

Adaptive Power Allocation and Control in Time-Varying Multi-Carrier MIMO Networks

Ioannis Stiakogiannakis, *Member, IEEE*, Panayotis Mertikopoulos, *Member, IEEE*, and Corinne Touati

Abstract

In this paper, we examine the fundamental trade-off between radiated power and achieved throughput in wireless multi-carrier, multiple-input and multiple-output (MIMO) systems that vary with time in an unpredictable fashion (e.g. due to changes in the wireless medium or the users' QoS requirements). Contrary to the static/stationary channel regime, there is no optimal power allocation profile to target (either static or in the mean), so the system's users must adapt to changes in the environment "on the fly", without being able to predict the system's evolution ahead of time. In this dynamic context, we formulate the users' power/throughput trade-off as an online optimization problem, and we provide a matrix exponential learning algorithm that leads to *no regret* – i.e. the proposed transmit policy is asymptotically optimal in hindsight, irrespective of how the system evolves over time. Furthermore, we also examine the robustness of the proposed algorithm under imperfect channel state information (CSI) and we show that it retains its regret minimization properties under very mild conditions on the measurement noise statistics. As a result, users are able to track the evolution of their individually optimum transmit profiles remarkably well, even under rapidly changing network conditions and high uncertainty. Our theoretical analysis is validated by extensive numerical simulations corresponding to a realistic network deployment, and providing further insights in the practical implementation aspects of the proposed algorithm.

Index Terms

Power allocation; MIMO; OFDMA; online optimization; no regret; matrix exponential learning.

I. INTRODUCTION

The wildfire spread of Internet-enabled mobile devices is putting existing wireless systems under enormous strain and is one of the driving forces behind the transition to next-generation mobile networks [1]. In this context, the

This research was supported in part by the French National Research Agency project NETLEARN (ANR-13-INFR-004) and by the European Commission in the framework of the FP7 Network of Excellence in Wireless COMMunications NEWCOM# (contract no. 318306). Part of this work was presented in Allerton 2014.

Ioannis Stiakogiannakis is with the Mathematical and Algorithmic Sciences Lab, France Research Center, Huawei Technologies Co. Ltd; this work was completed while he was with Inria, France. Panayotis Mertikopoulos is with the French National Center for Scientific Research (CNRS) and the Univ. Grenoble Alpes, LIG, F-38000 Grenoble, France; Corinne Touati is with Inria, France, and the Laboratoire d'Informatique de Grenoble, LIG, F-38000 Grenoble, France.

Email: ioannis.stiakogiannakis@huawei.com, panayotis.mertikopoulos@imag.fr, corinne.touati@inria.fr

Manuscript received March 7, 2015.

efficient control and allocation of radiated power comprises an indispensable aspect of wireless system design: for many applications (such as e-mail and voice calls), radiated power must be reduced to the bare minimum in order to preserve battery life; by contrast, for rate-hungry applications (such as multimedia streaming and video calling), it is crucial to optimize the allocation of the users' limited power across the network's degrees of freedom so as to maximize their throughput. In this way, wireless users are facing an important trade-off between radiated power and achieved throughput which must often be resolved in an adaptive and distributed manner, with minimal coordination between users.

In its most basic form, power control (PC) allows wireless links to achieve a target throughput while minimizing radiated power and the induced co-channel interference (CCI). Accordingly, power control has had a pivotal impact on wireless system design and operation ever since the early development stages of legacy wireless networks: starting with the pioneering work of Zander [2], Grandhi et al. [3], Foschini and Miljanic [4] and Yates [5], the design of efficient power control algorithms has given rise to a vast and extremely active corpus of literature – see e.g. [6] for a survey. Thus, in view of recent advances in MIMO technologies and the prolific deployment of orthogonal frequency-division multiple access (OFDMA) schemes, the envisioned transition to 5th generation (5G) mobile systems calls for power control algorithms tailored to networks with several degrees of freedom (spectral as well as spatial).

In this setting, most of the relevant literature has focused on maximizing the users' achievable transmission rate subject to their individual power constraints: [7–9] treat rate maximization as a constrained nonlinear optimization problem whereas [10–12] focus on multiple user interactions using game-theoretic methods; in a similar vein, [13–15] studied the power minimization problem subject to the users' rate requirements in multi-carrier multiple access channels (MACs), while [16] provided a two-layer framework for power minimization in MIMO–OFDMA systems. However, while the benefits of power control algorithms are relatively easy to assess in static networks, it is much harder to analyze their behavior in wireless systems that vary with time (e.g. due to user mobility, fading, temporal variations in the wireless medium, etc.). In the ergodic regime (where the users' channels follow a stationary ergodic process), [17, 18] provided power control algorithms that minimize the users' transmit power while achieving a minimum ergodic rate requirement. More recently, the authors of [19] studied the problem of ergodic rate maximization in fast-fading multi-carrier systems and they provided an efficient power allocation algorithm that allows users to attain the system's (ergodic) sum-capacity. However, when the wireless medium does not evolve according to an independent and identically distributed (i.i.d.) sequence of random variables, the efficient allocation and control of radiated power remains a very open issue.

In this paper, we drop all stationarity/i.i.d. assumptions and we focus squarely on wireless systems that evolve *arbitrarily* over time in terms of both channel conditions and user quality of service (QoS) requirements. In this framework, standard approaches based on linear programming (for static channels) and/or stochastic optimization (for the ergodic regime) are no longer relevant because there is no underlying optimization problem to solve – either static or in the mean. Instead, we treat power control as a *dynamically evolving optimization problem* and we employ techniques and ideas from online learning and optimization [20] to quantify how well the system's users can adapt to changes in the wireless medium.

The most widely used performance criterion in this setting is that of *regret minimization*, a seminal concept which was first introduced by Hannan [21] and which has since given rise to a vigorous literature at the interface of machine learning, optimization, statistics, and game theory – for a comprehensive survey, see e.g. [20, 22]. Specifically, in the language of game theory, the notion of regret compares a user’s cumulative payoff over a given time horizon to the cumulative payoff that he would have obtained by employing the *a posteriori* best possible action over the time horizon in question. Accordingly, in the context of power allocation and control, regret minimization corresponds to dynamic transmit policies that are asymptotically optimal in hindsight, *irrespective* of how the user’s environment and/or requirements evolve over time.

Regret minimization was recently used in [23] to study the transient phase of the Foschini–Miljanic (FM) power control algorithm in static environments and to propose alternative convergent power control schemes based on the notion of *swap regret* [24]. In [25], the authors considered a potential game formulation for the joint power control and channel allocation problem in cognitive radio (CR) networks and they employed a regret minimizing algorithm [26] to reach a Nash equilibrium state. The same problem was also examined in the context of infrastructureless wireless networks by the authors of [27] who formulated the problem as a potential game and provided a power control algorithm based on internal regret minimization that converges to the game’s unique correlated – and, hence, Nash – equilibrium. Finally, in a very recent paper, the authors of [28] employed online optimization methodologies to derive a dynamic transmit policy for online rate maximization in cognitive radio networks, but without attempting to control the users’ radiated power level.

Summary of results and paper outline

In this paper, we focus on multi-user MIMO–OFDMA systems that evolve arbitrarily over time (for instance, due to fading, intermittent user activity, changing QoS requirements, etc.), and we seek to provide an efficient power control and allocation scheme that allows users to balance their radiated power against their achieved throughput “on the fly”, based only on locally available (and possibly imperfect) CSI. In particular, we formulate the wireless users’ power minimization/throughput maximization trade-off as an online optimization problem and we derive a *no-regret* power control policy based on the method of matrix exponential learning (MXL) [29–31]. The proposed MXL algorithm is provably asymptotically optimal against the system’s evolution in hindsight; furthermore, it also enjoys the following desirable properties:

- *Distributedness*: users update their own power profiles based only on local information.
- *Asynchronicity*: there is no need for a global update timer to synchronize user updates.
- *Robustness*: the algorithm retains its properties even under imperfect CSI.
- *Statelessness*: transmitters do not need to know the network’s state and/or topology.

This work builds on (and significantly extends) our recent results on the regret minimization properties of the original Foschini–Miljanic dynamics in single-input and single-output (SISO), single-carrier systems that evolve continuously over time [32]. Compared to [32], the current paper represents an extension to multi-carrier systems with

several antennas (at both the transmitter and the receiver) and with imperfect feedback and channel state information at the transmitter (CSIT).

After presenting our wireless system model in Section II, the proposed algorithm for adaptive power control in MIMO-OFDMA systems is derived in Section III. Our main result therein is that the proposed algorithm leads to no regret; in addition, we examine the algorithm's behavior in the presence of imperfect CSI and we show that the algorithm retains its regret minimization properties almost surely, irrespective of the measurement noise level. Our theoretical analysis is supplemented by extensive numerical simulations in Section IV where we illustrate the power and throughput gains of the proposed power control algorithm under realistic network conditions.

II. SYSTEM MODEL AND PROBLEM FORMULATION

Consider a set $\mathcal{U} = \{1, \dots, U\}$ of wireless point-to-point connections formed over a set of orthogonal subcarriers $\mathcal{K} = \{1, \dots, K\}$; assume further that each connection $u \in \mathcal{U}$ comprises a transmit-receive pair (t_u, r_u) with M_u antennas at the transmitter and N_u antennas at the receiver. Thus, if $\mathbf{x}_k^u \in \mathbb{C}^{M_u}$ and $\mathbf{y}_k^u \in \mathbb{C}^{N_u}$ denote respectively the signals transmitted and received over connection u on subcarrier k , we obtain the familiar signal model:

$$\mathbf{y}_k^u = \mathbf{H}_k^{uu} \mathbf{x}_k^u + \sum_{v \neq u} \mathbf{H}_k^{vu} \mathbf{x}_k^v + \mathbf{z}_k^u \quad (1)$$

where $\mathbf{z}_k^u \in \mathbb{C}^{N_u}$ denotes the ambient noise over subcarrier k (including thermal, atmospheric and other peripheral interference effects) and $\mathbf{H}_k^{vu} \in \mathbb{C}^{N_u \times M_v}$ is the transfer matrix between t_v and r_u .

Unavoidably, the received signal \mathbf{y}_k^u is affected by the ambient noise and interference due to the transmissions of other connections on the same subcarrier, so we will write

$$\mathbf{w}_k^u = \sum_{v \neq u} \mathbf{H}_k^{vu} \mathbf{x}_k^v + \mathbf{z}_k^u \quad (2)$$

for the multi-user interference-plus-noise (MUI) at the receiver r_u of connection u (for a schematic representation, see Fig. 1); in this way, (1) attains the simpler form

$$\mathbf{y}_k^u = \mathbf{H}_k^{uu} \mathbf{x}_k^u + \mathbf{w}_k^u. \quad (3)$$

In particular, in what follows, we will focus on a specific connection $u \in \mathcal{U}$, so, for clarity, we will drop the index u altogether and we will write (3) even more compactly as:

$$\mathbf{y}_k = \mathbf{H}_k \mathbf{x}_k + \mathbf{w}_k. \quad (4)$$

In this context, assuming Gaussian input and noise and single user decoding (SUD) at the receiver (i.e. the multi-user interference by all other users is treated as additive noise), the transmission rate of the focal connection will be [33, 34]:

$$R(\mathbf{Q}) = \sum_{k \in \mathcal{K}} \left[\log \det \left(\mathbf{W}_k + \mathbf{H}_k \mathbf{Q}_k \mathbf{H}_k^\dagger \right) - \log \det \mathbf{W}_k \right], \quad (5)$$

where \mathbf{H}^\dagger denotes the Hermitian conjugate of \mathbf{H} and:

- $\mathbf{Q}_k = \mathbb{E}[\mathbf{x}_k \mathbf{x}_k^\dagger]$ is the $M \times M$ covariance matrix of the transmitted signal over subcarrier k .




