

Price Differentiation in the Kelly Mechanism*

Richard T. B. Ma, Dah Ming Chiu, John C. S. Lui, Vishal Misra, Dan Rubenstein

1. INTRODUCTION

Resource allocation [7, 4] and pricing [13, 1, 12] for network resources, e.g., bandwidth, have been studied extensively during the last decade. Among various proposed pricing and allocation mechanisms, the Kelly mechanism [7] stands out as a simple and scalable solution. Kelly et al [7, 6] first showed that it can be used as a *congestion pricing* mechanism to achieve *proportional fairness*. However, Johari and Tsisiklis [4] found that the resource allocation from the Kelly mechanism might induce an efficiency loss up to 25% of the social optimality. In this paper, we generalize the Kelly mechanism by designing a build-in price differentiation and show that the efficiency gap can be closed. In particular, we analyze the resource competition game under the generalized mechanism and show that any price differentiation induces a unique Nash equilibrium. We further reveal the relationship between the price differentiation and its resulting resource allocations and derive the optimal condition under which the social welfare reaches maximum. Our generalization extends the flexibility of the Kelly mechanism in two very different ways. First, it allows autonomous resource owners to use price differentiations so as to achieve individual objectives. Second, it provides a method for altruistic resource owners to make tradeoffs between user fairness (in terms of price differentiation) and system efficiency (in terms of social welfare).

1.1 Related Work

It was argued in [5] that the Kelly mechanism has low strategic flexibility (one-dimensional space) and low pricing flexibility (single price), and if we increase the strategic flexibility while preserving the single price restriction, the efficiency loss can be arbitrarily large. To achieve efficiency, many mechanisms have been designed by introducing pricing flexibility into the proportional allocation mechanism. Maheswaran and Basar [8, 9] modified the proportional allocation rule by adding a parameter ϵ to the total bids in the denominator and designed explicit price functions for players with different valuation functions. Nguyen and Vojnovic [11] introduced weights to the proportional allocation function and studied the revenue maximization problem for the resource provider. In our mechanism, the pricing flexibility is introduced as a build-in parameter and the proportional allocation remains the same; therefore, we view our mechanism as a generalization of the Kelly mechanism rather than an add-on pricing mechanism to the players.

*R. T. B. Ma is with Advanced Digital Sciences Center, Illinois at Singapore and National University of Singapore. D. M. Chiu and J. C. S. Lui are with Chinese University of Hong Kong. V. Misra and D. Rubenstein are with Columbia University.

2. RESOURCE ALLOCATION

We consider a set \mathcal{N} of rational users bidding for a share of divisible resource of capacity C . We define $N = |\mathcal{N}|$ as the number of users. Each $i \in \mathcal{N}$ has a valuation function $v_i(\cdot)$, where $v_i(d_i)$ defines the monetary utility to user i when she is given d_i amount of the resource. We make the same assumption as in [3, 5] as follows.

Assumption 1. Each valuation function $v_i(d_i)$ is concave, strictly increasing, and continuously differentiable over the domain $d_i \geq 0$, and the right directional derivative at 0, denoted $v_i'(0)$, is finite.

A common objective in resource allocation is to maximize the social welfare, i.e., the aggregate social utility of the system. In this context, it is to maximize the aggregate valuation of all users as the following optimization problem:

$$\text{Maximize} \quad \sum_{i=1}^N v_i(d_i) \quad (1)$$

$$\text{Subject to} \quad \sum_{i=1}^N d_i \leq C \quad \text{and} \quad d_i \geq 0, \forall i \in \mathcal{N}. \quad (2)$$

2.1 The Kelly Mechanism

In the Kelly mechanism [7], each user i submits a bid $w_i \geq 0$, which equals the payment for obtaining a share d_i of the resource. We denote u_i as the utility function of each user i , defined in a *quasi-linear* [10] environment as:

$$u_i(d_i) = v_i(d_i) - w_i,$$

which is the valuation of the allocated resource $v_i(d_i)$ less the payment w_i . The Kelly mechanism allocates the full capacity C among all users and the resource share d_i of each user i is proportional to her bid w_i . Mathematically, given a nonzero bid vector $\mathbf{w} = (w_1, w_2, \dots, w_N)$, the resource allocation vector $\mathbf{d} = (d_1, d_2, \dots, d_N)$ is defined by

$$d_i(\mathbf{w}) = \frac{w_i}{\sum_{j=1}^N w_j} C, \forall i \in \mathcal{N}. \quad (3)$$

