

Auction-Based Resource Allocation for Sharing Cloudlets in Mobile Cloud Computing

A-Long Jin, Wei Song, *Senior Member, IEEE*, and Weihua Zhuang, *Fellow, IEEE*

Abstract—Driven by pervasive mobile devices and ubiquitous wireless communication networks, mobile cloud computing emerges as an appealing paradigm to accommodate demands for running power-hungry or computation-intensive applications over resource-constrained mobile devices. Cloudlets that move available resources closer to the network edge offer a promising architecture to support real-time applications, such as online gaming and speech recognition. To stimulate service provisioning by cloudlets, it is essential to design an incentive mechanism that charges mobile devices and rewards cloudlets. Although auction has been considered as a promising form for incentive, it is challenging to design an auction mechanism that holds certain desirable properties for the cloudlet scenario. In this paper, we propose an incentive-compatible auction mechanism (ICAM) for the resource trading between mobile devices as service users (buyers) and cloudlets as service providers (sellers). ICAM can effectively allocate cloudlets to satisfy the service demands of mobile devices and determine the pricing. Both theoretical analysis and numerical results show that ICAM guarantees desired properties with respect to individual rationality, budget balance, truthfulness (incentive compatibility) for both buyers and sellers, and computational efficiency.

Index Terms—Mobile cloud computing, cloudlet, truthful double auction, incentive design.

I. INTRODUCTION

The past decade has witnessed an explosive growth of wireless communication networks, where a variety of smart mobile devices offer a plethora of applications. Nonetheless, the energy and resource constraints of mobile devices still limit the support of power-hungry or computation-intensive applications, even with the rapid progress of hardware technologies. In the mean time, cloud computing is achieving great success in empowering end users with rich experience by leveraging resource virtualization and sharing. Extending the success of cloud computing to the mobile domain, mobile cloud computing (MCC) creates a new appealing paradigm [1,2]. There have been many popular cloud-based mobile applications, e.g., deployed in Apple iCloud [3] and Amazon Silk [4]. By offloading power-hungry or computation-intensive tasks to clouds, MCC is expected to relax the local constraints of mobile devices in storage, energy, and networking [5].

Three typical MCC architectures are reviewed in [6], including the traditional centralized cloud [7], the recently emerged cloudlet [8], and the peer-based ad hoc mobile cloud [9].

This research was supported in part by Natural Sciences and Engineering Research Council (NSERC) of Canada. A-Long Jin and Wei Song are with the Faculty of Computer Science, University of New Brunswick, Fredericton, NB, Canada (emails: {along.jin, wsong}@unb.ca). Weihua Zhuang is with the Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON, Canada (email: wzhuang@uwaterloo.ca).




Fig. 1. Typical MCC architectures: Centralized cloud and cloudlets.

The ad hoc mobile cloud is a user-centric model which pools together a crowd of neighboring mobile devices for resource sharing. The other two larger-scale cloud architectures are illustrated in Fig.1. The centralized cloud hosts shared resources in remote data centers and acts as an agent between the original content providers and mobile devices. To access resources at the data centers, mobile devices often need to go through the backbone network. The long latency incurred to access the centralized cloud can be intolerable for interactive applications such as online gaming and speech recognition. Even with the acceleration of network speeds, the network resources will remain insufficient in a fairly long period to accommodate the soaring traffic demands. On the other hand, a cloudlet [8] is a trusted, resource-rich, Internet-connected computer or a cluster of computers, which can be utilized by mobile devices via a high-speed wireless local area network (WLAN). With such geographically distributed cloudlets, the close physical proximity can enable smoother interactions with the low one-hop communication latency. Thus, cloudlets offer an economical solution which can take advantage of content distribution close to the network edge.

We are particularly interested in the cloudlet architecture, which can complement the centralized cloud and accommodate communication-intensive or delay-sensitive applications. If the service demands of mobile devices can be satisfied by high-profile cloudlets in their vicinity, the mobile devices do not need to request resources from the centralized cloud, thereby balancing the workload and reducing the access latency. To achieve the potential benefits of cloudlets, many

practical issues need to be addressed, such as the deployment of cloudlets, reliability, business model, pricing and incentive design. For geographically distributed cloudlets, due to their spatial locations and distinct capabilities or hosted resources, mobile devices have different preferences over the cloudlets. For example, a mobile device may favor a cloudlet that provides a high level of quality of service (QoS), and associate a high valuation with the cloudlet. On the other hand, the cloudlets need to be motivated to share their resources, e.g., through gaining monetary values paid by the mobile devices for using the services. As seen, there exists a trade between the mobile devices requesting the services and the cloudlets providing such services.

Auction is a popular trading form that can efficiently distribute resources of sellers to buyers in a market at competitive prices. Auction theory [10] is a well-researched field in economics and has been applied to other domains, e.g., radio resource management in wireless communication systems [11]. An auction mechanism is expected to hold certain desirable properties, such as *individual rationality*, *budget balance*, and *system efficiency* [10]. Besides, *incentive compatibility* or *truthfulness* is another important aspect of auction design. Truthfulness is essential to resist market manipulation and ensure auction fairness. An auction mechanism is incentive-compatible or truthful if revealing the private valuation truthfully is always the dominant strategy for each participant to receive an optimal utility, no matter what strategies other participants are taking. In this work, it is critical to ensure truthfulness in the auction mechanism so that the allocation of cloudlets' resources is not interfered by untruthful behaviours that aim to boost a participant's own benefit. As such, the cloudlets' resources can be allocated to the mobile devices in need and satisfy their service demands to the utmost extent.

There are many existing auction mechanisms that satisfy some of the above properties but are not directly applicable to the cloudlet scenario. For example, the multi-round auctions studied in [12]–[14] are not suitable due to the high communication and computation overhead. We are particularly interested in double auction, in which *buyers* and *sellers* submit their bids and asks, respectively, to an *auctioneer* as an intermediate agent who hosts and directs the auction process, e.g., deciding the auction *commodity* allocation and the clearing price and payment. Well-known examples of double auction include McAfee double auction [15] and Vickrey-based auction [16]. Considering only homogeneous commodities, McAfee double auction can achieve three desired properties, i.e., individual rationality, budget balance, and truthfulness. The Vickrey-based auction proposed in [16] can be budget-balanced and efficient but not truthful simultaneously, according to [17]. A truthful double auction mechanism (TASC) is proposed in [18] for cooperative communications with heterogeneous trading commodities, i.e., services of relay nodes. Although TASC addresses a scenario similar to MCC with cloudlets, TASC cannot solve the resource sharing problem for cloudlets without losing some desired properties. Due to unique features of cloudlets, TASC cannot guarantee truthfulness for buyers even though it is still individual rational, budget-balanced, and truthful for sellers.

In this paper, we focus on designing an incentive-compatible auction mechanism (ICAM) to stimulate cloudlets to serve nearby mobile devices, so that the abundant resources of cloudlets are efficiently utilized to reduce the access latency of mobile devices for improved interactivity and balance the workload from the centralized cloud. ICAM ensures truthfulness for both buyers and sellers. In addition, ICAM is individually rational, budget-balanced, and *computationally efficient*. The computational efficiency requires that the auction outcome (allocation of commodities, and clearing price and payment) be computed in polynomial time. We provide rigorous analysis proving that the above desirable properties hold with ICAM. Numerical results verify that these properties are achieved with a reasonable *system efficiency*, which is another crucial property of auctions. Here, a higher system efficiency implies more mobile devices are successfully assigned to satisfactory cloudlets instead of resorting to the centralized cloud.

In the remainder of this paper, we first review related works in Section II. Section III provides the system model, problem formulation, and an example demonstrating the design challenges. Then, we introduce ICAM in Section IV and analyze its properties in Section V. Numerical results are presented in Section VI, followed by conclusions in Section VII.

II. RELATED WORKS

In this section, we give a brief review on related works in two groups, i.e., the incentive mechanisms specifically for mobile cloud computing in the networking literature, and more general auction mechanisms in the economics literature.

As a promising paradigm, mobile cloud computing has attracted considerable research attention and efforts. There have been a number of studies addressing various aspects of MCC, such as virtual machine migration [19], service enhancement with MCC [5], and emerging applications with MCC [20,21]. However, the research on incentive design for MCC is limited. In [22], cloud resources are categorized into several groups (e.g., processing, storage, and communications). Then, the resource allocation problem is formulated as a combinatorial auction with substitutable and complementary commodities. This combinatorial auction mechanism is not applicable for the cloudlet architecture since its key problem is the allocation of M resources of G groups in one MCC service provider to N users. In contrast, our system model with cloudlets focuses on distinct valuations of cloudlets to mobile users. Different from [22], we also consider computational efficiency and budget balance, which are critical to an auction mechanism.

Although auction theory has been widely studied in the economics literature, the existing auction mechanisms cannot be directly applied to the cloudlet scenario, since they fail to fully satisfy the required properties stated in Section I. One of the most well-known auction mechanisms is the truthful Vickrey-Clarke-Groves (VCG) auction [23]–[25]. In [16], Parkes *et al.* propose a Vickrey-based double auction, which achieves individual rationality and budget balance. The assignment between buyers and sellers is determined to maximize social welfare (system efficiency), while the player's utility equals the incremental contribution to the overall system, i.e., the

difference between the social welfare with and without a player's participation. However, the well-known result in [17] reveals that it is impossible to design a truthful, efficient, and budget-balanced double auction, even putting individual rationality aside. Therefore, the Vickrey-based double auction in [16] is only fairly efficient and fairly truthful.

