



Quantization for Decentralized Learning Under Subspace Constraints

Roula Nassif, Stefan Vlaski, Marco Carpentiero, Vincenzo Matta, Marc Antonini,
Ali H Sayed

► To cite this version:

Roula Nassif, Stefan Vlaski, Marco Carpentiero, Vincenzo Matta, Marc Antonini, et al.. Quantization for Decentralized Learning Under Subspace Constraints. IEEE Transactions on Signal Processing, 2023, 71, pp.2320-2335. <10.1109/TSP.2023.3287333>. <hal-04180842>

HAL Id: hal-04180842

<https://hal.science/hal-04180842v1>

Submitted on 14 Aug 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Quantization for decentralized learning under subspace constraints

Roula Nassif, Stefan Vlaski, Marco Carpentiero, Vincenzo Matta, Marc Antonini, Ali H. Sayed

Abstract—In this paper, we consider decentralized optimization problems where agents have individual cost functions to minimize subject to subspace constraints that require the minimizers across the network to lie in low-dimensional subspaces. This constrained formulation includes consensus or single-task optimization as special cases, and allows for more general task relatedness models such as multitask smoothness and coupled optimization. In order to cope with communication constraints, we propose and study an adaptive decentralized strategy where the agents employ *differential randomized quantizers* to compress their estimates before communicating with their neighbors. The analysis shows that, under some general conditions on the quantization noise, and for sufficiently small step-sizes μ , the strategy is stable both in terms of mean-square error and average bit rate: by reducing μ , it is possible to keep the estimation errors small (on the order of μ) without increasing indefinitely the bit rate as $\mu \rightarrow 0$ when variable-rate quantizers are used. Simulations illustrate the theoretical findings and the effectiveness of the proposed approach, revealing that decentralized learning is achievable at the expense of only a few bits.

Index Terms—Stochastic optimization, decentralized subspace projection, differential quantization, randomized quantizers, stochastic performance analysis, mixing parameter, decentralized learning, decentralized optimization.

I. INTRODUCTION

Mobile phones, wearable devices, and autonomous vehicles are examples of modern distributed networks generating massive amounts of data each day. Due to the growing computational power in these devices and the increasing size of the datasets, coupled with concerns over privacy, federated and decentralized training of statistical models have become desirable and often necessary [2]–[11]. In these approaches, each participating device (which is referred to as *agent* or *node*) has a local training dataset, which is never uploaded to the server. The training data is kept locally on the users' devices, which also serve as computational agents acting on the local data in order to update global models of interest. In applications where communication with a server becomes a bottleneck, *decentralized* topologies (where agents only communicate with their neighbors) are attractive alternatives to federated topologies (where a server connects with all remote devices). Compared to federated approaches [2]–[7],

decentralized implementations reduce the high communication cost on the central server since, in this case, the model updates are exchanged locally between agents without relying on a central coordinator [8]–[12].

In practice, there are several issues that arise in the implementations of decentralized algorithms due to the use of a communication network. For instance, in modern distributed networks comprising a massive number of devices, e.g., thousands of participating smartphones, communication can be slower than local computation by many orders of magnitude (due to limited resources such as energy and bandwidth). Designers are typically limited by an upload bandwidth of 1MB/s or less [2]. While there have been significant works in the literature on solving optimization and inference problems in a decentralized manner [8]–[39], with some exceptions [12], [27]–[39], the large majority of these works is not tailored to the specific challenge of limited communication capabilities. Moreover, the existing works on decentralized approaches are often designed only to solve consensus-based optimization problems. And many of these existing approaches rely mainly on quantization rules that assume that some quantities are represented with very high precision and neglect the associated quantization error in the theoretical analyses.

In this work, we do not make any assumptions about the high-precision representation of specific variables. Moreover, we consider quantization for decentralized learning under subspace constraints by studying the effects of quantization on the performance of the following decentralized stochastic gradient approach from [24], [25]:

$$\begin{cases} \psi_{k,i} = \mathbf{w}_{k,i-1} - \mu \widehat{\nabla_{\mathbf{w}_k} J_k}(\mathbf{w}_{k,i-1}) & (1a) \\ \mathbf{w}_{k,i} = \sum_{\ell \in \mathcal{N}_k} A_{k\ell} \psi_{\ell,i} & (1b) \end{cases}$$

where $\mu > 0$ is a small step-size parameter, \mathcal{N}_k is the neighborhood set of agent k (i.e., the set of nodes connected to agent k by a communication link or edge, including node k itself), $A_{k\ell}$ is an $M_k \times M_\ell$ matrix associated with link (k, ℓ) , $\mathbf{w}_k \in \mathbb{R}^{M_k}$ is the parameter vector at agent k , and $J_k(\mathbf{w}_k) : \mathbb{R}^{M_k} \rightarrow \mathbb{R}$ is a differentiable convex cost associated with agent k . This cost is usually expressed as the expectation of some loss function $L_k(\cdot)$ and written as $J_k(\mathbf{w}_k) = \mathbb{E} L_k(\mathbf{w}_k; \mathbf{y}_k)$, where \mathbf{y}_k denotes the random data at agent k (throughout the paper, random quantities are denoted in boldface). The expectation is computed relative to the distribution of the local data. In the stochastic-optimization framework, the statistical distribution of the data \mathbf{y}_k is usually unknown and, hence, the risks $J_k(\cdot)$ and their gradients $\nabla_{\mathbf{w}_k} J_k(\cdot)$ are unknown. In this case, and instead of using the true gradient, it is common to use approximate gradient

The work of R. Nassif was supported in part by ANR JCJC grant ANR-22-CE23-0015-01 (CEDRO project).

A short conference version of this work appears in [1].

R. Nassif and M. Antonini are with Université Côte d'Azur, I3S Laboratory, CNRS, France (email: {roula.nassif,marc.antonini}@unice.fr). S. Vlaski is with Imperial College London, UK (e-mail: s.vlaski@imperial.ac.uk). M. Carpentiero and V. Matta are with University of Salerno, Italy (e-mail: {mcarpentiero,vmatta}@unisa.it). A. H. Sayed is with the Institute of Electrical and Micro Engineering, EPFL, Switzerland (e-mail: ali.sayed@epfl.ch).

This paper has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the author. The material includes experimental results.

vectors such as $\widehat{\nabla_{w_k} J_k}(w_k) = \nabla_{w_k} L_k(w_k; \mathbf{y}_{k,i})$, where $\mathbf{y}_{k,i}$ represents the data realization observed at iteration i [8]. We note that a unique feature of algorithm (1) is the utilization of *matrix* valued combination weights, as opposed to scalar weighting as is commonly employed in conventional consensus and diffusion optimization [8], [13]. As explained in the next paragraph, this generalization allows the network to solve a broader class of multitask optimization problems beyond classical consensus. Multitask learning is suitable for network applications where regional differences in the data require more complex models and more flexible algorithms than single-task or consensus implementations. In multitask networks, agents generally need to estimate and track multiple distinct, though related, objectives. For instance, in distributed power state estimation problems, the local state vectors at neighboring control centers may overlap partially since the areas in a power system are interconnected [19]. Likewise, in weather forecasting applications, regional differences in the collected data distributions require agents to exploit the correlation profile in the data for enhanced decision rules [21]. Other multitask network applications include distributed minimum-cost flow [17], [18], distributed active noise control [16], and distributed sub-optimal beamforming [24].

Let N denote the total number of agents and let $\mathbf{w} = \text{col}\{w_1, \dots, w_N\}$ denote the M -th dimensional vector (where $M = \sum_{k=1}^N M_k$) collecting the parameter vectors from across the network. Let \mathcal{A} denote the $N \times N$ block matrix whose (k, ℓ) -th block is $A_{k\ell}$ if $\ell \in \mathcal{N}_k$ and 0 otherwise. It was shown in [24, Theorem 1] that, for sufficiently small μ and for a combination matrix \mathcal{A} satisfying:

$$\mathcal{A}\mathcal{U} = \mathcal{U}, \quad \mathcal{U}^\top \mathcal{A} = \mathcal{U}^\top, \quad \text{and} \quad \rho(\mathcal{A} - \mathcal{P}_\mathcal{U}) < 1, \quad (2)$$

where $\rho(\cdot)$ denotes the spectral radius of its matrix argument, \mathcal{U} is any given $M \times P$ full-column rank matrix (with $P \ll M$) that is assumed to be semi-unitary, i.e., its columns are orthonormal ($\mathcal{U}^\top \mathcal{U} = I_P$), and $\mathcal{P}_\mathcal{U} = \mathcal{U}\mathcal{U}^\top$ is the orthogonal projection matrix onto $\text{Range}(\mathcal{U})$, strategy (1) will converge in the mean-square-error sense to the solution of the following subspace constrained optimization problem :

$$\begin{aligned} \mathbf{w}^o = \arg \min_{\mathbf{w}} J^{\text{glob}}(\mathbf{w}) &\triangleq \sum_{k=1}^N J_k(w_k) \\ \text{subject to } \mathbf{w} &\in \text{Range}(\mathcal{U}). \end{aligned} \quad (3)$$

In particular, it was shown that $\limsup_{i \rightarrow \infty} \mathbb{E} \|\mathbf{w}_k^o - \mathbf{w}_{k,i}\|^2 = O(\mu)$ for all k , where \mathbf{w}_k^o is the k -th $M_k \times 1$ subvector of \mathbf{w}^o . As explained in [11], [26], [24, Sec. II], by properly selecting \mathcal{U} and \mathcal{A} , strategy (1) can be employed to solve different decentralized optimization problems such as (i) consensus or single-task optimization (where, as explained in Remark 1 further ahead, the agents' objective is to reach a consensus on the minimizer of the aggregate cost $\sum_{k=1}^N J_k(w)$) [8]–[10], [13], (ii) decentralized coupled optimization (where the parameter vectors to be estimated at neighboring agents are partially overlapping) [16], [17], [19], [20], [22], and (iii) multitask inference under smoothness (where the network parameter vector \mathbf{w} to be estimated is smooth w.r.t. the underlying network topology) [23], [24]. For instance, while projecting

onto the space spanned by the vector of all ones allows to enforce consensus across the network (see Remark 1), graph smoothness can in general be promoted by projecting onto the space spanned by the eigenvectors of the graph Laplacian corresponding to small eigenvalues (see the simulation section VI for an illustration). Besides being capable of solving single-task or consensus-based optimization problems, the constrained formulation (3) and its decentralized solution (1) are general enough to apply to a wide variety of network applications, including those listed in the previous paragraph.

The first step (1a) in algorithm (1) is the *self-learning* step corresponding to the stochastic gradient descent step on the individual cost $J_k(\cdot)$. This step is followed by the *social learning* step (1b) where agent k receives the intermediate estimates $\{\psi_{\ell,i}\}$ from its neighbors $\ell \in \mathcal{N}_k$ and combines them through $\{A_{k\ell}\}$ to form $\mathbf{w}_{k,i}$, which corresponds to the estimate of \mathbf{w}_k^o at agent k and iteration i . To alleviate the communication bottleneck resulting from the exchange of the intermediate estimates among agents over many iterations, quantized communication must be considered. In this paper, we study the effect of quantization on the convergence properties of the decentralized learning approach (1). First, we describe in Sec. II the class of randomized quantizers considered in this study. Then, we propose in Sec. III a *differential randomized quantization* strategy for solving problem (3) – see (28) further ahead. Compared with the unquantized version (1), the new approach consists of three steps where a quantization step is added. Interestingly, instead of exchanging compressed versions of the intermediate estimates $\{\psi_{\ell,i}\}$, and due to differential quantization, the agents in the proposed solution exchange compressed versions of the differences between subsequent iterates (which tend to have a reduced range when compared to the intermediate estimates). At the receiver side, prediction rules are implemented in order to reconstruct the intermediate estimates. This step leads to a new set of intermediate estimates $\{\phi_{\ell,i}\}$ which are then: (i) stored at the receiver in order to be used as *predictors* in the next iteration, and (ii) combined according to a modified version of the combination step. In the modified version, a *mixing* parameter γ is introduced in order to control the mixing speed of the algorithm. This allows to control the network stability in situations where quantization can lead to network instability. We establish in Sec. IV that, under some general conditions on the quantization noise and mixing parameter, and for sufficiently small step-sizes μ , the decentralized quantized approach is stable in the mean-square error sense. In addition to investigating the mean-square-error stability, we characterize the steady-state average bit rate of the proposed approach when *variable-rate* quantizers are used. The analysis shows that, by properly designing the quantization operators, the iterates generated by the quantized decentralized adaptive implementation lead to small estimation errors on the order of μ (as it happens in the *ideal case without quantization*), while concurrently guaranteeing a bounded average bit rate as $\mu \rightarrow 0$. While there exist several useful works in the literature that study decentralized learning approaches in the presence of differential quantization [12], [31]–[40], these works investigate standard consensus or single-task

optimization, and do not consider the subspace constrained formulation (3) or multitask variants. Moreover, with some exceptions that consider deterministic optimization [31], [39], [40], the analyses conducted in these works assume that some quantities (e.g., the norm or some components of the vector to be quantized) are represented with very high precision (e.g., machine precision) and the associated quantization error is neglected. On the other hand, the analysis in the current work is not limited to standard consensus optimization and does not assume a machine precision representation of some quantities – a detailed discussion of related work is provided in Sec. III-B.

Notation: All vectors are column vectors. Random quantities are denoted in boldface. Matrices are denoted in uppercase letters while vectors and scalars are denoted in lower-case letters. The symbol $(\cdot)^\top$ denotes matrix transposition. The operator $\text{col}\{\cdot\}$ stacks the column vector entries on top of each other. The operator $\text{diag}\{\cdot\}$ forms a matrix from block arguments by placing each block immediately below and to the right of its predecessor. The symbol \otimes denotes the Kronecker product. The $M \times M$ identity matrix is denoted by I_M . The abbreviation “w.p.” is used for “with probability”. The Gaussian distribution with mean m and covariance C is denoted by $\mathcal{N}(m, C)$. The notation $\alpha = O(\mu)$ signifies that there exist two positive constants c and μ_0 such that $|\alpha| \leq c\mu$ for all $\mu \leq \mu_0$. A vector of all zeros is denoted by 0 . For tables, the header is not counted as a row (i.e., “first row” means the first row after the header).

II. RANDOMIZED QUANTIZERS

In this paper, we study decentralized learning under subspace constraints in the presence of quantized communications. *Randomized quantizers* $\mathcal{Q}(\cdot)$ ¹ will be employed instead of deterministic quantizers. As explained in [33], deterministic quantizers can lead to severe estimation biases in inference problems. To overcome this issue, randomized quantizers $\mathcal{Q}(\cdot)$ are commonly used to compensate for the bias (on average, over time) [12], [31]–[34], [38]. This section is devoted to describing the class of randomized quantizers considered throughout the study.

For any deterministic input $x \in \mathbb{R}^L$ with L representing a generic vector length, the randomized quantizer $\mathcal{Q}(\cdot)$ is characterized in terms of a probability $\mathbb{P}[\mathcal{Q}(x) = y]$ for any y belonging to the set of output levels of the quantizer. We consider randomized quantizers $\mathcal{Q}(\cdot)$ satisfying the following general property, which as explained in the sequel, relaxes the condition on the mean-square error from [5], [12], [32]–[34], [38].

Property 1. (Unbiasedness and variance bound). *The randomized quantizer $\mathcal{Q}(\cdot)$ satisfies the following two conditions:*

$$\mathbb{E}[x - \mathcal{Q}(x)] = 0, \quad (4)$$

$$\mathbb{E}\|x - \mathcal{Q}(x)\|^2 \leq \beta_q^2 \|x\|^2 + \sigma_q^2, \quad (5)$$

for some $\beta_q^2 \geq 0$ and $\sigma_q^2 \geq 0$, and where the expectations are evaluated w.r.t. the randomness of $\mathcal{Q}(\cdot)$. When the quantizers

are applied to a random input x , conditions (4) and (5) become:

$$\mathbb{E}[x - \mathcal{Q}(x)|x] = 0, \quad (6)$$

$$\mathbb{E}[\|x - \mathcal{Q}(x)\|^2|x] \leq \beta_q^2 \|x\|^2 + \sigma_q^2. \quad (7)$$

□

Property 1 is satisfied by many randomized quantization operators of interest in decentralized learning. Table I further ahead lists some typical choices (a detailed comparison of the various schemes will be provided in Sec. III-B). Many existing works focus on studying decentralized learning approaches in the presence of randomized quantizers that satisfy the *unbiasedness* condition (4) and the *variance bound* (5) with the *absolute noise* term $\sigma_q^2 = 0$ [5], [12], [32]–[34], [38]. In contrast, the analysis in the current work is general and does not require σ_q^2 to be zero. As we will explain in Sec. III-B, neglecting the effect of σ_q^2 requires that some quantities (e.g., the norm of the vector to be quantized) are represented with no quantization error, in practice at the machine precision. In the following, we describe a useful framework for designing randomized quantizers that do not require high-precision quantization of specific variables.

