used to quantize each feature (see also pseudocode in Algorithm 1). Given a dataset sample $\mathbf{z} \in \mathbb{R}^d$, where $\|\mathbf{z}\|_2 \leq B$, we express it as $\mathbf{z} = \mathbf{z}^+ - \mathbf{z}^-$, where $\mathbf{z}^+ = \max(\mathbf{z}, 0)$ is a vector that captures the magnitude of the positive elements of \mathbf{z} (negative values are replaced by zero) and similarly $\mathbf{z}^- = \max(-\mathbf{z}, 0)$ for the negative elements. We create the non-negative concatenation vector $\tilde{\mathbf{z}} = [\mathbf{z}^+, \mathbf{z}^-]^T \in \mathbb{R}_+^{2d}$, that inherits the following properties from \mathbf{z} : (i) $\tilde{\mathbf{z}}$ has the same ℓ_2 norm as \mathbf{z} ; (ii) $\tilde{\mathbf{z}}$ implicitly captures the sign of the i -th element of \mathbf{z} by observing whether the i -th component of the subvector \mathbf{z}^+ is zero or not; (iii) If $\|\mathbf{z}\|_2 \leq B$, then $0 \leq \tilde{\mathbf{z}} \leq B$ elementwise. Then, for each coordinate i , we choose m equally spaced quantization levels in the interval $[0, B]$, where the i -th quantization level is $q_i = \frac{iB}{m-1}$, $i \in [m-1]$. Now, $\forall j \in [d]$, let

$$\mathbf{a}_j(\mathbf{z}) = \arg \max_{i \in [m-1]} \{q_i \mid q_i \leq |\mathbf{z}_j^+|\},$$

$$\mathbf{b}_j(\mathbf{z}) = \arg \max_{i \in [m-1]} \{q_i \mid q_i \leq |\mathbf{z}_j^-|\}. \quad (6)$$

The integer vectors $\mathbf{a} = [\mathbf{a}_1, \dots, \mathbf{a}_d]^T$ and $\mathbf{b} = [\mathbf{b}_1, \dots, \mathbf{b}_d]^T$ capture the indices of the levels that are just below the values in $\tilde{\mathbf{z}}$; DataQ maps the values in $\tilde{\mathbf{z}}$ to exactly the values indexed by \mathbf{a} and \mathbf{b} . In particular, \mathbf{z}^+ is quantized to $\mathbf{a} \frac{B}{m-1}$ and \mathbf{z}^- is quantized to $\mathbf{b} \frac{B}{m-1}$. Thus, As a result, we get that

$$\mathbf{z}_Q = \text{DataQ}(\mathbf{z}, m) = (\mathbf{a} - \mathbf{b}) \frac{B}{m-1}.$$

Communication Cost. The quantized point \mathbf{z}_Q generated by DataQ is uniquely represented by the integer vectors \mathbf{a} and \mathbf{b} . Thus, the communication cost equals the number of bits needed for lossless compression of \mathbf{a} and \mathbf{b} . Let us define the set \mathcal{S} to be

$$\mathcal{S} = \left\{ (\mathbf{v}^{(1)}, \mathbf{v}^{(2)}) \in \mathbb{N}^d \times \mathbb{N}^d \mid \|\mathbf{v}^{(1)}\|_1 + \|\mathbf{v}^{(2)}\|_1 \leq (m-1)^2 \right\}. \quad (7)$$

We will argue that for any (\mathbf{a}, \mathbf{b}) generated by DataQ as in (6), $(\mathbf{a}, \mathbf{b}) \in \mathcal{S}$ with probability 1. As a result, we only need at most $\lceil \log_2(|\mathcal{S}|) \rceil$ bits per sample to communicate (\mathbf{a}, \mathbf{b}) to the central node. We almost proved Theorem 1 (the complete proof is provided in Appendix C) that provides an upper bound on the communication cost.

Theorem 1. *The proposed DataQ algorithm satisfies the following statements: (i) For the integer vectors (\mathbf{a}, \mathbf{b}) uniquely defining \mathbf{z}_Q , we have that $(\mathbf{a}, \mathbf{b}) \in \mathcal{S}$ with probability one; (ii) DataQ uses at most $2 \log_2(m) +$*

TABLE I
COMPARISON OF DaQuSGD (WITH AND WITHOUT GRADIENT CORRECTION) WITH GRADIENT COMPRESSION SCHEMES.

METHOD	CONVERGENCE GUARANTEE	COMMUNICATION COST/ITERATION	LEARNER COMPUTES GRADIENT	NODES COMPUTES GRADIENT
DaQuSGD (DataQ)	$O(\frac{1}{\sqrt{n}} + \frac{1}{h})$	$O(d \log_2(h))$	✓	✗
DaQuSGD (DataQ + GradCorrQ)	$O(\frac{1}{\sqrt{n}})$	$O(d \log_2(h))$	✓	✓
MEMORYLESS GRADIENT COMPRESSION: QSGD [7], RATQ [9]	$O(\frac{1}{\sqrt{n}})$	$O(h)$	✗	✓

$\min\{2d \log_2(e^{\frac{2d+m^2}{2d}}), m^2 \log_2(e^{\frac{2d+m^2}{m^2}})\}$ bits per sample for communication; (iii) For the generated \mathbf{z}_Q , we have that $\|\mathbf{z} - \mathbf{z}_Q\|_\infty \leq \frac{B}{m-1}$ and $\|\mathbf{z}_Q\|_2 \leq (1 + \frac{\sqrt{d}}{m-1})B$.

Remark 1 (Bounded Second Moment). If the second moment satisfies $\mathbb{E}[\|\mathbf{z}\|_2^2] \leq B^2$, but $\|\mathbf{z}\|_2$ is not bounded almost surely, we can send $\|\mathbf{z}\|_2$ with full precision and then use DataQ to quantize $\mathbf{z}/\|\mathbf{z}\|_2$, with $B = 1$. This adds an overhead of sending (only) one scalar in full precision; moreover, we can now guarantee that $\mathbb{E}[\|\mathbf{z}_Q\|_2^2] \leq (1 + \frac{\sqrt{d}}{m-1})^2$ instead of the universal bound stated on $\|\mathbf{z}_Q\|_2$ in Theorem 1. Additionally, we have that $\mathbb{E}[\|\mathbf{z} - \mathbf{z}_Q\|_\infty] \leq \frac{B}{m-1}$.

Remark 2 (Splitting \mathbf{z}). From Theorem 1, we have that the dependency of the number of bits on the dataset dimension d is $\Omega(\log_2(d))$ when m is small. However, without splitting \mathbf{z} into $\mathbf{z}^+, \mathbf{z}^-$, we cannot represent the quantized values of \mathbf{z} , with a set of positive integers similar to \mathcal{S} , without directly sending the signs in \mathbf{z} which requires at least d bits. Note that for small values of d , splitting might require larger number of bits than without splitting, however, we are interested in how the number of bits grow with d .

Remark 3 (Quantization of gradient using DataQ). DataQ with stochastic quantization [7], [46]⁵ if applied to the gradient $g_{\mathbf{z}}(\mathbf{w})$ with $m = \sqrt{h} + 1$ would result in $\hat{g}_{\mathbf{z}}(\mathbf{w})$ that satisfies $\mathbb{E}[\hat{g}_{\mathbf{z}}(\mathbf{w})|g_{\mathbf{z}}(\mathbf{w})] = g_{\mathbf{z}}(\mathbf{w})$, $\|\hat{g}_{\mathbf{z}}(\mathbf{w}) - g_{\mathbf{z}}(\mathbf{w})\|_2 \leq 2B$, hence, achieves the optimal convergence rate of $O(\frac{1}{\sqrt{n}})$. This uses at most $\log_2(2h) + 2h \log_2(3e) = \mathcal{O}(h)$ bits, thus, achieving the communication lower bound in [9]. Unlike the results in [7], [41] which provide a guarantee on communication cost in terms of expectation, this provides a uniform upper bound on the required number of bits.

Remark 4 (Convergence of only DataQ). Note that, using (4) and Theorem 1, we get

$$\begin{aligned} \|g_{\mathbf{z}}(\mathbf{w}) - g_{\mathbf{z}_Q}(\mathbf{w})\|_2 &\stackrel{(4)}{\leq} C_z \|\mathbf{z} - \mathbf{z}_Q\|_2 \\ &\stackrel{\text{Th. 1}}{\leq} C_z B \frac{\sqrt{d}}{m-1}, \quad \forall \mathbf{w} \in \mathcal{W}, \end{aligned} \quad (8)$$

i.e., $g_{\mathbf{z}_Q}(\mathbf{w})$ is a biased estimate of $\nabla L(\mathbf{w})$, with a bounded bias given by (8). With this bias bound, we get the following convergence guarantee by appealing to the result in [47].

⁵Instead of mapping to a level below the feature, we map it either above or below with a probability depending on the distance.

Corollary 1. Let $\ell(\mathbf{w}, \mathbf{z})$ be a function satisfying the assumptions in Section II. By using the DataQ mechanism with parameter $m = h\sqrt{d}$, there exists a decaying learning rate η such that

$$\mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \left\| \nabla L(\mathbf{w}^{(i)}) \right\|_2^2 \right] \leq O \left(\frac{1}{\sqrt{n}} + \frac{1}{h} \right), \quad (9)$$

while using at most $O(d \log_2(h))$ bits per iteration.

The proof of Corollary 1 is included in Appendix D for completeness. The corollary implies that using DataQ guarantees convergence with rate $O(\frac{1}{\sqrt{n}})$ to a point with gradient $\nabla L(\mathbf{w}^{(n)})$ falling in a neighborhood of zero with radius $O(\frac{1}{h})$. Since h is typically large, in our experiments in Section VI, we are able to train models to sufficiently good performance using only DataQ.

To guarantee convergence to a local optimal point (with $\nabla L(\mathbf{w}) = 0$, instead of a neighborhood), in the following subsection, we introduce a gradient correction mechanism.

B. Gradient correction [GradCorrQ]

Algorithm 2 GradCorrQ(Δ, m)

Hyperparameters : C_z , Lipschitz constant of loss function in \mathbf{z} ; h , dimension of hypothesis space; d , dimension of dataset.

Initialize $\hat{\Delta} = \mathbf{0}^h$

Pick integer index $i^* \in [h]$ uniformly at random.

$t(i^*) \leftarrow$ binary representation of i^* using $\log_2(h)$ bits.

$p \leftarrow \frac{\Delta_{i^*}}{2C_z B \frac{\sqrt{d}}{m-1}} + \frac{1}{2}$.

$e_g \leftarrow \text{Bernoulli}(p)$

$\hat{\Delta}_{i^*} \leftarrow (2e_g - 1)C_z B h \frac{\sqrt{d}}{m-1}$.

Return $\hat{\Delta}, (t(i^*), e_g)$.

Comment: Using $(t(i^*), e_g)$, the central node can recreate $\hat{\Delta}$.

We here describe the gradient correction procedure GradCorrQ (see also the pseudocode in Algorithm 2) that augments the quantized gradient estimated from $\mathbf{z}_Q = \text{DataQ}(\mathbf{z}, m)$ to achieve $O(1/\sqrt{n})$ convergence.

