

Finite-Bit Quantization for Distributed Algorithms With Linear Convergence

Nicolò Michelusi^{ID}, Senior Member, IEEE, Gesualdo Scutari^{ID}, Fellow, IEEE, and Chang-Shen Lee^{ID}

Abstract—This paper studies distributed algorithms for (strongly convex) composite optimization problems over mesh networks, subject to quantized communications. Instead of focusing on a specific algorithmic design, a black-box model is proposed, casting linearly convergent distributed algorithms in the form of fixed-point iterates. The algorithmic model is equipped with a novel random or deterministic *Biased Compression* (BC) rule on the quantizer design, and a new *Adaptive encoding Non-uniform Quantizer* (ANQ) coupled with a communication-efficient encoding scheme, which implements the BC-rule using a *finite* number of bits (below machine precision). This fills a gap existing in most state-of-the-art quantization schemes, such as those based on the popular compression rule, which rely on communication of some scalar signals with negligible quantization error (in practice quantized at the machine precision). A unified communication complexity analysis is developed for the black-box model, determining the average number of bits required to reach a solution of the optimization problem within a target accuracy. It is shown that the proposed BC-rule preserves *linear* convergence of the unquantized algorithms, and a trade-off between convergence rate and communication cost under ANQ-based quantization is characterized. Numerical results validate our theoretical findings and show that distributed algorithms equipped with the proposed ANQ have more favorable communication cost than algorithms using state-of-the-art quantization rules.

Index Terms—Distributed algorithms, linear convergence, quantization, optimization methods.

I. INTRODUCTION

WE STUDY distributed optimization over a network of m agents modeled as an undirected (connected) graph. We consider mesh networks, that is, arbitrary topologies with no central hub connected to all the other agents, where

Manuscript received 22 July 2021; revised 26 April 2022; accepted 1 May 2022. Date of publication 25 May 2022; date of current version 21 October 2022. The work of Nicolò Michelusi was supported in part by the National Science Foundation under Grant CNS-2129015. The work of Gesualdo Scutari and Chang-Shen Lee was supported in part by the National Science Foundation under Grant CIF 1719205, in part by the Army Research Office under Grant W911NF1810238, and in part by the Office of Naval Research under Grant N00014-21-1-2673. (Corresponding author: Gesualdo Scutari.)

Nicolò Michelusi is with the School of Electrical, Computer and Energy Engineering, Arizona State University, Tempe, AZ 85287 USA (e-mail: nicolo.michelusi@asu.edu).

Gesualdo Scutari is with the School of Industrial Engineering, Purdue University, West Lafayette, IN 47907 USA (e-mail: gscutari@purdue.edu).

Chang-Shen Lee is with Bloomberg, New York City, NY 10022 USA (e-mail: ben3003neb@gmail.com).

Communicated by R. Venkataraman, Associate Editor for Machine Learning and Statistics, Communications, Signal Processing and Source Coding.

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TIT.2022.3176253>.

Digital Object Identifier 10.1109/TIT.2022.3176253

each agent can communicate with its immediate neighbors (master/worker architectures can be treated as a special case). The m agents aim at solving cooperatively the optimization problem

$$\min_{\mathbf{x} \in \mathbb{R}^d} \underbrace{\frac{1}{m} \sum_{i=1}^m f_i(\mathbf{x}) + r(\mathbf{x})}_{=F(\mathbf{x})}, \quad (P)$$

where each f_i is the local cost function of agent i , assumed to be smooth, convex, and known only to the agent; $r : \mathbb{R}^d \rightarrow [-\infty, \infty]$ is a non-smooth, convex (extended-value) function known to all agents, which may be used to force shared constraints or some structure on the solution (e.g., sparsity); and the global loss $F : \mathbb{R}^d \rightarrow \mathbb{R}$ (or in some cases each local loss f_i) is assumed to be strongly convex on the domain of r . This setting is fairly general and finds applications in several areas, including network information processing, telecommunications, multi-agent control, and machine learning (e.g., see [1]–[3]).

Since the functions f_i can only be accessed locally and routing local data to other agents is infeasible or highly inefficient, solving (P) calls for the design of distributed algorithms that alternate between a local computation procedure at each agent's side and some rounds of communication among neighboring nodes. While most existing works focus on *ad-hoc* solution methods, here we consider a *general* distributed algorithmic framework, encompassing algorithms whose dynamics are modeled by the fixed-point iteration

$$\mathbf{z}^{k+1} = \tilde{\mathcal{A}}(\mathbf{z}^k), \quad (1)$$

where \mathbf{z}^k is the updating variable at iteration k and $\tilde{\mathcal{A}}$ is a mapping that embeds the local computation and communication steps, whose fixed point typically coincides with solutions of (P). This model encompasses several distributed algorithms over different network architectures, each one corresponding to a specific expression of \mathbf{z} and $\tilde{\mathcal{A}}$ —see Sec. II for some examples.

By assuming that F is strongly convex, we explicitly target distributed schemes in the form (1) that converge to solutions of (P) at *linear rate*. Furthermore, since the cost of communications is often the bottleneck for distributed computing when compared with local (possibly parallel) computations (e.g., [4], [5]), we achieve communication efficiency by embedding the iterates (1) with quantized communication protocols.

Quantizing communication steps of distributed optimization algorithms have received significant attention in recent

years—see Sec. I-B for a comprehensive overview of the state-of-the-art. Here we only point out that existing distributed algorithms over *mesh networks* are applicable only to *unconstrained* instances of (P) and *smooth* objective functions (i.e., $r \neq 0$). Furthermore, even in these special instances, such schemes rely on quantization rules that subsume some scalar signals to be encoded with negligible error—the typical example is the renowned *compression* protocol [6]. While in practice this is achieved by quantizing at the *machine precision* (e.g., 32 or 64 bit floating-point), on the theoretical side, existing convergence analyses become elusive. In such cases, the assessment of their convergence is left to simulations.

The goal of this paper is to design a black-box quantization mechanism for the class of distributed algorithms (1) applicable to Problem (P) over mesh networks that preserves their linear convergence while employing communications quantized with *finite-bit* (below machine precision). As discussed next, this is an open problem.

A. Summary of Main Contributions

In a nutshell, we summarize our major contributions as follows.

1) *A Black-Box Quantization Model for (1)*: We propose a novel black-box model that introduces quantization in the communication steps of linearly convergent distributed algorithms cast in the form (1). Our approach paves the way to a *unified* design of quantization rules and analysis of their impact on the convergence rate of a gamut of distributed algorithms. This constitutes a major departure from the majority of existing studies focusing on ad-hoc algorithms and quantization rules, which in fact are special instances of our framework. Furthermore, our model brings for the first time quantization to distributed algorithms applicable to composite optimization problems (i.e., (P) with $r \neq 0$). This is particularly relevant in machine learning applications, where empirical risk minimization problems call for (non-smooth) regularizations or constraints to control the complexity of the solution (e.g., to enforce sparsity or low-rank structure) and avoid overfitting.

2) *Preserving Linear Convergence of (1) Under Quantization*: To enable quantization below the machine precision, we provide a novel *biased compression* rule (the BC-rule) on the quantizer design equipping the proposed black-box model, which allows to preserve *linear* convergence of the distributed algorithms while using a *finite* number of bits (below machine precision) and *without* altering their original tuning. Our condition encompasses several deterministic and random quantization rules as special cases, new and old [7]–[23].

3) *A Novel Finite-Bit Quantizer*: To implement the BC-rule, we also propose a novel *finite-bit* quantizer (below machine precision) fulfilling the BC-rule along with a communication-efficient bit-encoding/decoding rule which enables transmissions on digital channels. The resulting *Adaptive encoding Non-uniform Quantizer* (ANQ) adapts the number of bits of the output (discrete representation) based upon the input signal. By doing so, it achieves a more communication-efficient

design than existing quantizers that adopt a fixed number of bits based on a predetermined fixed [7]–[10], [12], [13], [15], [16], [18], [19], [23] or shrinking range [11], [14], [17], [20]–[22] of the input signal, or that rely on transmissions of scalar values at machine precision [6], [24]–[39].

4) *Communication Complexity*: We derive a unified complexity analysis for *any* distributed algorithm belonging to our black-box model, solving (P) (possibly with $r \neq 0$) over *mesh networks*. Specifically, we prove that, under suitable conditions and proper tuning (see Theorem 13), an ε -solution of (P) (using a suitable optimality measure) is achieved by any of such distributed algorithms in

$$\mathcal{O}\left(\frac{1}{1-\lambda} \log(d/\varepsilon)\right) \text{ iterations,}$$

using

$$\mathcal{O}\left(d \log\left(1 + \frac{1}{1-\lambda}\right)\right) \text{ bits/agent/iteration,}$$

where $\lambda \in (0, 1)$ is the convergence rate of the distributed algorithm our quantization method is applied to, which depends on the condition number of the agents' function as well as network parameters. Table I customizes the above result to a variety of state-of-the-art distributed schemes subject to quantization, using the explicit expression of λ (see Table II and Corollaries 14–17). This permits for the first time to benchmark several distributed algorithms under quantization, and compare them with other state-of-the-art quantization schemes (Table I). Notice that the proposed methods compare favorably with existing ones and, remarkably, they achieve ε -accuracy with the same number of iterations (in a \mathcal{O} -sense) as their unquantized counterparts.

5) *Numerical Evaluations*: Finally, we extensively validate our theoretical findings on smooth and non-smooth regularized linear and logistic regression problems. Among others, our evaluations show that 1) linear convergence of all distributed algorithms is preserved under finite-bit quantization based upon the proposed BC-rule; 2) as predicted by our analysis, the rate approaches that of their unquantized counterpart (implemented at machine precision) when a sufficient number of bits is used; 3) the proposed ANQ quantization outperforms, in terms of both convergence rate and communication cost, existing quantization schemes operating below machine precision—i.e., Q-Dual [22] and Q-NEXT [21]—or at machine precision—e.g., those relying on the conventional compression rule, such as LEAD [25] and COLD [39].

B. Related Works

The literature on distributed algorithms is vast; here, we review relevant works employing some form of quantization with linear convergence guarantees [20]–[22], [24], [30], [33], [36], [38], categorized into those relying on machine precision quantization and those relying on limited precision quantization.

1) *Machine Precision Quantization Schemes* [24], [30], [33], [36], [38]: Distributed algorithms employing quantization in the agents' communications are proposed in [24], [30], [33], [36], [38] for special instances of (P) with $r \equiv 0$ (i.e.,

TABLE I
COMMUNICATION COMPLEXITY OF THE PROPOSED QUANTIZED DISTRIBUTED ALGORITHMS VS. THE STATE OF THE ART

Algorithm	Problem, network	[Bits/agent/iteration] \times [Iterations to ε -accuracy]
State-of-the-art Quantized Distributed Algorithms		
Q-GD over star [22]	(P) with $r \equiv 0$, star	$\mathcal{O}\left(d \log(1 + \kappa d) \times \kappa \log(d/\varepsilon)\right)$
Q-NEXT [20], [21]	(P) with $r \equiv 0$, mesh	N/A
Q-Primal-Dual [22]	(P) with $r \equiv 0$, mesh	N/A
Primal-Dual (rand- K compr., Option-D) [24]**	(P) with $r \equiv 0$, mesh	$\mathcal{O}\left(d(B + \log d) \times \frac{\kappa}{1-\rho_2} \log(d/\varepsilon)\right)$
Primal-Dual (Option-D) [24], LEAD [25] (both dit- K compr.)**	(P) with $r \equiv 0$, mesh	$\mathcal{O}\left((d + B) \times \frac{\kappa}{1-\rho_2} \log(d/\varepsilon)\right)$
ANQ-embedded Distributed Algorithms (this work)		
GD over star	(P) with $r \equiv 0$, star	$\mathcal{O}\left(d \log(1 + \kappa) \times \kappa \log(d/\varepsilon)\right)$
(Prox-)EXTRA, (Prox-)NIDS	(P), mesh	$\mathcal{O}\left(d \log\left(\max\left\{1 + \kappa, \frac{1}{1-\rho_2}\right\}\right) \times \max\left\{\kappa, \frac{1}{1-\rho_2}\right\} \log(d/\varepsilon)\right)$
NIDS	(P) with $r \equiv 0$, mesh	$\mathcal{O}\left(d \log\left(\max\left\{1 + \kappa, \frac{1}{1-\rho_2}\right\}\right) \times \max\left\{\kappa, \frac{1}{1-\rho_2}\right\} \log(d/\varepsilon)\right)$
(Prox-)NEXT, (Prox-)DIGing	(P), mesh	$\mathcal{O}\left(d \log\left(\max\left\{1 + \kappa, \frac{1}{(1-\rho_2)^2}\right\}\right) \times \max\left\{\kappa, \frac{1}{(1-\rho_2)^2}\right\} \log(d/\varepsilon)\right)$
Primal-Dual	(P) with $r \equiv 0$, mesh	$\mathcal{O}\left(d \log\left(1 + \frac{\kappa}{1-\rho_2}\right) \times \frac{\kappa}{1-\rho_2} \log(d/\varepsilon)\right)$

Number of bits/agent/iteration times the number of iterations to achieve ε accuracy for the proposed ANQ applied to state of the art distributed algorithms solving various instances of (P) over star or mesh networks; κ is the condition number of each f_i , d is the dimension of \mathbf{x} , ρ_2 is the second largest eigenvalue of the gossip matrix \mathbf{W} used in these algorithms (note that $\rho_2 = 0$ for star-networks or fully-connected graphs); see also Corollaries 14–17. **: the complexities of these schemes are adapted from [24]; ρ and ρ_∞ in $\mathcal{O}(1/(1-\rho_2))$; the compression rate is chosen as $\Theta(1)$ to match the iterations to ε -accuracy of the ANQ-Primal-Dual, hence $K = \Theta(d)$ for rand- K and $S = 2^{K-1} - 1 = \Theta(\sqrt{d})$ for dit- K ; B is the number of bits used to encode each scalar with *negligible loss in precision* (as required by the convergence theory therein).

smooth and unconstrained optimization). In these schemes, quantization is implemented by compressing the signal $\mathbf{x} \in \mathbb{R}^d$ through a (random or deterministic¹) *compression operator* $\mathbf{x} \mapsto \mathcal{Q}(\mathbf{x})$, that satisfies the *compression rule*

$$\sqrt{\mathbb{E}[\|\mathcal{Q}(\mathbf{x}) - \mathbf{x}\|_2^2]} \leq \omega \|\mathbf{x}\|_2, \quad \text{for some } \omega > 0. \quad (2)$$

This rule subsumes the transmission of some scalar signals with negligible quantization errors, e.g., the norm of \mathbf{x} , requiring thus in theory infinite machine precision (see Corollary 9 for a formal proof). While quantizing at the machine precision is a viable strategy in practice, convergence guarantees of distributed algorithms relying on (2) become elusive—they are established only under negligible quantization errors of the transmitted norm signal. By generalizing the conventional compression rule, the proposed BC-rule overcomes this theoretical limitation and permits to explicitly model the communication cost of quantized communications below machine precision. As we will demonstrate numerically in Sec. VI, our explicit formulation of the communication cost allows to define more communication efficient quantization schemes than state-of-the-art algorithms relying on machine precision.

2) *Limited Precision Quantization Schemes* [20]–[22]: While finite-rate quantization has been extensively studied for average consensus schemes (e.g., [8], [9], [11], [12],

[14], [17]–[19], [23]), their extension to optimization algorithms over mesh networks is less explored [20]–[22]. Specifically, in our prior work [20], we equipped the NEXT algorithm [40], [41] with a finite-bit deterministic quantization to solve (P) with $r \equiv 0$; to preserve linear convergence, the quantizer shrinks its input range linearly. An expression of the convergence rate of the scheme in [20] has been later determined in [21] along with its scaling properties with respect to problem, network, and quantization parameters.

The closest paper to our work is [22], where the authors proposed a finite-bit quantization mechanism preserving linear convergence of some algorithms cast as (1). Yet, there are several key differences between [22] and our work. First, the convergence analysis in [22] is applicable only to algorithms solving smooth, unconstrained optimization problems, and thus not to Problem (P) with $r \not\equiv 0$. Second, linear convergence under finite-bit quantization is explicitly proved in [22] only for schemes whose updates utilize current iterate information, namely: gradient descent (GD) over star networks and the Primal-Dual algorithm in [42] over mesh networks. This leaves open the question of whether distributed algorithms using historical information—e.g., in the form of gradient tracking or dual variables—are linearly convergent under finite-bit quantization, and under which conditions; renowned examples include EXTRA [43], AugDGM [44], DIGing [45], Harnessing [46], and NEXT [40], [41]. Our work provides a positive answer to these open questions. Third, the communication complexity of the scheme in [22] is only provided

¹We treat compression rules using deterministic mappings \mathcal{Q}^k as special cases of the random ones; in this case, the expected value operator will just return the deterministic value argument.




Fig. 1. Examples of star network (a) versus mesh topology (b).

for star networks, not for mesh networks, which instead is a novel contribution of this work for a wide class of distributed algorithms—see Table I and Sec. V. Fourth, [22] proposed an ad-hoc deterministic quantization rule while the proposed BC-rule encompasses several deterministic and random quantizations (including that in [22] as a special case), possibly using a variable number of bits (adapted to the input signal). As a result, even when customized to the setting/algorithms in [22], the BC-rule leads to more communication-efficient schemes, both analytically (see Sec. V) and numerically (see Sec. VI).

C. Organization and Notation

The remainder of this paper is organized as follows. Sec. II introduces the proposed black-box model, which casts distributed algorithms in the form (1). Sec. III embeds quantized communications, introduces the proposed BC-rule, and analyzes the convergence properties. Sec. IV describes the proposed quantizer, the ANQ, and studies its communication cost. Sec. V investigates the communication complexity, and customizes the proposed framework and convergence guarantees to several existing distributed algorithms, equipping them with the ANQ rule. Sec. VI provides some numerical results, while Sec. VII draws concluding remarks. All the proofs of our results are presented in the appendix.

Notation: Throughout the paper, we will use the following notation. We denote by \mathbb{Z}, \mathbb{R} the set of integers and real numbers, respectively. For any positive integer a , we define $[a] \triangleq \{1, \dots, a\}$. We denote by $\mathbf{0}, \mathbf{1}$, and \mathbf{I} the vector of all zeros, the vector of all ones, and the identity matrix, respectively (of appropriate dimension). For vectors $\mathbf{c}_1, \dots, \mathbf{c}_m$ and a set $\mathcal{S} \subseteq [m]$, define $\mathbf{c}_{\mathcal{S}} \triangleq \{\mathbf{c}_i : i \in \mathcal{S}\}$. We use $\|\cdot\|$ to denote a norm in the Euclidean space (whose dimension will be clear from the context); when a specific norm is used, such as ℓ_2 or ℓ_∞ , we will append the associated subscript to $\|\cdot\|$. The i th eigenvalue of a real, symmetric matrix \mathbf{G} is denoted by $\rho_i(\mathbf{G})$, ordered in non-increasing order such that $\rho_1(\mathbf{G}) \geq \dots \geq \rho_i(\mathbf{G}) \geq \rho_{i+1}(\mathbf{G})$. We will use a superscript to denote iteration counters of sequences generated by the algorithms, for instance, x^k will denote the value of the x -sequence at iteration k ; we will instead use $(x)^k$ for the k th-power of x . Finally, asymptotic behaviors of functions are captured by the standard big- \mathcal{O} , Θ , and Ω notations: 1) $g(x) = \mathcal{O}(h(x))$ as $x \rightarrow x_0$ if and only if $\limsup_{x \rightarrow x_0} |g(x)/h(x)| \in [0, \infty)$; 2) $g(x) = \Omega(h(x))$ if and only if $h(x) = \mathcal{O}(g(x))$; and 3) $g(x) = \Theta(h(x))$ if and only if $g(x) = \mathcal{O}(h(x)) = \Omega(h(x))$.

We model a network of m agents as a fixed, undirected, connected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = [m]$ is the set of vertices (agents) and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of edges (communication links); $(i, j) \in \mathcal{E}$ if there is a link between agents i and j , so that the two can send information to each other. We let $\mathcal{N}_i = \{j : (i, j) \in \mathcal{E}\}$ be the set of neighbors of agent i , and assume that $(i, i) \in \mathcal{E}$, i.e., $i \in \mathcal{N}_i$. Master/workers architectures will be considered as special instances—see Fig. 1.

II. A GENERAL DISTRIBUTED ALGORITHMIC FRAMEWORK: EXACT COMMUNICATIONS

In this section, we cast distributed algorithms to solve (P) in the form (1). As a warm-up, we begin with schemes using only current information to produce the next update (cf. Sec. II-A). We then generalize the model to capture distributed algorithms using historical information via multiple rounds of communications between computation steps (cf. Sec. II-B).

A. Warm-Up: A Class of Distributed Algorithms

We cast distributed algorithms in the form (1) by incorporating computations and communications as two separate steps. We use state variable \mathbf{z}_i to capture local information owned by agent i (including optimization variables) and $\hat{\mathbf{c}}_i$ to denote the signal transmitted by agent i to its neighbors.² Similarly to [22], the updates of the z, \hat{c} -variables read: for agent $i \in [m]$,

$$\begin{aligned} \hat{\mathbf{c}}_i^k &= \mathcal{C}_i(\mathbf{z}_i^k), & \text{(communication step)} \\ \mathbf{z}_i^{k+1} &= \mathcal{A}_i(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^k), & \text{(computation step)} \end{aligned} \quad (\text{M0})$$

where the function $\mathbf{z}_i \mapsto \mathcal{C}_i(\mathbf{z}_i)$ models the processing on the local information \mathbf{z}_i^k at the current iterate, generating the signal $\hat{\mathbf{c}}_i^k$ transmitted to agent i 's neighbors; the function $(\mathbf{z}_i, \hat{\mathbf{c}}_{\mathcal{N}_i}) \mapsto \mathcal{A}_i(\mathbf{z}_i, \hat{\mathbf{c}}_{\mathcal{N}_i})$ produces the update of the agent i 's state variable \mathbf{z}_i , based upon the local information at iteration k , and the signals received by its neighbors in \mathcal{N}_i .

Some examples: The algorithmic model (M0) captures a variety of distributed algorithms that build updates using single rounds of communications; examples include the renewed DGD [47], NIDS [48], and the Primal-Dual scheme [42]. To show a concrete example, consider DGD, which aims at solving a special instance of (P) with $r \equiv 0$; agents' updates read

$$\mathbf{x}_i^{k+1} = \left(\sum_{j=1}^m w_{ij} \mathbf{x}_j^k \right) - \gamma \nabla f_i(\mathbf{x}_i^k), \quad i \in [m],$$

where \mathbf{x}_i^k is the local copy owned by agent i at iteration k of the optimization variables \mathbf{x} , γ is a step-size, and w_{ij} 's are nonnegative weights properly chosen and compliant with the graph \mathcal{G} (i.e., $w_{ij} > 0$ if $(i, j) \in \mathcal{E}$; and $w_{ij} = 0$ otherwise). It is not difficult to check that DGD can be cast in the form (M0) by letting

$$\begin{aligned} \mathbf{z}_i^k &= \hat{\mathbf{c}}_i^k = \mathbf{x}_i^k \\ \mathcal{A}_i(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^k) &= \left(\sum_{j=1}^m w_{ij} \hat{\mathbf{c}}_j^k \right) - \gamma \nabla f_i(\mathbf{x}_i^k). \end{aligned}$$

Despite its generality, model (M0) leaves out several important distributed algorithms, specifically, the majority of

²Dimensions of these vectors are algorithm-dependent and omitted for simplicity, and will be clear from the context.

schemes employing correction of the gradient direction based on past state information—these are the best performing algorithms to date. Examples include EXTRA [43], DIGing [45] and their proximal version, NEXT/SONATA [40], [41], [49], and the ABC framework [50], just to name a few. Consider, for instance, NEXT/SONATA:

$$\begin{cases} \mathbf{x}_i^{k+1} = \sum_{j \in \mathcal{N}_i} w_{ij} (\mathbf{x}_j^k - \gamma \mathbf{y}_j^k) \\ \mathbf{y}_i^{k+1} = \sum_{j \in \mathcal{N}_i} w_{ij} (\mathbf{y}_j^k + \nabla f_j(\mathbf{x}_j^{k+1}) - \nabla f_j(\mathbf{x}_j^k)). \end{cases} \quad (3)$$

Clearly, these updates do not fit model (M0): the update of the y -variable uses information from two iterations (k and $k+1$). This calls for a more general model, introduced next.

B. Proposed General Model (Using Historical Information)

We generalize the algorithmic model (M0) as follows: for all $i \in [m]$,

$$\begin{aligned} \hat{\mathbf{c}}_i^{k,1} &= \mathcal{C}_i^1(\mathbf{z}_i^k, \mathbf{0}_{\mathcal{N}_i}), \\ \vdots \\ \hat{\mathbf{c}}_i^{k,R} &= \mathcal{C}_i^R(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,R-1}), \\ \mathbf{z}_i^{k+1} &= \mathcal{A}_i(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,1}, \dots, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,R}), \end{aligned} \quad (\text{computation step}) \quad (M)$$

These updates embed $R \geq 1$ rounds of local communications, via the functions $(\mathbf{z}_i, \hat{\mathbf{c}}_{\mathcal{N}_i}^{s-1}) \mapsto \mathcal{C}_i^s(\mathbf{z}_i, \hat{\mathbf{c}}_{\mathcal{N}_i}^{s-1})$; the function $(\mathbf{z}_i, \hat{\mathbf{c}}_{\mathcal{N}_i}^1, \dots, \hat{\mathbf{c}}_{\mathcal{N}_i}^R) \mapsto \mathcal{A}_i(\mathbf{z}_i, \hat{\mathbf{c}}_{\mathcal{N}_i}^1, \dots, \hat{\mathbf{c}}_{\mathcal{N}_i}^R)$ updates the local state by using the $\hat{\mathbf{c}}_{\mathcal{N}_i}$ signals received from the neighbors during all R rounds of communications, along with \mathbf{z}_i . Stacking agents' state-variables \mathbf{z}_i , communication signals $\hat{\mathbf{c}}_i$, and mappings \mathcal{C}_i^s and \mathcal{A}_i into the respective vectors \mathbf{z} , $\hat{\mathbf{c}}$, \mathcal{C}^s and \mathcal{A} , we can rewrite (M) in the compact form

$$\begin{aligned} \hat{\mathbf{c}}^{k,0} &= \mathbf{0}, \\ \hat{\mathbf{c}}^{k,s} &= \mathcal{C}^s(\mathbf{z}^k, \hat{\mathbf{c}}^{k,s-1}), s \in [R], \\ \mathbf{z}^{k+1} &= \mathcal{A}(\mathbf{z}^k, \hat{\mathbf{c}}^{k,1}, \dots, \hat{\mathbf{c}}^{k,R}). \end{aligned} \quad (\text{computation step})$$

Absorbing the communication signals $\hat{\mathbf{c}}^{k,s}$ in the mapping \mathcal{A} , we can finally write the above system as a fixed-point iterate on the z -variables only:

$$\mathbf{z}^{k+1} = \tilde{\mathcal{A}}(\mathbf{z}^k) \triangleq \mathcal{A}\left(\mathbf{z}^k, \mathcal{C}^1(\mathbf{z}^k, \mathbf{0}), \dots, \mathcal{C}^R(\mathbf{z}^k, \mathcal{C}^{R-1}(\mathbf{z}^k, \dots, \mathcal{C}^1(\mathbf{z}^k, \mathbf{0}) \dots))\right). \quad (M')$$

Under suitable conditions, the iterates (M') converge to fixed-points $\mathbf{z}^\infty = \tilde{\mathcal{A}}(\mathbf{z}^\infty)$ of the mapping $\tilde{\mathcal{A}}$, possibly constrained to a set $\mathcal{Z} \ni \mathbf{z}^\infty$. The convergence rate depends on the properties of $\tilde{\mathcal{A}}$; here we focus on *linear* convergence, which can be established under the following standard condition.

