

Efficient and Truthful Bandwidth Allocation in Wireless Mesh Community Networks

Fabio Martignon, *Member, IEEE*, Stefano Paris, *Member, IEEE*, Ilario Filippini, *Member, IEEE*, Lin Chen, *Member, IEEE*, and Antonio Capone, *Senior Member, IEEE*

I. INTRODUCTION

WIRELESS mesh networks (WMNs) have emerged in recent years as a promising communication paradigm toward the cost-effective deployment of all-wireless network infrastructures [1]. Several operators have started using WMNs as a valuable technology to provide broadband Internet access in urban and rural areas, where the low return on investments cannot cover all costs to deploy more expensive wired solutions. With the aim of further reducing the overall maintenance costs and maximizing the profit, WMN operators have been

fostering the deployment of wireless mesh community networks (WMCNs) [2]. In WMCNs, a group of independent mesh routers owned by different individuals forms or extends a WMN to enhance the broadband connectivity, whose availability can be shared with other users not directly involved in the management of the community network.

In this context, we envision a marketplace scenario where an operator may lease the bandwidth of its wireless access network to a subset of customers in order to increase the network coverage of its WMN and provide access to other residential users through the customers' mesh client devices. The customers¹ who manage these mesh clients pay the network operator to exploit the access bandwidth, while they are rewarded directly by the residential users they serve. Note that both the operator and the customers gain from this agreement since the former can lease the bandwidth of its WMN, saving management and maintenance costs, while the latter can earn money by subleasing the purchased bandwidth to other residential users. Finally, the residential users that would not have been covered by the WMN operator (because of low payoffs) obtain a better Internet service. The proposed marketplace would therefore contribute to overcome the Digital Divide problem, improving the economical efficiency of public-private wireless partnerships like those analyzed in [3].

In order to be an attractive solution, the aforementioned bandwidth market managed by the WMN operator needs convincing allocation and payment mechanisms that should act as incentives for customers to participate and subscribe to the service. One of the main problems that might discourage a WMN operator from developing the bandwidth marketplace is the possibility that even few dishonest customers misbehave. Specifically, a customer could strategically bid false offers, thus manipulating the market as it prefers, in order to pay a lower price or rule out honest customers. These adversarial behaviors reduce the operator's revenue.

Motivated by the above analysis, we present in this paper an *economically efficient*² and resilient auction-based bandwidth allocation in WMNs. Our particular emphasis is on the resilience of the proposed mechanism against any actions of selfish customers that manipulate the bandwidth marketplace of the network scenario described above to obtain extra benefit.

Manuscript received March 09, 2012; revised January 28, 2013 and August 08, 2013; accepted November 24, 2013; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor K. Kar. This work was supported in part by the Italian MIUR and French ANR in the framework of the PRIN Gatecom and ANR Green-Dyspan projects.

F. Martignon is with the Laboratoire de Recherche en Informatique (LRI), Université Paris-Sud, 91405 Orsay, France, and also with Institut Universitaire de France (IUF), 75005 Paris, France (e-mail: fabio.martignon@lri.fr).

S. Paris is with the Laboratoire d'Informatique Paris Descartes, Université Paris-Descartes, 75270 Paris Cedex 06, France (e-mail: stefano.paris@parisdescartes.fr).

I. Filippini and A. Capone are with the Dipartimento di Elettronica Informazione e Bioingegneria, Politecnico di Milano, 20133 Milan, Italy (e-mail: ilario.filippini@polimi.it; antonio.capone@polimi.it).

L. Chen is with the Laboratoire de Recherche en Informatique (LRI), Université Paris-Sud, 91405 Orsay, France (e-mail: lin.chen@lri.fr).

Color versions of one or more of the figures in this paper are available online.

¹The *customers* are residential users that operate the mesh client devices. They connect directly to the provider's network and resell connectivity to other residential users.

²In the rest of the paper, the term "efficiency" refers to the economic efficiency when not otherwise specified.

To tackle this problem, we design an *optimal truthful auction* that forces each customer interested in leasing the available bandwidth to bid its real valuation of the required bandwidth demand.

More specifically, the approach consists in finding the optimal set of customers to be accepted by the operator (auction winners), whose traffic demands can be routed through the WMN, and the corresponding prices they have to pay for the leased service, which constitute the operator revenue. The optimal allocation and the pricing together ensure the truthfulness (also known as *incentive compatibility*) of the proposed auction scheme.

Despite the optimality and truthfulness of the developed auction mechanism, we show that finding such optimal allocation is NP-hard. Hence, we further propose a greedy algorithm that implements the auction and guarantees that bidding its real valuation is the best strategy for each participating customer. We also demonstrate theoretically that the proposed greedy algorithm satisfies the truthfulness property. Through extensive numerical studies, we show that the proposed greedy algorithm achieves a performance very close to the system optimum in a social perspective.

Existing research works, which investigate the use of auction theory to design efficient mechanisms for resource allocation, do not accurately capture the main features of wireless multihop networks and do not take into account the high computational time needed to carry out the auction. On the contrary, our scheme selects efficiently (i.e., in polynomial time) the winners considering both the *link utilization* necessary to satisfy the demands of the customers that participate to the auction and the *routing constraints* of wireless multihop transmission technologies.

In an effort to design an efficient marketplace for allocating the WMN's available bandwidth, our work makes the following unique contributions.

- We propose and analyze an innovative marketplace for the allocation of the WMN's available bandwidth to those customers who are willing to pay more for sharing the purchased bandwidth with other residential users.
- We propose a combinatorial truthful auction that maximizes the revenue of the WMN operator, which is resilient against any market manipulation and guarantees a fair allocation of the resources.
- We design a greedy algorithm to compute efficiently customer allocations and fair payments, which still guarantees that participating customers bid their real valuations. The proposed algorithm consists therefore in an alternative yet truthful auction mechanism.
- We perform a thorough numerical analysis of the proposed algorithms, including large-scale, real WiFi network scenarios (like the Google WiFi network [4]).

The rest of this paper is structured as follows. Section II discusses related work. Section III presents the communication and network models considered in our work. Section IV formulates the combinatorial auction as an optimization model, while Section V illustrates the greedy algorithm that we propose to efficiently compute the solution. The incentive compatibility property as well as the economic efficiency of the greedy algorithm are analyzed in Section VI. Section VII provides a

numerical evaluation of the proposed framework. Finally, conclusions are discussed in Section VIII.

II. RELATED WORK

Auction theory has been used to design efficient allocation mechanisms in several network contexts, such as cognitive radio networks, selfish routing, and resource allocation. Hereafter, we review the most relevant recent literature, highlighting the main differences with respect to our approach.

With the upcoming generation of cognitive radio networks, market-based auctions have been extensively studied as an efficient mechanism to dynamically sublease the unexploited licensed spectrum to secondary users and increase the revenue of the spectrum owner [5]–[11].

Auction theory has been exploited to design innovative traffic engineering techniques and routing protocols, both to enhance the utilization of unused network paths and force the collaboration of intermediate relaying nodes [12]–[19].

Ad Hoc-VCG [12] is a routing protocol based on the Vickrey–Clarke–Groves (VCG) auction, which guarantees that each intermediate node is refunded at least the true cost incurred to relay packets. The Commit algorithm [13] further develops this approach assuring that even the source node behaves correctly. iPass [20] adopts a similar approach, modeling the forwarding capability of each node as a market, where an auction process is used to determine the optimal price for the available resources. The performances of the previous incentive-based schemes are analytically evaluated by Jaramillo *et al.* in [14]; the analysis of their basic properties led to the design of DARWIN, a new protocol robust to imperfect measurements and collusion attacks. In [15] and [16], the truthful pricing mechanism proposed by Vickrey, Clarke, and Groves is used to solve a broad class of problems concerning the noncooperative behavior of intermediate nodes. Similar mechanisms are adopted in [17] and [18] to study and design innovative protocols for multicast transmissions in noncooperative networks, where each node exhibits selfish behavior. Specifically, the authors identify general properties to decide whether an incentive compatible mechanism can be defined on the top of any multicast protocol, and they present a solution to implement the proposed scheme in a distributed fashion. Zhong *et al.* in [21] exploit two solution concepts defined in game theory to consider also the collusion among network devices: They show that even if a Group Strategy-proof Equilibrium cannot be reached at the routing level, their proposed solutions reach Strong Nash Equilibria among network nodes, which are robust to deviations of any component of the colluding group.

We underline that the *tit-for-tat* strategy as well as its different variants such as the *generous tit-for-tat* (GTFT) [14] may not be robust in a wireless environment, and they can be exploited by adversaries to steer the system toward an inefficient equilibrium state.

Works sharing a similar approach to the solutions described in this paper have been recently proposed in [22] and [23]. In particular, Jain *et al.* in [22] present a mechanism for per-link bandwidth allocation of end-to-end paths in wired networks, whereas Fu *et al.* in [23] design an auction-based stochastic game for resource allocation of virtual operators in wireless cellular networks. However, these works do not accurately capture the main

TABLE I
BASIC NOTATION USED IN THE PAPER

Acronyms and Sets	
MC	Mesh Client
MAP	Mesh Access Point (m_a is the number of MAPs)
MR	Mesh Router
MGW	Mesh Gateway
\mathcal{N}	Set of Mesh Clients (i.e., Customers), $ \mathcal{N} = n$
\mathcal{M}	Set of WMN Devices (MR, MAP, and MGW), $ \mathcal{M} = m$
$\mathcal{M}_{C,i}$	Set of MRs operating as MAPs in the radio range of MC i
$\mathcal{G} \subset \mathcal{M}$	Set of MRs that act as MGWs
\mathcal{L}	Set of wireless links among MRs

Parameters	
C_j	Capacity of the wired link of MGW j
c_{jk}	Capacity of the wireless link (j, k)
d_i	Bandwidth demand of MC i
b_i	Bid offered by MC i for demand d_i
v_i	Real valuation of MC i for demand d_i
p_i	Price paid by MC i for demand d_i
u_i	Utility of MC i for demand d_i
$r_{ij}^{(max)}$	Maximum transmission rate of the wireless link established between nodes i and j
o_{ij}	Channel utilization of MAP j to satisfy the demand of MC i

Variables	
x_i	0-1 variable that indicates whether the demand d_i of MC i is satisfied
y_{ij}	0-1 variable that indicates if MC i is assigned to MAP j
f_{jk}	Flow variable which denotes the traffic flow routed on link (j, k)
f_j	Flow variable which denotes the traffic flow routed on wired link of MGW j

features of wireless multihop networks like the variable transmission rate that adapts to the channel condition and the limited capacity of the backbone network. Furthermore, these two latter solutions do not take into account the very large computational time needed to solve the considered auction in realistic network scenarios.