A. Uniform and non-uniform randomized quantizers

1) *Quantizers’ design:* Let $x \in \mathbb{R}^L$ denote the input vector to be quantized with x_j representing the j -th element of x . We consider a general quantization rule $\mathcal{Q} : \mathbb{R}^L \rightarrow \mathbb{R}^L$ of the following form – see Fig. 1 for an illustration:

$$[\mathcal{Q}(x)]_j = y_{n(x_j)}, \quad (8)$$

where $[\mathcal{Q}(x)]_j$ denotes the j -th element of $\mathcal{Q}(x)$ and $y_{n(x_j)}$ is the quantization output level (defined further ahead in (11)) associated with a realization of the random index $n(x_j) \in \{m, m+1\}$. The *probabilistic* rule to choose $n(x_j)$ is as follows:

$$n(x_j) = \begin{cases} m, & \text{w.p. } \frac{y_{m+1} - x_j}{y_{m+1} - y_m}, \\ m+1, & \text{w.p. } \frac{x_j - y_m}{y_{m+1} - y_m}. \end{cases} \quad (9)$$

Regarding the integer m in (9), and motivated by the so-called companding procedure [41], which can be conveniently described in terms of a non-linear function and its inverse, it is found according to (the dependence of m on x_j is left implicit in (9) for ease of notation):

$$m = \lfloor g(x_j) \rfloor, \quad (10)$$

where $\lfloor \cdot \rfloor$ denotes the floor function and $g : \mathbb{R} \rightarrow \mathbb{R}$ is some strictly increasing continuous function. Given the inverse function $h = g^{-1}$, the output level associated with an index m is defined by:

$$y_m = h(m). \quad (11)$$

By evaluating the expected value of $[\mathcal{Q}(x)]_j$ w.r.t. the quantizer randomness, we obtain:

$$\mathbb{E}[\mathcal{Q}(x)]_j \stackrel{(8),(9)}{=} \left(y_m \frac{y_{m+1} - x_j}{y_{m+1} - y_m} + y_{m+1} \frac{x_j - y_m}{y_{m+1} - y_m} \right) = x_j, \quad (12)$$

which establishes the unbiasedness condition (4) in Property 1.

Example 1. (Randomized uniform or dithered quantizer [40], [42]). By choosing:

¹Since the output of a randomized quantizer is random even for deterministic input, we use the boldface notation $\mathcal{Q}(\cdot)$ to refer to randomized quantizers.




Fig. 1. An illustration of the non-uniform randomized quantization scheme defined by (8)–(11). The real axis partition is constructed by choosing the functions $g(\cdot)$ and $h(\cdot)$ according to (15) with $a = 4.17$ and $b = 10$.

$$g(t) = \frac{t}{\Delta}, \quad h(t) = \Delta \cdot t, \quad (13)$$

we obtain the *uniform* quantizer described in Table I (second row) with a quantization step $\Delta > 0$. Condition (5) (with $\beta_q^2 = 0$ and $\sigma_q^2 = \frac{L\Delta^2}{4}$) can be established by letting $p = \frac{y_{m+1} - x_j}{y_{m+1} - y_m}$ and by evaluating the variance:

$$\begin{aligned} \mathbb{E}(x_j - [\mathcal{Q}(x)]_j)^2 &= (x_j - y_m)^2 p + (x_j - y_{m+1})^2 (1 - p) \\ &= \Delta^2 p(1 - p) \leq \frac{\Delta^2}{4}, \end{aligned} \quad (14)$$

where we replaced $y_{m+1} - y_m$ by Δ and we used the fact that $p(1 - p) \leq \frac{1}{4}$ for $p \in [0, 1]$. ■

Example 2. (Randomized logarithmic companding). For non-uniform quantizers, one popular choice is *logarithmic companding*, which corresponds to the following direct and inverse non-linear functions [41]:

$$g(t) = \text{sign}(t) a \ln(1 + b|t|), \quad h(t) = \text{sign}(t) \frac{1}{b} \left(e^{\frac{|t|}{a}} - 1 \right), \quad (15)$$

with $a > 0$ and $b > 0$. Two useful results were established in [31] regarding this choice. First, it was shown that, by setting the constants a and b according to:

$$a = \frac{1}{2 \ln(\omega + \sqrt{1 + \omega^2})}, \quad b = \frac{\omega}{\eta}, \quad (16)$$

we obtain the following bound on the variance:

$$\mathbb{E}\|x - \mathcal{Q}(x)\|^2 \leq \left(\omega \|x\| + \sqrt{L}\eta \right)^2. \quad (17)$$

The choice (16) gives the probabilistic ANQ rule reported in Table I (third row) with the non-linear functions given by:

$$\begin{cases} g(t) = \text{sign}(t) \frac{\ln(1 + \frac{\omega}{\eta}|t|)}{2 \ln(\omega + \sqrt{1 + \omega^2})}, \\ h(t) = \text{sign}(t) \frac{\eta}{\omega} \left[(\omega + \sqrt{1 + \omega^2})^{2|t|} - 1 \right]. \end{cases} \quad (18)$$

Second, it was shown in [31] that, for a fixed number of output levels, the choice in (16) maximizes the range of the input variable while still satisfying the error bound (17). This property justifies the efficiency of the choice (16) in terms of quantization bit rate.

Now, by taking the limit of $g(t)$ in (18) as $\omega \rightarrow 0$ and by applying l'Hôpital's rule, we get:

$$\lim_{\omega \rightarrow 0} g(t) = \text{sign}(t) \frac{|t|}{2\eta} = \frac{t}{2\eta}, \quad (19)$$

from which we conclude that the setting where $\omega = 0$ allows us to recover the uniform quantizer (13) of Example 1.

Finally, we can relate the quantizer parameters ω and η to the constants β_q^2 and σ_q^2 appearing in (5). In fact, by applying Jensen's inequality to (17), we can see that the bound in (5) is satisfied with:

$$\beta_q^2 = \frac{\omega^2}{\alpha}, \quad \sigma_q^2 = \frac{L\eta^2}{1 - \alpha}, \quad \text{for any } \alpha \in (0, 1). \quad (20)$$




Fig. 2. Illustration of the variable-rate coding scheme. The non-uniform randomized quantization scheme defined by (8)–(11) is used. The function $g(\cdot)$ is chosen according to (15) with $a = 4.17$ and $b = 10$. In this example, the codeword associated with the j -th element of the ordered set \mathcal{P}_b is given by the b -bit binary representation of the integer $j - 1$. For instance, the codewords associated with the first element in \mathcal{P}_3 (which is -7 according to (21)) and the sixth element in \mathcal{P}_4 (which is -10 according to (22)) are 000 and 0101, respectively. Subsequent integers are separated by a parsing symbol p .

2) *Variable-rate coding scheme and bit budget [31]*: Note that the quantization rule (8)–(11) maps a continuous variable to some integer by partitioning the real line into an infinite number of intervals. Thus, a fixed-rate quantizer (i.e., a quantizer that uses the same number of bits for any input) cannot represent all possible intervals of the partition. On the other hand, assuming a finite support for the quantizer input would require some boundedness assumptions (e.g., on the iterates, on the gradient) that are usually violated in stochastic optimization theory [43]. By following a standard approach in coding theory, we shall instead investigate the use of *variable-rate* quantizers which are able to adapt the bit rate based on the quantizer input, e.g., assigning more bits to larger inputs and less bits to smaller ones. In this work, we illustrate the main concepts by focusing on the variable-rate coding scheme proposed in [31] and described in the following. This rule has the advantage of not requiring any knowledge about the distribution of the variables to be quantized. This is particularly relevant since such knowledge is typically unavailable in learning applications.

Consider the following partition of the set of integers \mathbb{Z} :

$$\begin{aligned} \mathcal{P}_0 &= \{0\}, \quad \mathcal{P}_1 = \{-1, 1\}, \quad \mathcal{P}_2 = \{-3, -2, 2, 3\}, \\ \mathcal{P}_3 &= \{-7, -6, -5, -4, 4, 5, 6, 7\}, \quad \dots \end{aligned} \quad (21)$$

which can be written more compactly as:

$$\mathcal{P}_b = \begin{cases} \{0\}, & \text{if } b = 0 \\ \{-1, 1\}, & \text{if } b = 1 \\ \{-2^b + 1, \dots, -2^{b-1}, 2^{b-1}, \dots, 2^b - 1\}, & \text{if } b \geq 2. \end{cases} \quad (22)$$

Note that the ensemble of sets $\{\mathcal{P}_b | b = 0, 1, 2, \dots\}$ forms a partition of \mathbb{Z} where each set \mathcal{P}_b has cardinality equal to 2^b . Thus, an integer $n \in \mathcal{P}_b$ can be represented with a string of b bits. Accordingly, from (22), the number of bits to represent an integer $n \in \mathbb{Z}$ can be determined according to:

$$b = \lceil \log_2(|n| + 1) \rceil, \quad (23)$$

where $\lceil \cdot \rceil$ denotes the ceiling function. When the rule (8)–(11) is implemented to quantize an input vector $x \in \mathbb{R}^L$, a sequence of integers will be generated according (8) and (9). These integers must be encoded before being transmitted over the communication links—see Fig. 2 for an illustration. Thus, we need to encode sequences of integers n_1, n_2, \dots , with different integers belonging in general to different partitions \mathcal{P}_b . At the receiver side, and since the transmitted integers are unknown,

TABLE I

EXAMPLES OF QUANTIZERS $\mathcal{Q} : \mathbb{R}^L \rightarrow \mathbb{R}^L$ SATISFYING PROPERTY 1. FOR EACH SCHEME, WE REPORT THE QUANTIZATION RULE, THE PARAMETERS β_q^2 AND σ_q^2 IN (5), AND THE BIT-BUDGET. B_{HP} DENOTES THE NUMBER OF BITS REQUIRED TO ENCODE A SCALAR WITH HIGH PRECISION (TYPICAL VALUES FOR B_{HP} ARE 32 OR 64).

Quantizer name	Rule	β_q^2	σ_q^2	Bit-budget
No compression [24]	$\mathcal{Q}(x) = x$	0	0	LB_{HP}
Probabilistic uniform or dithered quantizer [40], [42]	$\mathcal{Q}(x)_j = \Delta \cdot \mathbf{n}(x_j)$ $\mathbf{n}(x_j) = \begin{cases} m, & \text{w.p. } \frac{(m+1)\Delta - x_j}{\Delta}, \\ m+1, & \text{w.p. } \frac{x_j - m\Delta}{\Delta}, \end{cases} \quad m = \lfloor \frac{x_j}{\Delta} \rfloor$ Δ is the quantization step	0	$L \frac{\Delta^2}{4}$	$\mathbf{r}(x)$ defined in (27)
Probabilistic ANQ [31]	$\mathcal{Q}(x)_j = y_{\mathbf{n}(x_j)}$ $\mathbf{n}(x_j) = \begin{cases} m, & \text{w.p. } \frac{y_{m+1} - x_j}{y_{m+1} - y_m}, \\ m+1, & \text{w.p. } \frac{x_j - y_m}{y_{m+1} - y_m}, \end{cases} \quad m = \left\lfloor \text{sign}(x_j) \frac{\ln(1 + \frac{\omega}{\eta} x_j)}{2 \ln(\omega + \sqrt{1 + \omega^2})} \right\rfloor$ $y_m = \text{sign}(m) \frac{\eta}{\omega} \left[(\omega + \sqrt{1 + \omega^2})^{2 m } - 1 \right]$ ω and η are two non-negative design parameters	$2\omega^2$	$2L\eta^2$	$\mathbf{r}(x)$ defined in (27)
Rand-c [34]	$[\mathcal{Q}(x)]_j = \begin{cases} \frac{L}{c} \cdot x_j, & \text{if } x_j \in \Omega_c \\ 0, & \text{otherwise} \end{cases} \quad (c \in \{1, \dots, L\})$ Ω_c is a set of c randomly selected coordinates	$(\frac{L}{c} - 1)$	0	$cB_{\text{HP}} + c\lceil \log_2(L) \rceil$
Randomized Gossip [12]	$\mathcal{Q}(x) = \begin{cases} \frac{x}{q}, & \text{w.p. } q \\ 0, & \text{w.p. } 1 - q \end{cases} \quad (q \in (0, 1])$	$(\frac{1}{q} - 1)$	0	$(LB_{\text{HP}})q$ (on average)
Gradient sparsifier [40]	$[\mathcal{Q}(x)]_j = \begin{cases} \frac{x_j}{q_j}, & \text{w.p. } q_j \\ 0, & \text{w.p. } 1 - q_j \end{cases}$ q_j is the probability that coordinate j is selected	$\frac{1}{\min_j \{q_j\}} - 1$	0	$\sum_{j=1}^L q_j (B_{\text{HP}} + \lceil \log_2(L) \rceil)$ (on average)
QSGD [5]	$[\mathcal{Q}(x)]_j = \ x\ \cdot \text{sign}(x_j) \cdot \frac{\mathbf{n}(x_j, x)}{s}$ $\mathbf{n}(x_j, x) = \begin{cases} m, & \text{w.p. } (m+1) - \frac{ x_j }{\ x\ } s, \\ m+1, & \text{w.p. } \frac{ x_j }{\ x\ } s - m, \end{cases} \quad m = \lfloor s \frac{ x_j }{\ x\ } \rfloor$ s is the number of quantization levels	$\min\left(\frac{L}{s^2}, \frac{\sqrt{L}}{s}\right)$	0	$B_{\text{HP}} + L + L\lceil \log_2(s) \rceil$

the decoder does not know the number of bits used for each integer. In order to solve this issue, the work [31] proposes to consider an overall encoder alphabet \mathcal{S} made of two binary digits plus a parsing symbol \mathbf{p} , namely, $\mathcal{S} = \{0, 1\} \cup \{\mathbf{p}\}$. In this way, we can encode a sequence of integers n_1, n_2, \dots , by first encoding each of the individual integers using a number of bits determined by the corresponding partition, and then using the parsing symbol to separate subsequent integers. For example, a string of the form (see also Fig. 2):

$$\underbrace{111\mathbf{p}}_{n_1} \underbrace{0101\mathbf{p}}_{n_2} \underbrace{10\mathbf{p}}_{n_3} \underbrace{\mathbf{p}}_{n_4} \quad (24)$$

corresponds to integer $n_1 \in \mathcal{P}_3$, followed by $n_2 \in \mathcal{P}_4$, $n_3 \in \mathcal{P}_2$, and $n_4 = 0$ (only the parsing symbol since partition $\mathcal{P}_0 = \{0\}$ requires 0 bits). Accounting for the parsing symbol, observe that the total number of ternary digits (considering the ternary alphabet \mathcal{S}) required to encode an integer n is:

$$1 + \lceil \log_2(|n| + 1) \rceil, \quad (25)$$

which corresponds to a number of bits equal to:

$$\log_2(3)(1 + \lceil \log_2(|n| + 1) \rceil) \quad (26)$$

since one ternary digit is equivalent to $\log_2(3)$ bits of information. Considering now the quantization rule (8)–(11) which is applied entrywise to a vector $x \in \mathbb{R}^L$, we find from (26) that the overall (random) bit budget is equal to:

$$\mathbf{r}(x) = \log_2(3) \sum_{j=1}^L (1 + \lceil \log_2(|\mathbf{n}(x_j)| + 1) \rceil). \quad (27)$$

