

Negotiation-Based Distributed Power Control in Wireless Networks with Autonomous Nodes

Vaggelis G. Douros, George C. Polyzos and Stavros Toumpis
Mobile Multimedia Laboratory, Department of Informatics
Athens University of Economics and Business
Patision 76, 10434, Athens, Greece
{douros, polyzos, toumpis}@aueb.gr

The efficient management of the radio spectrum is a key functionality in every type of wireless network. Wireless nodes generally have heterogeneous QoS targets, which sometimes cannot be satisfied for all of them due to the high interference levels that frequently arise, even in sparse topologies. In this work, we propose a distributed negotiation-based power control algorithm that aims at maximizing the number of nodes achieving their QoS targets. Our algorithm combines the influential Foschini-Miljanic power control algorithm with a bargaining-inspired phase (among the unsatisfied nodes only). In particular, all nodes are endowed with an initial budget; unsatisfied nodes randomly pick others to negotiate with for the level of their transmission powers; if a negotiation leads to an agreement, a node gives some (predefined) reward to the other and the latter reduces its power to the agreed level; the process is then repeated, using the updated budgets. Simulations show that, under various negotiation scenarios, our scheme is more efficient than previously proposed approaches that impose on the “weakest” nodes (those that are further from their targets) to turn off their power completely. More importantly, our scheme leads to a statistical rotation of the set of nodes that achieve their target, independently of the initial budget allocation, and hence is more fair.

Bargaining; Fairness; Interference;

I. INTRODUCTION

Power control, i.e., choosing the transmission power to achieve a specific target, is a widely applied dynamic radio resource management technique which is used for interference mitigation. This may lead to significant benefits, such as increasing the system capacity, decreasing the energy consumption and meeting QoS demands. One common property of all wireless networks is that entities that share a portion of the spectrum are competitors in the sense that their QoS targets, which are usually somehow related with the Signal-to-Interference-plus-Noise-Ratio (SINR), may not be simultaneously achievable. It is worth mentioning that this case arises quite often, even in small topologies [2].

In this work, we examine these challenging cases in an abstract context that captures various wireless setups: (i) different network operators that serve their clients through their Wi-Fi, (ii) owners of adjacent apartments that have installed Femtocell APs, (iii) a standard cellular network, where BSs/MNs transmit and interfere with each other, etc. A

game theoretic approach seems suitable to represent the interaction among these directly interfering entities.

The following analogy is one way of visualizing the problem: N couples of persons drink their coffees in a cafeteria and discuss with each other. As it is a noisy place (due to the loud music and the simultaneous conversations), it is difficult for the couples to communicate smoothly, even if they talk as loud as they can. If a couple is unsatisfied by the quality of its communication (the threshold for a successful communication may be different from couple to couple, depending on the nature of the discussion), it may decide to find another couple which is also unsatisfied to find out a solution that would be beneficial to both of them. So, this couple should try to “convince” the other to discuss more quietly (for example, they may treat them drinks).

The above problem corresponds to a weighted congestion game. In such games, each player chooses a resource over a finite set of resources and experiences a cost which depends on the number of players that use this resource (but the influence from different players is-generally-different) [6]. In this work, we investigate a setting in the context of a wireless network that consists of directly interfering entities that, apart from following a power control algorithm, may negotiate about the level of their transmission powers when not satisfied. Following our previous work in [2], we revisit our Bargaining Foschini-Miljanic algorithm (*BFM*), a joint power control and bargaining heuristic algorithm. *BFM* combines the famous SINR-based power control algorithm of Foschini and Miljanic (*FM*) [3] with a bargaining approach that aims at maximizing the number of entities that achieve their SINR targets, in cases where *FM* guarantees that it is impossible for all of them to achieve their SINR targets. More specifically, this work extends [2] in the following aspects: (i) we propose *DBFM*, a distributed version of *BFM*; (ii) *DBFM* generalizes *BFM*, as it permits negotiations among all entities that have not achieved their targets (instead of just two in *BFM*); (iii) we further examine and justify our findings on the long term fairness property of (*D*)*BFM* in conjunction with the initial budget allocation of the nodes; and (iv) we compare our algorithm with Utility-Based Power Control (*UBPC*) [7], a well-known non-cooperative power control algorithm, and show that, under various negotiation scenarios, our algorithm leads to better results in terms of achieved SINR per node.

