A Parallel Decomposition Method for Nonconvex Stochastic Multi-Agent Optimization Problems

Yang Yang, Gesualdo Scutari, Daniel P. Palomar, and Marius Pesavento

***Abstract*—This paper considers the problem of minimizing the expected value of a (possibly *nonconvex*) cost function parameterized by a random (vector) variable, when the expectation cannot be computed accurately (e.g., because the statistics of the random variables are unknown and/or the computational complexity is prohibitive). Classical stochastic gradient methods for solving this problem may suffer from slow convergence. In this paper, we propose a stochastic *parallel Successive Convex Approximation*-based (best-response) algorithm for general *nonconvex* stochastic sum-utility optimization problems, which arise naturally in the design of multi-agent networks. The proposed novel decomposition approach enables all users to update their optimization variables *in parallel* by solving a sequence of strongly convex subproblems, one for each user. Almost sure convergence to stationary points is proved. We then customize the algorithmic framework to solve the stochastic sum rate maximization problem over Single-Input-Single-Output (SISO) frequency-selective interference channels, multiple-input-multiple-output (MIMO) interference channels, and MIMO multiple-access channels. Numerical results corroborate that the proposed algorithms can converge faster than state-of-the-art stochastic gradient schemes while achieving the same (or better) sum-rates.**

***Index Terms*—Distributed algorithms, Multi-agent systems, stochastic optimization, successive convex approximation.**

I. Introduction

Wireless networks are composed of multiple users that may have different objectives and generate interference when no orthogonal multiplexing scheme is imposed to regulate the transmissions; examples are peer-to-peer networks, cognitive radio systems, and ad-hoc networks. A common design of such multi-user systems is to optimize the (weighted) sum of users' objective functions. This formulation however requires the knowledge of the system parameters, such as the users' channel states. In practice this information is either difficult to acquire (e.g., when the parameters are rapidly changing) or imperfect due to estimation and signaling errors. In such scenarios, it is convenient to focus on the optimization of the long-term performance of the network, measured in terms of the expected value of the sum-utility function, parametrized by the random system parameters. In this paper, we consider the frequently encountered difficult case that (the expected value of) the social function is *nonconvex* and the expectation cannot be computed (either numerically or in closed-form). Such a system design naturally falls into the class of stochastic optimization problems [2, 3].

Gradient methods for *unconstrained* stochastic *nonconvex* optimization problems have been studied in [4, 5, 6], where almost sure convergence to stationary points has been established, under some technical conditions; see, e.g., [5]. The extension of these methods to *constrained* optimization problems is not straightforward; in fact, the descent-based convergence analysis developed for unconstrained gradient methods no longer applies to their projected counterpart (due to the presence of the projection operator). Convergence of stochastic gradient *projection* methods has been proved only for *convex* objective functions [4, 7, 8].

To cope with nonconvexity, *gradient averaging* seems to be an essential step to resemble convergence; indeed, stochastic *conditional* gradient methods for *nonconvex* constrained problems hinge on this idea [9, 10, 11, 12]: at each iteration the new update of the variables is based on the average of the current and past gradient samples. Under some technical conditions, the average sample gradient eventually resembles the nominal (but unavailable) gradient of the (stochastic) objective function [9, 13]; convergence analysis can then be built on results from deterministic nonlinear programming.

Numerical experiments for large classes of problems show that plain gradient-like methods usually converge slowly. Some acceleration techniques have been proposed in the literature [8, 14], but only for *strongly convex* objective functions. Here we are interested in nonconvex (constrained) stochastic problems. Moreover, (proximal, accelerated) stochastic gradient-based schemes use only the first order information of the objective function (or its realizations); recently it was shown [15, 16, 17] that for deterministic nonconvex optimization problems exploiting the structure of the function by replacing its linearization with a "better" approximant can enhance empirical convergence speed. In this paper we aim at bringing this idea into the context of stochastic optimization problems.

Our main contribution is to develop a new broad algorithmic framework for the computation of stationary solutions of a wide class of *nonconvex* stochastic optimization problems, encompassing many multi-agent system designs of practical interest. The essential idea underlying the proposed approach is to decompose the original nonconvex *stochastic* problem into a sequence of (simpler) deterministic subproblems. In this case, the objective function is replaced by suitable chosen *sample convex* approximations; the subproblems can be then solved in a *parallel* and *distributed* fashion across the users. Other key features of the proposed framework are: i) it is very flexible in the choice of the approximant of the nonconvex objective function, which need not necessarily be its first order approximation, as in classical (proximal)

Y. Yang and M. Pesavento are with Communication Systems Group, Darmstadt University of Technology, Darmstadt, Germany. G. Scutari is with the Department of Industrial Engineering and Cyber Center (Discovery Park), Purdue University, West Lafayette, IN, USA. D. P. Palomar is with the Department of Electronic and Computer Engineering, Hong Kong University of Science and Technology, Hong Kong. Emails: <yang,pesavento>@nt.tu-darmstadt.de; gscutari@purdue.edu; palomar@ust.hk.

The work of Yang and Pesavento is supported by the Seventh Framework Programme for Research of the European Commission under grant number ADEL-619647. The work of Scutari is supported by the USA National Science Foundation under Grants CIF 1564044 and CAREER Award No. 1555850, and ONR Grant No. N00014-16-1-2244. The work of Palomar is supported by the Hong Kong RGC 16207814 research grant.

Part of this work has been presented at the 14th IEEE Workshop on Signal Processing Advances in Wireless Communications (SPAWC), Jun. 2013 [1].

gradient schemes; ii) it encompasses a gamut of algorithms that differ in cost per iteration, communication overhead, and convergence speed, *while all converging under the same conditions*; and iii) it can be successfully used to *robustify* the algorithms proposed in [15] for *deterministic* optimization problems, when only inexact estimates of the system parameters are available, which makes them applicable to more realistic scenarios. As illustrative examples, we customize the proposed algorithms to some resource allocation problems in wireless communications, namely: the sum-rate maximization problems over MIMO Interference Channels (ICs) and Multiple Access Channels (MACs). The resulting algorithms outperform existing (gradient-based) methods both theoretically and numerically.

The proposed decomposition technique hinges on successive convex approximation (SCA) methods, and it is a nontrivial generalization to *stochastic* (nonconvex) optimization problems of the solution method proposed in [15] for *deterministic* optimization problems. We remark that [15] is not applicable to stochastic problems wherein the expected value of the objective function cannot be computed analytically, which is the case for the classes of problems studied in this paper. In fact, as it shown also numerically (cf. Sec. IV-D), when applied to sample functions of stochastic optimization problems, the scheme in [15] may either not converge or converge to limit points that are not even stationary solutions of the stochastic optimization problem. Finally, since the scheme proposed in this paper is substantially different from that in [15], a further contribution of this work is establishing a new type of convergence analysis (see Appendix A) that conciliates random and SCA strategies, which is also of interest per se and could bring to further developments.

An SCA framework for stochastic optimization problems has also been proposed in a recent, independent submission [18]; however the proposed method differs from [18] in many features. Firstly, the iterative algorithm proposed in [18] is based on a majorization minimization approach, requiring thus that the convex approximation be a tight *global upper bound* of the (sample) objective function. This requirement, which is fundamental for the convergence of the schemes in [18], is no longer needed in the proposed algorithm. This represents a turning point in the design of distributed stochastic SCA-based methods, enlarging substantially the class of (large scale) stochastic nonconvex problems solvable using the proposed framework. Secondly, even when the aforementioned upper bound constraint can be met, it is not always guaranteed that the resulting convex (sample) subproblems are decomposable across the users, implying that a centralized implementation might be required in [18]; the proposed schemes instead naturally lead to a parallel and distributed implementation. Thirdly, the proposed methods converge under weaker conditions than those in [18]. Fourthly, numerical results on several test problems show that the proposed scheme outperforms [18], see Sec. IV.

Finally, within the classes of approximation methods for stochastic optimization problems, it is worth mentioning the Sample Average Approach (SAA) [18, 19, 20, 21]: the "true" (stochastic) objective function is approximated by an ensemble average. Then the resulting deterministic optimization problem is solved by an appropriate numerical procedure. When the original objective function is nonconvex, the resulting SSA problem is nonconvex too, which makes the computation of its global optimal solution at each step a difficult, if not impossible, task. Therefore SSA-based methods are generally used to solve stochastic *convex* optimization problems.

The rest of the paper is organized as follows. Sec. II formulates the problem along with some motivating applications. The novel stochastic decomposition framework is introduced in Sec. III; customizations of the framework to some representative applications are discussed in Sec. IV. Finally, Sec. VI draws some conclusions.

II. Problem Formulation

We consider the design of a multi-agent system composed of I users; each user i has his own strategy vector $\mathbf{x}_i$ to optimize, which belongs to the convex feasible set $\mathcal{X}_i \subseteq \mathbb{C}^{n_i}$. The variables of the remaining users are denoted by $\mathbf{x}_{-i} \triangleq (\mathbf{x}_j)_{j=1,j\neq i}^{I}$, and the joint strategy set of all users is the Cartesian product set $\mathcal{X} = \mathcal{X}_1 \times \ldots \times \mathcal{X}_I$.

The stochastic social optimization problem is formulated as:

$$\begin{array}{ll} \underset{\mathbf{x}\triangleq(\mathbf{x}_i)_{i=1}^I}{\text{minimize}} & U(\mathbf{x}) \triangleq \mathbb{E}\left[\sum_{j\in\mathcal{I}_f} f_j(\mathbf{x},\boldsymbol{\xi})\right] \\ \text{subject to} & \mathbf{x}_i \in \mathcal{X}_i, \quad i=1,\ldots,I, \end{array} \tag{1}$$

where $\mathcal{I}_f \triangleq \{1,\ldots,I_f\}$, with I_f being the number of functions; each cost function $f_j(\mathbf{x},\boldsymbol{\xi}) : \mathcal{X}\times\mathcal{D} \to \mathbb{R}$ depends on the joint strategy vector $\mathbf{x}$ and a random vector $\boldsymbol{\xi}$, defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, with $\Omega \subseteq \mathbb{C}^m$ being the sample space, $\mathcal{F}$ being the σ-algebra generated by subsets of Ω, and $\mathbb{P}$ being a probability measure defined on $\mathcal{F}$, which need not be known. Note that the optimization variables can be complex-valued; in such a case, all the gradients of real-valued functions are intended to be conjugate gradients [22, 23].

Assumptions: We make the following assumptions:

(a) Each $\mathcal{X}_i$ is compact and convex;
(b) Each $f_j(\mathbf{x},\boldsymbol{\xi})$ is continuously differentiable on $\mathcal{X}$, for any given $\boldsymbol{\xi}$, and the gradient is Lipschitz continuous with constant $L_{\nabla f_j(\boldsymbol{\xi})}$. Furthermore, the gradient of $U(\mathbf{x})$ is Lipschitz continuous with constant $L_{\nabla U} < +\infty$.

These assumptions are quite standard and are satisfied for a large class of problems. Note that the existence of a solution to (1) is guaranteed by Assumption (a). Since $U(\mathbf{x})$ is not assumed to be jointly convex in $\mathbf{x}$, (1) is generally nonconvex. Some instances of (1) satisfying the above assumptions are briefly listed next.

Example #1: Consider the maximization of the ergodic sum-rate over frequency-selective ICs:

$$\begin{array}{ll} \underset{\mathbf{p}_1,\ldots,\mathbf{p}_I}{\text{maximize}} & \mathbb{E}\left[\displaystyle\sum_{n=1}^{N}\sum_{i=1}^{I}\log\left(1+\frac{|h_{ii,n}|^2 p_{i,n}}{\sigma_{i,n}^2+\sum_{j\neq i}|h_{ij,n}|^2 p_{j,n}}\right)\right] \\ \text{subject to} & \mathbf{p}_i \in \mathcal{P}_i \triangleq \{\mathbf{p}_i : \mathbf{p}_i \geq \mathbf{0}, \mathbf{1}^T\mathbf{p}_i \leq P_i\}, \ \forall i, \end{array} \tag{2}$$

where $\mathbf{p}_i \triangleq \{p_{i,n}\}_{n=1}^N$ with $p_{i,n}$ being the transmit power of user i on subchannel (subcarrier) n, N is the number of parallel subchannels, P_i is the total power budget, $h_{ij,n}$ is the channel coefficient from transmitter j to receiver i on subchannel n, and $\sigma_{i,n}^2$ is the variance of the thermal noise over subchannel n at the receiver i. The expectation is taken over channel coefficients $(h_{ij,n})_{i,j,n}$.

Example #2: The maximization of the ergodic sum-rate over MIMO ICs also falls into the class of problems (1):

$$\begin{array}{ll} \underset{\mathbf{Q}_1,\ldots,\mathbf{Q}_I}{\text{maximize}} & \mathbb{E}\left[\sum_{i=1}^{I} \log\det\left(\mathbf{I}+\mathbf{H}_{ii}\mathbf{Q}_i\mathbf{H}_{ii}^H\mathbf{R}_i(\mathbf{Q}_{-i},\mathbf{H})^{-1}\right)\right] \\ \text{subject to} & \mathbf{Q}_i \in \mathcal{Q}_i \triangleq \{\mathbf{Q}_i : \mathbf{Q}_i \succeq \mathbf{0}, \text{Tr}(\mathbf{Q}_i) \leq P_i\}, \forall i, \end{array} \tag{3}$$

where $\mathbf{R}_i(\mathbf{Q}_{-i},\mathbf{H}) \triangleq \mathbf{R}_{\text{N}_i} + \sum_{j\neq i}\mathbf{H}_{ij}\mathbf{Q}_j\mathbf{H}_{ij}^H$ is the covariance matrix of the thermal noise $\mathbf{R}_{\text{N}_i}$ (assumed to be full rank) plus the multi-user interference, P_i is the total power budget, and the expectation in (3) is taken over the channels $\mathbf{H} \triangleq (\mathbf{H}_{ij})_{i,j=1}^I$.

Example #3: Another application of interest is the maximization of the ergodic sum-rate over MIMO MACs:

$$\begin{array}{ll} \underset{\mathbf{Q}_1,\ldots,\mathbf{Q}_I}{\text{maximize}} & \mathbb{E}\left[\log\det\left(\mathbf{R}_\text{N}+\sum_{i=1}^{I}\mathbf{H}_i\mathbf{Q}_i\mathbf{H}_i^H\right)\right] \\ \text{subject to} & \mathbf{Q}_i \in \mathcal{Q}_i, \forall i. \end{array} \tag{4}$$

This is a special case of (1) where the utility function is concave in $\mathbf{Q} \triangleq (\mathbf{Q}_i)_{i=1}^I$, $I_f = 1$, $\mathcal{I}_f = \{1\}$, and the expectation in (4) is taken over the channels $\mathbf{H} \triangleq (\mathbf{H}_i)_{i=1}^I$.

Example #4: The algorithmic framework that will be introduced shortly can be successfully used also to *robustify* distributed iterative algorithms solving *deterministic* (nonconvex) social problems, but in the presence of *inexact estimates* of the system parameters. More specifically, consider for example the following sum-cost minimization multi-agent problem:

$$\begin{array}{ll} \underset{\mathbf{x}}{\text{minimize}} & \sum_{i=1}^{I} f_i(\mathbf{x}_1,\ldots,\mathbf{x}_I) \\ \text{subject to} & \mathbf{x}_i \in \mathcal{X}_i, \quad i=1,\ldots,I, \end{array} \tag{5}$$

where $f_i(\mathbf{x}_i,\mathbf{x}_{-i})$ is uniformly convex in $\mathbf{x}_i \in \mathcal{X}_i$. An efficient distributed algorithm converging to stationary solutions of (5) has been recently proposed in [15]: at each iteration t, given the current iterate $\mathbf{x}^t$, every agent i minimizes (w.r.t. $\mathbf{x}_i \in \mathcal{X}_i$) the following convexified version of the social function:

$$f_i(\mathbf{x}_i,\mathbf{x}_{-i}^t) + \left\langle \mathbf{x}_i - \mathbf{x}_i^t, \sum_{j\neq i}\nabla_i f_j(\mathbf{x}^t)\right\rangle + \tau_i\left\|\mathbf{x}_i - \mathbf{x}_i^t\right\|^2,$$

where $\nabla_i f_j(\mathbf{x})$ stands for $\nabla_{\mathbf{x}_i^*} f_j(\mathbf{x})$, and $\langle \mathbf{a},\mathbf{b}\rangle \triangleq \Re\left(\mathbf{a}^H\mathbf{b}\right)$ ($\|\mathbf{a}\| = \sqrt{\langle \mathbf{a},\mathbf{a}\rangle}$). The evaluation of the above function requires the exact knowledge of $\nabla_i f_j(\mathbf{x}^t)$ for all $j \neq i$. In practice, however, only a noisy estimate of $\nabla_i f_j(\mathbf{x}^t)$ is available [24, 25, 26]. In such cases, convergence of pricing-based algorithms [15, 27, 28, 29] is no longer guaranteed. We will show in Sec. IV-C that the proposed framework can be readily applied, for example, to robustify (and make convergent), e.g., pricing-based schemes, such as [15, 27, 28, 29].

Since the class of problems (1) is in general nonconvex (possibly NP hard [30]), the focus of this paper is to design *distributed* solution methods for computing stationary solutions (possibly local minima) of (1). The major goal is to devise *parallel (nonlinear) best-response* schemes that converge even when the expected value in (1) cannot be computed accurately and only sample values of $\boldsymbol{\xi}$ are available.

