

Auction Mechanisms for Network Resource Allocation

Iordanis Koutsopoulos, George Iosifidis

► **To cite this version:**

Iordanis Koutsopoulos, George Iosifidis. Auction Mechanisms for Network Resource Allocation. WiOpt'10: Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks, May 2010, Avignon, France. pp.350-359, 2010. <inria-00502035>

HAL Id: inria-00502035

<https://hal.inria.fr/inria-00502035>

Submitted on 13 Jul 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Auction Mechanisms for Network Resource Allocation

Jordanis Koutsopoulos George Iosifidis

(Invited Paper)

Abstract—In the autonomic Internet of the future, auction mechanisms arise as key methods for realizing efficient resource allocation. The major asset of auctions is their obliviousness to node utilities, which renders them capable of achieving a desired resource allocation regime *without* knowledge of the utility functions of involved entities. Auctions can aid in addressing major research challenges in such autonomic settings, such as the need to cope with diverse and conflicting interests of network entities, the need to carry out resource allocation in a decentralized manner, the requirement for matching dynamic spatiotemporal patterns of demand and supply, and the need to operate under limited or no network state and utility information. In this survey paper, we delineate the main trends and challenges associated with auction design. We start from first principles auction design for maximum auctioneer revenue or maximum allocation efficiency for one or multiple indivisible items and for divisible resources. We gradually move to more composite models, those of position auctions for Internet advertisements and those arising in spectrum sharing in cognitive radio networks. We argue that some directions worth pursuing are: (i) the design of advanced auction models that capture multi-level interaction of involved entities, (ii) the employment of double auctions for multiple seller and buyer interaction, and (iii) the design of decentralized negotiation and resource trading mechanisms.

I. INTRODUCTION

The Internet of the future is pictured as a collection of diverse networked entities that engage transactions to accomplish certain objectives. These entities may be for example end-devices networked in ad-hoc manner, customer owned access points, residential femto-cell base stations or service providers and brokers and equipment belonging to them. The uncoordinated and inherently autonomic environment in which transactions occur creates several unprecedented challenges that need to be resolved in order to ensure viable and smooth operation.

First, future networks comprise diverse interacting *rational entities* with the natural propensity to solicit their own benefit and to strive to obtain maximum benefit from the network, while abstaining from any form of contribution to it. Entities are inclined towards misreporting local parameters that determine a socially optimal, global resource allocation regime. For example, they declare higher needs than real ones, in an effort to extract larger utility from apportioned resources. Such rational behaviors need to be understood through game theoretic

models and tamed through mechanisms that deter selfishness and promote good-will cooperation and truthfulness.

Second, the need for *decentralized optimal resource allocation* and control (where optimality may pertain to different operational objectives) is deemed more important than ever. Since the goods (resources) to be utilized reside ideally in a common pool and their management is to be realized autonomously without central brokers, self-management of network entities and resource trading arise as natural viable solutions. Third, there exists the requirement to *match resource supply and demand profiles* in time and space in the network. The dynamic, unpredictable spatiotemporal patterns of resource demand and availability, together with the need for online resource sharing and reallocation as requests arise and evolve, necessitates the adoption of dynamic, flexible resource management schemes.

Fourth and more importantly, control decisions have to be taken with *partial or no knowledge* of parameters of the associated optimization problem. Perfect global network state information may be too costly, impractical or simply impossible or meaningless to obtain. Rapid traffic load changes, interference, and topology or channel quality variations (the latter when it comes to wireless networks) render it difficult for individual nodes to obtain full view even of their own derived utility for different resource allocation regimes. Privacy concerns may also discourage a node from reporting its utility. Careful deliberation is needed to design mechanisms capable of handling such situations as well.

The Internet architecture is changing towards a federation of networked elements that come into spontaneous interaction, engage into mutual resource exchanges, reach agreements to promote their interests, and ultimately affect each other with their decisions. Furthermore, regulatory developments, liberalization of spectrum and of the regimes in which resource allocation takes place have intensified the trend towards consideration of a resource *market*. Each network entity possesses some amount of resource. All bring their resources to a common pool from where a certain resource allocation regime has to emerge. The network setting in which decentralized interactions occur obtains different twists depending on the nature of the resource. Below, we mention some of these paradigm settings, from the more classic to more contemporary ones:

- In *wireless ad-hoc networks*, nodes use their limited battery energy to transmit their own traffic to the next hop en route to the destination or to forward other nodes' traffic to their respective next hop. The underlying scarce resource

The authors are with the Department of Computer and Communications Engineering, University of Thessaly, Volos, GR 38221, Greece. Emails: {jordan, giosifid}@uth.gr

is energy and/or bandwidth, and the goal is to enforce a cooperation regime such that the total incurred cost in terms of energy or bandwidth consumption is minimized.

- In *wireless interference-limited regimes*, each user has a utility that is a function of received Signal-to-Interference plus Noise Ratio (SINR). The good to be allocated is transmit power, and this determines utility through mutual coupling due to interference. The objective is to reach a socially optimal objective.
- In *peer-to-peer networks*, peers exchange content over network overlays. The resource to be traded is the bandwidth of the link that connects the peer to the backbone, or equivalently the peer service time. Peers allocate their access link bandwidth either to their own downloads from others or to other peers that upload content from them. The goal is to achieve a network operating point that maximizes social welfare.
- In *disruption-tolerant (DT) or intermittently connected networks*, nodes store content in cases of interrupted connectivity and transmit it when link conditions allow it. The resource may be the cache memory for short-term storage, or the disk space for longer-term storage. Nodes need to decide how to allocate storage space to sporadically arising requests. The objective is to maximize utility or minimize delay in end-to-end data transfer.
- In *cognitive radio (CR) networks*, spectrum chunks need to be judiciously allocated among licensed primary providers or users, and secondary, unlicensed providers and users. The objectives may be to maximize revenue of primary providers or to reach a socially optimal spectrum allocation regime with high spectrum utilization.
- In *virtual server configurations*, the bottleneck resource is CPU. The problem is to migrate application processes from servers on which they physically run to virtual machines on other servers to improve overall server utilization. Virtual machines offer flexibility in hosting multiple processes on the same server but they consume CPU.
- In *distributed data centers*, data center storage space is again the critical resource to be shared among competing agents, but at a larger scale than in the DT example above.

Auction mechanisms migrated from economics and have been applied recently in network resource allocation [1], [2], [3, Ch.14]. The objectives range from revenue maximization for the auctioneer to social welfare maximization for the entities to which resources are allocated. Auction mechanisms start with a bidding from buyers, which is then mapped by the auctioneer to an allocation of good and a payment for each buyer. What makes auctions particularly attractive is that they are oblivious to individual buyer valuations about the resources to be auctioned. This leads to their ability to achieve a desired resource allocation regime *without* knowledge of the utility functions of involved entities. In envisioned network settings where operation under partial information is expected to be the rule, auction models for agent interaction and competition have a serious asset over more traditional market methods

such as pricing. Auctions can aid in addressing major research challenges, such as the need to cope with diverse and conflicting interests of network entities, the need to carry out resource allocation in a decentralized manner, the requirement for matching the dynamic spatiotemporal patterns of demand and supply in a network, and the need to operate under limited or no network state and node utility information.

