Postprint of: Konorski J., QoS provisioning in a single-cell ad hoc WLAN via queuing and stochastic games, SYSTEMS SCIENCE, Vol. 32, Iss. 3 (2007), pp. 69 - 77

ad hoc WLAN, CSMA/CA game, quality of service

Jerzy KONORSKI*

QoS PROVISIONING IN A SINGLE-CELL AD HOC WLAN VIA QUEUING AND STOCHASTIC GAMES

Selfish manipulation of the binary exponetial backoff scheme in an ad hoc IEEE 802.11 WLAN leads to a noncooperative CSMA/CA game with a payoff structure characteristic of a multiplayer Prisoners' Dilemma. For a simple QoS sensitivity model, assuming elastic traffic flows with a minimum bandwidth share requirement R, we modify the payoffs to define a QoS game. In an ideal scenario, WLAN stations take actions (switch to selfish play) sequentially, in which case for R large enough the game changes into a queuing game with multiple unfair Nash equilibria. Even if the stations are allowed to act simultaneously, these equilibria predict the outcome of the QoS game fairly accurately, depending on the schedule and limits on the number of actions. However, a station can selfishly depart from a common schedule of actions; to account for that, we propose a Markovian stochastic game model and prove the existence of a fair Markov perfect Nash equilibrium.

1. INTRODUCTION

Recent interest in performance models of ad hoc IEEE 802.11 WLANs, using CSMA/CA contention at the MAC level, has spawned a number of Markovian approximations under saturation load [2], [7], [8], [11], [15]. Novel techniques yield the distribution of bandwidth among WLAN stations if each station configures CSMA/CA individually in pursuit of a larger-than-fair bandwidth share, a scenario referred to as a *backoff attack* [7], [8]. Recall that backoff times upon frame collisions vary at random between 0 and $CW - 1$ slots, the *collision window* CW itself varying between a minimum $w_{\min}$ and maximum $w_{\max}$. Backoff attack distorts the bandwidth distribution in favor of stations that configure a smaller $w_{\min}$ and/or $w_{\max}$ than the IEEE standard prescribes.

Let $w_n = \langle w_{n,\min}, w_{n,\max}\rangle$ denote the CSMA/CA configuration at a WLAN station n. Selfish station behavior gives rise to a noncooperative *CSMA/CA game*, where each player (station) $n = 1,\dots,N$ configures $w_n \in W$ so as to maximize its own payoff (bandwidth share) b_n, and W is the set of feasible configurations. Received payoffs are determined by the configuration profile $(w_1,\dots,w_N)$. The likely outcomes of the CSMA/CA game are of interest, as they predict WLAN performance more realistically compared with cooperative models.

A recent game-theoretic study [10] establishes $w_s = \langle 2, 2\rangle$ as the strictly dominant configuration in W for each station, provided that $w_{n,\min} > 1$ i.e., the backoff mechanism is never disengaged. Given such a payoff structure, one envisages a scenario with two types of stations: *selfish* ones have the capability and desire to tamper with w_n and launch a backoff attack by

* Gdańsk University of Technology, ul. Narutowicza 11/12, 80-952 Gdańsk, Poland, jekon@eti.pg.gda.pl

configuring w_s, while *honest* ones have no such capability or desire and stick to a standard configuration e.g., w_h = <16, 1024>. Thus a *binary CSMA/CA game* arises with $W = \{w_s, w_h\}$. Let $b_h(N, x)$ and $b_s(N, x)$ denote the respective bandwidth shares, normalized with respect to the PHY-layer bandwidth, given that x out of the N stations are selfish. The following properties can be proved for any existing IEEE 802.11 setting: $b_s(N, x + 1) > b_h(N, x)$ for $x = 0,\ldots,N - 1$, $b_h(N, x)$ and $b_s(N, x)$ decrease in x, and $b_s(N, N) < b_h(N, 0)$. These findings make the binary CSMA/CA game an instance of N-station Prisoners' Dilemma [3], [17]. Consequently, all-$w_s = (w_s,\ldots\ w_s)$ is the unique *Nash equilibrium* (NE) of the game (a configuration profile where each station plays the best reply to the other stations' configuration profile, hence one from which no station wants to deviate unilaterally, as it would worsen its payoff). Our NE is characterized by "fair distribution of poverty" i.e., equal bandwidth shares summing up to a low bandwidth utilization; a suggestive way of quantifying this poverty is the comparison of *Nash* and *cooperative capacity* [9].

In this paper we ask how the stations' sensitivity to the obtained quality of service (QoS) may affect the payoff structure, and hence the outcome, of the binary CSMA/CA game. In Sec. 2 we discuss a *one-shot QoS game* with the payoffs given by $b_h(N, x)$ and $b_s(N, x)$. Using a simple model of QoS sensitivity we show that as the QoS requirements grow, the game changes from a Prisoners' Dilemma to a queuing game with multiple unfair Nash equilibria, at which a tradeoff between long-term bandwidth utilization and fairness occurs. In Sec. 3 we find that these Nash equilibria apply to the case when the stations take actions (switch to selfish play) sequentially; they predict the outcome fairly accurately even if the stations are allowed to act simultaneously, provided that some common probabilistic schedule of actions is observed. In Sec. 4 we formulate a (Markovian) *stochastic QoS game* model [14] to account for the case where any station is allowed to selfishly depart from the common schedule. If the stations' limits on the number of actions are equal and finite, we prove the existence of a fair *Markov perfect Nash equilibrium*. Sec. 5 concludes the paper.

2. ONE-SHOT QOS GAME

In our simplified QoS sensitivity model, QoS is associated with throughput only, neglecting other performance measures such as delay, jitter and loss. The model assumes that

- traffic flows generated by the application layer at the stations are "lower-bounded AMAP" (as-much-as-possible, or elastic) i.e., a station would eagerly consume any bandwidth share it obtains, provided it is above a certain minimum R below which the station is dissatisfied with the network service,
- R is common to all stations, and
- a station strives to fulfill its QoS requirement while keeping the transmission cost at a minimum; from the analysis of CSMA/CA [10] it follows that a station's frame transmission rate is significantly larger when playing w_s than w_h, therefore it is realistic to assume that a station playing w_s and obtaining a bandwidth share below R perceives the obtained bandwidth share as negative; let this conceptual "bandwidth share" be $b_C < 0$.

