Auctions for Divisible Resources: Price Functions, Nash Equilibrium and Decentralized Update Schemes *

Rajiv T. Maheswaran and Tamer Başar

Coordinated Science Laboratory, University of Illinois
1308 W. Main Street,
Urbana, Illinois 61801-2307, USA
`(maheswar, tbasar)@control.csl.uiuc.edu`

Abstract. One of the many emerging applications for software agents are their ability to serve as proxies for trade and bartering. This has led to the analysis and development of auction protocols for various goods. We consider agent-mediated allocation of computational and network resources through market mechanisms. Single-good and combinatorial auctions do not apply readily to these products, thus we propose a divisible auction that is proportionally fair and has low signaling and computational costs. The structure of the auction enables us to represent optimal responses as price functions. From this we are able to characterize agent valuations and prove the existence of a unique Nash equilibrium. We further develop a decentralized algorithm that allows the agents to converge to the operating point without sharing private information.

1 Introduction

The *information economy*, described by Kephart, *et al* [7] is the merging of traditional markets, the Internet and autonomous agents to form a new marketplace where agents serve as proxies for buyers, sellers and intermediaries. Evidence of this can already be seen on the Internet as web sites such as Yahoo!, Amazon.com and Ebay are hosts to auctions where participants have the ability to give simple agents information about their valuations and have them bid as proxies. There have been many technical developments motivated by this phenomenon such as Kasbah [4], an agent-mediated marketplace.

We focus on markets for network bandwidth and computational resources. Companies such as Arbinet, RateXchange and Band-X.com have introduced exchange markets and bandwidth trading via auctions. A current push is to move toward dynamic real-time bandwidth trading. Invisible Hand Networks has created a product that allows for distributed real-time auctioning of Internet bandwidth. Similar methods have been suggested for allocation of computational

---

* Research supported in part by the AFOSR MURI Grant AF-DC-5-36128, and in part by the NSF ITR Grant CCR 00-85917

resources. Spawn [17] and POPCORN [11] are two often referenced systems that use markets for computational resource allocation.

Mechanisms for the sale of a single good do not apply to computational resources and network bandwidth. It is rare that these resources are allocated totally to a single user. Even if an agent chooses to purchase an entire resource, we expect at some point that the resource will be partitioned and it is that exchange that we are interested in analyzing. There has been a lot of research in analyzing combinatorial auctions [5,12,18]. However, these methods are inappropriate for the goods we are considering because network and computational resources are rarely partitioned into well-defined bundles. Thus, we consider divisible auctions as a market mechanism. The most relevant recent work in this area is the Progressive Second Price auction proposed by Semret [14]. It extends Vickrey's second price framework to a divisible resource.

In this setting, the performance of each agent is affected by the actions of all other agents. The autonomy of agents creates an environment where each agent is acting to better its own utility. The nature of this negotiation and the attempt to find an operating point calls for game theory [2], which has been applied to many similar settings in communication networks [1,15]. Decentralized algorithms that allow agents to converge to Nash equilibrium operating points have been analyzed [9,10].

As an agent economy evolves, a multitude of agents will migrate through resources rapidly with allocations being recomputed in microseconds. The signaling load and the computation required to perform the allocation will have a significant impact on the implementability of the allocation scheme. We desire an auction that minimizes these costs. We also wish to address the problem of the lying auctioneer [13] inherent in second price auctions. Finally, we want to capture the notion of proportional fairness that has been advocated in network [6] and computational settings [8,16].

The rest of the paper is organized as follows. Section 2 describes the allocation mechanism. Section 3 presents our model of agents' utilities. Section 4 shows how we can model optimal responses as a price function. Section 5 discusses how we can obtain information about agents' valuations from price functions. In Sections 6, 7 and 8, we apply our results to agents given a task to complete a series of jobs under different constraints. In Section 9, we prove the existence of a Nash equilibrium. We develop decentralized algorithms that converge to equilibrium without sharing private information in Section 10. Section 11 summarizes our work and presents directions for future research.

2 Allocation Mechanism

We begin with N agents competing for a resource with fixed finite capacity. The resource is partitioned based on the relative signals or bids sent by the agents. We assume each agent submits a signal s_i to the resource. Because minimizing cost of communication is of interest to us, we restrict our analysis to one-dimensional signals. In designing our auction, we want our allocations to be *proportionally* *fair* by weight. This holds if the allocation x^* satisfies $\sum_{i=1}^{N} s_i(x_i - x_i^*)/x_i^* \leq 0$ for any x where $\sum_{i=1}^{N} x_i = 1$ where s_i denotes the weights. This can be achieved with the allocation $x_i(s) = s_i / \sum_j s_j$. We note that this not only satisfies the proportionally fair criterion for allocation in networks, but it also matches the proportional share allocations in many of the systems proposed for computational resources. It takes $O(N)$ operations to calculate the proposed allocation, which is the minimal computational cost for making variable allocations to N agents. The PSP auctions requires a computation load of $O(NlogN)$ to calculate its allocation. In addition the signals in PSP are two-dimensional (price and quantity), which doubles the signaling load. The cost for each agent is $c_i(s,x) = s_i$. In this auction, if the feedback from the resource is the sum of all bids, an agent can immediately verify if it has been given an accurate allocation. If an agent knows the received allocation x_i and its own bid s_i, any bid total suggested by the auctioneer other than s_i/x_i can be immediately identified as a signal of an inaccurate allocation or a lying auctioneer. Furthermore, under this cost structure, each agent pays the same price per unit resource received.

