Challenges in Auction Theory Driven Spectrum Management

George Iosifidis and Iordanis Koutsopoulos, University of Thessaly, Greece

ABSTRACT

Dynamic spectrum markets of the future will involve complex spectrum transactions among operators and users. The nontrivial dynamics in channel quality and spatiotemporal spectrum availability call for new ways of spectrum allocation. In this article, we advocate that auction theory can play a decisive role in shaping this evolving landscape since auctions are agnostic to user utility functions, and can build versatile and lightweight allocation mechanisms. Starting from auction preliminaries, we discuss how various intrinsic features of spectrum markets can be addressed with a modified version of auctions.

INTRODUCTION

Nowadays, auctions are one of the most popular methods used by both private agencies and governmental institutions for selling a variety of items, ranging from antiques to wireless spectrum. The common characteristic in all these cases is that the seller does not know the valuation of the items to potential buyers. In these settings, traditional market mechanisms such as pricing fail due to lack of this information. For example, if the seller sets very high prices for the items, these will probably remain unsold, while low prices will yield low revenue. On the other hand, auctions lead to allocation of items to the buyers with the highest valuations and at the same time to substantial increase of the revenue of the seller. Moreover, auctions require minimum interaction among sellers and buyers since the latter simply have to declare their preferences about auctioned items.

While initially auctions were designed empirically, in the last few decades game theory has been employed for their study. In 1961, Vickrey introduced the analysis of auctions as games of incomplete information [1]. In these games the players are the buyers who must select the appropriate bidding strategy in order to maximize their perceived utility (i.e., the value of the acquired items minus the payment to the seller). Each buyer is not aware of the valuation of the item for other buyers, and in some cases cannot even observe their actions (bids). From this perspective, the auction design is a mechanism design problem where the seller-buyers interaction rules must be selected so as to ensure the desirable equilibrium. In most cases, the objective is to allocate the item to the buyer with the highest valuation. Nevertheless, the selection of the proper auction remains an intricate task. Therefore, more often than not, specialized entities, i.e. companies or organizations, overtake the task of designing and running the auctions on behalf of sellers, which are the actual owners of the auctioned goodies.

The first spectrum auction was organized in New Zealand in 1990 for selling television spectrum bands [2]. Since then, many countries around the world have run auctions to sell spectrum bands. In most of the cases the result was remarkable, raising the revenue of the state to unexpected levels and granting the spectrum to the most interested buyers. Auctions evolved into a very popular method for allocating spectrum licenses. More important, in the last few years there has been an unprecedented demand for wireless services from an ever growing population of wireless users. This fact, coupled with the advent of new technologies such as cognitive radio, spurred discussions about the necessity for reforming the spectrum allocation policy. Today, it is common ground that spectrum regulation should be more dynamic and flexible. The state should grant spectrum licenses in different time and spatial ranges; moreover the license holders should be able to resell their idle spectrum channels [3]. This dynamic spectrum sharing model, at the timescale of channel or data flow dynamics, is expected to increase the utilization of spectrum. Spectrum will become a traded commodity in the *dynamic spectrum markets*.

In these markets, heterogeneous entities such as operators and users will interact, aiming to buy and resell spectrum. Usually, their objective will be to maximize their benefit, and hence they are expected to act selfishly. Moreover, these nodes will operate in a decentralized fashion and with limited information about the actual spectrum needs of other nodes. In this context, auction-based spectrum allocation mechanisms appear as an appropriate choice. First, auctions are agnostic to the utility functions of the bidders (i.e., the buyers' valuations for spectrum channels). Second, they require minimum interaction and hence are amenable to lightweight realization. Finally, they facilitate decentralized implementation. In the next sections we introduce the basic concepts of auction theory, and

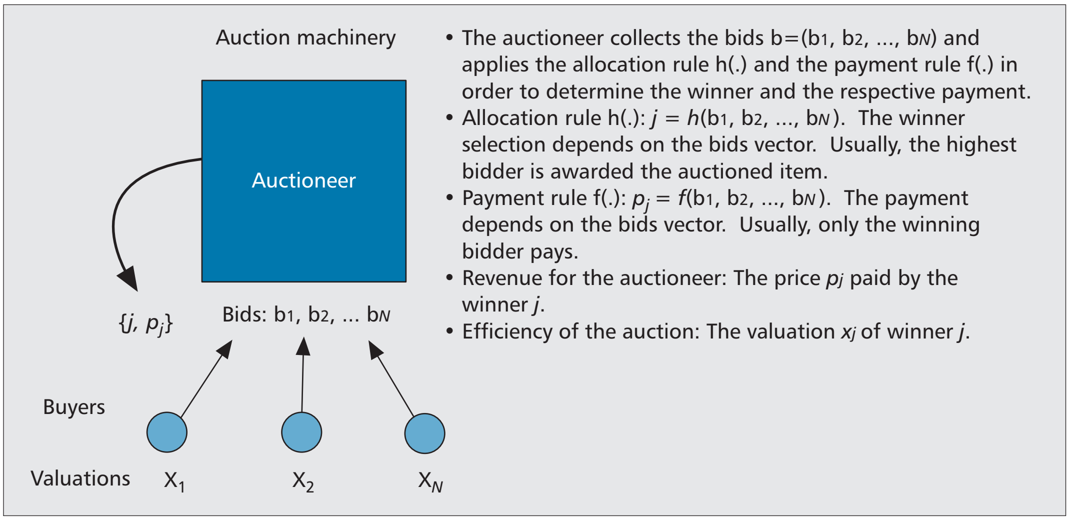

Figure 1. *Auction machinery. The winner and respective payment depend on the entire vector of bids. The revenue of the auctioneer is the payment of the winner. The efficiency of the auction is the total valuation of the allocated items for the winner(s).*

The components of every auction are the allocation rule, the payment rule and the bidding rule. The first determines the allocation of the auctioned item to buyers. Usually, the higher bidding buyer is awarded the item. The payment rule determines how much each bidder will pay.

present the properties and advantages that make them appropriate for selling spectrum. We also analyze the characteristics of emerging dynamic spectrum markets where auctions are expected to play crucial role. We discuss the challenges in designing these new auctions and present some auction types such as double auctions and hierarchical multilayer auctions.

