Quantized Decentralized Consensus Optimization

Amirhossein Reisizadeh, Aryan Mokhtari, Hamed Hassani, and Ramtin Pedarsani

Abstract—We consider the problem of decentralized consensus optimization, where the sum of n convex functions are minimized over n distributed agents that form a connected network. In particular, we consider the case that the communicated local decision variables among nodes are quantized in order to alleviate the communication bottleneck in distributed optimization. We propose the Quantized Decentralized Gradient Descent (QDGD) algorithm, in which nodes update their local decision variables by combining the quantized information received from their neighbors with their local information. We prove that under standard strong convexity and smoothness assumptions for local cost functions, QDGD achieves a vanishing mean solution error. To the best of our knowledge, this is the first algorithm that achieves vanishing consensus error in the presence of quantization noise. Moreover, we provide simulation results that show tight agreement between our derived theoretical convergence rate and the experimental results.

I. INTRODUCTION

Distributed optimization of a sum of convex functions has a variety of applications in different areas including decentralized control systems [1], wireless systems [2], sensor networks [3], networked multiagent systems [4], multirobot networks [5], and large scale machine learning [6]. In such problems, one aims to solve a consensus optimization problem to minimize $f(\mathbf{x}) = \sum_{i=1}^{n} f_i(\mathbf{x})$ cooperatively over n nodes or agents that form a connected network. The function $f_i(\cdot)$ represents the local cost function of node i that is only known by this node.

Distributed optimization has been largely studied in the literature starting from seminal works in the 80s [7], [8]. Since then, various algorithms have been proposed to address decentralized consensus optimization in multiagent systems. The most commonly used algorithms are decentralized gradient descent or gradient projection method [9]–[12], distributed alternating direction method of multipliers (ADMM) [13]–[15], decentralized dual averaging [16], [17], and distributed Newton optimization method [18], [19]. Furthermore, the decentralized consensus optimization problem has been considered in online or dynamic settings, where the dynamic cost function becomes an online regret function [20], [21].

Amirhossein Reisizadeh and Ramtin Pedarsani are with the Department of Electrical and Computer Engineering at University of California, Santa Barbara `reisizadeh@ucsb.edu`, `ramtin@ece.ucsb.edu`.

Aryan Mokhtari is with the Laboratory for Information and Decision Systems at Massachusetts Institute of Technology `aryanm@mit.edu`.

Hamed Hassani is with the Department of Electrical and Systems Engineering at University of Pennsylvania `hassani@seas.upenn.edu`.

A major bottleneck in achieving fast convergence in decentralized consensus optimization is limited communication bandwidth among nodes. As the dimension of input data increases (which is the current trend in large-scale distributed machine learning), a considerable amount of information must be exchanged among nodes, over many iterations of the consensus algorithm. This causes a significant communication bottleneck that can substantially slow down the convergence time of the algorithm [22], [23].

Quantized communication for the agents is brought into the picture for bounded and stable control systems [24]. Furthermore, consensus distributed averaging algorithms are studied under discretized message passing [25]. Motivated by the energy and bandwidth-constrained wireless sensor networks, the work in [26] proposes distributed optimization algorithms under quantized variables and guarantees convergence within a non-vanishing error. Deterministic quantization has been considered in distributed averaging algorithms [27] where the iterations converge to a neighborhood of the average of initials. However, randomized quantization schemes are shown to achieve the average of initials, in expectation [28]. The work in [29] also considers a consensus distributed optimization problem over a cooperative network of agents restricted to quantized communication. The proposed algorithm guarantees convergence to the optima within an error which depends on the network size and the number of quantization levels. Aligned with the communication bottleneck described earlier, [30] provides a quantized distributed load balancing scheme that converges to a set of desired states while the nodes are constrained to remain under maximum load capacities.

More recently, 1–Bit SGD [22] was introduced in which at each time step, the agents sequentially quantize their local gradient vectors by entry-wise signs while contributing the quantization error induced in previous iteration. Moreover, in [31], the authors propose the Quantized-SGD (QSGD), a class of compression scheme algorithms that is based on a stochastic and unbiased quantizer of the vector to be transmitted. QSGD provably provides convergence guarantees, as well a good practical performance. Recently, a different line of work has proposed the use of coding theoretic techniques to alleviate the communication bottleneck in distributed computation [32]–[35].

In this paper, our goal is to analyze the quantized decentralized consensus optimization problem, where node i transmits a quantized version of its local decision variable $Q(\mathbf{x}_i)$ to the neighboring nodes instead of the exact decision variable $\mathbf{x}_i$. Motivated by the stochastic quantizer proposed in [31], we consider unbiased and variance

bounded random quantizers $Q(\cdot)$, i.e. $\mathbb{E}\left[Q(\mathbf{x})|\mathbf{x}\right] = \mathbf{x}$ and $\mathbb{E}\left[\|Q(\mathbf{x}) - \mathbf{x}\|^2|\mathbf{x}\right] \leq \sigma^2$ for some fixed constant σ^2.

