

Dynamic Power Allocation Games in Parallel Multiple Access Channels

Panayotis Mertikopoulos^{*}
Department of Economics
École Polytechnique
91128 Palaiseau, France

Elena V. Belmega[†]
Laboratoire des Signaux et
Systèmes (LSS)
CNRS, Supélec,
Univ. Paris Sud 11
Palaiseau, France

Aris L. Moustakas[‡]
Department of Physics
University of Athens
15784, Athens, Greece

Samson Lasaulce[§]
Laboratoire des Signaux et
Systèmes (LSS)
CNRS, Supélec,
Univ. Paris Sud 11
Palaiseau, France

ABSTRACT

We analyze the distributed power allocation problem in parallel multiple access channels (MAC) by studying an associated non-cooperative game which admits an exact potential function. Even though games of this type have been the subject of considerable study in the literature [1–4], we find that the sufficient conditions which ensure uniqueness of Nash equilibrium points typically do not hold in this context. Nonetheless, we show that the parallel MAC game admits a unique equilibrium almost surely, thus establishing an important class of counterexamples where these sufficient conditions are not necessary. Furthermore, if the network’s users employ a distributed learning scheme based on the replicator dynamics, we show that they converge to equilibrium from almost any initial condition, even though users only have local information at their disposal.

1. INTRODUCTION

As a result of the massive scale at which wireless networks are deployed and operate, non-cooperative game theory is

rapidly becoming one of the main tools with which to describe and analyze distributed resource allocation problems in this context. The reason for this is simple: whereas solution concepts and centralized optimization protocols which depend on global information are very hard to justify or implement (especially in real time or in the presence of a large number of users), game theory offers a way to look at the problem from a more distributed and localized point of view which is often of great applicational relevance.

A prime example of this can be seen in the huge corpus of literature surrounding power allocation games in static Gaussian multi-user networks with the objective of reaching a Shannon-efficient state. The common characteristic of all these games is that the interference between multiple transmissions gives rise to non-trivial interactions between transmitters and imposes a bottleneck on the network performance: interference forces the power allocation policy of one user to depend on the power allocations of all other users. So, following [5], and given that the network users are left to optimally manage their resources on their own, the main questions that arise are *a*) whether there exist “equilibrium” allocations which are stable against unilateral deviations; *b*) whether these (Nash) equilibria are unique; and *c*) whether these equilibria can be reached by distributed (learning) algorithms which require only local information.

The two most important multi-user network models that have been studied from this perspective are the interference channel (IC) [6] and the multiple access channel (MAC) [7], two models which are inherently different from a communications point of view. On the one hand, the IC is composed of several non-cooperative transmitter-receiver pairs and the information-theoretic capacity region is still an open issue for this channel model; in fact, even in the simple case of single-input, single-output (SISO) two-user Gaussian IC only the achievable rates are known [6, 8, 9]. On the other hand, the MAC is composed of several transmitters and a single receiver which must decode the incoming messages,

^{*}panayotis.mertikopoulos@polytechnique.edu

[†]belmega@lss.supelec.fr. The work of this author was partially supported by the French L’Oréal program “For young women doctoral candidates in science” 2009.

[‡]arism@phys.uoa.gr

[§]lasaulce@lss.supelec.fr

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Valuetools 2010 Paris, France

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

and its capacity region is relatively well-understood [10–12], something which remains an open problem for the IC.

Perhaps the most general non-cooperative power allocation games studied in the context of static channels are those presented by Scutari et al. in a series of seminal papers [1–4] focusing on the static Gaussian IC where receivers employ the single-user decoding (SUD) scheme which treats incoming signals from other users as additive noise. There, the existence of a Nash equilibrium (in the “pure” sense of Rosen) is a consequence of the convexity properties of the users’ achievable rates and follows directly from Theorem 1 in [14]; in fact, under suitable (but stringent) conditions on the channel matrices, this equilibrium solution is unique.

Unfortunately, there are two issues with the approach of Scutari et al.: first, as the authors themselves admit, these sufficient conditions “may not be easy to check” [13, p. 1925] and, indeed, in most cases they are not (calculating the spectral radius of a matrix is very hard for large matrices). Secondly, these conditions are not necessary, so when they fail, the uniqueness issue is left wide open. In the specific case of two-user parallel IC, some progress has been made in [15], where the authors completely characterize the set of Nash equilibria. Depending on the geometric properties of the best-response functions (which are identical to the water-filling operators of [4]), the power allocation game may have one, two, three or an infinite number of Nash equilibria. Finally, in [16], assuming that the interference links in one of the bands are negligible, the game is shown to have strategic complementarities and the Nash set is studied using the super-modular property of the game.

A most interesting special case of these more general games consists of the parallel MAC power allocation games which are used to model uplink communication in multi-cellular wireless networks composed of several *nodes* (receivers, access points, base stations, etc.) that operate in orthogonal frequency bands. From a mathematical point of view, the results of [4] obviously apply to the MAC as well, but, as we shall see, the sufficient conditions of [1] are never met in the parallel MAC case, making them irrelevant to games of this type. To compensate for this, the authors of [17] considered two different power allocation games in parallel multiple access channels, depending on the users’ action sets: *i*) the users may distribute their available power among the wireless nodes; or *ii*) the users simply choose a node. There, for the first game (which is more relevant for realistic power allocation scenarios), the Nash equilibrium is argued to be unique, but the proof provided in [17] actually holds only under very restrictive conditions (otherwise, the authors’ strict convexity arguments break down).

In this paper, we analyze non-cooperative power allocation games in parallel multiple access channels with the standard assumption of single user decoding (SUD) at the receiver.¹ As in the more general MIMO MAC case, the parallel MAC game admits an exact potential (in the sense of [20]) whose extrema correspond to the system’s sum capac-

¹ More efficient decoding techniques such as successive interference cancellation can also be considered [18, 19], but optimality with respect to Shannon achievable rates will not concern us here; instead, the low level of signalling and decoder complexity of the SUD makes it more suitable for learning purposes. Furthermore, when using successive interference techniques, the exact potential property of the game is lost in general.

ity, and which can also be interpreted as the system achievable sum-rate if users were employing successive interference cancellation (SIC). Since this potential function is convex, the game’s Nash equilibria will correspond to the minima of the potential, so the game’s Nash set is necessarily convex and compact. However, we find that the game’s potential is, in general, *not* strictly convex (this was the mistake of [17]), so one would expect that uniqueness of Nash equilibria fails along with the sufficient conditions of [1]. Rather surprisingly, we find that this is not the case: *even though the conditions of [1] do not hold, the Nash equilibrium of the game is unique (a.s.)*.

As far as convergence to equilibrium is concerned, one of the main results of [21] is that if the transmitters know the local channel state and the overall interference-plus-noise covariance matrix, then the iterative sequential water-filling algorithm converges to the set of equilibrium points. On the other hand, asynchronous water-filling is harder to analyze because the sufficient conditions of [4] are typically not satisfied in the parallel MAC case. Finally, in a setting similar to our own (incorporating pricing but restricted to only one receiver), the authors of [22] have considered update algorithms which converge to equilibrium modulo certain conditions which do not always hold either.