Fig. 1. Example of a network with several active connections where we focus on a particular connection u between transmitter t_u and receiver r_u . The other active connections $v \in \mathcal{U} \setminus \{u\}$ cause co-channel interference to the focal connection $u \in \mathcal{U}$ which, together with the ambient subcarrier noise, is treated as additive colored noise.

- $\mathbf{Q} = \text{diag}(\mathbf{Q}_1, \dots, \mathbf{Q}_K)$ denotes the power profile of the focal transmitter over all subcarriers.
- $\mathbf{W}_k = \mathbb{E}[\mathbf{w}_k \mathbf{w}_k^\dagger]$ is the $N \times N$ MUI covariance matrix over subcarrier k .

In view of the above, let

$$\tilde{\mathbf{H}}_k = \mathbf{W}_k^{-1/2} \mathbf{H}_k \quad (6)$$

denote the user's *effective channel matrix* over subcarrier k . Then, Eq. (5) can be written as:

$$R(\mathbf{Q}) = \sum_{k \in \mathcal{K}} \log \det(\mathbf{I} + \tilde{\mathbf{H}}_k \mathbf{Q}_k \tilde{\mathbf{H}}_k^\dagger) \quad (7)$$

or, even more concisely:

$$R(\mathbf{Q}) = \log \det(\mathbf{I} + \tilde{\mathbf{H}} \mathbf{Q} \tilde{\mathbf{H}}^\dagger), \quad (8)$$

where the block-diagonal matrix $\tilde{\mathbf{H}} = \text{diag}(\tilde{\mathbf{H}}_1, \dots, \tilde{\mathbf{H}}_K)$ collects the user's effective channel matrices over all subcarriers $k \in \mathcal{K}$.

As we mentioned in the introduction, we focus on wireless users who seek to minimize their radiated power on the one hand while maximizing their transmission rate on the other. Thus, to account for this trade-off between transmit power and achieved throughput, we will consider the general power control objective:

$$\ell(\mathbf{Q}) = \text{tr}[\mathbf{Q}] - \phi(R(\mathbf{Q})) \quad (9)$$

where $\phi: \mathbb{R}_+ \rightarrow \mathbb{R}$ is a nondecreasing function of the user's achievable transmission rate $R(\mathbf{Q})$. By this token, $\ell(\mathbf{Q})$ can be interpreted as a “*loss function*” (or negative utility): higher values of $\ell(\mathbf{Q})$ indicate that the user is transmitting at very high power, at very low rate, or both, so he is incurring a “loss”. Accordingly, we will only assume that ϕ is Lipschitz continuous and concave: the former assumption is a mild technical requirement which we make for

simplicity, while the latter reflects the effects of “diminishing returns” on ever higher data rates (a rate increase from 1 bps to 2 bps is more impactful than an increase from 1,001 bps to 1,002 bps).

Remark 1. Utility-based formulations have a long history in the power control literature – see e.g. the recent papers [27, 35] for a related approach and [36, 37] for a similar formulation in terms of energy efficiency. Other possible approaches could involve achieving the Pareto frontier of the dual-objective trade-off between power minimization and throughput maximization; we focus on the specific model (9) on account of the model’s flexibility, generality and overall simplicity.

Remark 2. An important special case of the objective (9) concerns the scenario where the focal user seeks to minimize his transmit power $\text{tr}[\mathbf{Q}]$ subject to achieving a target transmission rate R^* . This classical formulation of power control can be recovered by considering a rate-adjustment function ϕ of the form $\phi(R) = f(R^* - R)$ with $f(r) = 0$ if $r \leq 0$ and $f(r) < 0$ otherwise – for instance, a standard choice would be to take $\phi(R) = -\lambda \cdot [R^* - R]_+$ for some $\lambda > 0$. In this way, when the target transmission rate is achieved (i.e. $R(\mathbf{Q}) \geq R^*$), the only term in the user’s loss function (9) is the user’s total transmit power $\text{tr}[\mathbf{Q}]$; otherwise, if the target transmission rate is not met, the user incurs an additional loss of at least $\phi'(0^-) \cdot (R^* - R(\mathbf{Q}))$.¹ In this way, the (positive) factor $\lambda = \phi'(0^-)$ represents the tolerance of the connection with respect to transmission rate deficits: smaller values of λ correspond to softer rate requirements, while, in the large λ limit, the loss function (9) stiffens to a hard constraint where no violations are tolerated.

In the above formulation, all sources of noise and co-channel interference by other users are collected in the effective channel matrix $\tilde{\mathbf{H}}_k$ of the focal connection; in this way, $\tilde{\mathbf{H}}_k$ collects all variables that are not under the direct control of the focal transmitter/receiver pair. As such, given that we make no assumptions on the behavior of the other connections in the network (or the evolution of the wireless medium itself), the matrix $\tilde{\mathbf{H}}$ may vary arbitrarily over time; our only assumptions will be as follows:

- (A1) $\tilde{\mathbf{H}}$ remains bounded for all time (e.g. due to RF circuit losses, antenna directivity, minimum distance between transmitter and receiver, etc.).
- (A2) The variability of $\tilde{\mathbf{H}}$ is such that standard results from information theory remain valid [33].

In this time-varying context, the throughput expression (5) becomes:

$$R(\mathbf{Q}; t) = \log \det [\mathbf{I} + \tilde{\mathbf{H}}(t) \mathbf{Q} \tilde{\mathbf{H}}^\dagger(t)], \quad (10)$$

where $\tilde{\mathbf{H}}(t)$ denotes the user’s effective channel matrix at time t . With this in mind, the user’s loss function at time t will be

$$\ell(\mathbf{Q}; t) = \text{tr}[\mathbf{Q}] - \phi(R(\mathbf{Q}; t); t). \quad (11)$$

We thus obtain the following online power control problem for MIMO–OFDMA systems:

$$\begin{aligned} & \text{minimize} && \ell(\mathbf{Q}; t), \\ & \text{subject to} && \mathbf{Q} \in \mathcal{X}, \end{aligned} \quad (\text{OPC})$$

¹Recall here that ϕ is assumed concave, so the user’s loss grows at least linearly with the rate deficit $R^* - R(\mathbf{Q})$.

where

$$\mathcal{X} = \{\mathbf{Q} : \mathbf{Q} \succeq 0, \text{tr}[\mathbf{Q}] \leq P\} \quad (12)$$

is the problem's state space and $P > 0$ denotes the user's maximum transmit power. More precisely, given that the user has no control over the effective channel matrices $\tilde{\mathbf{H}}$, the sequence of events that we envision is as follows:

- 1) At each update epoch $n = 1, 2, \dots$, the user selects a transmit power profile $\mathbf{Q}(n) \in \mathcal{X}$.
- 2) The user's loss $\ell(\mathbf{Q}(n); n)$ is determined by the state of the network and the behavior of all other users via the effective channel matrices $\tilde{\mathbf{H}}(n)$ at the time of the user's transmission.
- 3) The user selects a new transmit power profile $\mathbf{Q}(n+1) \in \mathcal{X}$ at stage $n+1$ in an effort to minimize the *a priori unknown* objective function $\ell(\mathbf{Q}; n+1)$ and the process repeats.

Needless to say, the key challenge in this dynamic framework is that the user does not know his objective function $\ell(\mathbf{Q}; n)$ ahead of time, so he must try to somehow adapt to the changing network conditions "on the fly" (recall that $\ell(\mathbf{Q}; n)$ depends at each stage n on the evolution of the environment and the choices of all other users). As a result, static solution concepts (such as Nash or correlated equilibria) are no longer relevant because, in general, there is no optimum system state to target – either static or in the mean.

Instead, given a time horizon T , we will compare the cumulative loss incurred by the user's power profile $\mathbf{Q}(n)$ for $n = 1, 2, \dots, T$, to the loss that the user would have incurred if he had chosen the *best possible transmit profile in hindsight*; specifically, we define the user's *regret* as:

$$\text{Reg}(T) = \max_{\mathbf{Q}^* \in \mathcal{X}} \sum_{n=1}^T [\ell(\mathbf{Q}(n); n) - \ell(\mathbf{Q}^*; n)]. \quad (13)$$

The seminal notion of regret was first introduced in a game-theoretic setting by Hannan [21] and it has since given rise to an extremely active field of research at the interface of optimization, statistics and machine learning – for a recent survey, see e.g. [20, 22].² The user's *average regret* is then defined as $T^{-1} \text{Reg}(T)$ and the goal of *regret minimization* is to devise a dynamic transmit policy $\mathbf{Q}(n)$ which is asymptotically optimal in hindsight, i.e. that leads to *no regret*:

$$\limsup_{T \rightarrow \infty} \text{Reg}(T)/T \leq 0, \quad \text{or, equivalently:} \quad \text{Reg}(T) = o(T), \quad (14)$$

irrespective of how the objective function (9) evolves over time.

Remark 3. Importantly, if the user's objective (9) does not vary with time (or if it varies in a stochastic fashion, following some i.i.d. process), a no-regret policy converges to the problem's static (or, respectively, average) solution [20]. Furthermore, if the user could predict the solution of (OPC) ahead of every stage $n = 1, 2, \dots, T$ in an oracle-like fashion, we would have $\text{Reg}(T) \leq 0$ in (13) for all T ; by this token, the no-regret requirement (14) is an indicator that $\mathbf{Q}(n)$ tracks the optimum solution $\mathbf{Q}^*(n)$ of (OPC) as it evolves over time.³

²The terminology stems from the fact that large positive values of $\text{Reg}(T)$ indicate that the user would have achieved a better power/rate trade-off in the past by employing some fixed \mathbf{Q}^* instead of $\mathbf{Q}(n)$, making him "regret" his choice.

³In the machine learning literature, there exist more sophisticated notions of regret (such as adaptive [38] or shifting [39] regret) that further quantify the quality of this tracking; due to space limitations however, we will focus our theoretical analysis almost exclusively on external regret minimization which requires less technical language to describe.

III. ADAPTIVE POWER CONTROL VIA EXPONENTIAL LEARNING

In this section, we derive an adaptive power control algorithm for the online optimization problem (OPC) based on the method of matrix exponential learning (MXL) [29–31]. We first consider the case where the transmitter has access to perfect channel state information (CSI); the case of measurement errors and imperfect CSIT is then discussed in Sec. III-B.