III. A Novel Parallel Stochastic Decomposition

The social problem (1) faces two main issues: i) the nonconvexity of the objective functions; and ii) the impossibility to estimate accurately the expected value. To deal with these difficulties, we propose a decomposition scheme that consists in solving a sequence of *parallel strongly convex* subproblems (one for each user), where the objective function of user i is obtained from $U(\mathbf{x})$ by replacing the expected value with a suitably chosen incremental *sample* estimate of it and linearizing the nonconvex part. More formally, at iteration t, a random vector $\boldsymbol{\xi}^t$ is realized,[1] and user i solves the following problem: given $\mathbf{x}^t \in \mathcal{X}$ and $\boldsymbol{\xi}^t \in \Omega$, let

$$\hat{\mathbf{x}}_i(\mathbf{x}^t,\boldsymbol{\xi}^t) \triangleq \underset{\mathbf{x}_i\in\mathcal{X}_i}{\arg\min}\ \hat{f}_i(\mathbf{x}_i;\mathbf{x}^t,\boldsymbol{\xi}^t), \tag{6a}$$

with the surrogate function $\hat{f}_i(\mathbf{x}_i;\mathbf{x}^t,\boldsymbol{\xi}^t)$ defined as

$$\begin{aligned} &\hat{f}_i(\mathbf{x}_i;\mathbf{x}^t,\boldsymbol{\xi}^t) \triangleq \\ &\quad \rho^t\sum_{j\in\mathcal{C}_i^t} f_j(\mathbf{x}_i,\mathbf{x}_{-i}^t,\boldsymbol{\xi}^t) + \rho^t\left\langle \mathbf{x}_i-\mathbf{x}_i^t, \boldsymbol{\pi}_i(\mathbf{x}^t,\boldsymbol{\xi}^t)\right\rangle \\ &\quad +(1-\rho^t)\left\langle \mathbf{x}_i-\mathbf{x}_i^t, \mathbf{f}_i^{t-1}\right\rangle + \tau_i\left\|\mathbf{x}_i-\mathbf{x}_i^t\right\|^2; \end{aligned} \tag{6b}$$

where the pricing vector $\boldsymbol{\pi}_i(\mathbf{x},\boldsymbol{\xi})$ is given by

$$\boldsymbol{\pi}_i\left(\mathbf{x}^t,\boldsymbol{\xi}^t\right) \triangleq \sum_{j\in\overline{\mathcal{C}}_i^t}\nabla_i f_j\left(\mathbf{x}^t,\boldsymbol{\xi}^t\right); \tag{6c}$$

and $\mathbf{f}_i^t$ is an accumulation vector updated recursively according to

$$\mathbf{f}_i^t = (1-\rho^t)\mathbf{f}_i^{t-1} + \rho^t\left(\boldsymbol{\pi}_i(\mathbf{x}^t,\boldsymbol{\xi}^t) + \sum_{j\in\mathcal{C}_i^t}\nabla_i f_j(\mathbf{x}^t,\boldsymbol{\xi}^t)\right), \tag{6d}$$

with $\rho^t \in (0,1]$ being a sequence to be properly chosen ($\rho^0 = 1$). Here $\boldsymbol{\xi}^0,\boldsymbol{\xi}^1,\ldots$ are realizations of random vectors defined on $(\Omega,\mathcal{F},\mathbb{P})$, at iterations $t = 0,1,\ldots$, respectively. The other symbols in (6) are defined as follows:

- In (6b): $\mathcal{C}_i^t$ is any subset of $\mathcal{S}_i^t \triangleq \{i \in \mathcal{I}_f : f_i(\mathbf{x}_i,\mathbf{x}_{-i}^t,\boldsymbol{\xi}^t)$ is convex on $\mathcal{X}_i$, given $\mathbf{x}_{-i}^t$ and $\boldsymbol{\xi}^t\}$; $\mathcal{S}_i^t$ is the set of indices of functions that are convex in $\mathbf{x}_i$;
- In (6c): $\overline{\mathcal{C}}_i^t$ denotes the complement of $\mathcal{C}_i^t$, i.e., $\overline{\mathcal{C}}_i^t \cup \mathcal{C}_i^t = \mathcal{S}_i^t$; thus, it contains (at least) the indices of functions $f_i(\mathbf{x}_i,\mathbf{x}_{-i}^t,\boldsymbol{\xi}^t)$ that are nonconvex in $\mathbf{x}_i$, given $\mathbf{x}_{-i}^t$ and $\boldsymbol{\xi}^t$;
- In (6c)-(6d): $\nabla_i f_j(\mathbf{x},\boldsymbol{\xi})$ is the gradient of $f_j(\mathbf{x},\boldsymbol{\xi})$ w.r.t. $\mathbf{x}_i^*$ (the complex conjugate of $\mathbf{x}_i$). Note that, since $f_j(\mathbf{x},\boldsymbol{\xi})$ is real-valued, $\nabla_{\mathbf{x}_i^*} f(\mathbf{x},\boldsymbol{\xi}) = \nabla_{\mathbf{x}_i^*} f(\mathbf{x},\boldsymbol{\xi})^* = (\nabla_{\mathbf{x}_i} f(\mathbf{x},\boldsymbol{\xi}))^*$.

Given $\hat{\mathbf{x}}_i(\mathbf{x}^t,\boldsymbol{\xi}^t)$, $\mathbf{x} \triangleq (\mathbf{x}_i)_{i=1}^I$ is updated according to

$$\mathbf{x}_i^{t+1} = \mathbf{x}_i^t + \gamma^{t+1}(\hat{\mathbf{x}}_i(\mathbf{x}^t,\boldsymbol{\xi}^t) - \mathbf{x}_i^t),\ i = 1,\ldots,K, \tag{7}$$

where $\gamma^t \in (0,1]$. Note that the iterate $\mathbf{x}^t$ is a function of the past history $\mathcal{F}^t$ of the algorithm up to iteration t (we omit this dependence for notational simplicity):

$$\mathcal{F}^t \triangleq \left\{\mathbf{x}^0,\ldots,\mathbf{x}^{t-1},\boldsymbol{\xi}^0,\ldots,\boldsymbol{\xi}^{t-1}\right\}.$$

Since $\boldsymbol{\xi}^0,\ldots,\boldsymbol{\xi}^{t-1}$ are random vectors, $\hat{\mathbf{x}}_i(\mathbf{x}^t,\boldsymbol{\xi})$ and $\mathbf{x}^t$ are random vectors as well.

The subproblems (6a) have an interesting interpretation: each user minimizes a *sample* convex approximation of the

[1] With slight abuse of notation, throughout the paper, we use the same symbol $\boldsymbol{\xi}^t$ to denote both the random vector $\boldsymbol{\xi}^t$ and its realizations.

Algorithm 1: Stochastic parallel decomposition algorithm

Data: $\boldsymbol{\tau} \triangleq (\tau_i)_{i=1}^I \geq \mathbf{0}$, $\{\gamma^t\}$, $\{\rho^t\}$, $\mathbf{x}^0 \in \mathcal{X}$; set $t = 0$.
(S.1): If $\mathbf{x}^t$ satisfies a suitable termination criterion: STOP.
(S.2): For all $i = 1, \ldots, I$, compute $\widehat{\mathbf{x}}_i(\mathbf{x}^t, \boldsymbol{\xi}^t)$ [cf. (6)].
(S.3): The random vector $\boldsymbol{\xi}^t$ is realized; update $\mathbf{x}^{t+1} = (\mathbf{x}_i^t)_{i=1}^I$ according to

$$\mathbf{x}_i^{t+1} = (1 - \gamma^{t+1})\mathbf{x}_i^t + \gamma^{t+1}\, \widehat{\mathbf{x}}_i(\mathbf{x}^t, \boldsymbol{\xi}^t),\ \forall i = 1, \ldots, I.$$

(S.4): For all $i = 1, \ldots, I$, update $\mathbf{f}_i^t$ according to (6d).
(S.5): $t \leftarrow t + 1$, and go to (S.1).

original nonconvex stochastic function. The first term in (6b) preserves the convex component (or part of it, if $C_i^t \subset S_i^t$) of the sample social function. The second term in (6b)—the vector $\boldsymbol{\pi}_i(\mathbf{x}, \boldsymbol{\xi})$—comes from the linearization of (at least) the nonconvex part. The vector $\mathbf{f}_i^t$ in the third term represents the incremental estimate of $\nabla_{\mathbf{x}^*} U(\mathbf{x}^t)$ (which is not available), as one can readily check by substituting (6c) into (6d):

$$\mathbf{f}_i^t = (1 - \rho^t)\mathbf{f}_i^{t-1} + \rho^t {\textstyle\sum_{j \in \mathcal{I}_f}} \nabla_i f_j(\mathbf{x}^t, \boldsymbol{\xi}^t). \tag{8}$$

Roughly speaking, the goal of this third term is to estimate on-the-fly the unknown $\nabla_{\mathbf{x}^*} U(\mathbf{x}^t)$ by its samples collected over the iterations; based on (8), such an estimate is expected to become more and more accurate as t increases, provided that the sequence ρ^t is properly chosen (this statement is made rigorous shortly in Theorem 1). The last quadratic term in (6b) is the proximal regularization whose numerical benefits are well-understood [31].

Given (6), we define the "best-response" mapping as: given $\boldsymbol{\xi} \in \otimes$,

$$\mathcal{X} \ni \mathbf{y} \mapsto \hat{\mathbf{x}}(\mathbf{y}, \boldsymbol{\xi}) \triangleq (\hat{\mathbf{x}}_i(\mathbf{y}, \boldsymbol{\xi}))_{i=1}^I. \tag{9}$$

Note that $\widehat{\mathbf{x}}(\bullet, \boldsymbol{\xi})$ is well-defined for any given $\boldsymbol{\xi}$ because the objective function in (6) is strongly convex with constant $\tau_{\min}$:

$$\tau_{\min} \triangleq \min_{i=1,\ldots,I} \{\tau_i\}. \tag{10}$$

The proposed decomposition scheme is formally described in Algorithm 1, and its convergence properties are stated in Theorem 1, under the following standard boundedness assumptions on the instantaneous gradient errors [24, 32].

Assumption (c): The instantaneous gradient is unbiased with bounded variance, that is, the following holds almost surely:

$$\mathbb{E}\big[\nabla U(\mathbf{x}^t) - {\textstyle\sum_{j \in \mathcal{I}_f}} \nabla f_j(\mathbf{x}^t, \boldsymbol{\xi}^t) \big| \mathcal{F}^t\big] = \mathbf{0}, \quad \forall t = 0, 1, \ldots,$$

and

$$\mathbb{E}\big[\big\|\nabla U(\mathbf{x}^t) - {\textstyle\sum_{j \in \mathcal{I}_f}} \nabla f_j(\mathbf{x}^t, \boldsymbol{\xi}^t)\big\|^2 \big| \mathcal{F}^t\big] < \infty, \quad \forall t = 0, 1, \ldots.$$

This assumption is readily satisfied if the random variables $\boldsymbol{\xi}^0, \boldsymbol{\xi}^1, \ldots$ are bounded and identically distributed.

Theorem 1. *Given problem (1) under Assumptions (a)-(c), suppose that $\tau_{\min} > 0$ in (6b) and the step-sizes $\{\gamma^t\}$ and $\{\rho^t\}$ are chosen so that*

$$\textbf{i)}\ \gamma^t \to 0,\ {\textstyle\sum_t} \gamma^t = \infty,\ \textstyle\sum (\gamma^t)^2 < \infty, \tag{11a}$$

$$\textbf{ii)}\ \rho^t \to 0,\ {\textstyle\sum_t} \rho^t = \infty,\ \textstyle\sum (\rho^t)^2 < \infty, \tag{11b}$$

$$\textbf{iii)}\ \lim_{t \to \infty} \gamma^t / \rho^t = 0, \tag{11c}$$

$$\textbf{iv)}\ \limsup_{t \to \infty} \rho^t \big({\textstyle\sum_{j \in \mathcal{I}_f}} L_{\nabla f_j(\boldsymbol{\xi}^t)}\big) = 0,\ \text{almost surely}. \tag{11d}$$

Then, every limit point of the sequence $\{\mathbf{x}^t\}$ generated by Algorithm 1 (at least one of such point exists) is a stationary point of (1) almost surely.

Proof: See Appendix A. ∎

On Assumption (c): The boundedness condition is in terms of the conditional expectation of the (random) gradient error. Compared with [18], Assumption (c) is weaker because in [18] it is required that every realization of the (random) gradient error must be bounded.

On Condition (11d): The condition has the following interpretation: all increasing subsequences of $\sum_{j \in \mathcal{I}_f} L_{\nabla f_j(\boldsymbol{\xi}^t)}$ must grow slower than $1/\rho^t$. We will discuss later in Sec. IV how this assumption is satisfied for specific applications. Note that if $\sum_{j \in \mathcal{I}_f} L_{\nabla f_j(\boldsymbol{\xi})}$ is uniformly bounded for any $\boldsymbol{\xi}$ (which is indeed the case if $\boldsymbol{\xi}$ is a bounded random vector), then (11d) is trivially satisfied.

On Algorithm 1: To the best of our knowledge, Algorithm 1 is the first *parallel best-response* (e.g., nongradient-like) scheme for nonconvex stochastic sum-utility problems in the form (1): all the users update *in parallel* their strategies (possibly with a memory) solving a sequence of *decoupled* (strongly) convex subproblems [cf. (6)]. It performs empirically better than classical stochastic gradient-based schemes at no extra cost of signaling, because the convexity of the objective function, if any, is better exploited. Numerical experiments on specific applications confirm this intuition; see Sec. IV. Moreover, by choosing different instances of the set $\mathcal{C}_i^t$ in (6b), one obtains convex subproblems that may exhibit a different trade-off between cost per iteration and convergence speed. Finally, it is guaranteed to converge under very weak assumptions (e.g., weaker than those in [18]) while offering some flexibility in the choice of the free parameters [cf. Theorem 1].

Diminishing stepsize rules: Convergence is guaranteed if a diminishing stepsize rule satisfying (11) is chosen. An instance of (11) is, e.g., the following:

$$\gamma^t = \frac{1}{t^\alpha},\ \rho^t = \frac{1}{t^\beta},\ 0.5 < \beta < \alpha \leq 1. \tag{12}$$

Roughly speaking, (11) says that the stepsizes γ^t and ρ^t, while diminishing (with γ^t decreasing faster than ρ^t), need not go to zero too fast. This kind of stepsize rules are of the same spirit of those used to guarantee convergence of gradient methods with error; see [33] for more details.

Implementation issues: In order to compute the best-response, each user needs to know $\sum_{j \in \mathcal{C}_i^t} f_j(\mathbf{x}_i, \mathbf{x}_{-i}^t, \boldsymbol{\xi}^t)$ and the pricing vector $\boldsymbol{\pi}_i(\mathbf{x}^t, \boldsymbol{\xi}^t)$. The signaling required to acquire this information is generally problem-dependent. If the problem under consideration does not have any specific structure, the most natural message-passing strategy is to communicate directly $\mathbf{x}_{-i}^t$ and $\boldsymbol{\pi}_i(\mathbf{x}^t, \boldsymbol{\xi}^t)$. However, in many specific applications significantly reduced signaling may be required; see Sec. IV for some examples. Note that the signaling is of the same spirit as that of pricing-based algorithms proposed in the literature for the maximization of deterministic sum-utility functions [15, 29]; no extra communication is required to update $\mathbf{f}_i^t$: once the new pricing vector $\boldsymbol{\pi}_i(\mathbf{x}^t, \boldsymbol{\xi}^t)$ is available, the recursive update (6d) for the "incremental" gradient is based on a local accumulation register keeping track of the last iterate $\mathbf{f}_i^{t-1}$. Note also that, thanks to the simultaneous nature of the proposed scheme, the overall communication overhead is expected to be less than that required to implement *sequential* schemes, such the deterministic schemes in [29].

A. Some special cases

We customize next the proposed general algorithmic framework to specific instances of problem (1) arising naturally in many applications.

1) Stochastic proximal conditional gradient methods: Quite interestingly, the proposed decomposition technique resembles classical stochastic conditional gradient schemes [4] when one chooses in (6b) $\mathcal{C}_i^t = \emptyset$, for all i and t, resulting in the following surrogate function:

$$\begin{aligned}\hat{f}_i(\mathbf{x}_i; \mathbf{x}^t, \boldsymbol{\xi}^t) &= \rho^t \left\langle \mathbf{x}_i - \mathbf{x}_i^t, \textstyle\sum_{j\in\mathcal{I}_f} \nabla_i f_j\left(\mathbf{x}^t, \boldsymbol{\xi}^t\right)\right\rangle \\ &\quad + (1-\rho^t)\left\langle \mathbf{x}_i - \mathbf{x}_i^t, \mathbf{f}_i^{t-1}\right\rangle + \tau_i \left\|\mathbf{x}_i - \mathbf{x}_i^t\right\|^2, \end{aligned} \tag{13}$$

with $\mathbf{f}_i^t$ updated according to (8). Note that traditional stochastic conditional gradient methods [9] do not have the proximal regularization term in (13). However, it is worth mentioning that, for some of the applications introduced in Sec. II, it is just the presence of the proximal term that allows one to compute the best-response $\hat{\mathbf{x}}_i(\mathbf{x}^t, \boldsymbol{\xi}^t)$ resulting from the minimization of (13) in closed-form; see Sec. IV-B.

2) Stochastic best-response algorithm for single (convex) functions: Suppose that the social function in (1) is a single function $U(\mathbf{x}) = \mathbb{E}\left[f(\mathbf{x}_1, \ldots, \mathbf{x}_I, \boldsymbol{\xi})\right]$, with $f(\mathbf{x}_1, \ldots, \mathbf{x}_I, \boldsymbol{\xi})$ convex in each $\mathbf{x}_i \in \mathcal{X}_i$ (but not necessarily jointly), for any given ξ. This optimization problem is a special case of the general formulation (1), with $I_f = 1$, $\mathcal{I}_f = \{1\}$ and $S_i^t = \{1\}$. Since $f(\mathbf{x}_1, \ldots, \mathbf{x}_I, \boldsymbol{\xi})$ is componentwise convex, a natural choice for the surrogate functions $\hat{f}_i$ is setting $\mathcal{C}_i^t = S_i^t = \{1\}$ for all t, resulting in the following

$$\begin{aligned}\hat{f}_i(\mathbf{x}_i; \mathbf{x}^t, \boldsymbol{\xi}^t) &= \rho^t f\left(\mathbf{x}_i, \mathbf{x}_{-i}^t, \boldsymbol{\xi}^t\right) \\ &\quad + (1-\rho^t)\left\langle \mathbf{x}_i - \mathbf{x}_i^t, \mathbf{f}_i^{t-1}\right\rangle + \tau_i \left\|\mathbf{x}_i - \mathbf{x}_i^t\right\|^2, \end{aligned} \tag{14}$$

where $\mathbf{f}_i^t$ is updated according to $\mathbf{f}_i^t = (1-\rho^t)\,\mathbf{f}_i^{t-1} + \rho^t \nabla_i f\left(\mathbf{x}^t, \boldsymbol{\xi}^t\right)$. Convergence conditions are still given by Theorem 1. It is worth mentioning that the same choice comes out naturally when $f(\mathbf{x}_1, \ldots, \mathbf{x}_I, \boldsymbol{\xi})$ is uniformly *jointly* convex; in such a case the proposed algorithm converges (in the sense of Theorem 1) to the *global optimum* of $U(\mathbf{x})$. An interesting application of this algorithm is the maximization of the ergodic sum-rate over MIMO MACs in (4), resulting in the first *convergent simultaneous* stochastic MIMO Iterative Waterfilling algorithm in the literature; see Sec. IV-C.

