Online Cooperative Resource Allocation at the Edge: A Privacy-Preserving Approach

Yuqing Li, Hok Chun Ng, Lin Zhang, Bo Li
Hong Kong University of Science and Technology
liyuqing@ust.hk, {hcngac, lzhangbv}@connect.ust.hk, bli@cse.ust.hk

Abstract*—Mobile edge computing provides a platform facilitating individual servers to pool their resources locally for cooperative computation. One fundamental problem in this new paradigm is how to effectively allocate crowdsourced edge resources to users competing in a highly unpredicted environment. This, apparently, cannot be realized without a *truthful open market*. On the other hand, enforcing truthfulness potentially incurs privacy problems. There have been efforts in differentially private auctions, in which *exponential mechanism*, designed for single-sided single-item auctions, is a common solution. However, such an approach is not applicable in two-sided combinatorial edge markets, further complicated by the extra *migration cost* on energy-constrained users often imposed by online allocation. In this paper, we propose OPTA, an online privacy-preserving truthful double auction mechanism for *dynamic resource cooperation* at the edge. Given *uncertainties in future market behaviors*, we harness competitive analysis by decomposing the online optimization into a series of single-round auctions such that their objectives are iteratively adjusted to capture the temporally-coupled nature of the problem. In each round, by jointly considering the features of exponential mechanism and greedy heuristic, we design a near-optimal allocation policy with efficiency and privacy guarantee. We further implement a critical-value pricing scheme for winners, realizing the truthfulness in expectation. Building upon the single-round results, our overall online algorithm achieves a provable competitive ratio. We validate the desirable properties of OPTA through theoretical analysis and extensive simulations.

I. Introduction

Mobile edge computing (MEC) has emerged as a new computing paradigm that supports resource-intensive, or/and delay-sensitive applications at the network edge, particularly those that cannot be served efficiently by conventional cloud computing platforms due to unpredictable latency and expensive bandwidth [1]-[6]. One of the most appealing features of this paradigm is availability of abundant computing resources located in the proximity of users, which applications can readily take advantage of. There have been attempts by cloud service providers to establish computational capabilities at the edge [7] (e.g., Amazon CloudFront [8]), but reaching its full potential remains an elusive goal. This motivate us to design an *open platform* [9], where individual servers (e.g., servers at base stations [3], micro datacenters in enterprises [6], or PCs in research labs [10]) can contribute their computing capacity. By leveraging lightweight virtualization techniques [3], these servers form a shared resource pool to offer flexible supply of distributed resources at low cost and with low latency.

978-1-7281-6992-7/20/$31.00 ©2020 IEEE

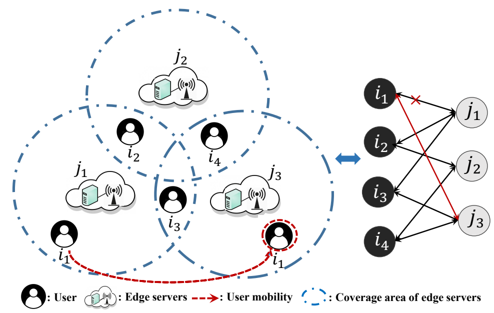

Fig. 1. Illustration of the open edge market for dynamic resource cooperation.

One of the intrinsic problems in MEC is how to optimally allocate the edge resources from the pool to satisfy diverse user demands in a competitive environment. This apparently cannot be materialized without the proper design of an efficient market [10]-[12]. Such issues have been extensively examined as a single-sided truthful auction market in cloud computing [13]-[14]. However, it is drastically different in an open edge platform, where both users and servers owned by different self-interested entities are strategic and non-cooperative in nature [15]. This forms a *truthful double auction-based open edge market* where a set of users (buyers) compete for resources from a set of edge servers (sellers) for computation offloading. The open platform is authorized to manage servers' payment and allocate the edge resources to users. In this way, a large number of potential edge servers can cooperatively provide services while economizing their valuable resources.

There are two fundamental challenges in the design and realization of an open edge market. The first comes from *system dynamics and uncertainties*. Edge servers operate in a dynamic and unpredictable environment. Different from existing auction models that focus on an isolated round [10]-[12], the unique dynamics in MEC systems call for new designs of the market running repetitively over time [16]-[19]. This raises the question of whether the connections between users and servers may persist over multiple rounds. If not, such cross-server migration would impose a considerable extra migration cost on users, e.g., energy consumption for connection establishment. In practice, most users are on a pre-allocated energy budget within a given time period, which makes the decisions across time intricately intertwined. Worse yet, the decisions have to be made without foreseeing future information, leading to sub-optimal allocation.

A Toy Example. We illustrate this problem using a toy example shown in Fig. 1. Edge servers j_1, j_2, j_3 cooperatively serve users i_1, i_2, i_3, i_4 within their coverage. First, each user submits a bid to compete for edge resources, and the auction comes out as that $(i_1, j_1), (i_2, j_2), (i_3, j_3)$ are winners. Given user mobility, assume after a while, i_1 moves to the coverage of j_3. If the task of this user is still offloaded to j_1, there may be allocation inefficiency or market failure. Thus, edge resources should be dynamically re-allocated to follow the *market dynamics*. That is, the market is expected to run repeatedly over time for an optimized user experience of MEC.

The second challenge is on *possible privacy leakage in multiple rounds of truthful bidding*. The desired market advocates that agents truthfully report their demands or supplies, which are private information. Although MEC platforms are usually considered trusted, there could exist honest-but-curious agents who follow the market mechanism strictly, but attempt to gain additional information [20]. It is generally known that the change in a single bid has the potential to alter the overall auction outcomes [21]-[22]. By regulating its own bid in multiple rounds and comparing the outcomes, a curious agent could infer bid information about others. With these information, one can learn the opponents' network configurations or business strategies for beneficial gain [23]. Therefore, the bid privacy issue, if not properly addressed, will discourage privacy-sensitive agents from participating in the market.

In this paper, for the first time, we develop an online privacy-preserving truthful double auction mechanism to facilitate dynamic resource allocation at the edge. Given uncertainties in future market behaviors, we harness competitive analysis by decomposing the online allocation problem into a series of single-round auctions. To address inter-round coupling, a scaled benefit specified for each bid-ask pair is introduced and iteratively adjusted based on user residual energy budget. We reformulate the $\mathcal{NP}$-hard single-round problem with a given scaled benefit vector, and design a heuristic truthful double auction via critical-value pricing. To alleviate bid privacy issue, we incorporate the notion of differential privacy[1] [25] such that any change in a single bid will not incur a significant change to the outcome. There have been efforts in differentially private auctions, in which *exponential mechanism* is a common choice [26]-[29]. Most existing solutions on pricing randomization assume uniform clearing pricing, making it not applicable in two-sided combinatorial edge markets. In contrast, we use exponential mechanism to randomize the winner selection by assigning a higher selection probability to the pair contributing more to social welfare. Building upon single-round results, our overall online algorithm achieves a provable competitive ratio.

Our main contributions are summarized as follows.

- We propose OPTA, an online privacy-preserving truthful market mechanism that facilitates individual servers to pool their resources locally for cooperative computation.
- We decompose the online optimization into a series of single-round auctions such that their objectives are iteratively adjusted to capture the temporally-coupled nature of the problem. In each round, we design an exponential based allocation policy by randomizing the winner selection, combined with critical-value pricing to form a differentially private truthful auction. We assign bid-ask pairs contributing more to social welfare with higher selection probabilities. This provides concrete guidelines on how to select winners efficiently and privately.
- We develop a competitive online algorithm which uses single-round auction as a building block towards an adaptive market mechanism dynamically allocating resources upon system changes. We theoretically prove that OPTA achieves a provable competitive ratio, differential privacy, truthfulness, approximate social welfare, individual rationality, budget balance and computational efficiency.

[1]Inference attacks fall into two major categories: membership inference and attribute inference [24]. In this work, we employ differential privacy techniques to protect bid privacy from membership inference attacks.

II. SYSTEM MODEL AND PROBLEM FORMULATION

A. Open Edge Market Model

Consider an online open edge market model consisting of a platform, a set $\mathcal{M} = \{1, \cdots, M\}$ of users (sellers) who pay for using edge resources for computation, and a set $\mathcal{N} = \{1, \cdots, N\}$ of servers (buyers) who have resources to share for profit. We use an agent to refer to either a user or a server. The time is discretized into slots, $\mathcal{T} = \{1, \cdots, T\}$, which is a much slower time scale than that of task execution [15]. In each slot, one round of resource trading is carried out, where the platform decides the allocation and price outcomes. The terms "slot" and "round" are used interchangeably. There may be other constraints (e.g., latency, geographic region) on whether a server is permitted to serve a user request [3]. For any slot $t \in \mathcal{T}$, we view the market model as an instance of bipartite graph $\mathcal{G}^t = \{\mathcal{M}, \mathcal{N}, \mathcal{E}^t\}$, where $(i, j) \in \mathcal{E}^t$ if server $j \in \mathcal{N}$ is permitted to serve user $i \in \mathcal{M}$. An illustration of a small-scale open edge market is shown in Fig. 1.

The proposed model can be generally applicable in the scenarios that require user interactions. For instance, in an app offering interactive gaming [30] services, servers are required to collect action information from the users, and then allocate resources to compute and generate new game states and videos for the users. Consider a MEC platform providing video streaming services [31], servers need to retrieve and process real-time camera frames from individual users, and respond with the object detection result. In such examples, users continuously generate requests, while servers receive rewards by completing user requests. This can be naturally represented by a market-based dynamic allocation.

B. Online Cooperative Resource Allocation Problem

In every slot, each user independently generates a task to be offloaded to the edge. Each task demands a combination of multiple types of resources [13]. Depending on the underlying MEC applications, these tasks may vary in resource demands [32]. Suppose there are K types of resources in set $\mathcal{K} = \{1, \cdots, K\}$, e.g., CPU, memory, bandwidth, etc. Let $\boldsymbol{\alpha}_i^t = \langle \alpha_{i,1}^t, \cdots, \alpha_{i,K}^t \rangle$ denote the desired resource bundle of user $i \in \mathcal{M}$ in slot t, where $\alpha_{i,k}^t \in [\alpha_{\min}, \alpha_{\max}]$ is the amount of type-k resource required by the task. We define the *bid* submitted by user i as $b_i^t = \langle \boldsymbol{\alpha}_i^t, v_i^t \rangle$, where valuation $v_i^t \in [v_{\min}, v_{\max}]$ is the maximum price it is willing to pay. Meanwhile, servers submit asks to the platform. Let $\boldsymbol{\beta}_j^t = \langle \beta_{j,1}^t, \cdots, \beta_{j,K}^t \rangle$ denote the idle resources that server $j \in \mathcal{N}$ can share, where $\beta_{j,k}^t \in [\beta_{\min}, \beta_{\max}]$. The *ask* submitted by server j can be specified as $a_j^t = \langle \boldsymbol{\beta}_j^t, \boldsymbol{c}_j^t \rangle$, where $\boldsymbol{c}_j^t = \langle c_{j,1}^t, \cdots, c_{j,K}^t \rangle$ with $c_{j,k}^t \in [c_{\min}, c_{\max}]$ being the minimum price at which it would sell per unit type-k resource. For user $i \in \mathcal{M}$, server $j \in \mathcal{N}$ has a bundle-specific cost towards its desired resource bundle, i.e., $c_{ij}^t = \sum_k c_{j,k}^t \alpha_{i,k}^t$.

