

Competitive Scheduling in Wireless Collision Channels with Correlated Channel State

Utku Ozan Candogan, Ishai Menache, Asuman Ozdaglar and Pablo A. Parrilo*

Abstract

We consider a wireless collision channel, shared by a finite number of mobile users who transmit to a common base station. Each user wishes to optimize its individual network utility that incorporates a natural tradeoff between throughput and power. The channel quality of every user is affected by global and time-varying conditions at the base station, which are manifested to all users in the form of a common channel state. Assuming that all users employ stationary, state-dependent transmission strategies, we investigate the properties of the Nash equilibrium of the resulting game between users. While the equilibrium performance can be arbitrarily bad (in terms of aggregate utility), we bound the efficiency loss at the best equilibrium as a function of a technology-related parameter. Under further assumptions, we show that sequential best-response dynamics converge to an equilibrium point in *finite* time, and examine how to exploit this property for better network usage.

1 Introduction

Wireless technologies are broadly used today for both data and voice communications. The transmission protocols of wireless devices need to cope with the scarce resources available, such as bandwidth and energy. Additional difficulties relate to the *dynamic* nature of wireless networks. For example, the mobility of terminals and the frequent change in their population introduce new challenges for routing protocols. An additional distinctive dynamic feature of wireless communications is the possible time variation in the channel quality between the sender and the receiver, an effect known as channel fading [3].

Motivated by scalability considerations, it has been advocated in recent years that mobiles should have the autonomy to adjust their transmission parameters (e.g., [7]). This leads to distributed network domains, in which “situation-aware” users take autonomous decisions regarding their network usage, based on the current network conditions and their individual preferences. As users generally do not coordinate their network actions, non-cooperative game-theory has become a natural theoretical framework for analysis and management (see, e.g., [1, 5], and [7] for a survey).

The general framework that will be considered in this work is that of users who obtain state information and may accordingly control their transmission parameters. A major issue in this context is whether such information is well exploited by self-interested users, or rather misused. More specifically, we consider a wireless collision channel, shared by a finite number of users who wish to optimally schedule their transmissions based on a natural tradeoff between throughput and power. The channel quality between each user and the base station is randomly time-varying and observed by the user prior to each transmission decision. The bulk of the research in the area has been carried under a simplified assumption that the channel state processes of different users are independent (see e.g., [2, 6]). In practice, however, there are global system effects, which simultaneously affect the quality of all transmissions (e.g., thermal noise at the base station, or common weather conditions). Consequently, a distinctive feature of our model is that the state processes of different users are *correlated*. As a baseline model, we consider in this paper the case of *full* correlation, meaning that all users observe the same state prior to transmission. A fully correlated state can have a positive role of a coordinating signal, in the sense that different states can be “divided” between different

*All authors are with the Laboratory for Information and Decision Systems, MIT. E-mails: {candogan, ishai, asuman, parrilo}@mit.edu

users. On the other hand, such state correlation increases the potential deterioration in system performance due to non-cooperation, as users might transmit simultaneously when good channel conditions are available.

Our results indicate that both the positive and negative effects of state correlation are possible. User interaction and unwillingness to give up the better-quality states lead to arbitrarily bad equilibria in terms of the aggregate utility. At the same time, there exist good-quality equilibria, whose performance-gap from the social welfare solution can be bounded as a function of a technology-related parameter. In some special cases, we establish that the game under consideration is a potential game [8]. We use this property and additional structure of our game, to establish the *finite-time* convergence of sequential best-response dynamics (although the underlying game is not a finite game). We further demonstrate through simulations that the convergence to an equilibrium can occur very fast. From a network management viewpoint, we indicate how the latter property can be employed to direct the network to good-quality equilibria.

The rest of the paper is organized as follows. The model and some basic properties thereof are presented in Section 2. The social welfare problem is defined and characterized in Section 3. We then proceed in Section 4 to study the efficiency loss under selfish user behavior, as compared to the socially optimal operating point. Section 5 studies the convergence properties of best-response dynamics. Conclusions are drawn in Section 6. Due to the page limitation in the current submission, most of the proofs are excluded from the main text and can be found in an accompanying technical report [4].

2 The Model and Preliminaries

We consider a wireless network, shared by a finite set of mobile users $\mathcal{M} = \{1, \dots, M\}$ who transmit at a fixed power level to a common base station over a shared collision channel. Time is slotted, so that each transmission attempt takes place within slot boundaries that are common to all. A transmission can be successful only if no other user attempts transmission simultaneously. Further details of our model are provided below.

