

Flexible Dynamic Spectrum Allocation in Cognitive Radio Networks Based on Game-theoretical Mechanism Design

José R. Vidal, Vicent Pla, Luis Guijarro, and Jorge Martinez-Bauset *

Universitat Politècnica de València, 46022 Valencia Spain
{jrvidal,vpla,lguijar,jmartine}@dcom.upv.es

Abstract. In this paper we present an approach based on game-theoretical mechanism design for dynamic spectrum allocation in cognitive radio networks. Secondary users (SU) detect when channels can be used without disrupting any primary user and try to use them opportunistically. When an SU detects a free channel, it estimates its capacity and sends the valuation of it to a central manager. The manager calculates a conflict-free allocation by implementing a truthful mechanism. The SUs have to pay for the allocation an amount which depends on the set of valuations, and they behave as benefit maximizers. We present and test two mechanisms implementing this idea which are proved to be truthful, and that are tractable and approximately efficient. We show the flexibility of these mechanisms by illustrating how they can be modified to achieve other objectives such as fairness and also how they can operate without really charging the SUs.

Keywords: Cognitive radio, spectrum sharing, game theory, mechanism design.

1 Introduction

Cognitive radio is the technology that enables dynamic spectrum access (DSA) networks to fully utilize the scarce spectrum resources. In DSA networks, users who have no spectrum licenses, known as SUs, are allowed to use the spectrum. In this paper, we will focus on DSA networks with *hierarchical* and *overlay* access [14]. In the hierarchical access model, SUs use spectrum that is licensed to primary users (PUs). As PUs have priority in using the spectrum, when SUs coexist with PUs, they have to perform real-time wideband monitoring of the licensed spectrum to be used in order to avoid harmful interference to PUs. In overlay access, also referred to as opportunistic spectrum access, SUs only use the licensed spectrum when PUs are not transmitting. In order not to interfere with the PUs, SUs need to sense the licensed frequency band and detect the spectrum opportunities.

* This work was supported by the Spanish government through projects TIN2008-06739-C04-02 and TIN2010-21378-C02-02, and by Universitat Politècnica de València through PAID-06-09.

The availability and quality of spectrum opportunities may change rapidly over time due to PUs activity and competition between SUs. Therefore, dynamic spectrum allocation and sharing schemes are needed to achieve flexible spectrum access in long-run scenarios. They should be able to adapt to the spectrum dynamics, e.g., channel variations, based on local observations. In centralized spectrum allocation, the opportunistic spectrum access of SUs is coordinated by a central element serving as a spectrum manager. The spectrum manager collects operation information from each SU, and allocates the spectrum resources to achieve certain objectives such as efficiency or fairness. This kind of spectrum sharing is non-cooperative, since SUs only aim at maximizing their own benefit, so SUs might exchange false information about their channel conditions in order to get more access opportunities to the spectrum. Therefore, cheat-proof spectrum sharing schemes should be developed to meet the objectives.

Dynamic spectrum sharing has been extensively studied from a game theoretical perspective [11]. Mechanism design is a game theoretical approach that can be applied to dynamically redistribute spectrum across several players (PUs and SUs) to meet their demands. Mechanisms aim at achieving the desired equilibrium by enforcing SUs to play truthfully, so that the spectrum resources are allocated according to reliable information. This is attained by means of payments, which are collected and redistributed by a trusted entity. In spectrum auction games, a specific form of mechanism design, a spectrum manager collects bids and allocates spectrum resources to SUs and charges them according to some rules. By multiplexing spectrum supply and demand in time and space, dynamic auctions can improve spectrum utilization or fairness.

Most of works on spectrum auctions focus on the scenario where one or more PUs want to lease spectrum to SUs, and a monetary gain for PUs is involved. In [12], an auction model is proposed where a PU announces a portion of its licensed band and a unit price, and SUs bid for the desired amount of bandwidth. In [3], SUs bid for a pool of homogeneous channels announcing a demand curve, from which the auction is cleared using an approximation algorithm. A belief-assisted double auction is proposed in [8], with collusion-resistant strategies based on the use of optimal reserve prices. Another solution for double auctions is presented in [15], where several PUs auction a channel each, while several SUs bid for just one of them, assuming that all the channels are homogeneous to the SUs.

Other works propose auctions for power allocation in spectrum sharing. In [2] sequential second price auctions are proposed assuming complete information. In [6] an auction model is proposed where the utility of each SU is defined as a function of the received signal-to-noise-and-interference ratio (SINR). SUs are charged a unit price for their received SINR, so that the auction achieves the desired social utility. Mechanism design is used in [13]. There the underlay model is assumed, so SUs may transmit in the licensed spectrum when PUs are also transmitting, and their power transmission is restricted by the interference temperature limit. The objective power allocation is calculated using channel infor-

mation obtained locally by SUs. A truthful mechanism with monetary transfers enforces SUs to reveal this information.

