



A distributed auction-based algorithm to allocate bandwidth over paths

Hoang-Hai Tran, Bruno Tuffin

► To cite this version:

Hoang-Hai Tran, Bruno Tuffin. A distributed auction-based algorithm to allocate bandwidth over paths. NET-COOP 2010 - 4th Workshop on Network Control and Optimization, Nov 2010, Ghent, Belgium. inria-00597139

HAL Id: inria-00597139

<https://inria.hal.science/inria-00597139>

Submitted on 31 May 2011

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.

A distributed auction-based algorithm to allocate bandwidth over paths

Hoang-Hai Tran
INRIA Rennes Bretagne Atlantique
Rennes, France 35042
Email: tran_hoang.hai@inria.fr

Bruno Tuffin
INRIA Rennes Bretagne Atlantique
Rennes, France 35042
Email: bruno.tuffin@inria.fr

Abstract—In the literature, Vickrey-Clark-Groves (VCG) double-sided auctions have been applied to inter-domain traffic exchange because they provide incentives to be truthful and lead to an efficient use of the network, among relevant properties of mechanism design. Unfortunately, the resulting resource allocation scheme is neither budget-balanced nor solvable in a decentralized way, two important properties. We present a different but more realistic auction-based algorithm for allocating bandwidth over paths to end users or ISPs, leading to a new budget-balanced pricing scheme for which allocations and charges can be computed in a decentralized way.

I. INTRODUCTION

In communication networks, there are several mechanism design-based approaches to allocate network bandwidth over time [1]. Pricing-based mechanisms have especially been considered as a relevant way to provide incentives to agents involved in networking games to act “properly”. Several papers have been devoted to the study of such situations [2], [3], [4], [5], [6]. We consider here inter-domain resource allocation problem where the network is made of Autonomous Systems (ASes) trying to deliver their traffic by routes going through independent ASes which have to be “rewarded”. Inter-domain allocation problem was first introduced in the context of *mechanism design* in [7], [4]. The authors formulated incentive inter-domain routing as a mechanism design problem using Vickrey-Clarke-Groves (VCG) auctions. Most works on pricing and resource allocation have followed up their proposition [8], [9], [10], [6], [11], [12] since VCG is the only mechanism [13] which provides a general way of constructing a dominant-strategy, incentive-compatible, individually rational, and efficient mechanism. However, it is often neglected in those works that VCG auction mechanisms for multiple heterogeneous items may not be computationally tractable and budget-balanced [14], [15]. Combinatorial double-sided auctions on the other hand have been studied by R. Jain and J. Walrand [6], [16] but in centralized control. Our pricing mechanism is inspired by those works but is different in several significant aspects.

In our paper, we formulate the inter-domain allocation problem in which the players are both buyers and sellers at the same time. Our main contribution is though to propose a new pricing rule such that all allocations and prices can be computed in a decentralized way, a key criterion for large networks often neglected in the literature. We look for an algorithm that is BGP-compatible such that each AS negotiates

only with its neighbors. We then illustrate numerically the impact of this scheme on the inter-domain pricing game and compare it with the previous propositions.

This paper is organized as follows. Section II presents a combinatorial double-sided auctions mathematical model applied to inter-domain. In Section III, we describe a forward and backward induction algorithm to find the globally optimal solution while allowing each agent to execute our algorithm asynchronously. Section IV illustrates the power of our algorithm thanks to simulations, as well as the convergence of the pricing game to an ϵ -Nash equilibrium. The comparison of efficiency (in terms of social welfare) and incentive compatibility obtained with VCG auctions and combinatorial double-sided auctions is also provided. Finally, Section V concludes the paper with a brief discussion of future work.