Let Δ be the error in computing the gradient using \mathbf{z}_Q defined as $\Delta = g_{\mathbf{z}}(\mathbf{w}) - g_{\mathbf{z}_Q}(\mathbf{w})$. The node that sent \mathbf{z}_Q to the learning agent can also communicate an estimate of Δ , by quantizing it as follows. The learning agent and the node

agree on a random number generator that generates $i^* \in [1 : h]$ uniformly at random. If there is no common randomness, then the node can send i^* to the agent using $\lceil \log_2(h) \rceil$ bits. The node also sends a single stochastic bit e_g , where e_g

$$\Pr\{e_g = 1\} = \frac{C_z B \frac{\sqrt{d}}{m-1} - \Delta_{i^*}}{2C_z B \frac{\sqrt{d}}{m-1}}. \quad (10)$$

This, in a way, uses random coordinate selection and stochastic quantization [7]. The agent constructs an estimate $\hat{\Delta}$ and uses it to create the quantized stochastic gradient as follows

$$\begin{aligned} \hat{g}_{\mathbf{z}}(\mathbf{w}) &= g_{\mathbf{z}_Q}(\mathbf{w}) + \hat{\Delta}, \text{ where} \\ \hat{\Delta}_i &= \begin{cases} (2e_g - 1)C_z B h \frac{\sqrt{d}}{m-1} & \text{if } i = i^*, \\ 0 & \text{otherwise.} \end{cases} \end{aligned} \quad (11)$$

Lemma 1, which is proved in Appendix E, shows the unbiasedness and variance properties of the gradient estimate. From this, Theorem 2, which summarizes the convergence guarantees for DaQuSGD follows.

Lemma 1. *The quantized stochastic gradient $\hat{g}_{\mathbf{z}}(\mathbf{w})$ is an unbiased estimate of $\nabla L(\mathbf{w})$ and satisfies $\|\hat{g}_{\mathbf{z}}(\mathbf{w}) - \nabla L(\mathbf{w})\|_2 \leq C_z B \left(2 + (h+1) \frac{\sqrt{d}}{m-1}\right)$.*

Theorem 2. (1) *Let \mathcal{W} be convex and let $\ell(\mathbf{w}, \mathbf{z})$ be a convex function satisfying the loss function assumptions in Section II. Let \mathcal{Z} be B -bounded and $\sup_{\mathbf{w} \in \mathcal{W}} \|\mathbf{w} - \mathbf{w}^{(0)}\|_2^2 = \tilde{D}^2$, where $\mathbf{w}^{(0)} \in \mathcal{W}$ is the initial model. Assume that the stochastic gradient descent uses the quantized gradients $\hat{g}_{\mathbf{z}}(\mathbf{w})$ obtained through DaQuSGD, with step size $\eta = \left(C_w + \frac{1}{\gamma}\right)^{-1}$, where $\gamma = \frac{\tilde{D}}{\hat{\sigma}} \sqrt{\frac{2}{n}}$. Then*

$$\begin{aligned} \mathbb{E} \left[L \left(\frac{1}{n} \sum_{i=1}^n \mathbf{w}^{(i)} \right) \right] - L(\mathbf{w}^*) &\leq \tilde{D} \hat{\sigma} \sqrt{\frac{2}{n}} + \frac{C_w \tilde{D}^2}{n}, \\ \text{with } \hat{\sigma} &= C_z B \left(2 + (h+1) \sqrt{d}/m - 1\right). \end{aligned}$$

(2) *Let $\ell(\mathbf{w}, \mathbf{z})$ be a function (possibly non-convex) satisfying the loss function assumptions in Section II, and $\|\nabla L(\mathbf{w})\|_2 \leq \tilde{D}$, $\forall \mathbf{w} \in \mathcal{W}^6$. Let $L(\mathbf{w}^{(0)}) - L(\mathbf{w}^*) = D_0$, where $\mathbf{w}^{(0)} \in \mathcal{W}$ is the initial model and \mathbf{w}^* is the optimal model. Then, DaQuSGD with step size $\eta = \min\{C_w^{-1}, \gamma\}$, where $\gamma = \frac{1}{\hat{\sigma}} \sqrt{\frac{2D_0}{nC_w}}$, satisfies*

$$\frac{1}{n} \sum_{i=1}^n \mathbb{E}[\|\nabla L(\mathbf{w}^{(i)})\|_2^2] \leq 2\hat{\sigma} \sqrt{\frac{2C_w D_0}{n}} + \frac{2D_0 C_w}{n}, \quad (12)$$

where $\hat{\sigma}$ is as before, and expectation is over the random selection of points, and randomness in DaQuSGD.

Using the definition of $\hat{\sigma}$ in Theorem 2 we can find the condition required for optimal convergence using DaQuSGD, which is stated by the following corollary.

Corollary 2. *For $m = h\sqrt{d}$, DaQuSGD achieves the optimal convergence rate, using at most $1 + \log_2(h) + 2\log_2(h\sqrt{d}) + 2d\log_2(e(1 + h^2/2))$ bits per iteration.*

⁶Note that the assumptions in Section II imply that $\|\nabla L(\mathbf{w})\|_2$ is bounded almost surely.




Fig. 2. Illustrative example for sample selection.

V. SELECTING SAMPLES TO TRANSMIT

In this part, we aim to further reduce the communication cost of DaQuSGD using sample selection. The high level incentive is the following: *If a data sample contributes minimally to the model learning, we do not transmit it.* This can intuitively be viewed as a generalization of support vectors in SVM; we want to only send samples that are necessary for designing the classifier, and we also want to decide which samples are necessary in an *online* way, *without the knowledge of the full dataset*. Note that, samples that are transmitted are still quantized using DaQuSGD.

Our sample selection method assumes terminal nodes have enough resources to perform the forward pass of the model that we want to train at the central node. First, for a new data sample in the i -th iteration, we apply the last version of the model available at the terminal node $\mathbf{w}^{(i-1)}$ on the new data sample \mathbf{z} to compute the loss $\ell(\mathbf{w}^{(i-1)}, \mathbf{z})$. Using this loss, we quantize and transmit the point to the central node only if the loss exceeds a threshold, i.e., if $\ell(\mathbf{w}^{(i-1)}, \mathbf{z}) > \ell_{th}^{(i)}$, where $\ell_{th}^{(i)}$ is the threshold in iteration i .

The intuition behind this loss thresholding approach is that for classification tasks, samples close to the classification boundary (points in P2 in Fig. 2) will typically have higher loss (important for the classifier design), while the samples that are well away from the boundary will exhibit smaller losses and thus have a minimal effects on the model (P1 and P3 in Fig. 2).

Formally, in the i -th iteration of the learning algorithm, by applying the selection based on thresholding, the algorithm minimizes the risk function $\hat{L}(\mathbf{w}) = \frac{1}{N} \sum_{i=1}^N \max\{\ell_{th}^{(i)}, \ell(\mathbf{w}, \mathbf{z}^{(i)})\}$. Hence, by decreasing the threshold $\ell_{th}^{(i)}$ as the algorithm progresses, the algorithm approaches the optimal point for a convex loss function (or a local optimal for non-convex loss function). Ideally, we want the loss of points that are not communicated to not be detrimental to the convergence rate of the algorithm. This cannot be guaranteed in general. The following theorem (proved in Appendix F) provides a sufficient condition on designing the thresholds $\{\ell_{th}^{(i)}\}_{i=1}^n$ such that our thresholding approach can converge with rate $\mathcal{O}(\frac{1}{\sqrt{n}})$, same as the vanilla SGD.

Theorem 3. *Let \mathcal{W} be convex, $L(\mathbf{w})$ be convex, C_w -smooth, and $\sup_{\mathbf{w} \in \mathcal{W}} \|\mathbf{w} - \mathbf{w}^{(0)}\|_2^2 = \tilde{D}^2$, where $\mathbf{w}^{(0)} \in \mathcal{W}$ is the initial hypothesis. Let $\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathcal{W}} L(\mathbf{w})$ and assume that $L(\mathbf{w}^*) \geq 0$. Assume that SGD is performed with stochastic gradients $\hat{g}(\mathbf{w})$ that satisfy (i) $\mathbb{E}[\hat{g}(\mathbf{w})] = \nabla L(\mathbf{w})$ (unbi-*

asedness), and (ii) $\mathbb{E}[\|\hat{g}(\mathbf{w}) - \nabla L(\mathbf{w})\|_2^2] \leq \tilde{B}^2$ (bounded variance). Additionally, assume that the loss thresholds satisfy that $\sum_{i=1}^n \sqrt{\ell_{th}^{(i)}} \leq \sqrt{n}$, then at iteration n , if the step size is $\eta = \min\{C_w^{-1}, (\sqrt{n}\tilde{B})^{-1}\}$, we have

$$\mathbb{E} \left[L \left(\frac{1}{n} \sum_{i=1}^n \mathbf{w}^{(i)} \right) \right] - L(\mathbf{w}^*) \leq \frac{\tilde{B}(\tilde{D}/2 + 1) + \tilde{D}\sqrt{2C_w}}{\sqrt{n}} + \frac{3C_w\tilde{D}\sqrt{2C_w}}{n\tilde{B}}. \quad (13)$$

Hence, this approach reduces communication cost, and computation cost (since backpropagation is only applied on a subset of the samples) without sacrificing the order of the convergence rate. In fact, we show numerically, in Section VI, that thresholding can provide an improvement in terms of how fast do we converge to a good performance model.

Remark 5 (Sample Selection + DaQuSGD). Theorem 3 assumes only the unbiasedness and bounded variance properties of the gradient estimates. Thus, we can apply the sample selection through thresholding on top of our DaQuSGD quantization approach described earlier in Section IV without a penalty in the order of convergence, since it generates gradient estimates with these properties at the central node.

VI. EXPERIMENTAL EVALUATION

We evaluate the practical gain DaQuSGD achieves in terms of communicated bits for training ResNet models on the CIFAR-10 and ImageNet datasets.

ImageNet Experimental Setup. For our ImageNet experiments, we train a ResNet-50 [48] ($h = 25, 557, 032$) on the ImageNet [49] dataset ($d = 150, 528$). We use a learning rate schedule with a base learning rate of 1e-1 with a piece-wise decay of 0.25 introduced every 30 epochs and batch size of 128. We consider a scenario with a single learning server and 10,000 distributed nodes, each holding a unique set of 128 ImageNet images locally. The collection of all disjoint image sets from all distributed nodes make up the ImageNet training set. The training was done using SGD with momentum of 0.9.

CIFAR-10 Experimental Setup. For CIFAR-10 ($d = 3, 072$), we train a ResNet-18 ($h = 11, 173, 962$). We use a learning rate schedule consisting of a base learning rate of 1e-3 with a piece-wise decay of 0.1 introduced at epoch 80 and batch size of 128. We assume a single central server that learns the classifier using data from 390 distributed nodes, where each distributed node stores 128 CIFAR images locally (totaling 49,920 total data points). Similar to the ImageNet models, the ResNet-18 networks were trained using SGD with momentum of 0.9.

DaQuSGD Implementation. We make some implementation adjustments to the DaQuSGD algorithm theoretically analyzed for convergence in Section IV. Our implementation differs as follows. First, instead of scaling each data vector of the batch by its own ℓ_2 -norm before applying DataQ, we scale all of them using $\max_{i \in [1:128]} \|\mathbf{z}^{(i)}\|_\infty$, the maximum feature value observed in the batch. This helps retain more values close to their original unquantized values and allows to

get away with communicating a single full-precision scaling factor instead of a factor for each datapoint.