A. Learning with perfect CSI

A key element in our approach will be the gradient $\mathbf{V} = \nabla_{\mathbf{Q}} \ell$ of the user's objective function (9). Specifically, if the rate-adjustment function ϕ is smooth,⁴ we readily get:

$$\mathbf{V} = \nabla_{\mathbf{Q}} \ell = \mathbf{I} - \phi'(R) \cdot \nabla_{\mathbf{Q}} R. \quad (15)$$

Some matrix calculus then yields:

$$\nabla_{\mathbf{Q}} R = \tilde{\mathbf{H}}^\dagger [\mathbf{I} + \tilde{\mathbf{H}} \mathbf{Q} \tilde{\mathbf{H}}^\dagger]^{-1} \tilde{\mathbf{H}}, \quad (16)$$

so the gradient of $\ell(\mathbf{Q}(n); n)$ at $\mathbf{Q}(n)$ will be:

$$\mathbf{V}(n) = \mathbf{I} - \phi'(R(\mathbf{Q}(n); n)) \cdot \tilde{\mathbf{H}}^\dagger(n) [\mathbf{I} + \tilde{\mathbf{H}}(n) \mathbf{Q}(n) \tilde{\mathbf{H}}^\dagger(n)]^{-1} \tilde{\mathbf{H}}(n). \quad (17)$$

Since the effective channel matrices $\tilde{\mathbf{H}}(n)$ are assumed bounded, $\mathbf{V}(n)$ will also be bounded for all n ; hence, we formally assume that there exists a positive constant V such that

$$\|\mathbf{V}(n)\| \leq V, \quad (18)$$

where $\|\mathbf{V}\| = \lambda_{\max}(\mathbf{V})$ denotes the ordinary spectral norm (spectral radius) of \mathbf{V} .

In view of the above, a first idea would be to update the user's power profile $\mathbf{Q}(n)$ along the direction of steepest descent indicated by $\mathbf{V}(n)$ [40]; however, this online gradient descent scheme would invariably violate the user's semidefiniteness constraint $\mathbf{Q} \succeq 0$, so it is not a viable transmit policy. Instead, inspired by the matrix regularization methods of [29–31], we propose an algorithm that tracks the direction of steepest descent in a dual, unconstrained space and then maps the result back to the problem's state space via matrix exponentiation. More precisely, assuming for the moment perfect CSIT, we will consider the matrix exponential learning scheme:

$$\begin{aligned} \mathbf{Y}(n) &= \mathbf{Y}(n-1) - \mathbf{V}(n), \\ \mathbf{Q}(n+1) &= P \frac{\exp(\eta n^{-1/2} \mathbf{Y}(n))}{1 + \text{tr}[\exp(\eta n^{-1/2} \mathbf{Y}(n))]}, \end{aligned} \quad (\text{MXL})$$

where $\eta > 0$ is a parameter that controls the user's learning rate and the recursion is initialized with $\mathbf{Y}(0) = 0$.

The recursion (MXL) will be the main focus of our paper, so some remarks are in order (for an algorithmic implementation, see Alg. 1):

Remark 1. Intuitively, the exponentiation step in (MXL) assigns more power to the spatial directions that perform well while the $n^{-1/2}$ factor keeps the eigenvalues of $\mathbf{Q}(n)$ from approaching zero too fast (note that $\mathbf{Y}(n)$ grows as $\mathcal{O}(n)$); the

⁴In the general Lipschitz case, it suffices to replace $\phi'(R)$ by any element of $[\phi'(0^-), \phi'(0^+)]$.

Algorithm 1: Matrix exponential learning (MXL)

```

parameter:  $\eta > 0$ 

/* Initialization */
1  $n \leftarrow 0; \mathbf{Y} \leftarrow \mathbf{0};$ 

2 repeat
3    $n \leftarrow n + 1;$ 
   /* Pre-transmission: Set Power */
4    $\mathbf{Q} \leftarrow P \frac{\exp(\eta n^{-1/2} \mathbf{Y})}{1 + \text{tr}[\exp(\eta n^{-1/2} \mathbf{Y})];}$ 
   /* Transmission */
   /* Post-transmission: Measure Rate and Effective Channel Matrices */
5    $R \leftarrow \log \det (\mathbf{I} + \tilde{\mathbf{H}} \mathbf{Q} \tilde{\mathbf{H}}^\dagger);$ 
6    $\mathbf{V} \leftarrow \mathbf{I} - \phi'(R) \cdot \tilde{\mathbf{H}}^\dagger (\mathbf{I} + \tilde{\mathbf{H}} \mathbf{Q} \tilde{\mathbf{H}}^\dagger)^{-1} \tilde{\mathbf{H}};$ 
7    $\mathbf{Y} \leftarrow \mathbf{Y} + \mathbf{V};$ 
8 until transmission ends;

```

trace normalization then ensures that $\mathbf{Q}(n)$ satisfies the feasibility constraints of (OPC) for all $n \geq 1$. In particular, as we show in Appendix A, the recursion (MXL) can be seen as a “primal-dual” online mirror descent (OMD) method [20] with a variable parameter [41]; for an in-depth discussion, see [20, 29–31, 41] and references therein.

Remark 2. From an implementation viewpoint, Algorithm 1 has the following desirable properties:

- (P1) It is *distributed*: each transmitter updates his own power profile based only on local CSI.
- (P2) It is *asynchronous*: the algorithm’s updates are event-based and can be performed without synchronization or any further signaling/coordination between connections.
- (P3) It is *agnostic*: transmitters do not need to know the status or geographical distribution of other connections in the network.
- (P4) It is *reinforcing*: each connection tends to minimize its individual loss.

Remark 3. In terms of feedback, Algorithm 1 requires that *a*) transmitters measure their achieved rates; and *b*) the receiver feeds back to the transmitter the received signal covariance $\mathbb{E}[\mathbf{y}\mathbf{y}^\dagger] = \mathbf{W} + \mathbf{H}\mathbf{Q}\mathbf{H}^\dagger$ (e.g. via broadcasting or over a duplex downlink). From a computational standpoint, it is then easy to see that the complexity of each iteration of Algorithm 1 is linear in the number of subcarriers K and polynomial in the number of transmit antennas M :⁵ in particular, since \mathbf{Y} is block-diagonal, fast Coppersmith–Winograd matrix multiplication [42] provides a worst-case complexity bound that is $\mathcal{O}(KM^{2.373})$ per iteration.

⁵We are implicitly assuming that $\phi'(R)$ can be calculated with very low cost – e.g. by means of a lookup table.

Our main theoretical result regarding the MXL power control algorithm (Alg. 1) is as follows:

Theorem 1. *The MXL algorithm (Alg. 1) leads to no regret in the online power control problem (OPC). In particular, the iterates of (MXL) enjoy the $\mathcal{O}(T^{-1/2})$ regret bound:*

$$\frac{1}{T} \text{Reg}(T) \leq \left(\frac{P \log(1 + KM)}{\eta} + \frac{\eta PV^2}{2} \right) \frac{1}{\sqrt{T}} + \frac{\eta PV^2}{4} \frac{1}{T}, \quad (19)$$

irrespective of how the system evolves over time.

Proof: See Appendix B. ■

The bound (19) is our main performance guarantee for Alg. 1, so we proceed with a few remarks:

Remark 4. Even though Theorem 1 focuses on a given connection $u \in \mathcal{U}$, the focal connection is still subject to interference from other connections in the network (the incurred interference is captured by the effective channel matrices $\widetilde{\mathbf{H}}_k$ which depend on the interfering users' transmit policies). In this light, Theorem 1 provides a worst-case performance guarantee which holds even in the presence of malicious users (jammers) that seek to shut down the focal connection.

On the other hand, a natural question that arises is whether users can meet more sophisticated criteria (such as reaching a globally efficient state or a Nash equilibrium) when they all follow the same algorithm and the wireless medium is otherwise static. In the MIMO multiple access channel (where all users transmit to a common receiver), it can be shown that the MXL algorithm leads to a socially optimum state; a more general treatment of this question (e.g. in the MIMO interference channel [10]) lies beyond the scope of this paper, so we delegate it to future work.

Remark 5. We should also note here that the first term of the bound (19) captures the dimensionality of the problem while the rest is an increasing function of the channel variability estimate V ; as such, the learning parameter η of Algorithm 1 can be fine-tuned to accelerate the algorithm's convergence to a no-regret state in terms of V . Specifically, the value of η which minimizes the dominant $\mathcal{O}(T^{-1/2})$ term of the regret bound (19) for a fixed time horizon T is:

$$\eta = V^{-1} \sqrt{2 \log(1 + KM)}. \quad (20)$$

In turn, this parameter choice leads to the optimized convergence rate:

$$\frac{1}{T} \text{Reg}(T) \leq PV \sqrt{2 \log(1 + KM)} \left(\frac{1}{\sqrt{T}} + \frac{1}{4T} \right). \quad (21)$$

The $\mathcal{O}(1/\sqrt{T})$ dependence of (21) is known to be asymptotically tight in the context of online optimization problems against an adversarial nature [20], while the $\mathcal{O}(\log KM)$ behavior represents a significant reduction in the dimensionality of the problem (which has $\mathcal{O}(KM)$ degrees of freedom). In fact, (21) becomes tight only in adversarial environments (e.g. induced by jamming), so, in practical situations, the user's regret minimization rate is considerably faster – cf. Section IV.

Remark 6. The agnostic initialization $\mathbf{Y}(0) = 0$ is a conservative choice reflecting the worst-case scenario where the user assumes bad channel conditions. Indeed, $\mathbf{Y}(0) = 0$ corresponds to initial transmit power equal to $P \cdot KM / (1 + KM) \sim P$ in the large K (or large M) limit; in this way, the user's transmit power will likely be reduced under

Algorithm 1 in the presence of good channel conditions. Hence, if the transmitter has some estimate of his expected channel conditions, it would be preferable to initialize power accordingly: if the user expects a good channel, initial power should be set lower (to save battery life); otherwise, if a bad channel is expected, initial transmit power should be set high so as to avoid very low transmission rates in the first few frames.

B. Adaptive power control with imperfect CSI

In practice, a major challenge occurs if the transmitters do not have access to perfect CSI with which to update the adaptive power control scheme (MXL). In particular, given that each user's gradient matrix \mathbf{V} is determined by his effective channel matrix $\tilde{\mathbf{H}}$, imperfect measurements of the users' channel or the multi-user interference-plus-noise (due e.g. to pilot contamination, undersampling or other factors) could have a catastrophic effect on the no-regret properties of the proposed scheme (MXL). Accordingly, our goal in this section will be to examine the robustness of (MXL) in the presence of measurement errors and observation noise.

To model errors of this kind, we assume that, at each update period $n = 1, 2, \dots$, the transmitter observes a noisy estimate $\hat{\mathbf{V}}(n)$ of the form

$$\hat{\mathbf{V}}(n) = \mathbf{V}(n) + \mathbf{Z}(n), \quad (22)$$

where the error process $\mathbf{Z}(n) = \text{diag}(\mathbf{Z}_1(n), \dots, \mathbf{Z}_K(n))$ satisfies the statistical hypotheses:

(H1) *Unbiasedness*:

$$\mathbb{E}[\mathbf{Z}(n) | \mathbf{Q}(n-1)] = 0. \quad (H1)$$

(H2) *Tame tails*:

$$\mathbb{P}(\|\mathbf{Z}(n)\| \geq z) \leq A/z^\alpha \quad \text{for some } A > 0 \text{ and for some } \alpha > 4. \quad (H2)$$

The unbiasedness assumption (H1) is a bare-bones assumption which simply boils down to asking that there is no biased, *systematic* error in the user's CSI measurements. Likewise, (H2) posits a fairly mild control on the probability of observing very high errors, and is satisfied by the vast majority of statistical error distributions (including for instance uniformly distributed, Gaussian, log-normal, Weibull and Lévy-type error processes); in particular, *we do not assume* that the measurement errors $\mathbf{Z}(n)$ are i.i.d., state-independent, or even a.s. bounded.

