Distributed Optimization in Multi-User MIMO Systems with Imperfect and Delayed Information

Pierre Coucheney
University of Versailles
PRiSM Laboratory
E-mail: pierre.coucheney@uvsq.fr

Bruno Gaujal
Inria
LIG Laboratory
E-mail: bruno.gaujal@inria.fr

Panayotis Mertikopoulos
French National Center for Scientific Research (CNRS)
LIG Laboratory
E-mail: panayotis.mertikopoulos@imag.fr

***Abstract*—In this paper, we analyze the problem of signal covariance optimization in Gaussian multiple-input, multiple-output (MIMO) channels under imperfect (and possibly delayed) channel state information. Starting from the continuous-time dynamics of matrix exponential learning, we develop a distributed optimization algorithm driven by a damping term which ensures the method's stability under stochastic perturbations and asynchronicities of arbitrary magnitude. As opposed to traditional water-filling methods, the algorithm's convergence properties (speed and accuracy) can be controlled by tuning the users' learning rate and/or the damping parameter. Accordingly, the algorithm converges arbitrarily close to an optimum signal covariance profile within a few iterations, even for large numbers of users and/or antennas per user; furthermore, the quality of the solution obtained remains robust in the presence of imperfect (or delayed) measurements and asynchronous user updates.**

***Index Terms*—Distributed optimization; imperfect CSI; matrix exponential learning; multiple access channels; MIMO; stochastic approximation.**

I. Introduction

The seminal prediction that the use of multiple antennas can lead to substantial performance gains in signal transmission and reception [1, 2] has made multiple-input and multiple-output (MIMO) technologies an integral component of most state-of-the-art wireless communication protocols, ranging from 3G LTE, 4G and HSPA+, to 802.11n WiFi and WiMax (to name but a few). Nonetheless, seeing as the radio spectrum is shared by all users, the intended receiver of a signal must still cope with unwarranted interference from a large number of transmitters, a factor which severely limits the capacity of the wireless system in question. On that account, and given that the theoretical performance limits of MIMO systems still elude us (even in basic network models such as the interference channel), a widespread approach is to use the mutual information for Gaussian input and noise as a performance metric, and to optimize (the covariance of) the input signal distribution of each transmitter in the presence of interference from all other users.

In this paper, we focus on uplink MIMO systems consisting of several non-cooperative (and mutually independent) Gaussian transmitters who upload data to a common receiver. This vector Gaussian multiple access channel (MAC) has attracted significant interest in the literature [2, 3], traditionally relying on water-filling (WF) techniques to achieve its capacity [3–5]. Unfortunately however, the convergence speed of iterative water-filling (IWF) methods decreases rapidly with the number of users (making such methods unsuitable for large networks), whereas the convergence of faster, simultaneous water-filling (SWF) methods [6] is conditional on certain "mild-interference" conditions which fail to hold even in simple 2-user parallel multiple access channels (PMACs) [7, 8].

To overcome these limitations, the authors of [9] proposed a so-called "matrix exponential learning" method whose convergence speed scales well with the number of users in the system. That said, just like water-filling, this method relies on perfect channel state information (CSI), an assumption which breaks down in rapidly evolving, unregulated networks. Consequently, a major challenge arises when this information can only be estimated in an *imperfect* manner, and/or when only *delayed* (and, hence, potentially obsolete) measurements are available: for instance, the analysis of [10] can be used to show that stochastic perturbations could lead the system to a *globally suboptimal* state with positive probability, even in the very simple case of a single user. Moreover, in the absence of a centralized scheduler, enforcing simultaneous user updates is all but impossible, so it is not clear whether methods that rely on synchronous decision-taking can be implemented in decentralized environments.

To address these issues, we introduce a dissipative variant of the matrix exponential learning method of [9] which penalizes zero eigenvalues in the users' signal covariance profile.[1] In so doing, the dynamics are stabilized (in the sense of stochastic approximation [11, 12]) and the system converges to a (nonsingular) signal covariance profile that is arbitrarily close to an optimum one. The powerful stochastic approximation techniques of [12] then allow us to show that the resulting algorithm converges very fast even for large numbers of users and/or antennas per user (in practice, within a few iterations), and this convergence remains robust even in the presence of arbitrarily large estimation errors and asynchronous updates.