Throughout this paper, the notation d_i and \mathbf{d} serve two purposes. When expressed without parenthesis, they are pre-determined resource allocation result or requirement. When followed by parenthesis, i.e., $d_i(\cdot)$ and $\mathbf{d}(\cdot)$, they represent the proportional allocation function and its value defined by Equation (3).

As a result of the Kelly mechanism, each user is charged the same unit price of the resource μ such that $\mu d_i = w_i$ for all users. This implicit unit price μ can be calculated as

$$\mu = \frac{\sum_{j=1}^N w_j}{C}. \quad (4)$$

2.2 The Generalized Kelly Mechanism

Rather than implementing a nondiscriminatory price μ under the Kelly mechanism, we consider a price differentiation among users. Our motivation of designing the price differentiation is to achieve different efficiency points for the social welfare defined as the objective function of (1). Under our generalization, we consider a strict positive price vector $\mathbf{p} = (p_1, p_2, \dots, p_N)$ as a parameter of the mechanism. Each user i submits a bid $t_i \geq 0$ to compete for resource. The resource allocation rule is the same proportional rule defined in Equation (3):

$$d_i(\mathbf{t}) = \frac{t_i}{\sum_{j=1}^N t_j} C, \quad \forall i \in \mathcal{N}. \quad (5)$$

The difference from the Kelly mechanism is that each user i pays $p_i t_i$ amount of money for the amount of shared resource $d_i(\mathbf{t})$ and therefore, obtains a utility of

$$u_i(\mathbf{t}, \mathbf{p}) = v_i(d_i(\mathbf{t})) - p_i t_i. \quad (6)$$

This generalized mechanism can be imagined as a process where users buy divisible tickets to compete for the resource. t_i is the number of tickets bought by user i and p_i is the monetary price for each ticket for user i . Like the Kelly mechanism, it allocates the fully capacity C among all users and the resource share $d_i(\mathbf{t})$ to each user i is proportional to the number of tickets bought: t_i . Although we do not differentiate tickets in resource allocation, the unit ticket price to users might be different. We denote $\mathcal{M}^{\mathbf{p}}$ as the mechanism associated with price vector \mathbf{p} . In particular, when $\mathbf{p} = \mathbf{1}$, the mechanism $\mathcal{M}^{\mathbf{1}}$ is equivalent to the Kelly mechanism.

Compared to the Kelly mechanism, the generalized mechanism achieves a similar virtual unit price ν in terms of tickets (measured as tickets per unit of resource) defined as

$$\nu = \frac{\sum_{j=1}^N t_j}{C}. \quad (7)$$

Consequently, the effective/real unit price for resource among users will be proportional to the price vector \mathbf{p} , because each user i 's real price becomes $p_i \nu$ (measured as dollars per unit of resource). Notice that although a pre-determined price is assigned to each user, the generalized mechanism inherits the simplicity/scalability of the Kelly mechanism in two ways: 1) the strategy space of the mechanism is still simple to be one-dimensional; and 2) only a single virtual price feedback, i.e., ν , is required to send to all users.

3. RESOURCE COMPETITION GAME

Kelly's original work [7] considered the *competitive equilibrium* (\mathbf{w}, μ) , a pair of a strategy profile and a unit price that satisfies:

$$v_i\left(\frac{w_i}{\mu}\right) - w_i \geq v_i\left(\frac{\hat{w}_i}{\mu}\right) - \hat{w}_i, \quad \forall \hat{w}_i \geq 0, \quad \forall i \in \mathcal{N}; \quad (8)$$

where μ is defined by Equation (4). The rationale of the competitive equilibrium relies on a price-taking assumption: the so-called market clearing price μ will not be affected by the action of a single player. Kelly proved that the competitive equilibrium solves the social welfare optimization problem (1)-(2).

