Adaptive Transmit Policies for Cost-Efficient Power Allocation in Multi-Carrier Systems

(Invited Paper)

Salvatore D'Oro*, Panayotis Mertikopoulos†‡, Aris L. Moustakas§¶ and Sergio Palazzo*

* CNIT Research Unit, University of Catania, Italy
† French National Center for Scientific Research (CNRS), LIG F-38000 Grenoble, France
‡ Univ. Grenoble Alpes, LIG F-38000 Grenoble, France
§ Department of Physics, University of Athens and Institute of Accelerating Systems and Applications (IASA), Athens, Greece
¶ École Supérieure d'Électricité (Supélec), Gif-sur-Yvette, France

Abstract—In this paper, we examine the problem of cost/energy-efficient power allocation in uplink multi-carrier orthogonal frequency-division multiple access (OFDMA) wireless networks. In particular, we consider a set of wireless users who seek to maximize their transmission rate subject to pricing limitations and we show that the resulting non-cooperative game admits a unique equilibrium for almost every realization of the system's channels. We also propose a distributed exponential learning scheme which allows users to converge to the game's equilibrium exponentially fast by using only local channel state information (CSI) and signal to interference-plus-noise ratio (SINR) measurements. Given that such measurements are often imperfect in practical scenarios, a major challenge occurs when the users' information is subject to random perturbations. In this case, by using tools and ideas from stochastic convex programming, we show that the proposed learning scheme retains its convergence properties irrespective of the magnitude of the observational errors.

I. Introduction

Ever since the early development stages of legacy wireless networks, power control has been an essential component of network design and operation, especially in decentralized environments where only local information is available at each mobile terminal [1]. As such, the introduction of fast and distributed power control algorithms (both closed- and open-loop) was one of the main improvements that were brought about in third generation CDMA-based cellular networks, in both single- and multi-carrier settings.

Controlling the transmitted power has two important purposes. The first is to minimize the interference of a given node to neighboring receivers, an issue of critical importance in future and emerging wireless network paradigms where cells are deployed at a massive scale – for instance, as in the case of femto-cell networks [2]. Due to their close proximity, neighboring users may create significant interference to one another, so care must be taken to choose a power allocation profile that maximizes the users' transmission rate while limiting their overall transmit power – otherwise, the situation could rapidly degenerate to a cascade of power increases.

Second, power control reduces the users' overall transmitted power. Mobile terminals are generally energy-constrained (e.g. due to the limitations of their power source or because of the cost of power consumption), so inefficient power allocation can bring about unnecessary losses in performance. As a result, the problem that arises is to derive distributed power allocation policies that maximize the users' transmission rate in energy-aware scenarios where transmission power also carries a commensurate cost. This objective is made more complicated by the fact that wireless users typically have conflicting interests and cannot be assumed to cooperate with each other for their collective benefit (in decentralized environments at least).

In view of the above, non-cooperative game theory has become an important tool to analyze the interactions between mobile users in wireless network scenarios where energy-efficient rate maximization is an issue – see e.g. [3–6] for applications to power control and [7–9] for power allocation problems. In this framework, the primary objective has been to develop policies and algorithms that wireless users can use to optimize their resources (power, bandwidth, etc.), so the main questions that arise are *a*) whether there exist "equilibrial" power allocation policies which are stable against unilateral deviations; *b*) whether these (Nash) equilibria are unique (and thus offer some predictive power with regard to the analysis of the system); and *c*) whether the system's users can reach such a state by means of distributed, adaptive methods that only require local (and readily available) information.

Optimizing power allocation in such a way has been investigated in both single- [4, 5, 10] and multi-carrier scenarios [6, 11, 12]. In particular, the authors of [3, 4, 10] investigated the role of pricing as an effective mechanism to measure the cost of power consumption, thus leading to an energy-efficient formulation where users seek to maximize their transmission rate while keeping their transmit power in check (see also the very recent paper [12] where the authors consider the problem of maximizing the users' transmission rate per unit of transmitted power subject to minimum rate requirements). Then, to reach an equilibrium state in such a setting, several distributed approaches have been proposed in the existing literature based on reaction functions [4], Gauss-Seidel and Jacobi update algorithms [3] or replicator-based learning [11].

In this paper, we consider the problem of cost/energy-efficient power allocation in uplink multi-carrier orthogonal frequency-division multiplexing (OFDM) networks, thus fusing and extending the cost-driven treatment of [4] for single-carrier systems with the analysis of [11] for multi-carrier systems in the absence of power consumption considerations. One straightforward scenario where power consumed by the transmitter needs to be priced in some way is related to the modern use of smartphones and other multi-purpose mobile devices. In such devices, the power available for wireless transmission is restricted by the usage of other time-dependent applications of varying priority, so one way to model this scenario is to add a priority-dependent penalty to the available power of the wireless transmission in each user.