II. INTER-DOMAIN RESOURCE ALLOCATION PROBLEM MODELING

The communication network is modeled as a graph $G = (V, L)$. There are n nodes in V , corresponding to the n agents, and links in L are the links between ASes. We assume that there is no capacity constraint. There is a set \mathcal{R} of routes. Each route $r \in \mathcal{R}$ is defined as an ordered list of nodes, each node appearing only once, and such that for two successive nodes in the list there exists a link $\ell \in L$ between those two nodes. Define also for node $v \in V$, $\mathcal{R}^S(v)$ the subset of routes starting at v and $\mathcal{R}^D(v)$ the subset of routes ending at v . Each provider v places buy-bids on a set of routes \mathcal{R}_v (note that routes are here initiated from v , but this can be extended without difficulty). For each route $r \in \mathcal{R}_v$, the bid is made of the maximum per unit price $\hat{c}_v^I(r)$ AS v is willing to pay and the maximum amount \hat{y}_r he is willing to get. He additionally places sell-bids $(\hat{c}_v^T(r), \hat{x}_v^T(r))$ for routes r AS v is on, where $\hat{c}_v^T(r)$ is the minimum unit price he wants to sell resource and $\hat{x}_v^T(r)$ the maximum amount he agrees to sell. Define \mathcal{R}_v^i this set of route going through v but not initiated from v , that is the routes where v has to transfer traffic.

The auctioneer designer seeks to determine allocations y_r on routes r as the solution of the following linear optimization problem

$$\max_y \sum_v \sum_{r \in \mathcal{R}_v} \hat{c}_v^I(r) y_r - \sum_v \sum_{r \in \mathcal{R}_v^i} \hat{c}_v^T(r) x_v^T(r) \quad (1)$$

such that

$$0 \leq y_r \leq \hat{y}_r, \forall r \quad (2)$$

$$y_r \leq \hat{x}_v^T(r) \quad \forall v, \forall r \in \mathcal{R}_v^i \quad (3)$$

$$x^T = Ay \leq \hat{x}^T, \quad (4)$$

where

$$A = (A_{vr})_{v,r} \text{ with } A_{vr} = 1 \text{ of } v \in r, 0 \text{ otherwise}$$

$$y = (y_r)_r$$

$$x^T = \left(\sum_r x_v^T(r) \right)_v.$$

The first constraint (2) says that the allocation flow should be matched with buyer's maximum demand \hat{y}_r and only non-negative allocations are allowed. The second constraint (3) means that at each node v the incoming traffic flow on each route r should not exceed the amount of traffic v is willing to sell. The third constraint says that total flow allocation x_v^T to seller v should equal the sum of flows on the routes through v .

Our new mechanism comes for the remark that (1) can hardly be solved in a decentralized way [15]. For this reason, we alternatively apply a modified double-sided auction. We consider routes r independently (thanks to the assumption about no capacity constraint, a relevant assumption in core networks), and therefore a single buyer with multiple sellers corresponding to intermediate nodes on the route. We will say that an allocation r^* is efficient for a source provider if it is a solution of the optimization problem

$$\arg \max_{r \in \mathcal{R}_i^S, y_r \geq 0} \hat{c}_i^I(r) y_r - \sum_{\forall r \in \mathcal{R}_i^S} \sum_{j \neq i} \hat{c}_j^T(r) x_j^T(r) \quad (5)$$

such that

$$0 \leq y_r \leq \hat{y}_r \quad (6)$$

$$x_j^T(r) \leq \hat{x}_j^T(r) \quad \forall j \in r \quad (7)$$

$$x_j^T(r) A_{jr} = y_r \quad \forall j, \forall r \in \mathcal{R}_i^S. \quad (8)$$

In Equation (5), $(\hat{c}_i^I(r), \hat{y}_r)$ denotes the buy-bid from source provider i on route r , $(\hat{c}_j^T(r), \hat{x}_j^T(r))$ denotes the sell-bids from intermediate nodes along route r . The allocation y_r is chosen in a way that each buyer's bid will be matched up with his maximum demand \hat{y}_r and seller's bid will be matched within their maximum supply $\hat{x}_j^T(r)$. In this optimization, the allocation is determined for a route or a set of routes for a single source-destination pair. The constraints have similar meaning as in the global optimization in (1) except for the last one saying that the actual flow should be the same on each intermediate nodes along route r .