Assumption 1: Let $\tilde{\mathcal{A}} : \mathcal{Z} \rightarrow \mathcal{Z}$; the following hold: (i) $\tilde{\mathcal{A}}$ admits a fixed-point \mathbf{z}^∞ ; and (ii) $\tilde{\mathcal{A}}$ is λ -pseudo-contractive on \mathcal{Z} w.r.t. some norm $\|\bullet\|$, that is, there exists $\lambda \in (0, 1)$ such that

$$\|\tilde{\mathcal{A}}(\mathbf{z}) - \mathbf{z}^\infty\| \leq \lambda \cdot \|\mathbf{z} - \mathbf{z}^\infty\|, \quad \forall \mathbf{z} \in \mathcal{Z}.$$

Without loss of generality, the norm $\|\bullet\|$ is scaled such that $\|\bullet\|_2 \leq \|\bullet\|^3$.

³This is always possible since $\|\bullet\|$ is a norm defined on a finite-dimensional field.

The following convergence result follows readily from Assumption 1 and [52, Ch. 3, Prop. 1.2].

Theorem 2: Let $\tilde{\mathcal{A}} : \mathcal{Z} \rightarrow \mathcal{Z}$ satisfy Assumption 1. Then: i) the fixed point \mathbf{z}^∞ is unique; and ii) the sequence $\{\mathbf{z}^k\}$ generated by the update (M') converges Q-linearly to \mathbf{z}^∞ w.r.t. the norm $\|\bullet\|$ at rate λ , i.e., $\|\mathbf{z}^{k+1} - \mathbf{z}^\infty\| \leq \lambda \cdot \|\mathbf{z}^k - \mathbf{z}^\infty\|$.

Discussion: Our model treats the underlying unquantized algorithm as a black-box with convergence rate λ , which depends on the optimization problem parameters—the smoothness and strong convexity constants L and μ of the agents' functions (often via the condition number κ)—and the network connectivity (spectral properties of \mathbf{W}). Table II collects some representative examples.

The algorithmic framework (M) encompasses a variety of distributed algorithms, while Theorem 2 captures their convergence properties; in addition to the schemes covered by (M0) as a special case when $R = 1$, (M) can also represent EXTRA [43] and its proximal version [50], NEXT [40], [41], [49] and its proximal version [50], DIGing [45] and its proximal version [50], and prox-NIDS [50]. Appendix D provides specific expressions for the mappings \mathcal{A} and \mathcal{C}^s for each of the above algorithms, along with their convergence properties under Theorem 2; here, we elaborate on the NEXT algorithm (3) as an example. It can be cast in the form (M) by using $R = 2$ rounds of communications and letting

$$\begin{aligned} \hat{\mathbf{c}}_i^k &= \begin{bmatrix} \mathbf{x}_i^k \\ \mathbf{y}_i^k \end{bmatrix}, \quad \hat{\mathbf{c}}_i^{k,1} = \mathbf{x}_i^k - \gamma \mathbf{y}_i^k, \\ \hat{\mathbf{c}}_i^{k,2} &= \mathbf{y}_i^k + \nabla f_i \left(\sum_{j \in \mathcal{N}_i} w_{ij} \hat{\mathbf{c}}_j^{k,1} \right) - \nabla f_i(\mathbf{x}_i^k), \text{ and} \\ \mathbf{z}_i^{k+1} &= \mathcal{A}_i(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,1}, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,2}) = \begin{bmatrix} \sum_{j \in \mathcal{N}_i} w_{ij} \hat{\mathbf{c}}_j^{k,1} \\ \sum_{j \in \mathcal{N}_i} w_{ij} \hat{\mathbf{c}}_j^{k,2} \end{bmatrix}. \end{aligned}$$

III. A GENERAL DISTRIBUTED ALGORITHMIC FRAMEWORK: QUANTIZED COMMUNICATIONS

In this section, we equip the distributed algorithmic framework (M) with quantized communications. The communication channel between any two agents is modeled as a noiseless digital channel: only quantized signals are received with no errors. This means that, in each of the communication rounds, the signals $\hat{\mathbf{c}}_j^{k,1}, \dots, \hat{\mathbf{c}}_j^{k,R}$, $j \in \mathcal{N}_i$, received by agent i may no longer coincide with the intended, unquantized ones $\mathcal{C}_j^1(\mathbf{z}_j^k, \mathbf{0}_{\mathcal{N}_j}), \dots, \mathcal{C}_j^R(\mathbf{z}_j^k, \hat{\mathbf{c}}_{\mathcal{N}_j}^{k,R-1})$, generated at the transmitter side of agents $j \in \mathcal{N}_i$. This calls for a proper encoding/decoding mechanism that transfers, via quantized communications, the aforementioned unquantized signals at the receiver sides with limited distortion. Here, we leverage differential encoding/decoding techniques [11] coupled with a novel quantization rule.

We begin by recalling the idea of quantized differential encoding/decoding in the context of a point-to-point communication—the same mechanism will be then embedded in the communication of the distributed multi-agent framework (M). Consider a transmitter-receiver pair; let \mathbf{c}^k be the unquantized information generated at iteration k , intended to be transferred to the receiver over the digital channel, and let $\hat{\mathbf{c}}^k$ be the estimate of \mathbf{c}^k , built using quantized information.

TABLE II
CONVERGENCE RATE OF REPRESENTATIVE LINEARLY CONVERGENT DISTRIBUTED ALGORITHMS CAST IN THE PROPOSED FRAMEWORK (M)

Algorithm	Problem	Network	Convergence rate λ
GD over star networks [51]	(P) with $r \equiv 0$	Star	$\frac{\kappa-1}{\kappa+1}$
(Prox-)EXTRA [50]	(P)	Mesh	$\max \left\{ \frac{\kappa}{\kappa+1}, \sqrt{\rho_2(\mathbf{W})} \right\}$
(Prox-)NIDS [50]	(P)	Mesh	$\max \left\{ \frac{\kappa-1}{\kappa+1}, \sqrt{\rho_2(\mathbf{W})} \right\}$
NIDS [48]	(P) with $r \equiv 0$	Mesh	$\max \left\{ \sqrt{1-\kappa^{-1}}, \sqrt{\rho_2(\mathbf{W})} \right\}$
(Prox-)NEXT [50]	(P)	Mesh	$\max \left\{ \frac{\kappa-1}{\kappa+1}, \sqrt{1-(1-\rho_2(\mathbf{W}))^2} \right\}$
(Prox-)DIGing [50]	(P)	Mesh	$\max \left\{ \frac{\kappa}{\kappa+1}, \sqrt{1-(1-\rho_2(\mathbf{W}))^2} \right\}$
Primal-Dual [22], [42]	(P) with $r \equiv 0$	Mesh	$\frac{\frac{\rho_1(\mathbf{L})}{\rho_{m-1}(\mathbf{L})} - \frac{1}{\kappa}}{\frac{\rho_1(\mathbf{L})}{\rho_{m-1}(\mathbf{L})} + \frac{1}{\kappa}}$

The convergence rates above are obtained under the tuning described in Appendix D; κ , L , and μ is the condition number, the smoothness constant and the strong convexity constant of each f_i , respectively, and \mathbf{W} is the gossip matrix used in the communication steps of the algorithms; for the Primal-Dual algorithm, \mathbf{L} is the graph Laplacian, defined in Appendix D.7.

The differential encoding/decoding rule reads: $\hat{\mathbf{c}}^0 = \mathbf{0}$, and for $k \geq 1$,

$$\begin{cases} \mathbf{q}^k = \mathcal{Q}^k(\mathbf{c}^k - \hat{\mathbf{c}}^{k-1}), \\ \hat{\mathbf{c}}^k = \hat{\mathbf{c}}^{k-1} + \mathbf{q}^k, \end{cases} \quad (4)$$

where \mathcal{Q}^k is the quantization operator (a map from real vectors to the set of quantized points), possibly dependent on iteration k . In words, at each iteration, the encoder quantizes the *prediction error* $\mathbf{c}^k - \hat{\mathbf{c}}^{k-1}$ rather than the current estimate \mathbf{c}^k , generating the quantized signal \mathbf{q}^k , which is then transmitted over the digital channel. The estimate $\hat{\mathbf{c}}^k$ of \mathbf{c}^k is then built from \mathbf{q}^k using a one-step prediction rule. Since \mathbf{q}^k is received unaltered, $\hat{\mathbf{c}}^k$ is identical at the transmitter's and receiver's sides. Note that, when quantization errors are negligible $\mathbf{q}^k = \mathcal{Q}^k(\mathbf{c}^k - \hat{\mathbf{c}}^{k-1}) \approx \mathbf{c}^k - \hat{\mathbf{c}}^{k-1}$, the estimate reads $\hat{\mathbf{c}}^k = \hat{\mathbf{c}}^{k-1} + \mathbf{q}^k \approx \hat{\mathbf{c}}^{k-1} + \mathbf{c}^k - \hat{\mathbf{c}}^{k-1} = \mathbf{c}^k$.

We can now introduce our distributed algorithmic framework using quantized communications, as described in Algorithm 1; it embeds the differential encoding/decoding rule (4) in each communication round of model (M). The fixed-point based formulation of Algorithm 1 then reads: for $i \in [m]$,

$$\begin{aligned} \mathbf{c}_i^{k,1} &= \mathcal{C}_i^1(\mathbf{z}_i^k, \mathbf{0}_{\mathcal{N}_i}), \\ \hat{\mathbf{c}}_i^{k,1} &= \hat{\mathbf{c}}_i^{k-1,1} + \mathcal{Q}_i^k(\mathbf{c}_i^{k,1} - \hat{\mathbf{c}}_i^{k-1,1}), \\ &\vdots \\ \mathbf{c}_i^{k,R} &= \mathcal{C}_i^R(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,R-1}), \\ \hat{\mathbf{c}}_i^{k,R} &= \hat{\mathbf{c}}_i^{k-1,R} + \mathcal{Q}_i^k(\mathbf{c}_i^{k,R} - \hat{\mathbf{c}}_i^{k-1,R}), \\ \mathbf{z}_i^{k+1} &= \mathcal{A}_i(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,1}, \dots, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,R}), \end{aligned} \quad \left. \begin{array}{l} \text{(multiple quantized communication rounds)} \\ \text{(computation step)} \end{array} \right\}$$

Stacking agents' state-variables \mathbf{z}_i , signals \mathbf{c}_i and $\hat{\mathbf{c}}_i$, and mappings \mathcal{C}_i^s , \mathcal{A}_i , and \mathcal{Q}_i^k into the respective vectors \mathbf{z} , \mathbf{c} , $\hat{\mathbf{c}}$, \mathcal{C}^s , \mathcal{A} , and \mathcal{Q}^k , we can rewrite the above steps in compact form as

$$\begin{aligned} \hat{\mathbf{c}}^{k,0} &= \mathbf{0}, \\ \mathbf{c}^{k,s} &= \mathcal{C}^s(\mathbf{z}^k, \hat{\mathbf{c}}^{k,s-1}), \\ \hat{\mathbf{c}}^{k,s} &= \hat{\mathbf{c}}^{k-1,s} + \mathcal{Q}^k(\mathbf{c}^{k,s} - \hat{\mathbf{c}}^{k-1,s}), \\ s &\in [R], \\ \mathbf{z}^{k+1} &= \mathcal{A}(\mathbf{z}^k, \hat{\mathbf{c}}^{k,1}, \dots, \hat{\mathbf{c}}^{k,R}). \end{aligned} \quad \left. \begin{array}{l} \text{(multiple quantized communication rounds)} \\ \text{(computation step)} \end{array} \right\} \quad (Q-M)$$

Algorithm 1 Distributed Algorithmic Framework With Quantized Communications

Initialization: $\hat{\mathbf{c}}^{-1,s} \triangleq \mathbf{0}$, for all $s \in [R]$; and $\mathbf{z}^0 \in \mathcal{Z}$. Set $k = 0$;

Iteration $k \rightarrow k + 1$

(S.1) : Multiple quantized communication rounds

for $s = 1, \dots, R$, each agent i :

- Computes $\mathbf{c}_i^{k,s} = \mathcal{C}_i^s(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,s-1})$ [with $\hat{\mathbf{c}}_i^{k,0} \triangleq \mathbf{0}$];
- Generates $\mathbf{q}_i^{k,s} = \mathcal{Q}_i^k(\mathbf{c}_i^{k,s} - \hat{\mathbf{c}}_i^{k-1,s})$ and broadcasts it to its neighbors $j \in \mathcal{N}_i$;
- Upon receiving the signals $\mathbf{q}_j^{k,s}$ from its neighbors $j \in \mathcal{N}_i$, it reconstructs $\hat{\mathbf{c}}_j^{k,s}$ as

$$\hat{\mathbf{c}}_j^{k,s} = \hat{\mathbf{c}}_j^{k-1,s} + \mathbf{q}_j^{k,s}, \quad j \in \mathcal{N}_i;$$

end

(S.2) : Computation Step

Each agent i updates its own \mathbf{z}_i^{k+1} according to

$$\mathbf{z}_i^{k+1} = \mathcal{A}_i(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,1}, \dots, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,R}).$$

Model (Q-M) paves the way to a unified design and convergence analysis of several distributed algorithms—all the schemes cast in the form (M)—employing quantization in the communications, as elaborated next.

A. Convergence Analysis

We begin by establishing sufficient conditions on the quantization mapping \mathcal{Q}^k and algorithmic functions \mathcal{A} and \mathcal{C} in (Q-M) to preserve linear convergence

1) *On the Quantization Mapping \mathcal{Q}^k :* A first critical choice is the quantizer \mathcal{Q} (we omit the dependence on k and i for notation simplicity), including both random and deterministic quantization rules (the latter as a special cases of the former). For random quantization, the function $\mathcal{Q}(\mathbf{x})$ is a random variable for any given $\mathbf{x} \in \mathbb{R}^d$, defined on a suitable probability space (generally dependent on \mathbf{x}). We define the following novel *biased compression rule* (BC-rule), for each agent i .

Definition 3 (Biased Compression Rule): Given $\mathbf{x} \in \mathbb{R}^d$, $\mathcal{Q}(\mathbf{x})$ (possibly, a random variable defined on a suitable probability space) satisfies the BC-rule with *bias* $\eta \geq 0$ and *compression rate* $\omega \in [0, 1)$ if

$$\sqrt{\mathbb{E}[\|\mathcal{Q}(\mathbf{x}) - \mathbf{x}\|_2^2]} \leq \sqrt{d}\eta + \omega\|\mathbf{x}\|_2, \quad \forall \mathbf{x} \in \mathbb{R}^d. \quad (5)$$

When $\mathcal{Q}(\mathbf{x})$ is a deterministic map, (5) reduces to

$$\|\mathcal{Q}(\mathbf{x}) - \mathbf{x}\|_2 \leq \sqrt{d}\eta + \omega\|\mathbf{x}\|_2, \quad \forall \mathbf{x} \in \mathbb{R}^d. \quad (6)$$

Roughly speaking, the bias η determines the basic spacing between quantization points, uniform across the entire domain. On the other hand, the compression term ω adds a non-uniform spacing between quantization points: points farther away from 0 have more separation.

The BC-rule encompasses and generalizes several existing compression and quantization rules proposed in the literature for specific algorithms, deterministic [6], [7], [9], [11], [13], [14], [16], [17], [20]–[22], [27]–[29], [32], [35], [38] and random [6], [8], [10], [12], [15], [18], [19], [23]–[26], [29]–[31], [33]–[39]. Specifically, **(i)** the compression rules proposed in [6], [24]–[26], [29]–[31], [33]–[39] (resp. [6], [27]–[29], [32], [35], [38]) can be interpreted as *unbiased* instances of (5) [resp. (6)], i.e., corresponding to $\eta = 0$. The proof of Lemma 7 in Sec. IV will show that such special instances theoretically require an infinite number of bits to encode quantized signals, even when the input \mathbf{x} is bounded. In practice, they are successfully implemented using finite bits at the machine precision (e.g., to encode some scalar quantities, such as the norm of the signal to be transmitted [36]). However, their convergence analyses tacitly assume that errors at machine precision are negligible. Hence, these schemes have no performance guarantees when implemented with limited (below machine) precision quantizations. On the other hand, as it will be seen in Sec. IV, the bias term $\eta > 0$ in the proposed BC-rule guarantees that quantized signals satisfying such rule can be encoded using a finite number of bits, hence can be implemented with limited precision. **(ii)** The quantization rules in [8], [10], [12], [15], [18], [19], [23] (resp. [7], [9], [11], [13], [14], [16], [17], [20]–[22]) are special instances of the BC-rule (5) [resp. (6)], with $\omega = 0$. While they can be implemented using a finite number of bits (provided that the signals to be quantized are uniformly bounded), they do not take advantage of the degree of freedom offered by the compression rate ω to improve communication efficiency, as shown numerically in Sec. VI-D.

2) On the Algorithmic Mappings \mathcal{A} and \mathcal{C}^s : Our analysis requires additional standard conditions on the mappings \mathcal{A} and \mathcal{C}^s to preserve linear convergence under quantization. Roughly speaking, the functions \mathcal{A} and \mathcal{C}^s should vary smoothly with respect to perturbations in their arguments, so that small quantization errors result in small deviations from the trajectory of the unquantized algorithm, as postulated next.⁴

⁴For the sake of notation, the constants L_A , L_C and L_Z defined in Assumptions 4 and 5 are assumed to be independent of the index s (communication round). Our convergence results can be readily extended to constants dependent on s .

Assumption 4: There exists a constant $L_A \geq 0$ such that, for every $s \in [R]$, it holds

$$\begin{aligned} & \|\mathcal{A}(\mathbf{z}, \mathbf{c}^1, \dots, \mathbf{c}^s, \dots, \mathbf{c}^R) - \mathcal{A}(\mathbf{z}, \mathbf{c}^1, \dots, \tilde{\mathbf{c}}^s, \dots, \mathbf{c}^R)\| \\ & \leq L_A \|\mathbf{c}^s - \tilde{\mathbf{c}}^s\|_2, \end{aligned} \quad (7)$$

for all $\mathbf{c}^s, \tilde{\mathbf{c}}^s \in \mathbb{R}^{md}$, uniformly with respect to $\mathbf{z} \in \mathcal{Z}$, and $\mathbf{c}^1, \dots, \mathbf{c}^{s-1}, \mathbf{c}^{s+1}, \dots, \mathbf{c}^R \in \mathbb{R}^{md}$.

Assumption 5: There exist constants $L_C, L_Z \geq 0$ such that

$$\|\mathcal{C}^s(\mathbf{z}, \mathbf{c}) - \mathcal{C}^s(\mathbf{z}, \mathbf{c}')\|_2 \leq L_C \|\mathbf{c} - \mathbf{c}'\|_2, \quad \forall \mathbf{c}, \mathbf{c}' \in \mathbb{R}^{md}, \quad (8)$$

$$\|\mathcal{C}^s(\mathbf{z}, \mathbf{c}) - \mathcal{C}^s(\mathbf{z}', \mathbf{c})\|_2 \leq L_Z \|\mathbf{z} - \mathbf{z}'\|_2, \quad \forall \mathbf{z}, \mathbf{z}' \in \mathcal{Z}, \quad (9)$$

uniformly with respect to $\mathbf{z} \in \mathcal{Z}$ and $\mathbf{c} \in \mathbb{R}^{md}$, respectively.

These assumptions are quite mild, and satisfied by a variety of existing distributed algorithms, as shown in Appendix D. We are now ready to introduce our main convergence result.

Theorem 6: Let $\{\mathbf{z}^k\}$ be the sequence generated by Algorithm 1 under Assumptions 1, 4, and 5, with \mathcal{Q}^k satisfying the BC-rule (5) with bias $\eta = \eta^0 \cdot (\sigma)^k$ and compression rate $\omega \in [0, \bar{\omega}(\sigma))$, for some $\sigma \in (\lambda, 1)$ and $\eta^0 > 0$, where $\bar{\omega}(\sigma)$ is defined as

$$\bar{\omega}(\sigma) \triangleq \frac{\sigma}{R} \cdot \frac{\sigma - \lambda}{\sigma - \lambda + 2L_A L_Z [R \max\{1, (2L_C)^{R-1}\}]^2}. \quad (10)$$

Then,

$$\sqrt{\mathbb{E}[\|\mathbf{z}^k - \mathbf{z}^\infty\|_2^2]} \leq V_0 \cdot (\sigma)^k, \quad k = 0, 1, \dots,$$

where V_0 is a positive constant, whose expression is given in (39), Appendix A.

Proof: See Appendix A. ■

Note that, when the deterministic instance of the BC-rule is used [see (6)], the convergence rate reads $\|\mathbf{z}^k - \mathbf{z}^\infty\|_2 \leq V_0 \cdot (\sigma)^k$, for all $k = 0, 1, \dots$

Theorem 6 shows that linear convergence is achievable when communications are quantized in distributed optimization, provided that the bias η and compression rate ω of the BC-rule are chosen suitably. The shrinking requirement on the bias η (linear at rate σ) is *not* restrictive, since it is consistent with the contraction dynamics of the iterates: the range of inputs to the quantizer vanishes in a similar fashion, a fact that guarantees that quantized values can be encoded using a uniformly bounded number of bits (see Theorem 12, Sec. V). Similarly, the bound on ω guarantees that quantization errors along the iterates do not accumulate excessively.

Theorem 6 certifies linear convergence in terms of number of iterations. However, the algorithm that uses quantized communications converges slower (with rate σ) than its unquantized counterpart (rate λ), revealing a tension between the amount of quantization/compression of the transmitted signals (measured by η and ω) and the resulting linear convergence rate σ : as we will see in the forthcoming sections, this tension entails a trade-off between faster convergence (closer to that of the unquantized algorithm) and communication cost, which we aim to characterize. Building on this result, we will also investigate the communication complexity of the schemes (Q-M) – the total number of bits needed to reach an ε -solution of problem (P). Since this analysis depends on the specific quantizer design, the next section introduces a novel quantizer

that efficiently implements the BC-rule, and a communication-efficient bit-encoding/decoding scheme. When embedded in (Q-M), the proposed quantization leads to linearly convergent distributed algorithms whose communication complexity compares favorably with that of existing ad-hoc schemes (V).

IV. NON-UNIFORM QUANTIZER WITH ADAPTIVE ENCODING/DECODING

As discussed in Sec. III, the BC-rule encompasses a variety of quantizer designs. In this section, we propose a quantizer that fulfills the BC-rule with minimum number of quantization points (Sec. IV-A). The quantizer is then coupled with a communication-efficient bit-encoding/decoding rule which enables transmission on the digital channel (Sec. IV-B). We refer to the proposed quantizer coupled with the encoding/decoding scheme as *Adaptive encoding Non-uniform Quantizer* (ANQ).

A. Quantizer Design

Since no information is assumed on the distribution of the input signal, a natural approach is to quantize each vector signal component-wise. We design such a scalar quantizer $\mathcal{Q} : [-\delta, \delta] \rightarrow \mathbb{Q}$ under the BC-rule by minimizing the number of quantization points $|\mathbb{Q}|$ for a fixed input range $[-\delta, \delta]$. Equivalently, we seek \mathcal{Q} that maximizes δ under the BC-rule, for a given number $N = |\mathbb{Q}|$ of quantization points. These designs are provided in Lemmas 7 and 8 for the deterministic and probabilistic cases, respectively. For convenience, we focus on the case of N odd; the case of N even is provided in Appendices B and B. We point out that the restriction of the input of \mathcal{Q} to $[-\delta, \delta]$, as opposed to the unconstrained domain in the BC-rule (Definition 3), is instrumental in formulating the quantizer's design as an optimization problem. As shown in Lemmas 7 and 8, the resulting optimized quantization points are independent of the specific choice of δ ; hence, the proposed quantizer can be used (component-wise) for input signals in \mathbb{R}^d .

Lemma 7 (Deterministic Quantizer): Let $\mathcal{Q} : [-\delta, \delta] \rightarrow \mathbb{Q}$. The maximum range δ that can be quantized using $|\mathbb{Q}| = N$ (odd) points under the BC-rule (6) with bias $\eta \geq 0$ and compression rate $\omega \in [0, 1)$ is

$$\delta(\eta, \omega, N) = \frac{q_{(N-1)/2} + q_{(N+1)/2}}{2}, \quad (11)$$

with quantization points

$$q_\ell = -q_{-\ell} = \frac{\eta}{\omega} \left[\left(\frac{1+\omega}{1-\omega} \right)^\ell - 1 \right], \quad \ell \geq 0. \quad (12)$$

The resulting optimal quantization rule reads: $x \mapsto \mathcal{Q}(x) = q_{\ell(x)}$, with

$$\ell(x) = \text{sign}(x) \cdot \left\lceil \frac{\ln(1-\omega) + \ln(1 + \frac{\omega}{\eta}|x|)}{\ln(1+\omega) - \ln(1-\omega)} \right\rceil. \quad (13)$$

Proof: See Appendix B.1. ■

Lemma 8 (Probabilistic Quantizer): For any given $x \in [-\delta, \delta]$, let $\mathcal{Q}(x) \in \mathbb{Q}$ be a random variable defined on a suitable probability space. The maximum range δ that can be

quantized using $|\mathbb{Q}| = N$ (odd) points under the BC-rule (5) with $\mathbb{E}[\mathcal{Q}(x)] = x$, bias $\eta \geq 0$ and compression rate $\omega \in [0, 1)$ is

$$\delta(\eta, \omega, N) = q_{(N-1)/2}, \quad (14)$$

with quantization points

$$q_\ell = -q_{-\ell} = \frac{\eta}{\omega} \left[\left(\sqrt{1 + (\omega)^2} + \omega \right)^{\ell} - 1 \right], \quad \ell \geq 0. \quad (15)$$

The resulting optimal quantization rule reads: $x \mapsto \mathcal{Q}(x) = q_{\ell(x)}$, with

$$\ell(x) = \begin{cases} \ell - 1, & \text{w.p. } \frac{q_\ell - x}{q_\ell - q_{\ell-1}}; \\ \ell, & \text{w.p. } \frac{x - q_{\ell-1}}{q_\ell - q_{\ell-1}}, \end{cases} \quad \text{and} \\ \ell = \text{sign}(x) \left\lceil \frac{\ln(1 + \frac{\omega}{\eta}|x|)}{2 \ln \left(\sqrt{1 + (\omega)^2} + \omega \right)} \right\rceil. \quad (16)$$

Proof: See Appendix B.2. ■

From these lemmas, one infers that quantization points under the BC-rule should be non-uniformly spaced—hence the name ANQ. Furthermore, the deterministic quantizer maps an input x to the nearest q_ℓ , whereas the probabilistic one maps x to one of the two nearest quantization points, selected randomly such that $\mathbb{E}[\mathcal{Q}(x)] = x$.

Note that, when specialized to the conventional compression rule that uses $\eta = 0$, the two lemmas above yield $\delta(0, \omega, N) = 0$ for any finite N , implying that infinite quantization points (hence number of bits) are required to encode signals. The next corollary formalizes this negative result.

Corollary 9 (Converse): The (unbiased, $\eta = 0$) compression rule (2) cannot be satisfied using a finite number of quantization points to quantize signals within a given range $[-\delta, \delta^d]$, for any finite $\delta > 0$. Therefore, the compression rules in [6], [24], [26]–[33], [36]–[38] theoretically require an infinite number of bits to encode quantized signals.

Proof: See Appendix B.3. □

Note that the index $\ell(x)$ is a sufficient information to reconstruct the quantization point $q_{\ell(x)}$ at the receiver. In the next section, we present a communication-efficient finite bit-encoding/decoding scheme to transmit $\ell(x)$ over the digital channel.

