

Parallel Stochastic Decomposition Algorithms for Multi-Agent Systems

Yang Yang[†], Gesualdo Scutari[‡], Daniel P. Palomar[†]

[†] Dept. of Electronic and Computer Eng., The Hong Kong Univ. of Sci. & Tech., Email: {yangyang.palomar}@ust.hk

[‡] Dept. of Electrical Eng., State Univ. of New York (SUNY), Buffalo, NY 14260, USA, Email: gesualdo@buffalo.edu

Abstract—In a stochastic optimization problem, the objective function is given in the form of the expectation with respect to some random variables. In many applications, this expectation cannot be computed accurately (e.g., because the statistics of the random variables are unknown). The common approach followed in the literature to deal with this issue is using stochastic gradient schemes, which however suffer from slow convergence. In this paper, we propose for the first time a class of provably convergent *Jacobi best-response* algorithms for general *nonconvex* stochastic sum-utility optimization problems, which arise naturally in the design of wireless multi-user interfering systems. The proposed novel decomposition enables all users to update their optimization variables *in parallel* by solving a sequence of strongly convex subproblems, one for each user. Finally, we customize our algorithms to solve the stochastic sum rate maximization problem over MIMO interference channels and multiple access channels. Numerical results show that our algorithms are much faster than state-of-the-art stochastic gradient schemes.

Index Terms—Multi-agent systems, parallel optimization, stochastic approximation.

I. INTRODUCTION

Wireless networks are composed of users that may have different objectives and generate interference, when no multiplexing scheme is imposed to regulate the transmissions; examples are peer-to-peer networks, cognitive radio systems, and ad-hoc networks. A usual and convenient way of designing such multi-user systems is by optimizing the “social function”, i.e., the (weighted) sum of the users’ objective functions.

This formulation however requires the knowledge of the system parameters, which in practice are difficult to estimate when they are changing rapidly or imperfect due to estimation errors. To enhance the robustness of the system against estimation errors and/or improve the performance under rapidly-changing parameters, we focus on optimizing the long-term performance of the system, measured as the expected value of the social function (the expectation is taken with respect to the random system parameters). More specifically, we address the frequent and difficult case in which (the expected value of) the social function is *nonconvex* and the expectation cannot be computed. Such a system design naturally falls into the class of stochastic optimization [1]–[3].

Gradient methods for *unconstrained* stochastic nonconvex optimization problems have been studied in [2], [4], [5], where

almost sure convergence to stationary points has been established. The *constrained* case is more involved; for instance, the convergence analysis developed for unconstrained gradient methods cannot be successfully applied. Stochastic gradient projection methods (for constrained optimization problems) have been proposed in [2], [6], [7], whose convergence however requires the convexity of the objective function.

To deal with nonconvexity, an additional step in the variable updating rule called *gradient averaging* seems to be essential; indeed, stochastic conditional gradient methods for nonconvex constrained problems hinge on this idea [2], [6], [8]. Roughly speaking, gradient averaging is needed to guarantee that the estimated sample gradient eventually resembles the real (but unavailable) gradient of the objective function [9], [10]. Despite the capability to deal with nonconvexity, stochastic conditional gradient methods usually converge slowly in practice.

A nonconvex and very general cross-layer optimization for wireless networks has been addressed in [11], where dual stochastic gradient projection methods are proposed. To implement such algorithms one needs to compute the global optimal solution of the Lagrangian minimization, which is however a nonconvex optimization problem. Despite its elegance, the approach of [11] is thus applicable only to small-scale problems, where a grid search over the primal variables is computationally affordable.

In this paper, we propose for the first time a novel stochastic *best-response*-based decomposition method that: i) converges almost surely to stationary points of a large class of (stochastic) nonconvex social problems, encompassing most sum-utility functions of practical interest; ii) decomposes well across the users, resulting in the parallel solution of (strongly) convex subproblems, one for each user; iii) exploits the structure of the sum-utility function (e.g., partial convexity, if any) better than stochastic conditional gradient methods; and iv) contains as special cases stochastic conditional gradient methods. Moreover, it can be easily particularized to well-known applications, such as sum-rate maximization problems over MIMO Interference Channels (ICs) and Multiple Access Channels (MACs), giving rise in a unified way to distributed simultaneous algorithms that outperform existing (gradient-based) methods both theoretically and numerically. The proposed decomposition technique hinges on successive convex approximation methods, and it is a generalization of [12] to stochastic optimization problems.

The work of Yang and Palomar is supported by Hong Kong RGC 617312. The work of Scutari is supported by NSF No. CNS-1218717 and NSF CAREER No. ECCS-1254739.

II. PROBLEM FORMULATION

We consider the design of a multi-user system composed of I coupled users, where each user i has a cost function $f_i(\mathbf{x}_i, \mathbf{x}_{-i}, \boldsymbol{\xi})$ that depends on his own strategy vector \mathbf{x}_i , belonging to the feasible set \mathcal{X}_i , the variables of the other users denoted by $\mathbf{x}_{-i} \triangleq (\mathbf{x}_j)_{j=1, j \neq i}^I$, and a random vector $\boldsymbol{\xi}$. The joint strategy set of all users is denoted by $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_I$. We also define $\mathcal{X}_{-i} = \mathcal{X}_1 \times \dots \times \mathcal{X}_{i-1} \times \mathcal{X}_{i+1} \times \dots \times \mathcal{X}_I$.

