

An Efficient Auction-based Mechanism for Mobile Data Offloading

Stefano Paris*, Fabio Martignon†, Ilario Filippini‡, and Lin Chen†

*LIPADE

Paris Descartes University
stefano.paris@parisdescartes.fr

†LRI

Paris-Sud University
{fabio.martignon, lin.chen}@lri.fr

‡DEIB

Politecnico di Milano
ilario.filippini@polimi.it



Abstract—The opportunistic utilization of third party WiFi access devices to offload customer traffic from the mobile network has recently gained momentum as a promising approach to increase the network capacity and simultaneously reduce the energy consumption of the Radio Access Network (RAN) infrastructure. To foster the opportunistic utilization of unexploited Internet connections, we propose a new and open market where a mobile operator can lease the bandwidth made available by third parties (residential users or private companies) through their access points to increase dynamically (and adaptively) the network capacity. We formulate the offloading problem as a reverse auction considering the most general case of *partial covering* of the traffic to be offloaded. We discuss the conditions (i) to offload the maximum amount of data traffic according to the capacity made available by third party access devices, (ii) to foster the participation of access point owners (individual rationality), and (iii) to prevent market manipulation (incentive compatibility). Finally, we propose three alternative greedy algorithms that efficiently solve the offloading problem, even for large-size network scenarios.

Index Terms—Data Offloading, Heterogeneous Networks, Mechanism Design, Auction.

1 INTRODUCTION

In recent years, the rapid growth of traffic demand required by content-rich Internet services accessed by mobile users through their smartphones has increased the pressure on mobile operators for upgrading their cellular networks. Consequently, mobile operators have increased the capacity of their radio access and backhaul networks through the development of new technologies and a pervasive deployment of new types of base stations. Nevertheless, mobile customers are experiencing a “bandwidth crunch” due to the steady growth of the demand required by real-time services and the limited capacity of the wireless access technology.

A promising solution to smoothly handle sudden peaks of bandwidth demand is represented by the opportunistic utilization of low-cost and low-power small access devices (either Small Base Stations, SBSs, or Access Points, APs) massively deployed over the macro-cell

areas by the operator or third party entities. Third party access devices can use the legacy transmission technology of the large cell, such as LTE or beyond, but also rely on existing technology such as WiFi, thus forming Heterogeneous Mobile Networks. Consequently, mobile operators could provide a better wireless access service without limiting the maximum traffic of their customers through a wise management of their resources and the opportunistic utilization of other access network technologies.

In this paper, we investigate innovative policies and mechanisms to foster the deployment of Heterogeneous Mobile Networks as a means for mobile operators to increase their network capacity without deploying additional base stations, thus reducing their capital expenditure (CAPEX). As any marketplace, the misbehavior of even few SBS/AP owners (either residential users or hotspot administrators) playing strategically might seriously affect the efficiency of the allocation mechanism used by the mobile operator, thus discouraging honest agents from participating to the market. This, in turn, reduces the maximum amount of traffic that can be offloaded and the potential CAPEX savings. To address this issue, we present a *reverse truthful auction* that forces each SBS/AP owner interested in leasing the unexploited bandwidth of its Internet connection to bid truthfully. More generally, we consider the case of partial/constrained data offloading that stems from the limited resources provided by third parties, showing that such a problem asks for a deep revision of the classical payment rules.

Our work makes the following contributions:

- We propose and analyze a combinatorial reverse auction to implement an innovative marketplace both for selecting the cheapest third party access devices and offloading the maximum amount of data traffic from the RAN.
- We show that a payment rule that considers only the variation of the objective function solving the ILP problem with and without the winner does not always ensure the individual rationality of the

F. Martignon is with the Laboratoire de Recherche en Informatique (LRI), Université Paris-Sud, France, and Institut Universitaire de France (IUF).

participants for the analyzed mobile data offloading problem.

- We present an innovative payment rule based on the Vickrey-Clarke-Groves (VCG) scheme and demonstrate that it guarantees both *individual rationality* and *incentive compatibility* (i.e., *truthfulness*). To the best of our knowledge, this is the first payment rule that considers explicitly the trade-off between the total cost and the gain of offloading data connections.
- Since the optimal reverse auction problem is NP-hard, we further propose three greedy algorithms that solve very efficiently (i.e., in polynomial time) the allocation problem, even for large network instances, while preserving the *truthfulness* property.
- We perform a thorough numerical analysis and comparative evaluation of the proposed optimal and greedy allocation algorithms, considering realistic network scenarios.

The paper is structured as follows: Section 2 discusses related work. Section 3 presents the system model considered in our work. Section 4 formulates the combinatorial reverse auction as an optimization problem, and presents our new payment rule that makes the auction individually rational and truthful. Section 5 describes the greedy algorithm to solve efficiently the problem, while Section 6 illustrates and analyzes numerical results. Finally, concluding remarks are discussed in Section 7.

2 RELATED WORK

In recent years, several research groups have investigated the benefit of opportunistically offloading 3G data traffic on WiFi access networks to improve the QoS experienced by mobile devices [1], [2], [3]. Wiffler [2] provides a middleware layer for delay tolerant applications to overcome the poor availability and performance of WiFi access technology in vehicular networks, showing the performance increase through different experimental scenarios. The Application Programming Interface proposed in [3] further develops this approach to improve the performance of the applications using opportunistic wireless networking.

These works shed lights on the benefits of opportunistically using multiple wireless connections to increase the throughput and reduce the latency experienced by data connections. However, they miss opportunities for optimizing communications, since they design user-centric approaches without exploiting the global vision of Heterogeneous Mobile Networks. More recently, works like [4], [5], [6], [7], [8], [9] leverage both on the global knowledge of the mobile network operator and the multiple access radio interfaces of 4G smart devices to design auction mechanisms that minimize the overall offloading cost. Nevertheless, these mechanisms fail to find a feasible solution in typical network scenarios, where only a *subset* of mobile customers connections can be offloaded to the surrounding

WiFi access networks without exceeding their overall capacity. Indeed, the payment rules designed in these works require the assignment of *all* mobile data connections (i.e., their complete covering) in order to guarantee individual rationality and truthfulness. In contrast, our work presents a new reverse auction tailored for the more general problem where the operator can offload only a *portion* of the overall traffic load generated by its customers.

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 [10], [11], [12], [13], [14], [15].

VERITAS [10] pinpoints the limits of conventional auctions and proposes a truthful and flexible mechanism requiring only polynomial complexity for solving the spectrum allocation problem. TRUST [11] further develops this approach to support multi-party spectrum trading through a truthful double spectrum auction based on the well-known McAfee mechanism. The work presented in [12] adopts a similar approach to model and solve a broad class of problems concerning the allocation of spectrum resources to primary and secondary users in cognitive radio networks, while [13] analyzes also the interplay among the spectrum broker, service providers which are interested in leasing spectrum bands, and end-users. In [14] the authors investigate a spectrum marketplace where the spectrum owner's uncertainty about the private valuations of spectrum bidders is modeled using a Bayesian approach.

Auction theory has also 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 [16], [17], [18], [19].

Finally, recent research has analyzed virtual network scenarios where several service providers compete among each other for using the resources owned and managed by a network operator [20], [21]. In particular, Jain et al. in [20] present a mechanism for per-link bandwidth allocation of end-to-end paths in wired network, whereas Fu et al. in [21] design an auction-based stochastic game for resource allocation of virtual operators in wireless cellular networks.

Unlike recent literature, our work envisions a new marketplace based on reverse auctions, where WiFi Access Points are exploited by mobile network operators to offload the traffic of their customers. This marketplace would reduce the installation and management costs for mobile network operators, as well as foster the development of Heterogeneous Mobile Networks. Furthermore, we explicitly consider the more general partial covering problem of data connections, proposing a new payment rule to address the limits of the previous schemes.

3 SYSTEM MODEL

This section presents the economic definitions and assumptions, as well as the network model we adopt in the design of our auction mechanisms. Let us refer to the Heterogeneous Mobile Network (HMN) sample scenario illustrated in Figure 1, which is composed of a mobile cellular network formed by four base stations and a set of WiFi Access Points connected to the Internet. To simplify the discussion, in the rest of the paper we consider only WiFi APs as third party devices rented by the mobile operator; however, we underline that the proposed mechanisms can be easily extended to consider other SBS access technologies like LTE or beyond. The mobile network is managed by a single operator that provides ubiquitous access to its mobile customers (MCs), while each participant to the trading marketplace (either a residential user or a hotspot operator) is the owner of a wireless Access Point. AP owners lease the unused capacity of their Internet connections made available through wireless access points, so that the mobile operator can rent the available APs' bandwidth to offload the data traffic of its customers when, for example, this latter exceeds the maximum capacity provided by the mobile network, or to save energy by switching some BSs off. Table 1 summarizes the notation used in the paper.

3.1 Economic Model

Each AP owner i has an unexploited capacity C_i of its Internet connection that he is willing to lease for a given price v_i , unknown to the operator. To this end, he submits to the operator the bid $[b_i, C_i]$, representing the price that i asks for leasing the capacity C_i of its AP to the operator.

