Reservation-Based Distributed Medium Access in Wireless Collision Channels

Ishai Menache* Nahum Shimkin†

December 13, 2008

Abstract

We consider an uplink wireless collision channel, shared by multiple mobile users. The medium access protocol incorporates channel reservation that relies on RTS (request-to-send) and CTS (clear-to-send) control packets. Consequently, collisions are reduced to the relatively short periods where mobiles request channel use. In our model, users individually schedule their channel requests, and the objective of each user is to minimize its own power investment subject to a minimum-throughput demand. Our analysis reveals that for feasible throughput demands, there exist exactly two Nash equilibrium points in stationary strategies, with one superior to the other uniformly over all users. We then show how this better equilibrium point can be obtained through distributed best-response mechanisms. Finally, we quantify and discuss the effect of the relative length of data and control periods on capacity, power and delay.

*Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139. E-mail: ishai@mit.edu

†Department of Electrical Engineering, Technion, Israel Institute of Technology, Haifa 32000, Israel. E-mail: shimkin@ee.technion.ac.il

1 Introduction

1.1 Background and Motivation

Current wireless networks consist of a relatively large number of users with heterogeneous Quality of Service (QoS) requirements (such as bandwidth, delay, and power). To reduce the management complexity, decentralized control of such networks is often to be preferred to centralized one. This requirement leads to distributed (or at least partially distributed) network domains, in which end-users take autonomous decisions regarding their network usage, based on their individual preferences. This framework is naturally formulated as a non-cooperative game, and has gained much interest in recent literature (e.g., see [8] for a recent collection of papers on game theory in communication systems).

In the context of wireless networks, self-interested user behavior can be harmful, as network resources (such as bandwidth) are often limited, and might be abused by a subset of greedy users. Consequently, a central question that arises in the design and management of networks is the following: What is the *right* degree of freedom that should be given to end-users during network operation? This dilemma is incorporated in mechanism or protocol design, as by restricting users to some protocol rules, an adequate performance level can be preserved, despite user selfishness.

In this paper, we examine a distributed access control mechanism, in which a reservation mechanism is applied to efficiently use the shared medium. Specifically, we adopt 802.11's virtual carrier sense mechanism (see [1] and [5] for a survey): A mobile station ready to transmit, first sends a short control packet, called RTS (Request to Send), which includes the source, destination and requested duration of the data transmission period. The base station then responds to the RTS request by sending an CTS (Clear to Send) packet with the same information as above.

As part of the reservation protocol, a mobile station that hears a CTS message addressed to another station will not use the channel until the end of the current transaction. This virtual carrier sense mechanism essentially reduces the probability of collision to the short duration of the RTS transmission. We note that non-interference with on-going transmissions may be seen to be in the self-interest of any legitimate user of the channel, and strictly so when the protocol mandates priority to interrupted transmissions. We do not consider here malicious users (jammers), whose sole interest is to interfere with the performance of other users.

Reservation-based mechanisms have been considered in recent wireless network literature, not only in the 802.11 context, but also in relation to mobile ad-hoc networks (see [5], [14] and references therein). Several recent papers have analyzed non-cooperative user-behavior in wireless systems, the bulk of which related to CDMA-based domains, which allow for multi-packet reception (see [13] for a survey). Wireless collision channels with self-interested users are studied as well, e.g., in [2, 3, 6]. These papers consider different user-utilities than the one studied here, and usually assume symmetry among users, which might not hold in practice. The basic framework of this paper is similar to our previous work in [10, 12, 11], where the first two papers consider rate-based equilibria in collision channels with fading, and the third considers the same model with power-level control. A more general capture channel model has been considered in [9]. None of the above references, however, has incorporated reservation mechanisms as part of the model, which is the central feature of this contribution.

1.2 Contribution and Paper organization

In this paper we study medium access control to a wireless channel while focusing on the effect of the RTS/CTS reservation mechanism, under the assumption that mobiles are free to schedule their individual channel requests. Distinctive features of our model include the following:

- The objective of each user is to minimize its average power investment subject to satisfying a given throughput demand. As we shall see, this inherently leads to the existence of *multiple* equilibrium points.
- We focus on the simplest scheduling strategy, in which each user sets a fixed (stationary) probability for issuing channel requests[1].

The resulting interaction between users may be viewed as a noncooperative game [4], the properties of which are the main focus of this work. We start by analyzing the equilibrium points of this game, focusing on their number and efficiency properties. Our analysis largely relies on the results of [12], after showing that the present model (with channel reservations) may be reduced to that of [12] through a simple transformation. We show

[1] We note however that the resulting (stationary) equilibrium points remain so even when the restriction to stationary strategies is removed. This follows since the best-response strategy for each user to a *stationary* strategy profile of the others is stationary as well, as may be readily verified.

that for (strictly) feasible throughput demands, there exists exactly two equilibrium points, with one strictly better than the other for *all* users (in terms of transmission attempt rate, hence average power). Based on our analysis, we focus on system *design and control* issues in two significant directions:

- Distributed best-response like mechanisms and their convergence properties.
- The effect of the data transmission period (relative to the control period) on basic performance characteristics, namely channel capacity, average power and transmission delay.

For the standard best-response mechanism, we establish strong convergence results to the *better* equilibrium point: Namely, the mechanism converges to the better equilibrium whenever started below the *worse* equilibrium. This significantly strengthes the analysis and results reported in [12], where convergence was established only for initial conditions below the *better* (i.e., lower) equilibrium. We also consider a simplified mechanism, the naive best response (NBR), that may be more plausible for simple-minded users. While the theoretical guarantees for this variant are weaker, in practice it achieves good performance.

The paper is organized as follows. The networking model, along with basic properties thereof, is presented in Section 2. In Section 3 we analyze the basic equilibrium properties of the underlying noncooperative game. Section 4 focuses on distributed mechanisms for leading the network to an efficient equilibrium point. In Section 5 we address the effect of the data transmission period. Conclusions and future research directions are outlined in Section 6.

2 The Model

Our model consists of a finite set of mobile users $\mathcal{I} = \{1, \ldots, n\}$ who connect to a common base station over a shared wireless channel. We focus on the uplink direction, where users transmit their data to the base station.

2.1 The Network Model

Medium Access Protocol. Time is slotted, and the access to the channel is obtained as follows. A mobile wishing to send data asks for transmission permission from the base

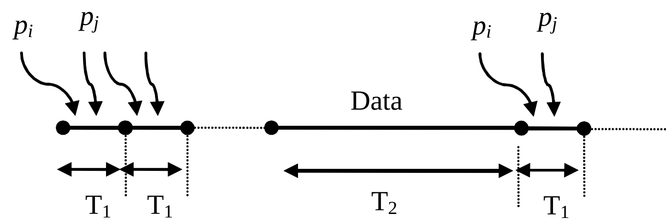

Figure 1: Illustration of the reservation-based medium access control. Each user requests the channel with some probability. In case of a single request, the channel is granted to the requesting user for a period of T_2 slots.

station by sending a short RTS packet. The base station in turn sends a CTS acknowledgement packet granting use of the channel to one station at most. The total duration of this RTS-CTS phase is a constant of T_1 slots duration (whether the base station sends a CTS or not). In case that the channel has not been assigned to any of the stations, the above step is repeated for the same duration of T_1 slots. Otherwise, the user that obtained the CTS occupies the channel for a fixed number of T_2 slots, in which its data is transmitted without interruption. See Figure 1 for a graphic illustration. Although the 802.11 standard allows for reserving the channel for a variable duration, we assume that all mobiles use a fixed-length data transmission interval. This assumption is commensurate with the worst-case scenario where all users always have packets to send, and therefore reserve that channel for the largest allowed duration at each transmission.

Reception Model. We assume the following.

- Simultaneous RTS transmissions of two or more users result in a collision. That is, in case of multiple requests, the base station is unable to recognize which of the users have requested the channel.

- During the data transmission period, the effective data rate of user i is a constant, say D_i, determined by the station characteristics (transmission power, distance from the base station, channel gain, etc.).

2.2 User Model

We associate with each user i a fixed throughput requirement ρ_i, which stands for the (minimal) data rate that this user wishes to maintain[2]. To simplify notations, we normalize the throughput requirement of each user by its effective data rate D_i, so that $0 < \rho_i \leq 1$. Further denote $\underline{\rho} = (\rho_1, \ldots, \rho_n)$ and $\rho = \sum_i \rho_i$. Note that the aggregate throughput ρ is actually upper bounded by $T_2/(T_1 + T_2)$, since each data transmission period must be preceded by a channel request phase.

We assume that a user always has packets to send, yet it may postpone its transmission requests in order to accommodate its required rate ρ_i (as we formally define below). Hence, each user i chooses a request probability p_i, which is the probability for sending an RTS frame in the designated time-period. We further assume that users transmit at a fixed power level (regardless of whether the transmission is an RTS packet or a data packet). Consequently, the strategy profile $\mathbf{p} = (p_1, \ldots, p_n)$ determines the average power investment of each user. To simplify notations, we shall normalize the per-slot power investment of each user's transmission to 1.

