

A NOVEL APPROACH FOR ALLOCATING NETWORK AND IT RESOURCES OFFERED BY DIFFERENT SERVICE PROVIDERS

Marina Bitsaki¹ and Manos Dramitinos²

¹Department of Applied Mathematics, University of Crete, Knossos Avenue, Heraklion
GR 71409 Greece

marina@tem.uoc.gr

²Department of Informatics, Athens University of Economics and Business, 76 Patision
Street, Athens GR 10434, Greece

mdramit@aueb.gr

ABSTRACT

IT TECHNOLOGIES ARE INCREASINGLY BEING USED TO OFFER COMPUTING RESOURCES OR INFORMATION IN DIFFERENT LOCATIONS USING NETWORK RESOURCES. IN THIS PAPER WE ADDRESS THE PROBLEM OF ALLOCATING RESOURCES TO ENABLE CUSTOMERS COMPOSE THEIR OWN DESIRED SERVICES. WE PROPOSE AN AUCTION MECHANISM THAT COMPRISES INDEPENDENT AUCTIONS PERFORMED BY EACH PROVIDER AND INVESTIGATE BIDDING STRATEGIES AND SOCIAL WELFARE.

KEYWORDS

Auctions, Broadband Networks, Services, Resource Allocation

1. INTRODUCTION

The pace of innovation in computer and networking technology allows many new IT services to emerge rapidly and often benefit from new network services that support faster connectivity and better performance. Centralized resources that are used to provide services to customers are impractical in a dynamic service-based system owned by multiple organizations. The new service oriented business environment enables service providers to dynamically integrate services across networks according to their customers' needs. This is achieved by exploiting IT resources that are easy to access no matter where they reside. For example, network service providers offer bandwidth to customers who are interested in buying content (e-books, music, streaming media) from web sites (content providers) or virtual machines (supporting operating systems or processes) from Amazon (grid service providers).

In order to keep pace with evolving business needs and increasing complexity in service oriented businesses, providers face the challenge of managing performance issues and network constraints that affect customers' satisfaction received by using the offered services. In other words, there is a need to provide cost effective connectivity to all locations around the world, while enabling the convergence of voice, data, and internet services. This complexity of allocating different kinds of resources seen as complementary goods and owned by different providers, makes traditional resource allocation approaches insufficient to reveal market information and optimize the system's performance.

In this paper, we address the problem of managing resources in systems that offer services in a distributed environment. We use a game theoretic approach to understand how customers' information flows from one system to another and how customers can become active and value

generators. In particular, we propose an auction mechanism that comprises independent auctions performed by each provider for resource allocation, without requiring their cooperation in the composition of a service that uses the auctioned resources. We adjust the VCG payment rule [1] in our approach and argue that sincere bidding is still an optimal strategy. We perform experiments applying our approach, in order to measure the degree of the exposure problem and thus the social welfare loss. The results show that the performance is generally good and the social welfare attained is close to the optimal one.

The structure of this paper is as follows. Section 2 overviews related work on combinatorial auctions. Section 3 describes the model and the proposed auction mechanism. Section 4 provides the experimental results that measure the social welfare loss and Section 5 provides some concluding remarks.

2. RELATED WORK

Much of the existing literature on auction theory analyzes the multi-object auction game. Often, in such auctions the goods and the bids for them are complementary, meaning that the value of a good to a bidder depends on which other goods he wins. Because of the possibility of such dependences in values, mechanism design theory suggests that single bids for combinations of goods are allowed. This has an impact on the complexity of the winner determination problem which is proved to be an NP-complete problem [2]. Many studies propose special structures in the allowed combinatorial bids in order to guarantee computational tractability ([3], [4]) and others propose designs in which bidders are forced to submit individual bids for each distinct good increasing mechanism complexity (rules difficult to analyze and understand) in order to eliminate the exposure problem and thus efficiency loss ([5], [6], [7]).

In [3], the problem of designing simultaneous multi-object auctions is addressed. Bidders are allowed to submit combinations of bids since their value depends on which objects they win. The authors identify different structures of combinatorial bids for which computational tractability is achieved. In [4], the authors study special settings in combinatorial auctions in which the winner determination problem can be solved in polynomial time. A graph with a desired property is constructed and bids are modelled as a connected set of vertices. The problem of constructing the graph and handling the winner determination is solved.

