Online Stochastic Buy-Sell Mechanism for VNF Chains in the NFV Market

Xiaoxi Zhang, *Student Member, IEEE*, Zhiyi Huang, *Member, IEEE*, Chuan Wu, *Senior Member, IEEE*, Zongpeng Li, *Senior Member, IEEE*, and Francis C. M. Lau, *Senior Member, IEEE*

***Abstract*— With the recent advent of network functions virtualization (NFV), enterprises and businesses are looking into network service provisioning through the service chains of virtual network functions (VNFs), instead of relying on dedicated hardware middleboxes. Accompanying this trend, an NFV market is emerging, where NFV service providers create VNF instances, assemble VNF service chains, and sell them for the use of customers, using resources (computing, bandwidth) that they own or rent from other resource suppliers. Efficient service chain provisioning and pricing mechanisms are still missing, to charge assembled service chains according to demand and the supply of resources at any time. We propose an online stochastic auction mechanism for on-demand service chain provisioning and pricing at an NFV provider. Our auction takes in buy bids for service chains from multiple customers and sell bids from various resource suppliers to supplement the NFV provider's geo-distributed resource pool, with resource occupation/contribution durations. We extend online primal-dual optimization framework for handling both buyers and sellers, with a new competitive analysis. The online mechanism maximizes the expected social welfare of the NFV ecosystem (the NFV provider, customers and resource suppliers) with a good competitive ratio as compared with the expected offline optimal social welfare, while guaranteeing truthfulness in bidding, individual rationality for both buyers and sellers, and polynomial time for computation. We evaluate our mechanism through trace-driven simulation studies, and demonstrate a close-to-offline-optimal performance in expected social welfare under realistic settings.**

***Index Terms*— Auction mechanism design, online algorithms, network functions virtualization.**

I. Introduction

Network functions virtualization (NFV) is a recent paradigm for running network functions, *e.g.*, intrusion detection systems (IDSs), firewalls, proxies and WAN optimizers, as software on virtual machine (VM) instances on industry standard servers. Today, a new form of network function typically emerges as a one-off solution to a specific need, implemented using dedicated hardware middleboxes and "patched" into the existing network infrastructure through manual installation, which are costly and difficult to maintain. NFV enables significant reduction in deployment cost and management overhead [1] through dynamic, automatic deployment of *virtualized network functions* (VNFs) [2]. VNFs are typically connected into *service chains* (an ordered set of VNFs) to provide *network service* [2].

More and more enterprises and businesses are resorting to VNFs for provisioning network services. Accompanying this trend, an NFV market is emerging, where dedicated NFV service providers create VNF instances, build VNF service chains, and offer them to customers on demand. The service chains are assembled using resources (computing and bandwidth) that an NFV provider owns or rents from other resource suppliers (*e.g.*, cloud providers). Customers can browse the VNFs available at an NFV provider, and specify the VNFs to compose their service chains. There are yet many open issues to resolve in order to enable an efficient NFV market, among which resource allocation optimization and pricing are the top priorities. Efficient service chain provisioning and pricing mechanisms are still missing.

We study an NFV market where the NFV provider owns a geo-distributed pool of resources, allocates resources to assemble service chains upon demand, and may purchase extra resources from resource suppliers if deemed appropriate. The customers bid for service chains while the resource suppliers sell their available resources to the NFV provider through sell bids. We seek an efficient online mechanism for the NFV provider to carry out dynamic service chain provisioning and pricing on the go. The design objective is to maximize the social welfare of all parties while guaranteeing truthfulness in bidding, individual rationality and polynomial time for online computation. Instead of resorting to a double auction, we novelly design the online auction such that both buy and sell bids can be similarly handled in a consistent fashion.

First, we characterize the interaction among the NFV provider, VNF service chain customers, and resource suppliers in an online model and design an efficient online mechanism to tackle the key issues of resource provisioning and pricing schemes. Such a practical three-party NFV market paradigm is new in the literature. Our auction guarantees truthfulness of both buyers and sellers, and obtains near-optimal expected social welfare with a competitive ratio of $1 - O(\epsilon)$, where ϵ can be arbitrarily close to 0. The key is to convert the social welfare maximization problem in the online stochastic model to a deterministic fractional program, exploiting the properties of the bid arrival process. The fractional program provides an upper-bound of the offline optimal social welfare in expectation, and facilitates our online algorithm design based on a primal-dual framework, by removing the time dimension indices of the dual prices.

Second, we extend the primal-dual framework to handle both buyers and sellers (with positive and negative values

Manuscript received April 30, 2016; revised September 30, 2016; accepted November 28, 2016. Date of publication January 31, 2017; date of current version March 31, 2017.

X. Zhang, Z. Huang, C. Wu, and F. C. M. Lau are with the Department of Computer Science, The Hong Kong University, China (e-mail: xxzhang2@cs.hku.hk; zhiyi@cs.hku.hk; cwu@cs.hku.hk; fcmlau@cs.hku.hk).

Z. Li is with the Department of Computer Science, University of Calgary, Calgary, AB T2N 1N4, Canada (e-mail: zongpeng@ucalgary.ca).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSAC.2017.2659259

of the input parameters). Technically, our model involves negative inputs (sell bids) in the underlying resource allocation problem (an integer linear program (ILP)) of the auction. Exploiting the auxiliary fractional linear program where the resource occupation/contribution durations are removed, we use a flat price independent of buy/sell price of each bidder for each type of resource, to handle both buy and sell bids for guaranteeing truthfulness. We apply online price learning as a subroutine that uses previous bids to learn the set of prices that determine future resource allocation and payments/rewards, and periodically update the prices as more bids are revealed, for better approaching the optimal social welfare.

Third, we carefully design the online price learning method to handle the departure of a bidder. To our knowledge, previous work on learning-based online primal dual algorithms (*e.g.*, [3]) do not deal with bid departures, which is a crucial part of our model – both service chain customers and resource contributors cease their resource occupation or supply after a specified period of time. By estimating the average resource demand/supply of a buyer/seller according to the valid duration of each bid, we learn future resource availability from the arrived bids and update average prices of resources for future time slots. Our analysis of competitive ratio novelly takes into account the time factor by connecting the expected resource consumption/contribution of each bid *at each time slot* with the average resource consumption/contribution estimation *over the entire time span*, as used in our price learning subroutine.

The rest of the paper is organized as follows. We discuss related work in Sec. II and define the problem model in Sec. III. Sec. IV presents our online auction design with theoretical analysis. Simulations are presented in Sec. V. Sec. VI concludes the paper.

II. Related Work

Early efforts on NFV focused on bridging the performance gap between specialized hardware and virtualized network functions running on VMs [4], [5], as well as designing management platforms for VNF deployment, traffic steering, and flow state migration across multiple instances of a VNF [6], [7]. These systems significantly facilitate the dynamic deployment and scaling of VNFs. A few recent studies investigate optimal placement and scaling of VNF instances and traffic routing in service chains for cost minimization. VNF-P [8] presents a one-time optimization model for VNF placement, considering hybrid deployment where part of the network service is provided by dedicated hardware and part by VNF instances, and designs a heuristic algorithm. Mehraghdam *et al.* [9] model a mixed integer quadratically constrained program (MIQCP) to pursue different optimization goals in VNF placement, without giving solution algorithms. These work mostly deal with one-off placement of the NFV service chains, ignoring the dynamic nature of an NFV system.

In contrast, we take a few important steps further: (1) we design an efficient online algorithm for dynamic service chain placement and pricing; (2) we study the rather new NFV market [10]; to the best of our knowledge, there is no existing study on mechanism design for the NFV market; (3) we consider dynamic resource pooling with both buy and sell of resources in the NFV market, which renders significant challenges in online mechanism design.

There has been a sustainable body of research on cloud auctions in both the offline and online settings [11], [12] [13]. Unlike the proposed online auction in this work, they cannot handle request departures or only involve buy bids for cloud resources. A very recent work [14] proposes an efficient VNF chain auction which guarantees truthfulness and achieves near-optimal social welfare in polynomial time. However, the problem they solved is in an offline setting and the technique can not be readily applied to online VNF auctions. An online version of double auction, where bids arrive and expire at different times, has been considered in the studies of *continuous double auction* (CDA) [15]. These studies aim to maximize the profit of the auctioneer or the number of items sold, but unfortunately provide no truthfulness guarantee. Wurman *et al.* [16] adopt and extend the monotonicity-based truthful characterization based on the work of Bredin and Parkes *et al.* [17] in developing their online truthful double auctions. Their methodology lacks economic efficiency (*e.g.*, social welfare maximization) guarantee, and can not be readily adapted to our NFV market. The reason is that the items sold or bought by agents in their problem are identical and cannot be released or reused after the agents depart. In contrast, buyers and sellers in our models have heterogenous demands/supplies of resource combinations, which could be reused/released after the occupation/contribution deadlines. Hajiaghayi *et al.* [18] design truthful online double auctions based on a bipartite matching algorithm. In their work, the auctioneer makes immediate and invokable decisions once an agent arrives, which is also one of the critical features that our mechanism provides. However, economic efficiency is not guaranteed and a strong assumption is made that all the sellers have fixed asking prices.

There also exist some theoretical work on online stochastic auctions. Wang *et al.* [19] design a learning-based algorithm for online resource allocation problems where the types (*e.g.*, bidding price, resource demands) of customers are i.i.d., without considering additional resource supplies. Agrawal *et al.* [3] design similar algorithms for solving online LPs where the columns of coefficient matrix of an LP arrive in a random order. Moreover, existing online stochastic algorithms typically assume that the total number of inputs is known in advance, while our mechanism works even when such knowledge (total number of bids) is estimated but not accurate.

III. Problem Model

A. The NFV Market

We consider an NFV market among three parties: (i) *Customers*, each requiring a service chain to process its data flow. (ii) *Resource suppliers* who can provide computing resources (CPU, RAM, disk) in the form of virtual machines (VM) and/or data transfer bandwidth between geo-dispersed VMs at certain times. (iii) An NFV service provider who owns geo-distributed computing resources and bandwidth in-between, installs VNFs, deploys and sells service chains according

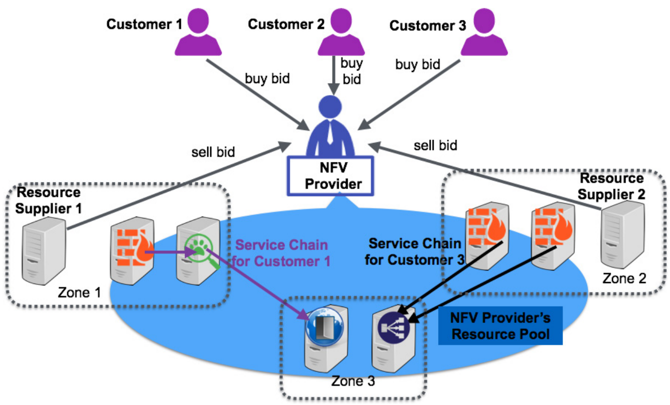

Fig. 1. NFV market.

to customer demand. It may also purchase resources from resource suppliers to supplement its resource pool. An illustration of the NFV market is given in Fig. 1.

The geographic span of the NFV provider's computing resources can be divided into zones in set $\mathcal{S}$, where a zone may correspond to one or multiple servers, or a datacenter. K types of VNFs in set $\mathcal{K}$ can be provided in each zone ($K = |\mathcal{K}|$). Each VNF can be deployed on multiple VMs and each of these VMs is referred to as an *instance of this VNF*. The NFV provider can maximally provision C_{ks} instances of type-k VNF in zone s, $\forall s \in \mathcal{S}, k \in \mathcal{K}$, using resources it owns in the respective zone. The NFV provider owns an upload bandwidth capacity of Θ_s^{out} from zone s to other zones, a download bandwidth capacity of Θ_s^{in} from other zones to zone s, and a bandwidth capacity of $L_{ss'}$ on the link from zone s to zone s', $\forall s \in \mathcal{S}, s' \in \mathcal{S}/\{s\}$. We assume that the bandwidth interconnecting computing resources in the same zone is always abundant. The entire lifespan of the system is $\mathcal{T} = [1, T]$.

Each customer i demands one service chain consisting of a sequence of VNFs in $\mathcal{K}$, with one or multiple instances needed for each VNF, *e.g.*, 2 firewall instances and 1 load balancer instance in the chain "Firewall→Load Balancer" for customer 3 in Fig. 1. Let d_{ik} denote the number of instances of VNF k that customer i requires in its service chain. $d_{ik} = 0$ for all VNFs (k's) which do not belong to the chain. Different VNFs and/or different instances of the same VNF can be deployed in different zones. Let $\hat{\pi}_i^{kk'}$ denote the bandwidth a customer needs for flow transfer between two instances of two consecutive VNFs k and k' in its chain.

A customer may have multiple options for geographic deployment of their VNF instances, *e.g.*, to deploy all the 2 firewalls and 1 load balancer in the same zone, or 2 firewalls in one zone and 1 load balancer in another. Let Γ_i denote the set of deployment options of customer i. We define a demand matrix $\mathbf{d}_{i\gamma} \in \mathbb{Z}^{K\times S}$ for each option $\gamma \in \Gamma_i$, where each element $d_{i\gamma}^{ks}$ is the number of instances of VNF k that customer i wants in zone s in this option, with $\sum_{s\in\mathcal{S}} d_{i\gamma}^{ks} = d_{ik}$. The service chain configuration (instance numbers and bandwidth in-between) and geographic deployment can be computed by the customer based on estimated flow arrival rates, processing capacity of each VNF instance, and any end-to-end flow delay requirement, which has been studied in [20], [21], and [9] and is orthogonal to the focus of this paper.

A customer submits its service chain demand as a buy bid. Let $v_{i\gamma}$ be bidder i's true valuation of option γ. The true values for different options indicate preferences of the customer for its options. $b_{i\gamma}$ is the corresponding bidding price customer i submits. Let τ_i be the usage duration of the service chain. Therefore, customer i's bid can hence be expressed as follows:

$$B_i = (\{b_{i\gamma}\}_{\gamma\in\Gamma_i}, \{d_{i\gamma}^{ks}\}_{k\in\mathcal{K},s\in\mathcal{S},\gamma\in\Gamma_i}, \{\pi_{i\gamma}^{ss'}\}_{s\neq s'\in\mathcal{S},\gamma\in\Gamma_i}, \tau_i). \tag{1}$$

Here, $\pi_{i\gamma}^{ss'}$ is the aggregate bandwidth demand of bidder i from s to s' in option γ, which can be readily computed according to a customer's flow demand $\hat{\pi}_i^{kk'}$ between VNFs and instance deployment strategies $d_{i\gamma}^{ks}$'s, *i.e.*, $\pi_{i\gamma}^{ss'} = \sum_{k\in\mathcal{K}}\sum_{k'\in\mathcal{K}/\{k\}} \hat{\pi}_i^{kk'} \max(d_{i\gamma}^{ks}, d_{i\gamma}^{k's'})$.