The underlying assumption of our model is that users are selfish and do not cooperate or coordinate in any manner in order to satisfy their throughput demands. Obviously, the transmission schedule of each user affects the throughput of all others. This state of affairs gives rise to a conflict situation, which we formalize as a non-cooperative game between the users. We are interested in the Nash equilibrium point of that game. To define the Nash equilibrium, we require some basic quantities, which are derived in the next subsection.

2.3 Basic Performance Measures

In this subsection we obtain explicit expressions for the throughput and power-investment averages, as a function of the stationary strategy profile $\mathbf{p} = (p_1, \ldots p_n)$.

Define

$$q_i \equiv q_i(\mathbf{p}) \stackrel{\triangle}{=} p_i \prod_{j \neq i} (1 - p_j) \,, \tag{1}$$

which is probability that user i alone transmits at a given RTS slot, and hence granted CTS

[2]We view here the throughput requirement as a lower bound mandated by the user application. Alternatively, this rate may be assigned (and policed) as an upper bound by the system as a control and management tool, leading to a similar model. In either case, the rate assignment procedure is exogenous to our model.

permission for the subsequent T_2 interval. Let $r_i(\mathbf{p})$ be user i's throughput (normalized by D_i) as determined by the request probabilities of all other users.

Proposition 1 *The long-term average throughput of user i is given by*

$$r_i(\mathbf{p}) = \frac{q_i T_2}{T_1 + \sum_j q_j T_2}. \tag{2}$$

Proof: Let X_k, $k = 1, 2, \ldots$ be the renewal process whose renewal points are the starting times of the user channel requests (i.e., transmission of an RTS packet by any user). Then $E(X_k) = T_1 + \sum_j q_j T_2$. The expected throughput of user i over the k-th renewal interval may be regarded as a reward R_k earned at the k-th renewal, so that (X_k, R_k) is a reward-renewal process [15]. The expected reward is given by $E(R_k) = q_i T_2$. Then, by Proposition 7.3 in [15],

$$r_i(\mathbf{p}) = \lim_{k\to\infty} \frac{\sum_{m=1}^{k} R(m)}{k} = \frac{E(R_k)}{E(X_k)} = \frac{q_i T_2}{T_1 + \sum_j q_j T_2}. \tag{3}$$

$\square$

Denote by $S_i(\mathbf{p})$ the (long-term) average transmission power investment of user i. For convenience, we normalize the average power and express it as a fraction of the instantaneous power of user i during transmission. Thus, S_i is identical to the fraction of time in which user i is transmitting (either data of control packets). Proceeding as in Proposition 1, we obtain the following expression.

Proposition 2 *The average power investment of user i is given by*

$$S_i(\mathbf{p}) = \frac{p_i \bar{T}_1 + q_i T_2}{T_1 + \sum_j q_j T_2}, \tag{4}$$

where $\bar{T}_1 < T_1$ is the duration of actual transmission of the RTS packet.

2.4 Game Formulation

A Nash equilibrium point (NEP) for our model is a vector of request probabilities $\mathbf{p} = (p_1, \ldots, p_n)$, which is self-sustaining in the sense that all throughput constraints are met, and neither user can lower its average energy investment by unilaterally modifying its transmission request probability. Formally,

Definition 2.1 (Equilibrium points) *A stationary strategy profile* $\mathbf{p} \triangleq (p_1, \ldots, p_n)$ *is a Nash equilibrium point if*

$$p_i \in \operatorname*{argmin}_{0 \le \tilde{p}_i \le 1} \{S_i(\tilde{p}_i, \mathbf{p}_{-i}) : r_i(\tilde{p}_i, \mathbf{p}_{-i}) \ge \rho_i\} \tag{5}$$

for each $i \in \mathcal{I}$*, where* S_i *and* r_i *are defined in (4) and (2), respectively, and* $\mathbf{p}_{-i}$ *stands for the vector* $\mathbf{p}$ *excluding its* i*-th element* p_i*.*

Noting that both functions S_i and r_i are strictly increasing in p_i, it follows that the above minimization over $\tilde{p}_i$ is equivalent to satisfying the second inequality with equality, namely $r_i(\mathbf{p}) = \rho_i$. We thus obtain the following equivalent characterization of a Nash equilibrium point $\mathbf{p}$ in our model:

$$r_i(\mathbf{p}) \equiv \frac{q_i(\mathbf{p})T_2}{T_1 + \sum_j q_j(\mathbf{p})T_2} = \rho_i \tag{6}$$

for every user i. We shall refer to (6) as the *equilibrium equations*.

3 Equilibrium Analysis

We start our analysis by reformulating the equilibrium equations in an equivalent form that is more convenient for our purpose. To that end, recall that

$$\rho = \sum_{i \in \mathcal{I}} \rho_i$$

(the aggregate throughput demand of all users), and let $\alpha = \frac{\rho}{1-\rho}$.

Proposition 3 *The equilibrium equations (6) are equivalent to the following set of equations:*

$$q_i(\mathbf{p}) = \tilde{\rho}_i\,, \quad i \in \mathcal{I} \tag{7}$$

where

$$\tilde{\rho}_i = \frac{\rho_i(1+\alpha)T_1}{T_2} = \frac{\rho_i T_1}{(1-\rho)T_2}\,. \tag{8}$$

Proof: Summing the equilibrium equations (6) over all users i, we obtain

$$\frac{\sum_i q_i T_2}{T_1 + \sum_j q_j T_2} = \rho. \tag{9}$$

Solving for $\sum_i q_i$ gives

$$\sum_i q_i = \frac{\alpha T_1}{T_2}. \tag{10}$$

Substituting (10) into (6) leads to (7). $\square$

The significance of the last proposition is that our equilibrium equations reduce to (a special case of) the model without reservations studied in [12], with the throughput requirements of the users set to the modified rates $\tilde{\rho}_i$. Consequently, the following results from [12] carry over to the current model. Let the *feasible throughput region* designate the set of throughput requirement vectors $(\rho_1, \ldots, \rho_n)$ for which there exists at least one equilibrium point.

Theorem 4 ([12]) *(i) The feasible throughput region is a non-empty and closed set.*

(ii) In every interior point of the feasible region, there exist exactly two Nash equilibria, denoted $\mathbf{p}^a$ and $\mathbf{p}^b$, and one of those (say $\mathbf{p}^a$) is uniformly better than the other ($\mathbf{p}^b$), in the sense that $p_i^a < p_i^b$ for every user i.

(iii) For the better equilibrium point $\mathbf{p}^a$, it holds that

$$\sum_i p_i^a < 1\,. \tag{11}$$

We note that any feasible throughput vector in the modified model must satisfy $\sum_i \tilde{\rho}_i \leq 1$ (fully-utilized channel), which is equivalent here to $\rho = \sum_i \rho_i \leq \frac{T_2}{T_1+T_2}$. Clearly, except for the case of a single user, the total throughput ρ will be smaller due to possible collisions. In fact, it was shown in [12] that for many small users, the maximal channel throughput approaches e^{-1}, which translates to $\rho \approx \frac{T_2}{eT_1+T_2}$ in the present model. Further properties of the throughput region will be discussed in Section 5.1.

Remark 1 *Let us briefly comment on the* computation *of the equilibria. As shown in [12], the two solutions of (7) can be computed by finding the zeros of a certain scalar and unimodal function. Therefore, computation of the Nash equilibria can be efficiently accomplished by standard scalar search techniques. In addition, the better-equilibrium point can be computed through simulating best-response dynamics, as established in Section 4.*

We next provide a simple expression for the average power in equilibrium.

Proposition 5 *The average power investment of user i in an equilibrium point $\mathbf{p}$ is given by*

$$S_i(\mathbf{p}) = \rho_i + \frac{\bar{T}_1}{T_1}(1-\rho)p_i \tag{12}$$

(with $\bar{T}_1$ as defined in Proposition 4).

Proof: Observe that the shared channel is either in reservation mode (where users send channel requests) or in data mode (where users send data packets). In every equilibrium point $\mathbf{p}$, each users consumes a fraction of ρ_i in data mode, which is the first summand in (12). Overall, ρ is the fraction of time in which the channel is in data mode, while a fraction of $1-\rho$ is used for reservations. Since user i's request probability is p_i, its corresponding power investment in that mode is given by $(1-\rho)p_i$, multiplied by the fraction $\frac{\bar{T}_1}{T_1}$ of time in which it actually transmits during the RTS period. □

Note that the expression (12) does not include the parameter T_2. However, the power investment is a function of p_i, which depends on this parameter through the equilibrium equations (7). Observe further that the power investment in equilibrium is linear in p_i. The immediate conclusions from both Theorem 4 and Proposition 5 are summarized below.

Corollary 6 *Let $\mathbf{p}^a$ and $\mathbf{p}^b$ be the two equilibrium points, such that $p_i^a < p_i^b$ for every user i. Then,*

(i) The power investment at $\mathbf{p}^a$ is strictly lower for each user, namely $S_i(\mathbf{p}^a) < S_i(\mathbf{p}^\mathbf{b})$ holds for every i.

(ii) $\sum_i S_i(\mathbf{p}^a) \leq 1$: the total power investment at the better equilibrium $\mathbf{p}^a$ is bounded by 1.