Instead of taking a water-filling approach, we present a learning scheme based on the replicator dynamics of evolutionary game theory [23] which only requires the players to know their channel coefficients and their rates. Dynamics of this sort have been studied extensively in finite Nash games (that is, games with multilinear payoff functions over a strategy space which is a product of simplices – see e.g. [24] for a survey) and in *continuous* population games [23, 25], but, in the case of finite *nonlinear* games (such as the one we have here), their properties are not as well understood. The first step in that direction consists of identifying the correct modified version of the users’ payoff functions which allows the replicator dynamics to behave well with respect to the solution concepts of the underlying game – in more “traditional” finite player games, this purpose is served by the payoffs that correspond to the pure strategies of the game, but here we have no such structure. Our main contribution is to then show that in parallel MAC power allocation games, the replicator dynamics converge to an equilibrium point unconditionally, *even in the zero-probability event where the game has multiple equilibria*.

Notational Conventions

Throughout this paper, we will use bold uppercase letters to denote matrices and a dagger “†” to denote the Hermitian transpose of a complex matrix.

If $\mathcal{S} = \{s_\alpha\}_{\alpha=1}^n$ is a finite set, we will denote by $K\mathcal{S}$ the disjoint union (categorical coproduct) $K\mathcal{S} \equiv \coprod_{k=1}^K \mathcal{S}$ of K copies of \mathcal{S} . Also, recall that the (real) vector space spanned by \mathcal{S} is defined as the space $\mathbb{R}^{\mathcal{S}} \equiv \text{Hom}(\mathcal{S}, \mathbb{R})$ of functions $x : \mathcal{S} \rightarrow \mathbb{R}$, equipped with the usual operations of addition and scalar multiplication of functions. The canonical basis $\{e_\alpha\}_{\alpha=1}^n$ of $\mathbb{R}^{\mathcal{S}}$ then consists of the indicator functions $e_\alpha : \mathcal{S} \rightarrow \mathbb{R}$ which take the value $e_\alpha(s_\alpha) = 1$ and vanish otherwise. Hence, under the natural identification $s_\alpha \mapsto e_\alpha$, we will use the index α to refer interchangeably to either s_α or e_α , depending on the context. Similarly, we will also identify the set $\Delta(\mathcal{S})$ of probability measures on \mathcal{S} with the standard $(n-1)$ -dimensional simplex of $\mathbb{R}^{\mathcal{S}}$:

$\Delta(\mathcal{S}) \equiv \{x \in \mathbb{R}^{\mathcal{S}} : \sum_{\alpha} x_{\alpha} = 1 \text{ and } x_{\alpha} \geq 0\}$.

Finally, as far as players and their strategies are concerned, we will consistently employ Latin indices for players (k, ℓ, \dots), while reserving Greek ones for their (“pure”) strategies (α, β, \dots).

2. THE SYSTEM MODEL

Following [17], the basic setup of our model is as follows: we have a set $\mathcal{K} = \{1, \dots, K\}$ of finitely many wireless (single-antenna) transmitters – the *players* of the game – that wish to connect to a network of wireless nodes $\mathcal{A} = \{1, \dots, A\}$ (for instance, a collection of base stations or access points). For simplicity, we are assuming that these nodes operate at distinct, non-interfering frequency bands, so that a user $k \in \mathcal{K}$ may split his transmitting power among the nodes $\alpha \in \mathcal{A}$ subject to the power constraint:

$$\sum_{\alpha} p_{k\alpha} \leq P_k, \quad (1)$$

where $p_{k\alpha}$ is the power with which user k transmits towards node α and P_k is the user’s maximum transmitting power. As a result, the *power allocation* of the k -th user will be represented by the point $p_k = \sum_{\alpha} p_{k\alpha} e_{\alpha} \in \mathbb{R}^{\mathcal{A}}$, while, in obvious notation, the corresponding *power profile* which collectively reflects all of the users’ power allocations will be represented by $p = (p_1, \dots, p_K) \in \mathbb{R}^{K\mathcal{A}}$.

Thus, under the standard assumption of single user decoding (SUD), the spectral efficiency of user k in the power profile p will be given by [17, 22]:

$$u_k(p) = \sum_{\alpha \in \mathcal{A}} u_{k\alpha}(p) = \sum_{\alpha \in \mathcal{A}} b_{\alpha} \log \left(1 + \frac{g_{k\alpha} p_{k\alpha}}{\sigma_{\alpha}^2 + \sum_{\ell \neq k} g_{\ell\alpha} p_{\ell\alpha}} \right), \quad (2)$$

where:

1. $b_{\alpha} = B_{\alpha}/B > 0$ is a normalized version of the bandwidth B_{α} of the node $\alpha \in \mathcal{A}$, rescaled to unity by the total bandwidth factor $B = \sum_{\alpha} B_{\alpha}$.
2. $g_{k\alpha} > 0$ is the channel gain of user k with respect to node α , assumed here to be static for the duration of the transmission, known to user k , and drawn from a continuous (and nonatomic) probability distribution on the positive real numbers – see also the relevant assumptions in [4, 13].
3. $\sigma_{\alpha}^2 > 0$ represents the noise level associated to node α (typically the variance of a Gaussian noise process).

Remark 1. It should be noted here that when the wireless users are spatially distributed, the set \mathcal{A} of wireless nodes need not be common to all users. As it turns out, it is not too hard to extend our analysis and results to this more general case, but, to keep our presentation as clear as possible, we will only consider the case where every user can reach every node.

Remark 2. We should also stress here that the channel gain coefficients $g_{k\alpha}$ are the only stochastic parameters in our model, and, in our static channel setting, they are given by $g_{k\alpha} = |h_{k\alpha}|^2$, where h is a realization of the continuous random matrix which describes the channel – see also [17]. So, unless explicitly mentioned otherwise, any probabilistic statement we make in this paper will refer to the probability law of the random variables $g_{k\alpha}$.

Now, as intuition would suggest (and as was shown rigorously in [17]), when the users’ utility is based solely on their spectral efficiency (2), it is clearly to the users’ best interest to transmit at the highest possible total power, i.e. satisfying (1) as an equality.² As a result, we obtain the following components of a normal form game \mathfrak{G} :

1. The set of *players* of \mathfrak{G} is $\mathcal{K} = \{1, \dots, K\}$.
2. The *strategy space* of player k is the (scaled) simplex $\Delta_k \equiv \{p_k \in \mathbb{R}^{\mathcal{A}} : p_{k\alpha} \geq 0 \text{ and } \sum_{\alpha} p_{k\alpha} = P_k\}$; as is customary, we will denote the game’s space of strategy profiles $p = (p_1, \dots, p_K)$ by $\Delta \equiv \prod_k \Delta_k$.
3. The players’ *payoffs* (or *utilities*) are given by the spectral efficiencies $u_k : \Delta \rightarrow \mathbb{R}$ of (2).