Importantly, under these mild hypotheses for the statistics of the measurement noise, we have:

Theorem 2. *The MXL algorithm (Alg. 1) run with imperfect observations satisfying (H1) and (H2) leads to no regret (a.s.); in particular, it enjoys the mean regret bound:*

$$\mathbb{E} \left[\frac{1}{T} \text{Reg}(T) \right] \leq \left(\frac{P \log(1 + KM)}{\eta} + \frac{\eta P \bar{V}^2}{2} \right) \frac{1}{\sqrt{T}} + \frac{\eta P \bar{V}^2}{4} \frac{1}{T}, \quad (23)$$

where $\bar{V}^2 = \sup_n \mathbb{E}[\|\hat{\mathbf{V}}(n)\|^2 | \mathbf{Q}(n-1)]^2$.

Proof: See Appendix C. ■

Remark 7. From an implementation perspective, we should note here that the mean bound (23) reduces to the deterministic bound (19) in the case of perfect CSI. Also, even though we have $\limsup_{T \rightarrow \infty} T^{-1} \text{Reg}(T) \leq 0$ (a.s.),

the realized regret of Alg. 1 may exceed the mean bound (19) with positive probability. By a concentration inequality argument [43], it is possible to estimate analytically the probability of such deviations in terms of the central moments of the error process, but this analysis would take us too far afield so we do not present it here.

Remark 8. Hypothesis (H2) implies that the error process \mathbf{Z} has finite (central) moments of up to fourth order – in fact, barring pathological examples, this requirement is essentially tantamount to (H2). The importance of fourth order moments has to do with the fact that we are using a variable learning parameter that decays as $n^{-1/2}$; by choosing a slower decay rate of the form $n^{-\gamma}$ for some $\gamma \in (0, 1/2)$, it is possible to relax Hypothesis (H2) down to second order moment control. However, given that (H2) already suffices for the framework at hand (and due to space limitations), we do not present this more general analysis here.

IV. NUMERICAL RESULTS

To validate the theoretical analysis of Section III, we conducted extensive numerical simulations over a wide range of design parameters and specifications. In what follows, we present a representative subset of these results, but the conclusions drawn remain valid in most typical mobile wireless environments.

Throughout this section, we consider a typical cellular OFDMA wireless network that occupies a 10 MHz band divided into 1024 subcarriers around a central frequency $f_c = 2.5$ GHz. We further assume that each cell employs a simple randomized access algorithm [44] to allocate subcarriers to the users it serves. In the following, we focus on $U = 4$ users that are located at different cells – served by different base stations (BSs) – and that have been allocated the same set of $K = 8$ subcarriers. We focus on the uplink (UL) case, so the receivers are assumed stationary whereas the transmitters may be either stationary or mobile, depending on the simulated scenario. Communication occurs over a time-division duplexing (TDD) scheme with frame duration $T_f = 5$ ms: specifically, transmission occurs during the UL subframe while receivers process the transmitted signal and provide feedback during the downlink (DL) subframe; upon reception of the feedback, transmitters update their transmit powers according to Algorithm 1, and the process repeats until transmission ends. For demonstration purposes, we simulated the case where each connection has a fixed rate requirement R_u^* which varies across connections $u \in \mathcal{U}$ so as to ensure diversity of QoS requirements (the users' tolerance and loss function is defined as indicated in Remark 2). For convenience, all simulation parameters are summarized in Table I.

For benchmarking purposes, the first simulated scenario focuses on the case where channels remain static during the transmission horizon. In Fig. (2a), we plot the evolution of the users' objective $\ell(\mathbf{Q}; n)$ under Algorithm 1: as can be seen, users quickly reach an optimal state corresponding to the minimum of their loss function (i.e. minimum transmit power subject to the users' rate requirements). In particular, as we see in Fig. (2b), even though all connections start with excessive transmit power (due to the algorithm's conservative initialization), they converge within 2 dB of their optimum transmit profile within a few frames (between 5 and 15, depending on the connection). Interestingly, we also see some slight power oscillations (of the order of 1 dB) that persist for a few frames after the initial ones: these are due to small violations of the users' rate requirements (due to the power updates of other users) that cause them to momentarily increase their transmit power. Similar oscillations are observed with respect to the achieved/target rate

TABLE I
OFDMA NETWORK SIMULATION PARAMETERS

Number of Cells	19
Cell Radius	1 km
Central Frequency	2.5 GHz
Available Bandwidth	10 MHz
Number of OFDM Subcarriers	1024
Subcarrier Spacing	10.9375 kHz
Transmit Antennas	2
Receive Antennas	2
Propagation Model	COST-Hata-Model
BS Antenna Height	32 m
MS Antenna Height	1.5 m
Shadowing	8.9 dB
AWGN Spectral Power Density	$N_0 = -174$ dBm/Hz
Receiver Noise Figure	7 dB
Frame duration	5 ms
Requested Bit Rate per User	{764.6, 113.7, 909.3, 1081.3} kbps
Maximum transmit power per User	{40.40, 41.10, 42.85, 45.58} dBm

gap $R(\mathbf{Q}; n)/R^*$ depicted in Fig. (2c): users quickly get within 2–5% of their target value, but they oscillate slightly for a few frames before converging. Finally, in Fig. (2d), we plot the user’s average regret $T^{-1} \text{Reg}(T)$ (solid lines) along with the theoretical bound predicted by Theorem 1 (dashed lines). To increase resolution, we plot the users’ regret in a logarithmic scale: in this way, the observed vertical drops to $-\infty$ correspond to the point where the users’ regret becomes negative (an indication of the number of frames required for the algorithm to converge). In tune with the above observations, we see that users only require a few frames to achieve a no-regret state.

The second simulated scenario examines the case of imperfect CSI. Specifically, in Fig. 3, we consider the same network realization as in Fig. 2, but we no longer assume that transmitters receive perfect CSI during the TDD feedback loop; instead, we assume imperfect channel state measurements and we plot the users’ power, rate and regret under Algorithm 1 with noisy observations. In particular, the transmitters’ CSI deviates from its corresponding mean value with standard deviation of $0.50V$ for connections 1 and 3, and $1.00V$ for connections 2 and 4, respectively. Due to this huge uncertainty, users are more conservative and tend to use up more power to achieve their rate requirements; however, after an initial sampling period (lasting a few tens of frames), they confidently reduce power and converge to an optimum rate/power trade-off (as evidenced by the minimization of their objective). A similar behavior is observed in Fig. (3b) which shows the evolution of the users’ throughput over time: even though there are more pronounced fluctuations over the first few frames, all connections eventually converge to their target rates. The main performance degradation is in the algorithm’s convergence time: as can be seen in Fig. 3c, Algorithm 1 takes longer to converge to a no-regret state, chiefly due to the regret generated during the algorithm’s training phase.

Finally, in Fig. 4, we simulate a realistic time-varying environment where the focal transmitters move at different

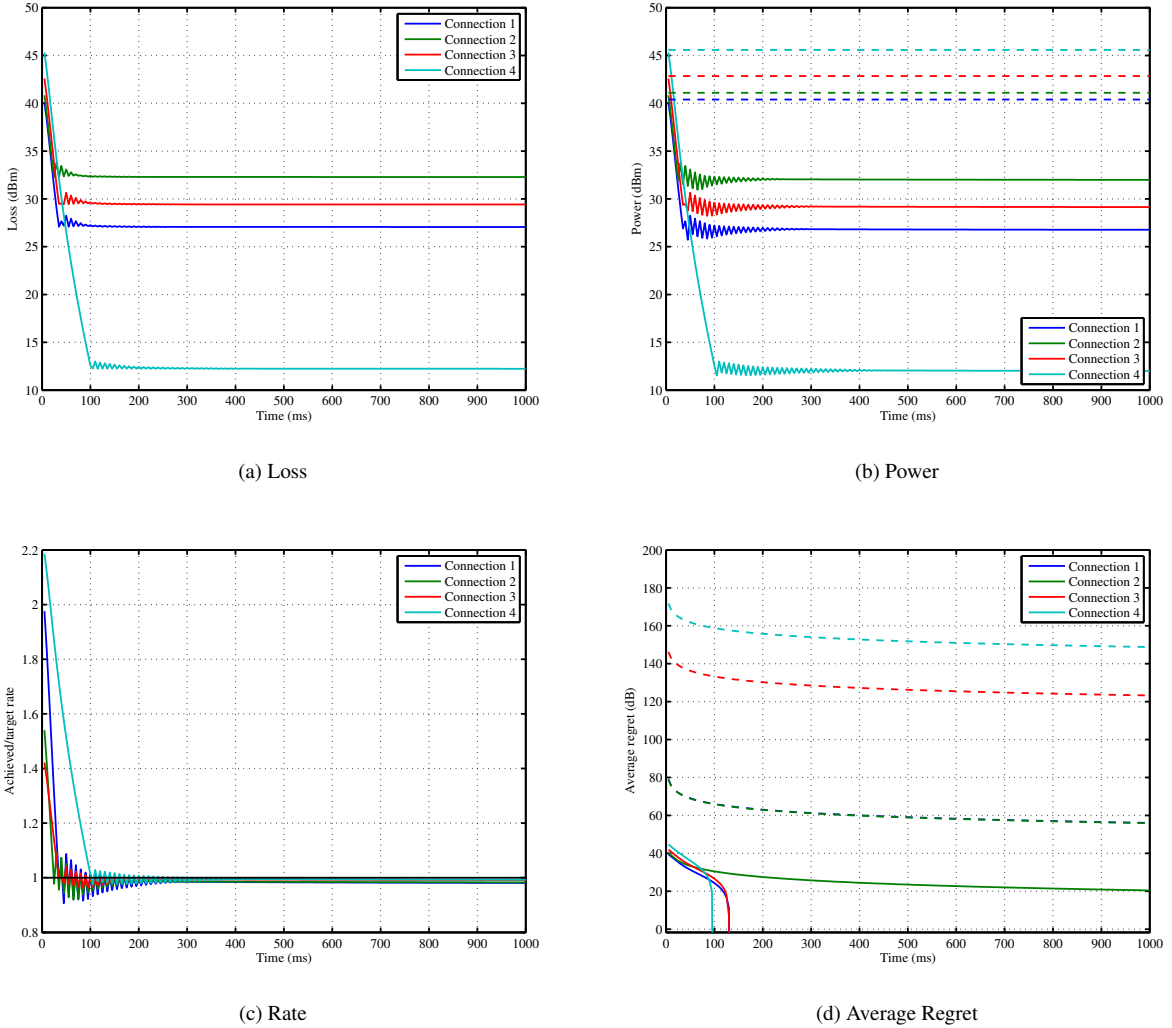


Fig. 2. Adaptive MIMO-OFDM power control under total power constraints for $U = 4$ connections with static channel conditions. Fig. 2a depicts the evolution of the users' objective function $\ell(\mathbf{Q}; n)$ under the online power control algorithm (MXL). Similarly, Fig. 2b shows the evolution of the users' total transmit power $\text{tr}[\mathbf{Q}(n)]$ (solid lines; dashed lines correspond to the users' maximum transmit power); Fig. 2c shows the achieved/target rate gap $r(\mathbf{Q}; n)/R^*$. Finally, the average regret $n^{-1} \text{Reg}(\mathbf{Q}^*; n)$ is plotted in Fig. 2d (solid lines), along with the theoretical bounds predicted by Theorem 1 (dashed lines); for simplicity, we only plot the positive part in the regret and we use a logarithmic dB scale for consistency.

speeds. For simulation purposes, we used the extended pedestrian A (EPA), extended vehicular A (EVA), and extended typical urban (ETU) channel models for pedestrian (2km/h), urban vehicular (30km/h) and high speed (130km/h) users respectively [45]. To illustrate the variability of the users' channels, we plot their evolving channel gains $\text{tr}[\mathbf{H}(n)\mathbf{H}^\dagger(n)]$ in Fig. (4a): as can be seen, channel variations are quite wide and become more profound for higher user velocities.