B. Adaptive Encoding Scheme

It remains to design an encoding/decoding scheme mapping the index $\ell(x)$ into a finite-bit representation, to be transmitted over the digital channel. To do so, we adopt an adaptive number of bits, based upon the value of $\ell(x)$, as detailed next.

We assume that a constellation $\mathbb{S} = [S] \cup \{0\}$ of $S + 1$ symbols is used, with $S \geq 2$ (this might be obtained as $\mathbb{S} \equiv (\tilde{\mathbb{S}})^w$, by concatenating sequences of w symbols from a smaller constellation $\tilde{\mathbb{S}}$). We use the symbol 0 to indicate the end of an information sequence, and the remaining S symbols in the set $[S]$ to encode the value of $\ell(x)$. With $\tilde{\mathcal{L}}_{-1} \equiv \emptyset$, let

$$\tilde{\mathcal{L}}_b \equiv \left\{ -\left\lceil \frac{S^{b+1} - 1}{2(S-1)} \right\rceil + 1, \dots, \left\lceil \frac{S^{b+1} - 1}{2(S-1)} \right\rceil \right\}, \quad b = 0, 1, \dots,$$

and

$$\mathcal{L}_b = \tilde{\mathcal{L}}_b \setminus \tilde{\mathcal{L}}_{b-1}, \quad b = 0, 1, \dots \quad (17)$$

It is not difficult to see that $\{\mathcal{L}_b : b = 0, 1, \dots\}$ creates a partition of \mathbb{Z} and $|\mathcal{L}_b| = (S)^b$. Therefore, a natural way to encode $\ell(x) \in \mathcal{L}_b$ is to use a unique sequence of b symbols from $[S]$, followed by 0 to mark the end of the information sequence. $\ell(x)$ is then encoded as $[s_1, \dots, s_b, 0] \in (\mathbb{S})^{b+1}$.

Upon receiving this sequence, the receiver can detect the start and end of the information symbols, and decode the associated $\ell(x)$ by inverting the symbol-mapping. The communication cost to transmit the index $\ell(x) \in \mathcal{L}_b$ is thus $b + 1$ (symbols), which leads to the following upper bound on the communication cost incurred by each agent i to quantize and encode a d -dimensional vector \mathbf{x} . Again, we focus on the case when N is odd; the other case is provided in the proof in Appendix C.4.

Lemma 10: The number of bits $C(\mathbf{x})$ required by the ANQ with bias $\eta \geq 0$ and compression rate $\omega \geq 0$ and constellation of $S + 1$ symbols to quantize and encode an input signal $\mathbf{x} \in \mathbb{R}^d$ is upper bounded by

(i) *Deterministic Quantizer:*

$$C(\mathbf{x}) \leq 3d \log_2(S + 1) + d \log_2(S + 1) \log_S \left(2 + \frac{\ln(1 - \omega) + \ln(1 + \frac{\omega \|\mathbf{x}\|_2}{\sqrt{d}\eta})}{\ln(1 + \omega) - \ln(1 - \omega)} \right) \text{ bits; } \quad (18)$$

(ii) *Probabilistic Quantizer With $\mathbb{E}[\mathcal{Q}(\mathbf{x})] = \mathbf{x}$:*

$$C(\mathbf{x}) \leq 3d \log_2(S + 1) + d \log_2(S + 1) \log_S \left(2 + \frac{\ln(1 + \frac{\omega \|\mathbf{x}\|_2}{\sqrt{d}\eta})}{2 \ln(\sqrt{1 + (\omega)^2} + \omega)} \right) \text{ bits, a.s.. } \quad (19)$$

Proof: See Appendix C.4. \square

Compared with existing deterministic quantizers [21], [22] that are special cases of the BC-rule (with $\omega = 0$), the proposed ANQ adapts the number of bits to the input signal—less bits for smaller input signals (mapped to smaller ℓ) and more bits for larger ones (mapped to larger ℓ)—rather than using a fixed number of bits determined by the worst-case input signal [21], [22]. This leads to more communication-efficient schemes, as will be certified by Theorem 13.

V. COMMUNICATION COMPLEXITY OF (Q-M) BASED ON THE ANQ RULE

We now study the communication complexity of the distributed schemes falling within the framework (Q-M) and using the ANQ to quantize communications. Our results complement Theorem 6 and are of two types: (i) first, we characterize the trade-off between the number of bits/agent required by the ANQ at each iteration and the linear convergence rate (Theorem 12); (ii) then, we investigate the communication complexity, characterizing the total number of bits/agent transmitted to achieve an ε -solution of (P) (Theorem 13). Since these results are applicable to any distributed algorithm within the setting of (Q-M), we finally customize (ii) to some specific instances (Sec. V-A).

Throughout this section, all the results stated in terms of \mathcal{O} -notation are meant asymptotically when $m, d \rightarrow \infty$. Also, the

following additional mild assumption is postulated, which is satisfied by a variety of existing algorithms, see Appendix D.

Assumption 11: The constants L_A, L_C, L_Z, R and the initial conditions $\|\mathcal{C}^s(\mathbf{z}^0, \mathbf{0})\|_2$ and $\|\mathbf{z}^0 - \mathbf{z}^\infty\|$ satisfy

$$\begin{aligned} L_A \cdot L_Z &= \mathcal{O}(1), \quad L_C = \mathcal{O}(1), \quad R = \mathcal{O}(1), \\ \|\mathbf{z}^0 - \mathbf{z}^\infty\| &= \mathcal{O}(\sqrt{md}) \\ \text{and } \|\mathcal{C}^s(\mathbf{z}^0, \mathbf{0})\|_2 &= \mathcal{O}(L_Z \sqrt{md}), \quad \forall s \in [R]. \end{aligned}$$

Our first result on the number of bits transmitted at each iteration to sustain linear convergence is summarized next.

Theorem 12: Instate the setting of Theorem 6, under the additional Assumption 11. Furthermore, suppose that the deterministic ANQ (or probabilistic ANQ with $\mathbb{E}[\mathcal{Q}(\mathbf{x})] = \mathbf{x}$) is used to quantize all the communications in (Q-M), with $\eta^0 = \Theta(L_Z(\sigma - \lambda))$ and ω such that $\frac{1 - \omega/\bar{\omega}(\sigma)}{\sigma - \lambda} = \Omega(1)$. Then, linear convergence $\sqrt{\mathbb{E}[\|\mathbf{z}^k - \mathbf{z}^\infty\|_2^2]} = \mathcal{O}(\sqrt{md} \cdot (\sigma)^k)$, $k = 0, 1, \dots$, is achieved with an average number of bits/agent at every iteration k of order

$$\mathcal{O}\left(d \cdot \log\left(1 + \frac{1}{\sigma(\sigma - \lambda)}\right)\right). \quad (20)$$

Proof: See Appendix C.1. \blacksquare

The following comments are in order.

(i) As expected, the faster the quantized algorithm (smaller σ), the larger the communication cost; in particular, when $\sigma \rightarrow \lambda$, the number of bits required to sustain linear convergence at rate σ grows indefinitely. In other words, an infinite number of bits is required if a quantized distributed scheme (Q-M) wants to match the convergence rate of its unquantized counterpart.

(ii) It is interesting to contrast the communication efficiency (bits transmitted per iteration) of the proposed model (Q-M) equipped with the ANQ with that of existing schemes applicable to special instances of (P) or network topologies. Specifically, [21], [22] study (P) with $r \not\equiv 0$ over mesh networks; the scheme therein converges linearly (under suitable tuning/assumptions) while using

$$\mathcal{O}\left(d \log\left(1 + \frac{\sqrt{md}}{\sigma(\sigma - \lambda)}\right)\right) \text{ bits/agent/iteration.}$$

The algorithm in [22] is also applicable to (P), with $r \not\equiv 0$, over star-networks; it uses

$$\mathcal{O}\left(d \log\left(1 + \frac{\sqrt{d}}{\sigma(\sigma - \lambda)}\right)\right) \text{ bits/agent/iteration.}$$

Both are less favorable than (20), due to the fact that the ANQ adapts the number of bits to the input signal rather than adopting a constant number of bits for any input signal.

(iii) Theorem 12 reveals a tension between convergence rate (the closer σ to λ , the faster the algorithm) and number of transmitted bits per iteration (the larger σ , the smaller the cost). In the following, we provide a favorable choice of σ that resolves this tension by characterizing the communication complexity, i.e., the total number of bits transmitted per agent to achieve a target ε -accuracy.

Theorem 13: Instate the setting of Theorem 12, with σ chosen so that

$$\frac{(1-\lambda)^2}{(1-\sigma)(\sigma-\lambda)} = \mathcal{O}(1).$$

Then, the average number of bits transmitted per agent to achieve $(1/m)\mathbb{E}[\|\mathbf{z}^k - \mathbf{z}^\infty\|_2^2] \leq \varepsilon$ scales as

$$\mathcal{O}\left(d \cdot \log\left(1 + \frac{1}{1-\lambda}\right) \frac{1}{1-\lambda} \log(d/\varepsilon)\right) \text{ bits/agent}, \quad (21)$$

achieved in

$$\mathcal{O}\left(\frac{1}{1-\lambda} \cdot \log(d/\varepsilon)\right) \text{ iterations} \quad (22)$$

and with

$$\mathcal{O}\left(d \cdot \log\left(1 + \frac{1}{1-\lambda}\right)\right) \text{ bits/agent/iteration}. \quad (23)$$

Proof: See Appendix C.3. ■

The following comments are in order.

(i) Intuitively, Theorem 13 provides a range of values of σ to balance the tension between convergence rate and communication cost per iteration, resulting from too small or too large values of σ . The condition of the theorem can be satisfied, e.g., by choosing $\sigma = (1 + \lambda)/2$.

(ii) The term $d \cdot \log(1 + 1/(1-\lambda))$ in (21) represents the number of bits/iteration/agent under the additional restriction on σ , as postulated by Theorem 13: the faster the unquantized algorithm (i.e., the smaller λ), the less bits are required.

(iii) The second term $(1-\lambda)^{-1} \log(d/\varepsilon)$ represents the total number of iterations required to achieve ε accuracy; remarkably, these are the same (in a \mathcal{O} -sense) as the unquantized algorithm, with $\log(d)$ capturing the gap of the initial point from the fixed point. As expected, the number of iterations increases as the unquantized algorithm slows down (λ increases), the dimension d increases, and/or the target error ε decreases.

Nest, we customize Theorem 13 to some distributed algorithms within (Q-M).

A. Special Cases of (Q-M) Using the ANQ Rule

1) *GD Over Star-Networks:* Our first case study is the GD algorithm solving (P) with $r \equiv 0$ over star networks. The unquantized scheme reads

$$\mathbf{x}^{k+1} = \mathbf{x}^k - \frac{\gamma}{m} \sum_{i=1}^m \nabla f_i(\mathbf{x}^k), \quad (24)$$

with $\gamma \in (0, 2/L)$. This is an instance of the algorithmic framework (M); therefore, we can employ quantization using (Q-M). When the ANQ is employed, a direct application of Theorem 13 leads to the following communication complexity for the quantized GD (we termed the algorithm ANQ-GD-star).

Corollary 14 (ANQ-GD-star, see Appendix D.1): Given (P) over a star-network, where $r \equiv 0$ and F is L -smooth and μ -strongly convex, thus with condition number $\kappa = L/\mu$, and unique minimizer \mathbf{x}^* , consider the ANQ-GD-star with

stepsize $\gamma = 2/(\mu + L)$ and tuning for the ANQ as in Theorem 13. Then, the average number of bits/agent for $\|\mathbf{x}^k - \mathbf{x}^*\|_2^2 \leq \varepsilon$ is of order

$$\mathcal{O}\left(d \cdot \log(1 + \kappa) \kappa \log(d/\varepsilon)\right), \quad (25)$$

achieved in

$$\mathcal{O}\left(\kappa \log(d/\varepsilon)\right) \text{ iterations} \quad (26)$$

and with

$$\mathcal{O}\left(d \cdot \log(1 + \kappa)\right) \text{ bits/agent/iteration}. \quad (27)$$

This behavior compares favorably with $\mathcal{O}(d\kappa \log(d(1 + \kappa)) \log(d/\varepsilon))$ bits/agent obtained in [22] using a deterministic quantizer, with fixed number of bits among agents and iterations. It also matches $\mathcal{O}(d\kappa \log(1 + \kappa) \log(d/\varepsilon))$ bits/agent obtained in [53] using the same deterministic quantizer as in [22]. However, [53] uses a central coordinator to optimize the number of bits used by each agent at every iteration. Note that the lower complexity bound provided in [54] for $\kappa = 1$ reads: $\Omega(d \log(d/\varepsilon))$ bits/agent, confirming the tightness of our result.

2) *Distributed Algorithms Employing Gradient Correction:* Our second example deals with distributed algorithms solving (P) (possibly, with $r \not\equiv 0$) over mesh networks. With a slight abuse of notation, below we denote by κ , L and μ the condition number, the smoothness constant and the strong convexity constant of each f_i , respectively. We consider the most popular schemes, employing gradient correction in the optimization direction—see Appendix D for a description of these algorithms. We denote by $\rho_2 \triangleq \rho_2(\mathbf{W})$ the second largest eigenvalue of the gossip matrix \mathbf{W} used in these algorithms (note that $\rho_2 = 0$ for star-networks or fully-connected graphs). The communication complexity of these algorithms when using the ANQ is summarized next—they follow readily from Theorem 13 and the convergence results of the unquantized algorithms in [50].

Corollary 15 (ANQ-(Prox-)EXTRA, ANQ-(Prox-)NIDS, and ANQ-NIDS over mesh networks, see Appendices D.2, D.3, and D.4): Consider Problem (P) over mesh networks, and the ANQ-(Prox-)EXTRA and ANQ-(Prox-)NIDS algorithms with stepsize $\gamma = 2/(L + \mu)$ and the ANQ-NIDS algorithm with stepsize $\gamma = 1/L$, where the ANQ is tuned as in Theorem 13. Then, the average number of bits/agent for $\|\mathbf{x}^k - \mathbf{x}^*\|_2^2 \leq \varepsilon$ is of the order of

$$\mathcal{O}\left(d \max\left\{\kappa, \frac{1}{1-\rho_2}\right\} \log\left(\max\left\{1 + \kappa, \frac{1}{1-\rho_2}\right\}\right) \log(d/\varepsilon)\right),$$

achieved in

$$\mathcal{O}\left(\max\left\{\kappa, \frac{1}{1-\rho_2}\right\} \log(d/\varepsilon)\right) \text{ iterations}$$

and with

$$\mathcal{O}\left(d \log\left(\max\left\{1 + \kappa, \frac{1}{1-\rho_2}\right\}\right)\right) \text{ bits/agent/iteration}.$$

Corollary 16 (ANQ-(Prox-)NEXT and ANQ-(Prox-)DIGing Over Mesh Networks, See Appendices D.5-D.6): Consider

Problem (P) over mesh networks as described above, and the ANQ-(Prox-)NEXT and ANQ-(Prox-)DIGing algorithms with stepsize $\gamma = 2/(L + \mu)$, where the ANQ is tuned as in Theorem 13. Then, the average number of bits/agent for $\|\mathbf{x}^k - \mathbf{x}^*\|_2^2 \leq \varepsilon$ is of the order of

$$\mathcal{O}\left(d \max\left\{\kappa, \frac{1}{(1-\rho_2)^2}\right\} \log\left(\max\left\{1+\kappa, \frac{1}{(1-\rho_2)^2}\right\}\right) \log(d/\varepsilon)\right),$$

achieved in

$$\mathcal{O}\left(\max\left\{\kappa, \frac{1}{(1-\rho_2)^2}\right\} \log(d/\varepsilon)\right) \text{ iterations}$$

and with

$$\mathcal{O}\left(d \log\left(\max\left\{1+\kappa, \frac{1}{(1-\rho_2)^2}\right\}\right)\right) \text{ bits/agent/iteration.}$$

Corollary 17 (ANQ-Primal-Dual over Mesh Networks, see Appendix D.7): Consider Problem (P) with $r \equiv 0$ over mesh networks as described above, and the ANQ-Primal-Dual algorithm with stepsize $\gamma = \frac{2L\mu}{\mu\rho_{m-1}(\mathbf{L}) + L\rho_1(\mathbf{L})}$, where the ANQ is tuned as in Theorem 13. Then, the average number of bits/agent for $\|\mathbf{x}^k - \mathbf{x}^*\|_2^2 \leq \varepsilon$ is of the order of⁵

$$\mathcal{O}\left(d \frac{\kappa}{1-\rho_2} \log\left(1 + \frac{\kappa}{1-\rho_2}\right) \log(d/\varepsilon)\right),$$

achieved in

$$\mathcal{O}\left(\frac{\kappa}{1-\rho_2} \cdot \log(d/\varepsilon)\right) \text{ iterations}$$

and with

$$\mathcal{O}\left(d \cdot \log\left(1 + \frac{\kappa}{1-\rho_2}\right)\right) \text{ bits/agent/iteration.} \quad (28)$$

It is interesting to compare the ANQ-Primal-Dual (Corollary 17) with [24], which applies quantization to a Primal-Dual scheme using the widely adopted compression rule. The number of iterations required for the Option-D scheme (the best performing one) in [24] to achieve ε -accuracy, using the rand- K or dit- K compression methods (see [24, Example 4]), is⁶

$$\mathcal{O}\left(\frac{d}{K} \frac{\kappa}{1-\rho_2} \log(d/\varepsilon)\right), \quad (29)$$

Therefore, to achieve the same number of iterations as ANQ-Primal-Dual (in a \mathcal{O} -sense), rand- K in [24] requires $K = \Theta(d)$, whereas dit- K in [24] requires $S = 2^{K-1} - 1 = \Theta(\sqrt{d})$,

⁵Note that $\frac{\rho_1(\mathbf{L})}{\rho_{m-1}(\mathbf{L})} = \mathcal{O}((1-\rho_2)^{-1})$.

⁶The asymptotic result in [24, Theorem 5] maps to our setting with the following modifications: ρ (the ratio between the largest and the second smallest eigenvalues of \mathbf{L}) and ρ_∞ (the ratio of the largest weight and the second smallest eigenvalue of \mathbf{L}) defined therein are both of order $\mathcal{O}(1/(1-\rho_2))$; when adopting the setting on the initial error from our Assumption 11 and the problem setting of (P), the parameter m therein becomes 1 and the numerator in the log term becomes d .

with a communication cost of (in bits/agent/iteration)

$$\begin{cases} \mathcal{O}(d + B) & (\text{dit-}K), \\ \mathcal{O}(d(B + \log d)) & (\text{rand-}K), \end{cases} \quad (30)$$

where B is the number of bits used to encode each scalar *with negligible loss in precision*.

Comparing the communication costs (28) and (30), the following remarks are in order: **(i)** If we neglect the B -dependence, our scheme has the same scaling behavior as dit- K and better scaling behavior than rand- K , as the problem dimension d increases. **(ii)** It is not clear how the number of bits B should be chosen to encode scalars in [24] as a function of the system parameters, in a \mathcal{O} -sense, in order to *make loss in precision truly negligible* (a condition required by the convergence analysis therein). On the other hand, our communication cost analysis, which focuses on quantization below machine precision, provides an explicit answer to this question.

VI. NUMERICAL RESULTS

In this section, we validate numerically our theoretical findings and compare different distributed algorithms using quantization. We consider two instances of (P): a linear regression and a logistic regression problem, both with F strongly convex. The communication network is modeled as an undirected graph of $m = 20$ agents, generated by the Erdos-Renyi model with edge activation probability of 0.6. We measure performance of the algorithms using the mean square error MSE^k and the network communication cost $\text{C}_{\text{cm}}(\varepsilon)$ incurred over the network to reach $\text{MSE}^k \leq \varepsilon$, defined as

$$\text{MSE}^k \triangleq \frac{\sum_{i=1}^m \|\mathbf{x}_i^k - \mathbf{x}^*\|_2^2}{m \|\mathbf{x}^*\|_2^2}, \quad \text{C}_{\text{cm}}(\varepsilon) \triangleq \sum_{k=0}^{k_\varepsilon} \sum_{s=1}^R \sum_{i=1}^m b_i^{k,s}, \quad (31)$$

where $k_\varepsilon \triangleq \min_{k \geq 0} \text{MSE}^k \leq \varepsilon$ is the number of iterations required to achieve ε -accuracy, \mathbf{x}^* is the optimal solution of (P), and $b_i^{k,s}$ is the number of bits used by the quantizer to encode the s th transmitted signal by agent i at iteration k .

A. Linear Regression Problem

1) Problem Setting: Consider the following linear regression problem over mesh networks:

$$f_i(\mathbf{x}) = \frac{1}{2} \|\mathbf{U}_i \mathbf{x} - \mathbf{v}_i\|_2^2 + \frac{0.01}{2} \|\mathbf{x}\|_2^2 \text{ and } r(\mathbf{x}) = \alpha \|\mathbf{x}\|_1, \quad (32)$$

where $\mathbf{U}_i \in \mathbb{R}^{20 \times 40}$ and $\mathbf{v}_i \in \mathbb{R}^{20 \times 1}$ are the feature vector and observation measurements, respectively, accessible only by agent i . The matrix \mathbf{U}_i is generated independently across agents, according to the model in [55], namely: $[\mathbf{U}_i]_{:,1} \sim \mathcal{N}(\mathbf{0}, \frac{1}{1-\beta^2} \mathbf{I})$ (first column), and for the other columns $q > 1$, $[\mathbf{U}_i]_{:,q} | [\mathbf{U}_i]_{:,q-1} \sim \mathcal{N}(\beta [\mathbf{U}_i]_{:,q-1}, \mathbf{I})$, where $\beta = 0.3$. In this way, each row of \mathbf{U}_i is a Gaussian random vector with zero mean and covariance depending on β : larger β generates more ill-conditioned covariance matrices. Then, letting $\mathbf{x}_0 \in \mathbb{R}^{40}$ be the ground truth vector, generated as a sparse vector with 70%




Fig. 2. Linear regression problem (32): MSE versus iterations. Solid curves and markers refer to algorithms implemented using machine precision and quantized communications, respectively. In the legend, for each scheme, we report the number of bits transmitted, per agent, per dimension, per iteration.

zero entries, and i.i.d. nonzero entries drawn from $\mathcal{N}(0, 1)$, we generate \mathbf{v}_i as $\mathbf{v}_i | (\mathbf{U}, \mathbf{x}_0) \sim \mathcal{N}(\mathbf{U}\mathbf{x}_0, 0.04\mathbf{I})$. We use μ, L for the strong convexity and smoothness parameters of each f_i , respectively.

We test several distributed algorithms considering either smooth ($\alpha = 0$) or non-smooth ($\alpha > 0$) instances of the least square problem (32). In fact, most of the existing quantization schemes are applicable only to smooth optimization problems. The free parameters of these algorithms are optimized based on the recommendations in the original papers, while optimizing numerical convergence with the smallest number of bits, unless otherwise stated; the weight matrix $\tilde{\mathbf{W}}$ used to mix the received signals is constructed according to the Metropolis-Hastings rule [56]; the number of bits transmitted by each scheme as reported in the figures is per agent, per dimension, per iteration. For each quantized algorithm, we choose $\sigma = 0.99 \cdot \lambda + 0.01$, where $\lambda = (\text{MSE}^{100}/\text{MSE}^{50})^{0.01}$ is the numerical estimate of the convergence rate of its unquantized counterpart (i.e., implemented at machine precision). In the simulations, all algorithms except those using LPQ are evaluated with 1 realization since they are deterministic algorithms. Those using the probabilistic quantizer LPQ, i.e., LEAD [25] and COLD [39], are averaged over 10 realizations, with fixed $\mathbf{U}_i, \mathbf{V}_i, \mathbf{x}^0$, and network topology.

2) *Smooth Linear Regression (Fig. 2a):* We begin by considering the smooth linear regression problem. We consider the following benchmark schemes, using machine precision (64-bit representation for each scalar):

- 1) Primal-Dual [42] with step-size $\gamma = 2L\mu/(\mu\rho_{m-1}(\mathbf{L}) + L\rho_1(\mathbf{L}))$ ([22, Proposition 2]), where \mathbf{L} is the graph Laplacian matrix associated with the graph.
- 2) NEXT [40], [41], [46] with step-size $\gamma = 0.0029$, manually tuned for fastest practical convergence.

3) NIDS [48], [50] with step-size $\gamma = \frac{2}{L+\mu}$ and mixing matrix $\mathbf{W} = [(1+\nu)\mathbf{I} + (1-\nu)\tilde{\mathbf{W}}]/2$, with $\nu = 0.001$.

In addition, we consider the following algorithms, that implement the above benchmark schemes using quantized communications:

- 4) Q-Dual [22] and Q-NEXT [21];
- 5) ANQ-Dual: this is the Primal-Dual algorithm [42] equipped with the proposed deterministic ANQ (see Appendix D.7), with $\eta^0 = 0.01$ and $\omega = \bar{\omega}/2$ [recall that $\bar{\omega}$ is defined in (10)];
- 6) ANQ-NEXT: this is the NEXT algorithm [40], [41], [46] quantized using the deterministic ANQ with $\eta^0 = 0.029$ and $\omega = \bar{\omega}/2$;
- 7) ANQ-NIDS: this is an instance of the NIDS algorithm [48], [50] equipped with the deterministic ANQ (see Appendix D.4) with parameters $\eta^0 = 0.1$ and $\omega = \bar{\omega}/2$
- 8) LEAD [25] and COLD [39], both implemented using the low-precision quantization (LPQ) [36]: to transmit a signal \mathbf{x} , the amplitude $\|\mathbf{x}\|_2$ is encoded at machine precision (64 bits), and 3 bits are adopted to encode each normalized element $x_i/\|\mathbf{x}\|_2$.

In Fig. 2a, we plot the MSE versus iteration index k . Remarkably, all algorithms, when equipped with the proposed ANQ, incur a negligible loss of convergence speed with respect to their machine precision counterpart. Comparing ANQ with the compression-based distributed algorithms LEAD [25] and COLD [39], we infer that the proposed ANQ-NIDS is faster while using less bits. Comparing ANQ with the state-of-the-art quantized algorithms, we notice that ANQ is more communication-efficient than Q-NEXT and Q-Dual, which instead use deterministic uniform quantizers with shrinking range: ANQ-NEXT (11.62 bits) and ANQ-Dual (2.65 bits) use less bits per iteration than Q-NEXT (28 bits)




Fig. 3. Logistic regression problem (33): MSE versus iterations. Solid curves and markers refer to algorithms implemented using machine precision and quantized communications, respectively. In the legend, for each scheme, we report the number of bits transmitted, per agent, per dimension, per iteration.

and Q-Dual (5 bits), respectively, while converging faster. Note that, with the parameters chosen as recommended in Theorem 6, all ANQ-based algorithms shown in the figure have convergence guarantees, while Q-Dual and Q-NEXT do not in the simulated setting. In fact, Q-Dual and Q-NEXT use 5 and 28 bits, respectively, which fall below the minimum number of bits that guarantee linear convergence, calculated to be 13 from [22, Theorem 1] and 78 from [21, Theorem 4], respectively.

3) *Non-Smooth Linear Regression* (Fig. 2b): We now move to the non-smooth instance of (32), with $\alpha = 10^{-4}$. To our knowledge, there is no existing quantized algorithms solving such instance of (32). Hence, we tested the following ANQ-based quantized algorithms:

- 1) ANQ-Prox-EXTRA: this is an instance of the Prox-EXTRA algorithm [50] equipped with the deterministic ANQ (see Appendix D.2) with parameters $\eta^0 = 6.67 \times 10^{-4}$ and $\omega = \bar{\omega}/2$;
- 2) ANQ-Prox-NEXT: this is the Prox-NEXT algorithm [50] equipped with the deterministic ANQ (see Appendix D.5) with parameters $\eta^0 = 2.34 \times 10^{-3}$ and $\omega = \bar{\omega}/2$;
- 3) ANQ-Prox-DIGing: this is the Prox-DIGing algorithm [50] equipped with the deterministic ANQ (see Appendix D.6) with parameters $\eta^0 = 3.05 \times 10^{-3}$ and $\omega = \bar{\omega}/2$;
- 4) ANQ-Prox-NIDS: this is the Prox-NIDS algorithm in [50] equipped with the deterministic ANQ (see Appendix D.3) with parameters $\eta^0 = 7.7 \times 10^{-4}$ and $\omega = \bar{\omega}/2$.