2.1 The Physical Network Model

Our model for the channel between each mobile (or *user*) and the base station is characterized by two basic elements.

a. Channel state process. We assume that the channel state between mobile m and the base station evolves as a stationary process $H^m(t)$, $t \in \mathbb{Z}_+$ (e.g., Markovian) taking values in a set $\mathcal{H}^m = \{1, 2, \dots, h^m\}$ of h^m states¹. The stationary probability distribution that mobile m would observe state $i \in \mathcal{H}^m$ at any time t is given by π_i^m .

b. Expected data rate. We denote by $R_i^m > 0$ the expected data rate (or simply, the rate) that user m can sustain at any given slot as a function of its current state $i \in \mathcal{H}^m$ (which is measured, say, in bits per second). We further denote by $\mathcal{R}^m = \{R_1^m, R_2^m, \dots, R_{h^m}^m\}$ the set of all possible data rates for user m , and define $\mathcal{R} = \mathcal{R}^1 \times \mathcal{R}^2 \times \dots \times \mathcal{R}^M$. For convenience, we assume that the expected data rate R_i^m strictly decreases in the state index i , so that $R_1^m > R_2^m > \dots > R_{h^m}^m$, i.e., state 1 represents the “best state” in which the highest rate can be achieved. A central assumption in this paper is that the state processes of different users are fully correlated, as we formalize below.

Assumption 1 (Full Correlation). *All users observe the same channel state $H(t)$ in any given time t . That is, for every mobile $m \in \mathcal{M}$: (i) $\mathcal{H}^m = \mathcal{H} = \{1, 2, \dots, h\}$, (ii) $\pi_i^m = \pi_i$ for every $i \in \mathcal{H}$ and (iii) $H^m(t) = H(t)$ for every $t \in \mathbb{Z}_+$ (where \mathcal{H} is the common state space, and $\pi = (\pi_1, \dots, \pi_h)$ is its stationary distribution).*

We emphasize that although all mobiles observe the same state, the corresponding rates R_i^m need *not* be equal across mobiles, i.e., in our general model we do *not* assume that $R_i^m = R_i^k$, $m, k \in \mathcal{M}$, $i \in \mathcal{H}$. Thus, our model allows for a heterogeneous mobile population, where each mobile may differ in its transmission power, distance from

¹Note that the actual channel quality may still take continuous value, which each user reasonably classifies into a finite number of information states. See [4] for further interpretation of our model.

the base station, and other factors, which result in different per-user rates. The case where the latter condition does hold will be considered as a special case in Section 5.

2.2 User Objective and Game Formulation

In this subsection we describe the user objective and the non-cooperative game which arises as a consequence of the user interaction over the collision channel. The basic assumption of our model is that users always have packets to send, yet they are free to determine their own transmission schedule in order to fulfill their objectives. Furthermore, users are unable to coordinate their transmission decisions. Our focus in this paper is on *stationary* transmission strategies, in which the decision whether to transmit or not can depend (only) on the current state². A formal definition is provided below.

Definition 1 (Stationary Strategies). *A stationary strategy for user m is a mapping $\sigma^m : \mathcal{H} \rightarrow [0, 1]^h$. Equivalently, σ^m is represented by an h -dimensional vector $\mathbf{p}^m = (p_1^m, \dots, p_h^m) \in [0, 1]^h$, where the i -th entry corresponds to user m 's transmission probability when the observed state is i .*

We use the term multi-strategy when referring to a collection of user strategies, and denote by the vector $\mathbf{p} = \{\mathbf{p}^1, \dots, \mathbf{p}^M\}$ the multi-strategy comprised of all users' strategies. The multi-strategy representing strategies of all users but the m th one is denoted by \mathbf{p}^{-m} . For a given multi-strategy \mathbf{p} , we define below the Quality of Service (QoS) measures that determine user performance. Let W^m be the (fixed) transmission power cost of user m per transmission attempt, and denote by $\tilde{P}^m(\mathbf{p}^m)$ its average power investment, as determined by its strategy \mathbf{p}^m . Then clearly, $\tilde{P}^m(\mathbf{p}^m) = W^m \sum_{i=1}^h \pi_i p_i^m$ for every user m . We normalize the latter measure by dividing it by W^m , and consider henceforth the normalized power investment, given by $P^m(\mathbf{p}^m) = \sum_{i=1}^h \pi_i p_i^m$. For simplicity, we shall refer to $P^m(\mathbf{p}^m)$ as the power investment of user m . We assume that each user m is subject to an individual power constraint $0 < \bar{P}^m \leq 1$, so that any user strategy \mathbf{p}^m should obey