Finally, we remark that, unlike existing works that use auctions as a technique to solve allocation problems, our proposed solutions select the winners considering the *resource utilization* of the requested services instead of their *raw requests*, by modeling both the rate adaptation mechanism operating at the MAC layer and the capacity limits due to the presence of interfering links within a WMN. As we verified in the network scenarios illustrated in Section VII, this permits to increase the revenue of the WMN operator since between two competitors, who demand the same bandwidth, our schemes select the one that requires the lowest resource utilization, leaving room for the allocation of additional customers.

III. SYSTEM MODEL

This section presents the communication and network models considered in our work, as well as the definitions and assumptions we adopt in the design of our auction mechanism.

Let us refer to the WMN scenario illustrated in Fig. 1, where the WMN is managed by a single operator that leases the bandwidth made available through its mesh access points (MAPs) to a subset of customers, which connect to the WMN through their mesh clients (MCs).

The mechanism we propose implements the bandwidth marketplace by allocating the available WMN capacity to a subset of customers, which in turn may sublease it to other residential users. Table I summarizes the basic notation used throughout the paper.




Fig. 1. Wireless mesh network scenario considered in this work. The WMN is managed by a single operator that leases the available bandwidth of the MAPs to customer MCs. Some mesh routers (MRs) act as mesh gateways (MGWs) to provide access to the Internet.

Each mesh client³ i has a bandwidth demand d_i that he wishes to satisfy by transmitting to one of the MAPs that cover it with their wireless signal. We assume, without loss of generality, that the term d_i accounts for the traffic demand of both the downlink and uplink since the wireless resource is a half-duplex channel. The uncertainty related to traffic description in 802.11 wireless systems can be broadly characterized by three parameters, namely: 1) its *burstiness*; 2) the *packet length distribution*; and 3) the *contention* level at the frame layer, which, in turn, is closely related to the collision probability.

The *first two parameters* are used by each MC to control all quality-of-service requirements that may affect its valuation by defining an equivalent flow bandwidth, as discussed in [24]. On

³In this paper, we use interchangeably the terms *customers* and *mesh clients* since the customers are the owners of mesh clients.

the other hand, the *contention* level, which is function of the traffic requirements of all selected bidders, is controlled by the WMN operator by computing the margin necessary to guarantee the requested effective bandwidth and avoid the throughput collapse caused by high contention on the wireless channel, using for example the model presented in [25]. Note that if we reconfigure the access scheme using contention-free MAC protocols (e.g., TDMA) exploiting architectures like those proposed in [26], such margin can in theory tend to zero.

To satisfy such demand, each buyer bids an offer b_i for its bandwidth demand to the WMN operator. This latter decides which MCs are served and the price that winners have to pay to exploit the available bandwidth.

We further assume that WMN devices (i.e., MRs, MAPs, and MGWs) are equipped with multiple radio interfaces and the operator designed its network to minimize intraflow and interflow interference effects according to optimization strategies like those proposed in [27]–[29]. Since MAPs use orthogonal channels, the different subsets of MCs assigned to each MAP do not interfere with each other.

We observe that the transmission rate and the channel utilization required to satisfy the MC's demand depend, clearly, on device technologies, but in particular on the distance between the mesh client and the mesh access point to which it is connected; hence, the allocation mechanism has a direct impact on the number of mesh clients that have the opportunity to exploit the available bandwidth. Therefore, the aim of the WMN operator is to increase its revenue by allocating the available bandwidth of its mesh access points to those mesh clients that are willing to pay the highest price for the channel utilization. To this end, we design a truthful auction that, in addition to maximizing the revenue of the WMN operator, prevents market distortion by forcing every mesh client to declare and bid its true valuation, $v_i = b_i$.

Each mesh client i submits its bid in the form (b_i, d_i) , where b_i represents the price that the buyer i is willing to pay for its bandwidth demand d_i . For each possible allocation of MC i to MAP $j \in \mathcal{M}$, the operator computes the corresponding channel utilization o_{ij} as the ratio between the required bandwidth demand d_i and the maximum achievable transmission rate of the wireless link that might connect MC i and MAP j , $r_{ij}^{(\max)}$, according to the following equation:

$$o_{ij} = \frac{d_i}{r_{ij}^{(\max)}}. \quad (1)$$

Note that $r_{ij}^{(\max)}$ can be easily obtained from the MAC layer through a scanning of the wireless channels, which is performed periodically by all network devices. Furthermore, our mechanism can be applied also to network scenarios with time-varying capacity of wireless links by simply considering the cumulative distribution of the transmission rate of any wireless link connecting MC i and MAP j .

Let us denote by p_i the price paid by user i when its demand is satisfied. Then, assuming a quasi-linear utility function for each customer i [30], we can define the utility of i , u_i , as the difference between its private valuation v_i and the price paid to exploit the bandwidth, p_i , according to the following expression:

$$u_i = \begin{cases} v_i - p_i, & \text{if } i\text{'s demand is satisfied} \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

According to (2), user i would obtain a positive utility only if its whole demand d_i is satisfied by the operator. Therefore, if i reports a smaller demand $d'_i < d_i$, its utility would be null since the operator will allocate to user i exactly d'_i . Obviously, when the demand of MC i is not satisfied, its utility is null since both the paid price and its valuation are null.

We focus on the practical scenario where the WMN operator has only a limited and imperfect knowledge about the real valuation that mesh clients are willing to pay for satisfying their traffic demand.

Mathematically, we apply the Myerson's work [31] and model the operator's uncertainty about the real valuation of any mesh client i as a continuous distribution function $F_i(x)$, which satisfies the regularity property, over a finite interval $x \in [a_i, b_i]$, with corresponding probability density function $f_i(x)$.

Note that the distribution describing the uncertainty of the WMN operator can be predicted exploiting the *a priori* knowledge of the system. Indeed, we can confidently suppose the use of an automatic system for implementing the bidding strategy of any customer (i.e., MC owner), like in Web-based marketplaces. In this system, a customer defines its preferences choosing appropriate lower and upper bounds on its bandwidth demands and valuations, so that the system can automatically bid on behalf of the customer.

The user valuation distribution modeling the operator's uncertainty can be assumed independent from quality parameters of the wireless link used to satisfy the demand, like the channel or traffic randomness, since such quality parameters can be simply controlled by defining an equivalent flow bandwidth as discussed in [24].

To design a revenue-maximization truthful auction, we therefore optimize over the virtual valuation function of customer i defined in

$$\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} \quad (3)$$

which we assume to be a monotone nondecreasing function. Note that the virtual valuation function represents the marginal revenue obtained by satisfying the demand of i [32].

Since we assume independence across the cumulative distribution functions modeling uncertainty on mesh clients valuations and *single-dimensional* settings, as illustrated in [31] and [33], such an auction can be implemented by assigning the bandwidth to the customers with the highest virtual valuations ϕ_i , provided they are nonnegative. The payment rule is as follows: The winners pay the smallest value b_k that would result in their winning, that is, the bid (thus the valuation v_k , as it is a truthful auction) of the first excluded customer. If the first excluded customer has a valuation v_k such that $\phi_k(v_k) < 0$, then the winner pays $b_r = \phi^{-1}(0)$, that is, the bid whose corresponding virtual valuation is 0. The value b_r serves as the *reservation price* for the auctioneer since he does not sell anything for bids below this value.

Finally, we underline that the problem formulation can be extended to consider a dynamic scenario where the WMN operator has the choice of reserving part of his bandwidth for a future sale in order to maximize the profit by anticipating the arrival of new mesh clients before the next auction round. However, the

analysis of the best strategies for reserving the optimal amount of bandwidth is out of the scope of this paper.

IV. OPTIMAL AND TRUTHFUL BANDWIDTH AUCTION

This section presents the combinatorial auction mechanism we propose to allocate the available access bandwidth of a WMN operator, maximizing its expected revenue. We formalize the optimal and truthful auction mechanism in two steps. First, we present a mixed integer linear programming (MILP) model that gives the optimal solution for the Optimal and Truthful Bandwidth Allocation Problem (OTBAP). Solving OTBAP, we obtain the assignment of MCs to MAPs that maximizes the expected revenue of the WMN operator. Then, we describe the algorithm that, exploiting the allocation of the MILP model, makes the auction truthful. This algorithm computes the price paid by MCs in such a way that the optimal strategy for each mesh client i is to bid its real valuation v_i .

Let \mathcal{N} denote the set of MCs, \mathcal{M} the set of MRs, and \mathcal{L} the set of wireless links (j, k) among MRs j and k such that the two MRs are in their reciprocal radio range (see Table I for reference). In particular, let us define $\mathcal{M}_{C,i}$ as the set of MRs operating as MAPs that are in the radio range of MC i and \mathcal{G} , $\mathcal{G} \subset \mathcal{M}$, as the set of MRs that act as gateways for the WMN to the wired backbone.