Auction Theory Fundamentals

Single-Object Auctions

In the simplest form of auction, there is a set of buyers who bid to obtain one or more items (sold in a package, i.e., as one), and an auctioneer who collects these bids and decides which buyer will get the item and how much s/he will pay. The components of every auction are the allocation rule, the payment rule, and the bidding rule. The first determines the allocation of the auctioned item to buyers. Usually, the higher bidding buyer is awarded the item. The payment rule determines how much each bidder will pay. For example, a winning bidder is charged with an amount equal to his/her bid or the second highest bid. The basic difference of auctions from other similar mechanisms such as pricing schemes is that the allocation and payment for each buyer depends not only on his/her bid, but also on the bids of other buyers. The bidding rules define the machinery of the auction: what bids are allowed, whether the bids are sealed or revealed to all participants in the auction, or whether the bidders are able to update their offers in subsequent rounds. Different combinations of these rules result in different auction schemes. The most fundamental criterion that drives the decisions of the auction designer is the objective of the auction. The objectives range from revenue maximization for the auctioneer to social welfare (efficiency) maximization. The latter is achieved when auctioned items are allocated to the buyers with highest valuations. Very often these objectives are conflicting. That is, an auction that maximizes the social welfare probably will not generate the highest revenue for the auctioneer, and a revenue maximizing auction is very likely to entail efficiency loss. Figure 1 shows the generic setup of an auction.

Let us consider the two basic auction schemes: the *first price* and *second price* (Vickrey) auction. Assume that there exist $i = 1, \ldots, N$ bidders who compete to obtain one item. Each bidder, independent of the others, attaches a value $x_i \in [0, w]$, $w \in R$, to the item. This value is private information, which is not available to the auctioneer or other bidders. However, the cumulative distribution function (cdf) of this valuation, $F_i(x_i) = P(X_i \leq x_i)$, is considered common knowledge. The bidders can be different or symmetric (i.e., $F_i(\cdot) = F(\cdot), \forall i$). Each potential buyer i submits a sealed bid b_i, $i = 1, \ldots, N$ to the auctioneer, which collects them and grants the item to the highest bidder. Hence, both auctions have the same allocation rule. However, in the first price auction the winner pays his/her bid, while in the Vickrey auction s/he pays the second highest bid (Fig. 2).

The different payment rules yield different properties for these auctions. When we refer to auction properties, we are mainly interested in the efficiency of the allocation and the revenue generated for the auctioneer. A prerequisite for efficient allocation is truthfulness (or incentive compatibility); that is, the auction induces the bidders to reveal their actual valuations, $b_i = x_i$, $\forall i$. On the other hand, revenue maximizing auctions grant the items to the bidder expected to pay higher, which is not always the one with the highest valuation. Although the first and second price auctions yield the same revenue under certain conditions (*revenue equivalence theorem* [4]), in general the former produces higher revenue and with higher probability (i.e., less risk). Therefore, when the revenue is the primary objective of the auctioneer, the first price auction is the proper choice. On the other hand, when the bidder values are independent, the Vickrey auction is always truthful and hence effi-

The efficiency and revenue of multi-unit auctions depend on the market setting and assumptions about the buyers. For example, the uniform price auction is efficient only if we assume that buyers have single-unit demand, while the Vickrey auction is also efficient for selling multiple units, but only if bidders' valuations are independent.

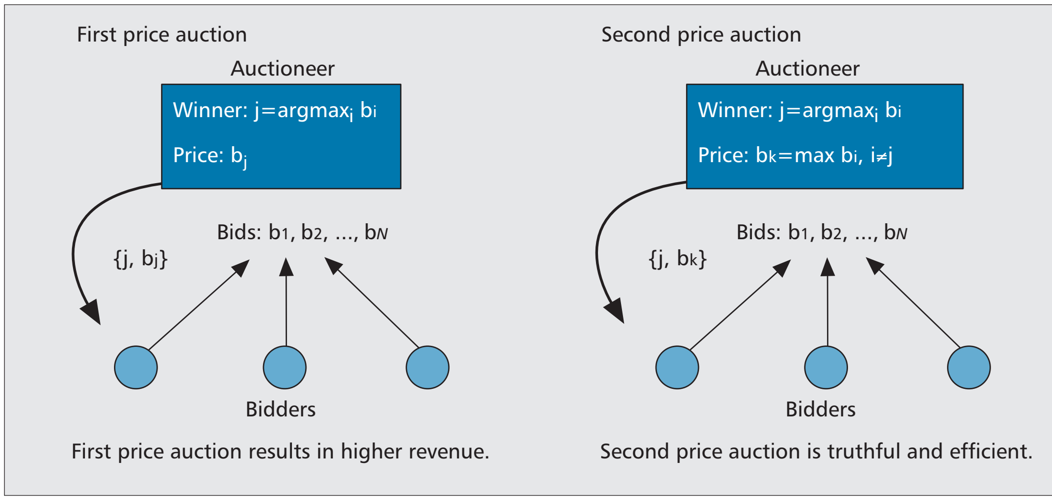

Figure 2. *Second price auction is always truthful and therefore efficient, since the auctioneer is able to allocate the item to the buyer with the highest valuation. First price auction results in higher revenue with less risk.*

cient, which is a very desirable property for auctioneers aiming at a socially optimal outcome.