In [15], McAfee double auction aims at a scenario with homogeneous commodities, where buyers have no preference over auction items. Each buyer (b_i) submits only one bid (D_i) and each seller (s_j) submits one ask (A_j). The auctioneer sorts the bids in a non-increasing order and the asks in a non-decreasing order to have $D_{i_l} \geq D_{i_{l+1}}$ and $A_{j_l} \leq A_{j_{l+1}}$, respectively. Let $D_{i_{m+1}}$ denote the smallest possible bid, and $A_{j_{m+1}}$ the largest possible ask. Then, the auctioneer determines the winning buyers $\{b_{i_1}, \dots, b_{i_k}\}$ and the winning sellers $\{s_{i_1}, \dots, s_{i_k}\}$, where k is the largest number such that $D_{i_k} \geq A_{j_k}$ and $D_{i_{k+1}} < A_{j_{k+1}}$. The auctioneer charges each winning buyer a clearing price P^b and rewards each winning seller a clearing payment P^s . Here, $P^b = P^s = P_o$ and $P_o = \frac{1}{2}(D_{i_{k+1}} + A_{j_{k+1}})$, if $A_{j_k} \leq P_o \leq D_{i_k}$; otherwise, $P^b = D_{i_k}$ and $P^s = A_{j_k}$. Although McAfee double auction can achieve three desirable economic properties, including individual rationality, budget balance, and truthfulness, the homogeneity of commodities in McAfee double auction limits its application to the cloudlet scenario of MCC, where the mobile devices as service buyers have preferences over the cloudlets as resource sellers.

In [18], Yang *et al.* propose a truthful double auction mechanism (TASC) for cooperative communications with heterogeneous trading commodities, i.e., services of relay nodes. In TASC double auction, there are two stages, namely, *Assignment* and *Winner-Determination & Pricing*. In the assignment stage, the auctioneer applies an assignment algorithm to determine the winning buyer candidates (source nodes), the winning seller candidates (relay nodes), and the mapping between these buyers and sellers. Depending on the design objective, the auctioneer can choose a different assignment algorithm. For example, the optimal relay assignment algorithm [26] can maximize the minimum QoS among all buyers; the maximum weighted matching algorithm [27] can maximize the overall QoS; and the maximum matching algorithm can maximize the number of successful trades (final matchings). In the winner-determination & pricing stage, TASC double auction tightly integrates the winner determination and the pricing operation. Based on the return of the assignment stage, the auctioneer applies McAfee double auction [15] to determine the winning buyers, the winning sellers, and the corresponding clearing price and payment. When TASC double auction is used in the cloudlet scenario, it can satisfy individual rationality, budget balance, and truthfulness for the sellers. However, we illustrate using an example in Section III-D that a buyer can bid untruthfully to improve its utility. Hence, TASC double auction cannot be applied to the MCC scenario of this study.

III. SYSTEM MODEL AND PROBLEM FORMULATION

A. Resource Allocation for Cloudlets

As depicted in Fig. 1, the cloudlets offer resource pools closer to the network edge. The close proximity of cloudlets can be

exploited to reduce the access overhead of mobile devices in energy consumption and communication latency. The cloudlets may value differently to mobile users depending on various factors [28], such as computation capability, communication cost, and wireless link performance (e.g., throughput, latency, and link variation). Such *valuation* of a mobile user toward a cloudlet varies with the channel conditions and is also associated with the service requirement. For instance, when a mobile user offloads a computation-intensive task, it values high a cloudlet with rich computing resources of memory and CPU capacity. In contrast, a mobile user with a real-time task prefers a cloudlet with a low communication latency, which requires large network bandwidth, high power level, and short physical distance.

On the other hand, the cloudlet can be paid for sharing resources as compensation for its computation and communication *cost*. Clearly, the trading between the cloudlets and the mobile devices should meet certain requirements to benefit both parties. The cloudlets need to be incentivized to provide the resources, and the demands of the mobile users should be satisfied. In particular, a cloudlet cannot be paid less than its cost, while the allocated resources of the cloudlet must fulfill a mobile user's service request. The more mobile users served by the cloudlets, the higher the resource utilization for cloudlets. To maximize the resource utilization, the incentive mechanism should properly assign the matching between the cloudlet's resources and the mobile users' demands.

B. Auction Model

Focusing on the MCC scenario with cloudlets in Fig. 1, we consider a discrete-time system so that in each time period mobile users submit their bids to a central controller, depending on the traffic arrivals and service demands. The asks of the cloudlets offering services in the vicinity are also collected. Then, we can design an incentive-compatible mechanism to allocate the resources of m cloudlets among n mobile devices. Similar to the single-round multi-item double auction model in [18], the mobile devices are *buyers* in this auction, while cloudlets are *sellers*. A control center closest to the participants can serve as the *auctioneer* to reduce the communication cost and delay. Considering the potential gain of serving mobile users by nearby cloudlets, the communication overhead in the auction procedure is affordable and worthwhile.

Considering a sealed-bid auction, each buyer (resp. seller) can submit its bid (resp. ask) privately to the auctioneer so that everyone has no information of other bids or asks.

- For each buyer $b_i \in \mathcal{B}$, $\mathcal{B} = \{b_1, b_2, \dots, b_n\}$, its bid vector is denoted by $\mathbf{D}_i = (D_i^1, D_i^2, \dots, D_i^m)$, where D_i^j is the bid for seller $s_j \in \mathcal{S}$, $\mathcal{S} = \{s_1, s_2, \dots, s_m\}$. The bid matrix consisting of the bid vectors of all buyers is defined as $\mathbf{D} = (\mathbf{D}_1; \mathbf{D}_2; \dots; \mathbf{D}_n)$.
- For all sellers in \mathcal{S} , the ask vector is denoted by $\mathbf{A} = (A_1, A_2, \dots, A_m)$, where A_j is the ask of seller $s_j \in \mathcal{S}$.

As seen, the asks of sellers do not differentiate among buyers since the sellers only aim at collecting payments for using their resources. In contrast, the bids of buyers differ with respect to sellers, as mobile devices have preferences over

cloudlets that vary in available resources and access overhead in energy consumption or communication latency.

Given $\mathcal{B}, \mathcal{S}, \mathbf{D}$ and \mathbf{A} , the auctioneer decides the winning buyer set $\mathcal{B}_w \subseteq \mathcal{B}$, the winning seller set $\mathcal{S}_w \subseteq \mathcal{S}$, the mapping between \mathcal{B}_w and \mathcal{S}_w , i.e., $\sigma : \{j : s_j \in \mathcal{S}_w\} \rightarrow \{i : b_i \in \mathcal{B}_w\}$, the price P_i^b that the winning buyer $b_i \in \mathcal{B}_w$ is charged, and the payment P_j^s that the winning seller $s_j \in \mathcal{S}_w$ is rewarded¹. To highlight the utilities for the particular matching between b_i and s_j , we also use P_{ij}^b and P_{ij}^s in certain cases to denote the price and payment, respectively.

In addition to the price and payment, the utilities of the buyers and sellers further depend on the valuations of the buyers toward the acquired services and the costs for providing such services by the sellers. Let V_i^j be the valuation to buyer b_i for having the service from seller s_j , and C_j be the cost to seller s_j for providing the service. The valuation vector of buyer b_i is denoted by $\mathbf{V}_i = (V_i^1, V_i^2, \dots, V_i^m)$. Given a buyer-seller mapping, $i = \sigma(j)$, the *utility* of buyer b_i and that of seller s_j are respectively defined as follows:

$$U_i^b = \begin{cases} V_i^j - P_i^b, & \text{if } b_i \in \mathcal{B}_w \\ 0, & \text{otherwise} \end{cases}$$

$$U_j^s = \begin{cases} P_j^s - C_j, & \text{if } s_j \in \mathcal{S}_w \\ 0, & \text{otherwise.} \end{cases}$$

Here, utility $U_i^b > 0$ means that mobile user b_i as a buyer is assigned to a cloudlet with a valuation greater than the charged price. Thus, U_i^b indicates the satisfaction level of the mobile user on the allocated cloudlet. On the other hand, utility U_j^s of cloudlet s_j as a seller represents the surplus of the received payment over its cost. In other words, U_j^s characterizes the profit of a cloudlet for sharing its resources. We also use U_{ij}^b and U_{ij}^s when necessary to emphasize that the utilities are with respect to the matching between buyer b_i and seller s_j .

Some important notations are summarized in Table I.

C. Desirable Properties and Design Objective

The auction model in Section III-B is represented by $\Psi = (\mathcal{B}, \mathcal{S}, \mathbf{D}, \mathbf{A})$. Accordingly, the auctioneer should follow an auction mechanism to determine the set of winning buyers \mathcal{B}_w , the set of winning sellers \mathcal{S}_w , the mapping σ between \mathcal{B}_w and \mathcal{S}_w , the set of clearing price \mathcal{P}_w^b charged to the winning buyers, and the set of clearing payment \mathcal{P}_w^s rewarded to the winning sellers. An effective auction mechanism should satisfy four desirable properties in the following.

- *Individual Rationality*: No winning buyer is charged more than its bid and no winning seller is rewarded less than its ask. With respect to the auction model Ψ , this means that for every winning matching between $b_i \in \mathcal{B}_w$ and $s_j \in \mathcal{S}_w$, we have $P_i^b \leq D_i^j$ and $P_j^s \geq A_j$.
- *Budget Balance*: The total price that the auctioneer charges all winning buyers is not less than the total payment that the auctioneer rewards all winning sellers,

¹To distinguish the price charged to *buyers* and the payment rewarded to *sellers*, we use b and s in the normal form as the superscript, respectively. The same naming routine is also applied to the utilities of buyers and sellers.