We can now introduce the decision variables used in our MILP model to solve OTBAP. Binary variables x_i , $i \in \mathcal{N}$, indicate which MCs win the auction, i.e., the buyers whose demands are satisfied by the allocation mechanism ($x_i = 1$ if the demand of MC i is satisfied, 0 otherwise). Binary variables y_{ij} , $i \in \mathcal{N}$, $j \in \mathcal{M}_{C,i}$, provide the assignment of MCs to MAPs ($y_{ij} = 1$ if MC i is assigned to MAP j , 0 otherwise). Finally, let variables f_{jk} , $(j, k) \in \mathcal{L}$, denote the traffic flow routed on link (j, k) and f_j , $j \in \mathcal{G}$, the traffic flow routed by mesh gateways toward the wired connection (note that this last value is null for mesh routers that do not act as gateways).

Given the above definitions and notation, the OTBAP of the combinatorial bandwidth auction amounts to the following mixed mathematical program:

$$\max \sum_{i \in \mathcal{N}} \phi_i(b_i) \cdot x_i \quad (4)$$

s.t.

$$\sum_{j \in \mathcal{M}_{C,i}} y_{ij} = x_i \quad \forall i \in \mathcal{N} \quad (5)$$

$$\sum_{\substack{i \in \mathcal{N}: \\ j \in \mathcal{M}_{C,i}}} y_{ij} o_{ij} \leq 1 \quad \forall j \in \mathcal{M} \quad (6)$$

$$\sum_{\substack{i \in \mathcal{N}: \\ j \in \mathcal{M}_{C,i}}} d_i y_{ij} + \sum_{\substack{k \in \mathcal{M}: \\ (j,k) \in \mathcal{L}}} (f_{kj} - f_{jk}) = 0 \quad \forall j \in \mathcal{M} \setminus \mathcal{G} \quad (7)$$

$$\sum_{\substack{i \in \mathcal{N}: \\ j \in \mathcal{M}_{C,i}}} d_i y_{ij} + \sum_{\substack{k \in \mathcal{M}: \\ (j,k) \in \mathcal{L}}} (f_{kj} - f_{jk}) = f_j \quad \forall j \in \mathcal{G} \quad (8)$$

$$f_{kj} + f_{jk} \leq c_{jk} \quad \forall (j, k) \in \mathcal{L} \quad (9)$$

$$f_j \leq C_j \quad \forall j \in \mathcal{G} \quad (10)$$

$$f_{jk}, f_q \geq 0 \quad \forall (j, k) \in \mathcal{L}, q \in \mathcal{G} \quad (11)$$

$$x_i, y_{ij} \in \{0, 1\} \quad \forall i \in \mathcal{N}, j \in \mathcal{M}_{C,i}. \quad (12)$$

Algorithm 1: Optimal and Truthful Bandwidth Auction

Input: $\mathcal{N}, \mathcal{M}, \mathcal{G}, \mathcal{L}, d_i, b_i$

Output: x_i, p_i, y_{ij}

- 1 Compute channel utilizations o_{ij} ;
 - 2 Compute virtual bids $\phi_i(b_i)$;
 - 3 $x_i \leftarrow$ Solve the MILP model (4)–(12);
 - 4 **foreach** $i \in \mathcal{N}$ **do**
 - if** $x_i = 1$ **then**
 - $\tilde{p}_i \leftarrow \max_{x_{-i}} \sum_{h \neq i} x_h \phi_i(b_h) - \max_{x_x} \sum_{h \neq i} x_h \phi_i(b_h)$;
 - $p_i \leftarrow \max \{ \phi^{-1}(0), \phi^{-1}(\tilde{p}_i) \}$;
 - else**
 - $p_i \leftarrow 0$;
 - end**
 - end**
-

The objective function (4) maximizes the expected revenue of the WMN operator obtained from the bandwidth auction.

Constraints (5) provide full coverage of all the mesh clients that win the auction. More specifically, if a mesh client i wins the bandwidth auction, then it must be associated only to one mesh access point among the set of those that cover it. These constraints also ensure that only the mesh clients that win the auction can be assigned to a mesh access point. Constraints (6) prevent the allocation of an overall bandwidth demand that cannot be satisfied by a mesh access point.

Constraints (7) and (8) define the flow balance at node j . The term $\sum d_i y_{ij}$ accounts for the total traffic that is assigned to mesh access point j , while the terms $\sum f_{kj}$ and $\sum f_{jk}$ represent the total incoming and outgoing traffic, respectively. The term f_j represents the traffic sent by mesh gateways to the wired backbone.

The set of constraints (9) ensures that the total traffic routed on a link established between two devices j and k does not exceed its capacity, denoted by c_{jk} , while (10) represents the capacity constraints for the wired backbone links, whose maximum capacity is denoted as C_j . Recall that in multichannel multiradio WMNs, wireless interfaces with directive antennas can be tuned to different channels to reduce interference effects on the backbone link capacity.

Finally, constraints (11) ensure the positiveness of the flow variables, while (12) ensures the integrality of the binary decision variables.

Having defined the MILP model representing the optimal auction, we now illustrate the algorithm that forces mesh clients to bid their real valuation.

Algorithm 1 describes the steps performed by the WMN operator to auction its available bandwidth. The algorithm receives as input the parameters that describe the network topology and mesh client bids; these latter are composed of the required demand d_i and the offered value b_i . It produces as output the allocation of mesh clients to mesh access points, y_{ij} , as well as the price p_i paid by each winning mesh client, $i \in \mathcal{N}$: $x_i = 1$, to exploit the required bandwidth.

The algorithm proceeds in four steps. In steps 1 and 2, mesh client demands are transformed into equivalent channel utilizations, and virtual valuations are computed using both the bids actually offered by mesh clients and the valuation distribution functions $F_i(x)$. Step 3 consists in solving the MILP model to

find the allocation that maximizes the expected revenue. Finally, in step 4, the operator computes the prices paid by the winners, which, according to Myerson [31], guarantees a truthful auction. In this step, the function $\max_{x_{-i}}$ represents the solution to the OTBAP maximization problem with the additional constraint $x_i = 0$ (i.e., the solution to OTBAP without considering MC i in the auction).

We observe that the mechanism implemented by Algorithm 1 satisfies the *ex-post individual rationality* property since both mesh clients and the WMN operator perceive nonnegative utility.

We next show that the optimal auction mechanism of Algorithm 1 is NP-hard, which motivates our further proposition of a *greedy*, yet truthful and efficient, allocation algorithm (see Section V) with polynomial-time complexity.

Proposition IV.1 (NP-Hardness of OTBAP): The optimal auction mechanism implemented by Algorithm 1 is NP-Hard.

Proof: We prove the proposition by showing that the Multiple Knapsack Problem (MKP) [34] can be reduced in polynomial time to the MILP model (4)–(12), which describes the OTBAP. To this end, we establish the following polynomial-time procedure to reduce any instance of the MKP to an equivalent instance of the OTBAP.

- 1) For each item i with weight w_i and value v_i of the MKP, we add a new MC $i \in \mathcal{N}$ with $d_i = w_i$ and $\phi(b_i) = v_i$.
- 2) For each knapsack j with capacity W_j of the MKP, we add a new MAP $j \in \mathcal{M}$ and a new MGW $g \in \mathcal{G}$ connected through the link $(j, g) \in \mathcal{L}$ with capacity $c_{jg} = W_j$. Furthermore, we set the capacity of the MGW wired link to $C_g = W_j$.
- 3) For each possible assignment (i, j) of item i to knapsack j in the MKP, we add a new channel utilization $o_{ij} = w_i/W_j$.

Since solving the MKP is at least as difficult as solving OTBAP (MKP \leq_P OTBAP) and MKP is NP-Hard, OTBAP is NP-Hard.

V. GREEDY BANDWIDTH AUCTION

As demonstrated in Section IV, the Optimal and Truthful Bandwidth Auction Problem is NP-Hard. Finding the exact system optimum can be thus extremely time-consuming, especially in large-scale, real wireless network scenarios as those analyzed in our numerical evaluation. Motivated by this observation, in the following we present a greedy algorithm to solve efficiently (i.e., in polynomial time) the bandwidth auction problem while preserving the *truthfulness* property.

We set out by presenting the greedy algorithm and describing its main phases. In Section VI, we then analyze its complexity and formally prove that revealing the real valuation of the available bandwidth is the best strategy for each MC participating in the auction.

Description of the Greedy Algorithm: The greedy auction is depicted in Algorithm 2. Its inputs and outputs are exactly the same as the optimal bandwidth auction illustrated in Algorithm 1. The algorithm is composed of two main phases: 1) *allocation* phase (step 1), which determines the winning mesh clients according to their virtual valuation per channel utilization; and 2) the *payment* phase (step 2), which establishes the price paid by each winner based on the best mesh client,

Algorithm 2: Greedy Bandwidth Auction

Input: $\mathcal{N}, \mathcal{M}, \mathcal{G}, \mathcal{L}, d_i, b_i$

Output: x_i, p_i, y_{ij}

- 1 $(x_i, y_{ij}) \leftarrow \text{Greedy_Allocation_Phase}(d_i, \phi_i(b_i), o_{ij})$;
 $\tilde{p}_c \leftarrow \frac{\phi(b_c)}{o_{c j_c}}$;
 - 2 **foreach** $i \in \mathcal{N} : x_i = 1$ **do**
 $\tilde{p}_i \leftarrow \tilde{p}_c \cdot \sum_{j \in \mathcal{M}_{C,i}} o_{ij} y_{ij}$;
 $p_i \leftarrow \max \{ \phi^{-1}(0), \phi^{-1}(\tilde{p}_i) \}$;
- end**
-

Algorithm 3: Greedy Allocation Phase (Step 1 of Algorithm 2)

Input: $d_i, \phi_i(b_i), o_{ij}$

Output: x_i, y_{ij}

Initialize: $o_j = 1, \forall j \in \mathcal{M}$;

$L \leftarrow \text{Sort} \left((i, j) \in \mathcal{N} \times \mathcal{M}_{C,i}, \frac{\phi_i(b_i)}{o_{ij}}, \text{"non-increasing"} \right)$;

while $L \neq \emptyset$ **do**

$(i, j) \leftarrow \text{Next}(L)$;

if $o_j - o_{ij} \geq 0 \wedge \text{FeasibleSol}(\tilde{x}, \tilde{y}, \tilde{d}) \wedge \phi_i(b_i) \geq 0$ **then**

$x_i \leftarrow 1, y_{ij} \leftarrow 1$;

$o_j \leftarrow o_j - o_{ij}$;

$\text{RemoveAll}(i, L)$;

end

end

$c \in \mathcal{N}$, whose demand is not satisfied. This latter is also referred to as *critical mesh client*, and its virtual bid per channel utilization as *critical value* ($\tilde{p}_c = \frac{\phi(b_c)}{o_{c j_c}}$). The value $o_{c j_c}$ is the lowest channel utilization among the links that the critical mesh clients can establish with the set of its covering MAPs to satisfy its traffic demand.

