

Convergence of Limited Communications Gradient Methods

Sindri Magnússon, Chinwendu Enyioha, Na Li, Carlo Fischione, and Vahid Tarokh

Abstract—Distributed optimization increasingly plays a central role in economical and sustainable operation of cyber-physical systems. Nevertheless, the complete potential of the technology has not yet been fully exploited in practice due to communication limitations posed by the real-world infrastructures. This work investigates fundamental properties of distributed optimization based on gradient methods, where gradient information is communicated using limited number of bits. In particular, a general class of quantized gradient methods are studied where the gradient direction is approximated by a finite quantization set. Sufficient and necessary conditions are provided on such a quantization set to guarantee that the methods minimize any convex objective function with Lipschitz continuous gradient and a nonempty and bounded set of optimizers. A lower bound on the cardinality of the quantization set is provided, along with specific examples of minimal quantizations. Convergence rate results are established that connect the fineness of the quantization and the number of iterations needed to reach a predefined solution accuracy. Generalizations of the results to a relevant class of constrained problems using projections are considered. Finally, the results are illustrated by simulations of practical systems.

I. INTRODUCTION

Recent advances in distributed optimization have enabled more economical and sustainable control and operation of cyber-physical systems. However, these systems usually assume the availability of high performing communication infrastructures, which is often not practically possible. For example, although large scale cyber-physical systems such as power networks are equipped with a natural communication infrastructure given by the power lines [1], such a communication network has a limited bandwidth. Instead, research efforts in distributed operation of power networks usually assume high data rates and low latency communication technologies that are, unfortunately, not affordable or available today. Another example is given by wireless sensor networks [2], where efficient usage of communication plays a central role. In fact, these networks are powered by battery sources for communication over wireless links; hence, they are constrained in how much transmission they engage. These communication limitations are especially harsh in underwater networks, where acoustic channels are generally used, which have strong bandwidth limits [3]. Light communications are also essential in coordinating data networks [4], where the control channels that support the data channels are obviously allocated limited bandwidth. Another relevant example is within the emerging

technology of extremely low latency networking or tactile internet [5], where information, especially for real time control applications, will be transmitted with very low latencies over wireless and wired networks. However, this comes at a cost of using short packets containing limited information.

In all the cyber-physical systems mentioned above, distributed optimization plays a central role. These systems are networks of nodes whose operations have to be optimized by local decisions yet the coordination information can only go through constrained communication channels. In this paper, we restrict ourselves to one of the most prominent distributed optimization methods, decomposition based on the gradient method, and we investigate the fundamental properties of such a method in terms of coordination limitations and optimality.

A. Related Literature

Decomposition methods in optimization have been widely investigated in wired/wireless communication [6]–[9], power networks [10], [11], and wireless sensor networks [12], among others. These methods are typically based on communicating gradient information from a set of source nodes to users, which then solve a simple local subproblem. The procedure can be performed using i) one-way communication where the source nodes estimate the gradient using available information [7], [13], [14] or ii) two-way communication where users and sources coordinate to evaluate the gradient. We investigate the performance of such methods using one-way communication where the number of bits per coordination step are limited.

Bandwidth constrained optimization has already received attention in the literature [15]–[20]. Initial studies are found in [15], where Tsitsiklis and Luo provide lower bounds on the number of bits that two processors need to communicate to (approximately) minimize the sum of two convex functions each of which is only accessed by one processor. More recently, the authors of [16] consider a variant of incremental gradient methods [21] over networks where each node projects its iterate to a grid before sending the iterate to the next node. Similar quantization ideas are considered [17]–[19] in the context of consensus-type subgradient methods [22]. The work in [20] studies the convergence of standard interference function methods for power control in cellular wireless systems where base stations send binary signals to the users optimizing the transmit radio power. Those papers consider only the original primal optimization problem, without introducing its dual problem, where quantized primal variables are communicated. However, many network and resource sharing/allocation optimization problems are naturally decomposed using duality theory, where it is the dual *gradients* that are communicated. This motivates our studies of limited communication gradient methods.

This work was supported by the VR Chromos Project, NSF grant No. 1548204, and NSF Career 1553407.

S. Magnússon and C. Fischione are with the Electrical Engineering School, KTH Royal Institute of Technology, Stockholm, Sweden. (e-mail: sindrim@kth.se; carlofi@kth.se)

C. Enyioha, N. Li, and V. Tarokh are with the School of Engineering and Applied Sciences, Harvard University, Cambridge, MA USA. E-mail: sindrim@kth.se. (email: cenyioha@seas.harvard.edu; nali@seas.harvard.edu; vahid@seas.harvard.edu)

B. Statement of Contributions

The main contribution of this paper is to investigate the convergence of gradient methods where gradients are communicated using limited bandwidth. We first consider gradient methods where each coordinate of the gradient is communicated using only one-bit per iteration. This setup is motivated by dual decomposition applications where a single entity maintains each dual variable, e.g., in the TCP control in [7] each dual variable is maintained by one flow line. Since dual problems are either unconstrained or constrained to the positive orthant \mathbb{R}_+^N ,¹ we consider both unconstrained problems and problems constrained to \mathbb{R}_+^N . We prove that when the step size $\gamma > 0$ of the gradient method is fixed, then the iterates converge approximately to the set of optimal solutions within some $\epsilon > 0$ accuracy in a finite number of iterations, where ϵ tends to 0 as γ converges to 0. Moreover, we provide an upper complexity bound on the number of iterations needed to reach any $\epsilon > 0$ accuracy. This upper bound grows proportionally to $1/\epsilon^2$ as ϵ goes to zero for unconstrained problems, and proportionally to $1/\epsilon^4$ for problems constrained by \mathbb{R}_+^N . We also prove that if the step-sizes $\gamma(t)$ are non-summable and converge to 0, then the iterates converge to the set of optimal solutions.

The second contribution of the paper is to investigate the convergence of more general class of quantized gradient methods (QGM), where the gradient direction is quantized at every iteration. We start by identifying necessary and sufficient conditions on the quantization so that the QGMs can minimize any convex objective function with Lipschitz continuous gradients and a nonempty and bounded set of optimizers. We show that the minimal quantizations that satisfy these conditions have the cardinality $N+1$, where N is the problem dimension. We prove that for fixed step-sizes $\gamma > 0$ the iterates converge approximately to the set of optimal solutions within some ϵ -accuracy, where ϵ converges to 0 as γ converges to 0. We provide an upper complexity bound on the number of iterations T needed to reach any solution accuracy $\epsilon > 0$. This upper bound depends on the fineness of the quantization. Moreover, we show that the solution accuracy $\epsilon > 0$ converges to zero at a rate proportional to $1/\sqrt{T}$ or $1/\sqrt{b}$ where T and b are the numbers of iterations and communicated bits, respectively. We show that when the step-sizes $\gamma(t)$ are non-summable and converge to zero then the iterates asymptotically converge to the set of optimal solutions.

A conference version of this work including parts of Section III and Section V appeared in [23], but without most of the proofs. The rest of the work is presented here for the first time. Our previous papers [24], [25] consider similar resource allocation problems as in this paper without communication constraints.

C. Notation

Vectors and matrices are represented by boldface lower and upper case letters, respectively. The set of real, positive real, and natural numbers, are denoted by \mathbb{R} , \mathbb{R}_+ , and \mathbb{N} ,

¹Depending on whether the primal problem has inequality or equality constraints.

respectively. The set of real and positive n vectors and $n \times m$ matrices are denoted by \mathbb{R}^n , \mathbb{R}_+^n , and $\mathbb{R}^{n \times m}$, respectively. Other sets are represented by calligraphy letters. We denote by \mathcal{S}^{N-1} , $\mathcal{S}^{N-1}(\mathbf{x}, R)$, and $\mathcal{B}^N(\mathbf{x}, R)$, respectively, the unit sphere in \mathbb{R}^N and the sphere and open ball centred at \mathbf{x} with radius R in \mathbb{R}^N . The superscript $(\cdot)^T$ stands for transpose. We let $[\mathbf{x}]_{\mathcal{X}}$ and $[\mathbf{x}]_+$ denote the projection of \mathbf{x} to the sets \mathcal{X} and \mathbb{R}_+ . $\|\cdot\|$ denotes the 2-norm. ∇f is the gradient of f . The distance between a vector $\mathbf{x} \in \mathbb{R}^N$ and a set $\mathcal{X} \subseteq \mathbb{R}^N$ is denoted by $\text{dist}(\mathbf{x}, \mathcal{X}) = \inf_{\mathbf{v} \in \mathcal{X}} \|\mathbf{v} - \mathbf{x}\|$.

II. PRELIMINARIES AND MOTIVATING EXAMPLES

In this paper we consider optimization problems of the form

$$\begin{aligned} & \underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} && f(\mathbf{x}), \\ & \text{subject to} && \mathbf{x} \in \mathcal{X}, \end{aligned} \quad (1)$$

where $f: \mathbb{R}^N \rightarrow \mathbb{R}$. We denote by f^* and $\mathcal{X}^* \subseteq \mathcal{X}$ the optimal value and the set of optimizers to Problem (1), respectively. We consider the following class of optimization problems:

Definition 1. Let $\mathcal{F}_L(\mathcal{X})$ denote the set of optimization problems of the form of Equation (1) where the function f is convex and differentiable with L -Lipschitz continuous gradient, \mathcal{X} is closed and convex set, and the optimal solution set \mathcal{X}^* is nonempty and bounded. We write $f \in \mathcal{F}_L(\mathcal{X})$ to indicate that the optimization Problem (1) with the objective function f in the class $\mathcal{F}_L(\mathcal{X})$.

For $f \in \mathcal{F}_L(\mathcal{X})$, it is well known that the gradient method

$$\mathbf{x}(t+1) = [\mathbf{x}(t) - \gamma(t)\nabla f(\mathbf{x}(t))]_{\mathcal{X}}, \quad (2)$$

converges to \mathcal{X}^* under appropriate step-size rules [26]. When only the gradient *direction* is known, the above iterates become

$$\mathbf{x}(t+1) = \left[\mathbf{x}(t) - \gamma(t) \frac{\nabla f(\mathbf{x}(t))}{\|\nabla f(\mathbf{x}(t))\|} \right]_{\mathcal{X}}, \quad (3)$$

where we set $\mathbf{x}(t+1) = \mathbf{x}(t)$ if $\|\nabla f(\mathbf{x}(t))\| = 0$. For appropriate diminishing step-size rules, the iterates converge to \mathcal{X}^* and for fixed step-size, the stopping condition $f(\mathbf{x}(t)) - f^* < \epsilon$ can be achieved for any $\epsilon > 0$ [27].

Problems on the form of Equation (1) often appear as dual problems used to decompose optimization problems with coupling constraints [8], [9]. There, a distributed solution approach is achieved by solving the dual problem using gradient methods as given in Equations (2) and (3). The dual gradient $\nabla(f(\mathbf{x}(t)))$ can often be measured over the course of the algorithm, as it is the constraint violation in the primal problem, given dual variable $\mathbf{x}(t)$ (see, e.g., [7], [13], [14]). To perform the gradient update in Equations (2) and (3) the gradient or gradient direction must be communicated, as illustrated in the following examples. However, since communication resources are scarce in many networks, we consider another variant of the gradient method in (2). That is

$$\mathbf{x}(t+1) = [\mathbf{x}(t) - \gamma(t)\mathbf{d}(t)]_{\mathcal{X}},$$

where $\mathbf{d}(t)$ is a quantized gradient direction coded using limited number of bits. Before introducing the details of our quantized methods we provide some application examples.

A. TCP Flow Control

Consider a communication network with N undirected links and S data sources. Let \mathcal{L} and \mathcal{S} denote the ordered sets $\{1, \dots, N\}$ and $\{1, \dots, S\}$. Denote the capacity of link $l \in \mathcal{L}$ by $c_l > 0$ and the transmission rate of source $s \in \mathcal{S}$ by $q_s \in [m_s, M_s]$, where m_s and M_s are upper and lower bound on the source. Source $s \in \mathcal{S}$ has utility function $U_s : [m_s, M_s] \rightarrow \mathbb{R}$. The data from source $s \in \mathcal{S}$ flows through a path consisting of links $\mathcal{L}_s \subseteq \mathcal{L}$ to its destination. We denote by $\mathcal{S}_l \subseteq \mathcal{S}$ the sources that use link $l \in \mathcal{L}$, i.e., $\mathcal{S}_l := \{s \in \mathcal{S} | l \in \mathcal{L}_s\}$. Then the TCP flow control is to find data rates q_s , $s \in \mathcal{S}$, that solve the following optimization problem [7], [28]–[30]

$$\begin{aligned} & \underset{q_1, \dots, q_S}{\text{maximize}} && \sum_{s=1}^S U_s(q_s) \\ & \text{subject to} && \sum_{s \in \mathcal{S}_l} q_s \leq c_l, \quad \text{for } l = 1, \dots, N \\ & && q_s \in [m_s, M_s], \quad \text{for } s = 1, \dots, S. \end{aligned} \quad (4)$$

For notational ease, we write $\mathbf{q} = (q_s)_{s \in \mathcal{S}}$, $\mathcal{Q} = \prod_{s \in \mathcal{S}} [m_s, M_s]$, $\mathbf{c} = (c_l)_{l \in \mathcal{L}}$, and $\mathbf{A} \in \mathbb{R}^{N \times S}$ where

$$\mathbf{A}_{ls} = \begin{cases} 1 & \text{if } l \in \mathcal{S}_s, \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

The dual problem of (4) is of the form (1) where $\mathcal{X} = \mathbb{R}_+^N$ and the dual function $f : \mathbb{R}^N \rightarrow \mathbb{R}$ is given by

$$f(\mathbf{x}) = \underset{\mathbf{q} \in \mathcal{Q}}{\text{maximize}} L(\mathbf{q}, \mathbf{x}) = L(\mathbf{q}(\mathbf{x}), \mathbf{x}),$$

where

$$\begin{aligned} L(\mathbf{q}, \mathbf{x}) &= \sum_{i=1}^M U_i(q_i) - \mathbf{x}^T (\mathbf{A}\mathbf{q} - \mathbf{c}), \\ q_i(\mathbf{x}) &= \underset{q_i \in [m_i, M_i]}{\text{argmax}} U_i(q_i) - q_i \sum_{l \in \mathcal{L}_i} x_l. \end{aligned} \quad (6)$$

The dual gradient is given by $\nabla f(\mathbf{x}) = \mathbf{c} - \mathbf{A}\mathbf{q}(\mathbf{x})$. The dual gradient is bounded since the set \mathcal{Q} is compact. Moreover, the set of optimal dual variables is bounded from Lemma 1 in [31] and the dual gradient $\nabla f(\cdot)$ is L -Lipschitz continuous from Lemma 3 in [7], provided that $U_i(\cdot)$ are strongly concave. Therefore, the dual iterates $\mathbf{x}(t)$ in Equation (2) or (3) and the associated primal iterates $\mathbf{q}(t) = \mathbf{q}(\mathbf{x}(t))$ converge to the optimal primal/dual solution of the optimization Problem (4), provided that $\gamma(t)$ are chosen properly.

Dual gradient methods are desirable because they entail distributed solution to Problem (4) since Subproblems (6) can be solved without any coordination between the sources \mathcal{S} . Moreover, the gradient component $\nabla_l f(\mathbf{x}) = c_l - \sum_{s \in \mathcal{L}(l)} q_s(x_l)$ can often be measured at the data link l since it is simply the difference between the link capacity, c_l , and the data transferred through the link [7]. Therefore, the algorithm can be accomplished using one-way communication where each iteration t consists of the following steps: (i) the links broadcast $\mathbf{x}(t)$ to the sources (ii) the sources solve the local Subproblem (6) and then transfer the source at the data rate $q_i(\mathbf{x}(t))$,

and (iii) the links measure the dual gradient $\nabla_l f(\mathbf{x}(t))$, the data flow through the line, to make the update (2) or (3).

The control channels used to coordinate communication networks are often bandwidth limited. Hence, it is not practically feasible to broadcast the real-valued vector $\nabla f(\mathbf{x}(t))$ to the users. The questions we address in this paper are these: can we still solve the optimization problem by communicating quantized versions of the gradient? And what are the trade-offs between optimality and quantization? This motivates our investigation of limited communication gradient methods.