Proof: (i) is immediate from (12) and Theorem 4(ii). Part (ii) follows by summing (12) over all users, using (11), and noting that $\bar{T}_1 < T_1$. □

It is interesting to note that by property (ii) the total (normalized) power at the better equilibrium is never higher than for the single-user, fully utilized channel case. This bound does not hold for the other equilibrium point.

The above results clearly motivate the study and design of distributed algorithms that converge to the *better* equilibrium $\mathbf{p}^a$. This is the subject of the next section.

4 Convergence to the Better Equilibrium

A Nash equilibrium point represents a strategically stable operating point, from which no user has incentive to deviate unilaterally. Still, the question of how the system arrives at an equilibrium needs to be addressed, especially when the users operate in a distributed environment with no cooperation or coordination by a central authority. Furthermore, observing that our system has two Nash equilibria with one strictly better than the other, it is important (from the system viewpoint, as well as for each individual user) to verify that the system converges to the better equilibrium rather than to the worse one.

We start by considering the natural best-response dynamics, where each user reacts to the observed system conditions by adjusting its strategy (i.e., the value of its request probability) to achieve its required throughput. We analyze the convergence of this scheme, and discuss the feasibility of its distributed implementation. We then consider a somewhat simplified scheme, the naive best response, which may seem more natural to the less sophisticated user. We conclude this section with simulation experiments that demonstrate the actual convergence properties of the two schemes.

4.1 Best Response Dynamics

A best-response strategy for a given user is generally its optimal response to a given set of strategies for all other users [4]. In our model, the best response for user i is the transmission request probability that equalizes the user's throughput with its throughput demand ρ_i. Observing (6), the best response $\tilde{p}_i$ of user i to a strategy profile $\mathbf{p} = (p_1, \ldots, p_n)$ is the probability $\tilde{p}_i$ that solves

$$r_i(\tilde{p}_i, \mathbf{p}_{-i}) = \rho_i \tag{13}$$

with $r_i(\mathbf{p})$ given by (2). An explicit expression for $\tilde{p}_i$ is given below in equation (15). If no solution exists (i.e., the user cannot satisfy its throughput requirement), we set $\tilde{p}_i = 1$.

The best-response (BR) dynamics we consider can now be described as follows: Each user updates its transmission probability from time to time through its best response map, so as to satisfy (13). The update times of each user need not be coordinated with other users. Furthermore, several users can modify their probabilities simultaneously.

This mechanism reflects what myopic, self-interested user would reasonably do: repeatedly observe the current network conditions and react to bring its own throughput to the required level.

The analysis of the best-response mechanism will be carried out under the following assumptions, that apply throughout this section.

Assumption 1

(i) Fixed Demands: *The user population and the users' throughput requirements* $\rho_1, \ldots, \rho_n$ *are fixed. Furthermore,* $\rho_1, \ldots, \rho_n$ *are within the feasible throughput region.*

(ii) Persistent Updates: *Each user updates its transmission probabilities using Eq. (13) at arbitrarily chosen time instants, and the number of updates is unbounded.*

Furthermore, in order to guarantee convergence to the better equilibrium point $\mathbf{p}^a$, we must impose restrictions on the initial strategy profile of the users, denoted $\mathbf{p}^0$. Such restrictions are clearly essential due to the inherent non-uniqueness of the equilibrium in our model: indeed, if we start at the *worse* equilibrium point $\mathbf{p}^b$, we will stay there indefinitely under BR. Even more seriously, for initial conditions above $\mathbf{p}^b$ we might observe throughput collapse, as some users increase their transmission probabilities to their maximal values of 1 (see the simulation section).

To specify the required condition, let us define the following sets of strategy profiles:

$$\Pi_1 = \{\mathbf{p} \in [0,1]^n : r_i(\mathbf{p}) \geq \rho_i \text{ for every } i, \text{ and } p_j < p_j^b \text{ for some } j\}$$
$$\Pi_0 = \{\mathbf{p} \in [0,1]^n : \mathbf{p} \leq \bar{\mathbf{p}} \text{ for some } \bar{\mathbf{p}} \in \Pi_1\}.$$

Our main convergence result is summarized below.

Theorem 7 (BR Convergence) *Let the initial request probabilities* $\mathbf{p}^0 = (p_1^0 \ldots, p_n^0)$ *satisfy* $\mathbf{p}^0 \in \Pi_0$*. Then the best response dynamics asymptotically converges to the* better *equilibrium point* $\mathbf{p}^a$*.*

The proof proceeds by applying a "sandwich" argument. We first show that BR converges monotonously (from below) to $\mathbf{p}^a$ when started at $\mathbf{p}^0 = 0$. We then show that it converges monotonously (from above) to $\mathbf{p}^a$ when started with $\mathbf{p}^0$ in Π_1. Finally we conclude by monotonicity of the BR that convergence must occur all for initial conditions between 0 and Π_1, namely for all $\mathbf{p}^0 \in \Pi_0$. The details are provided in the appendix.

The requirement that $\mathbf{p}^0 \in \Pi_0$ may be interpreted as a "slow start" requirement: the initial probabilities need to be set low enough to ensure the required convergence. This requirement may be satisfied by each user individually simply by selecting $p_i^0 \leq \rho_i$. However, as shown next, it applies much more broadly, to any strategy profile that is below the *worse* equilibrium point.

Corollary 8 (Convergence Region) *Let the initial strategy profile $\mathbf{p}^0$ be component-wise dominated by $\mathbf{p}^b$, namely $p_i^0 < p_i^b$ for al i. Then $\mathbf{p}^0 \in \Pi_0$, and consequently BR converges to $\mathbf{p}^a$ when started from such $\mathbf{p}^0$.*

The proof can be found in the appendix.

As mentioned, a slow-start requirement is essential to guarantee convergence to the better equilibrium. This requirement is quite mild here, since users are only required to start below their worse equilibrium probability. Still, an important question is whether the users would obey this requirement. The incentive to do so is the possibility that the system may be led to the worse equilibrium otherwise, or even worse to throughput collapse. As these possibilities are uniformly worse for all users, it is in the interest of each user to cooperate in a protocol that leads to the better operating point.

We next consider some practical implementation considerations of the BR mechanism in a distributed environment, where users are generally unaware of the detailed decisions and strategies of the others. Note that r_i can be written as

$$r_i(\mathbf{p}) = \frac{p_i f_i(\mathbf{p}_{-i})}{T_1/T_2 + p_i f_i(\mathbf{p}_{-i}) + (1-p_i) g_i(\mathbf{p}_{-i})} \tag{14}$$

where

$$f_i(\mathbf{p}_{-i}) = \prod_{j \neq i} (1 - p_j)\,, \quad g_i(\mathbf{p}_{-i}) = \sum_{j \neq i} p_j \prod_{k \neq i,j} (1 - p_k)\,.$$

Both f_i and g_i now have a clear meaning and can be estimated by the user without detailed information on other users. Indeed, f_i is the fraction of times where user i's transmission is successful, and g_i the fraction of successful transmissions over the channel when user i does not transmit. The latter can be counted by listening to the number of CTS's (or alternatively inferred from centralized channel utilization information provided by the base station). So, if f_i and g_i are known (for the current channel conditions), the best-response as defined in (13) is easily computed as

$$\tilde{p}_i = \frac{T_1/T_2 + g_i}{f_i(1-\rho_i)/\rho_i + g_i}\,. \tag{15}$$

We note that this aspect of estimating a global (rather than private) quantity, namely g_i, does not appear in the basic model without reservations; see [12].

It is evident that in any practical implementation of this (or similar) scheme the required estimates would be inaccurate and noisy. This, as well as other deviations from our idealized model, were not addressed in our analysis and would require separate treatment. However, the large convergence region obtained here indicates to the robustness of the BR mechanism to such deviations.

4.2 Naive Best-Response

As discussed above, computing the best response of each user requires dedicated sensing of the shared medium, as well the use of equation (14). In some cases, the users may be limited in their sensing abilities, or alternatively may be more simple-minded in their computational approach. We next provide an alternative and easier to implement mechanism, which still converges to the better equilibrium under plausible conditions.

The update rule we consider here is simple: Update the transmission request probability in proportion to the required increase (or decrease) in throughput. In more detail, let p_i denote the current request probability for user i, and $r_i(\mathbf{p})$ its current throughput. The new request probability is set to

$$\tilde{p}_i = \frac{\rho_i}{r_i(\mathbf{p})}\, p_i \,. \tag{16}$$

Again, if no solution exists (i.e., the right hand side in (16) exceeds 1), the user sets $\tilde{p}_i = 1$. While simple to implement, we emphasize that this update rule does *not* lead to the required rate ρ_i in a single step (as does the best-response update), even if the other users freeze their strategies. The reason is that it does not take into account the effect of p_i on the *denominator* of (14). This state of affairs, where users neglect their own effect on some system parameters, may be associated with the economic concept of *price-taking* users, as opposed to *price-anticipating* ones (see, e.g., [7] and references therein).