3) Stochastic pricing algorithms: Suppose that $I = I_f$ and each $S_i^t = \{i\}$ (implying that $f_i(\mathbf{x}_i, \mathbf{x}_{-i}, \boldsymbol{\xi})$ is uniformly convex on $\mathcal{X}_i$). By taking each $\mathcal{C}_i^t = \{i\}$ for all t, the surrogate function in (6b) reduces to

$$\begin{aligned}\hat{f}_i(\mathbf{x}_i; \mathbf{x}^t, \boldsymbol{\xi}^t) \triangleq\; &\rho^t f_i(\mathbf{x}_i, \mathbf{x}_{-i}^t, \boldsymbol{\xi}^t) + \rho^t \left\langle \mathbf{x}_i - \mathbf{x}_i^t, \boldsymbol{\pi}_i(\mathbf{x}^t, \boldsymbol{\xi}^t)\right\rangle \\ &+ (1-\rho^t)\left\langle \mathbf{x}_i - \mathbf{x}_i^t, \mathbf{f}_i^{t-1}\right\rangle + \tau_i \left\|\mathbf{x}_i - \mathbf{x}_i^t\right\|^2, \end{aligned} \tag{15}$$

where $\boldsymbol{\pi}_i(\mathbf{x}, \boldsymbol{\xi}) = \sum_{j\neq i} \nabla_i f_j(\mathbf{x}, \boldsymbol{\xi})$ and $\mathbf{f}_i^t = (1-\rho^t)\,\mathbf{f}_i^{t-1} + \rho^t(\boldsymbol{\pi}_i(\mathbf{x}^t, \boldsymbol{\xi}^t) + \nabla_i f_i(\mathbf{x}_i, \mathbf{x}_{-i}^t, \boldsymbol{\xi}^t))$. This is the generalization of the deterministic pricing algorithms [15, 29] to stochastic optimization problems. Examples of this class of problems are the ergodic sum-rate maximization problem over SISO and MIMO IC formulated in (2)-(3); see Sec. IV-A and Sec. IV-B.

4) Stochastic DC programming: A stochastic DC programming problem is formulated as

$$\begin{aligned} &\underset{\mathbf{x}}{\text{minimize}} && \mathbb{E}_{\boldsymbol{\xi}}\left[\textstyle\sum_{j\in\mathcal{I}_f}\left(f_j(\mathbf{x}, \boldsymbol{\xi}) - g_j(\mathbf{x}, \boldsymbol{\xi})\right)\right] \\ &\text{subject to} && \mathbf{x}_i \in \mathcal{X}_i, \quad i = 1, \ldots, I, \end{aligned} \tag{16}$$

where both $f_j(\mathbf{x}, \boldsymbol{\xi})$ and $g_j(\mathbf{x}, \boldsymbol{\xi})$ are uniformly convex functions on $\mathcal{X}$ for any given $\boldsymbol{\xi}$. A natural choice of the surrogate functions $\hat{f}_i$ for (16) is linearizing the concave part of the sample sum-utility function, resulting in the following

$$\begin{aligned}\hat{f}_i(\mathbf{x}_i; \mathbf{x}^t, \boldsymbol{\xi}^t) = \rho^t &\textstyle\sum_{j\in\mathcal{I}_f} f_j(\mathbf{x}_i, \mathbf{x}_{-i}^t, \boldsymbol{\xi}^t) \\ &+ \rho^t \left\langle \mathbf{x}_i - \mathbf{x}_i^t, \boldsymbol{\pi}_i(\mathbf{x}^t, \boldsymbol{\xi}^t)\right\rangle \\ &+ (1-\rho^t)\left\langle \mathbf{x}_i - \mathbf{x}_i^t, \mathbf{f}_i^{t-1}\right\rangle + \tau_i \left\|\mathbf{x}_i - \mathbf{x}_i^t\right\|^2, \end{aligned}$$

where $\boldsymbol{\pi}_i(\mathbf{x}, \boldsymbol{\xi}) \triangleq -\sum_{j\in\mathcal{I}_f} \nabla_i g_j(\mathbf{x}, \boldsymbol{\xi})$ and

$$\mathbf{f}_i^t = \left(1-\rho^t\right)\mathbf{f}_i^{t-1} + \rho^t(\boldsymbol{\pi}_i(\mathbf{x}^t, \boldsymbol{\xi}^t) + \textstyle\sum_{j\in\mathcal{I}_f} \nabla_i f_j(\mathbf{x}_i, \mathbf{x}_{-i}^t, \boldsymbol{\xi}^t)).$$

Comparing the surrogate functions (14)-(16) with (13), one can appreciate the potential advantage of the proposed algorithm over classical gradient-based methods: the proposed schemes preserves the (partial) convexity of the original sample function while gradient-based methods use only first order approximations. The proposed algorithmic framework is thus of the best-response type and empirically it yields faster convergence than gradient-based methods. The improvement in the practical convergence speed will be illustrated numerically in the next section.

IV. Applications

In this section, we customize the proposed algorithmic framework to some of the applications introduced in Sec. II, and compare the resulting algorithms with both classical stochastic gradient algorithms and state-of-the-art schemes proposed for the specific problems under considerations.. Numerical results clearly show that the proposed algorithms compare favorably on state-of-the-art schemes.

A. Sum-rate maximization over frequency-selective ICs

Consider the sum-rate maximization problem over frequency-selective ICs, as introduced in (2). Since the instantaneous rate of each user i,

$$r_i(\mathbf{p}_i, \mathbf{p}_{-i}, \mathbf{h}) = \sum_{n=1}^{N} \log\left(1 + \frac{|h_{ii,n}|^2\, p_{i,n}}{\sigma_{i,n}^2 + \sum_{j\neq i} |h_{ij,n}|^2\, p_{j,n}}\right),$$

is uniformly strongly concave in $\mathbf{p}_i \in \mathcal{P}_i$, a natural choice for the surrogate function $\hat{f}_i$ is the one in (15) wherein $r_i(\mathbf{p}_i, \mathbf{p}_{-i}, \mathbf{h}^t)$ is kept unchanged while $\sum_{j\neq i} r_j(\mathbf{p}_j, \mathbf{p}_{-j}, \mathbf{h}^t)$ is linearized. This leads to the following best-response functions

$$\hat{\mathbf{p}}_i(\mathbf{p}^t, \mathbf{h}^t) = \underset{\mathbf{p}_i \in \mathcal{P}_i}{\arg\max} \left\{ \rho^t \cdot r_i(\mathbf{p}_i, \mathbf{p}_{-i}^t, \mathbf{h}^t) + \rho^t \left\langle \mathbf{p}_i, \boldsymbol{\pi}_i^t \right\rangle + (1-\rho^t) \left\langle \mathbf{p}_i, \mathbf{f}_i^{t-1} \right\rangle - \tfrac{\tau_i}{2} \left\| \mathbf{p}_i - \mathbf{p}_i^t \right\|_2^2 \right\}, \tag{17a}$$

where $\boldsymbol{\pi}_i^t = \boldsymbol{\pi}_i(\mathbf{p}^t, \mathbf{h}^t) \triangleq (\pi_{i,n}(\mathbf{p}^t, \mathbf{h}^t))_{n=1}^N$ with

$$\begin{aligned}\pi_{i,n}(\mathbf{p}^t, \mathbf{h}^t) &= \textstyle\sum_{j\neq i} \nabla_{p_{i,n}} r_j(\mathbf{p}^t, \mathbf{h}^t) \\ &= -\textstyle\sum_{j\neq i} |h_{ji,n}^t|^2 \frac{\mathtt{SINR}_{j,n}^t}{(1+\mathtt{SINR}_{j,n}^t)\cdot \mathtt{MUI}_{j,n}^t}, \\ \mathtt{MUI}_{j,n}^t &\triangleq \sigma_{j,n}^2 + \textstyle\sum_{i\neq j} |h_{ji,n}^t|^2 p_{i,n}^t, \\ \mathtt{SINR}_{j,n}^t &= |h_{jj,n}^t|^2 p_{j,n}^t / \mathtt{MUI}_{j,n}^t.\end{aligned}$$

The variable $\mathbf{f}_i^t$ is updated according to $\mathbf{f}_i^t = (1-\rho^t)\,\mathbf{f}_i^{t-1} + \rho^t(\boldsymbol{\pi}_i^t + \nabla_{\mathbf{p}_i} r_i(\mathbf{p}^t, \mathbf{h}^t))$. Note that $\hat{\mathbf{p}}_i(\mathbf{p}^t, \mathbf{h}^t) \triangleq (\hat{p}_{i,n}(\mathbf{p}^t, \mathbf{h}^t))_{n=1}^N$ in (17a) can be computed in closed-form [15]:

$$\hat{p}_{i,n}(\mathbf{p}^t, \mathbf{h}^t) = \mathtt{WF}\big(\rho^t, \mathtt{SINR}_{i,n}^t / p_{i,n}^t, \tau_i, \rho^t \pi_{i,n}^t + (1-\rho^t) f_{i,n}^{t-1} + \tau_i^t p_{i,n}^t - \mu^\star\big), \tag{18}$$

where

$$\mathtt{WF}(a,b,c,d) = \frac{1}{2}\left[\frac{d}{c} - \frac{1}{b} + \sqrt{\left(\frac{d}{c} + \frac{1}{b}\right)^2 + \frac{4a}{c}}\right]^+,$$

and $\mu^\star$ is the Lagrange multiplier such that $0 \leq \mu^\star \perp \sum_{n=1}^N \hat{p}_{i,n}(\mathbf{p}^t, \mathbf{h}^t) - P_i \leq 0$, and it can be found efficiently using a standard bisection method.

The overall stochastic pricing-based algorithm is then given by Algorithm 1 with best-response mapping defined in (18); convergence is guaranteed under conditions i)-iv) in Theorem 1. Note that the theorem is trivially satisfied using stepsizes rules as required in i)-iii) [e.g., (12)]; the only condition that needs further consideration is condition iv). If $\limsup_{t\to\infty} \rho^t \big(\sum_{j\in\mathcal{I}_f} L_{\nabla f_j(\boldsymbol{\xi}^t)}\big) > 0$, we can assume without loss of generality (w.l.o.g.) that the sequence of the Lipschitz constant $\big\{\sum_{j\in\mathcal{I}_f} L_{\nabla f_j(\boldsymbol{\xi}^t)}\big\}$ is increasing monotonically at a rate no slower than $1/\rho^t$ (we can always limit the discussion to such a subsequence). For any $\bar{h} > 0$, define $p(\bar{h}) \triangleq \text{Prob}(|h_{ij,n}| \geq \bar{h})$ and assume w.l.o.g. that $0 \leq p(\bar{h}) < 1$. Note that the Lipschitz constant $L_{\nabla f_j(\boldsymbol{\xi})}$ is upper bounded by the maximum eigenvalue of the augmented Hessian of $f_j(\mathbf{x}, \boldsymbol{\xi})$ [34], and the maximum eigenvalue increasing monotonically means that the channel coefficient is becoming larger and larger (this can be verified by explicitly calculating the augmented Hessian of $f_j(\mathbf{x}, \boldsymbol{\xi})$; details are omitted due to page limit). Since $\text{Prob}(|h_{ij,n}^{t+1}| \geq |h_{ij,n}^t| \text{ for all } t \geq t_0) \leq \text{Prob}(|h_{ij,n}^{t+1}| \geq \bar{h} \text{ for all } t \geq t_0) = p(\bar{h})^{t-t_0+1} \underset{t\to\infty}{\longrightarrow} 0$, we can infer that the magnitude of the channel coefficient increasing monotonically is an event of probability 0. Therefore, condition (11d) is satisfied.

Numerical results. We simulated a SISO frequency selective IC under the following setting: the number of users is either five or twenty; equal power budget $P_i = P$ and white Gaussian noise variance $\sigma_i^2 = \sigma^2$ are assumed for all users; the SNR of each user $\mathtt{snr} = P/\sigma^2$ is set to 10dB; the instantaneous parallel subchannels $\mathbf{h}^t \triangleq (h_{ij,n}^t)_{i,j,n}$ are generated according to $\mathbf{h}^t = \mathbf{h} + \triangle\mathbf{h}^t$, where $\mathbf{h}$ (generated by MATLAB command `randn`) is fixed while $\triangle\mathbf{h}^t$ is generated at each t using $\delta \cdot$ `randn`, with $\delta = 0.2$ being the noise level.

We considered in Fig. 1 the following algorithms: i) the proposed stochastic best-response pricing algorithm (with $\tau_i = 10^{-8}$ for all i, $\gamma^1 = \rho^0 = \rho^1 = 1$, $\rho^t = 2/(t+2)^{0.6}$, and $\gamma^t = 2/(t+2)^{0.61}$ for $t \geq 2$). At each iteration, the users' best-responses have a closed-form solution, see (18); ii) the stochastic conditional gradient method [9] (with $\gamma^1 = \rho^0 = \rho^1 = 1$, $\rho^t = 1/(t+2)^{0.9}$, and $\gamma^t = 1/(t+2)^{0.91}$ for $t \geq 2$). In each iteration, a linear problem must be solved; iii) and the stochastic gradient projection method, proposed in [26] (with $\gamma^1 = 1$ and $\gamma^t = \gamma^{t-1}(1 - 10^{-3}\gamma^{t-1})$ for $t \geq 2$). At each iteration, the users' updates have a closed-form solution. Note that the stepsizes are tuned such that all algorithms can achieve their best empirical convergence speed.

In Fig. 1, for all the algorithms, we plot two merit functions versus the iteration index, namely: i) the ergodic sum-rate, defined as $\mathbb{E}_{\mathbf{h}}[\sum_{n=1}^N \sum_{i=1}^I r_i(\mathbf{p}^t, \mathbf{h})]$ (with the expected value estimated by the sample mean of 1000 independent realizations); and ii) the "achievable" sum-rate, defined as $\frac{1}{t}\sum_{m=1}^t \sum_{n=1}^N \sum_{i=1}^I r_i(\mathbf{p}^m, \mathbf{h}^m)$, which represents the sum-rate that is actually achieved in practice (it is the time average of the instantaneous (random) sum-rate). The experiment shows that for "small" systems (e.g., five active users), all algorithms perform quite well; the proposed scheme is just slightly faster. However, when the number of users increases (e.g., from 5 to 20), all other (gradient-like) algorithms suffer from slow convergence. Quite interestingly, the proposed scheme demonstrates also good scalability: the convergence speed is not notably affected by the number of users, which makes it applicable to more realistic scenarios. The faster convergence of proposed stochastic best-response pricing algorithm comes from a better exploitation of partial convexity in the problem than what more classical gradient algorithms do, which validates the main idea of this paper.

B. *Sum-rate maximization over MIMO ICs*

In this example we customize Algorithm 1 to solve the sum-rate maximization problem over MIMO ICs (3). Defining

$$r_i(\mathbf{Q}_i, \mathbf{Q}_{-i}, \mathbf{H}) \triangleq \log\det\left(\mathbf{I} + \mathbf{H}_{ii}\mathbf{Q}_i\mathbf{H}_{ii}^H \mathbf{R}_i(\mathbf{Q}_{-i}, \mathbf{H})^{-1}\right)$$

and following a similar approach as in the SISO case, the best-response of each user i becomes [cf. (15)]:

$$\hat{\mathbf{Q}}_i(\mathbf{Q}^t, \mathbf{H}^t) = \underset{\mathbf{Q}_i \in \mathcal{Q}_i}{\arg\max} \Big\{ \rho^t r_i(\mathbf{Q}_i, \mathbf{Q}_{-i}^t, \mathbf{H}^t) + \rho^t \big\langle \mathbf{Q}_i - \mathbf{Q}_i^t, \boldsymbol{\Pi}_i^t \big\rangle + (1-\rho^t)\big\langle \mathbf{Q}_i - \mathbf{Q}_i^t, \mathbf{F}_i^{t-1} \big\rangle - \tau_i \big\| \mathbf{Q}_i - \mathbf{Q}_i^t \big\|^2 \Big\}, \tag{19a}$$

where $\langle \mathbf{A}, \mathbf{B} \rangle \triangleq \Re(\text{tr}(\mathbf{A}^H \mathbf{B}))$; $\boldsymbol{\Pi}_i(\mathbf{Q}, \mathbf{H})$ is given by

$$\begin{aligned}\boldsymbol{\Pi}_i(\mathbf{Q}, \mathbf{H}) &= \textstyle\sum_{j\neq i} \nabla_{\mathbf{Q}_i^*} r_j(\mathbf{Q}, \mathbf{H}) \\ &= \textstyle\sum_{j\neq i} \mathbf{H}_{ji}^H \widetilde{\mathbf{R}}_j(\mathbf{Q}_{-j}, \mathbf{H})\,\mathbf{H}_{ji},\end{aligned} \tag{19b}$$

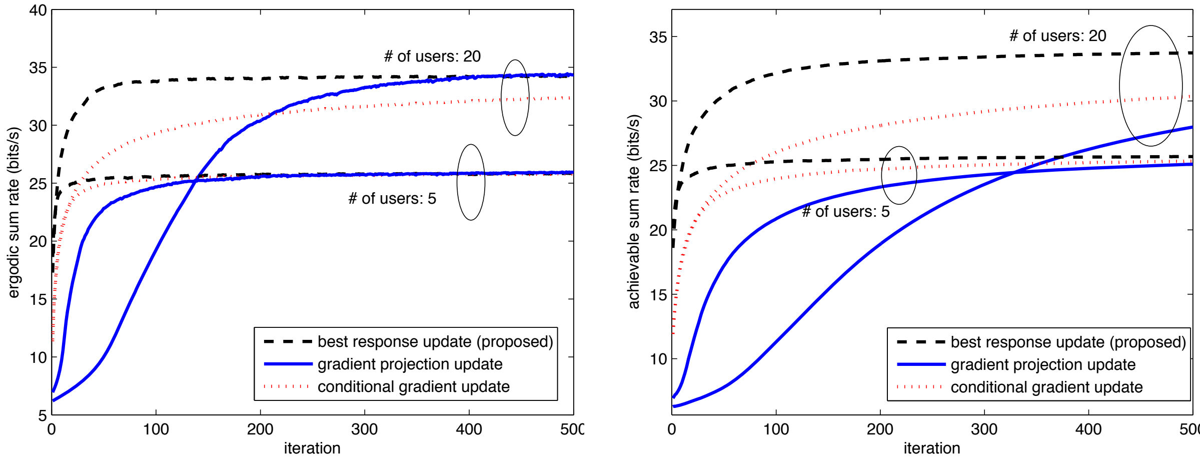

(a) ergodic sum-rate versus iterations (b) achievable sum-rate versus iterations

Figure 1. Sum-rate versus iteration in frequency-selective ICs.

with $r_j(\mathbf{Q},\mathbf{H}) = \log\det(\mathbf{I} + \mathbf{H}_{jj}\mathbf{Q}_j\mathbf{H}_{jj}^H\mathbf{R}_j(\mathbf{Q}_{-i},\mathbf{H})^{-1})$ and $\widetilde{\mathbf{R}}_j(\mathbf{Q}_{-j},\mathbf{H}) \triangleq \left(\mathbf{R}_j(\mathbf{Q}_{-j},\mathbf{H}) + \mathbf{H}_{jj}\mathbf{Q}_j\mathbf{H}_{jj}^H\right)^{-1} - \mathbf{R}_j(\mathbf{Q}_{-j},\mathbf{H})^{-1}$. Then $\mathbf{F}_i^t$ is updated by (6d), which becomes

$$\begin{aligned}\mathbf{F}_i^t &= (1-\rho^t)\,\mathbf{F}_i^{t-1} + \rho^t \textstyle\sum_{j=1}^I \nabla_{\mathbf{Q}_i^*} r_j(\mathbf{Q}^t,\mathbf{H}^t)\\ &= (1-\rho^t)\,\mathbf{F}_i^{t-1} + \rho^t \mathbf{\Pi}_i(\mathbf{Q}^t,\mathbf{H}^t)\\ &\quad + \rho^t(\mathbf{H}_{ii}^t)^H\left(\mathbf{R}_i^t + \mathbf{H}_{ii}^t\mathbf{Q}_i^t(\mathbf{H}_{ii}^t)^H\right)^{-1}\mathbf{H}_{ii}^t. \qquad (19c)\end{aligned}$$

We can then apply Algorithm 1 based on the best-response $\hat{\mathbf{Q}}(\mathbf{Q}^t,\mathbf{H}^t) = (\hat{\mathbf{Q}}_i(\mathbf{Q}^t,\mathbf{H}^t))_{i=1}^I$ whose convergence is guaranteed if the stepsizes are chosen according to Theorem 1.