B. Optimal Network Flow

Consider a directed network $(\mathcal{N}, \mathcal{E})$ where $\mathcal{N} = \{1, \dots, N\}$ and $\mathcal{E} = \{1, \dots, E\}$ denote the sets of nodes and edges, respectively. Let v_e denote the flow through the edge $e \in \mathcal{E}$. The flow through the network can then be expressed by the matrix $\mathbf{A} \in \mathbb{R}^{N \times E}$ defined as

$$\mathbf{A}_{ne} := \begin{cases} 1 & \text{if edge } e \text{ leaves node } n, \\ -1 & \text{if edge } e \text{ enters node } n, \\ 0 & \text{otherwise.} \end{cases}$$

Component $n \in \mathcal{N}$ of $\mathbf{A}\mathbf{v}$ indicates the flow injection/consumption at node n , where $\mathbf{v} = (v_1, \dots, v_E)$. We assume that the flow injection ($c_n > 0$) or consumption ($c_n \leq 0$) of node $n \in \mathcal{N}$ is $c_n \in \mathbb{R}$ and set $\mathbf{c} = (c_1, \dots, c_N) \in \mathbb{R}^N$. Then the Optimal Network Flow problem is [32]–[34]

$$\begin{aligned} & \underset{v_1, \dots, v_E}{\text{maximize}} && \sum_{e \in \mathcal{E}} -C_e(v_e) \\ & \text{subject to} && \mathbf{A}\mathbf{v} = \mathbf{c}, \end{aligned}$$

where $C_e : \mathbb{R} \rightarrow \mathbb{R}$ are cost functions of the flow through edge $e \in \mathcal{E}$. If C_e are μ -strongly convex then the dual gradient is L -Lipschitz continuous with $L = \lambda_{\max}(\mathbf{A}\mathbf{A}^T)/\mu$, see Lemma 1 in [33]. Then similar one-way communication dual decomposition algorithm can be performed as in Section II-A. In contrast to Section II-A, the dual problem is unconstrained, i.e., $\mathcal{X} = \mathbb{R}^N$. In addition, the dual variables are maintained by the nodes so it is the nodes that broadcast the dual gradients while the edges solve the local subproblems. Nevertheless, the dual gradient $\nabla f(\mathbf{x})$ can be measured at the nodes as the component $\nabla_n f(\mathbf{x})$ is simply the flow injection/consumption of node n for a given \mathbf{x} .

C. Task Allocation

Consider the problem of continuous allocation of N tasks between K machines. The sets of tasks and machines are denoted by $\mathcal{N} = \{1, \dots, N\}$ and $\mathcal{K} = \{1, \dots, K\}$, respectively. Let $\mathbf{c} \in \mathbb{R}^N$ denote the total amount of each task that needs to be completed. The amount of each task done by machine $k \in \mathcal{K}$ is represented by the vector $\mathbf{w}_k \in \mathcal{W}_k \subseteq \mathbb{R}^N$, where \mathcal{W}_k is a local constraint of machine k . Then the goal is to find the task allocation that minimizes the cost:

$$\begin{aligned} & \underset{\mathbf{w}_1, \dots, \mathbf{w}_K}{\text{maximize}} && \sum_{k \in \mathcal{K}} -C_k(\mathbf{w}_k) \\ & \text{subject to} && \sum_{k \in \mathcal{K}} \mathbf{w}_k = \mathbf{c}, \\ & && \mathbf{w}_k \in \mathcal{W}_k, \quad \text{for } k \in \mathcal{K}, \end{aligned} \quad (7)$$

where $C_k : \mathbb{R}^N \rightarrow \mathbb{R}$ is the cost of performing the different tasks on machine $k \in \mathcal{K}$. If C_k are μ -strongly convex then the dual gradient is L -Lipschitz continuous with $L = (K + 1)/\mu$, see Lemma 1 in [30]. Therefore, dual gradient methods (2) and (3) can solve the problem. If the task manager can measure the total amount done of each task, i.e., the dual gradient, then a similar one-way communication coordination scheme as in Sections II-A and II-B can solve Problem (7).

As shown later, the Lipschitz constant L will be used to characterize several complexity bounds. Since L on the dual gradient of the examples above is a function of primal problem parameters such as the number of users and the concavity parameter μ , those parameters affect the complexity bounds as well.

III. QUANTIZED GRADIENT DESCENT METHODS

We consider general quantized gradient methods of the form

$$\mathbf{x}(t+1) = [\mathbf{x}(t) - \gamma(t)\mathbf{d}(t)]_{\mathcal{X}}, \quad (8)$$

where $\mathbf{d}(t) \in \mathcal{D} \subseteq \mathcal{S}^{N-1}$ is a finite set of quantized gradient directions. The following relation is between the cardinality of \mathcal{D} and communicated bits of each Iteration (8).

Remark 1. *The set \mathcal{D} can be coded using $\log_2(|\mathcal{D}|)$ bits.*

We now introduce the quantization sets \mathcal{D} used in this paper.

A. Binary Quantization

In the application examples in Sections II-A, II-B, and II-C, each dual variable is associated to a single problem component, i.e., a link, user, or task, respectively. For example, in the TCP control example in Section II-A, the dual variable \mathbf{x}_n is associated with link n . Therefore, to achieve the dual gradient Algorithm (2) each link $l \in \mathcal{L}$ can measure its flow, i.e., the dual gradient component $\nabla_l f(\mathbf{x}(t))$, and then broadcast $\nabla_l f(\mathbf{x}(t))$ to the sources that use link l . However, it might be infeasible to broadcast the full dual gradient when bandwidth is limited. An alternative approach is to have the links broadcast a binary signal indicating whether the associated dual variable is to be increased or decreased, i.e., link l broadcasts $\text{sign}(\nabla_l f(\mathbf{x}(t)))$. Similarly, in the network flow problem in Section II-B, each node can measure the flow through the node and then broadcast a binary signal indicating the direction of the associated dual gradient component. This quantization can be formally expressed as follows.

Example 1 (Signs of the gradients). *Consider the quantized gradient method in Equation (8). Set*

$$\mathcal{D} = \left\{ (1/\sqrt{N})(e_1, e_2, \dots, e_N) \mid e_i \in \{-1, 1\} \right\},$$

and take $\mathbf{d}(t) = \text{sign}(\nabla f(\mathbf{x}(t)))$.

By using this binary quantization, we prove the convergence of the Iterates (8) when $\mathcal{X} = \mathbb{R}^N$ and $\mathcal{X} = \mathbb{R}_+^N$. Therefore, our results cover both the case when the optimization Problem (1) is a dual problem associated with equality and

inequality constrained primal problems.² Our results show that the Iterates (8) using the quantization in Example 1 converge (i) approximately to the set of optimal values when the step-sizes are fixed and (ii) asymptotically when the step-size are diminishing and non-summable. In section IV, we prove the convergence in the constrained case when $\mathcal{X} = \mathbb{R}_+^N$. The convergence in the unconstrained case $\mathcal{X} = \mathbb{R}^N$ is a special case of the more general convergence results in Section V.

B. Fundamental Limit: Proper Quantization

When the quantization in Example 1 is used in the TCP problem then there is no collaboration between the network links (or the nodes in the Network flow problem). As a result $|\mathcal{D}| = 2^N$ and $\log_2(2^N) = N$ bits are used to broadcast the quantized gradient direction per iteration. However, in many applications [8], [9], the dual problem is maintained by a single coordinator. Therefore, an interesting question is whether even fewer than N bits can be used per iteration when a single coordinator maintains the dual problem? In that case, what is the minimal quantization $|\mathcal{D}|$ so the Iterates (8) can solve the optimization Problem (1)? More generally, for what quantization sets \mathcal{D} do the Iterates (8) converge to optimal solution to the Problem (1)? To answer such questions we now formalize how a quantization set \mathcal{D} enables the Iterates (8) to solve the optimization Problem (1).

Definition 2. *Consider Iterations (8). A finite set \mathcal{D} is a proper quantization for the problem class $\mathcal{F}_L(\mathcal{X})$ if for every $f \in \mathcal{F}_L(\mathcal{X})$ and every initialization $\mathbf{x}(0) \in \mathcal{X}$ we can choose $\mathbf{d}(t) \in \mathcal{D}$ and $\gamma(t) \in \mathbb{R}_+$, for all $t \in \mathbb{N}$, such that $\lim_{t \rightarrow \infty} \text{dist}(\mathbf{x}(t), \mathcal{X}^*) = 0$.*

Using Definition 2 we investigate the following questions:

- Are there equivalent constructive conditions that can be used to determine whether \mathcal{D} is a proper quantization or to construct such quantization sets?
- What is the minimal quantization, i.e., size $|\mathcal{D}|$, for which \mathcal{D} is a proper quantization?
- What are the connections between the fineness of the quantization, i.e., the size of $|\mathcal{D}|$, and the possible convergence rate of the algorithm?

For the class $\mathcal{F}_L(\mathbb{R}^N)$ of unconstrained problems, the next few paragraphs answer all these questions. However, as shown in Section III-C, a proper quantization set \mathcal{D} for the class $\mathcal{F}_L(\mathbb{R}^N)$ might not be a proper quantization for $\mathcal{F}_L(\mathcal{X})$ when \mathcal{X} is a proper subset of \mathbb{R}^N .

Question A): Consider the following definition.

Definition 3. *The finite set \mathcal{D} is a θ -cover if $\theta \in [0, \pi/2)$ and for all $\mathbf{g} \in \mathcal{S}^{N-1}$ there is $\mathbf{d} \in \mathcal{D}$ such that $\text{ang}(\mathbf{g}, \mathbf{d}) = \cos^{-1}(\langle \mathbf{g}, \mathbf{d} \rangle) \leq \theta$. Equivalently, for all $\mathbf{g} \in \mathcal{S}^{N-1}$ there is $\mathbf{d} \in \mathcal{D}$ such that $\cos(\text{ang}(\mathbf{g}, \mathbf{d})) \geq \cos(\theta) > 0$.*

Informally, \mathcal{D} is a θ -cover if for any non-zero vector in \mathbb{R}^N , there exists some element in \mathcal{D} such that the angle between

²When the primal problem have equality constraint then the dual problem is unconstrained, so $\mathcal{X} = \mathbb{R}^N$. Otherwise, if the primal problem has inequality constraints then $\mathcal{X} = \mathbb{R}_+^N$.

them is smaller than or equal to θ . The following theorem shows that Definition 3 of θ -cover is actually equivalent to Definition 2 of proper quantization for the problem class $\mathcal{F}_L(\mathbb{R}^N)$.

Theorem 1. Consider a finite set $\mathcal{D} \subseteq \mathcal{S}^{N-1}$. \mathcal{D} is a proper quantization for the class $\mathcal{F}_L(\mathbb{R}^N)$ (Definition 2) if and only if there exists $\theta \in [0, 2\pi)$ such that \mathcal{D} is a θ -cover (Definition 3).

Proof: The proof is found in Appendix A. ■

Unlike Definition 2 of proper quantization, Definition 3 is *constructive* in the sense that it can be used to determine if a set \mathcal{D} is a proper quantization and to construct such sets. For example, we can use Definition 3 and Theorem 1 to show that the quantization scheme in Example 1 is a proper quantization for the problem class $\mathcal{F}_L(\mathbb{R}^N)$.

Lemma 1. The quantization in Example 1 is a θ -cover with $\cos(\theta) = 1/\sqrt{N}$.

The lemma follows from the fact that for any $\mathbf{x} \in \mathcal{S}^{N-1}$, if we choose $\mathbf{d} = (1/\sqrt{N})\text{sign}(\mathbf{x})$ then

$$\begin{aligned} \cos(\text{ang}(\mathbf{x}, \mathbf{d})) &= \langle \mathbf{x}, \mathbf{d} \rangle = \frac{1}{\sqrt{N}} \sum_{i=1}^N \mathbf{x}_i \cdot \text{sign}(\mathbf{x}_i) \\ &\geq \frac{1}{\sqrt{N}} \sum_{i=1}^N \mathbf{x}_i^2 = \frac{1}{\sqrt{N}} \|\mathbf{x}\|^2 = \frac{1}{\sqrt{N}}. \end{aligned}$$

Lemma 1 proves that the quantization in Example 1 is a proper quantization. We give some other interesting θ -covers now.

Example 2 (Minimal Proper Quantization: $|\mathcal{D}| = N + 1$). Set

$$\mathcal{D} = \{\mathbf{e}_1, \dots, \mathbf{e}_N, -\mathbf{1}/\sqrt{N}\}, \quad (9)$$

where \mathbf{e}_i is the i -th element of the normal basis and $\mathbf{1} = (1, \dots, 1) \in \mathbb{R}^N$. Clearly, $|\mathcal{D}| = N + 1$ so \mathcal{D} can be coded using only $\log_2(N + 1)$ bits. We show below in Theorem 2 that this is a minimal proper quantization as there does not exist proper quantization \mathcal{D} with $|\mathcal{D}| \leq N$. We show in Lemma 7 in the Appendix that \mathcal{D} is a θ -cover with

$$\cos(\theta) = \frac{1}{\sqrt{N^2 + 2\sqrt{N}(N - 1)}}. \quad (10)$$

Example 3 (Example in \mathbb{R}^2 : $|\mathcal{D}| = n$). For every $n \in \mathbb{N}$ set

$$\mathcal{D}_n = \left\{ \left[\begin{array}{c} \cos(2\pi k/n) \\ \sin(2\pi k/n) \end{array} \right] \in \mathbb{R}^2 \mid k = 0, 1, \dots, n - 1 \right\}.$$

Clearly, if $n \geq 3$, \mathcal{D}_n is a θ -cover with $\theta = \pi/n$.

Example 4 (Normal Basis: $|\mathcal{D}| = 2N$). Let $\mathcal{D} = \{\mathbf{e}_1, -\mathbf{e}_1, \mathbf{e}_2, -\mathbf{e}_2, \dots, \mathbf{e}_N, -\mathbf{e}_N\}$. $|\mathcal{D}| = 2N$ and \mathcal{D} is a θ -cover with $\cos(\theta) = 1/\sqrt{N}$ since for all $\mathbf{x} \in \mathcal{S}^{N-1}$, if we choose $\mathbf{d} = \text{sign}(\mathbf{x}_i)\mathbf{e}_i$ where $i = \text{argmax}_{j=1, \dots, N} |\mathbf{x}_j|$ then $\cos(\text{ang}(\mathbf{x}, \mathbf{d})) = \langle \mathbf{x}, \mathbf{d} \rangle = \mathbf{x}_i \cdot \text{sign}(\mathbf{x}_i) = |\mathbf{x}_i| \geq 1/\sqrt{N}$.

For constant $\theta \in (0, \pi/2)$, it can be of interest to find the θ -cover \mathcal{D} which has minimal cardinality $|\mathcal{D}|$. This problem has been investigated in the coding theory literature, [35], [36].

Question B): The minimal proper quantization for the problem class $\mathcal{F}_L(\mathbb{R}^N)$ is $|\mathcal{D}| = N + 1$. We already have a proper

quantization with $|\mathcal{D}| = N + 1$, see Example 2. The following result shows that there does not exist a quantization set \mathcal{D} with fewer elements than $N + 1$.

Theorem 2. Suppose that $\mathcal{D} \subseteq \mathcal{S}^{N-1}$ and $|\mathcal{D}| \leq N$. Then \mathcal{D} is not a proper quantization.

Proof: The proof is in Appendix B. ■

This result shows that the minimum data-rate needed for the algorithm to converge is $\log_2(N + 1)$ bits/iteration. To the best of our knowledge, there are no similar results on minimal quantizations for distributed optimization methods in the existing literature.

Question C): In Section V, we study the convergence of Iterates (8) for the problem class $\mathcal{F}_L(\mathbb{R}^N)$ when \mathcal{D} is a θ -cover. When the step-size is constant, i.e., $\gamma(t) = \gamma$, then we show that any solution accuracy $\|\nabla f(\mathbf{x})\| \leq \epsilon$ and $f(\mathbf{x}) - f^* \leq \epsilon$ can be achieved, for $\epsilon > 0$. We also give an upper bound on the number of iterations/bits needed to achieve that accuracy. An implication of the results is that after $T \in \mathbb{N}$ iterations the accuracy

$$\|\nabla f(\mathbf{x})\| \leq \frac{M}{\cos(\theta)\sqrt{T}},$$

can be reached using appropriate constant step-size choice, where $M > 0$ is some constant. Finally, we show how to choose the step-sizes so that every limit point of the algorithm is an optimizer of Problem (1).

C. Fundamental Limit: What If There Are Constraints?

We now show by simple examples why a θ -cover \mathcal{D} might not be a proper quantization for the problem class $\mathcal{F}_L(\mathcal{X})$ when the feasible set \mathcal{X} is a proper subset of \mathbb{R}^N . These examples are illustrated in Fig. 1a and Fig. 1b. Both figures demonstrate scenarios where a single step of Iteration (8) is taken from \mathbf{x} . In both figures $\mathcal{D} = \{\mathbf{d}_1, \mathbf{d}_2, \mathbf{d}_3\}$ is a θ -cover with $\theta = \pi/3$. The feasible region is depicted by gray color. The curves depict the contours of the objective function f . The dotted lines depict the angle θ .

Fig. 1a depicts a scenario where Iteration (8) may have a non-optimal stationary point, even though \mathcal{D} is a θ -cover. The point \mathbf{x} is a stationary point since $-\mathbf{d}_1$ is orthogonal to the constraint. However, \mathbf{x} is not an optimal solution of Problem (1), since the gradient $\nabla f(\mathbf{x})$ is not orthogonal to the constraint. This example shows that the equivalence established in Theorem 1 does not generalize to the constrained case. Fig. 1b shows that the Iterates (8) can go in the opposite direction of the optimal solution. The Iterate (8) is a not a descent direction; hence, the objective function value is increasing. The optimal solution of Problem (1) and the Iterate (8) are denoted by \mathbf{x}^* and $[\mathbf{x} - \gamma\mathbf{d}_1]_{\mathcal{X}}$, respectively.