A resource supplier sells its available resources to the NFV provider as VMs and bandwidth for VNF deployment. The resources owned by a supplier may be located in one zone or distributed across multiple zones. In general, each resource supplier i has 1 or more sell options in Γ_i. Similarly, the same B_i in (1) is used to express a sell bid from resource supplier i. Different from a buy bid where all quantities are non-negative, in a sell bid, we have $b_{i\gamma} \le 0$, $d_{i\gamma}^{ks} \le 0$, and $\pi_{i\gamma}^{ss'} \le 0$, $\forall k \in \mathcal{K}, s, s' \in \mathcal{S}, \gamma \in \Gamma_i$. It should be noted that $-b_{i\gamma} \ge 0$ denotes the asking price for option γ ($-v_{i\gamma}$ is the true valuation), $-d_{i\gamma}^{ks} \ge 0$ denotes the number of VM instances that i can provide in zone s to run VNF k, and $-\pi_{i\gamma}^{ss'} \ge 0$ denotes the bandwidth that i can provide from zone s to s'. τ_i is the duration during which i's resources can be used.

The NFV provider serves as an auctioneer to determine the winning bids and the service chain deployment. There are in total I bids arriving at different times ($I = |\mathcal{I}|$). Bid i arrives at t_i, and can be either a sell bid or a buy bid with any probability. Upon receiving a buy (sell) bid, the NFV provider makes the following decisions: (i) whether to accept the bid and if so, which buy (sell) option to accept, as indicated by binary variable $x_{i\gamma}$, with 1 indicating the acceptance of option γ of bid i, and 0 otherwise; (ii) the payment $\hat{p}_i$ to collect from the bidder i, where $\hat{p}_i > 0$ if the bidder is a customer, and $\hat{p}_i < 0$ if the bidder is a resource supplier ($-\hat{p}_i > 0$ is the reward from the NFV provider to the supplier).

B. Online Stochastic Model and Mechanism Design Goals

We model the arrival process of bids during $\mathcal{T} = [1, T]$ as a Poisson process with an arrival rate λ [22]. There are two key properties of a Poisson process [23]: (i) the total number of arrivals in $\mathcal{T}$, I, is a random variable following the Poisson distribution with an expectation of λT; (ii) the arrival time of each bid among I can be uniformly and independently mapped to $[1, T]$.[1] Exploiting (ii) above, we assume that the arrival time of bid i, *i.e.*, t_i, is uniformly and independently drawn within $[1, T]$, and the bids are indexed according to their order of arrival in any fixed realization of the arrival process. Each bid vector B_i, as defined in (1), is drawn independently from

[1]The algorithm design relies on the expectation of I and property (ii) but does not restrict that I follows Poisson distribution.

a set of bid types, $\mathcal{D}$, following an unknown distribution, *i.e.*, B_i's are i.i.d.[2]

We target the following properties in our NFV market mechanism design. (i) *Truthfulness in bidding price*: For any bidder, declaring its true valuation of the service chain to buy or resources to sell in its bid always maximizes its utility, regardless of any other bids. The utility function of bidder i is $u_i = \begin{cases} v_{i\gamma} - \hat{p}_i, & \text{if } \exists \gamma \in \Gamma_i, x_{i\gamma} = 1 \\ 0, & \text{if } x_{i\gamma} = 0, \forall \gamma \in \Gamma_i \end{cases}$ regardless of whether it is a sell or buy bid. (ii) *Computational efficiency*: Polynomial-time algorithms for resource allocation and payment calculation are needed for the auction to run efficiently in an online fashion. (iii) *Individual rationality*: Each bidder obtains a non-negative utility by participating in the auction. (iv) *Social welfare maximization in expectation*: Given a realization of the bid arrival process, the social welfare is the sum of aggregate utility of all customers and resource suppliers, $\sum_{i\in I}\sum_{\gamma\in\Gamma_i}(v_{i\gamma} - \hat{p}_i)x_{i\gamma}$, and the overall profit of the NFV provider. Since payments in the bidders' utilities and the NFV provider's profit cancel each other, the social welfare is $\sum_{i\in I}\sum_{\gamma\in\Gamma_i} v_{i\gamma}x_{i\gamma}$ (equals $\sum_{i\in I}\sum_{\gamma\in\Gamma_i} b_{i\gamma}x_{i\gamma}$ under truthful bidding). We aim to upper-bound the ratio of the *expected* social welfare achieved with our mechanism (over different realizations of the bid arrival process) to the expected offline optimal social welfare.

Under a fixed realization of the bid arrival process, the offline social welfare maximization and winner determination problem can be formulated as follows (also assuming truthful bidding is guaranteed). The expected offline optimal social welfare can be computed based on the objective value of this problem over different realizations of the bid arrival process.

$$P: \text{maximize: } \sum_{i\in I}\sum_{\gamma\in\Gamma_i} b_{i\gamma}x_{i\gamma} \tag{2}$$

subject to :

$$\sum_{\gamma\in\Gamma_i} x_{i\gamma} \le 1, \quad \forall i \in I \tag{2a}$$

$$\sum_{\substack{i\in I:\\ t_i\le t<t_i+\tau_i}}\sum_{\gamma\in\Gamma_i} d_{i\gamma}^{ks}x_{i\gamma} \le C_{ks}, \quad \forall k\in\mathcal{K},\ s\in\mathcal{S}, t\in\mathcal{T} \tag{2b}$$

$$\sum_{\substack{i\in I:\\ t_i\le t<t_i+\tau_i}}\sum_{\gamma\in\Gamma_i} \pi_{i\gamma}^{ss'}x_{i\gamma} \le L_{ss'}, \quad \forall s\in\mathcal{S},\ s'\in\mathcal{S}/\{s\},\ t\in\mathcal{T} \tag{2c}$$

$$\sum_{\substack{i\in I:\\ t_i\le t<t_i+\tau_i}}\sum_{\gamma\in\Gamma_i}\sum_{s'\in\mathcal{S}/\{s\}} \pi_{i\gamma}^{s's}x_{i\gamma} \le \Theta_s^{in}, \quad \forall s\in\mathcal{S},\ t\in\mathcal{T} \tag{2d}$$

$$\sum_{\substack{i\in I:\\ t_i\le t<t_i+\tau_i}}\sum_{\gamma\in\Gamma_i}\sum_{s'\in\mathcal{S}/\{s\}} \pi_{i\gamma}^{ss'}x_{i\gamma} \le \Theta_s^{out}, \quad \forall s\in\mathcal{S},\ t\in\mathcal{T} \tag{2e}$$

$$x_{i\gamma} \in \{0,1\}, \quad \forall i\in I,\ \gamma\in\Gamma_i \tag{2f}$$

Constraint (2a) indicates that at most one option is adopted for each bid i. (2b) guarantees that the overall number of instances of VNF k in any zone s at any time t, used in provisioning service chains (which are running at t) in all accepted bids, does not go beyond the capacity limit. (2c) ensures that the total traffic on each link (s, s') from all service chains using this link at t does not exceed the respective bandwidth capacity. Similarly, (2d) and (2e) guarantee that the total inbound and outbound bandwidth usage at each zone due to the deployed service chains at each time is no larger than the available download and upload bandwidth, respectively.

Note that in the left-hand-side summation in constraints (2b)-(2e), $d_{i\gamma}^{ks}$ and $\pi_{i\gamma}^{s's}$ are negative for an accepted option of a sell bid i. This is equivalent to adding the respective provisioned resources to the right-hand-side resource capacity in those constraints, *i.e.*, supplementing the resource pool of the NFV provider to serve buy bids.

The offline problem is established assuming complete knowledge of the system over its entire lifespan. In a dynamic system, with the arrival of bids, the variables and constraints emerge gradually. For example, on the arrival of bid i, a new constraint (2a) appears for this bid, and a set of new variables $x_{i\gamma}, \forall\gamma \in \Gamma_i$, are added to constraints (2b) – (2e). In the following, we design an online auction mechanism for the NFV provider to determine immediately whether to accept a bid i and which option to serve, as well as the bidder's payment if accepted.

IV. Online Auction Design

A. Primal-Dual Framework

We start with simplifying (2). Observe that constraints (2b) to (2e) are similar in the sense that they make sure the overall consumption of a resource (a VNF type, the bandwidth of a link, or the upload/download bandwidth in a zone) does not exceed the respective capacity at any time. We use m to index the generalized resource type and the total number of the generalized resource types (including all the resources) is equal to $(K+2)S + |\mathbb{E}|$, where $\mathbb{E} = \{(s,s'), \forall s \in \mathcal{S}, s' \in \mathcal{S}/\{s\}\}$ denotes the set of links connecting pairs of zones. Let $\mathcal{M}$ denote the set of generalized resource types and let $\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3, \mathcal{M}_4$ respectively denote the set of VNF types, the set of upload bandwidth at all the zones, the set of download bandwidth at all the zones, and the set of bandwidth of the links. We use C_m to denote the capacity of a generalized resource type $m \in \mathcal{M}$. Then (2) is equivalent to the following:

$$P: \text{maximize: } \sum_{i\in I}\sum_{\gamma\in\Gamma_i} b_{i\gamma}x_{i\gamma} \tag{3}$$

subject to :

$$\sum_{\gamma\in\Gamma_i} x_{i\gamma} \le 1, \quad \forall i \in I \tag{3a}$$

$$\sum_{\substack{i\in I:\\ t_i\le t<t_i+\tau_i}}\sum_{\gamma\in\Gamma_i} c_{i\gamma}^{m}x_{i\gamma} \le C_m, \quad \forall m\in\mathcal{M},\ t\in\mathcal{T} \tag{3b}$$

$$x_{i\gamma} \in \{0,1\}, \quad \forall\gamma\in\Gamma_i,\ i\in I \tag{3c}$$

[2]Assuming the overall system span T is much longer than the duration τ_i in all bid types in $\mathcal{D}$, we can safely assume that τ_i of different bids is drawn from the same range regardless of the bid arrival time, since the probability of $t_i + \tau_i > T$ is very small and not considering those extreme bids barely affects the overall social welfare in expectation.

where:

$$c_{i\gamma}^m = \begin{cases} d_{i\gamma}^{ks}, & \text{if } C_{ks} \to C_m, \quad \forall m \in \mathcal{M}_1 \\ \sum_{s' \in \mathcal{S}/\{s\}} \pi_{i\gamma}^{s's}, & \text{if } \Theta_s^{in} \to C_m, \quad \forall m \in \mathcal{M}_2 \\ \sum_{s' \mathcal{S}/\{s\}} \pi_{i\gamma}^{ss'}, & \text{if } \Theta_s^{out} \to C_m, \quad \forall m \in \mathcal{M}_3 \\ \pi_{i\gamma}^{ss'}, & \text{if } L_{ss'} \to C_m, \quad \forall m \in \mathcal{M}_4 \end{cases}$$

and $a \to b$ denotes the mapping from an original resource capacity a to the generalized resource capacity b. For example, $C_{ks} \to C_m$ means that C_{ks} is represented by C_m where m denotes the corresponding resource (VNF k in zone s) in set $\mathcal{M}_1$.

Let u_i and p_{mt} be the dual variables associated with 3(a) and 3(b), respectively. The dual program of (3) is (relaxing 3(c) to $x_{i\gamma} \geq 0$, noting that $x_{i\gamma} \leq 1$ is redundant due to 3(a)):

$$D: \text{minimize: } \sum_{t \in \mathcal{T}} \sum_{m \in \mathcal{M}} C_m p_{mt} + \sum_{i \in I} u_i \tag{4}$$

subject to :

$$u_i \geq b_{i\gamma} - \sum_{m \in \mathcal{M}} \sum_{t \in [t_i, t_i + \tau_i)} p_{mt} c_{i\gamma}^m, \quad \forall i \in I, \ \gamma \in \Gamma_i \tag{4a}$$

$$p_{mt} \geq 0, \quad \forall m \in \mathcal{M}, \ t \in \mathcal{T} \tag{4b}$$

$$u_i \geq 0, \quad \forall i \in I \tag{4c}$$

Our core idea of the online winner determination algorithm design, for social welfare maxmization, is as follows. We resort to the KKT conditions [24] of the offline primal and dual problems in (3) and (4) to maintain a feasible primal solution as well as a feasible dual solution online, which pursue the offline optimal solution. On the arrival of a sell or buy bid i, a new dual variable $u_i \geq 0$ appears, subject to constraints (4*a*a), that is, $u_i \geq b_{i\gamma} - \sum_{m \in \mathcal{M}} \sum_{t \in [t_i, t_i+\tau_i)} c_{i\gamma}^m p_{mt}$ for all $\gamma \in \Gamma_i$. Let $\tilde{\mathbf{p}}$ denote the offline optimal solution of dual variables $p_{mt}, \forall m \in \mathcal{M}, t \in \mathcal{T}$. The KKT conditions indicate that in the offline primal and dual solutions to (3) and (4), $x_{i\gamma}$ must be zero unless constraint (4a) is tight for option γ. Thus, temporarily assuming that we know $\tilde{p}_{mt}$, we can assign each u_i to be

$$u_i = \max\big\{0, \max_{\gamma \in \Gamma_i} \{b_{i\gamma} - \sum_{m \in \mathcal{M}} \sum_{t \in [t_i, t_i+\tau_i)} c_{i\gamma}^m \tilde{p}_{mt}\}\big\} \tag{5}$$

(the maximal of 0 and the right hand side (RHS) of constraints (4a)) and letting $\gamma_i = \text{argmax}_{\gamma \in \Gamma_i} \big\{b_{i\gamma} - \sum_{m \in \mathcal{M}} \sum_{t \in [t_i, t_i+\tau_i)} c_{i\gamma}^m \tilde{p}_{mt}\big\}$, making constraint (4a) tight for γ_i. Then, we let $x_{i\gamma_i} = 1$ if and only if $u_i > 0$, satisfying the *necessary condition* of the KKT conditions. The rationale for accepting a bid in this way is as follows: If we interpret $\tilde{p}_{mt}$ as the marginal price (payment) per unit of resource m at time t, then the second term on the RHS of (4a) becomes the total payment that bid i should pay for the requested service chain if it is a buy bid, or the inverse of the second term is the total reward that the bidder should receive if it is a sell bid. So the RHS of (4a) is the utility of bid i, assuming truthful bidding: valuation minus payment if it is a buy bid or reward minus valuation if it is a sell bid. Therefore, the above method effectively accepts bid i achieving positive utility, in the best option that maximizes its utility, and u_i is bid i's utility. In this way, we target utility maximization for each bidder (no matter whether it is a seller or buyer), which leads to truthfulness and social welfare maximization.