Given all submitted bids/asks, the platform makes allocation decisions $\boldsymbol{x}^t$, the element of which is a binary variable, i.e,

$$x_{ij}^t = \begin{cases} 1, & \text{if } i \text{ is allocated to } j \text{ in slot } t, \\ 0, & \text{otherwise.} \end{cases} \tag{1}$$

Here user $i \in \mathcal{M}$ being "allocated" to server $j \in \mathcal{N}$ means $(i, j) \in \mathcal{E}^t$ and the task of i is offloaded to j. Dynamic allocation can lead to the redistribution of user requests, which incurs extra "migration cost", in particular extra energy consumption. Take an interactive gaming service as an example, in which each server covers certain area by serving gaming requests. When a user moves out of a service coverage area, the service needs to be migrated or handed over to a different server. This consumes energy and incurs migration cost due to the connection re-establishment and etc. This can be troublesome for mobile devices running on battery. In order to capture the energy consumption, we propose to optimize the long-term performance, i.e., to derive the optimal allocation in multiple rounds instead of a single round by incorporating the energy efficiency into our formulation. Assume every user i has a total budget B_i, which is the upper bound imposed on its overall migration cost. Given the allocation sequences $\boldsymbol{x}_i = (\boldsymbol{x}_i^1, \cdots, \boldsymbol{x}_i^T)$, the overall migration cost of i is captured by $\sum_t \sum_j s_i(1 - x_{ij}^{t-1} x_{ij}^t)$, where s_i is per-time migration cost.

We study a non-profit platform with the goal of maximizing long-term social welfare. The online allocation problem can be formulated as social welfare maximization problem (SWM):

$$\max \quad \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{M}} \sum_{j \in \mathcal{N}} \omega_{ij}^t \cdot x_{ij}^t \tag{2}$$

$$\text{s.t.} \quad \sum_{j \in \mathcal{N}} x_{ij}^t \leq 1, \forall i \in \mathcal{M}, t \in \mathcal{T} \tag{2a}$$

$$\sum_{i \in \mathcal{M}} \alpha_{i,k}^t x_{ij}^t \leq \beta_{j,k}^t, \forall k \in \mathcal{K}, j \in \mathcal{N}, t \in \mathcal{T} \tag{2b}$$

$$\sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{N}} s_i(1 - x_{ij}^{t-1} x_{ij}^t) \leq B_i, \forall i \in \mathcal{M} \tag{2c}$$

$$x_{ij}^t \in \{0, 1\}, \quad \forall (i,j) \in \mathcal{E}^t, t \in \mathcal{T}, \tag{2d}$$

where $\omega_{ij}^t = v_i^t - c_{ij}^t$ denotes the net benefit of offloading the task of user i to server j; $(2a)$ suggests a task should be executed as a whole; $(2b)$ implies the allocated resources are constrained by server capacity; $(2c)$ enforces the long-term migration budget for i; $(2d)$ shows any user can only be served by permitted servers for guaranteed service quality.

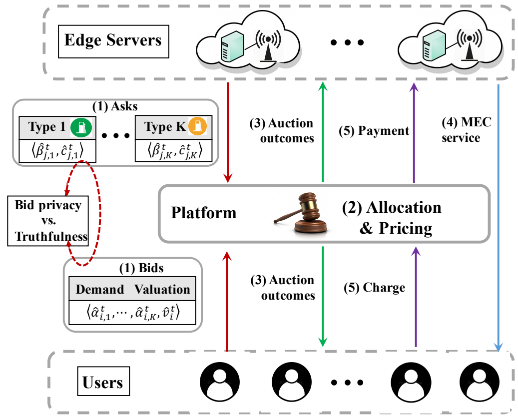

Fig. 2. General system workflow of the open edge market in round t.

C. Online Privacy-Preserving Truthful Double Auction

The optimal solution of SWM requires the full knowledge of all bids/asks, which are usually private information of agents. Even worse, if reported information are *untruthful*, the achieved social welfare would be far from the true optimum. In this scenario, online double auction becomes a desired approach. However, there are several challenges in applying online double auction: 1) system uncertainties require the decisions to be made without future bids/asks knowledge, often leading to sub-optimal allocation; 2) long-term energy consumption budget imposed on users complicates the decisions across slots; 3) multiple rounds of truthful bidding may cause bid privacy leakage, and discourage agent participation.

To address these issues, we develop OPTA, an online double auction mechanism to facilitate dynamic resource allocation with near-optimal efficiency and privacy guarantee. Fig. 2 shows the system workflow of the two-sided open edge market in each round. First, agents submit bids/asks to the platform, which may be different from real ones. We use $\hat{b}_i^t = \langle \hat{\boldsymbol{\alpha}}_i^t, \hat{v}_i^t \rangle$ and $\hat{a}_j^t = \langle \hat{\boldsymbol{\beta}}_j^t, \hat{\boldsymbol{c}}_j^t \rangle$ to denote the declarations of user i and server j in slot t. After that, the platform determines the winners $\mathcal{W}^t \subseteq \mathcal{E}^t$ and prices containing charge profile $\boldsymbol{q}^t = \langle q_1^t, \cdots, q_M^t \rangle$ for users and payment profile $\boldsymbol{p}^t = \langle p_1^t, \cdots, p_N^t \rangle$ for servers, where q_i^t, p_j^t are upper-bounded by $q_{\max}$, $p_{\max}$. $(i, j) \in \mathcal{W}^t$ suggests that the bid $\hat{b}_i^t$ and ask $\hat{a}_j^t$ win, and user i is allocated to server j. According to the auction outcomes, the platform requests resources from each winning server to serve its allocated users, and pays/charges them. Several favored properties of OPTA are expected to satisfy.

Definition 1 (Differential Privacy). *(revised from [26]). Denote the OPTA in slot t as a function $\Psi^t(\cdot)$ that maps input bids $\boldsymbol{b}^t$ and asks $\boldsymbol{a}^t$ to outcome $(\boldsymbol{x}^t, \boldsymbol{p}^t, \boldsymbol{q}^t)$. $\Psi^t(\cdot)$ gives (ϵ, δ)-differential privacy, if and only if for any two input sets, $(\boldsymbol{b}^t, \boldsymbol{a}^t)$ and $(\widetilde{\boldsymbol{b}}^t, \widetilde{\boldsymbol{a}}^t)$ differing in only one bid or one ask, and for any possible set of outcomes $\mathcal{S}^t \subseteq Range(\Psi^t)$, it satisfies*

$$\Pr[\Psi(\boldsymbol{b}^t, \boldsymbol{a}^t) \in \mathcal{S}^t] \leq e^{\epsilon} \cdot \Pr[\Psi(\widetilde{\boldsymbol{b}}^t, \widetilde{\boldsymbol{a}}^t) \in \mathcal{S}^t] + \delta. \tag{3}$$

Here privacy budget ϵ is to control privacy guarantee and δ is residual probability. For better privacy performance, ϵ and δ should be as close to 0 as possible. Differential privacy ensures that for each bid/ask, any arbitrary change of a single element would not greatly alter the outcome, making it hard for curious agents to infer bid information of others from the outcomes.

OPTA is also supposed to satisfy truthfulness, individual rationality (IR) and budget balance. Exact truthfulness is often too restrictive to be compatible with other desirable properties. Inspired by approximate truthfulness [21], we turn to a weaker but more practical concept of γ-truthfulness such that no agent can gain more than γ utility by bidding untruthfully. As for IR, it is expressed in terms of expected utility [33]. We first present the formal definition of agent utility.

Definition 2 (Utility for a Bid or an Ask). *For a user $i \in \mathcal{M}$, the utility for bid b_i^t is difference between true valuation for its desired resource bundle and charge of the platform, i.e.,*

$$u_i^t = \begin{cases} \sum_{j \in \mathcal{N}} v_i^t x_{ij}^t - q_i^t, & \text{if } i \text{ is a winner,} \\ 0, & \text{otherwise.} \end{cases} \tag{4}$$

For a server $j \in \mathcal{N}$, the utility for ask a_j^t is difference between payment from the platform and its overall actual cost, i.e.,

$$u_j^t = \begin{cases} p_j^t - \sum_{i \in \mathcal{M}} c_{ij}^t x_{ij}^t, & \text{if } j \text{ is a winner,} \\ 0, & \text{otherwise.} \end{cases} \tag{5}$$

Definition 3 (γ-Truthfulness). *Let b_i^t and a_j^t denote the truthful strategies for user $i \in \mathcal{M}$ and server $j \in \mathcal{N}$. For any small positive constant γ, OPTA achieves γ-truthfulness in expectation, if and only if, for any strategy $\hat{b}_i^t \neq b_i^t$,*

$$\mathbb{E}[u_i(b_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t)] \geq \mathbb{E}[u_i(\hat{b}_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t)] - \gamma, \tag{6}$$

and for any strategy $\hat{a}_j^t \neq a_j^t$,

$$\mathbb{E}[u_j(\hat{\boldsymbol{b}}^t, a_j^t, \hat{\boldsymbol{a}}_{-j}^t)] \geq \mathbb{E}[u_j(\hat{\boldsymbol{b}}^t, \hat{a}_j^t, \hat{\boldsymbol{a}}_{-j}^t)] - \gamma, \tag{7}$$

where $\hat{\boldsymbol{b}}_{-i}^t = \left\langle \hat{\boldsymbol{b}}_1^t, \cdots, \hat{\boldsymbol{b}}_{i-1}^t, \hat{\boldsymbol{b}}_{i+1}^t \cdots, \hat{\boldsymbol{b}}_M^t \right\rangle$ is the bid profile of users except i, and $\hat{\boldsymbol{a}}_{-j}^t = \langle \hat{\boldsymbol{a}}_1^t, \cdots, \hat{\boldsymbol{a}}_{j-1}^t, \hat{\boldsymbol{a}}_{j+1}^t \cdots, \hat{\boldsymbol{a}}_N^t \rangle$ denotes the ask profile of servers except j.

Definition 4 (Individual Rationality). *OPTA achieve IR only if all agents have non-negative expected utilities, i.e., $\mathbb{E}[u_i(\hat{\boldsymbol{b}}^t, \hat{\boldsymbol{a}}^t)] \geq 0$ for $i \in \mathcal{M}$ and $\mathbb{E}[u_j(\hat{\boldsymbol{b}}^t, \hat{\boldsymbol{a}}^t)] \geq 0$ for $j \in \mathcal{N}$.*

Definition 5 (Budget Balance). *OPTA achieves budget balance if the total charge from users is exactly sufficient to cover overall payment for servers, i.e., $\sum_i q_i^t = \sum_j p_j^t$.*

III. Single-Round Double Auction Mechanism

This paper addresses a general situation where the *market model varies over time under system dynamics*. We first reformulate single-round SWM and design sPTA, a single-round privacy-preserving truthful auction to determine the winners and prices. This will be a key step repeatedly invoked by the online auction mechanism we develop next, to derive the optimal outcomes of all slots jointly.

A. Single-Round SWM Problem

The single-round SWM problem is formulated as follows, which includes the same constraints related to the current slot from SWM, excludes budget constraints (2c), and modifies ω_{ij}^t to a scaled benefit $\hat{\omega}_{ij}^t$ (the rationale is detailed in Section IV).

$$\max \quad \sum_{i \in \mathcal{M}} \sum_{j \in \mathcal{N}} \hat{\omega}_{ij}^t \cdot x_{ij}^t \tag{8}$$

$$\text{s.t.} \quad \sum_{j \in \mathcal{N}} x_{ij}^t \leq 1, \forall i \in \mathcal{M} \tag{8a}$$

$$\sum_{i \in \mathcal{M}} \alpha_{i,k}^t x_{ij}^t \leq \beta_{j,k}^t, \forall k \in \mathcal{K}, j \in \mathcal{N} \tag{8b}$$

$$x_{ij}^t \in \{0, 1\}, \ \ \forall (i,j) \in \mathcal{E}^t. \tag{8c}$$

Consider a special case of single-round SWM by letting $N = K = 1$. The resulting problem is an instance of 0-1 knapsack problem, which is known to be $\mathcal{NP}$-hard [34]. Therefore, the single-round SWM problem is $\mathcal{NP}$-hard.