The price-taking assumption works well when the system has a large number of players such that a single bids impact on the aggregation is infinitesimal. However, with only a few players at presence, each players bid will have a large impact on μ as well as other players resource allocation. Taking that into consideration, a further step is to model price-anticipating players and regard the resource allocation mechanism as a competition game, through which each player i chooses her strategy t_i to maximize her utility

of u_i . More precisely, given a mechanism $\mathcal{M}^{\mathbf{p}}$, each user i tries to choose strategy t_i that maximizes the utility:

$$\text{Maximize } u_i(t_i; \mathbf{t}_{-i}, \mathbf{p}) = v_i(d_i(\mathbf{t})) - p_i t_i, \quad (9)$$

where \mathbf{t}_{-i} denotes the strategy profile of users other than i . A strategy profile \mathbf{t}^* is a *Nash equilibrium* of the resource competition game if for any user i , the following is satisfied:

$$u_i(t_i^*; \mathbf{t}_{-i}^*, \mathbf{p}) \geq u_i(\hat{t}_i; \mathbf{t}_{-i}^*, \mathbf{p}), \quad \forall \hat{t}_i \geq 0. \quad (10)$$

Hajek and Gopalakrishnan [2] showed that the Kelly mechanism (i.e., $\mathcal{M}^{\mathbf{1}}$) induces a unique Nash equilibrium. Johari and Tsitsiklis [5] showed that the worst efficiency loss of the Nash equilibrium relative to the social optimality, the solution to (1)-(2), is 25%.

3.1 Properties of $\mathcal{M}^{\mathbf{p}}$

We first show that for any strictly positive vector \mathbf{p} , the generalized mechanism $\mathcal{M}^{\mathbf{p}}$ induces a unique Nash equilibrium. The result is parallel to Theorem 2 of Johari and Tsitsiklis[4], originated from Hajek and Gopalakrishnan [2]. We define $\mathcal{P} = \{\mathbf{p} \mid p_i > 0, \forall i \in \mathcal{N}\}$ as the feasible set of price vectors.

Remark: With $p_i \leq 0$, a player can keep increasing t_i to increase utility, and therefore no Nash equilibrium exists.

Theorem 1 (UNIQUE NASH EQUILIBRIUM). *Suppose $N > 1$ and Assumption 1 holds. For any $\mathbf{p} \in \mathcal{P}$, there exists a unique Nash equilibrium $\mathbf{t} \geq 0$ for the resource competition game under $\mathcal{M}^{\mathbf{p}}$, and at least two components of \mathbf{t} are positive. In this case, the resource allocation vector \mathbf{d} defined by:*

$$d_i = d_i(\mathbf{t}) = \frac{t_i}{\sum_{j=1}^N t_j} C, \quad \forall i \in \mathcal{N},$$

is the unique solution to the following optimization problem:

$$\begin{aligned} & \text{Maximize} && \sum_{i=1}^N \hat{v}_i(d_i) \\ & \text{Subject to} && \sum_{i=1}^N d_i \leq C \quad \text{and} \quad d_i \geq 0, \quad \forall i \in \mathcal{N}, \quad \text{where} \\ & && \hat{v}_i(d_i) = \frac{1}{p_i} \left[\left(1 - \frac{d_i}{C}\right) v_i(d_i) + \left(\frac{d_i}{C}\right) \left(\frac{1}{d_i} \int_0^{d_i} v_i(z) dz\right) \right]. \end{aligned}$$

Theorem 1 states that for any $\mathbf{p} \in \mathcal{P}$, there is a corresponding Nash equilibrium. Thus, we denote $\mathbf{t}^{\mathbf{p}}$ as the unique Nash equilibrium of $\mathcal{M}^{\mathbf{p}}$ that satisfies:

$$u_i(t_i^{\mathbf{p}}; \mathbf{t}_{-i}^{\mathbf{p}}, \mathbf{p}) \geq u_i(\hat{t}_i; \mathbf{t}_{-i}^{\mathbf{p}}, \mathbf{p}), \quad \forall \hat{t}_i \geq 0. \quad (11)$$