As benchmark, we also included their unquantized counterparts, implemented at machine precision; in all these schemes, we used the weight matrix $\mathbf{W} = [(1+\nu)\mathbf{I} + (1-\nu)\tilde{\mathbf{W}}]/2$ with $\nu = 0.001$; the step-size is chosen according to [50], namely:

$\gamma = \frac{2\rho_m(\mathbf{W})}{L + \mu\rho_m(\mathbf{W})}$ for Prox-EXTRA, $\gamma = \frac{2\rho_m(\mathbf{W}^2)}{L + \mu\rho_m(\mathbf{W}^2)}$ for Prox-DIGing, and $\gamma = \frac{2}{L + \mu}$ for Prox-NEXT and Prox-NIDS.

Fig. 2b plots the MSE achieved by all the algorithms versus the iteration index. As predicted, all four quantized schemes converge linearly. Remarkably, all of the ANQ-equipped algorithms incur a negligible loss of convergence speed with respect to their machine precision counterparts, while using a fraction of the communication budget – only 14 bits per agent/dimension/iteration.

B. Logistic Regression

We now consider the distributed logistic regression problem using the MNIST dataset [57]. This is an instance of (P) with

$$f_i(\mathbf{x}) = \frac{0.01}{2} \|\mathbf{x}\|_2^2 + \frac{1}{3000} \sum_{p=1}^{3000} \ln \left(1 + \exp \left(-v_{i,p} \mathbf{u}_{i,p}^\top \mathbf{x} \right) \right),$$

$$\text{and } r(\mathbf{x}) = \alpha \|\mathbf{x}\|_1, \quad (33)$$

where $\mathbf{u}_{i,p} \in \mathbb{R}^{784 \times 1}$ and $v_{i,p} \in \{-1, 1\}$ are the feature vector and labels, respectively, only accessible by agent i . Here we implement the one-vs.-all scheme, i.e., the goal is to distinguish the data of label '0' from others. To generate $\mathbf{u}_{i,p}$, we first flatten each picture of size 28×28 in MNIST into a real feature vector of length $28 \times 28 = 784$, and then normalize it to unit l_2 norm. We then allocate equal number of feature vectors and labels to each agent. In the simulations, all algorithms except those using LPQ are evaluated with 1 realization since they are deterministic algorithms. Those using the probabilistic quantizer LPQ, i.e., LEAD [25] and COLD [39], are averaged over 10 realizations, with fixed feature/label allocations and network topology.

1) *Smooth Logistic Regression* (Fig. 3a): We begin by considering the smooth logistic regression problem (33), with $\alpha = 0$. We tested the same algorithms (with the same tuning) as described in Sec. VI-A for the smooth linear regression




Fig. 4. MSE versus iterations for different quantization rules applied to NIDS [48], [50].

problem. In Fig. 3a, we plot the MSE versus iteration index k . Consistently with the results in Fig. 2a, we notice the following facts. ANQ-NIDS achieves the fastest convergence, followed by ANQ-NEXT, ANQ-Dual, Q-Dual and Q-NEXT. Comparing our quantization method with existing ones on the same unquantized algorithm, we notice that the proposed ANQ is more communication-efficient than Q-NEXT and Q-Dual: ANQ-NEXT (6.28 bits) and ANQ-Dual (2.7 bits) use less bits per iteration than Q-NEXT (36 bits) and Q-Dual (5 bits), respectively, while at the same time converging faster. Comparing with the compression-based algorithms LEAD and COLD, it is shown that ANQ-NIDS achieves better convergence rate with less bits.

2) *Non-Smooth Logistic Regression (Fig. 3b):* We now consider the non-smooth instance of the logistic regression problem (33), with $\alpha = 10^{-4}$. We tested the same algorithms (with the same tuning) as described in Sec. VI-A for the non-smooth linear regression problem. Fig. 3b plots the MSE achieved by all the algorithms versus iterations k . The results confirm the trends already commented in Fig. 2b.

C. Comparison of Quantization Rules

From the above results, it is clear that ANQ-NIDS outperforms existing algorithms using dynamic quantization (DYQ)—including Q-Dual and Q-NEXT—as well as those employing low-precision quantization (LPQ)—such as LEAD and COLD. This advantage may be due to the *black-box* nature of ANQ: it can be applied to a variety of distributed algorithms, including those known in the literature to be the fastest ones. This is a significant advantage over ad-hoc quantization schemes, which are limited in their applicability to the specific algorithms they are designed for. Therefore, an interesting question is whether the performance superiority comes solely from the underlying (unquantized) algorithm, i.e., NIDS, or also from the quantization rule, i.e., ANQ. To answer this question, we compare the above three quantization rules, DYQ, LPQ, and ANQ on the same distributed algorithm, NIDS [48], [50].




Fig. 5. Network communication cost evaluation on non-smooth linear regression problem with $\alpha = 10^{-4}$ versus d , where λ remains fixed for all d in each algorithm.

Fig. 4a and Fig. 4b show the MSE versus iterations for ANQ-NIDS, LPQ-NIDS, and DYQ-NIDS solving the smooth linear regression and smooth logistic regression problems (32) and (33) ($r \equiv 0$), respectively. The parameters for DYQ-NIDS and LPQ-NIDS are selected to closely match the performance of NIDS with machine precision while using the smallest number of bits, whereas those for ANQ-NIDS are selected according to our analysis. Both figures show that ANQ-NIDS consistently requires less bits than the other quantization schemes. More precisely, ANQ-NIDS uses 25% less bits than DYQ-NIDS and 44% less than LPQ-NIDS in the linear regression case, and 50% less bits than DYQ-NIDS and 27% less than LPQ-NIDS in the logistic regression problem.

D. Communication Cost

We now study the effect of the dimension d on the communication cost for different algorithms solving the non-smooth




Fig. 6. Network communication cost evaluation of ANQ-NIDS on smooth linear regression problem.

linear regression problem (32), with $\alpha = 10^{-4}$. Note that the rate of a machine precision algorithm λ depends on both the weight matrix and the condition number $\kappa = L/\mu$, which depends itself on d . We chose the coefficient of the l_2 regularizer so as to make κ and thus λ remain fixed across different d . The rest of the settings are the same as in Fig. 2b. Fig. 5 plots the network communication cost $C_{cm}(\varepsilon)$ versus d as defined in (31), required to reach a target MSE-accuracy $\varepsilon = 10^{-8}$. We observe that C_{cm} scales roughly linearly with respect to the dimension for all algorithms, which is consistent with Theorem 13.

Finally, we investigate numerically the effect of σ and ω on the network communication cost as defined in (31), for a target MSE-accuracy $\varepsilon = 10^{-14}$. We consider the ANQ-NIDS algorithm with $\eta^0 = 0.001$, solving the smooth linear regression problem (32), with $\alpha = 0$. Fig. 6a plots the network communication cost versus σ with $\omega = \bar{\omega}/2$. Note that this figure justifies the discussion in Sec. V that σ should be chosen away from λ and 1 in order to save on communication cost. Fig. 6a plots the communication cost (31) versus ω with $\sigma = 0.99 \times \lambda + 0.01$. It can be seen that, by optimizing the compression rate ω , a saving of 15% in communication cost can be obtained over a quantization scheme that employs no compression ($\omega = 0$). This observation numerically supports our BC-rule, which generalizes the deterministic/probabilistic quantizers that have no compression term.

VII. CONCLUSION

In this paper, we propose a black-box model and unified convergence analysis for a general class of linearly convergent algorithms subject to quantized communications. The new black-box model encompasses composite optimization problems and distributed algorithms using historical information (e.g., EXTRA [43] and NEXT [40]), thus generalizing existing algorithmic frameworks, which are not applicable

to these settings. To enable quantization (below machine precision), we propose a novel *biased compression (BC-)rule* that preserves linear convergence of distributed algorithms while using a *finite* number of bits in each communication. As special instance of the BC-rule, we also proposed a new random or deterministic quantizer, the ANQ, coupled with a communication-efficient encoding scheme. We analyzed the communication cost of a gamut of distributed algorithms equipped with the ANQ (in a unified fashion), showing favorable performance analytically and numerically, both in terms of convergence rate and communication cost, with respect to state-of-the-art quantization rules (including uniform and compression-based ones) and ad-hoc distributed algorithms.

APPENDIX A LINEAR CONVERGE UNDER QUANTIZATION

In this appendix, we prove Theorem 6. We begin by introducing some preliminary results, whose proofs are deferred to Appendix A.3. Throughout this section, we make the blanket assumption that the conditions in Theorem 6 are satisfied. In particular, $\sigma \in (\lambda, 1)$ and $\omega \in [0, \bar{\omega}(\sigma)]$, with $\bar{\omega}(\sigma)$ defined in (10). Due to the possibly random nature of the quantizer, $\{\mathbf{z}^k, \mathbf{c}^{k,s}, \hat{\mathbf{c}}^{k,s}\}_{k \geq 0, s \in [R]}$ is a stochastic process defined on a proper probability space; we denote by $\mathcal{F}^{k,s}$ the σ -algebra generated by $\{\mathbf{z}^{k'}, \mathbf{c}^{k',s'}, \hat{\mathbf{c}}^{k',s'}\}_{k' < k, s' \in [R]} \cup \{\mathbf{z}^k, \mathbf{c}^{k,s}, \hat{\mathbf{c}}^{k,s-1}\}_{s' \leq s}$ ($\hat{\mathbf{c}}^{k,s}$ excluded).

1. Preliminaries

The idea of the proof is to show by induction that both the optimization error $\|\mathbf{z}^k - \mathbf{z}^\infty\|$ and the input to the quantizer, $\|\mathbf{c}^{k,s} - \hat{\mathbf{c}}^{k-1,s}\|_2$, are linearly convergent (in expectation) at rate σ , i.e.,

$$\sqrt{\mathbb{E}[\|\mathbf{z}^k - \mathbf{z}^\infty\|^2]} \leq V_0 \cdot (\sigma)^k, \quad (34)$$

$$\sqrt{\mathbb{E}[\|\mathbf{c}^{k,s} - \hat{\mathbf{c}}^{k-1,s}\|_2^2]} \leq F^s \cdot (\sigma)^k, \quad \forall s \in [R] \cup \{0\}, \quad (35)$$

where $F^0 = 0$, and $V_0, F^s, s \in [R]$ satisfy

$$V_0 \geq \max \left\{ \|\mathbf{z}^0 - \mathbf{z}^\infty\|, \frac{\sqrt{md}R\eta^0 + \omega\mathbf{F}^\top \mathbf{1}}{\sigma - \lambda} \tilde{L}_A \right\}, \quad (36)$$

$$F^s \geq \max \left\{ L_Z c^* + L_C \sqrt{md} \eta^0 + L_C (1 + \omega) F^{s-1}, \frac{\sqrt{md} \eta^0 (1 + L_C \sigma) + L_C \sigma (1 + \omega) F^{s-1} + L_Z (1 + \sigma) V_0}{\sigma - \omega} \right\}, \quad (37)$$

$\forall s \in [R]$, and we have defined $\mathbf{F} \triangleq (F^s)_{s \in [R]}$,

$$c^* \triangleq \frac{1}{L_Z} \max_{s=[R]} \|\mathcal{C}^s(\mathbf{z}^0, \mathbf{0})\|_2, \quad \tilde{L}_A \triangleq L_A \sum_{s=0}^{R-1} (L_C)^s. \quad (38)$$

The existence of such V_0 and $F^s, s \in [R]$, is proved in the following lemma.

Lemma 18: Let $\omega \in [0, \bar{\omega}(\sigma))$. Then, (36) and (37) are satisfied by

$$V_0 = \max \left\{ c^*, \|\mathbf{z}^0 - \mathbf{z}^\infty\| \right\} + \frac{L_A \psi \sqrt{md} R^2 \eta^0}{\sigma - \lambda} \cdot \frac{1 + \frac{R\omega}{\sigma} [(1 + L_C \sigma) \psi - 1]}{1 - \omega/\bar{\omega}}, \quad (39)$$

$$F^s = \frac{\sqrt{md} \eta^0 (1 + L_C \sigma) + 2L_Z V_0 (\sigma, \omega, \eta^0)}{\sigma - \omega} \times \sum_{s'=0}^{s-1} \left(\frac{2L_C}{1 - \omega/\sigma} \right)^{s'}, \quad s \in [R], \quad (40)$$

where $\psi \triangleq \max\{1, (2L_C)^{R-1}\}$.

Since the effect of $\|\mathbf{c}^{k,s} - \hat{\mathbf{c}}^{k-1,s}\|_2$ on $\|\mathbf{z}^k - \mathbf{z}^\infty\|$ is through quantization, we need the following bound on the quantization error (its proof follows readily by the Cauchy-Schwarz inequality).

Lemma 19: Let the quantizer \mathcal{Q}_i used by agent $i \in [m]$ satisfy the BC-rule (5) with bias $\eta \geq 0$ and compression rate $\omega \in [0, \bar{\omega}(\sigma))$. Then the following holds for the stack $\mathcal{Q} \triangleq [\mathcal{Q}_1, \dots, \mathcal{Q}_m]^\top$:

$$\sqrt{\mathbb{E}[\|\mathcal{Q}(\mathbf{x}) - \mathbf{x}\|_2^2 | \mathbf{x}]} \leq \sqrt{md} \eta + \omega \|\mathbf{x}\|_2, \quad (41)$$

$$\forall \mathbf{x} = [\mathbf{x}_1^\top, \dots, \mathbf{x}_m^\top]^\top \in \mathbb{R}^{md}.$$

A direct application of Lemma 19 leads to the following bound on the quantizer's input, which we use recurrently in the proofs:

$$\begin{aligned} & \sqrt{\mathbb{E}[\|\hat{\mathbf{c}}^{k,s} - \mathbf{c}^{k,s}\|_2^2 | \mathcal{F}^{k,s}]} \\ & \stackrel{(Q-M)}{=} \sqrt{\mathbb{E}[\|\mathcal{Q}^k(\mathbf{c}^{k,s} - \hat{\mathbf{c}}^{k-1,s}) - (\mathbf{c}^{k,s} - \hat{\mathbf{c}}^{k-1,s})\|_2^2 | \mathcal{F}^{k,s}]} \\ & \stackrel{(41)}{\leq} \sqrt{md} \eta^0 \cdot (\sigma)^k + \omega \|\mathbf{c}^{k,s} - \hat{\mathbf{c}}^{k-1,s}\|_2, \quad a.s., \end{aligned} \quad (42)$$

for all $s \in [R]$ and $k = 0, 1, \dots$, where we used $\eta^k = \eta^0 \cdot (\sigma)^k$.

The following lemma bounds the distortion introduced by quantization in one iteration of (Q-M).

Lemma 20: There holds: for all $k = 0, 1, \dots$,

$$\begin{aligned} & \sqrt{\mathbb{E}[\|\mathbf{z}^{k+1} - \mathbf{z}^\infty\|^2]} \leq \lambda \sqrt{\mathbb{E}[\|\mathbf{z}^k - \mathbf{z}^\infty\|^2]} \\ & + \tilde{L}_A \sqrt{md} R \eta^0 \cdot (\sigma)^k + \tilde{L}_A \omega \sum_{s=1}^R \sqrt{\mathbb{E}[\|\mathbf{c}^{k,s} - \hat{\mathbf{c}}^{k-1,s}\|_2^2]}, \end{aligned}$$

a.s., where \tilde{L}_A is defined in (38).

We conclude this section of preliminaries with the following useful result.

Lemma 21: Let $\{X_t : t \in [T]\} \subset \mathbb{R}$ be a collection of random variables. Then,

$$\sqrt{\mathbb{E} \left[\left(\sum_{t=1}^T X_t \right)^2 \right]} \leq \sum_{t=1}^T \sqrt{\mathbb{E}[X_t^2]}.$$

Proof: It can be proved by developing the square within the expectation on the left hand side expression, and by using $\mathbb{E}[X_t X_u] \leq \sqrt{\mathbb{E}[X_t^2]} \sqrt{\mathbb{E}[X_u^2]}$. \square

2. Proof of Theorem 6

We prove (34) and (35) by induction. Let V_0 and $F^s, s \in [R]$, satisfy (36) and (37). Since $\|\mathbf{z}^0 - \mathbf{z}^\infty\| \leq V_0$ (see (36)) and $\mathbf{c}^{0,0} = \hat{\mathbf{c}}^{-1,0} = \mathbf{0}$, (34) holds for $k = 0$ and (35) holds trivially for $k = 0$ and $s = 0$. We now use induction to prove that (35) holds for $k = 0$ and $s \in [R]$. Assume that (35) holds for $k = 0$ and $s < R$. Then, it follows that

$$\begin{aligned} & \|\mathbf{c}^{0,s+1} - \hat{\mathbf{c}}^{-1,s+1}\|_2 = \|\mathbf{c}^{0,s+1}\|_2 \\ & \stackrel{(a)}{\leq} \|\mathcal{C}^{s+1}(\mathbf{z}^0, \mathbf{0})\|_2 + \|\mathcal{C}^{s+1}(\mathbf{z}^0, \hat{\mathbf{c}}^{0,s}) - \mathcal{C}^{s+1}(\mathbf{z}^0, \mathbf{c}^{0,s})\|_2 \\ & \quad + \|\mathcal{C}^{s+1}(\mathbf{z}^0, \mathbf{c}^{0,s}) - \mathcal{C}^{s+1}(\mathbf{z}^0, \mathbf{0})\|_2 \\ & \stackrel{(8),(38)}{\leq} L_Z c^* + L_C \|\hat{\mathbf{c}}^{0,s} - \mathbf{c}^{0,s}\|_2 + L_C \|\mathbf{c}^{0,s} - \hat{\mathbf{c}}^{-1,s}\|_2, \quad a.s., \end{aligned}$$

where in (a) we used the triangle inequality and $\mathbf{c}^{0,s+1} = \mathcal{C}^{s+1}(\mathbf{z}^0, \hat{\mathbf{c}}^{0,s})$. Taking the conditional expectation on both sides and using Lemma 21 yield

$$\begin{aligned} & \sqrt{\mathbb{E}[\|\mathbf{c}^{0,s+1} - \hat{\mathbf{c}}^{-1,s+1}\|_2^2 | \mathcal{F}^{0,s}]} \\ & \leq L_Z c^* + L_C \sqrt{\mathbb{E}[\|\hat{\mathbf{c}}^{0,s} - \mathbf{c}^{0,s}\|_2^2 | \mathcal{F}^{0,s}]} + L_C \|\mathbf{c}^{0,s} - \hat{\mathbf{c}}^{-1,s}\|_2 \\ & \stackrel{(42)}{\leq} L_Z c^* + L_C \sqrt{md} \eta^0 + L_C (1 + \omega) \|\mathbf{c}^{0,s} - \hat{\mathbf{c}}^{-1,s}\|_2, \quad a.s.. \end{aligned}$$

Taking the unconditional expectation on both sides and using again Lemma 21 yield

$$\begin{aligned} & \sqrt{\mathbb{E}[\|\mathbf{c}^{0,s+1} - \hat{\mathbf{c}}^{-1,s+1}\|_2^2]} \\ & \leq L_Z c^* + L_C \sqrt{md} \eta^0 + L_C (1 + \omega) \sqrt{\mathbb{E}[\|\mathbf{c}^{0,s} - \hat{\mathbf{c}}^{-1,s}\|_2^2]} \\ & \stackrel{(35)}{\leq} L_Z c^* + L_C \sqrt{md} \eta^0 + L_C (1 + \omega) F^s \stackrel{(37)}{\leq} F^{s+1}, \end{aligned}$$

which completes the induction proof of (35) for $k = 0$ and $s \in [R]$.

Now, let us assume that (34) and (35) hold for a generic $k = 0, 1, \dots$; we prove that they hold at $k + 1$. We begin with (34). Invoking Lemmas 20, 21, and using the induction hypotheses (34) and (35) at k , yield

$$\begin{aligned} & \sqrt{\mathbb{E}[\|\mathbf{z}^{k+1} - \mathbf{z}^\infty\|^2]} \leq \lambda V_0 \cdot (\sigma)^k \\ & + \tilde{L}_A (\sqrt{md} R \eta^0 + \omega \mathbf{F}^\top \mathbf{1}) \cdot (\sigma)^k \leq V_0 \cdot (\sigma)^{k+1}, \end{aligned}$$

where the last inequality follows from the definition of V_0 in (36), which concludes the induction argument for (34).

We now prove that (35) holds for $k + 1$, by induction over $s \in [R]$. First, note that (35) holds trivially for $k + 1$ and

$s = 0$, since $\mathbf{c}^{k+1,0} = \hat{\mathbf{c}}^{k,0} = \mathbf{0}$. Now, assume that (35) holds at iteration $k + 1$ for $s < R$. Then,

$$\begin{aligned} & \|\mathbf{c}^{k+1,s+1} - \hat{\mathbf{c}}^{k,s+1}\|_2 \\ & \stackrel{(Q-M)}{=} \|\mathcal{C}^{s+1}(\mathbf{z}^{k+1}, \hat{\mathbf{c}}^{k+1,s}) - \mathcal{C}^{s+1}(\mathbf{z}^{k+1}, \mathbf{c}^{k+1,s}) \\ & \quad + \mathcal{C}^{s+1}(\mathbf{z}^{k+1}, \mathbf{c}^{k+1,s}) - \mathcal{C}^{s+1}(\mathbf{z}^k, \hat{\mathbf{c}}^{k,s}) + \mathbf{c}^{k,s+1} - \hat{\mathbf{c}}^{k,s+1}\|_2 \\ & \stackrel{(a)}{\leq} \|\mathcal{C}^{s+1}(\mathbf{z}^{k+1}, \hat{\mathbf{c}}^{k+1,s}) - \mathcal{C}^{s+1}(\mathbf{z}^{k+1}, \mathbf{c}^{k+1,s})\|_2 \\ & \quad + \|\mathcal{C}^{s+1}(\mathbf{z}^{k+1}, \mathbf{c}^{k+1,s}) - \mathcal{C}^{s+1}(\mathbf{z}^k, \hat{\mathbf{c}}^{k,s})\|_2 \\ & \quad + \|\mathbf{c}^{k,s+1} - \hat{\mathbf{c}}^{k,s+1}\|_2 \stackrel{(8),(9)}{\leq} L_C \|\hat{\mathbf{c}}^{k+1,s} - \mathbf{c}^{k+1,s}\|_2 \\ & \quad + L_C \|\mathbf{c}^{k+1,s} - \hat{\mathbf{c}}^{k,s}\|_2 + L_Z \|\mathbf{z}^{k+1} - \mathbf{z}^\infty\|_2 \\ & \quad + L_Z \|\mathbf{z}^k - \mathbf{z}^\infty\|_2 + \|\hat{\mathbf{c}}^{k,s+1} - \mathbf{c}^{k,s+1}\|_2, \text{ a.s..} \end{aligned}$$

Then, taking the expectation conditional on $\mathcal{F}^{k+1,s}$ and invoking Lemma 21 and (42) to bound $\sqrt{\mathbb{E}[\|\hat{\mathbf{c}}^{k+1,s} - \mathbf{c}^{k+1,s}\|_2^2 | \mathcal{F}^{k+1,s}]}$, yield

$$\begin{aligned} & \sqrt{\mathbb{E}[\|\mathbf{c}^{k+1,s+1} - \hat{\mathbf{c}}^{k,s+1}\|_2^2 | \mathcal{F}^{k+1,s}]} \leq L_C \sqrt{md\eta^0} \cdot (\sigma)^{k+1} \\ & \quad + L_C(1+\omega) \|\mathbf{c}^{k+1,s} - \hat{\mathbf{c}}^{k,s}\|_2 + L_Z \|\mathbf{z}^{k+1} - \mathbf{z}^\infty\|_2 \\ & \quad + L_Z \|\mathbf{z}^k - \mathbf{z}^\infty\|_2 + \|\hat{\mathbf{c}}^{k,s+1} - \mathbf{c}^{k,s+1}\|_2, \text{ a.s..} \end{aligned}$$

Now, taking the expectation conditional on $\mathcal{F}^{k,s+1} \subseteq \mathcal{F}^{k+1,s}$, invoking Lemma 21, and (42) to bound $\sqrt{\mathbb{E}[\|\hat{\mathbf{c}}^{k,s+1} - \mathbf{c}^{k,s+1}\|_2^2 | \mathcal{F}^{k,s+1}]}$, yield

$$\begin{aligned} & \sqrt{\mathbb{E}[\|\mathbf{c}^{k+1,s+1} - \hat{\mathbf{c}}^{k,s+1}\|_2^2 | \mathcal{F}^{k,s+1}]} \leq \sqrt{md\eta^0} (1 + L_C\sigma) \cdot (\sigma)^k \\ & \quad + L_C(1+\omega) \sqrt{\mathbb{E}[\|\mathbf{c}^{k+1,s} - \hat{\mathbf{c}}^{k,s}\|_2^2 | \mathcal{F}^{k,s+1}]} \\ & \quad + L_Z \sqrt{\mathbb{E}[\|\mathbf{z}^{k+1} - \mathbf{z}^\infty\|_2^2 | \mathcal{F}^{k,s+1}]} \\ & \quad + L_Z \sqrt{\mathbb{E}[\|\mathbf{z}^k - \mathbf{z}^\infty\|_2^2 | \mathcal{F}^{k,s+1}]} + \omega \|\mathbf{c}^{k,s+1} - \hat{\mathbf{c}}^{k-1,s+1}\|_2, \end{aligned}$$

a.s.. Taking the unconditional expectation and invoking Lemma 21 again yield

$$\begin{aligned} & \sqrt{\mathbb{E}[\|\mathbf{c}^{k+1,s+1} - \hat{\mathbf{c}}^{k,s+1}\|_2^2]} \leq \sqrt{md\eta^0} (1 + L_C\sigma) \cdot (\sigma)^k \\ & \quad + L_C(1+\omega) \sqrt{\mathbb{E}[\|\mathbf{c}^{k+1,s} - \hat{\mathbf{c}}^{k,s}\|_2^2]} + L_Z \sqrt{\mathbb{E}[\|\mathbf{z}^{k+1} - \mathbf{z}^\infty\|_2^2]} \\ & \quad + L_Z \sqrt{\mathbb{E}[\|\mathbf{z}^k - \mathbf{z}^\infty\|_2^2]} + \omega \sqrt{\mathbb{E}[\|\mathbf{c}^{k,s+1} - \hat{\mathbf{c}}^{k-1,s+1}\|_2^2]} \\ & \stackrel{(a)}{\leq} \sqrt{md\eta^0} (1 + L_C\sigma) \cdot (\sigma)^k + L_C(1+\omega) F^s \cdot (\sigma)^{k+1} \\ & \quad + \omega F^{s+1} \cdot (\sigma)^k + L_Z(1+\sigma)V_0 \cdot (\sigma)^k \stackrel{(37)}{\leq} F^{s+1} \cdot (\sigma)^{k+1}, \end{aligned}$$

where in (a) we used the induction hypotheses (35) (applied to the second and last terms) and (34) (applied to the third and fourth terms). This proves the induction for (35), and the theorem.