A method to further increase the revenue of an auction is to use reserve prices. A reserve price is the minimum price the auctioneer accepts to sell his/her item. If all bids are below this price, the item is not sold. Myerson was the first to systematically study the selection of reserve prices [5]. He applied concepts from mechanism design and proposed so-called *optimal auctions* that ensure the maximum expected revenue for selling one single item. For every bidder i who submits a bid b_i, the auctioneer calculates the optimal reservation price by using the cumulative distribution function $F_i(\cdot)$ and respective probability density function $f_i(\cdot)$ (Fig. 3). This reservation price is subtracted from the actual submitted bid in order to calculate the virtual valuation (virtual bid) ψ_i of the buyer. Given the virtual valuations of the bidders, an optimal auction simply allocates the item to the bidder with the maximum non-negative virtual value. The winner pays the minimum bid required to win the auction. Notice that optimal auctions can result in inefficient allocation for two reasons. First, if all virtual bids are negative, the item remains unsold despite the existence of positive actual valuations. Moreover, in the case of asymmetric bidders (i.e., $F_i(\cdot) \neq F_j(\cdot)$ for $i \neq j$), it is probable that the highest virtual bid will not represent the highest actual valuation [4].

Multi-Object and VCG Auctions

If the auctioneer wishes to sell more than one item, s/he organizes a multiple object auction. Such an auction can be *homogeneous* or *heterogeneous*, depending on whether auctioned items are identical or different. In the latter case, the items could be *independent*, *substitutes*, or *complements*. The buyers' valuation for every additional acquired item decreases for substitutes and increases for complements. If the items are auctioned one at a time, the auction is called *sequential*, while if items are sold in one shot the auction is *simultaneous*. Finally, the buyers bid only for one item if they have *single-unit demand*, or for bundles of items in case of *combinatorial* auctions. Apparently, the design space of multi-unit auctions is very rich.

Let us focus on homogeneous sealed-bid simultaneous auctions. Assume that there is one auctioneer with K identical items, and N buyers who bid for these items. Each bidder i submits a bid vector $\mathbf{b}^i = (b_1^i, b_2^i, \ldots, b_K^i)$, where b_m^i is the amount i is willing to pay for the mth item. The total amount bidder i is willing to pay for obtaining all items is $\Sigma_{j=1}^{K} b_j^i$. The most popular auction schemes in this category are the discriminatory-price, uniform-price, and Vickrey auctions. All of them have the same allocation rule, according to which the K highest bids are deemed "winning bids," and every bidder receives the items for which s/he was a winner. However, the payment rules of these auctions are different. The discriminatory auction is actually the multi-unit extension of the first price auction. That is, each bidder i pays an amount equal to his/her K_i winning bids, $\Sigma_{j=1}^{K_i} b_j^i$. In other words, different buyers pay different prices for different items. On the other hand, in the uniform-price auction, all K items are sold at the market clearing price, which is selected to equate total supply and total demand. Finally, in the multi-unit Vickrey auction, each buyer i who is awarded K_i items pays a price equal to the sum of the bids s/he has outbidden. These are the bids of the buyers that would have got the K_i items, if buyer i was absent.

The efficiency and revenue of multi-unit auctions depend on the market setting and assumptions about the buyers. For example, the uniform-price auction is efficient only if we assume that buyers have single-unit demand, while the Vickrey auction is efficient also for selling multiple units, but only if bidders valuations are independent. The only auction that ensures truth-telling and efficient allocation without any restrictions is the celebrated Vickrey-Clark-Groves (VCG) mechanism [4]. In VCG auctions, every buyer i pays a price that is equal to the externality s/he creates to the market. This is calculated as the difference between the sum of the valuations (social utility) of the

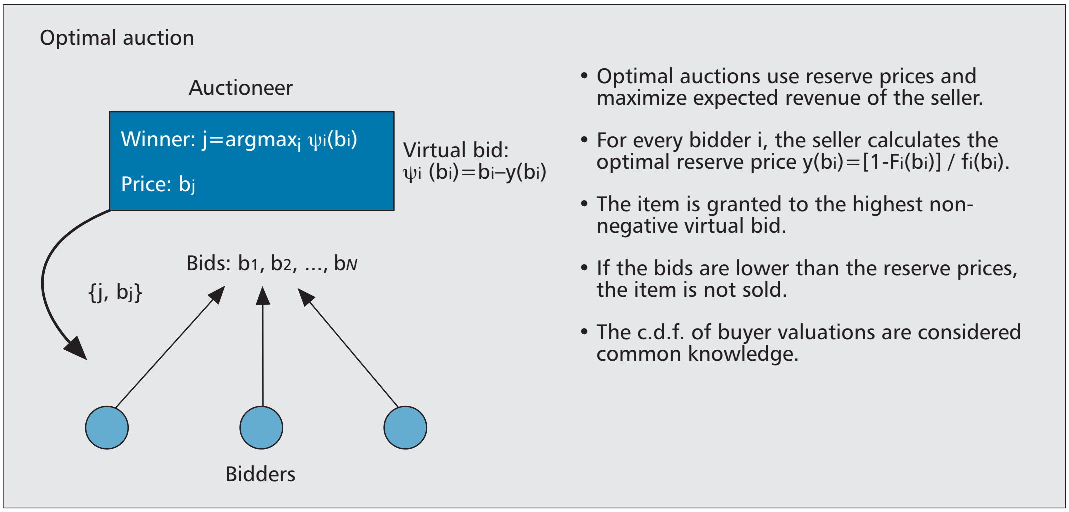

Figure 3. *The optimal auction uses virtual bids which stem from the actual submitted bids minus the optimal estimated reserve price for each bidder. If the winning virtual bid is less than zero, the item is not allocated. The optimal auction ensures the maximum possible expected revenue for the seller.*

