

A Game Theoretic Approach to Distributed Opportunistic Scheduling

Albert Banchs, *Senior Member, IEEE*, Andres Garcia-Saavedra, Pablo Serrano, *Member, IEEE*, and Joerg Widmer, *Senior Member, IEEE*

Abstract—Distributed Opportunistic Scheduling (DOS) is inherently more difficult than conventional opportunistic scheduling due to the absence of a central entity that knows the channel state of all stations. With DOS, stations use random access to contend for the channel and, upon winning a contention, they measure the channel conditions. After measuring the channel conditions, a station only transmits if the channel quality is good; otherwise, it gives up the transmission opportunity. The distributed nature of DOS makes it vulnerable to *selfish* users: by deviating from the protocol and using more transmission opportunities, a selfish user can gain a greater share of wireless resources at the expense of “well-behaved” users. In this paper, we address the problem of selfishness in DOS from a game theoretic standpoint. We propose an algorithm that satisfies the following properties: (i) when all stations implement the algorithm, the wireless network is driven to the optimal point of operation, and (ii) one or more selfish stations cannot obtain any gain by deviating from the algorithm. The key idea of the algorithm is to react to a selfish station by using a more aggressive configuration that (indirectly) punishes this station. We build on *multivariable control theory* to design a mechanism for punishment that is sufficiently severe to prevent selfish behavior yet not so severe as to render the system unstable. We conduct a game theoretic analysis based on *repeated games* to show the algorithm’s effectiveness against selfish stations. These results are confirmed by extensive simulations.

Index Terms—Contention-based channel access, distributed opportunistic scheduling, game theory, multivariable control theory, repeated games, selfish stations, wireless networks

I. INTRODUCTION

OPPORTUNISTIC scheduling techniques have been shown to significantly improve performance in wireless networks. These techniques take advantage of the fluctuations in the channel conditions of different wireless stations over time; by selecting the station with the best instantaneous channel for data transmission, opportunistic scheduling can

utilize wireless resources more efficiently. A key assumption of most opportunistic scheduling techniques [1], [2] is that the scheduler is centralized and has knowledge of the instantaneous channel conditions of all stations.

Distributed Opportunistic Scheduling (DOS) techniques [3]–[6] have been proposed only recently. In contrast to centralized schemes, with DOS each station has to make scheduling decisions without knowing the channel conditions of the other stations. Stations contend for the channel using random access with a given access probability. After successful contention, a station measures the channel and, if the channel conditions are poor (i.e., the instantaneous transmission rate is below a given threshold), the station gives up the transmission opportunity. This allows all stations to contend for the channel again, letting a station with better conditions win the contention, which increases the overall throughput. DOS techniques thus exploit both multi-user diversity across stations and time diversity across slots.

The lack of global channel information makes DOS systems very vulnerable to *selfish* users. By deviating from the above protocol and using a more aggressive configuration, a selfish user can easily gain a greater share of wireless resources at the expense of the other, well-behaved users. In this paper, we address the problem of selfishness in DOS from a game theoretic standpoint. In our formulation of the problem, the players are wireless stations that implement DOS and strive to obtain as great a share of resources as possible from the wireless network. We show that, in the absence of penalties, the wireless network naturally tends to either great unfairness or network collapse. Building on this result, we design a penalty mechanism in which any player who misbehaves will be punished by other players in such a way that there is no incentive to misbehave. A key challenge when designing such a penalty scheme is to carefully adjust the punishment inflicted on a misbehaving station. If the punishment is too light, a selfish station may still benefit from misbehaving. If it is too excessive, however, the punishment itself could be interpreted as misbehavior and trigger punishment from other stations, leading to an endless spiral of increasing punishments and ultimately throughput collapse. We address this challenge through a combination of game theory and control theory.

The most relevant prior work on DOS by Zheng et al. [3] lays the basic foundations of distributed opportunistic scheduling. The authors propose a mechanism based on optimal stopping theory and analyze its performance with well-behaved as well as selfish users. The aim of the algorithm is to maximize

Manuscript received June 28, 2011; revised March 23, 2012 and September 26, 2012; accepted November 04, 2012; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor M. Liu. Date of publication xxxx xx, xxxx; date of current version xxxx xx, xxxx. This work has been supported by the European Community’s Seventh Framework Programme (FP7-ICT-2009-5) under grant agreement no. 257263 (FLAVIA project) and by the Spanish Ministry of Economy and Competitiveness under grant agreement TIN2009-13992-C02-01 (QUARTET project).

A. Banchs is with the University Carlos III of Madrid and with the Institute IMDEA Networks (e-mail: banchs@it.uc3m.es).

A. Garcia-Saavedra and P. Serrano are with the University Carlos III of Madrid (e-mail: agsaaved@it.uc3m.es@it.uc3m.es; pablo@it.uc3m.es).

J. Widmer is with the Institute IMDEA Networks (e-mail: joerg.widmer@imdeia.org).

Digital Object Identifier xx.xxxx/TNET.xxxx.xxxxxx

the total throughput of the network. [4]–[6] extend the basic mechanism of [3] by analyzing the case of imperfect channel information [4], improving channel estimation through two-level channel probing [5], and incorporating delay constraints [6]. While our algorithm deals with the basic DOS mechanism of [3], it can be extended to incorporate the enhancements of [4]–[6]. The key contributions of our work are as follows:

- 1) We perform a joint optimization of both the transmission rate thresholds and the access probabilities, whereas [3] only optimizes the thresholds.
- 2) We provide a proportionally fair allocation that achieves a good tradeoff between total throughput and fairness. In contrast, [3] maximizes the total throughput of the network and, as a result, it risks starving stations with poor channel conditions.
- 3) We propose a simple algorithm based on control theory that guarantees stability and quick convergence to the optimal point of operation, in contrast to the comparatively complex heuristics of [3].
- 4) Our game theoretic analysis considers that users can selfishly configure both their access probability and transmission rate threshold, whereas the analysis of [3] assumes that selfish users only have control over the thresholds.
- 5) We use a penalty mechanism to force an optimal Nash equilibrium, whereas [3] introduces a pricing mechanism which may not be practical in many scenarios; furthermore, the performance of the pricing mechanism relies heavily on the cost parameter and it is suboptimal even for the best parameter setting.

Some of the concepts and tools used in this paper build on our previous works of [7] and [8]. In [7] we proposed an algorithm based on control theory to optimally adjust the configuration of DOS. In contrast to [7], in this paper our aim is to prevent selfish users from obtaining any benefit by deviating from the optimal configuration, which is a much more difficult problem. In [8] we designed an algorithm based on multivariable control theory to adjust the contention parameters of a WLAN. In this paper, we obtain the same linearized system as [8], and hence we use the corresponding part of the analysis from that paper; however, the purpose, algorithm and most of the analysis of this paper differ substantially from [8].

The remainder of the paper is organized as follows. In Section II, we present an analysis of our system and derive the optimal configuration of access probabilities and transmission rate thresholds. In Section III, we show that, in the absence of penalties, the wireless network tends to a highly undesirable resource allocation. We then propose an algorithm called *Distributed Opportunistic scheduling with distributed Control* (DOC) that avoids this by implementing a decentralized penalty mechanism to control selfish users. Section IV shows by means of control theory, that when all the stations implement DOC, the system is stable and converges to the optimal point of operation derived in Section II. In Section V, we conduct a game theoretic analysis of DOC to show that stations cannot obtain any gain by behaving selfishly. The performance of the proposed scheme is extensively eval-




Fig. 1. Example of channel contention.

uated through simulations in Section VI. Finally, Section VII provides some concluding remarks.

II. ANALYSIS AND OPTIMAL CONFIGURATION

In the following, we present our system model and analyze the throughput as a function of the access probabilities and transmission rate thresholds. We then compute the optimal configuration of these parameters for a proportionally fair throughput allocation, which is well known to provide a good tradeoff between total throughput and fairness.¹

A. System model

Our system model follows that of [3]–[6]. We consider a single-hop wireless network with N stations, where station i contends for the channel with an access probability p_i . We assume a collision model for the channel access, where a station contends successfully for the channel if no other station contends at the same time. Let τ denote the duration of a mini slot for channel contention, which can either be empty, contain a successful contention, or a collision.

As in [3]–[6], we assume that a station i obtains its local channel conditions after a successful contention. Let $R_i(\theta)$ denote the corresponding transmission rate at time θ . If $R_i(\theta)$ is small (indicating a poor channel), station i gives up this transmission opportunity and lets all the stations contend for the channel again. Otherwise, it transmits for a duration of \mathcal{T} . Fig. 1 depicts an example of such channel contention. Our model, like that of [3]–[6], assumes that $R_i(\theta)$ remains constant for the duration of a data transmission and that different observations of $R_i(\theta)$ are independent.² From [3], we have that the optimal transmission policy is a threshold policy: for a given threshold \bar{R}_i , station i only transmits after a successful contention if $R_i(\theta) \geq \bar{R}_i$.

B. Throughput analysis

The throughput r_i achieved by station i is a function of the parameters p_i and \bar{R}_i . Let l_i be the average number of bits that station i transmits following a successful contention and

¹The notion of *proportional fairness* was originally proposed by F. Kelly [9] and has been later applied to opportunistic scheduling [2]. According to [9], an allocation $\{r_1, \dots, r_N\}$ is *proportionally fair* if (i) it is feasible, and (ii) for any other feasible allocation, the aggregate of proportional changes is zero or negative. [9] shows that the proportionally fair allocation maximizes $\sum_i \log(r_i)$.

²The assumption that $R_i(\theta)$ remains constant during a data transmission is a standard assumption for the block-fading channel in wireless communications [10], [11]. The assumption that different observations are independent is justified in [3] through numerical calculations which show that in many practical scenarios the channel correlation between two adjacent successful contentions of a station is very small with a very high probability.