In [5], simultaneous clock auctions for many divisible goods that are offered by one seller are studied. The seller's objective is a combination of efficiency and revenue optimization. Natural linkages among goods (substitutes or complements) are also encountered. Two economic models are presented in which the allocation of goods is efficient and sincere bidding is proved to be an equilibrium strategy. In [6], complementary goods without bundling are auctioned by one seller. First price and second price auctions are examined in terms of equilibrium outcomes under the assumption that the value of the whole demanded bundle is higher than that of any incomplete bundle. In [7], Milgrom addresses the package allocation problem in which traders buy or sell combinations of goods. A survey of various mechanisms is presented emphasizing on the problems that arise when some of the goods are not substitutes.

Many studies have been proposed in the communication networks sector, dealing with the resource allocation problem ([8], [9], [10], [11], [12]). In [8], Cramton focuses on spectrum auction design where many heterogeneous items are auctioned. The package clock auction is analysed in which bidders bid on packages of lots. It is proved that efficiency is improved through truthful bidding. In [9], a bandwidth pricing mechanism that solves congestion problems in communication problems is presented. An efficient allocation is obtained when an

infinitely divisible resource is to be shared among users who are allowed to submit several bids in order to establish a connection. Lazar and Semret propose in [10] the Progressive Second

Price auction to perform bandwidth allocation in a network of arbitrary topology by running independent auctions in each link. In [11], the authors present a mechanism of a set of simultaneous multi-unit descending-price auctions for auctioning bandwidth on a network-wide basis to end users or ISPs that will utilize it for the same time period, attaining nearly efficient bandwidth allocation. In [12] a network auction mechanism is proposed wherein agents can express their bid for a whole path. The authors prove the existence of a Nash equilibrium which yields the efficient allocation.

The above studies assume that the distinct goods are sold by the same seller. In this paper we assume that each good is sold by a different seller and bidders submit independent bids (of their desired bundle) in each auction.

3. THE MODEL AND ANALYSIS

3.1. Assumptions

We begin our study by formulating the resource allocation problem of two service providers A and B who independently want to offer services to a common group of N customers in a specific time interval. In order to deliver a service, each provider sells divisible resources of finite capacity. In particular provider A sells good 1 with capacity C_1 (e.g. bandwidth) and provider B sells good 2 with capacity C_2 (e.g. computing power). The two types of resources are assumed to be complementary meaning that each customer may be willing to buy quantities of both of them in order to benefit from a combined service offered to him. (Such a customer is called inelastic in the sense that he is willing either to obtain all the bundle of the two goods or take nothing of them.)

Each provider's objective is to allocate his resources efficiently without revealing private information (e.g market power) to the other provider. We assume that each customer has a private value for each bundle of goods and that total demand of goods 1 and 2, D_1 and D_2 respectively for the examined time interval are publicly observed (symmetric information among bidders). Our goal is to achieve an efficient outcome and reveal correct signals about capacity demands for each good. We propose an auctioning mechanism in which independent sealed-bid auctions are performed for each good by each service provider. For simplicity reasons we have assumed two types of resources, though our model can be generalized to consider K types of resources.

3.2. Bidding rules

We assume 3 types of bidders: bidders of type 1 demand units of good 1, bidders of type 2 demand units of good 2 ("individual" bidders), and bidders of type 3 demand units of both goods 1 and 2 ("combinatorial" bidders).

Bidder's i private value is denoted by $v_j^i, i=1, \dots, N, j=1,2,12$ for good j , where the case $j=12$ represents the "combinatorial" bidder. Bidder's i bid in auction j is $\vec{s}_j^i = (q_j^i, b_j^i), i=1, \dots, N, j=1,2$, where q_j^i is the demanded quantity of good j and b_j^i is the total price for the demanded quantity. Combinatorial bidders submit two bids, one for each good. It is assumed that each bidder chooses bidding strategies that maximize his profit.