A slightly different form of the update rule in (16) will be useful. From equations (1) and (2), the rate of each user for given $\mathbf{p}$ can be written as

$$r_i(\mathbf{p}) = p_i R_i(\mathbf{p})\,, \tag{17}$$

where

$$R_i(\mathbf{p}) = \frac{T_2 \prod_{j\neq i}(1-p_j)}{T_1 + T_2 \sum_j q_j}\,. \tag{18}$$

R_i can be interpreted as the average rate obtained per channel request of user i (successful or not). Clearly, this quantity can be locally monitored by the user. The modified transmission request probability may now be computed as

$$\tilde{p}_i = \frac{\rho_i}{R_i(\mathbf{p})}\,. \tag{19}$$

This rule reflects the (false) assumption that $R_i(\mathbf{p})$ is not affected by p_i. It is easily seen to be equivalent to (16).

We shall refer to the asynchronous scheme defined by the update rule (16) or (19) as Naive Best-Response (NBR) mechanism. Its convergence will be established under more strict assumptions than before.

Theorem 9 (NBR Convergence) *Suppose the following conditions hold:*

(i) (Balanced Load): The better equilibrium point $\mathbf{p} = \mathbf{p}^a$ *satisfies*

$$\sum_{j\neq i} \frac{p_j}{1-p_j} < 1 \quad \text{for every } i \in \mathcal{I}\,. \tag{20}$$

(ii) (Slow Start): The initial request probabilities $\mathbf{p}^0$ *are set low enough so that for each user* i *we have (a)* $p_i^0 \leq p_i^a$*, and (b)* $r_i(\mathbf{p}^0) \leq \rho_i$*.*

Then NBR converges to better equilibrium point $\mathbf{p}^a$*.*

Proof: See the appendix.

This NBR convergence result depends on two conditions. These are merely sufficient conditions, and convergence may in fact hold more broadly (as illustrated in the experiments section). The 'slow start' requirement in (*ii*) is evidently stronger than in the BR case. However, it can still be satisfied by each user individually by setting p_i^0 to ρ_i. The other requirement imposed in (20) may in general be interpreted as a 'light traffic' condition, as discussed below. Interestingly, in the symmetric-user case this condition is automatically satisfied, which implies the following convergence result.

Corollary 10 (Symmetric Users) *Assume that all user demands are identical:* $\rho_i = \rho_1$ *for all* i*. Then the better equilibrium point* $\mathbf{p}^a$ *satisfies condition (20). Consequently, NBR converges* $\mathbf{p}^a$ *subject only to the slow start condition (ii).*

Proof: It may be readily verified that symmetric users imply symmetric strategies, namely $p_i = p$ at equilibrium. Therefore, condition (20) becomes $\frac{(n-1)p}{1-p} < 1$ or $p < \frac{1}{n}$. But since $\sum p_i < 1$ at the better equilibrium (by Theorem 4(iii)), then $p < \frac{1}{n}$ as required. $\square$

For the asymmetric case, we note that condition (20) on $\mathbf{p}^a$ is implied by the simpler requirement

$$\sum_{j\in\mathcal{I}} \frac{p_j^a}{1-p_j^a} < 1. \tag{21}$$

This may indeed be interpreted as a light-traffic condition. Moreover, we may expect this condition to hold when the system supports many small users, so that the request probability p_i of each is small. Indeed, recall from Theorem 4(iii) that the better-equilibrium probabilities must satisfy $\sum_i p_i^a < 1$. When each p_i^a is small, the latter sum closely approximates the sum in (21). This is commensurate with the observation that the *price-taking* approach, which is behind the NBR scheme, is natural when the effect of each user on the system is small.

4.3 Experiments

To conclude this section, we briefly examine by simulation the convergence properties of the BR and NBR mechanisms proposed above.

In our first set of experiments we illustrate the temporal evolution of the BR and NBR mechanism for a particular 3-user system with $T_2/T_1 = 6$ and throughput demand vector $\underline{\rho} = (0.6, 0.1, 0.05)$. The equilibrium probabilities for this system turn out to be $\mathbf{p}^a = (0.51, 0.147, 0.0797)$ and $\mathbf{p}^b = (0.75, 0.333, 0.2)$. The subset of users who update their request probabilities is chosen at random in every time-slot. Example runs for different values of the initial conditions vector $\mathbf{p}^0$ are shown in Figures 2 and 3, in which the evolution of p_3 is depicted. Figure 2 shows monotone convergence of the request probabilities in both algorithms: increasing for small $\mathbf{p}^0$ (left), and decreasing when $\mathbf{p}^0$ is chosen so that $r_i(\mathbf{p}^0) \geq \rho_i$ for every i (right). We observe that BR is slightly faster to converge than NBR, but the differences are mild. Further note that the (sufficient) convergence condition (20) for NBR is not satisfied here, since $\frac{p_1^a}{1-p_1^a} = \frac{0.51}{0.49} > 1$. However, NBR still converges for a wide set of initial conditions. Figure 3 provides an example for non-monotone convergence (left), and an example where the BR mechanism converges to the better equilibrium, while the NBR mechanism diverges (i.e., all probabilities are set to one, since the demands ρ_i cannot be satisfied at some point).

The second set of experiments examines in more detail the set of initial conditions for which convergence (to the better equilibrium) is obtained. The results are shown for a two-user system, with $T_2/T_1 = 7$, and $\underline{\rho} = (0.5, 0.25)$. The two equilibria are $\mathbf{p}^a = (0.37, 0.23)$

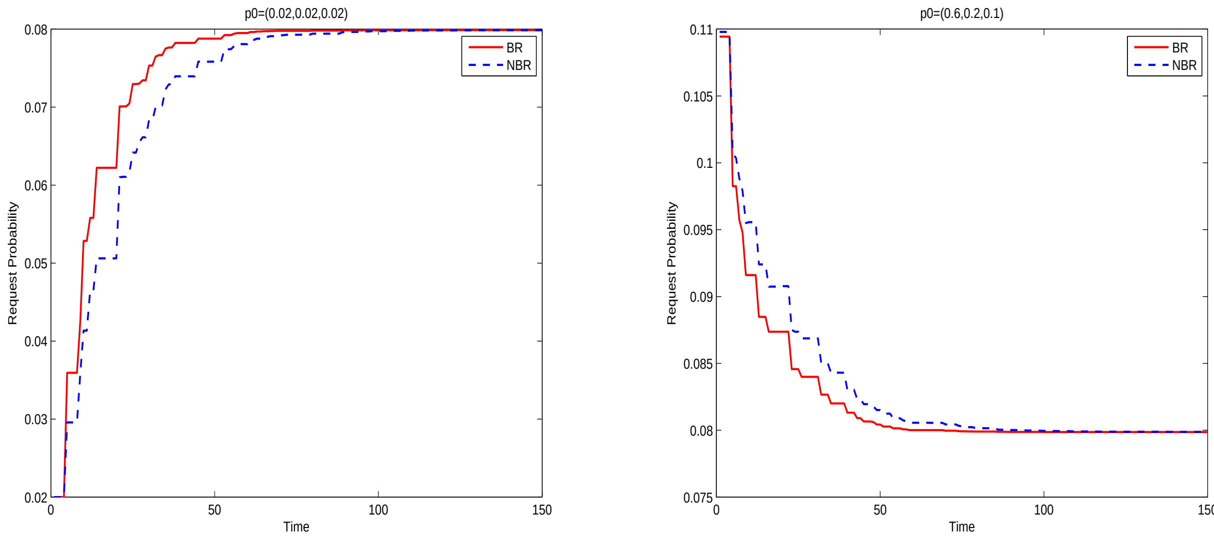

Figure 2: Experiment set 1. Convergence of p_3 under BR and NBR for different initial conditions, for a three-user system with $\underline{\rho} = (0.6, 0.1, 0.05)$. Monotone increasing convergence is obtained for $\mathbf{p}^0 = (0.02, 0.02, 0.02)$ (left) and decreasing for $(0.6, 0.2, 0.1)$ (right) in both dynamic schemes.

and $\mathbf{p}^b = (0.77, 0.63)$. Both mechanisms were run with initial conditions $\mathbf{p}^0$ that vary over a dense grid in the square $[0, 1]^2$. Figure 4 depicts the results. It may be seen that BR (left panel) converges for all initial conditions with $\mathbf{p}^0 < \mathbf{p}^b$, in accordance with Corollary 8. Outside this area, divergence occurs whenever $\mathbf{p}^0 > \mathbf{p}^b$, while in the remaining two regions convergence essentially depends on the order of user updates (recall that these are selected at random at each stage). Convergence to the *worse* equilibrium point has not been detected for any initial conditions (apart from $\mathbf{p}^0 = \mathbf{p}^b$). As for NBR, it may be viewed that convergence is not assured for all $\mathbf{p}^0 < \mathbf{p}^b$. However, the convergence region is still considerably larger than the theoretical guarantees in Theorem 9, and includes a wide margin beyond the better equilibrium point. Interestingly, NBR converges for some initial point in the region $\mathbf{p}^0 > \mathbf{p}^b$, while BR always diverges there. We observed a similar behavior for other network configurations that were tested.