T_i be the average time it holds the channel (including the time spent in contention). Then, the throughput of station i is

$$r_i = \frac{p_{s,i} l_i}{\sum_j p_{s,j} T_j + (1 - p_s) \tau} \quad (1)$$

where $p_{s,i}$ is the probability that a mini slot contains a successful contention of station i ,

$$p_{s,i} = p_i \prod_{j \neq i} (1 - p_j) \quad (2)$$

and p_s is the probability that the minislot contains a successful contention of any station in the system

$$p_s = \sum_i p_{s,i}. \quad (3)$$

Both l_i and T_i depend on \bar{R}_i . When a station contends successfully, it holds the channel for a time $\mathcal{T} + \tau$ if it transmits data and τ if it gives up the transmission opportunity. Thus, T_i can be computed as

$$T_i = \text{Prob}(R_i(\theta) < \bar{R}_i) \tau + \text{Prob}(R_i(\theta) \geq \bar{R}_i) (\mathcal{T} + \tau). \quad (4)$$

When the station uses the transmission opportunity, it transmits a number of bits given by $R_i(\theta) T_i$, which yields

$$l_i = \int_{\bar{R}_i}^{\infty} r T_i f_{R_i}(r) dr \quad (5)$$

where $f_{R_i}(r)$ is the pdf of $R_i(\theta)$.

Based on the above, we can compute r_i from $\mathbf{p} = \{p_1, \dots, p_N\}$ and $\bar{\mathbf{R}} = \{\bar{R}_1, \dots, \bar{R}_N\}$. In the following, we obtain the optimal configuration of these parameters to achieve proportional fairness.

C. Optimal p_i configuration

The problem of determining the configuration that provides proportional fairness can be formulated as the unconstrained optimization problem of finding the \mathbf{p} and $\bar{\mathbf{R}}$ configuration that maximizes $\sum_i \log(r_i)$. We start by computing the optimal configuration of \mathbf{p} . Let us define w_i as

$$w_i = \frac{p_{s,i}}{p_{s,1}} \quad (6)$$

where we take station 1 as reference. From the above equation we have that $p_{s,i} = w_i p_s / \sum_j w_j$. Substituting this into (1) yields

$$r_i = \frac{w_i p_s l_i}{\sum_j w_j p_s T_j + \sum_j w_j (1 - p_s) \tau}. \quad (7)$$

Following the results of [12], we approximate the optimal success probability p_s by $\mathcal{P} = (1 - 1/N)^{N-1}$, which is more accurate than the well-known approximation of $1/e$ for slotted random access.³ The numerical results provided in Section VI-A confirm the accuracy of the approximation. Substituting p_s by the constant value \mathcal{P} in (7) gives

$$r_i = \frac{w_i l_i}{\sum_j w_j T_j + \sum_j w_j (1/\mathcal{P} - 1) \tau}. \quad (8)$$

³[13] shows that the approximation of $1/e$ holds both for symmetric and asymmetric access probabilities, and it further shows that this approximation is accurate as long as the number of stations is sufficiently large and the access probabilities are sufficiently small.

The problem of determining the optimal \mathbf{p} configuration is equivalent to finding the w_i values that maximize $\sum_i \log(r_i)$, for r_i defined in (8). To obtain these w_i values, we impose

$$\frac{\partial \sum_j \log(r_j)}{\partial w_i} = 0 \quad (9)$$

which yields

$$\frac{1}{w_i} - N \frac{T_i + (1/\mathcal{P} - 1) \tau}{\sum_j w_j T_j + \sum_j w_j (1/\mathcal{P} - 1) \tau} = 0. \quad (10)$$

Combining this expression for w_i and w_j , we obtain

$$\frac{w_i}{w_j} = \frac{T_j + (1/\mathcal{P} - 1) \tau}{T_i + (1/\mathcal{P} - 1) \tau}. \quad (11)$$

From the above, the values of \mathbf{p} that solve the optimization problem are those that satisfy both $p_s = \mathcal{P}$ and (11). These values can be obtained by solving the following system of equations:

$$\sum_i p_i \prod_{j \neq i} (1 - p_j) = \mathcal{P} \quad (12)$$

$$\frac{p_i \prod_{j \neq i} (1 - p_j)}{p_1 \prod_{j \neq 1} (1 - p_j)} = \frac{T_i + \tau(1/\mathcal{P} - 1)}{T_1 + \tau(1/\mathcal{P} - 1)}, \quad i = 2, \dots, N. \quad (13)$$

As \mathcal{P} is only an approximation to the optimal p_s , the above system of equations has in fact two solutions. This can be seen as follows. From (13) we can express $\{p_i\}_{i=2, \dots, N}$ as a function of p_1 . With this, (12) becomes an equation with only one unknown (p_1). The left-hand side of this equation increases from 0 (for $p_1 = 0$) to a maximum value that is greater than \mathcal{P} and then decreases to 0 (for $p_1 = 1$). Hence, there are two distinct values of p_1 that solve (12). Taking these two values of p_1 and computing the corresponding values of $\{p_i\}_{i=2, \dots, N}$ in each case, we obtain the two solutions of the system of equations. For one of the solutions, all of the access probabilities are larger than the corresponding ones from the other solution; we select the solution with the larger access probabilities. As an exception to this, when all access probabilities are equal, the optimal p_s is exactly \mathcal{P} and the system has only one solution; in this case, we select this unique solution. We denote the selected solution by $\mathbf{p}^* = \{p_1^*, \dots, p_N^*\}$, and refer to these probabilities as the *optimal access probabilities*.

To determine \mathbf{p}^* above, the T_i values have to be computed for all stations. These depend on the optimal configuration of the thresholds $\bar{\mathbf{R}}$. In the following section, we compute the optimal $\bar{\mathbf{R}}$, which we denote by $\bar{\mathbf{R}}^* = \{\bar{R}_1^*, \dots, \bar{R}_N^*\}$.

D. Optimal \bar{R}_i configuration

In order to obtain the optimal configuration of $\bar{\mathbf{R}}$, we need to find the transmission threshold of each station that, given the \mathbf{p}^* computed above, optimizes the overall performance in terms of proportional fairness. This is given by the following theorem.

Theorem 1: Consider a station k that is alone in the network and contends for the channel with $p_k = \mathcal{P}$. Let \bar{R}_k^1 be the transmission rate threshold that optimizes the throughput

of this station under the assumption that different channel observations are independent. Then, $\bar{R}_k^* = \bar{R}_k^1$.

Proof: The proof is by contradiction. Assume there exists a configuration $\bar{\mathbf{R}}^*$ with $\bar{R}_k^* \neq \bar{R}_k^1$ for some station k that provides proportional fairness.

Let l_k^1 and T_k^1 be the values of l_k and T_k for the threshold \bar{R}_k^1 and l_k^* and T_k^* the corresponding values for \bar{R}_k^* . Since \bar{R}_k^1 maximizes r_k when station k is alone:

$$\frac{l_k^1}{T_k^1 + (1/\mathcal{P} - 1)\tau} > \frac{l_k^*}{T_k^* + (1/\mathcal{P} - 1)\tau}. \quad (14)$$

Consider a network with N stations that use configuration $\bar{\mathbf{R}}^*$. Given $\bar{\mathbf{R}}^*$, the \mathbf{p}^* that maximizes $\sum_i \log(r_i)$ is given by (12) and (13). This results in the following throughput for station k :

$$r_k^* = \frac{p_{s,k}^* l_k^*}{\sum_j p_{s,j}^* (T_j^* + (1/\mathcal{P} - 1)\tau)} = \frac{l_k^*}{N(T_k^* + (1/\mathcal{P} - 1)\tau)} \quad (15)$$

and for the other stations:

$$r_i^* = \frac{l_i^*}{N(T_i^* + (1/\mathcal{P} - 1)\tau)} \quad \forall i \neq k. \quad (16)$$

Let us now consider the alternative configuration \bar{R}_k^1 for station k and \bar{R}_i^* for the other stations. If we take the p_k^1 and p_i^1 configuration that satisfies (12) and (13) with this alternative configuration, we obtain the following throughput for station k :

$$r_k^1 = \frac{l_k^1}{N(T_k^1 + (1/\mathcal{P} - 1)\tau)} > r_k^* \quad (17)$$

and for the other stations:

$$r_i^* = \frac{l_i^*}{N(T_i^* + (1/\mathcal{P} - 1)\tau)} \quad \forall i \neq k. \quad (18)$$

With the above, we have found an alternative configuration that provides a higher throughput to station k and the same throughput to all other stations. This alternative configuration thus increases $\sum_i \log(r_i)$, which contradicts the initial assumption that the configuration $\bar{\mathbf{R}}^*$ provides proportional fairness. ■

Following the above theorem, the optimal configuration of the thresholds $\bar{\mathbf{R}}^*$ can be computed using *optimal stopping theory*. This is done in [3], which finds that the optimal threshold \bar{R}_i^* can be obtained by solving the following fixed point equation:

$$\mathbb{E} (R_i(\theta) - \bar{R}_i^*)^+ = \frac{\bar{R}_i^* \tau}{T\mathcal{P}}. \quad (19)$$

The above concludes the search for the optimal configuration. The key advantage of this configuration is that it allows each station to compute its \bar{R}_i^* based on *local information* only, and thus decouples the computation of \bar{R}_i^* from that of p_i^* . We use this finding to design a distributed mechanism for computing the optimal configuration, where each station uses a fixed $\bar{R}_i = \bar{R}_i^*$ obtained locally, together with an adaptive algorithm to determine the optimal p_i^* .

III. DOC ALGORITHM

In this section we propose an adaptive algorithm that satisfies the following properties: (i) when all stations implement the algorithm, it leads to the optimal configuration computed above, and (ii) a selfish station cannot obtain any gain by deviating from the algorithm. We first motivate our algorithm by showing that, in the absence of punishments, the system will naturally tend to a highly undesirable point of operation. We then present our algorithm, which uses punishments to drive the system to the optimal point of operation derived in the previous section.