The stochastic social optimization problem is formulated as:

$$\begin{aligned} & \underset{\mathbf{x} \triangleq \{\mathbf{x}_i\}}{\text{minimize}} && U(\mathbf{x}) \triangleq \mathbb{E} \left[\sum_{i=1}^I f_i(\mathbf{x}_i, \mathbf{x}_{-i}, \boldsymbol{\xi}) \right] \\ & \text{subject to} && \mathbf{x}_i \in \mathcal{X}_i, \quad i = 1, \dots, I, \end{aligned} \quad (1)$$

where the expectation is taken with respect to $\boldsymbol{\xi}$.

Assumptions: We make the following blanket assumptions:

- Each \mathcal{X}_i is closed and convex;
- Each $f_i(\mathbf{x}_i, \mathbf{x}_{-i}, \boldsymbol{\xi})$ is continuously differentiable on \mathcal{X} , and uniformly convex in $\mathbf{x}_i \in \mathcal{X}_i$ for any given $\mathbf{x}_{-i} \in \mathcal{X}_{-i}$ and $\boldsymbol{\xi}$;
- Function $U(\mathbf{x})$ is coercive on \mathcal{X} [i.e., $U(\mathbf{x}) \rightarrow \infty$ as $\|\mathbf{x}\| \rightarrow \infty$ while $\mathbf{x} \in \mathcal{X}$], and has a Lipschitz continuous gradient with constant $L_{\nabla U}$.

These assumptions are quite standard and are satisfied by a large class of problems [12]. For instance, the coercivity of $U(\mathbf{x})$ in (c) guarantees that (1) has a solution even when \mathcal{X} is unbounded (this assumption is trivially satisfied when \mathcal{X} is bounded). Note that $U(\mathbf{x})$ is not assumed to be jointly convex in \mathbf{x} . Some instances of (1) satisfying the above assumptions are briefly listed here.

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

$$\begin{aligned} & \underset{\mathbf{Q} \triangleq \{\mathbf{Q}_i\}}{\text{maximize}} && \mathbb{E} \left[\sum_{i=1}^I \log \det (\mathbf{I} + \mathbf{H}_{ii} \mathbf{Q}_i \mathbf{H}_{ii}^H \mathbf{R}_i^{-1}) \right] \\ & \text{subject to} && \mathbf{Q}_i \succeq \mathbf{0}, \text{Tr}(\mathbf{Q}_i) \leq P_i, \quad i = 1, \dots, I, \end{aligned} \quad (2)$$

where $\mathbf{R}_i(\mathbf{Q}_{-i}, \mathbf{H}) \triangleq \mathbf{R}_{n_i} + \sum_{j \neq i} \mathbf{H}_{ij} \mathbf{Q}_j \mathbf{H}_{ij}^H$ is the covariance matrix of the multi-user interference plus the thermal noise \mathbf{R}_{n_i} (assumed to be full-rank), and the expectation in (2) is taken over the channels $\mathbf{H} \triangleq (\mathbf{H}_{ij})_{i,j=1}^I$.

Example #2: Another application of interest is the maximization of the ergodic sum-rate over MIMO MACs:

$$\begin{aligned} & \underset{\mathbf{Q}}{\text{maximize}} && \mathbb{E} \left[\log \det \left(\mathbf{R}_n + \sum_{i=1}^I \mathbf{H}_i \mathbf{Q}_i \mathbf{H}_i^H \right) \right] \\ & \text{subject to} && \mathbf{Q}_i \succeq \mathbf{0}, \text{Tr}(\mathbf{Q}_i) \leq P_i, \quad i = 1, \dots, I. \end{aligned} \quad (3)$$

This is a special case of (1) where the utility function is jointly concave in $\{\mathbf{Q}_i\}_{i=1}^I$ for any given $\mathbf{H} \triangleq (\mathbf{H}_i)_{i=1}^I$.

Example #3: Consider the least square estimation problem with sparsity constraint [13]: given $\lambda > 0$,

$$\underset{\mathbf{x}}{\text{minimize}} \quad \mathbb{E} \left[\left\| \boldsymbol{\theta} - \sum_{i=1}^I \mathbf{A}_i \mathbf{x}_i \right\|_2^2 \right] + \lambda \cdot \|\mathbf{x}\|_1, \quad (4a)$$

where $\|\bullet\|_1$ is the ℓ_1 norm, and the expectation is over random matrices $\mathbf{A} = [\mathbf{A}_1 \dots \mathbf{A}_I]$ and random vectors $\boldsymbol{\theta}$. It is not difficult to see that (4a) can be equivalently rewritten as

$$\begin{aligned} & \underset{\mathbf{x}, \mathbf{z}}{\text{minimize}} && \mathbb{E} \left[\left\| \boldsymbol{\theta} - \sum_{i=1}^I \mathbf{A}_i \mathbf{x}_i \right\|_2^2 \right] + \lambda \cdot \mathbf{1}^T \mathbf{z} \\ & \text{subject to} && -\mathbf{z}_i \leq \mathbf{x}_i \leq \mathbf{z}_i, \quad i = 1, \dots, I, \end{aligned} \quad (4b)$$

which is an instance of (1) satisfying Assumptions (a)-(c).