IV. CONVERGENCE - BINARY QUANTIZATION

In this section, we investigate the convergence of the quantized gradient method

$$\mathbf{x}(t+1) = \left[\mathbf{x}(t) - \frac{\gamma(t)}{\sqrt{N}} \text{sign}(\nabla f(\mathbf{x}(t))) \right]^+, \quad (11)$$

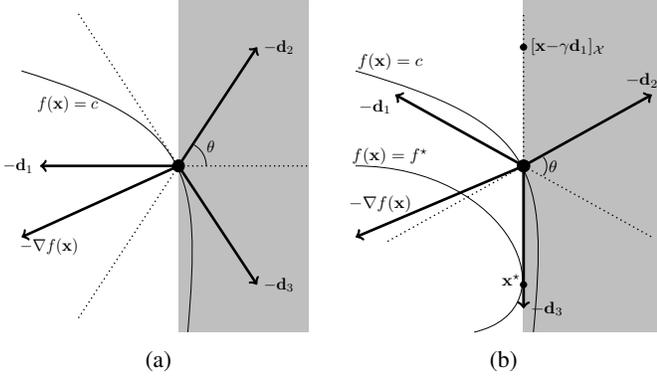


Fig. 1: An iteration of Equation (8) starting at \mathbf{x} . \mathbf{x} is in the middle of the figure marked by a big filled circle. The feasible set \mathcal{X} is marked by the shaded region. \mathbf{x}^* and f^* are the optimizers and optimal value. $\mathcal{D} = \{\mathbf{d}_1, \mathbf{d}_2, \mathbf{d}_3\}$ is a θ -cover with $\theta = \pi/3$. In (a) \mathbf{x} is a fixed point of the algorithm. In (b) the algorithm takes the step $[\mathbf{x} - \gamma \mathbf{d}_1]_{\mathcal{X}}$.

for solving the optimization Problem (1) when $\mathcal{X} = \mathbb{R}_+^N$. We make the additional assumption that the gradients ∇f are B -bounded, i.e., $\|\nabla f(\mathbf{x})\| \leq B$ for all $\mathbf{x} \in \mathbb{R}_+^N$.³ In the next section, which considers unconstrained problems, we allow the gradients to be unbounded.

In the analysis we take advantage of the following property of optimal solution \mathbf{x}^* .

Lemma 2. Consider optimization Problem (1) with $\mathcal{X} = \mathbb{R}_+^N$. For $\alpha > 0$ define the function

$$L_\alpha(\mathbf{x}) = \|\mathbf{x} - [\mathbf{x} - \alpha \nabla f(\mathbf{x})]^+\|. \quad (12)$$

Then $\mathbf{x} \in \mathcal{X}^*$ if and only if $L_\alpha(\mathbf{x}) = 0$.

Proof. The proof is found in Appendix C. \square

We investigate the convergence of the iterates in Equation (11) when the step-sizes are fixed in Section IV-A and when the step-sizes are diminishing in Section IV-B.

A. Constant Step-Size

In this section, we study the convergence of Iterates (11) when the step-size $\gamma(t)$ is constant, i.e., when $\gamma(t) = \gamma$ for all t . We show that the Iterates (11) can approximately solve optimization Problem (1) up to any ϵ -accuracy, provided that the step-size $\gamma > 0$ is small enough. By approximately solving (1), we mean that we can find $\mathbf{x} \in \mathbb{R}_+^N$ that approximately satisfies certain optimality conditions. In particular, we consider the following two types of optimality conditions for Problem (1):

$$\text{Type-1:} \quad L_\alpha(\mathbf{x}) \leq \epsilon, \quad (13)$$

$$\text{Type-2:} \quad f(\mathbf{x}) - f^* \leq \epsilon. \quad (14)$$

A point \mathbf{x} is an optimal solution to optimization Problem (1) if and only if (13) [or (14)] hold with $\epsilon = 0$. The Type-1 optimality condition is a generalization of the optimality condition

³This assumption is only needed in this section. The dual gradient is generally bounded if the primal problem is strongly convex and has bounded feasible set, see Proposition 6.1.1 in [26]. For example, the dual gradient is bounded in the TCP problem in Section II-A.

that $\|\nabla f(\mathbf{x})\| = 0$ for unconstrained problems. The Type-2 optimality condition simply state that the distance from the obtained objective value to the optimal value is less than ϵ . We now show that both Type-1 and Type-2 approximate optimality conditions can be reached in finite number of iterations.

1) *Stopping Condition of Type-1:* We start by showing the Type-1 optimality condition can be reached for all $\epsilon > 0$ in finite number of iterations. The following lemma is essential in proving the result.

Lemma 3. Suppose $f \in \mathcal{F}_L(\mathbb{R}_+^N)$ and ∇f is B -bounded. Suppose $\epsilon > 0$ and $\alpha \geq 0$ are such that

$$L_\alpha(\mathbf{x}) = \|\mathbf{x} - [\mathbf{x} - \alpha \nabla f(\mathbf{x})]^+\| \geq \epsilon. \quad (15)$$

Then the following holds

$$f\left(\left[\mathbf{x} - \frac{\gamma}{\sqrt{N}} \text{sign}(\nabla f(\mathbf{x}))\right]^+\right) \leq f(\mathbf{x}) - \bar{\delta}(\epsilon, \alpha, \gamma)$$

where

$$\bar{\delta}(\epsilon, \alpha, \gamma) = \left(\frac{2\epsilon^2}{L\alpha^2 B N^{3/2}} - \gamma\right) \frac{L}{2} \gamma.$$

Proof: The proof is provided in Appendix D. \blacksquare

The lemma shows that if $L_\alpha(\mathbf{x}) > 0$ for some $\mathbf{x} \in \mathbb{R}_+^N$, then objective function value $f(\mathbf{x})$ will decrease with Iterates (11), provided that the step-size $\gamma(t) > 0$ is small enough. We use this intuition to provide an upper bound on the number of iterations needed to achieve the Type-1 approximate optimality.

Theorem 3. Suppose $f \in \mathcal{F}_L(\mathbb{R}_+^N)$, ∇f is B -bounded, $\mathbf{x}(t)$ are generated by Equation (11), and define the set

$$\bar{\mathcal{X}}_\alpha(\epsilon) := \{\mathbf{x} \in \mathbb{R}_+^N \mid L_\alpha(\mathbf{x}) \leq \epsilon\}. \quad (16)$$

Then for any $\epsilon > 0$ and $\gamma \in (0, 2\epsilon^2/(\alpha^2 B N^{3/2}))$, there exists $T \in \mathbb{N} \cup \{0\}$ such that $\mathbf{x}(T) \in \bar{\mathcal{X}}_\alpha(\epsilon)$, with T bounded by

$$T \leq \left\lceil \frac{2(f(\mathbf{x}(0)) - f^*)\alpha^2 B N^{3/2}}{\gamma(2\epsilon^2 - L\gamma\alpha^2 B N^{3/2})} \right\rceil. \quad (17)$$

The upper bound in Equation (17) is minimized when the step-size $\gamma^* = \epsilon^2/(L\alpha^2 B N^{3/2})$ is used.

Proof: Let $\epsilon > 0$ be given and choose any $\gamma \in (0, 2\epsilon^2/(L\alpha^2 B N^{3/2}))$. Then from Lemma 3 we have for all $\mathbf{x}(t) \notin \bar{\mathcal{X}}_\alpha(\epsilon)$ that

$$f(\mathbf{x}(t+1)) \leq f(\mathbf{x}(t)) - \bar{\delta}(\epsilon, \alpha, \gamma), \quad (18)$$

where $\bar{\delta}(\epsilon, \alpha, \gamma) > 0$. By recursively applying Equation (18), it follows that if $\mathbf{x}(t) \notin \bar{\mathcal{X}}_\alpha(\epsilon)$ for all $t < s$ then

$$0 \leq f(\mathbf{x}(s)) - f^* \leq f(\mathbf{x}(0)) - f^* - s \bar{\delta}(\epsilon, \alpha, \gamma). \quad (19)$$

Therefore, there must exist $T \leq \lceil (f(\mathbf{x}(0)) - f^*)/\bar{\delta}(\epsilon, \alpha, \gamma) \rceil$ such that $\mathbf{x}(T) \in \bar{\mathcal{X}}_\alpha(\epsilon)$; otherwise, we can use Equation (19) with $s = \lceil (f(\mathbf{x}(0)) - f^*)/\bar{\delta}(\epsilon, \alpha, \gamma) \rceil + 1$ to get the contradiction that $f(\mathbf{x}(s)) < f^*$. By rearranging $T \leq \lceil (f(\mathbf{x}(0)) - f^*)/\bar{\delta}(\epsilon, \alpha, \gamma) \rceil$, we obtain the bound in Equation (17). The optimal step-size $\gamma^* = \epsilon^2/(L\alpha^2 B N^{3/2})$ comes by maximizing the denominator in Equation (17). \blacksquare

The theorem shows that the Type-1 approximate optimality condition can be reached in a finite number of iterations. Moreover, Equation (17) provides an upper complexity

bound on the algorithm, showing the number of iterations needed to reach any $\epsilon > 0$ accuracy. When the step-size is $\gamma^* = \epsilon^2 / (L\alpha^2 B N^{3/2})$ then the upper bound in Equation (17) increases proportionally to $1/\epsilon^4$ as ϵ goes to zero. In Section V-A, we show that this bound can be improved for unconstrained problems (where it increases proportionally to $1/\epsilon^2$ as ϵ goes to zero).

2) *Stopping Condition of Type-2*: We now show that the Type-2 approximate optimality can be reached for any ϵ accuracy. The following lemma is used to obtain the result.

Lemma 4. *Suppose that $f \in \mathcal{F}_L(\mathbb{R}_+^N)$, ∇f is B -bounded, and the iterates $\mathbf{x}(t)$ are generated by Equation (11). Then for all $\epsilon > 0$, $\alpha > 0$, $\gamma(t) \in (0, \bar{\gamma})$, with $\bar{\gamma} = 2\epsilon^2 / (\alpha^2 B N^{3/2})$, and $T \in \mathbb{N}$ such that $\mathbf{x}(T) \in \bar{\mathcal{X}}_\alpha(\epsilon)$*

$$f(\mathbf{x}(t)) \leq \bar{F}_\alpha(\epsilon) + \frac{L}{2}\bar{\gamma}^2, \text{ for all } t \geq T, \quad (20)$$

where $\bar{F}_\alpha : \mathbb{R}_+ \rightarrow \mathbb{R} \cup \{\infty\}$ is given by

$$\bar{F}_\alpha(\kappa) = \sup\{f(\mathbf{x}) \mid \mathbf{x} \in \bar{\mathcal{X}}_\alpha(\kappa)\}. \quad (21)$$

Moreover, there exists $\kappa > 0$ such that for all $\epsilon \in [0, \kappa]$ following holds (i) $\bar{\mathcal{X}}_\alpha(\epsilon)$ is bounded and (ii) $\bar{F}_\alpha(\epsilon) < \infty$. It also holds that $\lim_{\epsilon \rightarrow 0^+} \bar{F}_\alpha(\epsilon) = f^*$.

Proof: The proof is provided in Appendix E. ■

The lemma is useful in deriving Type-2 optimality conditions since it connects the results from Theorem 3 to the quantity $f(\mathbf{x}(t)) - f^*$ via the function $\bar{F}_\alpha(\kappa)$ defined in Equation (21). In particular, the lemma provides a bound on $f(\mathbf{x}(t))$ that depends on $\bar{F}_\alpha(\epsilon)$ and the step-size. Therefore, since $\lim_{\epsilon \rightarrow 0^+} \bar{F}_\alpha(\epsilon) = f^*$ we can enforce $f(\mathbf{x}(t)) - f^*$ to be arbitrarily small after some time T , i.e., for all $t \geq T$, by choosing a small enough step-size. The idea is now formalized.

Theorem 4. *Suppose that $f \in \mathcal{F}_L(\mathbb{R}_+^N)$, ∇f is B -bounded, and the iterates $\mathbf{x}(t)$ are generated by Equation (11). Then for any $\epsilon > 0$ there exists step-size $\gamma > 0$ and $T \in \mathbb{N}$ such that $f(\mathbf{x}(t)) - f^* \leq \epsilon$ for all $t \geq T$.*

Proof: The result follows directly from Lemma 4 and Theorem 3. ■

Theorem 4 proves that the Type-2 optimality condition [Eq. (14)] can be achieved for any $\epsilon > 0$ in a finite number of iterations. However, unlike for the Type-1 optimality condition, we could not provide an explicit step-size choice or a bound on the number of iterations needed to achieve the $\epsilon > 0$ accuracy.

B. Diminishing Step-Size

We now consider the diminishing step-size case. The following result shows that the step-sizes can be chosen so that Iterates (11) converge asymptotically to the optimal solutions of Problem (1).

Theorem 5. *Suppose that $f \in \mathcal{F}_L(\mathbb{R}_+^N)$, ∇f is B -bounded, and the iterates $\mathbf{x}(t)$ are generated by Equation (11) where*

$$\lim_{t \rightarrow \infty} \gamma(t) = 0 \text{ and } \sum_{t=0}^{\infty} \gamma(t) = \infty.$$

Then $\mathbf{x}(t)$ converges to the set of optimal solutions of the optimization Problem (1), i.e., $\lim_{t \rightarrow \infty} \text{dist}(\mathbf{x}(t), \mathcal{X}^) = 0$.*

Proof: Proof is provided in Appendix F. ■

The step-size choice in the theorem is also necessary to ensure that $\lim_{t \rightarrow \infty} \text{dist}(\mathbf{x}(t), \mathcal{X}^*) = 0$ holds for all $f \in \mathcal{F}_+(\mathbb{R}_+^N)$ with ∇f being B -bounded. To see why, consider the scalar function

$$f(x) = \begin{cases} 0.5(x-1)^2 & \text{if } |x-1| \leq 1 \\ |x-1| - 0.5 & \text{otherwise.} \end{cases} \quad (22)$$

Then f has the unique optimizer $x^* = 1$ and satisfies the assumptions of Theorem 5.

Let us first show that if $\lim_{t \rightarrow \infty} \gamma(t) \neq 0$ then $\lim_{t \rightarrow \infty} \text{dist}(x(t), \mathcal{X}^*) \neq 0$. If $\gamma(t) > 0$ does not converge to zero then there exists $I \in (0, 1)$ and a subsequence t_k such that $\gamma(t_k) \geq I$ for infinitely many $t \in \mathbb{N}$. Then either $|x(t_k + 1) - x^*|$ or $|x(t_k) - x^*|$ must be larger than or equal to $I/2$ for all $k \in \mathbb{N}$ because if $|x(t_k) - x^*| \leq I/2$ then $|x(t_k + 1) - x^*| \geq I/2$. Therefore, $|x(t) - x^*| \geq I/2$ for infinitely many $t \in \mathbb{N}$ so $\limsup_{t \rightarrow \infty} \text{dist}(x(t), \mathcal{X}^*) \geq I/2$. Let us next show that if $\sum \gamma(t) < \infty$ then $\lim_{t \rightarrow \infty} \text{dist}(x(t), \mathcal{X}^*) \neq 0$. Take $x(0) = 2 + \sum_{t=0}^{\infty} \gamma(t)$. Then $x(t) \geq \mathbf{x}(0) - \sum_{\tau=0}^t \gamma(\tau) \geq 2$ for all $t \in \mathbb{N}$, so $\liminf_{t \rightarrow \infty} \text{dist}(x(t), \mathcal{X}^*) \geq 1$.

The advantage of using diminishing step-sizes, as in Theorem 5, is that the algorithm can asymptotically converge to the set of optimal values. Moreover, diminishing step-size rules can be implemented even if problem parameters such as the Lipschitz constant L are unknown, unlike when the step-size is constant. On the other hand, it is more complicated to characterize the convergence rate, similar to Equation (17), when diminishing step-sizes are used.

V. CONVERGENCE - GENERAL QUANTIZATION

In the previous section, we studied quantized gradient methods where a particular quantization based on the sign of the gradient was used for constrained optimization problems. As we discussed in Section III-B, for unconstrained problems, a more general class of quantizations called θ -covers (Definition 3) ensures that the quantized gradient methods can minimize any $f \in \mathcal{F}_L(\mathbb{R}^N)$. In this section we formally prove this, i.e., if the quantization is a θ -cover then the quantized gradient methods converge (i) approximately to an optimal solution when the step-sizes are constant and (ii) asymptotically to an optimal solution when the step-sizes are non summable and converge to zero. Moreover, we study how the quantization fineness, i.e., θ , affects the algorithm convergence. We first consider the case when the step-sizes are fixed, i.e., $\gamma(t) = \gamma$, in subsection V-A. Then in subsection V-B we consider diminishing step-sizes.