3. Proof of Auxiliary Lemmas for Theorem 6

1) *Proof of Lemma 18:* It is not difficult to check that conditions (36) and (37) can be satisfied by choosing

$$\begin{aligned} V_0 & \geq \max \left\{ c^*, \|\mathbf{z}^0 - \mathbf{z}^\infty\|, \frac{\sqrt{md}R\eta^0 + \omega\mathbf{F}^\top\mathbf{1}}{\sigma - \lambda} \tilde{L}_A \right\}, \\ F^s & \geq \frac{\sqrt{md\eta^0}(1+L_C\sigma) + L_C\sigma(1+\omega)F^{s-1} + L_Z(1+\sigma)V_0}{\sigma - \omega}, \end{aligned}$$

$\forall s \in [R]$, where c^*, \tilde{L}_A are defined in (38). Moreover, since $\omega < \sigma < 1$, it is sufficient to choose

$$F^s = \frac{1}{\sigma} \frac{\sqrt{md\eta^0}(1+L_C\sigma) + 2L_C\sigma F^{s-1} + 2L_Z V_0}{1 - \omega/\sigma}, \quad \forall s \in [R].$$

Solving this expression recursively yields (40).

We now prove (39). We begin noting that F^s is a non-decreasing function of s , hence $F^s \leq F^R$. Moreover, F^R is an affine function of V_0 . Using the facts that $(\frac{2L_C}{1-\omega/\sigma})^s \leq \psi(1-\omega/\sigma)^{-s}$, where $\psi \triangleq \max\{1, (2L_C)^{R-1}\}$, and

$(1-\omega/\sigma)^{-s} \leq (1-\omega/\sigma)(1-R\omega/\sigma)^{-1}$, $\forall s \in [R-1] \cup \{0\}$, and $\omega < \bar{\omega}(\sigma) < \sigma/R$, we can upper bound F^s as

$$F^s \leq F^R \leq a_1 + V_0 a_2 \triangleq \bar{F}^R, \quad (43)$$

where

$$a_1 \triangleq \psi \sqrt{md\eta^0} (1 + L_C\sigma) \cdot \frac{R/\sigma}{1 - R\omega/\sigma}, \quad (44)$$

$$a_2 \triangleq 2L_Z \psi \cdot \frac{R/\sigma}{1 - R\omega/\sigma}. \quad (45)$$

Furthermore, since $\mathbf{F}^\top\mathbf{1} \leq RF^R \leq R(a_1 + V_0 a_2)$ and $\tilde{L}_A = L_A \sum_{s=0}^{R-1} (L_C)^s \leq L_A \psi R$, to satisfy (36), it is sufficient to choose V_0 as

$$V_0 \geq \max \left\{ c^*, \|\mathbf{z}^0 - \mathbf{z}^\infty\|, L_A \psi R^2 \frac{\sqrt{md\eta^0} + \omega(a_1 + V_0 a_2)}{\sigma - \lambda} \right\}.$$

Using $x \geq \max\{c, a + bx\} \Leftrightarrow x \geq \max\{c, a/(1-b)\}$, under $b < 1$, the above condition is equivalent to

$$V_0 \geq \max \left\{ c^*, \|\mathbf{z}^0 - \mathbf{z}^\infty\|, \frac{L_A \psi R^2 (\sqrt{md\eta^0} + \omega a_1)}{\sigma - \lambda - L_A \psi R^2 \omega a_2} \right\}, \quad (46)$$

as long as $L_A \psi R^2 \omega a_2 / (\sigma - \lambda) < 1$. Solving with respect to ω (note that a_2 is a function of ω), this condition is equivalent to $\omega \in [0, \bar{\omega}(\sigma)]$ with $\bar{\omega}(\sigma)$ given by (10), hence it holds by assumption. Substituting the values of a_1, a_2 in (46) and using $\max\{a, b\} \leq a + b$ (for $a, b \geq 0$), yields (39). \square .

2) *Proof of Lemma 20:* At iteration k , let ζ_s be defined as

$$\zeta^s = \mathcal{A}(\mathbf{z}^k, \hat{\mathbf{c}}^{k,1}, \dots, \hat{\mathbf{c}}^{k,s}, \hat{\mathbf{c}}_s^{s+1}, \dots, \hat{\mathbf{c}}_s^R),$$

where

$$\hat{\mathbf{c}}_s^s = \hat{\mathbf{c}}^{k,s} \quad \text{and} \quad \hat{\mathbf{c}}_s^{\ell+1} \triangleq \mathcal{C}^{\ell+1}(\mathbf{z}^k, \hat{\mathbf{c}}_s^\ell), \quad \forall \ell \geq s.$$

In other words, $\hat{\mathbf{c}}_s^\ell, \zeta^s$ are the communication signals at round ℓ and the updated computation state, respectively, obtained by applying the unquantized communication mapping after round s and the quantized one before round s . Clearly, $\mathbf{z}^{k+1} = \mathcal{A}(\mathbf{z}^k, \hat{\mathbf{c}}^{k,1}, \dots, \hat{\mathbf{c}}^{k,R}) = \zeta^R$ and $\tilde{\mathcal{A}}(\mathbf{z}^k) = \zeta^0$ (unquantized update of the computation state). It then follows that

$$\mathbf{z}^{k+1} = \zeta^R = \tilde{\mathcal{A}}(\mathbf{z}^k) + \sum_{s=1}^R (\zeta^s - \zeta^{s-1}), \quad \text{a.s..}$$

Invoking the triangle inequality yields

$$\|\mathbf{z}^{k+1} - \mathbf{z}^\infty\| \leq \|\tilde{\mathcal{A}}(\mathbf{z}^k) - \mathbf{z}^\infty\| + \sum_{s=1}^R \|\zeta^s - \zeta^{s-1}\|, \quad \text{a.s..} \quad (47)$$

We now study the second term. From the Lipschitz continuity of \mathcal{A} (Assumption 4) and the definition of ζ^s , it holds that

$$\|\zeta^s - \zeta^{s-1}\| \leq L_A \sum_{\ell=s}^R \|\tilde{\mathbf{c}}_s^\ell - \tilde{\mathbf{c}}_{s-1}^\ell\|_2, \quad a.s..$$

Furthermore,

$$\|\tilde{\mathbf{c}}_s^s - \tilde{\mathbf{c}}_{s-1}^s\|_2 = \|\hat{\mathbf{c}}^{k,s} - \mathcal{C}^s(\mathbf{z}^k, \hat{\mathbf{c}}^{k,s-1})\|_2 = \|\hat{\mathbf{c}}^{k,s} - \mathbf{c}^{k,s}\|_2$$

a.s., and, for $\ell > s$,

$$\begin{aligned} \|\tilde{\mathbf{c}}_s^\ell - \tilde{\mathbf{c}}_{s-1}^\ell\|_2 &= \|\mathcal{C}^\ell(\mathbf{z}^k, \tilde{\mathbf{c}}_s^{\ell-1}) - \mathcal{C}^\ell(\mathbf{z}^k, \tilde{\mathbf{c}}_{s-1}^{\ell-1})\|_2 \\ &\leq L_C \|\tilde{\mathbf{c}}_s^{\ell-1} - \tilde{\mathbf{c}}_{s-1}^{\ell-1}\|_2 \leq \dots \leq (L_C)^{\ell-s} \|\hat{\mathbf{c}}^{k,s} - \mathbf{c}^{k,s}\|_2, \quad a.s., \end{aligned}$$

where the last step follows from induction over ℓ . Replacing these bounds in (47), we finally obtain

$$\begin{aligned} &\|\mathbf{z}^{k+1} - \mathbf{z}^\infty\| \\ &\leq \|\tilde{\mathcal{A}}(\mathbf{z}^k) - \mathbf{z}^\infty\| + L_A \sum_{s=1}^R \sum_{\ell=0}^{R-s} (L_C)^\ell \|\hat{\mathbf{c}}^{k,s} - \mathbf{c}^{k,s}\|_2 \\ &\leq \lambda \|\mathbf{z}^k - \mathbf{z}^\infty\| + \tilde{L}_A \sum_{s=1}^R \|\hat{\mathbf{c}}^{k,s} - \mathbf{c}^{k,s}\|_2, \quad a.s., \end{aligned}$$

where \tilde{L}_A is defined in (38) and we used Assumption 1. Taking the expectation conditional on the filtration $\mathcal{F}^{k,s}$ while applying Lemma 21 and (42), starting from $s = R, R-1, \dots, 1$, it follows that

$$\begin{aligned} \sqrt{\mathbb{E}[\|\mathbf{z}^{k+1} - \mathbf{z}^\infty\|^2 | \mathcal{F}^{k,1}]} &\leq \lambda \|\mathbf{z}^k - \mathbf{z}^\infty\| + \tilde{L}_A \sqrt{md} R \eta^0 \cdot (\sigma)^k \\ &\quad + \tilde{L}_A \omega \sum_{s=1}^R \sqrt{\mathbb{E}[\|\mathbf{c}^{k,s} - \hat{\mathbf{c}}^{k-1,s}\|_2^2 | \mathcal{F}^{k,1}]}, \quad a.s.. \end{aligned}$$

Finally, taking unconditional expectation and using Lemma 21 concludes the proof.

APPENDIX B

DETERMINISTIC AND RANDOM QUANTIZER'S DESIGN

1. Proof of Lemma 7

Let $\mathcal{Q}(\bullet) : [-\delta, \delta]^d \rightarrow \mathbb{Q}^d$ be a component-wise quantizer, with the n th component quantizer $\mathcal{Q}_n(\bullet)$ mapping points in the interval $[-\delta, \delta]$ to discrete points in the set \mathbb{Q} . We assume that the same quantizer is applied across all n , since each component is optimized with the same range and number of quantization points. The goal is to define a quantizer \mathcal{Q} which satisfies the BC-rule within $\mathbf{x} \in [-\delta, \delta]^d$ with maximal range δ . To this end, a necessary and sufficient condition is

$$|\mathcal{Q}_n(x) - x| \leq \eta + \omega|x|, \quad \forall x \in [-\delta, \delta], \quad \forall n \in [d]. \quad (48)$$

The sufficiency can be proved using Cauchy-Schwarz inequality. To prove the necessity, assume that (48) is violated for some $x \in [-\delta, \delta]$, i.e., $|\mathcal{Q}_n(x) - x| > \eta + \omega|x|$, and let $\mathbf{x} = x\mathbf{1}$. It follows that

$$\begin{aligned} \|\mathcal{Q}(\mathbf{x}) - \mathbf{x}\|_2 &= \sqrt{d}|\mathcal{Q}_n(x) - x| > \sqrt{d}\eta + \omega\sqrt{d}|x| \\ &= \sqrt{d}\eta + \omega\|\mathbf{x}\|_2, \end{aligned}$$

implying the BC-rule is not satisfied at \mathbf{x} .

Hence, we now focus on the design of a component-wise quantizer \mathcal{Q}_n satisfying (48) with maximal range δ . In the following, we omit the dependence on n for convenience.

Assume that $N = |\mathbb{Q}|$ is odd (the case N even can be studied in a similar fashion, and is provided at the end of this proof for completeness), and let $\mathbb{Q} \triangleq \cup_{\ell=0}^{(N-1)/2} \{\tilde{q}_\ell, -\tilde{q}_\ell\}$ be the set of quantization points, with $0 = \tilde{q}_0 < \tilde{q}_1 < \dots, \tilde{q}_\ell < \tilde{q}_{\ell+1} < \dots$. Note that we restrict to a symmetric quantizer since the error metric is symmetric around 0 (the detailed proof on the optimality of symmetric quantizers is omitted due to space constraints). We then aim to solve

$$\begin{aligned} &\max_{\delta \geq 0, \mathbf{Q}} \quad \delta \\ \text{s.t.} \quad &|\mathcal{Q}(x) - x| \leq \eta + \omega x, \quad \forall x \in [0, \delta], \end{aligned} \quad (49)$$

where the constraint (48) is imposed only to $x \in [0, \delta]$ since the quantizer is symmetric around 0. Since the quantization error in (49) is measured in Euclidean distance, it is optimal to restrict the quantization points to $\mathbb{Q} \subset [-\delta, \delta]$ and to map the input to the nearest quantization point (ties may be resolved arbitrarily). Then, letting $\mathcal{X}_\ell = ((\tilde{q}_{\ell-1} + \tilde{q}_\ell)/2, (\tilde{q}_\ell + \tilde{q}_{\ell+1})/2]$, with $\tilde{q}_{-1} = 0$ and $\tilde{q}_{(N+1)/2} = 2\delta - \tilde{q}_{(N-1)/2}$, it follows that $[0, \delta] \equiv \cup_{\ell=0}^{(N-1)/2} \mathcal{X}_\ell$ and $\mathcal{Q}(x) = \tilde{q}_\ell, \forall x \in \mathcal{X}_\ell$. Therefore, the optimization problem (49) can be expressed equivalently as

$$\begin{aligned} &\max_{\delta \geq 0, \tilde{\mathbf{q}}} \quad \delta \\ \text{s.t.} \quad &(\tilde{q}_\ell - x)^2 \leq (\eta + \omega x)^2, \quad \forall x \in \mathcal{X}_\ell, \forall \ell = 0, 1, \dots, \frac{N-1}{2}, \\ &0 = \tilde{q}_0 \leq \dots \leq \tilde{q}_{(N+1)/2} = 2\delta - \tilde{q}_{(N-1)/2}. \end{aligned}$$

Equivalently,

$$\begin{aligned} &\max_{\delta \geq 0, \tilde{\mathbf{q}}} \quad \delta \\ \text{s.t.} \quad &\max_{x \in \mathcal{X}_\ell} (\tilde{q}_\ell - x)^2 - (\eta + \omega x)^2 \leq 0, \quad \forall \ell = 0, 1, \dots, \frac{N-1}{2}, \\ &0 = \tilde{q}_0 \leq \dots \leq \tilde{q}_{(N+1)/2} = 2\delta - \tilde{q}_{(N-1)/2}, \end{aligned}$$

and solving the maximization with respect to $x \in \mathcal{X}_\ell$ (note that the quadratic function is convex in x , hence it is maximized at the boundaries of \mathcal{X}_ℓ), we obtain

$$\begin{aligned} &\max_{\delta \geq 0, \tilde{\mathbf{q}}} \quad \frac{\tilde{q}_{(N-1)/2} + \tilde{q}_{(N+1)/2}}{2} \\ \text{s.t.} \quad &\tilde{q}_\ell \leq \tilde{q}_{\ell-1} \left(\frac{1+\omega}{1-\omega} \right) + \frac{2\eta}{1-\omega}, \quad \forall \ell \in [(N+1)/2], \\ &0 = \tilde{q}_0 < \tilde{q}_1 < \dots < \tilde{q}_{(N+1)/2}. \end{aligned}$$

Solving this problem with respect to $\tilde{\mathbf{q}}$ yields $q_0 = 0$ and

$$q_\ell = q_{\ell-1} \left(\frac{1+\omega}{1-\omega} \right) + \frac{2\eta}{1-\omega}, \quad \forall \ell \geq 1.$$

Solving by induction, we obtain q_ℓ as in (12), $\delta(\eta, \omega, N)$ as in (11), and

$$\ell(x) = \text{sign}(x) \cdot \min \left\{ \ell \geq 0 : \frac{q_\ell + q_{\ell+1}}{2} \geq |x| \right\},$$

yielding (13) after solving with the expression of q_ℓ . A similar technique can be proved for the case when N is even, yielding quantization points

$$q_\ell = -q_{-\ell} = \frac{\eta}{\omega} \left[\frac{(1+\omega)^\ell}{(1-\omega)^{\ell-1}} - 1 \right], \quad \forall \ell \geq 1$$

and $\delta(\eta, \omega, N) = \frac{q_{N/2} + q_{N/2+1}}{2}$, which concludes the proof.

2. Proof of Lemma 8

Using a similar technique as in Appendix B.1 when N is odd, using the fact that $[0, \delta] = \cup_{\ell \in [(N-1)/2]} [\tilde{q}_{\ell-1}, \tilde{q}_\ell]$ and $\delta = \tilde{q}_{N/2}$ it suffices to solve

$$\begin{aligned} \max_{\delta \geq 0, \tilde{\mathbf{q}}} & q_{(N-1)/2} \\ \text{s.t.} & \mathbb{E}[|\mathcal{Q}(x) - x|^2] \leq (\eta + \omega x)^2, \forall x \in [\tilde{q}_{\ell-1}, \tilde{q}_\ell], \forall \ell \in \left[\frac{N-1}{2} \right], \\ & \mathbb{E}[\mathcal{Q}(x)] = x, \quad 0 = \tilde{q}_0 \leq \dots \leq \tilde{q}_{(N-1)/2} = \delta. \end{aligned}$$

Furthermore, since $x \in [\tilde{q}_{\ell-1}, \tilde{q}_\ell]$ is mapped to $\tilde{q}_{\ell-1}$ w.p. $(\tilde{q}_\ell - x)/(\tilde{q}_\ell - \tilde{q}_{\ell-1})$ and to \tilde{q}_ℓ w.p. $(x - \tilde{q}_{\ell-1})/(\tilde{q}_\ell - \tilde{q}_{\ell-1})$ to satisfy $\mathbb{E}[\mathcal{Q}(x)] = x$, the problem can be expressed equivalently as

$$\begin{aligned} \max_{\delta \geq 0, \tilde{\mathbf{q}}} & q_{(N-1)/2} \\ \text{s.t.} & (x - \tilde{q}_\ell)(x - \tilde{q}_{\ell-1}) + (\eta + \omega x)^2 \geq 0, \forall x \in [\tilde{q}_{\ell-1}, \tilde{q}_\ell], \\ & \forall \ell \in \left[\frac{N-1}{2} \right], \quad 0 = \tilde{q}_0 \leq \dots \leq \tilde{q}_{(N-1)/2} = \delta, \end{aligned}$$

or equivalently

$$\begin{aligned} \max_{\delta \geq 0, \tilde{\mathbf{q}}} & q_{(N-1)/2} \\ \text{s.t.} & \min_{x \in [\tilde{q}_{\ell-1}, \tilde{q}_\ell]} (x - \tilde{q}_\ell)(x - \tilde{q}_{\ell-1}) + (\eta + \omega x)^2 \geq 0, \forall \ell \in \left[\frac{N-1}{2} \right], \\ & 0 = \tilde{q}_0 \leq \dots \leq \tilde{q}_{(N-1)/2} = \delta. \end{aligned}$$

Solving the minimization over $x \in [\tilde{q}_{\ell-1}, \tilde{q}_\ell]$ and solving with respect to $\tilde{\mathbf{q}}$ yields the following optimal quantization points: $q_0 = 0$ and

$$q_\ell = q_{\ell-1}(\sqrt{1 + (\omega)^2} + \omega)^2 + 2\eta(\sqrt{1 + (\omega)^2} + \omega), \quad \forall \ell \geq 1.$$

Solving by induction, we obtain q_ℓ as in (15), $\delta(\eta, \omega, N)$ as in (14), and the probabilistic quantization rule as in (16), with ℓ given by

$$\ell = \text{sign}(x) \cdot \min \left\{ \ell \geq 0 : q_\ell \geq |x| \right\},$$

yielding (16) after solving with the expression of q_ℓ .

A similar technique can be proved for the case when N is even, yielding the quantization points

$$q_\ell = \frac{\eta}{\omega} \left[\frac{(\sqrt{1 + (\omega)^2} + \omega)^{2\ell-1}}{\sqrt{1 + (\omega)^2}} - 1 \right], \quad \forall \ell \geq 1.$$

and $\delta(\eta, \omega, N) = q_{N/2}$, which concludes the proof.

1. Proof of Corollary 9

Let $\mathcal{Q}(\mathbf{x})$ be a generic deterministic or probabilistic quantizer with domain $[-\delta, \delta]^d$ and codomain $\mathbb{Q} \in \mathbb{R}^d$ with $|\mathbb{Q}| < \infty$, that satisfies the BC-rule with $\eta = 0$. It follows that

$$\begin{aligned} \omega \|\mathbf{x}\|_2 & \geq \sqrt{\mathbb{E}[\|\mathcal{Q}(\mathbf{x}) - \mathbf{x}\|_2^2]} \geq \min_{\mathbf{q} \in \mathbb{Q}} \|\mathbf{q} - \mathbf{x}\|_2 \\ & = \|\mathcal{Q}_{\text{det}}(\mathbf{x}) - \mathbf{x}\|_2, \quad \forall \mathbf{x} \in [-\delta, \delta]^d, \end{aligned} \quad (50)$$

where the lower bound is achievable by a deterministic quantizer that maps \mathbf{x} to the nearest quantization point, denoted as $\mathcal{Q}_{\text{det}}(\mathbf{x})$. Let $\mathcal{Q}_n(x) = \mathbf{e}_n^\top \mathcal{Q}_{\text{det}}(x \mathbf{e}_n)$ be the projection of \mathcal{Q}_{det} on its n th element, where \mathbf{e}_n is the n th canonical vector.

Since $\|\mathcal{Q}_{\text{det}}(x \mathbf{e}_n) - x \mathbf{e}_n\|_2 \geq |\mathcal{Q}_n(x) - x|$, from (50) it follows that

$$\omega \|\mathbf{x} \mathbf{e}_n\|_2 = \omega |x| \geq \|\mathcal{Q}_{\text{det}}(x \mathbf{e}_n) - x \mathbf{e}_n\|_2 \geq |\mathcal{Q}_n(x) - x|,$$

$\forall x \in [-\delta, \delta]$, hence \mathcal{Q}_n satisfies the BC-rule with $\eta = 0$ as well. Note that \mathcal{Q}_n is a scalar quantizer with $N_n \leq |\mathbb{Q}|$ quantization points. However, Lemma 7 dictates that $\delta = 0$ for this quantizer, hence the contradiction. We have thus proved the statement of Corollary 9 for both deterministic and probabilistic compression rules.

APPENDIX C COMMUNICATION COST ANALYSIS

1. Proof of Theorem 12

We first present some preliminary results instrumental in proving Theorem 12, whose proofs are deferred to Appendix C.2. The idea of the proof is to study the asymptotic behavior of an upper bound on the number of bits required per iteration, provided in the following lemma.

Lemma 22: Under the same setting as Theorem 6, and the proposed ANQ satisfying the BC rule, the average number of bits required per agent at the k th iteration, B^k , is upper bounded as

$$\mathbb{E}[B^k] \leq \log_2(S+1) \left[3dR + dR \log_S \left(3 + \frac{\bar{F}^R(\sigma, \omega, \eta^0)}{\sqrt{md\eta^0}} \right) \right]$$

bits, $\forall k \geq 0$, where $\bar{F}^R(\sigma, \omega, \eta^0)$ is defined in (43).

In addition, we need the following lemma to connect the asymptotic results of the logarithmic function and its argument.

Lemma 23: For positive functions f, g , it holds: $\ln f(x) = \mathcal{O}(\ln g(x))$ as $x \rightarrow x_0$ if $\liminf_{x \rightarrow x_0} g(x) > 1$, and $f(x) = \mathcal{O}(g(x))$ as $x \rightarrow x_0$.

We are now ready to prove the main theorem. From Lemma 22, the average number of bits per agent per iteration is upper bounded by

$$\mathbb{E}[B_k] \leq \log_2(S+1) \left[3dR + dR \log_S \left(3 + \frac{\bar{F}^R(\sigma, \omega, \eta^0)}{\sqrt{md\eta^0}} \right) \right],$$

$\forall k = 0, 1, \dots$, where

$$\bar{F}^R(\sigma, \omega, \eta^0) = a_1 + a_2 V_0,$$

with a_1, a_2 and V_0 defined in (44), (45) and (39), respectively. We want to prove that this is $\mathbb{E}[B_k] = \mathcal{O}(d \ln(1 + \frac{1}{\sigma(\sigma - \lambda)}))$ under Assumption 11 and conditions

$$\eta^0 = \Theta(L_Z(\sigma - \lambda)), \quad 1 - \omega/\bar{\omega}(\sigma) = \Omega(1).$$

Using the fact that $\bar{F}^R(\sigma, \omega, \eta^0)/\eta^0 = a_1/\eta^0 + a_2 V_0/\eta^0$, it is sufficient to show that $a_1/(\sqrt{md\eta^0}) = \mathcal{O}(1 + \frac{1}{\sigma})$ and $a_2 V_0/(\sqrt{md\eta^0}) = \mathcal{O}(1 + \frac{1}{\sigma(\sigma - \lambda)})$. In fact, using $R/\sigma \leq 1/\bar{\omega}(\sigma)$, we can bound

$$\begin{aligned} \frac{a_1}{\sqrt{md\eta^0}} & \leq \frac{1}{\sigma} (1 + L_C \sigma) \max\{1, (2L_C)^{R-1}\} \frac{R}{1 - \omega/\bar{\omega}(\sigma)}, \\ a_2 & \leq \frac{2L_Z}{\sigma} \max\{1, (2L_C)^{R-1}\} \frac{R}{1 - \omega/\bar{\omega}(\sigma)}. \end{aligned}$$

Clearly, $a_1/(\sqrt{md}\eta^0) = \mathcal{O}(1 + \frac{1}{\sigma})$ and $a_2 = \mathcal{O}(\frac{L_Z}{\sigma})$ since $L_C, R = \mathcal{O}(1)$ and $1 - \omega/\bar{\omega}(\sigma) = \Omega(1)$. We next study V_0 . From its expression in (39), we notice that $V_0 = \mathcal{O}(\sqrt{md})$ since $\omega/\sigma \leq 1$, $\max\{c^*, \|\mathbf{z}^0 - \mathbf{z}^\infty\|_2\} = \mathcal{O}(\sqrt{md})$, $\eta^0 = \Theta(L_Z(\sigma - \lambda))$, $L_A \cdot L_Z, L_C, R = \mathcal{O}(1)$, and $1 - \omega/\bar{\omega} = \Omega(1)$. Therefore, it follows that $a_2 V_0/(\sqrt{md}\eta^0) = \mathcal{O}(\frac{1}{\sigma(\sigma - \lambda)}) = \mathcal{O}(1 + \frac{1}{\sigma(\sigma - \lambda)})$, and the proof is completed by invoking Lemma 23.

2. Auxiliary Results for Theorem 12

3) *Proof of Lemma 22:* Let $\Delta\mathbf{c}_i^{k,s} \triangleq \mathbf{c}_i^{k,s} - \hat{\mathbf{c}}_i^{k-1,s}$ be the input to the quantizer for agent i , at iteration k and communication round s . We now study the average number of bits required for i) the deterministic quantizer, and ii) the probabilistic quantizer with $\mathbb{E}[\mathcal{Q}(x)] = x$.

i) *Deterministic Quantizer:* The average number of bits required is bounded as (see (18), one can also verify that the following also holds for even N)

$$b_i^{k,s} \leq \log_2(S+1) \left[3d + d \log_S \left(2 + \frac{\ln(1 + \frac{\omega \|\Delta\mathbf{c}_i^{k,s}\|_2}{\sqrt{d}\eta^0 \cdot (\sigma)^k})}{\ln(1 + \omega) - \ln(1 - \omega)} \right) \right]$$

bits. We now upper bound the argument inside the second logarithm. Since it is a decreasing function of ω , it is maximized in the limit $\omega \rightarrow 0$, yielding

$$\frac{\ln(1 + \frac{\omega \|\Delta\mathbf{c}_i^{k,s}\|_2}{\sqrt{d}\eta^0 \cdot (\sigma)^k})}{\ln(1 + \omega) - \ln(1 - \omega)} \leq \frac{\|\Delta\mathbf{c}_i^{k,s}\|_2}{2\sqrt{d}\eta^0 \cdot (\sigma)^k}.$$

With this upper bound, we can then upper bound the average number of bits per agent at communication round s , iteration k , as

$$\begin{aligned} \mathbb{E}[b^{k,s}] &\triangleq \frac{1}{m} \sum_{i=1}^m \mathbb{E}[b_i^{k,s}] \\ &\stackrel{(a)}{\leq} \log_2(S+1) \left[3d + d \log_S \left(2 + \frac{\sqrt{\mathbb{E}[\|\mathbf{c}^{k,s} - \hat{\mathbf{c}}^{k-1,s}\|_2^2]}}{2\sqrt{md}\eta^0 \cdot (\sigma)^k} \right) \right] \\ &\stackrel{(b)}{\leq} \log_2(S+1) \left[3d + d \log_S \left(2 + \frac{F^s(\sigma, \omega, \eta^0)}{2\sqrt{md}\eta^0} \right) \right] \\ &\stackrel{(c)}{\leq} \log_2(S+1) \left[3d + d \log_S \left(2 + \frac{\bar{F}^R(\sigma, \omega, \eta^0)}{2\sqrt{md}\eta^0} \right) \right], \end{aligned}$$

where (a) follows from Cauchy-Schwarz inequality, Jensen's inequality, and the definition of $\Delta\mathbf{c}^{k,s}$; (b) follows from (35); and (c) follows from $F^s \leq \bar{F}^R$ (see (43)).

ii) *Probabilistic Quantizer With $\mathbb{E}[\mathcal{Q}(x)] = x$:* Using the same technique as in i), along with the inequality $1 - 1/x \leq \ln(x) \leq x - 1$ for $x > 1$ to bound the argument inside the second logarithm of (19), we find the bound

$$\mathbb{E}[b^{k,s}] \leq \log_2(S+1) \left[3d + d \log_S \left(3 + \frac{\bar{F}^R(\sigma, \omega, \eta^0)}{\sqrt{md}\eta^0} \right) \right].$$

One can also verify that it also holds for even N .