In addition to the above examples, the framework we are going to introduce can be successfully used also to study convergence and robustness of distributed algorithms solving a deterministic social problem in the form of (1), where $\boldsymbol{\xi}$ is deterministic quantity but only a noisy estimate of (some) system parameters is available; some examples are [14]–[16]. We discuss this application in Section IV-C.

Since the class of problems (1) is in general nonconvex (possibly NP hard), the focus of this paper is to design *distributed* solution methods for computing stationary solutions (possibly local minima) of (1). Our major goal is to devise *simultaneous best-response* schemes that converge even when the expected value in (1) cannot be computed accurately.

III. A NOVEL PARALLEL STOCHASTIC DECOMPOSITION

The social problem (1) faces two main issues: i) the non-convexity 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 linearizing the nonconvex part and replacing the expected value with an incremental estimate of it. More formally, at iteration $t+1$, user i solves the following problem: given \mathbf{x}_{-i}^t and $\boldsymbol{\xi}^t$,

$$\begin{aligned} \widehat{\mathbf{x}}_i(\mathbf{x}^t, \boldsymbol{\xi}^t) = & \arg \min_{\mathbf{x}_i \in \mathcal{X}_i} \left\{ \epsilon^t f_i(\mathbf{x}_i, \mathbf{x}_{-i}^t, \boldsymbol{\xi}^t) \right. \\ & + \epsilon^t \langle \mathbf{x}_i - \mathbf{x}_i^t, \boldsymbol{\pi}_i(\mathbf{x}^t, \boldsymbol{\xi}^t) \rangle \\ & + (1 - \epsilon^t) \langle \mathbf{x}_i - \mathbf{x}_i^t, \mathbf{f}_i^{t-1} \rangle \\ & \left. + \frac{\tau_i}{2} \|\mathbf{x}_i - \mathbf{x}_i^t\|_2^2 \right\}, \end{aligned} \quad (5)$$

where $\langle \mathbf{x}, \mathbf{y} \rangle \triangleq \mathbf{x}^T \mathbf{y}$, the pricing vector $\boldsymbol{\pi}_i(\mathbf{x}, \boldsymbol{\xi})$ is defined as

$$\boldsymbol{\pi}_i(\mathbf{x}, \boldsymbol{\xi}) \triangleq \sum_{j \neq i} \nabla_i f_j(\mathbf{x}, \boldsymbol{\xi}),$$

and \mathbf{f}_i^t is updated recursively according to

$$\mathbf{f}_i^t = (1 - \epsilon^t) \mathbf{f}_i^{t-1} + \epsilon^t (\boldsymbol{\pi}_i(\mathbf{x}^t, \boldsymbol{\xi}^t) + \nabla_i f_i(\mathbf{x}^t, \boldsymbol{\xi}^t)), \quad (6)$$

where $\epsilon^t \in (0, 1]$ is a sequence to be properly chosen ($\epsilon^0 = 1$).

Note that \mathbf{x}^t is a random variable depending on \mathcal{F}^t , the past history of the algorithm up to iteration t :

$$\mathcal{F}^t \triangleq \{\mathbf{x}^0, \dots, \mathbf{x}^t, \boldsymbol{\xi}^0, \dots, \boldsymbol{\xi}^{t-1}, \gamma^0, \dots, \gamma^t, \epsilon^0, \dots, \epsilon^t\}; \quad (7)$$

therefore $\widehat{\mathbf{x}}(\mathbf{x}^t, \boldsymbol{\xi}^t)$ depends on \mathcal{F}^t as well (we omit this dependence for notational simplicity).

In (5), the first term of the objective function is the convex part of the (instantaneous) social function; the second term comes from the linearization of the nonconvex part, and the third term represents the incremental estimate of the gradient of the expected value $U(\mathbf{x}^t)$. Such an estimate is expected to be more and more accurate as t increases, provided that the sequence ϵ^t is properly chosen (this statement is made rigorous in Theorem 1). The last term is the proximal regularization whose numerical benefits are well-understood [17].

Given (5), we define the best-response mapping as

$$\mathcal{X} \ni \mathbf{y} \mapsto \widehat{\mathbf{x}}(\mathbf{y}, \boldsymbol{\xi}) \triangleq (\widehat{\mathbf{x}}_i(\mathbf{y}, \boldsymbol{\xi}))_{i=1}^I. \quad (8)$$

Note that $\widehat{\mathbf{x}}(\bullet, \boldsymbol{\xi})$ is well-defined for any given $\boldsymbol{\xi}$ because the objective function in (5) is strongly convex with constant τ_{\min} given by

$$\tau_{\min} \triangleq \min_{i=1, \dots, I} \{\tau_i + c_i\}, \quad (9)$$

where c_i is a nonnegative constant satisfying the following: for any given $\mathbf{x}_{-i} \in \mathcal{X}_{-i}$ and $\boldsymbol{\xi}$, $(\mathbf{x}_i^1 - \mathbf{x}_i^2)^T (\nabla_i f_i(\mathbf{x}_i^1, \mathbf{x}_{-i}, \boldsymbol{\xi}) - \nabla_i f_i(\mathbf{x}_i^2, \mathbf{x}_{-i}, \boldsymbol{\xi})) \geq c_i \|\mathbf{x}_i^1 - \mathbf{x}_i^2\|_2^2$, $\forall \mathbf{x}_i^1, \mathbf{x}_i^2 \in \mathcal{X}_i$. The existence of such a c_i follows from Assumption (b).