To that end, we provide a game-theoretic formulation for the problem of unilateral rate maximization subject to pricing limitations of a general form and we show that the resulting game admits a unique equilibrium for almost every realization of the system's channels. We then propose a distributed exponential learning scheme which converges to equilibrium very rapidly using only local channel state information (CSI) and signal to interference-plus-noise ratio (SINR) measurements. Importantly, by using powerful tools from stochastic convex programming [13], we are able to show that the algorithm retains its convergence properties even in the presence of imperfect measurements – and *irrespective* of the magnitude of the observational errors.

Paper Outline: Our paper is structured as follows: in Section II, we present our system model and prove that the associated non-cooperative game admits a unique equilibrium almost surely. In Sections III and IV we propose a distributed learning algorithm which allows users to converge to the game's equilibrium, even in the presence of (arbitrarily large) observational errors. Finally, our analysis is supplemented in Section V by extensive numerical simulations which validate the convergence results of the previous section and characterize the performance envelope of the system at its end-state (the efficiency of the game's equilibria in terms of power consumption and achieved transmission rate, etc.).

II. System Model

The network model that we will focus on consists of a set $\mathcal{K} = \{1, \dots, K\}$ of non-cooperative wireless (single-antenna) transmitters who communicate with a common receiver over a set $\mathcal{M} = \{1, \dots, M\}$ of non-interfering subcarriers (typically in the frequency domain if an OFDM scheme is employed). Focusing on the uplink case, the aggregate received signal y_μ over the μ-th subcarrier is then given by the familiar signal model:

$$y_\mu = \sum_{k \in \mathcal{K}} h_{k\mu} x_{k\mu} + z_\mu, \qquad (1)$$

where $x_{k\mu} \in \mathbb{C}$ denotes the transmitted signal of user k over subcarrier μ, $h_{k\mu} \in \mathbb{C}$ is the corresponding transfer coefficient (assumed fixed for the duration of the transmission) and $z_\mu \in \mathbb{C}$ is the noise in the channel, including thermal, atmospheric and other ambient effects – and modeled as a zero-mean Gaussian vector $z_\mu \sim \mathcal{CN}(0, \sigma_\mu^2)$ with non-singular covariance.

In this context, the average transmit power of user k on subcarrier μ is

$$p_{k\mu} = \mathbb{E}\left[|x_{k\mu}|^2\right], \qquad (2)$$

and we will be assuming that each user's *total* transmit power $p_k = \mathbb{E}[\mathbf{x}_k^\dagger \mathbf{x}_k] = \sum_\mu p_{k\mu}$ satisfies the constraint:

$$p_k = \sum_{\mu \in \mathcal{M}} p_{k\mu} \leqslant P_k, \qquad (3)$$

where P_k denotes the maximum transmit power of user $k \in \mathcal{K}$. Accordingly, the set of admissible power allocation vectors for user k will be

$$\mathcal{X}_k = \left\{ \mathbf{p}_k \in \mathbb{R}^{\mathcal{M}} : p_{k\mu} \geqslant 0 \text{ and } \sum_{\mu \in \mathcal{M}} p_{k\mu} \leqslant P_k \right\}, \qquad (4)$$

and the system's *state space* – viz. the space of all admissible power allocation profiles $\mathbf{p} = (\mathbf{p}_1, \dots, \mathbf{p}_K)$ – will be denoted by $\mathcal{X} = \prod_k \mathcal{X}_k$.

On that account, each user's achievable transmission rate will depend on his individual SINR

$$\mathsf{sinr}_{k\mu}(\mathbf{p}) = \frac{g_{k\mu} p_{k\mu}}{\sigma_\mu^2 + \sum_{\ell \neq k} g_{\ell\mu} p_{\ell\mu}}, \qquad (5)$$

where $g_{k\mu} = |h_{k\mu}|^2$ denotes the channel gain coefficient for user k over the μ-th subcarrier. Thus, in the single user decoding (SUD) regime – where interference by other users is treated as (possibly colored) noise – the maximum information transmission rate (achievable with random Gaussian codes) will be:

$$r_k(\mathbf{p}) = \sum_{\mu \in \mathcal{M}} \log\left(1 + \mathsf{sinr}_{k\mu}(\mathbf{p})\right). \qquad (6)$$

Given the form of this objective, each user will saturate the power constraint (3) and transmit with maximum possible power in order to maximize his throughput. In practical scenarios however, power consumption carries a commensurate cost, so we will instead consider the energy-aware utility model

$$u_k(\mathbf{p}) = r_k(\mathbf{p}) - c_k(p_k), \qquad (7)$$

where $p_k = \mathbb{E}[\mathbf{x}_k^\dagger \mathbf{x}_k] = \sum_{\mu \in \mathcal{M}} p_{k\mu}$ denotes the user's total transmit power and $c_k \colon [0, P_k] \to \mathbb{R}_+$ is a user-specific cost function measuring the impact of power consumption.

