A Random Search Framework for Convergence Analysis of Distributed Beamforming With Feedback

Che Lin, *Member, IEEE*, Venugopal V. Veeravalli, *Fellow, IEEE*, and Sean P. Meyn, *Fellow, IEEE*

Abstract—The focus of this work is on the analysis of transmit beamforming schemes with a low-rate feedback link in wireless sensor/relay networks, where nodes in the network need to implement beamforming in a distributed manner. Specifically, the problem of distributed phase alignment is considered, where neither the transmitters nor the receiver has perfect channel state information, but there is a low-rate feedback link from the receiver to the transmitters. In this setting, a framework is proposed for systematically analyzing the performance of distributed beamforming schemes. To illustrate the advantage of this framework, a simple adaptive distributed beamforming scheme that was recently proposed by Mudambai *et al.* is studied. Two important properties of the received signal magnitude function are derived. Using these properties and the systematic framework, it is shown that the adaptive distributed beamforming scheme converges both in probability and in mean. Furthermore, it is established that the time required for the adaptive scheme to converge in mean scales linearly with respect to the number of sensor/relay nodes.

Index Terms—Array signal processing, convergence of numerical methods, detectors, distributed algorithms, feedback communication, networks, relays.

I. INTRODUCTION

THE problem of distributed beamforming arises naturally in wireless sensor/relay networks. In a sensor network, sensors make estimates of a common observed phenomenon and reach a consensus using a local message passing algorithm. In a relay network, a source node intends to communicate with the destination node by passing the message to all relay nodes. In both settings, the sensor/relay nodes then serve as distributed transmitters and seek to convey a common message to the intended receiver. To preserve energy in this stage, transmit beamforming has emerged as a promising scheme due to its potential array gain and low-complexity. However, perfect channel state information (CSI) at the transmitter is required by conventional transmit beamforming schemes to generate beamforming coefficients and achieve phase alignment at the receiver end. This

requirement and the distributed nature of wireless sensor/relay networks make it difficult to implement transmit beamforming schemes in practice. Although obtaining perfect CSI may be too expensive from a practical point-of-view, partial CSI can be made available via a low-rate feedback link from the receiver to the transmitters. As a consequence, there has been increased interest in designing efficient schemes that achieve distributed phase alignment in the presence of a low-rate feedback link [1]–[6]. In this work, our goal is to provide a framework for systematically analyzing the performance of a general set of distributed beamforming schemes with such low-rate feedback.

To illustrate the advantages of our framework, we focus on the analysis of a recently proposed training scheme for distributed beamforming [1]–[3]. The proposed scheme is a simple adaptive algorithm using one bit of feedback information, and is attractive in practice since it is simple to implement. Naturally, one would expect a tradeoff in energy consumption due to possible slow convergence of such a simple adaptive scheme, but surprisingly, the scheme proposed in [1] converges rapidly and hence utilizes energy efficiently. The scheme adjusts the phases for all sensors simultaneously in each time slot to achieve phase alignment. This reduces the overhead significantly compared with direct channel estimation between each source node and the destination node. In fact, the convergence time of the scheme scales linearly with the number of nodes.

Although the scheme of [1] has many desirable features, the fundamental reasons behind the effectiveness of the scheme are unclear from previous work. In [2] and [3], the analyses of the convergence and linear scalability of distributed beamforming schemes have been based on model approximations, which may be loose for some cases. Assuming the stepsize approaches zero, stochastic approximation is used in [4] to show the convergence of the one-bit scheme in distribution. Furthermore, the authors proposed two more algorithms: the signed algorithm and the $\rho\%$ solution algorithm and proved the convergence of both algorithms via the same technique. A discrete version of the problem has been solved in [5], [6] by considering a simplified model with a binary channel and binary signaling.

In this work, instead of focusing on the convergence of a particular algorithm for a particular function, we seek a fundamental understanding into the convergence of distributed beamforming schemes more generally by studying them within the framework of local random search algorithms. Through this framework, we are able to provide a more comprehensive analysis of the fast convergence and linear scalability of the scheme proposed in [1]. In particular, our analysis does not involve approximation of any sort and hence makes statements on convergence and linear scalability in [2]–[4] more

rigorous. Further, we show that due to the special structure of the objective function considered in this problem, *any* adaptive distributed beamforming scheme that can be reformulated as a random search algorithm converges in probability.

We organize the paper as follows: In Section II, we introduce the system model and the received signal magnitude function, which is used as our metric to measure the beamforming array gain throughout the paper. In Section III, we propose a framework that allows for a systematic analysis of a general set of adaptive distributed beamforming schemes. Specifically, we reformulate this set of adaptive distributed beamforming schemes as random search algorithms via a general framework. This reformulation provides insights into the necessary condition for the convergence of the scheme proposed in [1]. These insights lead us to investigate the properties of the received signal magnitude function in Section IV. We further use these properties to prove the convergence of the local random search algorithm in probability and in mean, and provide simulations to validate our analysis. In Section V, we show that the time required for the algorithm to converge in mean scales linearly with the number of nodes. We also provide numerical results that validate our analysis. Finally, we conclude the paper in Section VI and suggest directions for future research.

II. SYSTEM SETUP

We consider the problem of distributed beamforming, where n_s transmitters seek to beamform a common message to one receiver in a distributed manner. We assume that each transmitter and the receiver is equipped with one antenna, and that the channels from the transmitters to the receiver experience frequency-flat, slow fading. The discrete-time, complex baseband system model over a coherence interval is given by

$$
\begin{aligned}
y[t] &= \sum_{i=1}^{n_s} h_i g_i[t] s[t] + w[t] \\
&= \sum_{i=1}^{n_s} a_i b_i[t] e^{j(\phi_i + \psi_i[t])} s[t] + w[t]
\end{aligned}
\tag{1}
$$

where $s[t] \in \mathbb{C}$ is the transmitted common message, $y[t] \in \mathbb{C}$ is the received signal, and $w[t] \sim \mathcal{CN}(0, \sigma^2)$ corresponds to the additive white Gaussian noise. For transmitter i, we denote the channel fading gains by $h_i = a_i e^{j\phi_i} \in \mathbb{C}$ and beamforming coefficients by $g_i[t] = b_i[t] e^{j\psi_i[t]} \in \mathbb{C}$. Note that $a_i \geq 0$, $b_i[t] \geq 0$, and $\phi_i \in [0, 2\pi]$, $\psi_i[t] \in [0, 2\pi]$ for all i and t since they are the corresponding magnitudes and phases of h_i and g_i, respectively. Moreover, a_i and ϕ_i are considered to be constant with time over the coherence interval due to the slow fading assumption. We assume an average power constraint on $s[t]$ given by $E[|s[t]|^2] \leq P$ for all t.

We assume a noncoherent communication model, where the realization of the channel is unknown at both the transmitters and receiver. There is, however, an error-free, zero-delay feedback link of finite capacity from the receiver to all transmitters conveying low-rate partial CSI in each time step.