Through the mechanisms proposed in this work, the operator selects both the access points (APs) and the subset of its mobile customers (MCs) whose data traffic is offloaded from the mobile network to the selected APs. We remark that, unlike classical optimization approaches, where the operator knows exactly the price of each device that it can install, in our scenario such




Fig. 1: Network scenario considered in this work. The MN is managed by a single operator that provides access to its customers (e.g., MC_j), while the unused capacity of wireless access devices (e.g., AP_i) is leased to the operator for data traffic offloading.

TABLE 1: Basic notation used in the paper.

Sets	
\mathcal{M}	Set of Mobile Customers, $ \mathcal{M} = m$
\mathcal{M}_i	Subset of MCs that are covered by AP i
\mathcal{A}	Set of Access Points (i.e., bidders), $ \mathcal{A} = n$
Parameters	
C_i	Capacity of the Internet connection offered by AP owner i
b_i	Bid offered by AP owner i for its Internet connection
v_i	Real valuation of AP owner i for its Internet connection
p_i	Price paid by the operator to AP owner i
u_i	Utility of AP owner i
d_j	Bandwidth demand of MC j
r_{ji}	Maximum transmission rate of the wireless link established between nodes $j \in \mathcal{M}$ and $i \in \mathcal{A}$
o_{ji}	Channel utilization of AP i to satisfy the bandwidth demand of MC j
Variables	
x_i	Binary variable that indicates if AP i wins the auction
y_{ji}	Binary variable that indicates if MC j is assigned to AP i

information is hidden, thus we may have $b_i \neq v_i$. To prevent market distortion, the mechanisms proposed in this paper incentivize AP owners to provide the true information about the private valuation of their APs ($b_i = v_i$) by ensuring that there are no benefits to lying.

Let us denote by $p_i \geq 0$ the price paid by the operator to AP owner i to exploit its available capacity C_i . Then, assuming a quasi-linear utility function for AP owner i , we can define the utility of i , u_i , as the difference between the price paid by the operator, p_i , and the private valuation v_i , according to Equation (1):

$$u_i = \begin{cases} p_i - v_i & \text{if AP } i \text{ is selected} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

The utility represents therefore the residual gain of owner i obtained from the leased capacity of its AP. Obviously, when AP i is not used, the utility of its owner is null, since both the paid price and the private valuation are null.

3.2 Network Model

We observe that the transmission rate and the channel utilization required to satisfy the bandwidth demand of the data traffic generated by mobile customers depend on the channel condition between the smartphone of the mobile customer and the access point to which it can be connected; hence, the allocation scheme influences the number of APs used for the mobile data offloading. Given the amount of traffic d_j of its mobile customer MC_j , the operator computes the vector of channel utilizations, $\vec{o}_j = [o_{j1} \ o_{j2} \ \dots \ o_{ji} \ \dots \ o_{jn}]$, where each pair (j, i) refers to a possible allocation of MC j to AP i , whereas n represents the number of APs in the network. Channel utilizations are computed as follows:

$$o_{ji} = \frac{d_j}{r_{ji}}, \quad (2)$$

where the element o_{ji} represents the channel utilization of AP i when it is used to offload the data traffic of

MC j , and it is computed as the ratio between the traffic demand d_j and the maximum achievable transmission rate of the wireless link that might connect MC j and AP i , r_{ji} . Note that this latter value can be easily obtained from the MAC layer through a scanning of the wireless channels, which is performed periodically by all network devices. To take into consideration the uncertainty related to traffic description in wireless systems and prevent throughput collapse caused by the contention level, we increase the MC bandwidth d_j by a fixed margin, which is computed according to the recent model presented in [22]. Finally, we observe that our model can be extended to consider other types of SBS by slightly modifying expression (2) according to the specific wireless access technology.

4 OPTIMAL AUCTION FOR MOBILE DATA OFFLOADING

This section presents the combinatorial reverse auction we propose to jointly select the wireless APs and the MCs data connections that can be offloaded from the cellular network when WiFi resources cannot satisfy the aggregated traffic demand (Section 4.1). We show that a payment rule, which considers only the variation of the objective function solving the ILP problem with and without the winner, provides no incentive to participate to the reverse auction under such general assumptions. Therefore, we define a *new payment rule* to guarantee individual rationality and truthfulness, as demonstrated in Section 4.2. The algorithm implementing the optimal auction is then presented in Section 4.3.

4.1 Optimal Allocation

Hereafter, we formalize the Integer Linear Programming (ILP) model which provides the optimal allocation for the auction, namely the APs to be purchased and the mobile data traffic that can be offloaded. We first describe the sets and variables used in our model, then we provide the ILP description of the problem.

Let \mathcal{M} denote the set of mobile customer devices (MCs), and \mathcal{A} the set of wireless access points (APs) whose owners participate to the reverse auction of the mobile operator. Let us define $\mathcal{M}_i \subseteq \mathcal{M}, i \in \mathcal{A}$ as the set of MCs that are covered by AP i (i.e., the MCs that are in the radio range of AP i).

We can now introduce the decision variables used in our ILP model. Binary variables $x_i, i \in \mathcal{A}$, indicate which residential users win the auction, i.e., the APs whose available capacity is exploited by the mobile operator to serve the extra-traffic of its MCs ($x_i = 1$ if the available capacity of AP i is used, 0 otherwise). Binary variables $y_{ji}, i \in \mathcal{A}, j \in \mathcal{M}$, provide the assignment of MCs to APs ($y_{ji} = 1$ if MC j is assigned to AP i , 0 otherwise).

Given the above definitions and notation, the reverse combinatorial auction problem with partial covering of

mobile customers can be stated as follows:

$$\min f(x, y) = \sum_{i \in \mathcal{A}} b_i \cdot x_i - c \cdot \sum_{i \in \mathcal{A}} \sum_{j \in \mathcal{M}_i} y_{ji} \quad (3)$$

s.t.

$$y_{ji} \leq x_i \quad \forall i \in \mathcal{A}, \forall j \in \mathcal{M}_i \quad (4)$$

$$\sum_{i \in \mathcal{A}} y_{ji} \leq 1 \quad \forall j \in \mathcal{M} \quad (5)$$

$$\sum_{j \in \mathcal{M}_i} y_{ji} o_{ji} \leq 1 \quad \forall i \in \mathcal{A} \quad (6)$$

$$\sum_{j \in \mathcal{M}_i} y_{ji} d_j \leq x_i C_i \quad \forall i \in \mathcal{A} \quad (7)$$

$$y_{ji} = 0 \quad \forall i \in \mathcal{A}, \forall j \notin \mathcal{M}_i \quad (8)$$

$$x_i, y_{ji} \in \{0, 1\} \quad \forall i \in \mathcal{A}, \forall j \in \mathcal{M}. \quad (9)$$

The first term of the objective function (3), $\sum_{i \in \mathcal{A}} b_i \cdot x_i$, represents the total cost paid by the operator to lease the APs used for the data offloading of its mobile network. The second term, $\sum_{i \in \mathcal{A}} \sum_{j \in \mathcal{M}_i} c \cdot y_{ji}$, aims at maximizing the offloading of data connections from the cellular to the rented WiFi networks. The parameter $c > 0$ is a trade-off value between these two opposing objectives, and it can be seen as the gain of the operator obtained by offloading the traffic of MC j to AP i . Constraints (4) are coherence constraints ensuring that only the access points that win the auction can be used to serve mobile customer connections. The set of constraints (5) ensures that mobile data connections are served using *at most* one leased access point.

Constraints (6) and (7) prevent the allocation of an overall traffic demand that cannot be satisfied by an access point, due to the maximum achievable transmission rate of the wireless channel and the limited capacity of the Internet connection made available by the residential user, while constraints (8) prevent the assignment of MCs to APs that are not in the reciprocal radio range. Note that the channel assignment of access points can be optimized in order to reduce interference effects among nearby devices by using classical coloring algorithms coupled with the IEEE 802.11k standard. Finally, constraints (9) ensure the integrality of the binary decision variables.

Since the operator aims at offloading its mobile network as much as possible, the parameter c should be set as pointed out by the following proposition.

Proposition 4.1. *In order to offload the maximum amount of traffic of Mobile Clients, the value of the parameter c must be greater than the maximum bid, namely $c > \max\{b_i\}$.*

In fact, it is easy to prove that when parameter $c > \max\{b_i\}$, we always get an improvement in terms of minimization of the objective function by selecting an additional AP h , since $b_h \leq \max\{b_i\} < c \cdot \sum_{j \in \mathcal{M}_h} y_{jh}$.