III. DECENTRALIZED LEARNING IN THE PRESENCE OF QUANTIZED COMMUNICATION

A. Differential randomized quantization approach

Motivated by the approaches proposed in [12], [31]–[39], we equip (1) with a quantization mechanism by proposing the following decentralized multitask learning approach—see Fig. 3:

$$\begin{cases} \psi_{k,i} = \mathbf{w}_{k,i-1} - \mu \widehat{\nabla_{\mathbf{w}_k} J_k}(\mathbf{w}_{k,i-1}) & (28a) \\ \phi_{k,i} = \phi_{k,i-1} + \mathcal{Q}_k(\psi_{k,i} - \phi_{k,i-1}) & (28b) \\ \mathbf{w}_{k,i} = (1 - \gamma)\phi_{k,i} + \gamma \sum_{\ell \in \mathcal{N}_k} A_{k\ell} \phi_{\ell,i} & (28c) \end{cases}$$

where $\gamma \in (0, 1]$ is a *mixing* parameter that tunes the degree of cooperation between agents, and $\mathcal{Q}_k(\cdot)$ is a *randomized* quantizer satisfying Property 1. Note that, since the quantizer characteristics can vary with k , the randomized quantizer becomes $\mathcal{Q}_k(\cdot)$ instead of $\mathcal{Q}(\cdot)$ with a subscript k added to \mathcal{Q} . Observe that *differential quantization* is used in (28b) to leverage possible correlation between subsequent iterates [33]. In this case, instead of communicating compressed versions of the estimates $\psi_{k,i}$, the *prediction* error $\psi_{k,i} - \phi_{k,i-1}$ is quantized at agent k and then transmitted [12], [31]–[39]. At each iteration i , agent k performs the *forward quantization* (see Fig. 3) by mapping the real-valued vector $\psi_{k,i} - \phi_{k,i-1}$ into a quantized vector $\delta_{k,i}$, sends $\delta_{k,i}$ to its neighbors through ideal communication links (i.e., it is assumed that node k can transmit perfectly and reliably $\delta_{k,i}$ to its neighbors), receives $\{\delta_{\ell,i}\}$ from its neighbors $\ell \in \mathcal{N}_k$, and performs the *reconstruction* (see Fig. 3) on each received vector $\delta_{\ell,i}$ by first decoding it and then computing $\{\phi_{\ell,i}\}$ according to step (28b):

$$\phi_{\ell,i} = \phi_{\ell,i-1} + \mathcal{Q}_\ell(\psi_{\ell,i} - \phi_{\ell,i-1}), \quad \ell \in \mathcal{N}_k. \quad (29)$$

Observe that implementing (29) requires storing the *predictors* $\{\phi_{\ell,i-1}\}_{\ell \in \mathcal{N}_k}$ by agent k . The reconstructed vectors $\{\phi_{\ell,i}\}$

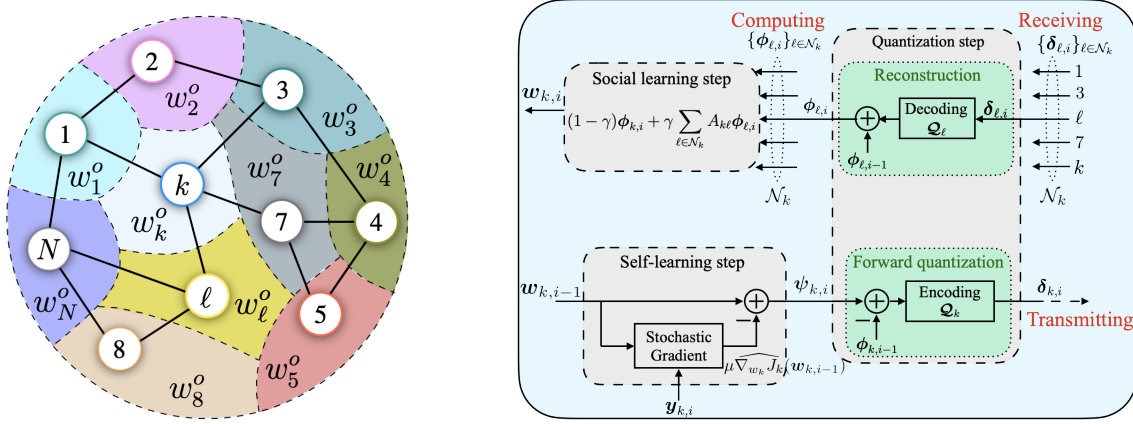


Fig. 3. (Left) An illustration of a multitask network [11]. The objective at agent k is to estimate w_k^o (of dimension $M_k \times 1$), the k -th subvector of w^o in (3). In this example, the neighborhood set of agent k is given by $\mathcal{N}_k = \{1, k, 3, \ell, 7\}$. (Right) The implementation of the quantized multitask approach (28) at agent k . The quantization step consists of (i) the *forward* step where agent k encodes the difference $\psi_{k,i} - \phi_{k,i-1}$ and sends the resulting vector $\delta_{k,i}$ (sequence of symbols or bits) to its neighbors, and (ii) the *reconstruction* step where agent k receives the encoded vectors $\{\delta_{\ell,i}\}_{\ell \in \mathcal{N}_k}$ from its neighbors and decodes them according to (29) to obtain $\{\phi_{\ell,i}\}_{\ell \in \mathcal{N}_k}$. The resulting vectors are then used in the social learning step (28c).

are then combined according to (28c) to produce the estimate $w_{k,i}$. As we will see in Secs. IV and V, in the presence of *relative* quantization noise (i.e., when $\beta_q^2 \neq 0$), the mixing parameter γ in (28c) will be used to control the network stability. However, in the absence of *relative* noise (i.e., when $\beta_q^2 = 0$), the quantization does not affect the network stability, and the parameter γ can be set to one.

Remark 1 (Quantized diffusion-type approach): Most prior literature on decentralized quantized learning focuses primarily on single-task consensus optimization where the objective at each agent is to estimate the *same* L -th dimensional vector w^o given by:

$$w^o = \arg \min_w \sum_{k=1}^N J_k(w). \quad (30)$$

For this reason, and in order to make the paper self-contained for readers interested in problem (30), we explain in this remark how to choose the blocks $\{A_{k\ell}\}$ in (28) in order to solve problem (30) as well. In fact, by setting in (3) $P = L$ and $\mathcal{U} = \frac{1}{\sqrt{N}}(\mathbf{1}_N \otimes I_L)$ where $\mathbf{1}_N$ is the $N \times 1$ vector of all ones, then solving problem (3) will be equivalent to solving the well-studied consensus problem (30). Different algorithms for solving (30) over strongly-connected networks have been proposed [8]–[10], [13]–[15]. By choosing an $N \times N$ doubly-stochastic matrix $A = [a_{k\ell}]$ satisfying:

$$a_{k\ell} \geq 0, \quad A\mathbf{1}_N = \mathbf{1}_N, \quad \mathbf{1}_N^\top A = \mathbf{1}_N^\top, \quad a_{k\ell} = 0 \text{ if } \ell \notin \mathcal{N}_k, \quad (31)$$

the *diffusion* strategy for instance would take the form [8]–[10]:

$$\begin{cases} \psi_{k,i} = w_{k,i-1} - \mu \widehat{\nabla_{w_k} J_k}(w_{k,i-1}) \\ w_{k,i} = \sum_{\ell \in \mathcal{N}_k} a_{k\ell} \psi_{\ell,i} \end{cases} \quad (32a) \quad (32b)$$

This strategy can be written in the form of (1) with $A_{k\ell} = a_{k\ell}I_L$ and $\mathcal{A} = A \otimes I_L$. It can be verified that, when A satisfies (31) over a strongly connected network, the matrix \mathcal{A} will satisfy (2). Consequently, by setting $A_{k\ell} = a_{k\ell}I_L$ in (28c) with a set of combination coefficients $\{a_{k\ell}\}$ satisfying (31), we

obtain a new *quantized diffusion-type approach* for solving the consensus optimization problem (30). \square

B. Related work

Table I provides a list of typical compression schemes with the corresponding quantization noise parameters β_q^2 and σ_q^2 . By comparing the reported schemes, we first observe that the “rand-c”, “randomized gossip”, and “gradient sparsifier” do not really quantize their input. These methods basically map a full vector into a sparse version thereof. In other words, these methods assume that the non-zero vector components are represented with very high (e.g., machine) precision, and that the overall compression gain lies in representing few entries of the input vector. For instance, under the “rand-c” scheme, c randomly selected components of the input vector are encoded with very high precision (32 or 64 bit are typical values for encoding a scalar), and then the resulting bits are communicated over the links in addition to the bits encoding the locations of the selected components. Such schemes should be more properly referred to as *compression operators*, rather than quantizers [12], [31], [34]. The idea behind sparsification operators is that, when the number of vector components is large, the gain resulting from encoding a *few* randomly selected components compensates for the high precision required to represent them. The QSGD scheme in Table I uses a different approximation rule. It assumes that the norm of the input vector is represented with high precision. In addition to encoding the norm, this rule requires L -bits to encode the signs of the vector components and $L\lceil \log_2(s) \rceil$ to encode the levels. It tends to be well-suited for high-dimensional settings where the number of components is large. As it can be observed from Table I, the sparsity-based schemes and the QSGD scheme have an *absolute noise* term $\sigma_q^2 = 0$. As explained in [31], neglecting the effect of σ_q^2 requires that some quantities (e.g., the norm of the input vector in QSGD) are represented with no quantization error². In comparison, the probabilistic uniform

²In particular, it is shown in [31, corollary 9] that the design of unbiased quantizers satisfying condition (7) with $\sigma_q^2 = 0$ using a finite number of quantization points is not feasible.

and ANQ schemes *do not make any assumptions on the high-precision quantization of specific variables* and have a *non-zero absolute noise term* σ_q^2 .

A detailed analysis of the decentralized strategy (28) in the presence of both the *relative* (captured through β_q^2) and the *absolute* (captured through σ_q^2) quantization noise terms will be conducted in Sec. IV. Consequently, our results apply to a general class of quantizers satisfying Property 1, including all those listed in Table I. Compared to prior works, we provide the following advances. First, the analysis results are novel compared to the works [5], [12], [32]–[36] since we do not make any assumptions about the high-precision quantization of specific variables. Although no such assumptions are made in the work [37], the analysis there is performed in the absence of the relative quantization and gradient noise terms. Our results are also novel in comparison to [31] since we consider: i) learning under *subspace constraints*; ii) a *stochastic* optimization setting with non-diminishing step-size; iii) combination policies that can lead to *non-symmetric* combination matrices; and iv) a network *global*, as opposed to local, strong convexity condition on the individual costs – see condition (35) further ahead. Moreover, the work [31] considers only deterministic optimization and does not deal with gradient approximation or mean-square error stability. In addition, it must be noted that even the works [5], [12], [32]–[37] do not consider our more general setting that simultaneously addresses i), ii), iii), and iv).

As the derivations in the next section will reveal, characterizing the behavior of the decentralized learning system under the aforementioned conditions is challenging due to at least two factors. First, the quantization noise that interferes with the operation of the algorithm. It is therefore important to assess its impact on the algorithm's performance and stability by extending the mean-square analyses of the works [8], [9], [21], [24], which consider decentralized learning in the absence of quantization. Second, as we will see later, when we remove the common assumption about the high precision quantization of specific variables (in which case we use variable-rate quantization as suggested in Sec. II-A2), then the quantizer resolution is required to increase as $\mu \rightarrow 0$ to guarantee small mean-square errors. It therefore becomes important to establish the bit rate stability of the network, namely, to show that the average bit rate remains finite as $\mu \rightarrow 0$. To the best of our knowledge, this analysis has not been carried out before. Existing works investigate either fixed-rate quantizers [12], [32]–[34] or variable-rate quantizers without streaming data [31].

In summary, we provide the following main contributions.

- We propose a decentralized strategy for learning and adaptation over networks under subspace constraints. This strategy is able to operate under finite communication rate by employing *differential randomized quantization*.
- We provide a detailed characterization of the proposed approach for a general class of quantizers and compression operators satisfying Property 1, both in terms of mean-square stability and communication resources.
- The analysis reveals the following useful conclusions. First, the quantization error does not impair the network

mean-square stability. Second, in the absence of the absolute quantization noise term, and at the expense of communicating some quantities with high precision, the iterates generated by the quantized approach (28) lead to small estimation errors on the order of μ , as it happens in the unquantized case (1). In the presence of the absolute noise term, the situation becomes more challenging. The analysis reveals that, to guarantee the $O(\mu)$ mean-square-error behavior, the absolute noise term must converge to zero as $\mu \rightarrow 0$. We prove that this result can be achieved with a bit rate that remains *bounded* as $\mu \rightarrow 0$, despite the fact that we are requiring an increasing precision as the step-size decreases. In particular, we illustrate one useful strategy (the variable-rate quantization scheme described in Sec. II-A) that achieves the aforementioned goals.

IV. MEAN-SQUARE AND BIT RATE STABILITY ANALYSIS

A. Modeling assumptions

In this section, we analyze strategy (28) with a matrix \mathcal{A} satisfying (2) by examining the average squared distance between $\mathbf{w}_{k,i}$ and \mathbf{w}_k^o , namely, $\mathbb{E}\|\mathbf{w}_k^o - \mathbf{w}_{k,i}\|^2$, under the following assumptions on the risks $\{J_k(\cdot)\}$, the gradient noise processes $\{\mathbf{s}_{k,i}(\cdot)\}$ defined by [8]:

$$\mathbf{s}_{k,i}(\mathbf{w}_k) \triangleq \nabla_{\mathbf{w}_k} J_k(\mathbf{w}_k) - \widehat{\nabla_{\mathbf{w}_k} J_k}(\mathbf{w}_k), \quad (33)$$

and the randomized quantizers $\{\mathcal{Q}_k(\cdot)\}$.

Assumption 1. (Conditions on individual and aggregate costs). *The individual costs $J_k(\mathbf{w}_k)$ are assumed to be twice differentiable and convex such that:*

$$\lambda_{k,\min} I_{M_k} \leq \nabla_{\mathbf{w}_k}^2 J_k(\mathbf{w}_k) \leq \lambda_{k,\max} I_{M_k}, \quad (34)$$

where $\lambda_{k,\min} \geq 0$ for $k = 1, \dots, N$. It is further assumed that, for any $\{\mathbf{w}_k \in \mathbb{R}^{M_k}\}$, the individual costs satisfy:

$$0 < \lambda_{\min} I_P \leq \mathcal{U}^\top \text{diag} \left\{ \nabla_{\mathbf{w}_k}^2 J_k(\mathbf{w}_k) \right\}_{k=1}^N \mathcal{U} \leq \lambda_{\max} I_P, \quad (35)$$

for some positive parameters $\lambda_{\min} \leq \lambda_{\max}$. \square

As explained in [24], condition (35) ensures that problem (3) has a unique minimizer \mathbf{w}^o .

Assumption 2. (Conditions on gradient noise). *The gradient noise process defined in (33) satisfies for $k = 1, \dots, N$:*

$$\mathbb{E}[\mathbf{s}_{k,i}(\mathbf{w}_{k,i-1}) | \{\phi_{\ell,i-1}\}_{\ell=1}^N] = 0, \quad (36)$$

$$\mathbb{E}[\|\mathbf{s}_{k,i}(\mathbf{w}_{k,i-1})\|^2 | \{\phi_{\ell,i-1}\}_{\ell=1}^N] \leq \beta_{s,k}^2 \|\tilde{\mathbf{w}}_{k,i-1}\|^2 + \sigma_{s,k}^2, \quad (37)$$

for some $\beta_{s,k}^2 \geq 0$ and $\sigma_{s,k}^2 \geq 0$. \square

As explained in [8]–[10], these conditions are satisfied by many cost functions of interest in learning and adaptation such as quadratic and regularized logistic risks. Condition (36) states that the gradient approximation should be unbiased conditioned on the predictors $\{\phi_{\ell,i-1}\}_{\ell=1}^N$. Condition (37) states that the second-order moment of the gradient noise should get smaller for better estimates, since it is bounded by the squared norm of the iterate.