In this survey paper, we attempt to delineate the main trends and challenges associated with auction design. In section II, we start from first principles auction design for maximum auctioneer revenue or maximum allocation efficiency for the case of one item. In section III we review multiple indivisible items and divisible resources. We gradually move to more composite auction models, namely those encountered in position auctions for internet advertisements (section IV) and those arising in spectrum sharing in cognitive radio networks (section V). In section VI we argue that some directions worth pursuing are: (i) the design of advanced auction models that capture multi-level interaction and dependence of involved entities, (ii) the employment of double auctions that model multiple seller and buyer interaction, and (iii) the design of decentralized negotiation and resource trading mechanisms.

II. SINGLE-ITEM AUCTIONS

In the simplest auction, there exists a seller that wants to sell one indivisible item and N rational potential buyers that bid to acquire the item, aiming at maximizing net payoff. Each bidder knows only his own valuation of the item and not those of others. Knowledge of other bidders valuations during the auction does not change the value of the item for a bidder. This model is one of *private values*. We consider the class of *standard auctions* where the item is given to the highest bidder.

The open ascending price or *English* auction is the most popular one. The auctioneer starts by announcing a low price and keeps increasing it in small steps as long as there are at least two interested bidders. The auction stops if there is only one bidder. In another variant, bidders progressively increase bid offers. The auction ends when only one bidder remains. That bidder wins the item, and he pays an amount equal to the price at which the second-last bidder dropped out. In the *Dutch* auction, the auctioneer starts by announcing a very high price at which none of the buyers is interested. He then progressively lowers the price until some bidder declares he is interested. That bidder wins the item at that price.

Consider now sealed-bid auctions in which a bidder does not see bids of others. In the sealed-bid *first price* auction, bidders submit bids in sealed envelopes. The buyer with the highest bid wins the item and pays the amount he bid. The sealed-bid *second price* (or Vickrey [4]) auction is similar, except that the highest bidder pays the second-highest bid. Under private values, the open Dutch auction is equivalent to the sealed-bid first price auction: bidding an amount in a first-price sealed-bid auction is the same as offering to buy at that amount in a Dutch auction provided the item is still available. An equivalence relationship holds also between the open English auction and the second-price sealed-bid auction [2].

The *Vickrey* auction has the desirable property that each bidder has no incentive not to bid its true valuation for the item. In other words, each bidder gains by truthfully declaring its true valuation for the item. To understand why this is true, consider bidder i with item valuation u_i . Let b denote the highest competing bid of others. Suppose first that $b < u_i$. If bidder i bids $b_i = u_i$, he wins with net payoff $u_i - b > 0$. He does not want to bid $b_i > u_i$, as then he wins but with negative net payoff. On the other hand, if his bid $b_i < u_i$ he reduces the chances of winning the auction and does not affect his net payoff if he wins (which is again $u_i - b$). So on average he reduces expected payoff. Now suppose $b > u_i$. If i 's bid is $b_i < u_i$ he does not win (payoff zero), while if $b_i > u_i$ bidder i may win the auction, but with net payoff $u_i - b < 0$. Thus again it is better to bid $b_i = u_i$ to lose the auction and have net payoff zero. Finally, if $b = u_i$, bidding $b_i = u_i$ does not make a difference from $b_i > u_i$ or $b_i < u_i$ (in all cases, net payoff is zero). Taking all into account, it is always to i 's benefit to bid $b_i = u_i$ regardless of the competing strategies.

A. Revenue and efficiency for some basic single-item auctions

Assume that each bidder $i = 1, \dots, N$ has valuation X_i for the item, where X_i is a random variable with cumulative distribution function (c.d.f) $F(\cdot)$ and probability density function (p.d.f) $f(\cdot)$, which are common knowledge to all, together with number N . Valuations are independent random variables. Functions $F(x) = \Pr(X_i < x)$ and $f(x) = F'(x)$ are the same for all $i = 1, \dots, N$ and defined at some interval $[0, w]$. This model is called one of symmetric bidders. Let x_i denote the realization of each X_i and let b_i be the bid of i . Assume that bidders are risk-neutral (see section II.C for a definition of risk-neutrality). Each bidder aims at optimizing its net payoff by adopting a bidding strategy $b_i(x_i)$, with $x_i \in [0, w]$.

1) *Second-price Auctions*: The net payoff $U_i(\cdot)$ of bidder i who participates in the auction with bid b_i is:

$$\phi_i(x_i, b_i) = \begin{cases} x_i - \max_{j \neq i} b_j, & \text{if } b_i > \max_{j \neq i} b_j \\ 0, & \text{else.} \end{cases}$$

It is optimal for each bidder to bid its valuation, i.e. $b_i(x_i) = x_i$ [2]. Let us compute the expected payment by a bidder. Fix a winner, say i . Call $X = X_i$ its random valuation, and let $Y = \max_{j \neq i} X_j$ be the second highest valuation (therefore, bid) that will be paid by i . Denote by $G(\cdot)$ and $g(\cdot)$ the c.d.f and p.d.f of Y . Suppose x is the winner valuation. We wish to compute the conditional p.d.f. of Y given that i wins, $g(y|Y < X, X = x)$. The conditional c.d.f. $G(y|x) = \Pr(Y \leq y|Y < X, X = x)$ is $\frac{\Pr(Y \leq y)}{\Pr(Y < x)}$ if $0 < y < x$, and it is 1, if $0 < x < y$. Then,

$$g(y|Y < X, X = x) = \frac{g(y)}{G(x)}, \quad 0 < y < x. \quad (1)$$

The conditional expected payment given that i wins is:

$$\mathbb{E}[Y|Y < X, X = x] = \frac{1}{G(x)} \int_0^x yg(y) dy.$$

The expected payment for valuation $X = x$ is therefore,

$$\mathbb{E}[Y|X = x] = \Pr(Y < x) \mathbb{E}[Y|Y < x] = \int_0^x yg(y) dy \quad (2)$$

with $g(y) = G'(y)$, $G(y) = F^{N-1}(y)$. One may also average over randomness of valuations to obtain the total average payment, $\mathbb{E}[Y] = \int_0^w \mathbb{E}[Y|X = x]f(x) dx$. Further, the total expected revenue of the seller is $N\mathbb{E}[Y]$.

2) *First-price Auctions*: The net payoff of a bidder is:

$$\phi_i(x_i, b_i) = \begin{cases} x_i - b_i, & \text{if } b_i > \max_{j \neq i} b_j \\ 0, & \text{else.} \end{cases}$$

In [2, Proposition 2.2] it is shown that the optimal symmetric bidding strategy is $b(x) = \mathbb{E}[Y|Y < X, X = x]$, where $Y = \max_{j \neq i} X_j$ as before. The expected payment to the seller for given winner valuation $X = x$, is $\Pr(Y < x) \mathbb{E}[Y|Y < X, X = x]$. This is equal to the expected payment for the second-price auction in (2). The same holds for the expected revenue to the seller.

It can be shown that the total expected revenue is equal to the expectation of the second highest valuation for both the first- and the second-price auction. This is known as the *Revenue Equivalence* principle and holds only for risk-neutral bidders and seller, and independent private valuations.