Based on our experiments, we may conclude the following: The BR mechanism outperforms the NBR mechanism in terms of both speed and robustness (i.e., resilience to initial conditions). However, the observed differences were not severe. While we have only

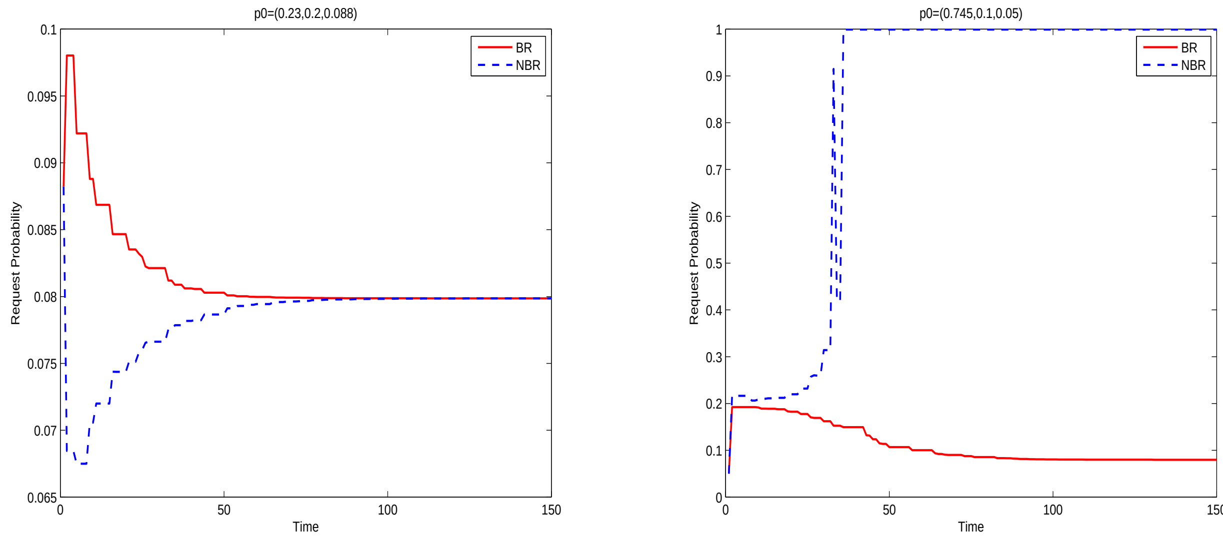

Figure 3: Experiment set 1. Left: Non-monotone convergence of p_3 for $\mathbf{p}^0 = (0.23, 0.2, 0.088)$. Right: BR converges while NBR diverges for $\mathbf{p}^0 = (0.745, 0.1, 0.05)$.

examined here systems with two and three users, we expect that in the may-user case the differences between the two schemes will only become smaller, provided that the relative size of an individual user decreases and hence its effect on the overall system diminish as well.

5 Effect of the Data Transmission Period

In this section we address the effect of the data transmission period T_2 on system performance. The relevant performance measures that we consider are capacity, power, and inter-packet delay. We note that the actual parameter that affects performance is the ratio T_2/T_1. However, we assume throughout that T_1 is fixed, as it is determined by the timing requirement of the RTS/CTS control sequence.

5.1 Capacity

Denote by $\underline{\rho} = (\rho_1, \ldots, \rho_n)$ the vector of throughput demands, and let Ω be the set of *feasible* throughput vectors $\underline{\rho}$ for which there exists at least one equilibrium point. Clearly

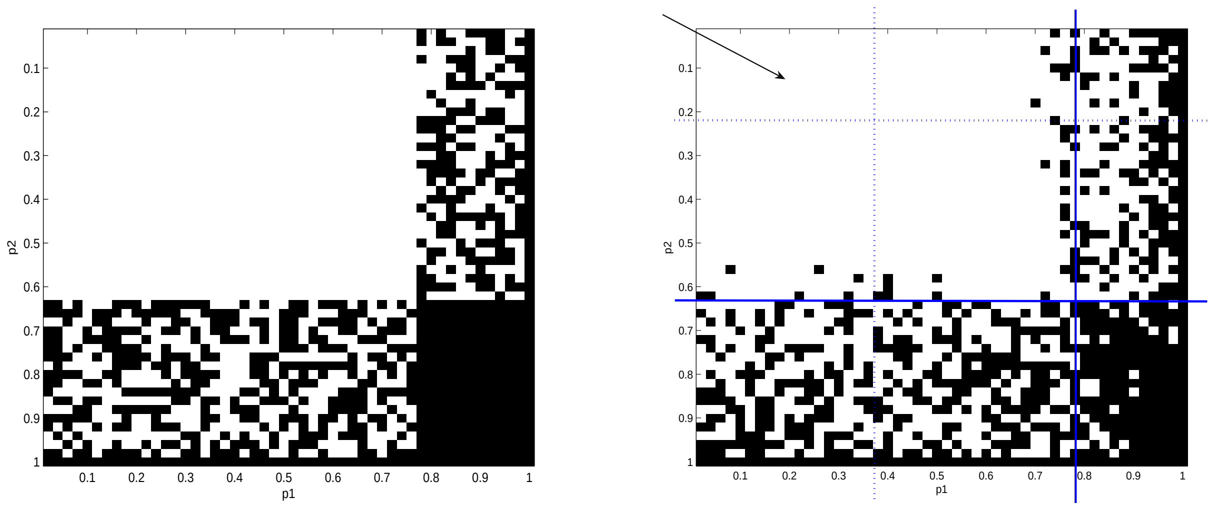

Figure 4: Convergence for different initial conditions for BR (left) and NBR (right). A white square stands for convergence to an equilibrium, while a black square represents no convergence. Marked on right panel are the coordinates of the worse equilibrium (solid line) and the better equilibrium (dotted line).

Ω depends on the data transmission period T_2. Figure 5 illustrates the set of feasible throughput demands for a simple two-user case. As mentioned before, $\Omega(T_2)$ is a closed and non-empty set.

We informally use the term "capacity" here in reference to the extent to which different throughput vectors can be accommodated by the network. The next proposition shows how the capacity of the channel increases in T_2.

Proposition 11 *Consider two data transmission periods T_2 and $\tilde{T}_2$, and let $\beta = \tilde{T}_2/T_2$ denote their ratio. Then the throughput vector $\underline{\rho} = (\rho_1, \ldots, \rho_n)$ is feasible in the T_2-system* if and only if *$\underline{\tilde{\rho}} = \gamma\underline{\rho}$ is feasible in the $\tilde{T}_2$-system, where*

$$\gamma = \frac{\beta}{1 + (\beta - 1)\rho} \tag{22}$$

and $\rho = \sum_i \rho_i$.

Proof: Let $\mathbf{p}^*$ denote an equilibrium strategy profile in the T_2-system. Then $\mathbf{p}^*$ satisfies the modified equilibrium equations (7). Fixing $\mathbf{p}^*$, it is easily verified that the same equation

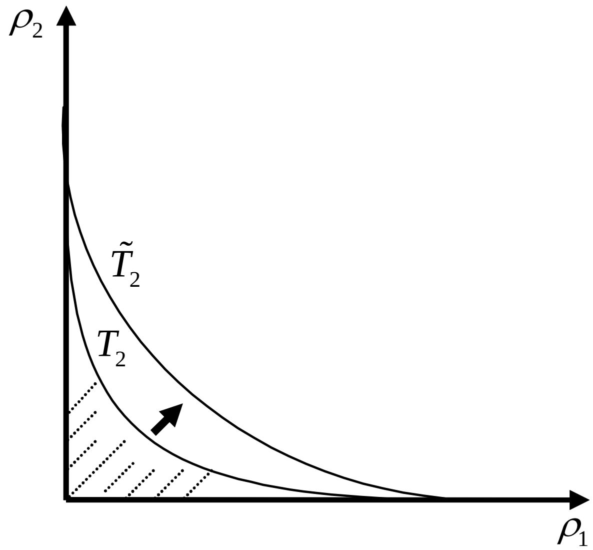

Figure 5: The set of feasible throughput demands for a two user network with two different data transmission periods $\tilde{T}_2 > T_2$. The upper-boundary increases due to the increase of the data transmission period.

is satisfied with $\tilde{T}_2$ and $\underline{\tilde{\rho}} = \gamma\underline{\rho}$ on the right hand side, hence $\mathbf{p}^*$ supports $\gamma\underline{\rho}$ in that case. The opposite implication follows similarly. □

The factor γ in equation (22) quantifies the gain in capacity that is obtained by increasing T_2 by a factor of β. Evidently, γ is increasing in β, and $\gamma > 1$ when $\beta > 1$. The gain is more significant at light loads ($\rho << 1$), and approaches its saturation value of $1/\rho$ when β becomes large. see Figure 6 for an graphical illustration of this relation.

5.2 Power and Delay

We next consider two important measures of performance from the user point of view: power and delay. We will provide some qualitative results for the former, and briefly comment on the latter.

