Pricing and QoS in Wireless Random Access Networks

Pavan Nuggehalli
Vanu, Inc.
Cambridge, MA 02139
Email:pavan@vanu.com

Jennifer Price
Department of Electrical and
Computer Engineering
University of Colorado at Colorado Springs
Colorado Springs, CO 80907
Email: jprice@eas.uccs.edu

Tara Javidi
Department of Electrical and
Computer Engineering
University of California San Diego
San Diego, CA 92093
Email: tara@ece.ucsd.edu

***Abstract*—In this paper, we examine the use of pricing for distributed, incentive-compatible and socially optimal resource allocation in a QoS-differentiated random-access wireless network. We argue that QoS mechanisms in wireless networks are susceptible to misuse by self-interested users. We first present a simple pricing scheme that leads to social optimality (*i.e.*, achieves QoS-differentiated proportional fairness) when users' utility functions are known to the AP. We then characterize the price of anarchy when users are strategic. Finally, the centerpiece of this paper is a pricing scheme that ensures socially optimal operation as a Nash equilibrium strategy among users whose utility functions are not known and who attempt to access the channel in a decentralized manner.**

I. Introduction

Many wireless medium access control (MAC) standards *e.g.*, IEEE 802.11e and 802.16 [1], [2], support differentiated Quality-of-Service (QoS) where users can either declare the QoS class of their applications or explicitly request for a certain level of performance. In the absence of any other mechanism, strategic users can exploit these QoS features to misrepresent their QoS class and obtain improved performance at the cost of other users [3]. For example, a user in an 802.11e network may declare her file transfer application as video to receive better throughput and bring down the throughput of other users. In this paper we propose a pricing scheme that achieves social optimality by maximizing the cumulative utility of the network. More specifically, we aim to achieve proportional fairness by assuming that each user's utility function is a scaled logarithm of her throughput, where the scale factor reflects the user's QoS requirement and willingness to pay.

We consider the uplink of a random access MAC where each user contends for channel access according to some user-chosen probability. We employ this somewhat simple MAC model in order to abstract out the essential features of QoS-enabled MAC and to avoid being overwhelmed by the complex details of more realistic MAC protocols. Note that, for window-based protocols such as IEEE 802.11, there is a direct correspondence between access probability and window size in steady-state [4], so the analysis here can be easily extended to such scenarios. Our model is even more applicable to slotted Aloha and its variants.

In commercially available implementations of the IEEE 802.11e standard, users can freely classify their application into four different priority levels or access categories (ACs) (commonly by modifying the appropriate IP TOS or VLAN tag header bits). As network interfaces become increasingly programmable, it is conceivable that users can gain deeper access to the MAC firmware and deviate from standard-mandated behavior. Our pricing scheme works as follows. Each user declares her utility function (*i.e.*, scale factor) to the AP. The AP then sets a price for each user which is a function of the user's declared utility function as well as her choice of access probability. In this distributed setting, pricing needs to play a comprehensive role: it should not only discourage strategic behavior but also propel the system to the socially optimal state in a distributed manner.

The contribution of this paper is as follows. First, we present a simple pricing scheme that leads to social optimality (*i.e.*, achieves proportional fairness) when users' utility functions are known to the AP. We then characterize the price of anarchy when users are strategic. The centerpiece of this paper is a pricing scheme that ensures socially optimal operation arises as a Nash equilibrium strategy among users whose utility functions are not known, and who attempt to access the channel in a decentralized manner. Implementing preference revelation as a dominant strategy in cases where the AP sets the price *and* determines the allocation of goods is usually achieved as a straight-forward application of the classical VCG mechanism (see [5]). However, it is remarkable that the distributed resource allocation problem, in which users are each responsible for choosing their own access probability, also admits a similar solution.

The remainder of this paper is organized as follows. Section II gives an overview of related work. In Section III, we introduce our model of the single-hop network and present a Kelly-type mechanism which ensures social optimality among price-taking users. In Section IV, we show that strategic users can significantly degrade the social utility of the system. In Section V, we provide the main result of the paper, a two-dimensional bid mechanism to ensure socially optimal allocation of access probability among strategic users. We conclude the paper in Section VI, where we discuss some areas of future research.