The utility/cost model above admits several interpretations, depending on one's point of view. Perhaps the most straightforward one is that of c_k representing the effective monetary cost of power consumption (whether the cost is paid up front or postponed to the moment where the battery of the wireless device will need to be recharged). Alternatively, from the viewpoint of energy efficiency, the cost function c_k could represent the user's adversity to transmit with higher power when not absolutely necessary. To keep things as general as possible, we will only consider c_k as a generic "price" function and assume that it is convex and increasing in p_k (in tune with standard economic assumptions).

With all this in mind, unilateral utility maximization leads to a non-cooperative game $\mathfrak{G} \equiv \mathfrak{G}(\mathcal{K}, \mathcal{M}, u)$ for cost-efficient power allocation defined as follows:

1) The set of *players* of $\mathfrak{G}$ comprises the set of wireless transmitters $\mathcal{K} = \{1, \dots, K\}$.

2) Each player's set of *actions* consists of the corresponding feasible power allocation profiles $\mathbf{p}_k \in \mathfrak{X}_k = \{\mathbf{p}_k \in \mathbb{R}^{\mathcal{M}} : p_{k\mu} \geq 0 \text{ and } \sum_{\mu \in \mathcal{M}} p_{k\mu} \leq P_k\}$.
3) Each player's *utility* $u_k \colon \mathfrak{X} \to \mathbb{R}$ is given by (7).

We will thus say that a power allocation profile $\mathbf{p} \in \mathfrak{X}$ is a *Nash equilibrium* of $\mathfrak{G}$ when

$$u_k(\mathbf{p}_k; \mathbf{p}_{-k}) \geq u_k(\mathbf{p}'_k; \mathbf{p}_{-k}), \tag{8}$$

for all $\mathbf{p}'_k \in \mathfrak{X}_k$ and for all $\mathbf{p}_{-k} \in \mathfrak{X}_{-k} \equiv \prod_{\ell \neq k} \mathfrak{X}_\ell$.

As the next lemma shows, an important property of the game $\mathfrak{G}$ is that the players' objectives are aligned along a (concave) *potential function* (in the sense of [14]):

Lemma 1. *The* (concave) *function*

$$\Psi(\mathbf{p}) = \sum_{\mu \in \mathcal{M}} \log\left(1 + \sum_k g_{k\mu} p_{k\mu}/\sigma_\mu^2\right) - \sum_{k \in \mathcal{K}} c_k(p_k) \tag{9}$$

is a potential function for $\mathfrak{G}$; more precisely:

$$u_k(\mathbf{p}_k; \mathbf{p}_{-k}) - u_k(\mathbf{p}'_k; \mathbf{p}_{-k}) = \Psi(\mathbf{p}_k; \mathbf{p}_{-k}) - \Psi(\mathbf{p}'_k; \mathbf{p}_{-k}), \tag{10}$$

for all $\mathbf{p}_k \in \mathfrak{X}_k$, $\mathbf{p}_{-k} \in \mathfrak{X}_{-k}$, and for all $k \in \mathcal{K}$.

Sketch of proof: The claim follows by carrying out the calculation at the right-hand side of (10). ∎

By exploiting the game's potential property, it is easy to see that the game's set of equilibria coincides with the set of maximizers of Ψ [15]; as such, we are led to the (nonlinear) concave maximization problem:

$$\begin{aligned} \text{maximize} \quad & \Psi(\mathbf{p}_1, \ldots, \mathbf{p}_k), \\ \text{subject to} \quad & p_{k\mu} \geq 0 \text{ and } \sum_{\mu \in \mathcal{M}} p_{k\mu} \leq P_k. \end{aligned} \tag{11}$$

Remark 1. The first term of the potential function Ψ is simply the system's sum rate under successive interference cancellation (SIC). As a result, maximizing Ψ over the set of feasible power allocation profiles $\mathbf{p} \in \mathfrak{X}$ is equivalent to maximizing the users' aggregate utility (sum rate minus aggregate cost) in a *centralized* environment where one can apply more sophisticated successive interference cancellation (SIC) techniques.

With Lemma 1 at hand, establishing the existence of Nash equilibria for $\mathfrak{G}$ is trivial; in fact, since the game's potential is concave, it follows that the set of Nash equilibria of the game is a convex subset of $\mathfrak{X}^*$ [15, 16]. As it turns out, this convex set is almost surely a singleton:

Proposition 1. *The cost-efficient power allocation game $\mathfrak{G}$ admits a unique Nash equilibrium for almost every realization of the channel coefficients $h_{k\mu}$.*

Sketch of proof: Due to lack of space, we will only sketch the basic elements of the proof following a technique introduced in [11]. First, note that Ψ is not *strictly* concave: $\Psi(\mathbf{p}) = \Psi(\mathbf{p}')$ whenever $\sum_k g_{k\mu} p_{k\mu}/\sigma_\mu^2 = \sum_k g_{k\mu} p'_{k\mu}/\sigma_\mu^2$ and $p_k = p'_k$. These liner relations define a convex subset of maximizers of Ψ which lie at the intersection of $\mathfrak{X}$ with an affine space of "degenerate" directions along which Ψ is constant. By employing the approach of [17] to represent the game as a (multi)graph, it can then be shown that, generically, this subspace can only intersect $\mathfrak{X}$ at a single point, as claimed. ∎

III. Distributed Learning Methods

The fact that the cost-efficient power allocation game $\mathfrak{G}$ admits a unique equilibrium for almost every channel realization is significant from the point of view of managing the system because it guarantees a unique stable solution. That said, it is far from clear how the system's users could actually reach this equilibrium state, so our goal in this section will be to provide a distributed, adaptive learning mechanism that can be employed by the system's users in order to reach this stable state.