Assumption 3. (Conditions on quantizers). *In step (28b) of the learning approach, each agent k at time i applies to*

the difference $\chi_{k,i} = \psi_{k,i} - \phi_{k,i-1}$ a randomized quantizer $\mathcal{Q}_k(\cdot)$ satisfying Property 1 with quantization noise parameters $\beta_{q,k}^2$ and $\sigma_{q,k}^2$. It is assumed that given the past history, the randomized quantization mechanism depends only on the quantizer input $\chi_{k,i}$. Consequently, from (6) and (7), we get:

$$\mathbb{E}[\chi_{k,i} - \mathcal{Q}_k(\chi_{k,i}) | \mathbf{h}_i] = \mathbb{E}[\chi_{k,i} - \mathcal{Q}_k(\chi_{k,i}) | \chi_{k,i}] = 0, \quad (38)$$

$$\mathbb{E}[\|\chi_{k,i} - \mathcal{Q}_k(\chi_{k,i})\|^2 | \mathbf{h}_i] = \mathbb{E}[\|\chi_{k,i} - \mathcal{Q}_k(\chi_{k,i})\|^2 | \chi_{k,i}] \leq \beta_{q,k}^2 \|\chi_{k,i}\|^2 + \sigma_{q,k}^2, \quad (39)$$

where \mathbf{h}_i is the vector collecting all iterates generated by (28) before the quantizer is applied to $\chi_{k,i}$, namely, $\left\{ \left\{ \phi_{\ell,j} \right\}_{j=0}^{i-1}, \left\{ \psi_{\ell,j} \right\}_{j=1}^i \right\}_{\ell=1}^N$. \square

Assumption 3 formalizes the quantization mechanism (which is in fact a standard mechanism in decentralized learning with randomized quantization) and highlights two fundamental relations (38) and (39) that will be useful in the technical derivations.

B. Network error vector recursion

In order to examine the evolution of the iterates $\mathbf{w}_{k,i}$ generated by algorithm (28) with respect to the minimizer $\mathbf{w}^o = \text{col}\{\mathbf{w}_k^o\}_{k=1}^N$ defined in (3), we start by deriving the network error vector recursion. Let $\tilde{\mathbf{w}}_{k,i} = \mathbf{w}_k^o - \mathbf{w}_{k,i}$, $\tilde{\psi}_{k,i} = \mathbf{w}_k^o - \psi_{k,i}$, and $\tilde{\phi}_{k,i} = \mathbf{w}_k^o - \phi_{k,i}$. Using (33) and the mean-value theorem [44, pp. 24], [8, Appendix D], we can express the stochastic gradient vector appearing in (28a) as follows:

$$\widehat{\nabla_{\mathbf{w}_k} J_k}(\mathbf{w}_{k,i-1}) = -\mathbf{H}_{k,i-1} \tilde{\mathbf{w}}_{k,i-1} + b_k - \mathbf{s}_{k,i}(\mathbf{w}_{k,i-1}) \quad (40)$$

where $\mathbf{H}_{k,i-1} \triangleq \int_0^1 \nabla_{\mathbf{w}_k}^2 J_k(\mathbf{w}_k^o - t\tilde{\mathbf{w}}_{k,i-1}) dt$ and $b_k \triangleq \nabla_{\mathbf{w}_k} J_k(\mathbf{w}_k^o)$. By subtracting \mathbf{w}_k^o from both sides of (28a) and by introducing the following network quantities:

$$\mathbf{b} \triangleq \text{col}\{b_1, \dots, b_N\}, \quad (41)$$

$$\mathbf{s}_i \triangleq \text{col}\{\mathbf{s}_{1,i}(\mathbf{w}_{1,i-1}), \dots, \mathbf{s}_{N,i}(\mathbf{w}_{N,i-1})\}, \quad (42)$$

$$\mathbf{H}_{i-1} \triangleq \text{diag}\{\mathbf{H}_{1,i-1}, \dots, \mathbf{H}_{N,i-1}\}, \quad (43)$$

$$\tilde{\mathbf{w}}_{i-1} \triangleq \text{col}\{\tilde{\mathbf{w}}_{1,i-1}, \dots, \tilde{\mathbf{w}}_{N,i-1}\}, \quad (44)$$

we can show that the network error vector $\tilde{\psi}_i = \text{col}\{\tilde{\psi}_{k,i}\}_{k=1}^N$ evolves according to:

$$\tilde{\psi}_i = (I_M - \mu \mathbf{H}_{i-1}) \tilde{\mathbf{w}}_{i-1} - \mu \mathbf{s}_i + \mu \mathbf{b}. \quad (45)$$

By subtracting \mathbf{w}_k^o from both sides of (28c), by replacing \mathbf{w}_k^o by $(1-\gamma)\mathbf{w}_k^o + \gamma\mathbf{w}_k^o$, and by using $\mathbf{w}_k^o = \sum_{\ell \in \mathcal{N}_k} A_{k\ell} \mathbf{w}_\ell^o$ [24, Sec. III-B], we obtain:

$$\tilde{\mathbf{w}}_{k,i} = (1-\gamma)\tilde{\phi}_{k,i} + \gamma \sum_{\ell \in \mathcal{N}_k} A_{k\ell} \tilde{\phi}_{\ell,i}. \quad (46)$$

From (46), we can show that the network error vector $\tilde{\mathbf{w}}_{i-1}$ in (44) evolves according to:

$$\tilde{\mathbf{w}}_{i-1} = (1-\gamma)\tilde{\phi}_{i-1} + \gamma \mathcal{A} \tilde{\phi}_{i-1} = \mathcal{A}' \tilde{\phi}_{i-1}, \quad (47)$$

where

$$\tilde{\phi}_i \triangleq \text{col}\{\tilde{\phi}_{1,i}, \dots, \tilde{\phi}_{N,i}\}. \quad (48)$$

$$\mathcal{A}' \triangleq (1-\gamma)I_M + \gamma \mathcal{A}, \quad (49)$$

By subtracting \mathbf{w}_k^o from both sides of (28b) and by adding and subtracting \mathbf{w}_k^o to the difference $\psi_{k,i} - \phi_{k,i-1}$, we can write:

$$\tilde{\phi}_{k,i} = \tilde{\phi}_{k,i-1} - \mathcal{Q}_k(\tilde{\phi}_{k,i-1} - \tilde{\psi}_{k,i}). \quad (50)$$

Let

$$\chi_{k,i} \triangleq \tilde{\phi}_{k,i-1} - \tilde{\psi}_{k,i}. \quad (51)$$

By introducing the quantization error vector:

$$\mathbf{z}_{k,i} \triangleq \chi_{k,i} - \mathcal{Q}_k(\chi_{k,i}), \quad (52)$$

we can write:

$$\tilde{\phi}_{k,i} = \tilde{\psi}_{k,i} + \mathbf{z}_{k,i}. \quad (53)$$

By combining (45), (47), and (53), we conclude that the network error vector $\tilde{\phi}_i$ in (48) evolves according to the following dynamics:

$$\tilde{\phi}_i = \mathcal{B}_{i-1} \tilde{\phi}_{i-1} - \mu \mathbf{s}_i + \mu \mathbf{b} + \mathbf{z}_i \quad (54)$$

where

$$\mathcal{B}_{i-1} \triangleq (I_M - \mu \mathbf{H}_{i-1}) \mathcal{A}', \quad (55)$$

$$\mathbf{z}_i \triangleq \text{col}\{\mathbf{z}_{k,i}\}_{k=1}^N. \quad (56)$$

In Sec. IV-C₂ we will first establish the boundedness of $\limsup_{i \rightarrow \infty} \mathbb{E}\|\tilde{\phi}_i\|^2$ and then we will use relation (47) to deduce boundedness of $\limsup_{i \rightarrow \infty} \mathbb{E}\|\tilde{\mathbf{w}}_i\|^2$. The analysis of recursion (54) is facilitated by transforming it to a convenient basis using the Jordan canonical decomposition of the matrix \mathcal{A}' defined in (49). Now, to exploit the eigen-structure of \mathcal{A}' , we first recall that a matrix \mathcal{A} satisfying the conditions in (2) (for a full-column rank semi-unitary matrix \mathcal{U}) has a Jordan decomposition of the form $\mathcal{A} = \mathcal{V}_\epsilon \Lambda_\epsilon \mathcal{V}_\epsilon^{-1}$ with [24, Lemma 2]:

$$\mathcal{V}_\epsilon = \left[\begin{array}{c|c} \mathcal{U} & \mathcal{V}_{R,\epsilon} \end{array} \right], \quad \Lambda_\epsilon = \left[\begin{array}{c|c} I_P & 0 \\ \hline 0 & \mathcal{J}_\epsilon \end{array} \right], \quad \mathcal{V}_\epsilon^{-1} = \left[\begin{array}{c|c} \mathcal{U}^\top & \\ \hline \mathcal{V}_{L,\epsilon}^\top \end{array} \right], \quad (57)$$

where \mathcal{J}_ϵ is a Jordan matrix with eigenvalues (which may be complex but have magnitude less than one) on the diagonal and $\epsilon > 0$ on the super-diagonal [24, Lemma 2], [8, pp. 510]. The parameter ϵ is chosen small enough to ensure $\rho(\mathcal{J}_\epsilon) + \epsilon \in (0, 1)$ [24]. Consequently, the matrix \mathcal{A}' in (49) has a Jordan decomposition of the form $\mathcal{A}' = \mathcal{V}_\epsilon \Lambda'_\epsilon \mathcal{V}_\epsilon^{-1}$ where:

$$\Lambda'_\epsilon = \left[\begin{array}{c|c} I_P & 0 \\ \hline 0 & \mathcal{J}'_\epsilon \end{array} \right], \quad \text{with } \mathcal{J}'_\epsilon \triangleq (1-\gamma)I_{M-P} + \gamma \mathcal{J}_\epsilon. \quad (58)$$

By multiplying both sides of (54) from the left by $\mathcal{V}_\epsilon^{-1}$ in (57), we obtain the transformed iterates and variables:

$$\mathcal{V}_\epsilon^{-1} \tilde{\phi}_i = \left[\begin{array}{c|c} \mathcal{U}^\top \tilde{\phi}_i \\ \hline \mathcal{V}_{L,\epsilon}^\top \tilde{\phi}_i \end{array} \right] \triangleq \left[\begin{array}{c} \bar{\phi}_i \\ \tilde{\phi}_i \end{array} \right], \quad (59)$$

$$\mathcal{V}_\epsilon^{-1} \mathbf{s}_i = \left[\begin{array}{c|c} \mathcal{U}^\top \mathbf{s}_i \\ \hline \mathcal{V}_{L,\epsilon}^\top \mathbf{s}_i \end{array} \right] \triangleq \left[\begin{array}{c} \bar{\mathbf{s}}_i \\ \tilde{\mathbf{s}}_i \end{array} \right], \quad (60)$$

$$\mathcal{V}_\epsilon^{-1} \mathbf{b} = \left[\begin{array}{c|c} \mathcal{U}^\top \mathbf{b} \\ \hline \mathcal{V}_{L,\epsilon}^\top \mathbf{b} \end{array} \right] \triangleq \left[\begin{array}{c} 0 \\ \tilde{\mathbf{b}} \end{array} \right], \quad (61)$$

$$\mathcal{V}_\epsilon^{-1} \mathbf{z}_i = \left[\begin{array}{c|c} \mathcal{U}^\top \mathbf{z}_i \\ \hline \mathcal{V}_{L,\epsilon}^\top \mathbf{z}_i \end{array} \right] \triangleq \left[\begin{array}{c} \bar{\mathbf{z}}_i \\ \tilde{\mathbf{z}}_i \end{array} \right], \quad (62)$$

where in (61) we used the fact that $\mathcal{U}^\top b = 0$ as shown in [24, Sec. III-B]. In particular, the transformed components $\bar{\phi}_i$ and $\check{\phi}_i$ evolve according to the recursions:

$$\bar{\phi}_i = (I_P - \mu \mathcal{D}_{11,i-1}) \bar{\phi}_{i-1} - \mu \mathcal{D}_{12,i-1} \check{\phi}_{i-1} + \bar{z}_i - \mu \bar{s}_i \quad (63)$$

$$\check{\phi}_i = (\mathcal{J}'_\epsilon - \mu \mathcal{D}_{22,i-1}) \check{\phi}_{i-1} - \mu \mathcal{D}_{21,i-1} \bar{\phi}_{i-1} + \check{z}_i + \mu \check{b} - \mu \check{s}_i \quad (64)$$

where

$$\mathcal{D}_{11,i-1} \triangleq \mathcal{U}^\top \mathcal{H}_{i-1} \mathcal{U}, \quad (65)$$

$$\mathcal{D}_{12,i-1} \triangleq \mathcal{U}^\top \mathcal{H}_{i-1} \mathcal{V}_{R,\epsilon} \mathcal{J}'_\epsilon, \quad (66)$$

$$\mathcal{D}_{21,i-1} \triangleq \mathcal{V}_{L,\epsilon}^\top \mathcal{H}_{i-1} \mathcal{U}, \quad (67)$$

$$\mathcal{D}_{22,i-1} \triangleq \mathcal{V}_{L,\epsilon}^\top \mathcal{H}_{i-1} \mathcal{V}_{R,\epsilon} \mathcal{J}'_\epsilon. \quad (68)$$

In the following, we shall establish the mean-square-error stability of algorithm (28). The analysis will reveal the influence of the step-size μ , the mixing parameter γ , and the quantization noise (through $\{\beta_{q,k}^2, \sigma_{q,k}^2\}$) on the network mean-square-error stability and performance, and will provide insights into the design of effective quantizers for decentralized learning under subspace constraints.

C. Mean-square-error stability

Theorem 1. (Mean-square-error stability). *Consider a network of N agents running the quantized decentralized strategy (28) under Assumptions 1, 2, and 3, with a matrix \mathcal{A} satisfying (2) for an $M \times P$ (with $P \ll M$) full-column rank semi-unitary matrix \mathcal{U} . Let $\gamma \in (0, 1]$ be such that:*

$$0 < \gamma < \min \left\{ 1, \frac{1 - (\rho(\mathcal{J}_\epsilon) + \epsilon)}{4v_1^2 v_2^2 \beta_{q,\max}^2 (\rho(I - \mathcal{J}_\epsilon) + \epsilon)^2} \right\}. \quad (69)$$

where $v_1 = \|\mathcal{V}_\epsilon^{-1}\|$, $v_2 = \|\mathcal{V}_\epsilon\|$, and $\beta_{q,\max}^2 \triangleq \max_{1 \leq k \leq N} \{\beta_{q,k}^2\}$. Then, the network is mean-square-error stable for sufficiently small step-size μ , namely, it holds that:

$$\limsup_{i \rightarrow \infty} \mathbb{E} \|w_k^o - w_{k,i}\|^2 = \delta O(\mu) + \bar{\sigma}_q^2 O(\mu^{-1}), \quad (70)$$

for $k = 1, \dots, N$ and where δ is a constant given by³:

$$\delta \triangleq 2v_1^2 \bar{\sigma}_s^2 + \frac{3\|\check{b}\|^2}{1 - \|\mathcal{J}'_\epsilon\|} + \beta_{q,\max}^2 v_1^2 v_2^2 (2v_1^2 \bar{\sigma}_s^2 + 12\|\check{b}\|^2), \quad (71)$$

$\bar{\sigma}_s^2 = \sum_{k=1}^N \sigma_{s,k}^2$, $\check{b} = O(1)$ is given by (61), and $\bar{\sigma}_q^2 = \sum_{k=1}^N \sigma_{q,k}^2$. Convergence to the steady-state value (70) is linear at a rate given by:

$$\rho(\Gamma) \leq \max \left\{ 1 - \mu \sigma_{11} + O(\mu^2), \right. \\ \left. \|\mathcal{J}'_\epsilon\| + \kappa \|I - \mathcal{J}'_\epsilon\|^2 + O(\mu) \right\} < 1, \quad (72)$$

for some positive constants σ_{11} and κ .