Note that the truthfulness property guaranteed by the payment scheme proposed for Algorithm 1 is no longer satisfied if the combinatorial auction is not solved to the optimality but only approximated, as shown in [35]. For this reason, we have modified the payment scheme of the greedy algorithm with respect to the optimal one (Algorithm 1), so that revealing the true valuation is still the dominant strategy for all the customers who participate to the approximated bandwidth allocation auction. In particular, the price paid by each winner is now proportional to its resource utilization and the unitary price that the *critical mesh client* is willing to pay.

We underline that the proposed payment and allocation schemes implemented by the greedy algorithm guarantee *ex-post individual rationality*, *efficiency*, and *truthfulness*, as we formally prove in Section VI.

The greedy allocation phase of step 1 is further detailed in Algorithm 3; it first sorts the list L of possible MC-MAP allocations in nonincreasing order of submitted virtual valuation per channel utilization, $\frac{\phi_i(b_i)}{o_{ij}}$. Then, each element of the sorted list is allocated only if its demand, or equivalently its channel utilization, can be satisfied by the corresponding MAP and routed toward any mesh gateway. Thus, $\text{FeasibleSol}(\tilde{x}, \tilde{y}, \tilde{d})$ verifies if the additional bandwidth demand of MC i assigned to MAP j can actually be routed through the WMN backbone toward mesh gateways, without violating the link capacities. To this end, we

develop a procedure to compute in polynomial time the maximum flow that can be routed over the wireless backbone using the well known Push-Relabel algorithm [36].

The choice of such an algorithm is motivated by its complexity, which depends on the number of networks edges; in sparse networks, like optimized multiradio and multichannel WMNs, this results in a faster execution with respect to other algorithms that depend only on the number of vertices.

$FeasibleSol(\tilde{x}, \tilde{y}, \tilde{d})$ in Algorithm 3 computes the maximum flow that can be routed over an extended version of the graph that represents the network topology. The node set of the extended graph V consists of the mesh routers set \mathcal{M} and two fictitious nodes, a source node s and an egress node t , $\{s, t\}$: $V = \mathcal{M} \cup \{s, t\}$. The arc set E comprises all original wireless links \mathcal{L} and includes a new arc from node s to each MAP that has been selected to satisfy the demand of the winners currently selected by the algorithm: $\exists (s, j) \in E, \sum_{i \in \mathcal{N}: j \in \mathcal{M}_{C,i}} x_i \cdot y_{ij} > 0$. The capacity of each arc leaving s is set to the sum of winners demands allocated to the corresponding MAP, i.e., $c(s, j) = \sum_{i \in \mathcal{N}: j \in \mathcal{M}_{C,i}} d_i y_{ij}$. If no mesh client has been allocated to a MAP, the capacity of the corresponding fictitious arc is null.

As for the fictitious node t representing the egress point, we create an additional arc from each MGW to t with a capacity equal to the maximum bandwidth of the wired link connecting the gateway node to the Internet: $\exists (g, t) \in E, \forall g \in \mathcal{G}$. If the maximum flow from s to t computed by $FeasibleSol(\tilde{x}, \tilde{y}, \tilde{d})$ on the extended graph is less than the sum of winners demands, the function returns *false* since the MC demand analyzed in the current iteration cannot be satisfied. Otherwise, the MC currently analyzed by Algorithm 3 can be selected as winner, provided that the virtual valuation is positive ($\phi_i(b_i) \geq 0$) and the available channel utilization of MAP j is enough to satisfy its bandwidth demand ($o_j - o_{ij} \geq 0$). To this end, throughout the iterations of Algorithm 3, the total utilization of each MAP (o_j) is updated and verified in order to keep the sum of allocated demands within the bandwidth limit of the access network formed by the MAPs.

$FeasibleSol(\tilde{x}, \tilde{y}, \tilde{d})$, which implements the Push-Relabel algorithm, has time complexity $O(m^4)$. In the worst case, the extended graph has a number of vertices $|V| = m + 2$ and a number of edges $|E| = l + 2m$, where l represents the number of original wireless links, since there are two fictitious nodes (s and t) in addition to the m WMN devices. In a fully connected network, we have $l = m^2$, and the relation $|V|^2|E| \leq m^4 + 4m^3 + 8m^2 + 8m$ holds, hence $FeasibleSol(\tilde{x}, \tilde{y}, \tilde{d}) \in O(m^4)$.

The feasibility evaluation of the solution obtained by selecting a new winner and performed by $FeasibleSol(\tilde{x}, \tilde{y}, \tilde{d})$ is illustrated in the example network scenario shown in Fig. 2, which illustrates the extended graph used by $FeasibleSol(\tilde{x}, \tilde{y}, \tilde{d})$ to evaluate the feasibility of the solution obtained by selecting a new MC in the sorted list L . Black circles represent all MCs selected as winners, while dashed circles depict the set of MCs assigned to the same MAP, which in turn is represented as a fictitious arc (s, j) from s to the corresponding MAP $j \in \{3, 4, 5\}$, with capacity equal to the overall bandwidth demand used by MCs inside the dashed circle. On the other hand, the capacity of the two fictitious arcs $\{(1, t); (2, t)\}$ is equal to the maximum transmission rate of the wired connections of nodes 1 and 2.




Fig. 2. Example of extended graph used by $FeasibleSol(\tilde{x}, \tilde{y}, \tilde{d})$ to verify the feasibility of the solution obtained by selecting a new MC as winner. The black circles represent the MCs already allocated by Algorithm 3. The capacity of each fictitious arc leaving s is equal to the sum of the MCs demands inside the corresponding dashed circle.

The maximum flow from s to t provides a lower bound on the maximum traffic that can be routed over the backbone. Hence, if the sum of all winners demands is lower than or equal to this value, the traffic can be safely transmitted over the backbone without violating the service bandwidth agreement with the winners.

Note that MC i might be satisfied by multiple MAPs $j \in \mathcal{M}_{C,i}$. However, once i is selected as winner and its demand assigned to the corresponding MAP, all remaining entries in the list representing alternative and feasible allocations are removed by the function $RemoveAll(i, L)$. Therefore, the function $RemoveAll(i, L)$ removes from the list L all elements $\frac{\phi_i(b_i)}{o_{ij}}$ representing the alternative allocation for MC i .

Remark: We observe that the proposed greedy mechanism implemented by Algorithm 2 has time complexity $O(nm^5)$. Indeed, each iteration of the loop in the greedy allocation phase is executed at most $n \cdot m$ times (with $m = |\mathcal{M}|$ and $n = |\mathcal{N}|$), which represent the maximum number of alternative assignments of MCs to MAPs in the worst-case scenario. Since each iteration of the loop has time complexity $O(m^4)$ (recall that $FeasibleSol(\tilde{x}, \tilde{y}, \tilde{d}) \in O(m^4)$), Algorithm 3, and thus Algorithm 2, terminates at most after $O(nm^5)$ steps.

We further observe that the backhaul section of a WMN is usually designed to satisfy the maximum bandwidth of the access network. That is, the channel assignment and the routing are jointly optimized assuming that the devices connected to the MAPs generate traffic at the maximum data rate (i.e., $\sum_{j \in \mathcal{M}_{C,i}} d_i y_{ij} = r_{ij}^{\max} \forall j \in \mathcal{M}$).

In such a case, we can speed up the computation by removing the feasibility check performed by $FeasibleSol(\tilde{x}, \tilde{y}, \tilde{d})$ in Algorithm 3 since we already know that the traffic served by MAPs can be routed through the backbone. Indeed, the complexity of the greedy algorithm would be substantially reduced, resulting in $O(nm^2)$ instead of $O(nm^5)$.

Illustrative Example: To better clarify the payment rule implemented by Algorithm 2, let us consider the illustrative example depicted in Fig. 3, where three MCs request as bandwidth demand $d_i = 24$ Mb/s and submit as virtual bids $\phi(b_1) = \phi(b_2) = 24$, and $\phi(b_3) = 12$ US dollars. The two dashed trapezoids represent the areas covered by the radio signals of the two




Fig. 3. Illustrative scenario to clarify the payment rule of the proposed greedy algorithm. Letters represent MAPs, whereas numbers identify MCs. The two dashed trapezoids represent the areas covered by the radio signals of the two MAPs. The number close to each wireless link (represented as solid line), which may connect an MC to an MAP, is its maximum transmission rate (in megabits per second).

TABLE II

CHANNEL UTILIZATIONS o_{ij} OF THE EXAMPLE DEPICTED IN FIG. 3, FOR ALL POSSIBLE ASSOCIATIONS OF MESH CLIENTS $i \in \{1, 2, 3\}$ TO MESH ACCESS POINT $j \in \{A, B\}$

	1	2	3
A	1	1	-
B	-	0.5	1

MAPs, whereas the solid lines represent the wireless links that may be established between MCs and MAPs. Their transmission rates (in megabits per second) are illustrated in the same figure, near the dashed edges. The corresponding channel utilizations (o_{ij}) are reported in Table II.

For the sake of simplicity, we assume, as in the *Remark* discussed above, that the backhaul network is designed to satisfy the maximum bandwidth of the access network.