$$P^m(\mathbf{p}^m) \leq \bar{P}^m. \quad (1)$$

The vector of power constraints is denoted by $\bar{\mathbf{P}} = (\bar{P}^1, \dots, \bar{P}^M)$. The second measure of interest is the mobile's average throughput, denoted by $T^m(\mathbf{p}^m, \mathbf{p}^{-m})$. The average throughput of every user m depends on the transmission success probability at any given state i , $\prod_{k \neq m} (1 - p_i^k)$. Hence,

$$T^m(\mathbf{p}^m, \mathbf{p}^{-m}) = \sum_{i=1}^h \pi_i R_i^m p_i^m \prod_{k \neq m} (1 - p_i^k). \quad (2)$$

Each user wishes to optimize a natural tradeoff between throughput and power, which is captured by maximizing the following utility function

$$u^m(\mathbf{p}^m, \mathbf{p}^{-m}) = T^m(\mathbf{p}^m, \mathbf{p}^{-m}) - \lambda^m P^m(\mathbf{p}^m), \quad (3)$$

subject to the power constraint (1), where $\lambda^m \geq 0$ is a user-dependent tradeoff coefficient. We use the notation $\lambda = (\lambda^1, \dots, \lambda^M)$ for the vector of all users' tradeoff coefficients; note that each game *instance* can now be formally described by the tuple $\mathcal{I} = \{\mathcal{M}, \mathcal{R}, \pi, \lambda, \bar{\mathbf{P}}\}$.

The term $\lambda^m P^m(\mathbf{p}^m)$ in (3) can be viewed as the power cost of the mobile. The user utility thus incorporates both a "hard" constraint on power consumption (in the form of (1)), but also accounts for mobile devices that do not consume their power abilities to the maximum extent, as energy might be a scarce resource, the usage of which needs to be evaluated against the throughput benefit. We note that the utility (3) accommodates the following two special cases:

1) *Fully "elastic" users.* By setting $\bar{P}^m = 1$, a user practically does not have a hard constraint on power usage.

²Yet, an equilibrium in stationary strategies remains an equilibrium when considering the larger space of general strategies, see [4].

Accordingly, the optimal operating point of the user is determined solely by the tradeoff between power and throughput, as manifested by the factor λ^m . The fully elastic user case has been considered in the wireless games literature in different contexts (see, e.g., [1]).

2) *Power-cost neutral users.* Consider a user with $\lambda^m = 0$. Such user is interested only in maximizing its throughput subject to a power constraint. This form of utility has been examined, e.g., in [2] and [5].

A strategy \mathbf{p}^m of user m is *feasible* if it obeys the power constraint (1). We denote by E^m the feasible strategy space of user m , $E^m = \{\mathbf{p}^m | P^m(\mathbf{p}^m) \leq \bar{P}^m, 0 \leq \mathbf{p}^m \leq 1\}$; the joint feasible action space is denoted by $E = \prod_{m \in \mathcal{M}} E^m$. A Nash equilibrium point (NE) for our model is a feasible multi-strategy from which no user can unilaterally deviate and improve its utility. Formally, a multi-strategy $\mathbf{p} \triangleq (\mathbf{p}^1, \dots, \mathbf{p}^M)$ is a Nash equilibrium point if $\mathbf{p}^m \in \operatorname{argmax}_{\tilde{\mathbf{p}}^m \in E^m} u^m(\tilde{\mathbf{p}}^m, \mathbf{p}^{-m})$, for every $m \in \mathcal{M}$. The existence of a pure Nash equilibrium point is guaranteed by the Kakutani fixed point theorem. In [4] we show that there are possibly infinitely many equilibria for our game.

We conclude this section by an example which demonstrates that the behavior of the system in an equilibrium can sometimes be counterintuitive.