In this dynamic setting, the main challenge for the users is to track the optimum signal covariance profile that balances their transmit power against their achieved throughput (i.e. that minimizes their loss) as this optimum profile




Fig. 3. Adaptive MIMO-OFDM power control for $U = 4$ connections with static channel conditions but imperfect CSIT (relative measurement errors depicted in each figure's legend). Figs. 3a and 3b respectively depict the evolution of the users' radiated power $\text{tr}[\mathbf{Q}(n)]$ and achieved/target rate gap $r(\mathbf{Q}; n)/R^*$ under the online power control algorithm (MXL). The users' average regret $n^{-1} \text{Reg}(\mathbf{Q}^*; n)$ is plotted in Fig. 3c (solid lines), along with the theoretical bounds predicted by Theorem 1 (dashed lines); for simplicity, we only plot the positive part of the regret and we use a logarithmic dB scale for consistency.

evolves over time. To that end, Fig. 4b shows that the users' radiated power under Algorithm 1 increases (to compensate for poor channel conditions) or decreases (when channel conditions are more favorable) in a way consistent with the evolution of the wireless medium (Fig. 4a). Dually, in Fig. 4c we plot the time average of the users' achieved/target rate ratio:⁶ as can be seen, users consistently achieve their target throughput, and their achieved/target throughput ratio converges to 1 over time (in practice, within a few frames for users that do not move at very high speeds). Furthermore,

⁶Time-averages are considered in order to weed out stochastic fluctuations (due to the users' changing fading environment) that could be potentially misleading.

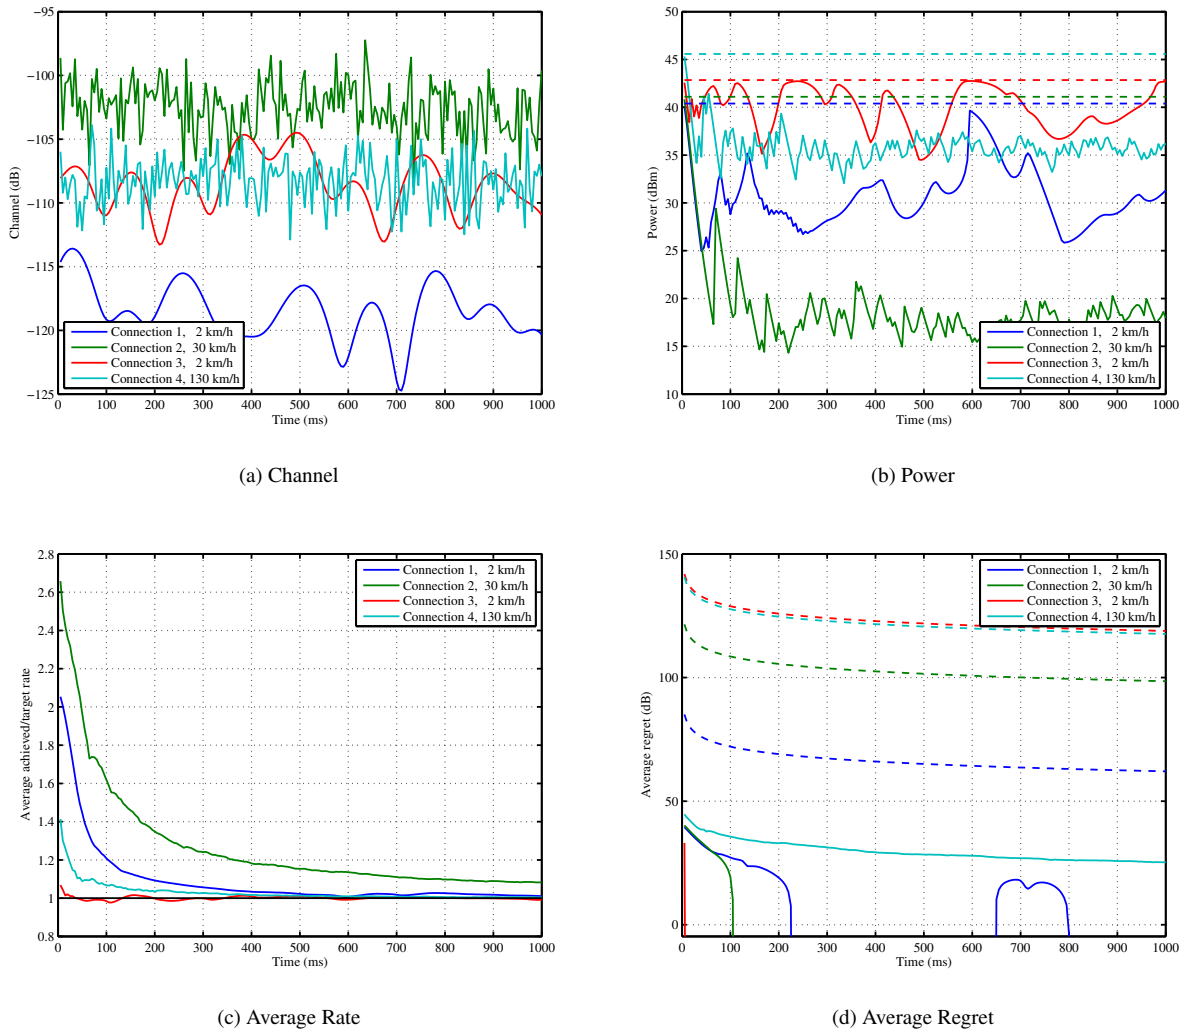


Fig. 4. Adaptive MIMO-OFDM power control for $U = 4$ connections with time-varying channel conditions corresponding to stationary receivers and mobile transmitters (average speed as in each figure's legend). For reference purposes, Fig. 4a depicts the evolution of the channel gains $\text{tr}[\mathbf{H}(t)\mathbf{H}^\dagger(t)]$ over time. Fig. 4b shows the evolution of the users' total transmit power under Algorithm 1 (dashed lines represent the users' maximum transmit power), while Fig. 4c shows the achieved/target rate gap $R(t)/R^*$. Finally, as in the static channel case, Fig. 4d shows the users' average regret $\text{Reg}(T)/T$: as predicted by Theorem 1, the users' regret quickly becomes negative, indicating that their transmit policy is asymptotically optimal in hindsight (for simplicity, we only plot the positive part of the regret and we use a logarithmic dB scale for consistency).

we see that connections with a softer tolerance for the satisfaction of their QoS requirements (e.g. Connection 1) are very aggressive in reducing transmit power when channel conditions seem to allow it, whereas connections that are less tolerant with respect to their QoS requirements (e.g. Connection 2) are more conservative and transmit at relatively high powers (resulting in higher rates) as a precaution against deep fading events.

Finally, as in the static channel case, Fig. 4d depicts the users' average regret over time: again, despite the pessimistic high-power initialization of Algorithm 1, the users' regret drops to the no-regret regime in just a few frames (much

faster than the $\mathcal{O}(1/T)$ bounds predicted by Theorem 1). The reason for this faster convergence is that the worst-case bounds of Theorem 1 only become relevant under very adverse (or adversarial) environments, occurring for example when users are being jammed by a third party: in standard mobility scenarios (such as the one simulated here), the evolution of the wireless medium is relatively tame from a statistical perspective, so users adapt to its variability much faster than in the adversarial regime.

V. CONCLUSIONS

In this paper, we examined the trade-off between radiated power and achieved throughput in wireless MIMO-OFDMA systems that evolve dynamically over time as the result of changing channel conditions and user QoS requirements. To account for the system's complete lack of stationarity (or any other type of average behavior that could allow the use of traditional solution concepts such as Nash/correlated equilibria), we provided a formulation based on online optimization and we derived an adaptive matrix exponential learning algorithm that leads to *no regret* – i.e. that is asymptotically optimal in hindsight, irrespective of how the wireless system varies with time. Importantly, the proposed algorithm requires only local CSIT and is robust with respect to measurement errors and imperfections: in particular, under fairly mild hypotheses for the uncertainty statistics, the proposed algorithm retains its regret minimization properties and converges to a no-regret state. As a result, thanks to the algorithm's no regret property, the system's users are able to track their optimal transmit power profile “on the fly”, even under randomly changing channel conditions and high uncertainty.

The proposed algorithmic framework can be readily extended to different precoding schemes (such as MMSE or ZF-type precoders), or to account for other transmission features such as spectral mask constraints, pricing, etc. Through judicious use of convexification techniques, it can also be applied to non-convex energy-efficiency objectives, such as the users' achieved throughput per Watt of radiated power; we intend to explore these directions in future work.

APPENDIX

TECHNICAL PROOFS

Our goal in this appendix is to prove the regret guarantees of (MXL) under both perfect and imperfect CSI (Theorems 1 and 2 respectively). Drawing on the approach of [41, 46], we will first establish the no-regret properties of Algorithm 1 in a continuous-time, “mean-field” setting, and we will then show that these properties descend to discrete time at the cost of an extra term in the algorithm's regret guarantees. The algorithm's robustness properties with respect to measurement noise and errors will then follow by using the theory of concentration inequalities.

For notational clarity and convenience, we will be suppressing the dependence on time whenever possible, and we will write e.g. $\dot{\mathbf{Q}}$ instead of $\frac{d}{dt}\mathbf{Q}(t)$ when there is no ambiguity.

A. No regret in continuous time

We begin by considering the following continuous-time analogue of the basic power control algorithm (MXL):

$$\begin{aligned} \dot{\mathbf{Y}} &= -\mathbf{V}, \\ \mathbf{Q} &= P \frac{\exp(\eta(t)\mathbf{Y})}{1 + \text{tr}[\exp(\eta(t)\mathbf{Y})]}, \end{aligned} \tag{MXL-c}$$

where $\eta(t) > 0$ is a smooth, nonincreasing learning parameter and the gradient matrix \mathbf{V} is defined as in (17). The following proposition shows that (MXL-c) leads to no regret in continuous time:

Proposition 3. *The learning scheme (MXL-c) guarantees the continuous-time regret bound:*

$$\max_{\mathbf{Q}^* \in \mathcal{X}} \int_0^T [\ell(\mathbf{Q}(t); t) - \ell(\mathbf{Q}^*; t)] dt \leq \frac{P \cdot \log(1 + KM)}{\eta(T)}, \tag{24}$$

for any measurable stream of effective channel matrices $\tilde{\mathbf{H}}(t)$, $t \geq 0$. In particular, if $\eta(t)$ satisfies the decay rate condition $\lim_{t \rightarrow \infty} t \cdot \eta(t) = \infty$, the learning scheme (MXL-c) leads to no regret.