3.3. Winner determination-Payment rule

Consider the set $S_j = \{\bar{s}_j^i, i = 1, \dots, N\}$, $j = 1, 2$ of bids in auction j . Each provider sorts the bids in each S_j according to price per unit $p_j^i = b_j^i / q_j^i$, $i = 1, \dots, N$ and assigns demanded quantities up to capacity constraint. The last winning bidder is assigned the remaining amount (might be less than his demand).

We apply the VCG payment rule [1], according to which each bidder is charged the social opportunity cost that his presence entails and ensures that the optimal strategy for a bidder is to bid the true values of the goods (sincere bidding). Formally, bidder i is charged $SW_{-i}(0, \bar{v}_j^{-i}) - SW_{-i}(\bar{v}_j)$ where SW_{-i} is the social welfare (sum of values of bidder that win in the efficient outcome) of bidders other than i , \bar{v}_j is the vector of bidders' reported values for auction j and $(0, \bar{v}_j^{-i})$ is the efficient outcome if bidder's i reported value were 0 and the other bidders' reports remain unchanged. Thus, each bidder pays the losing bids that would be winning if his own bid were set to 0.

In our context the auction outcome by applying the above rule is affected by the demand information "combinatorial" bidders possess and we need to investigate whether this assumption affects the optimality condition.

3.4. Bidding strategies

In this paragraph we will investigate the bidding strategies of all types of bidders under the above rules. We already know that sincere bidding (revealing the true values for the demanded quantity is an optimal strategy if the VCG rule is applied provided that each bidder has a privately known value for the good offered. Thus, the "individual" bidders' strategy is not affected by the presence of the "combinatorial" bidders. On the other hand a "combinatorial" bidder knows his own private value for the demanded bundle of goods but not the proportion of that value to each good separately. We propose the following bidding strategy and argue that it reserves the sincere bidding condition.

"Combinatorial" bidder i divides his total budget b^i (which equals v_{12}^i if he is sincere bidding)

according to the following rule: $b_1^i = \frac{D_1 / C_1}{D_1 / C_1 + D_2 / C_2} b^i$ and $b_2^i = \frac{D_2 / C_2}{D_1 / C_1 + D_2 / C_2} b^i$,

where $D_j = \sum_{i=1}^N q_j^i$, $j = 1, 2$. Factor D_i / C_i measures the competitiveness between the two goods

(the more aggressive the bidding is in one auction the higher the part of budget in this auction is) and forces "combinatorial" bidders to bid sincerely.

The reasoning behind this statement is as follows. Suppose that bidder i submits his whole budget b^i to the first auction and nothing to the second one ($\bar{s}_1^i = (q_1^i, b^i)$ and $\bar{s}_2^i = (q_2^i, 0)$).

Then he maximizes the probability of winning in auction 1 and minimizes the probability of winning in auction 2. He ends up with good 1 and makes negative profit since good 1 alone is useless to him (even if he has already paid for it). This is called the exposure problem that we will refer to below.

Suppose now that bidder i submits his whole budget b^i to both auctions ($(\bar{s}_1^i = (q_1^i, b^i))$ and $\bar{s}_2^i = (q_2^i, b^i)$) and lies about his true value. Then he maximizes the probability of winning in both auctions but maximizes the risk of paying more for the desired bundle than he is willing to (negative profit).

An intermediate case is when bidder i divides equally his budget to the two auctions ($\bar{s}_1^i = (q_1^i, b^i / 2)$ and $\bar{s}_2^i = (q_2^i, b^i / 2)$). Then he increases the probability of winning in the less competitive auction and decreases the probability of winning in the more competitive auction, again increasing the risk of obtaining only one of the two goods (negative profit).

Thus, bidding less or more than his true value or dividing randomly his budget in the two auctions, a “combinatorial” bidder can never increase his profit but may sometimes decrease it.