Example 1. Consider a game with two users, m, k , and two states 1, 2. Let $\pi_1 = \pi_2 = \frac{1}{2}$, $R_1^m = R_1^k = 8$, $R_2^m = R_2^k = 3$, $\lambda^m = \lambda^k = 1$ and $\bar{P}^m = 0.8$, $\bar{P}^k = 0.3$. The unique equilibrium of this game instance is given by $(p_1^m, p_2^m, p_1^k, p_2^k) = (1, 0.6, 0, 0.6)$. Observe that the total power investment at state 1 (0.5) is lower than the total power investment at state 2 (0.6).

The above example demonstrates some negative indications as to the *predictability* of the Nash equilibrium. Not only the number of equilibria is unbounded, but also we cannot rely on monotonicity results (such as total power investment increasing with the quality of the state) for equilibrium characterization. At the same time, these observations motivate the study of performance-loss bounds at *any* equilibrium point, and also of network dynamics that can converge to a predictable equilibrium point. Both directions would be examined in subsequent analysis.

3 Social Welfare and Threshold Strategies

In this section we characterize the optimal operating point of the network. Due to space constraints, we state here only the properties that are required for the equilibrium analysis. Additional characterization of the optimal solution can be found in [4].

An optimal multi-strategy in our system is a feasible multi-strategy that maximizes the aggregate user utility. Formally, \mathbf{p}^* is an optimal multi-strategy if it is a solution to the Social Welfare Problem (SWP), given by $\max_{\mathbf{p} \in E} u(\mathbf{p})$, where

$$u(\mathbf{p}) = \sum_m T^m(\mathbf{p}) - \lambda^m P^m(\mathbf{p}^m). \quad (4)$$

We note that (SWP) is a non-convex optimization problem (see [4] for a formal proof of this property). For a characterization of (SWP), we require the next two definitions.

Definition 2 (Partially and Fully Utilized States). Let \mathbf{p}^m be some strategy of user m . Under that strategy, state i is partially utilized by user m if $p_i^m \in (0, 1)$; state i is fully utilized by the user if $p_i^m = 1$.

Definition 3 (Threshold Policies). A strategy \mathbf{p}^m of user m is a threshold strategy, if the following conditions hold: (i) User m partially utilizes at most one state, and (ii) If user m partially utilizes exactly one state, then the power constraint (1) is active (i.e., met with equality). A multi-strategy $\mathbf{p} = (\mathbf{p}^1, \dots, \mathbf{p}^M)$ is a threshold multi-strategy if \mathbf{p}^m is a threshold strategy for every $m \in \mathcal{M}$.

The main result of this section is summarized below.

Theorem 1. There exists an optimal solution of (SWP) where all users employ threshold strategies.

Due to the non-convexity of (SWP), we cannot rely on first order optimality conditions for the characterization of the optimal solution. Nonetheless, Theorem 1 indicates that there always exist an optimal solution with some well-defined structure, which would be used in the next section for comparing the obtained performance at optimum, to an equilibrium performance.

4 Efficiency Loss

We proceed to examine the extent to which selfish behavior affects system performance. Specifically, we are interested in comparing the quality of the obtained equilibrium points to the centralized, system-optimal solution (defined in Section 3). Recently, there has been much work in quantifying the efficiency loss incurred by selfish behavior of users in networked systems (see [9] for a comprehensive review). The two concepts which are most commonly used in this context are the price of anarchy (PoA) and price of stability (PoS). The standard definitions of these terms consider all possible instances of the associated game. Recall that in our specific framework, a game instance is given by the tuple $\mathcal{I} = \{\mathcal{M}, \mathcal{R}, \pi, \lambda, \bar{\mathbf{P}}\}$. In [4], we show that without any restrictions on \mathcal{I} , even the best performing equilibrium can be arbitrarily bad. However, we note that for a *given* mobile technology, a game instance cannot obtain all possible parameter values. Specifically, π is determined by the technological ability of the mobiles to *quantize* the actual channel quality into a finite number of “information states”. Naturally, one may think of several measures for quantifying the quality of a given quantization. We represent the quantization quality by a single parameter $\pi_{max} \triangleq \max_{i \in \mathcal{H}} \pi_i$ (an implicit assumption here is that π_{max} decreases with the quantization quality, which is plausibly manifested by the number of states h). In addition, a specific wireless technology is obviously characterized by the power constraint \bar{P}^m . Again, we represent the power-capability of a given technology by a single parameter $P_{min} = \min_{m \in \mathcal{M}} \bar{P}^m$. Finally, we determine the *technological quality* of a set of mobiles through the scalar $Q = \frac{\pi_{max}}{P_{min}}$.