Our main contribution is to propose a Quantized Decentralized Gradient Descent (QDGD) method, which involves a novel way of updating the local decision variables by combining the quantized message received from the neighbors and the local information such that proper averaging is performed over the local decision variable and the neighbors' quantized vectors. We prove that under standard strong convexity and smoothness assumptions, for any unbiased and variance bounded quantizer, QDGD achieves a vanishing mean solution error: for all $i = 1, \ldots, n$ we obtain that $\mathbb{E}\left[\|\mathbf{x}_{i,T} - \mathbf{x}^*\|^2\right] \leq \mathcal{O}\left(\frac{1}{T^{1/4}}\right)$ for sufficiently large T, where $\mathbf{x}_{i,T}$ is the local decision variable of node i at iteration T and $\mathbf{x}^*$ is the optimal solution. To the best of our knowledge, this is the first decentralized gradient-based algorithm that achieves vanishing consensus error in the presence of non-vanishing quantization noise. We further provide simulation results that corroborate our theoretical results.

Notation. In this paper, we denote by $[n]$ the set $\{1, \cdots, n\}$ for any natural number $n \in \mathbb{N}$. The gradient of a function $f(\mathbf{x})$ is denoted by $\nabla f(\mathbf{x})$. For non-negative functions g and h of t, we denote $g(t) = \mathcal{O}(h(t))$ if there exist $t_0 \in \mathbb{N}$ and constant $c > 0$ such that $g(t) \leq ch(t)$ for all $t \geq t_0$. we use $\lceil a \rceil$ to indicate the least integer greater than or equal to a.

Paper Organization. The rest of the paper is organized as follows. In Section II, we precisely formulate the quantized decentralized consensus optimization problem. We provide the description of the Quantized Decentralized Gradient Descent algorithm in Section III. The main theorem of the paper is stated and proved in Section IV. We provide numerical studies in Section V. Finally, we conclude the paper and discuss future directions in Section VI.

II. PROBLEM FORMULATION

In this section we formally define the consensus optimization problem that we aim to solve. Consider a set of n nodes that communicate over a connected and undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where $\mathcal{V} = \{1, \cdots, n\}$ and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ denote the set of nodes and edges, respectively. We assume that nodes are only allowed to exchange information with their neighbors and use the notation $\mathcal{N}_i$ for the set of node i's neighbors. In our setting, we assume that each node i has access to a local convex function $f_i : \mathbb{R}^p \to \mathbb{R}$, and nodes in the network cooperate to minimize the aggregate objective function $f : \mathbb{R}^p \to \mathbb{R}$ taking values $f(\mathbf{x}) = \sum_{i=1}^n f_i(\mathbf{x})$. In other words, nodes aim to solve the optimization problem

$$\min_{\mathbf{x} \in \mathbb{R}^p} f(\mathbf{x}) = \min_{\mathbf{x} \in \mathbb{R}^p} \sum_{i=1}^n f_i(\mathbf{x}). \qquad (1)$$

We assume the local objective functions f_i are strongly convex and smooth, and, therefore, the aggregate function f is also strongly convex and smooth. In the rest of the paper, we use $\mathbf{x}^*$ to denote the unique minimizer of Problem (1).

In decentralized settings, nodes have access to a single summand of the global objective function f and to reach the optimal solution $\mathbf{x}^*$ communication with neighboring nodes is inevitable. To be more precise, nodes need to minimize their local objective functions, while they ensure that their local decision variables are equal to their neighbors'. This interpretation leads to an equivalent formulation of Problem (1). If we define $\mathbf{x}_i$ as the decision variable of node i, the alternative formulation of Problem (1) can be written as

$$\min_{\mathbf{x}_1, \ldots, \mathbf{x}_n \in \mathbb{R}^p} \sum_{i=1}^n f_i(\mathbf{x}_i)$$
$$\text{s.t.} \quad \mathbf{x}_i = \mathbf{x}_j, \qquad \text{for all } i, \ j \in \mathcal{N}_i. \qquad (2)$$

Since we assume that the underlying network is a connected graph, the constraint in (2) implies that any feasible solution should satisfy $\mathbf{x}_1 = \cdots = \mathbf{x}_n$. Under this condition the objective function values in (1) and (2) are equivalent. Hence, it follows that the optimal solutions of Problem (2) are equal to the optimal solution of Problem (1), i.e., if we denote $\{\mathbf{x}_i^*\}_{i=1}^n$ as the optimal solutions of Problem (2) it holds that $\mathbf{x}_1^* = \cdots = \mathbf{x}_n^* = \mathbf{x}^*$. Therefore, we proceed to solve Problem (2) which is naturally formulated for decentralized optimization in lieu of Problem (1).