B. Auction design objectives

A first meaningful criterion for performance assessment of auctions is the incurred auctioneer revenue. The auction should be designed so as to increase competition, inducing bidders to participate and submit high bids and increasing expected price at which the item is sold. Another criterion is *efficiency* of the auction. For one item, this is equivalent to allocating it to the buyer who values it most. This instance arises when a governmental institution auctions a public good, and it is sought to allocate it to the most appropriate bidder. For multiple indivisible goods or one divisible good, efficiency is equivalent to maximizing social welfare incurred by the allocation. Maximizing auctioneer revenue while achieving high efficiency of the allocation may be conflicting objectives [5].

A method to increase revenue is the adoption of a *reserve price*, namely a minimum (publicly announced) price at which the item is sold. One may counterbalance the risk of not selling the item with the higher and the payment if the item is sold to compute the optimal reserve price that maximizes expected revenue [2, Ch.2.5]. It can be shown that the expected gain from setting a small reserve price exceeds expected loss. However, one should keep in mind that any effort to maximize revenue may have undesired effect on allocation efficiency.

For many indivisible goods or one divisible good, fairness is another objective, which is related to certain properties of the vector of allocated quantities or the vector of obtained utilities. Other auction design objectives are promotion of truthful reporting of bidder valuations, bidder attraction, discouragement of collusion and simplicity of mechanism [6, Ch.3].

C. Some auction classification

1) *Private versus interdependent values*: In an auction with *private values*, each bidder knows its own valuation of the item, but he does not know those of other bidders. If a statistical model for valuations is used in an auction with private values, a bidder knows the probability distribution of his own valuation and of valuations of others. In any case, knowledge to a bidder about other bidders' valuations does not affect his own. In auctions of *interdependent values*, each bidder may have full or partial information about its own valuation of the item, however this valuation can be affected by information available to other bidders. A special case is the *common value* model, where the unknown valuation is common for all bidders.

2) *Risk-averse versus risk-neutral seller and/or bidders*: A seller (or bidder) is *risk-averse* if its utility function $U(\cdot)$ is concave. Assume that a seller runs the auction K times. Say at the i -th time, the item is sold at price (bidder payment) p_i and the utility to the seller is $U(p_i)$. Risk averseness in its simplest form means that:

$$\frac{1}{K} \sum_{i=1}^K U(p_i) \leq U\left(\frac{1}{K} \sum_{i=1}^K p_i\right). \quad (3)$$

That is, the average utility from repeating the auction N times (and possibly with different payments) is *less* than the total utility derived with the average payment at all auctions. Namely, payment variability around the mean payment reduces derived utility. A risk-averse seller prefers an auction with more balanced payments, even if this leads to smaller average payments. Similarly, a bidder is risk-averse if its average utility of the difference between valuation and the bid is less than the utility of the difference between the valuation and average bid. If v is the item valuation (assumed to be fixed), b_i are his bids at different auctions times i , and $U(\cdot)$ is the utility function, risk-averseness means,

$$\frac{1}{K} \sum_{i=1}^K U(v - b_i) \leq U\left(v - \frac{1}{K} \sum_{i=1}^K b_i\right). \quad (4)$$

A risk-averse bidder prefers to have small average net gains (and thus to bid high on average), rather than having variable net gains (and thus bids). He prefers to win more frequently by bidding high even if his average net gain is smaller.

On the other hand, a seller (or bidder) is *risk-neutral* if its utility function is linear. Then (3) and (4) hold with equality, and variability around the mean does not reduce utility. A first-price auction among risk-averse bidders leads to higher expected revenue for the seller than a second-price auction. First-price auctions are more preferable for risk-averse sellers as well [2, Ch.4.1].

Other variants also exist, depending on whether bidders have budget constraints or not, and whether the distributions of bidders' valuations are correlated (affiliated) or independent.

III. MULTIPLE OBJECT AUCTIONS

In multiple object auctions, multiple items are to be sold. These auctions are classified as *homogeneous* (or *multi-unit*)

and *heterogeneous*, depending on whether items are units of the same good, or they are different goods. Homogeneous auctions may be *uniform-price* or *discriminatory-price* ones, depending on whether identical items are sold at the same price or not for different bidders. If items are auctioned one at a time as single-item auctions, the auction is called *sequential*. If all items are sold simultaneously, the auction is called *simultaneous*. Finally, auctions are *individual* if bidders can bid only at one item, and *combinatorial* if bids are allowed to combinations of items [3, Ch.14.2]. Here, we focus on homogeneous auctions.

A. Homogeneous sealed-bid Multi-unit auctions

Consider a simultaneous auction of K identical items to N bidders. Bidders submit bids for acquiring one or more items. Each bidder i submits a bid vector $\mathbf{b}^i = (b_1^i, b_2^i, \dots, b_K^i)$, such that $b_1^i \geq b_2^i \geq \dots \geq b_K^i$, where b_1^i is the amount i is willing to pay for receiving one item, b_2^i is the additional amount he is willing to pay for obtaining two units, and so on. Hence, the total amount that bidder i is willing to pay for obtaining $M \leq K$ items is $\sum_{j=1}^M b_j^i$.

1) *Discriminatory-price auction*: In discriminatory-price auctions, the allocation is as follows. Bids \mathbf{b}^i , $i = 1, \dots, N$ are ordered in decreasing order. The K highest bids (K_i of which refer to bidder i) are selected, and the K items are allocated so that bidder i obtains K_i of them. Each bidder i pays an amount equal to the sum of his bids that are deemed to be winning, $\sum_{j=1}^{K_i} b_j^i$. Each item is sold at different price.

2) *Uniform-price auction*: In uniform-price auctions, all K items are sold at a single price (the market clearing price) so that the total demand is equal to total supply. First, the number of items \tilde{K}_i that bidder i wins is computed as follows. For each bidder i with bid vector \mathbf{b}^i (bids in decreasing order), let \mathbf{c}^{-i} be the K -vector of competing bids for i . This is the vector of the highest K bids out of the bids of bidders other than i , arranged in decreasing order. Bidder i gets \tilde{K}_i items if its highest bid exceeds the lowest of the competing ones, the second highest bid exceeds the second lowest of competing ones, and so on until the \tilde{K}_i -th highest bid, but this does not hold for the $(\tilde{K}_i + 1)$ -th highest bid. The market clearing price turns out to be the highest losing bid over all bidders, $p = \max_i b_{\tilde{K}_i+1}^i$. Note that for $K = 1$, the uniform-price auction reduces to the second-price sealed-bid auction.

3) *Vickey auction*: In the Vickrey sealed-bid multi-unit auction, the method to determine the number of items \tilde{K}_i each bidder will obtain is the one above for uniform-price auction. A bidder who wins \tilde{K}_i units pays the sum of the \tilde{K}_i highest *losing* bids in \mathbf{c}^{-i} . These are found by removing winning bids of other bidders from \mathbf{c}^{-i} and selecting the \tilde{K}_i remaining ones. It can be seen that the amount that bidder i pays is equal to the *externality* it causes to other bidders. The externality in this case is the additional amount that other bidders would pay in the allocation, had bidder i been absent.

B. Auctions for a divisible resource

Multi-unit auction models also capture auctions of a single divisible good. Each bidder i submits a continuous bid function

$b_i(x)$ that indicates the amount he is willing to pay for resource amount x . Such a scenario is encountered in network resource sharing, where the good may be link bandwidth, power, energy or another type of resource.