We underline that our model can be easily extended to consider other scenarios where, for example, the amount of allocated radio resources represents a more important metric to select the MCs that should be offloaded. Indeed, the objective function can be modified to consider as connection cost a parameter proportional to the radio

resources utilization of the Base Station, \hat{o}_j , as follows:

$$f(x, y) = \sum_{i \in \mathcal{A}} b_i \cdot x_i - \sum_{i \in \mathcal{A}} \sum_{j \in \mathcal{M}_i} c \cdot \hat{o}_j y_{ji} =$$

$$\sum_{i \in \mathcal{A}} b_i \cdot x_i - \sum_{i \in \mathcal{A}} \sum_{j \in \mathcal{M}_i} \hat{c}_j y_{ji}$$

$$\hat{o}_j = \frac{d_j}{\hat{r}_j}$$

where \hat{r}_j is the physical rate of the mobile terminal connection j , which depends on the radio resources allocated by the BS scheduling algorithm. Therefore, the value of \hat{r}_j can be easily obtained by the operator. Note that to offload the maximum number of connections, the condition stated in Proposition 4.1 must be modified as follows:

Proposition 4.2. *In order to offload the maximum amount of traffic of Mobile Clients, the minimum value among the parameters \hat{c}_j must be greater than the maximum bid, namely $\min\{\hat{c}_j\} > \max\{b_i\}$.*

We observe that the condition stated in Proposition 4.1 permits to achieve the highest energy savings, since the higher the number of offloaded connections, the larger is the set of BSs that can be switched off. Nonetheless, when the operator wants to limit the maximum offloading cost or have some guarantee on the price paid to third party APs, the cost c acts as a *reserve price*, excluding all those players that have submitted higher bids than the value that the operator puts on the MC connection. Therefore, the operator can choose the *reserve price* c to limit the maximum offloading cost, without affecting either the problem feasibility or the solution properties (i.e., individual rationality and truthfulness).

4.2 Payment Rule

Having defined the ILP model representing the optimal auction, we now illustrate the payment rule and the conditions that force AP owners to ask their real valuation for the utilization of the capacity that they make available through their access points. First, we demonstrate that a classical payment rule, which considers only the difference of the objective function minimized with and without the winner's presence, cannot be directly applied to the problem analyzed in this paper. Then, we propose a new payment rule that guarantees both individual rationality and truthfulness (incentive compatibility).

In reverse auctions, payment rules usually define the price paid to winner i as the *damage* that it causes to other participants, which can be computed as the difference between the optimal value of the objective function obtained with and without i participation.

Mathematically, let (x, y) be the solution to the ILP problem (3)-(9), and $f^{-i}(x, y)$ the value of the objective function without considering the bid for AP i , i.e., $f^{-i}(x, y) = \sum_{k \in \mathcal{A} \setminus \{i\}} b_k \cdot x_k - \sum_{k \in \mathcal{A} \setminus \{i\}} \sum_{j \in \mathcal{M}_k} c \cdot y_{jk}$.

Furthermore, let (x^{-i}, y^{-i}) denote the solution to the same problem without considering AP i (i.e., forcing

$x_i = 0$ as additional constraint to the original problem), and $f(x^{-i}, y^{-i}) = \sum_{k \in \mathcal{A} \setminus \{i\}} b_k \cdot x_k^{-i} - \sum_{k \in \mathcal{A} \setminus \{i\}} \sum_{j \in \mathcal{M}_k} c \cdot y_{jk}^{-i}$ the value of the corresponding objective function.

The price paid to winner i using the aforementioned rule is therefore equal to $p_i = f(x^{-i}, y^{-i}) - f^{-i}(x, y)$ ¹.

Theorem 4.3. *The payment scheme that considers only the difference in the objective function of the problem (3)-(9) caused by the winner's presence does not guarantee individual rationality when the connections covered by the winner i can be assigned to a more expensive AP $e \in \mathcal{A}$, i.e., $p_i = f(x^{-i}, y^{-i}) - f^{-i}(x, y) < v_i$.*

PROOF: To prove this theorem, we have to show that the winner of the auction $i \in \mathcal{A}$ is paid less than the value it asked for using its AP, i.e., $u_i < 0 \Leftrightarrow p_i < v_i$. To this end, we write the payment rule as follows:

$$p_i = f(x^{-i}, y^{-i}) - f^{-i}(x, y) =$$

$$= \sum_{k \in \mathcal{A} \setminus \{i\}} b_k \cdot x_k^{-i} - \sum_{k \in \mathcal{A} \setminus \{i\}} b_k \cdot x_k +$$

$$+ c \cdot \sum_{k \in \mathcal{A} \setminus \{i\}} \sum_{j \in \mathcal{M}} (y_{jk} - y_{jk}^{-i})$$
(10)

The absolute value of the term $q_i \triangleq \sum_{k \in \mathcal{A} \setminus \{i\}} \sum_{j \in \mathcal{M}} (y_{jk} - y_{jk}^{-i})$ represents the number of connections that can be offloaded even without i . Furthermore, $\sum_{k \in \mathcal{A} \setminus \{i\}} \sum_{j \in \mathcal{M}} y_{jk}^{-i} \geq \sum_{k \in \mathcal{A} \setminus \{i\}} \sum_{j \in \mathcal{M}} y_{jk}$, therefore $q_i \leq 0$. Indeed, according to Proposition 4.1, which forces the maximal covering of MCs connections, the number of connections assigned to i in the optimal solution y is at least equal to the one that can be offloaded without i in solution y^{-i} . Therefore, the enlargement of the solution space by the addition of a variable (i.e., x_i) can only increase the number of covered connections (recall, however, that the connections assigned to i are not considered in the payment rule (10)).

When the connections served by i can be assigned to a more expensive AP $e \in \mathcal{A}$, the relation $q_i = \sum_{k \in \mathcal{A} \setminus \{i\}} \sum_{j \in \mathcal{M}} (y_{jk} - y_{jk}^{-i}) < 0$ holds, thus $p_i = b_e + c \cdot q_i < 0 < v_i$, since $c > \max\{b_h\}$. \square

To better clarify the problem stated in Theorem 4.3, hereafter we present an example in which the individual rationality is not guaranteed. Let us refer to the network scenario illustrated in Figure 2, with three APs and two MCs. We assume that $b_1 > b_2 > b_3$, $C_1 = C_2 = C_3 = C$, and that $d_1 + d_2 > C$. The APs selected as winners are AP_2 and AP_3 . In order to determine the price paid to AP_3 according to the rule (10), we need to compute the optimal allocation and the corresponding value of the objective function with and without AP_3 . With AP_3 , the best solution (x, y) results in the assignments $y_{12} = y_{23} = 1$ (MC_1 is assigned to AP_2 , MC_2 to AP_3). The value of the objective function is equal to $f(x, y) = b_2 + b_3 - 2c$, hence $f^{-3}(x, y) = b_2 - c$.

On the contrary, without AP_3 , the best solution

1. Note that this rule is equivalent to the VCG scheme when all MCs can be offloaded even without i , as in [8].

(x^{-3}, y^{-3}) results in the assignments $y_{11}^{-3} = y_{22}^{-3} = 1$ (MC_1 is assigned to AP_1 , MC_2 to AP_2). The value of the objective function is equal in this case to $f(x^{-3}, y^{-3}) = b_1 + b_2 - 2c$.

The price paid to AP_3 according to the rule (10) is $p_3 = f(x^{-3}, y^{-3}) - f^{-3}(x, y) = b_1 - c < v_3$, hence $u_3 = p_3 - v_3 < 0$, as $c = \max\{b_i\} = b_1 = v_1 > v_3$. Therefore, the owner of AP_3 has no incentive to participate to the offloading market, since its utility is negative.




Fig. 2: Example scenario showing that the payment rule (10) does not guarantee individual rationality.

Theorem 4.3 and the example scenario point out that the payment rule (10) cannot be directly applied to the mobile data offloading problem with partial covering of mobile clients.

In order to guarantee individual rationality and make the payment acting as an incentive for the participation, we propose to modify the rule (10) adding a new term to the price paid to the winner that depends on the number of connections that its presence permits to offload, according to the following expression:

$$p_i = f(x^{-i}, y^{-i}) - f^{-i}(x, y) + c \cdot \sum_{j \in \mathcal{M}_i} y_{ji}. \quad (11)$$

Theorem 4.4 (Individual Rationality of (11)). *The payment rule defined in Equation (11) satisfies the individual rationality property, i.e., $\forall i \in \mathcal{A} : x_i = 1, p_i = f(x^{-i}, y^{-i}) - f^{-i}(x, y) + c \cdot \sum_{j \in \mathcal{M}_i} y_{ji} \geq v_i$.*

PROOF: To prove Theorem 4.4 we can observe that

$$\begin{aligned} p_i &= f(x^{-i}, y^{-i}) - f^{-i}(x, y) + c \cdot \sum_{j \in \mathcal{M}_i} y_{ji} = \\ &= \sum_{k \in \mathcal{A} \setminus \{i\}} b_k \cdot x_k^{-i} - \sum_{k \in \mathcal{A} \setminus \{i\}} b_k \cdot x_k + \\ &+ c \cdot \left[\sum_{j \in \mathcal{M}_i} y_{ji} + \sum_{k \in \mathcal{A} \setminus \{i\}} \sum_{j \in \mathcal{M}} (y_{jk} - y_{jk}^{-i}) \right] = \\ &= \left[\sum_{k \in \mathcal{A} \setminus \{i\}} b_k \cdot x_k^{-i} - \sum_{k \in \mathcal{A} \setminus \{i\}} b_k \cdot x_k \right] + c \cdot \left[\sum_{j \in \mathcal{M}_i} y_{ji} + q_i \right]. \end{aligned}$$

The first term $\left[\sum_{k \in \mathcal{A} \setminus \{i\}} b_k \cdot x_k^{-i} - \sum_{k \in \mathcal{A} \setminus \{i\}} b_k \cdot x_k \right]$ is always positive, since the optimal solution without i , (x^{-i}, y^{-i}) , always contains APs that are more expensive than the solution with i , (x, y) .