II. System Model

Consider a vector Gaussian multiple access channel consisting of a finite set of transmitters $k \in \mathcal{K} \equiv \{1, \dots, K\}$, each

[1] We should stress here that the resulting method can be extended to a wide class of nonlinear semidefinite programming problems; we chose to focus here on the MIMO MAC case for simplicity and concreteness.

equipped with m_k antennas, and each transmitting simultaneously to a base receiver with m_0 antennas. This system may then be represented by the familiar baseband model

$$\mathbf{y} = \sum_{k=1}^{K} \mathbf{H}_k \mathbf{x}_k + \mathbf{z}, \tag{1}$$

where $\mathbf{y} \in \mathbb{C}^{m_0}$ denotes the aggregate signal reaching the receiver, $\mathbf{x}_k \in \mathbb{C}^{m_k}$ is the message transmitted by user $k \in \mathcal{K}$, $\mathbf{H}_k \in \mathbb{C}^{m_0 \times m_k}$ is the associated $m_0 \times m_k$ (complex) channel matrix, and $\mathbf{z} \in \mathbb{C}^{m_0}$ is the noise in the channel, including thermal, atmospheric and other peripheral interference effects, and assumed to be a (zero-mean) circularly symmetric complex Gaussian random vector with non-singular covariance (taken equal to $\mathbf{I}$ after a change of basis).

In this setting, the average transmit power of user k will be

$$P_k = \mathbb{E}\,[\|\mathbf{x}_k\|^2] = \operatorname{tr}(\mathbf{Q}_k), \tag{2}$$

where the expectation is taken over the codebook of user k and $\mathbf{Q}_k$ denotes the corresponding signal covariance matrix:

$$\mathbf{Q}_k = \mathbb{E}\,[\mathbf{x}_k \mathbf{x}_k^\dagger]. \tag{3}$$

Hence, assuming successive interference cancellation (SIC) at the receiver, the maximum information transmission rate will be achieved for random Gaussian codes and will be given by

$$\Psi(\mathbf{Q}) = \log\det\left(\mathbf{I} + \textstyle\sum_k \mathbf{H}_k \mathbf{Q}_k \mathbf{H}_k^\dagger\right), \tag{4}$$

where $\mathbf{Q} = \operatorname{diag}(\mathbf{Q}_1, \dots, \mathbf{Q}_k)$ denotes the block-diagonal (direct) sum of the individual matrices $\mathbf{Q}_k$ [2]. In this way, we obtain the sum rate maximization problem [4, 5]:

$$\begin{aligned} &\text{maximize} && \Psi(\mathbf{Q}), \\ &\text{subject to} && \mathbf{Q}_k \in \mathcal{X}_k \quad (k = 1, \dots, K), \end{aligned} \tag{5}$$

where $\mathcal{X}_k = \{\mathbf{Q}_k \in \mathbb{C}^{m_k \times m_k} : \mathbf{Q}_k \succcurlyeq 0, \operatorname{tr}(\mathbf{Q}_k) = P_k\}$ is the set of feasible signal covariance matrices of user k that satisfy the power constraint $\operatorname{tr}(\mathbf{Q}_k) = P_k$.[2]

The sum rate problem (5) is traditionally solved by water-filling (WF) methods – whether iterative (IWF) [4] or simultaneous (SWF) [6]. In these methods, it is assumed that transmitters have perfect knolwedge of the channel matrices $\mathbf{H}_k$ and the aggregate signal-plus-noise covariance matrix

$$\mathbf{W} = \mathbf{I} + \textstyle\sum_\ell \mathbf{H}_\ell \mathbf{Q}_\ell \mathbf{H}_\ell^\dagger, \tag{6}$$

which can then be used to compute the user-specific matrices

$$\mathbf{W}_k = \mathbf{I} + \textstyle\sum_{\ell \neq k} \mathbf{H}_\ell \mathbf{Q}_\ell \mathbf{H}_\ell^\dagger = \mathbf{W} - \mathbf{H}_k \mathbf{Q}_k \mathbf{H}_k^\dagger. \tag{7}$$

As a result, the information requirements of IWF and SWF boil down to perfect channel state information at the transmitter (CSIT) and accurate measurements of the covariance matrix $\mathbf{W}$ at the receiver (who can e.g. broadcast this information via a dedicated radio channel). On the other hand, despite this similarity, IWF and SWF behave quite differently performance-wise: the former converges always (but slowly for large numbers of users), whereas the latter is faster but may fail to converge – even in 2-user parallel MACs [7, 8].