The combination of AMAP with a minimum bandwidth share can be justified e.g., by a station's intention to set up a number of variable-quality UDP traffic flows, with the requirement that at least one low-quality flow be carried through. In view of the foregoing assumptions, station n's payoff is the *perceived bandwidth share*:

$$b'_n = \begin{cases} f(b_n), \text{if } b_n \geq R \\ 0, \text{if } b_n < R \text{ and } w_n = w_\mathrm{h} \\ b_\mathrm{C}, \text{if } b_n < R \text{ and } w_n = w_\mathrm{s}, \end{cases} \tag{1}$$

where $f(\cdot)$ is any positive-valued and nondecreasing function. Unless stated otherwise we take $f(b) \equiv b$. With such a payoff function, we term our binary CSMA/CA game a *QoS game*. Here we only consider its one-shot version (where each station selects its CSMA/CA configuration only once) and study the Nash equilibria induced by (1) as functions of R. This subsumes QoS insensitivity as a special case $R = 0$. It is easy to see that the game remains a Prisoners' Dilemma as long as $R \leq b_\mathrm{s}(N, N)$, with the unique NE at all-w_s, and becomes trivial if $R > b_\mathrm{s}(N, 1)$ (the only NE being all-w_h with zero payoffs). For the more interesting range $b_\mathrm{s}(N, N) < R \leq b_\mathrm{s}(N, 1)$, define

$$x_\mathrm{NE}(R) = \max\{x \mid b_\mathrm{s}(N, x) \geq R\}. \tag{2}$$

Any CSMA/CA configuration profile with $x_\mathrm{NE}(R)$ selfish stations and $N - x_\mathrm{NE}(R)$ honest stations is a NE. Indeed, suppose that one of the honest stations switches from w_h to w_s, its bandwidth share becoming $b_\mathrm{s}(N, x_\mathrm{NE}(R) + 1)$. However, by (1) and the definition of $x_\mathrm{NE}(R)$, its perceived bandwidth share then becomes b_C, meaning a payoff decrease: the station now obtains a less-than-required bandwidth share "in exchange for" increased power expenditure. On the other hand, suppose that one of the selfish stations switches from w_s to w_h, its bandwidth share becoming $b_\mathrm{h}(N, x_\mathrm{NE}(R) - 1)$. This is less than the original bandwidth share (recall that $b_\mathrm{s}(N, x + 1) > b_\mathrm{h}(N, x)$), implying again that the station's payoff does not rise.

It is instructive to characterize the Nash equilibria in terms of fairness and bandwidth utilization relative to R. By its nature, $x_\mathrm{NE}(R)$ is a staircase (piecewise constant) function of R, with an initial plateau $x_\mathrm{NE}(R) = N$ corresponding to thea Prisoners' Dilemma, and with unit-size downward jumps at $R = b_\mathrm{s}(N, N)$, $b_\mathrm{s}(N, N - 1)$, …, $b_\mathrm{s}(N, 1)$. Thus $x_\mathrm{NE}(R)$ depends both on N and R in such a way that increased QoS requirements make for bandwidth monopolization by the selfish few.

A salient feature of a noncooperative wireless LAN setting that sets it apart from a cooperative setting is related to a tradeoff between long-term bandwidth utilization and fairness at a NE. We take the former in the sense of the sum of perceived bandwidth shares and the latter in the sense of the Jain fairness index of perceived bandwidth shares [6] i.e.,

$$b'_\Sigma = \sum_{n=1}^{N} b'_n, \tag{3}$$

$$J = \frac{1}{N} \cdot \frac{\left(\sum_{n=1}^{N} b'_n\right)^2}{\sum_{n=1}^{N} (b'_n)^2}. \tag{4}$$

Like $x_\mathrm{NE}(R)$, both the above quantities are staircase functions of R. Assume a basic access 54 Mb/s IEEE 802.11a setting with 1500-byte DATA frames and $N = 10$, and take the numerical values of bandwidth shares from the so-called *mixed model* [8]. Then the Jain index staircase begins at 100% near $R = 0$ (corresponding to ideal fairness at the unique NE of a Prisoners' Dilemma), and ends up at $1/N = 10\%$ for $R > b_\mathrm{s}(10, 2) = 17.8\%$. The bandwidth utilization staircase begins at the Nash capacity equal to $10 \cdot b_\mathrm{s}(10, 10) = 23\%$ and ends up at $b_\mathrm{s}(10, 1) = 65.2\%$. Numerical calculation reveals that $0 \approx b_\mathrm{h}(N, 1)) < b_\mathrm{s}(N, N)$, implying in turn that either $x_\mathrm{NE}(R) = N$ (if $R \leq b_\mathrm{s}(N, N)$) or $b_\mathrm{h}(N, x_\mathrm{NE}(R)) < R$ (otherwise). Thus in both cases, $J = x_\mathrm{NE}(R)/N$. Tradeoffs between long-term resource utilization and fairness, exemplified above, are familiar at the network layer (cf. the parking lot scenario [4]), but alien to known cooperative distributed MAC protocols, where high long-term bandwidth utilization and fairness are ensured, respectively, by optimum parameter configuration and station-wise symmetry.