The problem formulation in (2) suggests that each node i should minimize its local objective function f_i while keeping its decision variable $\mathbf{x}_i$ close to the decision variable $\mathbf{x}_j$ of its neighbors $j \in \mathcal{N}_i$. This goal can be achieved by exchanging local variables $\mathbf{x}_i$ among neighboring nodes to enforce consensus on the decision variables. Indeed, exchange of updated local vectors between the distributed nodes induces a potentially heavy communication load on the shared bus. To address this issue, we assume that each node provides a randomly quantized variant of its local updated variable to the neighboring nodes. That is, if we denote by $\mathbf{x}_i$ the decision variable of node i, then the corresponding quantized variant $\mathbf{z}_i = Q(\mathbf{x}_i)$ is communicated to the neighboring nodes, $\mathcal{N}_i$. Exchanging quantized vectors $\mathbf{z}_i$ instead of the true vectors $\mathbf{x}_i$ indeed reduces the communication burden at the cost of injecting noise to the information received by the nodes in the network. The main challenge in this setting is to ensure that nodes can still converge to the optimal solution of Problem (2), while they only have access to a quantized variant of their neighbors' true decision variables.

III. QDGD ALGORITHM

In this section, we propose a quantized gradient based method to solve the decentralized optimization problem in (2) and consequently the original problem in (1) in a fully decentralized fashion. To do so, consider $\mathbf{x}_{i,t}$ as the decision variable of node i at step t and $\mathbf{z}_{i,t} = Q(\mathbf{x}_{i,t})$ as the quantized version of the vector $\mathbf{x}_{i,t}$. In the proposed Quantized Decentralized Gradient Descent (QDGD) method, nodes update their local decision variables by combining the quantized information received from their neighbors with their local information. To formally state the update of QDGD, we first define w_{ij} as the weight that node i assigns

---

Algorithm 1 QDGD at node i

Require: Weights $\{w_{ij}\}_{j=1}^n$, total iterations T
  1: Set $\mathbf{x}_{i,0} = 0$
  2: **for** $t = 0, \cdots, T-1$ **do**
  3:    Send $\mathbf{z}_{i,t} = Q(\mathbf{x}_{i,t})$ to $j \in \mathcal{N}_i$ and receive $\mathbf{z}_{j,t}$
  4:    Compute $\mathbf{x}_{i,t+1}$ according to the update in (3)
  5: **end for**
  6: **return** $\mathbf{x}_{i,T}$

---

to node j. If nodes i and j are not neighbors then $w_{ij} = 0$, and if they are neighbors the weight $w_{ij} \geq 0$ is nonnegative. At each time step t, each node i sends its quantized $\mathbf{z}_{i,t}$ variant of its local vector $\mathbf{x}_{i,t}$ to its neighbors $j \in \mathcal{N}_i$ and receives their corresponding vectors $\mathbf{z}_{j,t}$. Then, using the received information it updates its local decision variable according to the update

$$\mathbf{x}_{i,t+1} = (1 - \varepsilon + \varepsilon w_{ii})\mathbf{x}_{i,t} + \varepsilon \sum_{j \in \mathcal{N}_i} w_{ij}\mathbf{z}_{j,t} - \alpha\varepsilon\nabla f_i(\mathbf{x}_{i,t}), \tag{3}$$

where ε and α are positive step-sizes. The update of QDGD in (3) shows that the updated decision variable $\mathbf{x}_{i,t+1}$ is evaluated by proper averaging over the local decision variable $\mathbf{x}_{i,t}$ and neighbors quantized vectors $\mathbf{z}_{j,t}$, and descending through the negative local gradient $\nabla f_i(\mathbf{x}_{i,t})$ with a proper stepsize. Note that quantized decision variables of the neighboring nodes contribute to the descent direction proportionally to step-size ε, unlike the noiseless local gradient which is scaled by $\alpha\varepsilon$. The steps of the proposed QDGD method are summarized in Algorithm 1.

Remark 1. The proposed QDGD algorithm can be interpreted as a variant of the decentralized (sub)gradient descent (DGD) method [9], [10] for quantized decentralized optimization (see Section IV). Note that the vanilla DGD method converges to a neighborhood of the optimal solution in the presence of quantization noise where the radius of convergence depends on the variance of quantization error [9], [10], [26], [29]. QDGD improves the inexact convergence of DGD by modifying the contribution of quantized information received from neighboring noise as described in update (3). In particular, as we show in Theorem 1, the sequence of iterates generated by QDGD converges to the optimal solution of Problem (1) in expectation.

Moreover, the proposed QDGD algorithm does not restrict the quantizer, except for few customary conditions. However, design of efficient quantizers has been taken into consideration. For instance, consider a low-precision representation specified by $\gamma \in \mathbb{R}$ and $b \in \mathbb{N}$. The range representable by scale factor γ and b bits is $\{-\gamma \cdot 2^{b-1}, \cdots, -\gamma, 0, \gamma, \cdots, \gamma \cdot (2^b - 1)\}$. For any $k\gamma \leq x < (k+1)\gamma$ in the representable range, the low-precision quantizer outputs

$$Q_{(\gamma,b)}(x) = \begin{cases} k\gamma & \text{w.p. } 1 - \frac{x - k\gamma}{\gamma}, \\ (k+1)\gamma & \text{w.p. } \frac{x - k\gamma}{\gamma}. \end{cases} \tag{4}$$

For any x in the range, the quantizer is unbiased

and variance bounded, i.e. $\mathbb{E}\left[Q_{(\gamma,b)}(x)\right] = x$ and $\mathbb{E}\left[\left\|Q_{(\gamma,b)}(x) - x\right\|^2\right] \leq \frac{\gamma^2}{4}$.