3.5. Social welfare loss - exposure problem

The proposed auction mechanism faces the well known ‘exposure problem’ [8]: it is described as the risk to obtain a bundle of goods, when combinatorial bidding is not allowed, that may lead to paying more for some individual goods than they are worth. This is due to the fact that we do not allow combinatorial bids and a bidder that desires a combination of goods might win in some auctions and lose in others, winning finally a bundle of goods that is smaller than the desired one. This has as a result a loss in social welfare so that the outcome might not be efficient. We investigate the magnitude of this problem in our approach by running experiments described in the next section.

4. EXPERIMENTAL RESULTS

This section presents the simulation methodology adopted in order to study the problem of this paper and provides some indicative simulation results. The section is concluded with a summary of the main findings from our study by means of simulations

4.1. Simulation Model and Methodology

In order to evaluate our auctioning scheme, we have implemented a specialized simulator in the Java programming language. We have conducted a plethora of simulations according to a very detailed model specifying:

- a) The bidding population mix, i.e. percentage of users bidding for resource 1, resource 2 and resources 1 and 2.
- b) The amount of users’ wealth, i.e. the per-unit price of each user’s bid.
- c) The quantity of units of resources sought.

In particular, for the bidders demanding units of just one resource, both the quantity and per-unit price of each user are individually randomly drawn from uniform distributions having support in $[1, q_{\max}^r]$ and $[1, p_{\max}^r]$ respectively. Note that the support of the distribution is in general different per resource r . This is justified by the fact that the two resources actually comprise different markets, where the expected willingness to pay of the users is in general different. For instance, the per-unit price of 1 Mbps of bandwidth is in general different than that of 1 virtual machine of a Grid computing provider.

For the bidders demanding – a generally different number of – units of both resources 1 and 2, the quantity of units sought per resource is drawn from two different uniform distributions, one per resource, having the same support with that of the bidders demanding just one resource; their wealth is drawn from a distribution having support which is in general different than that of the bidders demanding units of just one resource.

Also, note that since our approach relies on the splitting of the latter bidders' budget per resource based on the factor D_i/C_i , which measures the competitiveness between the two goods, it is also important to study the sensitivity of the simulation results on the accurate noise of this factor D_i/C_i . This is important since in cases where the providers are reluctant to announce (accurately) this factor to the market, the bidders may only deduce it by observing the two markets and the fluctuation of the auction cut-off prices in these markets. In this case, there may be a noise in the estimate of D_i/C_i , which will affect the budget splitting and ultimately the bidders' success in the two auctions. This is why in our simulation environment we have included a variable n that captures this effect of noise of estimate. If variable n is set to 1.0, then there is no noise in the estimate and essentially budget is split according to the D_i/C_i factor. A greater value of n implies that more of the budget is to be offered for the auction for resource 1 and less for the auction of resource 2, while the opposite that the user's bid for the sought units of resource 1 will be less compared to that that would be attained if the splitting had been done according to the factor D_i/C_i . In the simulations presented, unless explicitly specified otherwise, it is implicitly assumed that the splitting of budget is "accurately" done by means of the factor D_i/C_i .

Furthermore, note that the relations of the values used as the upper bound of the users' demand distributions compared to the respective capacity constraint of the resource, allows us to flexibly express a multitude of market conditions. For instance, assume the case where the value of q_{\max}^i is relatively close to that of C_i : this situation may arise in practise if few large ISPs compete in order to purchase bandwidth for Internet connectivity for their customers from the larger backbone providers of a region. Thus, this setting of parameters is suitable to depict an oligopoly market. On the contrary, a large difference of the values of these simulation model parameters may pertain to the case of a competitive market where many small users compete for the bandwidth of a broadband network.

Overall, we argue that our simulation model is both general enough and adequate to allow for a detailed and thorough study of the issues of interest of this paper. Finally, note that in order for our results to be unaffected by the noise of independent runs, we have adopted the approach of running a large number of experiments and subsequently averaging on the key performance metrics of interest.

Finally, it is also of great interest to assess the proximity of the performance of the individual auctions (assuming that bidders demanding units of both resources split their budget according to the D_i/C_i factor as explained earlier in this paper) to the best possible – in terms of social welfare – for each simulation run we also compute the social optimum outcome. This is done by means of a special algorithm that actually runs a combinatorial auction, which is the social optimum way of allocating the resources. Obviously, the exposure problem never arises under the combinatorial approach since the optimal allocation of resources is derived.