The list L , sorted according to the virtual bids per channel utilizations, is defined as follows:

$$L = \{ \langle (2, B), 48 \rangle, \langle (1, A), 24 \rangle, \langle (2, A), 24 \rangle, \langle (3, B), 12 \rangle \}$$

where the element $\langle (i, j), \frac{\phi(b_i)}{o_{ij}} \rangle$ represents the link (i, j) and the corresponding virtual bid per channel utilization.

Algorithm 2 selects Mesh Clients 1 and 2 as winners, which are assigned to MAPs A and B, respectively. The *critical value* is the virtual bid per channel utilization of MC 3, $\tilde{p}_c = \tilde{p}_3 = \frac{\phi(b_3)}{o_{3B}} = 12$. Therefore, the virtual prices paid by the two winners are $\tilde{p}_1 = \tilde{p}_3 \cdot o_{1A} = 12 < \phi(b_1)$ and $\tilde{p}_2 = \tilde{p}_3 \cdot o_{2B} = 6 < \phi(b_2)$, respectively.

VI. ANALYSIS OF THE GREEDY AUCTION

In this section, we analyze the key structural properties of our greedy auction. In particular, we formally prove that revealing the real valuation of the available bandwidth is the best strategy for each MC participating to the auction. Then, we analyze the economic efficiency of the greedy algorithm evaluating the performance gap with respect to the optimal auction.

A. Truthfulness Analysis

Having described the main phases of the greedy bandwidth auction, hereafter we prove formally that our mechanism satisfies the *truthfulness property*. We recall that an auction mechanism is truthful if the dominant strategy for each rational bidder is to declare always its private real valuation for the requested demand. This property guarantees that selfish bidders cannot benefit from cheating, preventing the strategic manipulation of

the marketplace, thus resulting in an efficient allocation of the available resources.

In order to prove that the auction implemented by our greedy algorithm is truthful, we have to show first that the allocation rule satisfies the following properties:

- 1) *monotonicity* of the allocation defined by the auction;
- 2) existence of a *critical value* for each winner mesh client, which determines if its demand is satisfied or not.

The following lemmas (VI.1–VI.3) prove that the allocation phase of Algorithm 2 satisfies the above properties and provide the basis to demonstrate in Theorem VI.1 that no bidder can unilaterally increase its utility by submitting a bid that is different from its private valuation.

Lemma VI.1: If the demand of mesh client i is satisfied when it bids b_i , then i 's demand is still satisfied if i increases its bid, $b'_i > b_i$.

Proof: Let L and L' be two sorted lists for virtual bids $\phi(b_i)$ and $\phi(b'_i)$ (corresponding to b_i and b'_i), respectively. Let us define $rank(i, L)$ a monotonic decreasing function of the MC i position in the list L . Since $\frac{\phi(b'_i)}{o_{ij}} > \frac{\phi(b_i)}{o_{ij}}$ (recall that $\phi(x)$ is a nondecreasing function), the sorting algorithm in the greedy allocation phase (Algorithm 3) moves i in a best position, i.e., $rank(i, L') > rank(i, L)$. Therefore, the rank of i can only increase if it submits a higher bid, and thus a higher virtual bid, resulting in a different order of the set of mesh clients whose demands are satisfied, which implies that if mesh client i is allocated by bidding b_i (i.e., $\phi(b_i)$), its demand is satisfied even with a higher bid $b'_i > b_i$ (i.e., $\phi(b'_i) > \phi(b_i)$). ■

Lemma VI.2: For each mesh client i , the greedy Algorithm 2 provides the critical value $\tilde{p}_c = \frac{\phi(b_c)}{o_{cj_c}}$ such that i 's demand is satisfied if i 's virtual bid per channel utilization is higher than \tilde{p}_c , whereas it is rejected if i 's virtual bid per channel utilization is lower than \tilde{p}_c .

Proof: The proof is straightforward since Algorithm 2 scans the list L in nonincreasing order of virtual bids per channel utilization $\left(\frac{\phi(b_i)}{o_{ij}} \right)$ until it cannot allocate more bandwidth demands or all mesh clients are satisfied. In the former case, the *critical value* is equal to the ratio $\frac{\phi(b_c)}{o_{cj_c}}$ of the first unsatisfied mesh client c , while in the latter case the *critical value* is null.

Lemma VI.3: The price paid by each winner i is lower than or equal to its submitted bid b_i .

Proof: To show that $p_i \leq b_i$, we need to demonstrate that the *critical value* times the channel utilization charged to winner i (i.e., the virtual price of i , \tilde{p}_i) is not greater than its virtual bid $\phi(b_i)$.

Each winner i is charged the *critical value* \tilde{p}_c times the channel utilization (o_{ij}) of the MAP j to which it has been assigned to satisfy its bandwidth demand, d_i . If all mesh clients can be satisfied, the *critical value* is zero, thus mesh clients are not charged for their bandwidth utilization. Otherwise, MC i pays virtually $\tilde{p}_i = \phi(b_c) \cdot \frac{o_{ij}}{o_{cj_c}} \leq \phi(b_i)$ since the mesh clients list L is sorted in nonincreasing order of the ratio $\frac{\phi(b_i)}{o_{ij}} \geq \frac{\phi(b_c)}{o_{cj_c}}$, where $\frac{\phi(b_c)}{o_{cj_c}}$ represents the ratio of the virtual bid and the utilization of the critical bidder (i.e., the first loser of the auction).

Finally, since $\phi(x)$ is a nondecreasing monotonic function, $p_i = \phi^{-1}(\tilde{p}_i) \leq b_i$.

TABLE III
AUCTION OUTCOMES AND UTILITY FOR MC i WHEN IT BIDS UNTRUTHFULLY

Outcome for MC i		Utility of MC i	
v_i	$b_i > v_i$	v_i	$b_i > v_i$
i wins	i wins	$u_i = v_i - p_i$	$u_i = v_i - p_i$
i loses	i wins	$u_i = 0$	$u_i = v_i - p_i < 0$
i wins	i loses	-	-
i loses	i loses	$u_i = 0$	$u_i = 0$

Outcome for MC i		Utility of MC i	
v_i	$b_i < v_i$	v_i	$b_i < v_i$
i wins	i wins	$u_i = v_i - p_i$	$u_i = v_i - p_i$
i loses	i wins	-	-
i wins	i loses	$u_i = v_i - p_i > 0$	$u_i = 0$
i loses	i loses	$u_i = 0$	$u_i = 0$

Theorem VI.1 (Truthfulness of Algorithm 2): Algorithm 2 implements a truthful auction.

Proof: We prove the theorem by showing that no mesh client can increase its utility by submitting a bid b_i different from its private valuation v_i . We underline that the utility perceived by MC i does not change by bidding either v_i or b_i since it is defined as $u_i(x) = v_i - p_i$.

We must consider two cases, namely: (A) $\phi(b_i) > \phi(v_i)$ (or equivalently, due to the nondecreasing monotonic property of the virtual function, $b_i > v_i$); and (B) $\phi(b_i) < \phi(v_i)$ (or equivalently, $b_i < v_i$). For each case, we must consider all possible four outcomes, detailed in the following and summarized in Table III.

Let us start with Case (A) by considering the following cases.

Case A.1: User i Wins Either by Bidding $\phi(b_i)$ or $\phi(v_i)$: If i wins by bidding either $\phi(b_i)$ or $\phi(v_i)$, then i is ranked in a better place in the list L when it submits $\phi(b_i)$. However, this changes only the order of the set of winners, but neither the *critical value* \tilde{p}_c nor the virtual price \tilde{p}_i , which are still given by the following expressions: $\tilde{p}_c = \frac{\phi(b_c)}{o_{c j_c}}$, and $\tilde{p}_i = \tilde{p}_c o_{i j}$. Hence the price paid by the winner does not vary, $p_i = \phi^{-1}(\tilde{p}_i)$. Therefore, the utility does not change: $u_i(b_i) = u_i(v_i)$.

Case A.2: User i Wins by Bidding $\phi(b_i)$ but Loses With $\phi(v_i)$: If i wins by submitting $\phi(b_i)$ but it loses with $\phi(v_i)$, then there exists a *critical value* $\tilde{p}_c = \frac{\phi(b_c)}{o_{c j_c}}$ such that $\frac{\phi(v_i)}{o_{i j}} < \frac{\phi(b_c)}{o_{c j_c}} < \frac{\phi(b_i)}{o_{i j}}$, thus $\phi(v_i) < \tilde{p}_i < \phi(b_i)$.

Due to the monotonic property, the private valuation of i is lower than the price that it pays when it submits $\phi(b_i)$, i.e., $v_i < \phi^{-1}(\tilde{p}_i) < b_i$. Therefore, the utility perceived by i is negative, $u_i(b_i) = v_i - \phi^{-1}(\tilde{p}_i) < 0$, hence it is better off losing the auction since in this latter case its utility is null, $u_i(v_i) = 0$.

Case A.3: User i Loses by Bidding $\phi(b_i)$ but Wins With $\phi(v_i)$: Due to the monotonic property, this case is impossible since by submitting a higher bid, and thus a higher virtual bid, user i will be placed in a better position of the sorted list L .

Case A.4: User i Loses Either by Bidding $\phi(b_i)$ or $\phi(v_i)$: If i loses by bidding both b_i and v_i , then its utility is always null: $u_i(b_i) = u_i(v_i) = 0$.

Similarly, for case (B) $\phi(b_i) < \phi(v_i)$, we can demonstrate that mesh client i cannot increase its utility by submitting a lower bid than its private valuation. ■

Table III summarizes all possible outcomes to show how MC i cannot increase its utility by bidding differently than its private valuation (i.e., $b_i \neq v_i$).

Since Algorithm 2 implements a truthful auction (which means that selfish bidders cannot benefit from manipulating their bids), a WMN operator can efficiently compute a solution for the auction problem, being assured that all MCs reveal the true valuation for their bandwidth demand. We emphasize that, even though our greedy algorithm provides a slightly suboptimal solution with respect to the optimal allocation for the bandwidth auction problem, according to [37], the most important properties necessary to prevent market manipulation are preserved by our proposal.