Of course, the game \mathfrak{G} defined in this way is not finite (in the original sense of [26]) because *a*) the players are not mixing over a finite set of possible actions; and *b*) even though the players’ strategy spaces happen to be simplices, their payoffs are not multilinear over them. On the other hand, since Δ is a convex polytope and the utilities u_k of the users are concave functions of their power allocations p_k , we immediately see that the game \mathfrak{G} is *concave* in the sense of Rosen [14]. Moreover, it was shown in [17] that \mathfrak{G} is actually an *exact potential* game,³ i.e. that it admits a (global) potential function $\Phi : \Delta \rightarrow \mathbb{R}$ such that:

$$u_k(p_{-k}; p'_k) - u_k(p_{-k}; p_k) = \Phi(p_{-k}; p_k) - \Phi(p_{-k}; p'_k), \quad (3)$$

for all players $k \in \mathcal{K}$, and for all power allocations $p_k, p'_k \in \Delta_k$ of user k and $p_{-k} \in \Delta_{-k} \equiv \prod_{\ell \neq k} \Delta_{\ell}$ of k ’s opponents $\mathcal{K}_{-k} \equiv \mathcal{K} \setminus \{k\}$.⁴

In fact, the authors of [17] provided the following explicit form for the potential function Φ :

$$\Phi(p) = - \sum_{\alpha} b_{\alpha} \log \left(\sigma_{\alpha}^2 + \sum_k g_{k\alpha} p_{k\alpha} \right). \quad (4)$$

For posterity, note here that Φ is itself convex, but not necessarily *strictly* so:⁵ indeed, any two power profiles $p, p' \in \Delta$ such that $\sum_k g_{k\alpha} p_{k\alpha} = \sum_k g_{k\alpha} p'_{k\alpha}$ for all $\alpha \in \mathcal{A}$ will also have $\Phi(p) = \Phi(p')$. This simple observation will be of crucial importance in determining the Nash set of the game, so we will pause here to introduce the concept of *degeneracy*.

To that end, let $T_p \Delta$ denote the tangent space of Δ at p . Since Δ is an affine polytope embedded in $\mathbb{R}^{K\mathcal{A}}$, it is easy to see that for every interior point $p \in \text{Int}(\Delta)$, $T_p \Delta$ will be isomorphic to the subspace which is “parallel” to the polytope Δ :

$$T_p \Delta \cong Z = \{z \in \mathbb{R}^{K\mathcal{A}} : \sum_{\alpha} z_{k\alpha} = 0 \text{ for all } k \in \mathcal{K}\}. \quad (5)$$


However, as we just noted, some of these $KA - K$ directions will be *degenerate* (or *redundant*), in the sense that the po-

²Of course, this need not be true if the cost of power consumption is too high [22], but we will not deal with this issue here.


³In the finite player sense of Monderer and Shapley [20], and not in the continuous sense of [25].

⁴The change of signs in (3) from [20] is deliberate. Our convention was chosen so as to conform with physics, where it is the *minima* of the potential function that are stable.

⁵This is precisely the subtle mistake that underlies the equilibrium uniqueness argumentation of [17].



(a) $\text{ind}(\mathfrak{G}) = 0$: generic (non-convex) level sets



(b) $\text{ind}(\mathfrak{G}) > 0$: degeneration into affine sets

Figure 1: The level sets (dashed blue lines) of the potential function Φ in a 2×2 game with and without degenerate directions (Figs. 1(b) and 1(a) respectively). Degeneracy essentially nullifies degrees of freedom and introduces redundant directions in the system.

tential Φ remains constant as we move along them. Specifically, the set of (almost surely independent) constraints

$$\sum_k g_{k\alpha} z_{k\alpha} = 0, \quad \alpha \in \mathcal{A}, \quad (6)$$

cuts itself a $(KA - A)$ -dimensional subspace W of \mathbb{R}^{KA} whose intersection with Z will correspond to the total of $K + A$ constraints:

$$a) \quad \sum_{\alpha} z_{k\alpha} = 0, \quad k \in \mathcal{K}; \quad (7a)$$

$$b) \quad \sum_k g_{k\alpha} z_{k\alpha} = 0, \quad \alpha \in \mathcal{A}. \quad (7b)$$

Of course, the K tangent space constraints (7a) are set in stone while the A degeneracy constraints (7b) depend on the realization of the channel gains $g_{k\alpha}$.⁶ It is thus possible (though improbable) that some of the constraints (7) are linearly dependent. To keep track of all this, we have:

Definition 1. The subspace $W \leq \mathbb{R}^{KA}$ defined by the constraints (7b) will be called the space of *degenerate* (or *redundant*) directions of the game \mathfrak{G} . Moreover, we define the *degeneracy* (or *redundancy*) *index* of \mathfrak{G} to be:

$$\text{ind}(\mathfrak{G}) \equiv \dim(W \cap Z), \quad (8)$$

where Z is the tangent space determined by the admissibility constraints (7a).

Example 1. As we just saw, $\text{ind}(\mathfrak{G}) = KA - K - A$ (a.s.), so there is no degeneracy in games with $K = 2$ users and $A = 2$ nodes. However, if the channel gains of the two users happen to be linearly dependent (a zero-probability event to be sure, but one which could be approximated reasonably well by strongly collocated users), then we can have degenerate directions even in a 2×2 game (see Fig. 2). In that

⁶These conditions are remarkably similar to the MIMO rank condition $\text{rank}(\mathbf{H}^\dagger \mathbf{H}) = \sum_{k=1}^K n_{t,k} \leq n_r + K$ (where $\mathbf{H} = [\mathbf{H}_1, \dots, \mathbf{H}_K]$ is the system's channel matrix) which ensures that there exists a unique Nash equilibrium [27].

case, the potential function Φ ceases to be strictly convex, so there is no a priori reason that the potential's minimum set will be a singleton.

3. EQUILIBRIUM ANALYSIS

In this section, our main goal will be to describe the Nash set of the game and, more specifically, to show that it consists (almost surely) of a unique equilibrium point which is located at a face of the strategy space Δ .

This problem has attracted considerable interest in the literature, where the papers by Scutari et al. [1–4] provide a set of sufficient conditions for uniqueness in more general interference channel scenarios, and, more recently, in [17], where the authors focus on the parallel MAC problem. Unfortunately, these approaches are problematic (for different reasons): on the one hand, we will see that the sufficient conditions of Scutari et al. [1–4] do not hold in our setting; on the other hand, the uniqueness proof of [17] is only valid under the extremely restricting condition that the game is *non-degenerate*, i.e. that $KA \leq K + A$ (otherwise, the potential Φ is not strictly convex).

Indeed, especially this last condition holds for very few scenarios, only one of which is (barely) non-trivial: *a*) when we have $K = 1$ user and an arbitrary number $A \geq 1$ of nodes (in which case the problem reduces to an optimization one solved by water-filling [28]); *b*) when an arbitrary number of $K \geq 1$ users transmits with the maximum possible power to a single node shared by all ($A = 1$; this was also the scenario studied by [22] who also introduced a linear pricing scheme to compensate for power costs); and *c*) in the $2x2$ case which is easy to solve directly. In spite of the above, our main result in this section is that the sufficient conditions of Scutari are actually far from necessary: *for (almost) any realization of the channel gain parameters $g_{k\alpha}$, there exists a unique Nash equilibrium.*

3.1 Nash Equilibrium Conditions

Since we have a finite number of players, the notion of Nash equilibrium takes the form of stability in the face of unilateral deviations. More specifically:

Definition 2. We will say that the power profile $q \in \Delta$ is at *Nash equilibrium* in the game \mathfrak{G} when

$$u_k(q) \geq u_k(q_{-k}; q'_k), \quad (9)$$

for all $k \in \mathcal{K}$, and for every deviation $q'_k \in \Delta_k$ of player k .