[2]The power constraint $\operatorname{tr}(\mathbf{Q}_k) = P_k$ is actually a special case of the more general inequality constraint $\operatorname{tr}(\mathbf{Q}_k) \leq P_k$. However, in the absence of power consumption considerations, every user will saturate this constraint in order to achieve higher transmission rates; as such, we will not treat this case here.

III. Entropy Adjusted Exponential Learning

To overcome the limitations of WF methods (and under the same information assumptions), the authors of [9] considered instead the dynamics of matrix exponential learning:

$$\dot{\mathbf{Y}}_k = \mathbf{V}_k, \qquad \mathbf{Q}_k = P_k \frac{\exp(\mathbf{Y}_k)}{\operatorname{tr}[\exp(\mathbf{Y}_k)]}, \tag{8}$$

where $\mathbf{Y}_k$ is an auxiliary Hermitian "scoring" matrix and

$$\mathbf{V}_k = \nabla_{\mathbf{Q}_k^*} \Psi = \mathbf{H}_k^\dagger \mathbf{W}^{-1} \mathbf{H}_k \tag{9}$$

denotes the (conjugate) derivative of Ψ w.r.t. $\mathbf{Q}_k$. By discretizing, we then obtain the learning scheme:

$$\mathbf{Y}_k(n) = \mathbf{Y}_k(n-1) + \gamma_n \mathbf{V}_k(n), \quad \mathbf{Q}_k(n+1) = P_k \frac{e^{\mathbf{Y}_k(n)}}{\operatorname{tr}[e^{\mathbf{Y}_k(n)}]}, \tag{10}$$

where γ_n is a variable step size satsifying the technical requirements $\sum_n \gamma_n = +\infty$ and $\sum_n \gamma_n^2 < +\infty$ (e.g. $\gamma_n = 1/n$) [11].

Intuitively, (10) reinforces the spatial directions that peform well by increasing the corresponding eigenvalues, so users converge to a solution of (5) within a few iterations [9]. In practice however, a major challenge occurs if the network's users only have acces to imperfect (and/or delayed) CSI; accordingly, our aim in the rest of this paper will be to develop a fast solution method for (5) that retains its convergence properties even in the presence of stochastically perturbed observations and/or asynchronicities (in either the updates or the observations).

A. *Learning in the Presence of Noise*

To account for as wide a range of measurement errors as possible, we will assume that at each update epoch $n = 1, 2, \dots$, the system's users can only measure a perturbed estimate

$$\hat{\mathbf{V}}_k(n) = \mathbf{V}_k(n) + \mathbf{\Xi}_k(n) \tag{11}$$

of $\mathbf{V}_k(n)$, with $\mathbf{\Xi}_k$ representing a random observational error (not necessarily i.i.d.). Specifically, our only assumption for $\mathbf{\Xi}_k$ will be that it is a bounded martingale difference, i.e. $\|\mathbf{\Xi}_k\| \leq \Sigma$ for some $\Sigma > 0$ and $\mathbb{E}\,[\mathbf{\Xi}_k(n) \,|\, \mathbf{\Xi}(n-1), \dots, \mathbf{\Xi}(1)] = 0$.

Under these assumptions, one could hope to apply the stochastic approximation techniques of [11, 12] to show that the resulting stochastic version of (10) converges to an optimum signal covariance profile. This, however, is not the case: in PMAC systems (where $\mathbf{H}_k$ and $\mathbf{Q}_k$ are jointly diagonalizable), the dynamics (8) boil down to the well-known *replicator dynamics* of evolutionary game theory [11], viz.

$$\dot{q}_{k\alpha} = q_{k\alpha}\left(V_{k\alpha} - \textstyle\sum_{\beta=0}^{m_k} q_{k\beta} V_{k\beta}\right), \tag{12}$$

where $q_{k\alpha}$ is the α-th diagonal eigenvalue of $\mathbf{Q}_k$ and $V_{k\alpha} = \frac{\partial \Psi}{\partial q_{k\alpha}}$ is the α-th eigenvalue of $\mathbf{V}_k$ [8]. As it turns out, the discrete replicator dynamics with linear costs are known to converge with positive probability to a global minimum of the objective function in stochastic environments [10], so (10) might lead to similarly unwarranted behavior.