II. Related Work

Of course, the use of pricing for resource allocation in communication networks is by no means a new topic. There

978-1-4244-2324-8/08/$25.00 © 2008 IEEE.

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE "GLOBECOM" 2008 proceedings.

is a large body of work addressing distributed implementation and the tragedy of commons problem among price-taking users, following from the work by Kelly [6]. However, such schemes are known to suffer from a loss of efficiency in the presence of strategic users [7], [8]. On the other hand, there is also a large body of work addressing the preference revelation problem in the presence of strategic users. The most common example of such a mechanism is the well-known VCG mechanism (see, e.g. [9], [10], [11]). Although this type of mechanism achieves implementation of the socially optimal allocation in dominant strategies, it requires users to report their entire (possibly multi-dimensional) utility function, and usually requires a centralized allocation of resources.

Recently, there has been a great deal of research on mechanisms that address the impact of strategic users, but either require only scalar bids [12] - [15] or do not require centralized resource allocation [15], [16]. The model proposed in our paper, while similar to that in [12], has significant differences that prohibit the application of the pricing schemes developed in [12] to our model. Specifically, [12] considers a centralized allocation model in which the only form of misbehavior available to users is lying about their type; once they have reported their (possibly false) type, the actual allocation of resources (in our case, the access probability) is controlled by the AP[1]. In contrast, our model considers a more general case where the actual allocation of resources is distributed - hence users can not only lie about their type, but they can also choose an action (*i.e.*, access probability) autonomously (subject, of course, to the price signal and their own self-interest).

The main difference between [15], [16] and our problem is in the formulation of the resource allocation problem. Since the authors in [15], [16] work with hard constraints, they are able to provide a Tatonnement-like process which provides a decentralized solution in an iterative manner. In our setting, however, users' utilities are a function of the access probabilities of all users in the system - hence we work with a penalty-type formulation (i.e. soft constraints). As a result, the pricing scheme in our setting cannot depend on the implementation of an iterative algorithm. Moreover, it can be shown that a two-part pricing scheme consisting of a penalty-based price (to address distributed implementation) and a mechanism-based price (to address the strategic behavior of users) is insufficient to guarantee a socially optimal solution [17]. We also note that, under our pricing scheme, if the users know the form of the pricing function a priori, then a single parameter announcement to each wireless client from the AP is sufficient to achieve the socially optimal resource allocation.

There are a number of other related works in this area. Pricing schemes for QoS in wired networks have been studied recently in [18], [19], and [20]. The authors in [21] develop Kelly-type pricing mechanisms that achieve rate stabilization in random access networks, but do not address the notion of strategic users. The authors in [22] address the impact of strategic users in 802.11 networks, but do so by modifying the backoff intervals of users that have been classified as having misbehaved, and do not include explicit pricing or utility maximization.

[1]Notice that although the mechanism described in [12] is called a VCG-Kelly mechanism, it does not address the notion of distributed allocation of resources. There, the "Kelly" component refers to the fact that the pricing scheme requires only scalar bids.

III. System Model

We consider the uplink of a system of N users communicating with a single access point (AP), in which the network is assumed to be saturated. Time is divided into slots, and the system employs a slotted-Aloha like random access MAC where user k contends for channel access with probability u_k. We emphasize that each user selects her own value of access probability to use. A transmission is successful if and only if there is a single transmission attempt - there is no carrier sensing, and we do not model explicit back-off. QoS differentiation is achieved when users with high priority applications transmit more often than those with low priority applications.

Let $\bar{u} = [u_1, u_2, \ldots, u_N]$ be the vector of contention probabilities for the N users. The saturation throughput of the k^{th} user is given by

$$\tau_k = u_k \prod_{j=1, j\neq k} (1 - u_j) \quad (1)$$

User demand is assumed to be elastic, and the utility of user k is given by $U_k(\tau_k)$. More specifically user utility is given by

$$U_k(\tau_k) = \theta_k \log \tau_k \quad (2)$$

Here θ_k can be thought of as a parameter representing the priority of user k's application as well as the user's willingness to pay and takes values in the range $[0, \infty)$.