II. REVIEW OF RELATED WORK

Consider a setting in which a set of links, all susceptible to thermal noise, share a common channel. We are interested in finding a “target” power vector \mathbf{P}^t so that the SINR for each link i is above or equal to a predefined target γ'_i . In [3], Foschini and Miljanic were the first who provided a distributed algorithm that can find \mathbf{P}^t every time that such a vector exists. Their work was subsequently simplified by Bambos [1] and Grandhi, Zander and Yates [4], who showed the following form of the Foschini-Miljanic formula, which includes a P_{\max} constraint and is applied by each link during every round k of the algorithm:

$$P_i(k+1) = \min \left\{ P_{\max}, \gamma'_i \frac{P_i(k)}{\text{SINR}_i(k)} \right\}. \quad (1)$$

Their algorithm, which we call the *FM* algorithm, is fully distributed, as there is no need for communication (and cooperation) among links. From (1), we can notice that when $P_i(k+1) = P_i(k)$, transmitter i has either achieved its SINR target ($\gamma'_i = \text{SINR}_i(k)$) or it is below its target and transmits with P_{\max} .

For the cases that a topology does not lead to a feasible solution (i.e., at least one link cannot reach its SINR target), some interesting approaches have been proposed in [7] and [8]. In [8], Zander proposed that, when the state is infeasible, the weakest link (the one that is furthest from its SINR target) should power off. We call this approach the *Trunc FM* power control algorithm, as $N-1$ links update their powers according to *FM*, whereas one link powers off. This is a partially distributed algorithm, as links should cooperate to find out the weakest one. In [7], Xiao, Shroff and Chong formulated a non-cooperative game where links adapt their SINR targets (which are now “soft”) to the channel conditions. As a link feels more interference, it decreases its SINR target and may even power off. Pricing of the transmission power is used as a cost function to encourage the links to adjust their targets. These policies, though effective, are both not appealing and quite difficult to implement in practice, as autonomous entities are not obliged to power on/off based on the instructions of an external entity.

Note the relation of our problem to weighted congestion games (e.g., [5]). The congestion games framework is well suited to model resource competition where the resulting payoff (e.g., the cost) is a function of the level of congestion (e.g., the number of concurrent users). However, a key assumption underlying the congestion game model is that all users have an equal impact on the congestion [6]. This is not true in wireless networks in general. A recent application of congestion games in the spectrum sharing problem is presented in [5]. They consider as resource a so-called spectrum-space block, where the space corresponds to the coverage of a BS and the spectrum corresponds to the channel that this BS has chosen to transmit into. However, the transmission power is considered constant. Instead of using such “artificial” resources, we believe that one should consider weighted congestion games, where the impact of each player onto others is different, having a direct equivalence with the problem that we treat here.

III. DBFM: THE DISTRIBUTED BARGAINING FOSCHINI-MILJANIC ALGORITHM

DBFM works on top of *FM* for the cases that a topology does not lead to a feasible solution. So, it takes as input the last state of *FM* where some transmitters have just achieved their SINR targets, whereas those that have failed transmit with P_{\max} . It is natural that the satisfied nodes are not interested in participating in any sort of negotiation, as they have achieved their targets. On the other hand, a transmitter that is below its target would definitely like to convince one or more transmitters to reduce their transmission powers so as to improve its SINR, as the interference that it perceives would decrease. For this reason, we consider that each link (sometimes, by “link” we mean “the owner of a link”) has some budget available, i.e., real or virtual currency, which is used during the negotiations among the unsatisfied entities.