3 Agent Utility

We assume that each agent has a valuation $v_i(x_i)$ for receiving an allocation x_i. This valuation may be a characterization of the estimated performance as a function of a given share of the resource. For example, it could be time to complete the processing of a job in a computational market, or the time to transmit some data given a particular share of network bandwidth obtained. Another derivation of the valuation could come from the estimated value of the sales that could be generated by obtaining a given share of the resource. This could be the case where agents act as brokers for computational resource and network bandwidth. We make the following assumptions about agents' valuations:

Assumption 1 *For all* $i \in \{1, \cdots, N\}$, $v_i(x_i)$ *is continuously differentiable* $v_i'(x_i) > 0$, *and* $v_i''(x_i) \leq 0 \quad \forall i \; x_i \in (0,1)$

The first assumption captures the notion that an agent's performance or marginal valuation of performance should not change dramatically given a miniscule change in allocation. The second assumption is made because with regard to computational and network resources, valuations tend to increase with allocation. The cases where valuations level off after a certain level of resource is achieved can be approximated with arbitrary precision by a strictly increasing valuation function. The third assumption captures the effect of diminishing returns.

Each agent's utility is the difference between the valuation and cost of its allocation, $U_i(s) = v_i(x_i(s)) - c_i(s,x)$. Substituting for allocation and cost, we have:

$$U_i(s) = U_i(s_i; s_{-i}) = v_i\left(\frac{s_i}{s_i + s_{-i} + \epsilon}\right) - s_i \tag{1}$$

where $s_{-i} = \sum_{j=1}^{N-1} s_j - s_i$ is the sum of the bids of all agents excluding the i-th agent and $s_N = \epsilon$ is a bid made by an agent representing the resource. By bidding ϵ, the resource declares a reservation value for its resource and prevents the possibility of agents colluding to purchase the resource for an arbitrarily small sum. The first order necessary condition for an interior solution is:

$$U_i'(s_i; s_{-i}) = v_i'\left(\frac{s_i}{s_i + s_{-i} + \epsilon}\right) \frac{s_{-i} + \epsilon}{(s_i + s_{-i} + \epsilon)^2} - 1 = 0.$$

This can be rewritten as follows:

$$v_i'\left(\frac{s_i}{s_i + s_{-i} + \epsilon}\right) - \frac{(s_i + s_{-i} + \epsilon)^2}{s_{-i} + \epsilon} = 0.$$

The LHS of the above equation is an increasing function of s_i as $v_i'(\cdot)$ is increasing in its argument, $s_i/(s_i + s_{-i} + \epsilon)$ is an increasing function of s_i and the second term has s_i only in the denominator. Thus, an interior solution exists if and only if the LHS is negative when $s_i = 0$. An agent will participate in the auction (submit a nonzero bid), if and only if:

$$v_i'(0) < s_{-i} + \epsilon. \tag{2}$$

If the condition is satisfied, there exists an interior extremal point. We can show $U_i''(s_i;_{-i}) < 0$, thus any interior extremal point is the unique value maximizing the agent's utility and is the agent's unique response when the total of all other agents' bids is s_{-i} and the resource bids ϵ. In the remainder of this paper, for notational simplicity, we will assume that the resource's bid is captured in the term s_{-i}.

4 Price Functions

Though the resource allocation is accomplished via an auction mechanism, each agent pays the same price per unit resource obtained. The auction can then be interpreted as a resource sold at a uniform price where the price is determined by the agents. The price per unit of the resource is $\sum_i s_i$, and each agent receives an allocation in proportion to that price. We can then define the *price* of a resource as the sum of all the bids, including the resource's bid, for that resource.

Next we define a *price function*, $p_i(x_i) : \Re \rightarrow \Re$ as the price at which the agent would choose an allocation of x_i. The price function represents the set of cost-allocation pairs which are the unique optimal responses of a given agent over a range of bids of other agents, i.e., $s_i = p_i(x_i)x_i$ is the unique optimal response to $s_{-i} = p_i(x_i)(1 - x_i)$. The inverse of the price function is the *demand function*, $d_i(p) : \Re \rightarrow \Re$, which is defined as the quantity of resource that the agent would desire if the price was p. This is again generated by an agent's unique optimal response in a way such that $s_i = d_i(p)p$ is the agent's reaction to $s_{-i} = (1 - d_i(p))p$. The price and demand functions are expected to be differentiable decreasing functions of their argument and the existence of one implies the existence of a well-defined inverse. One way to obtain these functions is to take the optimal response $s_i = f_i(s_{-i})$, substitute $s_{-i} = p - s_i$, and solve the fixed point equation $s_i = f_i(p - s_i)$. If a solution exists, one has s_i as a function of p. Then, making the substitution $s_i = px_i$, one can obtain an equation in terms of x_i and p from which the price and demand functions can be obtained. However, due to the nature of our auction, we can obtain the price function directly from an agent's valuation.