The second term $\sum_{j \in \mathcal{M}_i} y_{ji} + q_i = \sum_{j \in \mathcal{M}_i} y_{ji} -$

$\sum_{k \in \mathcal{A} \setminus \{i\}} \sum_{j \in \mathcal{M}} (y_{jk}^{-i} - y_{jk})$ represents the number of connections that cannot be covered without i (recall that $q_i \leq 0$).

For the same reasons explained in Theorem 4.3, we have $\sum_{j \in \mathcal{M}_i} y_{ji} \geq -q_i$, hence $\sum_{j \in \mathcal{M}_i} y_{ji} + q_i \geq 0$.

The utility of the AP owner i increases proportionally of the same quantity:

$$\begin{aligned} u_i &= p_i - v_i = \\ &= \sum_{k \in \mathcal{A} \setminus \{i\}} b_k \cdot x_k^{-i} - \sum_{k \in \mathcal{A} \setminus \{i\}} b_k \cdot x_k + c \cdot \left[\sum_{j \in \mathcal{M}_i} y_{ji} + q_i \right] - v_i. \end{aligned}$$

□

With our payment rule, the operator pays to the winners of the auction their contribution to the social welfare (i.e., the money that their presence permits to save) plus an additional incentive that depends on the connections that without their presence cannot be offloaded from the RAN, thus forcing to keep the Base Stations turned on.

Theorem 4.5 (Truthfulness of (11)). *The payment rule defined in Equation (11) satisfies the truthfulness property (incentive compatibility).*

PROOF: To prove this theorem, we have to show that $u(v_i) \geq u(v'_i)$, $\forall v'_i \neq v_i$, that is, an AP owner i cannot increase its utility by bidding untruthfully, namely $b_i = v'_i \neq v_i$.

Let (x, y) and (x', y') be the solutions to the problem (3)-(9), when the AP owner i declares v_i and v'_i , respectively. Furthermore, let (x^{-i}, y^{-i}) denote the solution to the same problem without considering the AP i (i.e., forcing $x_i = 0$ as additional constraint to the original problem). Note that $x_i^{-i} = x_i'^{-i}$.

The utility of i when it declares v_i , $u(v_i)$, is equal to:

$$\begin{aligned} u(v_i) &= p_i(v_i, x, y) - v_i = \\ &= \sum_{k \in \mathcal{A} \setminus \{i\}} v_k \cdot x_k^{-i} - \sum_{k \in \mathcal{A} \setminus \{i\}} \sum_{j \in \mathcal{M}} c \cdot y_{jk}^{-i} + \\ &- \sum_{k \in \mathcal{A} \setminus \{i\}} v_k \cdot x_k + \sum_{k \in \mathcal{A} \setminus \{i\}} \sum_{j \in \mathcal{M}} c \cdot y_{jk} + \\ &+ \sum_{j \in \mathcal{M}_i} c \cdot y_{ji} - v_i = \\ &= \sum_{k \in \mathcal{A} \setminus \{i\}} v_k \cdot x_k^{-i} - \sum_{k \in \mathcal{A} \setminus \{i\}} \sum_{j \in \mathcal{M}} c \cdot y_{jk}^{-i} + \\ &- \left(\sum_{k \in \mathcal{A}} v_k \cdot x_k - \sum_{k \in \mathcal{A}} \sum_{j \in \mathcal{M}} c \cdot y_{jk} \right), \end{aligned}$$

whereas, when it declares v'_i , the utility is equal to:

$$\begin{aligned} u(v'_i) &= p_i(v'_i, x', y') - v_i = \\ &= \sum_{k \in \mathcal{A} \setminus \{i\}} v_k \cdot x_k^{-i} - \sum_{k \in \mathcal{A} \setminus \{i\}} \sum_{j \in \mathcal{M}} c \cdot y_{jk}^{-i} + \\ &- \left(\sum_{k \in \mathcal{A} \setminus \{i\}} v_k \cdot x'_k + v_i - \sum_{k \in \mathcal{A} \setminus \{i\}} \sum_{j \in \mathcal{M}} c \cdot y'_{jk} - \sum_{j \in \mathcal{M}_i} c \cdot y'_{ji} \right). \end{aligned}$$

Since (x, y) is the solution that minimizes the objective function (3), $(x, y) = \arg \min_{x \in X, y \in Y} \sum_{i \in \mathcal{A}} b_i \cdot x_i - \sum_{i \in \mathcal{A}} \sum_{j \in \mathcal{M}_i} c \cdot y_{ji}$,

Algorithm 1: Optimal Reverse Auction

Input : $\mathcal{M}, \mathcal{A}, b_i, C_i, c, d_{ji}$
Output: x_i, p_i, y_{ji}
1 Compute channel utilizations o_{ji} ;
2 $x_i \leftarrow$ Solve the ILP model (3)-(9);
3 **foreach** $i \in \mathcal{A} : x_i = 1$ **do**
 $p_i = f(x^{-i}, y^{-i}) - f^{-i}(x, y) + c \cdot \sum_{j \in \mathcal{M}_i} y_{ji}$;
end

we have:

$$\sum_{k \in \mathcal{A}} v_k \cdot x_k - \sum_{k \in \mathcal{A}} \sum_{j \in \mathcal{M}} c \cdot y_{jk} \leq \sum_{k \in \mathcal{A} \setminus \{i\}} v_k \cdot x'_k + v_i - \sum_{k \in \mathcal{A} \setminus \{i\}} \sum_{j \in \mathcal{M}} c \cdot y'_{jk} - \sum_{j \in \mathcal{M}_i} c \cdot y'_{ji},$$

therefore $u(v_i) \geq u(v'_i)$, and the AP owner i cannot increase its utility by bidding unilaterally untruthfully. \square

We underline that when the APs provide enough capacity to offload the data traffic and any MC can be handled by multiple APs (i.e., the third party network provides enough capacity and redundancy to offload the MC traffic), the payment rule (10) is equivalent to (11), since there are no externalities due to traffic covering.

4.3 Optimal Algorithm

Hereafter, we illustrate the algorithm implementing the optimal mobile data offloading auction run by an operator to select the cheapest APs that are used to offload the data connections of the mobile customers from its RAN.

Algorithm 1 receives as input the parameters which describe the network topology and all offers from the APs' owners; these latter are composed of the capacity C_i made available through the APs and the cost b_i . It produces as output the subset of APs that will be used to offload the data traffic of the mobile terminals, ($i \in \mathcal{A} : x_i = 1$) and the price paid to their owners, p_i , as well as the assignment of the data connections to the selected APs ($(j, i) \in \mathcal{M} \times \mathcal{A} : y_{ji} = 1$).

The algorithm proceeds in 3 steps. In step 1, the demands of mobile customers connections are transformed into equivalent channel utilizations, using the achievable transmission rate of the links that can be established with all nearby APs, according to Equation (2). Step 2 consists in solving the ILP model to find the allocation that minimizes the objective function (3). Finally, in step 3, the operator computes the prices paid to the owners of the APs selected by the previous step according to our rule (11), which guarantees a truthful auction.

The optimal reverse auction problem detailed in Algorithm 1 is NP-hard. Indeed, it can be shown that the knapsack problem can be polynomially reduced to the problem (3)-(9). Therefore, an operator can hardly find a solution to reconfigure its mobile network on-the-fly, since the computation time necessary to solve large and real-life network instances increases very sharply with the network size and density. However, we observe

that in small-size network scenarios, where the set of covered mobile clients \mathcal{M}_i have minimal overlap, we can optimally solve the mobile offloading problem.

5 GREEDY AUCTION FOR MOBILE DATA OFFLOADING

In the following, we present three alternative versions of a very efficient algorithm to solve the allocation problem in polynomial time. Furthermore, we demonstrate that such algorithm preserves the *truthfulness* property, so that the proposed trading marketplace is robust against any cheating behavior attempted to rule out honest AP owners.

5.1 Greedy Algorithm

The greedy auction is summarized in Algorithm 2, and it is composed of two main phases: (1) the *allocation* phase, which selects the APs that are used to offload the maximum amount of data traffic generated by mobile customers, and (2) the *payment* phase, which establishes the price paid to each winner as a function of the first unused AP in the sorted list (the first loser). This latter is also referred to as *critical access point* for i (denoted by s), and the price asked by its owner as *critical value* for i , which will be denoted as p_s .