Algorithm 1 presents all steps of the *DBFM* algorithm for the particular case that it focuses on finding a feasible solution just for $N-1$ links (a so called “ $(N-1)$ -feasible” solution [2]). Each unsatisfied link independently chooses its status, i.e., whether it is interested in either receiving an offer to reduce its transmission power (“Seller”) or making an offer asking another node to reduce its transmission power (“Buyer”). Afterwards, it broadcasts its status and provided that there is at least one “Buyer” and at least one “Seller”, each “Buyer” B chooses one “Seller” S to start negotiations. It is clear that many “Buyers” may choose the same “Seller” to negotiate with. Following the ideas in [2], B computes the minimum reduction needed from S to achieve its SINR target in the next iteration of the algorithm. If this is possible, it makes an offer to S , which depends on both B ’s distance from its target and the percentage power reduction that it asks for. On the other hand, S collects all the offers that it receives and decides with a similar approach, which is the maximum power reduction (from those that it has been asked for) that is willing to accept. If this reduction is non-zero, S collects the rewards from *all* the “Buyers” that have asked for power reduction up to that level and reduces its transmission power to that level. If this is the case, “Buyers” whose offer has been accepted apply *FM* (and are guaranteed to achieve their SINR targets in the next iteration of the algorithm). If a Seller has either received no offer or did not accept any offer, it simply applies *FM*.

“Buyers” whose offer has been rejected reduce voluntarily their transmission power by a factor of c . The same policy is applied in the case that there are only “Buyers”, or only “Sellers”, as it is impossible for them to find somebody to negotiate. Though this voluntary reduction in power seems counterintuitive, we have argued that, under certain conditions, it may be proven beneficial for those that apply it. The reason is that unsatisfied nodes have another opportunity to negotiate (probably with a different node) in a following round, and this negotiation may be successful [2]. In any case, should they simply follow *FM*, chances are that all powers will remain at the same level, so the algorithm will unsuccessfully terminate at the last “if” statement of Algorithm 1.

Algorithm 1: The Distributed Bargaining Foschini-Miljanic Algorithm (DBFM)
 Input: Transmission Power Vector \mathbf{P}_{FM} of the last state of *FM*, Noise n , SINR Target γ'_i of each link, P_{\max} .

for $k=1$ to Max_Number_of_Iterations

if (# links that are below their SINR targets) = 1 break;
while (# links that have achieved their SINR targets) < N-1

1. Each link i that has achieved its SINR target applies *FM*
2. Each link i that is below γ'_i , independently decides whether it is a “Seller” or a “Buyer” and broadcasts its status to the network
3. if (# “Sellers” ≥ 2 && # “Buyers” = 0) || (# “Sellers” = 0 && # “Buyers” ≥ 2)
 - 3.1 Each “Seller” (“Buyer”) voluntarily reduces its power by a factor of c
4. if # “Sellers” ≥ 1 && # “Buyers” ≥ 1
 - 4.1 Each “Buyer” B_i selects at random one “Seller” S to negotiate with
 - 4.2 B_i computes the desired reduction Pred_{B_i}

$$\text{Pred}_{B_i} = \max \left\{ 0, \frac{P_{S_{\text{request}}}}{P_S(k)} \right\} \quad (2)$$

$P_{S_{\text{request}}}$ is the maximum possible transmission power of S

so that B_i can achieve its SINR target γ'_{B_i}

- 4.3 if $\text{Pred}_{B_i} = 0$, B_i makes no offer and voluntarily reduces its power by a factor of c
- 4.4 if $\text{Pred}_{B_i} > 0$, B_i is willing to offer $R_{B_i \rightarrow S}$ units from its current Budget $BUD_{B_i}(k)$ to S , where

$$R_{B_i \rightarrow S} = BUD_{B_i}(k) \frac{\text{SINR}_{B_i}(k)}{\gamma'_{B_i}(k)} \text{Pred}_{B_i} \quad (3)$$

- 4.5 S collects all the offers that receives and compares each offer with the quantity

$$R_{S \rightarrow B_i} = BUD_S(k) \frac{\text{SINR}_S(k)}{\gamma'_S(k)} \text{Pred}_{B_i} \quad (4)$$

- 4.6 if $R_{S \rightarrow B_i} \leq R_{B_i \rightarrow S}$, then S accepts the offer
 - 4.7 S computes the final reduction and updates its budget
$$\text{Pred}_{\text{final}} = \max \{ 0, \max \{ \text{Pred}_{B_i} \text{ s.t. } R_{S \rightarrow B_i} \leq R_{B_i \rightarrow S} \} \}$$

$$BUD_S(k+1) = BUD_S(k) + \sum R_{B_i \rightarrow S}, \text{ if } \text{Pred}_{\text{final}} \geq \text{Pred}_{B_i} \quad (5)$$
 - 4.8 if S has accepted at least one offer, its new power is $P_S(k+1) = \text{Pred}_{\text{final}} P_S(k)$, else it applies *FM*
 - 4.9 if B_i 's offer is accepted, B_i applies *FM*, else it voluntarily reduces its power by a factor of c
 - 4.10 if a “Seller” has not received any offer, it applies *FM*
5. if, for each i , $|P_i(k+1) - P_i(k)| \leq \epsilon$, ϵ : a small positive quantity break;