Our work differs from those cited above in two major aspects. Firstly, we address the problem of sharing spectrum opportunities between SUs without the intervention of PUs. These opportunities appear sparse in time and space, therefore we propose mechanisms to allocate them in real time, i.e., mechanisms that have to be run every time that a spectrum opportunity appears. We present two mechanisms (one deterministic and one randomized) for dynamic spectrum allocation which are truthful, computationally tractable and approximately efficient and that are simple enough for its use in real-time allocations. They have the flexibility to achieve long-run objectives other than efficiency, such as fairness, maintaining their properties. Secondly, we investigate how these mechanisms can operate without monetary transactions, which would make the solution easier to implement. If monetary gain is not involved in the spectrum allocation problem, SU payments will no longer be chosen as money, but as an alternative form of ‘virtual currency’. This internal currency will be managed by the spectrum manager, which will record the credit of every SU and will distribute it to them. In order that this currency retains its value for the SUs, the mechanism itself should consider the credit kept by any SU. That is, the choice and payment functions should depend on the credit so limiting what SUs spend, and a proper ‘cash flow’ should be redistributed to SUs. We investigate how this can be done and how it affects to the mechanisms properties.

The rest of the paper is organized as follows. In Section 2 we describe the model of spectrum sharing on networks with hierarchical and overlay access and with a central manager. In Section 3 we describe the mechanism design background applied to this problem, present two mechanisms based on the theoretical results and describe how these mechanisms can be modified to achieve fairness and to operate with a virtual currency. We show and discuss experimental results in Section 4 and conclude in Section 5.

2 Spectrum Sharing Model

We assume that the spectrum is divided into non-overlapping orthogonal channels and that SUs are able to detect when a channel can be used without interrupt any PU. When one of these opportunities, called *white spaces*, appears, SUs may try to use it opportunistically. Every white space might be used by one or several SUs, with the condition of not conflict between them. We model the interference between SUs according to the *protocol model* [4,7]. This model assumes that, for a given channel, SU i has a communication range R_i and a larger interference range R'_i . The channel can be used by SU i if the receiver is at a distance smaller than R_i . Two SUs i and j are in conflict if the distance between them is smaller than R'_i or smaller than R'_j , and in this case they cannot transmit simultaneously. In this network, R_i and R'_i depend on the transmitting power, which is imposed by the position and channel usage of nearby PUs. Additionally, the transmitting power and the distance to the receiver determines the

transmission rate, hence for a given white space, SU i will be able to transmit at B_i bps. With these assumptions, the conflicts for a given white space can be modelled by a *conflict graph* whose vertices correspond to the SUs which are able to use the channel. There is an edge between vertices i and j if SUs i and j are in conflict. This graph can be written in matrix form:

$$C = [c_{ij}]_{N \times N} \quad \text{with} \quad c_{ij} \in \{0, 1\}. \quad (1)$$

Where $c_{ij} = 1$ if SUs i and j are in conflict. The sharing problem in this scenario is, for each white space, to find an allocation compatible with the conflict graph:

$$X = [x_i]_{N \times 1} \quad \text{where} \quad x_i \in \{0, 1\} \quad \text{and} \quad x_i \cdot x_j = 0 \quad \text{if} \quad c_{ij} = 1. \quad (2)$$

Where $x_i = 1$ if the channel is allocated to SU i . The allocation cannot fulfil efficiency and fairness simultaneously, so the objective will be a trade-off between them. This objective should be achieved on long-term basis. Long-term efficiency is achieved if every white space is allocated efficiently. A single allocation is said to be efficient if

$$X_{\text{eff}} = \arg \max_X \sum_{i=1}^N x_i B_i. \quad (3)$$

However, fairness cannot be achieved for each white space, so long-run fairness should be defined as a time-average. If denote by X^j the allocation made at the j -th white space, the proportional fairness criterion for M consecutive allocations is

$$\mathbf{X}_{\text{fair}} = (X_{\text{fair}}^1, \dots, X_{\text{fair}}^M) = \arg \max_{\mathbf{X}} \sum_{j=1}^M \sum_{i=1}^N \log x_i^j B_i^j. \quad (4)$$

In our proposal we assume that there is a spectrum manager whose rules are abided by SUs, and there is a control channel dedicated to the communication between manager and SUs [5]. Every time that a white space appears, the SUs detect it and estimate B_i . Those SUs willing to use the channel send B_i to the manager. They also detect and communicate which neighbours they conflict with (c_{ij}), and from this the manager derives the conflict graph. The manager then will calculate the allocation according to the objectives, and communicate the allocation to SUs.