A. Constant Step-Size

Similar to Section IV-A, we consider the following two types of approximate optimality conditions:

$$\text{Type-1:} \quad \|\nabla f(\mathbf{x})\| \leq \epsilon, \quad (23)$$

$$\text{Type-2:} \quad f(\mathbf{x}(t)) - f^* \leq \epsilon. \quad (24)$$

1) *Stopping Condition of Type-1*: We start by showing that Type-1 approximate optimality can be achieved for any $\epsilon > 0$ in a finite number of iterations. Further, we provide a lower and upper bound on the number of iterations needed to achieve the ϵ -accuracy (that depends on θ). A key result used to obtain the result is the following lemma.

Lemma 5. *Suppose $f \in \mathcal{F}_L(\mathbb{R}^N)$, $\epsilon > 0$, $\theta \in [0, \pi/2)$, $\mathbf{x} \in \mathbb{R}^N$, $\|\nabla f(\mathbf{x})\| > \epsilon$, and $\mathbf{d} \in \mathcal{S}^{N-1}$ where $\text{ang}(\nabla f(\mathbf{x}), \mathbf{d}) \leq \theta$. Then*

$$f(\mathbf{x} - \gamma \mathbf{d}) \leq f(\mathbf{x}) - \delta(\epsilon, \gamma, \theta), \quad (25)$$

where

$$\delta(\epsilon, \gamma, \theta) = \left(\frac{2 \cos(\theta) \epsilon}{L} - \gamma \right) \frac{L}{2} \gamma. \quad (26)$$

Clearly, $\delta(\epsilon, \gamma, \theta) > 0$ when $\gamma \in (0, 2 \cos(\theta) \epsilon / L)$.

Proof. The proof is provided in Appendix G. \square

The lemma shows that if $\|\nabla f(\mathbf{x})\| > 0$ and $\text{ang}(\nabla f(\mathbf{x}), \mathbf{d}) \leq \theta < \pi/2$ for some $\mathbf{x} \in \mathbb{R}^N$ then the objective function value can be decreased by taking a step in the direction $-\mathbf{d}$, i.e., $f(\mathbf{x} - \gamma \mathbf{d}) \leq f(\mathbf{x})$ for small enough $\gamma > 0$. Therefore, if \mathcal{D} is a θ -cover then we can always find $\mathbf{d} \in \mathcal{D}$ and a step-size $\gamma > 0$ such that $f(\mathbf{x} - \gamma \mathbf{d}) \leq f(\mathbf{x})$. We now use this intuition from Lemma 5 to provide the upper and lower bounds on the number of iterations that are needed to reach the Type-1 optimality condition.

Theorem 6. *Suppose that $f \in \mathcal{F}_L(\mathbb{R}^N)$, \mathcal{D} is a θ -cover (Definition 3), the iterates $\mathbf{x}(t)$ are generated by Equation (8), and define the set*

$$\mathcal{X}(\epsilon) = \{\mathbf{x} \in \mathbb{R}^N \mid \|\nabla f(\mathbf{x})\| \leq \epsilon\}. \quad (27)$$

Then the following holds:

a) For any $\epsilon > 0$, if $\gamma \in (0, 2 \cos(\theta) \epsilon / L)$ then there exists $T \in \mathbb{N} \cup \{0\}$ such that $\mathbf{x}(T) \in \mathcal{X}(\epsilon)$, with T bounded by

$$T \leq \left\lceil \frac{2(f(\mathbf{x}(0)) - f^*)}{\gamma(2 \cos(\theta) \epsilon - L\gamma)} \right\rceil. \quad (28)$$

The upper bound in Equation (28) is minimized with the optimal step-size $\gamma^* = \cos(\theta) \epsilon / L$.

b) Given a fixed step-size $\gamma > 0$ and scalar $\kappa > 0$, if we choose

$$\epsilon(\kappa, \gamma) = \kappa + \gamma L / (2 \cos \theta) \quad (29)$$

then there exists $T \in \mathbb{N} \cup \{0\}$ such that $\mathbf{x}(T) \in \mathcal{X}(\epsilon(\kappa, \gamma))$, with T bounded by

$$T \leq \left\lceil \frac{f(\mathbf{x}(0)) - f^*}{\cos(\theta) \gamma \kappa} \right\rceil. \quad (30)$$

c) (Lower Bound on T) For any step-size $\gamma > 0$ and $\epsilon > 0$ if $\mathbf{x}(T) \in \mathcal{X}(\epsilon)$ then

$$\frac{\|\nabla f(\mathbf{x}(0))\| - \epsilon}{\gamma L} \leq T \quad (31)$$

Proof: a) The proof follows the same arguments as the proof of Theorem 3 using Lemma 5 in place of Lemma 3.

b) The result can be obtained by substituting $\epsilon(\kappa, \gamma)$ in Equation (28).

c) Using the fact that the gradient ∇f is L -Lipschitz continuous we have

$$\|\nabla f(\mathbf{x}(t)) - \nabla f(\mathbf{x}(t+1))\| \leq L \|\mathbf{x}(t) - \mathbf{x}(t+1)\| \leq L\gamma.$$

Therefore, using the triangle inequality, we have

$$\|\nabla f(\mathbf{x}(t))\| - L\gamma \leq \|\nabla f(\mathbf{x}(t+1))\| \quad \text{for all } t \in \mathbb{N}.$$

Recursively applying the inequality gives

$$\|\nabla f(\mathbf{x}(0))\| - L\gamma t \leq \|\nabla f(\mathbf{x}(t))\|.$$

Hence, $\|\nabla f(\mathbf{x}(t))\| \leq \epsilon$ can only hold when $t \geq (\|\nabla f(\mathbf{x}(0))\| - \epsilon) / (L\gamma)$. \blacksquare

Theorem 6-a) proves that if \mathcal{D} is a θ -cover then the Type-1 optimality condition [Eq. (23)] can be achieved with ϵ -accuracy in a finite number of iterations, for all $\epsilon > 0$. Moreover, the theorem gives an upper bound on the number of iterations needed to achieve such ϵ -accuracy. This bound decreases as θ decreases, i.e., as the quantization becomes finer. Even though the bound in Equation (28) is on the number of iteration, since $\log_2(|\mathcal{D}|)$ bits are communicated per iteration, the results shows that in the worst case scenario

$$\left\lceil \frac{2(f(\mathbf{x}(0)) - f^*)}{\gamma(2 \cos(\theta) \epsilon - L\gamma)} \right\rceil \log_2(|\mathcal{D}|) \text{ bits}$$

are needed to find $\mathbf{x} \in \mathbb{R}^N$ such that $\|\nabla f(\mathbf{x})\| \leq \epsilon$.

Theorem 6-b) demonstrates what ϵ -accuracy can be achieved for a given step-size. The parameter κ captures a trade-off between the ϵ -accuracy and the number of iterations needed to achieve that ϵ -accuracy. By optimizing over both γ and κ in Theorem 6-b) we can find an optimal bound on the accuracy ϵ that can be guaranteed in T iterations. That is to find γ and κ that solve the following optimization problem.

$$\begin{aligned} & \underset{\kappa, \gamma}{\text{minimize}} && \epsilon(\kappa, \gamma) = \kappa + \frac{L}{2 \cos(\theta)} \gamma \\ & \text{subject to} && \frac{f(\mathbf{x}(0)) - f^*}{\cos(\theta) \gamma \kappa} \leq T, \\ & && \gamma, \kappa > 0. \end{aligned} \quad (32)$$

Formally, this bound is given as follows.

Corollary 1. *For any $T \in \mathbb{N}$ we have:*

i) *The minimal bound $\epsilon(\kappa, \gamma)$ [Eq. (29)] achieved in T iterations, i.e., the solution to Problem (32), is*

$$\epsilon^* = \frac{\sqrt{2L(f(\mathbf{x}(0)) - f^*)}}{\cos(\theta) \sqrt{T}} \quad (33)$$

where the corresponding optimal γ and κ are

$$\gamma^* = \sqrt{\frac{2(f(\mathbf{x}(0)) - f^*)}{LT}} \quad \text{and} \quad \kappa^* = \frac{\sqrt{L(f(\mathbf{x}(0)) - f^*)}}{\cos(\theta) \sqrt{2T}}. \quad (34)$$

ii) *If the step-size γ is chosen as γ^* in Equation (34) then*

$$\|\nabla f(\mathbf{x}^*(T))\| \leq \epsilon^* \quad (35)$$

where ϵ^* is as in Equation (33) and

$$\mathbf{x}^*(T) = \underset{\mathbf{x} \in \{\mathbf{x}(t) \mid t=0, \dots, T\}}{\text{argmin}} \|\nabla f(\mathbf{x})\|. \quad (36)$$

Proof: i) First note that Problem (32) is convex, which can be seen by equivalently writing it as

$$\begin{aligned} \underset{\kappa, \gamma}{\text{minimize}} \quad & \kappa + \frac{L}{2 \cos(\theta)} \gamma \\ \text{subject to} \quad & \left(\frac{f(\mathbf{x}(0)) - f^*}{\cos(\theta)T} \right) \frac{1}{\gamma} - \kappa \leq 0, \\ & \gamma, \kappa > 0, \end{aligned} \quad (37)$$

and recalling that the reciprocal $1/\gamma$ is a convex function for $\gamma > 0$. It can be checked that κ^* and γ^* satisfy the KKT condition with the Lagrangian multiplier $\lambda^*=1$.

ii) Follows directly from part i). ■

In addition to minimizing the bound in Equation (29), Corollary 1 gives insights into the convergence of Iterations (8). For example, when T is fixed, then the upper bound in Equation (33) gets larger as θ decreases. As a result, when the quantization set \mathcal{D} becomes coarser then less accuracy can be ensured. Moreover, the results show that $\|\nabla f(\mathbf{x}^*(T))\|$ converges, in the worst case, at the rate $o(1/\sqrt{T})$ to 0.

In Corollary 1 we used the step-size γ^* given in Equation (34). To compute γ^* the optimal objective function value f^* is needed, which is usually not available prior to solving Problem (1). However, some bounds on the quantity $f(\mathbf{x}(0)) - f^*$ are often available. Any such upper bound $K \in \mathbb{R}$, with $K \geq f(\mathbf{x}(0)) - f^*$, can be used to obtain similar results as to those in Corollary 1 by replacing $f(\mathbf{x}(0)) - f^*$ by K .

Corollary 2. Take $T \in \mathbb{N}$ and $K \in \mathbb{R}$ such that $K \geq f(\mathbf{x}(0)) - f^*$. If we choose the step-size as $\gamma = 2K/(LT)$ then

$$\|\nabla f(\mathbf{x}^*(T))\| \leq \frac{\sqrt{2LK}}{\cos(\theta)\sqrt{T}}, \quad (38)$$

where $\mathbf{x}^*(t)$ is chosen as in Equation (36).

We next demonstrate how the convergence results translate to Type-2 stopping conditions [Eq. (24)].

2) *Stopping Condition of Type-2:* We now show that the Type-2 approximate optimality [Eq. (24)] can be achieved for any accuracy $\epsilon > 0$. The result is based on the following lemma.

Lemma 6. Suppose $f \in \mathcal{F}_L(\mathbb{R}^N)$, \mathcal{D} is a θ -cover, and the iterates $\mathbf{x}(t)$ are generated by Equation (8), then:

a) for any $\epsilon > 0$, $\gamma(t) \in (0, \bar{\gamma})$, where $\bar{\gamma} = 2 \cos(\theta)\epsilon/L$, and $T \in \mathbb{N}$ such that $\mathbf{x}(T) \in \mathcal{X}(\epsilon)$, the following holds

$$f(\mathbf{x}(t)) \leq F(\epsilon) + \frac{L}{2}\gamma^2, \quad \text{for all } t \geq T, \quad (39)$$

where $F: \mathbb{R}_+ \rightarrow \mathbb{R} \cup \{\infty\}$ is given by

$$F(\kappa) = \sup\{f(\mathbf{x}) | \mathbf{x} \in \mathcal{X}(\kappa)\}. \quad (40)$$

There exists $\kappa > 0$ such that for all $\epsilon \in [0, \kappa]$ (i) $\mathcal{X}(\epsilon)$ is bounded and $F(\epsilon) < \infty$. Moreover, $\lim_{\epsilon \rightarrow 0^+} F(\epsilon) = f^*$.

b) if f is μ -strongly convex then we have

$$F(\epsilon) \leq f^* + \epsilon^2/(2\mu). \quad (41)$$

Proof: The proof is provided in Appendix H. ■

Lemma 6 is useful in obtaining Type-2 approximate optimality as it connects the quantity $f(\mathbf{x}(t)) - f^*$ to Theorem 6 via the function $F(\epsilon)$ in Equation (40). In particular, Part a) of Lemma 6 bounds $f(\mathbf{x}(t))$ by a constant that depends on $F(\cdot)$, where $F(\epsilon)$ converges to f^* as ϵ converges to 0. Therefore, by using the dependence of $F(\epsilon)$ on $\mathcal{X}(\epsilon)$, defined in Equation (27), we can connect $f(\mathbf{x}(t)) - f^*$ to the convergence result in Theorem 6 to ensure that the Type-2 stopping condition can be achieved for any $\epsilon > 0$. Part b) of Lemma 6 then illustrates how the upper bound on $f(\mathbf{x}(t))$ depending on $F(\epsilon)$ can be further improved when f is μ -strongly convex. These ideas are formally illustrated in the following theorem.

Theorem 7. Suppose $f \in \mathcal{F}_L(\mathbb{R}^N)$, \mathcal{D} is a θ -cover, and $\mathbf{x}(t)$ are generated by Equation (8). Then for any $\epsilon > 0$:

- a) there exists a step-size $\gamma > 0$ and $T \in \mathbb{N}$ such that $f(\mathbf{x}(t)) - f^* < \epsilon$ for all $t \geq T$,
b) moreover, if f is μ -strongly convex and $\gamma \in (0, \bar{\gamma})$ where

$$\bar{\gamma} = \min \left\{ \frac{2 \cos(\theta) \sqrt{\mu\epsilon}}{L}, \sqrt{\frac{\epsilon}{L}} \right\}$$

then $f(\mathbf{x}(t)) - f^* \leq \epsilon$ for all $t \geq T$ where

$$T \leq \left\lceil \frac{2(f(\mathbf{x}(0)) - f^*)}{\gamma(2 \cos(\theta) \sqrt{\mu\epsilon} - L\gamma)} \right\rceil. \quad (42)$$

Proof: a) The result follows directly from Lemma 6-a) and Theorem 6-a).

b) Since $\gamma \in (0, 2 \cos(\theta) \sqrt{\mu\epsilon}/L)$ it follows from Theorem 6-a) there exists T bounded as in Equation (28) such that $\mathbf{x}(T) \in \mathcal{X}(\sqrt{\mu\epsilon})$. From Lemma 6-a), for all $t \geq T$ that

$$f(\mathbf{x}(t)) - f^* \leq F(\sqrt{\mu\epsilon}) - f^* + \frac{L}{2}\gamma^2 \leq \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon,$$

where the second inequality follows from Lemma 6-b) and that $\gamma \leq \sqrt{\epsilon/L}$. ■

Theorem 7 shows that when \mathcal{D} is a θ -cover the Type-2 optimality condition [Eq. (24)] can be achieved in finite number of iterations. For general functions, the theorem does not provide a step-size γ_ϵ that can achieve any particular ϵ -accuracy, even though such γ_ϵ always exist. This is challenging in general, as it can be difficult to bound the function $F(\cdot)$ for general convex functions $f(\cdot)$. Nevertheless, part b) of the proof shows that when f is μ -strongly convex, then a range of step-sizes that ensure a given $\epsilon > 0$ accuracy is provided. Moreover, when f is μ -strongly convex then we can obtain similar bound on number of iterations needed to achieve that ϵ -accuracy as in Equation (28) in Theorem 6-a).

B. Diminishing Step-Size

We now consider the diminishing step-size case. The following result shows that the step-sizes can actually be chosen so Iterates (8) converges to the optimal solution to Problem (1).