TABLE I
IMPORTANT NOTATIONS.

Symbol	Definition
b_i	Buyer (mobile device)
b_{ij}	Buyer b_i with positive valuation toward seller s_j
s_j	Seller (cloudlet)
n	Total number of buyers
m	Total number of sellers
\mathcal{B}	Set of buyers (mobile devices)
\mathcal{B}'	Extended set of buyers with positive valuations
\mathcal{S}	Set of sellers (cloudlets)
\mathbb{B}	Sorted buyer list of \mathcal{B}' in a descending order of positive valuations
\mathbb{S}	Sorted seller list of \mathcal{S} in an ascending order of asks
\mathcal{B}_c	Set of winning buyer candidates ($\mathcal{B}_c \subseteq \mathcal{B}$)
\mathcal{S}_c	Set of winning seller candidates ($\mathcal{S}_c \subseteq \mathcal{S}$)
\mathcal{B}_a	Set of winning buyers before elimination ($\mathcal{B}_a \subseteq \mathcal{B}_c$)
\mathcal{S}_a	Set of winning sellers before elimination ($\mathcal{S}_a = \mathcal{S}_c$)
\mathcal{B}_w	Set of winning buyers ($\mathcal{B}_w = \mathcal{B}_a$)
\mathcal{S}_w	Set of winning sellers ($\mathcal{S}_w \subseteq \mathcal{S}_a$)
$\hat{\sigma}(\cdot)$	Mapping function from the indices of \mathcal{S}_a to \mathcal{B}_a
$\sigma(\cdot)$	Mapping function from the indices of \mathcal{S}_w to \mathcal{B}_w
D_i^j	Bid of buyer b_i on seller s_j
\mathbf{D}_i	Bid vector of buyer b_i
\mathbf{D}	Bid matrix of all buyers
A_j	Ask of seller s_j
\mathbf{A}	Ask vector of all sellers
\mathbf{A}_{-j}	Ask vector of all sellers except s_j
V_i^j	Valuation of buyer b_i on service from seller s_j
\mathbf{V}_i	Valuation vector of buyer b_i
C_j	Cost of seller s_j for providing service
P_i^b	Price charged to buyer b_i
P_j^s	Payment rewarded to seller s_j
P_{ij}^b	Price charged to buyer b_i for service of seller s_j
P_{ij}^s	Payment rewarded to seller s_j with assigned buyer b_i
U_i^b	Utility of buyer b_i
U_j^s	Utility of seller s_j
U_{ij}^b	Utility of buyer b_i with assigned seller s_j
U_{ij}^s	Utility of seller s_j with assigned buyer b_i

so that there is no deficit for the auctioneer. That is, $\sum_{b_i \in \mathcal{B}_w} P_i^b \geq \sum_{s_j \in \mathcal{S}_w} P_j^s$.

- *Truthfulness or Incentive Compatibility*: We need to first give the definition of a *weakly dominant* strategy in the following. Based on this definition, we can further express the property of truthfulness or incentive compatibility.

Definition 1. For player i , strategy a_i *weakly dominates* strategy a'_i if the utilities satisfy $u_i(a_i, a_{-i}) \geq u_i(a'_i, a_{-i})$ for all partial action profiles a_{-i} of the other players except i . For player i , strategy a_i is weakly dominant if

TABLE II
AN ILLUSTRATIVE EXAMPLE.

(a) Bid matrix of 5 buyers.

	s_1	s_2	s_3	s_4	s_5	s_6	s_7
b_1	6	0	0	0	5	10	0
b_2	4	0	0	3	0	0	8
b_3	0	0	6	0	0	9	0
b_4	0	10	0	0	0	0	7
b_5	0	2	7	9	0	0	0

(b) Ask vector of 7 sellers.

Seller	s_1	s_2	s_3	s_4	s_5	s_6	s_7
Ask	3	2	5	6	4	1	7

it weakly dominates all other strategies of player i .

Then, an auction mechanism is truthful or incentive-compatible if playing (bidding or asking) truthfully is a weakly dominant strategy for each player (buyer or seller). In other words, no buyer can improve its utility by submitting a bid different from its true valuation, and no seller can improve its utility by submitting an ask different from its true cost. Specifically, it implies the following for our auction model: $\forall b_i \in \mathcal{B}$, U_i^b is maximized when the bidding $\mathbf{D}_i = \mathbf{V}_i$; and $\forall s_j \in \mathcal{S}$, U_j^s is maximized when the asking $A_j = C_j$.

- *Computational Efficiency*: The auction outcome, which includes the winning sets of buyers and sellers, their mapping, and the clearing price and payment, is tractable with a polynomial time complexity.

D. Technical Challenges

As discussed in Section II, the existing auction mechanisms cannot satisfy the preceding desirable properties when directly applied to the MCC scenario with heterogeneous cloudlets as auction commodities. The pioneer work in [18] provides a promising solution. Unfortunately, the following example shows that TASC double auction (i.e., the enhanced version in [18]) cannot guarantee truthfulness of buyers, although there is no problem with individual rationality, budget balance, and truthfulness of sellers.

To illustrate that buyers can gain higher utilities by bidding untruthfully, we consider a bid matrix of 5 buyers with true valuations in Table II(a), and the ask vector of 7 sellers with true costs in Table II(b). Suppose that the auctioneer uses the maximum weighted matching algorithm in the assignment stage to maximize the overall QoS. According to the assignment algorithm, the winning buyer candidates, the winning seller candidates and the mapping between them are shown in Fig. 2. Then, following the TASC strategy for winner-determination & pricing, we have the set of winning buyers $\mathcal{B}_w = \{b_1, b_4\}$, the set of winning sellers $\mathcal{S}_w = \{s_6, s_2\}$, the clearing price $\mathcal{P}_w^b = \{8\}$, and the clearing payment $\mathcal{P}_w^s = \{6\}$. The utility of b_3 is 0 since it is not within the winner set \mathcal{B}_w .

If buyer b_3 bids untruthfully by increasing its bid D_3^b from its true valuation 9 to $9 + \delta$ ($\delta > 1$), the new assignment result




Fig. 2. Assignment result with truthful bidding and asking: A bipartite graph of winning buyer candidates and winning seller candidates and their mapping.




Fig. 3. Different assignment result with untruthful bidding of buyer b_3 , which increases its bid D_3^b from its true valuation 9 to $9 + \delta$ ($\delta > 1$).

is shown in Fig. 3. The set of winning buyers becomes $\mathcal{B}_w = \{b_3, b_4\}$, while the set of winning sellers is still $\mathcal{S}_w = \{s_6, s_2\}$. The clearing price and payment remain unchanged according to TASC, i.e., $\mathcal{P}_w^b = \{8\}$ and $\mathcal{P}_w^s = \{6\}$. The new utility of b_3 becomes $9 - 8 = 1$. As seen, b_3 can improve its utility from 0 to 1 by bidding untruthfully. Hence, we cannot apply TASC double auction to the cloudlet scenario. In Section IV, we propose a new double auction mechanism, ICAM, which can guarantee truthfulness of both sellers and buyers, while holding the other desirable properties.

IV. PROPOSED AUCTION MECHANISM FOR CLOUDLETS

As discussed in Section II, the well-known Vickrey-based double auction [16] cannot simultaneously achieve truthfulness in addition to individual rationality and budget balance, while McAfee double auction [15] cannot be directly applied to the scenario with heterogeneous commodities. TASC double auction overcomes the limitation of McAfee double auction and accommodates heterogeneity. When TASC is applied to resource sharing with cloudlets, we have seen from the example in Section III-D that TASC is subject to the manipulation of untruthful buyers in the assignment stage.

In this section, we propose ICAM to resolve this problem. First of all, we change the sequence of the assignment stage and the winner-determination & pricing stage. In ICAM, the auctioneer first identifies the winning candidates. Then, each winning seller candidate is assigned to one winning buyer candidate. Also, the clearing price charged to each buyer candidate and the clearing payment rewarded to the seller candidate are determined accordingly. More importantly, ICAM can keep potentially multiple sellers for a single buyer until a new last stage. In the end, the new stage of winner elimination can guarantee that a winning buyer is assigned to only one winning seller.