As observed through Corollary 1 and Theorem 2: when GradCorrQ is not used, the convergence guarantee on $\nabla L(\mathbf{w})$ is worsened from being to zero to instead being to a ball of radius $O(\frac{1}{h})$ around zero. This is substantially small for $h > 10^7$ as in ResNet models. We also saw minimal performance gains in our experiments when GradCorr was used. Thus in our experiments, we opted to not use GradCorrQ, and instead use only DataQ. As a result, terminal nodes at most only need to calculate the loss function and do not need to run backpropagation.

When sample selection is applied, the threshold at the beginning of each epoch is set to be 0.2 of the average loss of the transmitted samples of the previous epoch.

Comparison to direct gradient quantization. On ImageNet, we compare the performance of DaQuSGD with sample selection to the state-of-the art gradient sparsification and quantization approach QTopK-SGD with memory [23] and momentum [50]. QTopK-SGD quantizes the top $K\%$ gradient values per model block using 16 bits each and performs error compensation by keeping memory of the error between true and compressed gradients.

The ResNet-50 model was trained on the ImageNet dataset for 90 epochs using our DaQuSGD approach and 150 epochs using the sparse QTopK-SGD approach.

Figure 3(a) shows the growth in cumulative communication budget during training versus the test accuracy (recorded at every epoch). For gradient quantization schemes, the number of bits communicated is the number of bits used to represent the quantized gradient that is generated at the local node using each of the considered schemes. Note that, this value is independent of the batch size in case of schemes that quantize gradients. For our proposed data quantization, each node communicates a number of bits dependent on its local batch, which is equal to the number of bits used to represent each point multiplied by the number of points that were allowed through by the sample selection module (or the full batch size if sample selection is not used).

The final accuracies for DaQuSGD with sample selection ($m=8$), DaQuSGD without sample selection ($m=8$), and QTopK-SGD ($K = 10\%$) are 73.63%, 73.48%, 72.16%, respectively. We see that our proposed approach offers a saving of up to a factor of 7 over QTopK-SGD, for the same accuracy.

Figure 3(b) illustrates that QTopK-SGD requires a higher communication budget to converge to the same accuracy for the DaQuSGD with sample selection. This higher communication budget is manifested through the larger number of training epochs (33% more than what DaQuSGD utilized) in order to achieve the same accuracy. Larger number of epochs translates to both a longer training time as well as additional computational cost. That is, DaQuSGD with sample selection does not require memory, and saves complexity in two ways: it allows to converge to the same top-1 accuracy using fewer epochs and it only requires to compute the gradient for a subset of samples at every epoch.

For training ResNet-18 over the CIFAR-10 dataset, our proposed model DaQuSGD and QSGD was trained for 100

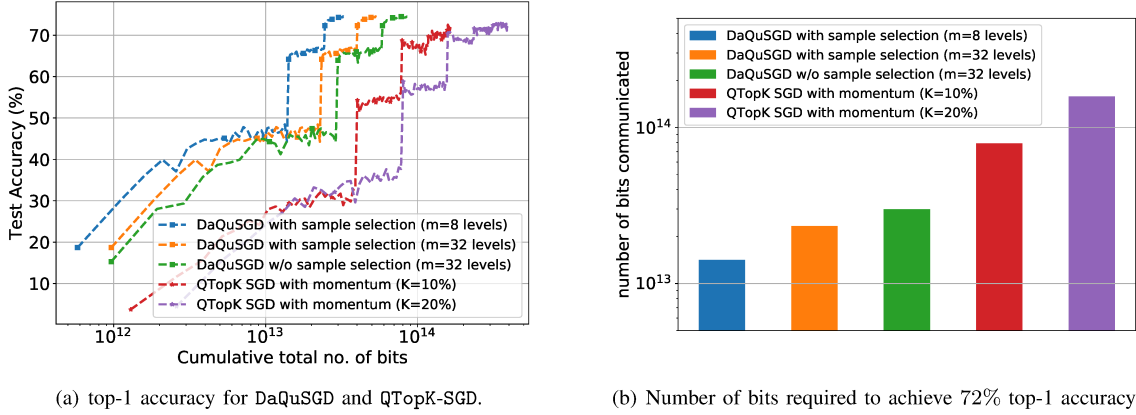


Fig. 3. Performance of DaQuSGD in comparison with QTopK-SGD when training ResNet-50 on ImageNet.

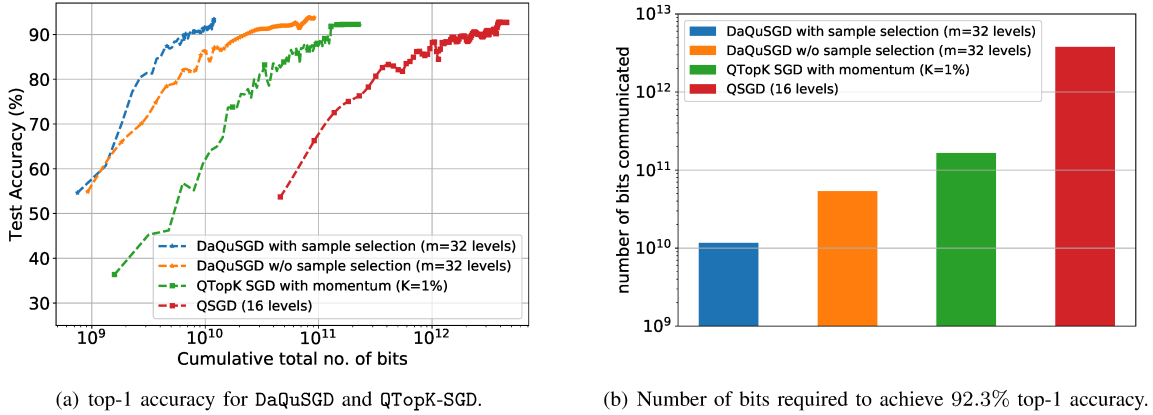


Fig. 4. Performance of DaQuSGD vs. QTopK-SGD and QSGD when training ResNet-18 on CIFAR-10.

epochs; QTopK-SGD was trained for 150 epochs. On CIFAR-10, we observed even more significant gains for our DaQuSGD approach over the QTopK-SGD and the QSGD [7] as seen in Fig. 4(a) and Fig. 4(b). In particular, our approach provided a factor of 14 reduction in the communication cost over QTopK-SGD and a factor of 300 over QSGD, to achieve the same top-1 accuracy. The final accuracies for DaQuSGD with sample selection, DaQuSGD without sample selection, QTopK-SGD and QSGD are 92.57%, 93.15%, 92.31%, 92.39%, respectively.

APPENDIX A SGD STANDARD CONVERGENCE THEOREMS

The central theorem used to prove convergence of stochastic gradient descent for smooth and convex risk functions is [42, Theorem 6.3], which modified to our notation is expressed as:

Theorem 4 ([42, Theorem 6.3]). *Let \mathcal{W} be convex, $L(\mathbf{w})$ be convex, C_w -smooth, and $\sup_{\mathbf{w} \in \mathcal{W}} \|\mathbf{w} - \mathbf{w}^{(0)}\|_2^2 = \tilde{D}^2$, where $\mathbf{w}^{(0)} \in \mathcal{W}$ is the initial hypothesis. Let $\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathcal{W}} L(\mathbf{w})$. Assume that stochastic gradient descent is performed with stochastic gradients $\hat{g}(\mathbf{w})$ that satisfy (i) $\mathbb{E}[\hat{g}(\mathbf{w})] = \nabla L(\mathbf{w})$ (unbiasedness), and (ii) $\mathbb{E}[\|\hat{g}(\mathbf{w}) - \nabla L(\mathbf{w})\|_2^2] \leq \tilde{B}^2$ (bounded variance), then at iteration n , if the step size is $\eta = \left(C_w + \frac{1}{\gamma}\right)^{-1}$ and $\gamma = \frac{\tilde{D}}{B} \sqrt{\frac{2}{n}}$,*

$$\mathbb{E} \left[L \left(\frac{1}{n} \sum_{i=1}^n \mathbf{w}^{(i)} \right) \right] - L(\mathbf{w}^*) \leq \tilde{D} \tilde{B} \sqrt{\frac{2}{n}} + \frac{C_w \tilde{D}^2}{n}. \quad (14)$$

For smooth non-convex risk functions, the stochastic gradient descent converges to a local optimal point. The main result we use in our analysis is from [15] is stated below.

Theorem 5 ([15]). *Let $L(\mathbf{w})$ be C_w -smooth, and $\|\nabla L(\mathbf{w})\|_2 \leq \tilde{D}$, $\forall \mathbf{w} \in \mathcal{W}$. Let $L(\mathbf{w}^{(0)}) - L(\mathbf{w}^*) = D_0$, where $\mathbf{w}^{(0)} \in \mathcal{W}$ is the initial hypothesis and \mathbf{w}^* is the optimal hypothesis. Assume that stochastic gradient descent is performed with stochastic gradients $\hat{g}(\mathbf{w})$ that satisfy (i) $\mathbb{E}[\hat{g}(\mathbf{w})] = \nabla L(\mathbf{w})$ (unbiasedness), and (ii) $\mathbb{E}[\|\hat{g}(\mathbf{w}) - \nabla L(\mathbf{w})\|_2^2] \leq \tilde{B}^2$ (bounded variance), then at iteration n with step size $\eta = \min\{C_w^{-1}, \gamma\}$ and $\gamma = \frac{1}{B} \sqrt{\frac{2D_0}{nC_w}}$, we have that*

$$\frac{1}{n} \sum_{t=1}^n \mathbb{E}[\|\nabla L(\mathbf{w}^{(t)})\|_2^2] \leq 2\tilde{B} \sqrt{\frac{2C_w D_0}{n}} + \frac{2D_0 C_w}{n}, \quad (15)$$

where the expectation is taken over the distribution of the stochastic gradient $\hat{g}(\mathbf{w})$.

APPENDIX B EXTENDED DISCUSSION FOR EXAMPLE 2 IN SECTION III

In this appendix, we show that the algorithm proposed for Example 2 in Section III achieves the optimal convergence rate with $\mathcal{O}(\log(h))$ bits and that $\Omega(\log(h))$ is a lower bound on the number of bits used by any SGD algorithm that achieves the optimal convergence rate for this example.

To show that the algorithm achieves the optimal convergence rate, it suffices to prove that it outputs an unbiased estimate of the gradient with a variance bounded by a constant. It is easy to see that any number in the set $\{1, \dots, k\}$ can be expressed as the sum of power of two of numbers from the set $\{1, \dots, \lceil \log_2(k) \rceil\}$, where each element appears in the sum at most once. With this, we now show that $\widehat{g}_{\mathbf{z}}(\mathbf{w})$ is an unbiased stochastic gradient. By definition, we have that

$$\begin{aligned}
\mathbb{E}[\widehat{g}_{\mathbf{z}}(\mathbf{w})_i | \mathbf{z}] &= \frac{1}{\sqrt{h}} \mathbb{E}[Q^d \circ Q^e(f'(\mathbf{w}^T \mathbf{v}(\mathbf{z}))) \\
&\quad \prod_{j=1}^{\lceil \log_2(k) \rceil} (Q^d \circ Q^e(\mathbf{z}^{2^j}))^{b_j(i) | \mathbf{z}}] \\
&\stackrel{[i]}{=} \frac{1}{\sqrt{h}} \mathbb{E}[Q^d \circ Q^e(f'(\mathbf{w}^T \mathbf{v}(\mathbf{z}))) | \mathbf{z}] \\
&\quad \prod_{j=1}^{\lceil \log_2(k) \rceil} \mathbb{E}[(Q^d \circ Q^e(\mathbf{z}^{2^j}))^{b_j(i) | \mathbf{z}}] \\
&\stackrel{[ii]}{=} \frac{1}{\sqrt{h}} \mathbb{E}[Q^d \circ Q^e(f'(\mathbf{w}^T \mathbf{v}(\mathbf{z}))) | \mathbf{z}] \\
&\quad \prod_{j=1}^{\lceil \log_2(k) \rceil} \mathbb{E}[Q^d \circ Q^e(\mathbf{z}^{2^j}) | \mathbf{z}]^{b_j(i)} \\
&\stackrel{[iii]}{=} \frac{1}{\sqrt{h}} f'(\mathbf{w}^T \mathbf{v}(\mathbf{z})) \mathbf{z}^{\sum_{j=1}^{\lceil \log_2(k) \rceil} b_j(i) 2^j} \\
&= \frac{1}{\sqrt{h}} f'(\mathbf{w}^T \mathbf{v}(\mathbf{z})) \mathbf{z}^i = g_{\mathbf{z}}(\mathbf{w})_i, \tag{16}
\end{aligned}$$

where: $[i]$ follows from the fact that conditioned on \mathbf{z} each of the quantized values are independent; $[ii]$ follows from the fact that $b_j(i) \in \{0, 1\}$ and therefore, we either we take the expectation over the variable or have the expectation over 1; $[iii]$ follows from the fact that the quantized values using $Q^d \circ Q^e$ are unbiased estimators of their unquantized values.