In particular, if q satisfies the strict version of the inequalities (9), then it will be called a *strict equilibrium* of \mathfrak{G} .

As is standard in convex potential games, to calculate the Nash set $\Delta^* \equiv \Delta^*(\mathfrak{G})$ of the game, we only need to look at the (necessarily convex) minimum set of the potential function Φ . To that end, the first order constrained Karush-Kuhn-Tucker minimization conditions [17] show that a power profile $q \in \Delta$ will be at Nash equilibrium if and only if:

$$a) \quad \lambda_k - \frac{b_\alpha g_{k\alpha}}{\sigma_\alpha^2 + \sum_\ell g_{\ell\alpha} q_{\ell\alpha}} \geq 0 \quad (10a)$$

$$b) \quad q_{k\alpha} \left(\lambda_k - \frac{b_\alpha g_{k\alpha}}{\sigma_\alpha^2 + \sum_\ell g_{\ell\alpha} q_{\ell\alpha}} \right) = 0, \quad (10b)$$

for all players $k \in \mathcal{K}$ and all nodes $\alpha \in \mathcal{A}$ (and with the obvious constraints imposed by the condition $q \in \Delta$).

An obvious observation that can be gleaned from the above is that if q is a Nash equilibrium, then either *a*) the support $\text{supp}(q_k) \equiv \{\alpha \in \mathcal{A} : q_{k\alpha} > 0\}$ of a user's power allocation is a singleton (i.e. the user only transmits to a single node); or *b*) we will have the "waterfilling" condition:

$$\frac{g_{k\alpha}}{g_{k\beta}} = \frac{r_\alpha}{r_\beta} \quad \text{for all } \alpha, \beta \in \text{supp}(q_k), \quad (11)$$

where r_α is the user-independent quantity given by:

$$r_\alpha^{-1} \equiv \frac{b_\alpha}{\sigma_\alpha^2 + \sum_\ell g_{\ell\alpha} q_{\ell\alpha}}. \quad (12)$$

In other words, if user k connects to more than one node and is at equilibrium, then he must be "waterfilling" the quantity $g_{k\alpha}/r_\alpha$ among the nodes that he employs.

A promising way to determine whether our game admits a *unique* equilibrium is to take advantage of the plethora of sufficient conditions that have been established in the literature for this purpose. In our setting, the condition which is easiest to check was the one proposed by Scutari et al. in [1, Equation (21)], and which takes the form:

$$\rho(\mathbf{S}_{\max}) < 1, \quad (\text{Cmax})$$

where $\rho(\mathbf{S}_{\max}) = \max\{|\lambda| : \lambda \in \text{eig}(\mathbf{S}_{\max})\}$ is the *spectral radius* (i.e. the eigenvalue with the largest modulus) of the $K \times K$ matrix $\mathbf{S}_{\max} = \{S_{k\ell}^{\max}\}$ defined as:

$$S_{k\ell}^{\max} = \begin{cases} 0, & k = \ell, \\ \max_\alpha \{g_{\ell\alpha}/g_{k\alpha}\}, & k \neq \ell. \end{cases} \quad (13)$$

However, since $\max_\alpha \{g_{\ell\alpha}/g_{k\alpha}\} = (\min_\alpha \{g_{k\alpha}/g_{\ell\alpha}\})^{-1} \geq (S_{\ell k}^{\max})^{-1}$, we immediately see that the entries of \mathbf{S}_{\max} satisfy the inequality $S_{k\ell}^{\max} S_{\ell k}^{\max} \geq 1$ for any distinct pair of users $k, \ell \in \mathcal{K}$. Hence, $\text{tr}(\mathbf{S}_{\max}^2)$ will be bounded from below by:

$$\text{tr}(\mathbf{S}_{\max}^2) = \sum_{k, \ell} S_{k\ell}^{\max} S_{\ell k}^{\max} \geq K(K-1), \quad (14)$$

and, by the spectral radius bounds of [29], we will have:

$$\rho(\mathbf{S}_{\max}) \geq \frac{|\text{tr}(\mathbf{S}_{\max})|}{S} + \sqrt{\frac{\text{tr}(\mathbf{S}_{\max}^2) - \text{tr}(\mathbf{S}_{\max})^2/S}{S(S-1)}}, \quad (15)$$

where $S = \text{rank}(\mathbf{S}_{\max})$.⁷ However, since $\text{tr}(\mathbf{S}_{\max}) = 0$ by definition, (15) gives $\rho(\mathbf{S}_{\max}) \geq 1$, so the sufficient condition (Cmax) fails.

Other sufficient conditions were put forth in [2–4] and [30], based on the matrices $\mathbf{S}(\alpha) = \{S_{k\ell}(\alpha)\}$ defined as:

$$S_{k\ell}(\alpha) = \begin{cases} 0, & k = \ell, \\ g_{\ell\alpha}/g_{k\alpha}, & k \neq \ell. \end{cases} \quad (16)$$

In particular, it was shown in [2] that if:

$$\rho(\mathbf{S}(\alpha)) < 1 \quad \text{for all } \alpha \in \mathcal{A}, \quad (\text{C1})$$

then the game \mathfrak{G} admits a unique Nash equilibrium; in a similar vein, the authors of [30] proposed the condition:

$$\mathbf{I} + \mathbf{S}(\alpha) \succcurlyeq 0 \quad \text{for all } \alpha \in \mathcal{A}, \quad (\text{C2})$$

where " $\succcurlyeq 0$ " signifies positive-definiteness.

Of these two conditions (C1) is *stronger* than (Cmax) in the sense that (Cmax) is sufficient for (C1). However, the same analysis as before shows that in the case of the $\mathbf{S}(\alpha)$ matrices, (14) holds as an equality, so we still get $\rho(\mathbf{S}(\alpha)) \geq 1$ for all $\alpha \in \mathcal{A}$, causing (C1) to fail. Similarly, even though the positive-definiteness condition (C2) is independent of (C1) and (Cmax), the definition of $\mathbf{S}(\alpha)$ yields $S_{k\ell}(\alpha) + S_{\ell k}(\alpha) \geq 2$ for all $k \neq \ell$. Consequently, the element with the largest modulus of the symmetrized matrix $\mathbf{I} + \frac{1}{2}(\mathbf{S}(\alpha) + \mathbf{S}^\dagger(\alpha))$ does not lie on the main diagonal, so the matrix $\mathbf{I} + \mathbf{S}(\alpha)$ cannot be positive-definite either.

Remark. Strictly speaking, condition (C1) was phrased in [2] in terms of a slightly different version of the matrix $\mathbf{S}(\alpha)$ where $S_{k\ell}(\alpha) = 0$ whenever the channel of α is "too bad" for either k or ℓ (in a sense made precise in [2]). In this more general setup, if α is "bad" for user k , then the k -th row and k -th column of $\mathbf{S}(\alpha)$ vanish, so the bound (14) is decreased to $(K-r)(K-r-1)$, where r is the number of zero rows and columns that were introduced in $\mathbf{S}(\alpha)$. However, this also reduces the rank of $\mathbf{S}(\alpha)$ accordingly, so, assuming that $\text{rank}(\mathbf{S}(\alpha)) \geq 2$, the bound (15) still gives $\rho(\mathbf{S}(\alpha)) \geq 1$.