As an aside, suppose the set W is extended by a third feasible action, "power-off." A station performs it whenever it has switched from w_h to w_s and observed that its obtained bandwidth share has not increased so much as to exceed R. Thus R is the minimum bandwidth share a station must obtain lest it should power off. Supposing further that initially all-w_h prevails, what course of play can be expected? Several colluding stations can exploit the underlying payoff structure to mount a variety of the backoff attack. One may call it the *cartel attack* because it is similar in spirit to certain objectionable business practices. In a simple two-phase scenario, $x_{NE}(R)$ stations simultaneously switch to w_s and subsequently wait until the other stations power off; when the $x_{NE}(R)$ stations have all the bandwidth to themselves, they switch back to w_h. In the process, their obtained bandwidth shares will have evolved from the fair share $b_h(N, 0)$ to $b_s(N, x_{NE}(R))$ after the first phase, and finally to $b_h(x_{NE}(R), 0)$. Numerical calculation shows that for most parameter settings, both phases entail an improvement of the obtained bandwidth share. As an example take the above described IEEE 802.11a setting with $N = 10$ and $R = 10\%$. Using the mixed model [8] we find $x_{NE}(R) = 4$ and the two-phase bandwidth share improvement proceeds from $b_h(10, 0) = 5.5\%$ to $b_s(10, 4) = 7.6\%$ to $b_h(4, 0) = 14.5\%$. For $N = 10$, Fig. 1 plots the first-phase and final bandwidth share of a cartel member (a station that remains selfish at the NE) against R. The two staircases begin at $b_s(10, 10)$ and $b_h(10, 0)$ near $R = 0$, and end up at $b_s(1, 1)$ and $b_h(1, 0)$ for $R > b_s(10, 2) = 17.8\%$. It is visible that the cartel attack is even more beneficial than an ordinary collusion-free backoff attack, the only exception being the case $x_{NE}(R) = 1$: a single selfish cartel member does better than an honest one on account of less backoff overhead. From the viewpoint of cartel members, increased QoS requirements imply higher bandwidth shares and fewer shareholders. In general, adding the "power-off" action for $R > 0$ steps up undesirable collusion incentives and brings about even more unfair Nash equilibria. Therefore we will further confine our interest to $W = \{w_h, w_s\}$.

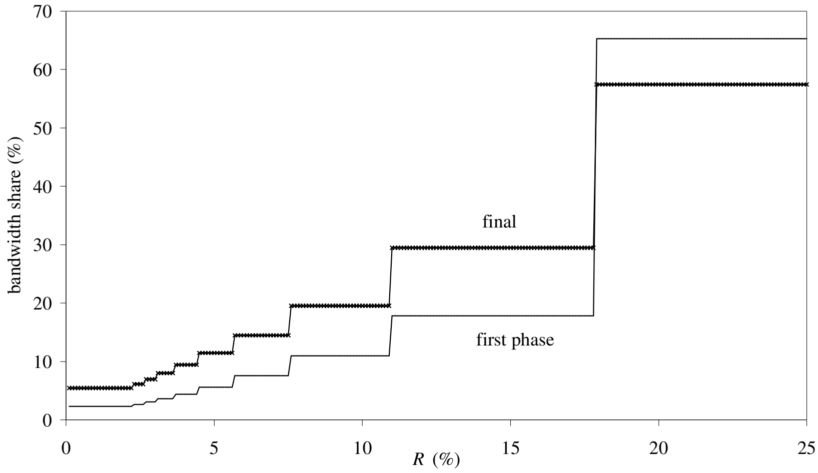

Fig. 1. First-phase and final bandwidth shares of selfish stations at a NE in the two-phase scenario ($N = 10$).

3. MULTISTAGE QoS GAME SCENARIOS

Unlike the QoS insensitive case ($R = 0$), the discussion in the previous section leaves a trace of doubt as to how and whether indeed the Nash equilibria described by (2) can be reached in a conceivable course of play starting from given initial conditions. The doubt may arise even in the context of the one-shot QoS game: since there are multiple (unfair) Nash equilibria, it may be conjectured that the outcome of the game depends on a schedule according to which the stations select their configurations. If $R \le b_h(N, 0)$ then certain schedules defined by the so-called SPELL [8] or CRISP [10] strategies lead to all-w_h, a satisfactory outcome for all. To explore beyond that, we assume $R > b_h(N, 0)$. It is therefore not possible for all the stations to fulfill their QoS requirements simultaneously.

Let the game start at all-w_h. Subsequently each station proceeds by tentatively switching to w_s at some moment, whereupon it checks if the obtained bandwidth share has exceeded R. If affirmative, the newly configured w_s is kept forever,[1] otherwise the station retreats to w_h forever. In an unfortunate albeit unlikely scenario, all the stations switch to w_s almost simultaneously, to find the obtained bandwidth share below R, and retreat to w_h, each ending up with a zero payoff. At the other extreme, stations switch to w_s sequentially, each being able to notice the effect of previous actions. The game then falls into the class of *queuing games* described in [12]. To paraphrase an illustrative example given there, imagine a group of airline passengers sitting in a departure hall and waiting for the boarding to commence. After a while, some of them cannot resist the urge to queue up at the counter. Early queuers have their pick of seats aboard the plane, but for late queuers this benefit is reduced and may be offset by the fatigue of standing instead of sitting. Thus a NE arises with a number of passengers standing and the other sitting. The same course of play in the QoS game leads to a NE with $x_{NE}(R)$ selfish stations. Let the time taken to decide whether to keep w_s forever or retreat to w_h be termed *switch window*. If switch windows are negligibly short, so that different stations' switch windows almost never overlap, the queuing game assumptions hold. Upon tentatively switching to w_s, each station retreats to w_h forever if $b_s(N, x + 1) < R$, where x is the number of stations that so far have decided to keep w_s forever. Note that each station only needs to tentatively switch to w_s once.[2]

In practice, however, a switch window must last long enough for a station to estimate the obtained bandwidth share and compare it with R. Therefore, switch windows may overlap. If at some instant of time, x stations out of N have decided to keep w_s forever, where x is such that $b_s(N, x + 2) < R < b_s(N, x + 1)$, then two other stations whose switch windows overlap will retreat to w_h forever upon a tentative switch to w_s whereas one of them could have kept w_s forever had their switch windows not overlapped. This is in fact a stand-off situation [5]. A form of insistence can remedy the problem: each station is allowed to switch to w_s repeatedly. Assume that to constrain the transmission cost of fulfilling its QoS objective, a station allows itself TS tentative switches to w_s at the maximum ($1 \le TS < \infty$). Two characteristics are subject to optimization, TS and the *switch interval* between tentative switches to w_s. Clearly, a very short switch interval is non-optimal in the sense of (1), as it may lead to frequent overlaps of switch windows and unnecessary retreats to w_h. To simplify the argument, we do not expand upon the optimization; neither do we consider for the moment that the tentative switches to w_s may be conditional on the observed past actions of other stations (this, along with the case $TS = \infty$, is discussed in the next section). Instead, we assume that switch intervals are geometrically distributed and that TS is fixed. Furthermore, let switch windows at all stations synchronize to the same stream of timeframes as in Fig. 2, and let P_s be the probability of a station's tentative switch in a given timeframe. The remaining number of tentative switches at

[1] For simplicity we consider an infinite-horizon game i.e., we do not set a specific time when the game ends.