Finally, we also observe that the reservation price $\phi^{-1}(0) = b_r$ set by the operator limits the impact of *collusion* on the revenue earned by the operator, thus discouraging the collusion among two or more selfish bidders. Note, however, that a complete analysis of the collusion on the performance of our mechanism is out of the scope of this paper and is left for future study.

B. Economic Efficiency Analysis

In this section, we analyze the economic efficiency of our greedy algorithm with respect to the optimal allocation computed by Algorithm 1 in the worst-case scenario by evaluating the *Price of Anarchy (PoA)*, which is defined as the ratio between the optimal and worst possible values of the Social Welfare, according to

$$PoA = \frac{SW^o}{SW^g} = \frac{\sum_{i \in \mathcal{N}} \phi(b_i) \cdot x_i^o}{\sum_{i \in \mathcal{N}} \phi(b_i) \cdot x_i^g} \quad (13)$$

where x_i^o and x_i^g represent the allocation computed by the optimal and greedy algorithms, respectively.

To provide better insight on the economic efficiency of our solution, we also analyze the *Revenue Ratio (RR)*, defined as the ratio between the revenues computed using the optimal and greedy algorithms, according to the following equation, since the operator would seek to minimize such performance metric:

$$RR = \frac{\sum_{i \in \mathcal{N}} p_i \cdot x_i^o}{\sum_{i \in \mathcal{N}} p_i \cdot x_i^g} \quad (14)$$

To illustrate the economic efficiency of the greedy approach, we consider a simple network scenario composed of one mesh access point, A , and two mesh clients ($\{1, 2\}$). The capacity C of the MAP's backhaul connection is large enough to accommodate the traffic transmitted over the wireless access interface. The channel utilizations of the two MCs are $o_{1A} = d_{1A}/r_{1A} = 1/D$ (D is a positive parameter) and $o_{2A} = d_{2A}/r_{2A} = 1$, while their bids are $b_1 = (1 + \epsilon)$ (ϵ is a small value larger than 0) and $b_2 = D$, respectively.

In this scenario, the allocation that maximizes the Social Welfare is the one that selects MC 2 ($SW^o = b_2 = D$). However, the greedy algorithm selects MC 1 since $b_1/o_1 > b_2/o_2$, and the Social Welfare is in this case equal to $SW^g = b_2 = 1 + \epsilon$. The ratio $\frac{SW^o}{SW^g}$ is therefore equal to $\frac{D}{1 + \epsilon}$, and the Price of Anarchy of the social welfare tends to infinity with D , $PoA = \lim_{D \rightarrow \infty} \frac{D}{1 + \epsilon} = \infty$.

Nevertheless, we underline that the revenue obtained using the greedy algorithm approaches that obtained using the optimal mechanism since the optimal revenue is only ϵ times larger than the revenue computed using the greedy algorithm. Indeed, the

RR defined in (14) is independent of D , and even in this limiting case it results

$$RR = \lim_{D \rightarrow \infty} \frac{R^o}{R^g} = \lim_{D \rightarrow \infty} \frac{1 + \epsilon}{1} = 1 + \epsilon. \quad (15)$$

We further underline that our allocation mechanism can be modified to bound the Price of Anarchy of the Social Welfare by simply fixing a minimum amount of bandwidth demand for any bidder that is willing to participate to the auction. In practice, as we illustrate in Section VII-B, the economic gap between the optimal and greedy solutions is almost null (the experimental Price of Anarchy we measured is, in fact, always lower than 1.05), thus proving the economic efficiency of our scheme.

VII. NUMERICAL RESULTS

In this section, we illustrate the numerical results obtained solving the bandwidth allocation auction using both the optimal and greedy algorithms detailed in previous sections. We first describe the results obtained in *randomly* generated network topologies, then we show the performance of our approaches using a *real-life* network, namely the Google Wifi network⁴ analyzed in [4].

A. Random Networks

Experimental Methodology: In this set of simulations, we consider typical WMCN topologies composed of 30, 60, 90, and 120 WMN devices (i.e., mesh routers, mesh access points, and mesh gateways) randomly scattered over an area of $1000 \times 1000 \text{ m}^2$, similarly to [38]. The ratios between the three different devices is fixed to 1:2 and 1:3 for MGWs:MRs and MGWs:MAPs, respectively.

In all the topologies, we vary the number of MCs, which participate to the bandwidth auction, from 400 to 1000. The bandwidth demands and bids are uniformly distributed in the range [1, 9] Mbps and [10, 30] monetary units (e.g., US dollars), respectively.

The channel capacity of both access and backbone links is defined according to the reception sensitivity of the Wistron CM9 commercial wireless cards (based on Atheros chipset)⁵. The path loss necessary to evaluate the sensitivity of the receiving node is computed according to the Friis propagation model. However, we underline that all the above assumptions do not affect the proposed algorithms, which can be used to solve any network scenario.

In order to gauge the performance of the proposed greedy algorithm (Section V) with respect to the optimal solution (Section IV), we consider the following metrics:

- *Revenue:* defined as the sum of the prices paid by all winners;
- *Social Welfare:* defined as the sum of the winner bids, $\sum_{i \in \mathcal{N}} b_i \cdot x_i$;
- *Winners:* This metric represents the number of winners selected among the mesh clients that participated to the auction. It provides an indication of the satisfaction of the customers;

- *Fairness:* We consider the *Jain's Fairness Index* [39], defined according to

$$\text{Jain's Fairness Index} = \frac{\left(\sum_{i=1}^w \rho_i \right)^2}{w \cdot \sum_{i=1}^w \rho_i^2} \quad (16)$$

where ρ_i represents the ratio between the paid price and the requested bandwidth demand of the winner i , $\rho_i = p_i/d_i$, whereas w represents the number of winners.

The *Jain's Fairness Index* therefore measures the spread of the price paid by winners per bandwidth unit, and it varies from $1/w$ (no fairness among winners) to 1 (perfect fairness).

For each network scenario, we performed 10 independent measurements, computing very narrow 95% confidence intervals. For the sake of clarity, the *Revenue* and the *Social Welfare* have been normalized with respect to the maximum value measured in the network topology composed of 120 WMN devices (about 12 500 monetary units).

Performance Evaluation: Fig. 4 shows the performance metrics measured in the network topologies composed of $m = 30, 60, 90,$ and 120 WMN devices as a function of the number of MCs, using the allocation mechanisms discussed in previous sections.

The curves identified with “R.o” and “R.g” represent the solutions obtained using Algorithms 1 (which could be computed only for $m = 30$) and 2, respectively. The curves “SW.o” and “SW.g” correspond to the Social Welfare. Finally, the remaining curve identified with “VCG” represents the revenue of solutions obtained according to the classical Vickrey–Clarke–Groves mechanism [30], which does not consider the bids distribution to compute the price paid by the winners. The performance gap between our mechanism and the VCG scheme illustrates the performance increase that can be obtained considering the additional information on the bids distribution.

Note that, due to the high computational complexity, we were able to solve the auction problem optimally only for the network scenario composed of 30 WMN devices. Nevertheless, even in this simple scenario, the maximum computational time we measured to solve the problem on a Pentium 4 with 3.0 GHz and 2 GB of RAM was approximately equal to 40 h. Conversely, the greedy approach takes always less than 30 s to find efficient allocations and the corresponding payments.

Fig. 4(a) illustrates the revenue earned by the operator auctioning its available bandwidth, when 30 WMN nodes (15 MAPs, 10 MRs, and 5 MGWs) are scattered randomly over the $1000 \times 1000\text{-m}^2$ square area.

As illustrated in the figure, the additional information provided by the virtual bids permits to increase the operator's expected revenue with respect to a mechanism that exploits only the MC bids. We can further notice that the auction implemented by the greedy algorithm well approaches the optimal revenue, and therefore it represents an effective and efficient solution for the computation of the prices paid by the MCs. In addition, the Social Welfare is always higher than the revenue earned when using Algorithms 1 and 2. Indeed, this value represents an ideal upper bound to the revenue since it can be achieved only assuming that all mesh clients behave honestly, submitting the price they are willing to pay for their bandwidth demands, even

⁴Available online at <http://wifi.google.com/>.

⁵Available online at <http://www.lri.fr/~fmartignon/CM9.pdf>

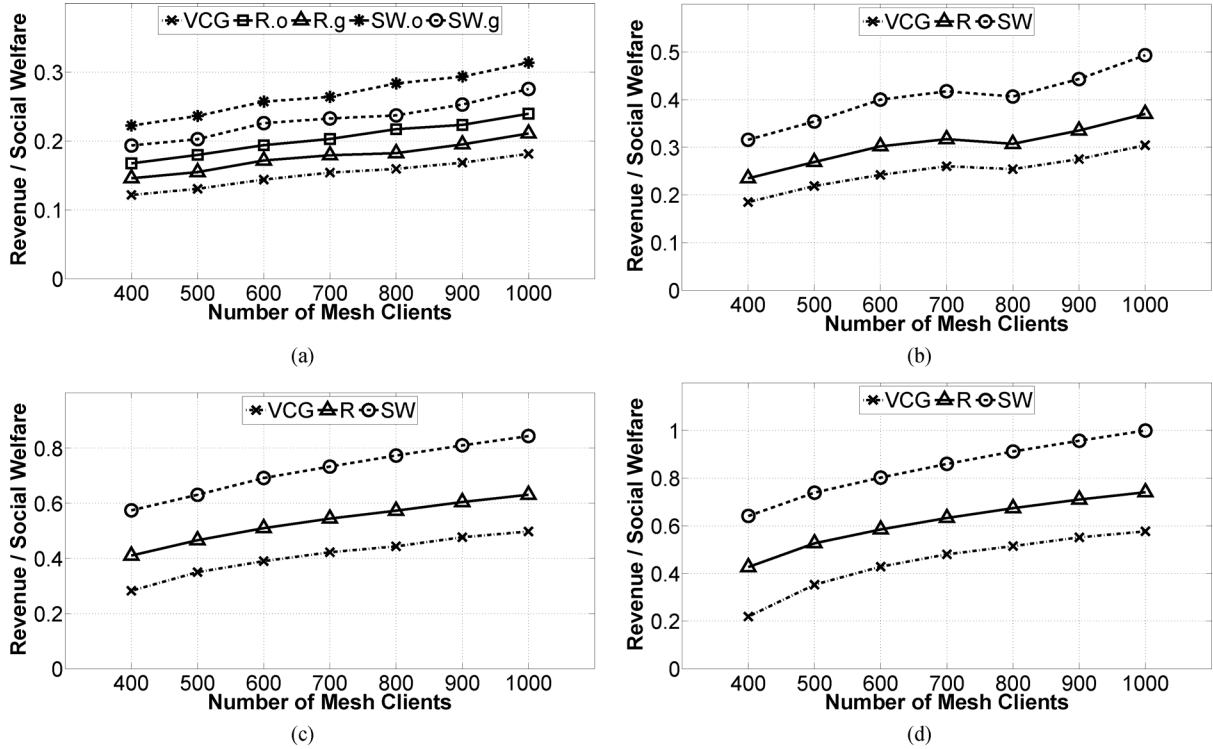


Fig. 4. Revenue and Social Welfare measured as a function of the number of mesh clients in scenarios with $m = 30, 60, 90,$ and 120 WMN devices. (a) $m = 30$. (b) $m = 60$. (c) $m = 90$. (d) $m = 120$.