IV. CONVERGENCE ANALYSIS

In this section, we prove that for sufficiently large number of iterations, the sequence of local iterates generated by QDGD converges to an arbitrarily precise approximation of the optimal solution of Problem (1). The following assumptions hold through out the analysis of the algorithm.

Assumption 1. Local objective functions f_i are differentiable and smooth with parameter L, i.e.,

$$\|\nabla f_i(\mathbf{x}) - \nabla f_i(\mathbf{y})\| \leq L\|\mathbf{x} - \mathbf{y}\|, \tag{5}$$

for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^p$.

Assumption 2. Local objective functions f_i are strongly convex with parameter μ, i.e.,

$$\langle \nabla f_i(\mathbf{x}) - \nabla f_i(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle \geq \mu\|\mathbf{x} - \mathbf{y}\|^2, \tag{6}$$

for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^p$.

Assumption 3. The random quantizer $Q(\cdot)$ is unbiased and has a bounded variance, i.e.,

$$\mathbb{E}\left[Q(\mathbf{x})|\mathbf{x}\right] = \mathbf{x}, \quad \text{and} \quad \mathbb{E}\left[\|Q(\mathbf{x}) - \mathbf{x}\|^2 |\mathbf{x}\right] \leq \sigma^2, \tag{7}$$

for any $\mathbf{x} \in \mathbb{R}^p$; and quantizations are carried out independently on distributed nodes.

Assumption 4. The weight matrix $W \in \mathbb{R}^{n \times n}$ with entries w_{ij} satisfies the following conditions

$$W = W^\top, \quad W\mathbf{1} = \mathbf{1}, \quad \text{null}(I - W) = \text{span}(\mathbf{1}). \tag{8}$$

The conditions in Assumptions 1 and 2 imply that the global objective function f is strongly convex with parameter μ and its gradients are Lipschitz continuous with constant L. Assumption 3 poses two customary conditions on the quantizer, that are unbiasedness and variance boundedness. Assumption 4 implies that weight matrix W is symmetric and doubly stochastic. The largest eigenvalue of W is $\lambda_1(W) = 1$ and all the eigenvalues belong to $(-1, 1]$, i.e. the ordered sequence of eigenvalues of W are $1 = \lambda_1(W) \geq \lambda_2(W) \geq \cdots \geq \lambda_n(W) > -1$. We denote by $1 - \beta$ the spectral gap associated to the stochastic matrix W, where $\beta = \max\{|\lambda_2(W)|, |\lambda_n(W)|\}$ is the second largest magnitude of the eigenvalues of matrix W. It is also customary to assume $\text{rank}(I - W) = n - 1$ such that $\text{null}(I - W) = \text{span}(\mathbf{1})$.

In the following theorem we show that the local iterations generated by QDGD converge to the global optima, as close as desired.

Theorem 1. Consider the distributed consensus optimization Problem (1) and suppose Assumptions 1– 4 hold. Then, for each node i, the expected deviation of the output of

Algorithm 1 from the solution to Problem (1) is upper bounded by

$$\mathbb{E}\left[\left\|\mathbf{x}_{i,T} - \mathbf{x}^*\right\|^2\right] \le \mathcal{O}\left(\frac{1}{T^{1/4}}\right), \qquad (9)$$

for $\varepsilon = \frac{c_1}{T^{1/2}}$, $\alpha = \frac{c_2}{T^{1/4}}$ and $T \ge T_0$, where c_1, c_2 and T_0 are positive constants independent of T.

Remark 2. Theorem 1 demonstrates that the proposed QDGD provides an approximation solution with vanishing deviation from the optimal solution, despite the fact that the quantization noise does not vanish by iteration.

To analyze the proposed method, we start by rewriting the update rule (3) as follows

$$\mathbf{x}_{i,t+1} = \mathbf{x}_{i,t} - \varepsilon_t\left((1 - w_{ii})\mathbf{x}_{i,t} - \sum_{j \ne i} w_{ij}\mathbf{z}_{j,t} + \alpha\nabla f_i(\mathbf{x}_{i,t})\right). \qquad (10)$$

The next step is to write the update (10) in a matrix form. To do so, we define the function $F : \mathbb{R}^{np} \to \mathbb{R}$ as $F(\mathbf{x}) = \sum_{i=1}^n f_i(\mathbf{x}_i)$ where $\mathbf{x}_i \in \mathbb{R}^p$ and $\mathbf{x} = [\mathbf{x}_1; \cdots; \mathbf{x}_n] \in \mathbb{R}^{np}$. It is easy to verify that the gradient of the function F is the concatenation of local gradients evaluated at the local variable, that is $\nabla F(\mathbf{x}_t) = [\nabla f_1(\mathbf{x}_{1,t}); \cdots; \nabla f_n(\mathbf{x}_{n,t})]$. We also define the matrix $\mathbf{W} = W \otimes I \in \mathbb{R}^{np \times np}$ as the Kronecker product of the weight matrix $W \in \mathbb{R}^{n \times n}$ and the identity matrix $I \in \mathbb{R}^{p \times p}$. Similarly, define $\mathbf{W}_D = W_D \otimes I \in \mathbb{R}^{np \times np}$, where $W_D = [w_{ii}] \in \mathbb{R}^{n \times n}$ denotes the diagonal matrix of the entries on the main diagonal of W. For the sake of consistency, we denote by the boldface $\mathbf{I}$ the identity matrix of size np. According to above definitions, we can write the concatenated version of (10) as follows,