From a mathematical viewpoint, the problem with (10) is that the scoring matrices $\mathbf{Y}_k$ may fail to remain bounded for all time, so the tracking techniques of stochastic approximation do not apply. Motivated by the analysis of [13] for finite games, we will thus consider the following *dissipative* variant of matrix exponential learning:

$$\mathbf{Y}_k(n+1) = \mathbf{Y}_k(n) + \gamma_n[\hat{\mathbf{V}}_k(n) - \tau\mathbf{Y}_k(n)],$$
$$\mathbf{Q}_k(n+1) = P_k\frac{\exp(\mathbf{Y}_k(n+1))}{\operatorname{tr}[\exp(\mathbf{Y}_k(n+1)]}, \tag{13}$$

where $\tau > 0$ is a damping parameter whose role is to keep the matrix scores bounded for all time:

Lemma 1. *The process* $\mathbf{Y}_k$ *of* (13) *is bounded* (*a.s.*).

The proof of this lemma[3] relies on a recursive argument that takes advantage of the fact that if an element of $\mathbf{Y}_k$ gets too big, the damping term $-\tau\mathbf{Y}_k$ will decrease it on the next iteration.[4] With this in mind, the general theory of stochastic approximation [11, 12] shows that the process defined by (13) will be an asymptotic pseudo-trajectory (APT) of the so-called mean dynamics:

$$\dot{\mathbf{Y}}_k = \mathbf{V}_k - \tau\mathbf{Y}_k, \qquad \mathbf{Q}_k = P_k\frac{\exp(\mathbf{Y}_k)}{\operatorname{tr}[\exp(\mathbf{Y}_k)]}, \tag{14}$$

i.e. the discrete-time recursion (13) will be attracted to the internally chain transitive (ICT) sets of the dynamics (14) [11].

Intuitively, this means that the discrete-time learning process (13) will converge to the same points as the continuous-time dynamics (14), if the latter converge. That said, the convergence properties of (14) are *not* the same as those of (8), so it is not clear if (14) (and, hence, (13)) will end up solving the rate maximization problem (5). Nevertheless, we have:

Theorem 1. *Any solution trajectory* $\mathbf{Q}(t)$ *of* (14) *converges to a signal covariance profile which lies within* $\varepsilon(\tau)$ *of a solution of* (5)*; moreover, the approximation error* $\varepsilon(\tau)$ *becomes vanishingly small as* $\tau \to 0$.

Sketch of proof:[3] Consider the so-called "free energy" function $F(\mathbf{Q}) = \Psi(\mathbf{Q}) - \tau h(\mathbf{Q})$, where $h(\mathbf{Q}) = \operatorname{tr}[\mathbf{Q}\log\mathbf{Q}]$ denotes (the negative of) the von Neumann quantum entropy of $\mathbf{Q}$. Letting $\mathbf{V}_\tau = \mathbf{V} - \tau\log\mathbf{Q}$, a lengthy calculation along the lines of [9] can be used to show that $\dot{F} = \int_0^1 \operatorname{tr}[\mathbf{Q}^{1-s}\mathbf{V}_\tau\mathbf{Q}^s\mathbf{V}_\tau]\,ds - \operatorname{tr}[\mathbf{Q}\mathbf{V}_\tau]^2 \geq 0$, i.e. F is Lyapunov for (14). Moreover, since F is strictly concave and steep (the gradient of F approaches infinity near the boundary $\operatorname{bd}(\mathcal{X})$ of $\mathcal{X}$), it follows that (14) will converge to the unique maximizer $\mathbf{Q}_\tau^\omega$ of F over $\mathcal{X}$; that this maximizer lies within vanishing distance of a solution of (5) is then a consequence of Berge's maximum theorem [15]. ∎

By combining Lemma 1 with Theorem 1, we then obtain:

Theorem 2. *The damped exponential learning process* (13) *with imperfect measurements given by* (11) *converges to a signal covariance profile which lies within* $\varepsilon(\tau)$ *of a solution of the rate maximization problem* (5). *Furthermore, the approximation error* $\varepsilon(\tau)$ *vanishes as* $\tau \to 0$, *irrespective of the magnitude of the measurement noise* $\mathbf{\Xi}_k$.

Sketch of proof:[3] By Lemma 1 and our assumptions on the noise, the iterates $\mathbf{Y}_k(n)$ and the error processes $\mathbf{\Xi}_k(n)$ will both be bounded (a.s.). Moreover, by Theorem 1, the dynamics (14) admit a unique rest point (which obviously has measure zero in the set $\mathcal{X}$ of feasible covariance profiles). The theorem then follows from general stochastic approximation arguments and an appeal to Sard's theorem – see e.g. [11, Theorem 5.7]. ∎

Remark. Theorem 2 guarantees that the learning scheme (13) converges arbitrarily close to a solution of the rate maximization problem (5), irrespective of the noise level. Importantly, the quality of the solution is not controlled by the quality of the measurements, but by the damping parameter τ: as such, even though imperfect CSI slows down the users' convergence, it will not otherwise impact the end-state of the algorithm – in stark contrast to water-filling methods (cf. Section IV).