Of course, this still leaves open a small window where the condition (C1) might be salvaged – namely the rare occurrence where the $\mathbf{S}(\alpha)$ matrices all have rank 1 or less. However, instead of focusing on this very special case, we note that even the extensive numerical simulations of [2] show that the sufficient condition (C1) almost never holds in the parallel MAC setting. Indeed, if we follow [2] and assume for simplicity that the transmitter-receiver distances are all equal ($d_{qr} = d_{rq}$ in their notation), then the "normalized interlink distance" becomes equal to 1 and Figure 1 of [2] reveals that (C1) fails almost surely.

We thus see that, despite their theoretical value, the sufficient conditions that have been established in the literature are quite problematic in the parallel MAC setting because they are typically never met (except possibly in some very

⁷Strictly speaking, (15) holds if $S \geq 2$, but we can trivially disregard the case $S < 2$, because $S = K$ almost surely and the one-user case holds little interest.

special cases). Therefore, in order to address the uniqueness issue in complete generality, we will need to develop a different set of tools.

3.2 Representing Power Profiles as Graphs

As we shall see, the “waterfilling” conditions (11) impose some pretty severe constraints on the structure of the equilibrium set Δ^* , because whenever a user waterfills between nodes, his channel gains must “split”, i.e. be of the form $g_{k\alpha} = \lambda_k r_\alpha$. This is actually best understood pictorially, by representing a power profile $p \in \Delta$ as a graph:

Definition 3. We will say that the (multi)graph $\mathcal{G} \equiv (\mathcal{V}, \mathcal{E})$ represents the power profile $p = (p_1, \dots, p_K) \in \Delta$ if:

1. $\mathcal{V} = \mathcal{A}$: the nodes of \mathcal{G} coincide with the network’s;
2. for each $k \in \mathcal{K}$, there is a node $\alpha \in \mathcal{A}$ (called the *hub* of user k in \mathcal{G}) to which the user assigns positive power $p_{k\alpha} > 0$, and which is joined by an edge of \mathcal{G} to every other node $\beta \in \text{supp}(p_k) \setminus \{\alpha\}$.

In simpler words, to represent a power profile $p \in \Delta$ as a graph, one merely has to take the set of wireless nodes as the set of the graph’s nodes, and then, for every user $k \in \mathcal{K}$, to pick a node which the user employs and connect it with an edge to every other node to which the user assigns positive power. Of course, depending on the choice of “hub” for each user $k \in \mathcal{K}$, one might end up with non-isomorphic graphs representing the same power profile p . However, this lack of uniqueness will not be important to us, so we will occasionally abuse Definition 3 by using $\mathcal{G}(p)$ to collectively denote *any* graph which represents the profile $p \in \Delta$.

In light of the above, we now state a few key lemmas and corollaries that will be crucial in our efforts to understand the structure of the equilibrium set Δ^* . The first one is an elegant structural property of equilibrium graphs:

Lemma 4. Let $\mathcal{G} \equiv \mathcal{G}(p)$ represent a power profile $p \in \Delta^*$ which is at Nash equilibrium. Then \mathcal{G} is almost surely a forest – that is, \mathcal{G} contains no cycles.

Proof. The intuitive idea behind this lemma is that if there is a cycle, then we can get a chain of fractions $g_{k_1, \alpha_1} / g_{k_0, \alpha_0}$, $g_{k_2, \alpha_2} / g_{k_1, \alpha_1}$, \dots , which will have a product equal to 1 because of the waterfilling condition (11). However, this represents a condition on the g ’s which occurs with zero probability, thus providing a contradiction.

To make this idea precise, assume that \mathcal{G} contains a cycle Γ denoted as a sequence of edges $\Gamma = (e_1, \dots, e_n)$.⁸ Since an edge can only be owned by a single player, this cycle gives rise to a sequence of players which we also denote by $\Gamma = (k_1, \dots, k_n)$.

So, if $(\alpha_0, \alpha_1, \dots, \alpha_n)$ is the corresponding sequence of nodes that Γ passes through (obviously, $\alpha_0 = \alpha_n$), then (11) gives:

$$\frac{g_{k_j, \alpha_j}}{g_{k_j, \alpha_{j-1}}} = \frac{r_{\alpha_j}}{r_{\alpha_{j-1}}}, \text{ for all } j = \{1, \dots, n\}. \quad (17)$$

Therefore, multiplying these n equations together, we get:

$$\frac{g_{k_1, \alpha_1} \dots g_{k_n, \alpha_n}}{g_{k_0, \alpha_0} g_{k_{n-1}, \alpha_{n-1}}} = \frac{r_{\alpha_1} \dots r_{\alpha_n}}{r_{\alpha_0} r_{\alpha_{n-1}}} = 1. \quad (18)$$

⁸Note that keeping track only of the nodes is not enough because two distinct edges might link the same pair of nodes.




Figure 2: A graph representing a power profile in a game with 3 users (red, blue and green) and 4 nodes (α, β, γ and δ). In the profile represented above, the red player uses α, β and δ , blue uses α, β and γ , and green employs β, γ and δ .

Since there are no cancellations in this last equation (recall that the nodes a_j , $j = 1, \dots, n - 1$ of Γ are all distinct), it will describe a measure zero submanifold of the space from which the channel coefficients g are drawn. As a result (18) only holds with probability zero and, hence, the assumption that \mathcal{G} contains a cycle is almost surely false. \square

As an easy corollary of the above lemma, we also get:

Corollary 5. If $p \in \Delta^*$ is an equilibrium power profile, then there are (a.s.) at most $A - 1$ instances of waterfilling (i.e. two nodes employed by the same player).

Proof. Simply note that a forest on A nodes can have at most $A - 1$ edges; our result then follows by recalling that an edge in this context simply represents an instance of waterfilling. \square

From a geometrical point of view, this shows that Nash equilibria can only live on the faces of the strategy space Δ :

Corollary 6. Let $p \in \Delta^*$ be an equilibrium power profile. Then p lies in the interior of an at most $(A - 1)$ -dimensional face of Δ (a.s.).

Proof. Since a user who waterfills between m nodes transmits with zero power towards the rest of the nodes, his power allocation p_k will belong to the interior of an $(m - 1)$ -dimensional face of Δ_k . The result then follows by combining this observation with Corollary 5. \square

We thus see that the Nash set Δ^* has to be contained in the interior of a face of Δ of dimension at most $A - 1$. We will now show that Δ^* is actually a singleton:

Theorem 7. The game \mathcal{G} has a unique Nash equilibrium (a.s.).

Proof. The basic idea of the proof is, essentially, a geometric one. Indeed, assume that there are two distinct equilibrium points, p and p' , sitting at the interior of some d -dimensional face ρ of Δ . By convexity, the linear segment spanned by p and p' will also belong to the Nash set Δ^* which implies that this line segment must (a.s.) lie in the subspace W of degenerate directions of the game.