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \varepsilon\left((\mathbf{W}_D - \mathbf{W})\mathbf{z}_t + (\mathbf{I} - \mathbf{W}_D)\mathbf{x}_t + \alpha\nabla F(\mathbf{x}_t)\right). \qquad (11)$$

As we discussed in Section II, the distributed consensus optimization Problem (1) can be equivalently written as Problem (2). The constraint in the latter restricts the feasible set to the consensus vectors, that is $\{\mathbf{x} = [\mathbf{x}_1; \cdots; \mathbf{x}_n] : \mathbf{x}_1 = \cdots = \mathbf{x}_n\}$. According to the discussion on rank of the weight matrix W, the null space of the matrix $I - W$ is $\mathrm{null}(I - W) = \mathrm{span}(\mathbf{1})$. Hence, the null space of $\mathbf{I} - \mathbf{W}$ is the set of all consensus vectors, i.e. $\mathbf{x} \in \mathbb{R}^{np}$ is feasible for problem (2) if and only if $(\mathbf{I} - \mathbf{W})\mathbf{x} = 0$, or equivalently $(\mathbf{I} - \mathbf{W})^{1/2}\mathbf{x} = 0$. Therefore, the alternative Problem (2) can be compactly represented as the following linearly-constrained problem,

$$\min_{\mathbf{x} \in \mathbb{R}^{np}} \quad F(\mathbf{x}) = \sum_{i=1}^n f_i(\mathbf{x}_i) \qquad (12)$$
$$\text{subject to} \quad (\mathbf{I} - \mathbf{W})^{1/2}\mathbf{x} = 0.$$

We denote by $\widetilde{\mathbf{x}}^* = [\mathbf{x}^*; \ldots; \mathbf{x}^*]$ the unique solution to (12).

Now, for given penalty parameter $\alpha > 0$, one can define the quadratic penalty function corresponding to the linearly constraint problem (12) as follows,

$$h_\alpha(\mathbf{x}) = \frac{1}{2}\mathbf{x}^\top(\mathbf{I} - \mathbf{W})\mathbf{x} + \alpha F(\mathbf{x}). \qquad (13)$$

Since $\mathbf{I} - \mathbf{W}$ is a positive semi-definite matrix and F is L-smooth and μ-strongly convex, therefore the function h_α is L_α-smooth and μ_α-strongly convex on $\mathbb{R}^{np}$ having $L_\alpha = 1 - \lambda_n(W) + \alpha L$ and $\mu_\alpha = \alpha\mu$. We denote by $\mathbf{x}_\alpha^*$ the unique minimizer of $h_\alpha(\mathbf{x})$, i.e.

$$\mathbf{x}_\alpha^* = \arg\min_{\mathbf{x} \in \mathbb{R}^{np}} h_\alpha(\mathbf{x}) = \arg\min_{\mathbf{x} \in \mathbb{R}^{np}} \frac{1}{2}\mathbf{x}^\top(\mathbf{I} - \mathbf{W})\mathbf{x} + \alpha F(\mathbf{x}). \qquad (14)$$

In the following, we link the solution of problem (14) to the local variable iterations provided by Algorithm 1. Specifically, for sufficiently large number of iterations T, we demonstrate that for proper choice of step-sizes, the expected squared deviation of $\mathbf{x}_T$ from $\mathbf{x}_\alpha^*$ vanishes sub-linearly.

Lemma 1. Consider the optimization Problem (14) and suppose Assumptions 1– 4 hold. Then, the expected deviation of the output of QDGD from the solution to Problem (14) is upper bounded by

$$\mathbb{E}\left[\left\|\mathbf{x}_T - \mathbf{x}_\alpha^*\right\|^2\right] \le \mathcal{O}\left(\frac{c_1 n \sigma^2 \|W - W_D\|_F^2}{\mu c_2} \frac{1}{T^{1/4}}\right), \qquad (15)$$

for $\varepsilon = \frac{c_1}{T^{1/2}}$, $\alpha = \frac{c_2}{T^{1/4}}$ and $T \ge T_1$, where c_1 and c_2 are positive constants independent of T, and $T_1 := \max\left\{\left\lceil\left(c_1 c_2 \mu\right)^{4/3}\right\rceil, \left\lceil\left(\frac{c_1(1 + c_2 L)^2}{c_2 \mu}\right)^4\right\rceil\right\}$.