Next, we give the detailed algorithms of ICAM, followed by a walk-through example. The properties of ICAM are analyzed

Algorithm 1 ICAM($\mathcal{B}, \mathcal{S}, \mathbf{D}, \mathbf{A}$).**Input:** $\mathcal{B}, \mathcal{S}, \mathbf{D}, \mathbf{A}$ **Output:** $\mathcal{B}_w, \mathcal{S}_w, \sigma, \mathcal{P}_w^b, \mathcal{P}_w^s$

- 1: $(\mathcal{B}_c, \mathcal{S}_c, D_{p_\varphi}^{q_\varphi}, A_{j_\phi}) \leftarrow \text{ICAM-WCD}(\mathcal{B}, \mathcal{S}, \mathbf{D}, \mathbf{A})$;
- 2: $(\mathcal{B}_a, \mathcal{S}_a, \hat{\sigma}, \mathcal{P}_a^b, \mathcal{P}_a^s) \leftarrow \text{ICAM-A\&P}(\mathcal{B}_c, \mathcal{S}_c, D_{p_\varphi}^{q_\varphi}, A_{j_\phi}, \mathbf{D})$;
- 3: $(\mathcal{B}_w, \mathcal{S}_w, \sigma, \mathcal{P}_w^b, \mathcal{P}_w^s) \leftarrow \text{ICAM-WE}(\mathcal{B}_a, \mathcal{S}_a, \hat{\sigma}, \mathcal{P}_a^b, \mathcal{P}_a^s, \mathbf{D})$;
- 4: **return** $(\mathcal{B}_w, \mathcal{S}_w, \sigma, \mathcal{P}_w^b, \mathcal{P}_w^s)$;

Algorithm 2 ICAM-WCD($\mathcal{B}, \mathcal{S}, \mathbf{D}, \mathbf{A}$).**Input:** $\mathcal{B}, \mathcal{S}, \mathbf{D}, \mathbf{A}$ **Output:** $\mathcal{B}_c, \mathcal{S}_c, D_{p_\varphi}^{q_\varphi}, A_{j_\phi}$

- 1: $\mathcal{B}_c \leftarrow \emptyset, \mathcal{S}_c \leftarrow \emptyset$;
- 2: Construct a set $\mathcal{B}' = \{b_{pq} : D_p^q > 0, b_p \in \mathcal{B}\}$ according to \mathbf{D} ;
- 3: Sort all buyers in \mathcal{B}' to obtain an ordered list $\mathbb{B} = \langle b_{p_1 q_1}, b_{p_2 q_2}, \dots, b_{p_x q_x} \rangle$ such that $D_{p_1}^{q_1} \geq D_{p_2}^{q_2} \geq \dots \geq D_{p_x}^{q_x}$;
- 4: Sort all sellers in \mathcal{S} to obtain an ordered list $\mathbb{S} = \langle s_{j_1}, s_{j_2}, \dots, s_{j_m} \rangle$ such that $A_{j_1} \leq A_{j_2} \leq \dots \leq A_{j_m}$;
- 5: Find the median ask A_{j_ϕ} of \mathbb{S} , where $\phi = \lceil \frac{m+1}{2} \rceil$;
- 6: Find the smallest φ , such that $D_{p_{\varphi+1}}^{q_{\varphi+1}} < A_{j_\phi}$;
- 7: $\mathcal{B}_c \leftarrow \mathbb{B}_\varphi$, where \mathbb{B}_φ is the sublist with first φ buyers in \mathbb{B} ;
- 8: **for** $b_{pq} \in \mathcal{B}_c$ **do**
- 9: **if** $A_q \geq A_{j_\phi}$ **then**
- 10: $\mathcal{B}_c \leftarrow \mathcal{B}_c \setminus \{b_{pq}\}$;
- 11: **else**
- 12: **if** $s_q \notin \mathcal{S}_c$ **then**
- 13: $\mathcal{S}_c \leftarrow \mathcal{S}_c \cup \{s_q\}$;
- 14: **end if**
- 15: **end if**
- 16: **end for**
- 17: **return** $(\mathcal{B}_c, \mathcal{S}_c, D_{p_\varphi}^{q_\varphi}, A_{j_\phi})$;

in Section V.

A. Details of ICAM

Following the preceding design rationale, we propose ICAM in Alg. 1, which includes three stages, namely, *winning candidate determination*, *assignment & pricing*, and *winner elimination*.

In the stage of winning candidate determination, Alg. 2 is used by the auctioneer to shortlist the buyer and seller candidates. Alg. 2 first constructs a new buyer set \mathcal{B}' from the original buyer set \mathcal{B} . Specifically, buyer $b_i \in \mathcal{B}$ becomes b_{ij} in \mathcal{B}' if $D_i^j > 0$. That is, a buyer can appear for a number of times with respect to the sellers for which the buyer has positive valuations. Then, \mathcal{B}' is ranked to \mathbb{B} in an ascending order of all positive bids (valuations), denoted by $\mathbb{D}' = \langle D_{p_1}^{q_1}, \dots, D_{p_x}^{q_x} \rangle$, where $x = |\mathcal{B}'|$. Seller set \mathcal{S} is sorted to \mathbb{S} in a descending order of \mathbf{A} , where the ordered list of \mathbf{A} is denoted by $\mathbb{A} = \langle A_{j_1}, \dots, A_{j_m} \rangle$. The ask of the median seller in \mathbb{S} , denoted by A_{j_ϕ} , where $\phi = \lceil \frac{m+1}{2} \rceil$, is used to find the smallest φ such that $D_{p_{\varphi+1}}^{q_{\varphi+1}} < A_{j_\phi}$. The two selected thresholds, $D_{p_\varphi}^{q_\varphi}$ and A_{j_ϕ} , are used to select winning candidates. Buyer b_{pq} is a winning buyer candidate in \mathcal{B}_c if $D_p^q \geq D_{p_\varphi}^{q_\varphi}$ and $A_q < A_{j_\phi}$. Seller s_q is a winning seller candidate in \mathcal{S}_c if $A_q < A_{j_\phi}$ and at least one winning buyer candidate bids for s_q with a positive bid. It is worth mentioning that ϕ is not limited to the median. As discussed later in Section VI-A, system efficiency varies with ϕ though other properties of ICAM stay the same with different ϕ . The reason for setting ϕ to the median in Alg. 2 is to balance the size of \mathcal{S}_c or \mathcal{B}_c , so as to achieve a reasonable system efficiency.

Algorithm 3 ICAM-A&P($\mathcal{B}_c, \mathcal{S}_c, D_{p_\varphi}^{q_\varphi}, A_{j_\phi}, \mathbf{D}$).**Input:** $\mathcal{B}_c, \mathcal{S}_c, D_{p_\varphi}^{q_\varphi}, A_{j_\phi}, \mathbf{D}$ **Output:** $\mathcal{B}_a, \mathcal{S}_a, \hat{\sigma}, \mathcal{P}_a^b, \mathcal{P}_a^s$

- 1: $\mathcal{B}_a \leftarrow \emptyset, \mathcal{S}_a \leftarrow \mathcal{S}_c, \mathcal{P}_a^b \leftarrow \emptyset, \mathcal{P}_a^s \leftarrow \emptyset$;
- 2: **for** $s_j \in \mathcal{S}_a$ **do**
- 3: $\mathcal{P}_j^s = A_{j_\phi}, \mathcal{P}_a^s \leftarrow \mathcal{P}_a^s \cup \{\mathcal{P}_j^s\}$;
- 4: $\mathcal{B}^j = \{b_{ij} : b_{ij} \in \mathcal{B}_c\}$;
- 5: **if** $|\mathcal{B}^j| = 1$ **then**
- 6: $\mathcal{B}_a \leftarrow \mathcal{B}_a \cup \{b_{ij}\}, \hat{\sigma}(j) = i$;
- 7: $\mathcal{P}_{ij}^b = D_{p_\varphi}^{q_\varphi}, \mathcal{P}_a^b \leftarrow \mathcal{P}_a^b \cup \{\mathcal{P}_{ij}^b\}$;
- 8: **else**
- 9: Sort \mathcal{B}^j to an ordered list \mathbb{B}^j such that $D_{i(1)}^j \geq D_{i(2)}^j \geq \dots \geq D_{p_\varphi}^{q_\varphi}$;
- 10: **if** the first t ($t \geq 2$) bids of \mathbb{B}^j are the same **then**
- 11: Randomly select a b_{ij} from the first t buyers of \mathbb{B}^j ;
- 12: **else**
- 13: Select the first buyer b_{ij} of \mathbb{B}^j with the highest bid;
- 14: **end if**
- 15: $\mathcal{B}_a \leftarrow \mathcal{B}_a \cup \{b_{ij}\}, \hat{\sigma}(j) = i$;
- 16: $\mathcal{P}_{ij}^b = D_{i(2)}^j, \mathcal{P}_a^b \leftarrow \mathcal{P}_a^b \cup \{\mathcal{P}_{ij}^b\}$;
- 17: **end if**
- 18: **end for**
- 19: **return** $(\mathcal{B}_a, \mathcal{S}_a, \hat{\sigma}, \mathcal{P}_a^b, \mathcal{P}_a^s)$;

Algorithm 4 ICAM-WE($\mathcal{B}_a, \mathcal{S}_a, \hat{\sigma}, \mathcal{P}_a^b, \mathcal{P}_a^s, \mathbf{D}$).**Input:** $\mathcal{B}_a, \mathcal{S}_a, \hat{\sigma}, \mathcal{P}_a^b, \mathcal{P}_a^s, \mathbf{D}$ **Output:** $\mathcal{B}_w, \mathcal{S}_w, \sigma, \mathcal{P}_w^b, \mathcal{P}_w^s$