All in all, each unsatisfied node has $M+1$ choices for the next level of its transmission power (where M is the number of unsatisfied nodes): (i) $M-1$ choices to reduce its transmission power at the level that each unsatisfied transmitter will ask for, (ii) one choice to apply *FM* if (a) it accepts no offer or (b) it

does not receive any offer or (c) it makes an offer and its offer is accepted and (iii) one choice to reduce voluntarily its transmission power (in case that it makes an offer and its offer is rejected). Though each “Buyer” chooses randomly a “Seller” to negotiate with, and the levels of the reward that each “Buyer” will offer as well as the percentage power reduction that it will ask for are defined by the above formulas (which is not the case for a classical bargaining scheme), we are inspired by the general bargaining idea and use the term ‘bargaining’ for our scheme.

IV. PERFORMANCE EVALUATION

A. Comparison with UBPC: A Case Study

In this section, we compare *DBFM* with *UBPC* [7] for the same topology and path loss model that the authors of [7] used in their simulations. There are six pairs of transmitters-receivers, where each one should satisfy the respective SINR targets (units are in dB): 12.5, 14, 17, 13.75, 13.5 and 13.

TABLE I LINK NEGOTIATIONS

Round	Buyer	Corresponding Seller	Buyer's Offer	Buyer's Request for % Power Reduction	Seller's Decision
1	Tx ₅	N/A	NO OFFER	NO OFFER	N/A
	Tx ₆	N/A	NO OFFER	NO OFFER	N/A
2	Tx ₅	Tx ₆	171.64	77.99	NO
3	Tx ₁	Tx ₅	819.37	86.74	YES
	Tx ₄	Tx ₅	NO OFFER	NO OFFER	NO
	Tx ₆	Tx ₅	156.24	19.38	YES
4	Tx ₅	Tx ₄	NO OFFER	NO OFFER	NO

TABLE I presents an example scenario that arises after the application of *DBFM* in that topology. It shows the negotiations among the unsatisfied links during each round of *DBFM*. The last column depicts the outcome of the negotiation (all links have an initial budget of 1000 units.). *N/A* means that there is no “Seller” so each “Buyer” (voluntarily) reduces its power by 100*c% (in our case, $c=0.1$). *NO* means that the “Buyer’s” offer is not accepted (or the “Buyer” is not interested in making an offer to that “Seller”). In that case, the “Buyer” reduces its power by 100*c%. *YES* means that the “Buyer’s” offer is accepted. So, the “Seller” should reduce its power at the maximum percentage power reduction that it accepted. In all other cases, nodes simply apply the *FM algorithm*.

Figure 1 shows the SINR evolution, whereas Figure 2 shows the evolution of the transmitter powers (powers are depicted in logarithmic scale). Note that the initial states correspond to the last state of *FM*, where Tx5 and Tx6 have failed to achieve their SINR targets (and they transmit at the maximum power). By following the above “rules”, we can easily trace the changes in transmission power in Figure 2. Note that, after the 2nd round, Tx1 and Tx4 fall below their SINR targets, though they had achieved their targets at the last state of *FM*. This happens as *FM* does not provide an “active link protection mechanism” as that of [1], so that a link that achieves its target would not necessarily retain it during the next iterations of the algorithm. In our scheme, the adoption of an admission control algorithm would not be desirable without having previously found an “(N-1)“-feasible solution. However, we could adopt it on top of *DBFM*, to protect the SINR targets of the links that lead to the “(N-1)“-feasible solution.