Proposition 1. *Given a valuation $v_i(x_i)$ that satisfies Assumption 1, there exists a corresponding differentiable decreasing price function characterized by $p_i(x_i) = v_i'(x_i)(1 - x_i)$.*

Proof. Let $f_i(s_i)$ be the i-th agent's unique optimal response. By the first order necessary condition, we have:

$$\begin{aligned} f_i(s_{-i}) + s_{-i} &= v_i'\left(\frac{f_i(s_{-i})}{f_i(s_{-i}) + s_{-i}}\right)\frac{s_{-i}}{f_i(s_{-i}) + s_{-i}} \\ &= v_i'\left(\frac{f_i(s_{-i})}{f_i(s_{-i}) + s_{-i}}\right)\left(1 - \frac{s_i}{f_i(s_{-i}) + s_{-i}}\right). \end{aligned}$$

By the definition of price, $p_i = f_i(s_{-i}) + s_{-i}$, and the allocation rule states $x_i = s_i/(f_i(s_{-i}) + s_{-i})$. Substituting this above, we have:

$$p_i(x_i) = v_i'(x_i)(1 - x_i).$$

We see that p_i is differentiable with derivative:

$$p_i'(x_i) = v_i''(x_i)(1 - x_i) - v_i'(x_i)$$

which is strictly negative given Assumption 1. Thus, the price function is decreasing. ■

This property of the auction lets us go directly from knowing an agent's valuation to the price function which is a transformation of its optimal response. We can obtain the optimal bid from the price function as follows:

$$s_i = f_i(s_{-i}) = f_i(s_{-i})\frac{f_i(s_{-i}) + s_{-i}}{f_i(s_{-i}) + s_{-i}} = x_i p_i(x_i) = v_i'(x_i)(1 - x_i)x_i$$

We see that $p_i(0) = v_i'(0)$ which states that if the price is greater than its largest marginal valuation, the agent will choose not to participate. This reflects the condition stated in Equation (2) derived from the first order necessary conditions. We also see that $p(1) = 0$, which states that the agent will purchase the entire resource if only the price is zero. This is equivalent to saying that the agent will demand the entire resource if the price was zero. This is a result of the structure of the auction where the only way an agent can obtain the entire resource is to be the only bidder in which case, the agent would make an arbitrarily small bid. This can never happen with the resource itself bidding ϵ. If the allocation was at equilibrium, the price per unit resource that the agent would be paying is less than its marginal utility by a factor of $(1 - x_i)$. This is the benefit gained by the agent for knowing its own effect on the price of the resource. Agents with larger allocations at equilibrium are able to scale their costs away from their marginal utility to a larger degree. In the case where there are many agents and each agent receives a small portion of the resource, i.e, $x_i \ll 1$, the prices being paid will be very close to the marginal valuations. The form of the price function also reflects the fact that shifting the valuation function by a constant will not change the optimal response, as the price function (and hence the agent's reaction function) depends only on the marginal valuation and not the absolute valuation.

5 Equivalent Valuations

A natural question to investigate is given a price function, under what conditions can one find an equivalent valuation function that captures the agent's behavior if its utility is modeled by Equation (1).

Proposition 2. *Let $p(x)$ be a positive continuous decreasing function of $x \in (0, 1)$ where $p(1) = 0$ and $p(0) = \lim_{x \searrow 0} p(x)$. If and only if $q(x) := p(x)/(1 - x)$ is differentiable, and $dq/dx \leq 0 \quad \forall\ x \in (0, 1)$, then the valuation function,*

$$v(x) = -\int_x^1 \frac{p(y)}{1-y} dy \tag{3}$$

satisfies Assumption 1, and produces the same optimal response characterized by $p(x)$.

Proof. This is straightforward to see as $v(x)$ is differentiable with $v'(x) = q(x) = p(x)/(1 - x)$ which is continuous as $p(x)$ is a continuous function. It is also straightforward to see that the price function generated by that valuation function is the same as the price function used in generating the valuation function, thus the optimal response that each characterizes will be the same. Because $p(x)$ is positive and $(1 - x)$ is positive for $x \in (0, 1)$, $v'(x) > 0$, we have $v'(x) > 0$ for $x \in (0, 1)$. We also have $v'(x) = dq/dx \leq 0$ for $x \in (0, 1)$ thus we have satisfied the conditions of Assumption 1. The "only if" results from the fact that if either $q(x)$ is not differentiable or $dq/dx > 0$ for some $x \in (0, 1)$, then either $v''(x)$ will not exist or will be positive for some $x \in (0, 1)$ violating the assumed properties of $v(x)$. ■

From Proposition 2, we note that the main factor in whether a price function has an equivalent concave valuation is whether $p(x)/(1 - x)$ is a decreasing functions as all the other conditions of Assumption 1 are satisfied through construction or by natural properties of $p(x)$. We can prove that it is necessary for a price function to lie beneath the line segment defined by the points $(0, p(0))$ and $(1, p(1))$, to have a concave valuation. Furthermore, every convex price function has a concave equivalent valuation. We apply the previous results to a specific agent task with varying objectives in the following sections.

6 Jobs in Series

We consider a scenario where agents are generated with a sequence of jobs to complete. The jobs require access to resources available at various nodes throughout the network. The agent attempts to purchase resources throughout the network to complete its set of jobs according to a given performance measure. Let us assume that the i-th agent has a sequence of jobs with K_i tasks, where q_i^k is the size of the job of the k-th task. Let C_i^k be the capacity of the resource providing the service needed by the k-th task of the i-th agent. We assume that every resource allocates its services using the proportionally fair auction. In this context, the bid will constitute a payment that the agent is willing to make per unit of time that it uses the resource. Let s_i^k be the bid of the i-th agent for the resource chosen for the k-th task on its itinerary and s_{-i}^k be the sum of the bids of other agents competing for that resource (which includes the bid ϵ^k made by the k-th resource). If the rate of service obtained by the i-th agent for its k-th task is x_k^i, and the time taken to complete the job is t_i^k, and the expense to the agent is e_i^k, we have

$$x_i^k = C_i^k \left(\frac{s_i^k}{s_i^k + s_{-i}^k} \right), \quad t_i^k = \frac{q_i^k (s_i^k + s_{-i}^k)}{C_i^k s_i^k}, \quad e_i^k = s_i^k t_i^k = \frac{q_i^k (s_i^k + s_{-i}^k)}{C_i^k}.$$