B. Exponential Mechanism

We first introduce the concept of exponential mechanism [35]-[36], that has been widely adopted to develop differentially private schemes. For any given outcome domain, the exponential mechanism denoted by $\epsilon_f^\epsilon(\mathcal{A})$ selects an outcome through a score function $f(\mathcal{A}, \mathcal{B})$, which maps a pair of input set $\mathcal{A}$ and possible outcome $\mathcal{B}$ to a real-valued score. This score captures how good $\mathcal{B}$ is for $\mathcal{A}$ compared with the optimal outcome. Let Δ be the sensitivity of f, denoting the largest change in scores on two input sets differing in only one element. The mechanism $\epsilon_f^\epsilon(\mathcal{A})$ gives $2\epsilon\Delta$-differential privacy, if it randomly generates the outcome based on the probability

$$\Pr\left[\epsilon_f^\epsilon(\mathcal{A}) = \mathcal{B}\right] \propto \exp(\epsilon f(\mathcal{A}, \mathcal{B})). \tag{9}$$

Accordingly, the outcomes with higher scores will be selected with exponentially larger probabilities, enabling the final outcome close to the optimum with respect to $f(\mathcal{A}, \mathcal{B})$. Moreover, the exponential mechanism leads to a highly near-optimal outcome with exponentially low probability.

Lemma 1. *The exponential mechanism $\epsilon_f^\epsilon(\mathcal{A})$ yields $2\epsilon\Delta$-differential privacy, when used to select an outcome $\mathcal{B} \in \mathcal{S}$. Let $\tilde{\mathcal{S}}$ be the subset of $\mathcal{S}$ achieving $f(\mathcal{A}, \mathcal{B}) = \max_{\mathcal{B}} f(\mathcal{A}, \mathcal{B})$. By Theorem 2.1 in [26], for any $\vartheta \geq 0$, $\epsilon_f^\epsilon(\mathcal{A})$ ensures that*

$$\Pr\left[f\left(\mathcal{A}, \epsilon_f^\epsilon(\mathcal{A})\right) < \max_{\mathcal{B}} f(\mathcal{A}, \mathcal{B}) - \frac{1}{\epsilon} \cdot \ln\frac{|\mathcal{S}|}{|\tilde{\mathcal{S}}|} - \frac{\vartheta}{\epsilon}\right] \leq e^{-\vartheta}.$$

C. Design Detail

To enforce differential privacy and truthfulness, sPTA utilizes exponential mechanism to randomize the allocation outcomes, and then charge/pay winners via critical-value pricing. To facilitate analysis, we first presume agents behave truthfully and show truthfulness is satisfied later.

1) Differentially Private Allocation: We first define a list of feasible bid-ask pairs in slot t, denoted by $\mathcal{L}^t = \{(i,j) | \hat{\omega}_{ij}^t \geq 0, \forall (i,j) \in \mathcal{E}^t\}$. Given $\mathcal{NP}$-hardness of single-round SWM, we harness heuristic techniques to design an allocation policy that approximates the optimal social welfare with computation efficiency. Intuitively, a more socially-efficient pair that generates greater social welfare with less resources is more

Algorithm 1: Differentially Private Allocation Policy in t

Input: $\mathcal{G}^t$, $\boldsymbol{b}^t$, $\boldsymbol{a}^t$, ϵ, Δ, δ
Output: $\boldsymbol{x}^t$, $\mathcal{W}^t$

1 $\mathcal{W}^t \leftarrow \emptyset$, $\boldsymbol{x}^t \leftarrow \boldsymbol{0}$; $\hat{\epsilon} \leftarrow \epsilon/(e\Delta\ln(e/\delta))$;
2 $\mathcal{L}^t \leftarrow \{(i,j)|\hat{\omega}_{ij}^t \geq 0, \forall (i,j) \in \mathcal{E}^t\}; \hat{\mathcal{L}}^t \leftarrow \mathcal{L}^t$;
3 **while** $|\hat{\mathcal{L}}^t| > 0$ **do**
4   **for** *each* $(i,j) \in \hat{\mathcal{L}}^t$ **do**
5   **if** $i \in \mathcal{W}^t$ & $\exists k \in \mathcal{K}, \beta_{j,k}^t < \alpha_{i,k}^t$ **then**
6             $\hat{\mathcal{L}}^t \leftarrow \hat{\mathcal{L}}^t \setminus \{(i,j)\}$;
7   Calculate allocation efficiency ρ_{ij}^t by (10);
8         $\Pr[\mathcal{W}^t \leftarrow \mathcal{W}^t \cup \{(i,j)\}] \leftarrow \frac{\exp(\hat{\epsilon}\rho_{ij}^t)}{\sum_{(i',j')\in\hat{\mathcal{L}}^t}\exp(\hat{\epsilon}\rho_{i'j'}^t)}$;
9   Randomly select (i,j) by $\Pr[\mathcal{W}^t \leftarrow \mathcal{W}^t \cup \{(i,j)\}]$;
10   **if** (i,j) *is selected* **then**
11         $x_{ij}^t \leftarrow 1$; $\hat{\mathcal{L}}^t \leftarrow \hat{\mathcal{L}}^t \setminus (i,j)$; $\mathcal{W}^t \leftarrow \mathcal{W}^t \cup \{(i,j)\}$;
12   **for** *each* $k \in \mathcal{K}$ **do**
13             $\beta_j^k \leftarrow \beta_{j,k}^t - \alpha_{j,k}^t$.

Algorithm 2: Critical-Value Pricing Policy in t

Input: $\boldsymbol{x}^t$, $\mathcal{W}^t$, $\mathcal{L}^t$, $\boldsymbol{b}^t$, $\boldsymbol{a}^t$
Output: $\boldsymbol{q}^t$, $\boldsymbol{p}^t$

1 $\boldsymbol{q}^t, \boldsymbol{p}^t \leftarrow \emptyset$; Re-order $\mathcal{L}^t$ in decreasing order of ρ_{ij}^t;
2 **for** *each winning pair* $(i,j) \in \mathcal{W}^t$ **do**
3     $\hat{\boldsymbol{q}}^c \leftarrow \emptyset$; $\hat{\boldsymbol{\beta}}_j^t \leftarrow \boldsymbol{\beta}_j^t$; $\mathcal{L}_{-i}^t \leftarrow \{(\hat{i},j)|(\hat{i},j) \in \mathcal{L}^t, \hat{i} \neq i\}$;
4   **if** $|\mathcal{L}_{-i}^t| < 1$ **then**
5         $\hat{\boldsymbol{q}}^c \leftarrow \hat{\boldsymbol{q}}^c \cup \{c_{ij}^t\}$;
6   **else**
7   **for** *each* $(\hat{i},j) \in \mathcal{L}_{-i}^t$ **do**
8   **if** $\hat{i} \notin \mathcal{W}^t$ & $\hat{\beta}_{j,k}^t \geq \alpha_{\hat{i},k}^t, \forall k \in \mathcal{K}$ **then**
9                 $\hat{\boldsymbol{q}}^c \leftarrow \hat{\boldsymbol{q}}^c \cup \{c_{ij}^t + \rho_{\hat{i}j}^t \cdot \sum_k \alpha_{i,k}^t/\beta_{j,k}^t\}$;
10                 $\hat{\beta}_{j,k}^t \leftarrow \hat{\beta}_{j,k}^t - \alpha_{i,k}^t, \forall k \in \mathcal{K}$;
11     $q_i^t \leftarrow \min(\hat{\boldsymbol{q}}^c)$; $\boldsymbol{q}^t \leftarrow \boldsymbol{q}^t \cup \{q_i^t\}$;
12 **for** *each winning server* $j \in \mathcal{W}^t$ **do**
13     $p_j^t \leftarrow \sum_i x_{ij}^t q_i^t$; $\boldsymbol{p}^t \leftarrow \boldsymbol{p}^t \cup \{p_j^t\}$.

preferred. For any feasible pair, we introduce the allocation efficiency metric to evaluate its contribution to social welfare:

$$\rho_{ij}^t = \frac{\hat{\omega}_{ij}^t}{\sum_k f_{j,k}^t \alpha_{i,k}^t}, \forall (i,j) \in \mathcal{L}^t. \tag{10}$$

The denominator determines the weight of bundle $\boldsymbol{\alpha}_i^t$ by homogenizing resources via relevance factor $f_{j,k}^t$, which captures scarcity of type-k resources for j. Intuitively, high $f_{j,k}^t$ means high resource scarcity, and low allocation efficiency. One simple measure of resource scarcity is $f_{j,k}^t = 1/\beta_{j,k}^t$. We employ exponential mechanism to randomize the allocation outcomes. To optimize social welfare, we associate the score function with allocation efficiency, making a pair contributing more assigned a higher selection probability. The platform chooses winners iteratively and maintains the remaining feasible set $\hat{\mathcal{L}}^t$. In each iteration, the selection probability for each feasible pair is proportional to the exponential of its score; and 0 otherwise. By normalizing, the selection probability for each pair is

$$\Pr[\mathcal{W}^t \leftarrow \mathcal{W}^t \cup \{(i,j)\}] = \begin{cases} \dfrac{\exp(\hat{\epsilon}\rho_{ij}^t)}{\sum\limits_{(i',j')\in\hat{\mathcal{L}}^t} \exp(\hat{\epsilon}\rho_{i'j'}^t)}, \text{if} (i,j) \in \hat{\mathcal{L}}^t, \\ 0, \quad \text{otherwise}, \end{cases}$$

where $\hat{\epsilon} = \epsilon/(e\Delta\ln(e/\delta))$. Here $\Delta = \alpha_{\max}(v_{\max} - 1)/\beta_{\min}$ is the sensitivity of ρ_{ij}^t, and $\epsilon, \delta > 0$ are parameters to balance privacy leakage and social welfare. The detailed algorithm is shown in Alg. 1. The platform constructs a list $\hat{\mathcal{L}}^t$ containing all feasible bid-ask pairs (line 2), and then iteratively selects winners (lines 3-13). In each iteration, redundant pairs are eliminated (lines 5-6) while each remaining pair is assigned a selection probability (lines 7-8). Afterwards, randomly select one pair (line 9) and update server capacity, feasible pair set and allocation outcome (lines 10-13). Repeat such randomized allocation until no feasible pairs can be selected.

2) Critical-Value Pricing: To guarantee truthfulness, we apply critical-value pricing to compute the charges/payments for winners. Critical-value pricing charges the winner the highest price of losing competitors [37]. If one bids higher than critical value, it wins; otherwise it loses. Given combinatorial nature of edge markets, we incorporate bundle-specific bidding into pricing. Different from existing critical-value pricing solutions, for any winning user, we need to recalculate the prices of losing competitors for its desired bundle, and select the smallest one. Alg. 2 is proposed to realize this goal. The platform first sorts all feasible pairs in descending order of allocation efficiency (line 1). For each winning user i, initialize the set $\hat{\boldsymbol{q}}^c$ of minimum prices required to outbid competitors and construct a new sorted list $\mathcal{L}_{-i}^t$ containing all initial feasible pairs, except for redundant ones with i (line 3). After searching for competitors among $\mathcal{L}_{-i}^t$, update the minimum price set and ongoing resource capacity (line 10). Repeat this process for all pairs in $\mathcal{L}_{-i}^t$. Given winning users' charges, the payment for each winning server can be obtained by accumulating the charges of associated users (lines 12-13).