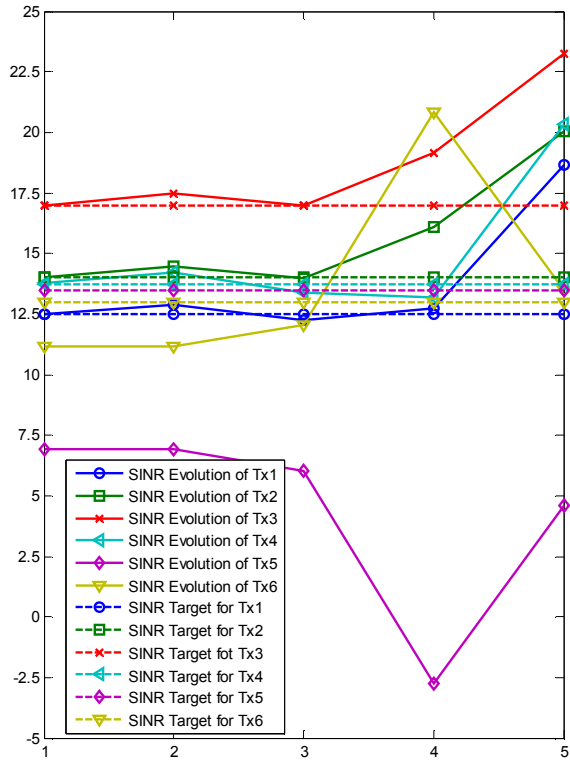


Figure 1 SINR evolution of links after application of *DBFM*. Initial SINRs are the ones of the last round of *FM*. Horizontal dashed lines correspond to the SINR targets of links. After 5 iterations, *BFM* stops and has led to a “(N-1)-feasible” solution as only Tx₄ is below its SINR target.

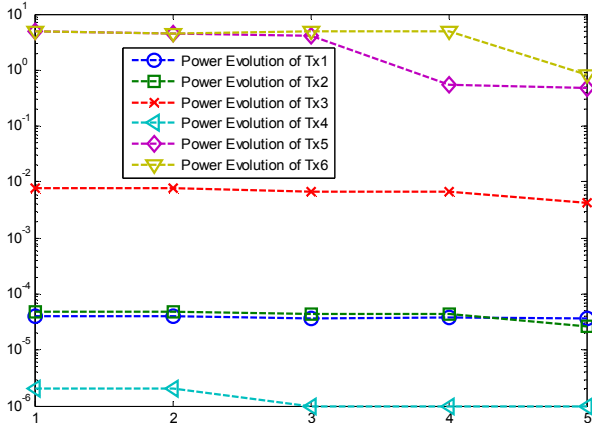


Figure 2 Power evolution of links after application of *DBFM*. Note that the vertical axis is in logarithmic scale. Initial powers are the ones of the last round of *FM*. Tx₅ and Tx₆ change their powers based on the negotiations that are presented in TABLE I, whereas the other nodes apply *FM*.

Our second observation concerns the SINR updates after the 3rd round. We can see that Tx₂ and Tx₃ overcome their targets without having offered any reward (corresponding to a couple in the cafeteria that improves its quality of communication without offering something to another). In addition, the SINR of Tx₆ becomes nearly double the target, even though Tx₆ has asked for the minimum reduction needed (19.38%) to simply “touch” its target. This happens because Tx₁ asked for a bigger reduction (86.74%) and since both offers got accepted, the SINR of Tx₆ is greatly increased

(corresponding to a couple that improves its quality of communication to a greater extent than its offer).

In Figure 3 we compare the performance of *UBPC* and that of *DBFM* in the same topology. Even though both algorithms lead to an “(N-1)-feasible” solution, *DBFM* (without enforcing a link to power off) clearly outperforms *UBPC* as 4 out of the 6 links achieve higher SINR. However, note that our scheme does not incorporate an explicit congestion control mechanism such as “soft” SINR targets or adaptive pricing that *UBPC* adopts. It is an open question whether players may avoid congestion by forming coalitions.

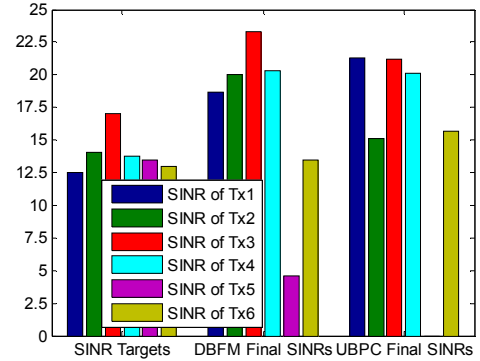


Figure 3 Comparison of the SINR that arise after the application of *DBFM* and *UBPC*. In both cases, 5 out of 6 links achieve their SINR targets. *DBFM* leads to a better SINR for 4 out of 6 links.