Our decomposition scheme is formally described in Algorithm 1, and its convergence properties are stated in Theorem 1. The analysis is carried out under the standard assumption on the boundedness of the gradient error [14], [18]:

Assumption (d):

$$\mathbb{E}[\|\nabla U(\mathbf{x}^t) - \sum_{i=1}^I \nabla f_i(\mathbf{x}^t, \boldsymbol{\xi}^t)\|_2^2 | \mathcal{F}^t] < \infty, \quad t = 0, 1, \dots$$

The above assumption is readily satisfied if, e.g., $\boldsymbol{\xi}$ is bounded. We omit the proof of Theorem 1 because of space limit, see [19].

Algorithm 1: Stochastic parallel decomposition algorithm

Data: $\boldsymbol{\tau} \triangleq (\tau_i)_{i=1}^I \geq \mathbf{0}$, $\{\gamma^t\}$, $\{\epsilon^t\}$, $\mathbf{x}^0 \in \mathcal{X}$; set $t = 0$.

(S.1): If \mathbf{x}^t satisfies a suitable termination criterion: STOP.

(S.2): For all $i = 1, \dots, I$, compute $\widehat{\mathbf{x}}_i(\mathbf{x}^t, \boldsymbol{\xi}^t)$ [cf. (5)].

(S.3): Update \mathbf{x}^{t+1} as $\mathbf{x}^{t+1} = (1 - \gamma^t) \mathbf{x}^t + \gamma^t \widehat{\mathbf{x}}(\mathbf{x}^t, \boldsymbol{\xi}^t)$.

(S.4): Update \mathbf{F}^t according to (6).

(S.5): $t \leftarrow t + 1$, and go to (S.1).

Theorem 1. Given problem (1) under Assumptions (a)-(d), suppose that $\tau_{\min} > 0$ and the stepsizes $\{\gamma^t\}$ and $\{\epsilon^t\}$ satisfy

$$\gamma^t \rightarrow 0, \quad \epsilon^t \rightarrow 0, \quad \lim_{t \rightarrow \infty} \gamma^t / \epsilon^t = 0,$$

$$\sum_{t=0}^{\infty} \gamma^t = \infty, \quad \sum_{t=0}^{\infty} \epsilon^t = \infty, \quad \sum_{t=0}^{\infty} (\epsilon^t)^2 < \infty. \quad (10)$$

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

On Algorithm 1: To our best knowledge, Algorithm 1 is the first Jacobi best-response scheme for nonconvex stochastic social problems: all the users update *in parallel* their strategies (possibly with a memory) solving a sequence of *decoupled* (strongly) convex subproblems (5). It is expected to perform better than classical stochastic gradient-based schemes at the cost of no extra signaling, because the convexity of the objective function, if any, is better exploited. Moreover, it is guaranteed to converge under the weakest assumptions available in literature while offering some flexibility in the choice of the free parameters [cf. Theorem 1].

Diminishing stepsize rules: In order to have convergence, a diminishing stepsize rule satisfying (10) is necessary. An instance of (10) is the following (see [19] for other rules):

$$\gamma^t = \frac{1}{(t+1)^\alpha}, \quad \epsilon^t = \frac{1}{(t+1)^\beta}, \quad 1 \geq \alpha > \beta > 0.5.$$

with $0.5 < \beta < \alpha \leq 1$. Roughly speaking, (10) says that the stepsizes γ^t and ϵ^t , while diminishing (with γ^t decreasing faster than ϵ^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 errors, see [4], [20].

Implementation issues: In order to compute the best-response, each user needs to estimate the pricing vector $\boldsymbol{\pi}_i(\mathbf{x}^t, \boldsymbol{\xi}^t)$. This requires in general some signaling exchange among (nearby) users, which is in the same spirit of that of interference-pricing algorithms proposed in the literature for the (deterministic) maximization of the sum-rate over ICs [12], [21]. Note 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, as in [21]. Once the new pricing vector $\boldsymbol{\pi}_i(\mathbf{x}^t, \boldsymbol{\xi}^t)$ is available, the recursive update (6) for the “incremental” gradient does not need any further signaling, but just a local accumulation register to keep track of the last iterate \mathbf{F}_i^{t-1} .