The cumulative system utility is given by

$$U = \sum_{k=1}^{N} \theta_k \log(\tau_k) \quad (3)$$

Note that the utility defined in (3) is non-positive because each $\tau_k <= 1$. One can add a sufficiently large positive constant to each user's utility function to make the utility function positive in the regime of interest. We do not explicitly include this constant in what follows.

Finally, the network is said to operate optimally if the access probabilities are chosen to maximize (3). Solving for the first order conditions (*i.e.*, taking the partial derivative with respect to u_k and equating it to 0), we find that the unique optimal solution is given by

$$u_k^* = \frac{\theta_k}{\sum_{j=1}^{N} \theta_j} \quad (4)$$

Global optimality is assured because the Hessian of (3) is negative semi-definite (a diagonal matrix with strictly negative diagonal entries), implying that the objective function is concave in $\bar{u}$.

A straightforward approach to reach the optimal state is to have the users convey the priority levels (θ_k) of the applications they are carrying to the AP. Suppose $\underline{\theta} = \{\theta_1, \ldots, \theta_N\}$. The AP can then calculate the optimal value of access probability for each user and inform all users. This approach will work only if users truthfully report their θ-values and subsequently adopt the access probability value suggested by the AP. In the next section, we examine what happens when these assumptions are violated.

IV. Non-Cooperative Game and the Price of Anarchy

In order to more carefully examine the notion of user incentives in the network described above, we formulate the reporting of θ-values and selection of access probabilities as a non-cooperative game. The game is structured such that the users know a priori that they will be charged for their access probability according to some known rule. Intuitively, the game consists of the following three stages:

1. *Parameter Announcements* - Each user reports a value $\hat{\theta}_k$, indicating its priority to the AP. The AP then broadcasts the price rate $\sum_{k=1}^{N} \hat{\theta}_k$ to all users. Note that, in general, $\hat{\theta}_k \neq \theta_k$.
2. *Choice of Access Probability* - Each user selects its access probability u_k in an attempt to maximize its own net utility (utility from achieved throughput minus the price charged by the AP).
3. *Price Selection* - The AP charges each user a price $p_k(\hat{\underline{\theta}}, u_k)$ according to a known rule.

Formally, we denote the access probability game by $F = [\mathcal{N}, \{\Sigma_k\}, \{U_k(\cdot)\}]$, where $\mathcal{N} = \{1, \ldots, N\}$ denotes the set of players, and $\Sigma_k = \{(\hat{\theta}_k, u_k) : \hat{\theta}_k \in [0, \infty), u_k \in [0, 1]\}$ denotes the strategy space for each player. (Since we are operating in an uncontrolled access mode, users have control over both their declared priority value and their access probability). We denote the net utility of user k by

$$U_k(\hat{\theta}_k, \hat{\theta}_{-k}, u_k, u_{-k}) = \theta_k \log\left(u_k \prod_{j \neq k}(1 - u_j)\right) - p_k(\hat{\theta}_k, \hat{\theta}_{-k}, u_k)$$

If we assume that users are truthful in reporting their θ-values but strategic in choosing their access probability, then we show that the following straightforward pricing scheme yields social optimality.

$$p_k(\hat{\underline{\theta}}, u_k) = u_k \times \sum_{j=1}^{N} \hat{\theta}_j \quad (5)$$

This scenario corresponds to a traditional Kelly-type setting where users are autonomous price-takers, i.e. they select their action strategically but do not lie about their utility function. This can be mapped to the access probability game by $F' = [\mathcal{N}, \{\Sigma_k\}, \{U_k(\cdot)\}]$, where $\mathcal{N} = \{1, \ldots, N\}$ denotes the set of players, and $\Sigma_k = \{(\hat{\theta}_k, u_k) : \hat{\theta}_k = \theta_k, u_k \in [0, 1]\}$ denotes the strategy space for each player.