The decision that faces the agent is how to balance its performance as measured by the time taken to complete its jobs and the cost of obtaining service. We first consider the following criterion:

$$\min \sum_{k=1}^{K_i} e_i^k + \alpha_i^k \sum_{k=1}^{K_i} t_i^k$$

where α_i^k represents the value to the i-th agent of the time taken to complete the k-th job relative to its cost. Let p_i^k be the price of the resource chosen by the i-th agent to complete its k-th task. The price is equivalent to the sum of all the bids made at that resource so $p_i^k = s_i^k + s_{-i}^k$. We can show that the optimal bids for the i-th agent are:

$$s_i^k = \frac{-\alpha_i^k + \sqrt{(\alpha_i^k)^2 + 4\alpha_i^k p_i^k}}{2}$$

which is the optimal response in terms of the price of the resource. The optimal bid at the k-th resource for the i-th agent can also be expressed as a price function:

$$p_i^k = \frac{\alpha_i^k (1 - x_i^k)}{(x_i^k)^2}.$$

Applying the results of Proposition 2, we find the equivalent valuation function:

$$v_i^k(x_i^k) = \frac{-\alpha_i^k}{x_i^k}.$$

We are able to extract a valuation function even though the objective was not a maximization of quasilinear utility. Here, the minimization problem can be decoupled with respect to each job. Then for each job, the problem can then be interpreted as a minimization problem for a job broken into several pieces (of arbitrarily small size) to be completed in series with the same objective. In the limit, the problem tends to an instantaneous optimization and thus, an instantaneous valuation of service.

7 Finite Budget

We now consider the situation where an agent is given an endowment, E_i, that it may not exceed as it attempts to minimize the total time taken to complete its jobs. There is no benefit for returning any of the endowment. This can be expressed as the following optimization problem:

$$\min \sum_{k=1}^{K_i} t_i^k \qquad \text{s.t.} \qquad \sum_{k=1}^{K_i} e_i^k \leq E_i.$$

We use Lagrangian methods, to determine the optimal ratios of bids, $s_i^k = s_i^j \sqrt{s^k_{-i}/s^j_{-i}}$. By substituting for optimal future bids ($k > 1$) in terms of the bid for the current resource ($k = 1$), we obtain the following optimal bid:

$$s_i^1 = \frac{E_i - \sum_{k \neq 1} \frac{q_i^k}{C_i^k} s^k_{-i} - \frac{q_i^1}{C^1} s^1_{-i}}{\frac{q_i^1}{C^1} + \sum_{k \neq 1} \frac{q_i^k}{C_i^k} \sqrt{\frac{s^k_{-i}}{s^1_{-i}}}} = \frac{\alpha^i - \beta^i s^1_{-i}}{\beta^i + \frac{\gamma^i}{\sqrt{s^1_{-i}}}},$$

where

$$\alpha_i := E_i - \sum_{k \neq 1} \frac{q_i^k}{C_i^k} s^k_{-i}, \quad \beta_i := \frac{q_i^1}{C^1}, \quad \gamma_i := \sum_{k \neq 1} \frac{q_i^k}{C_i^k} \sqrt{s^k_{-i}}.$$

This is the optimal bid for the first or current job of the i-th agent in terms of the demand of the resources for the jobs in its itinerary. For an agent to implement this strategy, it must have estimates of the demand at resources that it plans to visit in the future. This fits with the notion of a finite budget as one has to have an idea of how much money one will need in the future to know how much one can reasonably spend now. The reaction function is parameterized by the agent's beliefs about the future.

Intuitively, α_i represents the estimate of the maximum money available for the current job. If that amount is less than zero, the agent cannot afford to purchase service under the current state of the network. This is reflected in the optimal response, as a negative α_i would yield a negative s_i^k since β_i and s^1_{-i} are positive and agents are required to submit non-negative bids. In fact we see that, the optimal response will yield a negative bid whenever $s^1_{-i} > \alpha_i/\beta_i$.

There is intuition behind this as well. Because β_i represents the minimum time required to complete the current job, and α_i is the maximum money available for the current job, α_i/β_i is the largest amount of money per unit time that the i-th agent could spend or bid for this resource. If the other agents' total bids, s^1_{-i}, create a price that is greater than the i-th agent's spending limit, it will choose not to participate. The third parameter, γ_i is a factor that when divided by $\sqrt{s^1_{-i}}$ gives an estimate of the excess time (time above the minimum time to complete a job) necessary to complete the remaining tasks. Thus, the optimal response is the ratio of the excess money for all the jobs to the estimated excess time for all the jobs. We drop the superscript, with the knowledge that the bids are in reference to the current resource. Making the substitution, $s_{-i} = p_i - s_i$, we can solve for s_i as a function of p_i. By then making the substitution $s_i = p_i x_i$, we can solve for the demand function and the price function:

$$p_i = \frac{\alpha_i}{\beta_i} + \frac{1}{2\beta_i^2(1-x_i)}\left[\gamma_i^2 x_i^2 - \sqrt{\gamma_i^4 x_i^4 + 4\alpha_i\beta_i\gamma_i^2 x_i^2(1-x_i)}\right]$$