After knowing that each $\mathcal{M}^{\mathbf{p}}$ has a unique equilibrium, we ask the reverse question: Whether any resource allocation \mathbf{d} can always be realized as a Nash equilibrium of some mechanism $\mathcal{M}^{\mathbf{p}}$? If the answer is positive, the resource owner might be able to use \mathbf{p} as a control mechanism to achieve its desirable resource allocation \mathbf{d} for the system. Theorem 1 tells that at least two components of an equilibrium \mathbf{t} are positive. This implies that the resulting allocation under our generalized mechanism is *non-dictatorial*, meaning no single user can obtain the whole capacity C . We denote \mathcal{D} as the set of feasible non-dictatorial allocation defined as:

$$\mathcal{D} = \{\mathbf{d} \mid \sum_{j=1}^N d_j = C, \text{ and } 0 \leq d_j < C, \forall j \in \mathcal{N}\}.$$

Next, we show that any non-dictatorial allocation can be implemented as a Nash equilibrium of some mechanism $\mathcal{M}^{\mathbf{p}}$.

Theorem 2 (NASH IMPLEMENTATION). For any $\mathbf{d} \in \mathcal{D}$, there exists a mechanism $\mathcal{M}^{\mathbf{P}}$, whose unique Nash equilibrium $\mathbf{t}^{\mathbf{P}}$ induces the resource allocation of \mathbf{d} , i.e.,

$$d_i(\mathbf{t}^{\mathbf{P}}) = d_i, \quad \forall i \in \mathcal{N}.$$

In particular, if \mathbf{p} is defined by

$$p_i = h_i(d_i) = v'_i(d_i)(1 - \frac{d_i}{C}), \quad \forall i \in \mathcal{N}, \quad (12)$$

the Nash equilibrium strategy profile $\mathbf{t}^{\mathbf{P}}$ equals \mathbf{d} as well.

Although each $\mathcal{M}^{\mathbf{P}}$ induces a unique Nash equilibrium $\mathbf{t}^{\mathbf{P}}$, two different mechanisms might induce the same resource allocation and utilities for the players. We define the equivalence class of mechanisms as follows.

Definition 1. Two mechanisms $\mathcal{M}^{\mathbf{P}}$ and $\mathcal{M}^{\mathbf{Q}}$ are equivalent, if they induce the same resource allocation and utilities in equilibrium, i.e., $\mathbf{d}(\mathbf{t}^{\mathbf{P}}) = \mathbf{d}(\mathbf{t}^{\mathbf{Q}})$ and $\mathbf{u}(\mathbf{t}^{\mathbf{P}}, \mathbf{p}) = \mathbf{u}(\mathbf{t}^{\mathbf{Q}}, \mathbf{q})$ for all $i \in \mathcal{N}$.

Theorem 3 (LINEAR EQUIVALENCE). For any $\mathbf{p}, \mathbf{q} \in \mathcal{P}$ with $\mathbf{q} = k\mathbf{p}$ for some $k > 0$, the unique Nash equilibrium of $\mathcal{M}^{\mathbf{Q}}$ is $\mathbf{t}^{\mathbf{Q}} = \frac{1}{k}\mathbf{t}^{\mathbf{P}}$. Moreover, $\mathcal{M}^{\mathbf{Q}}$ is equivalent to $\mathcal{M}^{\mathbf{P}}$.

Theorem 3 states that when the price vector is scaled by a positive constant k , the resource allocation does not change in equilibrium. However, the strategy profile scales $1/k$ and keeps the prices unchanged. As a result, the user utilities will be the same in both equilibria too. A simple consequence of this theorem is as follows.

Corollary 1 (EQUIVALENT KELLY MECHANISMS). For any $k > 0$, the mechanism \mathcal{M}^k is equivalent to the Kelly mechanism \mathcal{M}^1 .