B. Asynchronous Updates and Delayed Information

Even though the learning scheme (13) with imperfect CSI converges arbitrarily close to an optimum signal covariance profile, it is not clear how it can be implemented in the absence of a centralized scheduler that could synchronize the users' update schedule. We will thus consider here a fully decentralized setting where each user updates his signal covariance matrix based on an individual timer, and independently of other users. In this case, the estimates for $\mathbf{V}_k$ that are calculated at each update might also suffer from delays due to asynchronicity, so Theorem 2 must be extended accordingly.

To account for all that, let n denote the n-th overall update epoch, let $\mathcal{K}_n \subseteq \mathcal{K}$ denote the subset of users who update at this epoch, and let $d_k(n)$ denote the number of steps elapsed between the update and measurement processes for the k-th user. The discretization (13) then becomes:

$$\mathbf{Y}_k(n+1) = \mathbf{Y}_k(n) + \gamma_{N_k(n)}\,\mathbb{1}(k \in \mathcal{K}_n)\cdot[\hat{\mathbf{V}}_k(n) - \tau\mathbf{Y}_k(n)],$$
$$\mathbf{Q}_k(n+1) = P_k\frac{\exp(\mathbf{Y}_k(n+1))}{\operatorname{tr}[\exp(\mathbf{Y}_k(n+1))]}, \tag{15}$$

where

$$\hat{\mathbf{V}}_k(n) = \mathbf{H}_k^\dagger(\mathbf{I} + \textstyle\sum_\ell \mathbf{H}_\ell\mathbf{Q}_\ell(n - d_\ell(n))\mathbf{H}_\ell^\dagger)^{-1}\mathbf{H}_k, \tag{16}$$

and $N_k(n)$ denotes the number of updates that have been performed by user k up to epoch n. As it turns out, this recursion is an asynchronous stochastic approximation of (14) in the sense of [12, Chap. 7], so, under some mild assumptions, it will converge (a.s.) to the same limit as the synchronous version (13). More precisely, we have:

Theorem 3. *If the delay functions* $d_k(n)$ *are bounded* (*a.s.*) *and the set of users* $\mathcal{K}_n$ *that update at step* n *is a homogeneous recurrent Markov chain (i.e. all users' update rates are finite and nonzero), then the conclusions of Theorem 2 continue to hold for the asynchronous learning process* (15) *with* $\gamma_n = 1/n$.

Sketch of proof:[3] Following [12, Thms. 2 and 3], the recursion (15) may be seen as a stochastic approximation of the rate-adjusted dynamics $\dot{\mathbf{Y}}_k = \eta_k[\mathbf{V}_k - \tau\mathbf{Y}_k]$, where $\eta_k = \lim_n N_k(n)/n > 0$ is the update rate of user k (recall that $\mathcal{K}_n$ is ergodic). The adjusted free energy $F = \Psi - \eta^{-1}\tau h$ is a Lyapunov function for these dynamics, so the claim follows in the same as way as that of Thm. 2. ∎

[3] Due to lack of space, the reader is referred to [14] for the full proofs.
[4] This is actually why the unadjusted process (10) may fail to converge.

IV. Numerical Results

In view of the above considerations, we obtain the following distributed algorithm (shown from the point of view of a single user):

Algorithm 1 Damped Exponential Learning (DXL).

Parameter: $\tau > 0$.
Initialize: $n \leftarrow 0$; $\mathbf{Y} \leftarrow 0$.
Repeat
  At each UpdateEvent
    $n \leftarrow n + 1$;
    calculate $\hat{\mathbf{V}} = \mathbf{H}^{\dagger}\mathbf{W}^{-1}\mathbf{H}$ based on latest observations;
    update score matrix $\mathbf{Y} \leftarrow \mathbf{Y} + \frac{1}{n}[\hat{\mathbf{V}} - \tau\mathbf{Y}]$;
    set $\mathbf{Q} \leftarrow P \exp(\mathbf{Y}) / \operatorname{tr}[\exp(\mathbf{Y})]$;
  until termination criterion is reached.