A. Motivation

If no constraints are imposed on the wireless network and stations are allowed to configure their $\{p_i, \bar{R}_i\}$ parameters to selfishly maximize their own benefit, the network will not naturally tend to the optimum configuration derived above. To show this, we model the wireless system as a static game in which each station can choose its configuration without suffering any penalty. The following theorem characterizes the Nash equilibria of this game.

Theorem 2: In the absence of penalties, there is at least one station that plays $p_i = 1$ in any Nash equilibrium.

Proof: The proof is by contradiction. Let us assume that there is a Nash equilibrium such that $p_j \neq 1 \forall j$.

If we consider one player i and take the partial derivative of its throughput r_i , we obtain

$$\frac{\partial r_i}{\partial p_i} = \frac{\prod_{j \neq i} (1 - p_j) l_i \hat{T}_{-i}}{(p_i \hat{T}_i + (1 - p_i) \hat{T}_{-i})^2} > 0 \quad (20)$$

where \hat{T}_i is the average duration during which the channel is occupied when station i transmits and \hat{T}_{-i} is the average duration of a transmission or an empty mini slot when station i does not transmit.

From the above, it can be seen that the throughput r_i is a strictly increasing function of p_i . It follows from this that $\{p_i, \bar{R}_i\}$, with $p_i \neq 1$, is not the best strategy for player i given the configuration of the other stations, since station i could obtain a higher throughput by increasing p_i to 1 and using the same \bar{R}_i . The configuration $\{p_i, \bar{R}_i\}$, with $p_i \neq 1$, is therefore not a Nash equilibrium, which contradicts our initial assumption. ■

Any of the above Nash equilibria are highly undesirable. If station i is the only one that plays $p_i = 1$, then player i achieves non-zero throughput while all other players have zero throughput. Conversely, if some other station j also plays $p_j = 1$, the result is a network collapse with all players obtaining zero throughput.

We conclude from the above that, in the absence of punishments, selfish behavior will severely degrade the performance of the wireless system. In the following, we propose an algorithm that addresses this problem by implementing a distributed punishment mechanism.

B. Rationale behind the algorithm

Before presenting the algorithm, we first discuss the rationale behind its design. This rationale relies heavily on the notion of *channel time* that a station obtains over a certain interval Θ , defined as

$$t_i(\Theta) = \sum_{k=1}^{n_i(\Theta)} (T_i^k(\Theta) + (1/\mathcal{P} - 1)\tau) \quad (21)$$

where $n_i(\Theta)$ is the number of successful contentions of station i in that period and $T_i^k(\Theta)$ is the duration of the k^{th} successful contention of the station in the interval. The above definition comprises the aggregate transmission time of the station plus a fixed overhead of $(1/\mathcal{P} - 1)\tau$ that is added every time the station accesses the channel.

An important observation that drives the design of our algorithm is that, with the configuration of Section II, all stations receive the same channel time on average, i.e., $t_i = t_j \forall i, j$ (where $t_i = \mathbb{E}[t_i(\Theta)]$). This can be seen as follows. From (21) we have

$$\begin{aligned} \frac{t_i}{t_j} &= \frac{\mathbb{E}[n_i(\Theta)] (\mathbb{E}[T_i^k(\Theta)] + (1/\mathcal{P} - 1)\tau)}{\mathbb{E}[n_j(\Theta)] (\mathbb{E}[T_j^k(\Theta)] + (1/\mathcal{P} - 1)\tau)} \\ &= \frac{p_{s,i}(T_i + (1/\mathcal{P} - 1)\tau)}{p_{s,j}(T_j + (1/\mathcal{P} - 1)\tau)}. \end{aligned} \quad (22)$$

Furthermore, from (13) we have $p_{s,i}(T_i + (1/\mathcal{P} - 1)\tau) = p_{s,j}(T_j + (1/\mathcal{P} - 1)\tau)$ and thus $t_i = t_j$.

When all stations use the optimal configuration, the overhead in the definition of channel time, $(1/\mathcal{P} - 1)\tau$, coincides with the average time between two successes. As a result, for an interval Θ of duration T_{interval} it holds that $\sum_i \mathbb{E}[t_i(\Theta)] = T_{\text{interval}}$. From this and $t_i = t_j$, we have that with the optimal configuration, all stations receive an average channel time of

$$\mathbb{E}[t_i(\Theta)] = T_{\text{interval}}/N \quad \forall i. \quad (23)$$

We define this average channel time as the *optimal channel time* and denote it by t^* (i.e., $t^* = T_{\text{interval}}/N$).

The last observation upon which our algorithm relies is that as long as a selfish station does not receive more channel time than t^* , it cannot increase its throughput. The throughput of a station with a given channel time and \bar{R}_i is equal to the throughput it would obtain if it were alone in the channel during this time with $p_i = \mathcal{P}$ and the same \bar{R}_i . From Theorem 1, we have that this throughput is maximized for the optimal transmission rate threshold \bar{R}_i^* . Therefore, as long as the station does not receive extra channel time, it will not be able to achieve a higher throughput.

Given these observations, we base our algorithm on the following principles: (i) if a given station i detects that another station k is receiving more channel time than itself, it considers station k to be selfish and indirectly punishes it by using a more aggressive configuration, and (ii) when punishing station k , the punishment needs to be severe enough to keep station k 's channel time below t^* so that station k does not benefit from misbehaving.

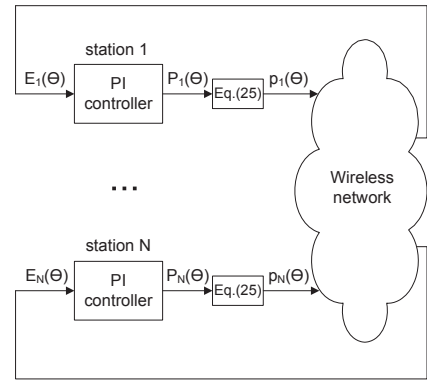


Fig. 2. DOC control system.

C. Algorithm design

The objective of DOC is to drive the system to the optimal configuration $\{\mathbf{p}^*, \bar{\mathbf{R}}^*\}$ obtained in Section II. As discussed in Section II-D, each station can locally compute its optimal configuration of \bar{R}_i independently of the configuration of the other stations. Therefore, with DOC each station maintains a fixed \bar{R}_i (equal to the optimal value) and implements an adaptive algorithm to configure its access probability p_i .

Time is divided into intervals of fixed length T_{interval} , and each station updates its access probability p_i at the beginning of every interval. We use the discrete variable Θ to refer to the different intervals, and $p_i(\Theta)$ to denote the value of p_i in a given interval Θ . The central idea behind DOC is that when a misbehaving station is detected, the other stations increase their access probabilities in subsequent intervals to prevent the selfish station from benefiting from its misbehavior.

A key challenge in DOC is to determine the appropriate reaction against a selfish station. If the reaction is not severe enough, a selfish station may benefit from misbehaving. However, if the reaction is too severe, the system may become unstable by entering an endless loop where all stations indefinitely increase their p_i to punish each other.

Control theory is a particularly suitable tool to address this challenge, as it helps guarantee the convergence and stability of adaptive algorithms. We use techniques from *multi-variable control theory* [14] for the design of the DOC algorithm. The algorithm is based on the classic system illustrated in Fig. 2, where each station runs an independent controller to compute its configuration. The controller that we have chosen for this paper is a *proportional-integral* (PI) controller, a well-known controller from classic control theory.

As shown in the figure, the PI controller of station i takes as input the error signal measured over an interval Θ , $E_i(\Theta)$, and provides as output the control signal $P_i(\Theta)$ for the next interval. The error signal indicates how far the system is from the desired point of operation. If the system is operating as desired, the error signals of all stations are zero; otherwise, the error signals are non-zero and the state of the system needs to change from its current point of operation to the desired one. To do this, the PI controller adjusts the control signal $P_i(\Theta)$, increasing it if $E_i(\Theta) > 0$ and decreasing it otherwise. In the following, we address the design of $P_i(\Theta)$ and $E_i(\Theta)$.

D. Control signal P_i

The DOC algorithm needs to adjust the access probability $p_i(\Theta)$ based on the control signal. To do this, there needs to be a one-to-one mapping between the control signal $P_i(\Theta)$ given by the controller and $p_i(\Theta)$. In addition, we design the system such that the $P_i(\Theta)$ values are the same for all stations at the optimal point of operation. This latter requirement is necessary to derive the conditions for stability in Section IV.

Based on the above requirements, we design $P_i(\Theta)$ as

$$P_i(\Theta) = \frac{p_i(\Theta)}{1 - p_i(\Theta)} (T_i + (1/\mathcal{P} - 1)\tau). \quad (24)$$

A station can therefore compute its $p_i(\Theta)$ from the control signal $P_i(\Theta)$ as

$$p_i(\Theta) = \frac{P_i(\Theta)}{T_i + (1/\mathcal{P} - 1)\tau + P_i(\Theta)}. \quad (25)$$

E. Error signal E_i

The design of the error signal $E_i(\Theta)$ has the following two goals: (i) selfish stations should not be able to obtain extra channel time from the wireless network by using a configuration different from the optimal, and (ii) as long as there are no selfish stations, $\mathbf{p}(\Theta)$ should converge to the optimal \mathbf{p}^* .

For the design of the error signal, DOC relies (like [15], [16]) on the broadcast nature of the wireless medium, which enables stations to overhear the transmissions of the other stations. In particular, in every interval Θ , each station measures (i) the channel time used by the other stations, $t_j(\Theta)$, and (ii) the average time (over the interval) that they hold the channel upon a successful contention, $T_j(\Theta) = \sum_{k=1}^{n_j(\Theta)} T_j^k(\Theta)/n_j(\Theta)$. Based on this, station i computes the error signal at the end of the interval as

$$E_i(\Theta) = \sum_{j \neq i} (t_j(\Theta) - t_i(\Theta)) - F_i(\Theta) \quad (26)$$

where $F_i(\Theta)$ is a function that we design below. The error signal $E_i(\Theta)$ consists of the following two components:

- The first component, $\sum_{j \neq i} t_j(\Theta) - t_i(\Theta)$, punishes selfish stations. If a station i receives less channel time than the other stations, this component will be positive and hence station i will increase its access probability $p_i(\Theta)$.
- The second component, $F_i(\Theta)$, drives the system to the desired point of operation in the absence of selfish behavior (i.e., when all stations receive the same channel time).