However, in the online problem, we do not know the offline optimal dual solution $\tilde{\mathbf{p}}$ of (4). Exploiting the stochastic bid arrival model, we hope to get an approximately optimal dual solution of the offline problem *in expectation* from the first $\epsilon \in (0, 1)$ fraction of bids, and to successively refine our dual solution as more bids arrive. We will show that our approximately optimal dual solution is sufficient for coordinating the primal winner determination online to approximately maximize social welfare.

B. An Online Auction With Stochastic Input

Upon the arrival of a new bid i, deciding whether to serve the bid and which option to use is equivalent to choosing a feasible assignment for the new primal variables $x_{i\gamma}$ of (3*a*). If the NFV provider decides to serve bid i in one of its options γ_i, then let $x_{i\gamma_i} = 1$; otherwise, $x_{i\gamma}$ will be zero for all options $\gamma \in \Gamma_i$. Besides winner determination, we should also determine how much to charge for each winning buy bid and how much to pay for each accepted sell bid.

Expected Offline Optimization Problem: The offline problem in (3) is defined with respect to a set of bids that have been realized from the underlying stochastic arrival process. Next, we describe the expected offline primal program (over the random realizations of bids) in (6) and the corresponding dual program in (7). We will refer to them as the *distribution instance programs*. The optimal objective value of the primal problem in (6) serves as an upper-bound of the expected social welfare of the offline optimal problem (3) in our competitive analysis (see Lemma 7) and guides us to design the online auction mechanism.

Let δ_j denote the probability that bid type j is drawn from distribution $\mathcal{D}$. Recall that the expected number of bids is λT. Thus, the expected number of times that type j is chosen among the realized bids is $\lambda T \delta_j$. In the distributed instance programs, we let $x_{j\gamma}$ denote the probability that a bid, *conditioned on* its type being j, is served in option γ, over random realizations of bids. Then, the contribution of type-j bids to the expected social welfare is $\lambda T \delta_j \sum_{\gamma \in \Gamma_i} b_{j\gamma} x_{j\gamma}$. With slight abuse of notation, we use j to denote a bid of type j instead of bid j, in $b_{j\gamma}$, $x_{j\gamma}$, $c_{j\gamma}^m$, Γ_j and u_j. Summing over all bid types we derive the objective function of (6).

Next, we revisit the capacity constraints. Constraint (6a) means that each bid of a specific type j is served in at most one option. If a bid of type j is served with option γ, then it consumes $c_{j\gamma}^m$ of type-m resource for a duration τ_j (out of T). So on average (over time), a bid of type j served with option γ consumes $\lambda T \delta_j \frac{\tau_i}{T} c_{j\gamma}^m x_{j\gamma}$ of type-m resource at each time. For example, in a $T = 2000$ time span, every 2 times lots 1 bid arrives in expectation, *i.e.,* $\lambda = 0.5$. Suppose the bid type j is drawn with probability of 0.1 each time, then the expected number of bids that are drawn as type j is $0.5 \times 2000 \times 0.1 = 100$ in total. If an option γ of a bid of type j requires

$c_{j\gamma}^m$ amount of resource of each m and occupies the resources for 1 time slot, then the expected amount of type-m resource occupied by the bid is $\frac{c_{j\gamma}^m}{2000}$ in each time slot if option γ is accepted, since the bid arrives with probability of $\frac{1}{T}$ in any time slot and only occupies resources for 1 time slot. Putting together, bids drawn as type j occupy $\frac{c_{j\gamma}^m}{20}$ amount of resource m in each time slot in expectation if option γ is accepted. Thus, constraint (6b) ensures that the average consumption of a resource does not exceed its capacity (expected capacity constraint). Note that this is a non-trivial relaxation of the capacity constraints in (3b) as we remove the time dimension. *A priori*, it is not clear whether the optimum of the relaxed program in (6) is much larger than that of (3). Fortunately, we show that it is possible to design an online algorithm based on (6), to obtain a $1 - O(\epsilon)$ fraction of the expected optimal social welfare, under the assumption that each bid does not consume a significant fraction of the overall capacity of any resource.

$$P^\delta : \text{maximize: } \sum_{j\in\mathcal{D}} \lambda T \delta_j \sum_{\gamma\in\Gamma_j} b_{j\gamma} x_{j\gamma} \quad (6)$$

subject to :

$$\sum_{\gamma\in\Gamma_j} x_{j\gamma} \le 1, \quad \forall j \in \mathcal{D} \quad (6a)$$

$$\sum_{j\in\mathcal{D}}\sum_{\gamma\in\Gamma_j} \lambda T \delta_j \frac{\tau_j}{T} c_{j\gamma}^m x_{j\gamma} \le C_m, \quad \forall m \in \mathcal{M} \quad (6b)$$

$$x_{j\gamma} \ge 0, \quad \forall j \in \mathcal{D},\ \gamma \in \Gamma_j \quad (6c)$$

The dual program of (6) is:

$$D^\delta : \text{minimize: } \sum_{m\in\mathcal{M}} C_m p_m + \sum_{j\in\mathcal{D}} \lambda T \delta_j u_j \quad (7)$$

subject to :

$$u_j \ge b_{j\gamma} - \sum_{m\in\mathcal{M}} p_m \frac{\tau_j}{T} c_{j\gamma}^m, \quad \forall j \in \mathcal{D},\ \gamma \in \Gamma_j \quad (7a)$$

$$p_m \ge 0, \quad \forall m \in \mathcal{M} \quad (7b)$$

$$u_j \ge 0, \quad \forall j \in \mathcal{D} \quad (7c)$$

After we remove the time dimension by using time averaged resource consumption in (6), it suffices to (approximately) derive price p_m for each type-m resource in (7).

To derive primal and dual solutions achieving a social welfare close to that of the expected offline problem (6), we resort to the KKT conditions to make decisions, similar to the idea discussed in Sec. IV-A. The obstacle is still that we do not know the offline optimal dual solution $\tilde{\mathbf{p}}$ in the online setting, and we cannot even derive the expected offline optimal $\tilde{\mathbf{p}}$ without any information of the distribution $\mathcal{D}$ of the bid types. Our idea is to learn an approximately optimal dual solution $\tilde{\mathbf{p}}$ of the offline problem based on the past bids, and progressively refine our dual solutions as time evolves. In particular, we divide bid arrivals in $[1, T]$ into $\log_2 \epsilon^{-1}$ stages, each marked by the arrival of a bid of index $2^\eta\lfloor\epsilon\lambda T\rfloor$. Here $\eta \in \{0, ..., \log_2 \epsilon^{-1} - 1\}$ indexes the stage. For each stage, we model an empirical formulation of (6) in P^η in (8) over the first $2^\eta\lfloor\epsilon\lambda T\rfloor$ bids in set $I_\eta = \{1, \dots, 2^\eta\lfloor\epsilon\lambda T\rfloor\}$, replacing the expectations over all bids in the objective and constraint (6b) with the sum over those sample bids, and accordingly shrinking the capacity limits by a factor of $\frac{I_\eta}{\lambda T}$, where $I_\eta = |I_\eta| = 2^\eta\lfloor\epsilon\lambda T\rfloor$. Let $\chi_\eta = \epsilon\sqrt{\frac{\lambda T}{I_\eta}} = \epsilon\sqrt{\frac{\lambda T}{2^\eta\epsilon\lambda T}} = \sqrt{\frac{\epsilon}{2^\eta}}$. Note that $\epsilon \le \chi_\eta \le \sqrt{\epsilon}$. We further shrink the capacity limits by a factor of $(1 - \epsilon\sqrt{\frac{\lambda T}{I_\eta}})$ to account for the sampling error and make sure overall resource consumption of bids is less than the capacity. Hence in (8), the modified resource capacity is $(1-\epsilon\sqrt{\frac{\lambda T}{I_\eta}})\frac{I_\eta}{\lambda T}C_m = (1-\chi_\eta)2^\eta\epsilon C_m$. Note that we convert bid type j in the expected program (6) back to bid i in (8). The dual of (8) is formulated in (9).

$$P^\eta : \text{maximize: } \sum_{i\in I_\eta}\sum_{\gamma\in\Gamma_i} b_{i\gamma} x_{i\gamma} \quad (8)$$

subject to :

$$\sum_{\gamma\in\Gamma_i} x_{i\gamma} \le 1, \quad \forall i \in I_\eta \quad (8a)$$

$$\sum_{i\in I_\eta}\sum_{\gamma\in\Gamma_i} \frac{\tau_i}{T} c_{i\gamma}^m x_{i\gamma} \le (1 - \epsilon\sqrt{\frac{\lambda T}{I_\eta}})\frac{I_\eta}{\lambda T} C_m, \quad \forall m \in \mathcal{M} \quad (8b)$$

$$x_{i\gamma} \in \{0, 1\}, \quad \forall i \in I_\eta, \gamma \in \Gamma_i \quad (8c)$$

$$D^\eta : \text{minimize: } \sum_{m\in\mathcal{M}} (1 - \epsilon\sqrt{\frac{\lambda T}{I_\eta}})\frac{I_\eta}{\lambda T} C_m p_m + \sum_{i\in I_\eta} u_i \quad (9)$$

subject to :

$$u_i \ge b_{i\gamma} - \sum_{m\in\mathcal{M}} p_m \frac{\tau_i}{T} c_{i\gamma}^m, \quad \forall i \in I_\eta,\ \gamma \in \Gamma_i \quad (9a)$$

$$p_m \ge 0, \quad \forall m \in \mathcal{M} \quad (9b)$$

$$u_i \ge 0, \quad \forall i \in I_\eta \quad (9c)$$

Upon the arrival of the $2^\eta\lfloor\epsilon\lambda T\rfloor^{\text{th}}$ bid, we solve the dual problem in (9) using all bids from bid 1 to bid $2^\eta\epsilon\lambda T$ to compute the dual variables, where the solution for $\mathbf{p}$ is denoted as $\mathbf{p}^\eta$. The dual problem is a linear program and can hence be exactly solved using an efficient algorithm such as the Karmarkar's Algorithm [25]. By iteratively solving (9) with more and more bids, we gradually learn an approximately optimal dual solution $\tilde{\mathbf{p}}$ of the offline problem. The intuition is that since the *types* of bids are i.i.d., drawn from an underlying distribution, the time-averaged resource consumption of the past bids approximately reflects the time-averaged resource consumption of all bids in expectation, especially when more past bids are accumulated.

With the learned dual solution $\mathbf{p}^\eta$, we make allocation and payment decisions upon each bid arrival: we let u_i be the maximal of 0 and the right hand side (RHS) of constraints (9a),

$$u_i = \max\left\{0, \max_{\gamma\in\Gamma_i}\left\{b_{i\gamma} - \sum_{m\in\mathcal{M}} \frac{\tau_i}{T} c_{i\gamma}^m p_m^\eta\right\}\right\} \quad (10)$$

and set $x_{i\gamma_i} = 1$ for $\gamma_i = \text{argmax}_{\gamma\in\Gamma_i}\{b_{i\gamma} - \sum_{m\in\mathcal{M}} \frac{\tau_i}{T} p_m^\eta c_{i\gamma_i}^m\}$, if $u_i > 0$ and $\sum_{i\in I}\sum_{\gamma\in\Gamma_i} c_{i\gamma}^m x_{i\gamma} + c_{i\gamma_i}^m < C_m, \forall m$ (no capacity of any resource would be exceeded if the bid is accepted in option γ_i), and compute the payment by $\hat{p}_i = \sum_{m\in\mathcal{M}} \frac{\tau_i}{T} p_m^\eta c_{i\gamma_i}^m$. However, note that our primal solution derived in this way may not be offline optimal as it may not satisfy the *sufficient condition* of the KKT conditions.

Algorithm 1: *Plastic*: An Online Auction With Price Learning

Input: S, K, $|\mathbb{E}|$, $\mathbf{C}$, ϵ, λ, T
Output: $\mathbf{x}, \hat{\mathbf{p}}$
Define: $I_\eta = 2^\eta \lfloor \epsilon\lambda T \rfloor$
Initialize: $\eta = 0$; $M = (K+2) \times S + |\mathbb{E}|$;
1 **while** *a new bid i arrives* **do**
2 **if** $i \le \lfloor \epsilon\lambda T \rfloor$ **then**
3 /*reject the first $\lfloor \epsilon\lambda T \rfloor$ bids*/
4 Reject bid i by setting $x_{i\gamma} = 0$ for all $\gamma \in \Gamma_i$;
5 **else**
6 $u_i = \max\left\{0, \max_{\gamma \in \Gamma_i}\left\{b_{i\gamma} - \sum_{m \in \mathcal{M}} \frac{\tau_i}{T} p_m^{\eta-1} c_{i\gamma}^m\right\}\right\}$;
7 $\gamma_i = \mathrm{argmax}_{\gamma \in \Gamma_i}\left\{b_{i\gamma} - \sum_{m \in \mathcal{M}} \frac{\tau_i}{T} p_m^{\eta-1} c_{i\gamma_i}^m\right\}$;
8 **if** $u_i > 0$ *and* $\sum_{i \in I} \sum_{\gamma \in \Gamma_i} c_{i\gamma}^m x_{i\gamma} + c_{i\gamma_i}^m < C_m, \forall m$ **then**
9 Accept bid i by setting $x_{i\gamma_i} = 1$, and $x_{i\gamma} = 0$ for all $\gamma \neq \gamma_i$;
10 Calculate the payment: $\hat{p}_i = \sum_{m \in \mathcal{M}} \frac{\tau_i}{T} p_m^{\eta-1} c_{i\gamma_i}^m$;
11 **if** $b_{i\gamma_i} > 0$ **then**
12 /*buy bid i*/
13 Assemble the service chain according to option γ_i and provide it to the customer;
14 Collect payment $\hat{p}_i$ from the customer;
15 **else**
16 /*sell bid i*/
17 Add the resources in option γ_i from the seller to the resource pool;
18 Send payment $-\hat{p}_i$ to the seller;
19 **end**
20 **else**
21 Reject bid i by setting $x_{i\gamma} = 0$ for all $\gamma \in \Gamma_i$;
22 **end**
23 **end**
24 **if** $i = I_\eta$ **then**
25 /*update prices*/
26 Exactly solve the dual linear program in ($8d$) and obtain $\mathbf{p}^\eta$;
27 Update $\eta = \eta + 1$;
28 **end**
29 **end**

The good news is that our solution does not deviate from the offline optimal solution much: the only difference between our solution and the offline optimum is that when $u_i = 0$, we choose $x_{i\gamma} = 0$ for all γ while the optimal solution may choose $x_{i\gamma_i} = 1$. That is, our solution is relatively conservative in the tie-breaking cases. Fortunately, the loss in social welfare can be bounded as we will show in Lemma 8, Lemma 9 and Lemma 10 in Sec. IV-C.2.