In contrast to the SISO case, the best-response in (19a) does not have a closed-form solution. A standard option to compute $\hat{\mathbf{Q}}(\mathbf{Q}^t,\mathbf{H}^t)$ is using general-purpose solvers for strongly convex optimization problems. By exploiting the structure of problem (19), we propose next an efficient iterative algorithm converging to $\hat{\mathbf{Q}}(\mathbf{Q}^t,\mathbf{H}^t)$, wherein the subproblems solved at each step have a closed-form solution.

Second-order dual method for problem (19a). To begin with, for notational simplicity, we rewrite (19a) in the following general form:

$$\begin{aligned}&\underset{\mathbf{X}}{\text{maximize}} && \rho\log\det(\mathbf{R}+\mathbf{H}\mathbf{X}\mathbf{H}^H) + \langle\mathbf{A},\mathbf{X}\rangle - \tau\left\|\mathbf{X}-\bar{\mathbf{X}}\right\|^2\\ &\text{subject to} && \mathbf{X}\in\mathcal{Q}, \qquad (20)\end{aligned}$$

where $\mathbf{R}\succ\mathbf{0}$, $\mathbf{A}=\mathbf{A}^H$, $\bar{\mathbf{X}}=\bar{\mathbf{X}}^H$ and $\mathcal{Q}$ is defined in (3). Let $\mathbf{H}^H\mathbf{R}^{-1}\mathbf{H} \triangleq \mathbf{U}\mathbf{D}\mathbf{U}^H$ be the eigenvalue/eigenvector decomposition of $\mathbf{H}^H\mathbf{R}^{-1}\mathbf{H}$, where $\mathbf{U}$ is unitary and $\mathbf{D}$ is diagonal with the diagonal entries arranged in decreasing order. It can be shown that (20) is equivalent to the following problem:

$$\underset{\tilde{\mathbf{X}}\in\mathcal{Q}}{\text{maximize}} \quad \rho\log\det(\mathbf{I}+\tilde{\mathbf{X}}\mathbf{D}) + \langle\tilde{\mathbf{A}},\tilde{\mathbf{X}}\rangle - \tau\left\|\tilde{\mathbf{X}}-\check{\mathbf{X}}\right\|^2, \quad (21)$$

where $\tilde{\mathbf{X}} \triangleq \mathbf{U}^H\mathbf{X}\mathbf{U}$, $\tilde{\mathbf{A}} \triangleq \mathbf{U}^H\mathbf{A}\mathbf{U}$, and $\check{\mathbf{X}} = \mathbf{U}^H\bar{\mathbf{X}}\mathbf{U}$. We now partition $\mathbf{D}\succeq\mathbf{0}$ in two blocks, its positive definite and zero parts ($\tilde{\mathbf{X}}$ is partitioned accordingly):

$$\mathbf{D} = \begin{bmatrix}\mathbf{D}_1 & \mathbf{0}\\ \mathbf{0} & \mathbf{0}\end{bmatrix} \quad \text{and} \quad \tilde{\mathbf{X}} = \begin{bmatrix}\tilde{\mathbf{X}}_{11} & \tilde{\mathbf{X}}_{12}\\ \tilde{\mathbf{X}}_{21} & \tilde{\mathbf{X}}_{22}\end{bmatrix}$$

where $\mathbf{D}_1\succ\mathbf{0}$, and $\tilde{\mathbf{X}}_{11}$ and $\mathbf{D}_1$ have the same dimensions. Problem (21) can be then rewritten as:

$$\underset{\tilde{\mathbf{X}}\in\mathcal{Q}}{\text{maximize}} \quad \rho\log\det(\mathbf{I}+\tilde{\mathbf{X}}_{11}\mathbf{D}_1) + \langle\tilde{\mathbf{A}},\tilde{\mathbf{X}}\rangle - \tau\left\|\tilde{\mathbf{X}}-\check{\mathbf{X}}\right\|^2, \quad (22)$$

Note that, since $\tilde{\mathbf{X}}\in\mathcal{Q}$, by definition $\tilde{\mathbf{X}}_{11}$ must belong to $\mathcal{Q}$ as well. Using this observation and introducing the slack variable $\mathbf{Y}=\tilde{\mathbf{X}}_{11}$, (22) is equivalent to

$$\begin{aligned}&\underset{\tilde{\mathbf{X}},\mathbf{Y}}{\text{maximize}} && \rho\log\det(\mathbf{I}+\mathbf{Y}\mathbf{D}_1) + \langle\tilde{\mathbf{A}},\tilde{\mathbf{X}}\rangle - \tau\left\|\tilde{\mathbf{X}}-\check{\mathbf{X}}\right\|^2\\ &\text{subject to} && \tilde{\mathbf{X}}\in\mathcal{Q},\ \mathbf{Y}=\tilde{\mathbf{X}}_{11},\ \mathbf{Y}\in\mathcal{Q}. \qquad (23)\end{aligned}$$

In the following we solve (23) via dual decomposition (note that the duality gap is zero). Denoting by $\mathbf{Z}$ the matrix of multipliers associated to the linear constraints $\mathbf{Y}=\tilde{\mathbf{X}}_{11}$, the (partial) Lagrangian function of (23) is:

$$\begin{aligned}L(\tilde{\mathbf{X}},\mathbf{Y},\mathbf{Z}) = {}&\rho\log\det(\mathbf{I}+\mathbf{Y}\mathbf{D}_1) + \langle\tilde{\mathbf{A}},\tilde{\mathbf{X}}\rangle\\ &-\tau\left\|\tilde{\mathbf{X}}-\check{\mathbf{X}}\right\|^2 + \langle\mathbf{Z},\mathbf{Y}-\tilde{\mathbf{X}}_{11}\rangle.\end{aligned}$$

The dual problem is then

$$\underset{\mathbf{Z}}{\text{minimize}} \quad d(\mathbf{Z}) = L(\tilde{\mathbf{X}}(\mathbf{Z}),\mathbf{Y}(\mathbf{Z}),\mathbf{Z}),$$

with

$$\tilde{\mathbf{X}}(\mathbf{Z}) = \underset{\tilde{\mathbf{X}}\in\mathcal{Q}}{\arg\max} \ -\tau\left\|\tilde{\mathbf{X}}-\check{\mathbf{X}}\right\|^2 - \langle\mathbf{Z},\tilde{\mathbf{X}}_{11}\rangle, \qquad (24)$$

$$\mathbf{Y}(\mathbf{Z}) = \underset{\mathbf{Y}\in\mathcal{Q}}{\arg\max}\ \rho\log\det(\mathbf{I}+\mathbf{Y}\mathbf{D}_1) + \langle\mathbf{Z},\mathbf{Y}\rangle. \qquad (25)$$

Problem (24) is quadratic and has a closed-form solution (see Lemma 2 below). Similarly, if $\mathbf{Z}\prec\mathbf{0}$, (25) can be solved in closed-form, up to a Lagrange multiplier which can be found efficiently by bisection; see, e.g., [29, Table I]. In our setting, however, $\mathbf{Z}$ in (25) is not necessarily negative definite. Nevertheless, the next lemma provides a closed-form expression of $\mathbf{Y}(\mathbf{Z})$ [and $\tilde{\mathbf{X}}(\mathbf{Z})$].

Lemma 2. *Given (24) and (25) in the setting above, the following hold:*

i) $\tilde{\mathbf{X}}(\mathbf{Z})$ in (24) is given by

$$\tilde{\mathbf{X}}(\mathbf{Z}) = \left[\check{\mathbf{X}} - \frac{1}{2\tau}\left(\mu^\star\mathbf{I} + \begin{bmatrix}\mathbf{Z} & \mathbf{0}\\ \mathbf{0} & \mathbf{0}\end{bmatrix}\right)\right]^+, \quad (26)$$

where $[\mathbf{X}]^+$ denotes the projection of $\mathbf{X}$ onto the cone of positive semidefinite matrices, and $\mu^\star$ is the multiplier such that $0 \leq \mu^\star \perp \text{tr}(\tilde{\mathbf{X}}(\mathbf{Z})) - P \leq 0$, which can be found by bisection;

ii) $\mathbf{Y}(\mathbf{Z})$ in (25) is unique and is given by

$$\mathbf{Y}(\mathbf{Z}) = \mathbf{V}\,[\rho\,\mathbf{I} - \mathbf{\Sigma}^{-1}]^+\,\mathbf{V}^H, \tag{27}$$

where $(\mathbf{V}, \mathbf{\Sigma})$ is the generalized eigenvalue decomposition of $(\mathbf{D}_1, -\mathbf{Z}+\mu^\star\mathbf{I})$, and $\mu^\star$ is the multiplier such that $0 \leq \mu^\star \perp \text{tr}(\mathbf{Y}(\mathbf{Z})) - P \leq 0$; $\mu^\star$ can be found by bisection over $[\underline{\mu}, \overline{\mu}]$, with $\underline{\mu} \triangleq [\lambda_{\max}(\mathbf{Z})]^+$ and $\overline{\mu} \triangleq [\lambda_{\max}(\mathbf{D}_1) + \lambda_{\max}(\mathbf{Z})/\rho]^+$.

Proof. See Appendix B. ■

Since $(\tilde{\mathbf{X}}(\mathbf{Z}), \mathbf{Y}(\mathbf{Z}))$ is unique, $d(\mathbf{Z})$ is differentiable , with conjugate gradient [22]

$$\nabla_{\mathbf{Z}^*} d(\mathbf{Z}) = \mathbf{Y}(\mathbf{Z}) - \tilde{\mathbf{X}}_{11}(\mathbf{Z}).$$

One can then solve the dual problem using standard (proximal) gradient-based methods; see, e.g., [34]. As a matter of fact, $d(\mathbf{Z})$ is twice continuously differentiable, whose augmented Hessian matrix [22] is given by [34, Sec. 4.2.4]:

$$\begin{aligned}\nabla^2_{\mathbf{Z}\mathbf{Z}^*} d(\mathbf{Z}) = &-[\mathbf{I} \;\; -\mathbf{I}]^H \cdot \\ &\left[\text{bdiag}(\nabla^2_{\mathbf{Y}\mathbf{Y}^*} L(\tilde{\mathbf{X}}, \mathbf{Y}, \mathbf{Z}), \nabla^2_{\tilde{\mathbf{X}}_{11}\tilde{\mathbf{X}}_{11}^*} L(\tilde{\mathbf{X}}, \mathbf{Y}, \mathbf{Z}))\right]^{-1} \cdot \\ &[\mathbf{I} \;\; -\mathbf{I}]\big|_{\tilde{\mathbf{X}}=\tilde{\mathbf{X}}(\mathbf{Z}), \mathbf{Y}=\mathbf{Y}(\mathbf{Z})},\end{aligned}$$

with

$$\begin{aligned}\nabla^2_{\mathbf{Y}\mathbf{Y}^*} L(\tilde{\mathbf{X}}, \mathbf{Y}, \mathbf{Z}) = &-\rho^2 \cdot (\mathbf{D}_1^{1/2}(\mathbf{I} + \mathbf{D}_1^{1/2}\mathbf{Y}\mathbf{D}_1^{1/2})^{-1}\mathbf{D}_1^{1/2})^T \\ &\otimes (\mathbf{D}_1^{1/2}(\mathbf{I} + \mathbf{D}_1^{1/2}\mathbf{Y}\mathbf{D}_1^{1/2})^{-1}\mathbf{D}_1^{1/2}),\end{aligned}$$

and $\nabla^2_{\tilde{\mathbf{X}}_{11}\tilde{\mathbf{X}}_{11}^*} L(\tilde{\mathbf{X}}\mathbf{Y}, \mathbf{Z}) = -\tau\mathbf{I}$. Since $\mathbf{D}_1 \succ \mathbf{0}$, it follows that $\nabla^2_{\mathbf{Z}\mathbf{Z}^*} d(\mathbf{Z}) \succ \mathbf{0}$ and the following second-order Newton's method can be used to update the dual variable $\mathbf{Z}$:

$$\text{vec}(\mathbf{Z}^{t+1}) = \text{vec}(\mathbf{Z}^t) - (\nabla^2_{\mathbf{Z}\mathbf{Z}^*} d(\mathbf{Z}^t))^{-1}\text{vec}(\nabla d(\mathbf{Z}^t)).$$

The convergence speed of the Newton's methods is typically fast, and, in particular, superlinear convergence rate can be achieved when $\mathbf{Z}^t$ is close to $\mathbf{Z}^\star$ [34, Prop. 1.4.1]. ■

As a final remark on efficient solution methods computing $\hat{\mathbf{Q}}_i(\mathbf{Q}^t, \mathbf{H}^t)$, note that one can also apply the *proximal* conditional gradient method as introduced in (13), which is based on a fully linearization of the social function plus a proximal regularization term:

$$\begin{aligned}\hat{\mathbf{Q}}_i(\mathbf{Q}^t, \mathbf{H}^t) &= \underset{\mathbf{Q}_i \in \mathcal{Q}}{\arg\max}\left\{\left\langle \mathbf{Q}_i - \mathbf{Q}_i^t, \mathbf{F}_i^t \right\rangle - \tau_i \left\|\mathbf{Q}_i - \mathbf{Q}_i^t\right\|^2\right\} \\ &= \left[\mathbf{Q}_i^t + \frac{1}{2\tau_i}(\mathbf{F}_i^t - \mu^\star\mathbf{I})\right]^+,\end{aligned} \tag{28}$$

where $\mu^\star$ is the Lagrange multiplier that can be found efficiently by the bisection method. Note that (28) differs from more traditional conditional stochastic gradient methods [9] by the presence of the proximal regularization, thanks to which one can solve (28) in closed-form [cf. Lemma 2].

The above examples, (19) and (28), clearly show the flexibility of the proposed scheme: choosing different instances of the set $\mathcal{C}_i^t$ leads to convex subproblems exhibiting a different trade-off between cost per iteration and practical convergence speed. Roughly speaking, when the number of iterations matters, one can opt for the approximation problem (19). On the other hand, when the cost per iteration is the priority, one can instead choose the approximation problem (28).

Practical implementations. The proposed algorithm is fairly distributed: once the pricing matrix $\mathbf{\Pi}_i$ is given, to compute the best-response, each user only needs to locally estimate the covariance matrix of the interference plus noise. Note that both the computation of $\hat{\mathbf{Q}}_i(\mathbf{Q}, \mathbf{H})$ and the update of $\mathbf{F}_i$ can be implemented locally by each user. The estimation of the pricing matrix $\mathbf{\Pi}_i$ requires however some signaling among nearby receivers. Interestingly, the pricing expression and thus the resulting signaling overhead necessary to compute it coincide with [29] (where a sequential algorithm is proposed for the deterministic maximization of the sum-rate over MIMO ICs) and the stochastic gradient projection method in [26]. We remark that the signaling to compute (19b) is lower than in [18], wherein signaling exchange is required twice (one in the computation of $\mathbf{U}_i$ and another in that of $\mathbf{A}_i$; see [18] for more details) in a single iteration to transmit among users the auxiliary variables which are of same dimensions as $\mathbf{\Pi}_i$.

Numerical Results. We considered the same scenario as in the SISO case [cf. Sec. IV-A] with the following differences: i) there are 50 users; ii) the channels are matrices generated according to $\mathbf{H}^t = \mathbf{H} + \triangle\mathbf{H}^t$, where $\mathbf{H}$ is given while $\triangle\mathbf{H}^t$ is realization dependent and generated by $\delta \cdot \texttt{randn}$, with noise level $\delta = 0.2$; and iii) the number of transmit and receive antennas is four. We simulate the following algorithms:

• The proposed stochastic best-response pricing algorithm (19) (with $\tau_i = 10^{-8}$ for all i) under two stepsizes rules, namely: *Stepsize 1* (empirically optimal): $\rho^t = 2/(t+2)^{0.6}$ and $\gamma^t = 2/(t+2)^{0.61}$ for $t \geq 2$; and *Stepsize 2*: $\rho^t = 2/(t+2)^{0.7}$ and $\gamma^t = 2/(t+2)^{0.71}$ for $t \geq 2$. For both stepsize rules we set $\gamma^1 = \rho^0 = \rho^1 = 1$. The best-response is computed using the second-order dual method, whose convergence has been observed in a few iterations;

• The proposed stochastic proximal gradient method (28) with $\tau = 0.01$ and same stepsize as the stochastic best-response pricing algorithm. The users' best-responses have a a closed-form expression;

• The stochastic conditional gradient method [9] (with $\gamma^1 = \rho^0 = \rho^1 = 1$ and $\rho^t = 1/(t+2)^{0.9}$ and $\gamma^t = 1/(t+2)^{0.91}$ for $t \geq 2$). In each iteration, a linear problem must be solved;

• The stochastic weighted minimum mean-square-error (SWMMSE) method [18]. The convex subproblems to be solved at each iteration have a closed-form solution.