An amount C of divisible resource is to be allocated among N users. Each user i is characterized by a strictly concave, increasing, continuous differentiable utility function $U_i(\cdot)$ which is only privately known to him but unknown to the allocation controller. Let x_i be the amount of good allocated to user i and $\mathbf{x} = (x_1, \dots, x_N)$ be an allocation vector. The social welfare maximization (SWM) problem is:

$$\max_{\mathbf{x} \geq \mathbf{0}} \sum_{i=1}^N U_i(x_i) \quad (5)$$

subject to:

$$\sum_{i=1}^N x_i = C. \quad (6)$$

If utility functions were known to the controller, the KKT conditions would give the necessary and sufficient conditions for the optimal allocation, $U_i'(x_i^*) = \lambda^*$ if $x_i^* > 0$, and $U_i'(0) \leq \lambda^*$ if $x_i = 0$, for $i = 1, \dots, N$, where λ^* is the optimal Lagrange multiplier for (6). Without loss of generality, assume $U_i(0) \rightarrow -\infty$, so that $x_i^* > 0$.

1) *Kelly mechanism*: Assume that the controller does not know utility functions $U_i(\cdot)$ but aims at socially optimal allocation. Consider the class of allocation mechanisms where each user submits a bid $b_i \geq 0$ for the amount he is willing to pay and is charged according to function $c(\cdot)$. The amount of allocated good, $x_i(b_i)$ is a function of their bid. Specifically, let $x_i(b_i) = b_i/\tilde{\lambda}$, where $\tilde{\lambda}$ is a price per unit of resource. We assume users are *price takers*, namely they do not consider the impact of their bid on the charge function $c(\cdot)$. It is reasonable to assume that each user is rational and casts his bid so as to maximize his net benefit, $U_i(x_i(b_i)) - c(b_i)$, namely his bid should satisfy:

$$U_i'(x_i^*) \frac{1}{\tilde{\lambda}} - c'(b_i) = 0. \quad (7)$$

Suppose the controller obtains bids \tilde{b}_i and makes the allocation according to the solution of the following problem (P):

$$\max_{\mathbf{x} \geq \mathbf{0}} \sum_{i=1}^N \tilde{b}_i \log x_i, \quad (8)$$

subject to $\sum_{i=1}^N x_i = C$, and $x_i \geq 0$, $i = 1, \dots, N$. The KKT conditions for this problem give:

$$\frac{\tilde{b}_i}{x_i} = \tilde{\lambda}, \quad (9)$$

where $\tilde{\lambda}$, $\tilde{\mathbf{x}}$, is the optimal Lagrange multiplier and the optimal solution respectively of (P). The goal is to equalize the solutions of optimization problems (SWM) and (P). It turns out that if each user is charged according to $c(b_i) = b_i$, then from (7), (9) it is $\lambda^* = \tilde{\lambda}$ and $x_i^* = \tilde{x}_i$, which gives $\tilde{b}_i = x_i^* U_i'(x_i^*)$.

Since the optimal solution to (P) should satisfy $\sum_{i=1}^N \tilde{x}_i = C$, by using (9) we get $\tilde{\lambda} = \lambda^* = \frac{1}{C} \sum_{i=1}^N \tilde{b}_i$. This is the market clearing price, set by the controller. Furthermore,

$$\tilde{x}_i = \frac{\tilde{b}_i}{\sum_{i=1}^N \tilde{b}_i} C, \quad (10)$$

namely the allocated amount to each user is proportional to its bid [7]. Therefore, socially optimal resource allocation can be achieved by bidding (where each user's bid is a single number), and an appropriate charging scheme.

Kelly *et al.* proposed this mechanism and showed that the problem above can be solved in a decentralized fashion [8]. The market clearing price $\lambda^{(n)}$ is iteratively computed at each step n by the auctioneer according to a standard dual algorithm. Essentially, it is increased or decreased, depending on whether the instantaneous allocation exceeds C or not. Then, each user adjusts its bid according to $U_i'(\frac{b_i}{\lambda^{(n)}}) = \lambda^{(n)}$. The dual price update together with the user response converges to the optimal solution of the network utility maximization problem. This algorithm is a distributed implementation of the bidding mechanism. The moral of the story is that for price-taking users, one-dimensional bids and appropriate charging lead to efficient allocation.

2) *Vickrey-Clark-Groves (VCG) mechanism*: Consider now achieving an efficient allocation if users are *price-anticipating*, namely they strategically adapt their bid by taking into account its impact on the price so that they maximize net profit. In that case, a game interaction emerges with certain efficiency loss. The setup is the same as the one above, and each user chooses his bid to maximize the quantity:

$$U_i\left(\frac{b_i}{\sum_{j=1}^N b_j} C\right) - b_i. \quad (11)$$

Notice that now user i explicitly understands that the price, $\frac{1}{C} \sum_{i=1}^N b_i$ depends also on its own bid b_i . A mechanism that guarantees an efficient allocation for selfish, price-anticipating users is the Vickrey-Clarke-Groves (VCG) mechanism [9],[10]. This is a generalization of the Vickrey mechanism for single item auctions. Here, the compromise is that the auctioneer requests each user to *reveal* its utility function. In the VCG mechanism, the amount charged to each user i is the externality it causes to others. This is the total utility reduction caused by i to all other users, and it is computed as follows. Let \mathbf{x}^* be the optimal solution to (SWM) problem, and let $\tilde{\mathbf{x}}$ be the optimal solution to the (SWM) problem *without* considering the effect of user i , namely to problem $\max_{\mathbf{x}} \sum_{j \neq i}^N U_j(x_j)$, such that $\sum_{j \neq i}^N x_j = C$. The charge to user i is:

$$p_i = \sum_{j \neq i}^N U_j(\tilde{x}_j) - \sum_{j \neq i}^N U_j(x_j^*). \quad (12)$$

In the VCG mechanism, declaration of the true utility function $U_i(\cdot)$ is the best strategy for each user [11, Ch.6]. Namely, a user i cannot do better by misreporting its utility function.

To see this, observe that the net profit for a user i that declares its true utility function is,

$$U_i(x_i^*) - p_i = \sum_{i=1}^N U_i(x_i^*) - \sum_{j \neq i} U_j(\bar{x}_j). \quad (13)$$

Suppose now that user i misreported its utility function and declared it as $\tilde{U}_i(\cdot)$ in an effort to get more profit. In that case, there would be a different solution (call it $\tilde{\mathbf{x}}$) to the (SWM) problem, and the profit of user i would be

$$U_i(\tilde{x}_i) - \tilde{p}_i = \sum_{i=1}^N U_i(\tilde{x}_i) - \sum_{j \neq i} U_j(\bar{x}_j). \quad (14)$$

If truthful reporting of utility were not optimal, (14) should exceed (13), which would mean $\sum_{i=1}^N U_i(\tilde{x}_i) > \sum_{i=1}^N U_i(x_i^*)$. This contradicts the fact that \mathbf{x}^* is the optimal solution of the (SWM) problem. Thus, truthful reporting is optimal under VCG. The VCG mechanism leads to efficient allocation. Its clear drawback is that each user needs to submit to the auctioneer its entire utility function, namely an infinitely dimensional vector, which renders the mechanism quite complex and burdensome in terms of information exchange.