We also assume that no SU can benefit from lying about the conflict graph. This is not strictly true but it is a reasonable assumption in most situations, because if SU i declares non-existing conflicts (it falsely declares $c_{ij} = 1$), this will reduce the set of possible allocations, so it will reduce its chances to obtain the channel. On the other side, if it hides a conflict (it falsely declares $c_{ij} = 0$), the resulting allocation may be useless for i and j . However, they could benefit from lying about the channel bitrate estimation; a higher declared value of B_i rises the value of $\sum_{i=1}^N x_i B_i$ if $x_i = 1$, so rising the manager evaluation of allocations including SU i . We propose a solution based on mechanism design for enforcing the SUs to tell the truth to the spectrum manager driven by self-interest.

3 Spectrum Sharing Based on Mechanism Design

We model the allocation procedure of each white space as a mechanism in which the players are the SUs and a spectrum manager that implements the game rules. Formally, a *mechanism* is a game in which the players do not know the utilities of the other players, and the rules of the game are designed such that the equilibrium of the game is guaranteed to have certain desired properties. It is defined by the tuple $(N, O, \Theta, p, u, A, M)$:

- N is a set of n players or, in the spectrum allocation problem, SUs.
- O is the set of possible outcomes and will include the allocations.
- $\Theta = \Theta_1 \times \dots \times \Theta_n$ is the set of possible SU types. SU i type $\theta_i \in \Theta_i$ is known only by SU i and determines its utility function. Here, θ_i is determined by B_i .
- p is the probability distribution of Θ . Here p depends on the position of SUs and interference restrictions.
- $u = (u_1, \dots, u_n)$, where $u_i : O \times \Theta \mapsto \mathbb{R}$ is the utility function of SU i depending on the outcome o and on its type θ_i .
- $A = A_1 \times \dots \times A_n$ is the action profile, where A_i is the set of actions available to SU i . In this problem, the action allowed to SU i is to declare its type θ_i , i.e., $A_i = \Theta_i$, what results in a so-called *direct mechanism*.
- $M : A \mapsto \Pi(O)$ maps each action profile to a distribution over outcomes. If the mechanism is deterministic, then $M : A \mapsto O$.

If we want the resulting allocations to have certain properties as efficiency or fairness, the mechanism should implement the corresponding social choice function $C : u \mapsto \Pi(O)$, i.e. the game must have an equilibrium a^* in which $M(a^*) = C(u)$. We also want the SUs to reveal truthfully their types, what will be fulfilled if the equilibrium is $a^* = (\theta_1, \dots, \theta_n)$. In this problem, the solution can be restricted to a so-called *quasilinear mechanism* [10], where the possible outcomes can be written as

$$O = X \times \mathbb{R}^n, \quad (5)$$

where X is a finite set. The outcome is then an allocation plus a vector of n real numbers. The i -th value of this vector is the price that SU i has to pay for the allocation. Thus, for a given vector of types $\theta \in \Theta$, the utility functions is

$$u_i(o, \theta) = u_i(x, \theta) - p_i, \quad (6)$$

where $u_i(x, \theta)$ is the *value* that allocation x has for SU i , and p_i is the price that SU i has to pay when the allocation is x . This mechanism has the property of *conditional utility independence*, because the value u_i depends only on SU i type and not on other SUs type. We refer to the value u_i as *valuation* of SU i for allocation x , and we write it as $v_i(x)$, considering implicit its dependence on θ_i . Then:

$$u_i(o, \theta) = v_i(x) - p_i. \quad (7)$$

Here $v_i(x)$ can be interpreted as the maximum amount that SU i would be willing to pay for allocation x . Let v_i denote the mapping that assigns a valuation $v_i(x)$ for each $x \in X$. Revealing type θ_i is equivalent to revealing v_i and the set of allowed actions for SU i is the set of possible values of v_i . Let \hat{v}_i denote the declared valuation of SU i , which might be different from v_i . The mechanism can be interpreted as an auction, being $\hat{v} = (\hat{v}_1 \cdots \hat{v}_n) \in V$ the bids, and $p = (p_1 \cdots p_n)$ the prices that the bidders have to pay. Given \hat{v} , the manager calculates allocation and payment from:

- A choice rule: $f : V \mapsto \Pi(X)$, or $f : V \mapsto X$ if deterministic.
- A payment rule: $p : V \mapsto \mathbb{R}^n$.