D. Theoretical Analysis

Theorem 1. *sPTA provides $(\widetilde{\epsilon}, \delta)$-differential privacy for both user valuation and server cost, where $\delta \leq (0, 1/2]$ and $\widetilde{\epsilon} = (e-1)\hat{\epsilon}\ln(\frac{e}{\delta})\frac{\beta_{\max}}{\alpha_{\min}} \cdot \max\{v_{\max} - v_{\min}, c_{\max} - c_{\min}\}$.*

Proof. Let $\{\boldsymbol{b}^t, \boldsymbol{a}^t\}$ and $\{\widetilde{\boldsymbol{b}}^t, \widetilde{\boldsymbol{a}}^t\}$ be two input sets differing in only one bid or one ask. We need to give an exponential upper-bound for $\frac{\Pr[\Psi(\boldsymbol{b}^t, \boldsymbol{a}^t) = \mathcal{W}^t]}{\Pr[\Psi(\widetilde{\boldsymbol{b}}^t, \widetilde{\boldsymbol{a}}^t) = \widetilde{\mathcal{W}}^t]}$, where $\mathcal{W}^t = \widetilde{\mathcal{W}}^t = \{w_1, \cdots, w_g, \cdots, w_l\}$ are two same ordered winner sets with length l and w_h is selected before w_g for any $g > h$. We have

$$\frac{\Pr[\Psi(\boldsymbol{b}^t, \boldsymbol{a}^t) = \mathcal{W}^t]}{\Pr[\Psi(\widetilde{\boldsymbol{b}}^t, \widetilde{\boldsymbol{a}}^t) = \widetilde{\mathcal{W}}^t]} = \frac{\Pr[\mathcal{W}^t \leftarrow \mathcal{W}^t = \mathcal{W}_1 \cup \{w_1\}] \cdots \Pr[\mathcal{W}^t \leftarrow \mathcal{W}^t = \mathcal{W}_l \cup \{w_l\}]}{\Pr[\widetilde{\mathcal{W}} \leftarrow \widetilde{\mathcal{W}}^t = \mathcal{W}_1 \cup \{w_1\}] \cdots \Pr[\widetilde{\mathcal{W}}^t \leftarrow \widetilde{\mathcal{W}}^t = \mathcal{W}_l \cup \{w_l\}]}$$
$$= \prod_{g=1}^{l} \frac{\exp(\hat{\epsilon}\rho_g)}{\exp(\hat{\epsilon}\widetilde{\rho}_g)} \cdot \prod_{g=1}^{l} \frac{\sum_{w_h \in \hat{\mathcal{L}}^t \setminus \mathcal{W}_g} \exp(\hat{\epsilon}\widetilde{\rho}_h)}{\sum_{w_h \in \hat{\mathcal{L}}^t \setminus \mathcal{W}_g} \exp(\hat{\epsilon}\rho_h)}.$$

• **Privacy for users.** Assume $\{\boldsymbol{b}^t, \boldsymbol{a}^t\}$ and $\{\widetilde{\boldsymbol{b}}^t, \widetilde{\boldsymbol{a}}^t\}$ only differ in the valuation of user i'. Then

$$\frac{\Pr[\Psi(\boldsymbol{b}^t,\boldsymbol{a}^t)=\mathcal{W}^t]}{\Pr[\Psi(\widetilde{\boldsymbol{b}}^t,\widetilde{\boldsymbol{a}}^t)=\widetilde{\mathcal{W}}^t]} = \prod_{g=1}^{l} \frac{\sum_{w_h} \exp\left(\frac{\hat{\epsilon}\left(\widetilde{v}_i^t - c_{ij}^t\right)}{\sum_k \alpha_{i,k}^t/\beta_{j,k}^t}\right)}{\sum_{w_h} \exp\left(\frac{\hat{\epsilon}\left(v_i^t - c_{ij}^t\right)}{\sum_k \alpha_{i,k}^t/\beta_{j,k}^t}\right)} \exp\left(\frac{\hat{\epsilon}\left(v_{i'}^t - \widetilde{v}_{i'}^t\right)}{\sum_k \alpha_{i',k}^t/\beta_{j',k}^t}\right).$$

Case 1: $\widetilde{v}_{i'}^t < v_{i'}^t$. Let $\Delta' = v_{\max} - v_{\min}$. The first term is less than 1, and we have

$$\frac{\Pr[\Psi(\boldsymbol{b}^t,\boldsymbol{a}^t)=\mathcal{W}^t]}{\Pr[\Psi(\widetilde{\boldsymbol{b}}^t,\widetilde{\boldsymbol{a}}^t)=\widetilde{\mathcal{W}}^t]} < \exp\left(\frac{\hat{\epsilon}\left(v_{i'}^t - \widetilde{v}_{i'}^t\right)}{\sum_k \alpha_{i',k}^t/\beta_{j',k}^t}\right) < \exp\left(\hat{\epsilon}\Delta' \frac{\beta_{\max}}{\alpha_{\min}}\right).$$

Case 2: $\widetilde{v}_{i'}^t > v_{i'}^t$. The second term is smaller than 1. Denote $\sigma_i^v = \widetilde{v}_i^t - v_i^t$. For all $\hat{\epsilon} \leq \alpha_{\min}/(\Delta'\beta_{\max})$, we obtain

$$\begin{aligned}\frac{\Pr[\Psi(\boldsymbol{b}^t,\boldsymbol{a}^t)=\mathcal{W}^t]}{\Pr[\Psi(\widetilde{\boldsymbol{b}}^t,\widetilde{\boldsymbol{a}}^t)=\widetilde{\mathcal{W}}^t]} &< \prod_{g=1}^{l} \frac{\sum_{w_h} \exp\left(\frac{\hat{\epsilon}\sigma_i^v}{\sum_k \alpha_{i,k}^t/\beta_{j,k}^t}\right)\exp\left(\frac{\hat{\epsilon}\left(v_i^t - c_{ij}^t\right)}{\sum_k \alpha_{i,k}^t/\beta_{j,k}^t}\right)}{\sum_{w_h} \exp\left(\frac{\hat{\epsilon}\left(v_i^t - c_{ij}^t\right)}{\sum_k \alpha_{i,k}^t/\beta_{j,k}^t}\right)} \\ &\leq \prod_{g=1}^{l} \mathbb{E}_{w_h}\left[1 + (e-1)\frac{\hat{\epsilon}\sigma_i^v}{\sum_k \alpha_{i,k}^t/\beta_{j,k}^t}\right] \\ &\leq \exp\left((e-1)\hat{\epsilon}\sum_{g=1}^{l} \mathbb{E}_{w_h}\left[\frac{\sigma_i^v}{\sum_k \alpha_{i,k}^t/\beta_{j,k}^t}\right]\right).\end{aligned}$$

We call a winner set q-good if $\sum_g \mathbb{E}_{w_h}\left[\frac{\sigma_i^v}{\sum_k \alpha_{i,k}^t/\beta_{j,k}^t}\right]$ is upper-bounded by $q \cdot \Delta' \frac{\beta_{\max}}{\alpha_{\min}}$, and q-bad otherwise. According to [25], the probability that $\mathcal{W}^t$ is q-bad is bounded by e^{1-q}. For any possible outcome set, $\mathcal{S}^t$, we can split it into $\hat{\mathcal{S}}^t = \{\mathcal{W}^t \in \mathcal{S}^t : \mathcal{W}^t \text{ is } (\ln(\frac{e}{\delta}))\text{-good}\}$ and $\widetilde{\mathcal{S}}^t = \mathcal{S}^t \backslash \hat{\mathcal{S}}^t$. Then

$$\begin{aligned}&\Pr[\Psi(\boldsymbol{b}^t, \boldsymbol{a}^t) \in \mathcal{S}^t] \\ &= \sum_{\mathcal{W}^t \in \hat{\mathcal{S}}^t} \Pr[\Psi(\boldsymbol{b}^t, \boldsymbol{a}^t) = \mathcal{W}^t] + \sum_{\mathcal{W}^t \in \widetilde{\mathcal{S}}^t} \Pr[\Psi(\boldsymbol{b}^t, \boldsymbol{a}^t) = \mathcal{W}^t] \\ &\leq \sum_{\mathcal{W}^t \in \hat{\mathcal{S}}^t} \exp\left((e-1)\hat{\epsilon}\Delta' \ln(\tfrac{e}{\delta})\tfrac{\beta_{\max}}{\alpha_{\min}}\right) \Pr[\Psi(\widetilde{\boldsymbol{b}}^t, \widetilde{\boldsymbol{a}}^t) = \mathcal{W}^t] + \delta \\ &= \exp\left((e-1)\hat{\epsilon}\Delta' \ln(\tfrac{e}{\delta})\tfrac{\beta_{\max}}{\alpha_{\min}}\right) \Pr[\Psi(\widetilde{\boldsymbol{b}}^t, \widetilde{\boldsymbol{a}}^t) \in \mathcal{S}^t] + \delta.\end{aligned}$$

Hence, $\Pr[\Psi(\boldsymbol{b}^t, \boldsymbol{a}^t) \in \mathcal{S}^t] \leq \exp(\widetilde{\epsilon}) \Pr[\Psi(\widetilde{\boldsymbol{b}}^t, \widetilde{\boldsymbol{a}}^t) \in \mathcal{S}^t] + \delta$ for both cases. sPTA is $(\widetilde{\epsilon}, \delta)$-differentially private for users.

• **Privacy for servers.** The proof is similar to the proof for user privacy. We thus omit it due to space limit. □

Theorem 2. *sPTA is γ-truthful for both agents, where $\gamma = (\hat{\epsilon} + \delta e^{-\hat{\epsilon}}) \cdot max\{v_{max} - c_{min}, p_{max} - c_{min}\}$.*

Proof. To facilitate analysis, we denote $b_i^t \succcurlyeq \hat{b}_i^t$ for user $i \in \mathcal{M}$ if i reports a higher valuation and requests fewer resources, i.e., $v_i^t \geq \hat{v}_i^t$ and $\sum_k \alpha_{i,k}^t f_{j,k}^t \leq \sum_k \hat{\alpha}_{i,k}^t f_{j,k}^t, \forall (i,j)$. Similarly, $a_j^t \succcurlyeq \hat{a}_j^t$ if $\sum_k \beta_{j,k}^t \geq \sum_k \hat{\beta}_{j,k}^t$ and $c_{ij}^t \leq \hat{c}_{ij}^t, \forall (i,j)$.

• **Truthfulness for users.** For user $i \in \mathcal{M}$, it gains utility $\mathbb{E}[u_i(b_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t)]$ when bidding truthfully and $\mathbb{E}[u_i(\hat{b}_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t)]$ otherwise. We prove $\mathbb{E}[u_i(b_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t)] \geq \mathbb{E}[u_i(\hat{b}_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t)] - \gamma$, which can be separated into two cases.