Lemma 1: If users are restricted to choose $\hat{\theta}_k = \theta_k$, then the optimal access probability u_k^* given by (4) emerges as a Nash equilibrium of non-cooperative access probability game F'.

Proof: We use the fact that user k chooses its access probability to maximize its net utility - i.e. it solves

$$u_k = \arg\max_u \left[\theta_k \log\left(u \prod_{j \neq k}(1 - u_j)\right) - u \sum_{j=1}^{N} \hat{\theta}_j\right] \quad (6)$$

The solution to (6) is $u_k = \frac{\theta_k}{\sum_{j=1}^{N} \hat{\theta}_j} = \frac{\theta_k}{\sum_{j=1}^{N} \theta_j} = u_k^*$. ■

Notice that the optimal access probabilities emerge as a Nash equilibrium of the game when users are assumed to be price-takers, i.e. when $\hat{\theta}_k = \theta_k$ is enforced. When users are strategic (have an interest and/or ability to hide information about their θ-values), however, they may benefit by lying about their θ-values. Here, we show that in the presence of strategic users, the pricing scheme given by (5) leads to sub-optimal performance.

Theorem 1: For the pricing scheme in (5), there exists a unique Nash equilibrium of F given by $(\hat{\theta}_k, u_k) = (0, 1)\ \forall\ k$.

Proof: Suppose

$$\phi_{\theta,u} = \left[\theta_k \log\left(u \prod_{j \neq k}(1 - u_j)\right) - u \times (\theta + \sum_{j \neq k} \hat{\theta}_j)\right]$$

User k will choose $\hat{\theta}_k$ and u_k to maximize

$$(\hat{\theta}_k, u_k) = \arg\max_{\theta > 0, 0 \leq u \leq 1} \phi_{\theta,u} \quad (7)$$

The partial derivative of (7) with respect to θ is non-positive. Therefore, all users will report their θ-value to be 0 and set their access probability to 1 since the price is 0. ■

Here, we see that the situation is somewhat worse than simply achieving sub-optimal performance. The pricing scheme given by (5) actually results in a classic tragedy of commons situation. Since pricing is based only on reported values, greedy users attempting to maximize their own utility will choose $u_k = 1$. If all users behave similarly (which we can expect rational users to do), then each user receives zero throughput and the system becomes unusable. In other words, the short-term interest of users maximizing their own utility results in an exhaustion of the resource.

Despite this negative result, it is possible for the AP to impose some order by mandating a lower limit on reported θ-values, say $\theta_{\min}$. This will result in at least some level of QoS differentiation (albeit sub-optimal).

Theorem 2: Suppose the AP enforces a rule such that $\hat{\theta}_k \geq \theta_{\min}\ \forall\ k$. Then there is a unique Nash equilibrium given by $\hat{\theta}_k = \theta_{\min}$, and $u_k = \min\{1, \frac{\theta_k}{N\theta_{\min}}\}$.

Proof: From the proof of theorem 1, it follows that each user will report the lowest possible value of θ. The rest of the proof follows by optimizing the access probability value u_k. ■

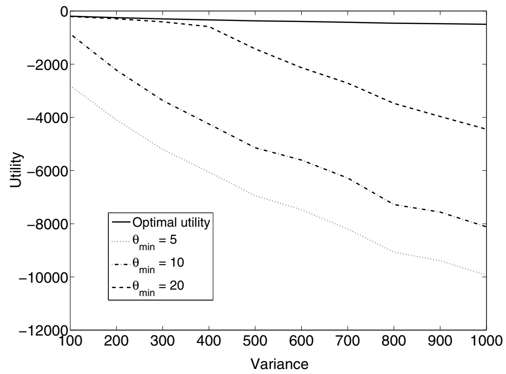

Fig. 1. System Utility in an Uncontrolled Access Mode Network with 4 Nodes

In Fig. 1, we plot the system utility achieved by a network operating in the uncontrolled access mode, as well as the optimal system utility as a function of the variance in θ-values, for a network with $N = 4$ nodes. Nodes were assigned θ values according to a uniform distribution, and results were averaged over 10000 runs for each value of variance. We see that, as expected, the system utility worsens as the lower limit on reported θ-values decreases. We also see that the system utility worsens as the variance of actual θ-values increases. In other words, the system utility actually becomes worse in networks where more QoS-differentiation is required!