From equation (12) we know that the average power S_i spent by a user increases linearly with its request probability p_i: this is just the overhead of RTS packages transmissions (and

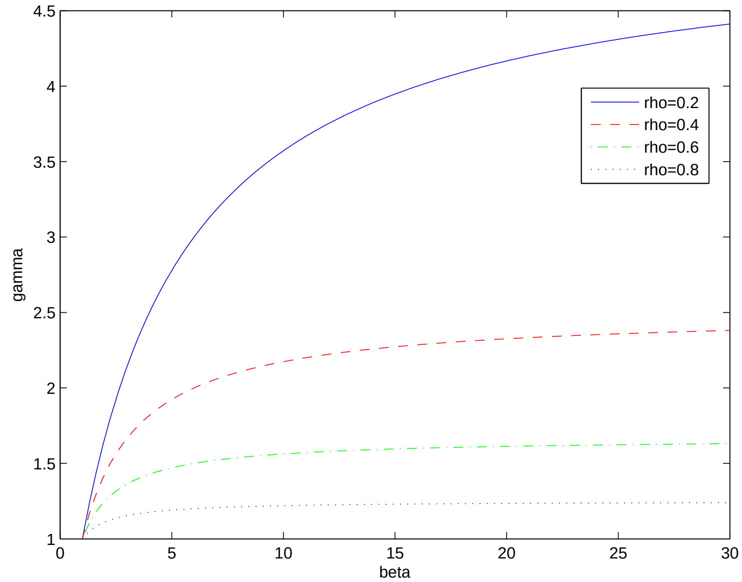

Figure 6: The capacity gain (γ) obtained by increasing T_2 by a factor of β. Graphs of γ vs. β are plotted according to equation (22) for several network loads.

their collisions). To clarify the dependence S_i on T_2, we still need to find the dependence of the *equilibrium* probabilities p_i on T_2. This forms the basis for the following qualitative result.

Proposition 12 *The average power investment of each user at the better equilibrium is decreasing in* T_2*.*

Proof: See the appendix.

This result appears quite intuitive. We emphasize however two aspects: first, power saving is common to all users; and second, this result depends on the system being in the *better* equilibrium point, and need not hold at the other equilibrium.

In Figure 7 we show the effect of the T_2 on the total power investment $\sum_i S_i(\mathbf{p}^a)$ at the better equilibrium $\mathbf{p}^a$. For simplicity, we focus on the symmetric user case, where $\rho_i = \frac{\rho}{n}$. We keep ρ fixed, and examine the performance for different number of users n. As expected, for any given T_2, the total power investment increases with n, as a large number of users waste more power on collisions. However, the gap between the curves for different n's becomes smaller as T_2 grows, since a large data transmission period diminishes the collision affect on the overall power cost.

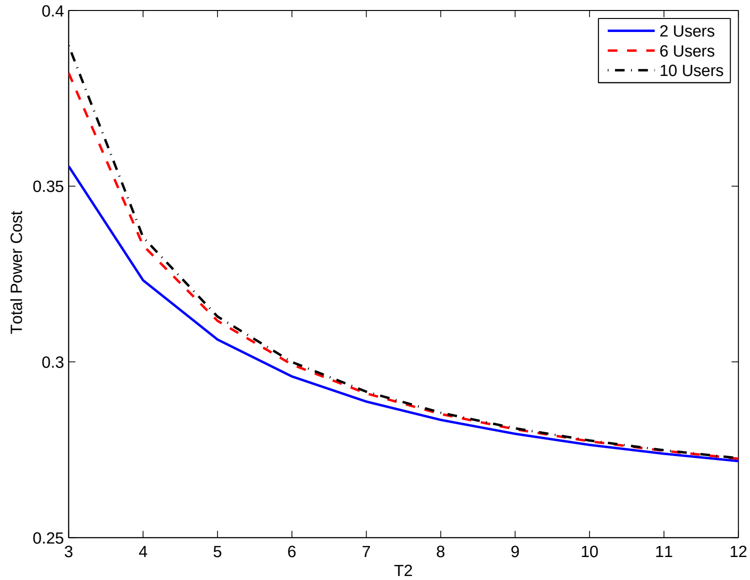

Figure 7: The overall power cost $\sum_i S_i(\mathbf{p}^a)$ as a function of T_2 for a fixed total demand of $\rho = 0.5$. Users are symmetric, hence $\rho_i = \frac{\rho}{n}$. We used $T_1 = 1$, $\bar{T}_1 = 0.5$ for this simulation.

.

Turning to consider a delay measure, we consider here the average time separation between two consecutive data transmission periods (namely, successive T_2 periods of the given users), which we denote by D_i. This delay may be especially relevant for real-time traffic, such as voice and video streaming. We shall not consider here other delay components, such as queueing delay, which requires a packet arrival process which is not part of our model.

It is easily verified that the delay of each user is directly proportional to T_2. More precisely,

Proposition 13 *The average delay D_i in equilibrium is given by*

$$D_i(T_2) = T_2/\rho_i. \tag{23}$$

Proof: Denote by f_i the frequency of data transmission intervals for user i. Then $f_i T_2 = \rho_i$. Noting that $D_i = 1/f_i$ leads to the stated relation. $\square$

To summarize, we have seen that increasing T_2 benefits the system capacity (i.e., sustainable channel throughput), as well as the power investment of each user. As may be expected, these benefits of increasing T_2 become marginal when T_2/T_1 is large. On the other hand, increasing T_2 leads to a proportional increase in the inter-packet delay. We shall not attempt here to quantify the tradeoff between these quantities, as T_2 may be limited by other considerations related to network management and control that fall outside our present framework. The results above provide however a basis for a more detailed analysis of this tradeoff in a non-cooperative setting.

6 Conclusion

Reservation mechanisms in random access channels are meant to reduce the collision overhead, and thus increase capacity and energy efficiency and reduce delay. However, it is not clear a-priory whether self-interested users would exploit or misuse this additional feature. The results in this paper indicate that substantial benefits are indeed associated with such mechanisms. However, care must be taken that the system remains in a favorable equilibrium point, out of several that may exist.

Our results show that two-equilibrium property observed in [12] carries over to the present model which incorporates a CTS/RTS-like reservation mechanism. We have further established here strong convergence results for the Best Response dynamics, showing that it converges to the better equilibrium whenever the initial users strategies are below the *worse* equilibrium. Clearly, however, to avoid divergence (or throughput collapse) due to stochastic effects one needs to supplement these simple schemes by collision avoidance and resolution mechanisms, such as 802.11's back-off mechanism [5]. The design and analysis analysis of such mechanisms within a non-cooperative framework in general, and within the framework of this paper in particular, poses a challenging and important direction for future research.

Acknowledgements

The research of Ishai Menache was supported by a Marie Curie Fellowship within the 7th European Community Framework Programme. This research was also supported by grant No. 2006379 from the United States–Israel Binational Science Foundation (BSF).

References

[1] IEEE 802.11 standards. Available from `http://standards.ieee.org/getieee802/802.11.html`.

[2] E. Altman, R. El-Azouzi, and T. Jimenez. Slotted aloha as a game with partial information. *Computer Networks*, 45(6):701–713, 2004.

[3] A. Fiat, Y. Mansour, and U. Nadav. Efficient contention resolution for selfish agents. In *Symposium on Discrete Computing (SODA)*, pages 179 – 188, 2007.

[4] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, Cambridge, MA, USA, 1991.

[5] Z. Hua, L. Ming, I. Chlamtac, and B. Prabhakaran. A survey of quality of service in IEEE 802.11 networks. *IEEE Wireless Communications*, 11:6 – 14, 2004.

[6] H. Inaltekin and S. Wicker. A one-shot random access game for wireless networks. In *International Conference on Wireless Networks, Communications and Mobile Computing*, pages 940– 945, 2005.

[7] R. Johari, S. Mannor, and J. N. Tsitsiklis. Efficiency loss in a network resource allocation game: The case of elastic supply. *IEEE Transactions on Automatic Control,*, 50(11):1712–1724, 2004.

[8] N. B. Mandayam, S. B. Wicker, J. Walrand, T. Basar, J. Huang, and D. P. Palomar (editors). Special issue on game theory in communication systems. *Journal of Selected Areas in Communications*, 26(7), 2008.

[9] I. Menache and N. Shimkin. Decentralized rate regulation in random access channels. In *Proceedings of INFOCOM*, 2008.

[10] I. Menache and N. Shimkin. Efficient rate-constrained Nash equilibrium in collision channels with state information. In *Proceedings of INFOCOM*, 2008.

[11] I. Menache and N. Shimkin. Noncooperative power control and transmission scheduling in wireless collision channels. In *Proceedings of ACM Sigmetrics*, 2008.

[12] I. Menache and N. Shimkin. Rate-based equilibria in collision channels with fading. *Journal of Selected Areas in Communications*, 26(7):1070–1077, 2008.

[13] 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.

[14] A. Muqattash and M. Krunz. CDMA-based MAC protocol for wireless ad hoc networks. In *MobiHoc '03: Proceedings of the 4th ACM international symposium on Mobile ad hoc networking & computing*, pages 153–164, 2003.

[15] S. M. Ross. *Introduction to Probability Models*. Academic Press, 2006.

A Proofs

This section establishes the convergence results of Section 4, as well as Proposition 12. We start with some lemmas that lead up to the proof of Theorem 7. We assume throughout that the strategy profile $\mathbf{p} = (p_1, \ldots, p_n)$ has all elements in $[0, 1]$.