Theorem 4 (ZERO-ALLOCATION EQUIVALENCE). For any $\mathbf{p} \in \mathcal{P}$, we define $\mathcal{E}^{\mathbf{P}} = \{i | t_i^{\mathbf{P}} = 0\}$ as the set of users that get zero allocation in equilibrium and define the set $\mathcal{Q}^{\mathbf{P}}$ as

$$\mathcal{Q}^{\mathbf{P}} = \{\mathbf{q} | q_i = p_i, \quad \forall i \notin \mathcal{E}^{\mathbf{P}}; \quad q_i \geq p_i, \quad \forall i \in \mathcal{E}^{\mathbf{P}}\}.$$

For any $\mathbf{q} \in \mathcal{Q}^{\mathbf{P}}$, the unique Nash equilibrium of $\mathcal{M}^{\mathbf{Q}}$ is $\mathbf{t}^{\mathbf{Q}} = \mathbf{t}^{\mathbf{P}}$. Moreover, $\mathcal{M}^{\mathbf{Q}}$ is equivalent to $\mathcal{M}^{\mathbf{P}}$.

Theorem 4 tells that for the users that get zero resource in an equilibrium, increasing their prices will not change the equilibrium. They will use the same strategy profile and obtain the same utilities.

Theorem 1 implies that any $\mathcal{M}^{\mathbf{P}}$ maps to a unique resource allocation \mathbf{d} . Theorem 3 and 4 imply that this mapping can be many-to-one. Theorem 2 states that the mapping from $\mathcal{M}^{\mathbf{P}}$ to \mathcal{D} is indeed onto. From Theorem 3, we know that the mechanisms that have linearly dependent \mathbf{p} vectors are equivalent. Without loss of generality, we focus on the domain $\hat{\mathcal{P}} \subseteq \mathcal{P}$ defined inside a unit simplex as $\hat{\mathcal{P}} = \{\mathbf{p} | \sum_{j=1}^N p_j = 1, \text{ and } p_j > 0, \quad \forall j \in \mathcal{N}\}$. The following theorem reveals a one-to-one and onto mapping relationship between \mathcal{D} and a subset of $\hat{\mathcal{P}}$.

Theorem 5 (MAPPING). There exists a connected set $\tilde{\mathcal{P}} \subseteq \hat{\mathcal{P}}$ such that, the mapping $f: \tilde{\mathcal{P}} \rightarrow \mathcal{D}$ defined by

$$f(\mathbf{p}) = \mathbf{d}(\mathbf{t}^{\mathbf{P}}), \quad \forall \mathbf{p} \in \tilde{\mathcal{P}}$$

is continuous and bijective. Particularly, if $n = 2$, then $\tilde{\mathcal{P}} = \hat{\mathcal{P}}$.

If we focus on the set $\hat{\mathcal{P}}$, we eliminate the linear dependency of the \mathbf{p} vectors addressed in Theorem 3. Theorem 5 states that there always exists a subset of $\hat{\mathcal{P}} \subseteq \tilde{\mathcal{P}}$ that maps to \mathcal{D} continuously and bijectively. The mapping from $\hat{\mathcal{P}}$ to \mathcal{D} might still be many-to-one, because there might still exist equivalent mechanisms addressed in Theorem 4. However, when $n = 2$, $\hat{\mathcal{P}}$ maps to \mathcal{D} bijectively, because no user gets zero allocation in equilibrium.

3.2 Interpretation and Illustration

In this subsection, we connect the properties of the generalized mechanism into a big picture and use two examples to illustrate.




Figure 1: Mapping between prices and resource allocation.

Figure 1 visualizes the relationship between the domain of mechanisms $\mathcal{M}^{\mathbf{P}}$ and the resulting resource allocations \mathbf{d} . \mathcal{T} defines the space of feasible bidding profiles. \mathcal{V} defines the space of valuation functions under Assumption 1. By Theorem 1 and 2, we know that each \mathbf{p} maps onto \mathcal{D} ; however, the mapping is many-to-one due to the linearly equivalent mechanisms addressed in Theorem 3. If we normalize every $\mathbf{p} \in \mathcal{P}$ into $\hat{\mathcal{P}}$, the mapping from $\hat{\mathcal{P}}$ to \mathcal{D} might still be many-to-one. After reducing the equivalent mechanisms addressed in Theorem 4, we finally obtain $\tilde{\mathcal{P}}$ that maps to \mathcal{D} bijectively. The exact set of $\tilde{\mathcal{P}}$ and the mapping to \mathcal{D} totally depend on the underlying valuation space \mathcal{V} of the users.