The fact that $\|\widehat{g}_{\mathbf{z}}(\mathbf{w})\|_2 \leq 1$ is obvious since the range of $Q^d \circ Q^e$ is in $\{-1, 1\}^k$.

Lower bound: For simplicity we only consider symmetric schemes. We will show that any symmetric quantization scheme that satisfies $\mathbb{E}[Q(\mathbf{v}(\mathbf{z})) | \mathbf{z}] = \mathbf{v}(\mathbf{z})$, $\forall \mathbf{z} \in \{\mathbf{z} \in \mathbb{R}^d | \|\mathbf{z}\| \leq 1\}$ uses at least $\log_2(d)$ bits. Consider a scheme that uses k bits, hence, for any \mathbf{z} with $\|\mathbf{z}\| \leq 1$, $Q(\mathbf{v}(\mathbf{z}))$ takes one of 2^k values denoted as e_1, \dots, e_{2^k} . Let $d_i = \mathbb{E}[Q^d(e_i)]$, where $Q^d(e_i)$ is the decoded value of e_i . Then, $\mathbb{E}[Q(\mathbf{v}(\mathbf{z})) | \mathbf{z}]$ is some convex combination of d_1, \dots, d_{2^k} . As a result, $\mathbb{E}[Q(\mathbf{v}(\mathbf{z})) | \mathbf{z}] = \mathbf{v}(\mathbf{z})$, $\forall \mathbf{z} \in \{\mathbf{z} \in \mathbb{R}^d | \|\mathbf{z}\| \leq 1\}$ implies that the set $\{\mathbf{z} \in \mathbb{R}^d | \|\mathbf{z}\| \leq 1\}$ is contained in the convex hull of the points d_1, \dots, d_{2^k} .

Now consider the matrix with columns $[\mathbf{v}(\frac{1}{n}), \mathbf{v}(\frac{2}{n}), \dots, \mathbf{v}(1)]$. This matrix has full rank, and clearly 2^k is lower bounded by the rank of this matrix. Hence, $k \geq \log_2(d)$.

APPENDIX C PROOF OF THEOREM 1

First, we reiterate Theorem 1 from the paper for readability.

Theorem. *The proposed DataQ algorithm satisfies the following statements: (1) For the integer vectors (a, b)*

uniquely defining \mathbf{z}_Q , we have that $(a, b) \in \mathcal{S}$ with probability one; (2) DataQ uses at most $2 \log_2(m) + \min\{2d \log_2(e^{\frac{2d+m^2}{2d}}, m^2 \log_2(e^{\frac{2d+m^2}{m^2}}))\}$ bits per sample for communication; (3) For the generated \mathbf{z}_Q , we have that $\|\mathbf{z} - \mathbf{z}_Q\|_\infty \leq \frac{B}{m-1}$ and $\|\mathbf{z}_Q\|_2 \leq (1 + \frac{\sqrt{d}}{m-1})B$.

1) **Proof that $(a, b) \in \mathcal{S}$:**

Recall that for datapoint \mathbf{z} , the definition of (a, b) is given by

$$\forall j \in \{1, 2, \dots, d\} :$$

$$a_j(\mathbf{z}) = \arg \max_{i \in \{0, 1, \dots, m-1\}} \left\{ q_i \mid q_i = \frac{iB}{m-1} \leq |\mathbf{z}_j^+| \right\}, \tag{17a}$$

$$b_j(\mathbf{z}) = \arg \max_{i \in \{0, 1, \dots, m-1\}} \left\{ q_i \mid q_i = \frac{iB}{m-1} \leq |\mathbf{z}_j^-| \right\}. \tag{17b}$$

Thus, given the upper bound B on $\|\mathbf{z}\|_2$, we have

$$\begin{aligned}
B^2 &\geq \|\mathbf{z}\|_2^2 = \|\mathbf{z}^+\|_2^2 + \|\mathbf{z}^-\|_2^2 \\
&\stackrel{(i)}{\geq} \sum_{j=1}^d \frac{B^2(a_j^2 + b_j^2)}{(m-1)^2} \stackrel{(i)}{\geq} \sum_{j=1}^d \frac{B^2(a_j + b_j)}{(m-1)^2}, \tag{17}
\end{aligned}$$

where (i) follows since $a_j, b_j \in \mathbb{N}$. From the above inequality, we now have that

$$\sum_{j=1}^d (a_j + b_j) \leq (m-1)^2$$

, which implies that $(a, b) \in \mathcal{S} = \left\{ (\mathbf{v}^{(1)}, \mathbf{v}^{(2)}) \in \mathbb{N}^d \times \mathbb{N}^d \mid \|\mathbf{v}^{(1)}\|_1 + \|\mathbf{v}^{(2)}\|_1 \leq (m-1)^2 \right\}$.

2) **Proof of the upper bound on number of bits:**

It suffices to show that $\log_2(|\mathcal{S}|) \leq 2 \log_2(m) + \min\{d \log_2(e^{\frac{2d+m^2}{2d}}, m^2 \log_2(e^{\frac{2d+m^2}{m^2}}))\}$. Note that $|\mathcal{S}|$ can be written as

$$\begin{aligned}
|\mathcal{S}| &= \sum_{q=0}^{(m-1)^2} \left| \left\{ (\mathbf{v}^{(1)}, \mathbf{v}^{(2)}) \in \mathbb{N}^d \times \mathbb{N}^d \right. \right. \\
&\quad \left. \left. \mid \|\mathbf{v}^{(1)}\|_1 + \|\mathbf{v}^{(2)}\|_1 \leq q \right\} \right|. \tag{19}
\end{aligned}$$

For a given integer q , the number of positive integral solutions to the equation $\sum_{i=1}^d (\mathbf{a}_i + \mathbf{b}_i) = q$ is a classical counting problem and its solution is given in closed form⁷ as $\binom{2d+q-1}{q}$ [51]. Thus, we can write $|\mathcal{S}|$ as

$$\begin{aligned}
|\mathcal{S}| &= \sum_{q=0}^{(m-1)^2} \binom{2d+q-1}{q} \leq \sum_{q=0}^{(m-1)^2} \binom{2d+q}{2d} \\
&\leq \sum_{q=0}^{(m-1)^2} \binom{2d+m^2}{2d} \\
&\stackrel{(i)}{\leq} m^2 \min \left\{ \binom{2d+m^2}{2d}, \binom{2d+m^2}{m^2} \right\} \\
&\stackrel{(ii)}{\leq} m^2 \min \left\{ \left(e^{\frac{2d+m^2}{2d}} \right)^{2d}, \left(e^{\frac{2d+m^2}{m^2}} \right)^{m^2} \right\}, \tag{20}
\end{aligned}$$

⁷The closed form relies on a standard approach in combinatorics called the ‘‘stars and bars method’’.

where: (i) is due to the symmetry of the binomial coefficient. Now by taking the logarithm of both sides, we get the intended upper bound for $\log_2(|S|)$; (ii) uses the upper bound on the binomial coefficient based on Sterling's bounding of the factorial.

3) **Proof that** $\|\mathbf{z}_j - \mathbf{z}_{Qj}\| \leq \frac{B}{m-1}$, $\|\mathbf{z}\|_2 \leq (1 + \frac{\sqrt{d}}{m-1})B$:

$\|\mathbf{z}_j - \mathbf{z}_{Qj}\|_\infty \leq \frac{B}{m-1}$ is directly due to the quantization scheme in DataQ since the distance between any two quantization levels is given by $\frac{B}{m-1}$.

Now to prove the bound on $\|\mathbf{z}\|_2$, note that

$$\begin{aligned} \|\mathbf{z}_{Qj}\|_2 &\stackrel{(i)}{\leq} \|\mathbf{z}\|_2 + \|\mathbf{z} - \mathbf{z}_{Qj}\|_2 \\ &\leq B + \sqrt{\|\mathbf{z} - \mathbf{z}_{Qj}\|_2^2} = B + \sqrt{\sum_{j=1}^d |\mathbf{z}_j - \mathbf{z}_{Qj}|^2} \\ &\leq B + \sqrt{\sum_{j=1}^d \frac{B^2}{(m-1)^2}} = \left(1 + \frac{\sqrt{d}}{m-1}\right) B, \end{aligned} \quad (21)$$

where (i) follows from the triangle inequality.

APPENDIX D PROOF OF COROLLARY 1

Here, we prove the convergence guarantee when only DataQ is applied, which is stated in Corollary 1. The proof is similar to the SGD proof in [47], [15].

Throughout the proof, we use the following notation: (1) The maximum difference in risk $L(\mathbf{w})$ is denoted as $D_L = \arg \max_{\mathbf{w}^{(1)}, \mathbf{w}^{(2)} \in \mathcal{W}} |L(\mathbf{w}^{(1)}) - L(\mathbf{w}^{(2)})|$; (2) the upper bound on the gradient $\|\nabla L(\mathbf{w})\|_2 \leq \tilde{D}$, $\forall \mathbf{w} \in \mathcal{W}$. We will assume that the learning rate used satisfies that $\eta \leq 1/C_w$.