V. Mechanism Design

In the previous section, we saw that constructing prices according to (5) was insufficient to align the selfish behavior of users with the socially optimal allocation of access probabilities. In this section, we introduce an alternate pricing scheme that addresses not only the preference revelation problem, but also eliminates the tragedy of commons phenomenon that can be observed in the uncontrolled mode when users are allowed to choose their own access probabilities. The basic structure of the game remains the same - users know a priori that they will be charged for their access probability according to some rule. In this case, however, the pricing function is:

$$\begin{aligned} p_k(\underline{\hat{\theta}}, u_k) &= \sum_{j\neq k} \hat{\theta}_j \log \left[\frac{\hat{\theta}_j}{\sum_{l\neq k} \hat{\theta}_l} \times \right. \\ & \left. \prod_{m\neq k, m\neq j} \left(1 - \frac{\hat{\theta}_m}{\sum_{l\neq k} \hat{\theta}_l}\right) \right] - \\ & \sum_{j\neq k} \hat{\theta}_j \log \left[\frac{\hat{\theta}_j}{\sum_{l\neq k} \hat{\theta}_l + \hat{\theta}} (1 - u_k) \times \right. \\ & \left. \prod_{m\neq k, m\neq j} \left(1 - \frac{\hat{\theta}_m}{\sum_{l\neq k} \hat{\theta}_l + \hat{\theta}}\right) \right] \qquad (8) \end{aligned}$$

We denote the non-cooperative game regulated by the pricing rule in (8) by $G = [\mathcal{N}, \{\Sigma_k\}, \{U_k(\cdot)\}]$, where $\mathcal{N}$ and Σ_k are as described in Section IV. The difference is in the net utility of user k, now given by

$$\begin{aligned} U_k(\hat{\theta}_k, \hat{\theta}_{-k}) &= \theta_k \log \left(u_k(\hat{\theta}_k, \hat{\theta}_{-k}) \prod_{j\neq k} (1 - u_j(\hat{\theta}_j, \hat{\theta}_{-j})) \right) \\ & \quad - p_k(\hat{\theta}_k, \hat{\theta}_{-k}, u_k(\hat{\theta}_k, \hat{\theta}_{-k})) \end{aligned}$$

Note that each user needs only two user-specific parameters from the AP to calculate price charged by the AP in this game.

Theorem 3: The optimal access probabilities $\underline{u}^*$ given by (4) emerge as a Nash equilibrium of the non-cooperative access probability game G.

Proof:

For simplicity, consider the strategy of the first user. This user first reports a priority-value $\hat{\theta}_1$ to the AP, then sets its access probability u_1. We first show that regardless of the value of $\hat{\theta}_1$ reported by the user, it will choose its access probability according to $u_1 = \frac{\theta_1}{\theta_1 + \sum_{j\neq 1} \hat{\theta}_j}$. This is readily seen by equating the partial derivative of user 1's utility with respect to u to 0, and observing that this point corresponds to a maxima. In addition, we note that this is true for all users - i.e. $u_k = \frac{\theta_k}{\theta_k + \sum_{j\neq k} \hat{\theta}_j}$.