Algorithm Procedure: Our online price learning auction, *Plastic*, is given in Alg. 1. For the first $\lfloor \epsilon\lambda T \rfloor$ bids, the price $\mathbf{p}^0$ is not in place; so we simply reject all of them (lines 2-4). Then from bid $\lfloor \epsilon\lambda T \rfloor + 1$ to bid $2\lfloor \epsilon\lambda T \rfloor$, we use the optimal value of $\mathbf{p}^0$ learned from the first $\lfloor \epsilon\lambda T \rfloor$ bids (by solving D^0 above) as the allocation threshold. We determine each utility variable u_i of bid $i \in [2^{\eta-1}\lfloor \lambda T \rfloor + 1, 2^\eta \lfloor \lambda T \rfloor]$ by using $p_m^{\eta-1}$ in place of $\tilde{p}_m$ as in (10) (also in line 6). If $u_i > 0$ and accepting the bid will not violate any resource capacity constraint, we accept the bid in the best option γ_i and compute the payment (lines 7-10). For a winning buy bid, we assemble the required service chain using the asked resources and charge bidder i the payment $\hat{p}_i$ (lines 11-14). For an accepted sell bid, we take the resources contributed by the seller and give it the reward $-\hat{p}_i$, since $c_{i\gamma}^m$ in a sell bid is non-positive and hence $\hat{p}_i \le 0$ (lines 15-18). On the arrival of bid $2^\eta \lfloor \epsilon\lambda T \rfloor$, we further aggregate the bids in I_η and solve the dual optimization problem in (9), and then update η to $\eta+1$ (lines 24-28). We repeat the process and the last time we update the price $\mathbf{p}^\eta$ is the arrival time of bid $2^{(\log_2 \epsilon^{-1})-1}\lfloor \epsilon\lambda T \rfloor$. For example, suppose $\lambda T = 2^{10}$ and $\epsilon = 2^{-6}$. We reject the first 2^4 bids and solve (9) under the input of the first 2^4 bids with $\eta = 0$ before the $(2^4+1)^{\text{th}}$ bid arrives. We use $\mathbf{p}^1$ solved from (9) as prices for the following bids from the $(2^4+1)^{\text{th}}$ to the $2^{5\text{th}}$ and solve (9) again under the input of the first 2^5 bids with $\eta = 1$. We periodically solve (9) to update prices every time the bid number doubles until the last bid arrives. Such prices serve as thresholds for filtering out low value bids and for calculating payment of each accepted bid. Note that our algorithm does not require any knowledge of the bid type distribution $\mathcal{D}$ (δ_i). Though it takes λ and T as input, their values can just be estimates. We will show in the simulations that inaccurate estimations of the total number of bids bring little impact on the performance of the algorithm.

Truthfulness, Individual Rationality and Polynomial Time:

Theorem 1: Our online auction Plastic in Alg. 1 guarantees truthfulness of both buyers and sellers, individual rationality of all bidders, as well as polynomial running time.

Proof (Truthfulness in Bidding Price:) Recall that in each stage we aggregate the bids arrived so far to solve dual prices for use in the next stage. Thus our adopted marginal prices for computing the payment/reward of bid i (a sell or buy bid) depend only on information of past bids, and are hence independent of bidding price of bid i. Moreover, in line 7 of Alg. 1, we always select the option to serve a winning bid that maximizes the bidder's utility, using the derived prices. Therefore our auction follows sequential posted price mechanisms [26] where a bidder cannot increase its utility by reporting a fake bidding price. *(Truthfulness in resource demand and duration:)* For a buyer, requesting less VNF instances/bandwidth or a shorter usage duration may lead to a failure in serving the customer's traffic flow. Thus no buyer would take the risk to do that. If a buyer reports higher-than-necessary resource demand or a longer duration, the payment will be larger, which decreases its utility since the resource prices are non-negative according to (9b). For a seller, its reported resource amounts and supply duration are no larger/longer than that he can provide, otherwise it cannot fulfill its promise if the bid is accepted. Inversely, reporting smaller resource amounts or a shorter contribution duration reduces its utility due to a lower reward it will get. *(Individual rationality:)* According to (10), the utility of a bidder (seller or buyer) is always non-negative. *(Polynomial running time:)* In the auction process, we exactly solve an LP , the dual problem in (9), for $\lfloor \log_2 \epsilon^{-1} \rfloor$ times in total. Solving the dual problem each time can be completed in polynomial time (e.g., using Karmarkar's Algorithm [25], the time complexity is $O(I^{3.5}N^2 \log N \log\log N)$ where each bid is encoded in N bits). To process a bid i, we sum up the payment for each requested type of resource by multiplying the $\frac{\tau_i}{T}$ fraction of the demand by the corresponding dual price for each option $\gamma \in \Gamma_i$. Then, we calculate the utility u_i and decide acceptance and payment of bid i by checking the utility. Since for each option the running time is $O(((K+2)S+|\mathbb{E}|)N)$ and the number of options is $O((K+2)S+|\mathbb{E}|)$, the time complexity for handling each bid is at most $O(((K+2)S+|\mathbb{E}|)^2N)$. Given the above, the time complexity of the online auction is $(\lfloor \log_2 \epsilon^{-1} \rfloor) O(I^{3.5}N^2 \log N \log\log N) + I\,O(((K+2)S+|\mathbb{E}|)^2N)$ which is at most $O(\lfloor \log_2 \epsilon^{-1} \rfloor I^{3.5}((K+2)S+|\mathbb{E}|)^2N^2 \log N \log\log N)$. □

C. Competitive Analysis

We next show that our algorithm achieves near optimal social welfare in expectation in three steps. (1) With high probability, the total number of bids coming in $\mathcal{T}$ is close to the expected number λT. Conditioned on (1), we can further show that (2) with high probability our algorithm does not over-allocate any resources, and (3) with high probability, our solution achieves a $1 - O(\epsilon)$ fraction of the offline optimal social welfare in expectation.

1) Feasibility of the Original Problem: Note that the solutions trivially satisfy constraint (3a) and (3c). Hence it remains to prove our solution $\mathbf{x}(\mathbf{p})$ satisfies constraint (3b) with high probability. Lemma 4 shows that with high probability, accepted bids of our algorithm consume on average (over time) at most a $1 - \frac{1}{2}\chi_\eta^2$ fractional of the capacity for any type of resource (*i.e.*, they satisfy (6b) even if we decrease the capacity by a factor of $1 - \frac{1}{2}\chi_\eta^2$). (ii) Lemma 5 shows that with high probability, accepted bids consume at most the maximal capacity for any type of resource *at any time t* (*i.e.*, they satisfy (3b)). We define the following random variables which will be useful for our analysis:

$$\mathbf{X}_{im} = \frac{\tau_i}{T} \sum_{\gamma \in \Gamma_i} c_{i\gamma}^m x_{i\gamma}(\mathbf{p}^\eta) \tag{11}$$

$$\mathbf{Y}_{im}(t) = \begin{cases} \sum_{\gamma \in \Gamma_i} c_{i\gamma}^m x_{i\gamma}(\mathbf{p}^\eta), & \text{if } t_i \le t < t_i + \tau_i \\ 0, & \text{otherwise} \end{cases} \tag{12}$$

Here $\mathbf{Y}_{im}(t)$ is the demand for resource m of bid i in the accepted option γ_i at time t, and in the following lemma, we show that $\mathbf{X}_{im}$ is at least the expectation of $\mathbf{Y}_{im}(t)$, which is used in the proof of Lemma 5 later. Let $\mathcal{A}$ denote the uniform distribution where the arrival time of each bid is independently drawn within $[1, T]$.

Lemma 1: The expectation of $\mathbf{Y}_{im}(\omega)$ *on t is no larger than* $\mathbf{X}_{im}$ *when* $\mathcal{A}$ *is a uniform distribution within* $[1, T]$.

Proof: Since $c_{i\gamma}^m, \tau_i$ are chosen independently of t_i, so we can fix $c_{i\gamma}^m$ and τ_i in the calculation of the expectation of $\mathbf{Y}_{im}(\omega)$ on $t_i \sim \mathcal{A}$. For any fixed $t \in [1, T]$, we have

$$\begin{aligned} E_{t_i \sim \mathcal{A}}[\mathbf{Y}_{im}(t)] &= \Pr[t_i \le t < t_i + \tau_i] \sum_{\gamma \in \Gamma_i} c_{i\gamma}^m x_{i\gamma}(\mathbf{p}^\eta) \\ &= \Pr[t - \tau_i < t_i \le t] \sum_{\gamma \in \Gamma_i} c_{i\gamma}^m x_{i\gamma}(\mathbf{p}^\eta) \end{aligned} \tag{13}$$

Since we have $0 \le t \le T$, then $\Pr[t - \tau_i < t_i \le t]$ has different values between two cases (see (13a) and (13b)). Note that the probability mass function of the random variable arrival time t_i for any fixed t is $\Pr[t_i = t] = \frac{1}{T}$ where $t \in [1, T]$. Thus we have

$$(13) = \begin{cases} \frac{t}{T} \sum_{\gamma \in \Gamma_i} c_{i\gamma}^m x_{i\gamma}(\mathbf{p}^\eta), \quad 0 \le t < \tau_i & (13a) \\ \frac{\tau_i}{T} \sum_{\gamma \in \Gamma_i} c_{i\gamma}^m x_{i\gamma}(\mathbf{p}^\eta), \quad \tau_i \le t \le T & (13b) \end{cases}$$

Given this, we have that $E[\mathbf{Y}_{im}(\omega)] < \mathbf{X}_{im}$ when $0 \le t < \tau_i$ and $E[\mathbf{Y}_{im}(\omega)] = \mathbf{X}_{im}$ when $\tau_i \le t \le T$, which gives the lemma in general. □

Lemma 2: With probability at least $1-\epsilon$*, the total number of bids arrived during* $[1, T]$ *is in the range of* $[(1-\frac{\epsilon}{2})\lambda T, (1+\frac{\epsilon}{2})\lambda T]$*, given* $\lambda T > \frac{4}{\epsilon^3}$.

Proof: According to Chebyshev's Inequality [27], we have that

$$\begin{aligned} \Pr[|I - \lambda T| \ge \frac{\epsilon}{2}\lambda T] &= \Pr[|I - \mathrm{E}[I]| \\ &\ge \frac{\epsilon}{2}\lambda T] \le \frac{\mathrm{Var}[I]}{(\frac{\epsilon}{2}\lambda T)^2} = \frac{4\lambda T}{\epsilon^2 \lambda^2 T^2} = \frac{4}{\epsilon^2 \lambda T} \end{aligned}$$

Thus we have

$$\begin{aligned} \Pr[(1-\frac{\epsilon}{2})\lambda T \le I \le (1+\frac{\epsilon}{2})\lambda T] &\ge 1 - \Pr[|I - \lambda T| \ge \frac{\epsilon}{2}\lambda T] \\ &\ge 1 - \frac{4}{\epsilon^2 \lambda T} \end{aligned} \tag{14}$$

If $\lambda T \ge 4/\epsilon^3$ holds, we have (14) $\ge 1-\epsilon$. □

Lemma 3: With probability at least $1-\epsilon$*, we have* $(1 - \frac{\chi_\eta}{2})\lambda T \le I \le (1 + \frac{\chi_\eta}{2})\lambda T$*, for any* $\eta \in \{0, ..., \log_{\frac{1}{2}} \epsilon - 1\}$.

Proof: Since $\min_\eta \chi_\eta = \sqrt{\frac{\epsilon}{2^{\log_{1/2}\epsilon}}} = \epsilon$, the lemma follows Lemma 2. □

Lemma 4: On the condition of $(1-\frac{\epsilon}{2})\lambda T \le I \le (1+\frac{\epsilon}{2})\lambda T$*, for any* $\eta \in \{0, 1, ..., \log_{\frac{1}{2}} \epsilon - 1\}$*, we have*

$$\sum_{i \in I_{\eta+1} \setminus I_\eta} \mathbf{X}_{im} \le (1 - \frac{\chi_\eta^2}{2}) 2^\eta \epsilon C_m, \quad \forall m \in \mathcal{M}$$

with probability at least $1-\epsilon$*, given* $\frac{\min_{m \in \mathcal{M}} C_m}{\max_{i,\gamma,m} |c_{i\gamma}^m|} \ge \frac{13M \log(2I^2/\epsilon)}{\epsilon^2}$.

Lemma 5: Conditioned on $(1 - \chi_\eta/2)\lambda T \le I \le (1 + \chi_\eta/2)\lambda T$ *and* $\sum_{i \in I} \mathbf{X}_{im} \le (1 - \frac{\chi_\eta^2}{2}) 2^\eta \epsilon C_m$*, we have*

$$\sum_{i \in I_{\eta+1} \setminus I_\eta} \mathbf{Y}_{im}(\omega) \le 2^\eta \epsilon C_m, \quad \forall m \in \mathcal{M}, \eta \in \{1, .., \log_{\frac{1}{2}} \epsilon - 1\},$$
$$\omega \in \{t_i, t_i + \tau_i\}_{i \in I}$$

with probability at least $1-\epsilon$*, given that* $\frac{\min_{m \in \mathcal{M}} C_m}{\max_{i \in I, \gamma \in \Gamma_i, m \in \mathcal{M}} |c_{i\gamma}^m|} \ge \frac{13M \log(2I^2/\epsilon)}{\epsilon^2}$ *for all* $m \in \mathcal{M}$.

Lemma 6: With probability at least $1-2\epsilon$*, we have*

$$\sum_{i\in I_{\eta+1}\setminus I_\eta} \mathbf{Y}_{im}(\omega) \le 2^\eta \epsilon C_m, \quad \forall\gamma\in\Gamma_i,\ m\in\mathcal{M},$$
$$\eta\in\{1,..,\log_{\frac{1}{2}}\epsilon-1\},\quad \omega\in\{t_i, t_i+\tau_i\}_{i\in I}$$

under the assumptions (i) $\lambda T \ge 4/\epsilon^3$*, and (ii)* $\frac{\min_{i\in\mathcal{M}} C_m}{\max_{i\in I,\gamma\in\Gamma_i,m\in\mathcal{M}} |c^m_{i\gamma}|} \ge \frac{13M\log(2I^2/\epsilon)}{\epsilon^2}$.