C. Other related work

For price-taking users, the work [12] introduced an auction mechanism for efficient capacity allocation to network flows. Users submit bids, they receive a resource amount that is proportional to their bid and they pay an amount equal to their bid. The authors show existence of Nash equilibrium and propose decentralized bidding mechanisms that converge to the equilibrium. On the other hand, for price anticipating users that strategically adapt their bid, a game interaction emerges with efficiency loss (price of anarchy) at most 25% [13].

The implementation of VCG-like auctions has become a central research topic. Various studies seek to overcome the difficulty by proposing a combination of VCG and the proportional allocation method. In the work [14], which was generalized in [15], the authors propose a class of one-dimensional bid and proportional allocation methods which lead to a unique efficient Nash equilibrium for the case of one divisible good. In [16], [17], users submit one-dimensional bids and are charged according to the rules of the VCG scheme. Specifically, in [16], a family of surrogate functions are used by the auctioneer, which play a role similar to that of $\sum_i b_i \log x_i$ in the Kelly mechanism. Efficiency of equilibrium is proved, to the expense of losing the truthful reporting incentive of users. In other works [18], [19], two-dimensional bids (a per-unit price and the maximum amount of resource the user is willing to buy) are submitted. The charging is performed as in VCG auctions, and the allocation is according to the total utility maximization problem. Some works generalize these mechanisms for multiple divisible goods, each of which may stand for link bandwidth in a network. Each user requires bandwidth on a set of links that constitute its path [20], [21]. In the work [22], simultaneous multi-unit Dutch auctions are run for each link, and the charging is made according to the VCG mechanism.

IV. SPONSORED SEARCH AUCTIONS

We now consider internet advertising (ad) auctions, initiated by web search engines. The terms “keyword” or “sponsored search” auctions are also coined. Here, bidders are advertisers who wish to have the advertisement of their enterprise appear on a user’s search results screen after the user types a related keyword. When they register their ad with the search engine, they provide keywords related to their ad. Following a keyword search by an internet user, the system finds a set of ads with keywords that match the user query. Advertisements appear in the search results as a ranked list. The user clicks on an advertisement and it is taken to the advertiser’s website. The advertiser then pays the search engine company for guiding the user to its web page. Advertisement positions (ranked slots) on search results are clearly of high importance to advertisements; the higher the ad is displayed on the list, the more probable it is that it will be clicked by the user, and the more likely it becomes that the advertiser will get some profit if the user buys the product or service. In ad auctions, ranked ad positions are auctioned to advertisers. An advertisement is considered successful if a user clicks on the respective ad link. Advertisers pay an amount each time a user clicks on their ad.

The underlying feature of previously presented auctions is that there exist two parties, the auctioneer and the bidders, who determine the rules of the auction. Bidders cast their bids, and the auctioneer determines the allocation of items and the payment. Clearly, ad auctions are different in that the auctioneer revenue and bidder payoff depend on a third entity, the internet user. Google, Yahoo! and other search engines auction advertising positions. Google [23] was the first to consider the dependency of position allocation and payment on user preferences.

Consider N advertisers who bid for $K < N$ ad slots for a specific keyword. Let u_i be the value of the ad for advertiser i , $i = 1, \dots, N$. Let b_i be the bid of advertiser i and p_i be the payment per click he will be charged. The auctioneer collects submitted bids and needs to decide which bidders will have their ads shown, in which order and the respective payments. Let c_{ij} denote the probability that the ad of advertiser i will be clicked by the user when in position j , $j = 1, \dots, K$. This is also called click-through rate (CTR) and can be calculated by the search engine based on history statistics with various methods [24, Ch.3]. CTR depends on the ad of advertiser i and the position j and can be assumed to be $c_{ij} = \alpha_i \beta_j$, where α_i an ad-dependent parameter, the per-ad CTR (the ratio of number of clicks received by the ad over the number of times the ad was displayed). It is $\alpha_i = \sum_{j=1}^K c_{ij} y_{ij}$, where y_{ij} is the probability that ad i is displayed in position j . Also, β_j is a position-dependent parameter, the per-position CTR. Higher ranked positions are more visible to users and attract attention more, so that $\beta_1 > \dots > \beta_K$.

The auction goes as follows. Each advertiser i chooses a bid b_i . Ads of advertisers appear in ad slots in decreasing order of their weighted bid, $b_i \alpha_i$. The advertiser in the k -th position, say with weighted bid $b_{(k)} \alpha_{(k)}$ pays a total amount equal to

the weighted bid of the advertiser in the next position $k + 1$, that is, total amount $b_{(k+1)}\alpha_{(k+1)}$. Hence, the amount paid per click is $p_{(k)} = b_{(k+1)}\alpha_{(k+1)}/\alpha_{(k)}$. The last ranked advertiser either pays a reserve price if $N < K$ or the amount of bid of the first omitted advertiser if $N > K$. This payment rule is a generalization of the one in Vickrey auction for one item, generalized to the setting where a set of ranked items are sold. Thus, the auction is often referred as *Generalized Second Price* (GSP) auction [25], [26]. The position allocation rule naturally ranks bidders in a decreasing order of expected revenues.

The probability that a user will click on an ad is a key factor to consider. Otherwise, less attractive ads will be displayed, and small revenue for the auctioneer will be incurred. Assume that bidders are risk-neutral. The net payoff for advertiser i when his ad is displayed in position j is $c_{ij}(u_i - p_i)$, where p_i is the payment per click. From the perspective of the auctioneer, the problem is to find the position allocation that maximizes expected revenue [24]:

$$\max_{\mathbf{X}} \sum_{i=1}^N \sum_{j=1}^K p_i c_{ij} x_{ij} \quad (15)$$

subject to $\sum_{j=1}^K x_{ij} \leq 1$, for $i = 1, \dots, N$, and $\sum_{i=1}^N x_{ij} = 1$ for $j = 1, \dots, K$, where \mathbf{X} is the $N \times K$ assignment matrix with $x_{ij} = 1$ if the i -th advertiser is allotted position j , and $x_{ij} = 0$ otherwise. If the auctioneer wants to maximize efficiency, he solves the assignment problem:

$$\max_{\mathbf{X}} \sum_{i=1}^N \sum_{j=1}^K u_i c_{ij} x_{ij} \quad (16)$$

subject to the allocation constraints above. The auctioneer does not know valuations u_i and uses submitted bids, which may differ from valuations. For both revenue maximization and allocation efficiency, it is crucial to consider CTRs. The allocation of a slot does not generate utility for the advertiser and revenue for the auctioneer unless the ad link is clicked.

A. Other related work

In the first ad auction by Overture, the payment scheme was the *Generalized First Price* (GFP) auction, where advertisers paid their bid. Soon it was realized that GFP is a non-truthful auction mechanism with unstable equilibria in terms of bidding strategies. GSP auctions were initially proposed by Google and also adopted by Yahoo!. GSP admit pure Nash equilibria but they do not have the desirable property of truthful declaration of valuation as the Vickrey auction does [25]. In [24, Ch. 3.8], it is shown that for risk-neutral users, the GSP auction produces the same expected revenue as VCG where payment is according to bidder externalities. In [27] the structure of payments of the GSP auction is exploited to show that finding the equilibrium position prices is closely related to finding a solution to the assignment problem [28] in a decentralized manner. In [29], a novel scheme termed the *laddered* auction ensures incentive compatibility and does not decrease revenue of the GSP auction under certain assumptions. A survey of static and dynamic aspects of position allocation is presented in [30, Ch.28].