In the absence of power considerations, [11] examined this problem by means of a continuous-time learning scheme based on the *replicator dynamics* of evolutionary game theory [18] driven by the users' so-called marginal utility functions $\frac{\partial r_k}{\partial p_{k\mu}}$. Unfortunately however, this approach cannot be applied in our case because replicator-driven techniques require the problem's state space to be a product of simplices – which, in our case, would amount to players saturating the total power constraint (7) by default.

To overcome this issue, let $p_{k,0}$ denote the *unused* power of user k, i.e.

$$p_{k,0} = P_k - p_k = P_k - \sum_{\mu \in \mathcal{M}} p_{k\mu}, \tag{12}$$

so that

$$\sum_{\alpha=0}^M p_{k\alpha} = P_k, \tag{13}$$

for all $k \in \mathcal{K}$. Accordingly, letting the index "0" denote a virtual, "unused" channel and writing $\mathcal{M}_0 = \mathcal{M} \bigcup \{0\} = \{0, 1, \ldots, M\}$ for the system's artificially augmented channel set, the concave problem (11) may be reformulated as:

$$\begin{aligned} \text{maximize} \quad & \Psi_0(\mathbf{p}_1, \ldots, \mathbf{p}_K), \\ \text{subject to} \quad & \mathbf{p}_k \in \Delta_k \equiv \{\mathbf{p}_k \in \mathbb{R}^{\mathcal{M}_0} : \\ & p_{k\alpha} \geq 0 \text{ and } \sum_{\alpha \in \mathcal{M}_0} p_{k\alpha} = P_k\}, \end{aligned} \tag{14}$$

where now $\mathbf{p}_k = (p_{k,0}, p_{k,1}, \ldots, p_{k,M})$ and

$$\begin{aligned} \Psi_0(\mathbf{p}_1, \ldots, \mathbf{p}_M) = \sum_{\mu=1}^M \log\left(1 + \sum_{k \in \mathcal{K}} g_{k\mu} p_{k\mu}/\sigma_\mu^2\right) \\ - \sum_{k \in \mathcal{K}} c_k(P_k - p_{k,0}), \end{aligned} \tag{15}$$

Hence, drawing on the analysis of [11] for power allocation problems with fixed transmit power p_k, we will consider here the *marginal utilities*:

$$v_{k\alpha} = \frac{\partial u_k}{\partial p_{k\alpha}} = \begin{cases} c'_k(P_k - p_{k,0}) & \text{if } \alpha = 0, \\ g_{k\mu}/(\sigma_\mu^2 + \sum_\ell g_{\ell\mu} p_{\ell\mu}) & \text{otherwise.} \end{cases} \tag{16}$$

Importantly, these marginal utilities can be calculated by each user with only local information at hand (such as SINR measurements). Indeed, $v_{k,0}$ only depends on the user's total transmit power $p_k = P_k - p_{k,0}$ so the same applies to the user's cost function c_k; furthermore, for $\mu = 1, \ldots, M$, some easy algebra yields

$$v_{k\mu} = \frac{1}{p_{k\mu}} \frac{\text{sinr}_{k\mu}}{1 + \text{sinr}_{k\mu}}, \tag{17}$$

so any learning scheme that relies on these marginal utilities may be implemented in a completely distributed fashion.

Remark 2. A case of particular interest is when the users' cost functions are linear, viz. $c_k = \lambda_k p_k$ where λ_k denotes the cost incurred by the user (monetary or otherwise) per Watt. In this context, the user's marginal cost $v_{k,0}$ will be:

$$v_{k,0} = \lambda_k, \tag{18}$$

so a higher price per Watt increases the user's tendency to allocate power to the "unused" channel.

In view of the above, we will consider the following exponential learning process in continuous time:

$$\dot{y}_{k\alpha} = v_{k\alpha},$$
$$p_{k\alpha} = P_k \frac{\exp(y_{k\alpha})}{\sum_{\beta \in \mathcal{M}_0} \exp(y_{k\beta})}. \tag{XL_0}$$

Of course, in the above formulation, "power" allocated to the virtual, "unused" channel 0 means "power unused due to cost considerations"; accordingly, by exploiting the properties of the exponential map, we may reformulate this learning scheme as:

$$\dot{z}_{k\mu} = v_{k\mu} - v_{k,0},$$
$$p_{k\mu} = P_k \frac{\exp(z_{k\mu})}{1 + \sum_{\nu \in \mathcal{M}} \exp(z_{k\nu})}. \tag{XL}$$

This last process admits the following reinforcement interpretation: each (actual) channel $\mu \in \mathcal{M}$ is scored by aggregating the difference between its marginal utility $v_{k\mu}$ and the marginal power consumption cost $v_{k,0}$, and power is allocated with exponential sensitivity to these cumulative performance scores.