if it is not their best strategy. On the contrary, we underline that our proposed solutions *ensure* that *all* mesh clients bid truthfully their valuations because it is their best strategy. The achieved revenue is approximately equal to 75% of the Social Welfare.

Greedy solutions illustrated in Fig. 4(b)–(d) confirm the trends observed for the network scenario composed of 30 WMN devices. Note how increasing the number of mesh clients guarantees higher revenues. This is due to the effect of the competition: Only mesh clients bidding more will be accepted.

Figs. 5 and 6 show respectively the number of winners and the Jain’s Fairness Index of $\rho_i = p_i/d_i$ obtained through the bandwidth auction implemented by Algorithms 1 and 2 as a function of the number of mesh clients that participate to the auction for the bandwidth allocation. It can be observed from Fig. 5 that the greedy algorithm selects a number of mesh clients very close to the value obtained using the optimal allocation algorithm (see the curves identified by labels “30.o” and “30.g”). In particular, the greedy algorithm leads to a performance gap always lower than 10%, for instance sizes when both algorithms can be run. The figure illustrates also the number of winners selected in the network scenarios composed of 60, 90, and 120 WMN devices. As expected, the higher the number of WMN devices, the higher the available network bandwidth, and the greater the number of mesh clients satisfied by the allocation algorithm implementing the auction.

Furthermore, Fig. 6 shows that the greedy scheme computes a solution with a slightly higher fairness than the optimal approach. This is due to the different allocation phases implemented by the two algorithms. The greedy scheme selects as winners the MCs with the highest ratios of $\rho_i = p_i/d_i$, even though they do not represent the best solution in terms of Social Welfare and Revenue (see Fig. 4). Therefore, the greedy




Fig. 5. Number of winners as a function of the number of mesh clients in scenarios with $m = 30, 60, 90,$ and 120 WMN devices.




Fig. 6. Comparative evaluation of the fairness (Jain’s Fairness Index of $\rho_i = p_i/d_i$) obtained using the optimal and the greedy algorithms in the network scenario with $m = 30$ WMN devices.

algorithm rules out the MCs with the worst ratios of $\rho_i = p_i/d_i$, which would increase the variability of the paid prices, but at the

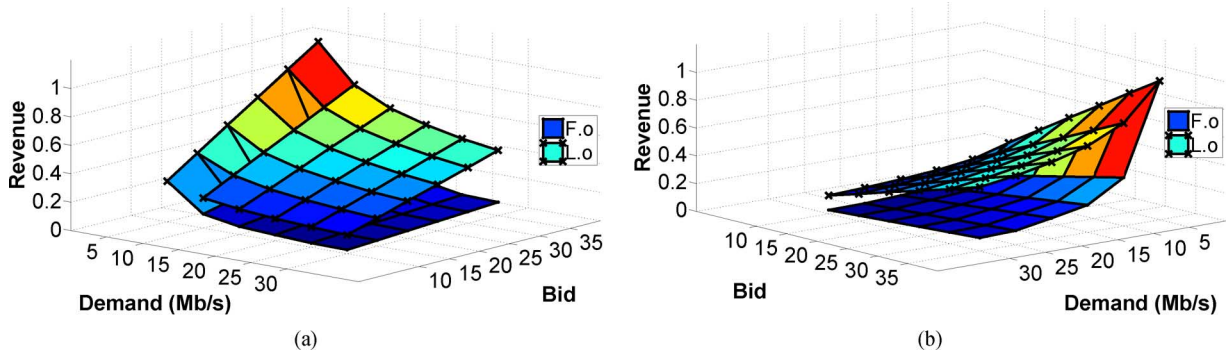


Fig. 7. Optimal revenue, (a) front and (b) rear, measured in the real-life network scenario (Google WiFi).

same time would permit to satisfy a higher number of MCs, thus increasing the overall Revenue earned by the WMN operator, as it can be observed from Figs. 4 and 5. We underline, however, that both algorithms compute fair solutions since the Jain’s Fairness Index of the ratio $\rho_i = p_i/d_i$ is very close to 1. Therefore, all MCs selected as winners pay almost the same price per bandwidth unit.

B. Real-Life WiFi Network

Experimental Methodology: The Google WiFi network is an operational, large-scale, wireless multihop network architecture deployed in Mountain View, CA, USA, by Google. In [4], the authors collected coverage measurements of 168 access points located in a 12-km² area. In particular, the authors measured the signal strength and the signal-to-noise ratio (SNR) from more than 75 000 client locations. We use the 168 access point locations as MAPs, and the 75 000 client positions as candidate MCs, respectively.

The channel capacity of the wireless link that can be established between the access point and each candidate mesh client is defined according to the measured SNR using the reception sensitivity thresholds of Atheros-based wireless cards.

In order to evaluate the effect of the demand and bid distributions on the operator’s revenue, we solve the auction using both the optimal and greedy algorithms varying the values of the requested bandwidth demand and the submitted bid, which are both drawn from uniform distributions. More specifically, mesh clients bids are distributed uniformly in the range $[0, 2x]$ monetary units, with $x \in \{10, 20, 25, 30, 35\}$, while we consider two different interval sizes for the demand distribution to simulate preference variability: *large* and *fixed* intervals which correspond to ranges $(0, 2y)$ Mb/s and $[y - 4, y + 4]$ Mb/s, respectively, with $y \in \{5, 10, 15, 20, 25, 30\}$.

As in the randomly generated network scenarios, we measure the *Revenue*, the number of *Winners*, and the Jain’s Fairness Index of the ratio $\rho_i = p_i/d_i$ of the allocations computed by the optimal and greedy algorithms. Furthermore, we compute the average utilization of the access network formed by all access points. For the sake of clarity, we do not show the *Social Welfare* in the following figures since it almost overlaps the *Revenue* computed with the optimal algorithm.

We underline that, in this scenario, we assume that the backhaul network has been designed to satisfy the maximum bandwidth of the access network in order to compare the performance of the optimal and greedy algorithms. Even with

this assumption, the computational time that we measured to solve the problem optimally on a Pentium 4 with 3.0 GHz and 2 GB of RAM was approximately equal to 5 h.

Performance Evaluation: Fig. 7 shows the *Revenue* obtained by the optimal algorithm in the Google WiFi network as a function of the average bid and demand submitted by the mobile clients. The curves identified with “L.o” and “L.g” represent the solutions obtained using Algorithms 1 and 2 in network scenarios with bandwidth demands drawn from a uniform distribution with *large* interval size. On the contrary, the curves “F.o” and “F.g” show the same performance metrics measured with uniformly distributed bandwidth demands with *fixed* interval size.

For the sake of clarity, in Fig. 7, we did not report the results measured by the greedy mechanism since they are very close to the values obtained using the optimal approach. Note, however, that Algorithm 2 selects a slightly lower number of winners than Algorithm 1, which in turn leads to a lower average utilization of the access points, as illustrated in Fig. 8(a) and (b). Indeed, the interval size of the bandwidth demand has a twofold effect on the performance of our algorithms.

- On the one hand, the higher the demand variability is, the higher the operator’s revenue, at the detriment of the fairness [see Fig. 8(c)]. Due to the larger interval size, there is a higher probability of having unbalanced bandwidth demands, which lead to widely distributed ratios of valuation per channel utilization. The operator can therefore select those clients with the highest such ratio, which increases its revenue, but results in a lower overall fairness.
- On the other hand, each decrease of the variability of the bandwidth demand produces a more fair allocation of the resources, at the expense of the total revenue earned by the operator and the utilization of its network, as illustrated in Fig. 8, since in this case the operator can only satisfy those clients who are willing to pay the highest price per bandwidth unit.

As expected, the lower the average demand requested by MCs is, the higher the revenue of the operator since this latter can satisfy a higher number of bidders with the fixed available capacity provided by its access points. Furthermore, we can clearly observe from Fig. 7 that the variation of the total *Revenue* is not equally affected by the average demand and the average bid. More specifically, the revenue gained by the operator benefits more from a combined decrease of the average demand (the y value) and its variability (the interval size) than an increase of




Fig. 8. (a) Number of winners, (b) network utilization and (c) Jain's Fairness Index measured in the real-life network scenario (Google WiFi) as a function of the requested demand. Each point is the average with respect to the bids.




Fig. 9. Cumulative distribution function of the experimental PoA and RR measured in the real-life network scenario (Google WiFi) with a large interval of the bandwidth demand distribution.

the average bid offered by MCs. Therefore, the operator should leverage on policies that favor large marketplaces formed by several MCs with low values of average demand to reach economic efficiency and increase its revenue.