We next address the design of the function $F_i(\Theta)$. In order to drive $\mathbf{p}(\Theta)$ to the desired \mathbf{p}^* when all stations receive the same channel time, we need $F_i(\Theta) > 0$ for $p_i(\Theta) > p_i^*$, such that in this case $p_i(\Theta)$ decreases, and $F_i(\Theta) < 0$ for $p_i(\Theta) < p_i^*$.

The design of $F_i(\Theta)$ should also prevent selfish stations from obtaining more channel time than t^* . In the following, we derive the conditions that $F_i(\Theta)$ needs to meet in order to satisfy this requirement. To derive these conditions, we assume that the system is in steady state, which implies that selfish

stations play with a static configuration. (In the analysis of Section V we show that DOC is also effective against selfish strategies that change the configuration over time.)

We first consider the case where one station k is selfish and all others are well-behaved and run the DOC algorithm. Since the PI controller drives the error signal $E_i(\Theta)$ to 0 in steady state, the following holds for all well-behaved stations:

$$F_i(\Theta) = \sum_{j \neq i} t_j(\Theta) - t_i(\Theta). \quad (27)$$

Summing $F_i(\Theta)$ over all stations except the selfish one yields:

$$\sum_{i \neq k} F_i(\Theta) = (N-1)t_k(\Theta) - \sum_{i \neq k} t_i(\Theta) = Nt_k(\Theta) - \sum_i t_i(\Theta). \quad (28)$$

If we combine the above with the requirement that the selfish station cannot gain, i.e., $t_k(\Theta) \leq t^*$, we obtain the following inequality,

$$\sum_{i \neq k} F_i(\Theta) \leq D(\Theta) \quad (29)$$

where $D(\Theta)$ is defined as the difference between the sum of channel times in optimal operation and the sum of channel times in the current interval, i.e., $D(\Theta) = Nt^* - \sum_i t_i(\Theta)$. Note that, if the current access probabilities are not optimal, $\sum_i t_i(\Theta)$ will be smaller than Nt^* . Hence, $D(\Theta)$ reflects the channel time lost due to non-optimal access probabilities.

The following upper bound on $F_i(\Theta)$ guarantees that (29) is satisfied, and thus ensures that a selfish station does not benefit from misbehaving:

$$F_i(\Theta) \leq \frac{1}{N-1} D(\Theta). \quad (30)$$

The intuition behind this upper bound is as follows. When a selfish station misbehaves, it receives more channel time than the well-behaved stations. This, however, moves the point of operation away from the optimal access probabilities, reducing the overall efficiency in terms of aggregate channel time. The above upper bound ensures that the additional channel time received by the selfish station does not outweigh the channel time it loses due to the overall loss of aggregate channel time. This guarantees that the selfish station does not receive more channel time and hence does not benefit from misbehaving.

We next consider the case of multiple selfish stations. In this case, the aggregate channel time received by the selfish stations must not exceed the aggregate channel time that they would receive in optimal operation, i.e., $\sum_{i=1}^m t_i(\Theta) \leq mt^*$ (where $\{1, \dots, m\}$ is the set of selfish stations). Following similar reasoning to that above, we obtain the upper bound

$$F_i(\Theta) \leq \frac{m}{N-m} D(\Theta). \quad (31)$$

Given all the above requirements, we design $F_i(\Theta)$ as:

$$F_i(\Theta) = \begin{cases} \min\left((N-1)D(\Theta), \frac{D(\Theta)}{N}\right), & p_i(\Theta) > p_i^{min} \\ \min\left((N-1)D(\Theta), -\frac{D(\Theta)}{N}, (N-1)\Delta\right), & p_i(\Theta) \leq p_i^{min} \end{cases} \quad (32)$$




Fig. 3. F_i as a function of $p_i(\Theta)$ when $t_i = t_j \forall i, j$.

where $\mathbf{p}^{\min} = \{p_1^{\min}, \dots, p_N^{\min}\}$ are the access probabilities that minimize $D = \mathbb{E}[D(\Theta)]$ subject to $t_i = t_j \forall i, j$, and Δ is the value that D takes at this point,

$$\Delta = D|_{\mathbf{p}=\mathbf{p}^{\min}}. \quad (33)$$

In order to compute \mathbf{p}^{\min} and Δ , the T_j of all stations are required. For these, we use the $T_j(\Theta)$ values measured over the current interval.

The above design satisfies all of our previous requirements:

- The term $D(\Theta)/N$ ensures that (30) and (31) are satisfied when $D(\Theta) > 0$ and the term $(N-1)D(\Theta)$ ensures that they are satisfied when $D(\Theta) < 0$. This provides the required protection against (one or more) selfish stations.
- As illustrated in Fig. 3, when all stations have the same expected channel time, the expected value of $F_i(\Theta)$ is positive for $p_i(\Theta) > p_i^*$ and negative otherwise. This ensures that $\mathbf{p}(\Theta)$ is driven to the desired \mathbf{p}^* .

The above design of the DOC algorithm is based on the assumption that the number of stations in the wireless network is fixed. In the following, we address the case of stations joining and leaving the network. With DOC, each station only keeps the state maintained by the PI controller, $\sum_{\Theta} \sum_{j \neq i} (t_j(\Theta) - t_i(\Theta)) + F_i(\Theta)$, which accounts for the deficit or surplus of the station's channel time over the other stations in the network. When a new station joins the wireless network, this station does not have a surplus or deficit, and therefore the other stations keep their state. The new station initializes the state of its PI controller such that its initial p_i corresponds to the optimal p_i^* . When a station leaves, the remaining stations keep their state: this ensures that the deficit accumulated by a selfish station is not reset if it leaves and rejoins the network.

This concludes the design of the algorithm. In the following two sections, we analytically evaluate its performance when all stations are well-behaved (Section IV) and when some stations misbehave (Section V).

IV. DOC ANALYSIS

In this section we analyze the performance of DOC when all stations are well-behaved. As stations do not obtain any benefit from misbehaving, it is to be expected that they will all play DOC, and therefore this is the most meaningful scenario for the performance analysis of the system. We first analyze the wireless system under steady state conditions and show that it is driven to the desired point of operation obtained in

Section II. We then conduct a transient analysis and derive sufficient conditions for stability.

A. Steady state analysis

Our analysis is based on the system model of Fig. 4. In this model, C represents the function implemented by the controllers, which computes the control signals $P_i(\Theta)$, taking the error signals $E_i(\Theta)$ as input. H represents the wireless system which provides the error signals $E_i(\Theta)$ based on the control signals $P_i(\Theta)$. In line with standard control theory [17], we model the randomness of the channel with the noise signals $W_i(\Theta)$ and let $E_i(\Theta)$ represent the expected value of the error signal for the given control signals $P_i(\Theta)$. Since the controller includes an integrator, there is no steady state error [17] and the steady state solution can be obtained from

$$E_i(\Theta) = 0 \quad \forall i. \quad (34)$$

Using (26) and (32), $E_i(\Theta)$ can be computed from $\mathbf{p}(\Theta)$. This enables (34) to be expressed as a system of equations in $\mathbf{p}(\Theta)$. The following theorem guarantees that the the solution of this system of equations is unique and shows that the unique stable point in steady state is the desired point of operation from Section II.

Theorem 3: The unique stable point of operation of the system in steady state is $\mathbf{p}(\Theta) = \mathbf{p}^*$.

Proof: Let us consider two stations i and j . From (34) we have $E_i(\Theta) - E_j(\Theta) = 0$, which yields

$$Nt_j(\Theta) + F_j(\Theta) - Nt_i(\Theta) - F_i(\Theta) = 0. \quad (35)$$

Note that $t_j(\Theta) > t_i(\Theta)$ implies $F_j(\Theta) \geq F_i(\Theta)$, and vice versa. This can be seen as follows: If $p_j(\Theta) > p_j^{\min}$ and $p_i(\Theta) > p_i^{\min}$, then $F_j(\Theta) = F_i(\Theta)$. If $p_j(\Theta) \leq p_j^{\min}$ and $p_i(\Theta) \leq p_i^{\min}$, then also $F_j(\Theta) = F_i(\Theta)$. If $p_j(\Theta) > p_j^{\min}$ and $p_i(\Theta) \leq p_i^{\min}$, then $F_j(\Theta) \geq F_i(\Theta)$. When $t_j(\Theta) > t_i(\Theta)$, we are in one of these three cases, and hence $F_j(\Theta) \geq F_i(\Theta)$. Combining this with (35) yields $t_i(\Theta) = t_j(\Theta) \forall i, j$. Substituting this into $E_i(\Theta) = 0$ yields $F_i(\Theta) = 0$. Given $t_i(\Theta) = t_j(\Theta)$, $F_i(\Theta)$ is an increasing function of $p_i(\Theta)$ that crosses 0 at $p_i(\Theta) = p_i^*$. Hence, the only $p_i(\Theta)$ that satisfies $F_i(\Theta) = 0$ is p_i^* . Since this holds for all i , the unique stable point of operation is $p_i(\Theta) = p_i^* \forall i$. ■

B. Stability analysis

We now conduct a stability analysis of DOC to configure the parameters of the PI controller. According to the definition of a PI controller [17], station i computes the value of P_i at interval Θ' as a function of the error values measured by the station in the current and previous intervals based on the following equation:

$$P_i(\Theta') = K_p E_i(\Theta') + K_i \sum_{\Theta=0}^{\Theta'-1} E_i(\Theta) \quad (36)$$

where K_p and K_i are the parameters of the controller that we need to configure.