Initially, the method used for this allocation procedure was comparative hearings, which in the 1980s was replaced by lotteries. However, it was soon realized that this process was profit-deficit, and a better choice was to directly sell spectrum to the buyers.

winners other than *i* when *i* does not participate and the sum of their valuations when *i* participates in the auction. In other words, each bidder pays an amount equal to the total utility decrease it causes to all other bidders.

There are also *reverse* or *procurement auctions* where some sellers (or producers) submit offer bids in order to provide a set of items to the auctioneer. An extension of multiple-object auctions are *share auctions* where the auctioned item is an infinitely divisible resource. In many cases there are more than one seller that can provide the same items to a set of buyers. In these auctions, known as *double-sided auctions*, the sellers compete for the buyers who are able to select the best offer. Figure 4 depicts the comparison of double- and single-sided auctions. Finally, other advanced topics in auction theory include *score auctions* [2], where bidders are sorted not only according to their bids but also with respect to their specific properties such as various quality metrics. A variation of a score auction is used by Internet search engines for selling advertisement slots in web pages that display the keyword search results of Internet users. In these *sponsored search auctions*, the bids are weighted with a scalar parameter that reflects the relevance of their advertisements to the specific keyword search [6, references therein].

Novel Challenges in Spectrum Auctions

In the first attempts to allocate wireless spectrum, specialized governmental institutions, such as the U.S. Federal Communications Commission (FCC), were appointed the task of acting as regulators and granting available spectrum licenses to interested parties. Initially, the method used for this allocation procedure was *comparative hearings*, which in the 1980s was replaced by *lotteries* [2]. However, it was soon realized that this process was profit-deficit, and a better choice was to directly sell spectrum to the buyers. Selling a public good such as spectrum is a very difficult task, mainly because its value is not known in advance. The seller cannot foresee which buyer values spectrum higher and how high the valuation is. In settings like this, where there is incomplete information about valuation of the goods, traditional market mechanisms such as pricing fail. On the other hand, auctions constitute a promising tool since they induce bidders to reveal their true valuations to the auctioneer, and require minimum communication between buyers and sellers.

In the first spectrum auction, which was organized in New Zealand in 1990, the organizers began with a second price auction, but due to low revenue they switched to a first price auction soon after. In 1994, the FCC organized the first U.S. spectrum auction for selling 10 narrowband personal communication services (PCS) licenses. Since then, the FCC has organized more than 70 auctions selling more than 31,000 licenses [2]. Similarly, in many countries all over the world, governments run auctions for selling spectrum licenses. While traditional auctions succeeded in increasing the revenue and efficiency of spectrum allocation, the unprecedented growth of demand for wireless communication services in the last decade has rendered necessary the reconsideration of this static allocation policy. Currently there are a variety of wireless services and a growing population of users with diverse requirements that vary in both temporal evolution and spatial location. In order to satisfy these demands, operators and users should have dynamic access to spectrum. Spectrum regulators soon realized the requirement for dynamic spectrum management [3]. Nowadays, it is common belief that there is a strong need for spectrum liberalization where spectrum owners will be able to resell, lease, or exchange their spectrum. With this dynamic spectrum sharing model, it is expected that spectrum utilization will be boosted, ensuring both user satisfaction and operator revenue. This twist in spectrum policy is leveraged by the advent of cognitive radio technology, *Spectrum allocation in dynamic spectrum markets requires the development of novel auction schemes. Unlike traditional auctions organized by state agencies, the auctioneer in these cases can be any operator or even a user who is willing to exchange his spectrum.*