- 1: $\mathcal{B}_w \leftarrow \mathcal{B}_a, \mathcal{S}_w \leftarrow \mathcal{S}_a, \sigma \leftarrow \hat{\sigma}, \mathcal{P}_w^b \leftarrow \mathcal{P}_a^b, \mathcal{P}_w^s \leftarrow \mathcal{P}_a^s$;
- 2: **for** any two buyers $b_{\sigma(\alpha)\alpha}, b_{\sigma(\beta)\beta} \in \mathcal{B}_w, \alpha \neq \beta$ **do**
- 3: **if** $\sigma(\alpha) = \sigma(\beta)$ **then**
- 4: $U_{\sigma(j)j}^b = D_{\sigma(j)}^j - P_{\sigma(j)j}^b, j = \{\alpha, \beta\}$;
- 5: **if** $U_{\sigma(\alpha)\alpha}^b = U_{\sigma(\beta)\beta}^b$ **then**
- 6: $j' \leftarrow$ randomly selected from $\{\alpha, \beta\}$;
- 7: **else**
- 8: $j' \leftarrow \arg \min_{j \in \{\alpha, \beta\}} \{U_{\sigma(j)j}^b\}$;
- 9: **end if**
- 10: $\mathcal{B}_w \leftarrow \mathcal{B}_w \setminus \{b_{\sigma(j')j'}\}, \mathcal{S}_w \leftarrow \mathcal{S}_w \setminus \{s_{j'}\}$;
- 11: $\mathcal{P}_w^b \leftarrow \mathcal{P}_w^b \setminus \{P_{\sigma(j')j'}^b\}, \mathcal{P}_w^s \leftarrow \mathcal{P}_w^s \setminus \{P_{j'}^s\}, \sigma(j') = \emptyset$;
- 12: **end if**
- 13: **end for**
- 14: **return** $(\mathcal{B}_w, \mathcal{S}_w, \sigma, \mathcal{P}_w^b, \mathcal{P}_w^s)$;

In the assignment & pricing stage, we tightly couple winner determination and pricing to prevent possible untruthful manipulation. As given in Alg. 3, the auctioneer first determines the winning buyer for each winning seller candidate s_j . If only one buyer candidate b_{ij} bids for s_j , then b_{ij} is added into the winning buyer set \mathcal{B}_a and charged a clearing price $D_{p_\varphi}^{q_\varphi}$. If more than one buyer candidate bids for s_j , the buyer candidate with the highest bid is added into the winning buyer set and charged a price of the second highest bid. Seller s_j is paid the median ask, A_{j_ϕ} . When there is a tie among the highest bids of buyer candidates, the auctioneer randomly selects a winning buyer from the candidates. For example, supposing $D_{p_\varphi}^{q_\varphi} = 3$ and $D_\alpha^j = D_\beta^j = 10$, the winning buyer for s_j can be either $b_{\alpha j}$ or $b_{\beta j}$, each with a 50% chance. If the next lower bid for s_j by $b_{\gamma j}$ is $D_\gamma^j = 5$, the winning buyer is charged 10 instead of 5, because the first two highest bids in the sorted list are both 10, i.e., $D_{i(1)}^j = D_{i(2)}^j = 10$. This is essential to avoid untruthful actions of buyers.

In the last stage, if a buyer in the original buyer set \mathcal{B} wins two or more sellers in \mathcal{S}_a , the auctioneer, depending on system

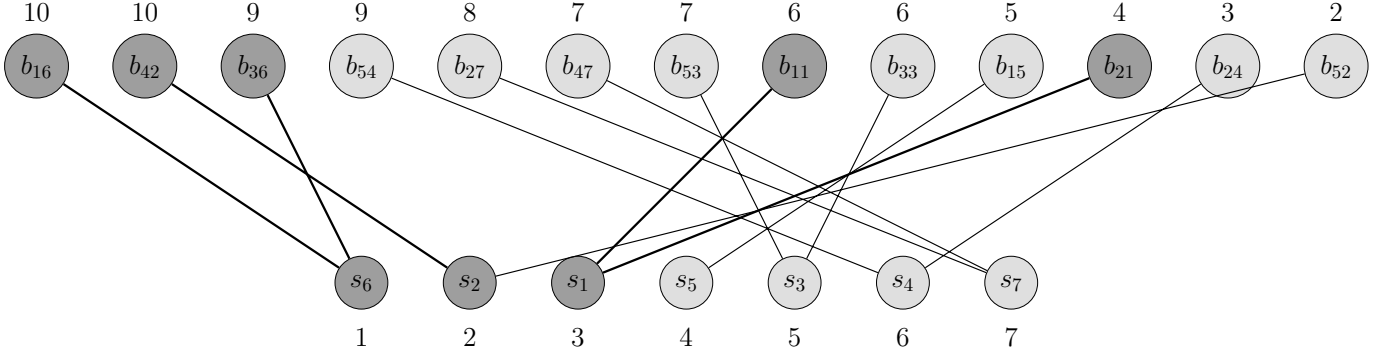


Fig. 4. Initial bipartite graph showing the ordered lists of the new buyer set and the seller set.

requirements, can choose only one seller for such a buyer using Alg. 4. For example, if both $b_{i\alpha}$ and $b_{i\beta}$ belong to \mathcal{B}_a , it means that b_i in the original buyer set \mathcal{B} wins two sellers, s_α and s_β . The auctioneer can select only one seller so that the corresponding buyer achieves the highest utility. Likewise, when there is a tie in terms of the achievable utilities, one seller is randomly selected. At the end of the winner elimination stage, every buyer $b_{\sigma(j)j} \in \mathcal{B}_w$ has a one-to-one mapping with only one winning seller $s_j \in \mathcal{S}_w$.

For the auction model in [18], each seller can be assigned to at most one buyer, while one buyer needs at most one seller. It is worth noting that our proposed auction mechanism can be easily modified to retain multiple winning sellers for one buyer by skipping the above elimination stage. Then, one buyer (a mobile device) is allowed to acquire resources from multiple sellers (cloudlets), e.g., for different resources of processing, storage, or networking.

B. A Walk-Through Example

Considering the bid matrix in Table II(a) and the ask vector in Table II(b), the following shows how ICAM works for the auctioneer to derive the auction outcome.

Winning candidate determination according to Alg. 2:

- Construct the new buyer set from original set \mathcal{B} : $\mathcal{B}' = \{b_{11}, b_{15}, b_{16}, b_{21}, b_{24}, b_{27}, b_{33}, b_{36}, b_{42}, b_{47}, b_{52}, b_{53}, b_{54}\}$;
- Sort buyers in \mathcal{B}' in a descending order to obtain: $\mathbb{B} = \{b_{16}, b_{42}, b_{36}, b_{54}, b_{27}, b_{47}, b_{53}, b_{11}, b_{33}, b_{15}, b_{21}, b_{24}, b_{52}\}$;
- Sort sellers in \mathcal{S} in an ascending order to obtain: $\mathbb{S} = \{s_6, s_2, s_1, s_5, s_3, s_4, s_7\}$;
- Based on \mathbb{B} and \mathbb{S} , construct an initial bipartite graph between \mathcal{B}' and \mathcal{S} as shown in Fig. 4;
- Decide two thresholds: $A_{j_\phi} = A_5 = 4$, $D_{p_\phi}^{q_\phi} = D_2^1 = 4$;
- Determine the set of winning buyer candidates: $\mathcal{B}_c = \{b_{16}, b_{42}, b_{36}, b_{11}, b_{21}\}$;
- Determine the set of winning seller candidates: $\mathcal{S}_c = \{s_6, s_2, s_1\}$.

According to the output of Alg. 2, a bipartite graph between \mathcal{B}_c and \mathcal{S}_c is constructed as shown in Fig. 5. Then, Alg. 3 is run to identify the winning buyers and sellers.

Assignment & pricing according to Alg. 3:

- The set of winning buyers: $\mathcal{B}_a = \{b_{16}, b_{42}, b_{11}\}$;
- The set of winning sellers: $\mathcal{S}_a = \{s_6, s_2, s_1\}$;




Fig. 5. Bipartite graph between winning candidates \mathcal{B}_c and \mathcal{S}_c .

- The assignment (mapping) between winning buyers and sellers (\mathcal{B}_a and \mathcal{S}_a): $\hat{\sigma}(\cdot) = \{\hat{\sigma}(6) = 1, \hat{\sigma}(2) = 4, \hat{\sigma}(1) = 1\}$;
- The clearing price charged to winning buyers: $\mathcal{P}_a^b = \{P_{16}^b = D_3^6 = 9, P_{42}^b = D_{p_\phi}^{q_\phi} = 4, P_{11}^b = D_2^1 = 4\}$;
- The clearing payment rewarded to winning sellers: $\mathcal{P}_a^s = \{P_6^s = P_2^s = P_1^s = A_{j_\phi} = 4\}$.

Alg. 3 returns the mapping between \mathcal{B}_a and \mathcal{S}_a , $\hat{\sigma}(\cdot) = \{\hat{\sigma}(6) = \hat{\sigma}(1) = 1, \hat{\sigma}(2) = 4\}$, which means that $b_1 \in \mathcal{B}$ wins two sellers (s_6 and s_1). If the auctioneer requires to keep only one winning seller for buyer b_1 , Alg. 4 is run to remove redundant sellers.

Winner elimination according to Alg. 4:

- Compute the utilities of buyer b_1 with respect to seller s_6 and s_1 , respectively: $U_{16}^b = D_1^6 - P_{16}^b = 10 - 9 = 1$, $U_{11}^b = D_1^1 - P_{11}^b = 6 - 4 = 2$;
- Since $U_{16}^b < U_{11}^b$, seller s_6 is eliminated so that a higher utility is provided to buyer b_1 by seller s_1 . Then, the set of winning buyers is obtained as: $\mathcal{B}_w = \{b_{16}, b_{42}, b_{11}\} \setminus \{b_{16}\} = \{b_{42}, b_{11}\} = \{b_4, b_1\}$;
- Update the set of winning sellers: $\mathcal{S}_w = \{s_6, s_2, s_1\} \setminus \{s_6\} = \{s_2, s_1\}$;
- Update the clearing price charged to winning buyers: $\mathcal{P}_w^b = \{P_{16}^b, P_{42}^b, P_{11}^b\} \setminus \{P_{16}^b\} = \{P_{16}^b = 9, P_{42}^b = 4\} = \{P_1^b = 9, P_4^b = 4\}$;
- Update the clearing payment rewarded to winning sellers: $\mathcal{P}_w^s = \{P_6^s, P_2^s, P_1^s\} \setminus \{P_6^s\} = \{P_2^s = 4, P_1^s = 4\}$;
- Update the final one-to-one mapping between winning buyers and sellers (\mathcal{B}_w and \mathcal{S}_w): $\sigma(\cdot) = \{\sigma(2) = 4, \sigma(1) = 1\}$.