Algorithm 2: Greedy Reverse Auction

Input : $\mathcal{M}, \mathcal{A}, b_i, C_i, d_{ji}, o_{ji}$
Output: x_i, p_i, y_{ji}
1 $L \leftarrow \text{Sort} \left(i \in \mathcal{A}, \frac{b_i}{|\mathcal{M}_i|}, \text{"non-decr"} \right)$;
2 $L \leftarrow L \setminus \text{last}(L)$;
 $U \leftarrow \mathcal{M}$;
3 **while** $L \neq \emptyset \wedge U \neq \emptyset$ **do**
 $i \leftarrow \text{Next}(L)$; $x_i \leftarrow 1$;
 $\mathcal{V}_i \leftarrow \text{Sort} (j \in \mathcal{M}_i, o_{ji}, \text{"non-decr"})$;
4 **while** $\sum_{j \in \mathcal{M}_i} y_{ji} o_{ji} \leq 1 \wedge \sum_{j \in \mathcal{M}_i} y_{ji} d_j \leq x_i C_i$ **do**
 $j = \text{Next}(\mathcal{V}_i)$;
 if $\sum_{h \in \mathcal{A}} y_{jh} = 0$ **then**
 $y_{ji} \leftarrow 1$;
 $U = U \setminus \{j\}$;
 end
 end
5 Refine_Assignment($\{j \in \mathcal{M}_i : y_{ji} = 0\}, \mathcal{A} \setminus L$);
 end
6 $s \leftarrow \text{Next}(L)$;
7 **foreach** $i \in \mathcal{A} : x_i = 1$ **do**
 $p_i \leftarrow \frac{b_s}{|\mathcal{M}_s|} |\mathcal{M}_i| = p_s |\mathcal{M}_i|$;
end

The greedy allocation phase (steps 1–6) sorts the set of APs that participate to the auction in ascending order, according to three alternative rules, as illustrated in Table 2. For the sake of clarity, we explain both phases of the greedy auction considering the first alternative (*greedy MC*, which we will also denote as G.1 for simplicity), namely the rule that sorts all APs in non-decreasing order of their submitted bids per number of covered MCs (i.e., the MCs that they may serve), $b_i/|\mathcal{M}_i|$. Each element of the sorted list is selected as winner until all available APs are selected or there exist MCs whose traffic has not

yet been offloaded to any AP. The assignment procedure in step 4 assigns to each AP $i \in \mathcal{A}$ selected as winner the maximum number of unsatisfied MCs in its radio range ($j \in \mathcal{M}_i : \sum_{h \in \mathcal{A}} y_{jh} = 0$) such that either the wireless channel is not saturated (i.e., its utilization is lower than 1) or the overall traffic demand does not exceed the capacity of the wired connection. Before performing a new iteration to select a new AP from the list L , the function *Refine_Assignment* in step 5 attempts to assign the remaining unsatisfied MCs to those APs that have been selected as winners in previous iterations. Indeed, previous winners may have enough spare capacity to serve also these unsatisfied MCs.

After selecting the winning APs opportunistically used by the operator and having performed the assignment of MCs to such APs, step 6 returns the *critical access point* $s \in \mathcal{A}$. AP s is the first unselected AP, or the last available AP of the sorted list L , which is removed in step 2 from the list to guarantee the incentive compatibility property. Eventually, AP s is used in step 7 to compute the prices paid by the operator to the winners for offloading its mobile network.

As we will demonstrate in the next section, Algorithm 2 implements a truthful auction. In fact, the allocation phase satisfies the monotonicity property (recall that the APs are sorted in non-decreasing order of their bid per number of covered mobile customers), and there exists a critical value which determines if the AP owners' bid are satisfied or not.

The proposed greedy auction implemented by Algorithm 2 has time complexity $O(n^2m)$ (with $m = |\mathcal{M}|$ and $n = |\mathcal{A}|$). Indeed, assuming that every summation has time complexity $O(1)$, each iteration k of the loop in the greedy allocation phase requires m operations for step 4 and $m \cdot |\mathcal{A} \setminus L| = m \cdot w_k$ assignment attempts within the function *Refine_Assignment*, where w_k represents the cumulative number of winners selected up to iteration k . Note that w_k has a unitary increase at each step k , and it takes value from 0 to $n - 1$. Therefore, the maximum number of iterations due to steps 4 and 5 executed throughout loop 3 cannot be larger than $\frac{n(n+1)}{2}$, thus resulting in $\frac{n(n+1)}{2} \cdot m$ total assignment operations.

We observe that the utilization of any sorting rule that does not affect the monotonicity property of the allocation phase still results in a truthful auction. Therefore, we design two alternative versions of the greedy auction that select the APs according to their price per channel utilization, as indicated in Table 2. Indeed, considering the resource utilization, which depends both on the traffic and the achievable rate of the MC connection, results in better performance. More specifically, the *greedy use* scheme (denoted by G.2) ranks APs according to their price per overall channel utilization, considering all MCs that can potentially be assigned to an AP. On the contrary, the *greedy max use* approach (G.3) computes the unitary price considering the subset of MCs whose aggregated demand can be satisfied by the AP access

capacity (or equivalently, whose aggregated channel utilization is lower than 1). In other words, the sorting and payment rules use the larger subset $\mathcal{O}_i = E \in \mathcal{P}(\mathcal{M}_i) : \sum_{j \in E} o_{ji} \leq 1$ ($\mathcal{P}(\mathcal{M}_i)$ is the partition set of \mathcal{M}_i). The *greedy use* rule aims at selecting as winners the APs that can potentially offload the highest portion of data traffic, whereas the *greedy max use* scheme leases the APs that can effectively satisfy the aggregated demand.

5.2 Truthfulness Analysis

Having described the main phases of the *greedy reverse auction*, hereafter we prove that our mechanism satisfies the *incentive compatibility* property (truthfulness). We recall that an auction mechanism is truthful if the dominant strategy for each rational bidder i is to declare always its real private valuation $b_i = v_i$. This property guarantees that selfish bidders cannot benefit from cheating, thus preventing the strategic manipulation of the marketplace.

The following lemmas (5.1 - 5.2) prove that the allocation phase of Algorithm 2 (steps 1–6) satisfies the *monotonicity* property and guarantees the existence of a *critical value* [23], which provide the basis to demonstrate Theorems 5.3 and 5.4.

Lemma 5.1. *If AP owner i is selected by the allocation algorithm when it bids b_i , then AP owner i is still selected if i decreases its bid b'_i , $b'_i < b_i$.*

PROOF: Let L and L' be two sorted lists corresponding to b_i and b'_i , respectively. Let us define $rank(i, L)$ as a monotonic decreasing function of AP owner i position in the list L . Since $\frac{b'_i}{|\mathcal{M}_i|} < \frac{b_i}{|\mathcal{M}_i|}$, the sorting algorithm in the greedy allocation phase (Algorithm 2) moves i in a better position, i.e., $rank(i, L') > rank(i, L)$. Therefore, the rank of i can only increase if AP owner i submits a lower bid (i.e., i offers a lower price), resulting in a different order of the set of access points that are selected as winners by the operator to offload the traffic of its mobile customers, which implies that if AP owner i wins by bidding b_i , it is selected even with a lower bid $b'_i < b_i$. \square

Lemma 5.2. *For each AP owner i , the greedy Algorithm 2 provides the critical value $p_s = \frac{b_s}{|\mathcal{M}_s|}$, which determines whether AP owner i is selected as winner of the reverse auction.*

PROOF: The proof is straightforward, since Algorithm 2 scans the list L of APs in non-decreasing order of their bids per number of covered mobile customers $\left(\frac{b_i}{|\mathcal{M}_i|}\right)$ until the maximum amount of data traffic generated by the mobile customers is satisfied or all but the last AP are selected as winners. The *critical value* is then equal to the ratio $p_s = \frac{b_s}{|\mathcal{M}_s|}$ submitted by the owner of the first unselected or the last AP owner s . \square

Note that if we do not exclude the last and most expensive AP from the auction, an AP owner may ask a high value for the utilization of its unexploited Internet

TABLE 2: Sorting and payment rules for the greedy auction

Greedy MC (G.1)	Greedy Use (G.2)	Greedy Max Use (G.3)
$L \leftarrow \text{Sort} \left(i \in \mathcal{A}, \frac{b_i}{ \mathcal{M}_i }, \text{"non-decr"} \right)$	$L \leftarrow \text{Sort} \left(i \in \mathcal{A}, \frac{b_i}{\sum_{j \in \mathcal{M}_i} o_{ji}}, \text{"non-decr"} \right)$	$L \leftarrow \text{Sort} \left(i \in \mathcal{A}, \frac{b_i}{\sum_{j \in \mathcal{O}_i} o_{ji}}, \text{"non-decr"} \right)$
$p_i \leftarrow \frac{b_s}{ \mathcal{M}_s } \mathcal{M}_i $	$p_i \leftarrow \frac{b_s}{ \sum_{j \in \mathcal{M}_s} o_{js} } \sum_{j \in \mathcal{M}_i} o_{ji} $	$p_i \leftarrow \frac{b_s}{ \sum_{j \in \mathcal{O}_s} o_{js} } \sum_{j \in \mathcal{O}_i} o_{ji} $

connection, $b_i \gg v_i$, being assured that its bid will be always satisfied. Therefore, the removal of the last AP ensures that all participants declare their real value v_i .