Proof. See Appendix A. \square

While expression (70) in Theorem 1 reveals the influence of the step-size μ , the quantization noise (captured by $\{\bar{\sigma}_q^2, \beta_{q,\max}^2\}$), and the gradient noise (captured by $\bar{\sigma}_s^2$) on the steady-state mean-square error, expression (69) reveals the

³Using relation (91) in Appendix A, and since $\gamma > 0$ and $\rho(\mathcal{J}_\epsilon) + \epsilon \in (0, 1)$, it can be verified that $1 - \|\mathcal{J}'_\epsilon\| \geq \gamma(1 - \rho(\mathcal{J}_\epsilon) - \epsilon) > 0$.

influence of the *relative quantization noise* term (captured by $\beta_{q,\max}^2$) on the network stability. One main conclusion stemming from expression (70) in Theorem 1 is that the mean-square-error contains an $O(\mu)$ term, which is classically encountered in the unquantized case, plus an $O(\mu^{-1})$ term, which can be problematic in the small step-size regime (if the quantity $\bar{\sigma}_q^2 = \sum_{k=1}^N \sigma_{q,k}^2$ is not coupled with μ). Since $\bar{\sigma}_q^2$ depends on the quantizers' absolute noise components $\{\sigma_{q,k}^2\}$, this issue can be addressed by ensuring that $\sigma_{q,k}^2 \propto \mu^2$. In this way, we would recover the small estimation error result $\limsup_{i \rightarrow \infty} \mathbb{E} \|w_k^o - w_{k,i}\|^2 = O(\mu)$ observed in the unquantized case.

However, this setup requires a careful inspection since small values of $\sigma_{q,k}^2$ imply small quantization errors, which might in principle require large bit rates. Consequently, in the small step-size regime ($\mu \rightarrow 0$), the bit rate might increase without bound when $\sigma_{q,k}^2 \propto \mu^2$. Our goal then becomes to find a quantization scheme that achieves the classical small estimation error result $\limsup_{i \rightarrow \infty} \mathbb{E} \|w_k^o - w_{k,i}\|^2 = O(\mu)$ while guaranteeing that the bit rate stays bounded as $\mu \rightarrow 0$. In the next theorem, we will be able to show that the variable-rate scheme illustrated in Sec. II-A achieves both objectives. This theoretical finding will be further illustrated in the simulation section VI-A.

D. Bit rate stability

Before establishing the bit rate stability, we recall that, at each iteration i and agent k , the quantizer input is given by $\psi_{k,i} - \phi_{k,i-1}$, which is equal to the vector $\chi_{k,i}$ in (51). Consequently, from (27), the bit rate at agent k and iteration i is given by:

$$r_{k,i} = \log_2(3) \sum_{j=1}^{M_k} (1 + \mathbb{E} [\lceil \log_2(|n(\chi_{k,i})_j| + 1) \rceil]), \quad (73)$$

where $[\chi_{k,i}]_j$ denotes the j -th entry of the vector $\chi_{k,i}$.

Theorem 2. (Bit rate stability). *Assume that each agent k employs the randomized quantizer with the non-linearities given by (18), and with the parameters ω_k and η_k scaling according to:*

$$\omega_k = c, \quad \eta_k \propto \mu, \quad (74)$$

where c is a constant independent of μ , and where the symbol \propto hides a proportionality constant independent of μ . First, under conditions (74), we have:

$$\sigma_{q,k}^2 \propto \mu^2. \quad (75)$$

Second, in steady-state, the average number of bits at agent k stays bounded as $\mu \rightarrow 0$, namely,

$$\limsup_{i \rightarrow \infty} r_{k,i} = O(1). \quad (76)$$

Proof. See Appendix B. \square

To be concrete, we deemed it useful to focus in Theorem 2 on the logarithmic companding rule (18). However, it should be noted that the proof in Appendix B can be extended to handle more general non-linear functions.

V. THEOREMS 1 AND 2: INSIGHTS AND OBSERVATIONS

The fundamental conclusion emerging from Theorems 1 and 2 is that, when compared with (1), and when the quantizers are properly designed, the quantized approach (28) can guarantee a finite average bit rate whilst still ensuring small estimation errors on the order of μ . As detailed in the following, there are further insights that can be gained from the theorems.

First, note that in the *absence of quantization*, the analysis allows us to recover the mean-square-error stability result established in [24, Theorem 1] where it was shown that $\limsup_{i \rightarrow \infty} \mathbb{E} \|w_k^o - w_{k,i}\|^2 = O(\mu)$ for all k and for $\gamma = 1$. To see this, we just set to 0 the quantization noise parameters $\{\beta_{q,\max}^2, \bar{\sigma}_q^2\}$ in Theorem 1.

Second, from our analysis, it is possible to explain the bit rate stability result. Consider, for simplicity, the dithered uniform quantizer from Table I. Then, setting $\sigma_{q,k}^2 \propto \mu^2$ is equivalent to requiring that the quantization step-size Δ is proportional to μ . Our analysis reveals that, when $\Delta \propto \mu$, the *differential* input $\chi_{k,i} = \psi_{k,i} - \phi_{k,i-1}$ in (28b) is on the order of μ at the steady-state⁴. In this setting, i) the *effective* range of the inputs used in the *differential* quantizer scheme vanishes as $O(\mu)$ when $\mu \rightarrow 0$, and ii) the quantizer resolution $\Delta \propto \mu$ scales proportionally to the *effective* range of the quantizer input. Notably, Theorem 2 reveals that the variable-rate scheme is able to adapt the number of bits to this effective range, in such a way that the expected bit rate stays bounded as $\mu \rightarrow 0$.

Third, we conclude from Theorem 1 that, while a value $\gamma = 1$ could lead to network instability when $\beta_{q,\max}^2 \neq 0$, this conclusion does not hold when $\beta_{q,\max}^2 = 0$. In other words, the *absolute quantization noise component does not affect the network stability*, and the network remains stable⁵ for sufficiently small step-size μ when $\gamma = 1$.

Fourth, by noting that step (28b) can be written alternatively as:

$$\phi_{k,i} = \psi_{k,i} - z_{k,i}, \quad (77)$$

with $z_{k,i}$ given by (52), it can be observed that the impact of quantization in (28c) is similar to the impact of exchanging information over the communication links in the presence of additive noise processes since step (28c) can be written alternatively as (for $\gamma = 1$):

$$w_{k,i} = \sum_{\ell \in \mathcal{N}_k} A_{k\ell} (\psi_{\ell,i} - z_{\ell,i}). \quad (78)$$

Therefore, in the absence of the relative quantization noise term (i.e., when $\beta_{q,\max}^2 = 0$), we obtain a setting similar to the one previously considered in the context of single-task [28] and multitask [29] estimation over mean-square-error networks in the presence of noisy links. However, notice that the analysis in [28] and [29] is limited to LMS diffusion algorithms, does not exploit the randomized quantizers design, and does not investigate the average bit rate stability.

Fifth, observe that in the *absence of absolute quantization noise term*, i.e., when $\bar{\sigma}_q^2 = 0$, we obtain:

⁴This can be seen by taking the limits on both sides of relation (111) and by using (127).

⁵In fact, from (113)–(117) and (121), when $\beta_{q,\max}^2 = 0$ and $\gamma = 1$, one can observe that a sufficiently small step-size μ can ensure the stability of the matrix Γ .

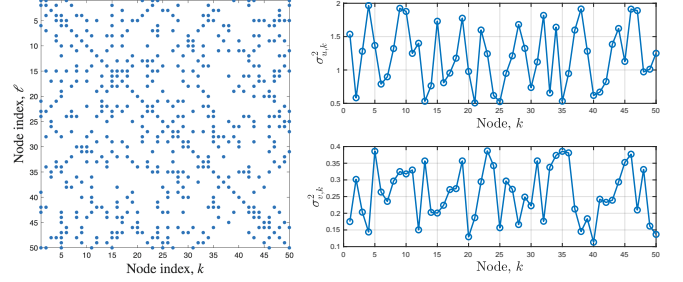


Fig. 4. Experimental setup. (Left) Link matrix ($N = 50$ nodes). (Right) Regressor and noise variances.

$$\limsup_{i \rightarrow \infty} \mathbb{E} \|w_k^o - w_{k,i}\|^2 = O(\mu). \quad (79)$$

Thus, the relative quantization noise component does not affect the $O(\mu)$ small estimation error result, but affects the network stability through condition (69). Regarding the bit rate, as it can be observed from Table I (column 5, rows 4–7), the reported quantization schemes are characterized by a fixed bit-budget that is independent of the quantizer input, but depends on the high-precision quantization parameter B_{HP} . The theoretical result (79) will be illustrated in Sec. VI-A by considering the QSGD quantizer of Table I (row 7).

Finally, and before discussing the stability condition (69), we note that the parameter γ allows to control $\rho(\mathcal{A}' - \mathcal{P}_u)$ (i.e., the spectral radius of the matrix $\mathcal{A}' - \mathcal{P}_u$), and consequently, the speed of convergence of $(\mathcal{A}')^i$ to \mathcal{P}_u [24, Lemma 2]. To see this, note that (by using the Jordan decomposition of the matrix \mathcal{A}'):

$$\mathcal{A}' - \mathcal{P}_u = \mathcal{V}_\epsilon \begin{bmatrix} 0 & 0 \\ 0 & \mathcal{J}'_\epsilon \end{bmatrix} \mathcal{V}_\epsilon^{-1}, \quad (80)$$

from which we obtain $\rho(\mathcal{A}' - \mathcal{P}_u) = \rho(\mathcal{J}'_\epsilon) \stackrel{(90)}{\leq} (1-\gamma) + \gamma \rho(\mathcal{J}_\epsilon)$ where $\rho(\mathcal{J}_\epsilon) \in (0, 1)$ and $\gamma \in (0, 1]$. Thus, the larger the *mixing* parameter γ is, the smaller $\rho(\mathcal{A}' - \mathcal{P}_u)$ tends to be, and consequently, the faster the convergence of $(\mathcal{A}')^i$ to \mathcal{P}_u will be. Now, returning to the stability condition (69), besides requiring $\gamma \in (0, 1]$, the mixing parameter γ must be chosen smaller than a value that is inversely proportional to $\beta_{q,\max}^2$. Thus, the larger $\beta_{q,\max}^2$ is, the tighter the upper bound in (69) is, and the smaller γ should be. If we consider for instance the QSGD quantizer from Table I, it can be observed that the smaller the number of levels s is, i.e., the smaller the number of used bits is, the larger $\beta_{q,k}^2$ is. Consequently, the larger $\beta_{q,k}^2$ is, the farther the agents (seeking to converge to $\text{Range}(\mathcal{U})$) can get from the subspace due to the quantization noise, and thus, the mixing parameter γ must be chosen small enough to ensure that the quantization noise perturbations will not affect the network stability.

VI. SIMULATION RESULTS

We apply strategy (28) to a network of $N = 50$ nodes with the link matrix shown in Fig. 4 (left). Each agent is subjected to streaming data $\{\mathbf{d}_k(i), \mathbf{u}_{k,i}\}$ assumed to satisfy a linear regression model of the form [8]:

$$\mathbf{d}_k(i) = \mathbf{u}_{k,i}^\top \mathbf{w}_k^* + \mathbf{v}_k(i), \quad (81)$$

for some unknown $L \times 1$ vector w_k^* with $v_k(i)$ denoting a zero-mean measurement noise and $L = 5$. A mean-square-error cost of the form $J_k(w_k) = \frac{1}{2}\mathbb{E}|\mathbf{d}_k(i) - \mathbf{u}_{k,i}^\top w_k|^2$ is associated with each agent k . The regressor and noise processes $\{\mathbf{u}_{k,i}, \mathbf{v}_k(i)\}$ are assumed to be zero-mean Gaussian with: i) $\mathbb{E}\mathbf{u}_{k,i}\mathbf{u}_{\ell,i}^\top = R_{u,k} = \sigma_{u,k}^2 I_L$ if $k = \ell$ and zero otherwise; ii) $\mathbb{E}\mathbf{v}_k(i)\mathbf{v}_\ell(i) = \sigma_{v,k}^2$ if $k = \ell$ and zero otherwise; and iii) $\mathbf{u}_{k,i}$ and $\mathbf{v}_k(i)$ are independent of each other. The variances $\sigma_{u,k}^2$ and $\sigma_{v,k}^2$ are illustrated in Fig. 4 (right). The signal $w^* = \text{col}\{w_1^*, \dots, w_N^*\}$ is generated by smoothing a signal w_o , which is randomly generated from the Gaussian distribution $\mathcal{N}(0.4 \times \mathbf{1}_{NL}, I_{NL})$, by a graph diffusion kernel with $\tau = 3$ – see [25, Sec. IV] for more details on the smoothing process. The matrix \mathcal{U} is generated according to $\mathcal{U} = U \otimes I_L$ where $U = [u_1 \ u_2]$, and u_1 and u_2 are the first two eigenvectors of the graph Laplacian \mathcal{L} . The Laplacian matrix is generated according to $\mathcal{L} = \text{diag}\{C\mathbf{1}_N\} - C$ where C is the $N \times N$ weighted adjacency matrix chosen such that the (k, ℓ) -th entry $[C]_{k\ell} = 0.1$ if $\ell \in \mathcal{N}_k$ and 0 otherwise. The combination matrix \mathcal{A} satisfying the conditions in (2) and having the same structure as the graph is found by following the same approach as in [45].

A. Effect of the small step-size parameter

In Fig. 5 (left), we report the network mean-square-deviation (MSD) learning curves:

$$\text{MSD}(i) = \frac{1}{N} \sum_{k=1}^N \mathbb{E} \|\mathbf{w}_k^o - \mathbf{w}_{k,i}\|^2 \quad (82)$$

for 3 different values of the step-size μ . The results were averaged over 100 Monte-Carlo runs. We used the probabilistic ANQ quantizer of Example 2 with $\omega_k = 0.25$ and $\eta_k = \frac{\mu}{\sqrt{2L}}$. This choice ensures that the quantizer settings of Theorem 2 are satisfied. We set $\gamma = 0.88$. We observe that, in steady-state, the network MSD increases by approximately 3dB when μ goes from μ_0 to $2\mu_0$. This means that the performance is on the order of μ , as expected from the discussion in Sec. V since in the simulations the absolute noise component is such that $\sigma_{q,k}^2 = \mu^2$. In Fig. 5 (middle), we report the average number of bits per node, per component, computed according to:

$$R(i) = \frac{1}{N} \sum_{k=1}^N \frac{1}{L} r_{k,i}, \quad (83)$$

where $r_{k,i}$ is the bit rate given by (73), which is associated with the encoding of the difference vector $\chi_{k,i} = \psi_{k,i} - \phi_{k,i-1}$ transmitted by agent k at iteration i according to (28b). As it can be observed, and thanks to the variable-rate quantization, a finite average number of bits is guaranteed (approximately 2.9 bits/component/iteration are required on average in steady-state). Moreover, it can be observed that the average number of bits scales coherently with the distortion scaling law induced by μ . That is, a smaller MSD in steady-state would require using a larger number of bits, and vice versa.

For comparison purposes, we report in Fig. 5 (right) the MSD learning curves when the QSGD quantizer of Table I (row 7) is employed instead of the probabilistic ANQ. Apart from the quantizer scheme, the same settings as above were assumed. For the QSGD scheme, we set the number of quantization levels s to 2. As it can be observed, this choice

allows us to compare the average number of bits for the ANQ and QSGD quantizers for similar values of steady-state MSD. From Table I (row 7, column 5), the bit-budget required to encode a 5×1 vector using the QSGD scheme ($s = 2$) is given by $B_{\text{HP}} + 10$. Now, by replacing B_{HP} by 32 (since we are performing the experiments on MATLAB 2022a which uses 32 bits to represent a floating number in single-precision), we find that the QSGD quantizer requires, at each iteration i , an average number of bits per node, per component, equal to $\frac{42}{5} = 8.4$, which is almost three times higher than the one obtained in steady-state when the probabilistic ANQ is used. This is expected since the QSGD scheme requires encoding the norm of the input vector with very high precision.

Note that, since the results of Theorems 1 and 2 hold for any $M \times P$ full-column rank semi-unitary matrix \mathcal{U} , similar observations will hold true when applying the quantized approach (28) to solve general constrained optimization problems of the form (3). In the supplementary material, we illustrate this fact by considering a simulation similar to the one considered in the current section VI-A, but instead we consider solving consensus optimization (30) by choosing $\mathcal{U} = \frac{1}{\sqrt{N}}(\mathbf{1}_N \otimes I_L)$.