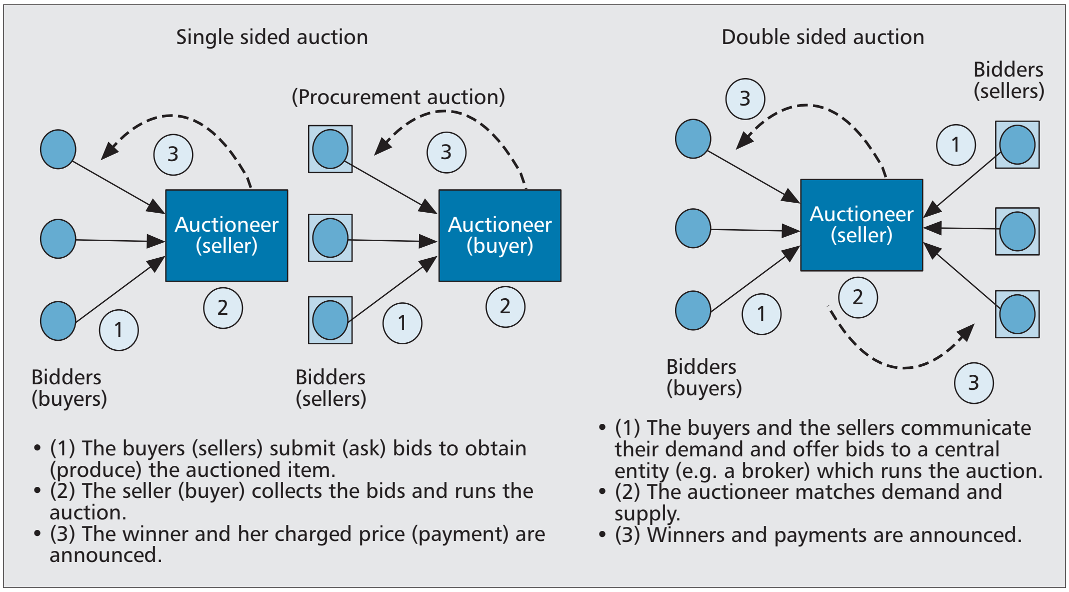

Figure 4. *Single and double sided auctions. In single-sided auctions the auctioneer sells or asks for an item (procurement auctions). In doublesided auctions there exist many buyers and many sellers interacting concurrently. A centralized entity collects the offer and the ask-bids and runs the auction algorithm to determine the allocation of the items and the respective payments.*

which enables the intelligent reconfiguration of nodes' transmission characteristics to adapt to the varying conditions of their surrounding environment.

In the new era of dynamic spectrum markets, the static spectrum allotment policy must be replaced by more flexible schemes. Spectrum channels will be granted in different time scales and for various spatial ranges to operators or directly to users. Channel allocation and spectrum access will be for either exclusive use (primary access) or low-cost secondary access. Regulators will organize auctions for selling spectrum licenses to the so- called primary operators (POs). Apart from serving primary users (PUs), the POs will exchange spectrum bands with each other and additionally lease unused bandwidth to secondary operators (SOs). The latter will be able to serve secondary users (SUs) in their range without the need to invest money for licenses. Secondary users are obliged to transmit without causing interference to the PUs and, more often than not, will have to compete with each other for spectrum access [7]. In these markets there will be many different scenarios for spectrum allocation [8]. The common denominator is the freedom of the various entities to trade spectrum at their own will and in the presence of limited information about spectrum demand. Auctions are expected to constitute the prime method for selling and redistributing the spectrum in these paradigms. In Fig. 5 a schematic representation of the spectrum sharing interactions between POs/SOs and users is depicted.

AUCTIONS IN DYNAMIC SPECTRUM MARKETS

Spectrum allocation in dynamic spectrum markets requires the development of novel auction schemes. Unlike traditional auctions organized by state agencies, the auctioneer in these cases can be any operator or even a user who is willing to exchange his/her spectrum. Lightweight mechanisms are needed, with minimum communication overhead among involved entities.

DYNAMIC SPECTRUM SHARING

Dynamic spectrum auctions differ in substantial ways from respective schemes where other assets are sold and from static spectrum auctions. First, in dynamic spectrum markets there are expected to be many networks that will cover small areas; therefore channel allocation will be a more complicated task than in traditional wireless networks. Spectrum can be reused by operators that are not in adjacent cells or, in general, nearby cells. Auctions should consider the spatial dimension and the fact that there may be many winners to which spectrum should be concurrently allocated. Second, spectrum bands *differ* in terms of quality due to inherent frequency selectivity of the wireless channel and the time-varying link quality. Third, *heterogeneity* and *unpredictability* of user demands and user mobility place additional challenges. Finally, the *small-scale* dynamic fashion in which spectrum allocation has to be accomplished renders the machinery of bidding, allocation, and payment challenging.

In this context, a common paradigm of spectrum auction is when a PO, such as a large broadcasting company, allocates its idle spectrum to a set of secondary operators, for example, mobile virtual network operators (MVNOs). Spectrum is divided in channels (spectrum chunks), and SOs at adjacent cells cannot use the same channel due to interference. The PO collects the bids and determines the channel allocation by solving an optimization problem subject to the interference constraints. The PO can use any type of multiple-object auction depending on its objective. For example, in [9]