Case 1: $\hat{b}_i^t \succcurlyeq b_i^t$, i.e., $\sum_k \hat{\alpha}_{i,k}^t f_{j,k}^t \leq \sum_k \alpha_{i,k}^t f_{j,k}^t$ or $\hat{v}_i^t \geq v_i^t$. If i only receives a subset of $\boldsymbol{\alpha}_i^t$, $v_i^t = 0$ and $u_i^t = v_i^t - q_i^t \leq 0$. It holds. We next consider $\hat{v}_i^t \geq v_i^t$. Given winner set $\mathcal{W}^t$, $\hat{v}_i^t \geq v_i^t$ for (i,j) corresponds to three cases: (1) $c_{ij}^t > \hat{v}_i^t \geq v_i^t$. Since $v_i^t - c_{ij}^t < \hat{v}_i^t - c_{ij}^t < 0$, i loses with b_i^t and $\hat{b}_i^t$, and $u_i^t = 0$; (2) $\hat{v}_i^t \geq c_{ij}^t > v_i^t$. i loses with b_i^t and is randomly selected with $\hat{b}_i^t$. If selected, $u_i^t < 0$ since $q_i^t \geq c_{ij}^t$; (3) $\hat{v}_i^t \geq v_i^t \geq c_{ij}^t$. i is randomly selected with both bids. If selected, $u_i^t \geq 0$ when $v_i^t \geq q_i^t$, and $u_i^t \leq 0$ otherwise. Hence, $\mathbb{E}_{\mathcal{W}^t}[u_i(b_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t)] = \mathbb{E}_{\mathcal{W}^t}[u_i(\hat{b}_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t)]$ when $v_i^t \geq q_i^t$. We have

$$\begin{aligned}&\mathbb{E}[u_i(b_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t)] \\ &= \sum_{\mathcal{W}^t \in \mathcal{S}^t} \mathbb{E}_{\mathcal{W}^t}[u_i(b_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t)] \Pr[\Psi(b_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t) = \mathcal{W}^t] \\ &\geq \sum_{\mathcal{W}^t \in \mathcal{S}^t \wedge q_i^t < v_i^t} \mathbb{E}_{\mathcal{W}^t}[u_i(b_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t)] \Pr[\Psi(b_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t) = \mathcal{W}^t] \\ &\geq \mathbb{E}[u_i(\hat{b}_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t)].\end{aligned}$$

Case 2: $b_i^t \succcurlyeq \hat{b}_i^t$, i.e., $\sum_k \hat{\alpha}_{i,k}^t f_{j,k}^t \geq \sum_k \alpha_{i,k}^t f_{j,k}^t$ or $\hat{v}_i^t \leq v_i^t$. i loses the incentive to submit a larger bundle, which may decrease the selection probability. We focus on $\hat{v}_i^t \leq v_i^t$. There are three cases: (1) $c_{ij}^t > v_i^t \geq \hat{v}_i^t$. Similarly, i loses with for both bids and $u_i^t = 0$; (2) $v_i^t \geq c_{ij}^t > \hat{v}_i^t$. i loses with $\hat{b}_i^t$ and is randomly selected with b_i^t. If selected, $u_i^t > 0$ when $v_i^t > q_i^t \geq c_{ij}^t$, and $u_i^t \leq 0$ otherwise; (3) $v_i^t \geq \hat{v}_i^t \geq c_{ij}^t$. i is randomly selected with both bids, where the selection probability with b_i^t is larger than that with $\hat{b}_i^t$. If selected, $u_i^t > 0$ when $v_i^t > q_i^t \geq c_{ij}^t$, and $u_i^t \leq 0$ otherwise. Hence, $\mathbb{E}_{\mathcal{W}^t}[u_i(b_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t)] \geq \mathbb{E}_{\mathcal{W}^t}[u_i(\hat{b}_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t)]$. Since $u_i^t \leq u_{\max}^I = v_{\max} - c_{\min}$, we have

$$\begin{aligned}&\mathbb{E}[u_i(b_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t)] \\ &\geq \frac{1}{e^{\hat{\epsilon}}} \sum_{\mathcal{W}^t \in \mathcal{S}^t} \mathbb{E}_{\mathcal{W}^t}[u_i(\hat{b}_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t)](\Pr[\Psi(\hat{b}_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t) = \mathcal{W}^t] - \delta) \\ &\geq (1 - \hat{\epsilon})\mathbb{E}[u_i(\hat{b}_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t)] - \delta e^{-\hat{\epsilon}} u_{\max}^I \\ &\geq \mathbb{E}[u_i(\hat{b}_i^t, \hat{\boldsymbol{b}}_{-i}^t, \hat{\boldsymbol{a}}^t)] - (\hat{\epsilon} + \delta e^{-\hat{\epsilon}}) \cdot (v_{\max} - c_{\min}).\end{aligned}$$

Above all, we conclude that sPTA is γ-truthful for users.

• **Truthfulness for servers.** The proof is similar to the proof for user truthfulness. We thus omit it due to space limit. □

Remark. Theorem 2 claims that no agent can gain in its expected utility with γ by bidding untruthfully. Hence, it is reasonable to consider agents in sPTA submit true bids/asks.

Theorem 3. *sPTA is of IR and budget balance.*

Proof. To verify IR, we show all agents receive non-negative expected utilities. Any loser is free of charge and here we only consider winners. For any user i with winning bid $\hat{b}_i^t$, $\mathbb{E}[u_i^t] = \mathbb{E}[v_i^t - q_i^t]$ under truthful bidding. By Alg. 2, $q_i^t = c_{ij}^t + \rho_{\hat{i}^*j}^t \sum_k \alpha_{i,k}^t/\beta_{j,k}^t$, where $\hat{i}^* = \arg\min_j\{\rho_{i'j}^t\}$. Since (i,j) is more competitive than $(\hat{i}^*, j)$ with a larger probability, $\mathbb{E}[\rho_{ij}^t - \rho_{\hat{i}^*j}^t] \geq 0$. Thus $\mathbb{E}[u_i] \geq \sum_k \frac{\alpha_{i,k}^t}{\beta_{j,k}^t} \mathbb{E}[\rho_{ij}^t - \rho_{\hat{i}^*j}^t] \geq 0$. For any server j with winning ask $\hat{a}_j^t$, $u_j^t = p_j^t - \sum_i x_{ij}^t c_{ij}^t$. Then, $\mathbb{E}[u_j^t] = \mathbb{E}\left[\sum_i x_{ij}^t (q_i^t - c_{ij}^t)\right] = \mathbb{E}\left[\sum_i x_{ij}^t \rho_{\hat{i}^*j}^t \sum_k \frac{\alpha_{i,k}^t}{\beta_{j,k}^t}\right] \geq 0$, i.e., IR holds. By Alg. 2, $\sum_j p_j^t = \sum_j \sum_i x_{ij}^t q_i^t = \sum_i q_i^t$. Thus, budget balance is guaranteed. □

Theorem 4. *With the probability of at least $1 - 1/L^{\mathcal{O}(1)}$, sPTA can compute an outcome with a minimum social welfare $\frac{1}{MK}\frac{\alpha_{min}}{\beta_{max}} R_s^{t,*} - \mathcal{O}(lnL)$. Here $R_s^{t,*}$ is the optimal social welfare of single-round SWM, and L is number of feasible pairs in $\mathcal{L}^t$.*

Proof. Let $\mathcal{W}^{t,*}$ be the optimal winner set with the maximum social welfare $R_s^{t,*}$. For sPTA, consider an arbitrary winner

Algorithm 3: Online Auction Algorithm

Input: $\mathcal{G}^t, \boldsymbol{b}^t, \boldsymbol{a}^t, \forall t,\ s_i, B_i, \forall i,\ \epsilon,\ \Delta,\ \delta$

Output: $\boldsymbol{x}^t, \boldsymbol{p}^t, \boldsymbol{q}^t, \forall t$

1 $\kappa_{ij}^t \leftarrow 0, J_i^t \leftarrow 0, \forall i, j, t;\ \hat{\eta} \leftarrow \max_i \frac{s_i}{B_i};\ \eta \leftarrow (1+\hat{\eta})^{\frac{1}{\hat{\eta}}}$;

2 **for** *each slot* $t \in \mathcal{T}$ **do**

3   **for** *each bid-ask pair* $(i,j) \in \mathcal{E}^t$ **do**

4     $s_{ij}^t \leftarrow s_i \cdot \mathbb{1}_{\{j \neq J_i^{t-1}\}}$;

5   **if** $\kappa_{ij}^{t-1} \geq 1 \cup \omega_{ij}^t \geq s_{ij}^t$ **then**

6       $\hat{\omega}_{ij}^t \leftarrow 0$;

7   **else**

8       $\hat{\omega}_{ij}^t \leftarrow \omega_{ij}^t - s_{ij}^t \kappa_{ij}^{t-1}$;

9   Run sPTA to obtain auction outcome $(\boldsymbol{x}^t, \boldsymbol{p}^t, \boldsymbol{q}^t)$;

10   **for** *each user* $i \in \mathcal{M}$ **do**

11   **if** $\exists j \in \mathcal{N}, x_{ij}^t = 1$ **then**

12   **if** $j = J_i^{t-1}$ **then**

13         $\kappa_{ij}^t \leftarrow \kappa_{ij}^{t-1}$;

14   **else**

15         $\kappa_{ij}^t \leftarrow \kappa_{ij}^{t-1}(1 + \frac{s_{ij}^t}{B_i}) + \frac{\omega_{ij}^t}{B_i(\eta-1)}$;

16       $J_i^t \leftarrow j$;

17   **else**

18       $J_i^t \leftarrow 0$;

set $\mathcal{W}^t = \{w_1, \cdots, w_h, \cdots, w_{|\mathcal{W}^t|}\}$ with social welfare R_s^t, where winners are numbered in order of being selected. For any $w_h \in \mathcal{W}^t$, define set $\mathcal{D}_h$ such that if $w_g \in \mathcal{D}_h$, w_h is selected before w_g and $w_g \in \mathcal{W}^{t,*}$ but $w_g \notin \mathcal{W}^t$ because of w_h, i.e., w_h blocks pairs in $\mathcal{D}_h$ from being selected. Any winner $w_g \in \mathcal{D}_h$ associated with (i', j') is selected after w_h. By Lemma 1 and taking $\vartheta = \mathcal{O}(\ln L)$, $v_j^t - c_{i'j'}^t \leq (v_i^t - c_{ij}^t)\frac{\sum_k \alpha_{i',k}^t/\beta_{j',k}^t}{\sum_k \alpha_{i,k}^t/\beta_{j,k}^t} + \mathcal{O}(\ln L)$ with a probability of at least $1 - 1/L^{\mathcal{O}(1)}$. Summing over $w_g \in \mathcal{D}_h$ yields $\sum_{w_g} v_j^t - c_{i'j'}^t \leq \sum_{w_g}\sum_k \frac{\alpha_{i,k}^t}{\beta_{j,k}^t}\left[\frac{v_i^t - c_{ij}^t}{\sum_k \alpha_{i,k}^t/\beta_{j,k}^t} + \mathcal{O}(\ln L)\right] \leq \frac{MK(v_i^t - c_{ij}^t)}{\sum_k \alpha_{i,k}^t/\beta_{j,k}^t} + \mathcal{O}(\ln L)$ with a probability of at least $1 - 1/L^{\mathcal{O}(1)}$. Since $\sum_k \frac{\alpha_{i,k}^t}{\beta_{j,k}^t} \geq \frac{\alpha_{\min}}{\beta_{\max}} \sum_{w_g}(v_j^t - c_{i'j'}^t) \leq (v_i^t - c_{ij}^t)\frac{MK\beta_{\max}}{\alpha_{\min}} + \mathcal{O}(\ln L)$, we have

$$\begin{aligned} R_s^{t,*} &= \sum_{w_g \in \mathcal{W}^{t,*}} \left(v_j^t - c_{i'j'}^t\right) \leq \sum_{w_g \in \cup_{w_h \in \mathcal{W}^t}\mathcal{D}_h} \left(v_j^t - c_{i'j'}^t\right) \\ &\leq \sum_{w_h \in \mathcal{W}^t} \left(v_i^t - c_{ij}^t\right) \cdot MK\frac{\beta_{\max}}{\alpha_{\min}} + \mathcal{O}(\ln L). \end{aligned}$$

Hence, the theorem follows. □

IV. The Online Auction Mechanism Design

We develop OPTA, an online double auction mechanism to solve SWM. The major challenge is lack of future knowledge. *Competitive analysis* is a desired approach, where neither exact values nor distribution of future input is known in advance [38]. Moreover, the long-term budget constraint $(2c)$ affects the allocation decision over time, and also leads to the changes of social welfare in different slots. If a user's budget runs out, some previously feasible pairs become invalid immediately. In this case, it can only be served by the current server, and cannot be switched to other more cost-effective ones. This negatively impacts the social welfare. An ideal situation is that the *budget of each user is sufficient for all slots*, and the platform can select the best decisions from all feasible pairs.