V. SPECTRUM AUCTIONS

Auctions are used for more than two decades now for licensing spectrum chunks [31]. Soon it was realized that there was a unbridged gap between increasing operator and user demands and the static spectrum supply. The latter led to serious spectrum underutilization in time and in spatial locations. The benefits of dynamic resource allocation actually manifested themselves first in dynamic channel allocation (DCA) schemes in cellular systems [32]. Instead of allocating a fixed set of channels (time slots or frequencies) to each cell and limit subscribers in that cell to these channels, all channels were brought in a common pool, and any user in any cell could be allocated any channel, provided that cochannel interference constraints were satisfied. DCA schemes could handle much larger load and could better control cell load variations than their fixed allocation counterparts.

Coexistence of wireless operators and services in overlapping spectrum bands brought dynamic spectrum access in stage [33]. Unused licensed spectrum bands of an operator are borrowed to others to match their high demands. In addition, license-exempt devices could opportunistically access unused licensed spectrum, owned by so called primary operators (POs). The latter gave rise to cognitive radio networks [34] and the IEEE 802.22 standard for underutilized TV spectrum. POs serve primary users and coordinate unused spectrum leasing to secondary operators (SOs). These in turn allocate spectrum to secondary users in their range. Recently, the concept of mobile virtual network (MVN) operators appeared. MVN operators serve their clientele with spectrum leased from primary operators. The goal is to improve spectrum utilization by allowing reselling and leasing of temporarily unoccupied spectrum.

Auctions are a promising solution for regulating spectrum allocation. The utility derived by an allocation for a certain entity is usually expressed in terms of transmission rate or delay. For spectrum allocation from a PO to SOs, the utility derived by the latter is difficult to estimate, since it depends on unpredictable demands, channel quality experienced by served users, or geographic range of the SO. The same problem arises if spectrum is auctioned by a SO to competing users. Time-varying demands and traffic loads, and interference from spectrum reuse makes it cumbersome to compute user utility.

Spectrum auctions substantially differ from conventional ones. First, spectrum can be *reused* by operators or users at sufficiently far locations. The auction should consider the spatial dimension and the fact that there may be many winners to which spectrum should be concurrently allocated. Spectrum reuse gives rise to *radio interference* due to mutual coupling of cochannel transmissions, and this in turn restricts the set of feasible channel allocations. Second, spectrum bands *differ* in terms of quality. Besides channel reuse, spectrum bands have different quality due to inherent frequency selectivity of the wireless channel, but also due to time-varying link quality and fading. Third, *heterogeneity* and *unpredictability* of user demands and mobility place additional challenges. Fourth, the *small-scale, online* fashion in which spectrum allocation needs

to operate renders the machinery of bidding, allocation and payment a challenging task. Finally, there exists a wide range of methods for the allocation, each with its own assets. For example, spectrum can be divided in variable size bundles. There may be one or many auctioneers who may even compete. Or, spectrum may be exchanged with double auctions (see section VI.B for more).

A. Related work

One of the first works on auctions for spectrum sharing amidst interference is [35]. Buyers are users that access a common band and transmit to a common receiver at the uplink. The controller auctions transmit power to transmitters so that a maximum total allowable received power constraint is satisfied, and potential primary transmission is not hindered. This latter constraint is similar to the allocation constraint $\sum_i x_i = C$ when allocating a divisible resource. This maximum received power is then allocated to users with an SINR or a power auction. An one-dimensional bidding method and a proportional power allocation achieve allocations close to those achieved by VCG auction. This work was extended in [36] for an interference channel with many transmitters and receivers. The method is clearly similar in spirit with that in [16]. In [37], sequential spectrum and power auctions are studied. At each round, the auctioneer considers each spectrum band and performs bandwidth auction by allocating different amounts to users in FDMA fashion or power auction by allocating transmit powers as in [35], [36]. The second price method is used. The worst-case efficiency loss due to selfish bidding is characterized for different utility functions.

In [38] a truthful, computationally manageable auction for spectrum allocation is proposed, which overcomes the burden of NP-Completeness of optimal allocation. A channel allocation based on a ranked list of bids is implemented, together with a charging scheme that considers the most crucial bidder in one's neighborhood. It is shown that auctions such as VCG and second price ones lose the truthfulness property when used for spectrum allocation. In [39] the authors present a secondary spectrum market where POs lease unused spectrum to SOs to maximize their revenue. SOs place bids for different number of channels in different cells. A VCG-like mechanism is adopted with different reservation prices for SOs. The notion of virtual valuation is introduced and used as weighting factor in the objective function. This assumes knowledge of distributions of user valuations. The main contribution of [40] where many identical channels are auctioned, is that bidders express the desired spectrum quantity at each per-unit price by using continuous piece-wise linear demand curves.

In [41] spectrum is auctioned to networks that can be primary or secondary on different bands. Each network submits bids for primary or secondary rights on each channel based on perceived valuations for each channel and traffic demands. The regulator allocates primary and secondary rights to each network on each channel based on bids. The goal is to find a channel configuration that maximizes revenue or social welfare. The

authors reduce the allocation problem for maximizing revenue to a maximum weighted matching one, if at most one secondary network is allowed on each channel. For more general scenarios they show NP-Completeness. In [42], an auction model of one or more bands to secondary users (SUs) is studied. Each SU can choose if it will participate in the auction and can decide which channel to sense (in case of multiple ones). There exists a participation cost and a channel sensing cost. At each stage, an SU decides on its actions based on bidding history and instantaneous channel valuation. For multiple channels the equilibrium is to choose a channel with probability proportional to the regret of not having chosen that channel in the past.

Situations with many spectrum providers and consumers are best modeled by double auctions. The problem for an outside auctioneer is to determine winning sellers and the channels they sell, determine winning buyers, allocate channels from sellers to buyers and compute the prices to pay to sellers and charge to buyers. The authors in [43] propose a variant of the McAfee double auction model [44] that is amenable to spectrum reuse. The work in [45] assumes that sellers can coalesce to increase their profit and studies competition of buyers in such scenarios. An overview of game theoretic mechanism design methods for spectrum allocation, including auctions and double auctions is presented in [46].

VI. FUTURE DIRECTIONS AND CHALLENGES

In the previous sections we described some basic auction models for resource allocation. The oftentimes diverse network architectural structures call for advanced auction models which capture salient features of the allocation problem. These features are not addressed by the traditional auction models described above. In this final section we outline some research directions on auction design which are worth pursuing.

A. Advanced auction models for resource allocation

One of the important facets in resource allocation is *heterogeneity* of involved entities, often expressed through certain hierarchies. For example, in spectrum allocation in cognitive radio networks, there exists a primary operator (PO) with a license for some spectrum chunks that serve the needs of its customers. In order to increase revenue, the operator may wish to lease surplus spectrum, namely spectrum that is unutilized by its own customers. The primary operator initiates an auction to decide how to lease spectrum to a set of secondary operators (SOs), each of which has a local range and clientele. At a first view, a conventional auction might seem to be sufficient. The PO could initiate an auction on spectrum chunks it wishes to allocate to SOs. Each SO submits a bid for each band or a set of bands, and subsequently the bid vector is mapped to an allocation and an payment rule. Discriminatory-price, uniform-price or Vickrey auctions for multiple object auctions could be used. In these auctions, the allocation and payment would depend only on the bid vector.