The settlement price is called *the reserved price* and is determined at each intermediate node. It is the ask-price of the matched seller which provides the corresponding route r^* to destination. Notice that if the solution of the optimization at each intermediate node is a set of routes $(r)^*$ from multiple sellers, the reserved price at this stage is the highest ask-price

among matched sellers and the declared cost of an intermediate node will be made of its own declared cost plus the reserved price from himself to destination in the backward induction process, see Equation (10) later on. The source provider's payment is finally determined as the total reserved prices at each backward stage, which is included in the cost of the matched neighbors of the source node, being

$$\bar{p}_{r^*} = \max(\hat{c}_j^T(r^*)), \forall j \in \mathcal{N}_i^D. \quad (9)$$

Using surplus maximization in each allocation stage, the individual buyer can satisfy its demand with the lowest cost while pricing mechanism remains incentive-compatible for all sellers (recalling that they are neighbors of the buyer) but the matched seller with the highest ask-price. The theoretical and numerical analysis on the incentive-compatible property will be given in the sections III and IV. The difference between our proposed payment scheme and [16] is that our payment rule determined step-by-step at each stage of the backward induction. The source provider's payment is the cost of its matched neighbors, and it includes all reserved cost of the remaining nodes. In [16], the payment rule is determined at the end of the allocation process: the source provider pays each intermediate node the highest ask-price among matched sellers. As a result, the settlement price in [16] is always more than or equal to the settlement price at each allocation stage, leading to the higher revenues for the sellers than our proposed pricing scheme. For this reason, the social welfare¹ in our proposed pricing scheme is relaxed as compare to the social welfare in centralized combinatorial double-sided auction in [16]. The simulation results in the next section will show the social welfare comparison of different pricing schemes.

III. FORWARD AND BACKWARD INDUCTION ALGORITHMS

We now propose a distributed algorithm to find the solution of (5). The algorithm is made of sequential stages between neighbor nodes in order to exchange bidding information. The ASes send messages to each of their neighbors; the source node initiates the allocation process and then other neighbor nodes relay the message according to the following steps. The forward induction is described in Algorithm 1, where the goal is to determine a set of possible paths (\mathcal{R}_{src}^S) to destination. More precisely:

- The source node initiates the request messages including its maximum demand \hat{y}_r for the route r , and sends to its neighbor nodes. During the allocation process, each node indeed communicates only to its neighbor nodes, i.e \mathcal{N}_i^D the set of neighbor nodes of node i .
- If a neighbor node j agrees to forward \hat{y}_r and the request message has not previously been received, indicated by $\text{visited}(j) = \text{false}$, it also initiates a request message including its own maximum demand $\hat{y}_r(r)$ to relay traffic on route r , $\hat{y}_r(j)$ ² and propagates this request to its

¹Aggregated utilities of the sellers

² $\hat{y}_r(j)$ may be different from \hat{y}_r of the source node, depending on its own demand traffic to destination on route r and its own maximum supplies.

neighbor nodes. Otherwise, it sends a reject message to previous hop. The procedure is repeated until all request messages reach the leaf nodes, which are nodes one-hop away from the destination.

Algorithm 1 Building all possible paths to destination node

```

repeat
  for all  $j \in \mathcal{N}_i^D$  do
    sending request  $(r, \hat{y}_r)$ 
    if visited( $j$ ) = false then
      Add its maximum demand  $\hat{y}_r(j)$  to request messages
      Forward request messages to its neighbors
      Set visited( $j$ ) = true
    else
      Send reject message to previous hop
    end if
  end for
end for
until  $j := \text{destination}$ 

```

When this algorithm stops, we proceed backward on the route started from leaf nodes, each leaf node k declares its bid, $(\hat{c}_k^T, \hat{x}_k^T)$, including it in ACK packets and sends it back to previous hop. The second step of Algorithm 2 is to find a flow (or maybe a set of flows) with minimal cost from neighbor bids at each intermediate buyer j (i.e. $\arg \min(\hat{c}_k^T \hat{x}_k^T), \forall k \in \mathcal{N}_j^D$). Each node announces its bid up on the tree. The price an intermediate node announces is the price at which it was reserved from its neighbor including its own transit price. If rational, an intermediate node should propose a price at least its reserved price as a transit price in order to have positive payoff, i.e.,

$$\hat{c}_k^T = \bar{c}_k^T + \bar{\rho}_k, \quad (10)$$

where $\bar{\rho}_k$ is its reserved price for its neighbors and $\bar{c}_k^T(r)$ is its transit price. Backward induction process is then repeated until it reaches back the source node. Figure 1 illustrates how the message passes through intermediate nodes. In our proposed algorithm, each node plays both the seller and the buyer role at the same time.