Towards these goals, we design a competitive online algorithm shown in Alg. 3, whose competitive ratio [18] is defined as the upper-bound ratio of the offline optimal social welfare derived by solving (2) exactly to the social welfare achieved by Alg. 3. Specifically, we introduce a regulation factor κ_{ij}^t for each feasible pair $(i,j) \in \mathcal{E}^t$ to adjust its net benefit ω_{ij}^t in slot t. Since the residual budget of user i initialized to B_i decreases over time, we scale up κ_{ij}^t if the task of i is migrated to another server, such that its value starts at 0 (line 1), increases with the decrease of residual budget and reaches 1 when the budget is exhausted (line 6). For ease of presentation, we use variable $J_i^t \in \{0\} \cup \mathcal{N}$ to store which server that user i is allocated in t, where 0 means i loses the auction (line 18). The new scaled benefit $\hat{\omega}_{ij}^t$ will be fed as input to sPTA to determine the winners and prices (line 9). Note that $\hat{\omega}_{ij}^t = 0$ if $\omega_{ij}^t \geq s_{ij}^t$ and $\kappa_{ij}^{t-1} \geq 1$ to guarantee allocation feasibility and budget constraint (see proof of Lemma 2). In this way, the bid from a user with a smaller residual budget will be assigned a smaller scaled benefit, reducing its chance to win. For each winner, κ_{ij}^t is updated if there exists a migration (lines 14-15), and remains unchanged otherwise (lines 12-13).

A. Theoretical Analysis

Lemma 2. *Alg. 3 gives a feasible solution to SWM in (2).*

Proof. sPTA computes a feasible solution to single-round SWM under constraints $(2a)$, $(2b)$, $(2d)$. To show the long-term migration budget constraint $(2c)$ is satisfied, we first prove

$$\kappa_{ij}^t \geq \frac{1}{\eta - 1}\left(\eta^{\frac{\sum_{\tau=1}^t \sum_{(i,j)\in\mathcal{W}^\tau} s_{ij}^\tau}{B_i}} - 1\right), \forall t \in \mathcal{T} \quad (11)$$

by induction. It is clear that (11) holds for $t = 0$. Suppose it holds for $t-1$. By Alg. 3, $\kappa_{ij}^t = \kappa_{ij}^{t-1}(1 + \frac{s_{ij}^t}{B_i}) + \frac{\omega_{ij}^\tau}{B_i(\eta-1)} \geq \frac{1}{\eta-1}\left(\eta^{\frac{\sum_{\tau=1}^{t-1}\sum_{(i,j)\in\mathcal{W}^\tau} s_{ij}^\tau}{B_i}}\left(1 + \frac{s_{ij}^t}{B_i}\right) - 1\right)$. We only need to prove $1 + \frac{s_{ij}^t}{B_i} \geq \eta^{s_{ij}^t/B_i}$. It holds since $\frac{s_{ij}^t}{B_i} \leq \hat{\eta}$, using the fact that $\frac{\ln(1+x)}{x} \geq \frac{\ln(1+y)}{y}, \forall 0 \leq x \leq y \leq 1$. For any user i, let t' denote the first time such that $\sum_{\tau=1}^{t'}\sum_{(i,j)\in\mathcal{W}^\tau} s_{ij}^\tau \geq B_i$. By (11), $\kappa_{ij}^{t'} \geq 1$. From lines 5-6 in Alg. 3, (i,j) will never be selected and constraint $(2c)$ holds. Hence, Alg. 3 respects all constraints of SWM and produces a feasible solution. □

Theorem 5. *OPTA is a truthful, differentially private, IR and budget-balanced online double auction that gives a $\frac{\eta}{\eta-1}\frac{\beta_{max}}{\alpha_{min}}MK$-competitive solution to SWM in polynomial time.*

Proof. • **Competitiveness.** Define $\Delta R^t = R^t - R^{t-1}$, where R^t is the objective value of SWM yielded by Alg. 3 after t slots. For any slot $t \in \mathcal{T}$, we have

$$
\begin{aligned}
\Delta R^t &= \sum_{(i,j)\in\mathcal{E}^t} \omega_{ij}^t x_{ij}^t = \sum_{(i,j)\in\mathcal{W}^t} \left(\hat{\omega}_{ij}^t + s_{ij}^t \kappa_{ij}^{t-1}\right) \\
&= \sum_{(i,j)\in\mathcal{W}^t} \hat{\omega}_{ij}^t + \sum_{(i,j)\in\mathcal{W}^t} \left(B_i(\kappa_{ij}^t - \kappa_{ij}^{t-1}) - \frac{\omega_{ij}^t}{\eta-1} \right) \\
&\geq \frac{1}{\mu} R_s^{t,*} + \sum_{(i,j)\in\mathcal{W}^t} B_i(\kappa_{ij}^t - \kappa_{ij}^{t-1}) - \frac{\Delta R^t}{\eta-1},
\end{aligned}
$$

where μ is the approximation ratio of sPTA, $\sum_{(i,j)} \hat{\omega}_{ij}^t$ is the objective value of single-round SWM derived by sPTA, and $R_s^{t,*}$ is its optimal solution. Summing both sides over $t \in \mathcal{T}$, rearranging the terms and using the fact that $\kappa_{ij}^t \geq \kappa_{ij}^{t-1}$ yield $\frac{\eta}{\eta-1} R^T \geq \frac{\eta}{\eta-1} R^T - \sum_t \sum_{(i,j)} B_i(\kappa_{ij}^t - \kappa_{ij}^{t-1}) \geq \frac{1}{\mu} R^{T,*}$, where R^T is the objective value of SWM achieved by Alg. 3, and $R^{T,*}$ is its optimal solution. We have $\frac{R^{T,*}}{R^T} \leq \frac{\eta\mu}{\eta-1}$. By Theorem 4, the competitive ratio of Alg. 3 is $\frac{\eta}{\eta-1} \frac{\beta_{\max}}{\alpha_{\min}} \cdot MK$.

• **Time complexity.** In Alg. 1, while loop terminates after at most MN iterations. In each iteration, select a pair in $\mathcal{O}(MNK)$ steps and perform updating in $\mathcal{O}(K)$ steps. The complexity of Alg. 1 is $\mathcal{O}(M^2N^2K)$. For Alg. 2, the sorting needs $\mathcal{O}(MN\log(MN))$. The first for loop for charging users runs $|\mathcal{W}^t| \leq M$ iterations. In each iteration, it verifies at most MN pairs and traverses competitive price set in $\mathcal{O}(M^2N)$ steps. The second for loop for paying servers takes $\mathcal{O}(N)$ time. The complexity of Alg. 2 is $\mathcal{O}(MN(\log(MN) + M))$. In Alg. 3, the outer for loop runs T slots. In each slot, besides running Alg. 1 and Alg. 2 to determine the winners and prices, there are two inner for loops. The former calculates scaled benefit in $\mathcal{O}(MN)$ steps and the latter updates variables in $\mathcal{O}(MN)$ steps. Thus, the complexity of outer for loop is $\mathcal{O}(\max\{MNK, \log(MN) + M\}TMN)$. Overall, the complexity of Alg. 3 is $\mathcal{O}(\max\{MNK, \log(MN) + M\}TMN)$.

• **Differential privacy.** It requires the following lemma.

Lemma 3. *Let $\Psi^t(\cdot)$ be an (ϵ^t, δ^t)-differentially private mechanism. If $\Psi^{\mathcal{T}}(\cdot) = (\Psi^1(\cdot), \cdots, \Psi^T(\cdot))$, then by Theorem 3.16 in [26], $\Psi^{\mathcal{T}}(\cdot)$ is $(\sum_t \epsilon^t, \sum_t \delta^t)$-differentially private.*

According to Theorem 1, in each slot t, sPTA is $(\widetilde{\epsilon}, \delta)$-differentially private for user valuation and server cost. By Lemma 3, OPTA satisfies $(\widetilde{\epsilon}T, \delta T)$-differential privacy.

• **Truthfulness, IR, budget balance.** By Theorem 2, sPTA is γ-truthful for both agents with given scaled benefit, which is determined based on bids/asks in Alg. 3 and only known to the platform, making OPTA with single-round sPTA truthful under solution feasibility [38]. Similarly, IR and budget balance are also guaranteed by Theorem 2, Theorem 3 and Lemma 2. □

B. Implementation Considerations

We next discuss how OPTA can be implemented. We use a representative example, BOINC [39], which is a volunteering platform facilitating individual servers to contribute idle resources for compute-intensive applications. Consider a gaming app provider using BOINC to assign user requests to servers. Given the proximity and capacity constraints, OPTA can be applied to find a near-optimal solution that dynamically partitions and allocates edge resources on demand without any future information, e.g., market dynamics. Users and servers participate in the auction as buyers and sellers by submitting asks and bids in terms of $\langle$resource demands, valuation$\rangle$ and $\langle$available capacities, cost$\rangle$. The provider iteratively choose winners from a set of feasible pairs based on the computed selection probability, and charges/pays winners via critical-value pricing. The gaming service follows the auction outcome, including computing new game states and generating game videos to stream to users. Finally, providers update market model for the next slot and then repeat the auction process.

There are other practical concerns that are worth further investigation. First, agents are endowed with ability to infer their valuations/costs from historical prices. Second, verifiability of service offerings is needed to ensure quality of service. Third, latencies incurred by dynamic resource adjustment are expected to be reduced through the hotplug technique [18], [40], which dynamically changes the amount of CPU/RAM/disk in use to produce customized resource bundles upon user request.

V. Performance Evaluation

We envision an open market for MEC system consisting of $M = 150$ users and $N = 40$ servers as default. We simulate a $500m \times 500m$ area with 40 servers distributed by homogeneous Poisson Point Process [15], whose coverage radius is set to $40m$. The user trajectory is generated by a random walk process [41]: in any slot, each user determiners its location for the next slot by randomly choosing a direction and a velocity uniformly distributed in $[0, 2\pi]$ and $[0, 1.5]$ m/s. Assume the market allows for trading $K = 3$ types of resources, e.g., CPU, RAM and bandwidth. Each server asks for the supplied sources, which are extracted from Amazon EC2 M5 instance types [42], while the resource cost $c_{j,k}^t$ uniformly distributes in $[0.1, 1]$. Since there are no real user request data publicly released by cloud providers, we artificially generate some user requests in the simulations, each of which has a limited resource demand $\alpha_{i,k}^t$ and a valuation v_i^t randomly generated in $[0.1, 12]$. For online auction, we set the length of each slot to 15 minutes and the time horizon to 28 slots, which is commonly used in previous studies [16]. The migration budget B_i is randomly valued in $[5, 20]$ unless otherwise specified, and per-time migration cost s_i uniformly distributes in $[3, 5]$. Additionally, privacy budget ϵ is set to 1 as default, and residual probability δ is set to 0.1.