B. Rate-distortion curves

In order to examine the performance of the proposed learning strategy, it is necessary to consider both the *attained learning error* (MSD) and the associated *bit expense*. This gives rise to a *rate-distortion* curve, where the rate quantifies the bit budget and the MSD quantifies the distortion. In Fig. 6, we illustrate the *rate-distortion* curves for probabilistic uniform and ANQ (for different values of the parameter ω). We set $\mu = 0.003$ and $\gamma = 0.88$. For the probabilistic uniform quantizer, each point of the rate/distortion curve corresponds to one value of the quantization step Δ . In the example, we selected 200 values of Δ , uniformly sampled in the interval $[0, 0.1]$. For each value of Δ (i.e., each point of the curve), the resulting MSD (distortion) and average number of bits/node/component (rate) were obtained by averaging the instantaneous mean-square-deviation $\text{MSD}(i)$ in (82) and averaging the number of bits $R(i)$ in (83) over 500 samples after convergence of the algorithm (the expectations in (82) and (73) are estimated empirically over 50 Monte Carlo runs). For the probabilistic ANQ scheme, we draw three curves corresponding to different values of the parameter ω . For each curve, the different rate/distortion pairs correspond to 200 values of the parameter η , uniformly sampled in the interval $[0, 0.5]$. The trade-off between rate and distortion can be observed from Fig. 6, namely, as the rate decreases, the distortion increases, and vice versa. For the ANQ, we further observe that the curves move away from the uniform case as ω increases, indicating a superior performance of the uniform probabilistic quantizer. In other words, the simulations show that in the considered example no advantage is obtained by employing non-uniform quantization. One useful interpretation for this behavior is as follows. In the theory of quantization, non-uniform quantizers are typically employed in a fixed-rate context, where the number of bits is determined in advance and does not depend on the input value. Under this setting, allocating high-resolution quantization intervals where the random variable is more likely to be observed provides

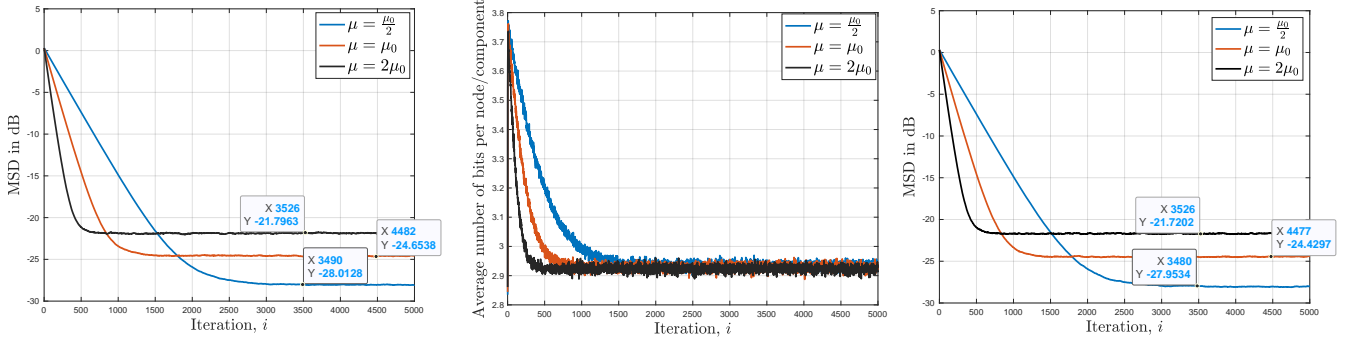


Fig. 5. Network performance w.r.t. w^o in (3) for three different values of the step-size ($\mu_0 = 0.003$). In the left and middle plots, the probabilistic ANQ quantizer of Example 2 is employed. (Left) Evolution of the MSD learning curves. (Middle) Evolution of the average number of bits per node, per component, when the variable-rate scheme described in Sec. II-A is used to encode the difference $\chi_{k,i} = \psi_{k,i} - \phi_{k,i-1}$ in (28b). (Right) Evolution of the MSD learning curves when the QSGD quantizer (with $s = 2$) of Table I is used.




Fig. 6. (Left) Rate-distortion curves for the probabilistic uniform and ANQ.

some advantage in terms of distortion, given the available *fixed* rate. However, in our case, we are considering a variable-rate quantizer that changes the number of bits depending on the input value. For example, when the quantizer input x belongs to a narrower quantization interval (i.e., smaller distortion), the variable-rate scheme allocates a higher number of bits. This somehow nullifies the distortion gain since we are allocating more bits. Therefore, allocating non-uniform intervals and then adapting the *variable* rate does not seem rewarding when compared to a uniform quantizer. This bears some similarity to what happens in the theory of quantization when one employs a fixed-rate quantizer followed by an entropy encoder. In this case, the possible advantages of non-uniform quantizers are nullified by the entropy encoder, and it is well known that (in the high resolution regime) the uniform quantizer is the best choice [46].

C. Tracking ability

To illustrate the tracking ability of the quantized approach (28), we modify w_k^* in (81) at time instants 1000 and 2000 by modifying w_o such that w_o is randomly generated from $\mathcal{N}(1.4 \times \mathbf{1}_{NL}, I_{NL})$ at $i = 1000$ and from $\mathcal{N}(2.4 \times \mathbf{1}_{NL}, I_{NL})$ at $i = 2000$. Apart from varying w_o (and, consequently, w^o in (3)) and from fixing the step-size μ to 0.006, the same settings as in Sec. VI-A were assumed. The resulting MSD learning curves and average number of bits/node/component curves are reported in Fig. 7 (left) and (right), respectively, for both the probabilistic ANQ and the QSGD quantizers. As it can be observed, approach (28) is able to track changes in the solution of the constrained optimization problem (3) despite quantization.




Fig. 7. Network performance in non-stationary environment (i.e., the solution w^o in (3) is changing over time due to changes in the individual costs $\{J_k(w_k)\}$) when the probabilistic ANQ quantizer of Example 2 and the QSGD quantizer of Table I are used. (Left) Evolution of the MSD learning curves. (Right) Evolution of the average number of bits per node, per component. The variable-rate scheme described in Sec. II-A is used in the probabilistic ANQ case.

VII. CONCLUSION

In this work, we considered inference problems over networks where agents have individual parameter vectors to estimate subject to subspace constraints that require the parameters across the network to lie in low-dimensional subspaces. The constrained subspace problem includes standard consensus optimization as a special case, and allows for more general task relatedness models such as multitask smoothness. To alleviate the communication bottleneck resulting from the exchange of the intermediate estimates among agents over many iterations, we proposed a new decentralized strategy relying on *differential randomized quantizers*. We studied its mean-square-error stability in the presence of both *relative* and *absolute* quantization noise terms. The analysis framework is general enough to cover many typical examples of probabilistic quantizers such as QSGD quantizer, gradient sparsifier, uniform or dithered quantizer, and non-uniform compandor. We showed that, for small step-sizes, and under some conditions on the quantizers (captured by the terms $\beta_{q,k}^2, \sigma_{q,k}^2$), on the network topology (captured by the eigendecomposition of the combination matrix \mathcal{A}), and on the mixing parameter γ , the decentralized quantized approach is able to converge in the mean-square-error sense within $O(\mu)$ from the solution of the constrained optimization problem, despite the gradient and quantization noises.

APPENDIX A

MEAN-SQUARE-ERROR ANALYSIS

We consider the transformed iterates $\bar{\phi}_i$ and $\check{\phi}_i$ in (63) and (64), respectively. Computing the second-order moment

of both sides of (63), we get:

$$\mathbb{E}\|\bar{\phi}_i\|^2 = \mathbb{E}\|(I_P - \mu\mathcal{D}_{11,i-1})\bar{\phi}_{i-1} - \mu\mathcal{D}_{12,i-1}\check{\phi}_{i-1}\|^2 + \mu^2\mathbb{E}\|\bar{s}_i\|^2 + \mathbb{E}\|\bar{z}_i\|^2, \quad (84)$$

where, from Assumptions 3 and 2 on the quantization and gradient noise processes, we used the fact that:

$$\mathbb{E}[\mathbf{x}_{i-1}^\top \bar{z}_i] = \mathbb{E}[\mathbb{E}[\mathbf{x}_{i-1}^\top \bar{z}_i | \mathbf{h}_i]] = \mathbb{E}[\mathbf{x}_{i-1}^\top \mathbb{E}[\bar{z}_i | \mathbf{h}_i]] = 0 \quad (85)$$

$$\begin{aligned} \mathbb{E}[\mathbf{x}_{i-1}^\top \bar{s}_i] &= \mathbb{E}\left[\mathbb{E}\left[\mathbf{x}_{i-1}^\top \bar{s}_i \middle| \{\phi_{\ell,i-1}\}_{\ell=1}^N\right]\right] \\ &= \mathbb{E}\left[\mathbf{x}_{i-1}^\top \mathbb{E}[\bar{s}_i | \{\phi_{\ell,i-1}\}_{\ell=1}^N]\right] = 0 \end{aligned} \quad (86)$$

$$\mathbb{E}[\bar{s}_i^\top \bar{z}_i] = \mathbb{E}[\mathbb{E}[\bar{s}_i^\top \bar{z}_i | \mathbf{h}_i]] = \mathbb{E}[\bar{s}_i^\top \mathbb{E}[\bar{z}_i | \mathbf{h}_i]] = 0 \quad (87)$$

with $\mathbf{x}_{i-1} = (I_P - \mu\mathcal{D}_{11,i-1})\bar{\phi}_{i-1} - \mu\mathcal{D}_{12,i-1}\check{\phi}_{i-1}$. Using similar arguments, we can also show that:

$$\begin{aligned} \mathbb{E}\|\check{\phi}_i\|^2 &= \mathbb{E}\|(\mathcal{J}'_\epsilon - \mu\mathcal{D}_{22,i-1})\check{\phi}_{i-1} - \mu\mathcal{D}_{21,i-1}\bar{\phi}_{i-1} + \mu\check{b}\|^2 \\ &\quad + \mu^2\mathbb{E}\|\check{s}_i\|^2 + \mathbb{E}\|\check{z}_i\|^2. \end{aligned} \quad (88)$$

Before proceeding, we show that, for small enough ϵ , the 2-induced matrix norm of \mathcal{J}'_ϵ in (58) satisfies $\|\mathcal{J}'_\epsilon\| \in (0, 1)$. This property will be used in the subsequent analysis. To establish the property, and by following similar arguments as in [8, pp. 516–517], we can first show that the block diagonal matrix \mathcal{J}'_ϵ , which is given by (58), satisfies:

$$\|\mathcal{J}'_\epsilon\|^2 \leq (\rho(\mathcal{J}'_\epsilon) + \gamma\epsilon)^2. \quad (89)$$

From (58), we can also show that:

$$\rho(\mathcal{J}'_\epsilon) \leq (1 - \gamma) + \gamma\rho(\mathcal{J}_\epsilon). \quad (90)$$

Using the fact that $\rho(\mathcal{J}_\epsilon) \in (0, 1)$ from [24, Lemma 2] and the fact that $\gamma \in (0, 1]$, we obtain $\rho(\mathcal{J}'_\epsilon) \in (0, 1)$. Since $\rho(\mathcal{J}'_\epsilon) + \gamma\epsilon$ is non-negative, by replacing (90) into (89), we obtain:

$$\begin{aligned} \|\mathcal{J}'_\epsilon\| &\leq (1 - \gamma) + \gamma\rho(\mathcal{J}_\epsilon) + \gamma\epsilon \\ &= 1 - \gamma(1 - \rho(\mathcal{J}_\epsilon) - \epsilon). \end{aligned} \quad (91)$$

This identity will also be used in the subsequent analysis. Regarding the block diagonal matrix $I - \mathcal{J}'_\epsilon$, which will appear in the following analysis, we can establish (by following similar arguments as in [8, pp. 516–517]):

$$\|I - \mathcal{J}'_\epsilon\|^2 \stackrel{(58)}{=} \|\gamma(I - \mathcal{J}_\epsilon)\|^2 = \gamma^2\|I - \mathcal{J}_\epsilon\|^2 \leq \gamma^2(\rho(I - \mathcal{J}_\epsilon) + \epsilon)^2, \quad (92)$$

where $\rho(I - \mathcal{J}_\epsilon) \in (0, 2)$ since $\rho(\mathcal{J}_\epsilon) \in (0, 1)$.

Now, returning to the error recursions (84) and (88), and using similar arguments as those used to establish inequalities (119) and (124) in [24, Appendix D], we can show that:

$$\begin{aligned} \mathbb{E}\|\bar{\phi}_i\|^2 &\leq (1 - \mu\sigma_{11})\mathbb{E}\|\bar{\phi}_{i-1}\|^2 + \frac{\mu\sigma_{12}^2}{\sigma_{11}}\mathbb{E}\|\check{\phi}_{i-1}\|^2 + \\ &\quad \mathbb{E}\|\bar{z}_i\|^2 + \mu^2\mathbb{E}\|\bar{s}_i\|^2, \end{aligned} \quad (93)$$

and

$$\begin{aligned} \mathbb{E}\|\check{\phi}_i\|^2 &\leq \left(\|\mathcal{J}'_\epsilon\| + \frac{3\mu^2\sigma_{22}^2}{1 - \|\mathcal{J}'_\epsilon\|}\right)\mathbb{E}\|\check{\phi}_{i-1}\|^2 + \\ &\quad \left(\frac{3\mu^2\sigma_{21}^2}{1 - \|\mathcal{J}'_\epsilon\|}\right)\mathbb{E}\|\bar{\phi}_{i-1}\|^2 + \left(\frac{3\mu^2}{1 - \|\mathcal{J}'_\epsilon\|}\right)\|\check{b}\|^2 + \mathbb{E}\|\check{z}_i\|^2 + \mu^2\mathbb{E}\|\check{s}_i\|^2 \end{aligned} \quad (94)$$

for some positive constant σ_{11} and non-negative constants σ_{12} , σ_{21} , and σ_{22} independent of μ . As it can be seen from (61), $\check{b} = \mathcal{V}_{L,\epsilon}^\top b$ depends on b in (41), which is defined in terms of the gradients $\{\nabla_{w_k} J_k(w_k^o)\}$. Since the costs $J_k(w_k)$ are twice differentiable, then $\|\check{b}\|^2$ is bounded and we obtain $\|\check{b}\|^2 = O(1)$.

For the gradient noise terms, by following similar arguments as in [8, Chapter 9], [24, Appendix D] and by using Assumption 2, we can show that:

$$\begin{aligned} \mathbb{E}\|\bar{s}_i\|^2 + \mathbb{E}\|\check{s}_i\|^2 &= \mathbb{E}\|\mathcal{V}_\epsilon^{-1} s_i\|^2 \leq v_1^2 \beta_{s,\max}^2 \mathbb{E}\|\tilde{w}_{i-1}\|^2 + v_1^2 \bar{\sigma}_s^2, \\ \text{where } v_1 &= \|\mathcal{V}_\epsilon^{-1}\|, \beta_{s,\max}^2 = \max_{1 \leq k \leq N} \beta_{s,k}^2, \text{ and } \bar{\sigma}_s^2 = \sum_{k=1}^N \sigma_{s,k}^2. \end{aligned} \quad (95)$$

Using expression (47) and the Jordan decomposition of the matrix \mathcal{A}' in (49), we obtain:

$$\begin{aligned} \mathbb{E}\|\bar{s}_i\|^2 + \mathbb{E}\|\check{s}_i\|^2 &\leq v_1^2 \beta_{s,\max}^2 \mathbb{E}\|\mathcal{A}' \tilde{\phi}_{i-1}\|^2 + v_1^2 \bar{\sigma}_s^2 \\ &\leq v_1^2 \beta_{s,\max}^2 \mathbb{E}\|\mathcal{V}_\epsilon \Lambda' (\mathcal{V}_\epsilon^{-1} \tilde{\phi}_{i-1})\|^2 + v_1^2 \bar{\sigma}_s^2 \\ &\leq v_1^2 \beta_{s,\max}^2 v_2^2 (\mathbb{E}\|\bar{\phi}_{i-1}\|^2 + \mathbb{E}\|\check{\phi}_{i-1}\|^2) + v_1^2 \bar{\sigma}_s^2, \end{aligned} \quad (96)$$

where $v_2 = \|\mathcal{V}_\epsilon\|$.