Figure 2: An example of 2 users with linear valuation.

In the first example, we have $N = 2$, $C = 10$, $v_1(d_1) = \theta_1 d_1$ and $v_2(d_2) = \theta_2 d_2$ for some positive constant θ_1 and θ_2 . Figure 2 illustrates the resource allocation for both users under various $\mathbf{p} \in \hat{\mathcal{P}}$. p_1 varies along x-axis and $1 - p_1$ corresponds to p_2 . On the y-axis, we plot $d_1(\mathbf{p})$, i.e., the resource allocation to user 1 in equilibrium. We can also easily identify the resource allocation to user 2 as $C - d_1(\mathbf{p})$ correspondingly in the figure. We observe that the resource allocation depends on both the price and the ratio of θ_1/θ_2 . Particularly, when $\theta_1 = \theta_2$, the resource allocation is inversely proportional to the price of the users.

In the second example, we have $N = 3$, and $v_1(x) = v_2(x) = v_3(x) = \theta x$ for some $\theta > 0$. Figure 3 illustrates the mapping from the price simplex $\hat{\mathcal{P}}$ to the non-dictatorial resource allocation set \mathcal{D} . In particular, there is a proper subset $\tilde{\mathcal{P}} = \{\mathbf{p} | 0 < p_i \leq \frac{1}{2}, \quad i = 1, 2, 3\}$ that maps to \mathcal{D} bijectively. The points in $\hat{\mathcal{P}} \setminus \tilde{\mathcal{P}}$ are mapped to the triangular boundary of \mathcal{D} , where one of users gets zero allocation. One can also check that the mapping from $\tilde{\mathcal{P}}$ to \mathcal{D} satisfies $d_i(\mathbf{p}) = (1 - 2p_i)C$ for $i = 1, 2, 3$.

3.3 Valuation Revelation and Optimality

By Theorem 2, any non-dictatorial resource allocation can be achieved as a Nash equilibrium. Thus, in theory, we can close the




Figure 3: An example of 3 users with linear valuation.

25% efficiency gap by choosing an appropriate price \mathbf{p} that maximizes the social welfare. In practice, however, we need to know the valuation functions that are private information and may not be disclosed by the users. Next, we try to derive the hidden valuations via observing the resulting bidding profile of the users as follows.

Theorem 6 (OBSERVABILITY OF MARGINAL UTILITY). *Let $d_i = d_i(\mathbf{t}^P)$ be the equilibrium allocation for i . If one can observe the bids \mathbf{t}^P , under the Nash equilibrium, the marginal utility of user i at the resource level of d_i can be derived as*

$$v'_i(d_i) = \frac{p_i}{C - d_i} \sum_{j=1}^N t_j^P. \quad (13)$$

Theorem 6 states that by observing the bidding profile, one can derive users' marginal utilities at the allocated resource level. In principle, the resource owner can vary the prices so as to observe the marginal utilities of the users at equilibrium, and then reveal their hidden valuation functions. However, this method might be tedious and slow. If the objective is to maximize the social welfare, we can explore the relationship between the optimal resource allocation and the corresponding prices directly.

Theorem 7 (CONDITION OF OPTIMALITY). *Suppose the optimality of problem (1)-(2) is a strictly positive resource allocation vector, i.e., $\mathbf{d}^* \in \mathcal{D}$ and $d_i^* > 0$ for all $i \in \mathcal{N}$. A vector $\mathbf{p}^* > 0$ induces the unique Nash equilibrium with the allocation \mathbf{d}^* if and only if the following condition is satisfied:*

$$p_i^* : p_j^* = C - d_i^* : C - d_j^*, \quad \forall i, j \in \mathcal{N}. \quad (14)$$

In particular, when $N = 2$, \mathbf{d}^* is achieved when both users incur the same amount of price, i.e., $p_1^* t_1^P = p_2^* t_2^P$.