With this optimal strategy for setting the access probability, it remains to show that user 1's net utility is maximized when it truthfully reports its θ-value. Since the goal is to show a Nash equilibrium strategy, we assume that $\hat{\theta}_j = \theta_j \ \forall \ j \neq k$. Thus, the net utility of user 1 as a function of its reported value $\hat{\theta}_1$ is given by

$$\begin{aligned} U_1(\hat{\theta}_1) &= \theta_1 \log \left[\frac{\theta_1}{\sum_{l=1}^N \theta_l} \prod_{j=2}^N \left(1 - \frac{\theta_j}{\sum_{l\neq 1} \theta_l + \hat{\theta}_1}\right) \right] \\ & \quad - \sum_{j=2}^N \theta_j \log \left[\frac{\theta_j}{\sum_{l\neq 1} \theta_l} \prod_{m\neq 1, m\neq j} \left(1 - \frac{\theta_m}{\sum_{l\neq 1} \theta_l}\right) \right] \\ & \quad + \sum_{j=2}^N \theta_j \log \left[\frac{\theta_j}{\sum_{l\neq 1} \theta_l + \hat{\theta}_1} (1 - u) \times \right. \\ & \quad \left. \prod_{m\neq 1, m\neq j} \left(1 - \frac{\theta_m}{\sum_{l\neq 1} \theta_l + \hat{\theta}_1}\right) \right] \qquad (9) \end{aligned}$$

To simplify notation, set $s = \sum_{l\neq 1} \theta_l$ and $s_j = \sum_{l\neq 1, l\neq j} \theta_l$. Differentiating U_1 with respect to $\hat{\theta}_1$ and collecting terms, we obtain

$$U_1(\hat{\theta}_1)' = \sum_{j=2}^N \frac{s_j + \theta_1}{s_j + \hat{\theta}_1} - (N-1) \frac{s + \theta_1}{s + \hat{\theta}_1} \qquad (10)$$

$U_1(\hat{\theta}_1)'$ is 0 if and only $\hat{\theta}_1 = \theta_1$ and $U''(\theta_1) < 0$, implying that user 1 maximizes its net utility by truthfully reported its θ-value.

We have shown that the strategy $u_k = \frac{\theta_k}{\theta_k + s}$ and $\hat{\theta}_k = \theta_k$ for all users k is a Nash equilibrium. This gives $u_k = \frac{\hat{\theta}_k}{\sum_{j=1}^N \theta_j} = \frac{\theta_k}{\sum_{j=1}^N \theta_j} = u_k^*$, and we are done. ■

Notice that although the pricing term given by (8) is similar in its form, it does not directly correspond to the well-known VCG mechanism. While each user's price depends on its own strategy $(\hat{\theta}, u_k)$, it does *not* depend on the access probabilities chosen by other users. In fact, it is precisely this distinction that allows us to implement such a pricing scheme in a distributed manner with a single round of feedback from the AP to each user. To see this, recall that the optimal access probability (from the perspective of a selfish user) is $u_k = \frac{\theta_k}{\theta_k + \sum_{j \neq k} \theta_k}$. When users are admitted to the system, they report a priority $\hat{\theta}$ to the AP. The AP then sends the value $\sum_{j=1}^{N} \hat{\theta}_j$ to each user. But notice that this, combined with a priori knowledge that prices will be charged according to (8), is all that users need to choose u_k in order to maximize their own selfish utility. This result is somewhat remarkable - a single price is sufficient to achieve *both* a socially optimal allocation of resources, and a simultaneous, distributed allocation of resources without the need for an iterative algorithm.

VI. Conclusions

In this paper we have identified the problem of strategic users exploiting QoS mechanisms to improve their performance at the cost of other users. We considered a simple yet general model of random access and showed that the effectiveness of straight-forward Kelly-like mechanisms depends on how much control users have over their access parameters. More specifically, when access parameters are controlled by the AP, the cost of strategic behavior is minimal. In stark contrast, allowing nodes to set their own parameters leads to the tragedy of the commons. The centerpiece of this paper is a pricing mechanism that ensures truthful behavior in both scenarios. We find it remarkable that such a distributed, one-shot, non-iterative pricing mechanism leads to optimal performance.

In the future, we intend to pursue two broad directions. First note that this paper has focused only on logarithmic utility functions. For more general utility functions, the theory provided here is not sufficient. The cumulative utility function (3) is, in general, not concave in the access probabilities even if the individual utility functions are concave functions of throughput. We would like to consider more general utility formulations in the future. Second, we want to apply the concepts developed in the paper to a real-life MAC protocol, namely the IEEE 802.11e standard. We would like to take advantage of recent work that analytically determines the saturation throughput in 802.11e networks and incorporate those models in our analysis.