We consider the efficiency loss for a given technological quality Q . Denote by \mathcal{I}_{Q_0} the subset of all game instances such that $Q = Q_0$. We define below the price of stability (PoS) and price of anarchy (PoA) as a function of Q .

Definition 4 (Price of Stability - Price of Anarchy). *For every game instance \mathcal{I} , denote by $N_{\mathcal{I}}$ the set of Nash equilibria, and let $p_{\mathcal{I}}^*$ be an optimal multi-strategy. Then for any fixed Q , the PoS and PoA are defined as $PoS(Q) = \sup_{\mathcal{I} \in \mathcal{I}_Q} \inf_{\mathbf{p} \in N_{\mathcal{I}}} \frac{u(p_{\mathcal{I}}^*)}{u(\mathbf{p})}$, $PoA(Q) = \sup_{\mathcal{I} \in \mathcal{I}_Q} \sup_{\mathbf{p} \in N_{\mathcal{I}}} \frac{u(p_{\mathcal{I}}^*)}{u(\mathbf{p})}$.*

We next provide upper and lower bounds for $PoS(Q)$ under the assumption that $Q < 1$. This assumption is justifiable, for example, when the number of states is relatively large.

Theorem 2. *Fix $Q < 1$. Let $\hat{\mathbf{p}}$ be some threshold multi-strategy, and let $u(\hat{\mathbf{p}})$ be the respective aggregate utility (4). Then (i) There exists an equilibrium point $\bar{\mathbf{p}}$ whose aggregate utility is not worse than $u(\hat{\mathbf{p}})(1 - Q)^2$. That is, $\frac{u(\bar{\mathbf{p}})}{u(\hat{\mathbf{p}})} \leq (1 - Q)^{-2}$. (ii) Consequently, $PoS(Q) \leq (1 - Q)^{-2}$.*

The key idea behind the proof is to start from a threshold multi-strategy $\hat{\mathbf{p}}$ and to reach an equilibrium point by some iterative process. At each step of the process we obtain the worst-case performance loss, which leads to the overall loss in the entire procedure. The above result implies that for P_{min} fixed, a finer quantization of the channel quality results in a better upper bound for the PoS, which approaches 1 as $\pi_{max} \rightarrow 0$. In [4], we establish a lower-bound of approximately $(1 - Q)^{-1}$. The gap between the upper and lower bound remains a subject of on-going work. We conclude this section by stating that the PoA is unbounded for any Q (see [4] for a proof).

Proposition 1. *For any given Q , $PoA(Q) = \infty$.*

The above result indicates that despite technological enhancements (which result in a low Q), the network can still arrive at bad-quality equilibria with unbounded performance loss. This negative result emphasizes the significance of mechanisms or distributed algorithms, which preclude such equilibria. We address this important design issue in the next section.

5 Best-Response Dynamics

A Nash equilibrium point for our system represents a strategically stable working point, from which no user has incentive to deviate unilaterally. In this section we address the question of if and how the system arrives at an equilibrium, which is of great importance from the system point of view.

5.1 Convergence Properties

The most natural mechanism for (distributed) convergence to an equilibrium relies on a user’s *best response*, which in general is a user’s strategy that maximizes its own utility, given the strategies of other users. In our game, a best response $\bar{\mathbf{p}}^m$ of user m to a multi-strategy \mathbf{p}^{-m} is given by $\bar{\mathbf{p}}^m \in BR^m(\mathbf{p}^{-m}) = \operatorname{argmax}_{\tilde{\mathbf{p}}^m \in E^m} u^m(\tilde{\mathbf{p}}^m, \mathbf{p}^{-m})$. An informal description of a general best-response mechanism is simple: *Each user updates its strategy from time to time through its best response.*

The best-response mechanism, described above in its most general form, is not guaranteed to converge to an equilibrium in our system without imposing additional assumptions. Our convergence analysis relies on establishing the existence of a *potential* function (see [8]) under a certain condition, referred to as the *rate alignment* condition.