Finally, for both the deterministic and probabilistic cases, the proof is completed by summing over $s \in [R]$ to get the average communication cost per agent at iteration k .

4) *Proof of Lemma 23:* If $f(x) = \mathcal{O}(g(x))$ as $x \rightarrow x_0$ and $\liminf_{x \rightarrow x_0} g(x) > 1$, then

$$\begin{aligned} \limsup_{x \rightarrow x_0} \left| \frac{\ln f(x)}{\ln g(x)} \right| &\leq 1 + \limsup_{x \rightarrow x_0} \left| \frac{\ln(f(x)/g(x))}{\ln g(x)} \right| \\ &\stackrel{(a)}{\leq} 1 + \frac{\limsup_{x \rightarrow x_0} |\ln(f(x)/g(x))|}{|\liminf_{x \rightarrow x_0} \ln g(x)|} < \infty, \end{aligned}$$

where (a) follows from $\liminf_{x \rightarrow x_0} g(x) > 1$. This completes the proof.

3. Proof of Theorem 13

Since $\sqrt{\mathbb{E}[\|\mathbf{z}^k - \mathbf{z}^\infty\|^2]} \leq V_0 \cdot (\sigma)^k$ (cf. Theorem 6), the ε -accuracy is achieved if $k[-\ln(\sigma)] \geq \ln(V_0/\sqrt{m\varepsilon})$, which yields $k(1 - \sigma) \geq \ln(V_0/\sqrt{m\varepsilon})$ since $-\ln(\sigma) \geq 1 - \sigma$. Hence, ε -accuracy is achieved if all conditions in Theorem 6 hold and $k \geq k_\varepsilon \triangleq \lceil \frac{1}{1-\sigma} \ln \frac{V_0}{\sqrt{m\varepsilon}} \rceil$. Hence, to compute the upper bound of the communication cost $\sum_{k=0}^{k_\varepsilon-1} \mathbb{E}[B^k]$, we need the upper bounds for $\mathbb{E}[B^k]$, $\frac{1}{1-\sigma}$ and V_0 . In the proof of Theorem 12, we found that, under Assumption 11 and the conditions $1 - \omega/\bar{\omega}(\sigma) = \Omega(1)$, $\eta^0 = \Theta(L_Z(\sigma - \lambda))$,

$$\mathbb{E}[B^k] = \mathcal{O}\left(d \log_2 \left(1 + \frac{1}{\sigma(\sigma - \lambda)}\right)\right), \quad V_0 = \mathcal{O}(\sqrt{md}).$$

Moreover, it can be shown that $1/\sigma \leq \frac{(1-\lambda)^2}{(1-\sigma)(\sigma-\lambda)}$ for $\sigma \in (\lambda, 1)$, and therefore

$$\frac{1}{\sigma(\sigma - \lambda)} \leq \frac{1}{(1-\lambda)} \left[\frac{(1-\lambda)^2}{(1-\sigma)(\sigma-\lambda)} \right]^2 \frac{1-\sigma}{1-\lambda} = \mathcal{O}\left(\frac{1}{1-\lambda}\right),$$

where we used the assumption in Theorem 13 that $\frac{(1-\lambda)^2}{(1-\sigma)(\sigma-\lambda)} = \mathcal{O}(1)$. It then follows from Lemma 23 that $\mathbb{E}[B^k] = \mathcal{O}(d \log_2(1 + \frac{1}{1-\lambda}))$. On the other hand, we can bound k_ε as

$$k_\varepsilon \leq \frac{1}{1-\lambda} \left(1 - \lambda + \frac{1-\lambda}{1-\sigma} \log_2 \frac{V_0}{\sqrt{m\varepsilon}} \right) = \mathcal{O}\left(\frac{1}{1-\lambda} \log_2 \left(\frac{d}{\varepsilon}\right)\right),$$

since $\frac{1-\lambda}{1-\sigma} = \mathcal{O}(1)$ and $V_0 = \mathcal{O}(\sqrt{md})$. Therefore, the communication cost satisfies

$$\begin{aligned} \sum_{k=0}^{k_\varepsilon-1} \mathbb{E}[B^k] &= \mathcal{O}\left(d k_\varepsilon \log_2 \left(1 + \frac{1}{1-\lambda}\right)\right) \\ &= \mathcal{O}\left(\frac{d}{1-\lambda} \log_2 \left(\frac{d}{\varepsilon}\right) \log_2 \left(1 + \frac{1}{1-\lambda}\right)\right), \end{aligned}$$

which completes the proof.

4. Proof of Lemma 10

Consider $\ell \geq 0$. Using (17), the number of information symbols required to encode ℓ is

$$b_\ell^* = \min \left\{ b \geq 0 : \ell \leq \left\lfloor \frac{(S)^{b+1} - 1}{2(S-1)} \right\rfloor \right\}.$$

Similarly, for $\ell < 0$,

$$b_\ell^* = \min \left\{ b \geq 0 : -\ell \leq \left\lceil \frac{(S)^{b+1} - 1}{2(S-1)} \right\rceil - 1 \right\}.$$

Since $\lfloor \frac{(S)^{b+1}-1}{2(S-1)} \rfloor \geq \lceil \frac{(S)^{b+1}-1}{2(S-1)} \rceil - 1 \geq \frac{(S)^{b+1}-1}{2(S-1)} - 1$, we can then upper bound $b_\ell, \ell \in \mathbb{Z}$, as

$$\begin{aligned} b_\ell^* &\leq \min\{b \geq 1 : 1 + 2(S-1)(1+|\ell|) \leq (S)^{b+1}\} \\ &= \min\{b \geq 1 : b \geq \log_S(2 - 1/S + 2(1 - 1/S)|\ell|)\} \\ &= \lceil \log_S(2 - 1/S + 2(1 - 1/S)|\ell|) \rceil. \end{aligned}$$

Using $\lceil x \rceil \leq x + 1$ We can then further upper bound

$$\begin{aligned} b_\ell^* &\leq 1 + \log_S(2 - 1/S + 2(1 - 1/S)|\ell|) \leq \log_S(2S + 2S|\ell|) \\ &\leq 2 + \log_S(1 + |\ell|). \end{aligned}$$

This yields the upper bound to the communication cost (including the termination symbol)

$$\bar{C}_{\text{comm}}(\ell) \leq 3 + \log_S(1 + |\ell|) \text{ symbols.} \quad (51)$$

Let $\mathbf{x} = (x_n)_{n=1}^d$. Note that for both the deterministic and probabilistic quantizers, we can express $\ell(x)$ as

$$|\ell(x)| \leq \left\lceil c_1 + c_2 \ln \left(1 + \frac{\omega}{\eta} |x| \right) \right\rceil,$$

for some $c_1 \leq 0$, $c_2 > 0$ (see (13) and (16) for a closed-form expression of c_1 and c_2). Invoking (51) and $\bar{C}_{\text{comm}}(\mathbf{x}) = \sum_{n=1}^d \bar{C}_{\text{comm}}(\ell_n)$ yields

$$\begin{aligned} C(\mathbf{x}) &\leq 3d + \sum_{n=1}^d \log_S \left(1 + \left\lceil c_1 + c_2 \ln \left(1 + \frac{\omega}{\eta} |x_n| \right) \right\rceil \right) \\ &\stackrel{(a)}{\leq} 3d + d \log_S \left(2 + c_1 + c_2 \ln \left(1 + \frac{\omega \|\mathbf{x}\|_2}{\sqrt{d}\eta} \right) \right) \text{ symbols} \\ &= \log_2(S+1) \left[3d + d \log_S \left(2 + c_1 + c_2 \ln \left(1 + \frac{\omega \|\mathbf{x}\|_2}{\sqrt{d}\eta} \right) \right) \right] \end{aligned}$$

bits, where (a) follows from $\lceil x \rceil \leq x + 1$, Jensen's inequality and Cauchy-Schwarz inequality, in order. Invoking the expressions of c_1 and c_2 from (13) and (16), respectively, yield the result for the deterministic and probabilistic quantizers with odd N . Similar techniques can be used to find $\ell(x)$ and thus the result for quantizers with even N .

APPENDIX D EXAMPLES OF (M)

In this section, we will show that (M) contains a gamut of distributed algorithms, corresponding to different choices of R, \mathcal{C}_i^s , and \mathcal{A}_i . Given (P), we will assume that each f_i is L -smooth and μ -strongly convex, and define $\tilde{\mathbf{x}}_i^* = \arg \min_{\mathbf{x}_i} f_i(\mathbf{x}_i)$.

Every distributed algorithm on mesh networks we will describe below alternates one step of optimization with possibly multiple rounds of communications. In each communication round, every agent i combines linearly the signals received by its neighbors using weights $(w_{ij})_{j \in \mathcal{N}_i}$; let $\mathbf{W} = (w_{ij})_{i,j=1}^m$. Consistently with the undirected graph \mathcal{G} , we will tacitly assume that \mathbf{W} is symmetric and doubly stochastic, i.e., $\mathbf{W} = \mathbf{W}^\top$ and $\mathbf{W} \cdot \mathbf{1} = \mathbf{1}$, with $w_{ij} > 0$ if $(j, i) \in \mathcal{E}$, and $w_{ij} = 0$ otherwise. We assume that the eigenvalues of \mathbf{W} are in $[\nu, 1]$, with $\nu > 0$.⁷ Note that this condition can be

⁷This assumption is also required in [50] for prox-EXTRA, prox-NEXT, prox-DIGing, and prox-NIDS to achieve $\|\mathbf{z}^k - \mathbf{z}^\infty\| = \mathcal{O}(\sqrt{md}(\lambda)^k)$.

achieved by design: in fact, given a doubly stochastic weight matrix $\tilde{\mathbf{W}}$, one can choose

$$\mathbf{W} = \frac{(1+\nu)}{2} \mathbf{I} + \frac{(1-\nu)}{2} \tilde{\mathbf{W}}, \quad (52)$$

for any given $\nu \in (0, 1]$. Note that, for any given \mathbf{z}^0 and \mathbf{z}^∞ with bounded entries, it holds $\|\mathbf{z}^0 - \mathbf{z}^\infty\|_2 = \mathcal{O}(\sqrt{md})$.

Finally, in the rest of this section, we will adopt the following notations: $\mathbf{x}^k = (\mathbf{x}_i^k)_{i=1}^m$, $\mathbf{y}^k = (\mathbf{y}_i^k)_{i=1}^m$, $\mathbf{w}^k = (\mathbf{w}_i^k)_{i=1}^m$, $\mathbf{x} = (\mathbf{x}_i)_{i=1}^m$, $\mathbf{y} = (\mathbf{y}_i)_{i=1}^m$, $\mathbf{w} = (\mathbf{w}_i)_{i=1}^m$ and $\tilde{\mathbf{x}}^* = (\tilde{\mathbf{x}}_i^*)_{i=1}^m$; $\mathbf{W} = \mathbf{W} \otimes \mathbf{I}_d$, and \mathbf{G}^\dagger is the pseudo-inverse of matrix \mathbf{G} . Given $\mathbf{x}^k = (\mathbf{x}_i^k)_{i=1}^m$, we also define $\nabla f(\mathbf{x}^k) \triangleq (\nabla f_i(\mathbf{x}_i^k))_{i=1}^m$. For any function $g : \mathbb{R}^d \rightarrow \mathbb{R}$ and positive semi-definite matrix \mathbf{G} , define $\|\mathbf{x}\|_{\mathbf{G}} \triangleq \sqrt{\mathbf{x}^\top \mathbf{G} \mathbf{x}}$ and

$$\text{prox}_{\mathbf{G}, g}(\mathbf{z}) \triangleq \arg \min_{\mathbf{z} \in \mathbb{R}^d} g(\mathbf{z}) + \frac{1}{2} \|\mathbf{z} - \mathbf{x}\|_{\mathbf{G}^{-1}}^2.$$

1. GD Over Star Networks [51]

Consider Problem (P) with $r \equiv 0$ over a master/workers system. The GD update

$$\mathbf{x}^{k+1} = \mathbf{x}^k - \frac{\gamma}{m} \sum_{i=1}^m \nabla f_i(\mathbf{x}^k), \quad (53)$$

with $\mathbf{x}^0 \in \mathbb{R}^d$, is implemented at the master node as follows: at iteration k , the server broadcasts \mathbf{x}^k to the m agents; each agent i then computes its own gradient $\nabla f_i(\mathbf{x}^k)$ and sends it back to the master; upon collecting all local gradients, the server updates the variable \mathbf{x}^{k+1} according to (53).

The GD (53) can be cast as (M) with $R = 1$ round of communications, using the following:

$$\begin{aligned} \mathbf{z} &= \mathbf{1}_m \otimes \mathbf{x}, \\ \hat{\mathbf{c}}_i^{k,1} &= \mathcal{C}_i^1(\mathbf{z}_i^k, \mathbf{0}) = \nabla f_i(\mathbf{x}^k), \end{aligned} \quad (54)$$

$$\mathbf{z}^{k+1} = \mathcal{A}(\mathbf{z}^k, \hat{\mathbf{c}}_{[m]}^{k,1}) = \mathbf{1}_m \otimes \mathbf{x}^k - \frac{\gamma}{m} \sum_{i=1}^m \mathbf{1}_m \otimes \hat{\mathbf{c}}_i^{k,1}. \quad (55)$$

We now show that the above instance of (M) satisfies Assumptions 1, 4, 5, and 11.

• On Assumption 1: Using [51] it is not difficult to check that, if $\gamma = 2/(\mu + L)$, then GD over star networks satisfies Assumption 1 with

$$\lambda = \frac{\kappa - 1}{\kappa + 1} < 1,$$

and the norm $\|\bullet\|$ defined as

$$\|\mathbf{z}\| = \sqrt{m} \|\mathbf{x}\|_2.$$

Note that $\|\mathbf{z}\| = \|\mathbf{z}\|_2$.

• On Assumptions 4, 5, and 11: Based on (54) and (55), the mappings \mathcal{A} and \mathcal{C} read

$$\mathcal{A}(\mathbf{z}, \mathbf{c}^1) = \mathbf{z} - \frac{\gamma}{m} \sum_{i=1}^m \mathbf{1}_m \otimes \mathbf{c}_i, \quad \mathcal{C}^1(\mathbf{z}, \mathbf{0}) = \nabla f(\mathbf{z}),$$

respectively; and $\mathcal{Z} = \{\mathbf{1}_m \otimes \mathbf{x} : \mathbf{x} \in \mathbb{R}^d\}$. Since

$$\mathcal{A}(\mathbf{z}, \mathbf{c}) - \mathcal{A}(\mathbf{z}, \mathbf{c}') = -\frac{\gamma}{m} \sum_{i=1}^m \mathbf{1}_m \otimes (\mathbf{c}_i - \mathbf{c}'_i),$$

we have

$$\|\mathcal{A}(\mathbf{z}, \mathbf{c}) - \mathcal{A}(\mathbf{z}, \mathbf{c}')\| = \frac{\gamma}{\sqrt{m}} \left\| \sum_{i=1}^m (\mathbf{c}_i - \mathbf{c}'_i) \right\|_2 \leq \gamma \|\mathbf{c} - \mathbf{c}'\|_2.$$

Hence, Assumption 4 holds with $L_A = \gamma$.

We now derive L_C and L_Z . Note that

$$\|\mathcal{C}^1(\mathbf{z}, \mathbf{0}) - \mathcal{C}^1(\mathbf{z}', \mathbf{0})\|_2 \leq L \|\mathbf{z} - \mathbf{z}'\|_2,$$

which implies that Assumption 5 holds with $L_C = 0$ and $L_Z = L$. Since $\gamma \leq 2/L$, it follows that $L_A \cdot L_Z = \mathcal{O}(1)$ and $\|\mathcal{C}^1(\mathbf{z}^0, \mathbf{0})\|_2 \leq L \|\mathbf{z}^0 - \tilde{\mathbf{x}}^*\|_2 = \mathcal{O}(L_Z \sqrt{md})$. Hence, Assumption 11 holds.

2. (Prox-)EXTRA [50]

The update of prox-EXTRA solving (P) reads

$$\begin{aligned} \mathbf{x}^k &= \text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}^k), \\ \mathbf{w}^{k+1} &= \hat{\mathbf{W}} \mathbf{x}^k - \gamma \nabla f(\mathbf{x}^k) - \mathbf{y}^k, \\ \mathbf{y}^{k+1} &= \mathbf{y}^k + (\mathbf{I} - \hat{\mathbf{W}}) \mathbf{w}^{k+1}, \end{aligned}$$

with $\mathbf{y}^0 = \mathbf{0}$ and $\mathbf{w}^0 \in \mathbb{R}^{md}$.

Prox-EXTRA can be cast as (M) with $R = 2$ rounds of communications with

$$\begin{aligned} \mathbf{z}^\top &= [\mathbf{y}^\top, \mathbf{w}^\top], \\ \hat{\mathbf{c}}_i^{k,1} &= \mathcal{C}_i^1(\mathbf{z}_i^k, \mathbf{0}) = \text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}_i^k), \end{aligned} \quad (56)$$

$$\hat{\mathbf{c}}_i^{k,2} = \mathcal{C}_i^2(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,1}) = \sum_{j \in \mathcal{N}_i} w_{ij} \hat{\mathbf{c}}_j^{k,1} - \gamma \nabla f_i(\hat{\mathbf{c}}_i^{k,1}) - \mathbf{y}_i^k, \quad (57)$$

$$\begin{aligned} \mathbf{z}_i^{k+1} &= \mathcal{A}_i(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,1}, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,2}) \\ &= \left[\begin{array}{c} \mathbf{y}_i^k + \sum_{j \in \mathcal{N}_i} w_{ij} (\hat{\mathbf{c}}_i^{k,2} - \hat{\mathbf{c}}_j^{k,2}) \\ \hat{\mathbf{c}}_i^{k,2} \end{array} \right]. \end{aligned} \quad (58)$$

We now show that the above instance of (M) satisfies Assumptions 1, 4, 5, and 11.

• On Assumption 1: Using [50, Theorem 18] it is not difficult to check that, if $\gamma = \frac{2\rho_m(\mathbf{W})}{L + \mu\rho_m(\mathbf{W})}$ and $\nu = \frac{\kappa}{\kappa+1}$, then prox-EXTRA satisfies Assumption 1 with

$$\begin{aligned} \lambda &= \max \left\{ \frac{\kappa - \rho_m(\mathbf{W})}{\kappa + \rho_m(\mathbf{W})} \frac{1}{\sqrt{\rho_m(\mathbf{W})}}, \sqrt{\rho_2(\mathbf{W})} \right\} \\ &\leq \max \left\{ \frac{\kappa}{\kappa+1} \sqrt{\frac{(\kappa+1)^3}{\kappa(\kappa+2)^2}}, \sqrt{\rho_2(\mathbf{W})} \right\} \\ &\leq \max \left\{ \frac{\kappa}{\kappa+1}, \sqrt{\rho_2(\mathbf{W})} \right\} < 1, \end{aligned}$$

(note that $\rho_m(\mathbf{W}) \geq \nu$) and the norm $\|\bullet\|$ defined as

$$\|\mathbf{z}\|^2 = \mathbf{y}^\top (\mathbf{I} - \hat{\mathbf{W}})^\dagger \mathbf{y} + \mathbf{w}^\top \hat{\mathbf{W}}^{-1} \mathbf{w}.$$

Note that $\|\mathbf{z}\|^2 \geq \|\mathbf{z}\|_2^2$, due to $\rho_i(\mathbf{W}) \in [\nu, 1]$, $i \in [m]$.

• On Assumptions 4, 5, and 11: Based on (56), (57) and (58), the mappings \mathcal{A} and \mathcal{C} read

$$\mathcal{A}(\mathbf{z}, \mathbf{c}^1, \mathbf{c}^2) = \left[\begin{array}{c} \mathbf{y} + (\mathbf{I} - \hat{\mathbf{W}}) \mathbf{c}^2 \\ \mathbf{c}^2 \end{array} \right],$$

$$\mathcal{C}^1(\mathbf{z}, \mathbf{0}) = \text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}), \text{ and } \mathcal{C}^2(\mathbf{z}, \mathbf{c}) = \hat{\mathbf{W}} \mathbf{c} - \gamma \nabla f(\mathbf{c}) - \mathbf{y},$$

respectively; and $\mathcal{Z} = \text{span}(\mathbf{I} - \hat{\mathbf{W}}) \times \mathbb{R}^{md}$, where we defined (with a slight abuse of notation) $\text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}) = (\text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}_i))_{i=1}^m$. Note that

$$\begin{aligned} \|\mathcal{A}(\mathbf{z}, \mathbf{c}^1, \mathbf{c}) - \mathcal{A}(\mathbf{z}, \mathbf{c}^1, \mathbf{c}')\|^2 &= \|\sqrt{\hat{\mathbf{W}}^{-1}}(\mathbf{c} - \mathbf{c}')\|_2^2 \\ + \|\sqrt{\mathbf{I} - \hat{\mathbf{W}}}(\mathbf{c} - \mathbf{c}')\|_2^2 &\leq (1 + \nu^{-1}) \|\mathbf{c} - \mathbf{c}'\|_2^2 \leq 3 \|\mathbf{c} - \mathbf{c}'\|_2^2, \end{aligned}$$

(note that $\nu^{-1} = 1 + \kappa^{-1} \leq 2$) and $\mathcal{A}(\mathbf{z}, \mathbf{c}^1, \mathbf{c}^2)$ is constant with respect to \mathbf{c}^1 , hence Assumption 4 holds with $L_A = \sqrt{3}$.

We next derive L_C and L_Z . Since the proximal mapping is non-expansive [58], it follows that

$$\begin{aligned} \|\mathcal{C}^1(\mathbf{z}, \mathbf{0}) - \mathcal{C}^1(\mathbf{z}', \mathbf{0})\|_2 &= \|\text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}) - \text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}')\|_2 \\ &\leq \|\mathbf{w} - \mathbf{w}'\|_2 \leq \|\mathbf{z} - \mathbf{z}'\|_2, \\ \|\mathcal{C}^2(\mathbf{z}, \mathbf{c}) - \mathcal{C}^2(\mathbf{z}', \mathbf{c}')\|_2 &= \|\hat{\mathbf{W}}(\mathbf{c} - \mathbf{c}') - \gamma(\nabla f(\mathbf{c}) - \nabla f(\mathbf{c}')) - (\mathbf{y} - \mathbf{y}')\|_2 \\ &\leq \|\mathbf{c} - \mathbf{c}'\|_2 + \gamma L \|\mathbf{c} - \mathbf{c}'\|_2 + \|\mathbf{z} - \mathbf{z}'\|_2. \end{aligned}$$

Hence, Assumption 5 holds with $L_C = 1 + \gamma L$ and $L_Z = 1$. Since $\gamma \leq 2/L$, it follows that $L_C = \mathcal{O}(1)$. For the initial conditions, we have

$$\begin{aligned} \|\mathcal{C}^1(\mathbf{z}^0, \mathbf{0})\|_2 &= \|\text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}^0)\|_2 \\ &\stackrel{(a)}{\leq} \|\text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}^0) - \text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}^\infty)\|_2 + \|\text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}^\infty)\|_2 \\ &\stackrel{(b)}{\leq} \|\mathbf{w}^0 - \mathbf{w}^\infty\|_2 + \|\mathbf{x}^\infty\|_2 = \mathcal{O}(\sqrt{md}), \\ \|\mathcal{C}^2(\mathbf{z}^0, \mathbf{0})\|_2 &= \|\mathbf{y}^0 + \gamma \nabla f(\mathbf{0})\|_2 \\ &\leq \|\mathbf{y}^0 - \mathbf{y}^\infty\|_2 + \|\mathbf{y}^\infty\|_2 + \gamma \|\nabla f(\mathbf{0})\|_2 \\ &\leq \|\mathbf{y}^0 - \mathbf{y}^\infty\|_2 + \|\mathbf{y}^\infty\|_2 + \gamma L \|\tilde{\mathbf{x}}^*\|_2 = \mathcal{O}(\sqrt{md}), \end{aligned}$$

where (a) follows from the triangle inequality; and (b) follows from the non-expansive property of the proximal mapping, and $\mathbf{x}^\infty = \text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}^\infty)$ at the fixed point.

Therefore, $L_A \cdot L_Z = \mathcal{O}(1)$, $L_C = \mathcal{O}(1)$, $\|\mathcal{C}^1(\mathbf{z}^0, \mathbf{0})\|_2 = \mathcal{O}(L_Z \sqrt{md})$ and $\|\mathcal{C}^2(\mathbf{z}^0, \mathbf{0})\|_2 = \mathcal{O}(L_Z \sqrt{md})$; hence Assumption 11 holds.

3. (Prox-)NIDS [50]

The update of prox-NIDS solving (P), reads

$$\begin{aligned} \mathbf{x}^k &= \text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}^k), \\ \mathbf{w}^{k+1} &= \hat{\mathbf{W}}(\mathbf{x}^k - \gamma \nabla f(\mathbf{x}^k)) - \mathbf{y}^k, \\ \mathbf{y}^{k+1} &= \mathbf{y}^k + (\mathbf{I} - \hat{\mathbf{W}}) \mathbf{w}^{k+1}, \end{aligned}$$

with $\mathbf{y}^0 = \mathbf{0}$ and $\mathbf{w}^0 \in \mathbb{R}^{md}$.

Prox-NIDS can be cast as (M) with $R = 2$ rounds of communications, using the following:

$$\mathbf{z}^\top = [\mathbf{y}^\top, \mathbf{w}^\top],$$

$$\mathcal{C}_i^1(\mathbf{z}_i^k, \mathbf{0}) = \text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}_i^k) - \gamma \nabla f_i(\text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}_i^k)), \quad (59)$$

$$\mathcal{C}_i^2(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,1}) = \sum_{j \in \mathcal{N}_i} w_{ij} \hat{\mathbf{c}}_j^{k,1} - \mathbf{y}_i^k, \quad (60)$$

$$\begin{aligned} \mathbf{z}_i^{k+1} &= \mathcal{A}_i(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,1}, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,2}) \\ &= \left[\begin{array}{c} \mathbf{y}_i^k + \sum_{j \in \mathcal{N}_i} w_{ij} (\hat{\mathbf{c}}_i^{k,2} - \hat{\mathbf{c}}_j^{k,2}) \\ \hat{\mathbf{c}}_i^{k,2} \end{array} \right]. \end{aligned} \quad (61)$$

We now show that the above instance of (M) satisfies Assumptions 1, 4, 5, and 11.