References

[1] *Draft Supplement to Part II: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Medium Access Control (MAC)Enhancements for Quality of Service (QoS)*, IEEE Std. 802.11e/D6.0, 2003.

[2] *IEEE Standard for Local and metropolitan area networks Part 16: Air Interface for Fixed and Mobile Broadband Wireless Access Systems Amendment for Physical and Medium Access Control Layers for Combined Fixed and Mobile Operation in Licensed Bands*, IEEE Std. 802.16e, 2005.

[3] P. Nuggehalli, M. Sarkar, and R. R. Rao, "QoS and selfish users: a MAC layer perspective," in *IEEE Globecom*, Nov. 2007.

[4] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE J. Sel. Areas Commun.*, vol. 18, pp. 535–547, Mar. 2000.

[5] R. B. Myerson, *Game Theory: Analysis of Conflict*. Cambridge, MA: Harvard University Press, 1991.

[6] F. P. Kelly, "Charging and rate control for elastic traffic," *European Transactions on Telecommunications*, vol. 8, pp. 33–37, 197.

[7] R. Johari and J. N. Tsitsiklis, "Efficiency loss in network resource alocation game," *Mathematics of Operations Research*, vol. 29, pp. 407–435, 2004.

[8] S. Sanghavi and B. Hajek, "Optimal allocation of a divisible good to strategic buyers," in *IEEE Conference on Decision and Control*, Paradise Island, Bahamas, Dec. 2004, pp. 2748–2753.

[9] M.-C. A, M. Whinston, and J. Green, *Microeconomic Theory*. Oxford University Press, 1995.

[10] D. Fudenberg and J. Tirole, *Game Theory*. The MIT Press, 1991.

[11] M. Jackson, "Mechanism theory," in *Encyclopedia of Life Support Systems*, U. Derigs, Ed. EOLSS Publishers, 2003.

[12] S. Yang and B. Hajek, "VCG-Kelly mechanisms for allocation of divisible goods: Adapting vcg mechanisms to one-dimensional signals," *IEEE J. Sel. Areas Commun.*, vol. 25, pp. 1237–1243, Aug. 2007.

[13] R. Jain and J. Walrand, "An efficient Nash-implementation mechanism for divisible resource allocation," *IEEE J. Sel. Areas Commun.*, Submitted.

[14] R. Johari and J. Tsitsiklis, "Communication requirements of VCG-like mechanisms in convex environments," in *Proc. Allerton Conference on Communication, Control and Computing*, Monticello, IL, Sep. 2005.

[15] R. Maheswaran and T. Basar, "Efficient signal proportional allocation (EPSA) mechanisms: Decentralized social welfare maximization for divisible resources," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 5, pp. 1000–1009, May 2006.

[16] T. Stoenescu and J. Ledyard, "A pricing mechanism which implements a network rate allocation problem in nash equilibria," *Submitted to the IEEE/ACM Transactions on Networking*.

[17] P. Nuggehalli, J. Price, and T. Javidi, "Pricing and qos in wireless random access networks," *Preprint*.

[18] P. Marbach, "Priority service and max-min fairness," in *Proc. IEEE INFOCOM*, New York, NY, Jun. 2002.

[19] L. He and J. Walrand, "Pricing differentiated internet services," in *Proc. IEEE INFOCOM*, Miami, FL, Mar. 2005.

[20] J. Shu and P. Varaiya, "Pricing network services," in *Proc. IEEE INFOCOM*, vol. 2, 2003, pp. 1221–1230.

[21] C. Yuen and P. Marbach, "Price-based rate control in random access networks," *IEEE/ACM Trans. Netw.*, vol. 13, no. 5, pp. 1027–1040, Oct. 2005.

[22] P. Kyasanur and N. H. Vaidya, "Selfish MAC layer misbehavior in wireless networks," *IEEE J. Sel. Areas Commun.*, vol. 4, pp. 502–516, Sep. 2005.