Proof: Suppose $\mathcal{N}$ is the event that $(1-\chi_\eta/2)\lambda T \le I \le (1+\chi_\eta/2)\lambda T$ and $\mathcal{F}$ is the event that $\sum_{i\in I_{\eta+1}\setminus I_\eta}\mathbf{X}_{im} \le (1-\frac{1}{2}\chi_\eta^2)C_m$ for all m, ω. Lemma 2 shows that $\mathcal{N}$ holds with probability at least $1-\epsilon$. Lemma 4 shows that conditioned on $\mathcal{N}$, $\mathcal{F}$ holds with probability at least $1-\epsilon$. Lemma 5 shows that conditioned on both $\mathcal{N}$ and $\mathcal{F}$, the conclusion of the lemma holds with probability at least $1-\epsilon$. Thus, we have

$$\begin{aligned}
&\Pr[\sum_{i\in I_{\eta+1}\setminus I_\eta} \mathbf{Y}_{im}(\omega) \le 2^\eta\epsilon C_m]\\
&\ge \Pr[\sum_{i\in I_{\eta+1}\setminus I_\eta} \mathbf{Y}_{im}(\omega) \le 2^\eta\epsilon C_m \mid \mathcal{F},\mathcal{N}]\Pr[\mathcal{F}\mid\mathcal{N}]\Pr[\mathcal{N}]\\
&\ge (1-\epsilon)^3 \ge 1-3\epsilon
\end{aligned}$$

□

2) Competitive Ratio in Social Welfare: Let OPT denote the offline optimal social welfare, *i.e.*, the optimal objective value of the offline problem in (3). The following lemma shows that the objective value of problem (6) serves as an upper-bound of the expected social welfare of the offline problem.

Lemma 7: The optimal objective value of (6) is an upper-bound of $E[OPT]$*, the expectation of the offline optimal social welfare, computed by solving the offline social welfare maximization problem in (3) over all possible realization of the bid arrival process.*

Proof: We can readily see that the average of optimal solutions of the offline problem in (3), computed over all possible realizations of bid arrival process, provides a feasible solution to the distribution instance program in (6). These solutions achieve an expected offline social welfare of $E[OPT]$. Hence the optimal social welfare of the distribution instance program can only be larger than $E[OPT]$. □

Assumption 1: The inputs of the program (8) are in general position, namely for any dual solution vector $\mathbf{p}$ *derived by solving the dual program (9), there can be at most* M *equations such that* $b_{i\gamma_i} = \sum_{m\in\mathcal{M}} p_m c^m_{i\gamma}$*, for all* $i\in I$ *and* γ_i *denotes the best option of bid* i*.*

Let $x_{i\gamma}(\mathbf{p}^\eta)$ denote the primal solution output by Alg. 1, which is a function of the price vector solved in (8). According to Alg. 1, we have

$$x_{i\gamma}(\mathbf{p}^\eta) = \begin{cases} 1, & \\ \quad\text{if } \gamma = \text{argmax}_{\gamma'\in\Gamma_i}\{b_{i\gamma'} - \sum_{m\in\mathcal{M}}\frac{\tau_i}{T}p^\eta_m c^m_{i\gamma'}\} & \\ \quad\text{and } b_{i\gamma} > \sum_{m\in\mathcal{M}}\frac{\tau_i}{T}p^\eta_m c^m_{i\gamma} & \\ 0, & \\ \quad\text{otherwise} & \end{cases} \tag{15}$$

Let $x^\eta_{i\gamma}$ denote the optimal solution of (8). Let $x_i(\mathbf{p}^\eta)$ and x^η_i be the sum of $x_{i\gamma}(\mathbf{p}^\eta)$ over $\gamma\in\Gamma_i$ and the sum of $x^\eta_{i\gamma}$ over $\gamma\in\Gamma_i$, respectively. The following lemma states that our solution $\mathbf{x}(\mathbf{p}^\eta)$, which is induced by dual price vector $\mathbf{p}^\eta$ solved in (9) at stage η, will deviate at most M from the optimal solution $\mathbf{x}^\eta$ of (8).

Lemma 8: $\sum_{i\in I} x^\eta_i - M \le \sum_{i\in I} x_{i\gamma_i}(\mathbf{p}^\eta) \le \sum_{i\in I} x^\eta_i$*, where* $\gamma_i = argmax_{\gamma'\in\Gamma_i}\{b_{i\gamma'} - \sum_{m\in\mathcal{M}}\frac{\tau_i}{T}p^\eta_m c^m_{i\gamma'}\}$*.*

Proof: Applying complementary slackness conditions to (8) and (9), the optimal solution of (8) satisfies that $x^\eta_{i\gamma} = 0$ if $b_{i\gamma_i} < \sum_{m\in\mathcal{M}} p^\eta_m c^m_{i\gamma}$. Compared with (15), if bid i satisfies $b_{i\gamma_i} = \sum_{m\in\mathcal{M}} p^\eta_m c^m_{i\gamma}$, the optimal solution $x^\eta_{i\gamma}$ could be 1 while our induced solution $x_{i\gamma}(\mathbf{p}^\eta)$ definitely equals to 0. In the remaining cases where $b_{i\gamma_i} < \sum_{m\in\mathcal{M}} p^\eta_m c^m_{i\gamma}$ or $b_{i\gamma_i} > \sum_{m\in\mathcal{M}} p^\eta_m c^m_{i\gamma}$, we have $x_{i\gamma}(\mathbf{p}^\eta) = x^\eta_{i\gamma}$. These imply that bids accepted by Alg. 1 are also accepted by the offline optimal solution while some of the accepted bids in the optimal solution are rejected by Alg. 1. Since there are at most M bids which satisfy $b_{i\gamma_i} = \sum_{m\in\mathcal{M}} p^\eta_m c^m_{i\gamma}$ according to Assumption 1, there are at most M bids that are accepted by the optimum while rejected by Alg. 1.

□

Lemma 9: Given $\frac{\min_{m\in\mathcal{M}} C_m}{\max_{i,\gamma,m}|\tau_i c^m_{i\gamma}|} \ge \frac{13M\log(2I^2/\epsilon)}{\epsilon^2}$*, we have*

$$\sum_{i\in I_{\eta+1}}\sum_{\gamma\in\Gamma_i} b_{i\gamma}x_{i\gamma}(\mathbf{p}^\eta)) \ge (1-2\chi_\eta-\epsilon)P^{\star\eta+1}(\mathbf{x}^{\eta+1}) \tag{16}$$

where $\sum_{i\in I_{\eta+1}}\sum_{\gamma\in\Gamma_i} b_{i\gamma}x_{i\gamma}(\mathbf{p}^\eta))$ *is the objective value of* $P^{\eta+1}$ *under the allocation vector* $x(\mathbf{p}^\eta)$ *output by Alg. 1, and* $P^{\star\eta+1}(\mathbf{x}^{\eta+1})$ *is the optimal objective value of (8) under optimal solution* $\mathbf{x}^{\eta+1}$*.*

Lemma 10: Given $\frac{\min_{m\in\mathcal{M}} C_m}{\max_{i,\gamma,m}|\tau_i c^m_{i\gamma}|} \ge \frac{13M\log(2I^2/\epsilon)}{\epsilon^2}$*, we have*

$$\sum_{i\in I_{\eta_{\max}+1}}\sum_{\gamma\in\Gamma_i} b_{i\gamma}x_{i\gamma}(\mathbf{p}^{\eta_{\max}})) \ge (1-2\chi_\eta-\epsilon)P^{\star\delta} \tag{17}$$

where $\sum_{i\in I_{\eta+1}}\sum_{\gamma\in\Gamma_i} b_{i\gamma}x_{i\gamma}(\mathbf{p}^\eta))$ *is the objective value of* $P^{\eta+1}$ *under the allocation vector* $x(\mathbf{p}^\eta)$ *output by Alg. 1, and* $P^{\star\delta}$ *is the optimal objective value of (6).*

Lemma 11: Let $(\mathbf{x}^\eta,\mathbf{p}^\eta,\mathbf{u}^\eta)$ *denote the optimal primal-dual solution of (8) and* $(\mathbf{x}^\delta,\mathbf{p}^\delta,\mathbf{u}^\delta)$ *denote the optimal primal-dual solution of the distribution instance program (6), we have*

$$P^{\star\eta}(\mathbf{x}^\eta) \le \frac{2^\eta\epsilon}{1-\epsilon/2}P^{\delta\star} \tag{18}$$

where $P^{\delta\star}$ *is the optimal objective value of* P^δ *in (6).*

Theorem 2: For any $\epsilon > 0$*, the online auction* $Plastic$ *in Alg. 1 is* $1-O(\epsilon)$ *competitive in expected social welfare with i.i.d. bid types and uniform bid arrival time distribution, as compared to the expected social welfare of offline problem (3), for all inputs such that*

$$\frac{\min_{m\in\mathcal{M}} C_m}{\max_{i,\gamma,m}|c^m_{i\gamma}|} \ge \frac{13M\log(2I^2/\epsilon)}{\epsilon^2} \tag{19}$$

V. Performance Evaluation

We perform trace-driven simulation studies to evaluate our online auction mechanism $Plastic$.

A. Simulation Setup

We set the number of zones to 13, according to the number of Google data centers [28]. Each customer demands a service chain containing $2-5$ VNFs, randomly picked among firewall, proxy, NAT and IDS. The rate of traffic flow to be handled by each service chain is produced by multiplying the average rate of HTTP requests in the Wikipedia trace [29] with a coefficient in [0.5, 1.5]. The number of instances of each VNF in each chain (*i.e.*, d_{ik}) is set based on the traffic rate and typical processing capacity of one instance of each VNF, as given in the table below [30]. For example, if the incoming traffic rate to a service chain "Firewall → IDS" is $800Mbps$, then the numbers of firewall and IDS instances are 1 and 2, respectively, since $900Mbps \times 1 \geq 800Mbps, 600Mbps \times 2 \geq 800Mbps$. We distribute the flow evenly to multiple instances of a VNF in the service chain. Hence the bandwidth demand between instances of two consecutive VNFs can be determined accordingly. Following the previous example, the bandwidth demand from a firewall instance to an IDS instance is $800Mbps/2$. The CPU demand of each instance is given in the table below while the RAM and disk capacities are set according to those of the respective VM instances on Amazon EC2 [31]. We independently randomly choose a zone for each instance in a required chain. By default, each customer submits $1-5$ options in its bid, each of which corresponds to a random VNF placement as done above. After setting all the resource demands ($c_{i\gamma}^m$), we normalize all $|c_{i\gamma}^m|$'s to be within [0, 1], by dividing each demand by $\max_{i\in I,\gamma\in\Gamma_i,m\in\mathcal{M}} |c_{i\gamma}^m|$. The bidding price of each option in each buy bid is the weighted sum of normalized demands, where the weights are randomly picked within [0, 1].

Network Function	CPU Required	Processing Capacity	AWS Instance Type
Firewall	4	900Mbps	m4.xlarge
Proxy	4	900Mbps	m4.xlarge
NAT	2	900Mbps	m4.large
IDS	8	600Mbps	m4.2xlarge

By default, we set $\lambda = 0.5$, and decide the total number of bids I according to the Poisson distribution with the expectation of λT. We independently and uniformly choose an arrival time within $[1, T]$ for each bid, to simulate a Poisson process. We will also show performance of our algorithm when λ, as an input to Alg. 1, is not accurately known, with the bid arrival process deviating from a Poisson process. We set each bid to be a buy bid with probability 0.9 and a sell bid with probability 0.1. For the sell bids, resource amounts and bidding prices are set similarly to how those in buy bids are set, except for making them negative values. The duration of each bid (τ_i) is uniformly randomly drawn from $[10, 10^3]$, which may not be much smaller than T.

For each zone, we set the upload/download bandwidth capacities of the NFV provider (Θ_s^{out} and Θ_s^{in}) to be the total outgoing traffic rate of all required service chains multiplied by a random number within [0.2, 1]. We set the bandwidth capacity of each link ($L_{ss'}$) as follows: suppose a zone is connected to n outgoing (incoming) links; then the bandwidth capacity of each of the links is set to be $\frac{1}{n}$ of the upload (download) capacity of the zone, multiplied by a random number within [0.6, 1]. In the default setting, we set $\frac{\min_{m\in\mathcal{M}} C_m}{\max_{i\in I,\gamma\in\Gamma_i,m\in\mathcal{M}} |c_{i\gamma}^m|}$, the lower-bound of the ratio of resource capacity to the corresponding demand in each option of each bid, to be 500. We decide the capacity of each VNF in each zone (C_{ks}), by multiplying $\frac{\min_{m\in\mathcal{M}} C_m}{\max_{i,m,\gamma} |c_{i\gamma}^m|}$ by a random number within [1.0, 1.5].

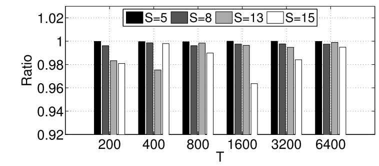

Fig. 2. Performance of $Plastic$ with different S.

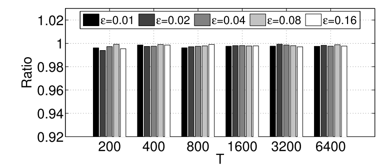

Fig. 3. Performance of $Plastic$ with differet ϵ.

B. Comparison With Offline Optimum

We estimate the expected offline social welfare by solving ($3a$) exactly using CPLEX for 50 times under different realizations of the bid arrival process, in each set of experiments below. We first compute the ratio of the average social welfare achieved by $Plastic$ (over different realizations of the bid process) over the expected offline optimal social welfare. Fig. 2 indicates that the ratio decreases slightly with the increase of the number of zones. Fig. 3 shows that the input parameter ϵ to our Alg. 1 does not influence the ratio much. It implies that very few bids need to be rejected in stage 1 in $Plastic$. Theorem 2 shows that the theoretical competitive ratio is inversely related to ϵ; our empirical studies reveal little impact of ϵ. In addition, recall the capacity assumption we made in (19). The default value 500 of $\frac{\min_{m\in\mathcal{M}} C_m}{\max_{i,\gamma,m} |c_{i\gamma}^m|}$ used in our experiments is much smaller than its lower bound in the assumption computed using M, I, ϵ in our experimental settings. This implies that although a lower bound of the capacity is required for our competitive analysis, in practice, even if the assumption is not obeyed, the average social welfare achieved by $Plastic$ is still close to the offline optimum.

In Fig. 4, we run $Plastic$ by feeding $\beta\lambda$ into the algorithm, instead of the actual arrival rate λ used to produce the bid arrivals, where β is the percentage in legends of the figure. The results show that the inaccurate estimation of λ (and hence λT needed in Alg. 1) does not affect the performance of our mechanism either. Fig. 5 illustrates that the ratio is positively correlated with resource capacity, which is consistent with our theoretical analysis that a larger $E[\frac{C_m}{|c_{i\gamma}^m|}]$ is more desirable. In Fig. 6, we compare the ratios achieved when the bid arrivals are produced following a Poisson process (the default) and

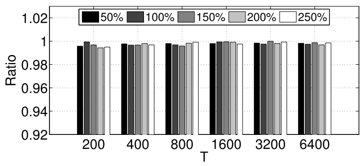

Fig. 4. Performance of *Plastic* with different levels of inaccurate estimation of λ.