A. Performance of Single-Round Auction sPTA

We implement the sPTA and compare it with three benchmarks: (1) OPTimal allocation (OPT) [37]: the optimal solution is obtained by solving (8) using an efficient MIP solver [43]; (2) Greedy allocation with maximum Social Welfare (Greedy-SW) [11]: users are allocated to servers to pursue maximum social welfare via greedy heuristic; (3) Greedy allocation with maximum Allocated pair number (Greedy-A) [3]: a greedy algorithm is performed to allocate as many users as possible. We conduct all simulations for 200 times and plot the mean values and standard deviations.

Social Welfare. Fig. 3 and Fig. 4 show the performance comparison of normalized social welfare. We observe all social welfare increases with the number of users or servers. Given fixed resources, we can allocate them more efficiently if there

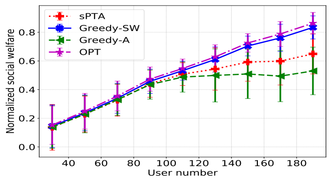

Fig. 3. Normalized social welfare vs. User number.

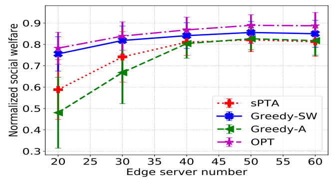

Fig. 4. Normalized social welfare vs. Edge server number.

TABLE I
AVERAGE COMPUTATION TIME FOR sPTA VS. BENCHMARKS

Algorithms	Average Computation Time
sPTA	1.7338 sec
Greedy-SW	1.5877 sec
Greedy-A	1.6861 sec
OPT	1546.146 sec

are more users, while with the increase of server number, more requests are satisfied. That's why we design a privacy-preserving double auction to attract more agents. Compared to Greedy-A, the social welfare that sPTA and Greedy-SW yield is much closer to the optimum since the curves of sPTA and Greedy-SW are much closer to that of OPT. sPTA and Greedy-SW jointly consider agent utility and demand-supply balance, which facilitates producing social-welfare maximizing sorting criterion. As for Greedy-A, failure to capture agent utility may lead to less efficient allocation. sPTA is inferior to Greedy-SW since the latter selects winners to pursue the maximal social welfare, while to enforce differential privacy, winners in sPTA are chosen by the selection probability.

User Satisfaction. We capture user satisfaction by the ratio of number of allocated requests and number of total requests. From Fig. 5, with increase of user number, user satisfaction decreases in all cases. The intuitive is that, given fixed edge resources, the available resources can't meet demands gradually as user number increases, leading to decrease of user satisfaction. Greedy-A greedily selects winners for maximum number of allocated pairs, yielding a satisfaction-maximizing sorting order. That's why user satisfaction in Greedy-A is the highest. Fig. 6 shows user satisfaction increases with server number. Given user set, larger server number means more resources to allocate. Thus, more pairs will be selected. As expected, Greedy-A achieves the highest user satisfaction.

Computation Efficiency. Table I shows all algorithms have different average computation time. As expected, OPT is very slow due to $\mathcal{NP}$-hardness of single-round SWM. Compared to two heuristics, sPTA is slightly slower because a random process is introduced to winner selection to enforce differential privacy. Taking Fig. 3 and Fig. 4 together, we find sPTA provides a good trade-off between performance and efficiency.

B. Performance of Online Auction OPTA

Offline/Online Ratio. We capture the competitiveness of OPTA by offline/online ratio defined as the ratio between offline optimal social welfare of (2) calculated by the MIP solver [43] and overall social welfare achieved by OPTA. The simulations are repeated for 20 times and the plots show the mean values and standard deviations. From Fig. 7, we observe that the ratio declines as server number increases. This is

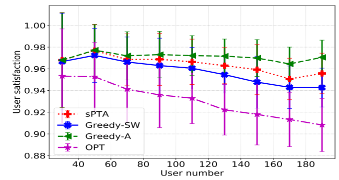

Fig. 5. User satisfaction vs. User number.

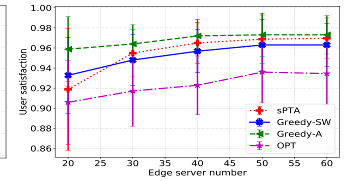

Fig. 6. User satisfaction vs. Edge server number.

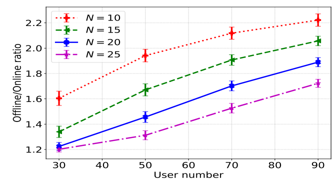

Fig. 7. Offline/online ratio vs. User number and edge server number.

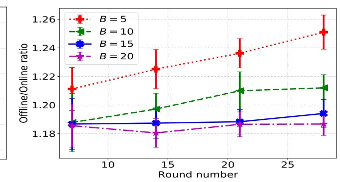

Fig. 8. Offline/online ratio vs. Round number and user budget.

reasonable since the platform can choose more cost-effective pairs. When more users compete for resources, the reduced selection probability degrades social welfare performance. That's why the ratio increases with user number. Fig. 8 illustrates the effect of round number and maximum user budget B on the ratio. For any given budget, the ratio always remains at similar levels. That is, OPTA can maintain performance stability regardless of the number of rounds it is applied. The budget increase makes users more likely to be switched to cheaper servers, increasing social welfare. Without future knowledge, OPTA is more sensitive to budget than offline solution. Thus, the ratio decreases as budget increases.

IR and Truthfulness. We only present the analysis results of valuation/cost, and that of demand/supply are similar. Assume users report truthful valuations. We conduct OPTA, and compare valuation and charge of each user. From Fig. 9, each user is charged no more than valuation. Hence, IR holds for users. We randomly pick a user and allow it to submit a bid whose valuation differs from the true one. Conduct OPTA on the untruthful dataset and present user utility and charge with varying reported valuations in Fig. 10. The utility and charge are zero when reporting a lower valuation, and remain unchanged when the reported valuation is no less than the true one. Thus, the user has no incentive to report a false valuation, i.e., OPTA is truthful for users. Suppose servers truthfully report costs. Implement OPTA, and compare server total cost and payment. Fig. 11 shows each server is paid no less than its cost. Thus, IR holds. We randomly pick a server and allow it to report false cost. We only consider one type of resource. Implement OPTA on this untruthful dataset, and illustrate its utility and payment with varying reported costs in Fig. 12. Similar to Fig. 10, OPTA is truthful for servers.

Privacy Preservation. We next illustrate the privacy performance of OPTA. Given a mechanism $\Psi^t(\cdot)$, let $(\boldsymbol{b}^t, \boldsymbol{a}^t)$ and $(\widetilde{\boldsymbol{b}}^t, \widetilde{\boldsymbol{a}}^t)$ be two input bid-ask profiles differing in only one element, $\Psi(\boldsymbol{b}^t, \boldsymbol{a}^t)$ and $\Psi(\widetilde{\boldsymbol{b}}^t, \widetilde{\boldsymbol{a}}^t)$ be the outcomes. By the definition of differential privacy, a good privacy-preserving scheme should keep the changes in outcomes as small as possible under a minor input change. To evaluate outcome changes,

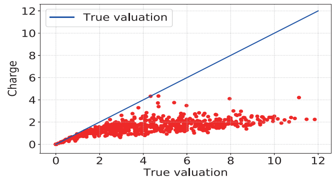

Fig. 9. IR for users.

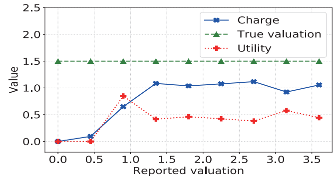

Fig. 10. User truthfulness.

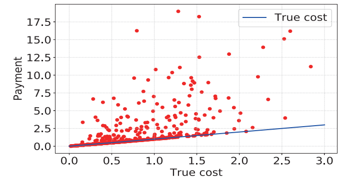

Fig. 11. IR for edge servers.

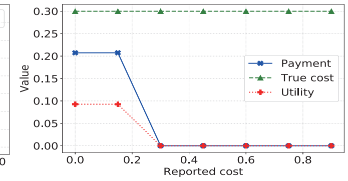

Fig. 12. Edge server truthfulness.

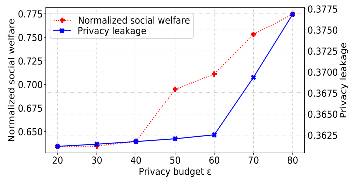

Fig. 13. Impact of privacy budget ϵ on social welfare and privacy leakage.

we're inspired by [29] to define the privacy leakage caused by OPTA as the Kullback-Leibler divergence of two outcome probability distributions based on $(\boldsymbol{b}^t, \boldsymbol{a}^t)$ and $(\widetilde{\boldsymbol{b}}^t, \widetilde{\boldsymbol{a}}^t)$, i.e.,

$$\pi^t = \sum_{\mathcal{W}^t \in \mathcal{S}^t} \Pr\left[\Psi(\boldsymbol{b}^t, \boldsymbol{a}^t) = \mathcal{W}^t\right] \ln\left(\frac{\Pr[\Psi(\boldsymbol{b}^t, \boldsymbol{a}^t) = \mathcal{W}^t]}{\Pr[\Psi(\widetilde{\boldsymbol{b}}^t, \widetilde{\boldsymbol{a}}^t) = \mathcal{W}^t]}\right).$$

Intuitively, the smaller π^t, the harder to distinguish the two inputs, i.e., a better privacy performance is achieved. Fig. 13 illustrates the social welfare and privacy leakage of OPTA with respect to privacy budget ϵ. We observe the social welfare increases with ϵ basically. This is due to the property of exponential based allocation: as ϵ increases, a pair with larger allocation efficiency is more likely to be selected, yielding higher social welfare. However, such increase in social welfare comes at a cost of large privacy leakage. The larger ϵ, the worse privacy performance, and thus the larger privacy leakage. Hence, there exists a trade-off between social welfare and privacy leakage. By carefully selecting ϵ, the system can achieve high social welfare with good differential privacy.

VI. Related Work

In the field of edge computing, various methods of improving its computing capacity have been proposed. A traditional approach is establishing facilities in proximity to users [7], but it may cause high construction costs and resource wastes. Instead, Zavodovski *et al.* [6] proposed open infrastructure for MEC, enabling individual servers to pool their resources locally for cooperative computation. There were similar works in the past [44]. For example, Vectordash is a commercial platform allowing GPU owners to rent out GPU instances [45]. The main question in this approach is how to efficiently allocate edge resources to users. Zavodovski *et al.* [9] present a truthful open market by bringing sharing economy to MEC. Given strategic behaviors of both agents, there has been growing interest in designing double auction schemes for MEC [9], [12]. Gao *et al.* [11] proposed a double auction-based allocation scheme, where users compete for resources from servers. These schemes focus on single-round auctions, making them simple and easy to implement. But in practice, the decisions must be made in real-time [15], because of: 1) continuously generated requests to be processed promptly; 2) market dynamics due to user mobility and random arrival/departure, time-varying bids/asks. Online auction [19] represents a natural solution for dynamic resource allocation. Chen *et al.* [16] designed an online auction market to incentivize cloudlets for edge emergency demand response. A closely related work [18] present an online auction framework for cloud resource provisioning. Our work differs from [18] in that: 1) they focus on VM trading under capacity constraints, while we consider a more complicated market under energy, proximity, capacity constraints; 2) they only address truthfulness, while we consider truthfulness, as well as system uncertainty and bid privacy.