The first thing that can be verified with respect to the exponential learning scheme (XL) is that it respects the constraints imposed by the users' power considerations: indeed, $p_{k\mu} \geqslant 0$ by definition and $\sum_\mu p_{k\mu} = P_k \sum_\mu \exp(z_{k\mu})/(1 + \sum_\nu \exp(z_{k\nu})) \leqslant P_k$ for any possible value of the performance scores $z_{k\mu}$. More importantly, as the next proposition shows, the learning scheme (XL) guarantees that users converge to a Nash equilibrium of the energy-efficient power allocation game $\mathfrak{G}$:

Proposition 2. *Let* $\mathbf{p}(t)$ *be the adaptive power allocation policy induced by the continuous-time learning scheme* (XL) *for some initialization* $z_{k\mu}(0)$ *of the channels' performance scores. Then, for almost every realization of the system's channel coefficients* $h_{k\mu}$, *we will have* $\lim_{t\to\infty} \mathbf{p}(t) = \mathbf{p}^*$ *where* $\mathbf{p}^*$ *denotes the game's* (*unique*) *Nash equilibrium.*

Moreover, $\mathbf{p}(t)$ *converges to* $\mathbf{p}^*$ *exponentially fast:*

$$D_{\mathrm{KL}}(\mathbf{p}^* \| \mathbf{p}(t)) = \mathcal{O}(e^{-ct}), \tag{19}$$

where $c > 0$ *and* $D_{\mathrm{KL}}(\mathbf{p}^*, \mathbf{p}) = \sum_{k,\mu} p_{k\mu}^* \log\left(p_{k\mu}^*/p_{k\mu}\right)$ *denotes the Kullback-Leibler divergence between* $\mathbf{p}^*$ *and* $\mathbf{p}$.

Sketch of proof: By decoupling the exponential learning scheme (XL), it can be shown that its solution trajectories $\mathbf{p}(t)$ satisfy an augmented version of the replicator equation of [11] with an extra strategy to account for the "unused power" channel $0 \in \mathcal{M}_0$. Our claim may then be proved by adapting the proof of Theorem 6 in [11]. ∎

IV. Algorithmic Implementation and Robustness

Despite its appealing convergence properties, (XL) is a dynamical system that evolves in continuous time, so it is not clear if it can be implemented as a bona fide, discrete-time learning algorithm. In this regard, there are two key challenges to overcome: *a)* to establish a properly discretized version of (XL) which retains its convergence in discrete time; and *b)* to ensure the algorithm's robustness in the presence of impefect CSI and noisy SINR observations.

To that end, we will work here with the following stochastic diminishing-step discretization of (XL):

$$z_{k\mu}(n) = z_{k\mu}(n-1) + \gamma_n[\hat{v}_{k\mu}(n) - \hat{v}_{k,0}(n)],$$
$$p_{k\mu}(n+1) = P_k \frac{\exp(z_{k\mu}(n))}{1 + \sum_\nu \exp(z_{k\nu}(n))}, \tag{20}$$

where $n = 1, 2, \ldots$, is the iteration counter of the process, γ_n is a variable step size whose role will be discussed below and $\hat{v}_{k\alpha}$ represents a perturbed version of the user's marginal utility at the n-th update period. In particular, to account for as wide a range of measurement errors as possible, we will assume that

$$\hat{v}_{k\mu}(n) = v_{k\mu}(n) + \xi_{k\mu}(n) \tag{21}$$

where the observational error $\xi_{k\mu}$ is a bounded martingale difference (not necessarily i.i.d.), i.e. $\|\xi_{k\mu}\| \leqslant \Sigma$ for some $\Sigma > 0$ and $\mathbb{E}[\xi_{k\mu}(n) \,|\, \xi_{k\mu}(n-1), \ldots, \xi_{k\mu}(1)] = 0$.