Remark 1. From an implementation point of view, DXL has the following desirable properties:

(P1) It is *distributed*: users have the same information requirements as in distributed water-filling.
(P2) It is *asynchronous*: there is no need for a global update timer to synchronize the network's users.
(P3) It is *stateless*: users do not need to know the state of the system (e.g. its topology).

Remark 2. The criteria that trigger an UpdateEvent could be arbitrary; that said, if there *is* a global update timer making an UpdateEvent occur simultaneously for all users, DXL boils down to the synchronized recursion (13).

Thanks to the analysis of the previous section, the convergence of DXL is guaranteed by Theorems 2 and 3; instead, in this section, our aim will be to assess the algorithm's performance via numerical simulations. To that end, we conducted extensive numerical simulations from which we illustrate here a selection of the most representative scenarios.

In Fig. 1, we simulated an uplink MIMO system consisting of a wireless receiver with 5 antennas and $K = 10$, 25, 50 or 100 transmitters, each with a random number m_k of transmit antennas between 2 and 6. Each user's channel matrix $\mathbf{H}_k$ was drawn from a complex Gaussian distribution at the outset of the transmission (but remained static once picked), and we then ran the DXL algorithm with synchronous updates (for benchmarking purposes) and a low damping parameter to ensure convergence to the system's capacity ($\tau = 10^{-3}$). The performance of the algorithm was assessed via the efficiency ratio $\operatorname{eff}(n) = [\Psi(n) - \Psi_{\min}]/[\Psi_{\max} - \Psi_{\min}]$ where $\Psi(n)$ denotes the users' sum rate at the n-th iteration of the algorithm and $\Psi_{\max}$ (resp. $\Psi_{\min}$) is the maximum (resp. minimum) value of Ψ. As can be seen in Fig. 1(a), even for large numbers of users, DXL effectively achieves the system's sum capacity within one or two iterations, a fact which represents a marked improvement over water-filling methods (Fig. 1(b)): on the one hand, IWF is significantly slower than DXL (it requires $\mathcal{O}(K)$ iterations to achieve the same performance level as the first iteration of DXL), whereas SWF fails to converge altogether.

The robustness of DXL is examined in Fig. 2 where we simulated an uplink MIMO system consisting of $K = 25$ transmitters with imperfect CSI and noisy measurements at the receiver. In particular, we plotted the efficiency of DXL over time for relative measurement error levels $\sigma = 10\%$ (low) and $\sigma = 100\%$ (high), and we ran the IWF and SWF algorithms with the same relative error levels for comparison. The performance of water-filling methods remains acceptable at low error levels (attaining 90–95% of the system's sum capacity); however, when the measurement noise gets higher, water-filling offers no perceptible advantage over the users' initial choice of covariance matrices. By contrast, DXL achieves the channel's capacity even when the error level is 100% – though, of course, its convergence speed is negatively affected.

Finally, to account for changing channel conditions, we also plotted the performance of DXL for fading channels following the well-known Jakes model [16]. More precisely, in Fig. 3, we simulated a system with 3 receive antennas, $K = 10$ users with 2 antennas each, transmitting at a frequency of $f = 2\,\text{GHz}$ and with average velocities of $v = 5\,\text{km/h}$ (pedestrian movement) corresponding to a channel coherence time of 108 ms. We then ran DXL with an update period of $\delta = 3\,\text{ms}$, and we plotted the achieved sum rate $\Psi(t)$ versus the maximum attainable sum rate $\Psi_{\max}(t)$ given the channel matrices $\mathbf{H}_k(t)$ at time t, and versus the sum rate that users could achieve by spreading their power uniformly over their antennas. As a result of its high convergence speed, DXL tracks the system's sum capacity remarkably well, despite the changing channel conditions; moreover, the sum rate difference between DXL and the uniform profile shows that this tracking is not an artifact of the system's sum capacity falling within a narrow band of what could be attained by spreading power uniformly over antennas.

V. Conclusions

In this paper, we introduced a damped exponential learning (DXL) process for rate maximization which achieves the sum capacity of uplink MIMO systems even under imperfect and delayed CSI. The DXL method converges arbitrarily close to the system's optimum transmit profile, and its convergence speed can be controlled by tuning the users' learning rate; as a result (and in contrast to traditional water-filling methods), DXL converges within a few iterations, even for large numbers of users and/or antennas per user, and even in the presence of arbitrarily large estimation errors or update asynchronicities. Importantly, DXL can be extended to a wide range of semidefinite problems; we focused here on the MIMO MAC for simplicity, but in the future, we aim to address more general channel models (such as the interference channel).