Similarly to the SISO ICs case, we consider both ergodic sum-rate and achievable sum-rate. In Fig. 2 we plot both objective functions versus the iteration index. It is clear from the figures that the proposed best-response pricing and proximal gradient algorithms outperform current schemes in terms of both convergence speed and achievable (ergodic or instantaneous) sum-rate. Note also that the best-response pricing algorithm is very scalable compared with the other

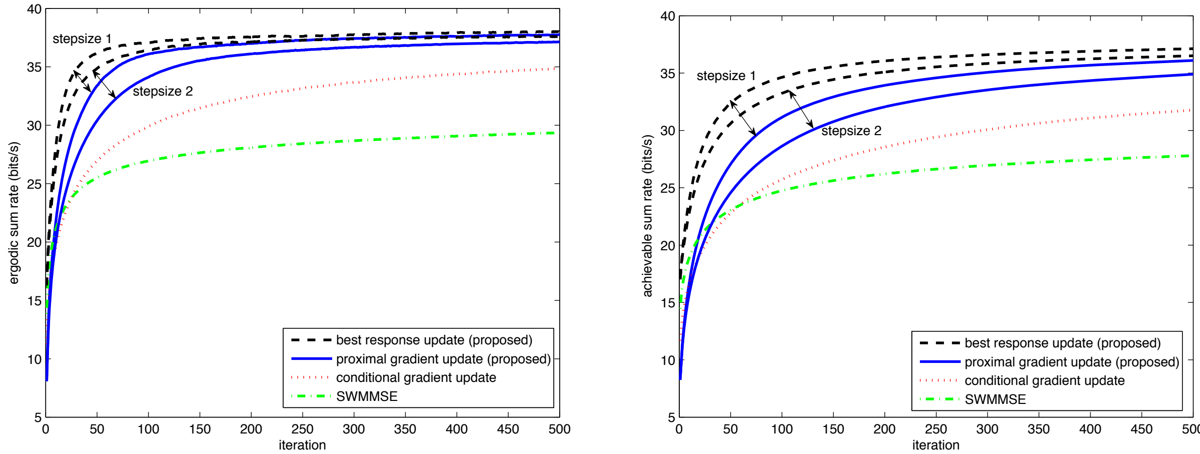

(a) ergodic sum-rate versus iterations

(b) achievable sum-rate versus iterations

Figure 2. Sum-rate versus iteration in a 50-user MIMO IC

algorithms. Finally, it is interesting to note that the proposed stochastic proximal gradient algorithm outperforms the conditional stochastic gradient method in terms of both convergence speed and cost per iteration. This is mainly due to the presence of the proximal regularization term in (19a).

Note that in order to achieve a satisfactory convergence speed, some tuning of the free parameters in the stepsize rules is typically required for all algorithms. Comparing the convergence behavior under two different sets of stepsize rules, we see from Fig. 2 (a) that, as expected, the proposed best-response pricing and proximal gradient algorithms under the faster decreasing Stepsize 2 converge slower than they do under Stepsize 1, but the difference is relatively small and the proposed algorithms still converge to a larger sum-rate in a smaller number of iterations than current schemes do. Hence this offers some extra tolerance in the stepsizes and makes the proposed algorithms quite applicable in practice.

C. Sum-rate maximization over MIMO MACs

In this example we consider the sum-rate maximization problem over MIMO MACs, as introduced in (4). This problem has been studied in [36] using standard convex optimization techniques, under the assumption that the statistics of CSI are available and the expected value of the sum-rate function in (4) can be computed analytically. When this assumption does not hold, we can turn to the proposed algorithm with proper customization: Define

$$r(\mathbf{H},\mathbf{Q}) \triangleq \log\det\left(\mathbf{R}_{\mathrm{N}}+\sum_{i=1}^{I}\mathbf{H}_i\mathbf{Q}_i\mathbf{H}_i^H\right).$$

A natural choice for the best-response of each user i in each iteration of Algorithm 1 is [cf. (14)]:

$$\begin{aligned}\hat{\mathbf{Q}}_i(\mathbf{Q}^t,\mathbf{H}^t) = \underset{\mathbf{Q}_i\in\mathcal{Q}_i}{\arg\max}\Big\{\rho^t\, r(\mathbf{H}^t,\mathbf{Q}_i,\mathbf{Q}_{-i}^t)\\ +(1-\rho^t)\big\langle\mathbf{Q}_i-\mathbf{Q}_i^t,\mathbf{F}_i^{t-1}\big\rangle-\tau_i\left\|\mathbf{Q}_i-\mathbf{Q}_i^t\right\|^2\Big\},\end{aligned} \quad (29)$$

and $\mathbf{F}_i^t$ is updated as $\mathbf{F}_i^t=(1-\rho^t)\mathbf{F}_i^{t-1}+\rho^t\,\nabla_{\mathbf{Q}_i^*}r(\mathbf{H}^t,\mathbf{Q}^t)$ while $\nabla_{\mathbf{Q}_i^*}r(\mathbf{H},\mathbf{Q})=\mathbf{H}_i^H(\mathbf{R}_{\mathrm{N}}+\sum_{i=1}^{I}\mathbf{H}_i\mathbf{Q}_i\mathbf{H}_i^H)^{-1}\mathbf{H}_i$.

Note that since the instantaneous sum-rate function $\log\det(\mathbf{R}_{\mathrm{N}}+\sum_{i=1}^{I}\mathbf{H}_i\mathbf{Q}_i\mathbf{H}^H)$ is jointly concave in $\mathbf{Q}_i$ for any $\mathbf{H}$, the ergodic sum-rate function is concave in $\mathbf{Q}_i$'s, and thus Algorithm 1 will converge (in the sense of Theorem 1) to the *global* optimal solution of (4). To the best of our knowledge, this is the first example of stochastic approximation algorithms based on best-response dynamics rather than gradient responses.

Numerical results. We compare the proposed best-response method (29) (whose solution is computed using the second-order dual method in Sec. IV-B) with the stochastic conditional gradient method [9], and the stochastic gradient projection method [8]. System parameters (including the stepsize rules) are set as for the MIMO IC example in Sec. IV-B. In Fig. 3 we plot both the ergodic sum-rate and the achievable sum-rate versus the iteration index. This figure clearly shows that Algorithm 1 outperforms the conditional gradient method and the gradient projection method in terms of convergence speed, and the performance gap is increasing as the number of users increases. This is because the proposed algorithm is a best-response type scheme, which thus explores the concavity of each user's rate function better than what gradient methods do. Note also that the proposed method exhibit good scalability properties.

D. Distributed deterministic algorithms with errors

The developed framework can also be used to robustify some algorithms proposed for the deterministic counterpart of the multi-agent optimization problem (1), when only noisy estimates of the users' objective functions are available. As a specific example, we show next how to robustify the *deterministic* best-response-based pricing algorithm proposed in [15]. Consider the deterministic optimization problem introduced in (5). The main iterate of the best-response algorithm [15] is given by (7) but with each $\widehat{\mathbf{x}}_i(\mathbf{x}^t)$ defined as [15]

$$\widehat{\mathbf{x}}_i(\mathbf{x}^t) = \underset{\mathbf{x}_i\in\mathcal{X}_i}{\arg\min}\left\{\begin{array}{l}\sum_{j\in\mathcal{C}_i}f_i(\mathbf{x}_i,\mathbf{x}_{-i}^t)+\left\langle\mathbf{x}_i-\mathbf{x}_i^t,\boldsymbol{\pi}_i(\mathbf{x}^t)\right\rangle\\ +\tau_i\left\|\mathbf{x}_i-\mathbf{x}_i^t\right\|^2,\end{array}\right\}, \quad (30)$$

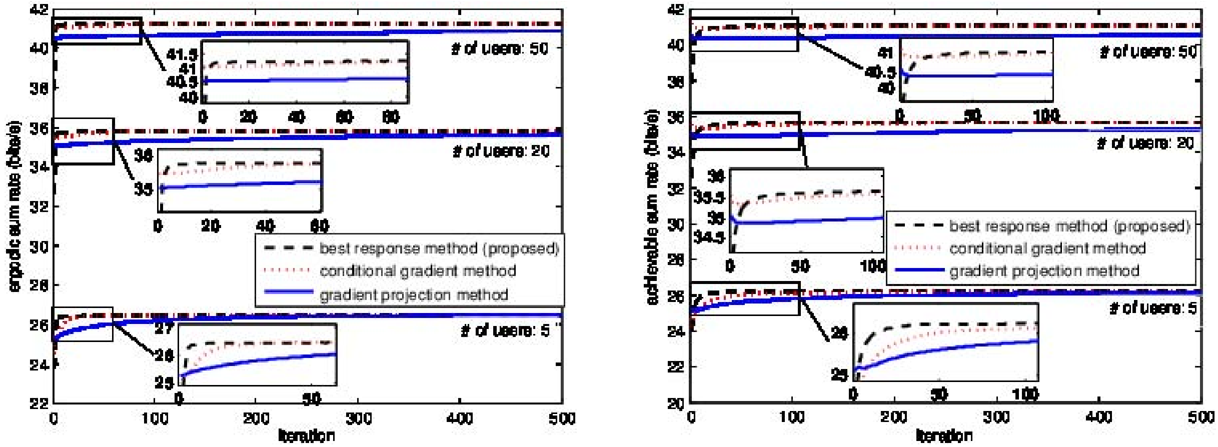

(a) ergodic sum-rate versus iterations

(b) achievable sum-rate versus iterations

Figure 3. Sum-rate versus iteration in MIMO MAC

where $\boldsymbol{\pi}_i(\mathbf{x}) = \sum_{j\in\overline{\mathcal{C}}_i}\nabla_i f_j(\mathbf{x})$. In many applications (see, e.g., [24, 25, 26]), however, only a noisy estimate of $\boldsymbol{\pi}_i(\mathbf{x})$ is available, denoted by $\widetilde{\boldsymbol{\pi}}_i(\mathbf{x})$. A heuristic is then to replace in (30) the exact $\boldsymbol{\pi}_i(\mathbf{x})$ with its noisy estimate $\widetilde{\boldsymbol{\pi}}_i(\mathbf{x})$. The limitation of this approach, albeit natural, is that convergence of the resulting scheme is no longer guaranteed.

If $\widetilde{\boldsymbol{\pi}}_i(\mathbf{x})$ is unbiased, i.e., $\mathbb{E}\left[\widetilde{\boldsymbol{\pi}}_i(\mathbf{x}^t)|\mathcal{F}^t\right] = \boldsymbol{\pi}_i(\mathbf{x}^t)$ [24, 25], capitalizing on the proposed framework, we can readily deal with estimation errors while guaranteeing convergence. In particular, it is sufficient to modify (30) as follows:

$$\widetilde{\mathbf{x}}_i(\mathbf{x}^t) = \underset{\mathbf{x}_i\in\mathcal{X}_i}{\arg\min}\left\{\textstyle\sum_{j\in\mathcal{C}_i} f_j(\mathbf{x}_i,\mathbf{x}_{-i}^t) + \rho^t\left\langle \mathbf{x}_i-\mathbf{x}_i^t, \widetilde{\boldsymbol{\pi}}_i(\mathbf{x}^t)\right\rangle\right.$$
$$\left. +(1-\rho^t)\left\langle \mathbf{x}_i-\mathbf{x}_i^t,\mathbf{f}_i^{t-1}\right\rangle + \tau_i\left\|\mathbf{x}_i-\mathbf{x}_i^t\right\|^2\right\}, \qquad (31)$$

where $\mathbf{f}_i^t$ is updated according to $\mathbf{f}_i^t = (1-\rho^t)\mathbf{f}_i^{t-1}+\rho^t\widetilde{\boldsymbol{\pi}}_i(\mathbf{x}^t)$. Algorithm 1 based on the best-response (31) is then guaranteed to converge to a stationary solution of (5), in the sense specified by Theorem 1.

As a case study, we consider next the maximization of the deterministic sum-rate over MIMO ICs in the presence of pricing estimation errors:

$$\begin{aligned}&\underset{\mathbf{Q}}{\text{maximize}} && \textstyle\sum_{i=1}^{I}\log\det(\mathbf{I}+\mathbf{H}_{ii}\mathbf{Q}_i\mathbf{H}_{ii}^H\mathbf{R}_i(\mathbf{Q}_{-i})^{-1})\\ &\text{subject to} && \mathbf{Q}_i\succeq\mathbf{0},\ \text{tr}(\mathbf{Q}_i)\le P_i,\quad i=1,\ldots,I.\end{aligned} \qquad (32)$$

Then (31) becomes:

$$\widehat{\mathbf{Q}}_i(\mathbf{Q}^t) = \underset{\mathbf{Q}_i\in\mathcal{Q}_i}{\arg\max}\Big\{\log\det\left(\mathbf{R}_i^t+\mathbf{H}_{ii}^t\mathbf{Q}_i(\mathbf{H}_{ii}^t)^H\right)$$
$$+\left\langle\mathbf{Q}_i-\mathbf{Q}_i^t,\rho^t\widetilde{\boldsymbol{\Pi}}_i^t+(1-\rho^t)\mathbf{F}_i^{t-1}\right\rangle - \tau_i\left\|\mathbf{Q}_i-\mathbf{Q}_i^t\right\|^2\Big\}, \qquad (33)$$

where $\widetilde{\boldsymbol{\Pi}}_i^t$ is a noisy estimate of $\boldsymbol{\Pi}_i(\mathbf{Q}^t,\mathbf{H})$ given by (19b)[2] and $\mathbf{F}_i^t$ is updated according to $\mathbf{F}_i^t = \rho^t\widetilde{\boldsymbol{\Pi}}_i^t + (1-\rho^t)\mathbf{F}_i^{t-1}$. Given $\widehat{\mathbf{Q}}_i(\mathbf{Q}^t)$, the main iterate of the algorithm becomes $\mathbf{Q}_i^{t+1} = \mathbf{Q}_i^t + \gamma^{t+1}\left(\hat{\mathbf{Q}}_i(\mathbf{Q}^t)-\mathbf{Q}_i^t\right)$. Almost sure convergence to a stationary point of the deterministic optimization problem (32) is guaranteed by Theorem 1. Note that if the channel matrices $\{\mathbf{H}_{ii}\}$ are full column-rank, one can also set in (33) all $\tau_i = 0$, and compute (33) in closed-form [cf. Lemma 2].

Numerical results. We consider the maximization of the deterministic sum-rate (32) over a 5-user MIMO IC. The other system parameters (including the stepsize rules) are set as in the numerical example in Sec. IV-B. The noisy estimate $\widetilde{\boldsymbol{\Pi}}_i$ of the nominal price matrix $\boldsymbol{\Pi}_i$ [defined in (19b)] is $\widetilde{\boldsymbol{\Pi}}_i^t = \boldsymbol{\Pi}_i + \Delta\boldsymbol{\Pi}_i^t$, where $\Delta\boldsymbol{\Pi}_i^t$ is firstly generated as $\Delta\mathbf{H}^t$ in Sec. IV-B and then only its Hermitian part is kept; the noise level δ is set to 0.05. We compare the following algorithms: i) the proposed robust pricing method−Algorithm 1 based on the best-response defined in (33); and ii) the plain pricing method as proposed in [15] [cf. (30)]. Note that the variable update in both algorithms has a closed-form solution. We also include as a benchmark the sum-rate achieved by the plain pricing method (30) when there is no estimation noise (i.e., perfect $\boldsymbol{\pi}_i(\mathbf{x})$ is available). In Fig. 4 we plot the deterministic sum-rate in (32) versus the iteration index t. As expected, Fig. 4 shows that the plain pricing method [15] is not robust to pricing estimation errors, whereas the proposed robustification preforms well. For instance, the rate achievable by the proposed method is about 50% larger than that of [15], and is observed to reach the benchmark value (achieved by the plain pricing method when there is no estimation noise). This is due to the fact that the proposed robustification filters out the estimation noise. Note that the limit point generated by the proposed scheme (33) is a stationary solution of the deterministic problem (32).

V. A more general SCA framework

The key idea behind the choice of the surrogate function $\hat{f}_i(\mathbf{x}_i;\mathbf{x}^t,\boldsymbol{\xi}^t)$ in (6) is to convexify the nonconvex part of the sample sum-utility function via partial linearization of $\sum_{j\in\overline{\mathcal{C}}_i^t} f_j(\mathbf{x}^t,\boldsymbol{\xi}^t)$. It is not difficult to show that one can generalize this idea and replace the surrogate $\hat{f}_i(\mathbf{x}_i;\mathbf{y}^t,\boldsymbol{\xi}^t)$ in (6) with a more general function. For example, one can use

[2] $\boldsymbol{\Pi}_i(\mathbf{Q},\mathbf{H})$ is always negative definite by definition [29], but $\widetilde{\boldsymbol{\Pi}}_i^t$ may not be so. However, it is reasonable to assume $\widetilde{\boldsymbol{\Pi}}_i^t$ to be Hermitian.

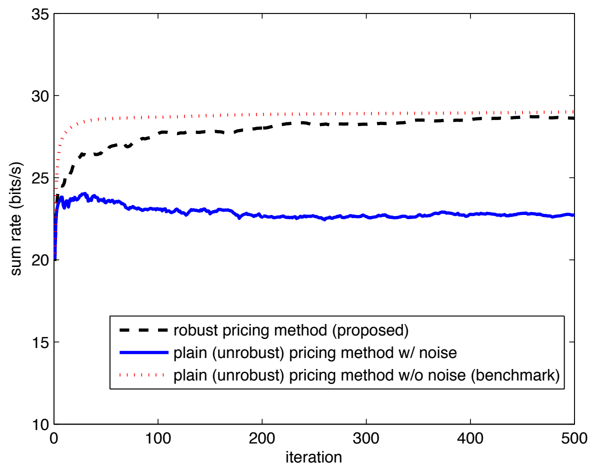

Figure 4. Maximization of deterministic sum-rate over MIMO IC under noisy parameter estimation: sum-rate versus iteration.

in Algorithm 1 the following sample best-response function

$$\hat{\mathbf{x}}_i(\mathbf{x}^t,\boldsymbol{\xi}^t)\triangleq\underset{\mathbf{x}_i\in\mathcal{X}_i}{\arg\min}\left\{\rho^t\tilde{f}_i(\mathbf{x}_i;\mathbf{x}^t,\boldsymbol{\xi}^t)+(1-\rho^t)\left\langle\mathbf{f}_i^{t-1},\mathbf{x}_i-\mathbf{x}_i^t\right\rangle\right\}. \tag{34}$$

where $\mathbf{f}_i^t = (1-\rho^t)\,\mathbf{f}_i^{t-1} + \rho^t\nabla_i f\left(\mathbf{x}^t,\boldsymbol{\xi}^t\right)$ [cf. (6)], and $\tilde{f}_i(\mathbf{x}_i;\mathbf{x}^t,\boldsymbol{\xi}^t)$ is any surrogate function satisfying the following technical conditions :

(A1) $\tilde{f}_i(\mathbf{x}_i;\mathbf{x}^t,\boldsymbol{\xi}^t)$ is uniformly strongly convex and continuously differentiable on $\mathcal{X}_i$ for all given $\mathbf{x}^t$ and $\boldsymbol{\xi}^t$;

(A2) $\nabla_{\mathbf{x}^t}\tilde{f}_i(\mathbf{x}_i;\mathbf{x}^t,\boldsymbol{\xi}^t)$ is Lipschitz continuous on $\mathcal{X}$;

(A3) $\nabla_i\tilde{f}_i(\mathbf{x}_i^t;\mathbf{x}^t,\boldsymbol{\xi}^t)=\sum_{j\in\mathcal{I}_f}\nabla_i f_j(\mathbf{x}^t,\boldsymbol{\xi}^t)$.