Lemma 1 guarantees convergence of the proposed iterations (3) to the solution of the later-defined Problem (14). Loosely speaking, Lemma 1 ensures that $\mathbf{x}_T$ is *close* to $\mathbf{x}_\alpha^*$ for large T. So, in order to capture the deviation of $\mathbf{x}_T$ from the global optima $\widetilde{\mathbf{x}}^*$, it suffices to show that $\mathbf{x}_\alpha^*$ is *close* to $\widetilde{\mathbf{x}}^*$, as well. The following lemma guarantees such argument.

Lemma 2. Consider the distributed consensus optimization Problem (1) and the problem defined in (14). If Assumptions 1, 2 and 4 hold, then the deviation of the two solutions is bounded as

$$\|\mathbf{x}_\alpha^* - \widetilde{\mathbf{x}}^*\| \le \mathcal{O}\left(\frac{c_2}{1 - \beta} \cdot \frac{1}{T^{1/4}}\right), \qquad (16)$$

for $\alpha = \frac{c_2}{T^{1/4}}$ and $T \ge T_2$, where c_2 is a positive constant independent of T and $T_2 := \max\left\{\left\lceil\left(\frac{c_2 L}{1 + \lambda_n(W)}\right)^4\right\rceil, \left\lceil c_2^4(\mu + L)^4\right\rceil\right\}$.

Proofs of Lemmas 1 and 2 are skipped to the long version of the paper [36]. Having set the main lemmas, now it is straightforward to prove Theorem 1.

Proof of Theorem 1. For the specified step-sizes ε and α and large enough iterations $T \ge T_0 := \max\{T_1, T_2\}$, Lemmas 1

and 2 are applicable and we have

$$
\begin{aligned}
\mathbb{E}\left[\|\mathbf{x}_T - \widetilde{\mathbf{x}}^*\|^2\right] &= \mathbb{E}\left[\|\mathbf{x}_T - \mathbf{x}_\alpha^* + \mathbf{x}_\alpha^* - \widetilde{\mathbf{x}}^*\|^2\right] \\
&\le 2\mathbb{E}\left[\|\mathbf{x}_T - \mathbf{x}_\alpha^*\|^2\right] + 2\|\mathbf{x}_\alpha^* - \widetilde{\mathbf{x}}^*\|^2 \\
&\le \mathcal{O}\left(\frac{1}{T^{1/4}}\right) + \mathcal{O}\left(\frac{1}{T^{1/2}}\right) \\
&= \mathcal{O}\left(\frac{1}{T^{1/4}}\right). \quad\quad (17)
\end{aligned}
$$

Since $\mathbb{E}\left[\|\mathbf{x}_{i,T} - \mathbf{x}^*\|^2\right] \le \mathbb{E}\left[\|\mathbf{x}_T - \widetilde{\mathbf{x}}^*\|^2\right]$ for any $i = 1, \ldots, n$, the inequality in (17) follows the claim of Theorem 1. $\qquad\square$

V. NUMERICAL EXPERIMENTS

In this section, we evaluate the performance of the proposed QDGD Algorithm on minimizing a distributed quadratic objective. We pictorially demonstrate the effect of quantization noise and graph topology on the relative expected error rate.

Consider the quadratic optimization problem

$$
\min_{\mathbf{x} \in \mathbb{R}^p} \; f(\mathbf{x}) = \sum_{i=1}^n \frac{1}{2}\mathbf{x}^\top \mathbf{A}_i \mathbf{x} + \mathbf{b}_i^\top \mathbf{x}, \quad\quad (18)
$$

where $f_i(\mathbf{x}) = \frac{1}{2}\mathbf{x}^\top \mathbf{A}_i \mathbf{x} + \mathbf{b}_i^\top \mathbf{x}$ denotes the local objective function of node $i \in [n]$. The unique solution to (18) is therefore $\mathbf{x}^* = -\left(\sum_{i=1}^n \mathbf{A}_i\right)^{-1}\left(\sum_{i=1}^n \mathbf{b}_i\right)$. We pick diagonal matrices $\mathbf{A}_i$ such that $p/2$ of the diagonal entries of each $\mathbf{A}_i$ are drawn from the set $\{1, 2, 2^2\}$ and the other $p/2$ diagonal entries are drawn from the set $\{1, 2^{-1}, 2^{-2}\}$, all uniformly at random. Entries of vectors $\mathbf{b}_i$ are randomly picked from the interval $(0, 1)$. In our settup, the graph of agents is a connected Erdös-Rényi with edge probability p_c. We set the edge weight matrix to be $W = I - \frac{2}{3\lambda_{\max}(L)}L$ where L is the Laplacian with $\lambda_{\max}(L)$ as its largest eigenvalue. In our simulations, we let an additive noise model the quantization error, i.e. $Q(\mathbf{x}) = \mathbf{x} + \eta$ where $\eta \sim \mathcal{N}(0, \frac{\sigma^2}{p}I_p)$.