In order to analyze our system from a control theoretic standpoint, we need to characterize the transfer functions C




Fig. 4. Control system.

and H in the system model of Fig. 4. The control and error signals in the figure are given by the following vectors in the z -domain [17]:

$$\mathbf{P}(z) = (P_1(z), \dots, P_N(z))^T \quad (37)$$

and

$$\mathbf{E}(z) = (E_1(z), \dots, E_N(z))^T. \quad (38)$$

Our control system consists of one PI controller in each station i that takes $E_i(z)$ as input and provides $P_i(z)$ as output. We can therefore express the relationship between $\mathbf{E}(z)$ and $\mathbf{P}(z)$ as follows

$$\mathbf{P}(z) = C \cdot \mathbf{E}(z) \quad (39)$$

where

$$C = \begin{pmatrix} C_{PI}(z) & 0 & 0 & \dots & 0 \\ 0 & C_{PI}(z) & 0 & \dots & 0 \\ 0 & 0 & C_{PI}(z) & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & C_{PI}(z) \end{pmatrix} \quad (40)$$

with $C_{PI}(z)$ being the z -transform of a PI controller [17],

$$C_{PI}(z) = K_p + \frac{K_i}{z-1}. \quad (41)$$

In order to characterize our wireless system with a transfer function H that takes $\mathbf{P}(z)$ as input and has $\mathbf{E}(z)$ as output, we proceed as follows. Equation (26) provides a nonlinear relationship between $\mathbf{E}(\Theta)$ and $\mathbf{P}(\Theta)$. To express this relationship as a transfer function, we linearize it at the optimal point of operation.⁴ We then study the linearized model and ensure its stability through appropriate choice of parameters. Note that the stability of the linearized model guarantees that our system is locally stable.⁵

We express the perturbations around the stable point of operation as follows:

$$\mathbf{P}(\Theta) = \mathbf{P}^* + \delta\mathbf{P}(\Theta) \quad (42)$$

where \mathbf{P}^* is the stable point of operation as given by (24) with $\mathbf{p}(\Theta) = \mathbf{p}^*$.

With the above, the perturbations of \mathbf{E} can be approximated by

$$\delta\mathbf{E}(\Theta) = H \cdot \delta\mathbf{P}(\Theta) \quad (43)$$

⁴This linearization provides a good approximation of the behavior of the system when it suffers small perturbations around the stable point of operation.

⁵A similar approach was used in [18] to analyze RED from a control theoretic standpoint.

where

$$H = \begin{pmatrix} \frac{\partial E_1(\Theta)}{\partial P_1(\Theta)} & \frac{\partial E_1(\Theta)}{\partial P_2(\Theta)} & \dots & \frac{\partial E_1(\Theta)}{\partial P_N(\Theta)} \\ \frac{\partial E_2(\Theta)}{\partial P_1(\Theta)} & \frac{\partial E_2(\Theta)}{\partial P_2(\Theta)} & \dots & \frac{\partial E_2(\Theta)}{\partial P_N(\Theta)} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial E_N(\Theta)}{\partial P_1(\Theta)} & \frac{\partial E_N(\Theta)}{\partial P_2(\Theta)} & \dots & \frac{\partial E_N(\Theta)}{\partial P_N(\Theta)} \end{pmatrix}. \quad (44)$$

To compute these partial derivatives we proceed as follows. The error signal $E_i(\Theta)$ can be expressed as

$$E_i(\Theta) = T_{interval} \sum_{j \neq i} \left(\frac{p_{s,j}(\Theta) (T_j + (\frac{1}{P} - 1) \tau)}{\sum_k p_{s,k}(\Theta) T_k + (1 - p_s(\Theta)) \tau} - \frac{p_{s,i}(\Theta) (T_i + (\frac{1}{P} - 1) \tau)}{\sum_k p_{s,k}(\Theta) T_k + (1 - p_s(\Theta)) \tau} \right) - F_i(\Theta). \quad (45)$$

The above can be rewritten as a function of $\mathbf{P}(\Theta)$ given by

$$E_i(\Theta) = T_{interval} \frac{\sum_{j \neq i} (P_j(\Theta) - P_i(\Theta))}{\sum_j P_j(\Theta) - \frac{p_s(\Theta)}{p_e(\Theta)} (\frac{1}{P} - 1) \tau + \frac{1 - p_s(\Theta)}{p_e(\Theta)} \tau} - F_i(\Theta) \quad (46)$$

where $p_e(\Theta) = \prod_j (1 - p_j(\Theta))$.

We start by showing that $\partial F_i(\Theta) / \partial P_i(\Theta) = 0$ at the stable point of operation. It follows from (32) that

$$\frac{\partial F_i(\Theta)}{\partial P_i(\Theta)} = 0 \iff \frac{\partial D(\Theta)}{\partial P_i(\Theta)} = 0. \quad (47)$$

$D(\Theta)$ can be expressed as

$$D(\Theta) = Nt^* - T_{interval} \frac{\sum_i p_{s,i}(\Theta) T_i + p_s(\Theta) (1/P - 1) \tau}{\sum_i p_{s,i}(\Theta) T_i + (1 - p_s(\Theta)) \tau}. \quad (48)$$

The partial derivative of $D(\Theta)$ can be computed as

$$\frac{\partial D(\Theta)}{\partial P_i(\Theta)} = \frac{\partial D(\Theta)}{\partial p_i(\Theta)} \frac{\partial p_i(\Theta)}{\partial P_i(\Theta)}. \quad (49)$$

Taking the partial derivative of (48) with respect to $p_i(\Theta)$ and evaluating it at the stable point of operation yields

$$\frac{\partial D(\Theta)}{\partial p_i(\Theta)} = T_{interval} \left(\frac{\tau/P}{\sum_i p_{s,i}(\Theta) T_i + (1/P - 1) \tau} \right) \frac{\partial p_s(\Theta)}{\partial p_i(\Theta)}. \quad (50)$$

Since $p_s(\Theta)$ has a maximum value at the stable point of operation, we have that $\partial p_s(\Theta) / \partial p_i(\Theta) = 0$, which yields $\partial D(\Theta) / \partial P_i(\Theta) = 0$ and hence

$$\frac{\partial F_i(\Theta)}{\partial P_i(\Theta)} = 0. \quad (51)$$

The partial derivative of $E_i(\Theta)$ evaluated at the stable point of operation can then be computed from (46) as

$$\left. \frac{\partial E_i(\Theta)}{\partial P_i(\Theta)} \right|_{\mathbf{P}(\Theta) = \mathbf{P}^*} = -(N-1) T_{interval} \frac{1}{\sum_j P_j^*}. \quad (52)$$

Using similar reasoning, we can see that

$$\left. \frac{\partial E_j(\Theta)}{\partial P_j(\Theta)} \right|_{\mathbf{P}(\Theta) = \mathbf{P}^*} = T_{interval} \frac{1}{\sum_j P_j^*}. \quad (53)$$

Substituting these expressions in matrix H yields

$$H = K_H \begin{pmatrix} -(N-1) & 1 & \dots & 1 \\ 1 & -(N-1) & \dots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \dots & -(N-1) \end{pmatrix} \quad (54)$$

where

$$K_H = T_{interval} \frac{1}{\sum_j P_j^*}. \quad (55)$$

Thus, the linearized system is fully characterized by the matrices C and H . The next step is to configure the K_p and K_i parameters. The following theorem provides sufficient conditions which $\{K_p, K_i\}$ must meet to ensure stability:

Theorem 4: The linearized system is guaranteed to be stable as long as K_p and K_i meet the following conditions:

$$K_i < K_p + \frac{1}{NK_H}, \quad K_i > 2K_p - \frac{1}{NK_H}. \quad (56)$$

Proof: The reader is referred to [8] for the proof of the theorem. Since [8] uses the same linearized system as this paper, the proof follows very closely that of [8]. ■

In addition to guaranteeing stability, our goal in the configuration of the $\{K_p, K_i\}$ parameters is to find the right balance between reaction time in transients and oscillations in steady state. To this end, we use the *Ziegler-Nichols* rules [17], which have been designed for this purpose. Following these rules (see [8] for a detailed description), we obtain the configuration:

$$K_p = \frac{0.4}{2NK_H}, \quad K_i = \left(\frac{1}{0.85 \cdot 2} \right) \frac{0.4}{2NK_H}. \quad (57)$$

The stability of the resulting configuration is guaranteed by the following corollary:

Corollary 1: The K_p and K_i configuration given by (57) is stable.

Proof: The proof follows from the fact that the configuration of (57) meets the conditions of Theorem 4. ■

Note that the above control theoretic analysis guarantees that the system will always converge to the desired point of operation regardless of the initial state. This implies that the system remains stable in the presence of any kind of perturbation. Such perturbations include, among others, transient selfish behavior or stations joining and leaving the network.

V. GAME THEORETIC ANALYSIS

In the previous section we have shown that, when all stations implement the DOC algorithm, they all play with $p_i = p_i^*$ and $\bar{R}_i = \bar{R}_i^*$, which leads to the optimal throughput allocation r_i^* obtained in Section II.⁶ In this section we conduct a game theoretic analysis to show that one or more stations cannot obtain any gain by deviating from DOC. In what follows, we say that a station is *honest* or well-behaved when it implements the DOC algorithm to configure its p_i and \bar{R}_i parameters, while we say that it is *selfish* or misbehaving when it plays a

strategy different from DOC to configure these parameters in order to obtain a greater share of wireless resources.

The game theoretic analysis conducted in this section assumes that users are *rational* and want to maximize their own benefit or *utility*, which is given by the throughput. Furthermore, it is reasonable to assume that the game is non-cooperative in that no binding agreements can be reached between the players as to their their future play [16]. The model is based on the theory of *repeated games* [19]. In repeated games, time is divided into stages and a player can take new decisions at each stage based on the observed behavior of the other players in the previous stages. This matches our algorithm, where time is divided into intervals and stations update their configuration at each interval.⁷ Like previous analyses on repeated games [15], [16], we consider an infinitely repeated game, which is a common assumption when the players do not know when the game will end.