Gradient averaging: A natural question at this point is whether the averaging on the gradient in the form of (6) is essential to guarantee convergence. The answer is positive, meaning that we cannot use in (5) just the current sample estimate of the gradient; this is briefly explained next. In the absence of gradient averaging (i.e., $\epsilon^t = 1$ for all t), the best-response in (5) would reduce to

$$\widehat{\mathbf{x}}_i(\mathbf{x}^t, \boldsymbol{\xi}^t) = \arg \min_{\mathbf{x}_i \in \mathcal{X}_i} \left\{ f_i(\mathbf{x}_i, \mathbf{x}_{-i}^t, \boldsymbol{\xi}^t) + \langle \mathbf{x}_i - \mathbf{x}_i^t, \boldsymbol{\pi}_i(\mathbf{x}^t, \boldsymbol{\xi}^t) \rangle + \frac{\tau_i}{2} \|\mathbf{x}_i - \mathbf{x}_i^t\|_2^2 \right\}, \quad (11)$$

which would lead to the following property for the best response $\widehat{\mathbf{x}}(\mathbf{x}^t, \boldsymbol{\xi}^t)$ [19]:

$$\langle \widehat{\mathbf{x}}(\mathbf{x}^t, \boldsymbol{\xi}^t) - \mathbf{x}^t, \sum_{i=1}^I \nabla f_i(\mathbf{x}^t, \boldsymbol{\xi}^t) \rangle \leq -\tau_{\min} \|\widehat{\mathbf{x}}(\boldsymbol{\xi}^t) - \mathbf{x}^t\|_2^2.$$

Differently from gradient-based updating directions, $\widehat{\mathbf{x}}(\mathbf{x}^t, \boldsymbol{\xi}^t) - \mathbf{x}^t$ with $\widehat{\mathbf{x}}(\mathbf{x}^t, \boldsymbol{\xi}^t)$ given by (11) is a nonlinear function of $\boldsymbol{\xi}^t$. This implies that $\mathbb{E}[\widehat{\mathbf{x}}(\mathbf{x}^t, \boldsymbol{\xi}^t) - \mathbf{x}^t | \mathcal{F}^t]$ is not a descent direction of $U(\mathbf{x})$ at $\mathbf{x} = \mathbf{x}^t$ and thus convergence is not guaranteed anymore, as instead happens for classical unconstrained stochastic gradient methods [4], [5].

Some special cases: Interestingly, the proposed decomposition technique reduces to the classical stochastic conditional gradient schemes (possibly with a proximal term) when $f_i(\mathbf{x}_i, \mathbf{x}_{-i}, \boldsymbol{\xi})$ is linearized with respect to \mathbf{x}_i as well [2]:

$$\widehat{\mathbf{x}}_i(\mathbf{x}^t, \boldsymbol{\xi}^t) = \arg \min_{\mathbf{x}_i \in \mathcal{X}_i} \left\{ \epsilon^t \langle \mathbf{x}_i - \mathbf{x}_i^t, \sum_{j=1}^I \nabla_i f_j(\mathbf{x}^t, \boldsymbol{\xi}^t) \rangle + (1 - \epsilon^t) \langle \mathbf{x}_i - \mathbf{x}_i^t, \mathbf{f}_i^{t-1} \rangle + \frac{\tau_i}{2} \|\mathbf{x}_i - \mathbf{x}_i^t\|_2^2 \right\}. \quad (12)$$

It turns out that the convergence conditions of Algorithm 1 contain as special cases those of classical stochastic conditional gradient. But, Algorithm 1 is able to exploit the convexity of the users’ objective functions, whereas gradient stochastic schemes “destroy” the “nice” structure of objective functions by linearizing the whole social function. Therefore

our algorithm is expected to be faster than its classical counterpart. This is in fact confirmed by the numerical results, see Section V.

Another interesting special case is when the objective function of all users are identical, i.e., $f_i(\mathbf{x}, \boldsymbol{\xi}) = f(\mathbf{x}, \boldsymbol{\xi})$ for all i , and $f(\mathbf{x}, \boldsymbol{\xi})$ is jointly convex on \mathcal{X} for any fixed $\boldsymbol{\xi}$, resulting in a convex social function $U(\mathbf{x}) = \mathbb{E}[f(\mathbf{x}, \boldsymbol{\xi})]$. To solve such a problem in a distributed and efficient way, we can use the proposed decomposition technique, which leads to the following users' best-response solution:

$$\begin{aligned} \hat{\mathbf{x}}_i(\mathbf{x}^t, \boldsymbol{\xi}^t) \triangleq & \arg \min_{\mathbf{x}_i \in \mathcal{X}_i} \left\{ \epsilon^t f(\mathbf{x}_i, \mathbf{x}_{-i}^t, \boldsymbol{\xi}^t) \right. \\ & + (1 - \epsilon^t) \langle \mathbf{x}_i - \mathbf{x}_i^t, \mathbf{f}_i^{t-1} \rangle \\ & \left. + \frac{\tau_i}{2} \|\mathbf{x}_i - \mathbf{x}_i^t\|_2^2 \right\}, \end{aligned} \quad (13)$$

where \mathbf{f}_i^t is updated according to $\mathbf{f}_i^t = (1 - \epsilon^t) \mathbf{f}_i^{t-1} + \epsilon^t \nabla_{\mathbf{x}_i} f(\mathbf{x}^t, \boldsymbol{\xi}^t)$. Convergence conditions are still given by Theorem 1. An interesting application of this algorithm is the maximization of the ergodic sum-rate over MIMO MACs in (3), resulting in the first *simultaneous* best-response algorithm in the literature.