Another line of prior work related to this paper is a series of differential privacy mechanisms [26]-[29] recently developed for privacy-preserving data analysis, which can be divided into two categories. The first is to add noise to the output [20], but which might not be feasible after adding noise. The second is to introduce randomization to the output via exponential mechanism [37]. There have been attempts in exponential based auction for differential privacy. Compared to pricing randomization [29] based on uniform pricing, randomizing the allocation [38] would be more applicable in combinatorial edge markets. Zhu *et al.* [35] designed a privacy-preserving spectrum market by using exponential method in winner selection. Lin *et al.* [36] randomized winner selection for privacy-preserving crowdsensing incentive mechanism. With respect to the differential privacy, our work is different from the existing exponential mechanisms in that: 1) they consider single-sided markets, while we address a double-side bidding market, ensuring privacy guarantee for both agents; 2) they only achieve single-round differential privacy, while we provide real-time privacy protection across multiple rounds.

VII. Conclusion

This paper presents an online privacy-preserving truthful double auction mechanism facilitating dynamic and collaborative resource allocation at the edge. We first design a differentially private truthful single-round auction, comprising of an exponential based allocation policy to select winners with efficiency and privacy guarantee, and a critical-value pricing policy to charge or pay winners. Building upon the single-round results, we propose a competitive online algorithm with a proven competitive ratio. Rigorous theoretical analysis and extensive simulations are performed, which verify the efficiency of the proposed method.

Acknowledgment

The research was support in part by RGC GRF grants under the contracts 16206417 and 16207818.

References

[1] M. Chiang and T. Zhang, "Fog and IoT: An Overview of Research Opportunities," *IEEE Internet of Things Journal*, vol. 3, no. 6, pp. 854-864, Dec. 2016.

[2] Y. Xiao and M. Krunz, "QoE and Power Efficiency Tradeoff for Fog Computing Networks with Fog Node Cooperation," in *Proc. of IEEE INFOCOM*, Apr. 2017.

[3] T. He, H. Khamfroush, S. Wang, T. L. Porta, and S. Stein, "It's Hard to Share: Joint Service Placement and Request Scheduling in Edge Clouds with Sharable and Non-sharable Resources," in *Proc. of IEEE ICDCS*, Jul. 2018.

[4] Y. Keshtkarjahromi, Y. Xing, H. Seferoglu, "Dynamic Heterogeneity-Aware Coded Cooperative Computation at the Edge," in *Proc. of IEEE ICNP*, Sept. 2018.

[5] Q. Liu and T. Han, "DARE: Dynamic Adaptive Mobile Augmented Reality with Edge Computing," in *Proc. of IEEE ICNP*, Sept. 2018.

[6] A. Zavodovski, N. Mohan, W. Wong and J. Kangasharju, "Open Infrastructure for Edge: A Distributed Ledger Outlook," in *The USENIX Workshop on Hot Topics in Edge Computing (HotEdge)*, Jul. 2019.

[7] Z. Xu, W. Liang, W. Xu, M. Jia, and S. Guo, "Efficient Algorithms for Capacitated Cloudlet Placements," *IEEE Transactions on Parallel and Distributed Systems*, vol. 27, no. 10, pp. 2866-2880, Oct. 2016.

[8] Amazon CloudFront. Available: https://aws.amazon.com/cn/cloudfront/.

[9] A. Zavodovski, S. Bayhan, N. Mohan, P. Zhou, W. Wong, and J. Kangasharju, "DeCloud: Truthful Decentralized Double Auction for Edge Clouds," in *Proc. of IEEE ICDCS*, Jul. 2019.

[10] D. T. Nguyen, L. B. Le, and V. K. Bhargava, "A Market-Based Framework for Multi-Resource Allocation in Fog Computing," *IEEE/ACM Transactions on Networking*, vol. 27, no. 3, pp. 1151-1164, Jun. 2019.

[11] G. Gao, M. Xiao, J. Wu, H. Huang, S. Wang, and G. Chen, "Auction-based VM Allocation for Deadline-Sensitive Tasks in Distributed Edge Cloud," *IEEE Transactions on Services Computing*, Mar. 2019.

[12] W. Sun, J. Liu, Y. Yue, and H. Zhang, "Double Auction-Based Resource Allocation for Mobile Edge Computing in Industrial Internet of Things," *IEEE Transactions on Industrial Informatics*, vol. 14, no. 10, pp. 4692-4701, Oct. 2018.

[13] G. V. Prasad, A. S. Prasad, and S. Rao, "A Combinatorial Auction Mechanism for Multiple Resource Procurement in Cloud Computing," *IEEE Transactions on Cloud Computing*, vol. 6, no. 4, pp. 904-914, Oct. 2018.

[14] M. M. Nejad, L. Mashayekhy, and D. Grosu, "A Family of Truthful Greedy Mechanisms for Dynamic Virtual Machine Provisioning and Allocation in Clouds," in *Proc. of IEEE CLOUD*, Jun. 2013.

[15] L. Chen, S. Zhou, and J. Xu, "Computation Peer Offloading for Energy-Constrained Mobile Edge Computing in Small-Cell Networks," *IEEE/ACM Transactions on Networking*, vol. 26, no. 4, pp. 1619-1632, Aug. 2018.

[16] S. Chen, L. Jiao, L. Wang, and F. Liu, "An Online Market Mechanism for Edge Emergency Demand Response via Cloudlet Control," in *Proc. of IEEE INFOCOM*, Apr. 2019.

[17] S. Gao, C. Courcoubetis, L. Duan, "Distributed Double Auctions for Large-scale Device-to-Device Resource Trading," in *Proc. of ACM MOBIHOC*, Jun. 2020.

[18] W. Shi, L. Zhang, C. Wu, Z. Li, and F. C. M. Lau, "An Online Auction Framework for Dynamic Resource Provisioning in Cloud Computing," *IEEE/ACM Transactions on Networking*, vol. 24, no. 4, pp. 2060-2073, Aug. 2016.

[19] M. H. Hajiesmaili, L. Deng, M. Chen, and Z. Li, "Incentivizing Device-to-Device Load Balancing for Cellular Networks: An Online Auction Design," *IEEE Journal on Selected Areas in Communications*, vol. 35, no. 2, pp. 265-279, Feb. 2017.

[20] Z. Wang, J. Hu, R. Lv, J. Wei, Q. Wang, D. Yang, and H. Qi, "Personalized Privacy-Preserving Task Allocation for Mobile Crowdsensing," *IEEE Transactions on Mobile Computing*, vol. 18, no. 6, pp. 1330-1341, Jun. 2019.

[21] Z. Chen, T. Ni, H. Zhong, S. Zhang, and J. Cui, "Differentially Private Double Spectrum Auction with Approximate Social Welfare Maximization," *IEEE Transactions on Information Forensics and Security*, vol. 14, no. 1, pp. 2805-2811, Nov. 2019.

[22] H. Jin, L. Su, B. Ding, K. Nahrstedt, and N. Borisov, "Enabling Privacy-Preserving Incentives for Mobile Crowd Sensing Systems," in *Proc. of IEEE ICDCS*, Jun. 2016.

[23] W. Jin, M. Li, L. Guo, L. Yang, "DPDA: A Differentially Private Double Auction Scheme for Mobile Crowd Sensing," in *Proc. of IEEE CNS*, May 2018.

[24] H. Zhao, J. Chi, Y. Tian, and G. J. Gordon, "Adversarial Privacy Preservation under Attribute Inference Attack," *arXiv preprint arXiv:1906.07902*, 2019.

[25] F. McSherry and K. Talwar, "Mechanism Design via Differential Privacy," in *Proc. of IEEE FOCS*, Oct. 2007.

[26] A. GuOPTA, K. Ligett, F. McSherry, A. Roth, and K. Talwar, "Differentially Private Combinatorial Optimization," in *Proc. of ACM-SIAM SODA*, Jan. 2010.

[27] C. Dwork and A. Roth, "The Algorithmic Foundations of Differential Privacy," *Theoretical Computer Science*, vol. 9, no. 3, pp. 211-407, 2014.

[28] M. Zhang, L. Yang, X. Gong, and J. Zhang, "Privacy-Preserving Crowdsensing: Privacy Valuation, Network Effect, and Profit Maximization," in *Proc. of IEEE GLOBECOM*, Dec. 2016.

[29] X. Jin and Y. Zhang, "Privacy-Preserving Crowdsourced Spectrum Sensing," *IEEE/ACM Transactions on Networking*, vol. 26, no. 3, pp. 1236-1249, Jun. 2018.

[30] Q. He, G. Cui, X. Zhang, F. Chen, S. Deng, H. Jin, Y. Li, and Y. Yang, "A Game-Theoretical Approach for User Allocation in Edge Computing Environment," *IEEE Transactions on Parallel and Distributed Systems*, vol. 31, no. 3, pp. 515-529, Mar. 2020.

[31] C. Wang, S. Zhang, Y. Chen, Z. Qian, J. Wu, and M. Xiao, "Joint Configuration Adaptation and Bandwidth Allocation for Edge-based Real-time Video Analytics," in *Proc. of IEEE INFOCOM*, Jul. 2020.

[32] E. Meskar and B. Liang, "Fair Multi-Resource Allocation with External Resource for Mobile Edge Computing," in *Proc. of IEEE INFOCOM WKSHPS*, Apr. 2018.

[33] M. Albert, V. Conitzer, and G. Lopomo, "Maximizing Revenue with Limited Correlation: The Cost of Ex-Post Incentive Compatibility," in *Proc. of AAAI*, Jan. 2015.

[34] H. Kellerer, U. Pferschy, and D. Pisinger, "Knapsack Problems," *Springer*, 2004.

[35] R. Zhu and K. G. Shin, "Differentially Private and Strategy-Proof Spectrum Auction with Approximate Revenue Maximization," in *Proc. of IEEE INFOCOM*, Apr. 2015.

[36] J. Lin, D. Yang, M. Li, J. Xu, and G. Xue, "Frameworks for Privacy-Preserving Mobile Crowdsensing Incentive Mechanisms," *IEEE Transactions on Mobile Computing*, vol. 17, no. 8, pp. 1851-1864, Aug. 2018.

[37] S. Chichin, Q. B. Vo, and R. Kowalczyk, "Towards Efficient and Truthful Market Mechanisms for Double-Sided Cloud Markets," *IEEE Transactions on Services Computing*, vol. 10, no. 1, pp. 37-51, Jan. 2017.

[38] R. Zhou, Z. Li, and C. Wu, "A Truthful Online Mechanism for Location-Aware Tasks in Mobile Crowd Sensing," *IEEE Transactions on Mobile Computing*, vol. 17, no. 8, pp. 1737-1749, Aug. 2018.

[39] D. P. Anderson, "BOINC: A System for Public-Resource Computing and Storage," in *Proc. of 5th IEEE/ACM International Workshop on Grid Computing*, Nov. 2001.

[40] KVM CPU Hotplug. Available: https://www.linux-kvm.org/page/CPUHotPlug.

[41] L. Wang, L. Jiao, J. Li, and M. Mühlhäuser, "Online Resource Allocation for Arbitrary User Mobility in Distributed Edge Clouds," in *Proc. of IEEE ICDCS*, Jun. 2017.

[42] Amazon EC2 Instance Types. Available: https://aws.amazon.com/ec2/instance-types/.

[43] CPLEX User's Manual: IBM ILOG CPLEX Optimization Studio. Available: https://www.ibm.com/support/knowledgecenter/SSSA5P_12.7.0/ilog.odms.studio.help/pdf/usrcplex.pdf

[44] K. Chard and K. Bubendorfer, "Co-Operative Resource Allocation: Building an Open Cloud Market Using Shared Infrastructure," *IEEE Transactions on Cloud Computing*, vol. 7, no. 1, pp. 183-195, Jan. 2019.

[45] Vectordash. Available: https://vectordash.com/.

[46] D. Li, Q. Yang, W. Yu, D. An, Y. Zhang, and W. Zhao, "Towards Differential Privacy-Based Online Double Auction for Smart Grid," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 971-986, Aug. 2019.