The goal of distributed beamforming is to pick the beamforming coefficients $\{g_i[t] = b_i[t] e^{j\psi_i[t]}\}$ to maximize the received SNR. In a noncoherent setting and with a low-rate feedback link, beamforming can only be achieved adaptively

through training. Without loss of generality, we assume that the signal $s[t]$ is constant during the training stage. Furthermore, we make the following two simplifications. First, we assume that each transmitter utilizes the same amount of energy for each transmission, i.e., that $b_i[t] = 1$ for all i and t, i.e., we do not optimize the beamforming gains, and we therefore set $s[t] = \sqrt{P}$. This assumption is justified for situations where the transmitters rely on a limited energy source (battery), and allowing them use different amounts of energy would cause some nodes to use up their energy before others. Secondly, we assume that the receiver can estimate the magnitude of the signal component[1] perfectly (without the noise term $w[t]$ in (1)). We therefore use the received signal magnitude as the metric for optimizing the beamforming phases.

The received signal magnitude can be expressed as

$$
\mathrm{Mag}\left(\theta_1[t], \ldots, \theta_{n_s}[t]\right) = \sqrt{P} \left| \sum_{i=1}^{n_s} a_i e^{j\theta_i[t]} \right|
\tag{2}
$$

where $\theta_i[t] = \phi_i + \psi_i[t]$ is the total received phase for sensor i. It is easy to see that $\mathrm{Mag}(\cdot)$ is maximized when the phases $\{\theta_i[t]\}$ are aligned, i.e., they are equal to each other (modulo 2π). Our goal is to study *adaptive distributed beamforming schemes* that achieve this phase alignment through the use of a low-rate feedback link from the receiver.

III. A FRAMEWORK FOR ANALYZING ADAPTIVE DISTRIBUTED BEAMFORMING SCHEMES

In this section, we introduce a framework for analyzing a general class of adaptive distributed beamforming schemes that can be reformulated as random search algorithms. Random search algorithms are well studied in the literature [7]–[9] as methods to maximize an unknown function via random sampling. Once an adaptive distributed beamforming scheme can be successfully reformulated as a random search algorithm, a systematic study of the convergence of such an adaptive scheme is possible.

A. Reformulation of Adaptive Distributed Beamforming Schemes as Random Search Algorithms

Adaptive distributed beamforming algorithms introduced in Section II seek to maximize $\mathrm{Mag}(\cdot)$ given in (2) with the help of a low-rate feedback link. At each step of the adaptation, the signal magnitude at the receiver is a sample of the function $\mathrm{Mag}(\cdot)$. Thus, from the receiver point of view, the problem of distributed phase alignment can be considered under the setting of the following problem:

Problem 1: Given a unknown function $f : \Theta \to \mathbb{R}, \Theta \subseteq \mathbb{R}^n$, where only samples of $f(\boldsymbol{\theta})$ are available for arbitrary $\boldsymbol{\theta} \in \Theta$, find the global maxima of f.

It is important to note that *Problem 1* is a global maximization problem in general if no special structure is assumed for the objective function f. To solve the maximization in *Problem 1*, one may be tempted to use gradient-based algorithms that are well-developed in the literature. Since it is possible for f to possess local maxima, conventional gradient-ascent methods would

---

[1]A good estimate of the received signal magnitude can be obtained directly when the noise is small, or by averaging over several time slots when the noise is not negligible.

fail in general. Besides, acquiring the gradient of the function f may be infeasible especially when the function itself is unknown. Hence, random search techniques [7]–[9] are more appropriate in this setting and can be described as follows.

A Random Search Algorithm:

- *Step zero*: Initialize the algorithm by choosing $\boldsymbol{\theta}[0] \in \Theta$.
- *Step one*: Generate a random perturbation $\boldsymbol{\delta}[t]$ from a probability measure μ_t that could be time-varying.
- *Step two*: Update the search point by $\boldsymbol{\theta}[t] = D(\boldsymbol{\theta}[t-1], \boldsymbol{\delta}[t])$, where the map D satisfies the condition $f(D(\boldsymbol{\theta}[t-1], \boldsymbol{\delta}[t])) \geq f(\boldsymbol{\theta}[t-1])$.

Clearly, for a random search algorithm, we require only function evaluations and control over the probability measure μ_t, which is used to sample the function. Any adaptive distributed beamforming scheme can be reformulated as a random search algorithm if each distributed transmitter initializes its phase as in *Step zero*, generates a random perturbation of phase as in *Step one*, and updates its new phase by the map D as in *Step two*. The low-rate feedback link is used to guarantee the condition $f(D(\boldsymbol{\theta}[t-1], \boldsymbol{\delta}[t])) \geq f(\boldsymbol{\theta}[t-1])$. Note that the unknown function f can be any objective function that we find fit for the distributed transmitters to optimize. This suggests that our framework can be used to analyze a more general function optimization problem over distributed networks. Note further that the probability measure μ_t for the sampling can be time-varying in general. The time-varying nature of the probability measure can be thought of as "*adaptive stepsize*" for distributed algorithms in the most general sense. In this sense, our framework can be used to analyze a large set of adaptive distributed algorithms.

B. One-Bit Adaptive Distributed Beamforming Scheme

To illustrate the advantage of our framework, we now analyze a one-bit adaptive distributed beamforming scheme recently proposed in [1]. Specifically, we reformulate this scheme as a local random search algorithm, which allows for its systematic analysis. We begin by describing the one-bit adaptive distributed beamforming scheme as follows:

A One-bit Adaptive Distributed Beamforming Scheme [1]:
- *Step zero:* Referring to (2) and noting that the i-th transmitter controls its beamforming phase $\psi_i[t]$, the algorithm is initialized by setting $\psi_i[0] = 0$, and hence $\theta_i[0] = \phi_i$ for transmitter i.
- *Step one:* In this step, a random perturbation $\delta_i[t]$ is generated at each distributed transmitter such that $\{\delta_i[t]\}_{i=1}^{n_s}$ are i.i.d. uniform random variables in $[-\delta_0, \delta_0]$ across time and transmitters, where δ_0 is a constant parameter. The random perturbation is added to the total phase of each transmitter. The distributed transmitters then use the perturbed total phases as their new total phases to transmit the training symbol.
- *Step two:* After receiving the training symbols, the receiver measures the received signal magnitude and compares it with the signal magnitude received in the previous time

slot. If the newly received signal magnitude is larger, the receiver feeds back a "*keep*" beacon to the transmitters. Otherwise, a "*discard*" beacon is sent to the transmitters. Note that the beacon is a broadcast from the receiver to all transmitters. Clearly, this feedback scheme only requires one bit of feedback information per time step. When a "*keep*" is received at the transmitters, each transmitter selects and keeps its newly updated total phase. Otherwise, the old phase is selected and the new phase discarded. This selection process is determined by whether the random perturbation increases or decreases the array gain for the adaptive distributed beamforming scheme. Specifically, the evolution of $\boldsymbol{\theta}[t]$ is given by

$$\boldsymbol{\theta}[t] = \begin{cases} \boldsymbol{\theta}[t-1] + \boldsymbol{\delta}[t], & \text{if } \boldsymbol{\delta}[t] \in \mathcal{K} \\ \boldsymbol{\theta}[t-1], & \text{if } \boldsymbol{\delta}[t] \notin \mathcal{K} \end{cases} \qquad (3)$$

where $\boldsymbol{\theta}[t] = [\theta_1[t], \ldots, \theta_{n_s}[t]]^T$, $\boldsymbol{\delta}[t] = [\delta_1[t], \ldots, \delta_{n_s}[t]]^T$, and $\mathcal{K} = \{\boldsymbol{\delta}[t] \, \| \, \text{Mag}(\boldsymbol{\theta}[t-1] + \boldsymbol{\delta}[t]) > \text{Mag}(\boldsymbol{\theta}[t-1])\}$.