Proof: We first note that the loss function $\ell(\mathbf{Q}; t)$ is convex with respect to \mathbf{Q} (to see this, simply recall that the Shannon rate function $R(\mathbf{Q}; t)$ is concave in \mathbf{Q} [47] while ϕ is assumed concave and nondecreasing). With this in mind, we obtain:

$$\ell(\mathbf{Q}(t); t) - \ell(\mathbf{Q}^*; t) \leq \text{tr}[(\mathbf{Q}(t) - \mathbf{Q}^*) \cdot \mathbf{V}(t)], \tag{25}$$

where $\mathbf{V}(t) = \nabla_{\mathbf{Q}(t)} \ell(\mathbf{Q}(t); t)$ denotes the gradient of $\ell(\cdot; t)$ evaluated at $\mathbf{Q}(t)$. Accordingly, to establish the no-regret bound (24) for (MXL), it suffices to show that

$$\int_0^T \text{tr}[(\mathbf{Q}(t) - \mathbf{Q}^*) \cdot \mathbf{V}(t)] dt \leq \frac{P \cdot \log(1 + KM)}{\eta(T)} \tag{26}$$

for all $\mathbf{Q}^* \in \mathcal{X}$.

To that end, (MXL) readily yields:

$$\begin{aligned} \int_0^T \text{tr}[(\mathbf{Q}(t) - \mathbf{Q}^*) \cdot \mathbf{V}(t)] dt &= \int_0^T \text{tr}[(\mathbf{Q}^* - \mathbf{Q}(t)) \cdot \dot{\mathbf{Y}}(t)] dt \\ &= \text{tr}[\mathbf{Y}(T) \cdot \mathbf{Q}^*] - \int_0^T \text{tr}[\mathbf{Q}(t)\dot{\mathbf{Y}}(t)] dt, \end{aligned} \tag{27}$$

where we have used the fact that $\mathbf{Y}(0) = 0$. To continue, note that the exponentiation step of (MXL) can be written more simply as:

$$\mathbf{Q} = P \cdot \nabla_{\mathbf{U}} \log [1 + \text{tr} \exp(\mathbf{U})], \tag{28}$$

where we have set $\mathbf{U} = \eta \mathbf{Y}$.⁷ As a result, with $\dot{\mathbf{U}} = \dot{\eta} \mathbf{Y} + \eta \dot{\mathbf{Y}}$, the integrand of the second term of (27) becomes:

$$\text{tr}[\mathbf{Q}\dot{\mathbf{Y}}] = \frac{1}{\eta} \text{tr}[\mathbf{Q}\dot{\mathbf{U}}] - \frac{\dot{\eta}}{\eta^2} \text{tr}[\mathbf{Q}\mathbf{U}] = \frac{P}{\eta} \frac{d}{dt} \log [1 + \text{tr} \exp(\mathbf{U})] - \frac{\dot{\eta}}{\eta^2} \text{tr}[\mathbf{Q}\mathbf{U}]. \tag{29}$$

⁷This is actually one of the main reasons behind the exponentiation step of (MXL).

Hence, after integrating (29) by parts (and recalling that $\mathbf{U}(0) = 0$), we get:

$$\begin{aligned} \int_0^T \text{tr}[\mathbf{Q}(t)\dot{\mathbf{Y}}(t)] dt &= \frac{P}{\eta(t)} \log[1 + \text{tr} \exp(\mathbf{U}(t))] \Big|_0^T \\ &+ P \int_0^T \frac{\dot{\eta}(t)}{\eta(t)^2} \log[1 + \text{tr}[\exp(\mathbf{U}(t))]] dt - \int_0^T \frac{\dot{\eta}(t)}{\eta(t)^2} \text{tr}[\mathbf{Q}(t)\mathbf{U}(t)] dt \\ &= \frac{P \log[1 + \text{tr} \exp(\mathbf{U}(T))]}{\eta(T)} - \frac{P \log(1 + KM)}{\eta(0)} \\ &+ \int_0^T \frac{\dot{\eta}(t)}{\eta(t)^2} [P \log[1 + \text{tr} \exp(\mathbf{U}(t))] - \text{tr}[\mathbf{Q}(t)\mathbf{U}(t)]] dt, \end{aligned} \quad (30)$$

where we have used the fact that $\mathbf{U}(0) = 0$ (implying in turn that $\text{tr} \exp(\mathbf{U}(0)) = KM$). Thus, combining all of the above, we obtain:

$$\begin{aligned} \int_0^T \text{tr}[(\mathbf{Q}(t) - \mathbf{Q}^*) \cdot \mathbf{V}(t)] dt &= \frac{P \log(1 + KM)}{\eta(0)} + \frac{\text{tr}[\mathbf{Q}^*\mathbf{U}(T)] - P \log[1 + \text{tr} \exp(\mathbf{U}(T))]}{\eta(T)} \\ &+ \int_0^T \frac{\dot{\eta}(t)}{\eta(t)^2} [\text{tr}[\mathbf{Q}(t)\mathbf{U}(t)] - P \log[1 + \text{tr} \exp(\mathbf{U}(t))]] dt, \end{aligned} \quad (31)$$

To proceed, we will require the inequality:

$$\text{tr}[\mathbf{A}\mathbf{X}] - \log[1 + \text{tr} \exp(\mathbf{X})] \leq \text{tr}[\mathbf{A} \log \mathbf{A}] + (1 - \text{tr} \mathbf{A}) \log(1 - \text{tr} \mathbf{A}) \quad (32)$$

valid for all Hermitian \mathbf{A}, \mathbf{X} , with $\mathbf{A} \geq 0$, $\text{tr} \mathbf{A} \leq 1$, and with equality holding if and only if

$$\mathbf{A} = \frac{\exp(\mathbf{X})}{1 + \text{tr} \exp(\mathbf{X})}. \quad (33)$$

To establish (32), it clearly suffices to show that the supremum of its LHS for fixed \mathbf{A} is precisely the RHS of (32). Accordingly, let

$$F(\mathbf{X}) = \log[1 + \text{tr} \exp(\mathbf{X})] - \text{tr}[\mathbf{A}\mathbf{X}], \quad (34)$$

so the maximizers of the LHS of (32) are given by the first-order stationarity condition $\nabla_{\mathbf{X}} F(\mathbf{X}) = 0$ (simply note that $F(\mathbf{X})$ is strictly concave in \mathbf{X}). By differentiating, we then obtain:

$$\nabla_{\mathbf{X}} F(\mathbf{X}) = \frac{\exp(\mathbf{X})}{1 + \text{tr} \exp(\mathbf{X})} - \mathbf{A}. \quad (35)$$

Thus, if $\mathbf{A} > 0$ and $\text{tr}[\mathbf{A}] < 1$, the equation $\nabla_{\mathbf{X}} F(\mathbf{X}) = 0$ always admits a (necessarily unique) solution given by:

$$\mathbf{X}^* = \log \mathbf{A} + \log(1 + \chi)\mathbf{I}, \quad (36)$$

with $\chi = \text{tr} \exp(\mathbf{X}^*)$. Moreover, setting $a = \text{tr} \mathbf{A}$ and tracing (35) readily yields $\chi = a/(1 - a)$, so, after some easy algebra, the maximum value of F will be:

$$F_{\max} = F(\mathbf{X}^*) = \text{tr}[\mathbf{A} \log \mathbf{A}] + (1 - a) \log(1 - a). \quad (37)$$

The above establishes (32) for the case $\mathbf{A} > 0$ and $\text{tr}[\mathbf{A}] < 1$; the boundary cases $\det \mathbf{A} = 0$ and/or $\text{tr}[\mathbf{A}] = 1$ then follow by continuity.

Thus, returning to (31), an immediate application of (32) gives:

$$\text{tr}[\mathbf{Q}^* \mathbf{U}(T)] - P \log[1 + \text{tr} \exp(\mathbf{U}(T))] \leq 0, \quad (38a)$$

$$\text{tr}[\mathbf{Q}(t) \mathbf{U}(t)] - P \log[1 + \text{tr} \exp(\mathbf{U}(t))] = P \cdot [\text{tr}[\mathbf{A}(t) \log \mathbf{A}(t)] + (1 - a(t)) \log(1 - a(t))], \quad (38b)$$

where we have set $\mathbf{A}(t) = \mathbf{Q}(t)/P$ and $a(t) = \text{tr} \mathbf{A}(t)$. As for (38b), its RHS can be expressed more concisely as the (negative) von Neumann quantum entropy of the augmented matrix $\mathbf{A}_0(t) = \text{diag}(a(t), \mathbf{A}(t))$, i.e.

$$\text{tr}[\mathbf{A}(t) \log \mathbf{A}(t)] + (1 - a(t)) \log(1 - a(t)) = \text{tr}[\mathbf{A}_0(t) \log \mathbf{A}_0(t)] \geq -\log(1 + KM), \quad (39)$$

where the last inequality simply corresponds to the maximum value of the von Neumann entropy (recall also that $\dim(\mathbf{A}_0) = 1 + KM$) [48]. Thus, substituting (38) and (39) back into (31), we finally obtain:

$$\int_0^T \text{tr}[(\mathbf{Q}(t) - \mathbf{Q}^*) \cdot \mathbf{V}(t)] dt \leq P \log(1 + KM) \left[\frac{1}{\eta(0)} - \int_0^T \frac{\dot{\eta}(t)}{\eta(t)^2} dt \right] = \frac{P \log(1 + KM)}{\eta(T)}, \quad (40)$$

where we have used the fact that $\dot{\eta} \leq 0$ (recall that η has been assumed nonincreasing). The regret bound (24) then follows by maximizing (25) over all $\mathbf{Q}^* \in \mathcal{X}$. ■

B. No regret in discrete time: the case of perfect CSI

We now return to the discrete-time process (MXL), written here in the more general form:

$$\begin{aligned} \mathbf{Y}(n) &= - \sum_{m=1}^n \mathbf{V}(m) \\ \mathbf{Q}(n+1) &= P \frac{\exp(\eta(n) \mathbf{Y}(n))}{1 + \text{tr}[\exp(\eta(n) \mathbf{Y}(n))]} \end{aligned} \quad (41)$$

with $\eta(n) = \eta n^{-1/2}$ for some positive parameter $\eta > 0$. To establish the regret bound (19) of Theorem 1, we will define an interpolated continuous-time process, use Proposition 3 to estimate the incurred regret in continuous time, and use a discrete-continuous comparison argument in order to bound the regret in discrete time.

Proof of Theorem 1: We begin by constructing a continuous-time interpolation of (MXL) and comparing it to its discrete-time analogue. To that end, consider the continuous-time processes $\mathbf{V}^c(t) = \mathbf{V}(\lceil t \rceil)$ and $\eta^c(t) = \eta(\lfloor t \rfloor)$ for all $t \geq 0$, with $\eta^c(0) = \eta(0) = \eta$ by convention.⁸ In this context, the continuous-time learning scheme (MXL-c) yields the processes:

$$\begin{aligned} \mathbf{Y}^c(t) &= - \int_0^t \mathbf{V}^c(s) ds \\ \mathbf{Q}^c(t) &= P \frac{\exp(\eta^c(t) \mathbf{Y}^c(t))}{1 + \text{tr}[\exp(\eta^c(t) \mathbf{Y}^c(t))]} \end{aligned} \quad (42)$$

whence we easily obtain:

$$\begin{aligned} \mathbf{Y}^c(n) &= - \int_0^n \mathbf{V}^c(s) ds = - \sum_{m=1}^n \int_{m-1}^m \mathbf{V}(m) ds = - \sum_{m=1}^n \mathbf{V}(m) = \mathbf{Y}(n), \\ \mathbf{Q}^c(n) &= \mathbf{Q}(n+1), \end{aligned} \quad (43)$$

⁸Note that $\mathbf{V}^c(t) = \mathbf{V}(n)$ and $\eta^c(t) = \eta(n-1)$ for all $t \in (n-1, n)$, i.e. $\mathbf{V}^c(t)$ precedes its discrete-time analogue, while $\eta^c(t)$ lags behind it; this one-step offset will be key in the rest of our proof.