Definition 5 (Rate Alignment Condition). *The set of user rates $\{R_i^m\}_{i \in \mathcal{H}, m \in \mathcal{M}}$ is said to be aligned if there exist per-user positive coefficients $\{c^m\}_{m \in \mathcal{M}}$ and per-state positive constants $\{R_i\}_{i \in \mathcal{H}}$ such that $R_i^m = c^m R_i$ for every $m \in \mathcal{M}$ and $i \in \mathcal{H}$. The rate alignment condition is satisfied if user rates are aligned.*

The coefficient c^m above reflects user m ’s relative quality of transmissions, which is affected mainly by its transmission power and location relative to the base station. While the rate alignment condition might not hold for general and heterogeneous mobile systems, a special case of interest which satisfies the rate alignment condition is the *symmetric-rate* case, i.e., $c^m = c$ for every $m \in \mathcal{M}$. Rate-symmetry is expected in systems where mobiles use the same technology (transmission power and coding schemes), and where “local” conditions, such as distance from the base station, are similar. In order to establish convergence, we require the following set of assumptions.

Assumption 2. (i) Rates are aligned (see Definition 5). (ii) At each time slot, only a single user m may update its strategy, if it strictly improves its utility (3). (iii) Users choose threshold strategies in any best-response update, and the updates take place in a round-robin manner³. (iv) The transmission-success probabilities $\prod_{k \neq m} (1 - p_i^k)$, $i \in \mathcal{H}$ are perfectly estimated by each user m before its strategy update.

Our convergence result is summarized below.

Theorem 3. *Let Assumption 2 hold. Then the best-response mechanism converges in finite time to an equilibrium point. The number of time-steps until convergence is upper bounded by $M(2e)^{M(h+1)}$.*

Proof. (outline) Under the rate-alignment condition, the game is a (weighted) potential game (see [8]), with a potential function given by $\phi(\mathbf{p}) = -\sum_{i=1}^h \pi_i R_i \prod_{m=1}^M (1 - p_i^m) - \sum_{i=1}^h \sum_{m=1}^M \pi_i \frac{\lambda^m}{c^m} p_i^m$, where c^m is user m ’s rate-coefficient (see Definition 5). The convergence of (sequential) best-response dynamics then follows as a property of potential games (see [8]). It can be easily shown that by restricting users to threshold strategies, the underlying game becomes a finite game (i.e., with a finite action space), with the same potential function as above. As such, the finite improvement property (FIP) in potential games holds (see [8]), and convergence is obtained in finite time. The above upper bound on the number of steps until convergence is obtained through a combinatorial argument regarding the possible number of threshold strategies. \square

We emphasize that the restriction to threshold strategies is commensurate with the user’s best interest. Not only there always exists such best-response strategy, but also it is reasonably easier to implement. We discuss next some important considerations regarding the presented mechanism and the assumptions required for its convergence.

³In [4], we consider relaxations of the round-robin ordering assumption (Assumption 2(iii)), and their consequences.

Assumption 2(ii) requires synchronization between the mobiles, which can be done centrally by the base station or by a supplementary distributed procedure. Observe that the schedule of updates is the only item that needs to be centrally determined. Assumption 2(iv) approximates a natural scenario where users update their transmission probabilities at much slower time-scales than their respective transmission rates. We emphasize that users need not be aware of the specific transmission probabilities p_i^m of other users. Indeed, in view of (3), only the transmission-success probabilities $\prod_{k \neq m} (1 - p_i^k)$, $i \in \mathcal{H}$ are required. These can be estimated by sensing the channel and keeping track of idle slots. A last comment relates to the rate-alignment condition. The convergence results in this section rely on establishing a potential function for the underlying game, which is shown to exist when rates are aligned. In an accompanying report [4], we show that in a system of three states or more, the alignment condition is not only sufficient, but also necessary for the existence of a potential. This suggests that novel methods would have to be employed for establishing dynamics convergence under more general assumptions.

5.2 Experiments

The objective in this subsection is to study through simulations the convergence properties of sequential best-response dynamics. More specifically, we wish to examine the dependence of convergence time on several factors, such as the number of users in the system, the number of states, and the technology factor Q . In all our experiments, we consider a relaxed version of Assumption 2, where the rate-alignment condition (Assumption 2(i)) is not enforced.

The specific setup for our simulations is as follows. We assume that $\pi_i = \frac{1}{h}$ for every $i \in \mathcal{H}$. For given Q , M and h , we construct a significant number of game instances (10000) by randomly choosing in each instance the power constraints \bar{P}^m , the tradeoff coefficient λ^m and the associated rates R_i^m for every $m \in \mathcal{M}$, $i \in \mathcal{H}$. We simulate each game instance, and examine the average convergence speed, measured in the number of round-robin iterations (i.e., each user updates its strategy exactly once in every iteration). Figure 1 presents the convergence speed results for two different values of Q , as a function of the number of users in the system. For each value of Q , we consider three different number of states h .