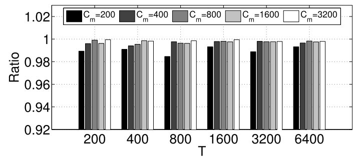

Fig. 5. Ratio of *Plastic* with different C_m.

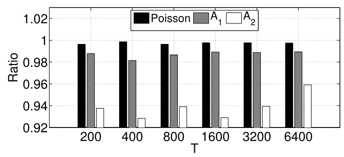

Fig. 6. Ratio of *Plastic* with different bid arrival processes.

two other arrival patterns ($\mathcal{A}_1$ and $\mathcal{A}_2$). Following $\mathcal{A}_1$, the total number of bids I is drawn from a deviated Poisson distribution according to the pattern shown in [23, Fig. 6], which violates the first property of a Poisson process that we mentioned at the beginning of Sec. III-B. Following $\mathcal{A}_2$, arrival time of each bid is chosen with a varying probability over $[1, T]$, set roughly according to [23, Fig. 3], which violates the second property in Sec. III-B. Fig. 6 show that the ratio is worse with $\mathcal{A}_2$ (the non-uniform bid arrival time distribution) than $\mathcal{A}_1$ (a different distribution of total number of bids). Nevertheless, the ratios achieved by our online auction in all these figures are above 0.92, which are indeed very close to the offline optimum.

C. Comparison With an Existing Scheme

We further compare *Plastic* with an online algorithm *HST*. *HST* is an extension of the algorithm from [32]: it uses a pre-determined price function (denoted by $p_{mt}(q_{mt})$) of the amount of currently allocated resource m (q_{mt}), for each type of resource at each time, to compute the payment for both buyers and sellers, which also serves as the threshold for winner determination; it does not include the initial bid rejection stage and price updates. The price function is as follows, where U_m and L_m are respectively the maximal and minimal demand for type-m resource in absolute value over all the bids.

$$p_{mt}(q_{mt}) = \frac{L_m}{2\max_{i\in I}\frac{T}{\tau_i}M}\left(\frac{2\max_{i\in I}\frac{T}{\tau_i}MU_m}{L_m}\right)^{\frac{q_{mt}}{C_m}}$$

We compare the performance of the algorithms by computing a similar ratio for *HST*, the average social welfare achieved by *HST* divided by the expected offline optimal social welfare.

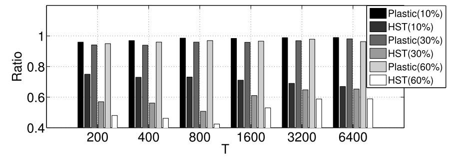

Fig. 7. Comparison with different percentages of sellers in all bids.

Fig. 7 shows that *Plastic* works much better than *HST* under different percentages of sell bids among all bids, since when utilization is extremely low, *HST* may mistakenly reject sellers who should have been accepted to supply resources for more later buyers. The results verify that our algorithm works better in an online auction where both buyers and sellers participate.

VI. Conclusion

This paper aims to design a novel online stochastic auction mechanism for the NFV market to provision and price service chains, and to purchase supplementary resources on the go. We leverage novel online primal-dual frameworks with a learning-based strategy embedded for obtaining the resource prices as the dual solutions. The prices are repeatedly updated to serve as thresholds for winner determination, and for the computation of the payment/reward for each winning buyer/seller. As new technical contributions, our mechanism design with competitive analysis significantly extends existing techniques by handling both buyers and sellers, each of which occupies the VNF service chain or supplies resources only for a limited period of time. We show that the proposed mechanism achieves truthfulness for both VNF service chain buyers and resource sellers, with a near-optimal social welfare in expectation. Trace-driven simulations further validate the performance of the online mechanism in different variations of the problem settings as used in theoretical analysis.

Appendix A
Proof of Lemma 4

Proof: For a fixed price $\mathbf{p}$ and m, we say a random sample I_η is "bad" for this $\mathbf{p}$ if and only if $\mathbf{p} = \mathbf{p}^\eta$, but $\sum_{i\in I}\frac{\tau_i}{T}c_{i\gamma}^m x_{i\gamma}(\mathbf{p}) > (1-\frac{\chi_\eta^2}{2})C_m$ for some m. First, we show that the probability of bad samples is small for every fixed $\mathbf{p}$ and m. Then we take union bound over all "distinct" prices and constraint dimension m of (8) that with high probability the learned price $\mathbf{p}^\eta$ will be such that $\sum_{i\in I_{\eta+1}\setminus I_\eta}\frac{\tau_i}{T}c_{i\gamma}^m x_{i\gamma}(\mathbf{p}^\eta) \le (1-\frac{\chi_\eta^2}{2})2^\eta\epsilon C_m$ for all m. For simplicity of notations, we define events $A = \{\sum_{i\in I_\eta}\mathbf{X}_{im} \le (1-\chi_\eta)2^\eta\epsilon C_m\}$, $B = \{\sum_{i\in I_{\eta+1}\setminus I_\eta}\mathbf{X}_{im} \ge (1-\frac{\chi_\eta^2}{2})2^\eta\epsilon C_m\}$, $D = \{\sum_{i\in I_{\eta+1}}\mathbf{X}_{im} \ge (1-\frac{\chi_\eta^2}{2})2^{\eta+1}\epsilon C_m\}$ and let D^c denote the complement of the event D. We aim to prove the probability that A and B occur simultaneously is at most ϵ under the condition of $(1-\frac{\epsilon}{2})\lambda T \le I \le (1+\frac{\epsilon}{2})\lambda T$. Since we have $\Pr[A,B] = \Pr[A,B,D]+\Pr[A,B,D^c]$. We calculate $\Pr[A,B,D]$ first and obtain $\Pr[A, B, D^c]$ similarly.

$$\Pr[A, B, D] \leq \Pr[A, D] \leq \Pr[A \mid D] \leq \Pr[\sum_{i \in I_\eta} \mathbf{X}_{im} - \frac{I_\eta}{I_{\eta+1}} \sum_{i \in I_{\eta+1}} \mathbf{X}_{im} \leq \alpha \mid D] \tag{20}$$

Where $\alpha = (1-\chi_\eta)\frac{I_\eta}{I}C_m - \frac{\lambda T}{2I}(1-\frac{\chi_\eta^2}{2})C_m$. Due to $\frac{\lambda T}{I} \geq \frac{1}{1+\frac{\chi_\eta}{2}}$, we further have

$$\alpha \leq 2^\eta \epsilon C_m(1 - \chi_\eta - \frac{1-\frac{\chi_\eta^2}{2}}{1+\chi_\eta/2}) \leq 2^\eta \epsilon C_m(\frac{-\chi_\eta^2/2}{1+\chi_\eta/2}) < 0 \tag{21}$$

We normalize $c_{i\gamma}^m$'s such that $|\mathbf{X}_{im}| \in [0, 1]$ and use $\frac{C_m}{\max_{i,\gamma,m}|c_{i\gamma}^m|}$ in place of C_m. We define random variables

$$\sigma^2(\mathbf{X}) = \frac{1}{I_{\eta+1}} \sum_{i \in I_{\eta+1}} (\mathbf{X}_{im} - \frac{1}{I_{\eta+1}} \sum_{i \in I_{\eta+1}} \mathbf{X}_{im})^2$$
$$\leq \frac{1}{2^\eta \epsilon I} \sum_{i \in I_{\eta+1}} \mathbf{X}_{im} = \frac{1}{2^\eta \epsilon I}(1 - \frac{\chi_\eta^2}{2})2^{\eta+1}\epsilon C_m \tag{22}$$

$$\Delta(\mathbf{X}) = \max_{i \in I} \mathbf{X}_{im} - \min_{i \in I} \mathbf{X}_{im} \leq 2 \tag{23}$$

According to Hoeffding-Berstein Inequality, we have

$$(20) \leq \exp(\frac{-\alpha^2}{2I_\eta \sigma^2(\mathbf{X}) + (-\alpha)\Delta(\mathbf{X})}) \tag{24}$$

Then taking (21), (22), (23) and $I_\eta/I \leq \frac{2^\eta \epsilon}{1-\chi_\eta/2}$, $\frac{\lambda T}{I} \leq \frac{1}{1-\frac{\chi_\eta}{2}}$ into (24), we have

$$(24) \leq \exp(\frac{-\alpha^2}{2\frac{I_\eta}{2^\eta \epsilon I}(1-\frac{\chi_\eta^2}{2})2^{\eta+1}\epsilon C_m + 2(-\alpha)})$$
$$\leq \exp(\frac{-2^{2\eta}\epsilon^2 C_m^2 \frac{\chi_\eta^2}{4(1+\chi_\eta/2)^2}}{\frac{1}{1-\chi_\eta/2}(1-\frac{\chi_\eta^2}{2})2^{\eta+1}\epsilon C_m + \frac{\chi_\eta^2}{2(1+\chi_\eta/2)^2}2^\eta \epsilon C_m})$$
$$\leq \exp(\frac{-2^\eta \epsilon C_m \frac{\chi_\eta^2}{4(1+\chi_\eta/2)^2}}{2(1+\chi_\eta/2) + \frac{\chi_\eta^2}{2(1+\chi_\eta/2)^2}})$$
$$\leq \exp(\frac{-2^{\chi_\eta}\epsilon C_m \chi_\eta^2}{8(1+\chi_\eta/2) + 2\chi_\eta^2}) \tag{25}$$

Putting $\chi_\eta < 1$, $\epsilon < \frac{1}{2}$, and $\chi_\eta = \sqrt{\frac{\epsilon}{2^\eta}}$ into (25), we have

$$(25) = \exp(\frac{-2^\eta \epsilon C_m \frac{\epsilon}{2^\eta}}{8(1+\frac{\chi_\eta}{2}) + 2\frac{\epsilon}{2^\eta}})$$
$$\leq \exp(\frac{-\epsilon^2 C_m}{12+2\epsilon}) \leq \exp(\frac{-\epsilon^2 C_m}{13}) \leq \frac{\epsilon}{2MI^M \log_{\frac{1}{2}}\epsilon} \tag{26}$$

where the last inequality holds under that $C_m/\max_{i\in I}|c_{i\gamma}^m| \geq \frac{13M\log(2I^2/\epsilon)}{\epsilon^2}$, and $(2I)^M \geq 2\log_2 1/\epsilon$ which follows $\lambda T \geq 4/\epsilon^3$ (Lemma 2) and $I \geq \frac{1}{1+\chi_\eta}\lambda T \geq 1/2\lambda T$. Similarly, we have $\Pr[A, B, D^c] \leq \frac{\epsilon}{2MI^M \log_{\frac{1}{2}}\epsilon}$. Then $\Pr[A, B] \leq \frac{\epsilon}{MI^M \log_{\frac{1}{2}}\epsilon}$. Then, we take a union bound over all "distinct" prices $\mathbf{p}$'s. Each distinct $\mathbf{p}$ is characterized by a unique separation of I ($\{b_{i\gamma}, \mathbf{c}_i\}_{i=1}^I$) points in M-dimensional space by a hyperplane. The total number of such distinct prices is at most I^M. Taking union bound over M resources and $\log_2(\epsilon^{-1})$ stages proves the lemma. □

Appendix B
Proof of Lemma 5

Proof: For a fixed $\mathbf{p}$, m, η, and ω, we have

$$\Pr[\sum_{i\in I_{\eta+1}\setminus I_\eta} \mathbf{Y}_{im}(\omega) \leq 2^\eta \epsilon C_m, \sum_{i\in I_{\eta+1}\setminus I_\eta} \mathbf{X}_{im} \geq (1-\frac{\chi_\eta^2}{2})2^\eta \epsilon C_m]$$
$$\leq \Pr[\sum_{i\in I_{\eta+1}\setminus I_\eta} \mathbf{Y}_{im}(\omega) \leq 2^\eta \epsilon C_m \mid \sum_{i\in I_{\eta+1}\setminus I_\eta} \mathbf{X}_{im} = (1-\frac{\chi_\eta^2}{2})2^\eta \epsilon C_m] \tag{27}$$

Similarly as $\sigma^2(\mathbf{X})$ and $\Delta(\mathbf{X})$, we have

$$\sigma^2(\mathbf{Y}) = \sum_{i\in I}(\mathbf{Y}_{im}(\omega) - \bar{\mathbf{Y}}_{kis}(\omega))^2 \leq 2\Delta(\mathbf{Y})$$
$$= \max_{i,m,\omega} \mathbf{Y}_{im}(\omega) - \min_{i,m,\omega} \mathbf{Y}_{im}(\omega) \leq 1 \tag{28}$$

where $\bar{\mathbf{Y}}_{im}(\omega)$ is the demand of type-m resource of i at time slot ω ($\omega \in \{t_i, t_i + \tau_i\}_{i\in I}$). Further we have that

$$(27) \leq \Pr[\sum_{i\in I_{\eta+1}\setminus I_\eta} \mathbf{Y}_{im}(\omega) - E_{t_i\sim\mathcal{A}}[\mathbf{Y}_{im}(\omega)] \geq \frac{\chi_\eta^2}{2}2^\eta \epsilon C_m \mid \mathcal{L}']$$
$$\leq \exp\Big(\frac{-(\frac{\chi_\eta^2}{2}2^\eta \epsilon C_m)^2}{2\sigma^2(\mathbf{Y}) + \frac{\chi_\eta^2}{2}2^\eta \epsilon C_m \Delta(\mathbf{Y})}\Big) \tag{29}$$

By the assumed lower bound on C_m, we have $\frac{\epsilon^2 C_m}{2} \geq 1$. Thus,

$$(29) \leq \exp(\frac{-\epsilon^4 C_m^2}{8\epsilon^2 C_m}) = \exp(\frac{-\epsilon^2 C_m}{8}) \leq \exp(-M\log(2I^2/\epsilon))$$
$$\leq \frac{\epsilon^M}{(2I)^M I^M \log_{\frac{1}{2}}\epsilon} \leq \frac{\epsilon}{2IMI^M \log_{\frac{1}{2}}\epsilon} \tag{30}$$

The last inequality holds due to $(2I)^M \geq 2IM$, when I and M are large. Taking union bound over I^M distinct prices, M resources, and $2I$ time, and over the $\log_2 \epsilon^{-1}$ stages proves the lemma. □

APPENDIX C
PROOF OF LEMMA 9

Proof: Define the auxiliary primal and dual programs as follows.