Using the bound (96) into (93) and (94), we obtain:

$$\begin{aligned} \mathbb{E}\|\bar{\phi}_i\|^2 &\leq (1 - \mu\sigma_{11} + \mu^2 v_1^2 \beta_{s,\max}^2 v_2^2) \mathbb{E}\|\bar{\phi}_{i-1}\|^2 + \\ &\quad \left(\frac{\mu\sigma_{12}^2}{\sigma_{11}} + \mu^2 v_1^2 \beta_{s,\max}^2 v_2^2\right) \mathbb{E}\|\check{\phi}_{i-1}\|^2 + \mathbb{E}\|\bar{z}_i\|^2 + \mu^2 v_1^2 \bar{\sigma}_s^2, \end{aligned} \quad (97)$$

and

$$\begin{aligned} \mathbb{E}\|\check{\phi}_i\|^2 &\leq \left(\|\mathcal{J}'_\epsilon\| + \frac{3\mu^2\sigma_{22}^2}{1 - \|\mathcal{J}'_\epsilon\|} + \mu^2 v_1^2 \beta_{s,\max}^2 v_2^2\right) \mathbb{E}\|\check{\phi}_{i-1}\|^2 + \\ &\quad \left(\frac{3\mu^2\sigma_{21}^2}{1 - \|\mathcal{J}'_\epsilon\|} + \mu^2 v_1^2 \beta_{s,\max}^2 v_2^2\right) \mathbb{E}\|\bar{\phi}_{i-1}\|^2 + \left(\frac{3\mu^2}{1 - \|\mathcal{J}'_\epsilon\|}\right) \|\check{b}\|^2 \\ &\quad + \mathbb{E}\|\check{z}_i\|^2 + \mu^2 v_1^2 \bar{\sigma}_s^2. \end{aligned} \quad (98)$$

Now, for the quantization noise, we have:

$$\mathbb{E}\|\bar{z}_i\|^2 + \mathbb{E}\|\check{z}_i\|^2 = \mathbb{E}\|\mathcal{V}_\epsilon^{-1} z_i\|^2 \stackrel{(56)}{\leq} v_1^2 \left(\sum_{k=1}^N \mathbb{E}\|z_{k,i}\|^2\right). \quad (99)$$

From (52) and Assumption 3, and since $\chi_{k,i} = \check{\phi}_{k,i-1} - \tilde{\psi}_{k,i}$, we can write:

$$\mathbb{E}\|z_{k,i}\|^2 \leq \beta_{q,k}^2 \mathbb{E}\|\check{\phi}_{k,i-1} - \tilde{\psi}_{k,i}\|^2 + \sigma_{q,k}^2, \quad (100)$$

and, therefore,

$$\mathbb{E}\|z_i\|^2 \leq \beta_{q,\max}^2 \mathbb{E}\|\check{\phi}_{i-1} - \tilde{\psi}_i\|^2 + \bar{\sigma}_q^2, \quad (101)$$

where $\beta_{q,\max}^2 = \max_{1 \leq k \leq N} \{\beta_{q,k}^2\}$ and $\bar{\sigma}_q^2 = \sum_{k=1}^N \sigma_{q,k}^2$. Since the analysis is facilitated by transforming the network vectors into the Jordan decomposition basis of the matrix \mathcal{A}' , we proceed by noting that the term $\mathbb{E}\|z_i\|^2$ can be bounded as follows:

$$\begin{aligned} \mathbb{E}\|z_i\|^2 &\stackrel{(101)}{\leq} \beta_{q,\max}^2 \mathbb{E}\|\mathcal{V}_\epsilon \mathcal{V}_\epsilon^{-1} (\check{\phi}_{i-1} - \tilde{\psi}_i)\|^2 + \bar{\sigma}_q^2 \\ &\leq \beta_{q,\max}^2 \mathbb{E}\|\mathcal{V}_\epsilon\|^2 \mathbb{E}\|\mathcal{V}_\epsilon^{-1} (\check{\phi}_{i-1} - \tilde{\psi}_i)\|^2 + \bar{\sigma}_q^2 \\ &\leq v_2^2 \beta_{q,\max}^2 [\mathbb{E}\|\bar{\chi}_i\|^2 + \mathbb{E}\|\check{\chi}_i\|^2] + \bar{\sigma}_q^2, \end{aligned} \quad (102)$$

where

$$\bar{\mathbf{x}}_i \triangleq \mathbf{U}^\top (\tilde{\boldsymbol{\phi}}_{i-1} - \tilde{\boldsymbol{\psi}}_i), \quad (103)$$

$$\check{\mathbf{x}}_i \triangleq \mathbf{V}_{L,\epsilon}^\top (\tilde{\boldsymbol{\phi}}_{i-1} - \tilde{\boldsymbol{\psi}}_i). \quad (104)$$

Therefore, by combining (99) and (102), we obtain:

$$\mathbb{E}\|\bar{\mathbf{z}}_i\|^2 + \mathbb{E}\|\check{\mathbf{z}}_i\|^2 \leq v_1^2 v_2^2 \beta_{q,\max}^2 [\mathbb{E}\|\bar{\mathbf{x}}_i\|^2 + \mathbb{E}\|\check{\mathbf{x}}_i\|^2] + v_1^2 \bar{\sigma}_q^2. \quad (105)$$

We focus now on deriving the network vector transformed recursions $\bar{\mathbf{x}}_i$ and $\check{\mathbf{x}}_i$. Subtracting $\tilde{\boldsymbol{\phi}}_{i-1}$ from both sides of (45) and using (47), we obtain:

$$\tilde{\boldsymbol{\phi}}_{i-1} - \tilde{\boldsymbol{\psi}}_i = (I_M - \mathcal{A}' + \mu \mathbf{H}_{i-1} \mathcal{A}') \tilde{\boldsymbol{\phi}}_{i-1} + \mu \mathbf{s}_i - \mu \mathbf{b}. \quad (106)$$

By multiplying both sides of (106) by $\mathbf{V}_{\epsilon}^{-1}$ and by using (59)–(61), (65)–(68), and the Jordan decomposition of the matrix \mathcal{A}' , we obtain:

$$\begin{bmatrix} \bar{\mathbf{x}}_i \\ \check{\mathbf{x}}_i \end{bmatrix} = \begin{bmatrix} \mu \mathbf{D}_{11,i-1} & \mu \mathbf{D}_{12,i-1} \\ \mu \mathbf{D}_{21,i-1} & I_{M-P} - \mathcal{J}'_{\epsilon} + \mu \mathbf{D}_{22,i-1} \end{bmatrix} \begin{bmatrix} \bar{\boldsymbol{\phi}}_{i-1} \\ \boldsymbol{\phi}_{i-1} \end{bmatrix} + \mu \begin{bmatrix} \bar{\mathbf{s}}_i \\ \check{\mathbf{s}}_i \end{bmatrix} - \mu \begin{bmatrix} \mathbf{0} \\ \check{\mathbf{b}} \end{bmatrix}. \quad (107)$$

Again, by using similar arguments as those used to establish inequalities (119) and (124) in [24, Appendix D] and by using Jensen's inequalities $\|x+y\|^2 \leq 2\|x\|^2 + 2\|y\|^2$ and $\|x+y+z\|^2 \leq 3\|x\|^2 + 3\|y\|^2 + 3\|z\|^2$, we can verify that:

$$\mathbb{E}\|\bar{\mathbf{x}}_i\|^2 \leq 2\mu^2 \sigma_{11}^2 \mathbb{E}\|\bar{\boldsymbol{\phi}}_{i-1}\|^2 + 2\mu^2 \sigma_{12}^2 \mathbb{E}\|\check{\boldsymbol{\phi}}_{i-1}\|^2 + \mu^2 \mathbb{E}\|\bar{\mathbf{s}}_i\|^2, \quad (108)$$

and

$$\begin{aligned} \mathbb{E}\|\check{\mathbf{x}}_i\|^2 &\leq 6\mu^2 \sigma_{21}^2 \mathbb{E}\|\bar{\boldsymbol{\phi}}_{i-1}\|^2 + \\ &2(\|I - \mathcal{J}'_{\epsilon}\|^2 + 3\mu^2 \sigma_{22}^2) \mathbb{E}\|\check{\boldsymbol{\phi}}_{i-1}\|^2 + 6\mu^2 \|\check{\mathbf{b}}\|^2 + \mu^2 \mathbb{E}\|\check{\mathbf{s}}_i\|^2. \end{aligned} \quad (109)$$

By combining expressions (108) and (109), we obtain:

$$\begin{aligned} \mathbb{E}\|\bar{\mathbf{x}}_i\|^2 + \mathbb{E}\|\check{\mathbf{x}}_i\|^2 &\leq (2\mu^2 \sigma_{11}^2 + 6\mu^2 \sigma_{21}^2) \mathbb{E}\|\bar{\boldsymbol{\phi}}_{i-1}\|^2 + \\ &(2\mu^2 \sigma_{12}^2 + 2\|I - \mathcal{J}'_{\epsilon}\|^2 + 6\mu^2 \sigma_{22}^2) \mathbb{E}\|\check{\boldsymbol{\phi}}_{i-1}\|^2 + \\ &6\mu^2 \|\check{\mathbf{b}}\|^2 + \mu^2 (\mathbb{E}\|\bar{\mathbf{s}}_i\|^2 + \mathbb{E}\|\check{\mathbf{s}}_i\|^2). \end{aligned} \quad (110)$$

Now, by using the bound (96) in (110), we obtain:

$$\begin{aligned} \mathbb{E}\|\bar{\mathbf{x}}_i\|^2 + \mathbb{E}\|\check{\mathbf{x}}_i\|^2 &\leq 6\mu^2 \|\check{\mathbf{b}}\|^2 + \mu^2 v_1^2 \bar{\sigma}_s^2 + \\ &(2\mu^2 \sigma_{11}^2 + 6\mu^2 \sigma_{21}^2 + \mu^2 v_1^2 \beta_{s,\max}^2 v_2^2) \mathbb{E}\|\bar{\boldsymbol{\phi}}_{i-1}\|^2 + \\ &(2\mu^2 \sigma_{12}^2 + 2\|I - \mathcal{J}'_{\epsilon}\|^2 + 6\mu^2 \sigma_{22}^2 + \mu^2 v_1^2 \beta_{s,\max}^2 v_2^2) \mathbb{E}\|\check{\boldsymbol{\phi}}_{i-1}\|^2. \end{aligned} \quad (111)$$

Using (111) and (105) in (97) and (98), we finally find that the variances of $\bar{\boldsymbol{\phi}}_i$ and $\boldsymbol{\phi}_i$ are coupled and recursively bounded as:

$$\begin{bmatrix} \mathbb{E}\|\bar{\boldsymbol{\phi}}_i\|^2 \\ \mathbb{E}\|\boldsymbol{\phi}_i\|^2 \end{bmatrix} \preceq \Gamma \begin{bmatrix} \mathbb{E}\|\bar{\boldsymbol{\phi}}_{i-1}\|^2 \\ \mathbb{E}\|\boldsymbol{\phi}_{i-1}\|^2 \end{bmatrix} + \begin{bmatrix} e + v_1^2 \bar{\sigma}_g^2 \\ f + v_1^2 \bar{\sigma}_q^2 \end{bmatrix}, \quad (112)$$

where Γ is the 2×2 matrix given by:

$$\Gamma = \begin{bmatrix} a & b \\ c & d \end{bmatrix}, \quad (113)$$

with

$$a = 1 - \mu \sigma_{11} + O(\mu^2) = 1 - \Theta(\mu), \quad (114)$$

$$b = 2v_1^2 v_2^2 \beta_{q,\max}^2 \|I - \mathcal{J}'_{\epsilon}\|^2 + O(\mu), \quad (115)$$

$$c = O(\mu^2), \quad (116)$$

$$d = \|\mathcal{J}'_{\epsilon}\| + 2v_1^2 v_2^2 \beta_{q,\max}^2 \|I - \mathcal{J}'_{\epsilon}\|^2 + O(\mu^2), \quad (117)$$

$$e \triangleq \mu^2 \left(6\|\check{\mathbf{b}}\|^2 v_1^2 v_2^2 \beta_{q,\max}^2 + v_1^2 \bar{\sigma}_s^2 v_1^2 v_2^2 \beta_{q,\max}^2 + v_1^2 \bar{\sigma}_s^2 \right), \quad (118)$$

$$f \triangleq \mu^2 \left(\frac{3\|\check{\mathbf{b}}\|^2}{1 - \|\mathcal{J}'_{\epsilon}\|} + 6v_1^2 v_2^2 \beta_{q,\max}^2 \|\check{\mathbf{b}}\|^2 + v_1^4 \bar{\sigma}_s^2 v_2^2 \beta_{q,\max}^2 + v_1^2 \bar{\sigma}_s^2 \right) \quad (119)$$

where the notation $g(\mu) = \Theta(\mu)$ signifies that there exists a positive constant c such that $\lim_{\mu \rightarrow 0} \frac{g(\mu)}{\mu} = c$.

If the matrix Γ is stable, i.e., $\rho(\Gamma) < 1$, then by iterating (112), we arrive at:

$$\limsup_{i \rightarrow \infty} \begin{bmatrix} \mathbb{E}\|\bar{\boldsymbol{\phi}}_i\|^2 \\ \mathbb{E}\|\boldsymbol{\phi}_i\|^2 \end{bmatrix} \preceq (I - \Gamma)^{-1} \begin{bmatrix} e + v_1^2 \bar{\sigma}_g^2 \\ f + v_1^2 \bar{\sigma}_q^2 \end{bmatrix}. \quad (120)$$

As we will see in the following, for some given data settings (captured by $\{\sigma_{11}^2, \sigma_{12}^2, \sigma_{21}^2, \sigma_{22}^2, \beta_{s,\max}^2, \bar{\sigma}_s^2\}$), small step-size parameter μ , network topology (captured by the matrix \mathcal{A} and its eigendecomposition $\{v_1^2, v_2^2, \mathcal{J}_{\epsilon}\}$), and quantizer settings (captured by $\{\beta_{q,\max}^2, \bar{\sigma}_q^2\}$), the stability of Γ can be controlled by the mixing parameter γ used in step (28c). As it can be observed from (58), (91), and (92), this parameter influences $\|\mathcal{J}'_{\epsilon}\|^2$ and $\|I - \mathcal{J}'_{\epsilon}\|^2$. Generally speaking, and since the spectral radius of a matrix is upper bounded by its 1-norm, the matrix Γ is stable if:

$$\rho(\Gamma) \leq \max\{|a| + |c|, |b| + |d|\} < 1. \quad (121)$$

Since $\sigma_{11} > 0$ and $\gamma \neq 0$, a sufficiently small μ can make $|a| + |c|$ strictly smaller than 1. For $|b| + |d|$, observe that if the mixing parameter γ is chosen such that:

$$\|\mathcal{J}'_{\epsilon}\| + 4v_1^2 v_2^2 \beta_{q,\max}^2 \|I - \mathcal{J}'_{\epsilon}\|^2 < 1, \quad (122)$$

then $|b| + |d|$ can be made strictly smaller than one for sufficiently small μ . It is therefore clear that the RHS of (121) can be made strictly smaller than one for sufficiently small μ and for a mixing parameter $\gamma \neq 0$ satisfying condition (122). Now, to derive condition (69) on the mixing parameter γ , we start by noting that the LHS of (122) can be upper bounded by:

$$\begin{aligned} \|\mathcal{J}'_{\epsilon}\| + 4v_1^2 v_2^2 \beta_{q,\max}^2 \|I - \mathcal{J}'_{\epsilon}\|^2 &\stackrel{(91),(92)}{\leq} 1 - \gamma(1 - \rho(\mathcal{J}_{\epsilon}) - \epsilon) + 4v_1^2 v_2^2 \beta_{q,\max}^2 \gamma^2 (\rho(I - \mathcal{J}_{\epsilon}) + \epsilon)^2. \end{aligned} \quad (123)$$

The upper bound in the above inequality is guaranteed to be strictly smaller than 1 if:

$$4v_1^2 v_2^2 \beta_{q,\max}^2 \gamma^2 (\rho(I - \mathcal{J}_{\epsilon}) + \epsilon)^2 - \gamma(1 - \rho(\mathcal{J}_{\epsilon}) - \epsilon) < 0. \quad (124)$$

Now, by using the above condition and the fact that γ must be in $(0, 1]$, we obtain condition (69).