In other words, we see that if there is not a unique Nash equilibrium, then the subspace W of degenerate directions intersects nontrivially with a d -dimensional face ρ whose interior contains an equilibrium. However, since $\dim(W) =$

$KA - A$ (a.s.), Corollary 6 gives $\dim(\rho) + \dim(W) \leq KA - A + A - 1 < KA$. On the other hand, it is well-known that two generic subspaces of a real vector space V intersect nontrivially if and only if the sum of their dimensions exceeds $\dim(V)$, so, since Δ is embedded in \mathbb{R}^{KA} , we may conclude that ρ and W intersect trivially with probability 1, a contradiction. \square

4. CONVERGENCE TO EQUILIBRIUM

Having determined the properties of the game's unique equilibrium point, our task in this section will be to present a decentralized learning scheme which allows users to converge to this equilibrium point (and to estimate the speed of this convergence).

Given that the structure of the game \mathfrak{G} does not adhere to the (multilinear) setting of Nash [26],⁹ the usual theory of evolutionary "random-matching" games does not apply either. This leaves on a rather unclear position on how to proceed, but since players invariably want to increase their rewards and an increase in payoff is equivalent to a decrease in potential, we will begin by considering the directional derivatives of the potential function Φ :

$$v_{k\alpha}(p) \equiv -\frac{\partial\Phi}{\partial p_{k\alpha}} = \frac{b_\alpha g_{k\alpha}}{\sigma_\alpha^2 + \sum_{\ell \neq k} g_{\ell\alpha} p_{\ell\alpha}}. \quad (19)$$

Clearly, if a player transmits with positive power to node α , then he will be able to calculate the gradient $v_{k\alpha}(p)$ in terms of the observables $p_{k\alpha}$ (the user's power allocation), $g_{k\alpha}$ (his channel gain coefficients), and the spectral efficiency $u_{k\alpha}(p) = b_\alpha \log\left(1 + \frac{g_{k\alpha} p_{k\alpha}}{\sigma_\alpha^2 + \sum_{\ell \neq k} g_{\ell\alpha} p_{\ell\alpha}}\right)$ (2) which user k observes at node α .¹⁰ As a result, any learning scheme which relies only on the $v_{k\alpha}$'s will be inherently distributed in the sense that it only requires information that is readily obtainable by the individual players.

With all this in mind, a particularly simple scheme to follow is that of the replicator dynamics [23] associated with the "marginal payoffs" $v_{k\alpha}$. More specifically, this means that the players update their power allocations according to the differential equation:

$$\frac{dp_{k\alpha}}{dt} = p_{k\alpha}(v_{k\alpha}(p(t)) - v_k(p(t))), \quad (20)$$

where v_k is just the user average $v_k(p) = P_k^{-1} \sum_{\beta} p_{k\beta} v_{k\beta}(p)$.

As usual, the rest points of (20) are characterized by the (waterfilling) property that, for every pair of nodes $\alpha, \beta \in \text{supp}(p)$ to which user k allocates positive power, we will also have $v_{k\alpha}(p) = v_{k\beta}(p)$. Hence, comparing this to the KKT conditions (10), we immediately see that the Nash equilibria of \mathfrak{G} are stationary in the replicator equation (20). This result is well-known in finite Nash games with multilinear payoffs [24] and in continuous population games [25], but the converse does not hold: for instance, every vertex of Δ is stationary in (20), so stationarity of (20) does not imply equilibrium.

Nevertheless, *only* Nash equilibria can be attracting, and, in fact, they attract almost every replicator solution orbit:

⁹Or even the continuous population models of [25]: there are no "node-specific" rewards in our problem like the "phenotype-specific" growth rates of evolutionary biology.

¹⁰Note that this is different from gradient techniques applied to the utility functions themselves, a practice which requires the utility functions to be known.

Theorem 8. *Let $q \in \Delta$ be the unique (a.s.) equilibrium of \mathfrak{G} . Then, every solution orbit of the replicator dynamics (20) which begins at finite Kullback-Leibler entropy from q will converge to it.*

Furthermore, even if the game does not admit a unique equilibrium, every interior trajectory still converges to a Nash equilibrium (and not merely to the Nash set of the game).

Remark. Recall that the Kullback-Leibler divergence (or relative entropy) of p with respect to q is [23]:

$$H_q(p) = \sum_k H_{q_k}(p_k) = \sum_{k,\alpha} q_{k\alpha} \log(q_{k\alpha}/p_{k\alpha}). \quad (21)$$

Clearly, $H_q(p)$ is finite if and only if p_k allocates positive power $p_{k\alpha} > 0$ to all nodes $\alpha \in \text{supp}(q)$ which are present in q_k ; more succinctly, the domain of H_q consists of all power allocations which are absolutely continuous w.r.t. q .

This convergence result (proved in Appendix A) is extremely powerful because it shows that the network's users will eventually settle down to a stable state which discourages unilateral deviations, even though they only have local information at their disposal. The only case that is left open in the above theorem is what happens if the initial K-L entropy of the solution orbit is infinite, i.e. if the users' initial power allocation does not support all of the nodes which are present in equilibrium. If this is the case, then the face-invariance property of the replicator dynamics ($p_{k\alpha}(t) = 0$ iff $p_{k\alpha} = 0$) will prevent the users from settling down to a Nash equilibrium. However, an easy analysis shows that if one takes the reduced game where each user only has access to the nodes to which he initially allocates positive power, then the users will actually converge to an equilibrium of this reduced game:


Proposition 9. *Let $p(0) = (p_1(0), \dots, p_K(0))$ be an initial power allocation profile in the game \mathfrak{G} and let $A_k = \text{supp}(p_k(0)) \subseteq A$. Then, if \mathfrak{G}^0 is a reduced version of \mathfrak{G} which is played over $\Delta \equiv \prod_k P_k \Delta(A_k)$ with payoffs*

$$u_k^0(p) = \sum_{k \in A_k} b_\alpha \log\left(1 + \frac{g_{k\alpha} p_{k\alpha}}{\sigma_\alpha^2 + \sum_{\ell \neq k} g_{\ell\alpha} p_{\ell\alpha}}\right), \quad (22)$$


the replicator dynamics (20) converge to the unique (a.s.) equilibrium of the reduced game \mathfrak{G}^0 .

5. CONCLUSIONS

In this paper, we studied distributed power allocation in parallel multiple access channels, modeling e.g. the problem of uplink communication in networks that consist of wireless receivers that operate in orthogonal frequency bands. Despite the fact that these games are special cases of the more general framework of [1–4], the sufficient conditions provided therein for uniqueness of equilibrium typically fail in our case. Nonetheless, we show that the game *does* admit a unique equilibrium by studying the properties of the game's exact potential function (and correcting a mistake of [17] in the process). Furthermore, by introducing a distributed learning scheme based on the replicator dynamics of evolutionary game theory, we show that users converge to the game's unique (a.s.) equilibrium. This result carries significant applicational potential because it ensures convergence to equilibrium even in decentralized settings where the users only have access to local information (in our case, the channel coefficients $g_{k\alpha}$ and the spectral efficiencies $u_{k\alpha}$).