Note however that the latter would be impossible to do in practise for a multitude of resources and possible bundles, while the approach of running individual auctions can be applied in practise for an arbitrarily large number of resources and users. An additional problem of the

combinatorial approach is that it relies on the cooperation of the different providers, while revenue sharing in this context cannot be generally done in an unambiguously fair and transparent way, since bids are submitted for combinations of resources that are owned by multiple providers.

The main key performance metrics in our study are the following:

Exposure problem percentage: The percentage of bidders demanding units of both resources 1 and 2 that managed to reserve the units demanded at exactly one of the two auctions; these bidders are henceforth referred to as “exposed bidders”. This percentage is computed as the ratio of the exposed bidders over the total number of bidders demanding units of both resources 1 and 2.

Social welfare percentage loss: Having computed the social welfare of the individual auctions and that of the combinatorial approach (which attains the maximum social welfare) we compute the percentile deviation from the maximum. Note that in the calculation of the social welfare under the independent auctions approach the utility of the exposed bidders is taken to be 0.

4.2. Simulations

Having presented the simulation model and methodology, we now proceed to present some **indicative** simulations that demonstrate the main results of our study.

The first set of experiments is obtained from 30 simulation runs where the simulation parameters are set to depict an oligopoly market. In particular, there are 20 bidders demanding at most 10 units of resource 1 whose capacity is 100 units, 10 bidders demanding at most 5 units of resource 2 whose capacity is 30 units and finally 5 bidders demanding at most 10 units of resource 1 and at most 5 units of resource 2. The maximum per unit price for the first category of bidders (i.e. the upper bound of the support of the uniform distribution from which the value realizations are randomly drawn) is set to 1000 cents, while that of the other two categories of bidders is set to 2000 cents.

Note that due to the fact that the bidders competing for resources in both auctions for resources 1 and 2 do not have a higher average per-unit price compared to the bidders competing only for units of resource 2, this is a scenario where the exposure problem is likely to arise since the wealth of the bidders demanding both types of resources does not suffice to guarantee their success in both the auctions where they have to participate.

We now proceed to present the results for this group of simulations:

The average percentile social welfare loss (the average is computed over the 30 simulation runs) is 1.48%. It is also worth noting that in 21 out of 30 simulations (70% of the simulation runs) it is actually 0, thus the individual auctions where bidders employ the budget split according to the way proposed in this paper attain the maximum social welfare. Figure 1 depicts the social welfare and the maximum social welfare per simulation run.




Figure 1. Social welfare plot in an oligopoly setting.

The average percentage of exposed bidders is 7.3%. Note that this percentage is relatively high due to the small number of bidders (set to 5) that demand resources in both auctions. Thus, even if just one bidder suffers from the exposure problem, this immediately results in a 20% percentage. However, in 70% of the simulation runs of this group of experiments, the exposure problem never appears. This is a very good result, given that – as explained earlier in this subsection – this specific scenario is inherently susceptible to the exposure problem. The percentage of exposed bidders per simulation run is depicted as Figure 2.




Figure 2. The percentage of bidders demanding units of both resources 1 and 2 that suffered from the exposure problem.

The second set of experiments is obtained from 30 simulation runs where the simulation parameters are set to depict a competitive market where the size of the individual bidders demand is small compared to the capacity of the resources auctioned. In particular, there are 100 bidders demanding at most 5 units of resource 1 whose capacity is 500 units, 100 bidders demanding at most 3 units of resource 2 whose capacity is 100 units and finally 25 bidders demanding at most 10 units of resource 1 and at most 5 units of resource 2. The maximum per unit price for the all the bidders (i.e. the upper bound of the support of the uniform distribution from which the value realizations are randomly drawn) is set to 1000 cents.

We now proceed to present the results for this group of simulations:

The average percentile social welfare loss (the average is computed over the 30 simulation runs) is 1.8%. Thus, it is once more very small, almost negligible given the market size of this simulation run. Figure 3 depicts the social welfare and the maximum social welfare per simulation run.




Figure 3. Social welfare plot in a competitive setting.