Theorem 8. Suppose that $f \in \mathcal{F}_L(\mathbb{R}^N)$, \mathcal{D} is a θ -cover, and that the iterates $\mathbf{x}(t)$ are generated by Equation (8). If the step-sizes $\gamma(t) \geq 0$ are chosen so that

$$\lim_{t \rightarrow \infty} \gamma(t) = 0 \quad \text{and} \quad \sum_{t=0}^{\infty} \gamma(t) = \infty$$

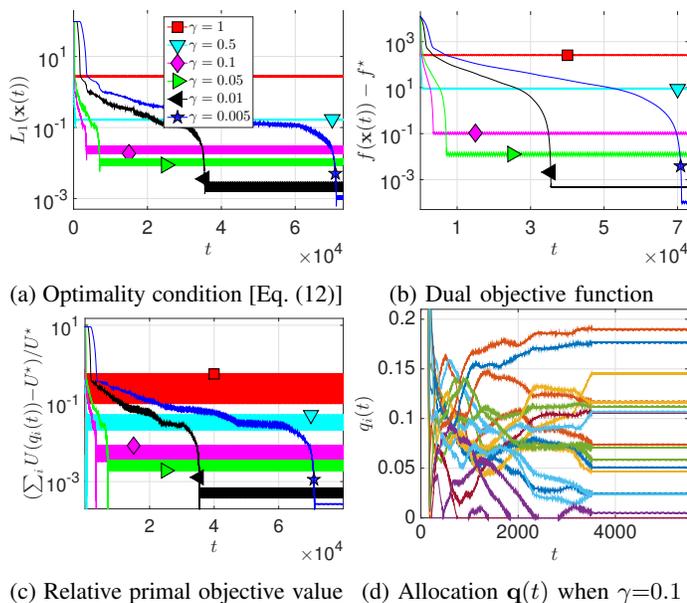


Fig. 2: Optimalty measure $L_1(\cdot)$ from Equation (12) (Fig. 2a), dual (Fig. 2b), and primal (Fig. 2c) objective function value at every iteration for different step-sizes. The resource allocations $q_i(t)$ (Fig. 2d) at every iteration when $\gamma = 0.1$.

then $\lim_{t \rightarrow \infty} \text{dist}(\mathbf{x}(t), \mathcal{X}^*) = 0$.

Proof: The proof is provided in Appendix I. ■

The step-size choice in the theorem is necessary to ensure that $\lim_{t \rightarrow \infty} \text{dist}(\mathbf{x}(t), \mathcal{X}^*) = 0$ for all $f \in \mathcal{F}_L(\mathbb{R}^N)$, consider the scalar function f defined in Equation (22).

Theorem 8 shows that when \mathcal{D} is a θ -cover then there exists a step-size rule such that every limit point of the quantized gradient methods is an optimal solution to Problem (1). A particular implication of this result is that every θ -cover is a proper quantization, see Definition 2. Therefore, Theorem 8 actually proves one direction of the equivalence established in Theorem 1.

VI. NUMERICAL ILLUSTRATIONS

We now illustrate how the studied algorithms perform on two of the application examples discussed in Section II. We compare the numerical performance with some of the theoretical results in the paper and with algorithms that use perfect communication with no quantization.

A. TCP Flow Control with Binary Feedback

We illustrate the convergence of Iterates (11) on the TCP Flow Control in Section II-A. We consider a network with $S = 20$ sources and $N = 100$ links. We use the same utility functions as in Experiment 1 of [7, Section VI-B], i.e., $U_s(q_s) = 1000 \log(1 + q_s)$, for $s \in \mathcal{S}$. Similarly as in [28], we generate the network matrix \mathbf{A} [Eq. (5)] randomly so each entry of \mathbf{A} is 1 with the probability 1/2 and 0 otherwise. We set $c_l = 1$ for all $l \in \mathcal{L}$. The local constraint of each source is $[0, 1]$. We consider the step-sizes $\gamma = 0.005, 0.01, 0.05, 0.1, 0.5$, and 1. Note that in the figures described below,

some of lines that appear to be thick lines actually show some fluctuations.

Fig. 2a depicts the optimalty measure $L_1(\mathbf{x}(t))$, see Equation (12) in Lemma 2. From Lemma 2, $\mathbf{x} \in \mathbb{R}_+^N$ is an optimal solution to Problem (1) if and only if $L_1(\mathbf{x}) = 0$. For all step-sizes γ , the measure $L_1(\mathbf{x}(t))$ converges to some small error floor and then fluctuates slightly there. For smaller step-sizes γ the optimalty measure $L_1(\mathbf{x}(t))$ converges to smaller values, roughly to 2.6, 0.16, 0.019, 0.009, 0.002, 0.001 for $\gamma = 1, 0.5, 0.01, 0.005, 0.001$, and 0.0005. These results show that the step-size choices γ in Theorem 3 are conservative.⁴ For example, to that ensure $\epsilon = 0.1$ accuracy in Theorem 3 the step-sizes should be $\gamma \in (0, 0.00002)$ but in for this example the step-size chooses $\gamma \leq 0.1$ achieve the $\epsilon = 0.1$ accuracy.

Figs. 2b and 2c depict the dual and primal objective function values at every iteration. The figures demonstrate a similar convergence behaviour of the primal/dual objective function values as in the optimalty measure $L_1(\mathbf{x}(t))$ in Fig. 2a. For the dual objective value these results agree with the results in Theorem 4. Finally, Fig. 2d illustrates the convergence of the data rate allocation to each source when γ . The results show that the all the data rate allocations converges after roughly 3500 iterations and then fluctuate slightly there.

B. Task Allocation

We now illustrate the performance of the quantized gradient methods on the Task Allocation Problem (4) from Section II-A with $K=4$ machines and $N=2$ tasks. For each machine $i=1, 2, 3, 4$ we have the cost function $C_i(\mathbf{q}_i) = a_i \mathbf{q}_{i,1}^2 + b_i \mathbf{q}_{i,2}^2$ where a_i and b_i are uniform random random variables on the interval $[1, 5]$. The private constraint of machine $i = 1, 2, 3, 4$ is $\mathcal{Q}_i = \{(x, y) \in \mathbb{R}^2 | x, y \geq 0, x + y \leq 3\}$. Clearly, $-C_i$ are strongly concave with concavity parameters $\mu_i = \min\{a_i, b_i\} \geq 1$. It can be verified that the dual gradient is L -Lipschitz continuous with $L = 4/\mu$, where $\mu = \min\{\mu_1, \dots, \mu_4\}$. The step-size is $\gamma = 0.1$ and the initialization is $\mathbf{x}(0) = (0, 0)$ (recall that \mathbf{x} is the dual variable). We use the quantization set \mathcal{D} from Example 3 when 2, 3, and 4 bits are communicated per iteration, i.e., when $|\mathcal{D}| = 4, 8, 16$, see Remark 1.

Fig. 3 depicts the norm of the gradient and the primal objective function at every iteration of the algorithm. The norm of the gradient $\|\nabla f\|$ reaches the accuracy $\epsilon = 0.1$ in roughly 51, 56, and 65 iterations using 204, 168, 130 bits when 2, 3, and 4 bits are communicated per iteration, respectively. We compare the results to Iterations (2) and (3) where no quantization is done, i.e., infinite bandwidth is used. Fig. 3a shows that by using 4 bits per iteration, the results achieved by QGM are almost as good as when the full gradient direction is communicated using Iterations (3). However, the QGMs do not perform as well as Iterations (2); this is to be expected, since in Iterations (2) the full direction and magnitude of the gradient is known. These results illustrate that we can dramatically reduce the number of bits communicated without sacrificing much in performance.

⁴The parameters in Theorem 3 for this problem are the parameters used are $\alpha = 1, B = \sqrt{N}, \mu = 250, \bar{S} = 58, \bar{L} = 17$ (see Lemma ??).

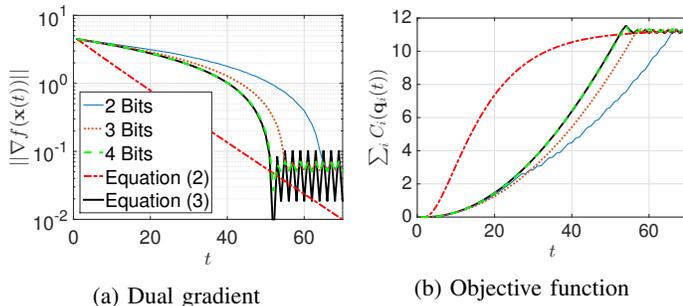


Fig. 3: The dual gradient (Fig. 3a) and the primal objective function value (Fig. 3a) at every iteration when 2, 3, and 4 bits are communicated per iteration. The results are compared with the algorithms in Equations (2) and (3).

VII. CONCLUSIONS AND FUTURE WORK

This paper studied gradient methods where the gradient direction is quantized at every iteration of the algorithm. Such methods are of interest, for example, in distributed optimization where the gradient can often be measured but has to be communicated to accomplish the algorithm. An instance of such a procedure is dual decomposition, where primal problems that are scattered between different entities are solved in a distributed fashion by performing a gradient descent on the dual problem, see Section II for examples. Our results show that a variant of the projected dual descent taking the sign of the dual gradients, i.e.,

$$\mathbf{x}(t+1) = [\mathbf{x}(t) - \gamma(t) \text{sign}(\nabla f(\mathbf{x}(t)))]_{\mathcal{X}}, \quad (43)$$

converges to the optimal solution under mild conditions on f when $\mathcal{X} = \mathbb{R}^N$ and $\mathcal{X} = \mathbb{R}_+^N$, i.e., for dual problems associated with primal problems with equality and inequality constraints, respectively. Therefore, when different entities maintain the components of the variable \mathbf{x} then each entity only needs to broadcast one bit per iteration to ensure convergence to the optimal solution. Our results also show that when a single entity maintains \mathbf{x} then the minimal quantization has cardinality $N+1$; for smaller quantizations there exists an optimization problem that the quantized gradient methods cannot solve. Therefore, only $\log_2(N+1)$ bits/iteration are communicated instead of N bits/iteration as when the components of \mathbf{x} are maintained by different entities in Equation (43). We also connect fineness of the quantization to the convergence rate of the algorithm to the available bandwidth (bits/iteration). The convergence rate improves as the bandwidth is increased.

Future work will consider how to additionally quantize the magnitude of the gradient to get a better trade-off between convergence rate and available bandwidth. Moreover, it is interesting to see if the results can be generalized to nonsmooth optimization problems.

APPENDIX A PROOF OF THEOREM 1

Proof: Let us start by showing by contradiction that \mathcal{D} being a proper quantization for the problem class $\mathcal{F}_L(\mathbb{R}^N)$ implies that there exists $\theta \in [0, \pi/2)$ such that \mathcal{D} is

a θ -cover. Suppose there does not exist such θ . Then $\min_{\mathbf{a} \in \mathcal{S}^{N-1}} \max_{\mathbf{d} \in \mathcal{D}} \cos(\text{ang}(\mathbf{a}, \mathbf{d})) \leq 0$ since the function $g(\mathbf{a}) = \max_{\mathbf{d} \in \mathcal{D}} \cos(\text{ang}(\mathbf{a}, \mathbf{d})) = \max_{\mathbf{d} \in \mathcal{D}} \langle \mathbf{a}, \mathbf{d} \rangle$ is continuous and \mathcal{S}^{N-1} is compact. Therefore, there exists $\mathbf{a} \in \mathcal{S}^{N-1}$ such that $\cos(\text{ang}(\mathbf{a}, \mathbf{d})) \leq 0$ for all $\mathbf{d} \in \mathcal{D}$. In particular, we have for all $\mathbf{d} \in \mathcal{D}$ that $\langle \mathbf{a}, \mathbf{d} \rangle = \cos(\text{ang}(\mathbf{a}, \mathbf{d})) \leq 0$. By choosing $\mathbf{x}(0) = \mathbf{a}$, using Iterations (8) and Cauchy-Schwarz inequality we conclude that for all $t \in \mathbb{N}$

$$\|\mathbf{x}(t)\| \geq \langle \mathbf{a}, \mathbf{x}(t) \rangle = \langle \mathbf{a}, \mathbf{a} \rangle - \sum_{i=0}^{t-1} \gamma(t) \langle \mathbf{a}, \mathbf{d}(t) \rangle \geq 1,$$

where the inequality follows from the fact that $\|\mathbf{a}\| = 1$ and that for all $\mathbf{d} \in \mathcal{D}$ we have $\langle \mathbf{a}, \mathbf{d} \rangle \leq 0$. If we choose $f(\mathbf{x}) = (L/2)\|\mathbf{x}\|^2$, then $f \in \mathcal{F}_L(\mathbb{R}^N)$ and f has the unique optimizer $\mathbf{x}^* = \mathbf{0}$, but $\text{dist}(\mathbf{x}(t), \mathcal{X}^*) = \|\mathbf{x}(t)\| \geq 1$, for all $t \in \mathbb{N}$. Since $\text{dist}(\mathbf{x}(t), \mathcal{X}^*) \geq 1$ for all $\mathbf{d}(t) \in \mathcal{D}$ and $\gamma(t) \in \mathbb{R}_+$, we can conclude that \mathcal{D} is not a proper quantization.

The fact that \mathcal{D} being a θ -cover implies that \mathcal{D} is a proper quantization, follows from Theorem 8 in Section V-B, where we showed that for all $f \in \mathcal{F}_L(\mathbb{R}^N)$ we can choose $\mathbf{d}(t) \in \mathcal{D}$ and $\gamma(t) \in \mathbb{R}_+$ such that $\lim_{t \rightarrow \infty} \text{dist}(\mathbf{x}(t), \mathcal{X}^*) = 0$. ■

APPENDIX B PROOF OF THEOREM 2

Proof: First consider the case where either $|\mathcal{D}| < N$ or $|\mathcal{D}| = N$ and the elements of \mathcal{D} are linearly dependent. Then $\text{Span}(\mathcal{D})$ is a proper subspace of \mathbb{R}^N , so there exists a normal $\mathbf{a} \in \mathcal{S}^{N-1}$ such that $\cos(\text{ang}(\mathbf{a}, \mathbf{d})) = \langle \mathbf{a}, \mathbf{d} \rangle \leq 0$ for all $\mathbf{d} \in \text{Span}(\mathcal{D})$. Since $\mathcal{D} \subseteq \text{Span}(\mathcal{D})$, \mathcal{D} is not a θ -cover for any $\theta \in [0, \pi/2)$ and the result follows from Theorem 1.

Let us next consider the other case, where $|\mathcal{D}| = N$ and the vectors of \mathcal{D} are linearly independent, i.e., $\text{Span}(\mathcal{D}) = \mathbb{R}^N$. Define $\mathbf{D} \in \mathbb{R}^{N \times N}$ such that for $i = 1, \dots, N$ row i in \mathbf{D} is the i -th element of \mathcal{D} , where the elements have some arbitrary order. Then \mathbf{D} is invertible and we can choose $\mathbf{a} = -\mathbf{D}^{-1}\mathbf{1}$ where $\mathbf{1} \in \mathbb{R}^N$ is a vector of all ones. Then we have for $i = 1, \dots, N$ that $\langle \mathbf{d}_i, \mathbf{a} \rangle = -\mathbf{d}_i \mathbf{D}^{-1} \mathbf{1} = -1$. Hence, as in the previous case, we get that $\langle \mathbf{a}, \mathbf{d} \rangle \leq 0$ for all $\mathbf{d} \in \mathcal{D}$ implying that \mathcal{D} can not be a θ -cover for any $\theta \in [0, \pi/2)$, and the result follows from Theorem 1. ■

APPENDIX C PROOF OF LEMMA 2

Proof: Since f and \mathcal{X} are convex, $\mathbf{x} \in \mathcal{X}^*$ if and only if the KKT optimality conditions hold for \mathbf{x} [32, Section 5.5]. It can be checked that since $\mathcal{X} = \mathbb{R}_+^N$ the optimal dual variable associated to $\mathbf{x} \in \mathcal{X}^*$ is $\boldsymbol{\lambda} = \nabla f(\mathbf{x})$. Therefore, the KKT conditions reduce to the following three conditions holding for all $i = 1, \dots, N$: (i) $\nabla_i f(\mathbf{x}) \mathbf{x}_i = 0$, (ii) $\nabla_i f(\mathbf{x}) \geq 0$, and (iii) $\mathbf{x}_i \geq 0$. We now show both directions of the proof.

First assume that $L(\mathbf{x}) = 0$. We show that (i), (ii), and (iii) hold, so $\mathbf{x} \in \mathcal{X}^*$. We have $\mathbf{x} = [\mathbf{x} - \alpha \nabla f(\mathbf{x})]^+$ or $\mathbf{x}_i = [\mathbf{x}_i - \alpha \nabla_i f(\mathbf{x})]^+$, for $i = 1, \dots, N$. So (i) holds because if $\mathbf{x}_i \neq 0$ then $\nabla_i f(\mathbf{x}) = 0$ and if $\nabla_i f(\mathbf{x}) \neq 0$ then $\mathbf{x}_i = 0$. Similarly, (ii) holds because if $\mathbf{x}_i \neq 0$ then $\nabla_i f(\mathbf{x}) = 0$ and if $\mathbf{x}_i = 0$ then $\nabla_i f(\mathbf{x}) \geq 0$. Finally, (iii) holds because $[\cdot]^+$ is the projection to \mathbb{R}_+ .