Additionally, recall that $\ell(\mathbf{w}, \mathbf{z})$ is C_w -smooth in \mathbf{w} and C_z -smooth in \mathbf{z} . Additionally, recall that at iteration j , we update $\mathbf{w}^{(j)}$ using $g_{\mathbf{z}^{(j)}}(\mathbf{w})$ which is the gradient computed from applying the model with parameters $\mathbf{w}^{(j)}$ on the quantized version of $\mathbf{z}^{(j)}$. Define, $\delta_j = g_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)}) - \nabla L(\mathbf{w}^{(j)})$. Now, from the smoothness of $\ell(\mathbf{w}, \mathbf{z})$, we have the following

$$\begin{aligned} &\mathbb{E}_{\mathbf{z}^{(j)}} [L(\mathbf{w}^{(j+1)}) - L(\mathbf{w}^{(j)})] \\ &\leq \nabla L(\mathbf{w}^{(j)})^T \mathbb{E}_{\mathbf{z}^{(j)}} [\mathbf{w}^{(j+1)} - \mathbf{w}^{(j)}] + \frac{C_w}{2} \mathbb{E}_{\mathbf{z}^{(j)}} [\|\mathbf{w}^{(j+1)} - \mathbf{w}^{(j)}\|_2^2] \\ &= -\eta \nabla L(\mathbf{w}^{(j)})^T \mathbb{E}_{\mathbf{z}^{(j)}} [g_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)})] + \frac{C_w \eta^2}{2} \mathbb{E}_{\mathbf{z}^{(j)}} \left[\left\| g_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)}) \right\|_2^2 \right] \\ &= -\eta \left\| \nabla L(\mathbf{w}^{(j)}) \right\|_2^2 - \eta \nabla L(\mathbf{w}^{(j)})^T \mathbb{E}_{\mathbf{z}^{(j)}} [\delta_j] \\ &\quad + \frac{C_w \eta^2}{2} \mathbb{E}_{\mathbf{z}^{(j)}} \left[\left\| \nabla L(\mathbf{w}^{(j)}) + \delta_j \right\|_2^2 \right] \\ &= -\eta \left\| \nabla L(\mathbf{w}^{(j)}) \right\|_2^2 - \eta \nabla L(\mathbf{w}^{(j)})^T \mathbb{E}_{\mathbf{z}^{(j)}} [\delta_j] \\ &\quad + \frac{C_w \eta^2}{2} \mathbb{E}_{\mathbf{z}^{(j)}} \left[\left\| \nabla L(\mathbf{w}^{(j)}) \right\|_2^2 + 2 \nabla L(\mathbf{w}^{(j)})^T \delta_j + \|\delta_j\|_2^2 \right] \\ &= -\eta \left(1 - \frac{\eta C_w}{2} \right) \left\| \nabla L(\mathbf{w}^{(j)}) \right\|_2^2 \\ &\quad + \eta (1 - \eta C_w) \nabla L(\mathbf{w}^{(j)})^T \mathbb{E}_{\mathbf{z}^{(j)}} [\delta_j] + \frac{C_w \eta^2}{2} \mathbb{E}_{\mathbf{z}^{(j)}} [\|\delta_j\|_2^2]. \end{aligned} \quad (22)$$

Reorganizing (22), we get that

$$\begin{aligned} &\|\nabla L(\mathbf{w}^{(j)})\|_2^2 \\ &\leq \frac{2}{\eta(2 - \eta C_w)} \left(L(\mathbf{w}^{(j)}) - L(\mathbf{w}^{(j+1)}) + \frac{\eta^2 C_w}{2} \mathbb{E}_{\mathbf{z}^{(j)}} [\|\delta_j\|_2^2] \right) \\ &\quad + \frac{2(1 - C_w \eta)}{(2 - C_w \eta)} \nabla L(\mathbf{w}^{(j)})^T \mathbb{E}_{\mathbf{z}^{(j)}} [\delta_j] \\ &\leq \frac{2}{\eta(2 - \eta C_w)} \left(L(\mathbf{w}^{(j)}) - L(\mathbf{w}^*) + \frac{\eta^2 C_w}{2} \mathbb{E}_{\mathbf{z}^{(j)}} [\|\delta_j\|_2^2] \right) \\ &\quad + \frac{2(1 - C_w \eta)}{(2 - C_w \eta)} \nabla L(\mathbf{w}^{(j)})^T \mathbb{E}_{\mathbf{z}^{(j)}} [\delta_j]. \end{aligned} \quad (23)$$

Hence,

$$\begin{aligned} &\|\nabla L(\mathbf{w}^{(j)})\|_2^2 \\ &\stackrel{(a)}{\leq} \frac{2}{\eta(2 - \eta C_w)} \left(L(\mathbf{w}^{(j)}) - L(\mathbf{w}^*) + \frac{\eta^2 C_w}{2} \mathbb{E}_{\mathbf{z}^{(j)}} [\|\delta_j\|_2^2] \right) \\ &\quad + \frac{2(1 - C_w \eta)}{(2 - C_w \eta)} \left\| \nabla L(\mathbf{w}^{(j)}) \right\|_2 \mathbb{E}_{\mathbf{z}^{(j)}} [\delta_j]_2, \end{aligned} \quad (24)$$

where, (a) follows from the Cauchy-Schwartz inequality and the fact that $\eta \leq 1/C_w$. Now, we note that

$$\begin{aligned} &\left\| \mathbb{E}_{\mathbf{z}^{(j)}} [\delta_j] \right\|_2 = \left\| \mathbb{E}_{\mathbf{z}^{(j)}} [g_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)})] - \nabla L(\mathbf{w}^{(j)}) \right\|_2 \\ &= \left\| \mathbb{E}_{\mathbf{z}^{(j)}} [g_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)}) - g_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)})] \right\|_2 \\ &\leq \mathbb{E}_{\mathbf{z}^{(j)}} \left[\left\| g_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)}) - g_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)}) \right\|_2 \right] \stackrel{(a)}{\leq} C_z B \frac{\sqrt{d}}{(m-1)}. \end{aligned} \quad (25)$$

where (a) follows from the bound in (8). Additionally, we also have that

$$\begin{aligned} &\mathbb{E}_{\mathbf{z}^{(j)}} [\|\delta_j\|_2^2] \\ &= \mathbb{E}_{\mathbf{z}^{(j)}} \left[\left\| \left(g_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)}) - g_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)}) \right) - \left(g_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)}) - \nabla L(\mathbf{w}^{(j)}) \right) \right\|_2^2 \right] \\ &\leq 3 \mathbb{E}_{\mathbf{z}^{(j)}} \left[\left\| g_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)}) - g_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)}) \right\|_2^2 + \left\| g_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)}) - \nabla L(\mathbf{w}^{(j)}) \right\|_2^2 \right] \\ &\stackrel{(a)}{\leq} 3 \left(C_z B \frac{d}{(m-1)^2} + \sigma^2 \right), \end{aligned} \quad (26)$$

where σ^2 is the bounded variance of the unquantized stochastic gradient and (a) follows from the bound in (8).

By setting $m = h\sqrt{d}$ in (25) and (26) and substituting in (24) we have that

$$\begin{aligned} &\|\nabla L(\mathbf{w}^{(j)})\|_2^2 \\ &\leq \frac{2}{\eta(2 - \eta C_w)} \left(L(\mathbf{w}^{(j)}) - L(\mathbf{w}^*) + \eta^2 C_w \overbrace{\frac{3}{2} \left(\frac{2C_z B}{h^2} + \sigma^2 \right)}^{\hat{\sigma}^2} \right) \\ &\quad + \frac{4(1 - C_w \eta)}{(2 - C_w \eta)} \left\| \nabla L(\mathbf{w}^{(j)}) \right\|_2 \frac{C_z B}{h}. \end{aligned} \quad (27)$$

By averaging across $j \in [n]$, we get

$$\begin{aligned} &\frac{1}{n} \sum_{j=1}^n \left\| \nabla L(\mathbf{w}^{(j)}) \right\|_2^2 \\ &\leq \frac{2}{n\eta(2 - \eta C_w)} \left(L(\mathbf{w}^{(0)}) - L(\mathbf{w}^*) + n \eta^2 C_w \hat{\sigma}^2 \right) \end{aligned}$$

$$+ \frac{4(1 - C_w \eta)}{(2 - C_w \eta)} \frac{C_z B}{h} \frac{1}{n} \sum_{j=1}^n \left\| \nabla L(\mathbf{w}^{(j)}) \right\|_2. \quad (28)$$

Hence,

$$\begin{aligned} & \frac{1}{n} \sum_{j=1}^n \left\| \nabla L(\mathbf{w}^{(j)}) \right\|_2^2 \\ & \leq \frac{2}{n\eta(2 - \eta C_w)} \left((L(\mathbf{w}^{(0)}) - L(\mathbf{w}^*)) + n \eta^2 C_w \hat{\sigma}^2 \right) + \frac{2C_z B}{h} \tilde{D} \\ & \leq \frac{2}{n\eta(2 - \eta C_w)} (D_L + n \eta^2 C_w \hat{\sigma}^2) + \frac{2C_z B}{h} \tilde{D} \\ & \stackrel{(a)}{\leq} \frac{2}{n\eta} (D_L + n \eta^2 C_w \hat{\sigma}^2) + \frac{2C_z B}{h} \tilde{D} \\ & = \frac{2}{n\eta} D_L + 2 \eta C_w \hat{\sigma}^2 + \frac{2C_z B}{h} \tilde{D}, \end{aligned} \quad (29)$$

where (a) follows from assuming $\eta \leq 1/C_w$. By setting $\eta = \min\{\frac{1}{C_w}, \frac{\sqrt{2D_L}}{\hat{\sigma}\sqrt{C_w n}}\}$, we get that

$$\begin{aligned} & \frac{1}{n} \sum_{j=1}^n \left\| \nabla L(\mathbf{w}^{(j)}) \right\|_2^2 \\ & \leq \frac{2D_L}{n} \max \left\{ C_w, \frac{\hat{\sigma}\sqrt{C_w n}}{\sqrt{2D_L}} \right\} \\ & \quad + 2 C_w \hat{\sigma}^2 \min \left\{ \frac{1}{C_w}, \frac{\sqrt{2D_L}}{\hat{\sigma}\sqrt{C_w n}} \right\} + \frac{2C_z B}{h} \tilde{D} \\ & \leq \frac{2D_L C_w}{n} + \frac{\hat{\sigma}\sqrt{2D_L C_w}}{\sqrt{n}} + 2 \frac{\hat{\sigma}\sqrt{2D_L C_w}}{\sqrt{n}} + \frac{2C_z B}{h} \tilde{D} \\ & \leq \frac{2D_L C_w}{n} + 3 \frac{\hat{\sigma}\sqrt{2D_L C_w}}{\sqrt{n}} + \frac{2C_z B}{h} \tilde{D} = O\left(\frac{1}{\sqrt{n}} + \frac{1}{h}\right). \end{aligned} \quad (30)$$

This concludes the proof of Corollary 1.