Matching the steps of the above one-bit adaptive scheme and those of a random search algorithm introduced in Section III-A, it is clear that the one-bit adaptive distributed beamforming algorithm can be regarded as a special case of the random search algorithm by setting

$$f = \text{Mag}(\cdot) \qquad (4)$$

$$n = n_s \qquad (5)$$

$$\Theta = [0, 2\pi]^{n_s} \qquad (6)$$

$$\mu_t = \mu \qquad (7)$$

$$D(\boldsymbol{\theta}[t-1], \boldsymbol{\delta}[t]) = \boldsymbol{\theta}[t-1] + 1_{\{\boldsymbol{\delta}[t] \in \mathcal{K}\}} \boldsymbol{\delta}[t] \qquad (8)$$

where $1_{\{\cdot\}}$ is the indicator function and μ is uniform on $[-\delta_0, \delta_0]^{n_s}$, which is a n_s-dimensional hypercube. Note that (8) is the same as the evolution described by (3).

Since the probability measure μ is nonzero only within a hypercube, with sides of length $2\delta_0$ and centered around $\boldsymbol{\theta}[t-1]$, the one-bit adaptive distributed beamforming scheme can be reformulated as a *local* random search algorithm. We emphasize again that we can use this framework to study more general adaptive distributed beamforming schemes. For example, the signed algorithm proposed in [4] can be reformulated as a local random search algorithm by simply replacing the indicator function in (8) by the signed function. The framework can also be used to study the case where the probability measure for sampling may be time-varying and with a support that spans the entire space Θ. We can also study adaptive distributed beamforming schemes with more than one bit of feedback information.

It is also interesting to note the connection between this local random search algorithm and simulated annealing [10]. Simulated annealing is a generic probabilistic algorithm that approximates the global optimal solution of a given function in a large search space. The algorithm uses a parameter T called the *temperature* to control the acceptance probability, i.e., the probability that the current state of the algorithm transitions to a new state. If we let $T \to 0$ and assume that the current state is only allowed to move to neighboring states, the simulated annealing procedure reduces to a local random search algorithm.

A local random search algorithm, however, does not necessarily converge in general. For example, if the unknown function possesses local maxima (that are not global maxima), the sequence $\{\boldsymbol{\theta}[t]\}_{t=0}^{\infty}$ is likely to be trapped in a local maximum if the local perturbation δ_0 is not large enough. Thus, a necessary condition for the convergence of local random search algorithms for arbitrary δ_0 is that there is no local maximum point for $\mathrm{Mag}(\cdot)$. With these points in mind, two questions arise naturally: a) does the reformulated local random search algorithm converge? b) if it does, is there a fundamental reason behind the convergence? In the following section, we investigate properties of the function $\mathrm{Mag}(\cdot)$ towards the goal of addressing these questions.

IV. CONVERGENCE OF THE DISTRIBUTED BEAMFORMING SCHEME

A. Properties of Received Signal Magnitude Function

The properties of the received signal magnitude function $\mathrm{Mag}(\cdot)$ do not depend on the time evolution of its arguments. We hence ignore the time dependence of $\theta[t]$ in this section. The following proposition states the first property of $\mathrm{Mag}(\cdot)$.

Proposition 1: For the received signal magnitude function $\mathrm{Mag}(\cdot)$ defined in (2), all local maxima are global maxima.

Proof: To facilitate analysis, we introduce a change of variables

$$\mathbf{x}_i := \begin{bmatrix} x_i^R \\ x_i^I \end{bmatrix} = \begin{bmatrix} \cos\theta_i \\ \sin\theta_i \end{bmatrix}$$

Equation (2) can be rewritten as

$$\mathrm{Mag}(\mathbf{x}_1, \ldots, \mathbf{x}_{ns}) = \sqrt{P} \left\| \sum_{i=1}^{n_s} a_i \mathbf{x}_i \right\|$$

where $\|\mathbf{x}_i\|^2 = 1$ for all $i = 1, \ldots, n_s$. The maximization of $\mathrm{Mag}(\cdot)$ can be rewritten as

$$\max_{\|\mathbf{x}_i\|^2=1, i=1,\ldots,n_s} \left\| \sum_{i=1}^{n_s} a_i \mathbf{x}_i \right\|^2. \tag{9}$$

In the following, we will show that all local maxima of this objective function correspond to complete phase alignment for all transmitters. That is, all local maximum points are global maximum points.

By relaxing the equality constraints to inequality constraints, the optimization problem in (9) is equivalent to

$$\max_{\|\mathbf{x}_i\|^2 \le 1, i=1,\ldots,n_s} \left\| \sum_{i=1}^{n_s} a_i \mathbf{x}_i \right\|^2 \tag{10}$$

This equivalence can be seen as follows: if $\mathbf{x}^*$ is a local maximum with an inactive constraint $\|\mathbf{x}_k^*\|^2 < 1$, by fixing all other variables $\{\mathbf{x}_j^*\}_{j \ne k}$, we obtain

$$\left\| \sum_{i=1}^{n_s} a_i \mathbf{x}_i^* \right\|^2 = \|a_k \mathbf{x}_k^* + \mathbf{c}\|^2$$
$$= \left(a_k x_k^{*R} + c^R \right)^2 + \left(a_k x_k^{*I} + c^I \right)^2$$

where $\mathbf{c} = [c^R c^I]^T$ is a constant vector depending on $\{\mathbf{x}_j^*\}_{j \ne k}$. Obviously, the above function can be improved by appropriately perturbing $\|\mathbf{x}_k^*\|$ according to the signs of c^R and c^I. This contradicts the fact that $\mathbf{x}^*$ is a maximum. Thus, all constraints are

active if $\mathbf{x}^*$ is a maximum point. This shows that the optimization problems (9) and (10) are equivalent.

Focusing on the optimization problem with relaxed constraints, the Lagrangian of (10) reads

$$L(\mathbf{x}, \boldsymbol{\lambda}) = -\|\mathbf{w}\|^2 + \sum_{i=1}^{n_s} \lambda_i(\|\mathbf{x}_i\|^2 - 1)$$

where $\mathbf{x} = [\mathbf{x}_1^T, \ldots, \mathbf{x}_{n_s}^T]^T$, $\boldsymbol{\lambda} = [\lambda_1, \ldots, \lambda_{n_s}]^T$, $\lambda_i \ge 0$ for all $i = 1, \ldots, n_s$, and $\mathbf{w} = \sum_{i=1}^{n_s} a_i \mathbf{x}_i$. By the Lagrange Multiplier Theorem, all local maxima satisfy

$$\nabla_{\mathbf{x}_i} L(\mathbf{x}, \boldsymbol{\lambda}) = -2a_i \mathbf{w}^T + 2\lambda_i \mathbf{x}_i^T = \mathbf{0}^T \tag{11}$$