IV. APPLICATIONS

A. Sum-rate maximization over MIMO ICs

We customize here Algorithm 1 to solve the sum-rate maximization problem over MIMO ICs (2). In particular, the best-response of each user i becomes [cf. (5)]

$$\begin{aligned} \hat{\mathbf{Q}}_i(\mathbf{Q}^t, \mathbf{H}^t) = & \arg \max_{\mathbf{Q}_i \in \mathcal{Q}_i} \left\{ \epsilon^t \log \det (\mathbf{R}_i^t + \mathbf{H}_{ii}^t \mathbf{Q}_i (\mathbf{H}_{ii}^t)^H) \right. \\ & + \epsilon^t \langle \mathbf{Q}_i - \mathbf{Q}_i^t, \mathbf{\Pi}_i^t \rangle + (1 - \epsilon^t) \langle \mathbf{Q}_i - \mathbf{Q}_i^t, \mathbf{F}_i^{t-1} \rangle \\ & \left. - \frac{\tau_i}{2} \|\mathbf{Q}_i - \mathbf{Q}_i^t\|_F^2 \right\}, \end{aligned} \quad (14)$$

where $\mathcal{Q}_i \triangleq \{\mathbf{Q}_i : \mathbf{Q}_i \succeq \mathbf{0}, \text{Tr}(\mathbf{Q}_i) \leq P_i\}$, \mathbf{R}_i^t is the short notation for $\mathbf{R}_i(\mathbf{Q}_{-i}^t, \mathbf{H}^t)$, $\langle \mathbf{A}, \mathbf{B} \rangle \triangleq \text{tr}(\mathbf{A}^H \mathbf{B})$, and $\|\bullet\|_F$ denotes the Frobenius norm. In (14), $\mathbf{\Pi}_i(\mathbf{Q}, \mathbf{H})$ is given by

$$\begin{aligned} \mathbf{\Pi}_i(\mathbf{Q}, \mathbf{H}) &= \sum_{j \neq i} \mathbf{H}_{ji}^H \tilde{\mathbf{R}}_j(\mathbf{Q}_{-j}, \mathbf{H}) \mathbf{H}_{ji} \\ \tilde{\mathbf{R}}_j(\mathbf{Q}_{-j}, \mathbf{H}) &\triangleq (\mathbf{R}_j(\mathbf{Q}_{-j}, \mathbf{H}) + \mathbf{H}_{jj} \mathbf{Q}_j \mathbf{H}_{jj}^H)^{-1} \\ &\quad - \mathbf{R}_j(\mathbf{Q}_{-j}, \mathbf{H})^{-1}, \end{aligned}$$

and \mathbf{F}_i^t is updated according to (6):

$$\begin{aligned} \mathbf{F}_i^t &= (1 - \epsilon^t) \mathbf{F}_i^{t-1} + \epsilon^t \mathbf{\Pi}_i(\mathbf{Q}^t, \mathbf{H}^t) \\ &\quad + \epsilon^t (\mathbf{H}_{ii}^t)^H (\mathbf{R}_i^t + \mathbf{H}_{ii}^t \mathbf{Q}_i^t (\mathbf{H}_{ii}^t)^H)^{-1} \mathbf{H}_{ii}^t. \end{aligned}$$

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 properly [cf. Theorem 1]. Note that if \mathbf{H}_{ii} is full-column rank, one can also choose $\tau_i = 0$.

The proposed algorithm is fairly distributed: once the pricing matrix $\mathbf{\Pi}_i$ is given, to compute his best-response, each user only needs to locally estimate the covariance matrix of the interference plus noise. The estimation of the pricing matrix $\mathbf{\Pi}_i$ requires however some signaling among nearby receivers.

Quite interestingly, the pricing expression and thus the resulting signaling overhead necessary to compute it coincide with [21], where a sequential algorithm is proposed for the deterministic maximization of the sum-rate over MIMO ICs.

B. Sum-rate maximization over MIMO MACs

In this subsection, we customize Algorithm 1 to solve the sum-rate maximization problem over MIMO MACs [cf. (3)]. In particular, the best response of each user i becomes:

$$\begin{aligned} \hat{\mathbf{Q}}_i(\mathbf{Q}^t, \mathbf{H}^t) = & \arg \max_{\mathbf{Q}_i \in \mathcal{Q}_i} \left\{ \epsilon^t \log \det (\mathbf{R}_i^t + \mathbf{H}_i^t \mathbf{Q}_i (\mathbf{H}_i^t)^H) \right. \\ & + (1 - \epsilon^t) \langle \mathbf{Q}_i - \mathbf{Q}_i^t, \mathbf{F}_i^{t-1} \rangle \\ & \left. - \frac{\tau_i}{2} \|\mathbf{Q}_i - \mathbf{Q}_i^t\|_F^2 \right\}, \end{aligned} \quad (15)$$

and \mathbf{F}_i^t is updated according to

$$\mathbf{F}_i^t = (1 - \epsilon^t) \mathbf{F}_i^{t-1} + \epsilon^t (\mathbf{H}_i^t)^H (\mathbf{R}_i^t + \mathbf{H}_i^t \mathbf{Q}_i^t (\mathbf{H}_i^t)^H)^{-1} \mathbf{H}_i^t.$$

Algorithm 1 based on the best-response (15) has the same desirable properties of those discussed for the sum-rate maximization over MIMO ICs.