A. Single selfish station

While the design of the DOC algorithm in Section III guarantees that a station cannot benefit from playing with a *fixed* selfish configuration, selfish stations might still benefit by *varying* their configuration over time. As an example, let us consider a naive algorithm that only takes into account the stations' behavior in the previous stage. While this algorithm may be effective against a fixed selfish configuration, it could easily be defeated by a selfish station that alternates between a selfish configuration ($p_k = 1, \bar{R}_k = 0$) and an honest one ($p_k = p_k^*, \bar{R}_k = \bar{R}_k^*$) at every other stage. Since this station would play selfish when all the others play honest, it would achieve a significantly higher throughput every other interval, thus benefiting from its misbehavior.

The above example shows that it is important to ensure that a selfish station cannot obtain any gain no matter how it varies its configuration over time. The following theorem confirms the effectiveness of DOC against any (fixed or variable) selfish strategy. The proof of the theorem relies on the integrator component of the PI controller, which keeps track of the aggregate channel time received by all stations and can thus be used to guarantee that this aggregate does not exceed a given amount.

Theorem 5: Let us consider a selfish station that uses a $p_k(\Theta)$ and $\bar{R}_k(\Theta)$ configuration that can vary over time. If all the other stations implement the DOC algorithm, the throughput received by this station will be no larger than r_k^* .

Proof: The PI controller computes P_i at a given interval Θ' according to the following expression:

$$P_i(\Theta') = P_i^{initial} + K_p \left(\sum_{j \neq i} (t_j(\Theta') - t_i(\Theta')) - F_i(\Theta') \right) + K_i \sum_{\Theta=0}^{\Theta'-1} \left(\sum_{j \neq i} (t_j(\Theta) - t_i(\Theta)) - F_i(\Theta) \right). \quad (58)$$

⁷Note that the game theoretic study conducted in Section III-A was based on static games instead of repeated ones. The reason is that we considered a system without penalties where a user does not react to the behavior of other users. Hence, we could model it as a static game where all players only make a single move at the beginning of the game.

⁶Since the throughput allocation $\{r_1^*, \dots, r_N^*\}$ maximizes $\sum_i \log(r_i)$, it is *Pareto optimal*. This follows from the fact that if there existed another feasible allocation that provided all stations with more throughput than r_i^* , this allocation would yield a larger $\sum_i \log(r_i)$.

With the above, $P_i(\Theta')$ stays between 0 and a given maximum value P_i^{max} . If at some point P_i reaches a P_i^{max} value such that $p_i = 1$, this will result in $t_j = 0$ for $j \neq i$ and $F_i > -(N-1)t_i$, which yields $E_i < 0$, and therefore P_i will decrease. Similarly, if at any time P_i reaches 0, then $t_i = 0$ and $F_i \leq 0$, which yields $E_i > 0$, and therefore P_i will increase.

Considering that $0 \leq P_i(\Theta') \leq P_i^{max}$, the above equation can be expressed as

$$\sum_{\Theta=0}^{\Theta'} \left(\sum_{j \neq i} (t_j(\Theta) - t_i(\Theta)) - F_i(\Theta) \right) = B_i \quad (59)$$

where B_i is a bounded value: $B_i = (P_i^{max} - P_i^{initial} + (K_i - K_p)E_i(\Theta'))/K_i$.

Let us consider the case in which there is a selfish station that changes its configuration over time and receives a channel time $t_k(\Theta)$. Equation (59) can be written as

$$\sum_{\Theta} t_k(\Theta) = \sum_{\Theta} \left((N-1)t_i(\Theta) - \sum_{j \neq i, k} t_j(\Theta) + F_i(\Theta) \right) + B_i. \quad (60)$$

Let us now consider a given interval Θ . From (30), we have

$$F_j(\Theta) \leq \frac{1}{N-1} \left(Nt^* - \sum_i t_i(\Theta) \right). \quad (61)$$

Summing the above expression over all $j \neq k$, we have $\sum_i t_i(\Theta) + \sum_{j \neq k} F_j(\Theta) \leq Nt^*$. As this is satisfied for all Θ ,

$$\sum_{\Theta} \left(\sum_i t_i(\Theta) + \sum_{j \neq k} F_j(\Theta) \right) \leq \sum_{\Theta} Nt^*. \quad (62)$$

Furthermore, by summing (60) over all $j \neq k$,

$$(N-1) \sum_{\Theta} t_k(\Theta) = \sum_{j \neq k} \sum_{\Theta} (t_j(\Theta) + F_j(\Theta)) + \sum_{j \neq k} B_j. \quad (63)$$

Adding the above two equations yields $N \sum_{\Theta} t_k(\Theta) \leq N \sum_{\Theta} t^* + \sum_{j \neq k} B_j$. If we consider a long period of time, the constant term $\sum_{j \neq k} B_j$ can be neglected, resulting in $\sum_{\Theta} t_k(\Theta) \leq \sum_{\Theta} t^*$. This means that the selfish station cannot receive more channel time using a selfish strategy than by playing DOC. Following the argument of Section III-B, this implies that it cannot obtain more throughput than it would by playing DOC, i.e., $r_k \leq r_k^*$, which proves the theorem. ■

The above theorem leads to Corollary 2.

Corollary 2: A state in which all stations play DOC (*All-DOC*) is a Nash equilibrium of the game.

Proof: According to Theorem 5, if all stations but one play DOC, the best response of this station is to play DOC as well since it cannot benefit from playing a different strategy. Thus, *All-DOC* is a Nash equilibrium. ■

The above shows that if all stations start playing with no previous history, none of them can benefit by deviating from DOC. In addition to this, in repeated games it is also important to ensure that, if at some point the game has a given history, a selfish station cannot exploit knowledge of this history by playing a different strategy from DOC. The following theorem confirms that *All-DOC* is a Nash equilibrium of any subgame (where a subgame is defined as the game resulting from

starting to play with a certain history). Therefore, a selfish station cannot benefit by deviating from DOC for any previous history of the game.

Theorem 6: All-DOC is a subgame perfect Nash equilibrium of the game.

Proof: Since the proof of Theorem 5 is not dependent on past history and can therefore be applied to any subgame, *All-DOC* is a Nash equilibrium of any subgame. ■

Note that, even though the above analysis assumes a fixed number of stations in the wireless network, it also holds for the case when the number of stations changes over time, as long as these changes occur over sufficiently long periods such that the constant term $\sum_{j \neq k} B_j$ is not significant.

B. Multiple selfish stations

The above results show the effectiveness of DOC against a single selfish station. In the following, we focus on the case of multiple selfish stations.

The following theorem shows that, by following a different strategy from DOC, multiple stations cannot gain any aggregate channel time.

Theorem 7: Let us consider a scenario with m selfish stations. If all other stations play DOC, the selfish stations cannot gain any aggregate channel time.

Proof: Without loss of generality, let us consider that stations $i = \{1, \dots, m\}$ are selfish. Applying a reasoning similar to Theorem 5 leads to $\sum_{i=1}^m \sum_{\Theta} t_i(\Theta) \leq m \sum_{\Theta} t^*$. As the left-hand side of this inequality is the aggregate channel time obtained by the selfish stations, and the right-hand side is the aggregate channel time that they would obtain if they played DOC, the theorem is proven. ■

According to the above theorem, it is possible for a selfish station to obtain some gain, but this will be at the expense of another selfish station that receives less channel time. Corollary 3 follows from this.

Corollary 3: Let us consider a scenario with m selfish stations. If all other stations play DOC and a selfish station k receives a throughput larger than r_k^* , then there exists another selfish station l that receives a throughput smaller than r_l^* .

Proof: If there is some station $k \in \{1, \dots, m\}$ for which $r_k > r_k^*$, this station must necessarily receive more channel time than it would if all stations played DOC. Since (according to Theorem 7) the selfish stations cannot gain any aggregate channel time, there must then exist some other station $l \in \{1, \dots, m\}$ that receives less channel time. For this station, it holds that $r_l < r_l^*$, which proves the corollary. ■

Based on the above, we argue that DOC is effective against multiple selfish stations, since two or more selfish stations cannot *simultaneously* benefit and therefore do not have any incentive to play a coordinated strategy different from DOC.

VI. PERFORMANCE EVALUATION

In this section we evaluate DOC by means of simulation to show that (i) in the absence of selfish stations, DOC provides optimal performance while remaining stable and reacting quickly to changes, and (ii) selfish stations cannot benefit by following a strategy different from DOC. Unless

otherwise stated, we assume that different observations of the channel conditions are independent. We also assume that the available transmission rate for a given SNR is given by the Shannon channel capacity: $R(h) = W \log_2(1 + \rho|h|^2)$ bits/s, where W is the channel bandwidth, ρ is the normalized average SNR and h is the random gain of Rayleigh fading. We implemented the DOC algorithm in OMNET++. In the simulations, we set $W = 10^7$, $\mathcal{T}/\tau = 10$ and the interval of the controller $T_{total} = 10^5\tau$. For all results, 95% confidence intervals are below 0.5%.

A. Validation of the optimal configuration

In order to assess the accuracy of the analysis of Section II, we compare the performance of the configuration computed in that section (*optimal configuration*) against the result of performing an exhaustive search over all possible configurations of $\{\mathbf{p}, \bar{\mathbf{R}}\}$ and selecting the best one (*exhaustive search*). We perform the following two experiments. First, we consider a wireless network with N stations, $N/2$ of them with a normalized SNR equal to ρ_1 and the other $N/2$ with a normalized SNR equal to ρ_2 . Fig. 5 shows the total throughput obtained by both approaches for different numbers of stations ($N = \{2, 4, \dots, 20\}$) and levels of heterogeneity ($\rho_1 = 1$, $\rho_2 = \{1, 5, 20, 100\}$). We observe that both approaches perform very closely, with a difference well below 0.5% in all cases. Next, we consider a wireless network with N stations, $N = \{2, 4, \dots, 20\}$, where the ρ_i of each station is randomly chosen in the range (ρ_1, ρ_2) , for $\rho_1 = 1$ and $\rho_2 = \{1, 5, 20, 100\}$. The results (not shown in a graph) confirm that also here the difference between the two approaches is well below 0.5% in all cases. We conclude from these two experiments that the analysis is very accurate.