We first consider a connected Erdös-Rényi graph of $n = 50$ nodes and connectivity probability of $p_c = 0.35$ and dimension $p = 20$. Fig. 1 shows the convergence rate corresponding to three values of quantization noise $\sigma^2 \in \{2, 20, 200\}$, compared to the theoretical upper bound derived in Theorem 1 in the logarithmic scale. As expected, Fig. 1 shows that the error rate linearly scales with the quantization noise; however, it does not saturate around a non-vanishing residual, regardless the variance. Moreover, Fig. 1 demonstrates that the convergence rate closely follows the upper bound derived in Theorem 1. For instance, for the plot corresponding to $\sigma^2 = 20$, the relative errors are evaluated as $e_{T_1}/e_0 = 0.0121$ and $e_{T_2}/e_0 = 0.0082$ for $T_1 = 3200$ and $T_2 = 12800$, respectively. Therefore, $e_{T_2}/e_{T_1} \approx 0.68$ which is upper bounded by $\left(\frac{T_1}{T_2}\right)^{1/4} \approx 0.7$.

To observe the effect of graph topology, quantization noise variance is fixed to $\sigma^2 = 2$ and we varied the connectivity ratio by picking three different values, i.e. $p_c \in \{0.1, 0.5, 1\}$ where $p_c = 1$ corresponds to the complete graph case. As


Fig. 1. Relative optimal squared error for three vales of quantization noise variance: $\sigma^2 \in \{2, 20, 200\}$, compared with the order of upper bound.


Fig. 2. Relative optimal squared error for three vales of graph connectivity ratio: $p_c \in \{0.1, 0.5, 1\}$, compared with the order of upper bound.

Fig. 2 depicts, for the same number of iterations, deviation from the optimal solution tends to increase as the graph is gets sparse. In other words, even noisy information of the neighbor nodes improves the gradient estimate for local nodes. It also highlights the fact that regardless of the sparsity of the graph, the proposed QDGD algorithm guarantees the consensus to the optimal solution for each local node, as long as the graph is connected.

VI. CONCLUSION

We proposed the QDGD algorithm to tackle the problem of quantized decentralized consensus optimization. The algorithm updates the local decision variables by combining the quantized messages received from the neighbors and the local information such that proper averaging is performed over the local decision variable and the neighbors' quantized vectors. We proved that the QGDG algorithm achieves a vanishing consensus error in mean-squared sense, and verified our theoretical results with numerical studies.

An interesting future direction is to establish a fundamental trade-off between the convergence rate of quantized consensus algorithms and the communication. More precisely, given a target convergence rate, what is the minimum number of bits that one should communicate in decentralized consensus? Another interesting line of research is to develop

novel source coding (quantization) schemes that have low computation complexity and are information theoretically near-optimal in the sense that they have small communication load and fast convergence rate.

REFERENCES

[1] Y. Cao, W. Yu, W. Ren, and G. Chen, "An overview of recent progress in the study of distributed multi-agent coordination," *IEEE Transactions on Industrial informatics*, vol. 9, no. 1, pp. 427–438, 2013.

[2] A. Ribeiro, "Ergodic stochastic optimization algorithms for wireless communication and networking," *IEEE Transactions on Signal Processing*, vol. 58, no. 12, pp. 6369–6386, 2010.

[3] M. Rabbat and R. Nowak, "Distributed optimization in sensor networks," in *Proceedings of the 3rd international symposium on Information processing in sensor networks*, pp. 20–27, ACM, 2004.

[4] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 215–233, 2007.

[5] W. Ren, R. W. Beard, and E. M. Atkins, "Information consensus in multivehicle cooperative control," *IEEE Control Systems*, vol. 27, no. 2, pp. 71–82, 2007.

[6] K. I. Tsianos, S. Lawlor, and M. G. Rabbat, "Consensus-based distributed optimization: Practical issues and applications in large-scale machine learning," in *Communication, Control, and Computing (Allerton), 2012 50th Annual Allerton Conference on*, pp. 1543–1550, IEEE, 2012.

[7] J. Tsitsiklis, D. Bertsekas, and M. Athans, "Distributed asynchronous deterministic and stochastic gradient optimization algorithms," *IEEE transactions on automatic control*, vol. 31, no. 9, pp. 803–812, 1986.

[8] J. N. Tsitsiklis, "Problems in decentralized decision making and computation.," tech. rep., MASSACHUSETTS INST OF TECH CAMBRIDGE LAB FOR INFORMATION AND DECISION SYSTEMS, 1984.

[9] A. Nedic and A. E. Ozdaglar, "Distributed subgradient methods for multi-agent optimization," *IEEE Transactions on Automatic Control*, vol. 54, no. 1, pp. 48–61, 2009.

[10] K. Yuan, Q. Ling, and W. Yin, "On the convergence of decentralized gradient descent," *SIAM Journal on Optimization*, vol. 26, no. 3, pp. 1835–1854, 2016.

[11] D. Jakovetić, J. Xavier, and J. M. Moura, "Fast distributed gradient methods," *IEEE Transactions on Automatic Control*, vol. 59, no. 5, pp. 1131–1146, 2014.

[12] S. S. Ram, A. Nedić, and V. V. Veeravalli, "Distributed stochastic subgradient projection algorithms for convex optimization," *Journal of optimization theory and applications*, vol. 147, no. 3, pp. 516–545, 2010.

[13] S. P. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.