An economy of agents with budget limits described in this section has been simulated in [3]. Plots of the price function for various levels of endowment can be seen in Figure 1(a). We can see that as endowments vary, the price curves can exceed the line segment defined by $(0, p(0))$ and $(1, p(1))$, which is a necessary condition for a concave equivalent valuation. To gain an understanding of this case, we investigate the resulting valuation when applying the transformation in Proposition 2. The second order condition for concavity of agent utility implies that if

$$v''(x)(1-x) < 2v'(x),$$

then x is a unique maximizing allocation if it meets the first order necessary condition $p(x) = v'(x)(1-x)$. We note that for all decreasing price functions, we have

$$p'(x) = v''(x)(1-x) - v'(x) < 0.$$

Given a decreasing price function, if we substitute the valuation function generated by Equation 3, into the second order condition, we have

$$\begin{aligned} \frac{p'(x)(1-x)+p(x)}{(1-x)^2}(1-x) &< \frac{2p(x)}{1-x} \\ p'(x) &< \frac{p(x)}{1-x}. \end{aligned}$$

Since the LHS is negative and the RHS is positive we know this is satisfied. The preceding show that a valuation function does not necessarily have to be concave for a unique maximizing response to exist. If a valuation is strictly convex, the effect is to push the agent into higher equilibrium allocations and higher equilibrium costs. This is counter to the concept of diminishing returns but the complementarities induced by the coupling of the jobs with the same finite budget forces the agent to deviate from a concave valuation. This gives us some intuition into the changes in the price functions as parameters vary.

As E_i increases, the agent has more money to spend and therefore can accept a higher price for each allocation. Thus, the price curves are uniformly higher as the endowment increases. Also, as an agent has more money to spend it is encouraged to purchase higher allocations which would explain why the curves get more concave (implying a convex valuation) as the endowment increases. As the current job size, q_i^1, increases and the budget remains the same, the agent cannot afford to spend as much on the current job, thus the price curves become uniformly lower. However, as the current job size increases with respect to the future jobs, its effect on overall performance increases as well, and even though the agent cannot spend as much money, it is encouraged to seek a higher allocation, thus the price curve becomes more concave as q_i^1 increases. Increasing future job sizes q_i^k and future demand s_{-i}^k for $k > 1$, both have the effect of increasing the importance of the future jobs and minimizing the importance of the first job. This is why we see the price functions get uniformly lower and progressively convex as we want the agent to spend less money and settle for lower allocations as it is necessary to save more of the finite endowment for the future.

This scenario shows the robustness of looking at an optimal response in the form of a price function as it can capture a greater range of agent valuations. It is important that the price functions are associated with valuations so that we can find a direct relation between motivation and action. By limiting ourselves to strictly concave valuations (and strictly convex valuations that satisfy the modified second order condition), we assure ourselves of unique responses. More complex forms of valuations is an open area for further research.

8 Finite Time

We consider another agent task where a sequence of jobs needs to be completed in a specified amount of time, T_i, while minimizing the cost accrued. This can be expressed as the following optimization problem:

$$\min \sum_{k=1}^{K} e_i^k \qquad \text{s.t.} \qquad \sum_{k=1}^{K} t_i^k \leq T_i.$$

Again, by Lagrangian methods, we are able to obtain the optimal ratios between bids, $s_i^k = s_i^j \sqrt{s_{-i}^k / s_{-i}^j}$. By substituting for optimal future bids ($k > 1$) in terms of the bid for the current resource ($k = 1$) in the constraint, we obtain the following relationship for the optimal bid:

$$\beta_i s_{-i}^1 + \gamma_i \sqrt{s_{-i}^1} = \alpha_i s_i^1$$

where

$$\alpha_i := T_i - \sum_{k=1}^{K} \frac{q_i^k}{C_i^k} \quad \beta_i := \frac{q_i^1}{C^1} \quad \gamma_i := \sum_{k\neq 1} \frac{q_i^k}{C_i^k}\sqrt{s_{-i}^k}$$

Dropping the superscript, and substituting $s_{-i} = p_i - s_i$, where p_i is the price of the current resource in the itinerary, we have:

$$\begin{aligned} \alpha_i s_i &= \beta_i(p_i - s_i) + \gamma_i\sqrt{p_i - s_i} \\ (\alpha_i + \beta_i)s_i - \beta_i p_i &= \gamma_i\sqrt{p_i - s_i}. \end{aligned}$$

We note that since the RHS of the previous equation is always positive, we require $s_i > p_i\beta/(\alpha+\beta)$, for a solution to exit. Making the substitution $s_i = p_i x_i$, we can solve for x_i to obtain the following demand function

$$x_i = \frac{\beta_i}{\alpha_i + \beta_i} + \frac{-\gamma_i^2 + \sqrt{\gamma_i^4 + 4\gamma_i^2(\alpha_i + \beta_i)\alpha_i p_i}}{2(\alpha_i + \beta_i)^2 p_i},$$

and solve for p_i to obtain the following price function

$$p_i = \frac{\gamma_i^2(1 - x_i)}{\left[(\alpha_i + \beta_i)x_i - \beta)\right]^2} = 0.$$

We note that $p_i(x_i)/(1 - x_i)$ is decreasing on $(\beta/(\alpha + \beta), 1)$. Thus, defining the valuation as in Equation (3), we have

$$v_i(x_i) = \left(\frac{\gamma_i^2}{\alpha_i + \beta_i}\right)\left(\frac{1}{\alpha_i} - \frac{1}{(\alpha_i + \beta_i)x_i - \beta_i}\right) \quad \forall\, x \in \left(\frac{\beta_i}{\alpha_i + \beta_i}, 1\right).$$

Sample plots of valuation, demand and price functions can be seen in Figure 1(b). We see that the equivalent valuation is a concave increasing function on a subinterval of the allocation space and meets the conditions of Assumption 1 on this subinterval. If we set the valuation to be $-\infty$ on $(0, \beta/(\alpha + \beta))$, the resulting optimal response would match the optimal response of the optimization problem stated in the beginning of this section. We see that the demand function does not go to zero as the price increases and the price function increases to infinite above an allocation of zero. The reason for this is that the agent is effectively inelastic with respect to allocation close to the minimal allocation requirement, and even exorbitant prices will not deter the agent. This is due to the lack of a constraint on the expenditure accrued. Nevertheless, this case of an inelastic agent can also be modeled with a demand function, price function and an equivalent instantaneous valuation.