The key approximation of the analysis is to assume that p_s is equal to $(1 - 1/N)^{N-1}$, which corresponds to the value of the optimal success probability with symmetric access probabilities [12]. By analyzing the access probabilities in the above experiments, we observe that all stations use similar access probabilities regardless of their channel conditions, which makes this approximation particularly accurate. For instance, even for the extreme case of $N = 2$ and $\rho_2 = 100$, the difference between the access probabilities of the two stations is only around 18%. This is a consequence of (11), whereby the ratio of the access probabilities of two stations depends on their T_i values; since the thresholds \bar{R}_i are set so that all stations have a similar probability of using a transmission opportunity, this implies that all T_i values are similar, and as a result, the access probabilities are also similar.

B. Throughput evaluation

For the throughput evaluation, we compare the performance of DOC to the following approaches: (i) the static optimal configuration obtained in Section II (*optimal configuration*), (ii) the configuration proposed in [3] (*DOS*), and (iii) an approach that does not perform opportunistic scheduling but always transmits after successful contention (*non-opportunistic*). We consider a scenario with $N = 10$ stations, half of them with a normalized SNR of $\rho_1 = 1$ and the other

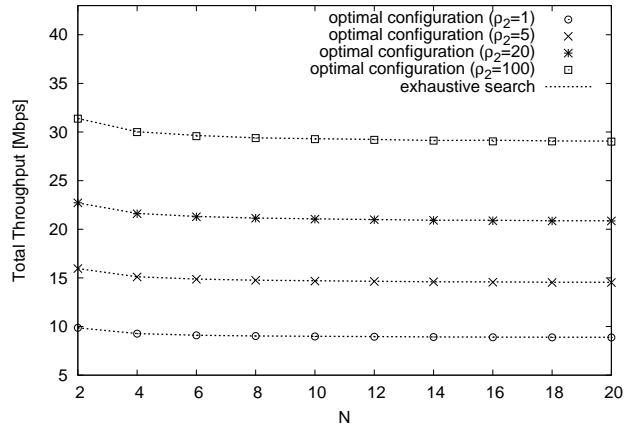


Fig. 5. Validation of the optimal configuration of Section II for different numbers of stations and levels of heterogeneity.




Fig. 6. Proportional fairness as a function of SNR ($\rho_1 = 1, 1 \leq \rho_2 \leq 10$).

half with a normalized SNR ρ_2 that varies from 1 to 10. Fig. 6 shows the proportional fairness metric, $\sum_i \log(r_i)$, as a function of ρ_2 . We observe that DOC performs at the same level as the benchmark given by the *optimal configuration*, while the other two approaches (*DOS* and *non-opportunistic*) provide a substantially lower performance.

For the above scenario with $\rho_2 = 4$, Fig. 7 depicts the individual throughput allocation of two stations (where r_1 is the throughput of a station with ρ_1 and r_2 that of a station with ρ_2). DOC is effective in driving the system to the optimal point of operation and provides the same throughput as the *optimal configuration*. In contrast, *DOS* exhibits a high degree of unfairness as it provides a much higher throughput to the station with high SNR. The *non-opportunistic* approach provides a reasonable degree of fairness but has lower throughput due to the lack of opportunistic scheduling. In conclusion, the proposed DOC algorithm provides a good tradeoff between overall throughput and fairness.

C. Selfish station with fixed configuration

We verify that a station cannot obtain more throughput with a selfish configuration than by playing DOC in a scenario with $N = 10$ stations, half of them with $\rho_1 = 1$ and the other half





Fig. 7. Throughput for heterogeneous SNRs ($\rho_1 = 1, \rho_2 = 4$).Fig. 9. Selfish station with fixed configuration for different N and ρ_2 values.Fig. 8. Throughput of a selfish station for fixed configurations of $\{p_k, \bar{R}_k\}$.

Fig. 10. Throughput of selfish station with different adaptive strategies.

(including the selfish station) with $\rho_2 = 4$. The selfish station uses a fixed configuration and all other stations implement DOC. Fig. 8 shows the throughput of the selfish station for different configurations $\{p_k, \bar{R}_k\}$ of the selfish station. This is compared to the throughput that the station would obtain if it played DOC, given by the horizontal line. We observe that none of the selfish configurations provides greater throughput than DOC.

Fig. 9 analyzes the impact of fixed selfish configurations for a range of different N and ρ_2 values. It shows the largest throughput that a selfish station can receive with a fixed configuration, which is obtained by performing an exhaustive search over the $\{p_k, \bar{R}_k\}$ space. This throughput is compared to that which station would receive if it played DOC. Again, we observe that the station never benefits from playing selfishly, which validates the design of the DOC algorithm.

D. Selfish station with variable configuration

According to Theorem 5, a selfish station cannot benefit from changing its configuration over time. To verify this, we evaluate the throughput obtained by a selfish station with different adaptive strategies. These strategies are inspired by the schemes used in [15] for a similar purpose. The underlying

principle of all of them is that the cheating station uses a selfish configuration to gain more throughput and, when it realizes that it is not obtaining more throughput, it assumes that it has been detected as selfish and switches back to the honest configuration to avoid being punished.

In particular, we consider the following strategies. The ‘adaptive p_k strategy’ fixes the \bar{R}_k configuration of the selfish station to its optimal value, $\bar{R}_k = \bar{R}_k^*$, and modifies the p_k configuration as follows: the station uses a selfish configuration of $p_k = 1$ as long as it obtains some gain, i.e. $r_k > r_k^*$. When r_k drops below r_k^* , the station switches to the honest configuration, $p_k = p_k^*$, and stays with this configuration as long as r_k remains below $0.95r_k^*$. It switches back to $p_k = 1$ when r_k exceeds $0.95r_k^*$. The ‘adaptive \bar{R}_k strategy’ fixes the p_k configuration to the optimal value, $p_k = p_k^*$, and modifies the \bar{R}_k configuration following a strategy similar to the one above: the station uses a selfish configuration of $\bar{R}_k = 0$ (i.e., it uses all transmission opportunities) as long as it obtains some gain and switches to the honest configuration when it stops benefiting. Finally, the ‘adaptive p_k and \bar{R}_k strategy’ follows a similar behavior to the previous ones but adapts the configuration of both p_k and \bar{R}_k .

Fig. 10 compares the throughput obtained with each of the above strategies to that obtained with DOC for different values




Fig. 11. Stability analysis of the parameters of the PI controller.

of N . As expected, when all other stations play DOC, a given station maximizes its payoff by playing DOC as well, as this results in a larger throughput for the station than any of the other strategies. This confirms the result of Theorem 5.

E. Multiple selfish stations

Corollary 3 states that all of the selfish stations cannot simultaneously benefit by deviating from DOC: if one or more of the selfish stations experience throughput gains, there must be other selfish stations that suffer some loss. To validate this result, we consider a network with $N = 10$ stations (two selfish), half of them (including one of the selfish stations) with $\rho_1 = 1$ and the other half (including the other selfish station) with $\rho_2 = 4$. We perform an exhaustive search over a wide range of $\{p_i, \bar{R}_i\}$ configurations of the two selfish stations. The results of this experiment show that there is no configuration that simultaneously improves the throughput of the two selfish stations, which confirms the result of Corollary 3.

F. Parameter setting of the PI controller

The main objective in the configuration of the K_p and K_i parameters proposed in Section IV is to achieve a good tradeoff between stability and reaction time.

To validate that our system guarantees stable behavior, we analyze the evolution of the throughput received by a station over time in a wireless network with $N = 10$ stations. Fig. 11 shows the throughput for the chosen setting (labeled “ K_p, K_i ”) and for a configuration of these parameters 10 times larger (labeled “ K_p*10, K_i*10 ”). We observe that with the proposed setting, the throughput only suffers minor deviations around its average value. In contrast, for a larger setting, it exhibits highly oscillatory, unstable behavior.

To investigate the speed with which the system reacts against selfish stations, we consider the following scenario. In a wireless network with $N = 10$ stations, initially all stations play DOC. After 50 intervals, one station becomes selfish and changes its access probability to $p_k = 1$. Fig. 12 shows the evolution of the throughput of the selfish station over time. We observe from the figure that with our setting (labeled “ K_p, K_i ”), the system reacts quickly, and after a few tens of intervals the selfish station no longer benefits from its behavior. In contrast, for a parameter setting 10 times smaller




Fig. 12. Speed of reaction provided by the parameters of the PI controller.

(labeled “ $K_p/10, K_i/10$ ”) the reaction is very slow and it takes almost 2000 intervals for the station to stop benefiting from its misbehavior.⁸

The results show that with a larger setting of $\{K_p, K_i\}$ the system suffers from instability while with a smaller one it reacts too slowly. Hence, the proposed setting provides a good tradeoff between stability and reaction time.

G. Impact of channel coherence time

Our channel model is based on the assumption that different observations of the channel conditions are independent. In order to understand the impact of this assumption, we repeat the experiment of Fig. 7 using *Jakes' channel model* [20] to obtain the different channel observations. The results, for a Doppler frequency of $f_D = 2\pi/100\tau$, are given in Fig. 13. We observe that the throughput obtained is slightly smaller than that of Fig. 7. This is due to the fact that when the channel is bad, a station does not transmit after a successful contention and therefore it takes (on average) a shorter time until the next successful contention of this station. As a result, a station accesses the channel more often when it is bad than when it is good, which introduces a bias that slightly reduces the throughput. Overall, the results are sufficiently similar to those of Fig. 7 to conclude that our assumption on the channel model only has a minor impact on the resulting performance.