(a) Global convergence to equilibrium



(b) Convergence to equilibrium in degenerate games

Figure 3: Convergence to equilibrium in 2x2 power allocation games (the dashed grey contours represent the level sets of the game's potential). If the game admits a unique equilibrium (which is almost always the case), then this equilibrium is (almost) globally attracting (Fig. 3(a)). However, even when the game has more than one equilibria (Fig. 3(b)), every interior replicator trajectory converges to an equilibrium point.

Since the parallel MAC game is a special case of the more general IC one, a natural question that arises is whether our analysis extends to this more general case as well. One immediate observation is that the convergence properties of the replicator dynamics are still valid in general convex potential games played over products of simplices, but since the MIMO MAC game is actually played over the polytope of non-negative definite channel matrices with constrained trace, it is not as easy to write a continuous-time equation there. Further directions to be explored include the speed of convergence of the replicator dynamics to equilibrium (which can be shown to be exponentially fast) and the robustness of the replicator dynamics under stochastic disturbances which reflect inaccuracies in the users' observations (e.g. of the channel coefficients $g_{k\alpha}$). However, a discussion of these issues would take us too far afield (and well beyond the space limitations of this paper), so we prefer to postpone them for the future.

APPENDIX

A. PROOF OF CONVERGENCE

This appendix is devoted to the proof of Theorem 8. The basic idea will be to show that the replicator dynamics (20) admit a Lyapunov function, i.e. a non-negative function $f: \Delta \rightarrow \mathbb{R}$ with $f(p) = 0$ if and only if $p = q$ and such that $\dot{f}(p) \equiv \sum_{k,\alpha} \partial f / \partial p_{k\alpha} \dot{p}_{k\alpha} < 0$ for all $p \neq q$.

A particularly appealing candidate is the game's own potential function Φ . Indeed, an easy differentiation of (4)

yields $\partial \Phi / \partial p_{k\alpha} = -\partial u_k / \partial p_{k\alpha}$, so we obtain:

$$\begin{aligned} \frac{d\Phi}{dt} &= -\sum_{k,\alpha} \frac{\partial u_{k\alpha}}{\partial p_{k\alpha}} \frac{dp_{k\alpha}}{dt} \\ &= \sum_{k,\alpha} v_{k\alpha}(p(t)) p_{k\alpha}(t) (v_{k\alpha}(p(t)) - v_k(p(t))) \\ &= \sum_k P_k \left[\sum_{\alpha} \frac{p_{k\alpha}(t)}{P_k} v_{k\alpha}^2(p(t)) - v_k^2(p(t)) \right] \leq 0, \end{aligned} \quad (23)$$

by Jensen's inequality (recall that $\sum_{\alpha} p_{k\alpha} = P_k$). Since this inequality is strict if all the $p_{k\alpha}$ are positive and $v_{k\alpha} \neq v_{k\beta}$ for $\alpha \neq \beta$, this proves convergence to equilibrium when the game only has a unique equilibrium and the game has no degeneracy.

To get the more general case (and, also, for independent interest), it is much more instructive to consider as a Lyapunov candidate the relative entropy H_q itself. Indeed, a simple differentiation gives:

$$\begin{aligned} \frac{dH_q}{dt} &= -\sum_{k,\alpha} \frac{q_{k\alpha}}{p_{k\alpha}(t)} \frac{dp_{k\alpha}}{dt} \\ &= -\sum_{k,\alpha} q_{k\alpha} (v_{k\alpha}(p(t)) - v_k(p(t))), \end{aligned} \quad (24)$$

and, after rearranging the last term, we get:

$$\frac{dH_q}{dt} = \sum_{k,\alpha} (p_{k\alpha}(t) - q_{k\alpha}) v_{k\alpha}(p(t)) \equiv -L_q(p(t)), \quad (25)$$

where

$$L_q(p) = -\sum_{k,\alpha} (p_{k\alpha} - q_{k\alpha}) v_{k\alpha}(p). \quad (26)$$

We are thus left to show that $L_q(p(t)) > 0$ and, to that end, the key observation is that L_q may be interpreted as a directional derivative of Φ . So, let us set $f(\theta) = \Phi(q + \theta z)$,

where $\theta \geq 0$ and z is a vector in the (solid) *tangent cone* $T_q^c \Delta$ of Δ at q :

$$T_q^c \Delta \equiv \{z \in Z : z_{k\alpha} \geq 0 \text{ for all } \alpha \text{ with } q_{k\alpha} = 0\}, \quad (27)$$

i.e. $T_q^c \Delta$ consists of those tangent directions $z \in Z$ which point towards the interior of Δ (recall that q might lie on the boundary of Δ). Clearly then, (26) may be rewritten as:

$$f'(\theta) = \sum_{k,\alpha} \frac{\partial \Phi}{\partial p_{k\alpha}} \Big|_{q+\theta z} z_{k\alpha} = \theta^{-1} L_q(q + \theta z), \quad (28)$$

for all sufficiently small $\theta > 0$ such that $q + \theta z \in \Delta$.

However, since q is the unique minimum of Φ (a.s.), $f(\theta)$ will be convex along any direction $z \in T_q^c \Delta$, so that $\theta f'(\theta) \geq f(\theta) - f(0)$. Hence, if $p = q + \theta z$ is an arbitrary point of Δ , equations (27) and (28) yield the growth estimate:

$$L_q(p) = \theta f'(\theta) \geq f(\theta) - f(0) = \Phi(p) - \Phi(q). \quad (29)$$

This last estimate shows that $L_q(p) \geq 0$ for all $p \neq q$, thus concluding our proof of Theorem 8 for the non-degenerate case (note that then $\Phi(p) - \Phi(q) > 0$ for all $q \neq p$).

To tackle the degenerate case, a semi-definite Lyapunov function (such as the game's potential Φ or the relative entropy H_q) is not enough because it ensures convergence to the set of minimum points and not to an actual point. Clearly, the replicator dynamics in degenerate games might, in principle, exhibit phenomena of this kind. However, there is much more at work in (20) than a single semi-definite Lyapunov function: there exists a whole *family* of such functions, one for each equilibrium q .¹¹

To take advantage of this, it will be useful to shift our point of view to the *evolution function* $\Theta(p, t)$ of the dynamics (20) which describes the solution trajectory that starts at p at time $t = 0$ and which satisfies the consistency condition:

$$\Theta(p, t + s) = \Theta(\Theta(p, t), s) \text{ for all } t, s \geq 0 \text{ and } p \in \Delta. \quad (30)$$

So, fix some initial condition $p \in \text{Int}(\Delta)$ (or, more generally, $p \in \Delta_q$ where Δ_q is the domain of the relative entropy function H_q) and let $p(t) = \Theta(p, t)$ be the corresponding solution orbit. If $q \in \Delta^*$ is Nash, then, in view of the above discussion, the function $V_q(t) \equiv H_q(\Theta(p, t))$ will be decreasing (though, perhaps, not strictly so) and will converge to some $m \geq 0$ as $t \rightarrow \infty$. It thus follows that $p(t)$ converges itself to the level set $H_q^{-1}(m)$.