B. On the Long Term Fairness of the Algorithms

The update rate of the transmission powers is very fast in modern types of networks (e.g., in a 3GPP WCDMA system, the rate is 1500 Hz). So, nodes that coexist in the same topology will definitely transmit at multiple time intervals (transmission rounds). This implies that a power control scheme should be fair in the following sense: if the same set of entities with the same SINR targets apply long-term the proposed scheme, the set of entities that satisfy their targets should vary over time. Unfortunately, *Trunc FM* and *UBPC* are by design unfair, as they always penalize the weakest node, which will never have the opportunity (not even to try) to transmit. We have shown in [2] that *BFM* fulfills to some degree the above definition of fairness. The idea was that during the (m+1)th transmission round, links reset all their parameters to the last state of *FM* except their budgets, which are the ones that arose after the application of *BFM* at the mth transmission round. So, the rewards that the unsatisfied links may have collected during negotiations of previous rounds could be used to increase their chances to make an offer that will get accepted to a following round. However, the number of transmission rounds was small (just 100) and the impact of the initial budget allocation was not taken into account.

Figure 4 presents the results of *DBFM* when applied in the topology of the previous section for 10000 transmission rounds. Every period of 100 transmission rounds, we count the number of times that each transmitter failed to achieve their target. In our example, these are Tx₅ and Tx₆. There is an average ratio 3:2 per period, i.e., 60 out of 100 transmission rounds Tx₅ powers off, whereas the remaining 40 Tx₆ powers off. Most importantly, even by looking at a random transmission period, this ratio is a good approximation of the

rotation of the satisfied nodes. This is an important advantage of *DBFM*, as all nodes get the opportunity to transmit their data. However, one may think that this result is strongly dependent on the initial budget of the nodes (which was set to 1000 units for all nodes). Further examples with random initial budgets for all nodes of the topology verify the same trend.

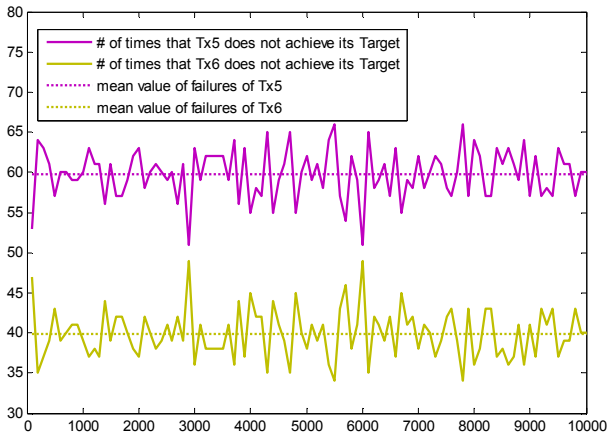


Figure 4 We apply *DBFM* for the same set of nodes, by resetting their parameters to the last state of *FM*. The budget at the $(m+1)^{\text{th}}$ round is the one that arises at the end of the m^{th} round. For every period of 100 transmission rounds, we count how many times Tx5 and Tx6 do not achieve their targets.

This is justified if we think about the dynamically adjusting mechanism that nodes follow when they either make or evaluate an offer. Each “Buyer” computes its offer in terms of *percentage* of its current budget, *not* as an absolute value of the form “I offer up to X units to ask for a reduction up to Y%”. Each “Seller” follows the same strategy too. This is a fair system as: (i) each “Buyer” offers the same percentage of its budget when it asks for the same percentage reduction (ii) each “Seller” rejects/ accepts an offer based on the reward that it would have offered had it asked for the same transmission power. This means that even the application of *DBFM* with small initial budgets is sufficient to lead to the same results.