Proposition 1: The backward induction process yields the minimal cost among possible routes at each intermediate node.

Proof: For each intermediate node j on route $r \in \mathcal{R}_j^S$, the chosen routes are based on local surplus maximization (5) which we recall

$$(r)^* = \arg \max_{r \in \mathcal{R}_i^S, y_r \geq 0} \hat{c}_j^I(r) y_r - \sum_{k \neq j} \hat{c}_k^T(r) x_k^T(r) \quad (11)$$

This problem is similar to minimizing the cost from preceding neighbor nodes on the built tree

$$\begin{aligned}
(r)^* &= \arg \min \sum_{k \in \mathcal{N}_j^D} \hat{c}_k^T x_k^T(r) \\
s.t. \quad &\sum_{k \in (r)^*} x_k^T \geq \hat{y}_r(j).
\end{aligned} \quad (12)$$

Algorithm 2 Finding the corresponding routes to destination backward on each intermediate node

```

repeat
  for all  $j \in \mathcal{T}_i$  do
    if  $j$  is next hop to destination then
      Attach its bid,  $(\hat{c}_j^T, \hat{x}_j^T)$ , to ACK
      Send ACK back to previous hop of  $j$ 
    else
      Find  $(r)^* = \arg \min(\hat{c}_k^T \hat{x}_k^T), \forall k \in \mathcal{N}_j^D$ 
      Subject to  $\sum_{k \in \mathcal{N}_j^D} \hat{x}_k^T \geq \hat{y}_r(j)$ 
      Attach its bid,  $(\hat{c}_j^T, \hat{x}_j^T)$ , to update messages of BGP protocol
      Send ACK back to previous hop of  $j$ 
    end if
  end for
until  $j := i$ ;

```




Fig. 1. Forward & backward induction illustration

Proposition 2: For given declared costs, our algorithm provides a decentralized solution of the surplus maximization (5).

Proof: With $(r)^*$ as the set of solution routes of the optimization following (5):

$$\arg \max_{r \in \mathcal{R}_i^S, y_r \geq 0} \hat{c}_i^I(r) y_r - \sum_{\forall r \in (r)^*} \sum_{\forall j \in r} \hat{c}_j^T(r) x_j^T(r). \quad (13)$$

In algorithm 2, the source node chooses a route with minimal cost or a set of such routes if demand is not fully met. Moreover, in Proposition 1, we have proved that each feedback route is formed by the route with minimal cost. As a result, we have

$$r^* = \arg \min \sum_{\forall r} \sum_{\forall j \in r} \hat{c}_j^T(r) x_j^T(r) \quad (14)$$

is the minimal aggregated valuation cost on intermediate nodes, providing the maximal value in (5). ■

Remark that this maximization property is in terms of the declared costs. Though, as we will see, since players' best

interest is not to bid truthfully, efficiency property will not be satisfied at the equilibrium point of the game on declared costs.

Proposition 3: The allocation outcome satisfies individual rationality.

Proof: The reserved price at each stage of an intermediate node is the highest ask-price among its matched sellers, which is $\bar{p}_i = \max(\hat{c}_j^T), \forall i \in r^*, \forall j \in \mathcal{N}_i^D$. Each intermediate node should pay the highest ask-price from its preceding neighbor and then receives a payment from its succeeding neighbors. Thus, the matched sellers never receive a payment less than their declared cost, leading to a non-negative payoff. The unmatched sellers get nothing, thus our algorithm verifies individual rationality. ■

Proposition 4: The mechanism is incentive-compatible except for the highest matched seller at each stage of backward algorithm in allocation process, assuming that each player has no information about others and is risk-averse.