All the convergence results presented so far are still valid (cf. Theorem 1). To the best of our knowledge, this is the first SCA framework for nonconvex stochastic optimization problems; it offers a lot of flexibility to tailor the surrogate function to individual problems, while guaranteeing convergence, which makes it appealing for a wide range of applications.

VI. Conclusions

In this paper, we have proposed a novel best-response-based solution method for general stochastic *nonconvex* multi-agent optimization problems and analyzed its convergence properties. The proposed novel decomposition enables all users to update their optimization variables *in parallel* by solving a sequence of strongly convex subproblems; which makes the algorithm very appealing for the distributed implementation in several practical systems. We have then customized the general framework to solve special classes of problems and applications, including the stochastic maximization of the sum-rate over frequency-selective ICs, MIMO ICs and MACs. Extensive experiments have provided a solid evidence of the superiority in terms of both achievable sum-rate and practical convergence of the proposed schemes with respect to to state-of-the-art stochastic-based algorithms.

Appendix

A. Proof of Theorem 1

We first introduce the following two preliminary results.

Lemma 3. *Given problem (1) under Assumptions (a)-(c), suppose that the stepsizes $\{\gamma^t\}$ and $\{\rho^t\}$ are chosen according to (11). Let $\{\mathbf{x}^t\}$ be the sequence generated by Algorithm 1. Then, the following holds*

$$\lim_{t\to\infty}\left\|\mathbf{f}^t-\nabla U(\mathbf{x}^t)\right\|=0, \qquad \text{w.p.1.}$$

Proof: This lemma is a consequence of [10, Lemma 1]. To see this, we just need to verify that all the technical conditions therein are satisfied by the problem at hand. Specifically, Condition (a) of [10, Lemma 1] is satisfied because $\mathcal{X}_i$'s are closed and bounded in view of Assumption (a). Condition (b) of [10, Lemma 1] is exactly Assumption (c). Conditions (c)-(d) come from the stepsize rules i)-ii) in (11) of Theorem 1. Condition (e) of [10, Lemma 1] comes from the Lipschitz property of ∇U from Assumption (b) and stepsize rule iii) in (11) of Theorem 1. ■

Lemma 4. *Given problem (1) under Assumptions (a)-(c), suppose that the stepsizes $\{\gamma^t\}$ and $\{\rho^t\}$ are chosen according to (11). Let $\{\mathbf{x}^t\}$ be the sequence generated by Algorithm 1. Then, there exists a constant $\hat{L}$ such that*

$$\left\|\hat{\mathbf{x}}(\mathbf{x}^{t_1},\boldsymbol{\xi}^{t_1})-\hat{\mathbf{x}}(\mathbf{x}^{t_2},\boldsymbol{\xi}^{t_2})\right\|\leq\hat{L}\left\|\mathbf{x}^{t_1}-\mathbf{x}^{t_2}\right\|+e(t_1,t_2),$$

and $\lim_{t_1,t_2\to\infty}e(t_1,t_2)=0$ *w.p.1.*

Proof: We assume w.l.o.g. that $t_2>t_1$; for notational simplicity, we define $\hat{\mathbf{x}}_i^t\triangleq\hat{\mathbf{x}}_i(\mathbf{x}^t,\boldsymbol{\xi}^t)$, for $t=t_1$ and $t=t_2$. It follows from the first-order optimality condition that [22]

$$\left\langle\mathbf{x}_i-\hat{\mathbf{x}}_i^{t_1},\nabla_i\hat{f}_i(\hat{\mathbf{x}}_i^{t_1};\mathbf{x}^{t_1},\boldsymbol{\xi}^{t_1})\right\rangle\geq 0, \tag{35a}$$

$$\left\langle\mathbf{x}_i-\hat{\mathbf{x}}_i^{t_2},\nabla_i\hat{f}_i(\hat{\mathbf{x}}_i^{t_2};\mathbf{x}^{t_2},\boldsymbol{\xi}^{t_2})\right\rangle\geq 0. \tag{35b}$$

Setting $\mathbf{x}_i=\hat{\mathbf{x}}_i(\mathbf{x}^{t_2},\boldsymbol{\xi}^{t_2})$ in (35a) and $\mathbf{x}_i=\hat{\mathbf{x}}_i(\mathbf{x}^{t_1},\boldsymbol{\xi}^{t_1})$ in (35b), and adding the two inequalities, we have

$$\begin{aligned}0&\geq\left\langle\hat{\mathbf{x}}_i^{t_1}-\hat{\mathbf{x}}_i^{t_2},\nabla_i\hat{f}_i(\hat{\mathbf{x}}_i^{t_1};\mathbf{x}^{t_1},\boldsymbol{\xi}^{t_1})-\nabla_i\hat{f}_i(\hat{\mathbf{x}}_i^{t_2};\mathbf{x}^{t_2},\boldsymbol{\xi}^{t_2})\right\rangle\\&=\left\langle\hat{\mathbf{x}}_i^{t_1}-\hat{\mathbf{x}}_i^{t_2},\nabla_i\hat{f}_i(\hat{\mathbf{x}}_i^{t_1};\mathbf{x}^{t_1},\boldsymbol{\xi}^{t_1})-\nabla_i\hat{f}_i(\hat{\mathbf{x}}_i^{t_1};\mathbf{x}^{t_2},\boldsymbol{\xi}^{t_2})\right\rangle\\&\quad+\left\langle\hat{\mathbf{x}}_i^{t_1}-\hat{\mathbf{x}}_i^{t_2},\nabla_i\hat{f}_i(\hat{\mathbf{x}}_i^{t_1};\mathbf{x}^{t_2},\boldsymbol{\xi}^{t_2})-\nabla_i\hat{f}_i(\hat{\mathbf{x}}_i^{t_2};\mathbf{x}^{t_2},\boldsymbol{\xi}^{t_2})\right\rangle.\end{aligned} \tag{36}$$

The first term in (36) can be lower bounded as follows:

$$\begin{aligned}&\left\langle\hat{\mathbf{x}}_i^{t_1}-\hat{\mathbf{x}}_i^{t_2},\nabla_i\hat{f}_i(\hat{\mathbf{x}}_i^{t_1};\mathbf{x}^{t_1},\boldsymbol{\xi}^{t_1})-\nabla_i\hat{f}_i(\hat{\mathbf{x}}_i^{t_1};\mathbf{x}^{t_2},\boldsymbol{\xi}^{t_2})\right\rangle\\&=\rho^{t_1}\textstyle\sum_{j\in\mathcal{C}_i^{t_1}}\left\langle\hat{\mathbf{x}}_i^{t_1}-\hat{\mathbf{x}}_i^{t_2},\right.\\&\qquad\qquad\left.\nabla_i f_j(\hat{\mathbf{x}}_i^{t_1},\mathbf{x}_{-i}^{t_1},\boldsymbol{\xi}^{t_1})-\nabla_i f_j(\mathbf{x}_i^{t_1},\mathbf{x}_{-i}^{t_1},\boldsymbol{\xi}^{t_1})\right\rangle\\&\quad-\rho^{t_2}\textstyle\sum_{j\in\mathcal{C}_i^{t_2}}\left\langle\hat{\mathbf{x}}_i^{t_1}-\hat{\mathbf{x}}_i^{t_2},\right.\\&\qquad\qquad\left.\nabla_i f_j(\hat{\mathbf{x}}_i^{t_1},\mathbf{x}_{-i}^{t_2},\boldsymbol{\xi}^{t_2})-\nabla_i f_j(\mathbf{x}_i^{t_2},\mathbf{x}_{-i}^{t_2},\boldsymbol{\xi}^{t_2})\right\rangle\\&\quad+\left\langle\hat{\mathbf{x}}_i^{t_1}-\hat{\mathbf{x}}_i^{t_2},\mathbf{f}_i^{t_1}-\mathbf{f}_i^{t_2}\right\rangle-\tau_i\left\langle\hat{\mathbf{x}}_i^{t_1}-\hat{\mathbf{x}}_i^{t_2},\mathbf{x}_i^{t_1}-\mathbf{x}_i^{t_2}\right\rangle\end{aligned} \tag{37a}$$

where in (37a) we used (8). Invoking the Lipschitz continuity of $\nabla f_j(\mathbf{x}_i^t, \mathbf{x}_{-i}^t, \boldsymbol{\xi}^t)$, we can get a lower bound for (37a):

$$\begin{aligned}
&\big\langle \hat{\mathbf{x}}_i^{t_1} - \hat{\mathbf{x}}_i^{t_2}, \nabla_i \hat{f}_i(\hat{\mathbf{x}}_i^{t_1}; \mathbf{x}^{t_1}, \boldsymbol{\xi}^{t_1}) - \nabla_i \hat{f}_i(\hat{\mathbf{x}}_i^{t_1}; \mathbf{x}^{t_2}, \boldsymbol{\xi}^{t_2}) \big\rangle \\
\geq\ & -\rho^{t_1} \textstyle\sum_{j\in\mathcal{C}_i^{t_1}} \big\|\hat{\mathbf{x}}_i^{t_1} - \hat{\mathbf{x}}_i^{t_2}\big\| \cdot \\
&\qquad \big\|\nabla_i f_j(\hat{\mathbf{x}}_i^{t_1}, \mathbf{x}_{-i}^{t_1}, \boldsymbol{\xi}^{t_1}) - \nabla_i f_j(\mathbf{x}_i^{t_1}, \mathbf{x}_{-i}^{t_1}, \boldsymbol{\xi}^{t_1})\big\| \\
& -\rho^{t_2} \textstyle\sum_{j\in\mathcal{C}_i^{t_2}} \big\|\hat{\mathbf{x}}_i^{t_1} - \hat{\mathbf{x}}_i^{t_2}\big\| \cdot \\
&\qquad \big\|\nabla_i f_j(\hat{\mathbf{x}}_i^{t_1}, \mathbf{x}_{-i}^{t_2}, \boldsymbol{\xi}^{t_2}) - \nabla_i f_j(\mathbf{x}_i^{t_2}, \mathbf{x}_{-i}^{t_2}, \boldsymbol{\xi}^{t_2})\big\| \\
& + \big\langle \hat{\mathbf{x}}_i^{t_1} - \hat{\mathbf{x}}_i^{t_2}, \mathbf{f}_i^{t_1} - \nabla_i U(\mathbf{x}^{t_1}) - \mathbf{f}_i^{t_2} + \nabla_i U(\mathbf{x}^{t_2}) \big\rangle \\
& + \big\langle \hat{\mathbf{x}}_i^{t_1} - \hat{\mathbf{x}}_i^{t_2}, \nabla_i U(\mathbf{x}^{t_1}) - \nabla_i U(\mathbf{x}^{t_2}) \big\rangle \\
& - \tau_i \big\langle \hat{\mathbf{x}}_i^{t_1} - \hat{\mathbf{x}}_i^{t_2}, \mathbf{x}_i^{t_1} - \mathbf{x}_i^{t_2} \big\rangle, \qquad (37b) \\
\geq\ & -\rho^{t_1} \Big(\textstyle\sum_{j\in\mathcal{I}_f} L_{\nabla f_j(\boldsymbol{\xi}^{t_1})}\Big) \big\|\hat{\mathbf{x}}_i^{t_1} - \hat{\mathbf{x}}_i^{t_2}\big\| \cdot \big\|\hat{\mathbf{x}}_i^{t_1} - \mathbf{x}_i^{t_1}\big\| \\
& -\rho^{t_2} \Big(\textstyle\sum_{j\in\mathcal{I}_f} L_{\nabla f_j(\boldsymbol{\xi}^{t_2})}\Big) \big\|\hat{\mathbf{x}}_i^{t_1} - \hat{\mathbf{x}}_i^{t_2}\big\| \cdot \big\|\hat{\mathbf{x}}_i^{t_1} - \mathbf{x}^{t_2}\big\| \\
& - \big\|\hat{\mathbf{x}}_i^{t_1} - \hat{\mathbf{x}}_i^{t_2}\big\| (\varepsilon^{t_1} + \varepsilon^{t_2}) \\
& - (L_{\nabla U} + \tau_{\max}) \big\|\hat{\mathbf{x}}_i^{t_1} - \hat{\mathbf{x}}_i^{t_2}\big\| \big\|\mathbf{x}^{t_1} - \mathbf{x}^{t_2}\big\| \qquad (37c) \\
\geq\ & -\rho^{t_1} \Big(\textstyle\sum_{j\in\mathcal{I}_f} L_{\nabla f_j(\boldsymbol{\xi}^{t_1})}\Big) C_x \big\|\hat{\mathbf{x}}_i^{t_1} - \hat{\mathbf{x}}_i^{t_2}\big\| \\
& -\rho^{t_2} \Big(\textstyle\sum_{j\in\mathcal{I}_f} L_{\nabla f_j(\boldsymbol{\xi}^{t_2})}\Big) C_x \big\|\hat{\mathbf{x}}_i^{t_1} - \hat{\mathbf{x}}_i^{t_2}\big\| \\
& - \big\|\hat{\mathbf{x}}_i^{t_1} - \hat{\mathbf{x}}_i^{t_2}\big\| (\varepsilon^{t_1} + \varepsilon^{t_2}) \\
& - (L_{\nabla U} + \tau_{\max}) \big\|\hat{\mathbf{x}}_i^{t_1} - \hat{\mathbf{x}}_i^{t_2}\big\| \big\|\mathbf{x}^{t_1} - \mathbf{x}^{t_2}\big\|, \qquad (37d)
\end{aligned}$$

where (37c) comes from the Lipschitz continuity of $\nabla f_j(\mathbf{x}_i^t, \mathbf{x}_{-i}^t, \boldsymbol{\xi}^t)$, with $\varepsilon^t \triangleq \|\mathbf{f}^t - \nabla U(\mathbf{x}^t)\|$ and $\tau_{\max} = \max_{1\leq i\leq I} \tau_i < \infty$, and we used the boundedness of the constraint set $\mathcal{X}$ ($\big\|\mathbf{x} - \mathbf{y}\big\| \leq C_x$ for some $C_x < \infty$ and all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$) and the Lipschitz continuity of $\nabla U(\mathbf{x})$ in (37d).

The second term in (36) can be bounded as:

$$\begin{aligned}
&\big\langle \hat{\mathbf{x}}_i^{t_1} - \hat{\mathbf{x}}_i^{t_2}, \nabla_i \hat{f}(\hat{\mathbf{x}}_i^{t_1}; \mathbf{x}^{t_2}, \boldsymbol{\xi}^{t_2}) - \nabla_i \hat{f}(\hat{\mathbf{x}}_i^{t_2}; \mathbf{x}^{t_2}, \boldsymbol{\xi}^{t_2}) \big\rangle \\
&= \rho^{t_2} \sum_{j\in\mathcal{C}_i^{t_2}} \big\langle \hat{\mathbf{x}}_i^{t_1} - \hat{\mathbf{x}}_i^{t_2}, \nabla_i f_j(\hat{\mathbf{x}}_i^{t_1}, \mathbf{x}_{-i}^{t_2}, \boldsymbol{\xi}^{t_2}) \big\rangle \\
&\quad - \rho^{t_2} \sum_{j\in\mathcal{C}_i^{t_2}} \big\langle \hat{\mathbf{x}}_i^{t_1} - \hat{\mathbf{x}}_i^{t_2}, \nabla_i f_j(\hat{\mathbf{x}}_i^{t_2}, \mathbf{x}_{-i}^{t_2}, \boldsymbol{\xi}^{t_2}) \big\rangle \\
&\quad + \tau_i \big\|\hat{\mathbf{x}}_i^{t_1} - \hat{\mathbf{x}}_i^{t_2}\big\|^2 \geq \tau_{\min} \big\|\hat{\mathbf{x}}_i^{t_1} - \hat{\mathbf{x}}_i^{t_2}\big\|^2,
\end{aligned} \qquad (38)$$

where the inequality follows from the definition of $\tau_{\min}$ and the (uniformly) convexity of the functions $f_j(\bullet, \mathbf{x}_{-i}^t, \boldsymbol{\xi}^t)$.