C. On the Number of “(N-1)-feasible” Solutions

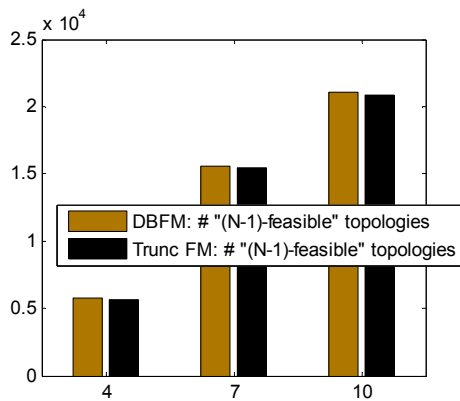


Figure 5 This figure presents the number of “(N-1)-feasible” solutions after the application of both *DBFM* and *Trunc FM* for 10000 different scenarios.

We now compare the number of “(N-1)-feasible” solutions that *DBFM* and *Trunc FM* find. Simulations parameters (path loss model, nodes distribution etc.) are provided in Table I of [2]. As shown in Figure 5, *DBFM* and *Trunc FM* find nearly

the same number of “(N-1)-feasible” solutions. This means that *DBFM*, even though satisfied nodes still receive interference from unsatisfied one(s), has the same performance with a policy that is not appealing as autonomous entities are obliged to follow the instructions of an external entity.

V. CONCLUSIONS AND FUTURE DIRECTIONS

This paper presents *DBFM*, a heuristic algorithm that extends and generalizes *BFM* [2], where we combined *FM* [3] with a bargaining approach, allowing negotiations among interfering links that have not achieved their SINR targets to decide upon their new transmission power levels. *(D)BFM* aims at minimizing the number of unsatisfied entities in cases where *FM* guarantees that it is impossible for all of them to achieve their SINR targets. Our new algorithm, *DBFM*, is a distributed approach towards that direction. Additionally, whereas in [2] only one negotiation is allowed per transmission round, *DBFM* evolves the negotiation mechanism providing the chance to all unsatisfied nodes to negotiate for their powers. Though we have not included a quantitative comparison of *DBFM* vs. *BFM*, *DBFM* outperforms *BFM* due to these extra features. Furthermore, we extend the discussion on the fairness property (i.e., the rotation of the links that satisfied their targets) that arises from the node negotiations.

Another contribution of our paper is that we discuss the correspondence of our model with a weighted congestion game. Due to the fundamental complexity of the wireless setting, applying the full machinery of weighted congestion games is a formidable task. As future work, we plan to further investigate that correspondence and apply findings of algorithmic game theory to our problem. Finally, we shall evaluate the robustness of our scheme considering that a percentage of the users are malicious.

ACKNOWLEDGMENT

Vaggelis G. Douros is supported by the HERAKLEITOS II Programme which is co-financed by the European Social Fund and National Funds through the Greek Ministry of Education.

REFERENCES

- [1] N. Bambos, “Toward power-sensitive network architectures in wireless communications: concepts, issues, and design aspects,” *IEEE Personal Communications*, vol. 5, no. 3, June 1998.
- [2] V.G. Douros, G.C. Polyzos and S. Toumpis, “A Bargaining Approach to Power Control in Networks of Autonomous Wireless Entities,” Proc. 8th ACM International Symposium on Mobility Management and Wireless Access (*MobiWAC*), Bodrum, Turkey, Oct. 2010.
- [3] G.J. Foschini and Z. Miljanic, “A simple distributed autonomous power control algorithm and its convergence,” *IEEE Trans. on Vehicular Technology*, vol. 42, no. 4, Nov. 1993.
- [4] S.A. Grandhi, J. Zander and R. Yates, “Constrained power control,” *Wireless Personal Communications (Kluwer)*, vol. 1, no. 4, Dec. 1994.
- [5] M. Liu and Y. Wu, “Spectrum sharing as congestion games,” Proc. 46th Annual Allerton Conference on Communication, Control, and Computing, Urbana-Champaign, IL, Sep. 2008.
- [6] P.N. Panagopoulou and P.G. Spirakis, “Algorithms for pure Nash equilibria in weighted congestion games,” *ACM Journal of Experimental Algorithmics*, vol. 11, no. 1, Feb. 2007.
- [7] M. Xiao, N.B. Shroff and E.K.P. Chong, “A utility-based power-control scheme in wireless cellular systems,” *IEEE/ACM Trans. on Networking*, vol. 11, no. 2, April 2003.
- [8] J. Zander, “Distributed cochannel interference control in cellular radio systems,” *IEEE Trans. on Vehicular Technology*, vol. 41, no. 3, Aug. 1992.