Proof: At each stage of the allocation process, an unmatched seller has no incentive to bid lower than his reservation cost because it leads to a negative utility. If he bids higher, he cannot change its settlement price but may end up getting unmatched. Thus, the best interest for them is to bid truthfully. The matched seller has no incentive to bid higher than its reservation cost because his action may end up getting unmatched (by increasing its valuation cost in optimization) due to the lack of information about competitors, which is not approved by the risk-averseness assumption. Thus, at each stage of allocation process when a intermediate node negotiates with its neighbors, it is incentive-compatible for all sellers except for the matched seller with the highest ask-price. The mechanism remains fairly incentive compatible which means the pricing mechanism are not fully satisfied truthful bidding to all agents involved in the game, similarly to the centralized algorithm provided in [16]. The further analysis on the *incentive compatibility* property which is described as the degree of manipulation (the amount of deviation from truthful bidding) compared to VCG auctions will be given later. ■

IV. ALGORITHM ANALYSIS

In this section, we first illustrate numerically our resource allocation algorithm for a random, single pair of nodes. We then investigate the behavior of nodes reacting to our allocation and pricing scheme when they can change their strategies periodically.

A. Numerical illustration

We consider the simple topology with the real cost of transit traffic under each node described in Figure 2. We assume that the source node starts the forward induction process by sending out its maximum demand, $\hat{y}_r = 5$ to its neighbor ASes 1 & 2. An intermediate node receives the request message and forwards it to its own neighbors including the maximum demand from source node. The bidding information can be sent thanks to *update messages* in the BGP protocol. When the request messages are forwarded to leaf nodes, which are ASes

3, 4, and 5, it ends the forward induction process because leaf nodes could transfer directly traffic flow to destination node. The result is a set of possible routes \mathcal{R}_{src}^S to destination node which is simply illustrated in Figure 3.




Fig. 2. A simple topology

During the backward induction process, leaf nodes submit their sell bids backward along the path as depicted in Figure 4 including the maximal prices per unit they are willing to sell and their maximal bandwidth supplies. After deciding the corresponding route to destination, each intermediate node declares its own bid backward to the source node which is described as follows:

- Node 1 receives one bid from leaf node 5, $(\hat{x}_5^T, \hat{c}_5^T) = (5, 2)$, and the corresponding route to destination node through AS 5 has a total cost of 10, being the amount times price per unit traffic. The highest ask-price for matched sellers is 2 per unit traffic because there is only one matched seller at this stage. In order to have a positive net payoff, we assume that it submits its own bid $(\hat{x}_1^T, \hat{c}_1^T) = (5, 3)$ which includes the price \bar{p}_1 reserved for its neighbor AS-5 plus its own transit price \bar{c}_1^T , and then sends it back to the source node.
- Node 2 receives two bids from its neighbors which are $(\hat{x}_4^T, \hat{c}_4^T) = (3, 2)$ and $(\hat{x}_3^T, \hat{c}_3^T) = (2, 3)$. Thus, it should allocate corresponding routes from both ASes 3 and 4 due to their limited bandwidth supplies. The matched sellers are both ASes 3 and 4 with the highest ask-price of 3 per unit traffic leading to the total route cost of 15. Similarly, it submits its bid $(\hat{x}_2^T, \hat{c}_2^T) = (5, 4)$ and then sends it back to the source node.
- Source node receives two bids from ASes 1 and 2, which are $(\hat{x}_1^T, \hat{c}_1^T) = (5, 4)$ and $(\hat{x}_2^T, \hat{c}_2^T) = (5, 3)$ respectively. The shortest cost route is the one from AS 1 having total cost of 15 and the highest ask-price from this matched seller is 3 per unit traffic. As a result, the route r^* is allocated through ASes 1 and 5. The payment imposed on source node by AS 1 is 3 per unit traffic and afterwards AS 1 should pay 2 per unit traffic to AS 5. It ends our algorithm.

Remark that our mechanism is strongly budget-balanced. Also the chosen route on each intermediate node is the one providing the minimal cost from feedback information of its neighbors. If AS 1 increases its own transit cost to 5 per traffic unit, the allocated routes would change to AS 2 because we




Fig. 3. Forward induction




Fig. 4. Backward induction

would have $\hat{x}_2^T \hat{c}_2^T < \hat{x}_1^T \hat{c}_1^T$. In that case, the result allocation $(r)^*$ would be to use multiple routes $\{(2, 3), (2, 4)\}$. The settlement price imposed on source node by node 2 would be 4 per traffic unit and on node 2 by leaf nodes 3 and 4 it would be 3 per traffic unit.