[14] W. Shi, Q. Ling, K. Yuan, G. Wu, and W. Yin, "On the linear convergence of the admm in decentralized consensus optimization.," *IEEE Trans. Signal Processing*, vol. 62, no. 7, pp. 1750–1761, 2014.

[15] A. Mokhtari, W. Shi, Q. Ling, and A. Ribeiro, "Decentralized quadratically approximated alternating direction method of multipliers," in *Signal and Information Processing (GlobalSIP), 2015 IEEE Global Conference on*, pp. 795–799, IEEE, 2015.

[16] J. C. Duchi, A. Agarwal, and M. J. Wainwright, "Dual averaging for distributed optimization: Convergence analysis and network scaling," *IEEE Transactions on Automatic Control*, vol. 57, no. 3, pp. 592–606, 2012.

[17] K. I. Tsianos, S. Lawlor, and M. G. Rabbat, "Push-sum distributed dual averaging for convex optimization," in *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, pp. 5453–5458, IEEE, 2012.

[18] E. Wei, A. Ozdaglar, and A. Jadbabaie, "A distributed newton method for network utility maximization–i: Algorithm," *IEEE Transactions on Automatic Control*, vol. 58, no. 9, pp. 2162–2175, 2013.

[19] A. Mokhtari, Q. Ling, and A. Ribeiro, "Network newton distributed optimization methods," *IEEE Transactions on Signal Processing*, vol. 65, no. 1, pp. 146–161, 2017.

[20] F. Yan, S. Sundaram, S. Vishwanathan, and Y. Qi, "Distributed autonomous online learning: Regrets and intrinsic privacy-preserving properties," *IEEE Transactions on Knowledge and Data Engineering*, vol. 25, no. 11, pp. 2483–2493, 2013.

[21] A. Mokhtari, W. Shi, Q. Ling, and A. Ribeiro, "A decentralized second-order method for dynamic optimization," in *Decision and Control (CDC), 2016 IEEE 55th Conference on*, pp. 6036–6043, IEEE, 2016.

[22] F. Seide, H. Fu, J. Droppo, G. Li, and D. Yu, "1-bit stochastic gradient descent and its application to data-parallel distributed training of speech dnns," in *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.

[23] M. Chowdhury, M. Zaharia, J. Ma, M. I. Jordan, and I. Stoica, "Managing data transfers in computer clusters with orchestra," *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 4, pp. 98–109, 2011.

[24] S. Yuksel and T. Basar, "Quantization and coding for decentralized lti systems," in *Decision and Control, 2003. Proceedings. 42nd IEEE Conference on*, vol. 3, pp. 2847–2852, IEEE, 2003.

[25] A. Kashyap, T. Basar, and R. Srikant, "Quantized consensus," *2006 IEEE International Symposium on Information Theory*, pp. 635–639, 2006.

[26] M. G. Rabbat and R. D. Nowak, "Quantized incremental algorithms for distributed optimization," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 4, pp. 798–808, 2005.

[27] M. El Chamie, J. Liu, and T. Başar, "Design and analysis of distributed averaging with quantized communication," *IEEE Transactions on Automatic Control*, vol. 61, no. 12, pp. 3870–3884, 2016.

[28] T. C. Aysal, M. Coates, and M. Rabbat, "Distributed average consensus using probabilistic quantization," in *Statistical Signal Processing, 2007. SSP'07. IEEE/SP 14th Workshop on*, pp. 640–644, IEEE, 2007.

[29] A. Nedic, A. Olshevsky, A. Ozdaglar, and J. N. Tsitsiklis, "Distributed subgradient methods and quantization effects," in *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*, pp. 4177–4184, IEEE, 2008.

[30] E. Gravelle and S. Martínez, "Quantized distributed load balancing with capacity constraints," in *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, pp. 3866–3871, IEEE, 2014.

[31] D. Alistarh, D. Grubic, J. Li, R. Tomioka, and M. Vojnovic, "Qsgd: Communication-efficient sgd via gradient quantization and encoding," in *Advances in Neural Information Processing Systems*, pp. 1707–1718, 2017.

[32] S. Li, M. A. Maddah-Ali, Q. Yu, and A. S. Avestimehr, "A fundamental tradeoff between computation and communication in distributed computing," *arXiv preprint arXiv:1604.07086*, 2016.

[33] K. Lee, M. Lam, R. Pedarsani, D. Papailiopoulos, and K. Ramchandran, "Speeding up distributed machine learning using codes," in *Information Theory (ISIT), 2016 IEEE International Symposium on*, pp. 1143–1147, IEEE, 2016.

[34] Y. H. Ezzeldin, M. Karmoose, and C. Fragouli, "Communication vs distributed computation: an alternative trade-off curve," *arXiv preprint arXiv:1705.08966*, 2017.

[35] S. Prakash, A. Reisizadeh, R. Pedarsani, and S. Avestimehr, "Coded computing for distributed graph analytics," in *2018 IEEE International Symposium on Information Theory (ISIT)*, pp. 1221–1225, 2018.

[36] A. Reisizadeh, A. Mokhtari, H. Hassani, and R. Pedarsani, "Quantized decentralized consensus optimization," *arXiv preprint arXiv:1806.11536*, 2018.