In this way, we obtain the following adaptive algorithm for cost-efficient power allocation in multicarrier systems:

Algorithm 1 Exponentially-Driven Power Allocation

Parameter: step size γ_n (default: $\gamma_n = 1/n$).
Initialize: $n \leftarrow 0$; scores $z_{k\mu} \leftarrow 0$ for all $k \in \mathcal{K}, \mu \in \mathcal{M}$.
Repeat
  $n \leftarrow n + 1$;
  for each user $k \in \mathcal{K}$ **do simultaneously**
    set transmit power $p_{k\mu} \leftarrow P_k \dfrac{\exp(z_{k\mu})}{1 + \sum_\nu \exp(z_{k\nu})}$;
    measure $\mathrm{sinr}_{k\mu}$;
    update marginal utilities: $v_{k\mu} \leftarrow \dfrac{1}{p_{k\mu}} \dfrac{\mathrm{sinr}_{k\mu}}{1 + \mathrm{sinr}_{k\mu}}$;
    update marginal power cost: $v_{k,0} \leftarrow c_k'(P_k - p_k)$;
    update scores: $z_{k\mu} \leftarrow z_{k\mu} + \gamma_n[v_{k\mu} - v_{k,0}]$;
    until termination criterion is reached.

By employing the stochastic optimization techniques of [13], it is then possible to show:

Proposition 3. *Let* γ_n *be a variable step size sequence such that* $\sum_n \gamma_n \to +\infty$ *and* $\sum_n \gamma_n^2 < +\infty$. *Then, for almost every realization of the system's channel coefficients* $h_{k\mu}$, *the iterates of Algorithm 1 with imperfect measurements given by* (21) *converge to the game's* (*unique*) *Nash equilibrium.*

Sketch of proof: Algorithm 1 can be seen as a greedy mirror descent method [13] for Ψ with respect to the L^1 norm on $\mathcal{X}$ and with the Shannon–Gibbs entropy as a "distance-generating function" in the sense of [13]. With this in mind, the

analysis of [13] can be used to show that $\mathbb{E}[\Psi_0(\mathbf{p}(n)) - \Psi_0(\mathbf{p}^*)] = \mathcal{O}\left(\sum_{m=1}^n \gamma_m^2 / \sum_{m=1}^n \gamma_m\right)$, where $\mathbf{p}^*$ is the game's unique equilibrium. In turn, this implies that $\mathbf{p}(n) \to \mathbf{p}^*$ and establishes our claim. ∎

V. Numerical Results

To validate the predictions of Section IV for the performance of Algorithm 1 in multicarrier wireless systems where power consumption carries a non-negligible cost, we conducted extensive numerical simulations from which we illustrate here a selection of the most representative scenarios.

Throughout this section, and unless explicitly stated otherwise, we will assume a population of $K = 10$ users and $M = 20$ subcarriers, while the channel gain coefficients $h_{k\mu}$ will be drawn randomly from $[0, 1]$. For simplicity, we also assume symmetric channels and users, i.e. $\sigma_\mu = 1$ for all $\mu \in \mathcal{M}$ and $P_k = 1$ W for all $k \in \mathcal{K}$.

With regards to the cost function c_k of the utility model (7), we will consider three distinct cases: no pricing (NP), linear pricing (LP), and nonlinear pricing (NLP). Specifically:

1) The NP model is defined trivially as $c_{\mathrm{NP}}(p) = 0$.
2) The LP model is defined as $c_{\mathrm{LP}}(p) = \lambda p$ for some $\lambda > 0$.
3) The NLP model is defined as $c_{\mathrm{NLP}}(p) = \lambda(e^{p/P} - 1)$.[1]

To develop some intuition, we first provide some general results on the evolution of the cost-aware power allocation game $\mathfrak{G}$ for $K = 2$ non-cooperative users that split their transmitting power over $M = 4$ channels, using only local SINR measurements. Specifically, in Figs. 1, we plotted the total allocated power (3) and achieved transmitted rate (6) for User 1 (solid lines) and User 2 (dashed lines), with channel gains $h_{1\mu}$ and $h_{2\mu}$ respectively. Without loss in generality, we took $h_{1\mu} > h_{2\mu}$, that is, the channel quality for User 1 is better that the one experienced by User 2.

In Fig. 1 we show how the users converge to the game's Nash equilibrium by means of the distributed exponential learning scheme (XL). As can be seen, User 1 achieves a higher transmission rate than User 2 due to the better channel status experienced by User 1 (which, correspondingly, requires User 1 to spend more power than User 2). Moreover, we also investigate the implication of pricing on game's outcome for different values of the pricing parameter λ. As can be seen, in the NP scenario, each user saturates his total allocated power as predicted by the form of the rate function (6); otherwise, in both the LP and NLP schemes, the cost of power consumption causes users to transmit at lower powers and the exponential learning algorithm (XL) converges to a more cost-efficient power allocation (depending on both the channel quality and the considered pricing model). As expected, the LP scheme allows higher rates than the NLP one, while higher values of the pricing parameter λ force users to allocate less power on each channel – thus reducing their individual transmission rates.

Fig. 2 shows how the users' average transmit power, achieved transmission rate and sum-rate evolve at each iteration of the

[1]Note that the constant term in $c_{NLP}(p)$ does not affect the equilibrium of the game as it represents a flat power rate for each user.


Fig. 1. Total allocated power and transmission rate of User 1 (solid line) and User 2 (dashed line) as a function of different pricing models and values of the pricing parameter λ.