Recall that one motivation for the proposed ICAM is to solve the problem illustrated by the example in Section III-D.




Fig. 6. Bipartite graph between winning candidates \mathcal{B}_c and \mathcal{S}_c , when b_3 deviates its bid D_3^6 from its true valuation 9 to $9 + \delta$ ($\delta > 1$).

Next, we briefly show how ICAM prevents such an untruthful buyer bidding, and leave the formal proof of truthfulness and other properties in Section V. Suppose similarly that buyer b_3 increases its bid D_3^6 from its truthful valuation 9 to $9 + \delta$, where $\delta > 1$. The winning candidates obtained from Alg. 2 will change to the bipartite graph in Fig. 6. As a result, b_3 needs to pay a price 10 to win seller s_6 , and its utility is $9 - 10 = -1 < 0$. Therefore, bidding truthfully should be the dominant strategy of b_3 .

V. ANALYSIS OF DESIRABLE PROPERTIES

In this section, we analyze the proposed auction mechanism ICAM with respect to the four desirable properties discussed in Section III-C. The following theorems prove that all four properties hold with ICAM. We leave the proof for truthfulness in the end, which requires complex and rigorous reasoning.

Theorem 1. *ICAM is computationally efficient.*

Proof. In the winning candidate determination stage, Alg. 2 involves at most nm buyers in the new buyer set \mathcal{B}' . Sorting the buyers in \mathcal{B}' takes $O(nm \log(nm))$ time, while sorting the sellers in \mathcal{S} takes $O(m \log m)$ time. In Line 7, there are at most $n \lceil \frac{m+1}{2} \rceil$ buyers in the winning candidate set \mathcal{B}_c . Hence, the for-loop (Line 8 – Line 16) has a time complexity $O(n \lceil \frac{m+1}{2} \rceil \cdot \lceil \frac{m+1}{2} \rceil) = O(nm^2)$. Note that the for-loop can also be improved to have a time complexity of $O(nm)$ with a space complexity of $O(m)$. Since we focus on the worst-case time complexity, Alg. 2 takes $O(nm \cdot (m + \log n))$ time.

In the assignment & pricing stage, Alg. 3 processes at most $n \lceil \frac{m+1}{2} \rceil$ buyers in \mathcal{B}_c and $\lceil \frac{m+1}{2} \rceil$ sellers in \mathcal{S}_c . Line 4 determines subset $\mathcal{B}^j \subseteq \mathcal{B}_c$ for the buyers with positive valuations toward seller $s_j \in \mathcal{S}_a$, which takes $O(n \lceil \frac{m+1}{2} \rceil) = O(nm)$ time. Taking advantage of the ordered list \mathbb{B} , we can sort \mathcal{B}^j without cost to obtain \mathbb{B}^j . Since there are at most n buyers in \mathbb{B}^j , it takes $O(n)$ time to determine the winning buyer for s_j . Hence, the for-loop (Line 2 – Line 18) costs $O(nm \cdot \lceil \frac{m+1}{2} \rceil) = O(nm^2)$. Thus, Alg. 3 takes $O(nm^2)$ time.

In the winner elimination stage, we know that set \mathcal{B}_a before elimination has a size $|\mathcal{B}_a| = |\mathcal{S}_a| \leq \lceil \frac{m+1}{2} \rceil$. Thus, the for-loop (Line 2 – Line 13) takes $O(\frac{|\mathcal{B}_a|(|\mathcal{B}_a|-1)}{2}) = O(m^2)$ time. Thus, Alg. 4 takes $O(m^2)$ time.

Therefore, the overall time complexity of ICAM in Alg. 1 is $O(nm \cdot (m + \log n))$. In other words, ICAM converges to the final assignment and pricing result in a polynomial time with respect to n and m . \square

Theorem 2. *ICAM is individually rational.*

Proof. For each winning seller $s_j \in \mathcal{S}_w \subseteq \mathcal{S}_c$, the payment rewarded to seller s_j is $P_j^s = A_{j_\phi} > A_j$ according to ICAM. Thus, the winning sellers satisfy individual rationality.

Next, consider the winning buyer set \mathcal{B}_a produced by Alg. 3. For each winning buyer $b_{ij} \in \mathcal{B}_a \subseteq \mathcal{B}_c$, there are two cases.

- In the first case, buyer b_{ij} wins s_j without competition, which means that b_{ij} is the only buyer in \mathcal{B}_c that bids for s_j . In this situation, we know that $P_{ij}^b = D_{p_\phi}^{q_\phi} \leq D_i^j$.
- In the second case, buyer b_{ij} wins s_j with competition, which means that more than one buyer in \mathcal{B}_c bids for s_j , and D_i^j is the highest. In this situation, b_{ij} is charged the second highest bid in \mathbb{B}^j . Obviously, $P_{ij}^b \leq D_i^j$.

Therefore, individual rationality also holds for the winning buyer set \mathcal{B}_a determined by Alg. 3.

If a winning buyer, $b_i \in \mathcal{B}$, wins multiple sellers, e.g., $b_{i\alpha} \in \mathcal{B}_a$ and $b_{i\beta} \in \mathcal{B}_a$, running Alg. 4 can eliminate redundant sellers and keep only one best seller for each winning buyer. Among all the sellers that buyer b_i wins, Alg. 4 simply keeps the seller, s_j (e.g., s_α or s_β), which gives b_i the highest utility. It is evident that this procedure does not change the charging price P_{ij}^b to the winning buyers. Thus, the buyers in \mathcal{B}_w after the winner elimination still satisfy individual rationality.

In summary, ICAM is individually rational. \square

Theorem 3. *ICAM is budget-balanced.*

Proof. After the winner elimination stage, every winning buyer $b_i \in \mathcal{B}_w$ has only one winning seller $s_j \in \mathcal{S}_w$. Considering this one-to-one mapping between \mathcal{B}_w and \mathcal{S}_w , we have $|\mathcal{B}_w| = |\mathcal{S}_w|$. For each matching $\sigma(j) = i$ between winning buyer b_i and assigned winning seller s_j , it is true that

$$P_{\sigma(j)}^b \geq D_{p_\phi}^{q_\phi} \geq A_{j_\phi} = P_j^s.$$

Then, it can be easily shown that

$$\sum_{b_i \in \mathcal{B}_w} P_i^b - \sum_{s_j \in \mathcal{S}_w} P_j^s = \sum_{s_j \in \mathcal{S}_w} (P_{\sigma(j)}^b - P_j^s) \geq 0$$

which completes the proof. \square

Before drawing a conclusion on truthfulness of ICAM, we first derive Lemma 1 and Lemma 2 in the following.

Lemma 1. *ICAM is truthful for sellers.*

Proof. Lemma 1 can be proved by Propositions 1-3, which are presented and proved in Appendix A (in the supplementary file). Let $k = \phi = \lceil \frac{m+1}{2} \rceil$, $\mathcal{S}_{l_1} = \mathcal{S}_{<k} \setminus \mathcal{S}_c$ and $\mathcal{S}_{l_2} = \mathcal{S}_a \setminus \mathcal{S}_w$. Then, according to Propositions 1-3, telling truth ($A_j = C_j$) is a weakly dominant strategy for each seller $s_j \in \mathcal{S}$ in ICAM, which completes the proof of Lemma 1. \square

Lemma 2. *ICAM is truthful for buyers.*

Proof. Similar to the proof of Lemma 1, we provide Propositions 4-7 in Appendix B (in the supplementary file), which lay the basis for Lemma 2. Following the notations therein and letting $D_\phi = A_\phi = A_{j_\phi}$, $\mathcal{B}_{l_1} = \mathcal{B}_\phi \setminus \mathcal{B}_c$ and $\mathcal{B}_{l_2} = \mathcal{B}_c \setminus \mathcal{B}_a$, we can draw a logical conclusion that telling truth is a weakly

dominant strategy for each buyer $b_i \in \mathcal{B}$ in ICAM. This proves Lemma 2. \square

Theorem 4. *ICAM is truthful (incentive-compatible).*

Proof. Lemma 1 and Lemma 2 together prove that ICAM is truthful (incentive-compatible). \square

According to Theorems 1-4, we can draw the final conclusion in Theorem 5.

Theorem 5. *ICAM is computationally efficient, individually rational, budget-balanced and truthful (incentive-compatible).*

As discussed in Section IV, ICAM also works when the elimination stage is skipped so that a mobile device (buyer) can acquire services from more than one cloudlet (seller). The four desired properties in Theorem 5 still hold.