The average percentage of exposed bidders is 16.36%. Note that this percentage is relatively high but this is mostly due to the fact that there are many high value bidders for the scarce resources 1 and 2 that bid just for units of one resource and whose willingness to pay is similar to that of the bidders desiring units of both resources. The percentage of exposed bidders per simulation run is depicted as Figure 2.

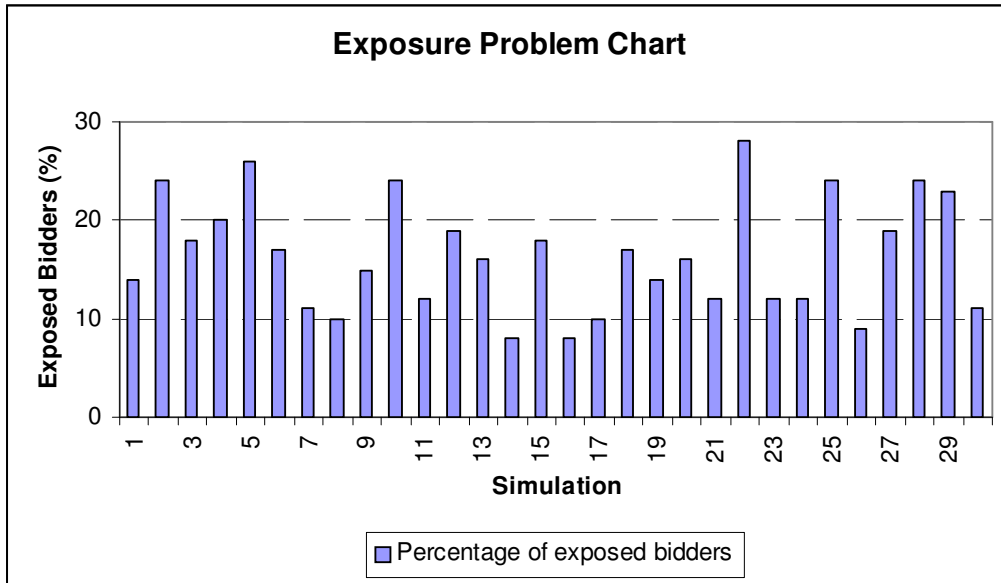


Figure 4. The percentage of bidders demanding units of both resources 1 and 2 that suffered from the exposure problem.

Finally, we have run a plethora of simulations to assess the impact of the noise of estimate of the noise in the estimate of the factor D_i / C_i that affects the bidding strategy of the bidders demanding units of both resources 1 and 2. In order to illustrate this, we keep the same setting with the previous group of simulations and introduce a noise in estimate of 20%. We then rerun a set of 30 simulations (keeping all other simulation parameters constant) in order to assess the impact of this noise in both the social welfare attained and the exposure problem incurred by the bidders.

The impact on the social welfare is small in terms of absolute numbers, i.e. there is no significant difference in the value of the percentile social welfare loss which in this case becomes 2,6%. This is due to the fact that the population of users that aim to reserve units in both resources 1 and 2 is relatively small compared to rival bidders; thus the impact of misallocation of goods due to the exposure problem is small in the total sum of all agents' utilities, i.e. the social welfare. Note however that the relative change is significant. So is the case with the percentage of the exposed bidders, which becomes 22.9%. This demonstrates the necessity of having an accurate estimate of the demand exhibited for the two resources but is also proof of the success of the budget splitting strategy proposed in this paper.

4.3. Main Findings

In this subsection, we briefly summarize the main findings of our simulations study. The main findings are as follows:

The performance of the proposed scheme is generally good and the social welfare attained is close to the optimal one.

- a) The exposure problem seldom arises and mostly pertains to bidders that would not be winners if a combinatorial approach were adopted.

- b) The proposed budget splitting method according to the D_i / C_i factor is successful.
- c) The accurate knowledge or alternatively an accurate estimate of the D_i / C_i factor further mitigates the exposure problem.
- d) Only large errors in the estimate of the D_i / C_i factor significantly reduce the performance of the individual auctions approach under the proposed budget splitting approach.