Suppose now that there exists some increasing sequence of times $t_n \rightarrow \infty$ such that $p_n \equiv p(t_n)$ does not converge to the Nash set $\Delta^* \equiv \Delta^*(\mathfrak{G})$. By compactness of Δ (and by descending to a subsequence if necessary), we may assume that $p_n = \Theta(p, t_n)$ converges to some $p^* \notin \Delta^*$ (but necessarily in $H_q^{-1}(m)$). Hence, for any $t > 0$:

$$\begin{aligned} H_q(\Theta(p, t_n + t)) &= H_q(\Theta(\Theta(p, t_n), t)) \\ &\rightarrow H_q(\Theta(p^*, t)) < H_q(p^*) = m \end{aligned} \quad (31)$$

where the (strict) inequality stems from the fact that $\dot{H}_q < 0$ outside Δ^* . On the other hand, $H_q(\Theta(p, t_n + t)) = V_q(t_n + t) \rightarrow m$, a contradiction.

Since the sequence t_n was arbitrary, this shows that $p(t)$ converges to the set Δ^* . So, let q' be a limit point of $p(t)$

with $p(t'_n) \rightarrow q'$ for some sequence of times $t'_n \rightarrow \infty$. Then, $V_{q'}(t'_n) = H_{q'}(p(t'_n))$ will converge to zero and, with $V_{q'}$ decreasing, we will have $\lim_{t \rightarrow \infty} V_{q'}(t) = 0$ as well. Seeing as $H_{q'}$ only vanishes at q' , we conclude that $p(t) \rightarrow q'$, i.e. every interior trajectory converges to equilibrium.

References

- [1] G. Scutari, D. P. Palomar, and S. Barbarossa, "Competitive design of multiuser MIMO systems based on game theory: a unified view," *IEEE J. Sel. Areas Commun.*, vol. 26, no. 7, pp. 1089–1103, September 2008.
- [2] —, "Optimal linear precoding strategies for wideband non-cooperative systems based on game theory—part I: Nash equilibria," *IEEE Trans. Signal Process.*, vol. 56, no. 3, pp. 1230–1249, March 2008.
- [3] —, "Optimal linear precoding strategies for wideband non-cooperative systems based on game theory—part II: algorithms," *IEEE Trans. Signal Process.*, vol. 56, no. 3, pp. 1250–1267, March 2008.
- [4] —, "Asynchronous iterative waterfilling for Gaussian frequency-selective interference channels," *IEEE Trans. Inf. Theory*, vol. 54, no. 7, pp. 2868–2878, July 2008.
- [5] S. Lasaulce, M. Debbah, and E. Altman, "Methodologies for analyzing equilibria in wireless games," *IEEE Signal Processing Magazine, Special issue on Game Theory for Signal Processing*, vol. 26, no. 5, pp. 41–52, Sep. 2009.
- [6] A. B. Carleial, "Interference channels," *IEEE Trans. Inf. Theory*, vol. 24, no. 1, pp. 60–70, Jan. 1978.
- [7] T. M. Cover and J. A. Thomas, "Elements of information theory," *Wiley Interscience*, 2006.
- [8] T. S. Han and K. Kobayashi, "A new achievable rate region for the interference channel," *IEEE Trans. Inf. Theory*, vol. 27, no. 1, pp. 49–60, Jan. 1981.
- [9] H. Sato, "The capacity of the Gaussian interference channel under strong interference," *IEEE Trans. Inf. Theory*, vol. 27, no. 6, pp. 786–788, Nov. 1981.
- [10] T. Cover, *Some advances in broadcast channels*. Advances in Communication Systems, vol. 4, Academic Press, 1975.
- [11] A. D. Wyner, "Recent results in Shannon theory," *IEEE Trans. Inf. Theory*, vol. 20, pp. 2–10, Jan. 1974.
- [12] R. S. Cheng and S. Verdú, "Gaussian multiaccess channels with ISI: Capacity region and multiuser waterfilling," *IEEE Trans. Inf. Theory*, vol. 39, no. 3, pp. 773–785, May 1993.
- [13] 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.
- [14] J. B. Rosen, "Existence and uniqueness of equilibrium points for concave n-person games," *Econometrica*, vol. 33, no. 3, pp. 520–534, 1965.

¹¹This is also the reason that the relative entropy is a more suitable Lyapunov candidate: the potential has the same value at the *entire* Nash set of the game, while the relative entropy with respect to a point only vanishes at the point itself.

- [15] E. V. Belmega, B. Djeumou, and S. Lasaulce, "Resource allocation games in interference relay channels," in *IEEE International Conference on Game Theory for Networks (GAMENETS)*, Istanbul, Turkey, May 2009, pp. 575–584.
- [16] R. Mochaourab and E. Jorswieck, "Resource allocation in protected and shared bands: Uniqueness and efficiency of Nash equilibria," in *Fourth International Conference on Performance Evaluation Methodologies and Tools (Valuetools)*, Pisa, Italy, Oct. 2009.
- [17] S. M. Perlaza, E. V. Belmega, S. Lasaulce, and M. Debbah, "On the base station selection and base station sharing in self-configuring networks," in *ValueTools '09: Proceedings of the 4th International Conference on Performance Evaluation Methodologies and Tools*, 2009.
- [18] E. V. Belmega, S. Lasaulce, M. Debbah, M. Jungers, and J. Dumont, "Power allocation games in wireless networks of multi-antenna terminals," *Springer Telecommunications Systems Journal*, vol. 2010, doi:10.1007/s11235-010-9305-3, May 2010.
- [19] E. V. Belmega, S. Lasaulce, and M. Debbah, "Power allocation games for MIMO multiple access channels with coordination," *IEEE Trans. Wireless Commun.*, vol. 8, pp. 3182–3192, Jun. 2009.
- [20] D. Monderer and L. S. Shapley, "Potential games," *Games and Economic Behavior*, vol. 14, no. 1, pp. 124 – 143, 1996.
- [21] W. Yu, W. Ree, 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, Jan. 2004.
- [22] 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.
- [23] J. W. Weibull, *Evolutionary Game Theory*. Cambridge, MA: MIT Press, 1995.
- [24] D. Fudenberg and D. K. Levine, *The Theory of Learning in Games*, ser. MIT Press Series on Economic Learning and Social Evolution. The MIT Press, 1998, vol. 2.
- [25] W. H. Sandholm, "Potential games with continuous player sets," *Journal of Economic Theory*, vol. 97, pp. 81–108, 2001.
- [26] J. F. Nash, "Non-cooperative games," *The Annals of Mathematics*, vol. 54, no. 2, pp. 286–295, September 1951.
- [27] E. V. Belmega, "On resource allocation games in MIMO wireless channels," Ph.D. dissertation, Université Paris Sud 11, 2010.
- [28] S. Boyd and L. Vandenberghe, "Convex optimization," *Cambridge University Press*, 2004.
- [29] B. G. Horne, "Lower bounds for the spectral radius of a matrix," *Linear Algebra and its Applications*, vol. 263, pp. 261–273, 1997.
- [30] Z.-Q. Luo and J.-S. Pang, "Analysis of iterative water-filling algorithms for multi-user power control in digital subscriber lines," *EURASIP J. Appl. Signal Process. [Online]*, May 2006.