Theorem 5.3 (Individual Rationality of Algorithm 2). *Each AP owner i selected as winner by the Greedy Algorithm is paid at least the price it asked for the utilization of the unexploited capacity of its Internet connection, $p_i \geq b_i$.*

PROOF: To show that $p_i \geq b_i$, we need to demonstrate that the *critical value* times the number of covered mobile customers paid to winner i is at least equal to its bid b_i . Each winner i is paid the *critical value* $\frac{b_s}{|\mathcal{M}_s|}$ times the number of mobile customers that it can cover (i.e., $|\mathcal{M}_i|$). Recall that the list L of APs is sorted in non-decreasing order of the ratio $\frac{b_i}{|\mathcal{M}_i|}$, therefore the relation $p_i = \frac{b_s}{|\mathcal{M}_s|} |\mathcal{M}_i| \geq b_i$ holds, since either s asked a higher price for using its AP (i.e., $b_s \geq b_i$) or the access network of its AP can be used to offload a lower traffic demand (i.e., $|\mathcal{M}_s| \leq |\mathcal{M}_i|$). \square

Theorem 5.4 (Truthfulness of Algorithm 2). *Algorithm 2 implements a truthful auction.*

PROOF: We prove the theorem by showing that no participant to the marketplace can increase its utility by asking a price b_i different from its private valuation v_i for the utilization of its AP. We underline that the utility of AP owner i does not change by bidding either v_i or b_i , since it is defined as $u_i(x) = p_i - v_i$. We must consider two cases, namely (A) $b_i < v_i$ (lower price), and (B) $b_i > v_i$ (higher price). For each case, we must consider all possible four outcomes, detailed in the following. Let us start with case (A) by considering the following cases.

A.1: AP owner i wins either by bidding b_i or v_i .

If AP owner i wins by bidding either b_i or v_i , then i is ranked in a better place in the list L when it submits b_i , since the list is sorted in non-decreasing order of the bids per number of covered mobile customers. However, this changes only the order of the set of winners, which does not affect the *critical value* p_s that is still given by the following expressions: $p_s = \frac{b_s}{|\mathcal{M}_s|}$. Hence, the price paid by the winner does not vary, $p_i = p_s \cdot |\mathcal{M}_i|$. Therefore, the utility does not change: $u_i(b_i) = u_i(v_i)$.

A.2: AP owner i wins by bidding b_i but loses with v_i .

If AP owner i wins by submitting b_i but loses with v_i , then there exists a *critical value* $p_s = \frac{b_s}{|\mathcal{M}_s|}$ such that $\frac{b_i}{|\mathcal{M}_i|} < \frac{b_s}{|\mathcal{M}_s|} < \frac{v_i}{|\mathcal{M}_i|}$.

Due to the monotonic property of the allocation algorithm, the private valuation of i is higher than the price paid when it submits b_i , i.e., $b_i < p_i < v_i$. Therefore, the utility perceived by i is negative, $u_i(b_i) = p_i - v_i < 0$,

hence it is better off losing the auction, since in this latter case its utility is null, $u_i(v_i) = 0$.

A.3: AP owner i loses by bidding b_i but wins with v_i .

Due to the monotonic property, this case is impossible, since by submitting a lower price, AP owner i will be placed in a better position of the sorted list L .

A.4: AP owner i loses either by bidding b_i or v_i .

If AP owner i loses by offering both b_i and v_i , due to the presence of cheaper access points, then its utility is always null: $u_i(b_i) = u_i(v_i) = 0$.

Similarly, for the case (B) $b_i > v_i$, we can demonstrate that AP owner i cannot increase its utility by asking a higher price than its private valuation for leasing its AP to the mobile operator. \square

Since Algorithm 2 implements a truthful auction (which means that selfish AP owners cannot benefit from manipulating their bids), a mobile operator can efficiently compute a solution for the reverse auction problem, being assured that all AP owners reveal their true price for leasing the available capacity of their APs.

Similar observations as those employed above can be easily formulated to prove the truthfulness of the other two sorting and payment rules.

We observe that the truthfulness of the greedy auction can be more easily demonstrated as in [24] under the assumption of sufficient capacity for offloading the whole traffic, thus ignoring capacity and covering constraints (6) and (7). Indeed, in this case by simply sorting the APs in increasing order of their bids per number of covered MCs, $\frac{b_i}{|\mathcal{M}_i|}$, we get the optimal solution that minimizes the overall offloading cost (i.e., the sum of the bids of the APs selected to offload the data traffic). Furthermore, to prove that the greedy allocation is truthful, we must also show that the bid per number of covered MCs of the first loser c , which is used as unitary price for the winners, satisfies the VCG rules with respect to the objective function $f(x, y) = \sum_{i \in \mathcal{A}} \frac{b_i}{|\mathcal{M}_i|} x_i$.

Let us assume without loss of generality that AP owners bids can be sorted according to their indexes as follows: $\frac{b_1}{|\mathcal{M}_1|} < \dots < \frac{b_k}{|\mathcal{M}_k|} < \frac{b_{k+1}}{|\mathcal{M}_{k+1}|} < \dots < \frac{b_n}{|\mathcal{M}_n|}$, and the first k out of n APs suffice to offload the whole MC traffic. The greedy allocation rule selects the first k APs and fixes the unitary price of each winner as $p_i = \frac{b_{k+1}}{|\mathcal{M}_{k+1}|}$, $1 \leq i \leq k$. On the other hand, we can easily see that whenever a winner i is removed from the AP set, we need to select also the $(k+1)$ th AP to offload the MCs' connections that were covered by i (the solution without i contains also AP $k+1$, i.e., $x_{k+1}^{-i} = 1$). Therefore, the unitary price computed according to the VCG payment scheme results $p_i = \sum_{k \in \mathcal{A} \setminus \{i\}} \frac{b_k}{|\mathcal{M}_k|} x_k^{-i} -$

$\sum_{k \in \mathcal{A} \setminus \{i\}} \frac{b_k}{|\mathcal{M}_k|} x_k = \frac{b_{k+1}}{|\mathcal{M}_{k+1}|} x_k$, which is exactly the value computed by the greedy algorithm.

5.3 Economic Efficiency Analysis

In the following, we quantify theoretically the economic efficiency gap between the greedy and optimal solutions. To this end, we consider a simple network scenario composed of one Mobile Client, A , and two Access Points ($\{1, 2\}$). The capacities made available by the two AP owners is large enough to accommodate the traffic transmitted over the wireless access interface by the MC. However, due to the different channel qualities, the utilization of the two access links, which can be exploited to offload the MC traffic, are $o_{A1} = d_{A1}/r_{A1} = 1/D$ (D is a positive parameter) and $o_{A2} = d_{A2}/r_{A2} = 1$, 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 minimizes the objective function $f(x, y)$ is the one that selects AP 1 and the corresponding social welfare is $SW^o = b_1 = 1 + \epsilon$. However, the greedy algorithm selects AP 2, since $b_2/o_2 < b_1/o_1$, and the social welfare is in this case equal to $SW^g = b_2 = D$. The ratio $\frac{SW^g}{SW^o}$ is therefore equal to $\frac{D}{1+\epsilon}$, and the Price of Anarchy tends to infinity with D :

$$PoA = \frac{SW^g}{SW^o} = \lim_{D \rightarrow \infty} \frac{D}{1+\epsilon} = \infty.$$

At the same time, we underline that the *economic gap* computed considering the *offloading cost*, which is the most important performance metric for the operator, is almost null, since the cost obtained using the greedy auction is only ϵ times larger than the cost computed using the optimal algorithm. Indeed, the Offloading Cost Ratio (*OCR*), which we define as the ratio between the greedy and optimal offloading costs, is independent of D , and even in this limiting case it results:

$$OCR = \frac{p^g(x, y)}{p^o(x, y)} = \frac{\sum_{i \in \mathcal{A}} p_i^g x_i^g}{\sum_{i \in \mathcal{A}} p_i^o x_i^o} = \lim_{D \rightarrow \infty} \frac{(1 + \epsilon)D}{D} = 1 + \epsilon.$$

where p_i^g and p_i^o are the prices computed with the greedy and optimal payment rules, respectively.

Note, however, that the conditions that lead to the example discussed above are hardly met in real network scenarios, since mobile operators do not handle such a small bandwidth granularity with their network equipment. Indeed, an infinitesimal channel utilization corresponds to offer an infinitesimal access bandwidth ($C_i = 1/\infty \geq \sum_{j \in \mathcal{M}_i} d_j y_{ji}$), which may not be realistically satisfied by any mobile operator. As a consequence, we can bound the PoA by simply fixing a minimum amount of access bandwidth that any bidder needs to provide in order to participate to the auction.

6 NUMERICAL RESULTS

This section presents the numerical results that illustrate the validity of the proposed approaches to implement the bandwidth trading marketplace for fostering mobile data offloading. More specifically, we aim at evaluating the impact of the device density and traffic load on the

performance of the mechanisms we designed for the mobile data offloading marketplace. We first describe the experimental methodology followed in our numerical analysis, then we analyze and discuss the performance achieved by the algorithms detailed in previous sections.

6.1 Experimental Methodology

For our numerical analysis, we refer to the scenarios designed within the FP7 European Project EARTH and described in [25]. More specifically, we extend the baseline reference deployment scenario composed of 7 cell sites, whose Inter-Site Distance (which specifies the distance between two sites) is fixed to 500 meters. Each macro Base Station (BS) installed on a central site serves 3 sectors, resulting in 21 sectors in total.

We vary the number of MCs and APs per sector in the ranges [2, 10] and [10, 15], respectively. Both MCs and the APs are placed randomly in the corresponding sector. Specifically, the MCs are deployed around each BS according to a bi-dimensional Gaussian distribution with standard deviation equal to approximately 160 meters, to take into account the proximity of MCs to the BSs. Indeed, cellular networks are usually designed considering the distribution of MCs; APs are instead scattered according to a uniform distribution inside each sector.