[2] This can be easily formalized using the general approach of *multistage games* [2]. Let σ denote the strategy "if the bandwidth share obtained upon tentatively switching to w_s exceeds R then keep w_s forever, otherwise keep w_h forever." Then the strategy profile all-σ turns out to be a *subgame perfect NE* of a multistage QoS game, leading to the outcome described by (2).

station n is controlled by a counter C_n, initially set to TS. By convention, let $C_n = -1$ mark a station that has decided to keep w_s forever, while $C_n = 0$ means that w_h has been retreated to forever.

Station n executes the following code, where random(1) returns a pseudorandom real number between 0 and 1:

```
<at the game start>
C_n := TS

<at the beginning of timeframe>
 if (C_n > 0) and (random(1) < P_s) then {
    tentatively switch to w_s;
    if b_n ≥ R then {C_n := −1; keep w_s forever}  else {C_n := C_n − 1; if C_n = 0 then keep w_h forever}
 }
```

In the simple scenario of Fig. 2, $N = 5$, $R = 10\%$, and $TS = 2$. From the mixed model [8] we get $b_s(5, 2) = 18.1\%$ and $b_s(5, 3) = 11.1\%$, implying $x_{NE}(R) = 3$. In timeframe 2, stations 4 and 5 start their switch windows and keep w_s forever, as they have obtained $b_s(5, 2) = 18.1\%$. Stations 1 through 3 start theirs in timeframe 5 and subsequently retreat to w_h since they each have obtained $b_s(5, 5) = 5.7\%$. Stations 1 and 2 tentatively switch again in timeframe 8 to obtain $b_s(5, 4) = 7.6\%$; since they have exhausted their tentative switch limit, they retreat to w_h forever. Station 3 has the good sense to wait until timeframe 10 to tentatively switch again, and keeps w_s forever upon obtaining $b_s(5, 3) = 11.1\%$. The final number of selfish stations thus equals $x_{NE}(R)$. Had station 3 tentatively switched again in timeframe 8, as did stations 1 and 2, there would be only $x_{NE}(R) - 1 = 2$ selfish stations in the end.

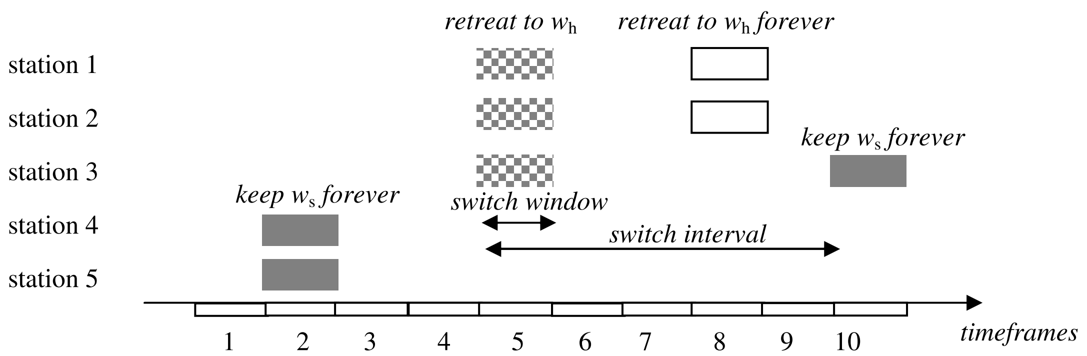

Fig. 2. Simple multistage scenario of the QoS game.

To see if there is much prospect of a station fulfilling its QoS requirement, and how frequently $x_{NE}(R)$ correctly predicts the outcome of the QoS game, we write down the dynamics of tentative switches and apply Monte Carlo simulation. Let $A_{P_s}(C_n)$ be a binary random variable indicating a station's tentative switch in a timeframe, given the current counter value:

$$A_{P_s}(C_n) = \begin{cases} 1, \text{if } C_n > 0 \text{ and random}(1) < P_s \\ 0, \text{otherwise.} \end{cases}$$

Denoting by $\boldsymbol{C}^k = (C_1^k, ..., C_N^k)$ the vector of the stations' counters in timeframe k and by $x(\boldsymbol{C}^k)$ the number of its "−1" elements, we reflect the above code as follows:

$$C_n^{k+1} = \begin{cases} C_n^k \cdot \mathbf{1}_{x(\boldsymbol{C}^k)+A_{P_s}(C_1^k)+...+A_{P_s}(C_N^k)>x_{\mathrm{NE}}(R)} -1, \text{if } A_{P_s}(C_n^k)=1 \\ C_n^k, \text{otherwise.} \end{cases} \quad (5)$$

A simulation run starts with $\boldsymbol{C}^0 = (TS,...,TS)$ and ends when $x(\boldsymbol{C}^K) = x_{\mathrm{NE}}(R)$ or $\boldsymbol{C}^K \leq \mathbf{0}$ for some K, in each case producing a final value of $x(\boldsymbol{C}^K)$. Fig. 3 depicts the average of this value as a percentage of N, obtained for N = 10, R = 10% and 20%, and various P_s and TS (Monte Carlo runs were repeated until the relative 95% confidence intervals narrowed below 5%). Note that this also represents the probability that a station is among those to eventually keep w_s forever, thus obtains a bandwidth share of R or more i.e., fulfills the QoS requirement. Furthermore, $x(\boldsymbol{C}^K) \leq x_{\mathrm{NE}}(R)$, hence the plot can be compared with $x_{\mathrm{NE}}(R)/N$, representing the ideal prediction of the outcome based on (2). Generally, the prediction improves as TS increases and worsens as R and P_s increase; the impact of increasing R is due to the difficulty of electing exactly one (or a small number) of stations to keep w_s forever. Regardless of R, the prediction tends to ideal as TS increases; for small enough P_s this happens quite rapidly.