$$\sum_{i=1}^{n_s} \lambda_i(\|\mathbf{x}_i\|^2 - 1) = 0 \tag{12}$$

$$\|\mathbf{x}_i\|^2 - 1 \le 0 \tag{13}$$

for all $i = 1, \ldots, n_s$. Let $\mathbf{x}^*$ be a local maximum and $\boldsymbol{\lambda}^*$ be the corresponding Lagrange multipliers. If $\lambda_i^* = 0$, (11) implies that $\mathbf{w} = \mathbf{0}$ since[2] $a_i > 0$. In this case, $\mathrm{Mag}(\mathbf{x}^*) = 0$ and this contradicts the fact that $\mathbf{x}^*$ is a local maximum, since we can always improve $\mathrm{Mag}(\cdot)$ by letting $\mathbf{x}_i^* = [\xi \ 0]^T$, $\xi \le 1$, and $\mathbf{x}_j = \mathbf{0}$ for all $j \ne i$. This leads to $\lambda_i > 0$ for all i. We hence have

$$\mathbf{x}_i^* = \frac{a_i}{\lambda_i^*} \mathbf{w} \tag{14}$$

$$\lambda_i^* = a_i \|\mathbf{w}\|. \tag{15}$$

The optimal solutions described by (14) and (15), however, also satisfy

$$\mathrm{Mag}(\mathbf{x}^*) = \sqrt{P} \left\| \sum_{i=1}^{n_s} a_i \frac{\mathbf{w}}{\|\mathbf{w}\|} \right\| = \sqrt{P} \sum_{i=1}^{n_s} a_i$$

and hence are global maxima. This completes our proof. ∎

Proposition 1 implies that the local random search algorithm cannot be trapped in a suboptimal local maximum since all local maxima are global maxima. Furthermore, it also suggests that the necessary condition for the convergence of random search algorithms is satisfied. While it is intuitively clear that the local random search algorithm should converge according to *Proposition* 1, it is to be noted that the condition is only necessary and may not be sufficient. We will provide a rigorous proof of the convergence of the local random search algorithm later. Now, we explore an additional property of $\mathrm{Mag}(\cdot)$ that explains the efficiency of the algorithm.

Another interesting property of $\mathrm{Mag}(\cdot)$ is that it is invariant under a common phase shift to all transmitters. That is

$$\mathrm{Mag}(\boldsymbol{\theta} + \theta_c \mathbf{e}) = \sqrt{P} \left| \sum_{i=1}^{n_s} a_i e^{j(\theta_i + \theta_c)} \right|$$
$$= \sqrt{P} \left| e^{j\theta_c} \sum_{i=1}^{n_s} a_i e^{j\theta_i} \right| = \mathrm{Mag}(\boldsymbol{\theta})$$

where $\mathbf{e}$ is a $n_s \times 1$ vector with all elements equal to one, and θ_c is a common phase shift that can depend on $\{\theta_i\}_{i=1}^{n_s}$. One possible choice for the common phase shift is to let $\theta_c(\theta_1, \ldots, \theta_{n_s})$

---

[2] Note that the case where $a_i = 0$ is not interesting since we can always reduce the dimension of the problem by ignoring $\mathbf{x}_i$

be such that the imaginary part within the modulus function is canceled, i.e.,

$$
\begin{aligned}
\mathrm{Mag}(\boldsymbol{\theta}) &= \mathrm{Mag}(\boldsymbol{\theta} + \theta_c(\theta_1, \ldots, \theta_{n_s})\mathbf{e}) \\
&= \sqrt{P} \sum_{i=1}^{n_s} a_i \cos(\theta_i + \theta_c(\theta_1, \ldots, \theta_{n_s})) \\
&= \sqrt{P} \sum_{i=1}^{n_s} a_i \cos \theta_i' = \mathrm{Mag}(\boldsymbol{\theta}')
\end{aligned}
$$

where $\boldsymbol{\theta}' = [\theta_1', \ldots, \theta_{n_s}']^T$. Note that in the shifted $\boldsymbol{\theta}'$ domain, the global maxima occur only when $\theta_i' = 0$ or $2k\pi$ for all i, where k is any integer. The shift-invariant property results in multiple global maxima for the function $\mathrm{Mag}(\cdot)$. In fact, all global maxima form a one-dimensional "ridge" since if $\boldsymbol{\theta}^*$ is a global maximum, $\bar{\boldsymbol{\theta}}$ with $\bar{\theta}_i = \theta_i^* + \theta_c$ is also a global maximum. This property leads to the rapid convergence of the local random search algorithm since converging to any of these global maximum points is adequate. We conclude this section by summarizing these two important properties of $\mathrm{Mag}(\cdot)$ as follows:

1) all local maxima are global maxima;
2) a common shift to its arguments does not change its value.

B. Proof of Convergence

Intuitively, *Property* 1 guarantees the convergence of any local random search algorithm. To make this precise, we introduce an ϵ-convergence region

$$
R_\epsilon = \{\boldsymbol{\theta} \in \Theta : \mathrm{Mag}(\boldsymbol{\theta}) > \mathrm{Mag}(\boldsymbol{\theta}^*) - \epsilon\} \quad (16)
$$

where $\boldsymbol{\theta}^*$ is the optimal total phase and satisfies $\mathrm{Mag}(\boldsymbol{\theta}^*) = \sqrt{P} \sum_{i=1}^{n_s} a_i$. We define the convergence of a random search algorithm in probability as follows.

Definition 1: A sequence $\{\boldsymbol{\theta}[t]\}_{t=0}^\infty$ generated by a random search algorithm is said to be convergent in probability if, given $\epsilon > 0$,

$$
\lim_{t \to \infty} \Pr[\boldsymbol{\theta}[t] \in R_\epsilon] = 1
$$

In other words, $\mathrm{Mag}(\boldsymbol{\theta}[t])$ converges to $\mathrm{Mag}(\boldsymbol{\theta}^*)$ in probability.

For the proof of convergence, we further derive a proposition stating that for any $\boldsymbol{\theta}$ outside of R_ϵ, there is a nonzero probability of improving $\mathrm{Mag}(\cdot)$ by applying a local perturbation to $\boldsymbol{\theta}$.

Proposition 2: For any given $\boldsymbol{\theta} \in \Theta \setminus R_\epsilon$ and $\delta_0 > 0$, there correspond $\gamma > 0$ and $0 < \eta \le 1$ such that

$$
\Pr[\mathrm{Mag}(\boldsymbol{\theta} + \boldsymbol{\delta}) - \mathrm{Mag}(\boldsymbol{\theta}) \ge \gamma] \ge \eta
$$

where $\boldsymbol{\delta}$ is a random vector with i.i.d. elements uniformly distributed over $[-\delta_0, \delta_0]$.