Consequently, for all $n \geq 1$ and for all $t \in (n-1, n)$, Hölder's inequality yields:

$$\begin{aligned} |\operatorname{tr}[(\mathbf{Q}^c(t) - \mathbf{Q}(n)) \cdot \mathbf{V}^c(t)]| &= |\operatorname{tr}[(\mathbf{Q}^c(t) - \mathbf{Q}^c(n-1)) \cdot \mathbf{V}^c(t)]| \\ &\leq \|\mathbf{V}^c(t)\| \cdot |\operatorname{tr}[\mathbf{Q}^c(t) - \mathbf{Q}^c(n-1)]| \leq V \cdot |\operatorname{tr}[\mathbf{Q}^c(t) - \mathbf{Q}^c(n-1)]| \end{aligned} \quad (44)$$

Using the analysis of [31], it can be shown that the map $\mathbf{U} \mapsto P \exp(\mathbf{U})/[1 + \operatorname{tr} \exp(\mathbf{U})]$ is $(P/2)$ -Lipschitz with respect to the spectral and nuclear norms (for the map's domain and codomain respectively). We may thus write:

$$\operatorname{tr}[\mathbf{Q}^c(t) - \mathbf{Q}^c(n-1)] \leq \frac{1}{2}P \|\eta^c(t)\mathbf{Y}^c(t) - \eta^c(n-1)\mathbf{Y}^c(n-1)\| = \frac{1}{2}\eta(n-1)P \|\mathbf{Y}^c(t) - \mathbf{Y}^c(n-1)\|. \quad (45)$$

Furthermore, by definition, we also have:

$$\|\mathbf{Y}^c(t) - \mathbf{Y}^c(n-1)\| = \left\| -\int_0^t \mathbf{V}^c(s) ds + \int_0^{n-1} \mathbf{V}^c(s) ds \right\| \leq \int_{n-1}^t \|\mathbf{V}^c(s)\| ds \leq V \cdot (t - n + 1) \quad (46)$$

and hence, by combining (44), (45) and (46), we get:

$$|\operatorname{tr}[(\mathbf{Q}^c(t) - \mathbf{Q}(n)) \mathbf{V}^c(t)]| \leq \frac{1}{2}P \|\mathbf{V}(n)\|^2 \eta(n-1) \cdot (t - n + 1). \quad (47)$$

Accordingly, with this discrete/continuous comparison result at hand, we get:

$$\begin{aligned} \left| \int_0^T \operatorname{tr}[\mathbf{Q}^c(t)\mathbf{V}^c(t)] dt - \sum_{n=1}^T \operatorname{tr}[\mathbf{Q}(n)\mathbf{V}(n)] \right| &= \left| \sum_{n=1}^T \int_{n-1}^n \operatorname{tr}[\mathbf{Q}^c(t)\mathbf{V}^c(t)] dt - \operatorname{tr}[\mathbf{Q}(n)\mathbf{V}(n)] \right| \\ &= \left| \sum_{n=1}^T \int_{n-1}^n [\operatorname{tr}(\mathbf{Q}^c(t)\mathbf{V}^c(t)) - \operatorname{tr}(\mathbf{Q}(n)\mathbf{V}(n))] dt \right| \\ &\leq \sum_{n=1}^T \int_{n-1}^n |\operatorname{tr}[\mathbf{Q}^c(t)\mathbf{V}^c(t)] - \operatorname{tr}[\mathbf{Q}(n)\mathbf{V}^c(t)]| dt \\ &\leq \sum_{n=1}^T \int_{n-1}^n \frac{P}{2} \|\mathbf{V}(n)\|^2 \eta(n-1) (t - n + 1) dt \leq \frac{PV^2}{4} \sum_{n=1}^T \eta(n-1), \end{aligned} \quad (48)$$

for all $T \geq 1$. Thus, using Proposition 3 and the convexity condition (25), we obtain:

$$\begin{aligned} \sum_{n=1}^T [\ell(\mathbf{Q}(n); n) - \ell(\mathbf{Q}^*; n)] &\leq \sum_{n=1}^T \operatorname{tr}[(\mathbf{Q}(n) - \mathbf{Q}^*) \mathbf{V}(n)] \\ &\leq \int_0^T \operatorname{tr}[\mathbf{Q}^c(t)\mathbf{V}^c(t)] dt + \frac{PV^2}{4} \sum_{n=1}^T \eta(n-1) - \int_0^T \operatorname{tr}[\mathbf{Q}^*\mathbf{V}^c(t)] dt \\ &= \int_0^T \operatorname{tr}[(\mathbf{Q}^c(t) - \mathbf{Q}^*) \mathbf{V}^c(t)] dt + \frac{PV^2}{4} \sum_{n=1}^T \eta(n-1) \\ &\leq \frac{P \log(1 + KM)}{\eta(T)} + \frac{PV^2}{4} \sum_{n=1}^T \eta(n-1), \end{aligned} \quad (49)$$

where we used (48) and the fact that $\sum_{n=1}^T \operatorname{tr}[\mathbf{Q}^*\mathbf{V}(n)] = \int_0^T \operatorname{tr}[\mathbf{Q}^*\mathbf{V}^c(t)] dt$ in the second line. Thus, substituting $\eta(n) = \min\{\eta n^{-1/2}, \eta\}$, the last term of (49) becomes

$$\sum_{n=1}^T \eta(n-1) = \eta + \sum_{n=1}^{T-1} \eta n^{-1/2} \leq \eta + \int_0^{T-1} \eta t^{-1/2} dt \leq \eta(1 + 2\sqrt{T}), \quad (50)$$

and our proof is completed by substituting in (49) and maximizing over all $\mathbf{Q}^* \in \mathcal{X}$. ■

C. No regret in discrete time: the case of imperfect CSI

To prove Theorem 2, we will use Eq. (48) to bound the user’s “virtual” regret with respect to the sequence of noisy gradient estimates $\hat{\mathbf{V}}(n)$, and we will then employ the Borel–Cantelli lemma to show that the user’s actual regret lies within a vanishing window of his “virtual” regret.

Proof of Theorem 2: As usual, the user’s regret is bounded by:

$$\text{Reg}(T) = \max_{\mathbf{Q}^* \in \mathcal{X}} \sum_{n=1}^T [\ell(\mathbf{Q}(n); n) - \ell(\mathbf{Q}^*; n)] \leq \max_{\mathbf{Q}^* \in \mathcal{X}} \sum_{n=1}^T \text{tr}[(\mathbf{Q}(n) - \mathbf{Q}^*) \cdot \mathbf{V}(n)], \quad (51)$$

so, for the first part of the theorem, it suffices to show that $\sum_{n=1}^T \text{tr}[(\mathbf{Q}(n) - \mathbf{Q}^*) \cdot \mathbf{V}(n)] = o(T)$ for all $\mathbf{Q}^* \in \mathcal{X}$. To that end, given that $\mathbf{V}(n) = \hat{\mathbf{V}}(n) - \mathbf{Z}(n)$, we have:

$$\text{tr}[(\mathbf{Q}(n) - \mathbf{Q}^*) \cdot \mathbf{V}(n)] = \text{tr}[(\mathbf{Q}(n) - \mathbf{Q}^*) \cdot \hat{\mathbf{V}}(n)] - \text{tr}[(\mathbf{Q}(n) - \mathbf{Q}^*) \cdot \mathbf{Z}(n)], \quad (52)$$

where $\mathbf{Q}(n)$ is defined via the stochastic recursion:

$$\begin{aligned} \mathbf{Y}(n) &= \mathbf{Y}(n-1) - \hat{\mathbf{V}}(n), \\ \mathbf{Q}(n+1) &= P \frac{\exp(\eta n^{-1/2} \mathbf{Y}(n))}{1 + \text{tr}[\exp(\eta n^{-1/2} \mathbf{Y}(n))]} \end{aligned} \quad (53)$$

Going back to the proof of Thm. 1, we may then use the last inequality of (48) to rewrite (49) as:

$$\sum_{n=1}^T \text{tr}[(\mathbf{Q}(n) - \mathbf{Q}^*) \cdot \hat{\mathbf{V}}(n)] \leq \frac{P \log(1 + KM)}{\eta(T)} + \frac{P}{4} \sum_{n=1}^T \eta(n-1) \|\hat{\mathbf{V}}(n)\|^2. \quad (54)$$

The last term of (54) can then be bounded as:

$$\begin{aligned} \sum_{n=1}^T \eta(n-1) \|\hat{\mathbf{V}}(n)\|^2 &\leq \sum_{n=1}^T \eta(n-1) [\|\mathbf{V}(n)\|^2 + 2\|\mathbf{V}(n)\| \cdot \|\mathbf{Z}(n)\| + \|\mathbf{Z}(n)\|^2] \\ &= V^2 \sum_{n=1}^T \eta(n-1) + \mathcal{O}\left(\sum_{n=1}^T \eta(n-1) \|\mathbf{Z}(n)\|^2\right), \end{aligned} \quad (55)$$

where we have used the triangle inequality in the first line. We now claim that

$$\frac{1}{T} \sum_{n=1}^T \eta(n-1) \|\mathbf{Z}(n)\|^2 \rightarrow 0 \quad \text{as } T \rightarrow \infty \text{ (a.s.)}. \quad (56)$$

Indeed, if we let $z(n) = \|\mathbf{Z}(n)\|$, Hypothesis (H2) implies that $\mathbb{P}(z(n) \geq n^{1/4-\varepsilon}) = \mathcal{O}(1/n^\beta)$ for some $\beta > 1$ and for all small enough $\varepsilon > 0$. We thus obtain:

$$\sum_{n=1}^{\infty} \mathbb{P}(z(n) \geq n^{1/4-\varepsilon}) = \mathcal{O}\left(\sum_{n=1}^{\infty} n^{-\beta}\right) = \mathcal{O}(1) < \infty, \quad (57)$$

and hence, by the Borel–Cantelli lemma, we conclude that

$$\mathbb{P}(z(n) \geq n^{1/4-\varepsilon} \text{ for infinitely many } n) = 0. \quad (58)$$

In turn, this implies that $z(n)^2 = \mathcal{O}(n^{1/2-2\varepsilon})$ almost surely, so, with $\eta(n) = n^{-1/2}$, we get:

$$\sum_{n=1}^T \eta(n-1) \|\mathbf{Z}(n)\|^2 = \mathcal{O}\left(\sum_{n=1}^T n^{-1/2} n^{1/2-2\varepsilon}\right) = \mathcal{O}\left(\sum_{n=1}^T 1/n^{2\varepsilon}\right) = o(T) \quad (\text{a.s.}). \quad (59)$$