Our objective is to design these rules so as to achieve truthfulness [10]. A truthful mechanism is in equilibrium when $\hat{v} = v$, and no SU can benefit from declaring a false valuation. A theoretical result says that, for non-restricted quasilinear preferences domains, the only existing truthful mechanism is the *weighted Vickrey Clarke Groves* (VCG) mechanism. This is a well known mechanism that implements efficiency and is not computationally tractable. However, here we need a mechanism able to implement other criteria as fairness, and simple to compute. Fortunately, the valuation setting described before belongs to the family of *single parameter valuations* [10], in which a valuation v_i is defined by a single value. Formally, for each SU i the set of allocations can be partitioned into a winning set W_i and a losing set,

$$v_i(x) = \begin{cases} v_i & \text{if } x \in W_i \\ 0 & \text{if } x \notin W_i \end{cases} \quad (8)$$

Here $x \in W_i$ if $x_i = 1$, i.e., SU i wins if the channel is allocated to it, regardless what happens to other SUs. Let \hat{v}_{-i} denote the vector of valuations of all SUs except SU i . The mechanism has good properties if the choice function f is monotone:

$$\forall \hat{v}_i, \forall \hat{v}'_i, \hat{v}'_i > \hat{v}_i, f(\hat{v}_i, \hat{v}_{-i}) \in W_i \Rightarrow f(\hat{v}'_i, \hat{v}_{-i}) \in W_i. \quad (9)$$

That is, if a SU wins with a given valuation, it wins also with all higher valuations. Given \hat{v}_{-i} , the critical value for SU i is defined as the minimum value of \hat{v}_i for which SU i wins:

$$c_i(\hat{v}_{-i}) = \sup_{f(v_i, \hat{v}_{-i}) \notin W_i} v_i. \quad (10)$$

A deterministic single parameter domain mechanism is truthful if and only if every winning bid pays $c_i(\hat{v}_{-i})$ plus a function independent on its valuation. The mechanism is said to be *normalized* if losing bidders pay 0. Every truthful mechanism can be turned into a normalized one. Thus, the payment rule can be expressed, without loss of generality, as

$$p_i(\hat{v}_i, \hat{v}_{-i}) = \begin{cases} c_i(\hat{v}_{-i}) + h_i(\hat{v}_{-i}) & \text{if } x \in W_i \\ 0 & \text{if } x \notin W_i \end{cases} \quad (11)$$

A randomized mechanism is truthful on expectation, if for all bidders, revealing its true valuation maximizes its expected benefit, i.e., $\forall i \forall v_i \forall v_{-i} \forall \hat{v}_i$:

$$E[v_i(f(v_i, \hat{v}_{-i})) - p_i(f(v_i, \hat{v}_{-i}))] \geq E[v_i(f(\hat{v}_i, \hat{v}_{-i})) - p_i(f(\hat{v}_i, \hat{v}_{-i}))]. \quad (12)$$

Let us denote by $\omega_i(\hat{v}_i, \hat{v}_{-i}) = Pr[f(\hat{v}_i, \hat{v}_{-i}) \in W_i]$ the probability that SU i wins. Then, the expected benefit for SU i is

$$u_i(\hat{v}_i, \hat{v}_{-i}) = v_i \omega_i(\hat{v}_i, \hat{v}_{-i}) - p_i(f(\hat{v}_i, \hat{v}_{-i})). \quad (13)$$

If $\omega_i(\hat{v}_i, \hat{v}_{-i})$ is monotonically non-decreasing in \hat{v}_i and v_i^0 is the valuation under which i cannot win, the truthfulness condition for a normalized mechanism is [1]

$$p_i(\hat{v}_i, \hat{v}_{-i}) = \hat{v}_i \omega_i(\hat{v}_i, \hat{v}_{-i}) - \int_{v_i^0}^{\hat{v}_i} \omega(t, \hat{v}_{-i}) dt + h_i(\hat{v}_{-i}). \quad (14)$$

Based on these results, we propose two simple truthful mechanisms for spectrum sharing: a deterministic one and a randomized one. As the experimental results in Sec. 4 show, the randomized one exhibits better properties. They are executed every time a white space appears. When SUs detect the white space, each of them estimates its valuation $v_i = B_i$. Then, each SU willing to use the channel sends a bid containing its declared valuation \hat{v}_i to the spectrum manager. The manager starts an auction when it receives the first bid and, after a fixed time interval, it closes the auction. The auction is cleared by applying the choice and paying rules to the vector \hat{v} , resulting in a set of winning SUs and a vector p of payments. Finally, the manager sends a message to the winners SUs.

3.1 A Deterministic Truthful Mechanism

The following mechanism approximates efficiency by giving priority to higher valuations. The resulting choice function is monotone and then every winning bid has a critical value. The payment function is the critical bid, so the mechanism is truthful. Furthermore, both choice and payment functions are computationally tractable.