The generality of these results if the units of multiple, more than 2, resources are auctioned, is left for future work.

5. CONCLUSIONS

In this paper we propose a novel approach to allocate resources from different locations, using independent sealed-bid auctions for each type of resources performed by different service providers for each type. We apply the VCG payment rule and investigate the bidding strategies that force bidders reveal their true values and split them in an almost efficient manner among the various auctions. We perform experiments to show the applicability of our approach and verify the effectiveness of the results. In particular, we measure the extend to which the exposure problem appears in our context and the social welfare percentage loss. We show that the performance of the proposed scheme regarding the above measures is good and the social welfare is close to the optimal one. Our methodology is generalized to the case of considering more than two different resources which is left for future work. Another direction for further research, is the investigation of optimal strategies for providers (that is, which mechanism to choose) in case the assumption of known demand is eliminated.

REFERENCES

- [1] Vijay Krishna, (2002) "Auction Theory", Academic Press.
- [2] D. Lehmann, R. Mueller & T. Sandholm, (2006) "The Winner Determination Problem", Chapter 12 of the book *Combinatorial Auctions*, Cramton, Shoham, and Steinberg, eds., MIT Press.
- [3] Michael, H. Rothkopf, Aleksandar Peke/ & Ronald, M. Harstad , (1998) "Computationally manageable combinatorial auctions", *Management Science*, Vol. 44, No. 8, pp1131-1147.
- [4] V. Conitzer, J. Derryberry & T. Sandholm, (2004) "Combinatorial Auctions with Structured Item Graphs", in *Proceedings of the National Conference on Artificial Intelligence (AAAI)*.
- [5] Lawrence M. Ausubel & Peter Cramton, (2004) "Auctioning many divisible goods", *Journal of the European Economic Association*, Vol. 2, No. 2-3, pp 408-493.
- [6] Hee-Su Kim, (1996) "Equilibrium and efficiency in auctions of complementary goods without bundling", *Economic Letters*, Vol. 52, pp 49-54.
- [7] Paul Milgrom, (2007) "Package Auctions and Exchanges", *Econometrica*, 75, 935-966.
- [8] Peter Cramton, (2009) "Spectrum auction design", working paper University of Maryland.
- [9] P. Maillé, B. Tuffin (2004), "[Multi-Bid Auctions for Bandwidth Allocation in Communication Networks](#)", **IEEE INFOCOM**, Hong-Kong.

- [10] A.A.Lazar and N.Semret (1998) "The PSP Auction Mechanism for Network Resource Sharing", Proceedings of 8th International Symposium on Dynamic Games and Applications, Maastricht, the Netherlands, July 5-8, pages 359-365.
- [11] M. Dramitinos, George D. Stamoulis and Costas Courcoubetis (2007), "[An Auction Mechanism for Allocating the Bandwidth of Networks to their Users](#)", Elsevier IEEE Computer Networks.
- [12] Antonis Dimakis, Rahul Jain and Jean Walrand (2006), "[Mechanisms for Efficient Allocation in Divisible Capacity Networks](#)," *IEEE CDC*.

Authors

Marina Bitsaki is a visiting professor in the Department of Applied Mathematics at the University of Crete in Greece. She received a Diploma in Mathematics (1996) from the National & Capodistrian University of Athens, Greece, a M.Sc. in Statistics and Operational Research (1998) from the University of Essex, Great Britain, and a PhD in Computer Science (2004) from University of Crete, Greece. She is currently a member of the Transformation Services Laboratory at University of Crete, Greece. Her current research interests include the development of pricing schemes applied in network services, auction theory, information economy.

Manos Dramitinos is a researcher of the Network Economics and Services Group of the Athens University of Economics and Business. He received his Diploma (1998) and MsC (2000) in Computer Science from the Computer Science Department of the University of Crete, Greece. He obtained his Ph.D from the Athens University of Economics and Business in 2006 and postdoc in 2009 from INRIA Rhône-Alpes. He specializes in auctions, pricing, fairness and QoS issues in networks. He has received several awards, served as reviewer in many mainstream conferences and journals and also poster chair in ACM PE-WASUN 2009.