To evaluate the number of APs that are used for offloading the amount of traffic served by a BS, we consider the two following use cases. In the first scenario, we evenly divide the maximum bandwidth of a BS sector (42 Mbps using 64 QAM dual-cell MIMO as suggested in [25]) among all MCs inside that sector. Such value provides an indication on the AP density necessary to switch off BSs (or put them in deep sleep/idle mode). In contrast, in the second scenario, we fix the network topology and vary the traffic load of MCs to investigate the impact of heterogeneous demands on our mechanisms. In contrast, the bids submitted by any AP owner i , b_i , are drawn from a uniform distribution with mean value equal to 5 monetary units (e.g., US dollars) and interval size twice the average, both to compare the overall offloading cost to the installation cost of additional BSs and evaluate the fairness of the payments in the worst case scenario. However, we underline that these assumptions does not affect the main findings on the performance of our algorithms.

The maximum achievable transmission rate of the access links that can be established between MC j and any of its surrounding APs i , r_{ji} , is defined according to the reception sensitivity of the Wistron CM9 commercial wireless cards based on Atheros chipset². The path loss, which is necessary to evaluate the sensitivity of the receiving node, is computed according to the Friis propagation model. To model the uncertainty related to traffic description in wireless systems caused by the contention level at the frame layer, we consider a fixed margin to compute the effective bandwidth necessary

2. <http://www.lri.fr/~fmartignon/CM9.pdf>

to satisfy MC demands and avoid throughput collapse. Indeed, according to recent mathematical models [22] for 802.11 networks, in real traffic conditions (i.e., in a non-saturated regime, where stations' demands are characterized by bursty data rates), both collision probability and overall aggregate throughput tend to reach a stable, constant value, for increasing traffic loads. In this work, we discounted the access bandwidth of all APs by 55%, increasing the traffic demand of all MCs by a factor equal to 2.22, which corresponds to 10 saturated stations according to [22]. We underline that all the above assumptions do not affect the proposed algorithms, which are general and can be used to solve any network scenario.

In order to evaluate the performance of the solutions proposed to implement the mobile data offloading marketplace, we consider the following metrics:

- *Cost*: defined as the sum of the prices paid by the operator to all winners.
- *Served MCs*: fraction of MCs whose connections can be completely offloaded on winning APs.
- *Winners*: fraction of AP owners selected as winners among the participants to the auction.
- *Fairness*: we consider the *Jain's Fairness Index (JFI)* [26], defined according to Equation (12):

$$\text{Jain's Fairness Index} = \frac{(\sum_{i=1}^w \rho_i)^2}{w \cdot \sum_{i=1}^w \rho_i^2} \quad (12)$$

where ρ_i represents the ratio between the paid price and the traffic demand served by AP i , $\rho_i = p_i / \sum_{j \in \mathcal{M}_i} y_{ji} \cdot d_j$, whereas $w = \sum_{i \in \mathcal{A}} x_i$ represents the number of winners. The *Jain's Fairness Index* measures the spread of the price per unit of traffic paid by the operator to its winners, and varies from $1/w$ (no fairness) to 1 (perfect fairness).

For each network scenario we perform 100 independent measurements, computing very narrow 95% confidence intervals.

6.2 Analysis of Device Density

We first evaluate the effect of the number of MCs within each sector on the performance of our mechanisms, in order to evaluate rural, suburban, and urban scenarios. Specifically, we consider three different density levels, namely *low*, *medium* and *high*, corresponding respectively to 2, 4, and 6 MCs within each BS sector. In all scenarios, we vary the number of APs within each sector in the [10, 15] range.

We have further considered scenarios with 8 and 10 MCs within each BS sector. However, for the sake of brevity, we omit these results since they are very close to those observed with 6 MCs per sector.

Figures 3, 4, and 5 show the performance metrics of our four mechanisms as a function of the number of APs inside a BS sector for the *low-density*, *mid-density*, and *high-density* scenarios, respectively. The curves identified by labels "O.", "G.1", "G.2" and "G.3" illustrate,

respectively, the performance metrics computed using the optimal and greedy algorithms with the three sorting rules defined in Table 2 (i.e., G.1, APs sorted according to their *bids per number of covered MCs*, G.2, *bids per channel utilization*, and G.3, *bids per maximum channel utilization*). For the sake of clarity, the cost has been normalized with respect to the maximum value obtained over all the three scenarios.

In particular, Figures 3(a), 4(a), and 5(a) show the *overall cost* paid by the operator to offload the data traffic of its MCs with the proposed mechanisms. It can be observed that the *greedy use* scheme (G.2) well approaches the optimal solution in all scenarios. The slightly lower cost achieved by G.2 in Figure 3(a) is due to the additional contribution of the optimal payment rule (11). In contrast, the optimal solution always achieves the lowest value for the objective function, which represents the social welfare in our auction. Interestingly, the *greedy max use* approach (G.3) produces very different solutions with respect to the *greedy use* scheme. Specifically, in the *low-density* scenario the cost increases as a function of the APs, since the sorting rule sets approximately the same unitary price for all APs ignoring their positions, which instead can lead to lower costs as illustrated by the *greedy use* curve. However, as long as the MC density increases, the *greedy max use* scheme approaches the cost obtained using the *greedy use* scheme, since the capacity of the BS sector is spread among more MCs and their utilization of the spare APs capacity gets similar for the allocation rules of the two corresponding greedy algorithms. The *greedy MC* solution (G.1) provides similar results to the *optimal* algorithm in the *low-density* scenario, since the sorting rule provides higher ranks to those APs that are selected by the optimal allocation. However, considering only the number of covered MCs while completely ignoring their resource utilization results in higher costs when the MC density increases, because the lower MC demand can be better offloaded to closest APs (recall that the BS capacity is evenly distributed among MCs).

We further emphasize that all proposed solutions achieve high fairness, since the JFI, which we omit for the sake of brevity, is always higher than 0.85. In particular, the gap between optimal and greedy JFI values is negligible (the greedy curves are almost always overlapped to the optimal curve, and only in the worst case the JFI gap reaches 10%). Therefore, the data offloading price paid by the operator is almost independent of the AP selected by our proposed mechanisms.

Figures 3(b), 4(b), and 5(b) show the fraction of served MCs whose traffic demand can be offloaded onto WiFi APs. All schemes satisfy approximately the same number of MCs (all curves are practically overlapped). It can be further observed that the higher density has a positive effect on the number of offloaded MCs connections, as illustrated in the *mid* and *high-density* scenarios (Figures 4(b) and 5(b)). Indeed, the higher the MC density within a BS sector, the lower the amount of data traffic of each MC connection. Moreover, the higher density




Fig. 3: Performance metrics measured in the *low-density* scenario (2 MCs in each of the 21 sectors).




Fig. 4: Performance metrics measured in the *medium-density* scenario (4 MCs in each of the 21 sectors).




Fig. 5: Performance metrics measured in the *high-density* scenario (6 MCs in each of the 21 sectors).

increases the proximity among MCs and APs, thus increasing the transmission rate that can be used on the links established among these devices.

Figures 3(c), 4(c), and 5(c) show the fraction of APs that are selected as winners to offload the traffic from the mobile network. In the *low-density* scenario, while all schemes offload the same amount of data traffic, the greedy algorithms that sort the APs according to the channel utilization select a larger number of winners with respect to the *optimal* and *greedy MCs* solutions. Similarly to previous metrics, by spreading the BS capacity among a larger set of MCs we can reduce the number of APs necessary to satisfy the same amount of aggregated demand, thus increasing the competition among the APs that participate to the auction. As the curves G.2 and G.3 show, when we increase the number of MCs, the number of winners selected by the *greedy use* and *greedy max use* solutions decreases down to 20%, dropping from 100% in the *low-density* scenario to 60% in the *high-density* scenario. We can finally observe that the AP density contributes to reduce the number of winners selected by all greedy schemes, with a gain that ranges from 10% to 25% when the number of available APs in

each sector varies from 10 to 15.

6.3 Analysis of Traffic Load

The second set of simulated scenarios, whose results are depicted in Figures 6, 7 and 8, aims at evaluating the effect of the traffic load heterogeneity on the performance of our proposed schemes. To this end, within each of the 21 sectors, we randomly place 6 MCs. The MC traffic demand, d_j , is distributed uniformly in the range $[x - 0.4, x + 0.4]$ Mbps, with $x = \{6, 7, 8\}$ Mbps: this corresponds to three traffic load scenarios that we denote, respectively, with *underload*, *peak-load* and *overload* (the aggregated bandwidth in every sector is equal to 36, 42 and 48 Mbps). Furthermore, we vary the number of APs in the $[5, 15]$ range, generating 100 device deployments and 100 different demand distributions for each network scenario. Due to the high computational time, we solve the optimal auction only for a subset of network instances. Since the results confirm the trends and gaps obtained in the scenarios described in Section 6.2, we only show the curves obtained using the greedy auctions. In particular, the *greedy use* algorithm (G.2) achieves the best performance among the greedy algorithms in terms




Fig. 6: Performance metrics measured in the *underload* scenario (average traffic demand equal to 6 Mbps).




Fig. 7: Performance metrics measured in the *peak-load* scenario (average traffic demand equal to 7 Mbps).