APPENDIX E PROOF OF LEMMA 1

Lemma. *The quantized stochastic gradient $\hat{g}_z(\mathbf{w})$ is an unbiased estimate of $\nabla L(\mathbf{w})$ and satisfies $\|\hat{g}_z(\mathbf{w}) - \nabla L(\mathbf{w})\|_2 \leq C_z B \left(2 + (h+1)\frac{\sqrt{d}}{m-1}\right)$.*

Proof. Recall that $\Delta = g_z(\mathbf{w}) - g_{z_Q}(\mathbf{w})$, and that, we have

$$\hat{\Delta}_i = \begin{cases} 0 & \text{with probability } 1 - \frac{1}{h} \\ C_z B h \frac{\sqrt{d}}{m-1} & \text{with probability } \frac{\Delta_i + C_z B \frac{\sqrt{d}}{m-1}}{2C_z B \frac{\sqrt{d}}{m-1}} \frac{1}{h} \\ -C_z B h \frac{\sqrt{d}}{m-1} & \text{with probability } \frac{C_z B \frac{\sqrt{d}}{m-1} - \Delta_i}{2C_z C \frac{\sqrt{d}}{m-1}} \frac{1}{h}. \end{cases} \quad (31)$$

Hence, we can prove the unbiasedness property of $\hat{\Delta}$ by direct computation as follows

$$\begin{aligned} \mathbb{E}[\hat{\Delta}_i | \Delta] &= \mathbb{E}[\hat{\Delta}_i | \Delta_i] \\ &= 0 + \left(C_z B h \frac{\sqrt{d}}{m-1} \right) \frac{\Delta_i + C_z B \frac{\sqrt{d}}{m-1}}{2C_z B \frac{\sqrt{d}}{m-1}} \frac{1}{h} \\ & \quad + \left(-C_z B h \frac{\sqrt{d}}{m-1} \right) \frac{C_z B \frac{\sqrt{d}}{m-1} - \Delta_i}{2C_z C \frac{\sqrt{d}}{m-1}} \frac{1}{h} \\ &= \left(2C_z B h \frac{\sqrt{d}}{m-1} \right) \frac{2\Delta_i}{2C_z B \frac{\sqrt{d}}{m-1} h} = \Delta_i. \end{aligned} \quad (32)$$

Next, we can prove the unbiasedness of $\hat{g}_z(\mathbf{w})$ as follows

$$\begin{aligned} \mathbb{E}[\hat{g}_z(\mathbf{w}) | \mathbf{z}] &= \mathbb{E}_{\mathbf{z}_Q} [\mathbb{E}[\hat{g}_z(\mathbf{w}) | \mathbf{z}, \mathbf{z}_Q]] \\ & \stackrel{(i)}{=} \mathbb{E}_{\mathbf{z}_Q} \left[\mathbb{E}[g_{z_Q}(\mathbf{w}) + \hat{\Delta} | \mathbf{z}, \mathbf{z}_Q] \right] \\ & \stackrel{(ii)}{=} \mathbb{E}_{\mathbf{z}_Q} \left[\mathbb{E}[g_{z_Q}(\mathbf{w}) + \mathbb{E}[\hat{\Delta} | \Delta, \mathbf{z}, \mathbf{z}_Q] | \mathbf{z}, \mathbf{z}_Q] \right] \\ & \stackrel{(iii)}{=} \mathbb{E}_{\mathbf{z}_Q} [\mathbb{E}[g_{z_Q}(\mathbf{w}) + \Delta | \mathbf{z}, \mathbf{z}_Q]] \\ & \stackrel{(iv)}{=} \mathbb{E}_{\mathbf{z}_Q} [\mathbb{E}[g_z(\mathbf{w}) | \mathbf{z}, \mathbf{z}_Q]] \\ &= g_z(\mathbf{w}), \end{aligned} \quad (33)$$

where: (i) follows from the definition of $\hat{g}_z(\mathbf{w})$; (ii) follows by the tower property of expectation and the fact that $\hat{\Delta}$ is constructed by stochastic quantization of Δ ; (iii) follows from the unbiasedness of $\hat{\Delta}$ in (32); (iv) is due to the fact that $\Delta = g_z(\mathbf{w}) - g_{z_Q}(\mathbf{w})$.

To prove the bound on the variance, we note that

$$\begin{aligned} \|\hat{g}_z(\mathbf{w}) - \nabla L(\mathbf{w})\|_2 & \stackrel{(i)}{\leq} \|g_{z_Q}(\mathbf{w}) - \nabla L(\mathbf{w})\|_2 + \|\hat{\Delta}\|_2 \\ & \stackrel{(i)}{\leq} \|g_z(\mathbf{w}) - \nabla L(\mathbf{w})\|_2 + \|g_z(\mathbf{w}) - g_{z_Q}(\mathbf{w})\|_2 + \|\hat{\Delta}\|_2 \\ & \stackrel{(ii)}{\leq} \|g_z(\mathbf{w}) - \nabla L(\mathbf{w})\|_2 + C_z B \frac{\sqrt{d}}{m-1} + \|\hat{\Delta}\|_2 \\ & \stackrel{(iii)}{\leq} \|g_z(\mathbf{w}) - \nabla L(\mathbf{w})\|_2 + C_z B \frac{\sqrt{d}}{m-1} + C_z B h \frac{\sqrt{d}}{m-1} \\ &= \left\| \frac{1}{N} \sum_{i=1}^N (g_{z^{(i)}}(\mathbf{w}) - g_z(\mathbf{w})) \right\|_2 \\ & \quad + C_z B \frac{\sqrt{d}}{m-1} + C_z B h \frac{\sqrt{d}}{m-1} \\ & \stackrel{(iv)}{\leq} \frac{1}{N} \sum_{i=1}^N \|g_{z^{(i)}}(\mathbf{w}) - g_z(\mathbf{w})\|_2 \\ & \quad + C_z B \frac{\sqrt{d}}{m-1} + C_z B h \frac{\sqrt{d}}{m-1} \\ & \stackrel{(v)}{\leq} \frac{1}{N} \sum_{i=1}^N C_z \|\mathbf{z}^{(i)} - \mathbf{z}\|_2 + C_z B \frac{\sqrt{d}}{m-1} + C_z B h \frac{\sqrt{d}}{m-1} \\ & \leq 2C_z B + C_z B \frac{\sqrt{d}}{m-1} + C_z B h \frac{\sqrt{d}}{m-1}, \end{aligned} \quad (34)$$

where: (i) is due to the triangle inequality; (ii) follows from the C_z -Lipschitz continuity of g_z and the definition of `DataQ`; (iii) follows from (31); (iv) follows due to the convexity of the norm; (v) follows from the C_z -Lipschitz continuity of $g_z(\mathbf{w})$ in \mathbf{z} . \square

APPENDIX F PROOF OF THEOREM 3

In this part we prove the convergence of SGD when our sample selection scheme through thresholding is applied. We denote the stochastic gradient with loss thresholding at the j -th iteration to be $\tilde{g}_{z^{(j)}}(\mathbf{w}^{(j)})$ which is given by

$$\tilde{g}_{z^{(j)}}(\mathbf{w}^{(j)}) = \hat{g}_{z^{(j)}}(\mathbf{w}^{(j)}) \mathbb{I} \left(\ell(\mathbf{w}^{(j)}, \mathbf{z}^{(j)}) \leq \ell_{th}^{(j)} \right),$$

where $\mathbb{I}(\cdot)$ is the indicator function. In the following we present a simple proof of our result by adapting a standard proof of SGD convergence. From convexity of $L(\mathbf{w})$, we have that

$$L(\mathbf{w}^{(j)}) \leq L(\mathbf{w}^*) + \nabla L(\mathbf{w}^{(j)})^T (\mathbf{w}^{(j)} - \mathbf{w}^*). \quad (35)$$

It is well known that for convex function, smoothness implies a quadratic upper bound [42], i.e., we have that

$$\begin{aligned} L(\mathbf{w}^{(j+1)}) &\leq L(\mathbf{w}^{(j)}) + \nabla L(\mathbf{w}^{(j)})^T (\mathbf{w}^{(j+1)} - \mathbf{w}^{(j)}) \\ &\quad + \frac{C_w}{2} \|\mathbf{w}^{(j+1)} - \mathbf{w}^{(j)}\|_2^2 \\ &\stackrel{(a)}{=} L(\mathbf{w}^{(j)}) - \eta \nabla L(\mathbf{w}^{(j)})^T (\tilde{g}_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)})) \\ &\quad + \frac{\eta^2 C_w}{2} \|\tilde{g}_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)})\|_2^2, \end{aligned} \quad (36)$$

where (a) follows from the definition of the model update rule using SGD in (2) with gradient $\tilde{g}_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)})$.

Taking the expectation of both sides over the randomness in $\tilde{g}_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)})$, we get

$$\begin{aligned} \mathbb{E} [L(\mathbf{w}^{(j+1)})] &\leq L(\mathbf{w}^{(j)}) - \eta \nabla L(\mathbf{w}^{(j)})^T (\mathbb{E} [\tilde{g}_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)})]) \\ &\quad + \frac{\eta^2 C_w}{2} \mathbb{E} [\|\tilde{g}_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)})\|_2^2], \end{aligned} \quad (37)$$

Note that by definition of $\tilde{g}_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)})$, we have that the expectation of the computed gradient is $\mathbb{E} [\tilde{g}_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)})] = \nabla L(\mathbf{w}^{(j)}) - \Delta_j$, where $\Delta_j = \frac{1}{N} \sum_{i=1}^N \nabla \ell(\mathbf{w}, \mathbf{z}^{(i)}) \mathbb{I}(\ell(\mathbf{w}^{(j)}, \mathbf{z}^{(i)}) \leq \ell_{th}^{(j)})$. Additionally, we have that

$$\mathbb{E} \left[\left(\tilde{g}_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)}) \right)^2 \right] \leq \mathbb{E} \left[\left(\hat{g}_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)}) \right)^2 \right] = \|\nabla L(\mathbf{w})\|_2^2 + \tilde{B}^2, \quad (38)$$

Thus, by substituting this bounds on first and second moments of $\tilde{g}_{\mathbf{z}^{(j)}}$ in (37), and using the fact that $\eta \leq \frac{1}{C_w}$, we have that

$$\begin{aligned} \mathbb{E} [L(\mathbf{w}^{(j+1)})] &\leq L(\mathbf{w}^{(j)}) - \eta \nabla L(\mathbf{w}^{(j)})^T (\nabla L(\mathbf{w}^{(j)}) - \Delta_j) \\ &\quad + \frac{\eta^2 C_w}{2} \|\nabla L(\mathbf{w}^{(j)}) - \Delta_j\|_2^2 + \frac{\eta^2 C_w}{2} \tilde{B}^2 \\ &\quad - \frac{\eta^2 C_w}{2} \|\Delta_j\|_2^2 + \eta^2 C_w \Delta_j^T \nabla L(\mathbf{w}^{(j)}) \\ &= L(\mathbf{w}^{(j)}) - \eta \left(1 - \frac{\eta C_w}{2}\right) \|\nabla L(\mathbf{w}^{(j)}) - \Delta_j\|_2^2 + \frac{\eta^2 C_w}{2} \tilde{B}^2 \\ &\quad - \eta \left(1 - \eta C_w\right) \Delta_j^T \nabla L(\mathbf{w}^{(j)}) + \eta \left(1 - \frac{\eta C_w}{2}\right) \|\Delta_j\|_2^2 \\ &\leq L(\mathbf{w}^{(j)}) - \frac{\eta}{2} \|\nabla L(\mathbf{w}^{(j)}) - \Delta_j\|_2^2 \\ &\quad - \eta \left(1 - \eta C_w\right) \Delta_j^T \nabla L(\mathbf{w}^{(j)}) + \frac{\eta}{2} (\tilde{B}^2 + 2\|\Delta_j\|_2^2). \end{aligned} \quad (39)$$

Combining (35), (39), we get

$$\begin{aligned} \mathbb{E} [L(\mathbf{w}^{(j+1)})] &\leq L(\mathbf{w}^*) + \nabla L(\mathbf{w}^{(j)})^T (\mathbf{w}^{(j)} - \mathbf{w}^*) \\ &\quad - \frac{\eta}{2} \|\nabla L(\mathbf{w}^{(j)}) - \Delta_j\|_2^2 + \frac{\eta}{2} (\tilde{B}^2 + 2\|\Delta_j\|_2^2) \\ &\quad - \eta \left(1 - \eta C_w\right) \Delta_j^T \nabla L(\mathbf{w}^{(j)}). \end{aligned} \quad (40)$$