$$P^{\psi}\text{ : maximize: } \sum_{i\in I_{\eta+1}}\sum_{\gamma\in\Gamma_i} b_{i\gamma}x_{i\gamma} \tag{31}$$

subject to :

$$\sum_{\gamma\in\Gamma} x_{i\gamma} \le 1, \quad \forall i \in I_{\eta+1} \tag{31a}$$

$$\sum_{i\in I_{\eta+1}}\frac{\tau_i}{T}\sum_{\gamma\in\Gamma_i} c^m_{i\gamma}x_{i\gamma} \le \psi_m, \quad \forall m\in\mathcal{M} \tag{31b}$$

$$x_{i\gamma}\in\{0,1\}, \quad \forall\gamma\in\Gamma_i,\ i\in I_{\eta+1} \tag{31c}$$

$$\text{where : } \psi_m = \begin{cases} \sum_{i\in I_{\eta+1}}\frac{\tau_i}{T}\sum_{\gamma\in\Gamma_i} c^m_{i\gamma}x_{i\gamma}(\mathbf{p}^\eta), \\ p^\eta_m > 0 & (31d)\\ \max\{\sum_{i\in I_{\eta+1}}\frac{\tau_i}{T}\sum_{\gamma\in\Gamma_i} c^m_{i\gamma}x_{i\gamma}(\mathbf{p}^\eta), \\ 2^{\eta+1}\epsilon C_m\},\ p^\eta_m = 0 & (31e)\end{cases}$$

Its dual program is:

$$D^{\psi}\text{ : minimize: } \sum_{m\in\mathcal{M}}\psi_m p_m + \sum_{i\in I_{\eta+1}} u_i \tag{31}$$

subject to :

$$u_i \ge b_{i\gamma} - \sum_{m\in\mathcal{M}} p_m\frac{\tau_i}{T}c^m_{i\gamma}, \quad \forall\gamma\in\Gamma_i, i\in I_{\eta+1} \tag{32a}$$

$$p_m \ge 0, \quad \forall m\in\mathcal{M} \tag{32b}$$

$$u_i \ge 0, \quad \forall i\in I_{\eta+1} \tag{32c}$$

Note that $\{x_{i\gamma}(\mathbf{p}^\eta)\}_{i\in I,\gamma\in\Gamma_i}$ and $\mathbf{p}^\eta$ satisfy the complementarity slackness conditions [24], and hence are optimal primal and dual solutions of the constructed linear programs (31) and (32). The remaining proof shows that with probability at least $1-\epsilon$, $(1-2\chi_\eta-\epsilon)\mathbf{x}^\delta$ is a feasible solution to the constructed program (31), and then the lemma follows. First, we show that with probability $1-\epsilon$, we have $\psi_m \ge (1-2\chi_\eta-\epsilon)2^{\eta+1}\epsilon C_m, \forall m\in\mathcal{M}$. Recall that $\mathbf{p}^\eta$ is the optimal dual solution of program (9). Let $\mathbf{x}^\eta$ be the optimal primal solution of program (8) (Recall $\mathbf{x}(\mathbf{p}^\eta)$ is the induced solution by optimal dual prices of P^η in (9)). Then, by the complementarity conditions [24], if $p^\eta_m > 0$, the m-dimension constraint must be satisfied by equality, *i.e.*, $\sum_{i\in I_\eta}\frac{\tau_i}{T}c^m_{i\gamma}x^\eta_{i\gamma} = (1-\chi_\eta)2^\eta\epsilon C_m$. Then, given the observation made in Lemma 8, and the condition $C_m \ge \frac{M}{\epsilon^2} \ge \frac{M}{2^\eta\epsilon}$ which follows the assumption of $\frac{\min_{m\in\mathcal{M}} C_m}{\max_{i,\gamma,m}|c_{im\gamma}|} \ge \frac{13M\log(\frac{2I^2}{\epsilon})}{\epsilon^2}$ given $\epsilon \le 1$, we have for any m,

$$\sum_{i\in I_\eta}\frac{\tau_i}{T}c^m_{i\gamma}x_{i\gamma}(\mathbf{p}^\eta) \ge \sum_{i\in I_\eta}\frac{\tau_i}{T}c^m_{i\gamma}x^\eta_{i\gamma} - M \ge (1-\chi_\eta-\epsilon)2^\eta\epsilon C_m. \tag{33}$$

Next, using the *Hoeffding-Bernstein's Inequality*, we show that with probability at least $1-\epsilon$, $\forall m\in\mathcal{M}$,

$$\psi_m = \sum_{i\in I_{\eta+1}}\frac{\tau_i}{T}\sum_{\gamma\in\Gamma_i} c^m_{i\gamma}x_{i\gamma}(\mathbf{p}^\eta) \ge (1-2\chi_\eta-\epsilon)2^{\eta+1}\epsilon C_m. \tag{34}$$

The proof of (34) is as follows. If $p^\eta_m = 0$, then by definition we have $\psi_m \ge C_m$ which satisfies (34). It remains to reason about the case where $p^\eta_m > 0$ that, for any fixed m, the probability of $\sum_{i\in\mathcal{D}} I_{\eta+1}\delta_i\frac{\tau_i}{T}c^m_{i\gamma}x_{i\gamma}(\mathbf{p}^\eta) < (1-2\chi_\eta-\epsilon)2^{\eta+1}\epsilon C_m$ is at most ϵ. Recall the definition of $\mathbf{X}_{im}$ in (11). We have

$$\begin{aligned}\sigma^2(\mathbf{X}) &\le \frac{1}{2^{\eta+1}\epsilon I}\sum_{i\in I_{\eta+1}}\mathbf{X}_{im} = \frac{1}{I}(1-2\chi_\eta-\epsilon)C_m,\\ \Delta(\mathbf{X}) &\le 2\max_{i\in I_{\eta+1}}|\mathbf{X}_{im}| \le 2\end{aligned} \tag{35}$$

According to (33), we have

$$\sum_{i\in I_\eta}\frac{\tau_i}{T}\sum_{\gamma\in\Gamma_i} c^m_{i\gamma}x_{i\gamma}(\mathbf{p}^\eta) \ge (1-\chi_\eta-\epsilon)2^\eta\epsilon C_m \tag{36}$$

Define events

$$\mathcal{G} = \{\sum_{i\in I_{\eta+1}}\frac{\tau_i}{T}\sum_{\gamma\in\Gamma_i} c^m_{i\gamma}x_{i\gamma}(\mathbf{p}^\eta) \le (1-2\chi_\eta-\epsilon)2^{\eta+1}\epsilon C_m\} \tag{37}$$

$$\mathcal{G}' = \{\sum_{i\in I_{\eta+1}}\frac{\tau_i}{T}\sum_{\gamma\in\Gamma_i} c^m_{i\gamma}x_{i\gamma}(\mathbf{p}^\eta) = (1-2\chi_\eta-\epsilon)2^{\eta+1}\epsilon C_m\} \tag{38}$$

Then, for a fixed m, and any distinct price vector $\mathbf{p}$, when $\mathbf{p} = \mathbf{p}^\eta$, we have

$$\begin{aligned}&\Pr[\sum_{i\in I_\eta}\mathbf{X}_{im} \ge (1-\chi_\eta-\epsilon)2^\eta\epsilon C_m, \mathcal{G}]\\ &\quad\le \Pr[\sum_{i\in I_\eta}\mathbf{X}_{im} \ge (1-\chi_\eta-\epsilon)2^\eta\epsilon C_m \mid \mathcal{G}']\end{aligned} \tag{39}$$

Let $\alpha = 2^\eta\epsilon(1-\chi_\eta-\epsilon-\frac{\lambda T}{I}(1-2\chi_\eta-\epsilon))$. Then we have

$$(39) \le \Pr[\sum_{i\in I_\eta}\mathbf{X}_{im} - \frac{\lambda T}{2I}\sum_{i\in I_{\eta+1}}\mathbf{X}_{im} \ge \alpha C_m \mid \mathcal{G}] \tag{40}$$

Since $\alpha = 2^\eta\epsilon(1-\chi_\eta-\epsilon-\frac{\lambda T}{I}(1-2\chi_\eta-\epsilon))$, and $\frac{1}{1+\chi_\eta/2} \le \frac{\lambda T}{I} \le \frac{1}{1-\chi_\eta/2}$, we have $2^\eta\epsilon\frac{(3-\epsilon)\chi_\eta-\chi^2_\eta}{2+\chi_\eta} \ge \alpha \ge 2^\eta\epsilon\frac{\chi_\eta(1+\epsilon)+\chi^2_\eta}{2-\chi_\eta} > \frac{2\epsilon^3}{1-\epsilon} > 0$. Putting (35) and $\mathbf{X}_{im}$ in (11) into (40), we have that (40) is at most

$$\begin{aligned}&\exp(\frac{-\alpha^2C^2_m}{2I_\eta\sigma^2(\mathbf{X}) + \alpha C_m\Delta(\mathbf{X})})\\ &\quad\le \exp(\frac{-\alpha^2C^2_m}{\frac{2^\eta\epsilon\lambda TC_m}{I}(1-2\chi_\eta-\epsilon) + 2\alpha C_m})\\ &\quad\le \exp(\frac{-2^\eta\epsilon(\frac{\chi_\eta(1+\epsilon)+\chi^2_\eta}{2-\chi_\eta})^2C^2_m}{\frac{(1-2\chi_\eta-\epsilon)C_m}{1-\frac{\chi_\eta}{2}} + \frac{2^\eta C_m[(3-\epsilon)\chi_\eta-\chi^2_\eta]}{2+\chi_\eta}})\end{aligned} \tag{41}$$

Given $\epsilon \leq \chi_\eta \leq \sqrt{\epsilon}$ and $1 \leq 2^\eta \leq \frac{1}{\epsilon}$, we have

$$2^\eta\epsilon(\frac{\chi_\eta(1+\epsilon)+\chi_\eta^2}{2-\chi_\eta})^2 \geq 2^\eta\epsilon(\frac{(\sqrt{\frac{\epsilon}{2^\eta}}+\frac{\epsilon}{2^\eta})}{2-\sqrt{\frac{\epsilon}{2^\eta}}})^2$$

$$\geq 2^\eta\epsilon(\frac{\sqrt{\frac{\epsilon}{2^\eta}}}{2-\epsilon})^2 \geq \frac{\epsilon^2}{4};$$

$$\frac{1}{1-\frac{\chi_\eta}{2}}(1-2\chi_\eta-\epsilon) \leq \frac{1-3\epsilon}{1-\frac{\sqrt{\epsilon}}{2}} \leq 2; \quad 2^\eta\frac{(3-\epsilon)\chi_\eta-\chi_\eta^2}{2+\chi_\eta}$$

$$\leq \frac{2^\eta 3\epsilon^{\frac{1}{2}}}{2+\sqrt{\frac{\epsilon}{2^\eta}}} \leq \frac{3\epsilon^{\frac{1}{2}}}{\frac{2}{\epsilon}+1} \leq 1 \tag{42}$$

Putting (42) into (41), we have that (41) $\leq \exp(\frac{-\epsilon^2 C_m/4}{2+1}) \leq \frac{\epsilon}{2IMI^M\log_{\frac{1}{2}}\epsilon}$, where the last inequality holds under the assumption that $C_m/\max_{i\in I}|c_{i\gamma}^m| \geq \frac{13M\log(2I^2/\epsilon)}{\epsilon^2}$, and $(2I)^M \geq \log_2 1/\epsilon$ which follows $\lambda T \geq 4/\epsilon^3$ and $I \geq \frac{1}{1+\chi_\eta}\lambda T \geq 1/2\lambda T$. Then we take union bound over each distinct $\mathbf{p}$ and the number of capacity constraints $2IM$, we have proved (34) in the case $p_m^\eta > 0$. Given that (i) constraints (8a) and (8c) are the same as (31a) and (31c) and (ii) constraints (8b) and (31b) which only differ in RHS are associated by the result of (34), we have that $(1-2\chi_\eta-\epsilon)\mathbf{x}^{\eta+1}$ is a feasible solution to (31). Therefore, the optimal objective value of (31), *i.e.,* $P^\psi(\mathbf{x}(\mathbf{p}^\eta))$, is at least the objective value of (31) under the solution of $(1-2\chi_\eta-\epsilon)\mathbf{x}^{\eta+1}$, *i.e.,* $P^\psi((1-2\chi_\eta-\epsilon)\mathbf{x}^{\eta+1})$. Let $\eta_{\max} = \log_2\epsilon^{-1}-1$ denote the maximal value in the range that η follows. We have the following lemma. □

Appendix D
Proof of Lemma 10

Proof: According to Lemma 9, we have (34) for all $\eta \in \{0,..,\log_2\epsilon^{-1}-1\}$. Thus we get that $\psi_m = \sum_{i\in I_{\eta_{max}+1}}\frac{\tau_i}{T}\sum_{\gamma\in\Gamma_i}c_{i\gamma}^m x_{i\gamma}(\mathbf{p}^{\eta_{\max}}) \geq (1-2\chi_{\eta_{\max}}-\epsilon)2^{\eta_{\max}+1}\epsilon C_m$. Since we have defined $I_{\eta+1} = I$ when $\eta = \eta_{max}$. And the expectation of the number of bids is equal to I. Thus, similarly as the end of Lemma 9, we have (i) constraints (6a) and (6c) are the same as (31a) and (31c) *in expectation* and (ii) constraints (6b) and the expectation of (31b) only differ in RHS are associated by the result of (34), we have that $(1-2\chi_\eta-\epsilon)\mathbf{x}^\delta$ is a feasible solution to (31) where $\mathbf{x}^\delta$ is the optimal solution of P^δ. Therefore, the optimal objective value of (31), *i.e.,* $P^\psi(\mathbf{x}(\mathbf{p}^\delta))$, is at least the objective value of (31) under the solution of $(1-2\chi_\eta-\epsilon)\mathbf{x}^\delta$, *i.e.,* $P^\psi((1-2\chi_\eta-\epsilon)\mathbf{x}^\delta)$. □