Fig. 2. Evolution of the average allocated power, sum-rate and average transmission rate for the LP (solid line) and NLP (dashed line) models for different network configurations.

learning algorithm (XL). As expected, the end-state of the algorithm depends quite strongly on the number of users and available channels: as could be expected, best performance is achieved in the uncontested regime where the number of available channels is higher than the number of users trying to access them, i.e. $K/M < 1$. Fig. 2 also shows how the NLP model affects the power allocation process; in fact, since NLP leads to a sharp increase of the transmission cost for higher powers, it follows that the corresponding loss in transmission rate is not negligible compared to the LP scheme.

In Fig. 3 we plot the sum-rate, average allocated power and its respective cost as a function of different pricing models and values of the pricing parameter λ for different network configurations. The most interesting result is that three distinct regions can be identified: a) For λ below a certain threshold λ_l, the transmission cost in negligible compared to the contribution

Fig. 3. Sum-rate, average allocated power and respective cost as a function of the pricing parameter λ for different channel configurations.

Fig. 4. Evolution of the equilibration rate for different pricing models and values of the step-size γ_n.

of the transmission rate in the utility function, so pricing does not impact the system's performance at equilibrium; *b)* in the second region, say $\lambda_l < \lambda < \lambda_u$, the average allocated power and the sum-rate decrease with λ, thus leading to a nontrivial trade-off between achievable transmission rate and the cost of power consumption; finally, for large $\lambda > \lambda_u$ the transmission cost is so high that it ends up dominating each user's utility function, so users remain relatively quiet due to the high cost of power consumption. An important result is that the average cost paid by users for each transmission is maximized at $\lambda = \lambda_l$ (recall that for $\lambda \leqslant \lambda_l$ the exact value of λ does not affect network performance and users don't mind paying a small cost in order to maximize their throughput). This result could be interesting and useful in all of those scenarios where the receiver of the uplink channel, e.g., the network operator which sells its channels by applying a pricing model, wants to maximize its revenues while maintaining high network performance.[2] Fig. 3 also shows that LP performs better than NLP in terms of the users' transmission rate, precisely because users incur a higher cost under the NLP model (so users will allocate less power on each available channel thus decreasing their achievable transmission rates).

We also investigated the impact of different values of the ratio between the number of channels M and the number of users K on the system's overall performance. As shown in Fig. 3, the congested regime $K > M$ leads to worse aggregate throughput values, as expected; in fact, an increase in the number of users reduces the SINR of each user transmitting on the same channel. On the other hand, if $K \leqslant M$ there is a reduction in the multiuser interference, so higher transmission rates can be achieved for all users.

Finally, to assess the algorithm's convergence speed, we plotted the system's equilibration level (EQL), defined as follows:

$$\text{EQL}(n) = \frac{\Psi_0(\mathbf{p}(n)) - \Psi_{0,\min}}{\Psi_{0,\max} - \Psi_{0,\min}} \qquad (22)$$

[2]Note that the revenue/profit maximization occurs when $\lambda = \lambda_l$ but, this problem lies beyond the scope of the current paper.

where $\mathbf{p}(n)$ is the users' power profile at the n-th iteration of the exponential learning (XL) and $\Psi_{0,\min}$ (resp. $\Psi_{0,\min}$) denotes the maximum (resp. minimum) value of the game's potential Ψ_0 [11]. In Fig. 4 we plot the equilibration rate of (XL) for different pricing models and values of the step-size parameter γ_n. In the case of a diminishing step size, our exponential learning scheme converges to equilibrium more slowly than if a constant step size is used, and the algorithm's convergence speed increases with γ_n. In Fig. 4 we further investigate the impact of different pricing strategies in the game. When power consumption carries no cost (the NP regime), the algorithm converges very fast to an equilibrium point which saturates the users' power constraint; otherwise, the algorithm's convergence to equilibrium is slower in the LP than in the NLP case, a phenomenon which implies that the optimal choice of step size for the algorithm depends delicately on the users' pricing scheme.[3]

In Fig. 5 we plot the game's equilibration level (EQL) under the linear pricing model for $\gamma_n = 1.25, 2.5$ and different values of M and K: as expected, the convergence rate with $\gamma_n = 2.5$ is faster compared to the one achieved for $\gamma_n = 1.25$. The most interesting result concerns the scalability of the algorithm: as a matter of fact, the algorithm converges to equilibrium within a number of iterations that is roughly independent of the underlying network configuration. To better understand the algorithm's scalability, we plotted in Fig. 6 the number of iterations needed to reach the equilibrium for different network configurations as a function of the step-size parameter. Specifically, we considered $K/M = 0.5$ and varied M to study the impact of the number of users on the algorithm's convergence rate: Fig. 6 shows that the proposed distributed algorithm scales well with the number of users, especially if higher values of the step-size

[3]Importantly, the convergence of Algorithm 1 with a variable step-size $\gamma_n = 1/n$ is guaranteed by Proposition 3; this result does not apply to the constant step-size case, but, nonetheless, the algorithm converged to the game's (unique) Nash equilibrium in all our simulations.