By adding and subtracting $\frac{1}{2\eta} \|\mathbf{w}^{(j)} - \mathbf{w}^*\|_2^2$ to complete square we get that

$$\mathbb{E} [L(\mathbf{w}^{(j+1)})] \leq L(\mathbf{w}^*) + \frac{1}{2\eta} \|\mathbf{w}^{(j)} - \mathbf{w}^*\|_2^2$$

$$\begin{aligned} &- \frac{1}{2\eta} \|\mathbf{w}^{(j)} - \mathbf{w}^* - \eta(\nabla L(\mathbf{w}^{(j)}) - \Delta_j)\|_2^2 \\ &\quad + \Delta_j^T (\mathbf{w}^{(j)} - \mathbf{w}^*) - \eta(1 - \eta C_w) \Delta_j^T \nabla L(\mathbf{w}^{(j)}) \\ &\quad + \frac{\eta}{2} (\tilde{B}^2 + 2\|\Delta_j\|_2^2) \\ &\leq L(\mathbf{w}^*) - \frac{1}{2\eta} \mathbb{E} \left[\|\mathbf{w}^{(j)} - \mathbf{w}^* - \eta \tilde{g}_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)})\|_2^2 \right] \\ &\quad + \frac{1}{2\eta} \|\mathbf{w}^{(j)} - \mathbf{w}^*\|_2^2 + \Delta_j^T (\mathbf{w}^{(j)} - \mathbf{w}^*) \\ &\quad + \eta^2 C_w \Delta_j^T \nabla L(\mathbf{w}^{(j)}) + \eta (\tilde{B}^2 + \frac{1}{2} \|\Delta_j\|_2^2) \\ &\stackrel{(a)}{=} L(\mathbf{w}^*) - \frac{1}{2\eta} \mathbb{E} \left[\|\mathbf{w}^{(j+1)} - \mathbf{w}^*\|_2^2 \right] + \frac{1}{2\eta} \|\mathbf{w}^{(j)} - \mathbf{w}^*\|_2^2 \\ &\quad + \Delta_j^T (\mathbf{w}^{(j)} - \mathbf{w}^*) + \eta^2 C_w \Delta_j^T \nabla L(\mathbf{w}^{(j)}) \\ &\quad + \eta (\tilde{B}^2 + \frac{1}{2} \|\Delta_j\|_2^2) \\ &\stackrel{(b)}{\leq} L(\mathbf{w}^*) - \frac{1}{2\eta} \mathbb{E} \left[\|\mathbf{w}^{(j+1)} - \mathbf{w}^*\|_2^2 - \|\mathbf{w}^{(j)} - \mathbf{w}^*\|_2^2 \right] \\ &\quad + \|\Delta_j\|_2 \|\mathbf{w}^{(j)} - \mathbf{w}^*\|_2 + \eta \|\Delta_j\|_2 \|\nabla L(\mathbf{w}^{(j)})\|_2 \\ &\quad + \eta (\tilde{B}^2 + \frac{1}{2} \|\Delta_j\|_2^2) \\ &= L(\mathbf{w}^*) - \frac{1}{2\eta} \mathbb{E} \left[\|\mathbf{w}^{(j+1)} - \mathbf{w}^*\|_2^2 - \|\mathbf{w}^{(j)} - \mathbf{w}^*\|_2^2 \right] + \eta \tilde{B}^2 \\ &\quad + \|\Delta_j\|_2 \left(\|\mathbf{w}^{(j)} - \mathbf{w}^*\|_2 + \eta \|\nabla L(\mathbf{w}^{(j)})\|_2 + \frac{\eta}{2} \|\Delta_j\|_2 \right), \end{aligned} \quad (41)$$

where: (a) followed from the update rule of $\mathbf{w}^{(j+1)}$ from $\mathbf{w}^{(j)}$ and $g_{\mathbf{z}^{(j)}}(\mathbf{w}^{(j)})$; (b) from Cauchy-Schwartz inequality. We now want to find bounds for the terms $\|\Delta_j\|_2$, $\|\mathbf{w}^{(j)} - \mathbf{w}^*\|_2$ and $\|\nabla L(\mathbf{w}^{(j)})\|_2$. Note that from our assumptions, we have that $\|\mathbf{w}^{(j)} - \mathbf{w}^*\|_2 \leq \tilde{D}$.

For a point \mathbf{z} , let $\mathbf{w}'_{\mathbf{z}}$ be the minimizer of $\ell(\cdot, \mathbf{z})$, then we have that

$$\begin{aligned} \|\nabla \ell(\mathbf{w}, \mathbf{z}) - \nabla \ell(\mathbf{w}'_{\mathbf{z}}, \mathbf{z})\|_2^2 &\stackrel{(a)}{\leq} 2C_w (\nabla \ell(\mathbf{w}'_{\mathbf{z}}, \mathbf{z}))^T (\mathbf{w}'_{\mathbf{z}} - \mathbf{w}) \\ &\quad + \ell(\mathbf{w}, \mathbf{z}) - \ell(\mathbf{w}'_{\mathbf{z}}, \mathbf{z}) \\ &\stackrel{(b)}{\leq} 2C_w (\nabla \ell(\mathbf{w}'_{\mathbf{z}}, \mathbf{z}))^T (\mathbf{w}'_{\mathbf{z}} - \mathbf{w}) + \ell(\mathbf{w}, \mathbf{z}), \end{aligned}$$

where (a) is a consequence of the smoothness of $\ell(\mathbf{w}, \mathbf{z})$ and (b) follows from our assumption that $\ell(\mathbf{w}, \mathbf{z}) \geq 0$. By making use of the above equation and the fact that $\nabla \ell(\mathbf{w}'_{\mathbf{z}}, \mathbf{z}) = 0$ for a given \mathbf{z} , we get that $\|\nabla \ell(\mathbf{w}, \mathbf{z})\|_2^2 \leq 2C_w L(\mathbf{w}, \mathbf{z})$. Using this bound on $\|\nabla \ell(\mathbf{w}, \mathbf{z})\|_2^2$, we have that

$$\begin{aligned} \|\Delta_j\|_2 &= \left\| \frac{1}{N} \sum_{i=1}^N \nabla \ell(\mathbf{w}, \mathbf{z}^{(i)}) \mathbb{I}(\ell(\mathbf{w}^{(j)}, \mathbf{z}^{(i)}) \leq \ell_{th}^{(j)}) \right\|_2 \\ &\leq \frac{1}{N} \sum_{i=1}^N \left\| \nabla \ell(\mathbf{w}, \mathbf{z}^{(i)}) \right\|_2 \mathbb{I}(\ell(\mathbf{w}^{(j)}, \mathbf{z}^{(i)}) \leq \ell_{th}^{(j)}) \\ &\leq \frac{1}{N} \sum_{i=1}^N \left(\sqrt{2C_w \ell(\mathbf{w}, \mathbf{z}^{(i)})} \right) \mathbb{I}(\ell(\mathbf{w}^{(j)}, \mathbf{z}^{(i)}) \leq \ell_{th}^{(j)}) \\ &\leq \sqrt{2C_w \ell_{th}^{(j)}}. \end{aligned} \quad (42)$$

Finally, note that the smoothness property implies that $\|\nabla \ell(\mathbf{w}, \mathbf{z})\|_2 \leq 2C_w \tilde{D}$, which using the same steps as in (42)

implies that $\|\Delta_j\|_2 \leq 2C_w\tilde{D}$. Additionally, it also implies that $\|\nabla L(\mathbf{w})\|_2 \leq 2C_w\tilde{D}$.

By substituting these upper bounds in (41), we get that

$$\begin{aligned} \mathbb{E} \left[L(\mathbf{w}^{(j+1)}) \right] &\leq L(\mathbf{w}^*) - \frac{1}{2\eta} \mathbb{E} \left[\|\mathbf{w}^{(j+1)} - \mathbf{w}^*\|_2^2 \right. \\ &\quad \left. - \|\mathbf{w}^{(j)} - \mathbf{w}^*\|_2^2 \right] + \eta\tilde{B}^2 \\ &\quad + \|\Delta_j\|_2 \left(\|\mathbf{w}^{(j)} - \mathbf{w}^*\|_2 + \eta\|\nabla L(\mathbf{w}^{(j)})\|_2 + \frac{\eta}{2}\|\Delta_j\|_2 \right) \\ &= L(\mathbf{w}^*) - \frac{1}{2\eta} \mathbb{E} \left[\|\mathbf{w}^{(j+1)} - \mathbf{w}^*\|_2^2 - \|\mathbf{w}^{(j)} - \mathbf{w}^*\|_2^2 \right] \\ &\quad + \left(\sqrt{2C_w\ell_{th}^{(j)}} \right) \left(\tilde{D} + 2\eta C_w\tilde{D} + \eta C_w\tilde{D} \right) + \eta\tilde{B}^2 \\ &= L(\mathbf{w}^*) - \frac{1}{2\eta} \mathbb{E} \left[\|\mathbf{w}^{(j+1)} - \mathbf{w}^*\|_2^2 - \|\mathbf{w}^{(j)} - \mathbf{w}^*\|_2^2 \right] \\ &\quad + \tilde{D}\sqrt{2C_w\ell_{th}^{(j)}}(1 + 3\eta C_w) + \eta\tilde{B}^2. \end{aligned} \quad (43)$$

And from convexity and $\sum_{j=1}^n \sqrt{\ell_{th}^{(j)}} \leq \sqrt{n}$, we have

$$\begin{aligned} \mathbb{E} \left[\frac{1}{n} \sum_{j=1}^n L(\mathbf{w}^{(j)}) \right] &\leq L(\mathbf{w}^*) + \frac{\|\mathbf{w}^{(0)} - \mathbf{w}^*\|_2^2}{2n\eta} \\ &\quad + \frac{\tilde{D}\sqrt{2C_w}(1 + 3\eta C_w)}{\sqrt{n}} + \eta\tilde{B}^2. \end{aligned} \quad (44)$$

Substituting $\eta = \frac{1}{B\sqrt{n}}$, we get the result.