Fig. 8: Performance metrics measured in the *overload* scenario (average traffic demand equal to 8 Mbps).

of cost paid by the mobile operator. The offloading cost obtained using the greedy schemes (see Figures 6(a), 7(a) and 8(a)) keeps increasing as long as the APs density achieves a knee point, where the leasing cost slightly decreases due to the higher competition. We further observe that the variability of the traffic demand increases the gap between the *greedy use* auction (G.2) and the other approximated solutions, due to the suboptimal allocations implemented by the different schemes.

Figures 6(b), 7(b), and 8(b) show the fraction of MCs that are offloaded from the mobile network on the leased WiFi APs. While all greedy auctions need approximately 7 APs per sector to offload all data connections in the underload scenario, every increase of 1 Mbps of the average data traffic demand (i.e., peak-load and overload scenarios) requires two additional APs to satisfy all MCs.

The curves illustrated in Figures 6(c), 7(c), and 8(c), which represent the fraction of winners selected by the proposed auctions, follow a trend similar to the overall cost paid by the operator. Specifically, the increasing part of the curves obtained using the greedy algorithms is due to the low AP density. In such cases, some APs are too far from the MCs and cannot be used to serve the MCs data traffic, forcing the algorithm to select all APs

that permit to offload the greatest portion of traffic. The greedy algorithms keep selecting APs as long as their density achieves the point where the WiFi capacity is enough to serve all MCs and additional APs are useless. Nonetheless, we underline that, while there exists an optimal number of APs for offloading the whole data traffic, the higher is the number of APs, the higher is the competition, thus increasing the economic efficiency of the mechanism.

7 CONCLUSION

This paper proposed a new trading marketplace where mobile operators can rent the bandwidth of Internet connections made available by third party WiFi Access Points to offload the data traffic of their mobile customers.

The offloading problem was formulated as a combinatorial auction, and an innovative payment rule was designed to guarantee both individual rationality and truthfulness for realistic scenarios in which only part of the data traffic can be offloaded.

In order to solve efficiently (i.e., in polynomial time) the offloading problem for large-scale network scenarios,

we also proposed a greedy algorithm, with two alternative versions of the allocation phase, that preserves the *truthfulness* property.

Numerical results demonstrate that the proposed schemes well capture the economical and networking essence of the problem, thus representing a promising solution to implement a trading marketplace for next-generation access networks composed of heterogeneous systems.

ACKNOWLEDGMENT

This work was partially supported by French ANR in the framework of the Green-Dyspan project.

REFERENCES

- [1] M. Buddhikot, G. Chandranmenon, S. Han, Y.W. Lee, S. Miller, and L. Salgarelli. Integration of 802.11 and Third-Generation Wireless Data Networks. *IEEE INFOCOM*, pages 503–512, 2003.
- [2] A. Balasubramanian, R. Mahajan, and A. Venkataramani. Augmenting Mobile 3G using WiFi. *ACM MobiSys*, pages 209–222, 2010.
- [3] B.D. Higgins, A. Reda, T. Alperovich, J. Flinn, T.J. Giuli, B. Noble, and D. Watson. Intentional Networking: Opportunistic Exploitation of Mobile Network Diversity. *ACM MobiCom*, pages 73–84, 2010.
- [4] X. Zhuo, W. Gao, G. Cao, and S. Hua. An Incentive Framework for Cellular Traffic Offloading. *IEEE Trans. on Mobile Computing*, pages 206–215, 2013.
- [5] X. Zhuo, W. Gao, G. Cao, and Y. Dai. Win-Coupon: An Incentive Framework for 3G Traffic Offloading. *IEEE ICNP*, 2011.
- [6] W. Dong, S. Rallapalli, R. Jana, L. Qiu, K.K. Ramakrishnan, L. Razoumov, Y. Zhangn, and T.W. Cho. iDEAL: Incentivized Dynamic Cellular Offloading via Auctions. *IEEE INFOCOM*, pages 755–763, 2013.
- [7] Y. Chen, J. Zhang, Q. Zhang, and J. Jia. A Reverse Auction Framework for Access Permission Transaction to Promote Hybrid Access in Femtocell Network. *IEEE INFOCOM*, pages 2761–2765, 2012.
- [8] S. Paris, F. Martignon, I. Filippini, and A. Capone. A Truthful Auction for Access Point Selection in Heterogeneous Mobile Networks. *IEEE ICC*, 2012.
- [9] Y. Chen, J. Zhang, Q. Zhang, and J. Jia. A Reverse Auction Framework for Access Permission Transaction to Promote Hybrid Access in Femtocell Network. *IEEE INFOCOM*, pages 2761–2765, 2012.
- [10] X. Zhou, S. Gandhi, S. Suri, and H. Zheng. eBay in the Sky: Strategy-proof Wireless Spectrum Auctions. *ACM MobiCom*, pages 2–13, 2008.
- [11] X. Zhou and H. Zheng. TRUST: A General Framework for Truthful Double Spectrum Auctions. *IEEE INFOCOM*, pages 999–1007, 2009.
- [12] G.S. Kasbekar and S. Sarkar. Spectrum Auction Framework for Access Allocation in Cognitive Radio Networks. *ACM MobiHoc*, 2009.
- [13] S. Sengupta and M. Chatterjee. An Economic Framework for Dynamic Spectrum Access and Service Pricing. *IEEE/ACM Trans. on Networking*, 17(4):1200–1213, 2009.
- [14] J. Jia, Q. Zhang, Q. Zhang, and M. Liu. Revenue Generation for Truthful Spectrum Auction in Dynamic Spectrum Access. *ACM MobiCom*, pages 3–12, 2009.
- [15] Q. Wang, B. Ye, S. Lu, and S. Guo. A Truthful QoS-Aware Spectrum Auction with Spatial Reuse for Large-Scale Networks. *IEEE Trans. on Parallel and Distributed Systems*, PP(99):1–1, 2013.
- [16] 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. on Mobile Computing*, 7(1):19–33, 2008.
- [17] J. Jaramillo and R. Srikant. DARWIN: Distributed and Adaptive Reputation Mechanism for Wireless Ad hoc Networks. *ACM MobiCom*, 2007.

- [18] Y. Wu, S. Tang, P. Xu, and X.Y. Li. Dealing With Selfishness and Moral Hazard in Non-Cooperative Wireless Networks. *IEEE Trans. on Mobile Computing*, 9(3):420–434, 2009.
- [19] X. Su, G. Peng, and S. Chan. Multi-Path Routing and Forwarding in Non-Cooperative Wireless Networks. *IEEE Trans. on Parallel and Distributed Systems*, PP(99):1–1, 2013.
- [20] R. Jain and J. Walrand. An Efficient Mechanism for Network Bandwidth Auction. *IEEE NOMS*, pages 227–234, 2008.
- [21] F. Fu and U.C. Kozat. Wireless Network Virtualization as a Sequential Auction Game. *IEEE INFOCOM*, pages 1945–1953, 2010.
- [22] D. Malone, K. Duffy, and D. Leith. Modeling the 802.11 Distributed Coordination Function in Nonsaturated Heterogeneous Conditions. *IEEE/ACM Trans. on Networking*, 15(1):159–172, 2007.
- [23] N. Nisan, T. Roughgarden, E. Tardos, and V.V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [24] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords. *American Economic Review*, 97(1):242–259, 2007.
- [25] A. Ambrosy, G. Auer, Blume O., M. Caretti, and et al. D2.2 - Definition and Parameterization of Reference System and Scenarios. *INFSO-ICT-247733 EARTH*, 2010.
- [26] R. Jain. *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling*. Wiley - Interscience, 1991.

Stefano Paris is an Assistant Professor at the Department of Computer Science (LIPADE) at Paris Descartes University. He received his M.S. degree in Computer Engineering from University of Bergamo in 2007, and the Ph.D. in Information Engineering from Politecnico di Milano in 2011. His main research interests include topics related to optimization, performance evaluation and security of wireless networks. He is a member of the IEEE Communication Society.

Fabio Martignon received the M.S. and the Ph.D. degrees in telecommunication engineering from the Politecnico di Milano in October 2001 and May 2005, respectively. He has been associate professor at University of Bergamo, and he is now Full Professor in LRI (Laboratory for Computer Science) at Paris-Sud University, and member of Institut Universitaire de France. His current research activities include cognitive radio networks, content-centric networks, network planning and game theory applications to networking problems.

Ilario Filippini received the B.S. and M.S. degrees in telecommunication engineering and Ph.D. degree in information engineering from the Politecnico di Milano, Milan, Italy, in 2003, 2005, and 2009, respectively. From February to August 2008, he worked as a Visiting Researcher with the Department of Electrical and Computer Engineering, The Ohio State University, Columbus. He is currently an Assistant Professor with the Electronics and Information Department, Politecnico di Milano. His research activities include networking and optimization issues, in particular wireless multihop network planning, optimization, and protocol design.

Lin Chen received his B.E. degree in Radio Engineering from Southeast University, China in 2002 and the Engineer Diploma from Telecom ParisTech, Paris in 2005. He also holds a M.S. degree of Networking from the University of Paris 6. He currently works as Assistant Professor in the department of computer science of the University of Paris-Sud XI. His main research interests include modeling and control for wireless networks, security and cooperation enforcement in wireless networks and game theory.