Combining the inequalities (36), (37d) and (38), we have

$$\begin{aligned}
\big\|\hat{\mathbf{x}}_i^{t_1} - \hat{\mathbf{x}}_i^{t_2}\big\| \leq\ & (L_{\nabla U} + \tau_{\max}) \tau_{\min}^{-1} \big\|\mathbf{x}^{t_1} - \mathbf{x}^{t_2}\big\| \\
& + \tau_{\min}^{-1} C_x \rho^{t_1} \Big(\textstyle\sum_{j\in\mathcal{I}_f} L_{\nabla f_j(\boldsymbol{\xi}^{t_1})}\Big) \\
& + \tau_{\min}^{-1} C_x \rho^{t_2} \Big(\textstyle\sum_{j\in\mathcal{I}_f} L_{\nabla f_j(\boldsymbol{\xi}^{t_2})}\Big) \\
& + \tau_{\min}^{-1} (\varepsilon^{t_1} + \varepsilon^{t_2}),
\end{aligned}$$

which leads to the desired (asymptotic) Lipschitz property:

$$\big\|\hat{\mathbf{x}}^{t_1} - \hat{\mathbf{x}}^{t_2}\big\| \leq \hat{L} \big\|\mathbf{x}^{t_1} - \mathbf{x}^{t_2}\big\| + e(t_1, t_2),$$

with $\hat{L} \triangleq I\, \tau_{\min}^{-1} (L_{\nabla U} + \tau_{\max})$ and

$$\begin{aligned}
e(t_1, t_2) \triangleq\ & I\, \tau_{\min}^{-1} \Big((\varepsilon^{t_1} + \varepsilon^{t_2}) + \\
& + C_x \big(\rho^{t_1} \textstyle\sum_{j\in\mathcal{I}_f} L_{\nabla f_j(\boldsymbol{\xi}^{t_1})} + \rho^{t_2} \sum_{j\in\mathcal{I}_f} L_{\nabla f_j(\boldsymbol{\xi}^{t_2})} \big) \Big).
\end{aligned}$$

In view of Lemma 3 and (11d), it is easy to check that $\lim_{t_1\to\infty, t_2\to\infty} e(t_1, t_2) = 0$ w.p.1. ∎

Proof of Theorem 1. Invoking the first-order optimality conditions of (6), we have

$$\begin{aligned}
& \rho^t \big\langle \mathbf{x}_i^t - \hat{\mathbf{x}}_i^t, \textstyle\sum_{j\in\mathcal{C}_i^t} \nabla_i f_j(\hat{\mathbf{x}}_i^t, \mathbf{x}_{-i}^t, \boldsymbol{\xi}^t) + \boldsymbol{\pi}_i(\mathbf{x}^t, \boldsymbol{\xi}^t) \big\rangle \\
& \quad + (1 - \rho^t) \big\langle \mathbf{x}_i^t - \hat{\mathbf{x}}_i^t, \mathbf{f}_i^{t-1} \big\rangle + \tau_i \big\langle \mathbf{x}_i^t - \hat{\mathbf{x}}_i^t, \hat{\mathbf{x}}_i^t - \mathbf{x}_i^t \big\rangle \\
& = \rho^t \textstyle\sum_{j\in\mathcal{C}_i^t} \big\langle \mathbf{x}_i^t - \hat{\mathbf{x}}_i^t, \nabla_i f_j(\hat{\mathbf{x}}_i^t, \mathbf{x}_{-i}^t, \boldsymbol{\xi}^t) - \nabla_i f_j(\mathbf{x}_i^t, \mathbf{x}_{-i}^t, \boldsymbol{\xi}^t) \big\rangle \\
& \quad + \big\langle \mathbf{x}_i^t - \hat{\mathbf{x}}_i^t, \mathbf{f}_i^t \big\rangle - \tau_i \big\|\hat{\mathbf{x}}_i^t - \mathbf{x}_i^t\big\|^2 \geq 0,
\end{aligned}$$

which together with the convexity of $\sum_{j\in\mathcal{C}_i^t} f_j(\bullet, \mathbf{x}_{-i}^t, \boldsymbol{\xi}^t)$ leads to

$$\big\langle \hat{\mathbf{x}}_i^t - \mathbf{x}_i^t, \mathbf{f}_i^t \big\rangle \leq -\tau_{\min} \big\|\hat{\mathbf{x}}_i^t - \mathbf{x}_i^t\big\|^2. \qquad (39)$$

It follows from the descent lemma on U that

$$\begin{aligned}
U(\mathbf{x}^{t+1}) \leq\ & U(\mathbf{x}^t) + \gamma^{t+1} \big\langle \hat{\mathbf{x}}^t - \mathbf{x}^t, \nabla U(\mathbf{x}^t) \big\rangle \\
& + L_{\nabla U} (\gamma^{t+1})^2 \big\|\hat{\mathbf{x}}^t - \mathbf{x}^t\big\|^2 \\
=\ & U(\mathbf{x}^t) + \gamma^{t+1} \big\langle \hat{\mathbf{x}}^t - \mathbf{x}^t, \nabla U(\mathbf{x}^t) - \mathbf{f}^t + \mathbf{f}^t \big\rangle \\
& + L_{\nabla U} (\gamma^{t+1})^2 \big\|\hat{\mathbf{x}}^t - \mathbf{x}^t\big\|^2 \\
\leq\ & U(\mathbf{x}^t) - \gamma^{t+1} (\tau_{\min} - L_{\nabla U} \gamma^{t+1}) \big\|\hat{\mathbf{x}}^t - \mathbf{x}^t\big\|^2 \\
& + \gamma^{t+1} \big\|\hat{\mathbf{x}}^t - \mathbf{x}^t\big\| \big\|\nabla U(\mathbf{x}^t) - \mathbf{f}^t\big\|, \qquad (40)
\end{aligned}$$

where in the last inequality we used (39). Let us show by contradiction that $\liminf_{t\to\infty} \big\|\hat{\mathbf{x}}^t - \mathbf{x}^t\big\| = 0$ w.p.1. Suppose $\liminf_{t\to\infty} \big\|\hat{\mathbf{x}}^t - \mathbf{x}^t\big\| \geq \chi > 0$ with a positive probability. Then we can find a realization such that at the same time $\big\|\hat{\mathbf{x}}^t - \mathbf{x}^t\big\| \geq \chi > 0$ for all t and $\lim_{t\to\infty} \big\|\nabla U(\mathbf{x}^t) - \mathbf{f}^t\big\| = 0$; we focus next on such a realization. Using $\big\|\hat{\mathbf{x}}^t - \mathbf{x}^t\big\| \geq \chi > 0$, the inequality (40) is equivalent to

$$\begin{aligned}
& U(\mathbf{x}^{t+1}) - U(\mathbf{x}^t) \leq \\
& -\gamma^{t+1} \Big(\tau_{\min} - L_{\nabla U} \gamma^{t+1} - \tfrac{1}{\chi} \|\nabla U(\mathbf{x}^t) - \mathbf{f}^t\| \Big) \big\|\hat{\mathbf{x}}^t - \mathbf{x}^t\big\|^2.
\end{aligned} \qquad (41)$$

Since $\lim_{t\to\infty} \big\|\nabla U(\mathbf{x}^t) - \mathbf{f}^t\big\| = 0$, there exists a t_0 sufficiently large such that

$$\tau_{\min} - L_{\nabla U} \gamma^{t+1} - \frac{1}{\chi} \big\|\nabla U(\mathbf{x}^t) - \mathbf{f}^t\big\| \geq \bar{\tau} > 0, \quad \forall\, t \geq t_0. \qquad (42)$$

Therefore, it follows from (41) and (42) that

$$U(\mathbf{x}^t) - U(\mathbf{x}^{t_0}) \leq -\bar{\tau} \chi^2 \textstyle\sum_{n=t_0}^t \gamma^{n+1}, \qquad (43)$$

which, in view of $\sum_{n=t_0}^\infty \gamma^{n+1} = \infty$, contradicts the boundedness of $\{U(\mathbf{x}^t)\}$. Therefore it must be $\liminf_{t\to\infty} \|\hat{\mathbf{x}}^t - \mathbf{x}^t\| = 0$ w.p.1.

We prove now that $\limsup_{t\to\infty} \|\hat{\mathbf{x}}^t - \mathbf{x}^t\| = 0$ w.p.1. Assume $\limsup_{t\to\infty} \|\hat{\mathbf{x}}^t - \mathbf{x}^t\| > 0$ with some positive probability. We focus next on a realization along with $\limsup_{t\to\infty} \|\hat{\mathbf{x}}^t - \mathbf{x}^t\| > 0$, $\lim_{t\to\infty} \big\|\nabla U(\mathbf{x}^t) - \mathbf{f}^t\big\| = 0$, $\liminf_{t\to\infty} \big\|\hat{\mathbf{x}}^t - \mathbf{x}^t\big\| = 0$, and $\lim_{t_i, t_2\to\infty} e(t_1, t_2) = 0$, where $e(t_1, t_2)$ is defined in Lemma 4. It follows from $\limsup_{t\to\infty} \|\hat{\mathbf{x}}^t - \mathbf{x}^t\| > 0$ and $\liminf_{t\to\infty} \big\|\hat{\mathbf{x}}^t - \mathbf{x}^t\big\| = 0$ that there exists a $\delta > 0$ such that $\|\triangle \mathbf{x}^t\| \geq 2\delta$ (with $\triangle \mathbf{x}^t \triangleq \hat{\mathbf{x}}^t - \mathbf{x}^t$) for infinitely many t and also $\|\triangle \mathbf{x}^t\| < \delta$ for infinitely many t. Therefore, one can always find an infinite set of indexes, say $\mathcal{T}$, having the following properties: for any $t \in \mathcal{T}$, there exists an integer $i_t > t$ such that

$$\begin{aligned} &\|\triangle \mathbf{x}^t\| < \delta, \quad \left\|\triangle \mathbf{x}^{i_t}\right\| > 2\delta, \\ &\delta \leq \|\triangle \mathbf{x}^n\| \leq 2\delta, \quad t < n < i_t. \end{aligned} \tag{44}$$

Given the above bounds, the following holds: for all $t \in \mathcal{T}$,

$$\begin{aligned} \delta &\leq \left\|\triangle \mathbf{x}^{i_t}\right\| - \left\|\triangle \mathbf{x}^t\right\| \\ &\leq \left\|\triangle \mathbf{x}^{i_t} - \triangle \mathbf{x}^t\right\| = \left\|(\hat{\mathbf{x}}^{i_t} - \mathbf{x}^{i_t}) - (\hat{\mathbf{x}}^t - \mathbf{x}^t)\right\| \\ &\leq \left\|\hat{\mathbf{x}}^{i_t} - \hat{\mathbf{x}}^t\right\| + \left\|\mathbf{x}^{i_t} - \mathbf{x}^t\right\| \\ &\leq (1+\hat{L})\left\|\mathbf{x}^{i_t} - \mathbf{x}^t\right\| + e(i_t, t) \\ &\leq (1+\hat{L})\textstyle\sum_{n=t}^{i_t-1}\gamma^{n+1}\left\|\triangle \mathbf{x}^n\right\| + e(i_t, t) \\ &\leq 2\delta(1+\hat{L})\textstyle\sum_{n=t}^{i_t-1}\gamma^{n+1} + e(i_t, t), \end{aligned} \tag{45}$$

implying that

$$\liminf_{\mathcal{T}\ni t\to\infty} \textstyle\sum_{n=t}^{i_t-1}\gamma^{n+1} \geq \bar{\delta}_1 \triangleq \dfrac{1}{2(1+\hat{L})} > 0. \tag{46}$$

Proceeding as in (45), we also have: for all $t \in \mathcal{T}$,

$$\begin{aligned} \left\|\triangle \mathbf{x}^{t+1}\right\| - \left\|\triangle \mathbf{x}^t\right\| &\leq \left\|\triangle \mathbf{x}^{t+1} - \triangle \mathbf{x}^t\right\| \\ &\leq (1+\hat{L})\gamma^{t+1}\left\|\triangle \mathbf{x}^t\right\| + e(t, t+1), \end{aligned}$$

which leads to

$$(1+(1+\hat{L})\gamma^{t+1})\left\|\triangle \mathbf{x}^t\right\| + e(t,t+1) \geq \left\|\triangle \mathbf{x}^{t+1}\right\| \geq \delta, \tag{47}$$

where the second inequality follows from (44). It follows from (47) that there exists a $\bar{\delta}_2 > 0$ such that for sufficiently large $t \in \mathcal{T}$,

$$\left\|\triangle \mathbf{x}^t\right\| \geq \frac{\delta - e(t,t+1)}{1 + (1+\hat{L})\gamma^{t+1}} \geq \bar{\delta}_2 > 0. \tag{48}$$

Here after we assume w.l.o.g. that (48) holds for all $t \in \mathcal{T}$ (in fact one can always restrict $\{\mathbf{x}^t\}_{t\in\mathcal{T}}$ to a proper subsequence).

We show now that (46) is in contradiction with the convergence of $\{U(\mathbf{x}^t)\}$. Invoking (40), we have: for all $t \in \mathcal{T}$,

$$\begin{aligned} &U(\mathbf{x}^{t+1}) - U(\mathbf{x}^t) \\ &\quad\leq -\gamma^{t+1}\left(\tau_{\min} - L_{\nabla U}\gamma^{t+1}\right)\left\|\hat{\mathbf{x}}^t - \mathbf{x}^t\right\|^2 \\ &\qquad + \gamma^{t+1}\delta\left\|\nabla U(\mathbf{x}^t) - \mathbf{f}^t\right\| \\ &\quad\leq -\gamma^{t+1}\left(\tau_{\min} - L_{\nabla U}\gamma^{t+1} - \frac{\left\|\nabla U(\mathbf{x}^t) - \mathbf{f}^t\right\|}{\delta}\right) \\ &\qquad\cdot\left\|\hat{\mathbf{x}}^t - \mathbf{x}^t\right\|^2 + \gamma^{t+1}\delta\left\|\nabla U(\mathbf{x}^t) - \mathbf{f}^t\right\|^2, \end{aligned} \tag{49}$$

and for $t < n < i_t$,

$$\begin{aligned} &U(\mathbf{x}^{n+1}) - U(\mathbf{x}^n) \\ &\quad\leq -\gamma^{n+1}\left(\tau_{\min} - L_{\nabla U}\gamma^{n+1} - \frac{\left\|\nabla U(\mathbf{x}^n) - \mathbf{f}^n\right\|}{\left\|\hat{\mathbf{x}}^n - \mathbf{x}^n\right\|}\right) \\ &\qquad\cdot\left\|\hat{\mathbf{x}}^n - \mathbf{x}^n\right\|^2 \\ &\quad\leq -\gamma^{n+1}\left(\tau_{\min} - L_{\nabla U}\gamma^{n+1} - \frac{\left\|\nabla U(\mathbf{x}^n) - \mathbf{f}^n\right\|}{\delta}\right) \\ &\qquad\cdot\left\|\hat{\mathbf{x}}^n - \mathbf{x}^n\right\|^2, \end{aligned} \tag{50}$$

where the last inequality follows from (44). Adding (49) and (50) over $n = t+1, \ldots, i_t - 1$ and, for $t \in \mathcal{T}$ sufficiently large (so that $\tau_{\min} - L_{\nabla U}\gamma^{t+1} - \delta^{-1}\left\|\nabla U(\mathbf{x}^n) - \mathbf{f}^n\right\| \geq \hat{\tau} > 0$ and $\left\|\nabla U(\mathbf{x}^t) - \mathbf{f}^t\right\| < \hat{\tau}\bar{\delta}_2^2/\delta$), we have

$$\begin{aligned} &U(\mathbf{x}^{i_t}) - U(\mathbf{x}^t) \\ &\overset{(a)}{\leq} -\hat{\tau}\textstyle\sum_{n=t}^{i_t-1}\gamma^{n+1}\left\|\hat{\mathbf{x}}^n - \mathbf{x}^n\right\|^2 + \gamma^{t+1}\delta\left\|\nabla U(\mathbf{x}^t) - \mathbf{f}^t\right\| \\ &\overset{(b)}{\leq} -\hat{\tau}\,\bar{\delta}_2^2\textstyle\sum_{n=t+1}^{i_t-1}\gamma^{n+1} - \gamma^{t+1}\left(\hat{\tau}\bar{\delta}_2^2 - \delta\left\|\nabla U(\mathbf{x}^t) - \mathbf{f}^t\right\|\right) \\ &\overset{(c)}{\leq} -\hat{\tau}\,\bar{\delta}_2^2\textstyle\sum_{n=t+1}^{i_t-1}\gamma^{n+1}, \end{aligned} \tag{51}$$

where (a) follows from $\tau_{\min} - L_{\nabla U}\gamma^{t+1} - \delta^{-1}\left\|\nabla U(\mathbf{x}^n) - \mathbf{f}^n\right\| \geq \hat{\tau} > 0$; (b) is due to (48); and in (c) we used $\left\|\nabla U(\mathbf{x}^t) - \mathbf{f}^t\right\| < \hat{\tau}\bar{\delta}_2^2/\delta$. Since $\{U(\mathbf{x}^t)\}$ converges, it must be $\liminf_{\mathcal{T}\ni t\to\infty}\sum_{n=t+1}^{i_t-1}\gamma^{n+1} = 0$, which contradicts (46). Therefore, it must be $\limsup_{t\to\infty}\left\|\hat{\mathbf{x}}^t - \mathbf{x}^t\right\| = 0$ w.p.1.

Finally, let us prove that every limit point of the sequence $\{\mathbf{x}^t\}$ is a stationary solution of (1). Let $\mathbf{x}^\infty$ be the limit point of the convergent subsequence $\{\mathbf{x}^t\}_{t\in\mathcal{T}}$. Taking the limit of (35) over the index set $\mathcal{T}$, we have

$$\begin{aligned} &\lim_{\mathcal{T}\ni t\to\infty}\left\langle \mathbf{x}_i - \hat{\mathbf{x}}_i^t, \nabla_i \hat{f}_i\left(\hat{\mathbf{x}}_i^t; \mathbf{x}^t, \boldsymbol{\xi}^t\right)\right\rangle \\ &= \lim_{\mathcal{T}\ni t\to\infty}\big\langle \mathbf{x}_i - \hat{\mathbf{x}}_i^t, \\ &\quad \mathbf{f}_i^t + \tau_i\left(\hat{\mathbf{x}}_i^t - \mathbf{x}_i^t\right) \\ &\quad +\rho^t\textstyle\sum_{j\in\mathcal{C}_i^t}\left(\nabla_i f_j(\hat{\mathbf{x}}_i^t, \mathbf{x}_{-i}^t, \boldsymbol{\xi}^t) - \nabla_i f_j(\mathbf{x}_i^t, \mathbf{x}_{-i}^t, \boldsymbol{\xi}^t)\right)\big\rangle \\ &\quad = \left\langle \mathbf{x}_i - \mathbf{x}_i^\infty, \nabla U(\mathbf{x}_i^\infty)\right\rangle \geq 0,\ \forall\, \mathbf{x}_i \in \mathcal{X}_i, \end{aligned} \tag{52}$$

where the last equality follows from: i) $\lim_{t\to\infty}\left\|\nabla U(\mathbf{x}^t) - \mathbf{f}^t\right\| = 0$ [cf. Lemma 3]; ii) $\lim_{t\to\infty}\left\|\hat{\mathbf{x}}_i^t - \mathbf{x}^t\right\| = 0$; and iii) the following

$$\begin{aligned} &\left\|\rho^t\textstyle\sum_{j\in\mathcal{C}_i^t}(\nabla_i f_j(\hat{\mathbf{x}}_i^t, \mathbf{x}_{-i}^t, \boldsymbol{\xi}^t) - \nabla_i f_j(\mathbf{x}_i^t, \mathbf{x}_{-i}^t, \boldsymbol{\xi}^t))\right\| \\ &\leq C_x\rho^t\textstyle\sum_{j\in\mathcal{I}_f}L_{\nabla f_j(\boldsymbol{\xi}^t)} \underset{t\to\infty}{\longrightarrow} 0, \end{aligned} \tag{53}$$

where (53) follows from the Lipschitz continuity of $\nabla f_j(\mathbf{x}, \boldsymbol{\xi})$, the fact $\|\hat{\mathbf{x}}_i^t - \mathbf{x}_i^t\| \leq C_x$, and (11d).