REFERENCES

- [1] D. P. Kumar, T. Amgoth, and C. S. R. Annamalai, "Machine learning algorithms for wireless sensor networks: A survey," *Information Fusion*, vol. 49, pp. 1–25, 2019.
- [2] P. Nayak, G. Swetha, S. Gupta, and K. Madhavi, "Routing in wireless sensor networks using machine learning techniques: Challenges and opportunities," *Measurement*, vol. 178, p. 108974, 2021.
- [3] E. U. Warriach and K. Tei, "Fault detection in wireless sensor networks: A machine learning approach," in *2013 IEEE 16th International Conference on Computational Science and Engineering*. IEEE, 2013, pp. 758–765.
- [4] F. Seide, H. Fu, J. Droppo, G. Li, and D. Yu, "1-bit stochastic gradient descent and its application to data-parallel distributed training of speech dnns," in *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.
- [5] N. Strom, "Scalable distributed dnn training using commodity gpu cloud computing," in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [6] W. Wen, C. Xu, F. Yan, C. Wu, Y. Wang, Y. Chen, and H. Li, "Terngrad: Ternary gradients to reduce communication in distributed deep learning," in *Advances in neural information processing systems*, 2017, pp. 1509–1519.
- [7] D. Alistarh, D. Grubic, J. Li, R. Tomioka, and M. Vojnovic, "Qsgd: Communication-efficient sgd via gradient quantization and encoding," in *Advances in Neural Information Processing Systems*, 2017, pp. 1709–1720.
- [8] V. Gandikota, R. K. Maity, and A. Mazumdar, "vqsgd: Vector quantized stochastic gradient descent," *arXiv preprint arXiv:1911.07971*, 2019.
- [9] P. Mayekar and H. Tyagi, "Ratq: A universal fixed-length quantizer for stochastic optimization," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 1399–1409.
- [10] H. Touvron, A. Vedaldi, M. Douze, and H. Jégou, "Fixing the train-test resolution discrepancy," in *Advances in Neural Information Processing Systems*, 2019, pp. 8252–8262.
- [11] E. Real, A. Aggarwal, Y. Huang, and Q. V. Le, "Regularized evolution for image classifier architecture search," in *Proceedings of the aaai conference on artificial intelligence*, vol. 33, 2019, pp. 4780–4789.
- [12] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT*. Association for Computational Linguistics, 2019, pp. 4171–4186.
- [13] A. L. Perryman, D. Inoyama, J. S. Patel, S. Ekins, and J. S. Freundlich, "Pruned machine learning models to predict aqueous solubility," *ACS omega*, vol. 5, no. 27, pp. 16562–16567, 2020.
- [14] A. L. Perryman, T. P. Stratton, S. Ekins, and J. S. Freundlich, "Predicting mouse liver microsomal stability with pruned machine learning models and public data," *Pharmaceutical research*, vol. 33, no. 2, pp. 433–449, 2016.
- [15] S. Ghadimi and G. Lan, "Stochastic first-and zeroth-order methods for nonconvex stochastic programming," *SIAM Journal on Optimization*, vol. 23, no. 4, pp. 2341–2368, 2013.
- [16] J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, M. Mao, M. Ranzato, A. Senior, P. Tucker, K. Yang *et al.*, "Large scale distributed deep networks," in *Advances in neural information processing systems*, 2012, pp. 1223–1231.
- [17] D. Yu, A. Eversole, M. Seltzer, K. Yao, Z. Huang, B. Guenter, O. Kuchaiev, Y. Zhang, F. Seide, H. Wang *et al.*, "An introduction to computational networks and the computational network toolkit," Microsoft Technical Report MSR-TR-2014-112, Tech. Rep., 2014.
- [18] C. M. De Sa, C. Zhang, K. Olukotun, and C. Ré, "Taming the wild: A unified analysis of hogwild-style algorithms," in *Advances in neural information processing systems*, 2015, pp. 2674–2682.
- [19] D. Alistarh, T. Hoefler, M. Johansson, N. Konstantinov, S. Khirirat, and C. Renggli, "The convergence of sparsified gradient methods," in *Advances in Neural Information Processing Systems*, 2018, pp. 5973–5983.
- [20] Y. Lin, S. Han, H. Mao, Y. Wang, and B. Dally, "Deep gradient compression: Reducing the communication bandwidth for distributed training," in *International Conference on Learning Representations*, 2018.
- [21] S. U. Stich, J.-B. Cordonnier, and M. Jaggi, "Sparsified sgd with memory," in *Advances in Neural Information Processing Systems*, 2018, pp. 4447–4458.
- [22] J. Wu, W. Huang, J. Huang, and T. Zhang, "Error compensated quantized sgd and its applications to large-scale distributed optimization," in *International Conference on Machine Learning*, 2018, pp. 5325–5333.
- [23] D. Basu, D. Data, C. Karakus, and S. Diggavi, "Qsparse-local-sgd: Distributed sgd with quantization, sparsification and local computations," in *Advances in Neural Information Processing Systems*, 2019, pp. 14695–14706.
- [24] H. Zhang, J. Li, K. Kara, D. Alistarh, J. Liu, and C. Zhang, "Zipml: Training linear models with end-to-end low precision, and a little bit of deep learning," in *International Conference on Machine Learning*, 2017, pp. 4035–4043.
- [25] W. W. Ng, D. S. Yeung, and I. Cloete, "Input sample selection for rbf neural network classification problems using sensitivity measure," in *SMC'03 Conference Proceedings. 2003 IEEE International Conference on Systems, Man and Cybernetics. Conference Theme-System Security and Assurance (Cat. No. 03CH37483)*, vol. 3. IEEE, 2003, pp. 2593–2598.
- [26] X. Xu, T. Liang, J. Zhu, D. Zheng, and T. Sun, "Review of classical dimensionality reduction and sample selection methods for large-scale data processing," *Neurocomputing*, vol. 328, pp. 5–15, 2019.
- [27] B. Mirzasoleiman, J. Bilmes, and J. Leskovec, "Coresets for data-efficient training of machine learning models," in *International Conference on Machine Learning*. PMLR, 2020, pp. 6950–6960.
- [28] P. Zhao and T. Zhang, "Stochastic optimization with importance sampling for regularized loss minimization," in *international conference on machine learning*. PMLR, 2015, pp. 1–9.
- [29] L. A. Jaeckel, *The infinitesimal jackknife*. Bell Telephone Laboratories, 1972.
- [30] F. R. Hampel, "The influence curve and its role in robust estimation," *Journal of the american statistical association*, vol. 69, no. 346, pp. 383–393, 1974.
- [31] R. D. Cook, "Detection of influential observation in linear regression," *Technometrics*, vol. 19, no. 1, pp. 15–18, 1977.
- [32] P. W. Koh and P. Liang, "Understanding black-box predictions via influence functions," in *International Conference on Machine Learning*, 2017, pp. 1885–1894.
- [33] S. Basu, P. Pope, and S. Feizi, "Influence functions in deep learning are fragile," *arXiv preprint arXiv:2006.14651*, 2020.

- [34] M. Kuchnik and V. Smith, "Efficient augmentation via data subsampling," in *International Conference on Learning Representations*, 2018.
- [35] B. Settles, "Active learning literature survey," 2009.
- [36] —, "Curious machines: Active learning with structured instances," Ph.D. dissertation, University of Wisconsin–Madison, 2008.
- [37] S. Tong, *Active learning: theory and applications*. Stanford University, 2001.
- [38] C. E. Monteleoni, "Learning with online constraints: shifting concepts and active learning," 2006.
- [39] A. Gersho and R. M. Gray, *Vector quantization and signal compression*. Springer Science & Business Media, 2012, vol. 159.
- [40] J. Bernstein, Y.-X. Wang, K. Azizzadenesheli, and A. Anandkumar, "signsgd: Compressed optimisation for non-convex problems," in *International Conference on Machine Learning*, 2018, pp. 560–569.
- [41] A. T. Suresh, X. Y. Felix, S. Kumar, and H. B. McMahan, "Distributed mean estimation with limited communication," in *International Conference on Machine Learning*, 2017, pp. 3329–3337.
- [42] S. Bubeck *et al.*, "Convex optimization: Algorithms and complexity," *Foundations and Trends® in Machine Learning*, vol. 8, no. 3-4, pp. 231–357, 2015.
- [43] S. Shalev-Shwartz and S. Ben-David, *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- [44] A. Agarwal, P. L. Bartlett, P. Ravikumar, and M. J. Wainwright, "Information-theoretic lower bounds on the oracle complexity of stochastic convex optimization," *IEEE Transactions on Information Theory*, vol. 5, no. 58, pp. 3235–3249, 2012.
- [45] A. F. Aji and K. Heafield, "Sparse communication for distributed gradient descent," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 2017, pp. 440–445.
- [46] R. M. Gray and T. G. Stockham, "Dithered quantizers," *IEEE Transactions on Information Theory*, vol. 39, no. 3, pp. 805–812, 1993.
- [47] N. Agarwal, A. T. Suresh, F. Yu, S. Kumar, and H. B. McMahan, "cpsgd: communication-efficient and differentially-private distributed sgd," in *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 2018, pp. 7575–7586.
- [48] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [49] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
- [50] N. Singh, D. Data, J. George, and S. Diggavi, "Squarm-sgd: Communication-efficient momentum sgd for decentralized optimization," *arXiv preprint arXiv:2005.07041*, 2020.
- [51] W. Feller, *An introduction to probability theory and its applications, vol 2*. John Wiley & Sons, 2008.



Osama A. Hanna is a Ph.D. candidate in the Electrical and Computer Engineering Department at the University of California, Los Angeles (UCLA). He received his BS and MS degrees in electrical engineering from the Faculty of Engineering Cairo University and Nile University in Egypt in 2014 and 2018 respectively. He received the Award of Excellence from Cairo University in 2014. He received the Masters Fellowship and a graduate Research Assistantship from Nile University for the years 2014-2018. He received the Electrical and Computer

Engineering Department Fellowship from UCLA for the year 2018/2019. His research interests are machine learning, information theory and algorithms.



Yahya H. Ezzeldin is currently a Postdoctoral Research Associate in the Electrical and Computer Engineering department at USC. He received his B.S. and M.S. degrees in Electronics and Communications Engineering from Alexandria University in 2011 and 2014, respectively. He received his Ph.D. degree in Electrical and Computer Engineering from UCLA in 2020. His current research interests include information theory, federated learning and distributed optimization. He worked as a machine learning platform engineer with Intel Corporation in the summer of 2018. He received the 2020-2021 Distinguished Ph.D. Dissertation Award in Signals and Systems from the Electrical and Computer Engineering Department at UCLA. He is also the recipient of the UCLA University Fellowship in 2014, the Henry Samueli Fellowship in 2016 and the Dissertation Year Fellowship at UCLA in 2019.



Christina Fragouli is a Professor in the Electrical and Computer Engineering Department at UCLA. She received the B.S. degree in Electrical Engineering from the National Technical University of Athens, Athens, Greece, and the M.Sc. and Ph.D. degrees in Electrical Engineering from the University of California, Los Angeles. She has worked at the Information Sciences Center, AT&T Labs, Florham Park New Jersey, and the National University of Athens. She also visited Bell Laboratories, Murray Hill, NJ, and DIMACS, Rutgers University. Between 2006–2015 she was an Assistant and Associate Professor in the School of Computer and Communication Sciences, EPFL, Switzerland. She is an IEEE fellow, has served as an Information Theory Society Distinguished Lecturer, and as an Associate Editor for IEEE Communications Letters, for Elsevier Journal on Computer Communication, for IEEE Transactions on Communications, for IEEE Transactions on Information Theory, and for IEEE Transactions on Mobile Communications. She has also served in several IEEE committees, and received awards for her work. Her research interests are in network information flow, network security and privacy, wireless networks and bioinformatics.



Suhas Diggavi is currently a Professor of Electrical and Computer Engineering at UCLA. His undergraduate education is from IIT, Delhi and his PhD is from Stanford University. He has worked as a principal member research staff at AT&T Shannon Laboratories and directed the Laboratory for Information and Communication Systems (LICOS) at EPFL. At UCLA, he directs the Information Theory and Systems Laboratory.

His research interests include information theory and its applications to several areas including machine learning, security & privacy, wireless networks, data compression, cyber-physical systems, bio-informatics and neuroscience; more information can be found at <http://licos.ee.ucla.edu>.

He has received several recognitions for his research including the 2013 IEEE Information Theory Society & Communications Society Joint Paper Award, the 2013 ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc) best paper award, the 2006 IEEE Donald Fink prize paper award among others. He was selected as a Guggenheim fellow in 2021. He also received the 2019 Google Faculty Research Award and 2020 Amazon faculty research award. He served as a IEEE Distinguished Lecturer and also currently serves on board of governors for the IEEE Information theory society. He is a Fellow of the IEEE.

He has been an associate editor for IEEE Transactions on Information Theory, ACM/IEEE Transactions on Networking, IEEE Communication Letters, a guest editor for IEEE Selected Topics in Signal Processing and in the program committees of several IEEE conferences. He has also helped organize IEEE and ACM conferences including serving as the Technical Program Co-Chair for 2012 IEEE Information Theory Workshop (ITW), the Technical Program Co-Chair for the 2015 IEEE International Symposium on Information Theory (ISIT) and General co-chair for Mobihoc 2018. He has 8 issued patents.