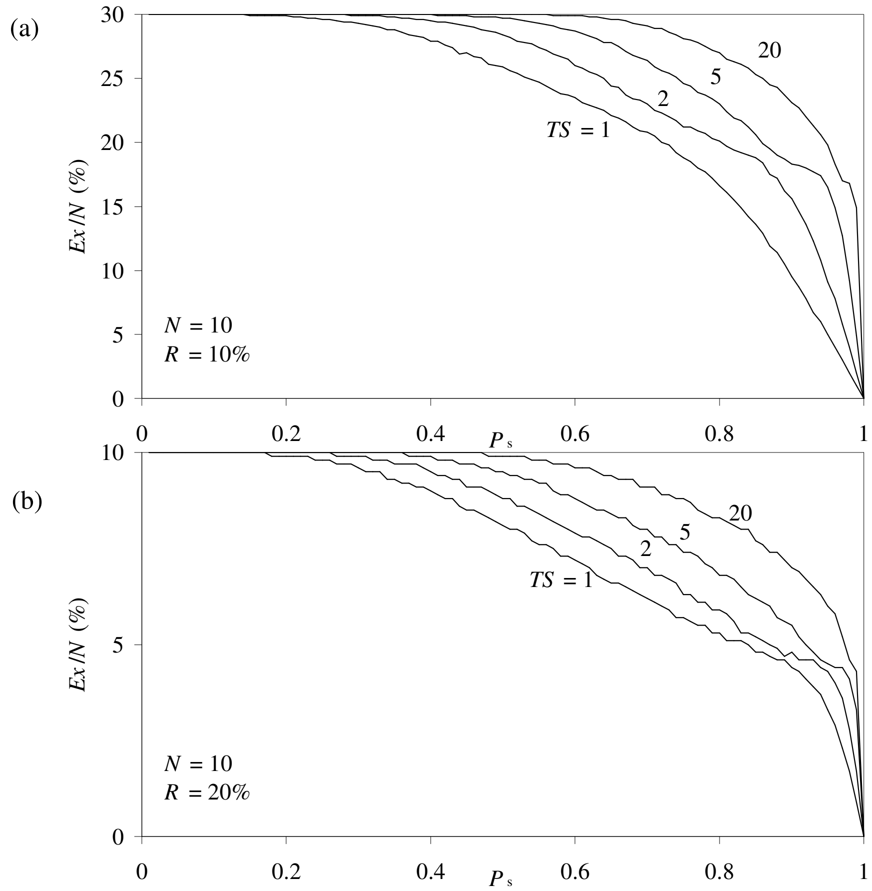

Fig. 3. Probability of fulfilling QoS requirement; (a) R = 10% ($x_{\mathrm{NE}}(R)/N$ = 30%), (b) R = 20% ($x_{\mathrm{NE}}(R)/N$ = 10%).

4. STOCHASTIC QoS GAME

Let us now address the case where the probability of a tentative switch to w_s by a station is conditional on the current counter vector $\boldsymbol{C}$. Thus each station is assumed to know $\boldsymbol{C}$ at all times, an idealization requiring the knowledge of N and the number of stations that switch to w_s in a given timeframe (these might be inferred via coarse profile observability as proposed in [10], if the profile observability threshold $x^* = N$). Supposing that each station maximizes a utility of its own, a (Markovian) *stochastic game* arises, whose likely outcome is a *Markov perfect NE*. Both these notions are special cases of a multistage game and a subgame perfect NE, respectively; a stochastic game model was applied to an ALOHA game [13]. For clarity, we alter the wording and notation of definitions from [3], [14] to fit in with our QoS game.

Definition:

(i) Consider an N-station infinite-horizon multistage game whose stages correspond to QoS game timeframes, with $W = \{w_h, w_s\}$, stage payoffs given by (1), and the set of play paths Π; a play path up to timeframe k has the form $\boldsymbol{\pi}^k = (\boldsymbol{w}^1,\ldots, \boldsymbol{w}^k)$, where $\boldsymbol{w}^k = (w_1^k,...,w_N^k)$ is the configuration profile in timeframe k. A (Markovian) *stochastic game* in addition defines: (1) a function *state*: $\Pi \to SS$ that maps play paths onto *game states*, the state space SS being countable and $\boldsymbol{C}$ being a generic game state (current counter vector), and (2) station n's strategy in the form of a function σ_n: $SS \to [0, 1]$ such that $\sigma_n(\boldsymbol{C}) = Pr[w_n^k = w_s \mid state(\boldsymbol{\pi}^{k-1}) = \boldsymbol{C}]$ i.e., $\sigma_n(\boldsymbol{C})$ is the probability of a tentative switch to w_s in timeframe k conditional on the current game state ($\sigma_n(\boldsymbol{C}) = 0$ if $C_n = 0$). Define the stations' utilities as guaranteed asymptotic averages of stage payoffs, dependent on the strategy profile ($\sigma_1,\ldots, \sigma_N$) and the initial game state:

$$U_n(\sigma_1,...,\sigma_N \mid (TS,...,TS)) = \liminf_{k\to\infty} E_{\mu(\sigma_1,...,\sigma_N \mid (TS,...,TS))} b_n'^k, \tag{6}$$

where $\mu^k(\sigma_1,\ldots, \sigma_N \mid \boldsymbol{C})$ is the probability distribution of stage k payoffs given the initial game state $\boldsymbol{C}$ and $\liminf_{k\to\infty} a_k$, defined as $\lim_{k\to\infty} \inf\{a_k, a_{k+1}, \ldots\}$, exists for any bounded sequence $(a_k)_{k=1,2,\ldots}$.