- Choice function. Given a list of bids:
 1. Order the list from highest to lowest value of valuations.
 2. From the beginning of the list until the end of the list:
Allocate the channel to bidder i if it does not conflict with any preceding winning bidder: $x_i = 1$ if $c_{ij}x_j = 0, \forall j$ preceding i in the list.
- Payment rule. Keep the order of the bid list above. For each bid i :
 1. Remove bid i from the list.
 2. Start a virtual allocation x' with $x'_j = 0 \forall j$.
 3. From the beginning of the list until the end of the list or until a bidder j conflicting with i virtually wins:
Allocate virtually the channel to bidder k if it does not conflict with any preceding virtual winner: $x'_k = 1$ if $c_{kl}x'_l = 0, \forall l$ preceding k in the list.
Then:

- If the end of the list is reached, bidder i wins and pays 0.
 - If $\hat{v}_i > \hat{v}_j$, bidder i wins and pays the critical value $c_i = \hat{v}_j$.
 - If $\hat{v}_i < \hat{v}_j$, bidder i losses and pays 0.
4. Insert bid i into the list at its original position.

3.2 A Randomized Truthful Mechanism

A randomization of the previous mechanism results in a more flexible choice function. The resulting winning probability function $\omega_i(\hat{v}_i, \hat{v}_{-i})$ is monotonically non-decreasing. Equation (14) is applied to obtain the payment function, so ensuring the truthfulness of the mechanism. Both choice and payment functions are computationally tractable.

- Choice function. Given a list of bids:
 1. For each bid j , calculate a random value k_j with uniform distribution between 0 and \hat{v}_j .
 2. Order the list from highest to lowest value of k_j .
 3. From the beginning of the list until the end of the list:
Allocate the channel to each bidder if it does not conflict with any preceding winning bidder: $x_i = 1$ if $c_{ij}x_j = 0, \forall j$ preceding i in the list.
- Payment rule. Keep the order of the bid list. For each bid i :
 1. Remove bid i from the list.
 2. Start a virtual allocation x' with $x'_j = 0 \forall j$.
 3. From the beginning of the list until the end of the list or until a bidder j conflicting with i virtually wins:
Allocate virtually the channel to bidder k if it does not conflict with any preceding virtually winner: $x'_k = 1$ if $c_{kl}x'_l = 0, \forall l$ preceding k in the list.
Then:
 - If the end of the list is reached, i wins and pays 0.
 - If $\hat{v}_i > k_j$, bidder i wins with probability¹:

$$\omega_i(\hat{v}_i, k_{-i}) = 1 - \frac{k_j}{\hat{v}_i}, \quad (15)$$

and pays

$$p_i = \hat{v}_i \omega_i(\hat{v}_i, k_{-i}) - \int_{k_j}^{\hat{v}_i} \omega_i(t, k_{-i}) dt = k_j \ln \frac{\hat{v}_i}{k_j}. \quad (16)$$

- If $\hat{v}_i < k_j$, bidder i cannot win and pays 0.
4. Insert bid i into the list at its original position.

¹ Here the function $\omega_i(\hat{v}_i, k_{-i})$ has been used instead of $\omega_i(\hat{v}_i, \hat{v}_{-i})$. However, the proof of the condition expressed by (14) holds for both functions [1].

3.3 Fairness

Unfairness has two main causes. Firstly, all SUs may have not the same opportunities because physical restrictions and those SUs reporting higher valuations, if true, are favoured. This may be alleviated by limiting the budget of SUs, as described in the next section. Secondly, the mechanism itself does not treat all the bidders the same way. It is easy to see that bidders having more competition (more neighbours conflicting with) have less chances to win and when winning, they pay a higher price.

The previous mechanisms can be modified to compensate unfairness caused by the degree of competition. A simple way to do it is by weighting the valuations by a function increasing with the number of edges in the conflict graph. Simulation results are shown in Sec. 4. Clearly, this does not affect the truthfulness, because the original conditions still hold. The choice function is still monotonically non-decreasing and the payment function for the deterministic mechanism is the critical value and, for the randomized mechanism is given by

$$p_i = \hat{v}_i \omega_i(\hat{v}_i, k_{-i}) - \int_{\frac{k_j}{w_i}}^{\hat{v}_i} \omega_i(t, k_{-i}) dt = \frac{k_j}{w_i} \ln \frac{w_i \hat{v}_i}{k_j}. \quad (17)$$

3.4 Credit Restriction and Redistribution

So far we have considered that SUs behave as profit maximizers who value a unit of transmission rate the same as a monetary unit. Consequently, they obtain the maximum profit when the rate they get minus the money they pay is maximum. This approach is valid if SUs are effectively charged for the spectrum. Another possibility is that what they pay was not real money, but a virtual currency unit internal to the system. This could be useful if what we want is just to fairly share the spectrum, not to trade with it. This would also make the system simpler by avoiding monetary transactions. Furthermore, a limitation of the credit offered to SUs would help to achieve fairness.