Adding (52) over $i = 1, \ldots, I$, we get the desired first-order optimality condition: $\left\langle \mathbf{x} - \mathbf{x}^\infty, \nabla U(\mathbf{x}^\infty)\right\rangle \geq 0$, for all $\mathbf{x} \in \mathcal{X}$. Therefore $\mathbf{x}^\infty$ is a stationary point of (1). ■

B. *Proof of Lemma 2*

We prove only (27). Since (25) is a convex optimization problem and $\mathcal{Q}$ has a nonempty interior, strong duality holds for (25) [37]. The dual function of (25) is

$$d(\mu) = \max_{\mathbf{Y}\succeq\mathbf{0}}\{\rho\log\det(\mathbf{I}+\mathbf{Y}\mathbf{D}_1) + \left\langle \mathbf{Y}, \mathbf{Z} - \mu\mathbf{I}\right\rangle\} + \mu P, \tag{54}$$

where $\mu \in \{\mu : \mu \succeq 0, d(\mu) < +\infty\}$. Denote by $\mathbf{Y}^\star(\mu)$ the optimal solution of the maximization problem in (54), for any given feasible μ. It is easy to see that $d(\mu) = +\infty$ if $\mathbf{Z} - \mu\mathbf{I} \succeq \mathbf{0}$, so μ is feasible if and only if $\mathbf{Z} - \mu\mathbf{I} \prec \mathbf{0}$, i.e.,

$$\mu\begin{cases} \geq \underline{\mu} = [\lambda_{\max}(\mathbf{Z})]^+ = 0, & \text{if } \mathbf{Z} \prec \mathbf{0}, \\ > \underline{\mu} = [\lambda_{\max}(\mathbf{Z})]^+, & \text{otherwise}, \end{cases}$$

and $\mathbf{Y}^\star(\mu)$ is [29, Prop. 1]

$$\mathbf{Y}^\star(\mu) = \mathbf{V}(\mu)[\rho\mathbf{I} - \mathbf{D}(\mu)^{-1}]^+\mathbf{V}(\mu)^H,$$

where $(\mathbf{V}(\mu), \mathbf{\Sigma}(\mu))$ is the generalized eigenvalue decomposition of $(\mathbf{D}_1, -\mathbf{Z} + \mu\mathbf{I})$. Invoking [38, Cor. 28.1.1], the uniqueness of $\mathbf{Y}(\mathbf{Z})$ comes from the uniqueness of $\mathbf{Y}^\star(\mu)$ that was proved in [39].

Now we prove that $\mu^\star \leq \overline{\mu}$. First, note that $d(\mu) \geq \mu P$. Based on the eigenvalue decomposition $\mathbf{Z} = \mathbf{V_Z}\mathbf{\Sigma_Z}\mathbf{V_Z}^H$, the following inequalities hold:

$$\begin{aligned}\text{tr}((\mathbf{Z} - \mu\mathbf{I})^H\mathbf{X}) &= \text{tr}(\mathbf{V_Z}(\mathbf{\Sigma_Z} - \mu\mathbf{I})\mathbf{V_Z}^H\mathbf{X})\\ &\leq (\lambda_{\max}(\mathbf{\Sigma_Z}) - \mu)\text{tr}(\mathbf{X}),\end{aligned}$$

where $\lambda_{\max}(\mathbf{\Sigma_Z}) = \lambda_{\max}(\mathbf{Z})$. In other words, $d(\mu)$ is upper bounded by the optimal value of the following problem:

$$\max_{\mathbf{Y}\succeq\mathbf{0}} \quad \rho \log\det(\mathbf{I} + \mathbf{Y}\mathbf{D}_1) + (\lambda_{\max}(\mathbf{Z}) - \mu)\text{tr}(\mathbf{Y}) + \mu P. \tag{55}$$

When $\mu \geq \overline{\mu}$, it is not difficult to verify that the optimal variable of (55) is $\mathbf{0}$, and thus $\mathbf{Y}^\star(\mu) = \mathbf{0}$. We show $\mu^\star \leq \bar{\mu}$ by discussing two complementary cases: $\bar{\mu} = 0$ and $\bar{\mu} > 0$.

If $\bar{\mu} = 0$, $d(\bar{\mu}) = d(0) = \mu P = 0$. Since $\mathbf{Y}^\star(0) = \mathbf{0}$ and the primal value is also 0, there is no duality gap. From the definition of saddle point [37, Sec. 5.4], $\bar{\mu} = 0$ is a dual optimal variable.

If $\bar{\mu} > 0$, $d(\mu) \geq \mu P > 0$. Assume $\mu^\star > \overline{\mu}$. Then $\mathbf{Y}^\star(\mu^\star) = \mathbf{0}$ is the optimal variable in (25) and the optimal value of (25) is 0, but this would lead to a non-zero duality gap and thus contradict the optimality of $\mu^\star$. Therefore $\mu^\star \leq \overline{\mu}$.

References

[1] Y. Yang, G. Scutari, and D. P. Palomar, "Parallel stochastic decomposition algorithms for multi-agent systems," *2013 IEEE 14th Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, pp. 180–184, Jun. 2013.

[2] H. Robbins and S. Monro, "A Stochastic Approximation Method," *The Annals of Mathematical Statistics*, vol. 22, no. 3, pp. 400–407, Sep. 1951.

[3] H. J. Kushner and G. Yin, *Stochastic approximation and recursive algorithms and applications*, 2nd ed. Springer-Verlag, 2003, vol. 35.

[4] B. Polyak, *Introduction to optimization*. Optimization Software, 1987.

[5] D. P. Bertsekas and J. N. Tsitsiklis, "Gradient convergence in gradient methods with errors," *SIAM Journal on Optimization*, vol. 10, no. 3, pp. 627–642, 2000.

[6] J. N. Tsitsiklis, D. P. Bertsekas, and M. Athans, "Distributed asynchronous deterministic and stochastic gradient optimization algorithms," *IEEE Transactions on Automatic Control*, vol. 31, no. 9, pp. 803–812, Sep. 1986.

[7] Y. Ermoliev, "On the method of generalized stochastic gradients and quasi-Fejer sequences," *Cybernetics*, vol. 5, no. 2, pp. 208–220, 1972.

[8] F. Yousefian, A. Nedić, and U. V. Shanbhag, "On stochastic gradient and subgradient methods with adaptive steplength sequences," *Automatica*, vol. 48, no. 1, pp. 56–67, Jan. 2012.

[9] Y. Ermoliev and P. I. Verchenko, "A linearization method in limiting extremal problems," *Cybernetics*, vol. 12, no. 2, pp. 240–245, 1977.

[10] A. Ruszczyński, "Feasible direction methods for stochastic programming problems," *Mathematical Programming*, vol. 19, no. 1, pp. 220–229, Dec. 1980.

[11] ——, "A merit function approach to the subgradient method with averaging," *Optimization Methods and Software*, vol. 23, no. 1, pp. 161–172, Feb. 2008.

[12] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro, "Robust Stochastic Approximation Approach to Stochastic Programming," *SIAM Journal on Optimization*, vol. 19, no. 4, pp. 1574–1609, Jan. 2009.

[13] A. M. Gupal and L. G. Bazhenov, "Stochastic analog of the conjugant-gradient method," *Cybernetics*, vol. 8, no. 1, pp. 138–140, 1972.

[14] B. T. Polyak and A. B. Juditsky, "Acceleration of Stochastic Approximation by Averaging," *SIAM Journal on Control and Optimization*, vol. 30, no. 4, pp. 838–855, Jul. 1992.

[15] G. Scutari, F. Facchinei, P. Song, D. P. Palomar, and J.-S. Pang, "Decomposition by Partial Linearization: Parallel Optimization of Multi-Agent Systems," *IEEE Transactions on Signal Processing*, vol. 62, no. 3, pp. 641–656, Feb. 2014.

[16] F. Facchinei, G. Scutari, and S. Sagratella, "Parallel Selective Algorithms for Nonconvex Big Data Optimization," *IEEE Transactions on Signal Processing*, vol. 63, no. 7, pp. 1874–1889, Nov. 2015.

[17] A. Daneshmand, F. Facchinei, V. Kungurtsev, and G. Scutari, "Hybrid Random/Deterministic Parallel Algorithms for Convex and Nonconvex Big Data Optimization," vol. 63, no. 13, pp. 3914–3929, August 2015.

[18] M. Razaviyayn, M. Sanjabi, and Z.-Q. Luo, "A stochastic successive minimization method for nonsmooth nonconvex optimization with applications to transceiver design in wireless communication networks," Jun. 2013. [Online]. Available: http://arxiv.org/abs/1307.4457

[19] S. M. Robinson, "Analysis of Sample-Path Optimization," *Mathematics of Operations Research*, vol. 21, no. 3, pp. 513–528, Aug. 1996.

[20] J. Linderoth, A. Shapiro, and S. Wright, "The empirical behavior of sampling methods for stochastic programming," *Annals of Operations Research*, vol. 142, no. 1, pp. 215–241, Feb. 2006.

[21] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online Learning for Matrix Factorization and Sparse Coding," *The Journal of Machine Learning Research*, vol. 11, pp. 19–60, 2010.

[22] G. Scutari, F. Facchinei, J.-S. Pang, and D. P. Palomar, "Real and Complex Monotone Communication Games," *IEEE Transactions on Information Theory*, vol. 60, no. 7, pp. 4197–4231, Jul. 2014.

[23] A. Hjørungnes, *Complex-valued matrix derivatives with applications in signal processing and communications*. Cambridge: Cambridge University Press, 2011.

[24] J. Zhang, D. Zheng, and M. Chiang, "The Impact of Stochastic Noisy Feedback on Distributed Network Utility Maximization," *IEEE Transactions on Information Theory*, vol. 54, no. 2, pp. 645–665, Feb. 2008.

[25] M. Hong and A. Garcia, "Averaged Iterative Water-Filling Algorithm: Robustness and Convergence," *IEEE Transactions on Signal Processing*, vol. 59, no. 5, pp. 2448–2454, May 2011.

[26] P. Di Lorenzo, S. Barbarossa, and M. Omilipo, "Distributed Sum-Rate Maximization Over Finite Rate Coordination Links Affected by Random Failures," *IEEE Transactions on Signal Processing*, vol. 61, no. 3, pp. 648–660, Feb. 2013.

[27] J. Huang, R. A. Berry, and M. L. Honig, "Distributed interference compensation for wireless networks," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 5, pp. 1074–1084, 2006.

[28] C. Shi, R. A. Berry, and M. L. Honig, "Monotonic convergence of distributed interference pricing in wireless networks," in *2009 IEEE International Symposium on Information Theory*. IEEE, Jun. 2009, pp. 1619–1623.

[29] S.-J. Kim and G. B. Giannakis, "Optimal Resource Allocation for MIMO Ad Hoc Cognitive Radio Networks," *IEEE Transactions on Information Theory*, vol. 57, no. 5, pp. 3117–3131, May 2011.

[30] Z.-Q. Luo and S. Zhang, "Dynamic Spectrum Management: Complexity and Duality," *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, no. 1, pp. 57–73, Feb. 2008.

[31] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and distributed computation: Numerical methods*. Prentice Hall, 1989.

[32] S. Sundhar Ram, A. Nedić, and V. V. Veeravalli, "Incremental Stochastic Subgradient Algorithms for Convex Optimization," *SIAM Journal on Optimization*, vol. 20, no. 2, pp. 691–717, Jan. 2009.

[33] K. Srivastava and A. Nedić, "Distributed Asynchronous Constrained Stochastic Optimization," *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, no. 4, pp. 772–790, Aug. 2011.

[34] D. P. Bertsekas, *Nonlinear programming*. Athena Scientific, 1999.

[35] D. P. Bertsekas, A. Nedic, and A. E. Ozdaglar, *Convex Analysis and Optimization*. Athena Scientific, 2003.

[36] R. Zhang, M. Mohseni, and J. Cioffi, "Power Region for Fading Multiple-Access Channel with Multiple Antennas," in *2006 IEEE International Symposium on Information Theory*, no. 1, 2006, pp. 2709–2713.

[37] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge Univ Pr, 2004.

[38] R. T. Rockafellar, *Convex Analysis*, 2nd ed. Princeton, NJ: Princeton Univ. Press, 1970.

[39] W. Yu, W. Rhee, S. Boyd, and J. Cioffi, "Iterative Water-Filling for Gaussian Vector Multiple-Access Channels," *IEEE Transactions on Information Theory*, vol. 50, no. 1, pp. 145–152, Jan. 2004.

Yang Yang received the B.S. degree in School of Information Science and Engineering, Southeast University, Nanjing, China, in 2009, and the Ph.D. degree in Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology. From Nov. 2013 to Nov. 2015 he had been a postdoctoral research associate at the Communication Systems Group, Darmstadt University of Technology, Darmstadt, Germany. He joined Intel Deutschland GmbH as a research scientist in Dec. 2015.

His research interests are in distributed solution methods in convex optimization, nonlinear programming, and game theory, with applications in communication networks, signal processing, and financial engineering.

Daniel P. Palomar (S'99-M'03-SM'08-F'12) received the Electrical Engineering and Ph.D. degrees from the Technical University of Catalonia (UPC), Barcelona, Spain, in 1998 and 2003, respectively.

He is a Professor in the Department of Electronic and Computer Engineering at the Hong Kong University of Science and Technology (HKUST), Hong Kong, which he joined in 2006. Since 2013 he is a Fellow of the Institute for Advance Study (IAS) at HKUST. He had previously held several research appointments, namely, at King's College London (KCL), London, UK; Stanford University, Stanford, CA; Telecommunications Technological Center of Catalonia (CTTC), Barcelona, Spain; Royal Institute of Technology (KTH), Stockholm, Sweden; University of Rome "La Sapienza", Rome, Italy; and Princeton University, Princeton, NJ. His current research interests include applications of convex optimization theory, game theory, and variational inequality theory to financial systems, big data systems, and communication systems.

Dr. Palomar is an IEEE Fellow, a recipient of a 2004/06 Fulbright Research Fellowship, the 2004 and 2015 (co-author) Young Author Best Paper Awards by the IEEE Signal Processing Society, the 2002/03 best Ph.D. prize in Information Technologies and Communications by the Technical University of Catalonia (UPC), the 2002/03 Rosina Ribalta first prize for the Best Doctoral Thesis in Information Technologies and Communications by the Epson Foundation, and the 2004 prize for the best Doctoral Thesis in Advanced Mobile Communications by the Vodafone Foundation.

He is a Guest Editor of the IEEE Journal of Selected Topics in Signal Processing 2016 Special Issue on "Financial Signal Processing and Machine Learning for Electronic Trading" and has been Associate Editor of IEEE Transactions on Information Theory and of IEEE Transactions on Signal Processing, a Guest Editor of the IEEE Signal Processing Magazine 2010 Special Issue on "Convex Optimization for Signal Processing," the IEEE Journal on Selected Areas in Communications 2008 Special Issue on "Game Theory in Communication Systems," and the IEEE Journal on Selected Areas in Communications 2007 Special Issue on "Optimization of MIMO Transceivers for Realistic Communication Networks."

Gesualdo Scutari (S'05–M'06–SM'11) received the Electrical Engineering and Ph.D. degrees (both with Hons.) from the University of Rome "La Sapienza," Rome, Italy, in 2001 and 2005, respectively. He is an Associate Professor with the Department of Industrial Engineering, Purdue University, West Lafayette, IN, USA, and he is the Scientific Director for the area of Big-Data Analytics at the Cyber Center (Discovery Park) at Purdue University. He had previously held several research appointments, namely, at the University of California at Berkeley, Berkeley, CA, USA; Hong Kong University of Science and Technology, Hong Kong; University of Rome, "La Sapienza," Rome, Italy; University of Illinois at Urbana-Champaign, Urbana, IL, USA. His research interests include theoretical and algorithmic issues related to big data optimization, equilibrium programming, and their applications to signal processing, medical imaging, machine learning, and networking. Dr. Scutari is an Associate Editor of the IEEE TRANSACTIONS ON SIGNAL PROCESSING and he served as an Associate Editor of the IEEE SIGNAL PROCESSING LETTERS. He served on the IEEE Signal Processing Society Technical Committee on Signal Processing for Communications (SPCOM). He was the recipient of the 2006 Best Student Paper Award at the International Conference on Acoustics, Speech and Signal Processing (ICASSP) 2006, the 2013 NSF Faculty Early Career Development (CAREER) Award, the 2013 UB Young Investigator Award, the 2015 AnnaMaria Molteni Award for Mathematics and Physics from the Italian Scientists and Scholars in North America Foundation (ISSNAF), and the 2015 IEEE Signal Processing Society Young Author Best Paper Award.

Marius Pesavento (M'00) received the Dipl.-Ing. and M.Eng. degrees from Ruhr-UniversitÃd't Bochum, Germany, and McMaster University, Hamilton, ON, Canada, in 1999 and 2000, respectively, and the Dr.-Ing. degree in electrical engineering from Ruhr-UniversitÃd't Bochum in 2005. Between 2005 and 2007, he was a Research Engineer at FAG Industrial Services GmbH, Aachen, Germany. From 2007 to 2009, he was the Director of the Signal Processing Section at mimoOn GmbH, Duisburg, Germany. In 2010, he became an Assistant Professor for Robust Signal Processing and a Full Professor for Communication Systems in 2013 at the Department of Electrical Engineering and Information Technology, Darmstadt University of Technology, Darmstadt, Germany. His research interests are in the area of robust signal processing and adaptive beamforming, high-resolution sensor array processing, multiantenna and multiuser communication systems, distributed, sparse and mixed-integer optimization techniques for signal processing and communications, statistical signal processing, spectral analysis, parameter estimation. Dr. Pesavento was a recipient of the 2003 ITG/VDE Best Paper Award, the 2005 Young Author Best Paper Award of the IEEE TRANSACTIONS ON SIGNAL PROCESSING, and the 2010 Best Paper Award of the CROWNCOM conference. He is a member of the Editorial board of the EURASIP Signal Processing Journal, an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING. He currently is serving the 2nd term as a member of the Sensor Array and Multichannel (SAM) Technical Committee of the IEEE Signal Processing Society (SPS).