What is *not* actually captured by the auction model above is the fact that the appropriateness of a specific frequency

allocation depends also on the *experience* of the end user which will be served by the SO through frequency allocation. In fact, since end users served by the SO will most likely be clients of the PO as well, the PO would like to devise an auction that considers *user feedback* besides SO valuations. If viewed within the classical auction model, this enhanced model corresponds to one where auctioned items are not simply allocated to one of the bidders, but they are used or sold further. In other words, there exists an end-user *substrate* in addition to the auctioneer and the bidder substrates, and feedback from the user substrate is important to consider.

Clearly, it would not make much sense to allocate a frequency to a SO which would in turn assign it to a user for which the specific frequency is of low quality, either due to interference from excessive frequency reuse or due to limited range of that SO. The frequency allocation should *modulate* the SO bid with user experience. With an appropriate frequency selection, end users would be satisfied and willing to use the PO for other services as well. The PO would benefit if users preferred him instead of other POs. The SO would also benefit if users with a freedom of choice selected him over other SOs. This versatility of users to select among several SOs gives yet another degree of freedom to the user and an interesting twist to the problem. In addition, novel concepts of efficiency might need to be defined in that case.

Other instances of hierarchical resource allocation paradigms emerge in the case of leasing of excess storage space by a storage facility operator to a set of interested proxy parties which themselves need to serve their clients. In general, any instance of resource allocation where involved entities reside in different substrates that interact with each other falls within context of this model.

B. Double auctions

In autonomic networks comprising self-interested nodes, each with different needs and utilities, each entity has some resource in its possession and engages in transactions with others to achieve its objectives. The migration from classic all-to-one client-server architectures to decentralized multi-party interaction creates the need for regulating resource exchange through certain rules, in much the same way rules are devised to guide the economic market to desirable operating points. In such network settings, each node can be at the same time resource provider and consumer. For example, in cognitive radio networks, service providers with relatively low spectrum demand from their users may wish to sell surplus spectrum to other providers, which are themselves in need of spectrum due to higher demand. In peer-to-peer overlay networks, peers use their access link bandwidth to download content from other peers or to let other peers upload content from them. The traded resource is link bandwidth or equivalently service time. Whenever a peer does not have high content demands, it may lease link bandwidth to others which are interested in content that the peer has. Along the same lines, energy, power, storage space or CPU speed resources can be traded. At a macroscopic

level, the goal is to match spatiotemporally varying demand and supply patterns by resource trading.

In *double auctions*, potential buyers contend to obtain resources by placing bids to several sellers. They may place different bids to different sellers for different goods. Sellers that wish to sell (part of) their resource announce ask bids at which they sell it. Multiple sellers compete to attract buyers to buy their resource. Efficient and viable network operation relies precisely on node synergy and multi-lateral resource trading. Nodes face the dilemma of devoting their limited resource to their own benefit and thereby directly gaining utility, versus acting altruistically, with the anticipation to be aided themselves when they need to. For instance, in wireless ad-hoc networks, nodes use their limited battery energy to transmit their own traffic to the next hop en route to the destination or to forward other nodes' traffic to their respective next hop. The underlying scarce resource is energy. A node clearly benefits only if it uses its energy to transmit its own traffic.

The main challenge is to understand the dynamics of interactions of multiple entities and to manage a large volume of resource requests across several geographic areas by opportunistically exploiting spectrum surpluses and needs. In this way, resource supply and demand will be multiplexed in time and space to achieve best resource utilization. First, designing appropriate double auction models entails devising contention resolution schemes through novel allocation and payment methods for selling and buying resources. In double auctions, there exist multiple buyers and sellers, and therefore attempting to maximize revenue for all is likely to lead to a tragedy of commons. Instead, what is needed is a careful reconsideration of the *properties* of the operating points and the real-life scenarios that they capture. For instance, *resource reciprocity* is a meaningful such property which should be reflected by the operating point. Second, the *interdependent provider and consumer roles* and the way the one restrains the other need to be taken explicitly into account. That leads to an inherent constraint imposed in buying and selling resource and has repercussions in bidding strategies for buying and selling items. A first attempt to understand the dual provider-consumer role through double auctions appeared in [47].

Furthermore, *behavioral profiles* of involved entities need further study. Risk-averse, risk-neutral or risk-seeking entities lead to different resource exchanges and operating points. The impact of the portion of risk-averseness or risk-neutrality needs to be studied. Another interesting question concerns the impact of knowledge about the risk of some entities onto strategic behavior of others.

C. Negotiation and trading mechanisms

Auction mechanisms involve a seller and multiple buyers. There exists limited interaction among the seller and buyers and no interaction among buyers. These constraints can be relaxed if we allow negotiation. Different parties that engage in negotiation have conflicting interests, and the goal is to reach a mutually desirable resource sharing regime. The entities

are in principle cooperative since they agree to participate in negotiation. While in negotiation, an entity attempts to influence the outcome by making offers to others. Other involved parties may reject or accept the offer, or they may make alternative offers themselves. The simpler instance of negotiation arises between a seller and a buyer [48]. Different assumptions can be made about the information a party has about attributes of the other. There already exists mature theory in economics (see [12], [49] and references therein) but no related work yet in network resource allocation. Still, negotiation and decentralized trading capture precisely the sporadic interaction among networked entities. The design of negotiation mechanisms and the study of properties of the operating points to which such agreements naturally lead warrant investigation.

VII. ACKNOWLEDGMENTS

The authors acknowledge support of the European Commission through the STREP project NADA (FP7-ICT-223850) and the NoE project EuroNF (FP7-ICT-216366).