In this last approach, the manager should redistribute the payments and record the credit of each SU. The amount of credit given to each SU must be limited to approximately the amount of credit they would spend when it received the fair share of the spectrum. Credit limitation is the way to get them to value the credit and still behave as profit maximizers. To achieve this, the mechanism has to be modified in two ways:

- SUs need to receive a flow of credit along the time. For system stability, the total amount of credit received by them should be equal to the total amount of their payments. If this property holds for each auction, it is called *budgeted balance*. Here we only need it to hold in the long term, and this is achieved if in every auction the expected value of what each bidder receives equals the expected value of what it pays.
- SUs should be punished when they run out of credit. However, SUs with negative credit cannot be excluded from auctions, because they must still get

their share of the redistribution and so recover credit for successive auctions. Instead, when a SU has low credit, the mechanism would grant it a lower winning probability.

There are several solutions for the redistribution problem that preserve truthfulness [9]. All of them satisfy the condition that what a bidder receives is independent of its valuation. Based on this idea, we have tested the following redistribution method. For every auction, a parallel auction is created with the same bidders and random valuations. Every bidder receives what it would pay in this parallel auction. If random valuations have a random distribution which approximates the distribution of real valuations, and if bidders play truthfully, long term budget balance is achieved.

For credit limitation, we have added to the previous mechanisms another weighting factor q_i which depends on credit. Valuations are weighted by q_i and the payment functions are modified accordingly. If we consider a single auction, this change preserves the truthfulness of the mechanism, as it did the weighting for fairness described in Sec. 3.3. However, this does not hold if we consider successive auctions, because q_i depends on the credit of SU i , which in turn depends on \hat{v}_i . As a consequence the valuation on the current auction has an effect not only in the benefit obtained in this auction, but also in what happens in successive auctions. For this reason, it cannot be assured that SUs could not obtain a long term benefit from lying. Because successive auctions are not independent, this issue should be studied as a repeated game.

To alleviate this problem, the influence of credit in the mechanism should be as small as possible, that is, it should be $q_i = 1$ most of time while $q_i < 1$ is only required when SU i is close to run out of credit. Then, if the redistribution policy works properly, SUs will work most of time with $q_i = 1$. When a SU overbids and spends all its credit it is punished. On the other side, a SU cannot benefit from saving credit because, once it reaches enough credit to have $q_i = 1$, having more does not increase its winning probability. Simulation results shown in Sec. 4 suggest that the mechanism still behaves reasonably well, though truthfulness cannot be guaranteed in all situations.

4 Experimental Results and Discussion

We have evaluated the properties of previous mechanisms by means of discrete-event simulations. We have simulated a static configuration of 6 SUs with the conflict graph of Fig. 1, located into the interference range of a PU which conveys traffic bursts whose inter-arrival time is exponentially distributed with mean 20 time units and their duration is exponentially distributed with mean 10 units. The PU has a pool of 10 channels, and allocates traffic bursts randomly. Idle periods of these channels are detected by SUs as white spaces. SUs estimations of the bit rate of a white space is uniformly distribution between 0 and 2 bit-rate units. Every white space all SUs send a bid containing its valuation of the channel to the central manager, which executes the mechanism. Then it sends a message

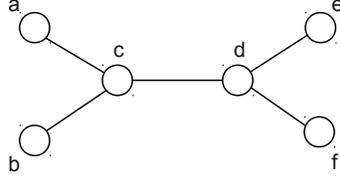


Fig. 1. Conflict graph of simulation set.

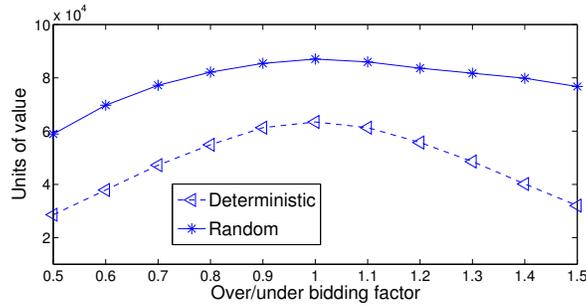


Fig. 2. Total benefit of SU c as a function of its h .

to the winners who occupy the channel until it is used by the PU again. The credit available for each SU is recorded and updated by the spectrum manager. We assume that SUs are backlogged and their objective is to send as much traffic as possible. All simulation runs have a length of 10^7 time units, what have been checked to yield statistically significant results

The plot in Fig. 2 illustrates the truthfulness of both mechanisms without credit restriction and confirms the theoretical result. Here, total benefit of SU c is the sum of the bit-rate of all the channels it won minus the total price it paid for them. This is plotted as a function of an *over/under bidding factor* h which measures the relation between the declared valuation and the true valuation: $\hat{v}_c = hv_c$. When $h > 1$, c it is overbidding and when $h < 1$ it is underbidding. The rest of SUs are truthful. It can be seen that SU c maximizes its benefit when $h = 1$, i.e., when it declares the truth.