Proof: From *Proposition 1*, all local maxima are global maxima for the function $\mathrm{Mag}(\cdot)$. This implies that for all $\boldsymbol{\theta} \notin R_\epsilon$ and all $\delta_0 > 0$, there exists a point $\boldsymbol{\theta}_u \in S_{\boldsymbol{\theta}}$ and a constant $\gamma(\boldsymbol{\theta}) > 0$ such that

$$
\mathrm{Mag}(\boldsymbol{\theta}_u) - \mathrm{Mag}(\boldsymbol{\theta}) \ge 2\gamma(\boldsymbol{\theta}) \quad (17)
$$

where the set $S_{\boldsymbol{\theta}}$ is a hypercube of length $2\delta_0$ centered around $\boldsymbol{\theta}$ given by

$$
S_{\boldsymbol{\theta}} = \{\boldsymbol{\omega} \in \Theta : \boldsymbol{\omega} = \boldsymbol{\theta} + \boldsymbol{\delta}, \ \boldsymbol{\delta} \in [-\delta_0, \delta_0]^{n_s}\}
$$

The continuity of $\mathrm{Mag}(\cdot)$ implies that there exists $\sigma(\boldsymbol{\theta}_u) > 0$ such that for all $\boldsymbol{\xi} \in T := \{\boldsymbol{\omega} \in \Theta : \|\boldsymbol{\omega}\| \le \sigma(\boldsymbol{\theta}_u)\}$, we have

$$
|\mathrm{Mag}(\boldsymbol{\theta}_u + \boldsymbol{\xi}) - \mathrm{Mag}(\boldsymbol{\theta}_u)| \le \gamma(\boldsymbol{\theta}) \quad (18)
$$

Combining (17) and (18), we arrive at a lower bound

$$
\begin{aligned}
\mathrm{Mag}(\boldsymbol{\theta}_u + \boldsymbol{\xi}) - \mathrm{Mag}(\boldsymbol{\theta}) &= \mathrm{Mag}(\boldsymbol{\theta}_u + \boldsymbol{\xi}) - \mathrm{Mag}(\boldsymbol{\theta}_u) \\
&\quad + \mathrm{Mag}(\boldsymbol{\theta}_u) - \mathrm{Mag}(\boldsymbol{\theta}) \\
&\ge -\gamma(\boldsymbol{\theta}) + 2\gamma(\boldsymbol{\theta}) = \gamma(\boldsymbol{\theta})
\end{aligned}
$$

Referring to (4) for the definition of μ, the above lower bound leads to

$$
\Pr[\mathrm{Mag}(\boldsymbol{\theta} + \boldsymbol{\delta}) - \mathrm{Mag}(\boldsymbol{\theta}) \ge \gamma(\boldsymbol{\theta})] \ge \mu(T) =: \eta(\boldsymbol{\theta}).
$$

Note that $\mu(T)$ is a function of $\boldsymbol{\theta}$, since $\boldsymbol{\theta}_u$ is a function of $\boldsymbol{\theta}$. We complete the proof of the proposition by letting

$$
\begin{aligned}
\gamma &= \inf_{\boldsymbol{\theta} \in \Theta \setminus R_\epsilon} \gamma(\boldsymbol{\theta}) \\
\eta &= \inf_{\boldsymbol{\theta} \in \Theta \setminus R_\epsilon} \eta(\boldsymbol{\theta}).
\end{aligned}
$$

$\blacksquare$

Note that the proof of this proposition can easily be generalized for any local random perturbation $\boldsymbol{\delta}$. Since before the sequence reaches the ϵ-convergence region, there is always a nonzero probability of improving $\mathrm{Mag}(\cdot)$ for each time step, the convergence of the sequence is to be expected. A simple deterministic analogue is the convergence of a monotonically nondecreasing function.

The probabilistic nature of the algorithm complicates the proof. This will become clear in the proof of our next theorem.

Theorem 1: For the function $\mathrm{Mag}(\cdot)$ defined in (2), let $\{\boldsymbol{\theta}[t]\}_{t=1}^\infty$ be a sequence generated by the local random search algorithm described in (4)–(8). Then the resulting sequence converges in probability, i.e., given $\epsilon > 0$

$$
\lim_{t \to \infty} \Pr[\boldsymbol{\theta}[t] \in R_\epsilon] = 1
$$

Proof: By *Proposition 2*, we know that given any time t

$$
\begin{aligned}
\Pr[\{\mathrm{Mag}(\boldsymbol{\theta}[t-1] + \boldsymbol{\delta}[t]) &- \mathrm{Mag}(\boldsymbol{\theta}[t-1]) \ge \gamma\} \text{ or} \\
&\{\boldsymbol{\theta} \in R_\epsilon\}] \ge \bar{\eta}
\end{aligned}
$$

where $\bar{\eta} = \min\{\Pr[\boldsymbol{\theta} \in R_\epsilon], \eta\}$. Since Θ is compact and $\mathrm{Mag}(\cdot)$ is continuous, there always exists a positive integer p, such that

$$
p\gamma > \mathrm{Mag}(\boldsymbol{\theta}_1) - \mathrm{Mag}(\boldsymbol{\theta}_2), \ \forall \boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \Theta.
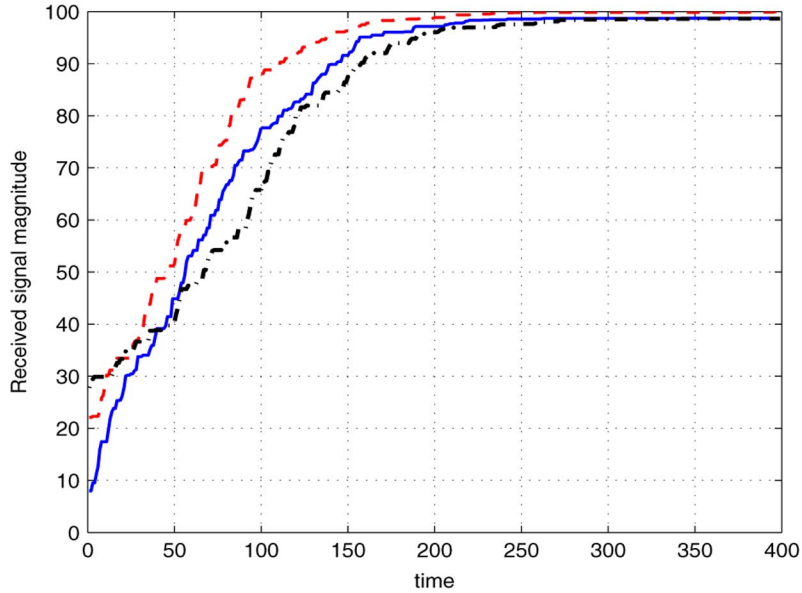$$

Fig. 1. Evolutions of sequences generated by the adaptive distributed beamforming scheme.

The probability that the sequence lies in R_ϵ after p time steps is hence lower bounded by

$$\Pr[\boldsymbol{\theta}[p] \in R_\epsilon] \geq \bar{\eta}^p$$

since $\{\boldsymbol{\delta}[t]\}_{t=0}^{\infty}$ are independent across time. This leads to $\Pr[\boldsymbol{\theta}[p] \notin R_\epsilon] \leq 1 - \bar{\eta}^p$ and

$$\Pr[\boldsymbol{\theta}[pm] \in R_\epsilon] = 1 - \Pr[\boldsymbol{\theta}[pm] \notin R_\epsilon] \geq 1 - (1 - \bar{\eta}^p)^m$$

for all $m = 1, 2, \ldots$. The lower bound is still valid if we let the sequence progress ℓ time steps further, i.e.,

$$\Pr[\boldsymbol{\theta}[pm + \ell] \in R_\epsilon] \geq 1 - (1 - \bar{\eta}^p)^m$$

for all $m = 1, 2, \ldots, \ell = 0, \ldots, p - 1$. We complete the proof by letting $m \to \infty$. Note that by invoking the Borel–Cantelli lemma [11], we can show that the sequence converges almost surely as well. ∎

Theorem 1 states that the local random search algorithm in (4)–(8) converges in probability, and hence also provides a proof of convergence for the one-bit adaptive distributed beamforming scheme in (3). In particular, *Theorem 1* implies the convergence of the sequence $\{\text{Mag}(\boldsymbol{\theta}[t])\}_{t=0}^{\infty}$ in probability. Since the sequence is nonnegative and monotonically nondecreasing, we can conclude that $\{\text{Mag}(\boldsymbol{\theta}[t])\}_{t=0}^{\infty}$ also converges in mean by the Monotone Convergence Theorem [11]. Further, by properly generalizing *Proposition* 2, it is possible to show that *any* adaptive distributed beamforming scheme that can be reformulated a local random search algorithm and seeks to maximize *any* objective function that satisfies *Property* 1 converges in probability.