(ii) ($\sigma_1,\ldots, \sigma_N$) is a *Markov perfect NE* (MPNE) of the stochastic game if

$$U_n(\sigma_1,...,\sigma_n,...,\sigma_N \mid \boldsymbol{C}) \geq U_n(\sigma_1,...,\sigma',...,\sigma_N \mid \boldsymbol{C}), \tag{7}$$

for $n = 1,\ldots,N$, and any station n strategy σ' and $\boldsymbol{C} \in SS$. That is, starting from any game state, each station's strategy is a best reply to the other stations' strategy profile. The MPNE is *fair* if $\sigma_1 = \ldots = \sigma_N$ and *degenerate* if $\sigma_n(\boldsymbol{C}) \equiv 1$.

Game states in successive timeframes constitute a Markov chain whose transition matrix is determined by ($\sigma_1,\ldots, \sigma_N$) in a manner similar to (5). State $\boldsymbol{C} = \boldsymbol{0}$ and states with $x(\boldsymbol{C}) = x_{NE}(R)$ are absorbing (if TS is finite then so is the number of stages). For each state $\boldsymbol{C}$, let the random variable $NS_{\sigma_1,...,\sigma_N}(\boldsymbol{C})$ represent the game state in the next timeframe. $U_n(\sigma_1,...,\sigma_N \mid \boldsymbol{C})$ is the solution of the following linear system with unknowns $u(\boldsymbol{C})$, stemming from the law of total expectation:

$$u(\boldsymbol{C}) = Eu(NS_{\sigma_1,...,\sigma_N}(\boldsymbol{C})), \; \boldsymbol{C} \in SS, \tag{8}$$

where the expectation is taken with respect to the probabilities of $NS_{\sigma_1,\ldots,\sigma_N}(\boldsymbol{C})$.

Proposition: For $R > b_{\mathrm{h}}(N, 0)$,

(i) if $TS < \infty$ then there exists a fair non-degenerate MPNE all-σ.
(ii) if $TS = \infty$ then there exists no fair MPNE.

Proof: To prove (i), first exclude the case $\sigma_m(\boldsymbol{C}) \equiv 1$ for $m \neq n$ from the considerations. Indeed, it cannot be part of a fair MPNE: station n using $\sigma_n(\boldsymbol{C}) \equiv 1$ then receives a utility of $b_{\mathrm{C}} < 0$, whereas using e.g., $\sigma_n(\boldsymbol{C}) \equiv P_{\mathrm{s}} \in (0, 1)$ yields $b_{\mathrm{s}}(N, 1)\cdot(1 - (P_{\mathrm{s}})^{TS}) \geq 0$. Therefore, if a fair MPNE exists, it is non-degenerate. To show that it does exist, we use a fixed-point argument. In brief, if $\sigma_m \equiv \sigma$ for all $m \neq n$, then the correspondence between σ and the set Σ_n of station n strategies fulfilling (7) is the solution of a classical dynamic programming problem i.e.,

$$\Sigma_n = \arg\max_{\sigma_n} U_n(\sigma_n, \sigma \,|\, \boldsymbol{C}) \tag{9}$$

subject to (8). Note that (8) is linear in σ_n and U_n is continuous in σ_n. Therefore the point-to-set mapping (9) is convex-valued and, by Berge's Theorem of the Maximum [1], upper hemi-continuous. Consequently, by Kakutani's theorem [16], it has a fixed point corresponding to a fair MPNE of the QoS game.

If $TS = \infty$, it is easy to notice that U_n has a discontinuity at $\sigma_m(\boldsymbol{C}) \equiv 1$ (as seen from the reasoning below). Hence, Berge's and Kakutani's theorems do not apply. To prove (ii), notice that the game state specification may be reduced to $x(\boldsymbol{C})$, the number of "−1" components in $\boldsymbol{C}$, whereupon (8) becomes uniform for all $n = 1,\ldots,N$. However, $u(x(\boldsymbol{C}))$ depends on whether $C_n = -1$, 0, or > 0; to avoid this ambiguity define $U_n(\sigma_1,\ldots,\sigma_N \,|\, \boldsymbol{C})$ for $\boldsymbol{C}$ with $C_n > 0$ only. Then $x(\boldsymbol{C}) = 0,\ldots,x_{\mathrm{NE}}(R)$ are only relevant and obviously, $u(x_{\mathrm{NE}}(R)) = 0$. Letting $\sigma_n = \sigma'$ and $\sigma_m \equiv \sigma$ for all $m \neq n$ and recalling (5) we rewrite (8) as follows:

$$\begin{aligned} u(x) = u(x)\cdot&\left[\sigma'(x)B(x_{\mathrm{NE}}(R)-x) + \sum_{j=x_{\mathrm{NE}}(R)-x+1}^{N-x-1} B(j)\right] + \\ &\sum_{j=0}^{x_{\mathrm{NE}}(R)-x-1} B(j)\big[(1-\sigma'(x))\cdot u(x+j) + \sigma'(x)\cdot b_{\mathrm{s}}(N, x+j+1)\big], \end{aligned} \tag{10}$$

where $B(j) = \binom{N-x-1}{j}\cdot(\sigma(x))^j\cdot(1-\sigma(x))^{N-x-1-j}$. Taking (10) at $x = x_{\mathrm{NE}}(R) - 1$ we have

$$u(x_{\mathrm{NE}}(R)-1) = \frac{\sigma'(x)\cdot B(0)\cdot b_{\mathrm{s}}(N, x_{\mathrm{NE}}(R))}{1 - \sum_{j=2}^{N-x_{\mathrm{NE}}(R)} B(j) - \sigma'(x)\cdot\big[B(1)-B(0)\big]}. \tag{11}$$