For the second term of (52), let $\xi(n) = \text{tr}[(\mathbf{Q}(n) - \mathbf{Q}^*) \cdot \mathbf{Z}(n)]$. Then, given that $\mathbf{Q}(n)$ is fully determined by $\mathbf{Q}(n-1)$ and $\mathbf{Z}(n-1)$, it follows that $\mathbb{E}[\xi(n) | \mathbf{Q}(n-1)] = 0$, i.e. $\xi(n)$ is a martingale difference sequence; as a result, we get

$\lim_{T \rightarrow \infty} T^{-1} \sum_{n=1}^T \xi(n) = 0$ by the strong law of large numbers for martingale differences – see e.g. Theorem 2.18 in [43]. Combining this with (59), we then get

$$\sum_{n=1}^T \text{tr}[(\mathbf{Q}(n) - \mathbf{Q}^*) \cdot \mathbf{V}(n)] = o(T) \quad (\text{a.s.}), \quad (60)$$

i.e. (53) leads to no regret, as claimed. The mean bound (23) is then obtained by taking expectations on both sides of (54) and recalling that $\mathbb{E}[\hat{\mathbf{V}}(n)|\mathbf{Q}(n-1)] = \mathbf{V}(n)$. ■

REFERENCES

- [1] J. G. Andrews, S. Buzzi, W. Choi, S. Hanly, A. Lozano, A. C. K. Soong, and J. C. Zhang, “What will 5G be?” *IEEE J. Sel. Areas Commun.*, vol. 32, no. 6, pp. 1065–1082, June 2014.
- [2] J. Zander, “Performance of optimum transmitter power control in cellular radio systems,” *IEEE Trans. Veh. Technol.*, vol. 41, no. 1, pp. 57–62, Feb. 1992.
- [3] S. Grandhi, R. Vijayan, and D. Goodman, “Distributed power control in cellular radio systems,” *IEEE Trans. Commun.*, vol. 42, no. 234, pp. 226–228, 1994.
- [4] G. Foschini and Z. Miljanic, “A simple distributed autonomous power control algorithm and its convergence,” *IEEE Trans. Veh. Technol.*, vol. 42, no. 4, pp. 641–646, Nov. 1993.
- [5] R. Yates, “A framework for uplink power control in cellular radio systems,” *IEEE J. Sel. Areas Commun.*, vol. 13, no. 7, pp. 1341–1347, 1995.
- [6] M. Chiang, P. Hande, T. Lan, and C. W. Tan, “Power control in wireless cellular networks,” *Foundations and Trends in Networking*, vol. 2, no. 4, pp. 381–533, 2007.
- [7] Z. Shen, J. Andrews, and B. Evans, “Optimal power allocation in multiuser OFDM systems,” in *IEEE Global Telecommunications Conference, 2003. GLOBECOM '03*, vol. 1, Dec. 2003, pp. 337–341 Vol.1.
- [8] 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.
- [9] M. Kobayashi and G. Caire, “An iterative water-filling algorithm for maximum weighted sum-rate of gaussian MIMO-BC,” *IEEE J. Sel. Areas Commun.*, vol. 24, no. 8, pp. 1640–1646, Aug. 2006.
- [10] G. Scutari, D. P. Palomar, and S. Barbarossa, “The MIMO iterative waterfilling algorithm,” *IEEE Trans. Signal Process.*, vol. 57, no. 5, pp. 1917–1935, May 2009.
- [11] Y. Yang, G. Scutari, P. Song, and D. P. Palomar, “Robust MIMO cognitive radio under interference temperature constraints,” *IEEE J. Sel. Areas Commun.*, vol. 31, no. 11, pp. 2465–2483, November 2013.
- [12] E. V. Belmega, S. Lasaulce, and M. Debbah, “Power allocation games for MIMO multiple access channels with coordination,” *IEEE Trans. Wireless Commun.*, vol. 8, no. 5, pp. 3182–3192, June 2009.
- [13] J. Oh, S.-J. Kim, and J. Cioffi, “Optimum power allocation and control for OFDM in multiple access channels,” in *Vehicular Technology Conference, 2004. VTC2004-Fall. 2004 IEEE 60th*, vol. 2, Sep. 2004, pp. 774–778 Vol. 2.
- [14] D. Yu and J. Cioffi, “SPC10-2: Iterative water-filling for optimal resource allocation in OFDM multiple-access and broadcast channels,” in *IEEE Global Telecommunications Conference, 2006. GLOBECOM '06*, Nov. 2006, pp. 1–5.
- [15] G. Wunder and T. Michel, “Optimal resource allocation for parallel gaussian broadcast channels: Minimum rate constraints and sum power minimization,” *IEEE Trans. Inf. Theory*, vol. 53, no. 12, pp. 4817–4822, Dec. 2007.
- [16] M. Moretti, L. Sanguinetti, and X. Wang, “Resource allocation for power minimization in the downlink of THP-based spatial multiplexing MIMO-OFDMA systems,” *IEEE Trans. Veh. Technol.*, to appear.
- [17] N. Bambos, S. C. Chen, and G. J. Pottie, “Channel access algorithms with active link protection for wireless communication networks with power control,” *IEEE/ACM Trans. Netw.*, vol. 8, no. 5, pp. 583–597, Oct. 2000.
- [18] T. Holliday, N. Bambos, P. Glynn, and A. Goldsmith, “Distributed power control for time varying wireless networks: Optimality and convergence,” in *in Proceedings: Allerton Conference on Communications, Control, and Computing*, 2003.
- [19] P. Mertikopoulos, E. V. Belmega, A. L. Moustakas, and S. Lasaulce, “Distributed learning policies for power allocation in multiple access channels,” *IEEE J. Sel. Areas Commun.*, vol. 30, no. 1, pp. 96–106, January 2012.

- [20] S. Shalev-Shwartz, "Online learning and online convex optimization," *Foundations and Trends in Machine Learning*, vol. 4, no. 2, pp. 107–194, 2011.
- [21] J. Hannan, "Approximation to Bayes risk in repeated play," in *Contributions to the Theory of Games, Volume III*, ser. Annals of Mathematics Studies, M. Dresher, A. W. Tucker, and P. Wolfe, Eds. Princeton, NJ: Princeton University Press, 1957, vol. 39, pp. 97–139.
- [22] N. Cesa-Bianchi and G. Lugosi, *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- [23] J. Dams, M. Hoefer, and T. Kesselheim, "Convergence time of power-control dynamics," *IEEE J. Sel. Areas Commun.*, vol. 30, no. 11, pp. 2231–2237, Dec. 2012.
- [24] A. Blum and Y. Mansour, "From external to internal regret," *J. Mach. Learn. Res.*, vol. 8, pp. 1307–1324, Dec. 2007.
- [25] B. Latifa, Z. Gao, and S. Liu, "No-regret learning for simultaneous power control and channel allocation in cognitive radio networks," in *Computing, Communications and Applications Conference (ComComAp), 2012*, Jan. 2012, pp. 267–271.
- [26] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire, "Gambling in a rigged casino: The adversarial multi-armed bandit problem," in *36th Annual Symposium on Foundations of Computer Science, 1995. Proceedings*, Oct. 1995, pp. 322–331.
- [27] S. Maghsudi and S. Stanczak, "Joint channel selection and power control in infrastructureless wireless networks: A multi-player multi-armed bandit framework," *IEEE Trans. Veh. Technol.*, vol. PP, no. 99, pp. 1–1, 2014.
- [28] P. Mertikopoulos and E. V. Belmega, "Transmit without regrets: online optimization in MIMO-OFDM cognitive radio systems," *IEEE J. Sel. Areas Commun.*, vol. 32, no. 11, Nov. 2014.
- [29] K. Tsuda, G. Rätsch, and M. K. Warmuth, "Matrix exponentiated gradient updates for on-line Bregman projection," *Journal of Machine Learning Research*, vol. 6, pp. 995–1018, 2005.
- [30] 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.
- [31] S. M. Kakade, S. Shalev-Shwartz, and A. Tewari, "Regularization techniques for learning with matrices," *The Journal of Machine Learning Research*, vol. 13, pp. 1865–1890, 2012.
- [32] I. Stiakogiannakis, P. Mertikopoulos, and C. Touati, "No regrets: Distributed power control under time-varying channels and QoS requirements," in *2014 52st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, Oct. 2014.
- [33] E. Telatar, "Capacity of multi-antenna Gaussian channels," *European Transactions on Telecommunications*, vol. 10, no. 6, pp. 585–595, Nov. 1999.
- [34] H. Bolcskei, D. Gesbert, and A. Paulraj, "On the capacity of OFDM-based spatial multiplexing systems," *IEEE Trans. Commun.*, vol. 50, no. 2, pp. 225–234, 2002.
- [35] Z. Wang, L. Jiang, and C. He, "Optimal price-based power control algorithm in cognitive radio networks," *IEEE Trans. Wireless Commun.*, vol. 13, no. 11, pp. 5909–5920, Nov. 2014.
- [36] C. Saraydar, N. Mandayam, and D. Goodman, "Pricing and power control in a multicell wireless data network," *IEEE J. Sel. Areas Commun.*, vol. 19, no. 10, pp. 1883–1892, 2001.
- [37] F. Meshkati, A. J. Goldsmith, H. V. Poor, and S. C. Schwartz, "A game-theoretic approach to energy-efficient modulation in CDMA networks with delay QoS constraints," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 6, pp. 1069–1078, August 2007.
- [38] E. Hazan and C. Seshadri, "Efficient learning algorithms for changing environments," in *ICML '09: Proceedings of the 26th International Conference on Machine Learning*, 2009.
- [39] N. Cesa-Bianchi, P. Gaillard, G. Lugosi, and G. Stoltz, "Mirror descent meets fixed share (and feels no regret)," in *Advances in Neural Information Processing Systems*, 989–997, Ed., vol. 25, 2012.
- [40] M. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in *ICML '03: Proceedings of the 20th International Conference on Machine Learning*, 2003.
- [41] J. Kwon and P. Mertikopoulos, "A continuous-time approach to online optimization," 2014, <http://arxiv.org/abs/1401.6956>.
- [42] A. M. Davie and A. J. Stothers, "Improved bound for complexity of matrix multiplication," *Proceedings of the Royal Society of Edinburgh, Section: A Mathematics*, vol. 143, no. 2, pp. 351–369, 4 2013.
- [43] P. Hall and C. C. Heyde, *Martingale Limit Theory and Its Application*, ser. Probability and Mathematical Statistics. New York: Academic Press, 1980.
- [44] I. N. Stiakogiannakis and D. I. Kaklamani, "A radio resource management framework for multi-user multi-cell OFDMA networks based on game theory," *Wireless Personal Communications*, vol. 69, no. 2, pp. 745–770, Mar. 2013.

- [45] "User equipment (UE) radio transmission and reception," 3GPP, Technical Specification 36.101 V12.4.0, Jun. 2014. [Online]. Available: <http://www.3gpp.org>
- [46] S. Sorin, "Exponential weight algorithm in continuous time," *Mathematical Programming*, vol. 116, no. 1, pp. 513–528, 2009.
- [47] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge University Press, 2004.
- [48] E. Carlen, "Trace inequalities and quantum entropy: An introductory course," in *Entropy and the Quantum*, ser. Contemporary Mathematics. AMS, Mar. 2009, vol. 529, pp. 73–140.