C. Distributed deterministic algorithms with errors

The developed framework can be useful in the context of deterministic optimization as well to study the robustness of best-response based-algorithms to noisy estimates of system parameters. Consider the following optimization problem:

$$\begin{aligned} \underset{\mathbf{x}}{\text{minimize}} \quad & \sum_{i=1}^I f_i(\mathbf{x}_i, \mathbf{x}_{-i}) \\ \text{subject to} \quad & \mathbf{x}_i \in \mathcal{X}_i, \quad i = 1, \dots, I. \end{aligned} \quad (16)$$

The deterministic counterpart of Algorithm 1 is [12]:

$$\begin{aligned} \hat{\mathbf{x}}_i(\mathbf{x}^t) = & \arg \min_{\mathbf{x}_i \in \mathcal{X}_i} \left\{ f_i(\mathbf{x}_i, \mathbf{x}_{-i}^t) \right. \\ & \left. + \langle \mathbf{x}_i - \mathbf{x}_i^t, \boldsymbol{\pi}_i(\mathbf{x}^t) \rangle + \frac{\tau_i}{2} \|\mathbf{x}_i - \mathbf{x}_i^t\|_2^2 \right\}, \quad (17) \\ \mathbf{x}_i^{t+1} = & \mathbf{x}_i^t + \gamma^t (\hat{\mathbf{x}}_i(\mathbf{x}^t) - \mathbf{x}_i^t), \quad i = 1, \dots, I. \end{aligned}$$

where $\boldsymbol{\pi}_i(\mathbf{x}) = \sum_{j \neq i} \nabla_{\mathbf{x}_i} f_j(\mathbf{x})$. In many applications (see, e.g., [14]–[16]), only a noisy estimate $\tilde{\boldsymbol{\pi}}_i(\mathbf{x})$ is available due to error-prone measurement schemes; iterate (17) then becomes

$$\begin{aligned} \tilde{\mathbf{x}}_i(\mathbf{x}^t) = & \arg \min_{\mathbf{x}_i \in \mathcal{X}_i} \left\{ f_i(\mathbf{x}_i, \mathbf{x}_{-i}^t) \right. \\ & \left. + \langle \mathbf{x}_i - \mathbf{x}_i^t, \tilde{\boldsymbol{\pi}}_i(\mathbf{x}^t) \rangle + \frac{\tau_i}{2} \|\mathbf{x}_i - \mathbf{x}_i^t\|_2^2 \right\}. \end{aligned} \quad (18)$$

Convergence of this scheme is, however, no longer guaranteed. Capitalizing on the proposed framework, we can readily deal with estimation errors while guaranteeing convergence. In particular, it is sufficient to modify the iterate (18) as:

$$\begin{aligned} \tilde{\mathbf{x}}_i(\mathbf{x}^t) = & \arg \min_{\mathbf{x}_i \in \mathcal{X}_i} \left\{ f_i(\mathbf{x}_i, \mathbf{x}_{-i}^t) + \epsilon^t \langle \mathbf{x}_i - \mathbf{x}_i^t, \tilde{\boldsymbol{\pi}}_i(\mathbf{x}^t) \rangle \right. \\ & \left. + (1 - \epsilon^t) \langle \mathbf{x}_i - \mathbf{x}_i^t, \mathbf{f}_i^{t-1} \rangle + \frac{\tau_i}{2} \|\mathbf{x}_i - \mathbf{x}_i^t\|_2^2 \right\}, \end{aligned} \quad (19)$$

where \mathbf{f}_i^t is updated according to

$$\mathbf{f}_i^t = (1 - \epsilon^t) \mathbf{f}_i^{t-1} + \epsilon^t \tilde{\boldsymbol{\pi}}_i(\mathbf{x}^t).$$

Algorithm 1 based on the best-response (19) is then guaranteed to converge to a stationary solution of (16), in the sense specified by Theorem 1.

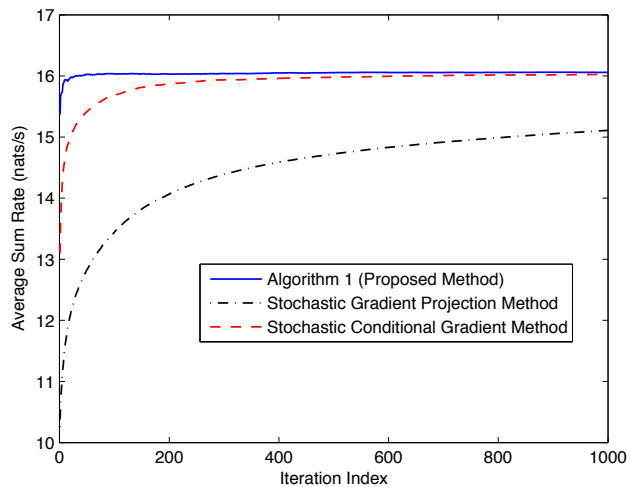


Figure 1. Average sum-rate versus iteration number achieved by Algorithm 1 and stochastic gradient methods.