VI. NUMERICAL RESULTS


In this section, we present numerical results to validate the properties of ICAM analyzed in Section V. In addition, we evaluate the performance of ICAM in terms of system efficiency. As seen in the proof in Section V and the appendices (in the supplementary file), ICAM guarantees individual rationality, budget balance, truthfulness for both buyers and sellers, and computational efficiency. The proof does not set any presumption on the bids of buyers or the asks of sellers. Thus, the conclusions are valid for any possible data sets of the bids and asks.

Because there are no existing statistics on service demands of mobile users or resource costs of real cloudlets [28], for generality, we randomly generate the bids of buyers and the asks of sellers according to uniform distributions within $(0, V_{max}]$ and within $(0, 1]$, respectively. Intuitively, V_{max} will affect the auction outcome, and even the parameter ϕ that is used to determine the auction thresholds and winning candidates. In the following, we first illustrate the impact of ϕ and its variation with V_{max} , so that the numerical results thereafter will be mainly based on fixed ϕ and V_{max} . At the end of this section, we also relax the setting of uniformly distributed bids and asks to investigate the sensitivity of system efficiency on such statistics.


A. Impact of Parameter ϕ

In the winning candidate determination stage of ICAM, buyer and seller candidates are selected based on the ϕ -th ask of the ascending ordered list of all sellers' asks, $A_{j\phi}$. The candidate sets, \mathcal{B}_c and \mathcal{S}_c , are determined in Alg. 2. Intuitively, a larger value of ϕ results in a smaller set for \mathcal{B}_c . On the other hand, \mathcal{S}_c can be too small if ϕ is too small. The candidate sets directly affect the auction outcome. Fig. 7 shows the impact of ϕ on the performance of ICAM with different values of V_{max} , with 100 buyers and 100 sellers.

Fig. 7(a) shows the number of successful trades (N_{ST}) versus ϕ . The variation therein is due to the opposite effects of ϕ on the sizes of \mathcal{B}_c and \mathcal{S}_c . When ϕ is too small or too large, the size of \mathcal{S}_c or \mathcal{B}_c is too small, respectively. As a result, the number of successful trades (i.e., matchings between winning



(a) Number of successful trades vs. ϕ .



(b) Total valuation of winning buyers vs. ϕ .

Fig. 7. Impact of ϕ on the performance of ICAM.

buyers and sellers) is small. In addition, examining the peak points of the curves with different V_{max} , we find out that the optimal value of ϕ that attains the highest N_{ST} increases with a larger V_{max} . The highest N_{ST} also increases accordingly. This is because a larger ϕ can be selected when V_{max} increases so as to enlarge \mathcal{B}_c and \mathcal{S}_c . Thus, the highest N_{ST} increases with a larger V_{max} .

Fig. 7(b) shows the impact of ϕ on the total valuation of winning buyers, with different V_{max} . It is clear that Fig. 7(b) exhibits a similar trend as Fig. 7(a). The reason is that the total valuation is proportional to the number of successful trades.

Given the observations in Fig. 7, we can see that ϕ should be adapted to V_{max} for the best performance. In the following experiments, since V_{max} is fixed to 1, we set $\phi = \lceil \frac{m+1}{2} \rceil$ based on the observations in Fig. 7. The relative difference between the number of buyers (n) and the number of sellers (m) may also affect the selection of ϕ . In fact, the performance of ICAM can be improved when the optimal ϕ is selected according to different values of n and m .

B. Computational Efficiency

To confirm our analysis on time complexity in Theorem 1, we obtain the computation time of ICAM with different settings in Table III. For each setting, we randomly generate 1000 instances and average the results. All the tests run on a

TABLE III
COMPUTATION TIME.

$n = 100$	m	50	100	150	200	250	300
	Time (ms)	0.4	0.7	1.0	1.2	1.5	1.8
$m = 100$	n	50	100	150	200	250	300
	Time (ms)	0.4	0.7	0.9	1.1	1.2	1.5

Windows PC with 3.16 GHz Intel® Core™2 Duo processor and 4 GB memory. As seen, ICAM is subject to a polynomial computation time with respect to n and m , which are the numbers of buyers and sellers, respectively.

C. Individual Rationality

To validate Theorem 2 regarding individual rationality of ICAM, we present the bids and prices of winning buyers in Fig. 8(a), and the payments and asks of winning sellers in Fig. 8(b). Clearly, each winning buyer is charged a price not higher than its bid, while each winning seller receives a payment not less than its ask from the auctioneer. Therefore, ICAM is individually rational. The results demonstrate that the winning mobile users and cloudlets that are successfully matched gain positive utilities, i.e., benefit from using or providing the demanded resources. The winning cloudlets receive sufficient compensations as incentive to share their resources. On the other hand, the winning mobile users are allocated the demanded resources and pay no more than their valuations toward these resources. Thus, the mobile users are also stimulated to request resources from the cloudlets instead of the centralized cloud.


D. Budget Balance

Theorem 3 proves that ICAM is budget-balanced, which means that the total price charged to the winning buyers is not less than the total payment rewarded to the winning sellers. Fig. 9 shows the total price and the total payment with different settings. Here, we fix the number of buyers to 100, and vary the number of sellers from 50 to 150 with an increment of 10. As seen, the total price from the winning buyers is always greater than the total payment to the winning sellers. Therefore, the auctioneer conducts the auction without a deficit, and is thus inclined to assist in the resource allocation for cloudlets.


E. Truthfulness

To verify truthfulness of ICAM, we randomly pick two buyers and two sellers to examine how their utilities change when they bid or ask different values. The results are depicted in Fig. 10.

Fig. 10(a) shows a case that buyer b_{i_α} wins the seller s_{j_α} and gains utility $U_{i_\alpha}^b = 0.2131$ when it bids truthfully with $D_{i_\alpha}^{j_\alpha} = V_{i_\alpha}^{j_\alpha} = 0.7444$. It can be seen that buyer b_{i_α} cannot improve its utility no matter what other bids it takes. Fig. 10(b) shows a different scenario that buyer b_{i_β} does not win the seller s_{j_β} when it bids truthfully with $D_{i_\beta}^{j_\beta} = V_{i_\beta}^{j_\beta} = 0.6841$. Thus, b_{i_β} achieves zero utility ($U_{i_\beta}^b = 0$) without having the service.



(a) Bids & prices of winning buyers.



(b) Payments & asks of winning sellers.

Fig. 8. Individual rationality of ICAM.

Fig. 10(b) shows the utility cannot be greater than zero even when b_{i_β} bids untruthfully.

Fig. 10(c) shows an example with winning seller s_{j_η} that asks truthfully with $A_{j_\eta} = C_{j_\eta} = 0.1706$ and achieves utility $U_{j_\eta}^s = 0.3546$. As seen, the utility with a truthful ask is the highest among all possible asks. Fig. 10(d) shows that seller s_{j_ξ} loses when asking truthfully with $A_{j_\xi} = C_{j_\xi} = 0.8564$ and thus obtains zero utility ($U_{j_\xi}^s = 0$). For all other asks, the achievable utility is either zero or negative, but cannot be more than zero.

In summary, ICAM guarantees truthfulness for both buyers and sellers since the utility cannot be improved by bidding or asking untruthfully. Thus, ICAM can be freed from the interference of untruthful participants (cloudlets and mobile users) that try to strategize over others.

F. System Efficiency

Both the theoretical proof in Section V and the numerical results show that ICAM is computationally efficient, individually rational, budget-balanced and truthful. System efficiency is another important metric for an auction mechanism. Unfortunately, it has been shown in [17] that a double auction is impossible to achieve truthfulness, budget balance, and system efficiency simultaneously. Depending on the system requirement, we can evaluate system efficiency in terms of




Fig. 9. Budget balance of ICAM.




Fig. 10. Truthfulness of buyers and sellers with ICAM.

the number of successful trades or the total valuation of winning buyers. Usually, the total valuation is proportional to the number of successful trades, which has been observed in Fig. 7. Hence, we focus on the number of successful trades (N_{ST}) in the following to evaluate system efficiency.

Fig. 11 compares the number of successful trades among three different auction mechanisms when the bids and asks are uniformly distributed within $(0, 1]$. In the optimal strategy, the auctioneer maximizes N_{ST} with complete information in matching the buyers and sellers. As seen, N_{ST} increases with the number of sellers, which is intuitive since more sellers can better satisfy the diverse demands of buyers. ICAM achieves around 50% of the system efficiency of the optimal strategy, where the loss is mainly due to the cost of maintaining truthfulness. Moreover, ICAM outperforms TASC in system efficiency in addition to completing the truthfulness guarantee of TASC. The higher system efficiency of ICAM is attributed to the fact that ICAM involves much more winning buyer candidates in the assignment & pricing stage, and removes very few winning players in the winner elimination stage. Therefore, ICAM can achieve all the desirable properties while maintaining a reasonable system efficiency.

To further investigate the sensitivity of the auction mech-




Fig. 11. System efficiency with uniformly distributed bids and asks.