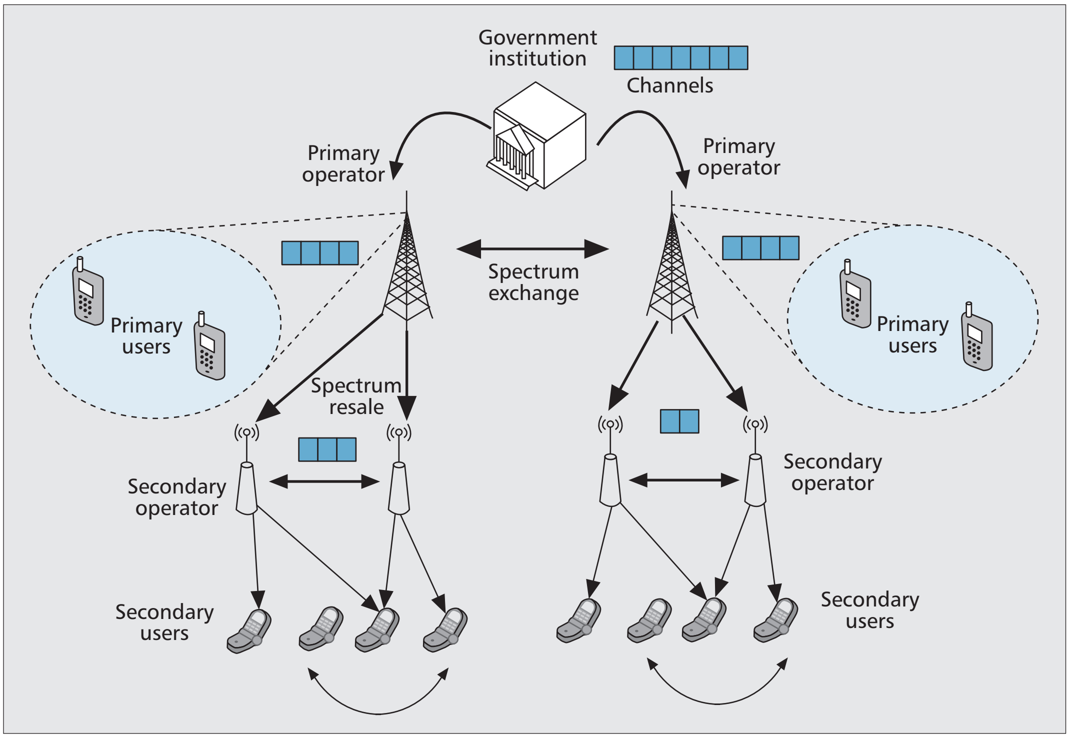

Figure 5. *Dynamic spectrum sharing: a government institution allocates spectrum channels to a set of primary operators. The POs serve their users and resale/lease unutilized spectrum to a set of secondary operators. These, in turn, provide services to a set of secondary users. Operators and users can exchange spectrum channels to satisfy their dynamic varying spectrum needs.*

In these double spectrum auctions, every entity becomes both a buyer and a seller, and submits requests for new channels while at the same time offering its idle spectrum. The challenge is to manage this large volume of spectrum requests across several geographic areas by opportunistically exploiting spectrum surpluses.

an extension of the optimal auction is proposed that is suitable for maximizing the revenue of the seller when it auctions multiple channels. However, due to spectrum reuse, these auctions either lose their incentive compatibility property (truthfulness) or have exponential complexity. To alleviate this problem, the authors of [10] propose a suboptimal auction scheme that is computationally efficient but reduces the revenue of the seller.

Similar issues arise when the SOs organize auctions to sell spectrum channels to SUs. These auctions can be one-shot or repeated. In the latter case, SUs can update their bidding strategy in order to increase their benefit from channel usage. Each channel can be allocated to more than one SU. In this case, the SUs share the channel and communicate through a protocol such as carrier sense multiple access (CSMA). Another proposition is to have a central agency selling spectrum to the PUs and SUs [11]. Each user is able to ask for a primary license if it has inelastic spectrum demands or a secondary spectrum license if it prefers a low-cost service. The agency allocates spectrum channels so as to increase efficiency or revenue. In this setting, the optimal access allocation problem is NP-complete, but under certain conditions and relaxations can be solved with polynomial-time algorithms using either dynamic programming techniques or tools from graph theory.

A characteristic of dynamic spectrum markets is that they aim to facilitate spectrum exchange among different entities. Each PO may lease or even exchange its channels to other POs. Similarly, SOs will be able to directly interact and satisfy their dynamic needs by redistributing the spectrum they have leased from POs. Finally, even users may exchange their channel access licenses. These many-to-many interactions can be captured by double auction mechanisms. Spectrum owners sell their channels to a set of buyers who are able to select the most attractive offers in terms of cost and channel quality. We note that double auctions presume the existence of a central entity that will collect all bids and run the auction algorithm in order to match requests and offers.

In these double spectrum auctions, every entity becomes both a buyer and a seller, and submits requests for new channels while at the same time offering its idle spectrum. The challenge here is to manage this large volume of spectrum requests across several geographic areas by opportunistically exploiting spectrum surpluses. This way, spectrum utilization will increase, and at the same time the dynamic spectrum needs of the operators and users will be satisfied. Devising truthful and efficient double auctions is an intricate task. In [12] the authors explain that the requirement for spectrum reuse results in non-truthful auctions. In other words, there is a trade-off between the economic robustness of these auctions and the efficiency of spectrum allocation. To address this trade-off they propose a double auction scheme, based on the *McAfee* double auction, where the bidders are grouped, and each group is allocated the same channel. Finally, in some networks where there are no central nodes, such as ad hoc networks, it is necessary to employ methods for distributed execution of the auction