References

[1] G. J. Foschini and M. J. Gans, "On limits of wireless communications in a fading environment when using multiple antennas," *Wireless Personal Communications*, vol. 6, pp. 311–335, 1998.

[2] I. E. Telatar, "Capacity of multi-antenna Gaussian channels," *European Transactions on Telecommunications and Related Technologies*, vol. 10, no. 6, pp. 585–596, 1999.

[3] R. S. Cheng and S. Verdú, "Gaussian multiaccess channels with ISI: capacity region and multiuser water-filling," *IEEE Trans. Inf. Theory*, vol. 39, no. 3, pp. 773–785, May 1993.

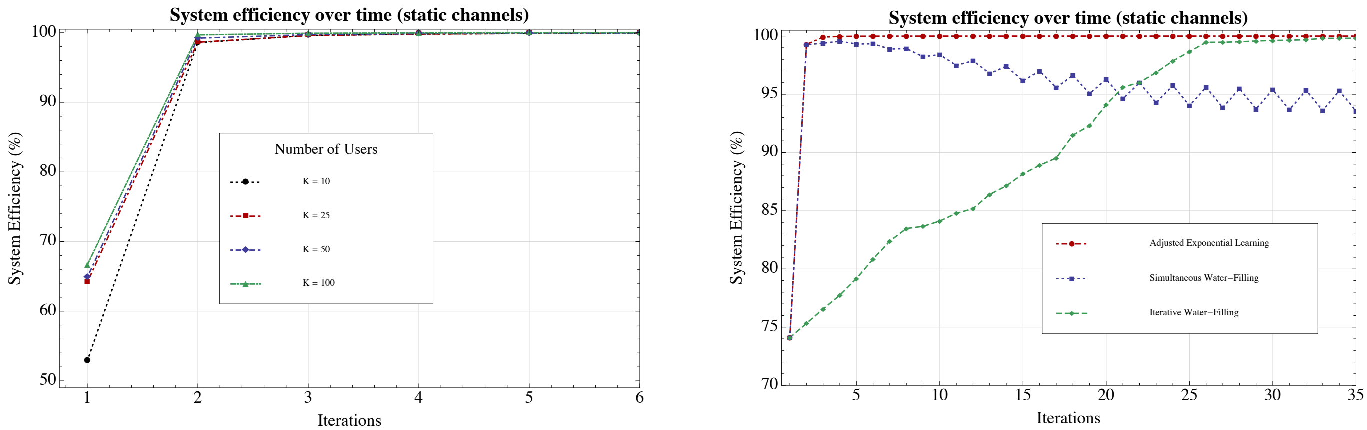

(a) The performance of DXL for different numbers of users.

(b) Adjusted exponential learning vs. water-filling for $K = 25$ users.

Fig. 1. The convergence speed of DXL as a function of the number of users compared to water-filling techniques.

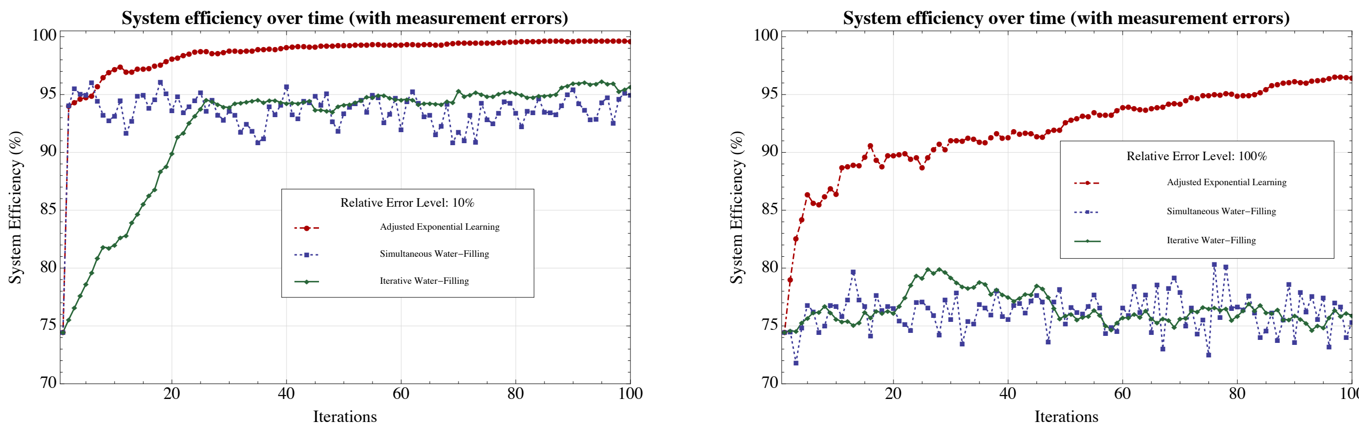

(a) Learning with an average relative error level of 10%.