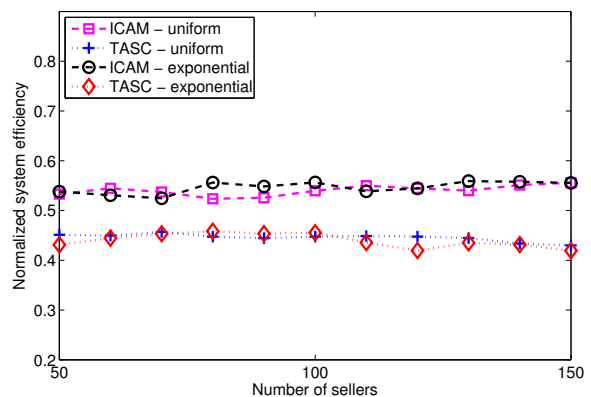


Fig. 12. Normalized system efficiency with uniformly or exponentially distributed bids and asks.

anisms to the random distributions of bids and asks, Fig. 12 shows the system efficiency normalized with respect to that of the optimal strategy. Here, we show the normalized system efficiency when the bids and asks are uniformly or exponentially distributed with the same mean. As seen, both ICAM and TASC are insensitive to the statistics of bids and asks, while ICAM maintains a stable improvement over TASC. Also, it is noticed that the normalized system efficiency only fluctuates slightly with the number of sellers. This is because the system efficiency of ICAM, TASC, and the optimal strategy all increases with the number of sellers.

VII. CONCLUSIONS AND FUTURE WORKS


In this paper, we focus on a promising paradigm of MCC with cloudlets that provide resources to nearby mobile devices. Due to spatial locations of cloudlets and their distinct capabilities or hosted resources, the cloudlets offer heterogeneous valuations toward mobile devices. The mobile users can acquire services from different cloudlets to maximize their utilities. To improve resource utilization of cloudlets, we have proposed a double auction mechanism ICAM, which coordinates the resource trading between mobile devices as service users (buyers) and cloudlets as service providers (sellers). ICAM can effectively allocate the cloudlets' resources among mobile users to satisfy their service demands, while maintaining

the desirable properties, including computational efficiency, individual rationality, budget balance, and truthfulness for both buyers and sellers. We have provided rigorous proof on these properties of ICAM and confirmed the analysis with extensive simulation results.

There are still many open issues and this research can be extended in the following aspects. As shown in Fig. 12, the system efficiency of ICAM and TASC is around 50% of that of the optimal strategy. Thus, more efforts are needed to further improve the system efficiency of the auction mechanisms while maintaining other desirable properties. In addition, more sophisticated features can be incorporated into the system model. For example, we can distinguish the types of services available at each cloudlet. Some cloudlets may only provide storage service, while other cloudlets may provide computing and networking services. Thus, the mobile user needs to assign its service requests to the “compatible” cloudlets. With such a system model, it can be much more challenging to design an auction mechanism with the desirable properties.

REFERENCES


- [1] H. Dinh, C. Lee, D. Niyato, and P. Wang, “A survey of mobile cloud computing: Architecture, applications, and approaches,” *Wireless Communications and Mobile Computing*, vol. 13, pp. 1587–1611, 2013.
- [2] N. Fernando, S. W. Loke, and W. Rahayu, “Mobile cloud computing: A survey,” *Future Generation Computer Systems*, vol. 29, pp. 84–106, 2013.
- [3] Apple Inc., “Apple iCloud,” <https://www.icloud.com/>.
- [4] Amazon.com, “Amazon Silk,” <http://amazonsilk.wordpress.com/>.
- [5] S. Abolfazli, Z. Sanaei, E. Ahmed, A. Gani, and R. Buyya, “Cloud-based augmentation for mobile devices: Motivation, taxonomies, and open challenges,” *IEEE Communications Surveys & Tutorials*, vol. 16, no. 1, pp. 337–368, 2014.
- [6] F. Liu, P. Shu, H. Jin, L. Ding, J. Yu, D. Niu, and B. Li, “Gearing resource-poor mobile devices with powerful clouds: Architectures, challenges, and applications,” *IEEE Wireless Communications*, vol. 20, no. 3, pp. 14–22, Jun. 2013.
- [7] Z. Liu, Y. Feng, and B. Li, “Socialize spontaneously with mobile applications,” in *Proc. IEEE INFOCOM*, 2012, pp. 1942–1950.
- [8] M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies, “The case for VM-based cloudlets in mobile computing,” *IEEE Pervasive Computing*, vol. 8, no. 4, pp. 14–23, 2009.
- [9] D. Huang, T. Xing, and H. Wu, “Mobile cloud computing service models: A user-centric approach,” *IEEE Communications Surveys & Tutorials*, vol. 27, no. 5, pp. 6–11, 2013.
- [10] V. Krishna, *Auction Theory*, 2nd ed. Academic Press, Aug. 2009.
- [11] Y. Zhang, C. Lee, D. Niyato, and P. Wang, “Auction approaches for resource allocation in wireless systems: A survey,” *IEEE Communications Surveys & Tutorials*, vol. 15, no. 3, pp. 1020–1041, 2013.
- [12] L. M. Ausubel, “An efficient dynamic auction for heterogeneous commodities,” *American Economic Review*, vol. 96, no. 3, pp. 602–629, 2006.
- [13] G. Demange, D. Gale, and M. Sotomayor, “Multi-item auctions,” *Journal of Political Economy*, vol. 94, pp. 863–872, 1986.
- [14] D. Mishra and R. Garg, “Descending price multi-item auctions,” *Journal of Mathematical Economics*, vol. 42, pp. 161–179, 2006.
- [15] R. P. McAfee, “A dominant strategy double auction,” *Journal of Economic Theory*, vol. 56, pp. 434–450, Apr. 1992.
- [16] D. Parkes, J. Kalagnanam, and M. Eso, “Achieving budget-balance with Vickrey-based payment schemes in exchanges,” in *Proc. International Joint Conference on Artificial Intelligence*, 2001, pp. 1161–1168.
- [17] R. B. Myerson and M. A. Satterthwaite, “Efficient mechanisms for bilateral trading,” *Journal of Economic Theory*, vol. 29, no. 2, pp. 265–281, Apr. 1983.
- [18] D. Yang, X. Fang, and G. Xue, “Truthful auction for cooperative communications,” 2011, <http://inside.mines.edu/~djyang/doc/MOBIHOC2011-TASC.pdf>.
- [19] L. Gkatzikis and I. Koutsopoulos, “Migrate or not? Exploiting dynamic task migration in mobile cloud computing systems,” *IEEE Wireless Communications*, vol. 20, no. 3, pp. 24–32, Jun. 2013.
- [20] Y. Xu and S. Mao, “A survey of mobile cloud computing for rich media applications,” *IEEE Wireless Communications*, vol. 20, no. 3, pp. 46–53, Jun. 2013.
- [21] L. Zhou, Z. Yang, J. Rodrigues, and M. Guizani, “Exploring blind online scheduling for mobile cloud multimedia services,” *IEEE Wireless Communications*, vol. 20, no. 3, pp. 54–61, Jun. 2013.
- [22] Y. Zhang, D. Niyato, and P. Wang, “An auction mechanism for resource allocation in mobile cloud computing systems,” in *Proc. International Conference on Wireless Algorithms, Systems, and Applications*, 2013, pp. 76–87.
- [23] W. Vickrey, “Counterspeculation, auctions, and competitive sealed tenders,” *The Journal of Finance*, vol. 16, pp. 8–37, 1961.
- [24] E. H. Clarke, “Multipart pricing of public goods,” *Public Choice*, vol. 11, pp. 17–33, 1971.
- [25] T. Groves, “Incentives in teams,” *Econometrica*, vol. 41, no. 4, pp. 617–631, Jul. 1973.
- [26] Y. Shi, S. Sharma, Y. T. Hou, and S. Kompella, “Optimal relay assignment for cooperative communications,” in *Proc. ACM MobiHoc*, 2008.
- [27] D. West, *Introduction to Graph Theory*, 2nd ed. Prentice Hall, 1999.
- [28] C. Lee, P. Wang, and D. Niyato, “A real-time group auction system for efficient allocation of cloud Internet applications,” *IEEE Transactions on Services Computing*, 2014, in press.



A-Long Jin received his B.Eng. degree in communications engineering from Nanjing University of Posts and Telecommunications, Nanjing, China, in 2012. He received his M.Sc. degree in computer science from University of New Brunswick, Fredericton, NB, Canada, in 2015. His research interests include cooperative wireless networks, mobile cloud computing, game theory for wireless networks, and machine learning.



Wei Song (M’09-SM’14) received the Ph.D. degree in electrical and computer engineering from University of Waterloo, Waterloo, ON, Canada, in 2007. In 2009, she joined the Faculty of Computer Science, University of New Brunswick, Fredericton, NB, Canada, where she is now an Associate Professor. Her current research interests include mobile cloud computing, cooperative wireless networking, energy-efficient wireless networks, and device-to-device communications. She received a *UNB Merit Award* in 2014, a *Best Student Paper Award* from IEEE CCNC 2013, a *Top 10% Award* from IEEE MMSP 2009, and a *Best Paper Award* from IEEE WCNC 2007. She is the Communications/Computer Chapter Chair of IEEE New Brunswick Section. She is also an editor for *IEEE Transactions on Vehicular Technology* and *Wireless Communications and Mobile Computing* (Wiley).



Weihua Zhuang (M’93-SM’01-F’08) has been with the Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON, Canada, since 1993, where she is a Professor and a Tier I Canada Research Chair in Wireless Communication Networks. She received the *Outstanding Performance Award* 4 times since 2005 from the University of Waterloo, and the *Premier’s Research Excellence Award* in 2001 from the Ontario Government. Dr. Zhuang was the Editor-in-Chief of *IEEE Transactions on Vehicular Technology* (2007-2013), and the Technical Program Symposia Chair of IEEE GLOBECOM 2011. She is also a Fellow of the Canadian Academy of Engineering (CAE), a Fellow of the Engineering Institute of Canada (EIC), and an elected member in the Board of Governors of the IEEE Vehicular Technology Society.