Appendix E
Proof of Lemma 11

Proof: Compare the dual programs (7) with (9). Since any realization of bid $i \in I_\eta$ can be found in the distribution $\mathcal{D}$, thus $(\mathbf{p}^\delta, \mathbf{u}^\delta)$ is a feasible solution to D^η. Further, since $(\mathbf{p}^\eta, \mathbf{u}^\eta)$ is the optimal solution to the minimization problem D^η, we have that $D^\eta(\mathbf{p}^\eta, \mathbf{u}^\eta) \leq D^\eta(\mathbf{p}^\delta, \mathbf{u}^\delta)$. Combined with weak duality, we have that $P^\eta(\mathbf{x}^\eta) \leq D^\eta(\mathbf{p}^\eta, \mathbf{u}^\eta) \leq D^\eta(\mathbf{p}^\delta, \mathbf{u}^\delta)$. Then we have

$$\begin{aligned}
E[D^\eta(\mathbf{p}^\eta, \mathbf{u}^\eta)] &\leq E[D^\eta(\mathbf{p}^\delta, \mathbf{u}^\delta)] \\
&= E[\sum_{m\in\mathcal{M}}(1-\epsilon\sqrt{\tfrac{\lambda T}{I_\eta}})\tfrac{I_\eta}{\lambda T}C_m p_m^\delta + \sum_{j\in I_\eta}u_j^\delta] \\
&\leq E[\sum_{m\in\mathcal{M}}(1-\sqrt{\tfrac{\epsilon}{2^\eta}})\tfrac{I_\eta}{\lambda T}C_m p_m^\delta + \sum_{j\in I_\eta}u_j^\delta] \\
&\leq \tfrac{I_\eta}{\lambda T}\sum_{m\in\mathcal{M}}C_m p_m^\delta + \sum_{j\in\mathcal{D}}I_\eta\delta_j u_j^\delta
\end{aligned} \tag{43}$$

Since $I_\eta = 2^\eta\epsilon\lambda T$ when $\eta = \{0,\cdots,\log_2\epsilon^{-1}-1\}$, and we have $(1-\epsilon/2)\lambda T \leq I \leq (1+\epsilon/2)\lambda T$, thus $\frac{I_\eta}{I} \leq 2^\eta\epsilon\frac{\lambda T}{I} \leq \frac{2^\eta\epsilon}{1-\epsilon/2}$. Therefore, (43) $\leq \frac{2^\eta\epsilon}{1-\epsilon/2}\big(\sum_{m\in\mathcal{M}}C_m p_m^\delta + \sum_{i\in\mathcal{D}}I\delta_j u_j^\delta\big) = \frac{2^\eta\epsilon}{1-\epsilon/2}D^\delta(\mathbf{p}^\delta, \mathbf{u}^\delta) = \frac{2^\eta\epsilon}{1-\epsilon/2}P^\delta(\mathbf{x}^\delta)$ Note that $P^\delta(\mathbf{x}^\delta) = P^{\delta\star}$, which the lemma follows. □

Appendix F
Proof of Theorem 2

Proof: Using Lemma 6, Lemma 9 and Lemma 1, we have that with probability of $(1-2\epsilon)\times(1-\epsilon) \geq 1-3\epsilon$, the events $\sum_{i\in I_{\eta+1}\setminus I_\eta}\mathbf{Y}_{im}(\omega) \leq 2^\eta\epsilon C_m$, $\sum_{i\in I_{\eta+1}}\sum_{\gamma\in\Gamma_i}b_{i\gamma}x_{i\gamma}(\mathbf{p}^\eta)) \geq (1-2\chi_\eta-\epsilon)P^{\star\eta+1}(\mathbf{x}^{\eta+1})$, $\sum_{i\in I_{\eta_{\max}+1}}\sum_{\gamma\in\Gamma_i}b_{i\gamma}x_{i\gamma}(\mathbf{p}^{\eta_{\max}})) \geq (1-2\chi_{\eta_{max}}-\epsilon)P^{\star\delta}$, and $P^{\star\eta}(\mathbf{x}^\eta) \leq \frac{2^\eta\epsilon}{1-\epsilon/2}P^{\delta\star}$ happen simultaneously for all $\eta \in \{0,...,\log_{\frac{1}{2}}\epsilon-1\}, m \in \mathcal{M}, \gamma \in \Gamma_i$. Let $\boldsymbol{\Omega}$ denote the event that the four events happen simultaneously. Then we have

$$\begin{aligned}
&E[\sum_\eta\sum_{i\in I_{\eta+1}\setminus I}\sum_{\gamma\in\Gamma_i}b_{i\gamma}x_{i\gamma}(\mathbf{p}^\eta) \mid \boldsymbol{\Omega}] \\
&\geq \sum_\eta E[\sum_{i\in I_{\eta+1}}b_{i\gamma_i}x_{i\gamma_i}(\mathbf{p}^\eta) \mid \boldsymbol{\Omega}] \\
&\quad - \sum_\eta E[\sum_{i\in I_\eta}b_{i\gamma_i}x_{i\gamma_i}(\mathbf{p}^\eta) \mid \boldsymbol{\Omega}] \\
&\geq \sum_\eta(1-2\chi_\eta-\epsilon)E[P^{\star\eta+1}(\mathbf{x}^{\eta+1}) \mid \boldsymbol{\Omega}] \\
&\quad - \sum_\eta E[P^{\star\eta}(\mathbf{x}^\eta) \mid \boldsymbol{\Omega}] \\
&\geq P^{\delta\star} - \frac{1}{\Pr[\omega]}(E[P^{\star\eta_{\min}}(\mathbf{x}^{\eta_{\min}})] \\
&\quad - \sum_\eta(2\chi_\eta+\epsilon)E[P^{\star\eta+1}(\mathbf{x}^{\eta+1})]) \\
&\geq (1-\epsilon)P^{\delta\star} - \frac{1}{1-3\epsilon} \\
&\quad \times \frac{(2\chi_\eta+\epsilon)(2^{\eta_{\min}}\epsilon+\sum_\eta 2^{\eta+1}\epsilon)P^{\delta\star}}{1-\frac{\epsilon}{2}} \qquad (44) \\
&\geq P^{\delta\star} - \frac{1}{1-3\epsilon}13\epsilon P^{\delta\star}\frac{1}{1-\frac{\epsilon}{2}} \\
&\quad \times (\sum_\eta 2^\eta\epsilon = 1-\epsilon, \sum_\eta\chi_\eta 2^\eta\epsilon \leq 2.5\epsilon) \qquad (45)
\end{aligned}$$

$$
\begin{aligned}
& E[\sum_{\eta} \sum_{i \in I_{\eta+1} \setminus I_{\eta}} b_{i\gamma_i} x_{i\gamma_i}(\mathbf{p}^{\eta})] \geq Pr[\Omega] \\
& \quad \times E[\sum_{\eta} \sum_{i \in I_{\eta+1} \setminus I} \sum_{\gamma \in \Gamma_i} b_{i\gamma} x_{i\gamma}(\mathbf{p}^{\eta}) \mid \Omega] \qquad (46) \\
& \geq (1-3\epsilon)\big(P^{\delta\star} - \frac{13\epsilon}{(1-3\epsilon)(1-\frac{\epsilon}{2})}\big) \\
& \geq (1-21\epsilon)P^{\delta\star} \quad (\epsilon \leq \frac{1}{2}) \qquad (47)
\end{aligned}
$$

□

References

[1] *Network Functions Virtualization—White Paper*, accessed on 2015. [Online]. Available: https://goo.gl/q2uBgy

[2] ETSI. *Network Functions Virtualisation (NFV): Terminology for Main Concepts in NFV*, accessed on 2016. [Online]. Available: http://goo.gl/P8zzsL

[3] S. Agrawal, Z. Wang, and Y. Ye, "A dynamic near-optimal algorithm for online linear programming," *Oper. Res.*, vol. 62, no. 4, pp. 876–890, 2014.

[4] J. Martins *et al.*, "ClickOS and the art of network function virtualization," in *Proc. USENIX Symp. NSDI*, 2014, pp. 459–473.

[5] V. Sekar, N. Egi, S. Ratnasamy, M. K. Reiter, and G. Shi, "Design and implementation of a consolidated middlebox architecture," in *Proc. 9th USENIX Symp. NSDI*, 2012, p. 24.

[6] Z. A. Qazi, C.-C. Tu, L. Chiang, R. Miao, V. Sekar, and M. Yu, "SIMPLE-fying middlebox policy enforcement using SDN," in *Proc. ACM SIGCOMM*, 2013, pp. 27–38.

[7] J. Sherry, S. Hasan, C. Scott, A. Krishnamurthy, S. Ratnasamy, and V. Sekar, "Making middleboxes someone else's problem: Network processing as a cloud service," in *Proc. ACM SIGCOMM*, 2012, pp. 13–24.

[8] H. Moens and F. de Turck, "VNF-P: A model for efficient placement of virtualized network functions," in *Proc. 10th Int. Conf. Netw. Service Manage. (CNSM)*, Nov. 2014, pp. 418–423.

[9] S. Mehraghdam, M. Keller, and H. Karl, "Specifying and placing chains of virtual network functions," in *Proc. IEEE CloudNet*, Oct. 2014, pp. 7–13.

[10] D. Dietrich, A. Abujoda, and P. Papadimitriou, "Network service embedding across multiple providers with nestor," in *Proc. IFIP Netw.*, May 2015, pp. 1–9.

[11] S. Zaman and D. Grosu, "Combinatorial auction-based allocation of virtual machine instances in clouds," *J. Parallel Distrib. Comput.*, vol. 73, no. 4, pp. 495–508, 2013.

[12] Q. Wang, K. Ren, and X. Meng, "When cloud meets eBay: Towards effective pricing for cloud computing," in *Proc. IEEE INFOCOM*, Mar. 2012, pp. 936–944.

[13] W. Shi, L. Zhang, C. Wu, Z. Li, and F. C. M. Lau, "An online auction framework for dynamic resource provisioning in cloud computing," in *Proc. ACM SIGMETRICS*, 2014, pp. 71–83.

[14] S. Gu, Z. Li, C. Wu, and C. Huang, "An efficient auction mechanism for service chains in the NFV market," in *Proc. IEEE INFOCOM*, Apr. 2016, pp. 1–9.

[15] P. R. Wurman, W. E. Walsh, and M. P. Wellman, "Flexible double auctions for electronic commerce: Theory and implementation," *Decision Support Syst.*, vol. 24, no. 1, pp. 17–27, 1998.

[16] J. Bredin and D. C. Parkes, "Models for truthful online double auctions," in *Proc. Uncertainty Artif. Intell.*, Jul. 2005, pp. 50–59.

[17] M. T. Hajiaghayi, "Online auctions with re-usable goods," in *Proc. ACM Conf. Electron. Commerce*, 2005, pp. 165–174.

[18] S. Wang, P. Xu, X. Xu, S. Tang, X. Li, and X. Liu, "TODA: Truthful online double auction for spectrum allocation in wireless networks," in *Proc. IEEE Symp. New Frontiers Dyn. Spectr.*, Apr. 2010, pp. 1–10.

[19] *Near-Optimal Online Algorithms for Dynamic Resource Allocation Problems*. accessed on 2016. [Online]. Available: http://goo.gl/Ur9bde

[20] R. Cohen, L. Lewin-Eytan, J. S. Naor, and D. Raz, "Near optimal placement of virtual network functions," in *Proc. IEEE INFOCOM*, May 2015, pp. 1346–1354.

[21] M. F. Bari, S. R. Chowdhury, R. Ahmed, and R. Boutaba. (Mar. 2015). "On orchestrating virtual network functions in NFV." [Online]. Available: https://arxiv.org/abs/1503.06377

[22] H. Yu, D. Zheng, B. Y. Zhao, and W. Zheng, "Understanding user behavior in large-scale video-on-demand systems," in *Proc. ACM SIGOPS/EuroSys*, 2006, pp. 333–344.

[23] *Poisson Point Process*. accessed on 2016. [Online]. Available: https://en.wikipedia.org/wiki/Poisson_point_process

[24] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.

[25] *Karmarkar Algorithm*. accessed on 2016. [Online]. Available: https://en.wikipedia.org/wiki/Karmarkar%27s_algorithm

[26] S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan, "Multi-parameter mechanism design and sequential posted pricing," in *Proc. ACM STOC*, 2010, pp. 311–320.

[27] *Chebyshev's Inequality*. accessed on 2016. [Online]. Available: https://en.wikipedia.org/wiki/Chebyshev%27s_inequality

[28] *Google Data Center Bandwidth Capacity*. accessed on 2017. [Online]. Available: http://goo.gl/zJvCTH

[29] G. Urdaneta, G. Pierre, and M. van Steen, "Wikipedia workload analysis for decentralized hosting," *Comput. Netw.*, vol. 53, no. 11, pp. 1830–1845, Jul. 2010.

[30] M. F. Bari, S. R. Chowdhury, R. Ahmed, and R. Boutaba. (2015). "On orchestrating virtual network functions in NFV." [Online]. Available: https://arxiv.org/abs/1503.06377

[31] *Amazon EC2 Instance Types*. [Online]. Available: http://aws.amazon.com/ec2/instance-types

[32] Z. Huang and A. Kim, "Welfare maximization with production costs: A primal dual approach," in *Proc. ACM SODA*, 2015, pp. 59–72.

Xiaoxi Zhang (S'15) received the B.E. degree from the Department of Electronics and Information Engineering, Huazhong University of Science and Technology, in 2013. She is currently pursuing the Ph.D. degree with the Department of Computer Science, The University of Hong Kong, China. Her research interests include cloud computing, network economics, online algorithms, and online learning.

Zhiyi Huang (M'16) received the B.E. degree from the Department of Computer Science and Technology, Tsinghua University, China, in 2008, and the Ph.D. degree from the Department of Computer and Information Science, University of Pennsylvania, USA, in 2013. He is currently an Assistant Professor of Computer Science with The University of Hong Kong, China. His research interests include theoretical computer science, algorithmic game theory, differential privacy, and online algorithms.

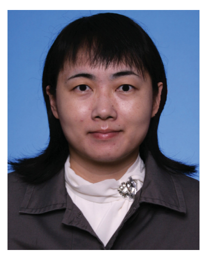

Chuan Wu (SM'16) received the B.E. and M.E. degrees from the Department of Computer Science and Technology, Tsinghua University, China, in 2000 and 2002, respectively, and the Ph.D. degree from the Department of Electrical and Computer Engineering, University of Toronto, Canada, in 2008. She is currently an Associate Professor with the Department of Computer Science, The University of Hong Kong, China. Her research interests include cloud computing and online/mobile social network.

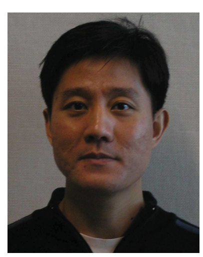

Zongpeng Li (SM'16) received the B.E. degree in computer science and technology from Tsinghua University, China, in 1999, and the M.S. degree in computer science and the Ph.D. degree in electrical and computer engineering from the University of Toronto in 2001 and 2005, respectively. He is currently a Professor with the Department of Computer Science, University of Calgary, Canada. His research interests are in computer networks, particularly in network optimization, multicast algorithm design, network game theory, and network coding.

Francis C. M. Lau (SM'17) received the Ph.D. degree from the Department of Computer Science, University of Waterloo. He is currently a Professor in computer science with The University of Hong Kong, China. His research interests include computer systems, networks, programming languages, and application of computing in arts. He is currently the Editor-in-Chief of the *Journal of Interconnection Networks*.