Now assume $\mathbf{x} \in \mathcal{X}^*$ so (i), (ii), and (iii) above hold. If $\mathbf{x}_i = 0$ for some $i = 1, \dots, N$ then $\mathbf{x}_i = \lceil \mathbf{x}_i - \alpha \nabla_i f(\mathbf{x}) \rceil^+$ since by (ii) $\nabla_i f(\mathbf{x}) \geq 0$. Otherwise, if $\mathbf{x}_i > 0$, then $\nabla_i f(\mathbf{x}) = 0$ by (i) so $\mathbf{x}_i = \lceil \mathbf{x}_i - \alpha \nabla_i f(\mathbf{x}) \rceil^+$. ■

APPENDIX D PROOF OF LEMMA 3

Proof: For all $\beta \in [0, 1]$ we have

$$\frac{\beta\epsilon}{\sqrt{N}} \leq \frac{\beta}{\sqrt{N}} \|\mathbf{x} - \lceil \mathbf{x} - \alpha \nabla f(\mathbf{x}) \rceil^+\| \quad (44)$$

$$\leq \beta \|\mathbf{x} - \lceil \mathbf{x} - \alpha \nabla f(\mathbf{x}) \rceil^+\|_\infty \quad (45)$$

$$= \beta \|\mathbf{x}_i - \lceil \mathbf{x}_i - \alpha \nabla_i f(\mathbf{x}) \rceil^+\| \quad (46)$$

$$\leq \|\mathbf{x}_i - \lceil \mathbf{x}_i - \alpha \beta B \text{sign}(\nabla_i f(\mathbf{x})) \rceil^+\| \quad (47)$$

where, in Equation (46), $i = \text{argmax}_j |x_j(t) - \lceil x_j(t) - \alpha \nabla_j f(x(t)) \rceil^+|$, Equation (44) comes from the bound in Equation (15), Equation (45) comes from the equivalence of the 2- and ∞ -norms, Equation (46) comes by using the definition of i , Equation (47) comes from Equations (67) and (68) in Lemma 8 in Appendix J and that $\|\nabla f\| \leq B$. By taking $\beta = \gamma/(\alpha B \sqrt{N})$ in Equation (47) we get

$$\frac{\epsilon\gamma}{\alpha B N} \leq \|\mathbf{x}_i - \lceil \bar{\mathbf{x}}_i \rceil^+\|, \quad (48)$$

where $\bar{\mathbf{x}} = \mathbf{x} - (\gamma/\sqrt{N}) \text{sign}(\nabla f(\mathbf{x}))$. Moreover, from Equation (46) and the nonexpansiveness of projections, see [26, Proposition B.11-c)], we have

$$\frac{\epsilon}{\alpha \sqrt{N}} \leq \|\nabla_i f(\mathbf{x})\|. \quad (49)$$

Therefore, we have

$$\langle \nabla f(\mathbf{x}), \mathbf{x} - \lceil \bar{\mathbf{x}} \rceil^+ \rangle = \sum_{j=1}^N \nabla_j f(\mathbf{x}) (\mathbf{x}_j - \lceil \bar{\mathbf{x}}_j \rceil^+), \quad (50)$$

$$\geq \nabla_i f(\mathbf{x}) (\mathbf{x}_i - \lceil \bar{\mathbf{x}}_i \rceil^+), \quad (51)$$

$$\geq \frac{\epsilon^2 \gamma}{\alpha^2 B N^{3/2}}, \quad (52)$$

where Equation (51) comes from the fact that every component of the sum is nonnegative, see Equation (69) in Lemma 8 in Appendix J, and Equation (52) comes from using the bound in Equations (48) and (49). Inequality (52) and the descent lemma [37, eq. (2.1.6)] yield

$$\begin{aligned} f(\lceil \bar{\mathbf{x}} \rceil^+) &\leq f(\mathbf{x}) - \langle \nabla f(\mathbf{x}), \mathbf{x} - \lceil \bar{\mathbf{x}} \rceil^+ \rangle + \frac{L}{2} \|\mathbf{x} - \lceil \bar{\mathbf{x}} \rceil^+\|^2, \\ &\leq f(\mathbf{x}) - \frac{\epsilon^2 \gamma}{\alpha^2 B N^{3/2}} + \frac{L}{2} \gamma^2, \end{aligned}$$

where the last term comes from the fact that $\|\mathbf{x} - \lceil \bar{\mathbf{x}} \rceil^+\| \leq \gamma$ by the non-expansiveness of the projection. ■

APPENDIX E PROOF OF LEMMA 4

Proof: Step 1: We prove by induction that Equation (20) holds for all $t \geq T$. When $t = T$ then Equation (20) holds because $f(\mathbf{x}(T)) \leq \bar{F}_\alpha(\epsilon)$ by definition of \bar{F}_α , see Equation (21). Now suppose Equation (20) holds for $t \geq T$.

We will show that (20) also holds for $t + 1$. Consider first the case when $\mathbf{x}(t) \in \bar{\mathcal{X}}_\alpha(\epsilon)$. Then from [37, eq. (2.1.6)]

$$\begin{aligned} f(\mathbf{x}(t+1)) &\leq f(\mathbf{x}(t)) - \langle \nabla f(\mathbf{x}(t)), \mathbf{x}(t) - \lceil \bar{\mathbf{x}}(t) \rceil^+ \rangle \\ &\quad + \frac{L}{2} \|\mathbf{x}(t) - \lceil \bar{\mathbf{x}}(t) \rceil^+\|^2, \\ &\leq \bar{F}_\alpha(\epsilon) + \frac{L}{2} \gamma^2 \end{aligned}$$

where $\bar{\mathbf{x}}(t) = \mathbf{x}(t) - (\gamma(t)/\sqrt{N}) \text{sign}(\nabla f(\mathbf{x}(t)))$, the second inequality comes from the fact that (i) that $f(\mathbf{x}(t)) \leq \bar{F}_\alpha(\epsilon)$ since $\mathbf{x}(t) \in \bar{\mathcal{X}}_\alpha(\epsilon)$, (ii) that the inner product term is non-negative because every term of the sum [Eq. (50)] is non-negative following Equation (69) in Lemma 8, and (iii) that $\|\mathbf{x}(t) - \lceil \bar{\mathbf{x}}(t) \rceil^+\| \leq \bar{\gamma}$ because of the non-expansiveness of the projection $\lceil \cdot \rceil^+$, see [26, Proposition B.11-c)]. Otherwise, if $\mathbf{x}(t) \notin \bar{\mathcal{X}}_\alpha(\epsilon)$ then $f(\mathbf{x}(t+1)) \leq f(\mathbf{x}(t))$ by Lemma 3, yielding the result.

Step 2: We will prove that there exists $\kappa > 0$ such that (i) $\bar{\mathcal{X}}_\alpha(\epsilon)$ is bounded set and (ii) $\bar{F}_\alpha(\epsilon) < \infty$ for all $\epsilon \in [0, \kappa]$. Part (i) follows directly from Lemma 9 in Appendix J. To prove part (ii), note that $\bar{\mathcal{X}}_\alpha(\epsilon)$ is closed set and also bounded for all $\epsilon \in [0, \kappa]$ for some $\kappa > 0$ from part (i). In particular, $\bar{\mathcal{X}}_\alpha(\epsilon)$ is compact set for all $\epsilon \in [0, \kappa]$ so the supremum in (21) is attained and hence $\bar{F}_\alpha(\epsilon) < \infty$.

Step 3: We will prove that $\lim_{\epsilon \rightarrow 0^+} \bar{F}_\alpha(\epsilon) = f^*$. In particular, we show that \bar{F}_α is continuous at 0 which implies the result, since $\bar{F}_\alpha(0) = f^*$. Take any sequence $(\epsilon_k)_{k \in \mathbb{N}}$ in \mathbb{R}_+ such that $\lim_{k \rightarrow \infty} \epsilon_k = 0$. Then there exists $K \in \mathbb{N}$ and a sequence $(\mathbf{x}(k))_{k \in \mathbb{N}}$ such that $f(\mathbf{x}(k)) = \bar{F}_\alpha(\epsilon)$ holds for all $k \geq K$, since $\bar{\mathcal{X}}_\alpha(\epsilon)$ is compact for all $\epsilon \in [0, \kappa]$, where κ is chosen as in **Step 2**. Moreover, by the definition of $\bar{\mathcal{X}}_\alpha(\epsilon)$ we have that $\lim_{k \rightarrow \infty} L_\alpha(\mathbf{x}(k)) = 0$. Now since L_α is a continuous function we can conclude that for every limit point $\bar{\mathbf{x}}$ of $(\mathbf{x}(k))_{k \in \mathbb{N}}$ it holds that $L_\alpha(\bar{\mathbf{x}}) = 0$, i.e., $\bar{\mathbf{x}} \in \mathcal{X}^*$ or $f(\bar{\mathbf{x}}) = f^*$. Since $f(\bar{\mathbf{x}}) = f^*$ holds for every limit point of $(\mathbf{x}(k))_{k \in \mathbb{N}}$ and f is continuous we can conclude that $\lim_{k \rightarrow \infty} f(\mathbf{x}(k)) = \lim_{k \rightarrow \infty} \bar{F}_\alpha(\epsilon_k) = f^*$. ■

APPENDIX F PROOF OF THEOREM 5

Proof: Step 1: We will prove by contradiction that for any $\alpha > 0$

$$I := \liminf_{t \rightarrow \infty} L_\alpha(\mathbf{x}(t)) = 0. \quad (53)$$

Suppose, to the contrary that $I > 0$. Choose T such that $L_\alpha(\mathbf{x}(t)) \geq I/2$ and $\gamma(t) < \min\{1, I^2/(4L\alpha^2 B N^{3/2})\}$ for all $t \geq T$. Then by Lemma 3 we get for all $t \geq T$

$$f(\mathbf{x}(t+1)) \leq f(\mathbf{x}(t)) - \frac{I^2 \gamma(t)}{4\alpha^2 B N^{3/2}} + \frac{L}{2} \gamma(t)^2 \quad (54)$$

$$= f(\mathbf{x}(t)) - \frac{I^2 \gamma(t)}{8\alpha^2 B N^{3/2}} + \left(\gamma(t) - \frac{I^2}{4L\alpha^2 B N^{3/2}} \right) \frac{L}{2} \gamma(t)$$

$$\leq f(\mathbf{x}(t)) - \frac{I^2 \gamma(t)}{8\alpha^2 B N^{3/2}}. \quad (55)$$

Since Equation (55) holds for all $t \geq T$ we obtain

$$f(\mathbf{x}(t)) \leq f(\mathbf{x}(T)) - \frac{I^2}{8\alpha^2 B N^{3/2}} \sum_{\tau=T}^{t-1} \gamma(\tau), \text{ for } t \geq T. \quad (56)$$

Since $\gamma(t)$ is non-summable Equation (56) implies that $\lim_{t \rightarrow \infty} f(\mathbf{x}(t)) = -\infty$, which contradicts the fact that \mathcal{X}^* is non-empty. Therefore we can conclude that $I = 0$, [Eq. (53)].

Step 2: We will prove that $\lim_{t \rightarrow \infty} f(\mathbf{x}(t)) = f^*$. Let $\epsilon > 0$ be given. Choose $\kappa > 0$ such that $\bar{F}_\alpha(\kappa) < f^* + \epsilon/2$, where $\bar{F}_\alpha(\kappa)$ is defined in Equation (21) of Lemma 4, such κ exists since $\lim_{\epsilon \rightarrow 0^+} F(\epsilon) = f^*$. Now choose T such that $\mathbf{x}(T) \in \bar{\mathcal{X}}_\alpha(\kappa)$ [Eq. (16)] and for all $t \geq T$ it holds that $\gamma(t) \leq \bar{\gamma} := \sqrt{\epsilon/L}$, such T exists because of Equation (53) and that $\lim_{t \rightarrow \infty} \gamma(t) = 0$. Then from Equation (20) in Lemma 4 we have for all $t \geq T$ that

$$f(\mathbf{x}(t)) - f^* \leq \bar{F}_\alpha(\kappa) - f^* + \frac{L}{2} \bar{\gamma}^2 \leq \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon.$$

Step 3: We will prove that the sequence $\mathbf{x}(t)$ is bounded. Take $\kappa > 0$ such that $\bar{F}_\alpha(\kappa) < \infty$, where \bar{F}_α is defined in Equation (21), such κ exists by Lemma 4. From Equation (53), $\mathbf{x}(t) \in \bar{\mathcal{X}}_\alpha(\kappa)$ holds for infinitely many $t \in \mathbb{N}$. Let $\mathbf{x}(t_k)$, with $k \in \mathbb{N}$, be the subsequence of all $\mathbf{x}(t) \in \bar{\mathcal{X}}_\alpha(\kappa)$. Choose $T \in \mathbb{N}$ such that $\gamma(t) \leq \epsilon^2/(L\alpha^2BN^{3/2})$ for all $t \geq T$. Then, by following the same steps as used to obtain Equations (55) and (56) and using the fact that $f(\mathbf{x}(t_k)) \leq \bar{F}_\alpha(\kappa)$, we have for every $k \in \mathbb{N}$ such that $t_k \geq T$ and all $t \in \mathbb{N}$ such that $t_k < t < t_{k+1}$ that

$$f(\mathbf{x}(t)) \leq \bar{F}_\alpha(\kappa) - \frac{\epsilon^2}{2\alpha^2BN^{3/2}} \sum_{\tau=t_k}^t \gamma(\tau).$$

Therefore, since $f^* \leq f(\mathbf{x}(t))$ we have that

$$\sum_{\tau=t_k}^{t_{k+1}} \gamma(\tau) \leq \frac{2\alpha BN^{3/2}}{\epsilon^2} (\bar{F}_\alpha(\kappa) - f^*). \quad (57)$$

We also have from Lemma 4 that $\bar{\mathcal{X}}_\alpha(\kappa)$ is bounded so there exists $A \in \mathbb{R}_+$ such that $\|\mathbf{x}\| \leq A$ for all $\mathbf{x} \in \bar{\mathcal{X}}_\alpha(\kappa)$. As a result, for all $t \geq T$ and $k(t) = \max\{k \in \mathbb{N} | t_k \leq t\}$ we get

$$\|\mathbf{x}(t)\| \leq \sum_{\tau=t_k(t)}^{t-1} \|\mathbf{x}(\tau+1) - \mathbf{x}(\tau)\| + \|\mathbf{x}(t_k)\| \leq \sum_{\tau=t_k(t)}^t \gamma(\tau) + A$$

where the first inequality comes by writing $\mathbf{x}(t)$ as a telescoping series starting at $\mathbf{x}(t_k)$ together with the triangle inequality and the second inequality comes from the relation

$$\|\mathbf{x}(t+1) - \mathbf{x}(t)\| = \left\| \lceil \bar{\mathbf{x}}(t) \rceil^+ - \mathbf{x}(t) \right\| \leq \gamma(t)$$

for all $t \in \mathbb{N}$, where $\bar{\mathbf{x}}(t) = \mathbf{x}(t) - (\gamma(t)/\sqrt{N}) \text{sign}(\nabla f(\mathbf{x}(t)))$. Thus, from Equation (57), we can conclude that the sequence $\mathbf{x}(t)$ is bounded.