9 Nash Equilibrium

We now address the question of whether there is an allocation of the resource at a cost where all agents participating in the auction are satisfied. In the language

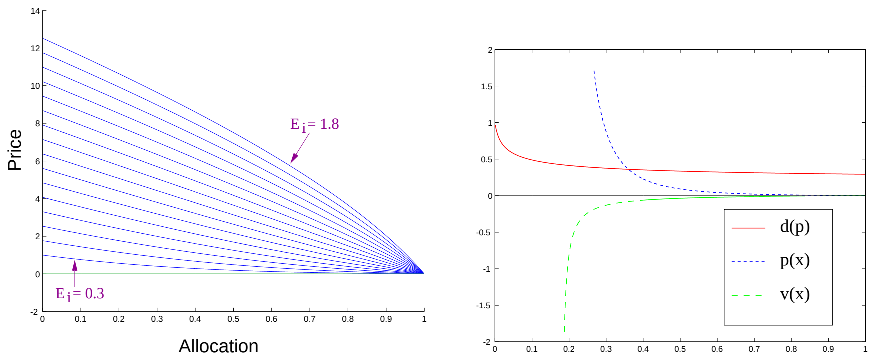

(a) (b)

Fig. 1. (a) Price Functions for Finite Budget Agents with Varying Endowments. (b) Price, Demand and Valuation Functions for Finite Time Agent.

of game theory, we ask whether there is a set of bids $\{s_i^*\}_{i=1}^N$, where N is the number of agents competing for the desired resource such that no single agent wishes to deviate from its bids given that the other agents' bids remain the same. This state, a Nash equilibrium, occurs if no agent can improve its utility by changing its bid under current market conditions, i.e.,

$$s_i^* = \arg\max_{s_i} U_i(s_i; s_{-i}^*) \quad \forall\ i\ \in \{1, \cdots, N\}$$

where s_{-i}^* implies $s_j = s_j^*$, $\forall\ j \neq i$. Because every agent's optimal response is captured in its price and demand function, we can use these as tools to evaluate the existence of a Nash equilibrium.

We find it useful to work in the space of demand functions. Due to the structure of our auction, the total of the allocated resources will always be one. If the optimal demand total at a particular price is greater or less than one, the allocation made to at least one agent will be unsatisfactory. Thus, it is equivalent to ask whether there is a price (or bid total) where the total demand of all the agents at that price is equal to one. Valid demand functions for elastic agents are assumed to be decreasing functions of price that go to zero as the price tends to infinite.

Proposition 3. *Given demand functions* $\{d_i(\theta)\}_{i=1}^N$*, where* $\sum_{i=1}^N d_i(0) > 1$, $lim_{\theta\to\infty} d_i(\theta) = 0, \forall\ i$*, and* $d_i(\theta_1) > d_i(\theta_2)\ \forall \theta_1, \theta_2$ *such that* $p_1 < p_2$ *for* $i = 1, \ldots, N$*, there exists a unique value* θ^* *such that* $\sum_{i=1}^N d_i(\theta^*) = 1$.

Proof. Let $\bar{d}(\theta) = \sum_{i=1}^N d_i(\theta)$. Then $\bar{d}(\theta)$ is a continuously decreasing function whose maximum is $\bar{d}(0) > 1$. We also have $\lim_{\theta\to\infty} \bar{d}(\theta) = 0$ which implies that for some $\bar{\theta}$ sufficiently large, $\bar{d}(\bar{\theta}) < 1$. Applying the Intermediate Value Theorem for $\bar{d}(\theta)$ on $[0, \bar{\theta}]$, we know that there exists at least one θ^* such that

$\bar{d}(\theta^*) = \sum_{i=1}^{N} d_i(\theta^*) = 1$. Let us assume that there are at least two values of θ where $\bar{d}(\theta) = 1$. Let us choose two of these values as θ_1^* and θ_2^*, where $\theta_1^* < \theta_2^*$. Then, we have $d_i(\theta_1^*) > d_i(\theta_2^*)$ $\forall i = 1, \ldots, N$, which implies that $\bar{d}(\theta_1^*) > \bar{d}(\theta_2^*)$. But we have $\bar{d}(\theta_1^*) = \bar{d}(\theta_2^*) = 1$, which is a contradiction and thus we can have only one θ where $\bar{d}(\theta) = \sum_{i=1}^{K} d_i(\theta) = 1$. ■

A graphical representation of the Nash equilibrium can be seen in Figure 2(a). By working in the space of demand functions, we can use the property that the demands are decreasing to easily see that there is a unique Nash equilibrium. Uniqueness of the Nash equilibrium is significant as we have a single stable operating point. Thus, given any set of agents there is a unique set of bids that yield an allocation where each agent is satisfied. The bids can be characterized in terms of the demand functions and Nash equilibrium price, θ^*, as $\{s_i : s_i = d_i(\theta^*)\ \theta^*\}_{i=1}^{N}$. The condition $\sum_{i=1}^{N} d_i(0) > 1$ is satisfied as $p(1) = 0 \ \Rightarrow d(0) = 1$ for each agent, unless the marginal valuation at one is infinite which will not occur for any reasonable valuation. We also requires that $N > 2$, and this is always satisfied as we have the bids of the resource and at least one agent requesting service. The price, θ^* determines which agents receive service as any agent with $d(\theta^*) = 0$, will have a zero bid as its optimal response.