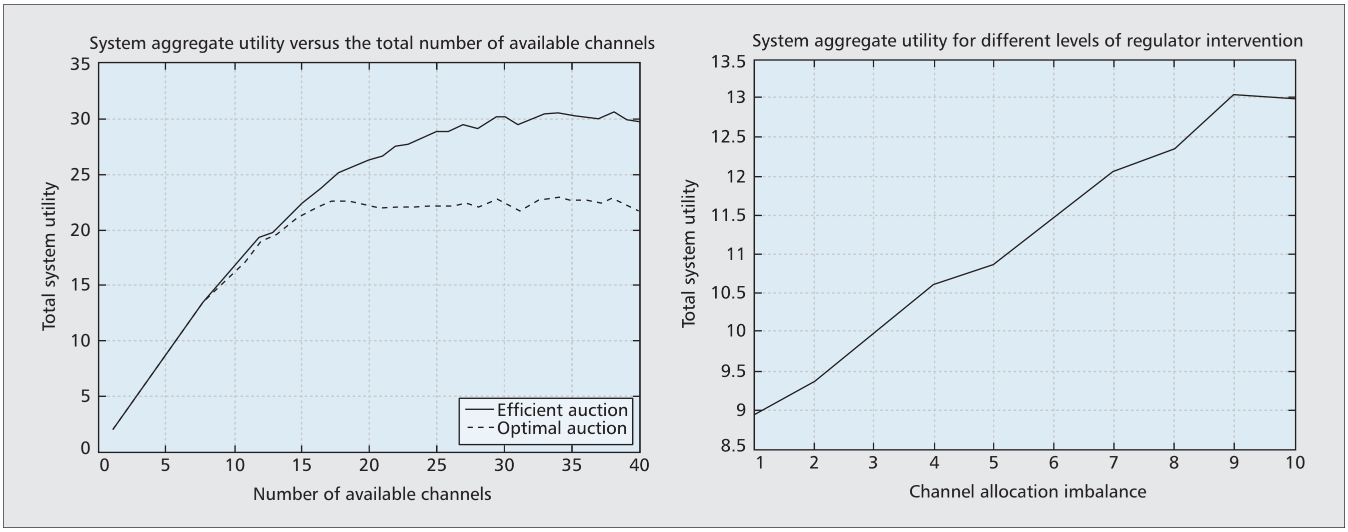

Figure 6. *Hierarchical successive spectrum allocation induces efficiency loss for the network due to the selfish revenue maximizing behavior of operators in the second layer. The regulator can decrease this loss by allocating more channels to the socially aware operators: a) The upper curve represents the system aggregate utility when the channels are allocated directly and efficiently by the regulator to the 60 users. The lower, dotted curve, depicts the system utility when each operator receives 40 channels and resells them using an optimal auction to his 30 users. b) System aggregate utility for different values of the channel allocation imbalance parameter.*

mechanisms. For example, in [13] a distributed algorithm based on primal-dual Lagrange decomposition is proposed for the implementation of the auction mechanism.

Hierarchical Spectrum Allocation

An important feature of dynamic spectrum markets is the heterogeneity of involved entities that interact through certain hierarchies. In particular, in these markets spectrum distribution becomes a multilayer resource allocation process. In every layer, a set of entities (operators or brokers) request some spectrum, which they may resell (as a whole or parts) to operators or users in another layer. The interaction of the entities between two successive layers affects the utility of entities in other layers. In other words, there is an interdependence among layers. However, the different objectives of operators and users in the different layers, as well as their different transmission capabilities and spectrum needs can result in inefficient spectrum allocation or revenue loss. Apparently, there is a need for mechanisms that will align the incentives of the various entities and enable their coordination. Clearly, traditional two-party auction schemes are inadequate to capture and address issues that arise in this context.

An example of such a multilayer interaction is presented in Fig. 5. Consider a government agency that organizes an auction to allocate spectrum chunks to a set of POs, and assume that these POs are allowed to lease or resell the spectrum they acquire to SOs. The agency is expected to be a socially aware entity that is interested in maximizing spectrum utilization (i.e., the efficiency of spectrum allocation). Therefore, it will run a truth-telling auction such as Vickrey or VCG to allocate the spectrum to the POs. However, the POs are market entities and hence are expected to resell their spectrum so as to maximize their revenue. We already explained that auctions which maximize revenue, such as the optimal mechanism of Myerson, can cause efficiency loss [4]. Clearly, the objectives of the POs and the agency are not aligned, and the successive auctions will not ensure the initial goal of efficient spectrum utilization.

Obviously, the most suitable solution for these spectrum allocation problems would be to have a central regulator that would continuously redistribute the spectrum to both primary and secondary users, and ensure maximum possible efficiency. Since most probably this may not be realizable, one seeks other mechanisms to ensure the overall objective (efficiency or revenue) in the presence of these spectrum allocation hierarchies. An idea in this direction is to use *score* auctions [2]. The seller at the top of each hierarchy can run an auction to allocate spectrum, but does not consider only the bids of buyers. Instead, s/he will use additional criteria that capture the behavior of the bidders when they act as sellers for the next-layer market. For example, the top auctioneer can enable an enhanced auction mechanism where s/he considers feedback from the bottom-layer buyers. Sponsored search auctions constitute a simple instance of this paradigm. These auction schemes will provide incentives to the middle-layer nodes to comply with the goal of the top-layer auctioneer and reduce their revenue in favor of the benefit of the bottom-layer nodes.

Another instance of a three-layer hierarchical interaction is among POs, SOs, and SUs (Fig. 5). Again, each PO wishes to sell unused spectrum to a set of SOs. The SOs submit bids, and the POs determine the payments and channel allocation. The appropriateness of this allocation depends on the *experience* of the SUs. Obviously, it is not desirable to allocate certain channels to SOs that would in turn assign them to users for which the specific frequency is of low quality, due to either interference from excessive frequency reuse or limited range of the specific SOs. In other words, the PO has to determine the most suitable SO for each specific channel. Therefore, it is imperative to have a mechanism that will enable the coordination of the entities in all the layers of this hierarchy. Moreover, in some settings the SUs may be also clients of the POs [11]; hence, the latter have an additional incentive to take into account the users' preferences. In order to ensure overall efficient allocation of the channels, the POs should consider the *feedback* from users about the quality of the services they received from each SO. This feedback can be used to modulate the bids of the SOs when they request spectrum, in the spirit of score and sponsored search auctions.