Lemma 14 *The effective throughput $r_i(\mathbf{p})$ of user i is (i) strictly increasing in p_i, and (ii) strictly decreasing in p_j for each $j \neq i$.*

Proof: Recalling equations (1) and (2), it follows after some algebra that

$$r_i(\mathbf{p}) = \left(\frac{T_1}{q_i T_2} + \sum_{j \neq i} \frac{h(p_j)}{h(p_i)} + 1 \right)^{-1}, \tag{24}$$

where $h(p) = \frac{p}{1-p}$. Since $h(p)$ is strictly increasing in p, the second summand in (24) obviously strictly decreases in p_i and strictly increases in p_j. Recalling again that $q_i = p_i \prod_{j \neq i}(1 - p_j)$, the first summand in (24) has the same monotonicity properties. Taking the inverse establishes the required properties. $\square$

Let $\mathrm{BR}_i(\mathbf{p})$ denote the best-response set of user i to a strategy profile $\mathbf{p}$, that is, $\mathrm{BR}_i(\mathbf{p})$ is the set of probabilities $\bar{p}_i \in [0, 1]$ that satisfy the throughput requirement (13) for that user. When $\mathrm{BR}_i(\mathbf{p})$ is a singleton we shall use the same notation to denote its value.

Lemma 15 *For each user i,*
(i) The best-response set $\mathrm{BR}_i(\mathbf{p})$ is either a singleton or else empty.
(ii) The best-response map is monotone increasing in $\mathbf{p}$. That is, if $\mathbf{p}^1 \leq \mathbf{p}^2$ (component-wise) and $\mathrm{BR}_i(\mathbf{p}^2) \neq \emptyset$, then $\mathrm{BR}_i(\mathbf{p}^1) \neq \emptyset$, and $\mathrm{BR}_i(\mathbf{p}^1) \leq \mathrm{BR}_i(\mathbf{p}^2)$.

Proof: Both claims follow directly from the definition of the best-response and the monotonicity properties of $r_i(\mathbf{p})$ that were established in Lemma 14. $\square$

We henceforth consider the BR scheme under the assumptions of Theorem 7. The following notation will be used. Starting with some initial strategy vector $\mathbf{p}^0$, let t_k denote the time of the k-th BR update, where at each update some subset of users simultaneously modify their transmission probability according to the BR map. Thus, $p_i^{k+1} = \mathrm{BR}_i(\mathbf{p}^k)$ if user i updates at time t^k, and $p_i^{k+1} = p_i^k$ otherwise. In case $\mathrm{BR}_i(\mathbf{p}^k)$ is empty we set it to 1 (always attempt transmission), and the process effectively terminates.

Lemma 16 *Assume that* $\mathbf{p}^0 = 0$. *Then the strategy profile sequence* $\mathbf{p}^k$ *converges monotonously to* $\mathbf{p}^a$.

Proof: We first show by induction that $\mathbf{p}^k \leq \mathbf{p}^a$ and $\mathbf{p}^k \leq \mathbf{p}^{k+1}$. Indeed, this trivially holds for $\mathbf{p}^0 = 0$. Assume this holds true for some $k \geq 0$. Note that $\mathrm{BR}(\mathbf{p}^a) = \mathbf{p}^a$ by definition of the equilibrium. Hence, by monotonicity of the best response (Lemma 15), it follows that $\mathbf{p}^{k+1} = \mathrm{BR}(\mathbf{p}^k) \leq \mathrm{BR}(\mathbf{p}^a) = \mathbf{p}^a$. Similarly, monotonicity of BR and $\mathbf{p}^k \leq \mathbf{p}^{k+1}$ imply that $\mathbf{p}^{k+1} \leq \mathbf{p}^{k+2}$.

It follows that $\{\mathbf{p}^k\}$ is a monotone increasing sequence, and bounded above by $\mathbf{p}^a$. Thus, $\mathbf{p}^k$ converges to a finite limit $\mathbf{p}^\infty \leq \mathbf{p}^a$. Since Assumption 1(ii) (persistent updates) is in effect, a standard continuity argument may be invoked to show that this limit is a fixed point of the best response map, namely an equilibrium point. But since $\mathbf{p}^a$ is the better (i.e., smallest) equilibrium point, it follows that $\mathbf{p}^\infty = \mathbf{p}^a$. $\square$

We note that the above lemma (and proof) applies to any initial profile such that $\mathbf{p}^0 \leq \mathbf{p}^a$ and $r_i(\mathbf{p}^0) \leq \rho_i$. However, this is not needed for the present argument.

Lemma 17 *Assume that* $\mathbf{p}^0 \in \Pi_1$. *Then the strategy profile sequence* $\mathbf{p}^k$ *converges monotonously (from above) to* $\mathbf{p}^a$.

Proof: We will show by induction that $r_i(\mathbf{p}^k) \geq \rho_i$ and $\mathbf{p}^{k+1} \leq \mathbf{p}^k$ hold for all i and k. For $k = 0$, $r_i(\mathbf{p}^0) \geq \rho_i$ holds by definition of π_1. To show that $\mathbf{p}^1 \leq \mathbf{p}^0$, consider any user i that updates its strategy at $k = 0$. As such a user reduces its throughput from $r_i(\mathbf{p}^0) \geq \rho_i$ to ρ_i, it follows by Lemma 14(i) (monotonicity of $r_i(\mathbf{p})$ in p_i) that $p_i^1 \leq p_i^0$ for such a user. Since the other users do not modify their strategies, we have $\mathbf{p}^1 \leq \mathbf{p}^0$.

Assume next that the above holds for some $k \geq 0$. Then $\mathbf{p}^{k+1} \leq \mathbf{p}^k$ follows exactly as argued for $k = 0$ above. It remains to show that $r_i(\mathbf{p}^{k+1}) \geq \rho_i$ for each i. Consider a user i that did not modify its strategy as time k, namely $p_i^{k+1} = p_i^k$. Since $r_i(\mathbf{p}^k) \geq \rho_i$ and $\mathbf{p}^{k+1} \leq \mathbf{p}^k$, it follows from Lemma 14(ii) (monotonicity of $r_i(\mathbf{p})$ in p_j) that $r_i(\mathbf{p}^{k+1}) \geq r_i(\mathbf{p}^k) \geq \rho_i$. Consider next a user i that modified its strategy to p^{k+1} according to BR. After that change, we have $r_i(\tilde{\mathbf{p}}^k) = \rho_i$ where $\tilde{\mathbf{p}}^k$ denotes $\mathbf{p}^k$ with p_i^k replaced by p_i^{k+1}. Next, allowing other users to modify their probabilities as well, we obtain as before that $r_i(\mathbf{p}^{k+1}) \geq r_i(\tilde{\mathbf{p}}^k) = \rho_i$. □

Proof of Theorem 7: Let $\mathbf{p}^0 \in \Pi_0$. Then there exists $\hat{\mathbf{p}} \in \Pi_1$ so that $\mathbf{p}^0 \leq \hat{\mathbf{p}}$. Let $\{\mathbf{p}^k\}$ be a BR sequence started from $\mathbf{p}^0$. We now consider two additional BR sequences, one started at $\check{\mathbf{p}}^0 = 0$, and denoted $\{\check{\mathbf{p}}^k\}$, and the other started at $\hat{\mathbf{p}}^0$ and denoted $\{\hat{\mathbf{p}}^k\}$. We further let the set of users who modify their strategy at each given step be the same in all three sequences. Noting that $\check{\mathbf{p}}^0 = 0 \leq \mathbf{p}^0 \leq \hat{\mathbf{p}}^0$, it follows by Lemma 15 (monotonicity of the best response) that $\check{\mathbf{p}}^k \leq \mathbf{p}^k \leq \hat{\mathbf{p}}^k$ for all k. However, by Lemmas 16 and 17 the two outer sequences converge to $\mathbf{p}^a$. Hence so does the middle sequence $\{\mathbf{p}^k\}$. □

Proof of Corollary 8: Let $\mathbf{p}^0$ be any strategy profile such that $\mathbf{p}^0 < \mathbf{p}^b$. We will show that $\mathbf{p}^0 \in \Pi_0$, namely that $\mathbf{p}^0 \leq \mathbf{p}$ for some $\mathbf{p} \in \Pi_1$. To that end, it suffices to show that there exist vectors $\mathbf{p} \in \Pi_1$ arbitrarily close to $\mathbf{p}^b$. In fact, we will show that there exists $\mathbf{p}$ arbitrarily close to $\mathbf{p}^b$ with $\mathbf{p} < \mathbf{p}^b$ and $r_i(\mathbf{p}) > \rho_i$ for each i (which together imply that $\mathbf{p} \in \Pi_1$ as required). For that purpose, we use a similar construction as the one used in Theorem 3 of [12] to establish the two-equilibria property.