We have also tested the mechanisms with credit restriction an redistribution as described in Sec. 3.4. Here the weighting factor q_i is set to 1 when $credit_i > 50$, to 0 when $credit_i < -50$, and it is varied linearly between 0 and 1 when $-50 < credit_i < 50$. The initial credit is 100. By doing this, SUs are forced to keep its credit not far under 0, that is, they cannot spend much more than they receive. Since credit is not exceeded nor accumulated, now the benefit is the obtained bit-rate, i.e., the mean traffic. Some results for SU c are shown in Figs. 3 and 4. We have run simulations varying h in other SUs and with other conflict graphs obtaining similar results not shown here.

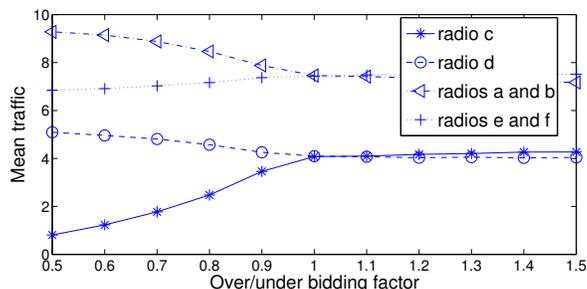


Fig. 3. Mean traffic of SU c as a function of its h . Deterministic mechanism.

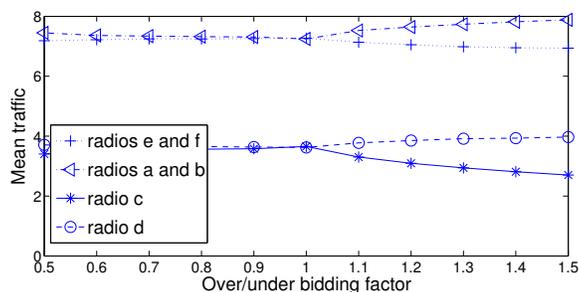


Fig. 4. Mean traffic of SU c as a function of its h . Randomized mechanism.

It can be seen that, with the deterministic mechanism, SU c can benefit from overbidding. This does not happen with the random mechanism. The reason of this different behaviour can be found in Fig. 5, which plots the unit price that SU c pays as a function of h . For the deterministic mechanism, the price grows up to a maximum for $h = 1$, because the payment function depends only on the critical value. Therefore, if a bid wins, the price does not depend on the bid value, that is, two winning bids with different valuations pay the same. In contrast, with the randomized mechanism the price function is given by (16) and depends on the valuation, so the unit price grows as it can be seen in the plot. This characteristic of the payment function adds stability to the random mechanism, because makes it more expensive to overbid.

The capability of the mechanisms to compensate the unfairness due to disadvantage in the conflict graph, as described in Sec. 3.3, has been evaluated in an experiment in which valuations \hat{v}_i are multiplied by a competition weight $w_i = e_i^{1.5}$, where e_i is the number of edges related to vertex i in the conflict graph of SU i . Figure 6 shows the result for SU c with the random mechanism. It can be seen that, compared with Fig. 4, a much more fair share is achieved.

We have also evaluated the efficiency of these mechanisms by comparing the sum of the mean traffic of the 6 SUs without fairness compensation with that obtained by another mechanism with an efficient choice function implemented

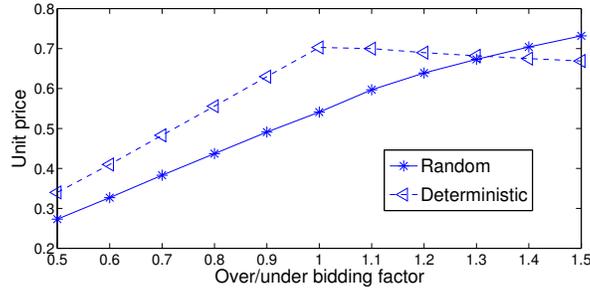


Fig. 5. Unit price paid by SU c as a function of its h .