Finally, we consider a representative three-layer hierarchical spectrum market and run numerical experiments. At the top of the hierarchy there is one socially aware spectrum regulator which allocates 40 channels to each of two revenue maximizing operators, and each operator in turn serves 30 users. User needs are private information, but their distribution is uniform and known to operators and the regulator. In Fig. 6a we plotted the overall system efficiency (i.e., the sum of the winning users' valuations vs. the number of allocated channels) for two scenarios:

- When channels are allocated by operators using an *optimal auction* scheme that maximizes their expected revenue
- When channels are allocated directly by the regulator to the users using an *efficient auction* such as the Vickrey auction

The intervention of the operators in the spectrum distribution process is shown to introduce significant efficiency loss, which increases with the number of channels. If one of the two operators is socially aware while the other is a selfish revenue maximizing entity, a simple method to increase the overall efficiency of the hierarchical allocation is to have the regulator allocate more channels to the efficient operator and fewer to the selfish operator. In Fig. 6b we have plotted the efficiency of the system when the regulator allocates 20 channels to the two operators, each one serving 30 users. Initially, the operators receive an equal number of channels. Next, the regulator allocates more channels to the efficient operator and fewer to the selfish operator by properly tuning a scalar parameter we call *channel allocation imbalance*. Finally, when the imbalance value is 10, all 20 channels are allocated to the socially aware operator, which increases the overall efficiency.

Conclusions

We have explained the contribution of auctions in spectrum allocation in dynamic spectrum markets. The complex spectrum transactions among network entities, the spatiotemporal channel and flow dynamics, as well as the multiple hierarchies among entities necessitate a fresh look at auction design for spectrum management.

References

[1] W. Vickrey, "Counterspeculation, Auctions and Competitive Sealed Tenders," *J. Finance*, vol. 16, no. 1, 1961, pp. 8–37.

[2] P. Milgrom, *Putting Auction Theory to Work*, Cambridge Univ. Press, 2004.

[3] J. M. Peha, and S. Panichpapiboon, "Real-Time Secondary Markets for Spectrum," *Telecommun. Policy*, vol. 28, no. 7, Aug. 2004, pp. 603–18.

[4] V. Krishna, *Auction Theory*, 2nd ed., Academic Press, 2010.

[5] R. B. Myerson, "Optimal Auction Design," *Mathematics of Operations Research*, vol. 6, 1981, pp. 58–73.

[6] I. Koutsopoulos and G. Iosifidis, "Auction Mechanisms for Network Resource Allocation," *Proc. Wksp. Resource Allocation in Wireless Networks*, Avignon, France, 2010, pp. 554–63.

[7] Q. Zhao, and B. M. Sadler, "A Survey of Dynamic Spectrum Access," *IEEE Sig. Proc. Mag.*, vol. 24, no. 3, May 2007, pp. 79–89.

[8] I. F. Akyildiz *et al.*, "NeXt Generation/Dynamic Spectrum Access/Cognitive Radio Wireless Networks: A Survey," *Computer Networks J.*, vol. 50, no. 13, Sept. 2006, pp. 2127–59.

[9] S. H. Chun, and R. J. La, "Auction Mechanism for Spectrum Allocation and Profit Sharing," *Proc. Int'l. Conf. Game Theory for Networks*, Istanbul, Turkey, 2009, pp. 498–507.

[10] J. Jia *et al.*, "Revenue Generation for Truthful Spectrum Auction in Dynamic Spectrum Access," *Proc. ACM Mobihoc,* 2009.

[11] G. S. Kasbekar and S. Sarkar, "Spectrum Auction Framework for Access Allocation in Cognitive Radio Networks," *IEEE/ACM Trans. Net.*, vol. 18, no. 6, 2010, pp. 1841–54.

[12] X. Zhou, and H. Zheng, "TRUST: A General Framework for Truthful Double Spectrum Auctions," *Proc. 31st IEEE INFOCOM*, Rio De Janeiro, Brazil, 2009, pp. 999–1007.

[13] G. Iosifidis and I. Koutsopoulos, "Double Auction Mechanisms for Resource Allocation in Autonomous Networks," *IEEE JSAC*, vol. 28, no. 1, Jan. 2010, pp. 95–102.

Biographies

George Iosifidis (giosifid@inf.uth.gr) received his Diploma in electronics and telecommunications engineering from the Greek Air Force Academy, Athens, in 2000 and his M.Sc. from the Department of Computer and Communications Engineering of the University of Thessaly in 2007. He is currently working toward his Ph.D. degree at the same university. His current research interests are in the area of wireless and wireline networking, with emphasis on distributed resource allocation algorithms for peer-to-peer and cognitive radio networks. Since 2000 he has been serving as an electronics and telecommunications engineering officer in the Greek Air Force.

Iordanis Koutsopoulos [S'99, M'03] is an assistant professor at the Department of Computer and Communications Engineering, University of Thessaly, Greece, since 2005. He is also affiliated with the Institute for Telematics and Informatics of the Center for Research and Technology Hellas (CERTH). He obtained his Diploma in electrical and computer engineering from the National Technical University of Athens, Greece, in 1997, and M.Sc. and Ph.D. degrees in electrical and computer engineering from the University of Maryland, College Park (UMCP) in 1999 and 2002, respectively. From 1997 to 2002 he was a Fulbright Fellow and a research assistant with the Institute for Systems Research (ISR) of UMCP. He held internship positions with Hughes Network Systems, Germantown, Maryland, Hughes Research Laboratories LLC, Malibu, California, and Aperto Networks Inc., Milpitas, California, in 1998, 1999, and 2000, respectively. For the summer of 2005 he was a visiting scholar with the University of Washington, Seattle. For the period 2005–2007 he was awarded a Marie Curie International Reintegration Grant (IRG). His research interests are in the area of network control and optimization, with emphasis on wireless network cross-layer design and performance analysis, peer-to-peer, and sensor networks.

An important feature of dynamic spectrum markets is the heterogeneity of involved entities which interact through certain hierarchies. In particular, in these markets spectrum distribution becomes a multi-layer resource allocation process.