10 Decentralized Bidding Algorithm

Knowing that there is a unique Nash equilibrium, the natural question that follows is how to arrive at that allocation. If the demand functions of all the agents, $\{d_i(\theta)\}_{i=1}^{N}$, were communicated to the resource, it could calculate the equilibrium allocation by a binary search over θ and enforce it immediately. However, this would add a significant signaling load. Also, if the resource is operating as a profit maker as opposed to a mediator, agents would not want to reveal their private information. Thus, it would be desirable if the agents could reach the Nash equilibrium allocation in a decentralized manner.

We assume that each agent is aware of the share of the resource that it currently receives. Also, the resource can provide the current price (or equivalently, the total of all bids) for the resource. If, at time slot n, the i-th agent bids s_i^n, then the price would be $\sum_{i=1}^{N} s_i^n$ and the i-th agent would receive $s_i^n / \sum_{i=1}^{N} s_i^n$ of the resource. This feedback from the resource prevents the possibility of the lying auctioneer that exists in second price auctions. Any agent can verify the price being announced by the resource as being valid by comparing it to the ratio of its bid to its allocation, which are both known to the agent. To obtain a viable decentralized algorithm, We seek a set of update policies $\{f_i\}_{i=1}^{N}$ such that if $s_i^{n+1} = f_i(s^n)$, where $s^n = [s_1^n\ s_2^n\ \cdots\ s_N^n]$, then $\lim_{n\to\infty} s_i^n = s_i^* = d_i(\theta^*)\ \theta^*,\ i = 1, \ldots, N$, where θ^* is a Nash equilibrium price. After bids are made by all the agents, the i-th agent will receive a feedback pair $(\theta, x_i) = \left(\sum_{i=1}^{N} s_i, s_i / \sum_{i=1}^{K} s_i\right)$ which denotes the congestion for that current time slot and the service rate received. The agent knows that if this pair does not lie on the curve $(\theta, d_i(\theta))$ or equivalently $(p_i(x_i), x_i)$, the current price-allocation pair is not optimal. Thus, any viable update algorithm must project

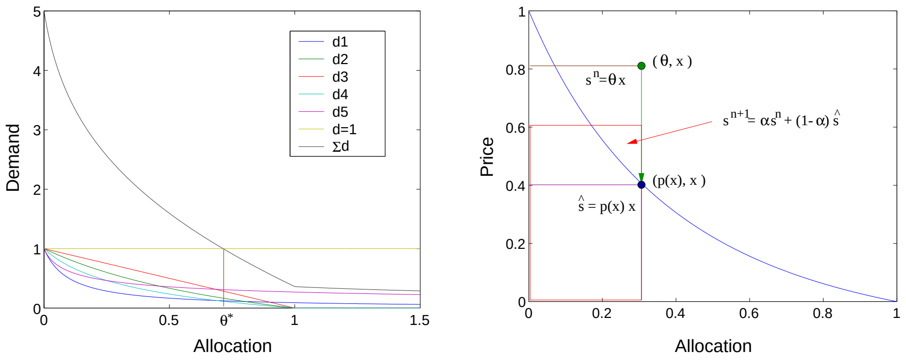

(a) Nash Equilibrium

(b) Relaxed Update

Fig. 2. Agent Equilibrium and Decentralized Allocation Mechanism

the feedback pair to a point on the demand or price curve. If a user receives allocation x_i for some bid, it would project to the point $(p_i(x_i), x_i)$ on the price curve and the corresponding bid would be $s = p_i(x_i)\ x_i$.

We propose an update algorithm where each agent projects the feedback point vertically onto the price function. This method has the advantage that the agent does not even require the feedback of the current price of the resource. Each agent projects its allocation to the price it would desire for the current allocation and makes the appropriate bid. This further reduces the signaling load required by the auction. It also eliminates the need to worry about the truthfulness of the auctioneer. Here, we present a relaxed version of the update scheme:

$$s_i^{n+1} = \alpha_i (s_i^n/\bar{s}^n) p_i(s_i^n/\bar{s}^n) + (1 - \alpha_i) s_i^n \tag{4}$$

where $\alpha_i \in (0, 1]$, which also covers the unrelaxed case $(\alpha_i = 1)$. A graphical interpretation of the relaxed update scheme can be seen in Figure 2(b). This update scheme depends on knowing only the received allocation, $s_i^n/\bar{s}_i^n$ and the previous bid s_i^n. The relaxed version of the update scheme requires no additional signaling and only requires that each agent store its last bid in memory. As α_i approaches zero, the time to convergence will delay as bids change more slowly. Thus, we desire to find the largest α_i that the i-th agent should use that will make the algorithm stable. Let us define $q_i := p_i(x_i^*) + x_i^* p_i'(x_i^*)/\theta^*$, where x_i^* is the equilibrium allocation for the i-th agent, $p_i'(\cdot)$ is the derivative of the price function $p_i(\cdot)$, and θ^* is the equilibrium bid total. We assume that all agents are restricted to those whose price or demand functions ensure that the agents are neither infinitely sensitive nor completely insensitive to the price at equilibrium. This will be satisfied if $q_i \neq -\infty$ and $q_i \neq 1$. We can show that if α_i is chosen such that $\alpha_i < 2/(1 - q_i)$, $\forall i$, then then update scheme is locally stable. Simulation have shown that the algorithm is globally stable under the same conditions.