(b) Learning with an average relative error level of 100%.

Fig. 2. The robustness of DXL under imperfect CSI: in contrast to water-filling, DXL converges even in the presence of very large measurement errors.

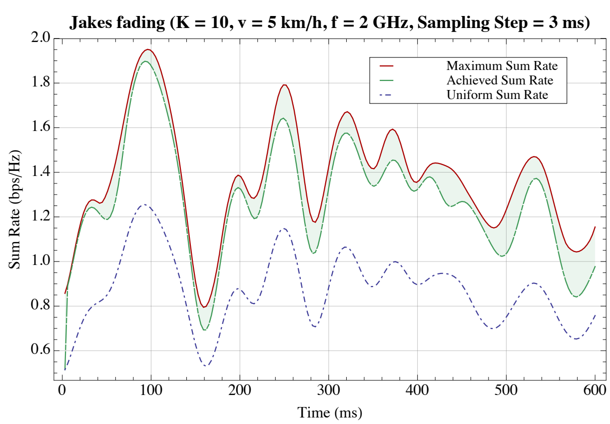

Fig. 3. The performance of DXL under changing channel conditions.

[4] W. Yu, W. Rhee, S. Boyd, and J. M. Cioffi, "Iterative water-filling for Gaussian vector multiple-access channels," *IEEE Trans. Inf. Theory*, vol. 50, no. 1, pp. 145–152, 2004.

[5] G. Scutari, D. P. Palomar, and S. Barbarossa, "Competitive design of multiuser MIMO systems based on game theory: a unified view," *IEEE J. Sel. Areas Commun.*, vol. 26, no. 7, pp. 1089–1103, September 2008.

[6] ——, "Simultaneous iterative water-filling for Gaussian frequency-selective interference channels," in *ISIT '06: Proceedings of the 2006 International Symposium on Information Theory*, 2006.

[7] P. Mertikopoulos, E. V. Belmega, A. L. Moustakas, and S. Lasaulce, "Dynamic power allocation games in parallel multiple access channels," in *ValueTools '11: Proceedings of the 5th International Conference on Performance Evaluation Methodologies and Tools*, 2011.

[8] ——, "Distributed learning policies for power allocation in multiple access channels," *IEEE J. Sel. Areas Commun.*, vol. 30, no. 1, pp. 96–106, January 2012.

[9] P. Mertikopoulos, E. V. Belmega, and A. L. Moustakas, "Matrix exponential learning: Distributed optimization in MIMO systems," in *ISIT '12: Proceedings of the 2012 IEEE International Symposium on Information Theory*, 2012, pp. 3028–3032.

[10] D. Lamberton, G. Pagès, and P. Tarrès, "When can the two-armed bandit algorithm be trusted?" *The Annals of Applied Probability*, vol. 14, no. 3, pp. 1424–1454, 2004.

[11] M. Benaïm, "Dynamics of stochastic approximation algorithms," *Séminaire de probabilités de Strasbourg*, vol. 33, 1999.

[12] V. S. Borkar, *Stochastic approximation*. Cambridge University Press and Hindustan Book Agency, 2008.

[13] P. Coucheney, B. Gaujal, and P. Mertikopoulos, "A class of robust entropy-driven learning procedures in games," Inria, Tech. Rep., 2013. [Online]. Available: http://hal.inria.fr/hal-00790815

[14] ——, "Distributed optimization in multi-user MIMO systems with imperfect and delayed information," Inria, Tech. Rep. RR-8426, 2013. [Online]. Available: http://hal.inria.fr/hal-00918762

[15] C. Berge, *Topological Spaces*. New York: Dover, 1997.

[16] G. Calcev, D. Chizhik, B. Göransson, S. Howard, H. Huang, A. Kogiantis, A. F. Molisch, A. L. Moustakas, D. Reed, and H. Xu, "A wideband spatial channel model for system-wide simulations," *IEEE Trans. Veh. Technol.*, vol. 56, no. 2, p. 389, March 2007.