Finally, in order to show the gap of the economic efficiency of our greedy mechanism with respect to the optimal solution, we experimentally computed the PoA and the RR in the Google WiFi network scenario, with bandwidth demands drawn from a uniform distribution with large interval size (i.e., $(0, 2y)$ Mb/s with $y \in \{5, 10, 15, 20, 25, 30\}$). Indeed, this scenario contains bids per bandwidth demands that may negatively affect the PoA , as discussed in Section VI-B.

As illustrated in Fig. 9, the experimental Price of Anarchy we measured is always lower than 1.05, thus proving the economic efficiency of our greedy scheme.

VIII. CONCLUSION

In this paper, we proposed two effective mechanisms to allocate the available bandwidth of a WMN operator to those customers who are willing to pay the higher price for satisfying their bandwidth demand. We first formulated the allocation mechanism as a combinatorial auction, which guarantees that all customers reveal their real valuation of the required bandwidth. Then, we proposed a greedy algorithm that finds efficient allocations in polynomial time even for large-scale, real network scenarios while maintaining the truthfulness property.

We evaluated our solutions in several large-scale network topologies generated both randomly and based on real-life deployments, like the Google WiFi scenario. Numerical results show that the greedy algorithm performs very close to the optimal combinatorial auction, thus representing an efficient, fair, and practical alternative for solving the auction of the proposed bandwidth marketplace.

The analysis performed using real wireless traces suggests to design market policies that force MC owners to lower their bandwidth requirements rather than increasing their offers to maintain the same level of service. In addition to improving the operator profit, this permits to enhance the overall system satisfaction and fairness.

REFERENCES

- [1] I. F. Akyildiz, X. Wang, and W. Wang, "Wireless mesh networks: A survey," *Comput. Netw.*, vol. 47, no. 4, pp. 445–487, 2005.
- [2] N. Nandiraju, D. Nandiraju, L. Santhanam, B. He, J. Wang, and D. P. Agrawal, "Wireless mesh networks: Current challenges and future directions of web-in-the-sky," *IEEE Wireless Commun.*, vol. 14, no. 4, pp. 79–89, Aug. 2007.
- [3] P. A. Frangoudis, G. C. Polyzos, and V. P. Kemerlis, "Wireless community networks: An alternative approach for nomadic broadband network access," *IEEE Commun. Mag.*, vol. 49, no. 5, pp. 206–213, May 2011.
- [4] J. Robinson, R. Swaminathan, and E. Knightly, "Assessment of urban-scale wireless networks with a small number of measurements," in *Proc. ACM MobiCom*, Sep. 2008, pp. 187–198.
- [5] X. Zhou, S. Gandhi, S. Suri, and H. Zheng, "eBay in the sky: Strategy-proof wireless spectrum auctions," in *Proc. ACM MobiCom*, 2008, pp. 2–13.
- [6] X. Zhou and H. Zheng, "TRUST: A general framework for truthful double spectrum auctions," in *Proc. IEEE INFOCOM*, 2009, pp. 999–1007.
- [7] J. Jia, Q. Zhang, Q. Zhang, and M. Liu, "Revenue generation for truthful spectrum auction in dynamic spectrum access," in *Proc. ACM MobiCom*, 2009, pp. 3–12.
- [8] G. S. Kasbekar and S. Sarkar, "Spectrum auction framework for access allocation in cognitive radio networks," in *Proc. ACM MobiHoc*, 2009, pp. 1841–1854.
- [9] S. Sengupta and M. Chatterjee, "An economic framework for dynamic spectrum access and service pricing," *IEEE/ACM Trans. Netw.*, vol. 17, no. 4, pp. 1200–1213, Aug. 2009.
- [10] L. Chen, S. Iellamo, M. Coupechoux, and P. Godlewski, "Spectrum auction with interference constraint for cognitive radio networks with multiple primary and secondary users," *Wireless Netw.*, vol. 17, no. 5, pp. 1355–1371, 2011.
- [11] D. Yang, X. Fang, and G. Xue, "Truthful auction for cooperative communications," in *Proc. ACM MobiHoc*, 2011, pp. 89–98.
- [12] L. Anderegg and S. Eidenbenz, "Ad hoc-VCG: A truthful and cost-efficient routing protocol for mobile Ad hoc networks with selfish agents," in *Proc. ACM MobiCom*, 2003, pp. 245–259.
- [13] S. Eidenbenz, G. Resta, and P. Santi, "The COMMIT protocol for truthful and cost-efficient routing in Ad hoc networks with selfish nodes," *IEEE Trans. Mobile Comput.*, vol. 7, no. 1, pp. 19–33, Jan. 2008.
- [14] J. Jaramillo and R. Srikant, "DARWIN: Distributed and adaptive reputation mechanism for wireless Ad hoc networks," in *Proc. ACM MobiCom*, 2007, pp. 87–98.
- [15] W. Wang and X. Y. Li, "Low-cost routing in selfish and rational wireless Ad hoc networks," *IEEE Trans. Mobile Comput.*, vol. 5, no. 5, pp. 596–607, May 2006.
- [16] Y. Wu, S. Tang, P. Xu, and X. Y. Li, "Dealing with selfishness and moral hazard in noncooperative wireless networks," *IEEE Trans. Mobile Comput.*, vol. 9, no. 3, pp. 420–434, Mar. 2009.

- [17] W. Wang, X. Li, and Y. Wang, "Low-cost truthful multicast in selfish and rational wireless Ad Hoc networks," in *Proc. ACM MobiCom*, 2004, pp. 534–536.
- [18] W. Z. Wang, X. Y. Li, Z. Sun, and Y. Wang, "Design multicast protocols for non-cooperative networks," in *Proc. IEEE INFOCOM*, 2005, vol. 3, pp. 1596–1607.
- [19] S. Zhong and F. Wu, "A collusion-resistant routing scheme for non-cooperative wireless Ad hoc networks," *IEEE/ACM Trans. Netw.*, vol. 18, no. 2, pp. 582–595, Apr. 2010.
- [20] K. Chen and K. Nahrstedt, "iPass: An incentive compatible auction scheme to enable packet forwarding service in MANET," in *Proc. IEEE ICDCS*, 2004, pp. 534–542.
- [21] S. Zhong and F. Wu, "On designing collusion-resistant routing schemes for non-cooperative wireless Ad hoc networks," in *Proc. ACM MobiCom*, 2007, pp. 278–289.
- [22] R. Jain and J. Walrand, "An efficient mechanism for network bandwidth auction," in *Proc. IEEE NOMS*, 2008, pp. 227–234.
- [23] F. Fu and U. C. Kozat, "Wireless network virtualization as a sequential auction game," in *Proc. IEEE INFOCOM*, 2010, pp. 1945–1953.
- [24] R. Guerin, H. Ahmadi, and M. Naghshineh, "Equivalent capacity and its application to bandwidth allocation in high-speed networks," *IEEE J. Sel. Areas Commun.*, vol. 9, no. 7, pp. 968–981, Sep. 1991.
- [25] D. Malone, K. Duffy, and D. Leith, "Modeling the 802.11 distributed coordination function in nonsaturated heterogeneous conditions," *IEEE/ACM Trans. Netw.*, vol. 15, no. 1, pp. 159–172, Feb. 2007.
- [26] G. Bianchi, P. Gallo, D. Garlisi, F. Giuliano, F. Gringoli, and I. Tinirello, "MAClets: Active MAC protocols over hard-coded devices," in *Proc. ACM CoNEXT*, 2012, pp. 229–240.
- [27] M. Alicherry, R. Bhatia, and L. E. Li, "Joint channel assignment and routing for throughput optimization in multi-radio wireless mesh networks," in *Proc. ACM MobiCom*, 2005, pp. 58–72.
- [28] S. Avallone, I. F. Akyildiz, and G. Ventre, "A channel and rate assignment algorithm and a layer-2.5 forwarding paradigm for multi-radio wireless mesh networks," *IEEE/ACM Trans. Netw.*, vol. 17, no. 1, pp. 267–280, Feb. 2009.
- [29] E. Amaldi, A. Capone, M. Cesana, I. Filippini, and F. Malucelli, "Optimization models and methods for planning wireless mesh networks," *Comput. Netw.*, vol. 52, no. 11, pp. 2159–2171, 2008.
- [30] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*. Cambridge, U.K.: Cambridge Univ. Press, 2007.
- [31] R. B. Myerson, "Optimal auction design," *Math. Oper. Res.*, vol. 6, no. 1, pp. 58–73, 1981.
- [32] S. Chawla, J. D. Hartline, and R. Kleinberg, "Algorithmic pricing via virtual valuations," in *Proc. ACM Conf. Electron. Commerce*, 2007, pp. 243–251.
- [33] Y. Cai, C. Daskalakis, and S. M. Weinberg, "On optimal multidimensional mechanism design," *SIGecom Exchanges*, vol. 10, no. 2, pp. 29–33, 2011.
- [34] C. Chekuri and S. Khanna, "A PTAS for the multiple Knapsack problem," in *Proc. ACM-SIAM Symp. Discrete Algor.*, 2000, pp. 213–222.
- [35] D. Lehmann, L. I. O'Callaghan, and Y. Shoham, "Truth revelation in approximately efficient combinatorial auctions," *J. ACM*, vol. 49, no. 5, pp. 577–602, 2002.
- [36] A. V. Goldberg and R. E. Tarjan, "A new approach to the maximum-flow problem," *J. ACM*, vol. 35, no. 4, pp. 921–940, 1988.
- [37] P. Klemperer, "What really matters in auction design," *J. Econ. Perspect.*, vol. 16, no. 1, pp. 169–189, 2002.
- [38] F. Martignon, S. Paris, I. Filippini, and A. Capone, "Efficient bandwidth allocation in wireless community networks," in *Proc. IFIP Wireless Days*, 2011, pp. 1–6.
- [39] R. Jain, *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling*. New York, NY, USA: Wiley-Interscience, 1991.