REFERENCES

- [1] P. Klemperer, "Auction theory: A guide to the literature", *Journal of Econ. Surveys*, vol.13, no.3, pp.227-286, July 1999.
- [2] V. Krishna, *Auction theory (2nd ed.)*, Academic Press, 2010.
- [3] C. Courcoubetis and R. Weber, *Pricing Communication Networks: Economics, Technology and Modeling*, Wiley, 2003.
- [4] W. Vickrey, "Counterspeculation, auctions and competitive sealed tenders", *Journal of Finance*, Vol. 16, No. 1, pp.8-37, 1961.
- [5] R.B. Myerson, "Optimal auction design", *Math. of Oper. Research*, vol.6, no.1, pp.5873, Feb. 1981.
- [6] P. Klemperer, *Auctions: Theory and practice*, Princeton University Press, 2004.
- [7] F.P. Kelly, "Charging and rate control for elastic traffic", *European Trans. Telecom.*, vol.8, pp.3337, 1997.
- [8] F.P. Kelly, A. Maulloo, and D. Tan, "Rate control for communication networks: Shadow prices, proportional fairness and stability", *J. Oper. Res. Soc.*, vol.49, no.3, pp.237-252, 1998.
- [9] E. Clark, "Multipart pricing of public goods", *Public Choice*, vol.2, pp. 19-33, 1973.
- [10] T. Groves, "Incentives in teams", *Econometrica*, vol.41, no. 4, 1973, pp. 617-631.
- [11] S. Shakkottai and R. Srikant, "Network Optimization and Control", *Foundations and Trends in Networking*, NoW Publishers, 2007.
- [12] R. Maheswaran and T. Basar, "Nash equilibrium and decentralized negotiation in auctioning divisible resources", *J. Group Decision and Negotiation*, vol.13, no.2, 2003.
- [13] R. Johari, and J. Tsitsiklis, "Efficiency loss in a network resource allocation game", *Math. of Oper. Res.*, Vol. 29, No. 3, pp. 407-435, Aug. 2004.
- [14] R. Maheswaran, and T. Basar, "Efficient signal proportional allocation (ESPA) mechanisms: Decentralized social welfare maximization for divisible resources", *IEEE J. Select. Areas Commun.*, vol. 24, no. 5, pp. 1000-1009, 2006.
- [15] S. Yang and B. Hajek, "Revenue and stability of a mechanism for efficient allocation of a divisible good", *submitted*, Oct. 2006.
- [16] S. Yang, and B. Hajek, "VCG-Kelly mechanisms for allocation of divisible goods: Adapting VCG Mechanisms to one-dimensional signals", *IEEE J. Select. Areas Commun.*, vol. 25, 2007, pp. 1237-1243.
- [17] R. Johari, and J. Tsitsiklis, "Communication requirements of VCG-like mechanisms in convex environments", in *Proc. Allerton Conf. on Commun., Control, Comput.*, 2005.
- [18] A. Dimakis, R. Jain, and J. Walrand, "Mechanisms for efficient allocation in divisible capacity networks", in *Proc. IEEE Conf. Decision Contr.*, San Diego, 2006, pp.1264-1269.
- [19] N. Semret, "Market mechanisms for network resource sharing", *PhD thesis*, Columbia University, 1999.
- [20] A. Lazar and N. Semret, "The progressive second price auction mechanism for network resource sharing", in *Proc. Int. Symp. on Dynamic Games Applic.*, 1997.
- [21] M. Bitsaki, G.D. Stamoulis and C. Courcoubetis, "A new strategy for bidding in the network-wide progressive second price auction for bandwidth", in *Proc. CoNEXT*, 2005.
- [22] M. Dramitinos, G.D. Stamoulis and C. Courcoubetis, "An auction mechanism for allocating the bandwidth of networks to their users", *Elsevier Computer Networks*, vol.51, pp.4979-4996, Dec. 2007.
- [23] www.adwords.google.com
- [24] Y. Narahari, D. Garg, R. Narayanam, H. Prakash, "Game theoretic problems in network economics and mechanism design Solutions", *Advanced Information and Knowledge Processing*, Springer, 2009.
- [25] B. Edelman, M. Ostrovsky, and M. Schwarz, "Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords", in *Proc. 2nd Workshop Sponsored Search Auctions*, 2006.
- [26] H. Varian, "Online ad auctions", *American Econ. Review*, vol.99, no.2, pp.430-434, May 2009.
- [27] H. Varian, "Position auctions", *Inter. J. Industr. Organization*, vol.25, pp.1163-1178, Oct. 2006.
- [28] L.S. Shapley and M.Shubik, "The assignment game I: The core", *Intern. J. Game Theory*, vol.1, no.1, pp.111-130, Dec. 1971.
- [29] G. Aggarwal, A. Goel, and R. Motwani, "Truthful auctions for pricing search keywords", in *Proc. 7th ACM Conf. Electr. Commerce (EC)*, 2006.
- [30] T. Roughgarden, E. Tardos, V. Vazirani, *Algorithmic Game Theory*, Cambridge University Press, 2007.
- [31] P. Milgrom, "Putting Auction Theory to Work", *Cambridge University Press*, 2004.
- [32] I. Katzela and M. Naghshineh, "Channel assignment schemes for cellular mobile telecommunication systems: A comprehensive survey", *IEEE Pers. Commun.*, vol.3, no.3, pp.1031, Jun. 1996.
- [33] E. M. Noam, "Taking the Next Step Beyond Spectrum Auctions: Open Spectrum Access", *IEEE Commun. Mag.*, vol.33, no.12, pp.66-73, 1996.
- [34] I.F. Akyildiz, W.Y. Lee, M.C. Vuran, and S. Mohanty, "NeXt Generation/Dynamic Spectrum Access/Cognitive Radio Wireless Networks: A Survey", *Elsevier Computer Networks*, September 2006.
- [35] J. Huang, R. Berry, and M. L. Honig, "Auction-based spectrum sharing", *ACM Mob. Net. and Applic. (MONET)*, vol.11, pp.405418, 2006.
- [36] J. Huang, R.A. Berry and M.L. Honig, "Auction mechanisms for distributed spectrum sharing", in *Proc. Allerton Conf. Commun. Control Comp.*, 2004.
- [37] J. Bae, E. Beigman, R. Berry, M.L. Honig and R. Vohra, "Sequential bandwidth and power auctions for spectrum sharing", *IEEE J. Select. Areas Commun.*, vol.26, no.7, pp.1193-1203, Sept. 2008.
- [38] X. Zhou, S. Gandhi, S. Suri, H. Zheng, "eBay in the Sky: strategy-proof wireless spectrum auctions", in *Proc. ACM MOBICOM*, 2008.
- [39] J. Jia, Q. Zhang, Q. Zhang, and M. Liu, "Revenue generation for truthful spectrum auction in dynamic spectrum access", in *Proc. Mobihoc*, 2009.
- [40] S. Gandhi, C. Buragohain, L. Cao, H. Zheng, S. Suri, "A general framework for wireless spectrum auctions", in *Proc. 2nd IEEE Symp. Dyn. Spectr. Access Net. (DySPAN)*, 2007.
- [41] G.S. Kasbekar, and S. Sarkar, "Spectrum auction framework for access allocation in cognitive radio networks", in *Proc. ACM Mobihoc*, 2009.
- [42] Z. Han, R. Zheng, V. H. Poor, "Repeated auctions with learning for spectrum access in cognitive radio networks", in *Proc. Allerton Conf. Commun., Control, Comp.*, 2009.
- [43] X. Zhou, and H. Zheng, "TRUST: A general framework for truthful double spectrum auctions", in *Proc. IEEE INFOCOM*, 2009.
- [44] R.P. McAfee, "A dominant strategy double auction" *Journal Econ. Theory*, vol.56, no.2, pp.434-450, April 1992).
- [45] S.H. Chun, and R.J. La, "Auction mechanism for spectrum allocation and profit sharing", in *Proc. Int. Conf. Game Theory for Networks, GameNets*, 2009.
- [46] B. Wang, Y. Wu, Z. Ji, K.J.R. Liu and T.C. Clancy, "Game theoretic mechanism design methods", *IEEE Signal Proc. Mag.*, vol.25, no.6, pp.74-84, Nov. 2008.
- [47] G. Iosifidis and I. Koutsopoulos, "Double auction mechanisms for resource allocation in autonomous networks", *IEEE J. Select. Areas Commun.*, vol.28, no.1, pp.95-102, Jan. 2010.
- [48] R.B. Myerson, "Analysis of two bargaining problems with incomplete information", in *Game Theoretic Models of Bargaining*, Cambridge University Press, 1985.
- [49] J. Watson, "Alternating-offer bargaining with two-sided incomplete information", *Review of Econ. Studies*, vol.65, no.3, pp.573-594, Jul. 1998.