Fig. 5. Evolution of the equilibration rate for the LP model for different values of the step-size γ_n and network configurations.


Fig. 6. Evolution of the equilibration rate for different pricing models and network configurations as a function of the step-size γ_n.

γ_n are used in the learning process. More precisely, even for a large number of users, the algorithm converges to the game's equilibrium within a few iterations, and this convergence rate is approximately independent of the exact number of users.

VI. Conclusions

In this paper we examined the problem of transmit power allocation in energy constrained in uplink multi-carrier networks. Assuming that mobile terminals can split their power over several non-interfering channels, we investigated the implication of energy-driven limitations in multi-carrier power allocation scenarios (reflecting e.g. the cost of power consumption, limited battery life, or the use of high-priority power-hungry applications in a fixed power device). In the case where the effective cost of power is given by a general convex pricing function, the resulting non-cooperative power allocation game admits a unique equilibrium for almost every channel realization.

Furthermore, we also proposed a distributed dynamic transmit policy based on exponentially learning which allows users to reach the game's equilibrium exponentially fast, even with (arbitrarily) imperfect CSI; in practice, our numerical analysis shows that users reach an equilibrium within a few iterations of the algorithm. In future work, we plan to extend this result to the study of temporally fluctuating price constraints so as to account for fading, user mobility and other variability phenomena.

References

[1] G. J. Foschini and Z. Miljanic, "A simple distributed autonomous power control algorithm and its convergence," *IEEE Trans. Veh. Technol.*, vol. 42, pp. 641–646, November 1993.

[2] V. Chandrasekhar, J. G. Andrews, and A. Gatherer, "Femtocell networks: A survey," *IEEE Commun. Mag.*, vol. 46, no. 9, pp. 59–67, 2008.

[3] G. Scutari, S. Barbarossa, and D. Palomar, "Potential games: A framework for vector power control problems with coupled constraints," in *Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on*, vol. 4, May 2006, pp. IV–IV.

[4] T. Alpcan, T. Başar, R. Srikant, and E. Altman, "CDMA uplink power control as a noncooperative game," *Wireless Networks*, vol. 8, pp. 659–670, 2002.

[5] C.-W. Sung and W. S. Wong, "A noncooperative power control game for multirate CDMA data networks," *IEEE Trans. Wireless Commun.*, vol. 2, no. 1, pp. 186–194, January 2003.

[6] F. Meshkati, M. Chiang, H. V. Poor, and S. C. Schwartz, "A game-theoretic approach to energy-efficient power control in multicarrier CDMA systems," *IEEE J. Sel. Areas Commun.*, vol. 24, pp. 1115–1129, 2006.

[7] S. T. Chung, S. J. Kim, J. Lee, and J. M. Cioffi, "A game-theoretic approach to power allocation in frequency-selective Gaussian interference channels," in *ISIT '03: Proceedings of the 2003 International Symposium on Information Theory*, 2003.

[8] G. A. Gupta and S. Toumpis, "Power allocation over parallel Gaussian multiple access and broadcast channels," *IEEE Trans. Inf. Theory*, vol. 52, no. 7, pp. 3274–3282, 2006.

[9] F. Meshkati, H. V. Poor, and S. C. Schwartz, "Energy-efficient resource allocation in wireless networks: an overview of game-theoretic approaches," *IEEE Signal Process. Mag.*, vol. Special Issue on Resource-Constrained Signal Processing, Communications and Networking, May 2007.

[10] C. U. Saraydar and D. G. Narayan B. Mandayam, "Efficient power control via pricing in wireless data networks," *IEEE Trans. Commun.*, vol. 50, no. 2, pp. 291–303, February 2002.

[11] 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.

[12] G. Bacci, E. V. Belmega, and L. Sanguinetti, "Distributed energy-efficient power optimization in cellular relay networks with minimum rate constraints," in *ICASSP '14: Proceedings of the 2014 International Conference on Acoustics, Speech and Signal Processing*, to appear.

[13] A. S. Nemirovski, A. Juditsky, G. G. Lan, and A. Shapiro, "Robust stochastic approximation approach to stochastic programming," *SIAM Journal on Optimization*, vol. 19, no. 4, pp. 1574–1609, 2009.

[14] D. Monderer and L. S. Shapley, "Potential games," *Games and Economic Behavior*, vol. 14, no. 1, pp. 124 – 143, 1996.

[15] A. Neyman, "Correlated equilibrium and potential games," *International Journal of Game Theory*, vol. 26, pp. 223–227, 1997.

[16] W. H. Sandholm, *Population Games and Evolutionary Dynamics*, ser. Economic learning and social evolution. Cambridge, MA: MIT Press, 2010.

[17] 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.

[18] P. D. Taylor and L. B. Jonker, "Evolutionary stable strategies and game dynamics," *Mathematical Biosciences*, vol. 40, no. 1-2, pp. 145–156, 1978.