Figure 1: Convergence speed as a function of the number of users.

As seen from Figure 1, the average number of iterations required for convergence is less than three on average. We emphasize that all game instances converge without requiring the rate-alignment condition, indicating the possibility to exclude this condition in future analysis of best-response convergence. We observe that all graph curves initially increase as a function of the number of users, and at some point gradually decrease until reaching a fixed number of iterations. This interesting phenomenon can be intuitively justified as follows: When the number of users is relatively small, there is less competition on each state, and convergence is fast. At the other extreme, when the number of users is larger than some threshold, then there are more users who can fully utilize states at the first iteration (see Definition 2), thereby decreasing the competition at subsequent iterations and leading to faster convergence.

We conclude this section by briefly discussing possible means for obtaining high-quality equilibria in terms of the aggregate utility (4). Theorem 2(i) suggests that if the system is initiated at some threshold multi-strategy, then the equilibrium performance cannot deviate by much, compared to the performance at the initial working

point. Consequently, one may consider an iterative *hybrid* algorithm, in which a network-management entity forces some initial working-point (a threshold multi-strategy), waits enough time until convergence, and if the equilibrium performance is unsatisfactory, enforces a different working point, until reaching a satisfactory equilibrium. Such an algorithm relies on the fast convergence to an equilibrium (demonstrated in all our simulations), which allows to consider numerous initial working points in plausible time-intervals. Further discussion can be found in [4].

6 Conclusion

We have considered in this paper a wireless network game, where mobiles interact over a shared collision channel. The novelty in our model is the state correlation assumption, which incorporates the effect of global time-varying conditions on performance. In general, the correlated state can be exploited by the users for time-division of their transmission, which would obviously increase the system capacity. However, we have shown that under self-interested user behavior, the equilibrium performance can be arbitrarily bad. Nevertheless, the efficiency loss at the best equilibrium can be bounded as a function of a technology parameter, which accounts both for the mobiles power limitations and the level of discretization of the underlying channel quality. Importantly, we have shown that best-response dynamics may converge to an equilibrium in finite time, and empirically verified that such dynamics converge fairly fast. We briefly note several extensions and open directions: The convergence analysis of best-response dynamics under more general conditions is of great interest. As mentioned, new tools rather than the use of a potential function seem to be necessary. Another challenging direction is to obtain a tight bound on the price of stability, and examine how the price of anarchy can be bounded while fixing other game parameters besides the technological quality. At a higher level, one may consider the *partial* correlation case, in which a user reacts to a channel state that incorporates both global and local temporal conditions.

References

- [1] T. Alpcan, T. Başar, R. Srikant, and E. Altman. CDMA uplink power control as a noncooperative game. *Wireless Networks*, 8:659–670, 2002.
- [2] E. Altman, K. Avrachenkov, G. Miller, and B. Prabhu. Discrete power control: Cooperative and non-cooperative optimization. In *INFOCOM*, pages 37–45, 2007.
- [3] E. Biglieri, J. Proakis, and S. Shamai. Fading channels: Information-theoretic and communications aspects. *IEEE Transactions on Information Theory*, 44(6):2619–2692, 1998.
- [4] U. O. Candogan, I. Menache, A. Ozdaglar, and P. A. Parrilo. Competitive scheduling in wireless collision channels with correlated channel state. Technical Report 2804, LIDS, MIT, November 2008. Available from <http://www.mit.edu/~candogan/publications/CorrFadGame.pdf>.
- [5] Z. Luo and J. Pang. Analysis of iterative waterfilling algorithm for multiuser power control in digital subscriber lines. *EURASIP J. Appl. Signal Process.*, 2006(1):1–10.
- [6] I. Menache and N. Shimkin. Efficient Rate-Constrained Nash Equilibrium in Collision Channels with State Information. *INFOCOM 2008*, pages 403–411, 2008.
- [7] F. Meshkati, H. V. Poor, and S. C. Schwartz. Energy-efficient resource allocation in wireless networks. *IEEE Signal Processing Magazine*, 24(3):58–68, 2007.
- [8] D. Monderer and L. S. Shapley. Potential Games. *Games and Economic Behavior*, 14(1):124–143, 1996.
- [9] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.