Theorem 7 states that for any pair of users i and j , the optimal price ratio should equal the ratio of $C - d_i^* : C - d_j^*$. This result not only gives us a way to verify the optimality without knowing hidden valuations, but it also provides a hint for the resource owner to see which user is over-allocated between any pair of users.

4. CONCLUSIONS AND FUTURE WORK

In this paper, we design a resource allocation mechanism that generalizes the Kelly mechanism via a build-in price differentiation. This generalization inherits multiple desirable properties of the Kelly mechanism, both mathematical (e.g. uniqueness of Nash equilibrium) and operational (e.g. simple virtual price feedback). By controlling the resource prices, the generalized mechanism can close up any efficiency gap and maximize social welfare. Actually, any non-dictatorial resource allocation can be realized as the unique Nash equilibrium of a price vector. Moreover, by observing

the bidding profile in equilibrium, one can also derive the underlying valuation functions of the users. This generalization largely extends the flexibility of the Kelly mechanism such that resource owners can choose resource prices to make tradeoffs between user fairness system efficiency. Some of the future directions include:

- In order to maximize social welfare in practice, feed-back control mechanisms might be designed based on the observed marginal utilities of the users and the optimality condition of the maximization under equilibrium.
- Adaptiveness of the mechanism can be studied under dynamic user arrival/departure.
- Instead of social welfare, a resource owner might be interested in profit. An orthogonal direction is to consider the revenue maximization from the resource owner's perspective.
- Extensions can be made for a multi-resource framework where users' utility depends on different types of resource.

In conclusion, the proposed generalization of the Kelly mechanism opens a whole spectrum of distributed resource pricing and allocation mechanisms that achieve various tradeoffs between user fairness and system efficiency, we believe that it will suit a larger range of future applications.

5. REFERENCES

- [1] M. Falkner, M. Devetsikiotis, and I. Lambadaris. An overview of pricing concepts for broadband IP networks. *IEEE Communications Surveys*, 3(2), 2000.
- [2] B. Hajek and G. Gopalakrishnan. Do greedy autonomous systems make for a sensible Internet? *Presented at the Conference on Stochastic Networks, Stanford University*, 2002.
- [3] R. Johari. Efficiency loss in market mechanisms for resource allocation. *PhD thesis, Massachusetts Institute of Technology*, June 2004.
- [4] R. Johari and J. N. Tsitsiklis. Efficiency loss in a network resource allocation game. *Mathematics of Operation Research*, 29(3):407–435, 2004.
- [5] R. Johari and J. N. Tsitsiklis. Efficiency of scalar-parameterized mechanisms. *Operation Research*, 2008.
- [6] F. Kelly, A. Maulloo, and D. Tan. Rate control in communication networks: shadow prices, proportional fairness and stability. In *Journal of the Operational Research Society*, volume 49, 1998.
- [7] F. P. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8:33–37, 1998.
- [8] R. T. Maheswaran and T. Basar. Nash equilibrium and decentralized negotiation in auctioning divisible resources. *Group Decision and Negotiation*, 12(5):361–395, 2003.
- [9] R. T. Maheswaran and T. Basar. Efficient signal proportional allocation (ESPA) mechanisms: Decentralized social welfare maximization for divisible resources. *IEEE Journal on Selected Areas of Communications*, 24(5):1000–1009, 2006.
- [10] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic theory*. Oxford University Press, 1995.
- [11] T. Nguyen and M. Vojnovic. Weighted proportional allocation. In *Proceedings of ACM Sigmetrics*, June 2011.
- [12] A. Odlyzko. Paris metro pricing for the Internet. *Proceedings of ACM Conference on Electronic Commerce (EC'99)*, pages 140–147, 1999.
- [13] S. Shenker, D. Clark, D. Estrin, and S. Herzog. Pricing in computer networks: Reshaping the research agenda. *Telecommunications Policy*, 20(3):183–201, 1996.