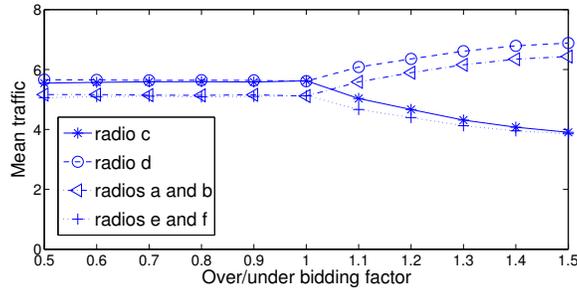


Fig. 6. Mean traffic of SU c as a function of its h with fairness compensation.

by exhaustive search. Results are plot in Fig. 7. It can be can see that both mechanism closely approximate efficiency.

5 Conclusions

We present two truthful and low complexity mechanisms for real-time spectrum allocations. We show how they can be modified to implement social fairness, maintaining its properties. We also show how they can work when a virtual currency instead of money is used, by controlling the credit of the SUs, making the choice and payment functions dependent on the credit, and redistributing cash to SUs. However, when the mechanisms are dependent on credit, although they are still truthful on a single run, on repeated runs truthfulness does not hold, because successive runs become dependent. Although truthfulness cannot be guaranteed in all situations, our experimental results shown that under certain conditions they still behave truthfully.

We are currently working in defining which are the conditions that the credit restriction and the redistribution policy have to fulfil so that the resulting repeated mechanism is truthful.

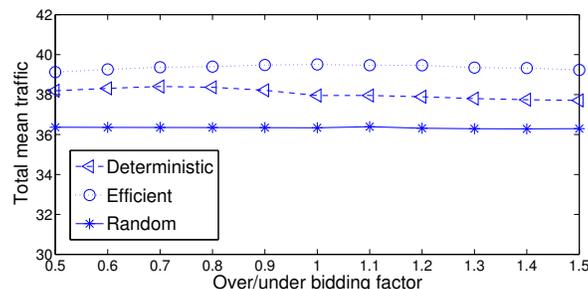


Fig. 7. Mean traffic conveyed by all SUs as a function of h of SU c .

References

1. Archer, A., Tardos, E.: Truthful mechanisms for one-parameter agents. In: IEEE symposium on Foundations of Computer Science (FOCS '01). pp. 482–491. IEEE Computer Society (2001)
2. Bae, J., Beigman, E., Berry, R.A., Honig, M.L., Vohra, R.V.: Sequential bandwidth and power auctions for distributed spectrum sharing. *IEEE Journal on Selected Areas in Communications* 26(7), 1193–1203 (2008)
3. Gandhi, S., Buragohain, C., Cao, L., Zheng, H., Suri, S.: Towards real-time dynamic spectrum auctions. *Comput. Netw.* 52, 879–897 (2008)
4. Gupta, P., Kumar, P.R.: The capacity of wireless networks. *IEEE Transactions on Information Theory* 46, 388–404 (2000)
5. Han, C., Wang, J., Yang, Y., Li, S.: Addressing the control channel design problem: OFDM-based transform domain communication system in cognitive radio. *Computer Networks* pp. 795–815 (2008)
6. Huang, J., Berry, R.A., Honig, M.L.: Auction-based spectrum sharing. *Mob. Netw. Appl.* 11, 405–418 (2006)
7. Jain, K., Padhye, J., Padmanabhan, V., Qiu, L.: Impact of interference on multi-hop wireless network performance. *Wireless networks* 11(4), 471–487 (2005)
8. Ji, Z., Liu, K.: Belief-assisted pricing for dynamic spectrum allocation in wireless networks with selfish users. In: *SECON '06*. vol. 1, pp. 119–127 (2006)
9. Moulin, H., Shenker, S.: Strategyproof sharing of submodular costs: budget balance versus efficiency. *Economic Theory* 18(3), 511–533 (2001)
10. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V. (eds.): *Algorithmic game theory*. Cambridge University Press (2007)
11. Wang, B., Wu, Y., Liu, K.R.: Game theory for cognitive radio networks: An overview. *Comput. Netw.* 54, 2537–2561 (2010)
12. Wang, X., Li, Z., Xu, P., Xu, Y., Gao, X., Chen, H.H.: Spectrum sharing in cognitive radio networks: an auction-based approach. *Trans. Sys. Man Cyber. Part B* 40, 587–596 (2010)
13. Wu, Y., Wang, B., Liu, K.J.R., Clancy, T.C.: Repeated open spectrum sharing game with cheat-proof strategies. *Trans. Wireless. Comm.* 8, 1922–1933 (2009)
14. Zhao, Q.: A survey of dynamic spectrum access: signal processing, networking, and regulatory policy. In: *IEEE Signal Processing Magazine*. pp. 79–89 (2007)
15. Zhou, X., Zheng, H.: TRUST: A general framework for truthful double spectrum auctions. In: *INFOCOM*. pp. 999–1007 (2009)