In Fig. 1, we illustrate the evolution of the sequences generated by the local random search algorithm from different initial points. The initial points are generated randomly from a uniform distribution over Θ. Only three sample paths of the sequence are included in the figure since similar behaviors can be observed for other sample paths. For each iteration, the random perturba-

tion δ_i for the ith transmitter is a uniform random variable over $[-\delta_0, \delta_0]$, where $\delta_0 = \pi/30$. Note that we use the same channel coefficients to generate these sequences since the focus here is on the effect of different initial points. In particular, the channel coefficients are randomly generated from i.i.d. $\mathcal{CN}(0, 1)$ in the beginning of the simulation, and remain fixed afterwards.

From the figure, we observe the rapid convergence of the local random search algorithm, irrespective of where it is initialized. We emphasize again that the fast convergence results follow from the two important properties for the function $\text{Mag}(\cdot)$ as discussed in Section IV-A. *Property* 1 guarantees the convergence of the local search algorithm; *Property* 2 results in multiple global maxima for the function $\text{Mag}(\cdot)$ and hence the fast convergence of the algorithm. The simulations provide a partial validation of our proof since we would expect the convergence to fail from some initial points if there were nonoptimal local maxima for $\text{Mag}(\cdot)$. It is to be noted that the convergence of the local random search algorithm does not guarantee that it is the most efficient scheme in terms of the number of function evaluations, and hence the most efficient scheme in terms of energy. However, the algorithm does have a desirable scaling property, i.e., the time required for the algorithm to converge in mean scales linearly with the number of transmitters. This is the topic of the following section.

V. SCALING LAW

Due to the probabilistic nature of the local random search algorithm, we defined convergence in probability in Section IV-B and showed that the local random search algorithm converges. For the analysis of the scaling law, however, we can only show convergence in mean, which is defined as follows:

Definition 2: A sequence $\{\boldsymbol{\theta}[t]\}$ generated by a random search algorithm is said to converge in mean if for each $\epsilon > 0$ there exists $t_N \geq 0$ such that

$$E_{\{\boldsymbol{\delta}[\tau]\}_{\tau=0}^{t} \mid \mathbf{a}, \boldsymbol{\theta}[0]}[\text{Mag}(\boldsymbol{\theta}[t])] > \text{Mag}(\boldsymbol{\theta}^*) - \epsilon = \sqrt{P} \sum_{i=1}^{n_s} a_i - \epsilon$$
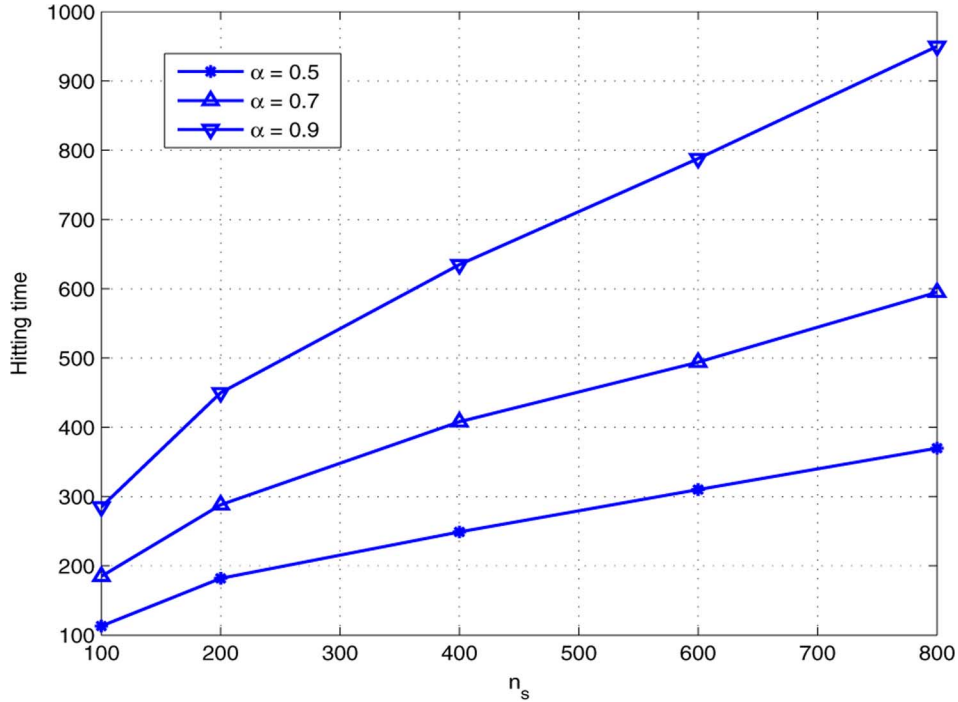
Fig. 2. Hitting time for the adaptive distributed beamforming scheme with different values of α.

for all $t \geq t_N$, where $\mathbf{a} = [a_1, \ldots, a_{n_s}]^T$. That is, $\mathrm{Mag}(\boldsymbol{\theta}[t])$ converges to $\mathrm{Mag}(\boldsymbol{\theta}^*)$ in mean.

In this section, our goal is to find the time required for the local random search algorithm to converge in mean, starting from any initial point. In other words, we are interested in finding the *hitting time*[3] of the random search algorithm, and determining its behavior as a function of the number of transmitters. Specifically, we derive an upper bound on the hitting time of the local random search algorithm as a function of n_s. Note that the study of the hitting time makes sense only if the sequence indeed converges in mean, which we established in Section IV-B.

To facilitate analysis, we define the increment function of $\mathrm{Mag}(\cdot)$ at time τ as

$$
I[\tau] = [\mathrm{Mag}(\boldsymbol{\theta}[\tau]) - \mathrm{Mag}(\boldsymbol{\theta}[\tau - 1])]^+
$$
$$
= [\mathrm{Mag}(\boldsymbol{\theta}[\tau - 1] + \boldsymbol{\delta}[\tau]) - \mathrm{Mag}(\boldsymbol{\theta}[\tau - 1])]^+ \quad (19)
$$

where $[x]^+ = \max(x, 0)$. We then rewrite the received signal magnitude function at any given time $k_0 n_s$ as

$$
\mathrm{Mag}(\boldsymbol{\theta}[k_0 n_s]) = \sum_{\tau=1}^{k_0 n_s} I[\tau] + \mathrm{Mag}(\boldsymbol{\theta}[0])
$$
$$
=: \sum_{\tau=1}^{k_0 n_s} I[\tau] + c_0 \quad (20)
$$

where k_0 is a positive integer and $c_0 \geq 0$.