In comparison with the centralized version in [16], the allocated routes are the same than with our algorithm but the settlement price imposed on source node is the highest ask-price from matched sellers which is 2 per traffic unit for both ASes 1 and 5, i.e., a different value. In case of a single allocated route from source to destination, an intermediate node receives as a payment exactly what he has bid.

B. Game on declared costs

We wish now to investigate numerically how ASes can play with their declared costs when they can change their bids over time instead of having a one-shot game, in order to improve their revenue. We consider the same network topology than in Section IV-A where the true reservation costs of nodes are randomly taken from a uniform distribution over the integer set $\{2, \dots, 8\}$, while the initial declared cost are also randomly chosen, by adding a random cost uniformly chosen in $\{0, 1, 2\}$ to their true cost. Assuming the bid process is played round by round, the nodes restart the allocation process after each round and update their declared transit cost (always taken larger than or equal to their true cost to ensure a non-negative revenue) based on previous history. To summarize the behavior, the matched sellers increase by a small amount their current declared cost to try to increase their revenue during next round, while the unmatched sellers decrease their current declared cost, if possible, to try to get matched by the mechanism

during next round. The amount added or subtracted at each round is randomly and uniformly taken within the set $[0, 0.5]$.




Fig. 5. Payoff of nodes over allocation process through a single route




Fig. 6. Payoff of nodes over allocation process through multiple routes

Figures 5 and 6 display the payoffs of the different nodes during the simulation. The results illustrate a convergence to a periodic behavior. This oscillatory behavior is due to the fact that, starting from round 6 in the allocation process illustrated in Figure 5, the matched sellers get unmatched next round if they increase by a small amount their declared cost. For the unmatched sellers, their best strategies are to bid close to their true reservation cost to become selected, and whenever they try to increase their utilities, they get unmatched next round again. In order to avoid such oscillations, the network manager can impose a bid fee $\epsilon > 0$ for changing costs after a given number of rounds, bid sufficiently high that no one has an interest in deviating from its current bid. Another option is to fix the number of rounds to a finite value unknown from the ASes, so that they cannot play strategically in terms of it. For example, if we chose a bid fee ϵ defined as the maximal gain of nodes from previous oscillation $\epsilon = \frac{\max(u_{osc}) - \min(u_{osc})}{3}$ since round 10 of allocation process, we have found a so-called ϵ -Nash equilibrium which is illustrated in Figure 7, where an ϵ -Nash equilibrium is a bid profile such that no player can improve his utility by more than ϵ .

³A bid fee ϵ is chosen as the maximal gain of nodes from previous oscillation, where u_{osc} is the net pay-off during oscillation.




Fig. 7. ϵ -Nash equilibrium

We can remark in this game that ASes declare costs larger than the true cost. It illustrates the difference with VCG auctions described in [4], [8], where it is a dominant strategy for nodes to bid truthfully. We illustrate in Figure 8 the impact of untruthful bidding by displaying the cumulative distribution of social welfare for our proposed pricing scheme, VCG auctions and the combinatorial double-sided auctions[16], given the random true cost choices. The social welfare is also taken with our proposed pricing algorithm at the ϵ -Nash equilibrium. It shows that even at ϵ -Nash equilibrium, the efficiency is relaxed too, as compared with the combinatorial double-sided auction and VCG auctions, but not to a large extent.

Nodes	True reservation cost	Average deviation from truthful bidding
1	7	16%
2	7	1.5%
3	3	14%
4	6	5%
5	4	18%

TABLE I
AVERAGE DEVIATION FROM TRUTHFUL BIDDING.




Fig. 8. Empirical cumulative distribution of the social welfare.

V. CONCLUSION

In this paper, we have proposed a BGP-compatible, decentralized resource allocation algorithm based on double-sided

auction. To have a decentralized and strongly budget-balanced allocation pricing scheme, two key properties, we have relaxed efficiency and incentive compatibility properties as compared to VCG auctions [4], [8] and combinatorial double-sided auctions [16]. We give some simulation results and illustrate that the system outcome reaches an ϵ -Nash equilibrium in a finite number of rounds if we charge a bid fee ϵ for any change of declared cost. This outcome is illustrated to be a good trade-off between complexity and efficiency.