Step 4: We will prove $\lim_{t \rightarrow \infty} \text{dist}(\mathbf{x}(t), \mathcal{X}^*) = 0$ by contradiction. Suppose that there exists $\epsilon > 0$ and a subsequence $\mathbf{x}(t_k)$ such that $\text{dist}(\mathbf{x}(t_k), \mathcal{X}^*) \geq \epsilon$ for all $k \in \mathbb{N}$. Then since $\mathbf{x}(t)$ is bounded, so we can without loss of generality restrict $\mathbf{x}(t_k)$ to a convergent subsequence to some point $\bar{\mathbf{x}}$, so $\lim_{k \rightarrow \infty} \mathbf{x}(t_k) = \bar{\mathbf{x}}$. Now since f is continuous and $\lim_{t \rightarrow \infty} f(\mathbf{x}(t)) = f^*$ we can conclude that $f(\bar{\mathbf{x}}) = f^*$ and $\bar{\mathbf{x}} \in \mathcal{X}^*$. Then $\lim_{t \rightarrow \infty} \bar{\mathbf{x}}(t_k) = \bar{\mathbf{x}} \in \mathcal{X}^*$ contradicts that $\text{dist}(\mathbf{x}(t_k), \mathcal{X}^*) \geq \epsilon$ for all $k \in \mathbb{N}$. ■

APPENDIX G PROOF OF LEMMA 5

Proof: By using that the gradients of f are L -Lipschitz continuous, we can apply the descent lemma (see for example [37, eq. (2.1.6)] or [26, Proposition A.24]). The descent lemma states that for all γ we have

$$f(\mathbf{x} - \gamma \mathbf{d}) \leq f(\mathbf{x}) - \langle \nabla f(\mathbf{x}), \mathbf{d} \rangle \gamma + \frac{L}{2} \|\mathbf{d}\|^2 \gamma^2, \quad (58)$$

$$= f(\mathbf{x}) + \left(\frac{L}{2} \gamma - \langle \nabla f(\mathbf{x}), \mathbf{d}(t) \rangle \right) \gamma, \quad (59)$$

$$\leq f(\mathbf{x}) + \left(\frac{L}{2} \gamma - \cos(\theta) \epsilon \right) \gamma \quad (60)$$

$$= f(\mathbf{x}) - \delta(\epsilon, \gamma, \theta) \quad (61)$$

where Equation (59) comes from that $\|\mathbf{d}\|=1$, Equation (60) comes from that $\text{ang}(\nabla f(\mathbf{x}), \mathbf{d}) \leq \theta$, $\|\nabla f(\mathbf{p})\| \geq \epsilon$, since $\mathbf{x} \notin \mathcal{X}(\epsilon)$, and $\langle \nabla f(\mathbf{x}), \mathbf{d} \rangle = \|\nabla f(\mathbf{x})\| \cos(\text{ang}(\mathbf{d}, \nabla f(\mathbf{x})))$. ■

APPENDIX H PROOF OF LEMMA 6

Proof: a) The result can be proved using **Steps 1, 2, and 3** used to prove Lemma 4, using $\|\nabla f(\cdot)\|$, $\mathcal{X}(\epsilon)$, $F(\epsilon)$, and Lemma 5 in place of $L_\alpha(\cdot)$, $\bar{\mathcal{X}}_\alpha(\alpha)$, $\bar{F}_\alpha(\epsilon)$ and Lemma 3, respectively.

b) For any $\mathbf{x} \in \mathcal{X}(\epsilon)$ [37, eq. (2.1.19) in Theorem 2.1.10]

$$f(\mathbf{x}) \leq f^* + \frac{1}{2\mu} \|\nabla f(\mathbf{x})\|^2 \leq f^* + \frac{\epsilon^2}{2\mu},$$

where we have used that $\nabla f(\mathbf{x}^*) = \mathbf{0}$ for all $\mathbf{x}^* \in \mathcal{X}^*$. ■

APPENDIX I PROOF OF THEOREM 8

Proof: The results can be proved using **Steps 1, 2, 3, and 4** used to prove Theorem 5. The main difference is that here **Step 1** is to prove that $\liminf_{t \rightarrow \infty} \|\nabla f(\mathbf{x}(t))\| = 0$, instead of $\liminf_{t \rightarrow \infty} L_\alpha(\mathbf{x}(t)) = 0$ as in the proof of Theorem 5. Moreover, here we use $\|\nabla f(\cdot)\|$, $\mathcal{X}(\epsilon)$, $F(\epsilon)$, Lemma 5 and Lemma 6-a) in place of $L_\alpha(\cdot)$, $\bar{\mathcal{X}}_\alpha(\alpha)$, $\bar{F}_\alpha(\epsilon)$, Lemma 3, and Lemma 4, respectively. ■

APPENDIX J ADDITIONAL LEMMAS

Lemma 7. Consider \mathcal{D} defined in Equation (9) in Example 2 of Section III-B. \mathcal{D} is a θ -cover with the θ in Equation (10).

Proof: We show that for θ defined in Equation (10) it holds for any $\mathbf{x} \in \mathcal{S}^{N-1}$ that there exists $\mathbf{d} \in \mathcal{D}_1$ such that Equation (3) holds.

First consider the case where $\mathbf{x}_j \geq \cos(\theta)$ for some component j . Then for $\mathbf{e}_j \in \mathcal{D}_1$ we get $\cos(\text{ang}(\mathbf{x}, \mathbf{e}_j)) = \langle \mathbf{x}, \mathbf{e}_j \rangle = \mathbf{x}_j \geq \cos(\theta)$. Therefore, we finalize the proof by showing that if $\mathbf{x} \in \mathcal{S}^{N-1}$ and $\mathbf{x}_i \leq \cos(\theta)$ for $i = 1, \dots, N$ then

$$\cos \left(\text{ang} \left(\mathbf{x}, -\frac{1}{\sqrt{N}} \mathbf{1} \right) \right) = \frac{-1}{\sqrt{N}} \sum_{i=1}^N \mathbf{x}_i \geq \cos(\theta),$$

Without loss of generality, let the components of \mathbf{x} be ordered so that $\mathbf{x}_i \geq 0$ if $i = 1, \dots, K$ and $\mathbf{x}_i < 0$ if $i = K+1, \dots, N$, where K is the number of positive components of \mathbf{x} . Then

$$\frac{-1}{\sqrt{N}} \sum_{i=1}^N \mathbf{x}_i \geq -\frac{1}{\sqrt{N}} \left(\sum_{i=1}^K \mathbf{x}_i - \sqrt{1 - \sum_{i=1}^K \mathbf{x}_i^2} \right) \quad (62)$$

$$\geq -\frac{1}{\sqrt{N}} \left(K \cos(\theta) - \sqrt{1 - K \cos(\theta)^2} \right), \quad (63)$$

where Equation (62) comes by using that $\sum_{i=1}^N \mathbf{x}_i^2 = 1$ and the inequality between the 1 and 2 norm, i.e.,

$$\sum_{i=K+1}^N |\mathbf{x}_i| \geq \sqrt{\sum_{i=K+1}^N \mathbf{x}_i^2} = \sqrt{1 - \sum_{i=1}^K \mathbf{x}_i^2},$$

and Equation (63) comes by noting that (62) is decreasing and that $\mathbf{x}_i \leq \cos(\theta)$ for all i . Now, by inserting our choice of $\cos(\theta)$ from Equation (10) in Equation (63) we get

$$\frac{-1}{\sqrt{N}} \sum_{i=1}^N \mathbf{x}_i \geq -\frac{K - \sqrt{N^2 + 2\sqrt{N}(N-1)} - K}{\sqrt{N} \sqrt{N^2 + 2\sqrt{N}(N-1)}} \quad (64)$$

$$\geq -\frac{N-1 - \sqrt{N^2 + 2\sqrt{N}(N-1)} - (N-1)}{\sqrt{N} \sqrt{N^2 + 2\sqrt{N}(N-1)}} \quad (65)$$

$$= \frac{\sqrt{N}}{\sqrt{N} \sqrt{N^2 + 2\sqrt{N}(N-1)}} = \cos(\theta) \quad (66)$$

where the Equation (65) comes from the fact that Equation (64) is decreasing in K and $K \leq N-1$, and the Equation (66) comes by using that $N^2 + 2\sqrt{N}(N-1) - (N-1) = ((N-1) + \sqrt{N})^2$. ■

Lemma 8. For all $\beta \in [0, 1]$, $z \in \mathbb{R}$ and $x, \alpha_1, \alpha_2 \in \mathbb{R}_+$ with $\alpha_1 \leq \alpha_2$ following holds

$$\beta|x - [x - z]^+| \leq |x - [x - \beta z]^+|, \quad (67)$$

$$|x - [x - \alpha_1 z]^+| \leq |x - [x - \alpha_2 z]^+|, \quad (68)$$

$$0 \leq z(x - [x - \alpha_1 z]^+). \quad (69)$$

Proof: We first prove (67). Direct inspection shows that

$$\phi_1(x, z, \beta) := \beta|x - [x - z]^+| = \begin{cases} \beta|z| & \text{if } x \geq z \\ \beta x & \text{if } x \leq z \end{cases} \quad (70)$$

$$\phi_2(x, z, \beta) := |x - [x - \beta z]^+| = \begin{cases} \beta|z| & \text{if } x \geq \beta z \\ x & \text{if } x \leq \beta z. \end{cases} \quad (71)$$

Therefore, for $z \in \mathbb{R}_+$ we have

$$\begin{aligned} \phi_1(x, z, \beta) &= \beta|z| = \phi_2(x, z, \beta) & \text{if } x \in [z, \infty), \\ \phi_1(x, z, \beta) &= \beta x \leq \beta|z| = \phi_2(x, z, \beta) & \text{if } x \in [\beta z, z], \\ \phi_1(x, z, \beta) &= \beta x \leq x = \phi_2(x, z, \beta) & \text{if } x \in [0, \beta z]. \end{aligned}$$

So $\phi_1(x, z, \beta) \leq \phi_2(x, z, \beta)$ for all $x, z \in \mathbb{R}_+$ and $\beta \in [0, 1]$ which yields (67).

Equation (68) follows directly from using (71). To prove (69), we use the fact that $\text{sign}(z)[x - \alpha_1 z]^+ \leq$

$\text{sign}(z)x$ or by rearranging $0 \leq \text{sign}(z)(x - [x - \alpha_1 z]^+)$. By multiplying $|z|$ on both sides we obtain (69). ■

Lemma 9. Suppose \mathcal{X}^* is bounded. Then:

(i) There exists $\kappa > 0$ such that $\mathcal{X}(\epsilon)$ defined in Equation (27) is bounded for all $\epsilon < \kappa$.

(ii) If $\|\nabla f(\mathbf{x})\| \leq B$ for all $\mathbf{x} \in \mathbb{R}_+^N$, then there exists $\kappa > 0$ such that $\bar{\mathcal{X}}_\alpha(\epsilon)$ in Equation (16) is bounded for all $\epsilon < \kappa$.

Proof: (i) Take any $\mathbf{x}^* \in \mathcal{X}^*$ and choose $R > 0$ so that $\mathcal{X}^* \subseteq \mathcal{B}^N(\mathbf{x}^*, R)$. Take $\kappa_1 > 0$ given by

$$\kappa_1 = \frac{1}{L} \min_{\mathbf{x} \in \mathcal{S}^{N-1}(\mathbf{x}^*, R)} \|\nabla f(\mathbf{x})\|^2. \quad (72)$$

Note that such a κ_1 exists since $\mathcal{S}^{N-1}(\mathbf{x}^*, R)$ is compact and $\kappa > 0$ since $\mathcal{S}^{N-1}(\mathbf{x}^*, R) \cap \mathcal{X}^*$ is empty. Moreover, using [37, (2.1.8) in Theorem 2.1.5], that $\nabla f(\mathbf{x}^*) = \mathbf{0}$, and (72), we have for all $\mathbf{x} \in \mathcal{S}^{N-1}(\mathbf{x}^*, R)$ that

$$\langle \nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle \geq (1/L) \|\nabla f(\mathbf{x})\|^2 \geq \kappa_1. \quad (73)$$

We now show that for all $\mathbf{x} \in \mathbb{R}^N \setminus \mathcal{B}^N(\mathbf{x}^*, R)$ we have $\|\nabla f(\mathbf{x})\| \geq \kappa$, where $\kappa = \kappa_1/R$. Take some $\mathbf{x} \in \mathbb{R}^N \setminus \mathcal{B}^N(\mathbf{x}^*, R)$ and let $\bar{\mathbf{x}}$ denote the unique point in the intersection of the line segment $[\mathbf{x}^*, \mathbf{x}]$ and $\mathcal{S}^{N-1}(\mathbf{x}^*, R)$, such a $\bar{\mathbf{x}}$ exists because $\mathbf{x} \in \mathbb{R}^N \setminus \mathcal{B}^N(\mathbf{x}^*, R)$ and $\mathbf{x}^* \in \mathcal{B}^N(\mathbf{x}^*, R)$. Consider now the function $G : [0, \infty) \rightarrow \mathbb{R}^N$ with

$$G(\tau) = \nabla f(\mathbf{x}^* + \tau(\bar{\mathbf{x}} - \mathbf{x}^*)). \quad (74)$$

Clearly, $G(0) = \nabla f(\mathbf{x}^*) = \mathbf{0}$, $G(1) = \nabla f(\bar{\mathbf{x}})$ and there exists $\hat{\tau} \geq 1$ such that $G(\hat{\tau}) = \nabla f(\mathbf{x})$. By using that gradients of convex functions are monotone, i.e., for all $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^N$ it holds that $\langle \nabla f(\mathbf{x}_1) - \nabla f(\mathbf{x}_2), \mathbf{x}_1 - \mathbf{x}_2 \rangle \geq 0$, we conclude that for $\tau_1, \tau_2 \in \mathbb{R}_+$ with $\tau_1 \geq \tau_2$ it holds that $\langle G(\tau_1) - G(\tau_2), (\tau_1 - \tau_2)(\bar{\mathbf{x}} - \mathbf{x}^*) \rangle \geq 0$. Rearranging this,

$$\langle G(\tau_1), (\bar{\mathbf{x}} - \mathbf{x}^*) \rangle \geq \langle G(\tau_2), (\bar{\mathbf{x}} - \mathbf{x}^*) \rangle, \quad \text{for } \tau_1 \geq \tau_2. \quad (75)$$

By combining (73) and (75) we get that

$$\langle G(\hat{\tau}), (\bar{\mathbf{x}} - \mathbf{x}^*) \rangle \geq \langle G(1), (\bar{\mathbf{x}} - \mathbf{x}^*) \rangle \geq \kappa_1. \quad (76)$$

Hence, by the Cauchy-Schwarz inequality we have $\|\nabla f(\mathbf{x})\|R = \|G(\hat{\tau})\|R \geq \kappa_1$ and by rearranging we get $\|\nabla f(\mathbf{x})\| \geq \kappa_1/R = \kappa$. Since $\|\nabla f(\mathbf{x})\| \geq \kappa$ holds for all $\mathbf{x} \in \mathbb{R}^N \setminus \mathcal{B}^N(\mathbf{x}^*, R)$ we can conclude that $\mathcal{X}(\epsilon)$ is bounded for $\epsilon < \kappa$.

(ii) We prove the result by contradiction and suppose $\bar{\mathcal{X}}_\alpha(\kappa)$ is unbounded for all $\kappa > 0$. Then there exists a sequence $\mathbf{x}^k \in \mathbb{R}_+^N$ such that $\lim_{k \rightarrow \infty} \|\mathbf{x}^k\| = \infty$ and $\lim_{k \rightarrow \infty} L_\alpha(\mathbf{x}^k) = 0$. The proof follows from the following steps:

Step 1: We will show that there exists $\bar{\kappa} > 0$ and R such that $\|\nabla f(\mathbf{x})\| \geq \bar{\kappa}$ holds for all $\mathbf{x} \in \mathbb{R}_+^N$ and $\|\mathbf{x} - \mathbf{x}^*\| \geq R$. If there exists $\mathbf{x}^* \in \mathcal{X}^*$ such that $\|\nabla f(\mathbf{x}^*)\| = 0$, then the result follows from part (i). Therefore, without loss of generality, suppose we can take $\mathbf{x}^* \in \mathcal{X}^*$ with $\|\nabla f(\mathbf{x}^*)\| > 0$. Then the set $\mathcal{J} := \{j = 1, \dots, N, |\nabla_j f(\mathbf{x}^*)| \neq 0\}$ is nonempty. We also have, using the KKT conditions [32, Section 5.9.2], that $\mathbf{x} \in \mathcal{X}^*$ if and only if the following three conditions hold (A) $\mathbf{x} \in \mathbb{R}_+^N$, (B) $\nabla f(\mathbf{x}) \geq 0$ for $i = 1 \dots N$, and (C) $\nabla_i f(\mathbf{x})\mathbf{x}_i = 0$ for $i = 1 \dots N$.⁵

⁵The Lagrangian multiplier associated with $\mathbf{x} \in \mathcal{X}^*$ is $\lambda = \nabla f(\mathbf{x})$.

We first show that $\langle \nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle > 0$ for all $\mathbf{x} \in \mathbb{R}_+^N \setminus \mathcal{X}^*$. Consider first the case when $\mathbf{x} \in \mathbb{R}_+^N \setminus \mathcal{X}^*$ and $\mathbf{x}_j > 0$ for some $j \in \mathcal{J}$. Then we have

$$\langle \nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle \geq \langle \nabla f(\mathbf{x}^*), \mathbf{x} - \mathbf{x}^* \rangle \geq \sum_{i=1}^N \nabla_i f(\mathbf{x}^*) \mathbf{x}_i > 0,$$

where the first inequality comes by the monotonicity of ∇f , the second inequality comes by the optimality condition (C), and the final inequality comes by the optimality condition (B), the fact that $\nabla_i f(\mathbf{x}^*) > 0$ for all $j \in \mathcal{J}$, and that $\mathbf{x}_j > 0$ for some $j \in \mathcal{J}$. Consider next the case when $\mathbf{x} \in \mathbb{R}_+^N \setminus \mathcal{X}^*$ and $\mathbf{x}_j^* = 0$ for all $j \in \mathcal{J}$. Then $\nabla_i f(\mathbf{x}) \neq \nabla_i f(\mathbf{x}^*)$ for some $i \notin \mathcal{J}$ because otherwise the optimality conditions (A), (B), and (C) hold for \mathbf{x} so $\mathbf{x} \in \mathcal{X}^*$. In particular, $\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{x}^*)\| > 0$. Therefore, we have [37, eq. (2.1.8)]

$$\langle \nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle \geq \langle \nabla f(\mathbf{x}^*), \mathbf{x} - \mathbf{x}^* \rangle + \frac{1}{L} \|\nabla f(\mathbf{x}) - \nabla f(\mathbf{x}^*)\|^2 > 0,$$

where the final inequality comes by that $\langle \nabla f(\mathbf{x}^*), \mathbf{x} - \mathbf{x}^* \rangle \geq 0$ for all $\mathbf{x} \in \mathbb{R}_+^N$ and that $\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{x}^*)\| > 0$.