From *Proposition 2*, we have that for any given τ such that $\boldsymbol{\theta}[\tau - 1] \notin R_\epsilon$ and any local random perturbation $\boldsymbol{\delta}[\tau]$, there correspond $\gamma > 0$ and $0 < \eta \leq 1$ such that

$$
\mathrm{Pr}[\mathrm{Mag}(\boldsymbol{\theta}[\tau - 1] + \boldsymbol{\delta}[\tau]) - \mathrm{Mag}(\boldsymbol{\theta}[\tau - 1]) \geq \gamma] \geq \eta
$$

[3]The hitting time in this work is defined as the time required for the algorithm to converge in mean.

Thus, we have

$$
E_{\boldsymbol{\delta}[\tau] \mid \mathbf{a}, \boldsymbol{\theta}[\tau-1]}[I[\tau]]
$$
$$
\geq \gamma \mathrm{Pr}[\mathrm{Mag}(\boldsymbol{\theta}[\tau - 1] + \boldsymbol{\delta}[\tau]) - \mathrm{Mag}(\boldsymbol{\theta}[\tau - 1]) \geq \gamma]
$$
$$
\geq \gamma \eta > 0
$$

for any τ such that $\boldsymbol{\theta}[\tau - 1] \notin R_\epsilon$. Referring to (19)–(20), we obtain

$$
E_{\{\boldsymbol{\delta}[\tau]\}_{\tau=0}^{k_0 n_s} \mid \mathbf{a}, \boldsymbol{\theta}[0]}[\mathrm{Mag}(\boldsymbol{\theta}[k_0 n_s])]
$$
$$
= \sum_{\tau=1}^{k_0 n_s} E_{\boldsymbol{\delta}[\tau] \mid \mathbf{a}, \boldsymbol{\theta}[\tau-1]}[I[\tau]] + c_0
$$
$$
\geq k_0 n_s \gamma \eta + c_0 \geq \sqrt{P} \sum_{i=1}^{n_s} a_i
$$

where the last inequality follows by choosing $k_0 = \lceil \frac{\sqrt{P} \max_i \{a_i\}}{\gamma \eta} \rceil$. This implies that the hitting time for the local random search algorithm is at most $k_0 n_s$, from any initial point. Hence, the hitting time for the algorithm scales linearly with the number of transmitters.

In our simulations, we say that the sequence converges to the α fraction of the global maxima if $\mathrm{Mag}(\boldsymbol{\theta}[t]) \geq \alpha \mathrm{Mag}(\boldsymbol{\theta}^*)$. We assume that channel coefficients are i.i.d. complex Gaussian variables $\mathcal{CN}(0, 1)$, and use the origin as our initial point. We set $\delta_0 = \pi/90$ for all our simulations. Fig. 2 demonstrates the hitting time required for the adaptive distributed beamforming scheme to converge in a relative sense when $\alpha = 0.5, 0.7$, and 0.9. It is clear that the hitting time increases as α increases. The scaling law for the hitting time with respect to n_s, however, is the same for all values of α. Indeed, we observe linear scaling for all values of α. This observation confirms our theoretical analysis.
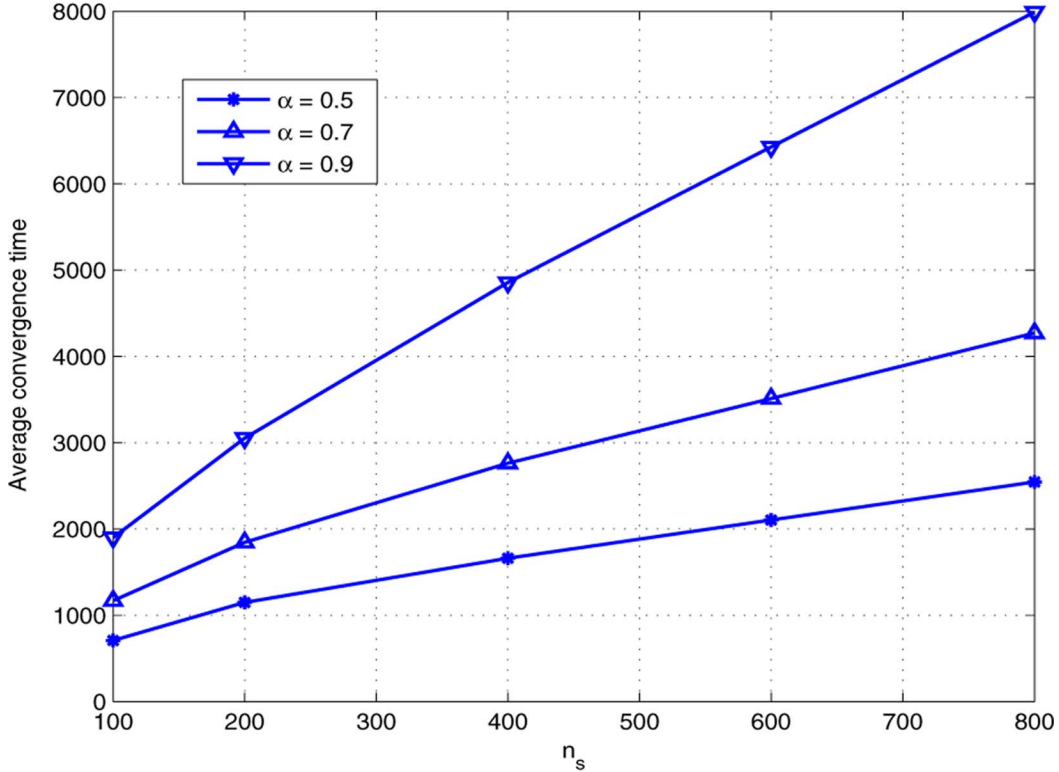
Fig. 3. Average convergence time for the adaptive distributed beamforming scheme with different values of α.

Fig. 3 shows the average convergence time for the adaptive distributed beamforming scheme to within a fraction of the globally maximum value $\alpha \mathrm{Mag}(\boldsymbol{\theta}^*)$, for different values of α. It is important to note the difference between the hitting time and the average convergence time. Since our algorithm is probabilistic in nature, the convergence time is essentially a random variable and each run of the algorithm provides a sample for this random variable. Fixing the number of transmitters n_s, we obtain the average convergence time by averaging over a hundred samples of this random variable, while the hitting time is obtained by comparing $E[\mathrm{Mag}(\boldsymbol{\theta}[t])]$ with $\alpha \mathrm{Mag}(\boldsymbol{\theta}^*)$. From Fig. 3, we observe the same linear scaling behavior for the average convergence time. We expect this property for the average convergence time can be shown in a similar manner.