Recall that the two equilibria $\mathbf{p}^a$ and $\mathbf{p}^b$ are obtained as the solutions of the equilibrium equations (6), or equivalently as solutions of the modified equilibrium equations (7), namely $q_i(\mathbf{p}) = \tilde{\rho}_i$ for all i. Let us write these equations for the scaled throughput demands $\gamma\tilde{\rho}_i$, namely $q_i(\mathbf{p}) = \gamma\tilde{\rho}_i$ for all i, where $\gamma > 0$ is a scaling parameter. Now, for each $p_1 \in [0, 1)$, these equations can be uniquely solved for $(p_2, \ldots, p_n)$ and γ, with p_j, $j \geq 2$ increasing in p_1. This becomes apparent after noting that division of the j-th equation by the first gives

$$\tilde{\rho}_1 \frac{p_j}{1 - p_j} = \tilde{\rho}_j \frac{p_1}{1 - p_1},$$

with $\frac{x}{1-x}$ a strictly increasing function of x in $[0, 1)$. Furthermore, it has been shown in the proof of Theorem 3 of [12] that γ is a continuous and strictly unimodular function of p_1, with $\gamma = 1$ obtained at the two equilibria values $p_1^a < p_2^b$, while $\gamma > 1$ holds in between. Thus, as p_1 approaches p_1^b from below, denoted $p_1 \nearrow p_1^b$, we obtain that $\mathbf{p} \nearrow \mathbf{p}^b$, $\gamma \searrow 1$, and $q_i(\mathbf{p}) \searrow \tilde{\rho}_i$. Finally, substituting $q_i(\mathbf{p}) = \gamma\tilde{\rho}_i$ in (6) (2) we obtain that

$$r_i(\mathbf{p}) = \frac{\gamma\tilde{\rho}_i T_2}{T_1 + \gamma\tilde{\rho} T_2} = \frac{\gamma\rho_i}{(1 - \rho) + \gamma\rho}$$

(cf. (8) for the last equality), so that as $\gamma \searrow 1$ we have $r_i(\mathbf{p}) \searrow \rho_i$. To summarize, for p_1 approaching p_1^b from below we obtain a strategy profile $\mathbf{p}$ converging to $\mathbf{p}^b$ with $r_i(\mathbf{p}) > \rho_i$ for each i, as desired. □

We turn next to the proof of Theorem 9. Recall the definition of the naive best-response in equation (17), namely $\bar{p}_i = \rho_i / R_i(\mathbf{p})$. The following properties of R_i will be instrumental in our proof.

Lemma 18 *The function $R_i(\mathbf{p})$ is:*
(i) Strictly decreasing in p_j for each $j \neq i$.
(ii) Decreasing in p_i if (and only if) $\sum_{j\neq i} \frac{p_j}{1-p_j} \leq 1$.

Proof: Part (i) follows immediately from Lemma 14, after observing that $R_i(\mathbf{p}) = r_i(\mathbf{p})/p_i$ (Eq. 17). As for part (ii), using (17) and (1) we can write $R_i(\mathbf{p})$ as

$$R_i(\mathbf{p})^{-1} = \frac{T_1}{T_2 \prod_{j\neq i}(1-p_j)} + p_i + (1-p_i) \sum_{j\neq i} \frac{p_j}{1-p_j} \,.$$

The derivative in p_i obviously equals $1 - \sum_{j\neq i} \frac{p_j}{1-p_j}$, which implies the required monotonicity. □

We note that condition (20) is equivalent to the required inequality in Lemma 18(ii). The next lemma establishes the key properties of the naive best response updates.

Lemma 19 *Suppose $\mathbf{p}^a$ satisfies the condition (20). Let $\mathbf{p}$ be a strategy profile that satisfies the slow-start requirements for NBR, namely: (a) $\mathbf{p} \leq \mathbf{p}^a$, (b) $r_i(\mathbf{p}) \leq \rho_i$ for each i.*
Let $\bar{\mathbf{p}}$ be obtained from $\mathbf{p}$ by letting some user i employ a naive best-response, namely $\bar{p}_i = \rho_i / R_i(\mathbf{p})$, and $\bar{p}_j = p_j$ for $j \neq i$. Then $\bar{p}_i \geq p_i$, and $\bar{\mathbf{p}}$ satisfies requirements (a) and (b) above.

Proof: Note first that $\mathbf{p} \leq \mathbf{p}^a$ implies that $\mathbf{p}$ satisfies (20) as well, since the function $\frac{x}{1-x}$ is strictly increasing in $x \in [0, 1]$. Recall that $r_i(\mathbf{p}) = p_i R_i(\mathbf{p})$. Since $r_i(\mathbf{p}) \leq \rho_i$, it follows that $\bar{p}_i = \rho_i / R_i(\mathbf{p}) \geq p_i$. Next, to establish property (a), we need to show that $\bar{p}_i \leq p_i^a$. This is done by comparing $\bar{p}_i$ with the (exact) best response $\tilde{p}_i = BR_i(\mathbf{p})$ of user i to the strategy profile $\mathbf{p}$. Recall that $\tilde{p}_i$ satisfies (13), namely $r_i(\tilde{p}_i, \mathbf{p}_{-i}) = \rho_i$, or equivalently $\tilde{p}_i R_i(\tilde{p}_i, \mathbf{p}_{-i}) = \rho_i$. Now, it was argued in the proof of Theorem 7 that $p_i \leq \tilde{p}_i \leq p_i^a$. Therefore, by Lemma 18(ii), it follows that $R_i(\tilde{p}_i, \mathbf{p}_{-i}) \leq R_i(p_i, \mathbf{p}_{-i}) \equiv R_i(\mathbf{p})$. This, in turn, implies that

$$\bar{p}_i = \frac{\rho_i}{R_i(\mathbf{p})} \leq \frac{\rho_i}{R_i(\tilde{p}_i, \mathbf{p}_{-i})} = \tilde{p}_i \leq p_i^a \,.$$

We finally need to establish property (b). Noting that $\bar{p}_i \geq p_i$ (as shown above), it follows by Lemma 14(ii) and our choice of $\mathbf{p}$ that $r_j(\bar{\mathbf{p}}) \leq r_j(\mathbf{p}) \leq \rho_j$ for each $j \neq i$. Hence, it remains only to show that $r_i(\bar{\mathbf{p}}) \leq \rho_i$. But this readily follows from $\bar{p}_i \leq \tilde{p}_i$, which implies that $r_i(\bar{\mathbf{p}}) \equiv r_i(\bar{p}_i, \mathbf{p}_{-i}) \leq r_i(\tilde{p}_i, \mathbf{p}_{-i}) = \rho_i$. $\square$

Proof of Theorem 9: Let $\{\mathbf{p}^k\}$ be the sequence of strategy profiles obtained by the naive best-response scheme. By our assumptions $\mathbf{p}^0$ satisfied the requirements of Lemma 19. It may be seen that the conclusions of this lemma hold true even if several users update their strategy simultaneously. Therefore $\mathbf{p}^1 \geq \mathbf{p}^0$, and $\mathbf{p}^1$ satisfies the requirements of Lemma 19 as well. Proceeding by induction, and arguing as in the proof of Lemma 16, we may now establish that $\mathbf{p}^k$ converges to $\mathbf{p}^a$, as claimed. $\square$

Proof of Proposition 12: Fix $\rho = (\rho_1, \ldots \rho_n)$ and consider two different data transmission periods $\hat{T}_2$ and $\bar{T}_2$ such that $\hat{T}_2 > \bar{T}_2$. Then the equilibrium equations (7) for $\hat{T}_2$ and $\bar{T}_2$ are $q_i = \hat{\rho}_i$ and $q_i = \bar{\rho}_i$, where $\hat{\rho}_i < \bar{\rho}_i$ for every user i. We shall refer to the vectors $\hat{\rho} = (\hat{\rho}_1, \ldots, \hat{\rho}_n)$ and $\bar{\rho} = (\bar{\rho}_1, \ldots, \bar{\rho}_n)$ as modified demands (to distinguish these quantities from the actual demands ρ). Proposition 12 now follows immediately from the next lemma by noting (12).

Lemma 20 *Let $\hat{\rho}$ and $\bar{\rho}$ be two modified demand vectors such that $\hat{\rho} < \bar{\rho}$ (component-wise), and let $\hat{\mathbf{p}}$ and $\bar{\mathbf{p}}$ denote the request probabilities at the respective better equilibria. Then $\hat{\mathbf{p}} < \bar{\mathbf{p}}$.*

Proof: For the proof, we track the best response dynamics with $\mathbf{p}^0 = 0$ and parallel updates (where all users update their probabilities at every k), which are guaranteed to converge to an equilibrium point by Theorem 7. We next show that $\bar{\mathbf{p}}^{\mathbf{k}} \geq \hat{\mathbf{p}}^k$ for every k, thus also at the limit. Note that since $r_i(\bar{p}_i^1, \mathbf{0}) = \bar{\rho}_i \geq r_i(\hat{p}_i^1, \mathbf{0}) = \hat{\rho}_i$, then by the monotonicity of r_i, $\tilde{p}_i^1 \geq p_i^1$ for every i. At the next iteration, $r_i(\bar{p}_i^2, \bar{\mathbf{p}}_{-i}^1) = \bar{\rho}_i \geq r_i(\hat{p}_i^1, \hat{\mathbf{p}}_{-i}^1) = \hat{\rho}_i$. Since $\bar{\mathbf{p}}_{-i}^1 \geq \hat{\mathbf{p}}_{-i}^1$, it follows that $\bar{p}_i^2 \geq \hat{p}_i^2$ for every i. The same argument carries over to subsequent iteration, thus it is valid also at the limit. $\square$