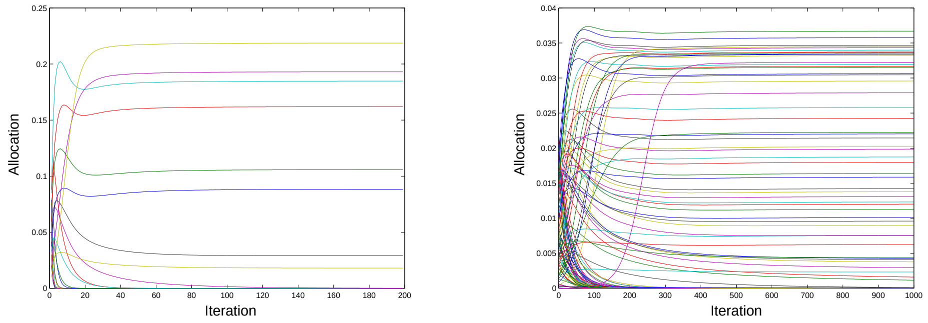

(a) 20 Agents

(b) 100 Agents

Fig. 3. Sample Evolution of Bids with Relaxed Update

Sample bid evolutions are shown in Figures 3(a) and 3(b). We can interpret q_i as an indicator of the price sensitivity of the i-th agent at equilibrium. The unrelaxed algorithm will converge if $|q_i| < 1 \ \forall \ i$. The relaxation softens the effects of agents with high price sensitivities, making the agent less aggressive in its bidding changes and effectively dampens the price sensitivity of the agent.

11 Conclusion

We have analyzed an auction for divisible resources that has a low cost of signaling and computation, enables agents to verify the auctioneer's prices and maintains proportionally fair allocation. We have shown that characterizing optimal responses as price functions allows us to capture valuations for wide classes of agent problems. We further show that the auction mechanism has a unique Nash equilibrium and develop decentralized algorithms that converge to the equilibrium without needing resource feedback or sharing private information.

In agent economies for computational and network resources, divisible auctions are appropriate mechanisms for allocation. This is a field with many areas open for investigation. We have studied the effects of coalition formation and extended the auction to a multiple resource setting. The seller revenue problem and generalization of divisible auction mechanisms are some of the issues to be addressed in the future.

References

1. E. Altman, T. Başar, and R. Srikant. Nash equilibria for combined flow control and routing in networks: asymptotic behavior for a large number of users. *IEEE Transactions on Automatic Control*, 2002 (to appear).

2. T. Başar and G. J. Olsder. *Dynamic Noncooperative Game Theory.* Academic Press, London, second edition, 1995.
3. J. Bredin, R. T. Maheswaran, O. Ç. İmer, T. Başar, D. Kotz, and D. Rus. A game-theoretic formulation of multi-agent resource allocation. In *Proceedings of the 4th International Conference on Autonomous Agents*, pages 349–356, Barcelona, Spain, 2000.
4. A. Chavez and P. Maes. Kasbah: an agent marketplace for buying and selling goods. In *Proceedings of the 1st International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology*, pages 75–90, London, April 1996.
5. S. de Vries and R. Vohra. Combinatorial auctions: a survey. *INFORMS Journal of Computing*, (to appear).
6. F. Kelly, A. Maulloo, and D. Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operations Research Society*, 49(3):237–252, March 1998.
7. J. O. Kephart, J. E. Hanson, and A. R. Greenwald. Dynamic pricing by software agents. *Computer Networks*, 32(6):731–752, May 2000.
8. U. Maheshwari. Charge-based proportional scheduling. Technical Memorandum MIT/LCS/TM-529, MIT Laboratory for Computer Science, January 1995.
9. R. T. Maheswaran and T. Başar. Multi-user flow control as a Nash game: performance of various algorithms. In *Proceedings of the 37th IEEE Conference on Decision and Control*, pages 1090–1095, Tampa, FL, December 16-18 1998.
10. R. T. Maheswaran and T. Başar. Decentralized network resource allocation as a repeated noncooperative market game. In *Proceedings of the 40th IEEE Conference on Decision and Control*, pages 4565–4570, Orlando, FL, December 4-7 2001.
11. O. Regev and N. Nisan. The popcorn market: An online market for computational resources. In *Proceedings of the 1st International Conference on Information and Computational Economies*, pages 148–157, Charleston, SC, October 1998.
12. T. Sandholm. An algorithm for optimal winner determination in combinatorial auctions. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, pages 542–547, Stockholm, July 1999.
13. T. W. Sandholm. Limitations of the Vickrey auction in computational multiagent systems. In *Proceedings of the 2nd International conference on Multi-Agent Systems*, pages 299–306, Kyoto, Japan, December 1996.
14. N. Semret. *Market Mechanisms for Network Resource Sharing.* PhD thesis, Columbia University, 1999.
15. S. J. Shenker. Making greed work in networks: a game-theoretic analysis of switch service disciplines. *IEEE/ACM Transactions on Networking*, 3(6):819–831, December 1995.
16. I. Stoica, H. Abdel-Wahab, K. Jeffay, S. K. Baruah, J. E. Gehrke, and C. G. Plaxton. A proportional share resource allocation algorithm for real-time time-shared systems. In *Proceedings of the 17th IEEE Real-Time Systems Symposium*, pages 288–299, Washington, D.C., December 1996.
17. C. A. Waldspurger, T. Hogg, B. A. Huberman, J. O. Kephart, and W. S. Stornetta. Spawn: a distributed computational economy. *IEEE Transactions on Software Engineering*, 18(2):103–117, February 1992.
18. M. Wellman, W. Walsh, P. Wurman, and J. Mackie-Mason. Auction protocols for decentralized scheduling. *Games and Economic Behavior*, 35:271–303, 2001.