VI. ACKNOWLEDGMENT

This work was done in the framework of the INRIA and Alcatel-Lucent Bell Labs Joint Research Lab on Self Organized Networks.

REFERENCES

- [1] E. Altman, T. Boulogne, R. El-Azouzi, T. Jiménez, and L. Wynter, "A survey on networking games in telecommunications," *Computers and Operations Research*, vol. 33, no. 2, pp. 286–311, 2006.
- [2] W. Wang, X. yang Li, and Y. Wang, "Truthful multicast routing in selfish wireless networks," in *In Proceedings of MobiCom*. ACM Press, 2004, pp. 245–259.
- [3] J. Crowcroft, R. Gibbens, F. Kelly, and S. Östring, "Providing incentives in providerless networks," *Ad Hoc Networks*, vol. 2, no. 3, pp. 283 – 289, 2004.
- [4] J. Feigenbaum, C. Papadimitriou, R. Sami, and S. Shenker, "A BGP-based mechanism for lowest-cost routing," *Distributed Computing*, vol. 18, no. 1, pp. 61–72, 2005.
- [5] X. Su, S. Chan, and G. Peng, "Auction in multi-path multi-hop routing," *Comm. Letters*, vol. 13, no. 2, pp. 154–156, 2009.
- [6] R. Jain and J. Walrand, "An efficient Nash-implementation mechanism for network resource allocation," *Automatica*, 2010.
- [7] N. Nisan, A. Ronen, and A. Ronen, "Algorithmic mechanism design," *Games and Economic Behavior*, vol. 35, pp. 166–196, 2001.
- [8] J. Feigenbaum, V. Ramachandran, and M. Schapira, "Incentive-compatible Interdomain routing," in *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*. New York, NY, USA: ACM, 2006, pp. 130–139.
- [9] S. Eidenbenz, G. Resta, and P. Santi, "COMMIT: a sender-centric truthful and energy-efficient routing protocol for ad hoc networks with selfish nodes," in *Parallel and Distributed Processing Symposium, 2005. Proceedings. 19th IEEE International*, 2005, p. 10 pp. [Online]. Available: <http://dx.doi.org/10.1109/IPDPS.2005.142>
- [10] L. Anderegg and S. Eidenbenz, "Ad hoc-VCG: a truthful and cost-efficient routing protocol for mobile ad hoc networks with selfish agents," in *MobiCom '03: Proceedings of the 9th annual international conference on Mobile computing and networking*. New York, NY, USA: ACM, 2003, pp. 245–259.
- [11] B. Sichao Yang, Hajek, "VCG-kelly mechanisms for allocation of divisible goods: Adapting vcg mechanisms to one-dimensional signals," *Selected Areas in Communications, IEEE Journal on*, vol. 25, no. 6, pp. 1237–1243, August 2007.
- [12] J. Hershberger and S. Suri, "Vickrey pricing in network routing: Fast payment computation," in *In Proc. of the 42nd IEEE Symposium on Foundations of Computer Science*, 2001, pp. 252–259.
- [13] H. Moulin and S. Shenker, "Strategyproof sharing of submodular access costs: Budget balance versus efficiency," Duke University, Department of Economics, Working Papers 96-31, 1996. [Online]. Available: <http://ideas.repec.org/p/duk/dukeec/96-31.html>
- [14] P. Maille and B. Tuffin, "Why VCG auctions can hardly be applied to the pricing of inter-domain and ad hoc networks," *3rd EuroNGI Conference on Next Generation Internet Networks*, pp. 36–39, May 2007.
- [15] H.-H. Tran and B. Tuffin, "Inter-domain pricing: Challenges and possible approaches," *International Journal of Network Management*, To appear 2010.
- [16] R. Jain and P. Varaiya, "Efficient market mechanisms for network resource allocation," *4th IEEE Conference and European Control Conference on Decision and Control, CDC-ECC '05*, pp. 1056–1061, Dec. 2005.