If $\sigma(x) \neq 1$, this is maximized by $\sigma'(x) = 1$, otherwise the right-hand side of (11) is zero for all $\sigma'(x)$. Hence, $\sigma'(x) = \sigma(x) = 1$ is the only candidate for a fair MPNE, yielding $u(x_{\mathrm{NE}}(R) - 1) = b_{\mathrm{C}}$. Further, taking $x = x_{\mathrm{NE}}(R) - 2$ yields

$$
\begin{aligned}
&u(x_{\mathrm{NE}}(R)-2) \\
&= \frac{\sigma'(x)\cdot\left[B(0)b_{\mathrm{s}}(N,x_{\mathrm{NE}}(R)-1)+B(1)b_{\mathrm{s}}(N,x_{\mathrm{NE}}(R))\right]+(1-\sigma'(x))B(1)\cdot u(x_{\mathrm{NE}}(R)-1)}{1-\sum_{j=3}^{N-x_{\mathrm{NE}}(R)}B(j)-\sigma'(x)\cdot\left[B(1)-B(0)\right]}. \\
&= \frac{\sigma'(x)\cdot\left[B(0)b_{\mathrm{s}}(N,x_{\mathrm{NE}}(R)-1)+B(1)b_{\mathrm{s}}(N,x_{\mathrm{NE}}(R))\right]+(1-\sigma'(x))B(1)\cdot b_{\mathrm{C}}}{1-\sum_{j=3}^{N-x_{\mathrm{NE}}(R)}B(j)-\sigma'(x)\cdot\left[B(1)-B(0)\right]},
\end{aligned}
\tag{12}
$$

to which the same reasoning applies. Solving (10) successively for descending x one easily finds that the only candidate for a fair MPNE is a degenerate $\sigma_n(x) \equiv 1$, leaving each station with a utility of b_{C}. However, if $\sigma_m(x) \equiv 1$ for all $m \neq n$ then using $\sigma_n(x) \equiv 0$ yields station n a zero utility, which is higher than b_{C}. Therefore no fair MPNE exists.

In the case $TS < \infty$, finding a fair MPNE is difficult except for very small TS and N, due to the prohibitive dimensionality of the state space. Recall that from the viewpoint of any given station n, the other stations are anonymous; hence, all distinct game states in SS_n can be encoded in the form $\boldsymbol{C} = (C_n, N_{-1}, N_0, \ldots, N_{TS})$, where $0 < C_n \leq TS$ and N_i is the number of stations other than n whose counters currently read i (with $N_{-1} < x_{\mathrm{NE}}(R)$ and $N_{-1} + N_0 + \ldots + N_{TS} = N - 1$). It follows that SS_n contains $TS \cdot \sum_{j=0}^{x_{\mathrm{NE}}(R)-1} \binom{TS+N-j}{N-1-j}$ states; for $N = 10$ and $TS = 10$ this gives $1.7 \cdot 10^6$ if $x_{\mathrm{NE}}(R) = 1$ and $2.4 \cdot 10^6$ if $x_{\mathrm{NE}}(R) = 2$. Another difficulty is that not all distinct game states are observable to station n, which is only able to record the *number* of stations tentatively switching to w_{s} in successive stages. The search for a fair MPNE must therefore rely on further simplifications.

For illustration, let us confine feasible station m strategies to $\sigma_m(\boldsymbol{C}) \equiv P_{\mathrm{s}}$ ($m \neq n$) and station n strategies to $\sigma_n(\boldsymbol{C}) \equiv P'_{\mathrm{s}}$. Fig. 4 depicts $U_n(P_{\mathrm{s}}, P'_{\mathrm{s}} | (TS,\ldots,TS))$ against P'_{s} for $N = 10$, $R = 10\%$ (i.e., $x_{\mathrm{NE}}(R)/N = 30\%$), and $TS = 2$ and $TS = 20$. In (1), we have taken $f(b) \equiv 1$; notice that by (6), station n utility is then the probability of ultimately fulfilling the QoS requirement; hence the values in Fig. 4 can easily be compared with those in Fig. 3. Black dots are drawn at $P'_{\mathrm{s}} = P_{\mathrm{s}}$, so that a fair MPNE corresponds to the curve (drawn thick) whose maximum coincides with the dot. As P_{s} decreases, the curve maxima move to the right. For $TS = 2$, a fair MPNE can be clearly observed (Fig. 4a, $P'_{\mathrm{s}} = P_{\mathrm{s}} = 0.48$); for $TS = 20$, MPNE occurs at $P'_{\mathrm{s}} = P_{\mathrm{s}} = 0.625$ (Fig. 4b). For larger TS, decreasing P_{s} causes the value at $P'_{\mathrm{s}} = 1$ compete ever more with the local maximum at $P'_{\mathrm{s}} < P_{\mathrm{s}}$ (ultimately the local maximum vanishes, rendering $P'_{\mathrm{s}} = 1$ station n's best reply. This explains part (ii) of the Proposition.

Note that an edge over the other stations in terms of the limit TS coaxes station n into forcing tentative switches to w_{s} ($P'_{\mathrm{s}} = 1$) irrespective of TS, as illustrated by the dashed curves.

5. CONCLUSION

To illustrate the nature of games arising when selfish WLAN stations become QoS sensitive, we have considered a simple model of AMAP (elastic) traffic flows with a minimum bandwidth share requirement, R. We have demonstrated that as R increases, the game changes from a Prisoners' Dilemma to a kind of a queuing game. One of interesting consequences of increasing R is the tradeoff between long-term bandwidth utilization and fairness, unfamiliar in symmetric distributed MAC protocols. Using a plausible multistage game scenario we have argued that a simple analysis of one-shot QoS game can predict the outcome of the stations' interaction in most cases. For multistage QoS game scenarios we have proposed a stochastic game model and proved the existence of a fair Markov perfect Nash equilibrium for finite TS.