We further investigate whether, in the above scenario, a station with $\rho_2 = 4$ could obtain more throughput by using a selfish configuration. While the station obtains 1.752 Mbps with DOC, it can obtain up to 1.757 Mbps with a selfish configuration. Note that this increase is not due to the DOC design, as no other configuration provides the selfish station with more channel time, but rather due to the fact that the transmission rate threshold of [3] is not truly optimal under *Jakes' channel model*. In any case, the throughput gain of the selfish station is negligible.

⁸We note that, while the analysis of Section IV guarantees stability when all stations run DOC, our system is also stable when some of the stations are selfish. This is shown by the experiment of Fig. 12 where, after one of the stations turns selfish, the others increase their access probability to a value that ensures the selfish station does not have any gain. The system then remains stable at this point of operation.




Fig. 13. Performance with Jakes' channel model.




Fig. 14. No throughput gain for the selfish station with stations joining and leaving the network.

H. Stations joining and leaving the network

To assess the effectiveness of DOC with stations joining and leaving the network, we perform the following experiment. We consider a wireless network with 5 stations, one of which is a selfish station. After 1000 intervals, 5 additional stations join the wireless network, stay for I intervals and then leave. The initial 5 stations stay for another 1000 intervals. The selfish station plays with a configuration $\{p_1, \bar{R}_1\}$ when there are 5 stations in the network and a configuration $\{p_2, \bar{R}_2\}$ when there are 10 stations. We obtain these configurations by performing an exhaustive search over all possible configurations and selecting the $\{p_1, \bar{R}_1\}$ and $\{p_2, \bar{R}_2\}$ values that provide the selfish station with the largest average throughput. Fig. 14 shows the average throughput obtained by the station with this selfish strategy compared to the throughput it would obtain if it played DOC. The results confirm that the selfish station cannot obtain any gain by deviating from DOC.

VII. CONCLUSIONS

Recently proposed Distributed Opportunistic Scheduling (DOS) techniques provide throughput gains in wireless networks that do not have a centralized scheduler. One of the problems of these techniques, however, is that they are vulnerable to malicious users who may configure their parameters to obtain a greater share of the wireless resources. In this paper we address this problem by proposing a novel algorithm that

prevents such throughput gains from selfish behavior. With our approach, upon detecting a selfish user, stations react by using a more aggressive parameter configuration that indirectly punishes the selfish station. Such an adaptive algorithm has to carefully adjust the reaction against a selfish station in order to prevent the system from becoming unstable. A key aspect of the paper is that we use tools from *control theory* combined with *game theory* to design our algorithm: by conducting a control theoretic analysis, we show that when all stations run DOC the system converges to the desired configuration, and by conducting game theoretic analysis, we show that selfish stations cannot benefit from playing a different strategy.

ACKNOWLEDGEMENTS

The authors would like to thank the anonymous reviewers and Dr. S. Murphy for their valuable feedback.

REFERENCES


- [1] M. Andrews *et al.*, "Providing quality of service over a shared wireless link," vol. 39, no. 2, pp. 150–154, Feb. 2001.
- [2] P. Viswanath, D. Tse, and R. Laroia, "Opportunistic beamforming using dumb antennas," *IEEE Trans. Inf. Theory*, vol. 48, no. 6, pp. 1277–1294, Jun. 2002.
- [3] D. Zheng, W. Ge, and J. Zhang, "Distributed opportunistic scheduling for ad hoc networks with random access: An optimal stopping approach," *IEEE Trans. Inf. Theory*, vol. 55, no. 1, pp. 205–222, Jan. 2009.
- [4] D. Zheng *et al.*, "Distributed opportunistic scheduling for ad hoc communications with imperfect channel information," *IEEE Trans. Wireless Commun.*, vol. 7, no. 12, pp. 5450–5460, Dec. 2008.
- [5] P. S. C. Thejaswi *et al.*, "Distributed opportunistic scheduling with two-level probing," *IEEE/ACM Trans. Netw.*, vol. 18, no. 5, pp. 1464–1477, Oct. 2010.
- [6] S. Tan, D. Zheng, J. Zhang, and J. R. Zeidler, "Distributed opportunistic scheduling for ad-hoc communications under delay constraints," in *Proc. IEEE INFOCOM*, Mar. 2010, pp. 2874–2882.
- [7] A. Garcia-Saavedra, A. Banchs, P. Serrano, and J. Widmer, "Distributed opportunistic scheduling: A control theoretic approach," in *Proc. IEEE INFOCOM*, Mar. 2012, pp. 540–548.
- [8] P. Patras, A. Banchs, P. Serrano, and A. Azcorra, "A control-theoretic approach to distributed optimal configuration of 802.11 w lans," *IEEE Trans. Mob. Comput.*, vol. 10, no. 6, pp. 897–910, Jun. 2011.
- [9] F. Kelly, "Charging and rate control for elastic traffic," *Eur. Trans. Telecommun.*, vol. 8, no. 1, pp. 33–37, Jan. 1997.
- [10] G. Holland, N. Vaidya, and P. Bahl, "A rate-adaptive mac protocol for multi-hop wireless networks," in *Proc. ACM MOBICOM*, Jul. 2001, pp. 236–251.
- [11] B. Sadhegi, V. Kanodia, A. Sabharwal, and E. Knightly, "Opportunistic media access for multirate ad hoc networks," in *Proc. ACM MOBICOM*, Sep. 2002, pp. 24–35.
- [12] I. Menache and N. Shimkin, "Rate-based equilibria in collision channels with fading," *IEEE J. Sel. Areas Commun.*, vol. 26, no. 7, pp. 1070–1077, Sep. 2008.
- [13] P. Gupta, Y. Sankarasubramaniam, and A. Stolyar, "Random-access scheduling with service differentiation in wireless networks," in *Proc. IEEE INFOCOM*, Mar. 2005, pp. 1815–1825.
- [14] T. Glad and L. Ljung, *Control Theory: Multivariable and Nonlinear Methods*. London: Taylor & Francis, 2000.
- [15] M. Cagalj, S. Ganeriwal, I. Aad, and J.-P. Hubaux, "On selfish behavior in csma/ca networks," in *Proc. IEEE INFOCOM*, Mar. 2005, pp. 2513–2524.
- [16] J. Konorski, "A game-theoretic study of csma/ca under a backoff attack," *IEEE/ACM Trans. Netw.*, vol. 16, no. 6, pp. 1167–1178, Dec. 2006.
- [17] G. F. Franklin, J. D. Powell, and M. L. Workman. Reading, MA: Addison-Wesley, 1990.
- [18] C. Hollot, V. Misra, D. Towsley, and W. B. Gong, "A control theoretic analysis of red," in *Proc. IEEE INFOCOM*, Apr. 2001, pp. 1510–1519.
- [19] D. Fudenberg and J. Tirole, *Game Theory*. Cambridge, MA: MIT Press, 1991.
- [20] W. C. Jakes, *Microwave Mobile Communications*. New York: Wiley & Sons, 1975.



Albert Banchs (M'04–SM'12) received his degree in telecommunications engineering from the Polytechnic University of Catalonia in 1997, and his Ph.D. degree from the same university in 2002. He received a national award for the best Ph.D. thesis on broadband networks. He was a visiting researcher at ICSI, Berkeley, in 1997, worked for Telefonica I+D, Spain, in 1998, and for NEC Europe Ltd., Germany, from 1998 to 2003. He has been with the University Carlos III of Madrid since 2003. Since 2009, he also has a double affiliation as Deputy Director of the

institute IMDEA Networks.

Albert Banchs has authored over 80 publications in peer-reviewed journals and conferences and holds six patents. He is area editor for *Computer Communications* and has been senior and associate editor for *IEEE Communications Letters* and guest editor for *IEEE Wireless Communications*, *Computer Networks* and *Computer Communications*. He has served on the TPC of a number of conferences and workshops including IEEE INFOCOM, IEEE ICC and IEEE GLOBECOM, and has been TPC chair for European Wireless 2010, IEEE HotMESH 2010 and IEEE WoWMoM 2012. He is senior member of IEEE.



Andres Garcia-Saavedra received his B.Sc. degree in Telecommunications Engineering from University of Cantabria in 2009, and his M.Sc. degree in Telematics Engineering from University Carlos III of Madrid (UC3M) in 2010. He currently holds the position of Teaching Assistant and pursues his Ph.D. in the Telematics Department of UC3M. His work focuses on performance evaluation and energy efficiency of wireless networks.



Pablo Serrano (M'09) received his degree in telecommunication engineering and his Ph.D. degree from University Carlos III of Madrid (UC3M) in 2002 and 2006, respectively. He has been with the Telematics Department of UC3M since 2002, where he currently holds the position of associate professor. In 2007 he was a visiting researcher at the Computer Network Research Group at the University of Massachusetts, Amherst. His current work focuses on performance evaluation of wireless networks. He has authored over 40 scientific papers in peer-reviewed

international journals and conferences. He also serves as TPC member of several international conferences, including IEEE GLOBECOM and IEEE INFOCOM.



Joerg Widmer (M'06–SM'10) is a Chief Researcher at Institute IMDEA Networks in Madrid, Spain. His research expertise covers computer networks and distributed systems, ranging from MAC layer design, sensor networking, and network coding to transport protocols and Future Internet architectures. From June 2005 to July 2010, he was manager of the Ubiquitous Networking Research Group at DOCOMO Euro-Labs in Munich, Germany, leading several projects in the area of mobile and cellular networks. Before joining DOCOMO Euro-Labs, he

worked as post-doctoral researcher at EPFL, Switzerland, focusing on ultra-wide band communication and network coding.

Joerg Widmer received his M.S. and Ph.D. degrees in computer science from the University of Mannheim, Germany in 2000 and 2003, respectively. In 1999 and 2000 he was a visiting researcher at the International Computer Science Institute in Berkeley, CA, USA. He has authored more than 100 conference and journal papers and three IETF RFCs. He also holds several patents, serves on the editorial board of *IEEE TRANSACTIONS ON COMMUNICATIONS*, and regularly participates in the program committees of several major conferences. He is senior member of IEEE and ACM.