V. SIMULATIONS

In this section, we consider a MIMO MACs composed of 4 users, while each user is equipped with 2 transmit antennas and 4 receive antennas. The estimated channel matrix is $\mathbf{H} + \Delta\mathbf{H}$ (with $\|\Delta\mathbf{H}\| = 0.1\|\mathbf{H}\|_F$) where \mathbf{H} is the (given) perfect channel matrix while $\Delta\mathbf{H}$ is generated randomly in each iteration by matlab command `randn`. The covariance matrix of the thermal noise is assumed to be the identity matrix.

We compare Algorithm 1 with the popular stochastic conditional gradient method (SCGM) and stochastic gradient projection method (SGPM) in terms of average sum-rate:

$$\frac{1}{t} \sum_{\nu=1}^t \log \det \left(\mathbf{I} + \sum_{i=1}^I \mathbf{H}_i^\nu \mathbf{Q}_i^\nu (\mathbf{H}_i^\nu)^H \right) \Big|_{\mathbf{H}_i^\nu = \mathbf{H}_i + \Delta\mathbf{H}_i^\nu}$$

versus iteration index t . This metric represents the actual average sum transmission rate, and is thus of great practical interest. Since Algorithm 1, SCGM and SGPM are iterative algorithms, we compare them under the same initial point (uniform power allocation) and similar stepsize rules ($\alpha = 0.52$ and $\beta = 0.51$ for Algorithm 1 and SCGM, and $\gamma^t = 1/(t+1)^{0.51}$ for SGPM).

Figure 1 clearly shows that Algorithm 1 outperforms SCGM and SGPM in terms of convergence speed, and the performance gap is notable especially in beginning iterations. This is because the proposed algorithm is a best-response type scheme, which thus explores the convexity of each users' rate function better than what SCGM and SGPM do. Algorithm 1 and SCGM outperform SGPM because the former two are directly based on ascent direction argument.

VI. CONCLUSIONS

In this paper, we have proposed a novel parallel optimization method with provable convergence to compute the stationary solutions of general stochastic nonconvex optimization problems. The distributed algorithm is based on the successive convex approximation idea, and can be implemented among

users in parallel by solving a sequence of (strongly) convex optimization subproblems, one for each user. The memory and signaling burden is maintained at a very low level, making the algorithm very appealing for the distributed design of several multi-agent systems. As case study, we specialized our framework to the stochastic maximization of the sum-rate over MIMO ICs and MACs, showing the superiority of the resulting best-response algorithm with respect to state-of-the-art stochastic gradient schemes.

REFERENCES

- [1] H. Robbins and S. Monro, "A Stochastic Approximation Method," *The Annals of Mathematical Statistics*, vol. 22, no. 3, pp. 400–407, Sep. 1951.
- [2] B. Polyak, *Introduction to optimization*. Optimization Software, 1987.
- [3] H. J. Kushner and G. Yin, *Stochastic approximation and recursive algorithms and applications*, 2nd ed. Springer-Verlag, 2003, vol. 35.
- [4] 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.
- [5] 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.
- [6] Y. Ermoliev, "On the method of generalized stochastic gradients and quasi-Fejer sequences," *Cybernetics*, vol. 5, no. 2, pp. 208–220, 1972.
- [7] 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.
- [8] A. Ruszczyński, "Feasible direction methods for stochastic programming problems," *Mathematical Programming*, vol. 19, no. 1, pp. 220–229, Dec. 1980.
- [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. M. Gupal and L. G. Bazhenov, "Stochastic analog of the conjugant-gradient method," *Cybernetics*, vol. 8, no. 1, pp. 138–140, 1974.
- [11] A. Ribeiro, "Ergodic Stochastic Optimization Algorithms for Wireless Communication and Networking," *IEEE Transactions on Signal Processing*, vol. 58, no. 12, pp. 6369–6386, Dec. 2010.
- [12] G. Scutari, F. Facchinei, P. Song, D. P. Palomar, and J.-S. Pang, "Decomposition by Partial Linearization: Parallel Optimization of Multi-Agent Systems," Feb. 2013, submitted to *IEEE Transactions on Signal Processing*. [Online]. Available: <http://arxiv.org/abs/1302.0756>
- [13] S. Barbarossa, S. Sardellitti, and P. Di Lorenzo, "Distributed detection and estimation in wireless sensor networks," to appear on *E-Reference Signal Processing*, R. Chellapa and S. Theodoridis, Eds., Elsevier, 2013.
- [14] 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.
- [15] 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.
- [16] 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.
- [17] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and distributed computation: Numerical methods*. Prentice Hall, 1989.
- [18] 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.
- [19] Y. Yang, G. Scutari, and D. P. Palomar, "Stochastic parallel decomposition algorithms of multi-user systems," 2013, in preparation.
- [20] 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.
- [21] 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.