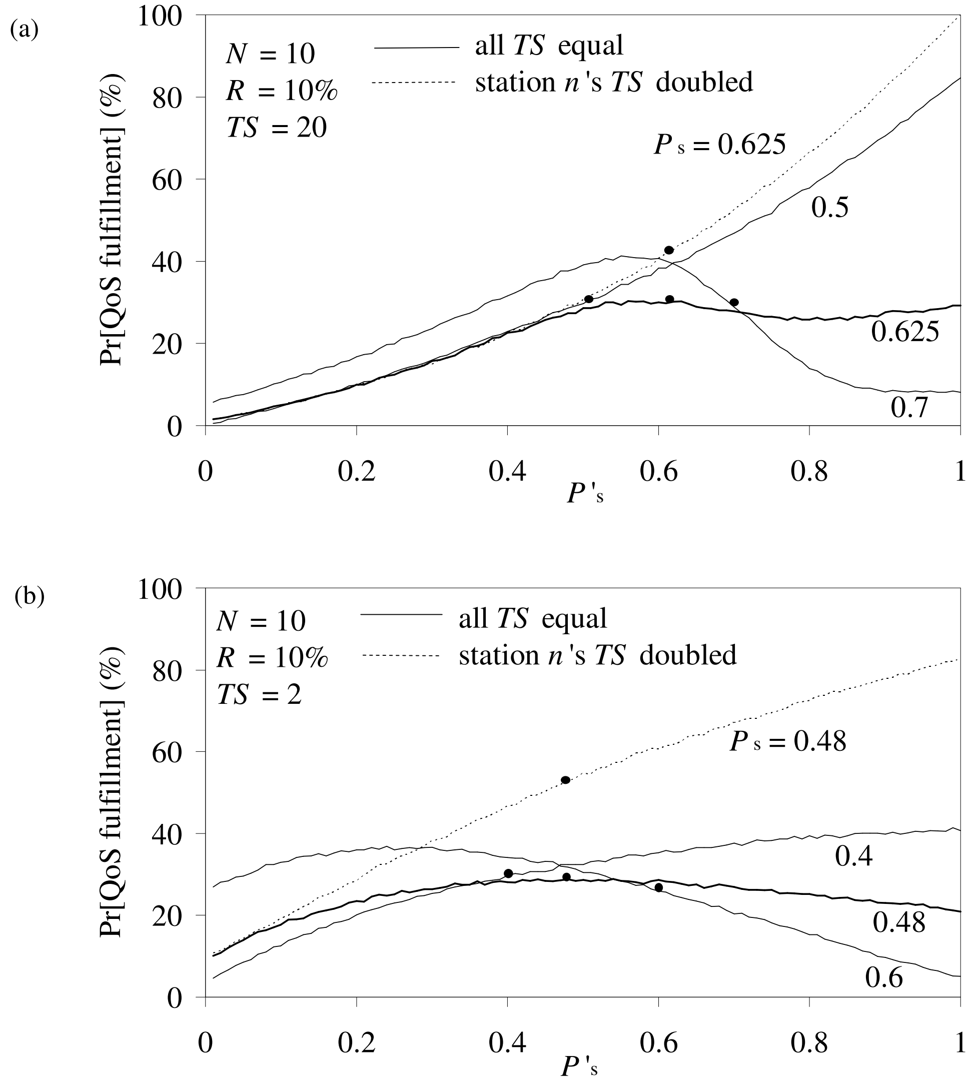

Fig. 4. Illustration of an MPNE; (a) TS = 2, (b) TS = 20

ACKNOWLEDGMENT

This work was supported by the Ministry of Education and Science, Poland, under Grant 1599/T11/2005/29.

REFERENCES

[1] BERGE C., *Topological Spaces*, Macmillan, New York 1963.

[2] BIANCHI G., *Performance Analysis of the IEEE 802.11 Distributed Coordination Function*, IEEE J. on Selected Areas in Commun., Vol. 18 (2000), No. 3, pp. 535-547.

[3] FUDENBERG D., TIROLE J., *Game Theory*, MIT Press 1991.

[4] GEVROS P., CROWCROFT J., KIRSTEIN P., BHATTI S., *Congestion Control Mechanisms and the Best Effort Service Model*, IEEE Network, Vol. 15 (2001), No. 3, pp. 16-26.

[5] GROVER G.A., JAFFE J.M., *Stand-off and Stand-off Resolution in Deadlock-free Networks*, IBM Tech. Rep. 7/30/85, 1985.

[6] JAIN R., *Fairness: How to measure quantitatively?*, ATM Forum/94-0881, Sept. 1994.

[7] KONORSKI J., *Solvability of a Markovian Model of an IEEE 802.11 LAN under a Backoff Attack*, IEEE Comp. Soc. Press 2005.

[8] KONORSKI J., *Playing CSMA/CA Game to Deter Backoff Attacks in Ad Hoc Wireless LANs*, LNCS 3738, Springer-Verlag 2005, pp. 127-140.

[9] KONORSKI J., *Nash Capacity of an Ad Hoc IEEE 802.11 LAN under Backoff Attack*, Proc. 12th Polish Teletraffic Symp., Poznań 2005.

[10] KONORSKI J., *A Game-Theoretic Study of CSMA/CA under a Backoff Attack*, to appear in ACM/IEEE Trans. on Networking.

[11] KUMAR A., ALTMAN E., MIORANDI D., GOYAL M., *New Insights from a Fixed-Point Analysis of Single Cell IEEE 802.11 WLANs*, Proc. IEEE INFOCOM 2005, Miami 2005.

[12] McCAIN R.A., *Game Theory: A Nontechnical Introduction to the Analysis of Strategy*, Thomson SouthWestern 2003.

[13] MacKENZIE A.B., WICKER S.B., *Stability of Multipacket Slotted ALOHA with Selfish Users and Perfect Information*, Proc. IEEE INFOCOM 2003, San Francisco 2003.

[14] NEYMAN A., SORIN S., *Stochastic games and applications*, Kluwer Academic Publishers 2003.

[15] RAMAIYAN V., KUMAR A., ALTMAN E., *Fixed Point Analysis of Single Cell IEEE 802.11e WLANs: Uniqueness, Multistability and Throughput Differentiation*, Proc. SIGMETRICS'05, Banff, Canada 2005.

[16] RASMUSSEN E., *Games and Information: An Introduction to Game Theory*, 3rd ed., Blackwell Publ. 2001.

[17] YAO X., *Evolutionary Stability in the n-Person Iterated Prisoners' Dilemma*, BioSystems, Vol. 39 (1996), pp. 189-197.