• On Assumption 1: Using [50, Theorem 18] it is not difficult to check that, if $\gamma = 2/(\mu + L)$, then prox-NIDS satisfies Assumption 1 with

$$\lambda = \max \left\{ \frac{\kappa - 1}{\kappa + 1}, \sqrt{\rho_2(\tilde{\mathbf{W}})} \right\} < 1,$$

and the norm $\|\bullet\|$ defined as

$$\|\mathbf{z}\|^2 = \mathbf{y}^\top \hat{\mathbf{W}}^{-1} (\mathbf{I} - \hat{\mathbf{W}})^\dagger \hat{\mathbf{W}}^{-1} \mathbf{y} + \mathbf{w}^\top \hat{\mathbf{W}}^{-1} \mathbf{w}.$$

Note that $\|\mathbf{z}\|^2 \geq \|\mathbf{z}\|_2^2$.

• On Assumptions 4, 5, and 11: Based on (59), (60), and (61), the mappings \mathcal{A} and \mathcal{C} read

$$\mathcal{A}(\mathbf{z}, \mathbf{c}^1, \mathbf{c}^2) = \begin{bmatrix} \mathbf{y} + (\mathbf{I} - \hat{\mathbf{W}})\mathbf{c}^2 \\ \mathbf{c}^2 \end{bmatrix},$$

$$\mathcal{C}^1(\mathbf{z}, \mathbf{0}) = \text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}) - \gamma \nabla f(\text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w})), \text{ and}$$

$$\mathcal{C}^2(\mathbf{z}, \mathbf{c}) = \hat{\mathbf{W}}\mathbf{c} - \mathbf{y},$$

respectively; and $\mathcal{Z} = \text{span}(\mathbf{I} - \hat{\mathbf{W}}) \times \mathbb{R}^{md}$. Note that

$$\begin{aligned} \|\mathcal{A}(\mathbf{z}, \mathbf{c}_1, \mathbf{c}) - \mathcal{A}(\mathbf{z}, \mathbf{c}_1, \mathbf{c}')\| &= \|\hat{\mathbf{W}}^{-1}(\mathbf{c} - \mathbf{c}')\|_2 \\ &\leq \nu^{-1} \|\mathbf{c} - \mathbf{c}'\|_2, \end{aligned}$$

and $\mathcal{A}(\mathbf{z}, \mathbf{c}^1, \mathbf{c}^2)$ is constant with respect to \mathbf{c}^1 . It follows that Assumption 4 holds with $L_A = \nu^{-1}$.

We next derive L_C and L_Z . Using the non-expansive property of the proximal mapping, it holds

$$\begin{aligned} \|\mathcal{C}^1(\mathbf{z}, \mathbf{0}) - \mathcal{C}^1(\mathbf{z}', \mathbf{0})\|_2 &= \|(\text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}) - \text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}')) \\ &\quad - \gamma(\nabla f(\text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w})) - \nabla f(\text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}')))\|_2 \\ &\leq (1 + \gamma L) \|\text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}) - \text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}')\|_2 \\ &\leq (1 + \gamma L) \|\mathbf{w} - \mathbf{w}'\|_2 \leq (1 + \gamma L) \|\mathbf{z} - \mathbf{z}'\|_2, \\ \|\mathcal{C}^2(\mathbf{z}, \mathbf{c}) - \mathcal{C}^2(\mathbf{z}', \mathbf{c}')\|_2 &= \|\hat{\mathbf{W}}(\mathbf{c} - \mathbf{c}') - (\mathbf{y} - \mathbf{y}')\|_2 \\ &\leq \|\hat{\mathbf{W}}(\mathbf{c} - \mathbf{c}')\|_2 + \|\mathbf{y} - \mathbf{y}'\|_2 \leq \|\mathbf{c} - \mathbf{c}'\|_2 + \|\mathbf{z} - \mathbf{z}'\|_2, \end{aligned}$$

which implies that Assumption 5 holds with $L_C = 1$ and $L_Z = 1 + \gamma L$. Since $\gamma \leq 2/L$, it follows that $L_Z = \mathcal{O}(1)$. For the initial conditions, using $\mathbf{x}^\infty = \text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}^\infty)$ for the fixed point, we have

$$\begin{aligned} \|\mathcal{C}^1(\mathbf{z}^0, \mathbf{0})\|_2 &\leq \|\text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}^0) - \text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}^\infty) + \mathbf{x}^\infty\|_2 \\ &\quad + \gamma \|\nabla f(\text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}^0)) - \nabla f(\text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}^\infty)) + \nabla f(\mathbf{x}^\infty)\|_2 \\ &\leq \|\mathbf{w}^0 - \mathbf{w}^\infty\|_2 + \|\mathbf{x}^\infty\|_2 + \gamma L \|\mathbf{w}^0 - \mathbf{w}^\infty\|_2 \\ &\quad + \gamma L \|\mathbf{x}^0 - \tilde{\mathbf{x}}^*\|_2 = \mathcal{O}(\sqrt{md}) \end{aligned}$$

and $\|\mathcal{C}^2(\mathbf{z}^0, \mathbf{0})\|_2 \leq \|\mathbf{y}^0 - \mathbf{y}^\infty\|_2 + \|\mathbf{y}^\infty\|_2 = \mathcal{O}(\sqrt{md})$.

Therefore, $L_A \cdot L_Z = \mathcal{O}(1)$, $L_C = \mathcal{O}(1)$, $\|\mathcal{C}^1(\mathbf{z}^0, \mathbf{0})\|_2 = \mathcal{O}(L_Z \sqrt{md})$, and $\|\mathcal{C}^2(\mathbf{z}^0, \mathbf{0})\|_2 = \mathcal{O}(L_Z \sqrt{md})$; hence Assumption 11 holds.

4. NIDS [48]

The update of NIDS solving (P) with $r \equiv 0$, reads

$$\begin{aligned} \mathbf{x}^{k+1} &= \mathbf{x}^k - \gamma \nabla f(\mathbf{x}^k) \\ &\quad - \gamma \left[\mathbf{y}^k + c(\mathbf{I} - \check{\mathbf{W}})(\mathbf{x}^k - \gamma \nabla f(\mathbf{x}^k) - \gamma \mathbf{y}^k) \right], \\ \mathbf{y}^{k+1} &= \mathbf{y}^k + c(\mathbf{I} - \check{\mathbf{W}})(\mathbf{x}^k - \gamma \nabla f(\mathbf{x}^k) - \gamma \mathbf{y}^k), \end{aligned}$$

with $\mathbf{x}^0 \in \mathbb{R}^{md}$ and $\mathbf{y}^0 = \mathbf{0}$, where the step-size satisfies $\gamma < 2/L$, $\check{\mathbf{W}} = \tilde{\mathbf{W}} \otimes \mathbf{I}$, $\tilde{\mathbf{W}}$ is a doubly stochastic weight matrix satisfying the communication topology, and $c > 0$ satisfies $c\gamma(\mathbf{I} - \check{\mathbf{W}}) \preceq \mathbf{I}$.

NIDS [48] can be cast as (M) with $R = 1$ round of communications, using the following:

$$\mathbf{z}^\top = [\mathbf{x}^\top, \gamma \mathbf{y}^\top],$$

$$\mathcal{C}_i^1(\mathbf{z}_i^k, \mathbf{0}) = \mathbf{x}_i^k - \gamma \nabla f_i(\mathbf{x}_i^k) - \gamma \mathbf{y}_i^k, \quad (62)$$

$$\begin{aligned} \mathbf{z}_i^{k+1} &= \mathcal{A}_i(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,1}) \\ &= \left[\begin{array}{l} \mathbf{x}_i^k - \gamma \nabla f_i(\mathbf{x}_i^k) - \gamma \mathbf{y}_i^k - \gamma c \sum_{j \in \mathcal{N}_i} \tilde{w}_{ij} (\hat{\mathbf{c}}_i^{k,1} - \hat{\mathbf{c}}_j^{k,1}) \\ \gamma \mathbf{y}_i^k + \gamma c \sum_{j \in \mathcal{N}_i} \tilde{w}_{ij} (\hat{\mathbf{c}}_i^{k,1} - \hat{\mathbf{c}}_j^{k,1}) \end{array} \right]. \end{aligned} \quad (63)$$

We now show that the above instance of (M) satisfies Assumptions 1, 4, 5, and 11.

• On Assumption 1: Using [48] it is not difficult to check that, if $\gamma = 1/L$ and $c = \frac{1}{2\gamma}$, then NIDS satisfies Assumption 1 with

$$\lambda = \max \left\{ \sqrt{1 - \kappa^{-1}}, \sqrt{\frac{1 + \rho_2(\tilde{\mathbf{W}})}{2}} \right\},$$

and the norm $\|\bullet\|$ defined as

$$\|\mathbf{z}\|^2 = \|\mathbf{x}\|_2^2 + \frac{1}{\gamma} \|\gamma \mathbf{y}\|_{c^{-1}(\mathbf{I} - \check{\mathbf{W}})^\dagger}^2,$$

which satisfies $\|\mathbf{z}\|^2 \geq \|\mathbf{z}\|_2^2$. Note that, with \mathbf{W} defined as in (52), we can further bound

$$\lambda \leq \max \left\{ \sqrt{1 - \kappa^{-1}}, \sqrt{\rho_2(\mathbf{W})} \right\} < 1.$$

• On Assumptions 4, 5, and 11: Based on (62) and (63), the mappings \mathcal{A} and \mathcal{C} read

$$\begin{aligned} \mathcal{A}(\mathbf{z}, \mathbf{c}^1) &= \begin{bmatrix} \mathbf{x} - \gamma \nabla f(\mathbf{x}) - \gamma \mathbf{y} - \gamma c(\mathbf{I} - \check{\mathbf{W}})\mathbf{c}^1 \\ \gamma \mathbf{y} + \gamma c(\mathbf{I} - \check{\mathbf{W}})\mathbf{c}^1 \end{bmatrix} \quad \text{and} \\ \mathcal{C}^1(\mathbf{z}, \mathbf{0}) &= \mathbf{x} - \gamma \nabla f(\mathbf{x}) - \gamma \mathbf{y}, \end{aligned}$$

respectively; and $\mathcal{Z} = \mathbb{R}^{md} \times \text{span}(\mathbf{I} - \check{\mathbf{W}})$. Note that

$$\begin{aligned} \|\mathcal{A}(\mathbf{z}, \mathbf{c}) - \mathcal{A}(\mathbf{z}, \mathbf{c}')\|^2 &= (c\gamma)^2 (\mathbf{c} - \mathbf{c}')^\top (\mathbf{I} - \check{\mathbf{W}})^2 (\mathbf{c} - \mathbf{c}') \\ &\quad + c\gamma (\mathbf{c} - \mathbf{c}')^\top (\mathbf{I} - \check{\mathbf{W}}) (\mathbf{c} - \mathbf{c}') \leq 2\|\mathbf{c} - \mathbf{c}'\|_2^2. \end{aligned}$$

It follows that Assumption 4 holds with $L_A = \sqrt{2}$.

We next derive L_C and L_Z . Using the non-expansive property of the proximal mapping, it holds

$$\begin{aligned} \|\mathcal{C}^1(\mathbf{z}, \mathbf{0}) - \mathcal{C}^1(\mathbf{z}', \mathbf{0})\|_2 &= \|(\mathbf{x} - \mathbf{x}') - \gamma [\nabla f(\mathbf{x}) - \nabla f(\mathbf{x}')] - \gamma (\mathbf{y} - \mathbf{y}')\|_2 \\ &\leq \|\mathbf{x} - \mathbf{x}'\|_2 + \|\gamma \mathbf{y} - \gamma \mathbf{y}'\|_2 + \gamma L \|\mathbf{x} - \mathbf{x}'\|_2 \\ &\leq (\sqrt{2} + \gamma L) \|\mathbf{z} - \mathbf{z}'\|_2, \end{aligned}$$

which implies that Assumption 5 holds with $L_C = 1$ and $L_Z = \sqrt{2} + \gamma L$. Since $\gamma = 1/L$, it follows that $L_Z = \mathcal{O}(1)$. For the initial conditions, we have

$$\begin{aligned} \|\mathcal{C}^1(\mathbf{z}^0, \mathbf{0})\|_2 &\leq \|\mathbf{x}^0 - \mathbf{x}^\infty + \mathbf{x}^\infty\|_2 + \gamma \|\nabla f(\mathbf{x}^0) - \nabla f(\tilde{\mathbf{x}}^*)\|_2 \\ &\leq \|\mathbf{x}^0 - \mathbf{x}^\infty\|_2 + \|\mathbf{x}^\infty\|_2 + \gamma L \|\mathbf{x}^0 - \tilde{\mathbf{x}}^*\|_2 = \mathcal{O}(\sqrt{md}). \end{aligned}$$

Therefore, $L_A \cdot L_Z = \mathcal{O}(1)$, $L_C = \mathcal{O}(1)$, $\|\mathcal{C}^1(\mathbf{z}^0, \mathbf{0})\|_2 = \mathcal{O}(L_Z \sqrt{md})$; hence Assumption 11 holds.

5. (Prox-)NEXT [50]

The update of prox-NEXT solving (P) reads

$$\begin{aligned}\mathbf{x}^k &= \text{prox}_{\gamma\mathbf{I},r}(\mathbf{w}^k), \\ \mathbf{w}^{k+1} &= \hat{\mathbf{W}}^2(\mathbf{x}^k - \gamma\nabla f(\mathbf{x}^k)) - \mathbf{y}^k, \\ \mathbf{y}^{k+1} &= \mathbf{y}^k + (\mathbf{I} - \hat{\mathbf{W}})^2\mathbf{w}^{k+1},\end{aligned}$$

with $\mathbf{y}^0 = \mathbf{0}$ and $\mathbf{w}^0 \in \mathbb{R}^{md}$.

Prox-NEXT can be cast as (M) with $R = 4$ rounds of communications, using the following definitions:

$$\begin{aligned}\mathbf{z}^\top &= [\mathbf{y}^\top, \mathbf{w}^\top], \\ \hat{\mathbf{c}}_i^{k,1} &= \mathcal{C}_i^1(\mathbf{z}_i^k, \mathbf{0}) = \text{prox}_{\gamma\mathbf{I},r}(\mathbf{w}_i^k) - \gamma\nabla f_i(\text{prox}_{\gamma\mathbf{I},r}(\mathbf{w}_i^k)),\end{aligned}\quad (64)$$

$$\hat{\mathbf{c}}_i^{k,2} = \mathcal{C}_i^2(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,1}) = \sum_{j \in \mathcal{N}_i} w_{ij} \hat{\mathbf{c}}_j^{k,1}, \quad (65)$$

$$\hat{\mathbf{c}}_i^{k,3} = \mathcal{C}_i^3(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,2}) = \sum_{j \in \mathcal{N}_i} w_{ij} \hat{\mathbf{c}}_j^{k,2} - \mathbf{y}_i^k, \quad (66)$$

$$\hat{\mathbf{c}}_i^{k,4} = \mathcal{C}_i^4(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,3}) = \sum_{j \in \mathcal{N}_i} w_{ij} (\hat{\mathbf{c}}_i^{k,3} - \hat{\mathbf{c}}_j^{k,3}), \quad (67)$$

$$\begin{aligned}\mathbf{z}_i^{k+1} &= \mathcal{A}_i(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,1}, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,2}, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,3}, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,4}) \\ &= \begin{bmatrix} \mathbf{y}_i^k + \sum_{j \in \mathcal{N}_i} w_{ij} (\hat{\mathbf{c}}_i^{k,4} - \hat{\mathbf{c}}_j^{k,4}) \\ \hat{\mathbf{c}}_i^{k,3} \end{bmatrix}.\end{aligned}\quad (68)$$

We now show that the above instance of (M) satisfies Assumptions 1, 4, 5, and 11.

• On Assumption 1: Using [50, Theorem 18] it is not difficult to check that, if $\gamma = 2/(\mu + L)$, then prox-NEXT satisfies Assumption 1 with

$$\lambda = \max \left\{ \frac{\kappa - 1}{\kappa + 1}, \sqrt{1 - (1 - \rho_2(\hat{\mathbf{W}}))^2} \right\} < 1,$$

and the norm $\|\bullet\|$ defined as

$$\|\mathbf{z}\|^2 = \mathbf{y}^\top \hat{\mathbf{W}}^{-2} ((\mathbf{I} - \hat{\mathbf{W}})^2)^\dagger \hat{\mathbf{W}}^{-2} \mathbf{y} + \|\hat{\mathbf{W}}^{-2} \mathbf{w}\|_{\mathbf{I} - (\mathbf{I} - \hat{\mathbf{W}})^2}^2.$$

Note that $\|\mathbf{z}\|^2 \geq \|\mathbf{z}\|_2^2$.

• On Assumptions 4, 5, and 11: Based on (64)-(68), the mappings \mathcal{A} and \mathcal{C} read

$$\mathcal{A}(\mathbf{z}, \mathbf{c}^1, \mathbf{c}^2, \mathbf{c}^3, \mathbf{c}^4) = \begin{bmatrix} \mathbf{y} + (\mathbf{I} - \hat{\mathbf{W}})\mathbf{c}^4 \\ \mathbf{c}^3 \end{bmatrix},$$

$$\begin{aligned}\mathcal{C}^1(\mathbf{z}, \mathbf{0}) &= \text{prox}_{\gamma\mathbf{I},r}(\mathbf{w}) - \gamma\nabla f(\text{prox}_{\gamma\mathbf{I},r}(\mathbf{w})), \quad \mathcal{C}^2(\mathbf{z}, \mathbf{c}) = \hat{\mathbf{W}}\mathbf{c}, \\ \mathcal{C}^3(\mathbf{z}, \mathbf{c}) &= \hat{\mathbf{W}}\mathbf{c} - \mathbf{y}, \quad \text{and} \quad \mathcal{C}^4(\mathbf{z}, \mathbf{c}) = (\mathbf{I} - \hat{\mathbf{W}})\mathbf{c},\end{aligned}$$

respectively; and $\mathcal{Z} = \text{span}(\mathbf{I} - \hat{\mathbf{W}}) \times \mathbb{R}^{md}$. Note that

$$\begin{aligned}\|\mathcal{A}(\mathbf{z}, \mathbf{c}^1, \mathbf{c}^2, \mathbf{c}, \mathbf{c}^4) - \mathcal{A}(\mathbf{z}, \mathbf{c}^1, \mathbf{c}^2, \mathbf{c}', \mathbf{c}^4)\|^2 \\ = (\mathbf{c} - \mathbf{c}')^\top \hat{\mathbf{W}}^{-2} [2\hat{\mathbf{W}}^{-1} - \mathbf{I}] (\mathbf{c} - \mathbf{c}') \\ \leq \nu^{-2} (2\nu^{-1} - 1) \|\mathbf{c} - \mathbf{c}'\|_2^2 \leq \nu^{-4} \|\mathbf{c} - \mathbf{c}'\|_2^2\end{aligned}$$

and

$$\begin{aligned}\|\mathcal{A}(\mathbf{z}, \mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3, \mathbf{c}) - \mathcal{A}(\mathbf{z}, \mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3, \mathbf{c}')\| \\ = \|\hat{\mathbf{W}}^{-2}(\mathbf{c} - \mathbf{c}')\|_2 \leq \nu^{-2} \|\mathbf{c} - \mathbf{c}'\|_2.\end{aligned}$$

Moreover, $\mathcal{A}(\mathbf{z}, \mathbf{c}^1, \mathbf{c}^2, \mathbf{c}^3, \mathbf{c}^4)$ is constant with respect to $\mathbf{c}^1, \mathbf{c}^2$. Therefore, Assumption 4 holds with $L_A = \nu^{-2}$.

We next derive L_C and L_Z . Using the non-expansive property of the proximal operator, it follows that

$$\begin{aligned}\|\mathcal{C}^1(\mathbf{z}, \mathbf{0}) - \mathcal{C}^1(\mathbf{z}', \mathbf{0})\|_2 &= \|(\text{prox}_{\gamma\mathbf{I},r}(\mathbf{w}) - \text{prox}_{\gamma\mathbf{I},r}(\mathbf{w}'))\|_2 \\ &\quad - \gamma(\nabla f(\text{prox}_{\gamma\mathbf{I},r}(\mathbf{w})) - \nabla f(\text{prox}_{\gamma\mathbf{I},r}(\mathbf{w}')))\|_2 \\ &\leq (1 + \gamma L) \|\text{prox}_{\gamma\mathbf{I},r}(\mathbf{w}) - \text{prox}_{\gamma\mathbf{I},r}(\mathbf{w}')\|_2 \\ &\leq (1 + \gamma L) \|\mathbf{w} - \mathbf{w}'\|_2 \leq (1 + \gamma L) \|\mathbf{z} - \mathbf{z}'\|_2, \\ \|\mathcal{C}^2(\mathbf{c}, \mathbf{z}) - \mathcal{C}^2(\mathbf{c}', \mathbf{z}')\|_2 &= \|\hat{\mathbf{W}}(\mathbf{c} - \mathbf{c}')\|_2 \leq \|\mathbf{c} - \mathbf{c}'\|_2, \\ \|\mathcal{C}^3(\mathbf{c}, \mathbf{z}) - \mathcal{C}^3(\mathbf{c}', \mathbf{z}')\|_2 &\leq \|\hat{\mathbf{W}}(\mathbf{c} - \mathbf{c}')\|_2 + \|\mathbf{y} - \mathbf{y}'\|_2 \\ &\leq \|\mathbf{c} - \mathbf{c}'\|_2 + \|\mathbf{z} - \mathbf{z}'\|_2, \\ \|\mathcal{C}^4(\mathbf{c}, \mathbf{z}) - \mathcal{C}^4(\mathbf{c}', \mathbf{z}')\|_2 &= \|(\mathbf{I} - \hat{\mathbf{W}})(\mathbf{c} - \mathbf{c}')\|_2 \leq \|\mathbf{c} - \mathbf{c}'\|_2,\end{aligned}$$

which implies that Assumption 5 holds with $L_C = 1$ and $L_Z = 1 + \gamma L$. Since $\gamma \leq 2/L$, it follows that $L_Z = \mathcal{O}(1)$. For the initial conditions, using the fixed point $\mathbf{x}^\infty = \text{prox}_{\gamma\mathbf{I},r}(\mathbf{w}^\infty)$, we have

$$\begin{aligned}\|\mathcal{C}^1(\mathbf{z}^0, \mathbf{0})\|_2 &\leq \|\text{prox}_{\gamma\mathbf{I},r}(\mathbf{w}^0) - \text{prox}_{\gamma\mathbf{I},r}(\mathbf{w}^\infty) + \mathbf{x}^\infty\|_2 \\ &\quad + \gamma \|\nabla f(\text{prox}_{\gamma\mathbf{I},r}(\mathbf{w}^0)) - \nabla f(\text{prox}_{\gamma\mathbf{I},r}(\mathbf{w}^\infty)) + \nabla f(\mathbf{x}^\infty)\|_2 \\ &\leq (1 + \gamma L) \|\mathbf{w}^0 - \mathbf{w}^\infty\|_2 + \|\mathbf{x}^\infty\|_2 + \gamma L \|\mathbf{x}^\infty - \tilde{\mathbf{x}}^*\|_2 = \mathcal{O}(\sqrt{md})\end{aligned}$$

and $\|\mathcal{C}^2(\mathbf{z}^0, \mathbf{0})\|_2 = \|\mathcal{C}^3(\mathbf{z}^0, \mathbf{0})\|_2 = \|\mathcal{C}^4(\mathbf{z}^0, \mathbf{0})\|_2 = 0$.

Therefore, $L_A \cdot L_Z = \mathcal{O}(1)$, $L_C = \mathcal{O}(1)$, $\|\mathcal{C}^i(\mathbf{z}^0, \mathbf{0})\|_2 = \mathcal{O}(L_Z \sqrt{md})$, $\forall i = 1, \dots, 4$; hence Assumption 11 holds.

6. (Prox-)DIGing [50]

The update of prox-DIGing solving (P), reads

$$\begin{aligned}\mathbf{x}^k &= \text{prox}_{\gamma\mathbf{I},r}(\mathbf{w}^k), \\ \mathbf{w}^{k+1} &= \hat{\mathbf{W}}^2 \mathbf{x}^k - \gamma \nabla f(\mathbf{x}^k) - \mathbf{y}^k, \\ \mathbf{y}^{k+1} &= \mathbf{y}^k + (\mathbf{I} - \hat{\mathbf{W}})^2 \mathbf{w}^{k+1},\end{aligned}$$

with $\mathbf{y}^0 = \mathbf{0}$ and $\mathbf{w}^0 \in \mathbb{R}^{md}$.

Prox-DIGing can be cast as (M) with $R = 4$ rounds of communications, using the following definitions:

$$\mathbf{z}^\top = [\mathbf{y}^\top, \mathbf{w}^\top],$$

$$\hat{\mathbf{c}}_i^{k,1} = \mathcal{C}_i^1(\mathbf{z}_i^k, \mathbf{0}) = \text{prox}_{\gamma\mathbf{I},r}(\mathbf{w}_i^k), \quad (69)$$

$$\hat{\mathbf{c}}_i^{k,2} = \mathcal{C}_i^2(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,1}) = \sum_{j \in \mathcal{N}_i} w_{ij} \hat{\mathbf{c}}_j^{k,1}, \quad (70)$$

$$\begin{aligned}\hat{\mathbf{c}}_i^{k,3} &= \mathcal{C}_i^3(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,2}) \\ &= \sum_{j \in \mathcal{N}_i} w_{ij} \hat{\mathbf{c}}_j^{k,2} - \gamma \nabla f_i(\text{prox}_{\gamma\mathbf{I},r}(\mathbf{w}_i^k)) - \mathbf{y}_i^k,\end{aligned}\quad (71)$$

$$\hat{\mathbf{c}}_i^{k,4} = \mathcal{C}_i^4(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,3}) = \sum_{j \in \mathcal{N}_i} w_{ij} (\hat{\mathbf{c}}_i^{k,3} - \hat{\mathbf{c}}_j^{k,3}), \quad (72)$$

$$\begin{aligned}\mathbf{z}_i^{k+1} &= \mathcal{A}_i(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,1}, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,2}, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,3}, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,4}) \\ &= \begin{bmatrix} \mathbf{y}_i^k + \sum_{j \in \mathcal{N}_i} w_{ij} (\hat{\mathbf{c}}_i^{k,4} - \hat{\mathbf{c}}_j^{k,4}) \\ \hat{\mathbf{c}}_i^{k,3} \end{bmatrix}.\end{aligned}\quad (73)$$

We now show that the above instance of (M) satisfies Assumptions 1, 4, 5, and 11.

• On Assumption 1: Using [50, Theorem 18] and following similar derivations as for (Prox-)EXTRA, it is not difficult

to check that, if $\gamma = \frac{2(\rho_m(\mathbf{W}))^2}{L+\mu(\rho_m(\mathbf{W}))^2}$ and $\nu = \sqrt{\frac{\kappa}{\kappa+1}}$, then prox-DIGing satisfies Assumption 1 with

$$\begin{aligned}\lambda &= \max \left\{ \frac{\kappa - (\rho_m(\mathbf{W}))^2}{\kappa + (\rho_m(\mathbf{W}))^2} \frac{1}{\sqrt{\rho_m(\mathbf{W})(2 - \rho_m(\mathbf{W}))}}, \right. \\ &\quad \left. \sqrt{1 - (1 - \rho_2(\mathbf{W}))^2} \right\} \\ &\leq \max \left\{ \frac{\kappa}{\kappa+1}, \sqrt{1 - (1 - \rho_2(\mathbf{W}))^2} \right\} < 1,\end{aligned}$$

(note that $\rho_m(\mathbf{W}) \geq \nu$) and the norm $\|\bullet\|$ defined as

$$\|\mathbf{z}\|^2 = \frac{1}{2\nu - \nu^2} \left(\mathbf{y}^\top ((\mathbf{I} - \hat{\mathbf{W}})^2)^\dagger \mathbf{y} + \|\mathbf{w}\|_{\mathbf{I} - (\mathbf{I} - \hat{\mathbf{W}})^2}^2 \right).$$

Note that $\|\mathbf{z}\|^2 \geq \|\mathbf{z}\|_2^2$.