VI. Concluding Remarks and Future Work

In this work, we have proposed a framework that allows for a systematic analysis of adaptive distributed beamforming schemes in sensor/relay networks. We used this framework to study the convergence and scaling law of a recently proposed one-bit adaptive distributed beamforming scheme [1]. We first reformulated the one-bit adaptive scheme as a local random search algorithm. This reformulation provided insights into the convergence of the one-bit adaptive scheme, and led us to investigate the fundamental properties of the received signal magnitude function $\mathrm{Mag}(\cdot)$. We identified two important properties of the function that contribute to the rapid convergence of the algorithm. First, all local maxima are global maxima. This prevents any local random search algorithm from being trapped in nonoptimal local maximum points. Secondly, the $\mathrm{Mag}(\cdot)$

function is invariant under a common shift to its arguments. This property results in multiple global maximum points for $\mathrm{Mag}(\cdot)$ and hence the rapid convergence of the algorithm. Based on these properties, we have shown the convergence of the algorithm, both in probability and in mean. We further provided an upper bound on the hitting time of the algorithm, and demonstrated that the hitting time scales linearly with the number of sensor/relay nodes. This linear scaling is desirable, especially when the network is densely populated. We have also provided simulations that validate our analysis.

It is important to note that the effectiveness of the one-bit adaptive distributed beamforming scheme depends critically on the properties of the function $\mathrm{Mag}(\cdot)$. Maximizing $\mathrm{Mag}(\cdot)$ is equivalent to maximizing the received SNR if there is no error in obtaining the common message, which is true in the training stage since the common message is simply fixed and known to the receiver. On the other hand, if adaptation is being performed blindly (without training) it would be necessary to consider the possibility of errors in common message. The corresponding objective function may then not possess the same desirable properties as $\mathrm{Mag}(\cdot)$, e.g., the objective function may possess local maxima that are not global maxima. Much work needs to be done to understand how our results can be applied in this more complicated scenario. One thing that is clear, however, is that we will need to develop new algorithms that exploit the global structure of the new objective function since local algorithms can be trapped in local maxima. Our general framework for studying adaptive beamforming algorithms is even more useful in this context since it connects the problem to a well-studied field of global optimization algorithms.

REFERENCES

[1] R. Mudumbai, B. Wild, U. Madhow, and K. Ramchandran, "Distributed beamforming using 1 bit feedback: From concept to realization," in *Proc. Allerton Conf. Communication, Control, and Computing*, Monticello, IL, Sep. 2006.

[2] R. Mudumbai, J. P. Hespanha, U. Madhow, and G. Barriac, "Distributed transmit beamforming using feedback control," *IEEE Trans. Inform. Theory*, vol. 56, no. 1, pp. 411–426, Jan. 2010.

[3] R. Mudumbai, "Energy efficient wireless communication using distributed beamforming," Ph.D. dissertation, Univ. California, Santa Barbara, CA, Dec. 2007.

[4] J. Bucklew and W. Sethares, "Convergence of a class of decentralized beamforming algorithms," *IEEE Trans. Signal Process.*, vol. 56, no. 6, pp. 2280–2288, Jun. 2008.

[5] J. Thukral and H. Bolcskei, "Distributed spatial multiplexing with 1-bit feedback," in *Proc. Allerton Conf. Communication, Control, and Computing*, Monticello, IL, Sep. 2007.

[6] M. Johnson, M. Mitzenmacher, and K. Ramchandran, "Distributed beamforming with binary signaling," in *Proc. IEEE Int. Symp. Information Theory*, Toronto, ON, Canada, Jul. 2008.

[7] F. J. Solis and R. J.-B. Wets, "Minimization by random search techniques," *J. Math. Oper. Res.*, vol. 6, pp. 19–30, Feb. 1981.

[8] L. Shi and S. Olafsson, "Nested partitions method for global optimization," *J. Math. Oper. Res.*, vol. 48, pp. 390–407, May 2000.

[9] Z. B. Tang, "Adaptive partitioned random search to global optimization," *IEEE Trans. Autom. Control*, vol. 39, no. 11, pp. 2235–2244, Nov. 1994.

[10] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, no. 4598, pp. 671–680, 1983.

[11] R. A. Durrett, *Probability: Theory and Examples*, 2nd ed. Pacific Grove, PA: Duxbury Press, 1995.

Che Lin (S'02–M'08) received the B.S. degree in electrical engineering from National Taiwan University, Taipei, Taiwan, R.O.C., in 1999, the M.S. degree in electrical and computer engineering in 2003, the M.S. degree in math in 2008, and the Ph.D. degree in electrical and computer engineering in 2008, all from the University of Illinois at Urbana-Champaign.

Since 2008, he has been with National Tsing Hua University, where he is currently an Assistant Professor. He holds a U.S. patent, which has been included in the 3GPP LTE standard. His research interests include feedback systems, distributed algorithms in networks, practical MIMO code design in general networks, optimization theory, and information theory.

Dr. Lin received a two-year Vodafone Graduate Fellowship in 2006 and the E. A. Reid fellowship Award in 2008.

Venugopal V. Veeravalli (S'86–M'92–SM'98–F'06) received the B.Tech. degree (Silver Medal Honors) in 1985 from the Indian Institute of Technology, Bombay, the M.S. degree in 1987 from Carnegie-Mellon University, Pittsburgh, PA, and the Ph.D. degree in 1992 from the University of Illinois at Urbana-Champaign, all in electrical engineering.

He joined the University of Illinois at Urbana-Champaign in 2000, where he is currently a Professor in the Department of Electrical and Computer Engineering, and a Research Professor in the Coordinated Science Laboratory. He served as a Program Director for Communications Research at the U.S. National Science Foundation, Arlington, VA from 2003 to 2005 and has previously held academic positions at Harvard University, Rice University, and Cornell University. His research interests include distributed sensor systems and networks, wireless communications, detection and estimation theory, and information theory.

Dr. Veeravalli was on the Board of Governors of the IEEE Information Theory Society from 2004 to 2007. He was an Associate Editor for Detection and Estimation for the IEEE TRANSACTIONS ON INFORMATION THEORY from 2000 to 2003 and an Associate Editor for the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS from 1999 to 2000. Among the awards he has received for research and teaching are the IEEE Browder J. Thompson Best Paper Award, the National Science Foundation CAREER Award, and the Presidential Early Career Award for Scientists and Engineers (PECASE). He is a distinguished Lecturer for the IEEE Signal Processing Society for 2010–2011.

Sean P. Meyn (S'85–M'87–SM'95–F'02) received the B.A. degree in mathematics (*summa cum laude*) from the University of California at Los Angeles in 1982 and the Ph.D. degree in electrical engineering from McGill University, Montreal, QC, Canada, in 1987.

After a two-year postdoctoral fellowship at the Australian National University, Canberra, he and his family moved to the Midwest. He is now a Professor in the Department of Electrical and Computer Engineering, and a Research Professor in the Coordinated Science Laboratory, both at the University of Illinois at Urbana-Champaign, where he is Director of the Decision and Control Lab. He has held visiting positions at universities all over the world, including the Indian Institute of Science, Bangalore (1997–1998), where he was a Fulbright Research Scholar. During his latest sabbatical during the 2006–2007 academic year, he was a Visiting Professor at MIT and United Technologies Research Center (UTRC). His research interests include stochastic processes, control and optimization, complex networks, and information theory. His current funding is provided by the National Science Foundation, the Department of Energy, AFOSR, and DARPA. He is coauthor (with R. Tweedie) of the monograph *Markov Chains and Stochastic Stability* (London, U.K.: Springer-Verlag, 1993) and author of *Control Techniques for Complex Networks* (Cambridge, MA: Cambridge University Press, 2007).

Dr. Meyn received jointly (with R. Tweedie) the 1994 ORSA/TIMS Best Publication In Applied Probability Award.