Now take $R > 0$ such that $\mathcal{X}^* \subseteq \mathcal{B}^N(\mathbf{x}^*, R)$. Then since $\mathcal{S}^{N-1}(\mathbf{x}^*, R) \cap \mathbb{R}_+^N$ is compact, there exists $\kappa_1 = \min_{\mathbf{x} \in \mathcal{S}^{N-1}(\mathbf{x}^*, R) \cap \mathbb{R}_+^N} \langle \nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle > 0$. We can now follow same arguments as in the proof of part (i) to show that $\|\nabla f(\mathbf{x})\| \geq \bar{\kappa}$ where $\bar{\kappa} = \kappa_1/R$.

Step 2: We will show that the following inequality holds for all $\mathbf{x} \in \mathcal{R}_+^N \setminus \mathcal{B}^N(\mathbf{x}^*, R)$,

$$\cos(\text{ang}(\nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}^*)) \geq \frac{\bar{\kappa}}{B}, \quad (77)$$

where R and $\bar{\kappa}$ are defined as in **Step 1**. Take some $\mathbf{x} \in \mathbb{R}_+^N \setminus \mathcal{B}^N(\mathbf{x}^*, R)$. Similarly as in part (i), let $\bar{\mathbf{x}}$ denote the unique point in the intersection of the line segment $[\mathbf{x}^*, \mathbf{x}]$ and $\mathcal{S}^{N-1}(\mathbf{x}^*, R)$. Moreover, take $\hat{\tau} > 1$ such that $\mathbf{x} = \mathbf{x}^* + \hat{\tau}(\bar{\mathbf{x}} - \mathbf{x}^*)$ and define $G : [0, \infty) \rightarrow \mathbb{R}^N$ as in Equation (74). Then by rearranging (76) and multiplying both sides with $1/R$

$$\cos(\text{ang}(\nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}^*)) \geq \frac{\kappa_1}{R} \frac{1}{\|G(\hat{\tau})\|} \geq \frac{\bar{\kappa}}{B},$$

where κ_1 and $\bar{\kappa}$ are defined as in part (i) and **Step 1**.

Step 3: We will show that the subsequence \mathbf{x}^k can be restricted so that (a) $\lim_{k \rightarrow \infty} \nabla f(\mathbf{x}^k) = \mathbf{f}$ for some $\mathbf{f} \in \mathbb{R}^N$ and (b) for each component $i = 1, \dots, N$ either $\lim_{k \rightarrow \infty} \mathbf{x}_i^k = 0$ or $\mathbf{x}_i \geq W$, for some $W > 0$. We first show (a). Since ∇f is bounded by B , the sequence $\nabla f(\mathbf{x}^k)$ is bounded. Therefore, we can restrict the sequence \mathbf{x}^k so that $\nabla f(\mathbf{x}^k)$ is a convergent subsequence with $\lim_{k \rightarrow \infty} \nabla f(\mathbf{x}^k) = \mathbf{f}$. To show (b), for each component $i = 1, \dots, N$ we restrict the sequence \mathbf{x}^k so that $\mathbf{x}_i^k \geq W_i$ if \mathbf{x}_i^k does not converge to 0 and taking $W = \max W_i$. Set

$$\mathcal{I} = \{i = 1, \dots, N \mid \lim_{k \rightarrow \infty} \mathbf{x}_i^k = 0\}. \quad (78)$$

Step 4: We will prove that $\mathbf{f}_i = 0$ for $i \notin \mathcal{I}$. We prove the result by contradiction. Without loss of generality, suppose $\mathbf{f}_i > 0$ for some $i \notin \mathcal{I}$, the case when $\mathbf{f}_i < 0$ follows same arguments. Then there exists $K \in \mathbb{N}$ such that $\nabla_i f(\mathbf{x}^k) \geq \eta_0 := \mathbf{f}_i/2 > 0$ for all $k \geq K$. This, together with that $\mathbf{x}_i^k \geq W$ implies that $|\mathbf{x}_i^k - \lceil \mathbf{x}_i^k - \nabla_i f(\mathbf{x}^k) \rceil_+| \geq \min\{\eta_0, W\} =: \eta > 0$.

Therefore, $L_\alpha(\mathbf{x}^k) \geq \eta$ for all $k \geq K$, contradicting that $\lim_{t \rightarrow \infty} L_\alpha(\mathbf{x}^k) = 0$.

Step 5: We will prove contradiction when \mathcal{I} is empty. From **Step 4**, we have $\mathbf{f} = \mathbf{0}$. However, we also have from **Step 1** that $\|\nabla f(\mathbf{x})\| \geq \bar{\kappa}$ for all $\mathbf{x} \in \mathbb{R}_+^N \setminus \mathcal{B}^N(\mathbf{x}^*, R)$. Since $\mathcal{B}^N(\mathbf{x}^*, R)$ is bounded and $\lim_{t \rightarrow \infty} \|\mathbf{x}(t)\| = \infty$, we have that $\|\mathbf{f}\| \geq \bar{\kappa} > 0$, which contradicts that $\mathbf{f} = \mathbf{0}$.

Step 6: We will prove contradiction when \mathcal{I} is nonempty. Consider the sequence $\mathbf{z}^k = (\mathbf{x}^k - \mathbf{x}^*) / (\|\mathbf{x}^k - \mathbf{x}^*\|)$. Since \mathbf{z}^k is bounded, we can restrict the subsequence \mathbf{x}^k so that \mathbf{z}^k has a convergent subsequence, with the limit \mathbf{z} . We have $\mathbf{z}_i = 0$ for $i \in \mathcal{I}$ and $\|\mathbf{z}\| = 1$, since $\|\cdot\|$ is continuous function and \mathbf{z}^k convergent sequence. Therefore, as both $\nabla f(\mathbf{x}^k)$ and \mathbf{z}^k are convergent sequences and the inner product $\langle \cdot, \cdot \rangle$ is continuous function, the sequence $\langle \nabla f(\mathbf{x}^k), \mathbf{z}^k \rangle$ is convergent and has the limit $\langle \mathbf{f}, \mathbf{z} \rangle = 0$. However, for all $\mathbf{x}^k \in \mathcal{R}_+^N \setminus \mathcal{B}^N(\mathbf{x}^*, R)$

$$\langle \nabla f(\mathbf{x}^k), \mathbf{z}^k \rangle = \|\nabla f(\mathbf{x}^k)\| \cos(\text{ang}(\nabla f(\mathbf{x}^k), \mathbf{z}^k)) \geq \frac{\bar{\kappa}^2}{B},$$

where the inequality comes by Equation (77) in **Step 2**. ■

REFERENCES

- [1] S. Galli, A. Scaglione, and Z. Wang, "For the grid and through the grid: The role of power line communications in the smart grid," *Proceedings of the IEEE*, vol. 99, no. 6, pp. 998–1027, 2011.
- [2] C. Fischione, "Fast-lipschitz optimization with wireless sensor networks applications," *IEEE Transactions on Automatic Control*, vol. 56, no. 10, pp. 2319–2331, Oct 2011.
- [3] J. Partan, J. Kurose, and B. N. Levine, "A survey of practical issues in underwater networks," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 11, no. 4, pp. 23–33, 2007.
- [4] D. P. Bertsekas, R. G. Gallager, and P. Humblet, *Data networks*. Prentice-Hall International New Jersey, 1992, vol. 2.
- [5] G. Durisi, T. Koch, and P. Popovski, "Toward massive, ultrareliable, and low-latency wireless communication with short packets," *Proceedings of the IEEE*, vol. 104, no. 9, pp. 1711–1726, Sept 2016.
- [6] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate control for communication networks: Shadow prices, proportional fairness and stability," *The Journal of the Operational Research Society*, vol. 49, no. 3, pp. 237–252, 1998.
- [7] S. Low and D. Lapsley, "Optimization flow control. I. basic algorithm and convergence," *IEEE/ACM Transactions on Networking*, vol. 7, no. 6, pp. 861–874, Dec 1999.
- [8] D. Palomar and M. Chiang, "A tutorial on decomposition methods for network utility maximization," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 8, pp. 1439–1451, Aug 2006.
- [9] —, "Alternative distributed algorithms for network utility maximization: Framework and applications," *IEEE Transactions on Automatic Control*, vol. 52, no. 12, pp. 2254–2269, Dec 2007.
- [10] N. Li, L. Chen, and S. H. Low, "Optimal demand response based on utility maximization in power networks," in *Power and Energy Society General Meeting, 2011 IEEE*, 2011, pp. 1–8.
- [11] L. Chen, N. Li, S. H. Low, and J. C. Doyle, "Two market models for demand response in power networks," *IEEE SmartGridComm*, vol. 10, pp. 397–402, 2010.
- [12] R. Madan and S. Lall, "Distributed algorithms for maximum lifetime routing in wireless sensor networks," *IEEE Transactions on Wireless Communications*, vol. 5, no. 8, pp. 2185–2193, Aug 2006.
- [13] C. Zhao, U. Topcu, N. Li, and S. Low, "Design and stability of load-side primary frequency control in power systems," *IEEE Transactions on Automatic Control*, vol. 59, no. 5, pp. 1177–1189, May 2014.
- [14] E. Dall'Anese, S. V. Dhople, and G. B. Giannakis, "Photovoltaic inverter controllers seeking ac optimal power flow solutions," *IEEE Transactions on Power Systems*, vol. 31, no. 4, pp. 2809–2823, July 2016.
- [15] J. N. Tsitsiklis and Z. Q. Luo, "Communication complexity of convex optimization," *Journal of Complexity*, vol. 3, pp. 231–243, 1987.
- [16] M. Rabbat and R. Nowak, "Quantized incremental algorithms for distributed optimization," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 4, pp. 798–808, April 2005.

- [17] A. Nedic, A. Olshevsky, A. Ozdaglar, and J. Tsitsiklis, "Distributed subgradient methods and quantization effects," in *IEEE Conference on Decision and Control (CDC)*, Dec 2008, pp. 4177–4184.
- [18] Y. Pu, M. N. Zeilinger, and C. N. Jones, "Quantization design for unconstrained distributed optimization," in *American Control Conference (ACC)*. IEEE, 2015, pp. 1229–1234.
- [19] P. Yi and Y. Hong, "Quantized subgradient algorithm and data-rate analysis for distributed optimization," *IEEE Transactions on Control of Network Systems*, vol. 1, no. 4, pp. 380–392, 2014.
- [20] J. Herdtner and E. Chong, "Analysis of a class of distributed asynchronous power control algorithms for cellular wireless systems," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 3, pp. 436–446, March 2000.
- [21] A. Nedic and D. P. Bertsekas, "Incremental subgradient methods for nondifferentiable optimization," *SIAM Journal on Optimization*, vol. 12, no. 1, pp. 109–138, 2001.
- [22] A. Nedic and A. Ozdaglar, "Distributed subgradient methods for multi-agent optimization," *IEEE Transactions on Automatic Control*, vol. 54, no. 1, pp. 48–61, Jan 2009.
- [23] S. Magnússon, K. Heal, C. Enyioha, N. Li, C. Fischione, and V. Tarokh, "Convergence of limited communications gradient methods," in *American Control Conference (ACC)*. IEEE, 2016.
- [24] C. Enyioha, S. Magnússon, K. Heal, N. Li, C. Fischione, and V. Tarokh, "Robustness analysis for an online decentralized descent power allocation algorithm," in *IEEE Information Theory and Applications Workshop (ITA)*, February 2016.
- [25] S. Magnússon, C. Enyioha, K. Heal, N. Li, C. Fischione, and V. Tarokh, "Distributed resource allocation using one-way communication with applications to power networks," in *IEEE Annual Conference on Information Sciences and Systems (CISS)*, March 2015, pp. 1–6.
- [26] D. P. Bertsekas, *Nonlinear Programming: 2nd Edition*. Athena Scientific, 1999.
- [27] N. Z. Shor, *Minimization Methods for Non-Differentiable Functions*. Springer, 1985.
- [28] 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, Sept 2013.
- [29] —, "A distributed Newton method for network utility maximization II: Convergence," *IEEE Transactions on Automatic Control*, vol. 58, no. 9, pp. 2176–2188, Sept 2013.
- [30] A. Beck, A. Nedic, A. Ozdaglar, and M. Teboulle, "An $O(1/k)$ gradient method for network resource allocation problems," *IEEE Transactions on Control of Network Systems*, vol. 1, no. 1, pp. 64–73, March 2014.
- [31] A. Nedic and A. Ozdaglar, "Approximate primal solutions and rate analysis for dual subgradient methods," *SIAM Journal on Optimization*, vol. 19, no. 4, pp. 1757–1780, 2009.
- [32] S. Boyd and L. Vandenberghe, *Convex Optimization*. New York, NY, USA: Cambridge University Press, 2004.
- [33] E. Ghadimi, I. Shames, and M. Johansson, "Multi-step gradient methods for networked optimization," *IEEE Transactions on Signal Processing*, vol. 61, no. 21, pp. 5417–5429, Nov 2013.
- [34] M. Zargham, A. Ribeiro, A. Ozdaglar, and A. Jadbabaie, "Accelerated dual descent for network flow optimization," *IEEE Transactions on Automatic Control*, vol. 59, no. 4, pp. 905–920, April 2014.
- [35] A. D. Wyner, "Random packings and coverings of the unit n -sphere," *The Bell System Technical Journal*, vol. 46, no. 9, pp. 2111–2118, Nov 1967.
- [36] P. Solé, "The covering radius of spherical designs," *European journal of combinatorics*, vol. 12, no. 5, pp. 423–431, 1991.
- [37] Y. Nesterov, *Introductory Lectures on Convex Optimization*. Springer, 2004.

Sindri Magnússon received the B.Sc. degree in Mathematics from University of Iceland, Reykjavík Iceland, in 2011, and the Masters degree in Mathematics from KTH Royal Institute of Technology, Stockholm Sweden, in 2013. He is currently pursuing the PhD in Electrical Engineering at KTH Royal Institute of Technology, Stockholm, Sweden. He has held a six months visiting research position at Harvard University, Cambridge, MA (2015) sponsored by the Engblom Foundation. His research interests include distributed optimization, both theory and applications.

Chinwendu Enyioha received the B.Sc. degree in Mathematics from Gardner-Webb University (GWU), Boiling Springs, NC, and the PhD degree in Electrical and Systems Engineering from the University of Pennsylvania, Philadelphia, PA, in 2014. He is currently a Postdoctoral Research Fellow in the School of Engineering and Applied Sciences at Harvard. Prior to arriving Harvard, he was a Postdoctoral Researcher in the GRASP Lab at the University of Pennsylvania. Dr. Enyioha is a Fellow of the Ford Foundation, was named a William Fontaine Scholar at the University of Pennsylvania and has received the Mathematical Association of America Patterson award. His research lies in the area of design, optimization and control of distributed networked systems, with applications to power systems and Cyber-Physical networks.

Na Li received her B.S. degree in mathematics and applied mathematics from Zhejiang University in China and her PhD degree in Control and Dynamical systems from the California Institute of Technology in 2013. She is an Assistant Professor in the School of Engineering and Applied Sciences in Harvard University. She was a postdoctoral associate of the Laboratory for Information and Decision Systems at Massachusetts Institute of Technology. She was a Best Student Paper Award finalist in the 2011 IEEE Conference on Decision and Control. She received NSF CAREER Award in 2016 and Air Force Young Investigator Award in 2017. Her research lies in the design, analysis, optimization and control of distributed network systems, with particular applications to power networks and systems biology/physiology.

Carlo Fischione is currently a tenured Associate Professor at KTH Royal Institute of Technology, Electrical Engineering, Stockholm, Sweden. He received the Ph.D. degree in Electrical and Information Engineering (3/3 years) in May 2005 from University of LAquila, Italy, and the Laurea degree in Electronic Engineering (Laurea, Summa cum Laude, 5/5 years) in April 2001 from the same University. He has held research positions at Massachusetts Institute of Technology, Cambridge, MA (2015, Visiting Professor); Harvard University, Cambridge, MA (2015, Associate); University of California at Berkeley, CA (2004-2005, Visiting Scholar, and 2007-2008, Research Associate). His research interests include optimization with applications to wireless sensor networks, networked control systems, wireless networks, security and privacy. He received or co-received a number of awards, including the best paper award from the IEEE Transactions on Industrial Informatics (2007). He is Member of IEEE (the Institute of Electrical and Electronic Engineers), and Ordinary Member of DASP (the academy of history Deputazione Abruzzese di Storia Patria).

Vahid Tarokh received the PhD in electrical engineering from the University of Waterloo, Ontario, Canada. He then worked at AT&T Labs-Research and AT&T wireless services until August 2000, where he was the head of the Department of Wireless Communications and Signal Processing.

In September 2000, he joined the Department of Electrical Engineering and Computer Sciences (EECS) at MIT as an associate professor. In June 2002, he joined Harvard University, where he is a Professor of Applied Mathematics. He has received a Guggenheim Fellowship in Applied Mathematics and holds 4 honorary degrees.