• On Assumptions 4, 5, and 11: Based on (69)-(73), the mappings \mathcal{A} and \mathcal{C} read

$$\begin{aligned}\mathcal{A}(\mathbf{z}, \mathbf{c}^1, \mathbf{c}^2, \mathbf{c}^3, \mathbf{c}^4) &= \begin{bmatrix} \mathbf{y} + (\mathbf{I} - \hat{\mathbf{W}})\mathbf{c}^4 \\ \mathbf{c}^3 \end{bmatrix}, \\ \mathcal{C}^1(\mathbf{z}, \mathbf{0}) &= \text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}), \quad \mathcal{C}^2(\mathbf{z}, \mathbf{c}) = \hat{\mathbf{W}}\mathbf{c}, \\ \mathcal{C}^3(\mathbf{z}, \mathbf{c}) &= \hat{\mathbf{W}}\mathbf{c} - \gamma \nabla f(\text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w})) - \mathbf{y}, \quad \text{and} \\ \mathcal{C}^4(\mathbf{z}, \mathbf{c}) &= (\mathbf{I} - \hat{\mathbf{W}})\mathbf{c},\end{aligned}$$

respectively; and $\mathcal{Z} = \text{span}(\mathbf{I} - \hat{\mathbf{W}}) \times \mathbb{R}^{md}$. Note that

$$\begin{aligned}\|\mathcal{A}(\mathbf{z}, \mathbf{c}_1, \mathbf{c}_2, \mathbf{c}, \mathbf{c}^4) - \mathcal{A}(\mathbf{z}, \mathbf{c}_1, \mathbf{c}_2, \mathbf{c}', \mathbf{c}^4)\|^2 \\ = \frac{1}{2\nu - \nu^2} (\mathbf{c} - \mathbf{c}')^\top [\mathbf{I} - (\mathbf{I} - \hat{\mathbf{W}})^2] (\mathbf{c} - \mathbf{c}') \leq \frac{\|\mathbf{c} - \mathbf{c}'\|_2^2}{2\nu - \nu^2},\end{aligned}$$

and

$$\|\mathcal{A}(\mathbf{z}, \mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3, \mathbf{c}) - \mathcal{A}(\mathbf{z}, \mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3, \mathbf{c}')\|^2 = \frac{\|\mathbf{c} - \mathbf{c}'\|_2^2}{2\nu - \nu^2}.$$

Moreover, $\mathcal{A}(\mathbf{z}, \mathbf{c}^1, \mathbf{c}^2, \mathbf{c}^3, \mathbf{c}^4)$ is constant with respect to $\mathbf{c}^1, \mathbf{c}^2$. Therefore, Assumption 4 holds with $L_A = 1/\sqrt{\nu} \leq 2^{1/4}$.

We next derive L_C and L_Z . We have

$$\begin{aligned}\|\mathcal{C}^1(\mathbf{z}, \mathbf{0}) - \mathcal{C}^1(\mathbf{z}', \mathbf{0})\|_2 &\leq \|\mathbf{w} - \mathbf{w}'\|_2 \leq \|\mathbf{z} - \mathbf{z}'\|_2, \\ \|\mathcal{C}^2(\mathbf{z}, \mathbf{c}) - \mathcal{C}^2(\mathbf{z}', \mathbf{c}')\|_2 &= \|\hat{\mathbf{W}}(\mathbf{c} - \mathbf{c}')\|_2 \leq \|\mathbf{c} - \mathbf{c}'\|_2, \\ \|\mathcal{C}^3(\mathbf{z}, \mathbf{c}) - \mathcal{C}^3(\mathbf{z}', \mathbf{c}')\|_2 &\leq \|\mathbf{c} - \mathbf{c}'\|_2 + \gamma L \|\mathbf{w} - \mathbf{w}'\|_2 \\ &\quad + \|\mathbf{y} - \mathbf{y}'\|_2 \leq \|\mathbf{c} - \mathbf{c}'\|_2 + \sqrt{1 + (\gamma L)^2} \|\mathbf{z} - \mathbf{z}'\|_2, \\ \|\mathcal{C}^4(\mathbf{z}, \mathbf{c}) - \mathcal{C}^4(\mathbf{z}', \mathbf{c}')\|_2 &= \|(\mathbf{I} - \hat{\mathbf{W}})(\mathbf{c} - \mathbf{c}')\|_2 \leq \|\mathbf{c} - \mathbf{c}'\|_2,\end{aligned}$$

which implies that Assumption 5 holds with $L_C = 1$ and $L_Z = \sqrt{1 + (\gamma L)^2}$. Since $\gamma \leq 2/L$, it follows that $L_Z = \mathcal{O}(1)$. For the initial conditions, using $\text{prox}_{\gamma \mathbf{I}, r}(\mathbf{w}^\infty) = \mathbf{x}^\infty$ for the fixed point, we have

$$\begin{aligned}\|\mathcal{C}^1(\mathbf{z}^0, \mathbf{0})\|_2 &\leq \|\mathbf{w}^0 - \mathbf{w}^\infty\|_2 + \|\mathbf{x}^\infty\|_2 = \mathcal{O}(\sqrt{md}), \\ \|\mathcal{C}^3(\mathbf{z}^0, \mathbf{0})\|_2 &\leq \gamma L (\|\mathbf{w}^0 - \mathbf{w}^\infty\|_2 + \|\mathbf{x}^\infty - \tilde{\mathbf{x}}^*\|_2) = \mathcal{O}(\sqrt{md}),\end{aligned}$$

and $\|\mathcal{C}^2(\mathbf{z}^0, \mathbf{0})\|_2 = \|\mathcal{C}^4(\mathbf{z}^0, \mathbf{0})\|_2 = 0$. Therefore, $L_A \cdot L_Z = \mathcal{O}(1)$, $L_C = \mathcal{O}(1)$, $\|\mathcal{C}^i(\mathbf{z}^0, \mathbf{0})\|_2 = \mathcal{O}(L_Z \sqrt{md})$, $\forall i = 1, \dots, 4$; hence Assumption 11 holds.

7. Primal-Dual Algorithm [22], [42]

Let $\mathbf{L} = (l_{ij})_{i,j=1}^m$ be the Laplacian matrix associated with the 0-1 adjacency matrix of \mathcal{G} , i.e., $l_{ii} = |\mathcal{N}_i \setminus \{i\}|$, $i \in [m]$; and $l_{ij} = -\mathbb{1}\{(i, j) \in \mathcal{E}\}$, $i \neq j \in [m]$; and $\hat{\mathbf{L}} = \mathbf{L} \otimes \mathbf{I}_d$.

The Primal-Dual algorithm solving (P) with $r \equiv 0$ reads [22], [42]

$$\begin{aligned}\mathbf{x}_i^k &= \arg \min_{\mathbf{x}_i} f_i(\mathbf{x}_i) + \mathbf{x}_i^\top \mathbf{y}_i^k, \\ \mathbf{y}_i^{k+1} &= \mathbf{y}_i^k + \gamma \sum_{j \in \mathcal{N}_i} l_{ij} \mathbf{x}_j^k,\end{aligned}$$

with $\mathbf{y}_i^0 = \mathbf{0}$.

The Primal-Dual algorithm can be cast in the form (M), with $R = 1$ round of communications, using the following:

$$\begin{aligned}\mathbf{z} &= \mathbf{y}, \\ \mathcal{C}_i^1(\mathbf{z}_i^k, \mathbf{0}) &= \arg \min_{\mathbf{x}_i} f_i(\mathbf{x}_i) + \mathbf{x}_i^\top \mathbf{y}_i^k,\end{aligned}\quad (74)$$

$$\mathbf{z}_i^{k+1} = \mathcal{A}_i(\mathbf{z}_i^k, \hat{\mathbf{c}}_{\mathcal{N}_i}^{k,1}) = \mathbf{y}_i^k + \gamma \cdot \sum_{j \in \mathcal{N}_i} l_{ij} \hat{\mathbf{c}}_j^{k,1}. \quad (75)$$

We now show that the above instance of (M) satisfies Assumptions 1, 4, 5, and 11.

• On Assumption 1: Define $\mathbf{M} = \sqrt{\Sigma} \mathbf{Q}$, where $\hat{\mathbf{L}} = \mathbf{Q}^\top \Sigma \mathbf{Q}$ is the eigenvalue decomposition of $\hat{\mathbf{L}}$, with Σ being diagonal with elements sorted in descending order; and let $\hat{\mathbf{M}}$ be the matrix containing the non-zero rows of \mathbf{M} . Using [22] it is not difficult to check that, if $\gamma = \frac{2L\mu}{\mu\rho_{m-1}(\hat{\mathbf{L}}) + L\rho_1(\hat{\mathbf{L}})}$, the Primal-Dual algorithm satisfies Assumption 1 with

$$\lambda = \frac{\frac{\rho_1(\mathbf{L})}{\rho_{m-1}(\mathbf{L})} - \frac{1}{\kappa}}{\frac{\rho_1(\mathbf{L})}{\rho_{m-1}(\mathbf{L})} + \frac{1}{\kappa}} < 1,$$

and the norm $\|\bullet\|$ defined as

$$\|\mathbf{z}\| = \sqrt{\rho_1(\mathbf{L})} \left\| (\mathbf{M} \mathbf{M}^\top)^{-1} \hat{\mathbf{M}} \mathbf{z} \right\|_2.$$

Note that $\|\mathbf{z}\| \geq \|\mathbf{z}\|_2$.

• On Assumptions 4, 5, and 11: Based on (74) and (75), the mappings \mathcal{A} and \mathcal{C} read

$$\mathcal{A}(\mathbf{z}, \mathbf{c}) = \mathbf{z} + \gamma \cdot \hat{\mathbf{L}} \mathbf{c}, \quad \mathcal{C}^1(\mathbf{z}, \mathbf{0}) = \begin{bmatrix} \arg \min_{\mathbf{x}} f_1(\mathbf{x}) + \mathbf{x}^\top \mathbf{z}_1 \\ \vdots \\ \arg \min_{\mathbf{x}} f_m(\mathbf{x}) + \mathbf{x}^\top \mathbf{z}_m \end{bmatrix},$$

respectively; and $\mathcal{Z} = \text{span}(\mathbf{L})$. Note that

$$\mathcal{A}(\mathbf{z}, \mathbf{c}) - \mathcal{A}(\mathbf{z}, \mathbf{c}') = \gamma \hat{\mathbf{L}}(\mathbf{c} - \mathbf{c}').$$

It follows that

$$\begin{aligned}\|\mathcal{A}(\mathbf{z}, \mathbf{c}) - \mathcal{A}(\mathbf{z}, \mathbf{c}')\| &= \gamma \sqrt{\rho_1(\mathbf{L})} \|\hat{\mathbf{M}}(\mathbf{c} - \mathbf{c}')\|_2 \\ &\leq \gamma \rho_1(\mathbf{L}) \|\mathbf{c} - \mathbf{c}'\|_2,\end{aligned}$$

which implies that Assumption 4 holds with $L_A = \gamma \rho_1(\mathbf{L})$. We now derive L_C and L_Z . Since $\mathcal{C}_i^1(\mathbf{z}_i, \mathbf{0}) = \arg \min_{\mathbf{x}_i} f_i(\mathbf{x}_i) + \mathbf{x}_i^\top \mathbf{z}_i$, it follows that $\mathbf{z}_i = -\nabla f_i(\mathcal{C}_i^1(\mathbf{z}_i, \mathbf{0}))$, and strong convexity of f_i implies

$$\|\mathbf{z}_i' - \mathbf{z}_i\|_2 \geq \mu \|\mathcal{C}_i^1(\mathbf{z}_i', \mathbf{0}) - \mathcal{C}_i^1(\mathbf{z}_i, \mathbf{0})\|_2.$$

It readily follows that Assumption 5 holds with $L_C = 0$ and $L_Z = 1/\mu$. For the initial conditions, since $\mathbf{z}^0 = \mathbf{0}$ we have $\mathcal{C}^1(\mathbf{z}^0, \mathbf{0}) = \tilde{\mathbf{x}}^*$ and $\|\mathcal{C}^1(\mathbf{z}^0, \mathbf{0})\|_2 = \|\tilde{\mathbf{x}}^*\|_2 = \mathcal{O}(\sqrt{md})$.

From the expression of γ , it is straightforward to see that $\gamma \leq \frac{2\mu}{\rho_1(\mathbf{L})}$, so that $L_A \cdot L_Z = \mathcal{O}(1)$, $L_C = \mathcal{O}(1)$. Furthermore, for all the objective functions of (P) such that $\mu = \mathcal{O}(1)$, we also have $\|\mathcal{C}^1(\mathbf{z}^0, \mathbf{0})\|_2 = \mathcal{O}(L_Z\sqrt{md})$, hence Assumption 11 holds. For instance, this is the typical case in machine learning problems where a regularization $\mu/2\|\mathbf{x}\|^2$ is added to the objective function to enforce strong convexity, with $\mu = \mathcal{O}(1)$.

REFERENCES

- [1] D. K. Molzahn *et al.*, “A survey of distributed optimization and control algorithms for electric power systems,” *IEEE Trans. Smart Grid*, vol. 8, no. 6, pp. 2941–2962, Nov. 2017.
- [2] A. Nedić, J.-S. Pang, G. Scutari, and Y. Sun, *Multi-Agent Optimization*, 1st ed. Cham, Switzerland: Springer, 2018.
- [3] T. Yang *et al.*, “A survey of distributed optimization,” *Annu. Rev. Control*, vol. 47, pp. 278–305, May 2019.
- [4] R. Bekkerman, M. Bilenko, and J. Langford, *Scaling up Machine Learning: Parallel and Distributed Approaches*. Cambridge, U.K.: Cambridge Univ. Press, 2011.
- [5] X. Lian, C. Zhang, H. Zhang, C.-J. Hsieh, W. Zhang, and J. Liu, “Can decentralized algorithms outperform centralized algorithms? A case study for decentralized parallel stochastic gradient descent,” in *Proc. 31st NeurIPS*, Dec. 2017, pp. 5330–5340.
- [6] A. Koloskova, S. U. Stich, and M. Jaggi, “Decentralized stochastic optimization and gossip algorithms with compressed communication,” in *Proc. 36th ICML*, Jun. 2019, pp. 3478–3487.
- [7] A. Kashyap, T. Başar, and R. Srikant, “Quantized consensus,” *Automatica*, vol. 43, no. 7, pp. 1192–1203, 2007.
- [8] 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, Oct. 2008.
- [9] A. Nedić, A. Olshevsky, A. Ozdaglar, and J. N. Tsitsiklis, “On distributed averaging algorithms and quantization effects,” *IEEE Trans. Autom. Control*, vol. 54, no. 11, pp. 2506–2517, Nov. 2009.
- [10] S. Kar and J. M. F. Moura, “Distributed consensus algorithms in sensor networks: Quantized data and random link failures,” *IEEE Trans. Signal Process.*, vol. 58, no. 3, pp. 1383–1400, Mar. 2010.
- [11] T. Li, M. Fu, L. Xie, and J.-F. Zhang, “Distributed consensus with limited communication data rate,” *IEEE Trans. Autom. Control*, vol. 56, no. 2, pp. 279–292, Feb. 2011.
- [12] R. Rajagopal and M. J. Wainwright, “Network-based consensus averaging with general noisy channels,” *IEEE Trans. Signal Process.*, vol. 59, no. 1, pp. 373–385, Jan. 2011.
- [13] J. Lavaei and R. M. Murray, “Quantized consensus by means of gossip algorithm,” *IEEE Trans. Autom. Control*, vol. 57, no. 1, pp. 19–32, Jan. 2012.
- [14] D. Thanou, E. Kokopoulou, Y. Pu, and P. Frossard, “Distributed average consensus with quantization refinement,” *IEEE Trans. Signal Process.*, vol. 61, no. 1, pp. 194–205, Jan. 2013.
- [15] S. Zhu, Y. C. Soh, and L. Xie, “Distributed parameter estimation with quantized communication via running average,” *IEEE Trans. Signal Process.*, vol. 63, no. 17, pp. 4634–4646, Sep. 2015.
- [16] M. El Chamie, J. Liu, and T. Basar, “Design and analysis of distributed averaging with quantized communication,” *IEEE Trans. Autom. Control*, vol. 61, no. 12, pp. 3870–3884, Dec. 2016.
- [17] H. Li, G. Chen, T. Huang, and Z. Dong, “High-performance consensus control in networked systems with limited bandwidth communication and time-varying directed topologies,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 28, no. 5, pp. 1043–1054, Sep. 2017.
- [18] C.-S. Lee, N. Michelusi, and G. Scutari, “Topology-agnostic average consensus in sensor networks with limited data rate,” in *Proc. 51st Asilomar Conf. Signals, Syst., Comput.*, Oct. 2017, pp. 553–557.
- [19] C.-S. Lee, N. Michelusi, and G. Scutari, “Distributed quantized weight-balancing and average consensus over digraphs,” in *Proc. IEEE Conf. Decis. Control (CDC)*, Dec. 2018, pp. 5857–5862.
- [20] C.-S. Lee, N. Michelusi, and G. Scutari, “Finite rate quantized distributed optimization with geometric convergence,” in *Proc. 52nd Asilomar Conf. Signals, Syst., Comput.*, Oct. 2018, pp. 1876–1880.
- [21] Y. Kajiyama, N. Hayashi, and S. Takai, “Linear convergence of consensus-based quantized optimization for smooth and strongly convex cost functions,” *IEEE Trans. Autom. Control*, vol. 66, no. 3, pp. 1254–1261, Mar. 2021.
- [22] S. Magnússon, H. Shokri-Ghadikolaei, and N. Li, “On maintaining linear convergence of distributed learning and optimization under limited communication,” *IEEE Trans. Signal Process.*, vol. 68, pp. 6101–6116, 2020.
- [23] C.-S. Lee, N. Michelusi, and G. Scutari, “Finite rate distributed weight-balancing and average consensus over digraphs,” *IEEE Trans. Autom. Control*, vol. 66, no. 10, pp. 4530–4545, Oct. 2021.
- [24] D. Kovalev, A. Koloskova, M. Jaggi, P. Richtárik, and S. U. Stich, “A linearly convergent algorithm for decentralized optimization: Sending less bits for free!” in *Proc. 24th AISTATS*, Apr. 2021, pp. 4087–4095.
- [25] X. Liu, Y. Li, R. Wang, J. Tang, and M. Yan, “Linear convergent decentralized optimization with compression,” in *Proc. 9th ICLR*, May 2021.
- [26] D. Alistarh, D. Grubic, J. Li, R. Tomioka, and M. Vojnovic, “QSGD: Communication-efficient SGD via gradient quantization and encoding,” in *Proc. 31st NeurIPS*, Dec. 2017, pp. 1709–1720.
- [27] D. Alistarh, T. Hoefler, M. Johansson, N. Konstantinov, S. Khirirat, and C. Renggli, “The convergence of sparsified gradient methods,” in *Proc. 32nd NeurIPS*, Dec. 2018, pp. 5973–5983.
- [28] P. Karimireddy, Q. Rebjock, S. Stich, and M. Jaggi, “Error feedback fixes SignSGD and other gradient compression schemes,” in *Proc. 36th ICML*, Jul. 2019, pp. 3252–3261.
- [29] S. U. Stich, J.-B. Cordonnier, and M. Jaggi, “Sparsified SGD with memory,” in *Proc. 32nd NeurIPS*, Dec. 2018, pp. 4447–4458.
- [30] H. Tang, S. Gan, C. Zhang, T. Zhang, and J. Liu, “Communication compression for decentralized training,” in *Proc. 32nd NeurIPS*, Dec. 2018, pp. 7652–7662.
- [31] X. Zhang, J. Liu, Z. Zhu, and E. S. Bentley, “Compressed distributed gradient descent: Communication-efficient consensus over networks,” in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, Apr. 2019, pp. 2431–2439.
- [32] S. Zheng, Z. Huang, and J. Kwok, “Communication-efficient distributed blockwise momentum SGD with error-feedback,” in *Proc. 33rd NeurIPS*, Dec. 2019.
- [33] A. Beznosikov, S. Horváth, P. Richtárik, and M. Safaryan, “On biased compression for distributed learning,” 2020, *arXiv:2002.12410*.
- [34] E. Gorbunov, F. Hanzely, and P. Richtárik, “A unified theory of SGD: Variance reduction, sampling, quantization and coordinate descent,” in *Proc. 23rd AISTATS*, Apr. 2020, pp. 680–690.
- [35] S. U. Stich, “On communication compression for distributed optimization on heterogeneous data,” 2020, *arXiv:2009.02388*.
- [36] H. Taheri, A. Mokhtari, H. Hassani, and R. Pedarsani, “Quantized decentralized stochastic learning over directed graphs,” in *Proc. 37th ICML*, Jul. 2020, pp. 9324–9333.
- [37] F. Haddadpour, M. M. Kamani, A. Mokhtari, and M. Mahdavi, “Federated learning with compression: Unified analysis and sharp guarantees,” in *Proc. 24th AISTATS*, Apr. 2021, pp. 2350–2358.
- [38] Y. Liao, Z. Li, K. Huang, and S. Pu, “Compressed gradient tracking methods for decentralized optimization with linear convergence,” 2021, *arXiv:2103.13748*.
- [39] J. Zhang, K. You, and L. Xie, “Innovation compression for communication-efficient distributed optimization with linear convergence,” 2021, *arXiv:2105.06697*.
- [40] P. Di Lorenzo and G. Scutari, “NEXT: In-network nonconvex optimization,” *IEEE Trans. Signal Inf. Process. Netw.*, vol. 2, no. 2, pp. 120–136, Jun. 2016.
- [41] Y. Sun, G. Scutari, and A. Daneshmand, “Distributed optimization based on gradient tracking revisited: Enhancing convergence rate via surrogation,” *SIAM J. Optim.*, vol. 32, no. 2, pp. 354–385, 2022.
- [42] C. A. Uribe, S. Lee, A. Gasnikov, and A. Nedić, “A dual approach for optimal algorithms in distributed optimization over networks,” *Optim. Methods Softw.*, vol. 36, no. 1, pp. 171–210, Jan. 2021.
- [43] W. Shi, Q. Ling, G. Wu, and W. Yin, “Extra: An exact first-order algorithm for decentralized consensus optimization,” *J. Optim.*, vol. 25, no. 2, pp. 944–966, May 2015.

- [44] J. Xu, S. Zhu, Y. C. Soh, and L. Xie, "Convergence of asynchronous distributed gradient methods over stochastic networks," *IEEE Trans. Autom. Control*, vol. 63, no. 2, pp. 434–448, Feb. 2018.
- [45] A. Nedić, A. Olshevsky, and W. Shi, "Achieving geometric convergence for distributed optimization over time-varying graphs," *SIAM J. Optim.*, vol. 27, no. 4, pp. 2597–2633, 2017.
- [46] G. Qu and N. Li, "Harnessing smoothness to accelerate distributed optimization," *IEEE Trans. Control Netw. Syst.*, vol. 5, no. 3, pp. 1245–1260, Sep. 2018.
- [47] A. Nedic and A. Ozdaglar, "Distributed subgradient methods for multi-agent optimization," *IEEE Trans. Autom. Control*, vol. 54, no. 1, pp. 48–61, Jan. 2009.
- [48] Z. Li, W. Shi, and M. Yan, "A decentralized proximal-gradient method with network independent step-sizes and separated convergence rates," *IEEE Trans. Signal Process.*, vol. 67, no. 17, pp. 4494–4506, Sep. 2019.
- [49] G. Scutari and Y. Sun, "Distributed nonconvex constrained optimization over time-varying digraphs," *Math. Program.*, vol. 176, nos. 1–2, pp. 497–544, Feb. 2019.
- [50] J. Xu, Y. Tian, Y. Sun, and G. Scutari, "Distributed algorithms for composite optimization: Unified framework and convergence analysis," *IEEE Trans. Signal Process.*, vol. 69, pp. 3555–3570, 2021.
- [51] Y. Nesterov, *Introductory Lectures on Convex Optimization: A Basic Course*, 1st ed. New York, NY, USA: Springer, 2014.
- [52] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Belmont, MA, USA: Athena Scientific, 1989.
- [53] D. Alistarh and J. H. Korhonen, "Towards tight communication lower bounds for distributed optimisation," in *Proc. 35th NeurIPS*, 2021, pp. 7254–7266.
- [54] P. Davies, V. Gurunanthan, N. Moshrefi, S. Ashkboos, and D. Alistarh, "New bounds for distributed mean estimation and variance reduction," in *Proc. 9th ICLR*, May 2021.
- [55] A. Agarwal, S. Negahban, and M. J. Wainwright, "Fast global convergence rates of gradient methods for high-dimensional statistical recovery," in *Proc. 24th NeurIPS*, Dec. 2010, pp. 37–45.
- [56] L. Xiao and S. Boyd, "Fast linear iterations for distributed averaging," *Syst. Control Lett.*, vol. 53, no. 1, pp. 65–78, 2004.
- [57] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
- [58] J.-J. Moreau, "Proximité et dualité dans un espace Hilbertien," *Bull. De La Société Mathématique De France*, vol. 93, pp. 273–299, Apr. 1965.

Nicolò Michelusi (Senior Member, IEEE) received the B.Sc. and M.Sc. degrees (Hons.) from the University of Padova, Italy, in 2006, and 2009, respectively, the M.Sc. degree in telecommunications engineering from the Technical University of Denmark, Denmark, in 2009, as part of the T.I.M.E. double degree program, and Ph.D. degree from the University of Padova in 2013. From 2013 to 2015, he was a Post-Doctoral Research Fellow with the Ming-Hsieh Department of Electrical Engineering, University of Southern California, Los Angeles, CA, USA, and from 2016 to 2020, he was an Assistant Professor with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, USA. He is currently an Assistant Professor with the School of Electrical, Computer and Energy Engineering, Arizona State University, Tempe, AZ, USA. His research interests include 5G wireless networks, millimeter-wave communications, stochastic optimization, distributed optimization, and federated learning over wireless. He received the NSF CAREER Award in 2021. He was the Co-Chair for the Distributed Machine Learning and Fog Network Workshop at the IEEE INFOCOM 2021, the Wireless Communications Symposium at the IEEE GLOBECOM 2020, the IoT, M2M, Sensor Networks, and Ad-Hoc Networking Track at the IEEE VTC 2020, and the Cognitive Computing and Networking Symposium at the ICNC 2018. He served as an Associate Editor for the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS from 2016 to 2021 and a reviewer for several IEEE journals.

Gesualdo Scutari (Fellow, IEEE) received the electrical engineering and Ph.D. degrees (Hons.) from the University of Rome "La Sapienza" Rome, Italy, in 2001 and 2005, respectively. He is currently a Professor with the School of Industrial Engineering, Purdue University, West Lafayette, IN, USA. His research interests include continuous and distributed optimization, equilibrium programming, and their applications to signal processing and machine learning. He was a recipient of the 2013 NSF CAREER Award, the 2015 IEEE Signal Processing Society Young Author Best Paper Award, and the 2020 IEEE Signal Processing Society Best Paper Award. He is a Senior Area Editor of the IEEE TRANSACTIONS ON SIGNAL PROCESSING and an Associate Editor of *SIAM Journal on Optimization*.

Chang-Shen Lee received the B.Sc. degree in electrical engineering and computer science and the M.Sc. degree in communications engineering from the National Chiao Tung University, Hsinchu, Taiwan, in 2012 and 2014, respectively, and the Ph.D. degree in electrical and computer engineering from Purdue University, West Lafayette, IN, USA, in 2021. He is currently a Senior Software Engineer at Bloomberg, New York City, NY, USA.