Under condition (69), $\rho(\Gamma) < 1$, and consequently, the matrix Γ is stable. Moreover, it holds that:

$$(I - \Gamma) = \begin{bmatrix} \Theta(\mu) & O(1) \\ O(\mu^2) & \Theta(1) \end{bmatrix}, \quad (125)$$

and:

$$(I - \Gamma)^{-1} = \begin{bmatrix} \Theta(\mu^{-1}) & O(\mu^{-1}) \\ O(\mu) & \Theta(1) \end{bmatrix}. \quad (126)$$

Now, using (118), (119), and (126) into (120), we arrive at:

$$\limsup_{i \rightarrow \infty} \begin{bmatrix} \mathbb{E} \|\bar{\phi}_i\|^2 \\ \mathbb{E} \|\check{\phi}_i\|^2 \end{bmatrix} \preceq \begin{bmatrix} \beta_{q,\max}^2 \bar{\sigma}_s^2 O(\mu) + \beta_{q,\max}^2 \|\check{b}\|^2 O(\mu) + \bar{\sigma}_s^2 O(\mu) + \|\check{b}\|^2 O(\mu) + \bar{\sigma}_q^2 O(\mu^{-1}) \\ O(\mu^2) + \bar{\sigma}_q^2 O(1) \end{bmatrix} \quad (127)$$

By noting that:

$$\begin{aligned} \limsup_{i \rightarrow \infty} \mathbb{E} \|\tilde{\mathbf{w}}_i\|^2 &\stackrel{(47)}{=} \limsup_{i \rightarrow \infty} \mathbb{E} \|\mathcal{A}' \tilde{\phi}_i\|^2 \\ &= \limsup_{i \rightarrow \infty} \mathbb{E} \|\mathcal{V}_\epsilon \Lambda'_\epsilon \mathcal{V}_\epsilon^{-1} \tilde{\phi}_i\|^2 \\ &\leq \limsup_{i \rightarrow \infty} \|\mathcal{V}_\epsilon\|^2 \|\Lambda'_\epsilon\|^2 [\mathbb{E} \|\bar{\phi}_i\|^2 + \mathbb{E} \|\check{\phi}_i\|^2], \end{aligned} \quad (128)$$

and by using (127), we can finally conclude (70).

APPENDIX B BIT RATE STABILITY

Equation (75) follows straightforwardly from (20). Regarding (76), we first note from (18) that $g(t) \geq 0$ for non-negative arguments, and that:

$$|\mathbf{n}(t)| = \mathbf{n}(t) \leq \lfloor g(t) \rfloor + 1 \leq g(t) + 1, \quad \text{for } t \geq 0 \quad (129)$$

where the first inequality follows from (9) and (10), and the fact that the random index $\mathbf{n}(t)$ can be either equal to $\lfloor g(t) \rfloor$ or to the integer $\lfloor g(t) \rfloor + 1$. The second inequality follows from the definition of the floor operator. Likewise, for $t < 0$, we obtain:

$$\mathbf{n}(t) \geq \lfloor g(t) \rfloor = \lfloor -g(-t) \rfloor = -\lceil g(-t) \rceil \geq -g(-t) - 1. \quad (130)$$

Combining (129) and (130), we obtain:

$$|\mathbf{n}(t)| \leq g(|t|) + 1, \quad (131)$$

and, consequently, we can write:

$$\lceil \log_2(|\mathbf{n}(t)| + 1) \rceil \leq \log_2(|\mathbf{n}(t)| + 1) + 1 \leq \log_2(g(|t|) + 2) + 1. \quad (132)$$

Now, from (18), we can write:

$$\log_2(g(|t|) + 2) = \log_2 \left(\frac{\ln \left(1 + \frac{\omega}{\eta} \sqrt{t^2} \right)}{2 \ln(\omega + \sqrt{1 + \omega^2})} + 2 \right), \quad (133)$$

which can be verified to be a concave function w.r.t. to the argument t^2 . Therefore, if we consider a random input \mathbf{t} in (132), take the expectation, and use Jensen's inequality, we get:

$$\mathbb{E} [\lceil \log_2(|\mathbf{n}(\mathbf{t})| + 1) \rceil] \leq \log_2 \left(\frac{\ln \left(1 + \frac{\omega}{\eta} \sqrt{\mathbb{E}[\mathbf{t}^2]} \right)}{2 \ln(\omega + \sqrt{1 + \omega^2})} + 2 \right) + 1. \quad (134)$$

By choosing $\mathbf{t} = [\mathbf{x}_{k,i}]_j$ in (134) and by using the fact that $([\mathbf{x}_{k,i}]_j)^2 \leq \|\mathbf{x}_{k,i}\|^2$, we can upper bound the individual summand in (73) by the quantity:

$$2 + \log_2 \left(\frac{\ln \left(1 + \frac{\omega}{\eta} \sqrt{\mathbb{E} \|\mathbf{x}_{k,i}\|^2} \right)}{2 \ln(\omega + \sqrt{1 + \omega^2})} + 2 \right). \quad (135)$$

By taking the limit superior of (111) as $i \rightarrow \infty$ and by using (127), we obtain $l_k = \limsup_{i \rightarrow \infty} \mathbb{E} \|\mathbf{x}_{k,i}\|^2 = O(\mu^2)$. Consequently, we find that the limit superior of (135) as $i \rightarrow \infty$ is:

$$2 + \log_2 \left(\frac{\ln \left(1 + \frac{\omega}{\eta} \sqrt{l_k} \right)}{2 \ln(\omega + \sqrt{1 + \omega^2})} + 2 \right). \quad (136)$$

where we have used the fact that the function in (135) is continuous and increasing in the argument $\mathbb{E} \|\mathbf{x}_{k,i}\|^2$ and, hence, the limit superior is preserved. Under the conditions in (74), the above quantity is $O(1)$, and consequently, the rate in (73) satisfies (76).

REFERENCES

- [1] R. Nassif, S. Vlaski, M. Antonini, M. Carpentiero, V. Matta, and A. H. Sayed, "Finite bit quantization for decentralized learning under subspace constraints," in *Proc. Eur. Signal Process. Conf.*, Sep. 2022, pp. 1–5.
- [2] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. Y. Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. Int. Conf. Artif. Intell. Stat.*, Ft. Lauderdale, FL, USA, 2017, vol. 54, pp. 1273–1282.
- [3] T. Li, A. K. Sahu, A. S. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," *IEEE Signal Process. Mag.*, vol. 37, pp. 50–60, May 2020.
- [4] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, . . . , and S. Zhao, "Advances and open problems in federated learning," *Found. Trends Mach. Learn.*, vol. 14, no. 1–2, pp. 1–210, 2021.
- [5] D. Alistarh, D. Grubic, J. Li, R. Tomioka, and M. Vojnovic, "QSGD: Communication-efficient SGD via gradient quantization and encoding," in *Proc. Adv. Neural Inf. Process. Syst.*, Long Beach, CA, USA, 2017, pp. 1709–1720.
- [6] M. Aledhari, R. Razzak, R. M. Parizi, and F. Saeed, "Federated learning: A survey on enabling technologies, protocols, and applications," *IEEE Access*, vol. 8, pp. 140699–140725, Jul. 2020.
- [7] V. Smith, C.-K. Chiang, M. Sanjabi, and A. S. Talwalkar, "Federated multi-task learning," in *Proc. Adv. Neural Inf. Process. Syst.*, Long Beach, CA, USA, Dec. 2017, vol. 30.
- [8] A. H. Sayed, "Adaptation, learning, and optimization over networks," *Found. Trends Mach. Learn.*, vol. 7, no. 4–5, pp. 311–801, 2014.
- [9] A. H. Sayed, "Adaptive networks," *Proc. IEEE*, vol. 102, no. 4, pp. 460–497, 2014.
- [10] A. H. Sayed, S.-Y. Tu, J. Chen, X. Zhao, and Z. Towfic, "Diffusion strategies for adaptation and learning over networks," *IEEE Signal Process. Mag.*, vol. 30, no. 3, pp. 155–171, May 2013.
- [11] R. Nassif, S. Vlaski, C. Richard, J. Chen, and A. H. Sayed, "Multitask learning over graphs: An approach for distributed, streaming machine learning," *IEEE Signal Process. Mag.*, vol. 37, no. 3, pp. 14–25, 2020.
- [12] A. Koloskova, S. Stich, and M. Jaggi, "Decentralized stochastic optimization and gossip algorithms with compressed communication," in *Proc. Int. Conf. Mach. Learn.*, 2019, vol. 97, pp. 3478–3487.
- [13] A. Nedic and A. Ozdaglar, "Distributed subgradient methods for multi-agent optimization," *IEEE Trans. Automat. Contr.*, vol. 54, no. 1, pp. 48–61, Jan. 2009.
- [14] D. P. Bertsekas, "A new class of incremental gradient methods for least squares problems," *SIAM J. Optim.*, vol. 7, no. 4, pp. 913–926, 1997.
- [15] A. G. Dimakis, S. Kar, J. M. F. Moura, M. G. Rabbat, and A. Scaglione, "Gossip algorithms for distributed signal processing," *Proc. IEEE*, vol. 98, no. 11, pp. 1847–1864, 2010.
- [16] J. Plata-Chaves, A. Bertrand, M. Moonen, S. Theodoridis, and A. M. Zoubir, "Heterogeneous and multitask wireless sensor networks – Algorithms, applications, and challenges," *IEEE J. Sel. Top. Signal Process.*, vol. 11, no. 3, pp. 450–465, 2017.
- [17] J. F. C. Mota, J. M. F. Xavier, P. M. Q. Aguiar, and M. Püschel, "Distributed optimization with local domains: Applications in MPC and network flows," *IEEE Trans. Automat. Contr.*, vol. 60, no. 7, pp. 2004–2009, 2015.
- [18] R. Nassif, C. Richard, A. Ferrari, and A. H. Sayed, "Diffusion LMS for multitask problems with local linear equality constraints," *IEEE Trans. Signal Process.*, vol. 65, no. 19, pp. 4979–4993, 2017.

- [19] V. Kekatos and G. B. Giannakis, "Distributed robust power system state estimation," *IEEE Trans. Power Syst.*, vol. 28, no. 2, pp. 1617–1626, 2013.
- [20] S. A. Alghunaim and A. H. Sayed, "Distributed coupled multiagent stochastic optimization," *IEEE Trans. Automat. Contr.*, vol. 65, no. 1, pp. 175–190, 2020.
- [21] R. Nassif, S. Vlaski, C. Richard, and A. H. Sayed, "Learning over multitask graphs—Part I: Stability analysis," *IEEE Open Journal of Signal Processing*, vol. 1, pp. 28–45, 2020.
- [22] A. K. Sahu, D. Jakovetić, and S. Kar, "*CTRF*: A distributed random fields estimator," *IEEE Trans. Signal Process.*, vol. 66, no. 18, pp. 4980–4995, 2018.
- [23] J. Chen, C. Richard, and A. H. Sayed, "Multitask diffusion adaptation over networks," *IEEE Trans. Signal Process.*, vol. 62, no. 16, pp. 4129–4144, 2014.
- [24] R. Nassif, S. Vlaski, and A. H. Sayed, "Adaptation and learning over networks under subspace constraints—Part I: Stability analysis," *IEEE Trans. Signal Process.*, vol. 68, pp. 1346–1360, 2020.
- [25] R. Nassif, S. Vlaski, and A. H. Sayed, "Adaptation and learning over networks under subspace constraints—Part II: Performance analysis," *IEEE Trans. Signal Process.*, vol. 68, pp. 2948–2962, 2020.
- [26] P. Di Lorenzo, S. Barbarossa, and S. Sardellitti, "Distributed signal processing and optimization based on in-network subspace projections," *IEEE Trans. Signal Process.*, vol. 68, pp. 2061–2076, 2020.
- [27] A. Nedic, A. Olshevsky, A. Ozdaglar, and J. N. Tsitsiklis, "Distributed subgradient methods and quantization effects," in *Proc. IEEE Conf. Decis. Control*, Cancun, Mexico, Dec. 2008, pp. 4177–4184.
- [28] X. Zhao, S.-Y. Tu, and A. H. Sayed, "Diffusion adaptation over networks under imperfect information exchange and non-stationary data," *IEEE Trans. Signal Process.*, vol. 60, no. 7, pp. 3460–3475, 2012.
- [29] R. Nassif, C. Richard, J. Chen, A. Ferrari, and A. H. Sayed, "Diffusion LMS over multitask networks with noisy links," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Shanghai, China, 2016, pp. 4583–4587.
- [30] D. Thanou, E. Kokiopoulou, Y. Pu, and P. Frossard, "Distributed average consensus with quantization refinement," *IEEE Trans. Signal Process.*, vol. 61, no. 1, pp. 194–205, 2013.
- [31] N. Michelusi, G. Scutari, and C.-S. Lee, "Finite-bit quantization for distributed algorithms with linear convergence," *IEEE Trans. Inf. Theory*, vol. 68, no. 11, pp. 7254–7280, Nov. 2022.
- [32] M. Carpentiero, V. Matta, and A. H. Sayed, "Distributed adaptive learning under communication constraints," Available as arXiv:2112.02129, Dec. 2021.
- [33] M. Carpentiero, V. Matta, and A. H. Sayed, "Adaptive diffusion with compressed communication," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2022, pp. 1–5.
- [34] D. Kovalev, A. Koloskova, M. Jaggi, P. Richtarik, and S. Stich, "A linearly convergent algorithm for decentralized optimization: Sending less bits for free!," in *Proc. Int. Conf. Artif. Intell. Stat.*, Virtual, 2021, pp. 4087–4095.
- [35] N. Singh, D. Data, J. George, and S. Diggavi, "SPARQ-SGD: Event-triggered and compressed communication in decentralized optimization," *IEEE Trans. Automat. Contr.*, vol. 68, no. 2, pp. 721–736, 2023.
- [36] N. Singh, D. Data, J. George, and S. Diggavi, "SQuARM-SGD: Communication-efficient momentum SGD for decentralized optimization," in *Proc. IEEE Int. Symp. Inf. Theory*, Melbourne, Victoria, Australia, 2021, pp. 1212–1217.
- [37] H. Tang, S. Gan, C. Zhang, T. Zhang, and J. Liu, "Communication compression for decentralized training," in *Proc. Adv. Neural Inf. Process. Syst.*, Montréal, Canada, 2018, pp. 7663–7673.
- [38] H. Taheri, A. Mokhtari, H. Hassani, and R. Pedarsani, "Quantized decentralized stochastic learning over directed graphs," in *Proc. Int. Conf. Mach. Learn.*, Jul. 2020, vol. 119, pp. 9324–9333.
- [39] X. Zhang, J. Liu, Z. Zhu, and E. S. Bentley, "Compressed distributed gradient descent: Communication-efficient consensus over networks," in *Proc. IEEE Int. Conf. Commun.*, Paris, France, 2018, pp. 2431–2439.
- [40] A. Reisizadeh, A. Mokhtari, H. Hassani, and R. Pedarsani, "An exact quantized decentralized gradient descent algorithm," *IEEE Trans. Signal Process.*, vol. 67, no. 19, pp. 4934–4947, 2019.
- [41] A. Gersho and R. M. Gray, *Scalar Quantization I: Structure and Performance*, pp. 133–172, Springer US, Boston, MA, 1992.
- [42] T. C. Aysal, M. J. Coates, and M. G. Rabbat, "Distributed average consensus with dithered quantization," *IEEE Trans. Signal Process.*, vol. 56, no. 10, pp. 4905–4918, 2008.
- [43] L. M. Nguyen, P. H. Nguyen, M. Van Dijk, P. Richtarik, K. Scheinberg, and M. Takáč, "SGD and Hogwild! Convergence without the bounded gradients assumption," in *Proc. Int. Conf. Mach. Learn.*, Stockholm, SWE, 2018.
- [44] B. T. Polyak, *Introduction to Optimization*, Optimization Software, New York, 1987.
- [45] R. Nassif, S. Vlaski, and A. H. Sayed, "Distributed inference over networks under subspace constraints," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Brighton, UK, 2019, pp. 5232–5236.
- [46] H. Gish and J. Pierce, "Asymptotically efficient quantizing," *IEEE Trans. Inf. Theory*, vol. 14, no. 5, pp. 676–683, 1968.