Decentralized Network Resource Allocation As A Repeated Noncooperative Market Game [1]

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

Market-based methods are an emerging paradigm for controlling large decentralized systems. We introduce in this paper a bidding mechanism for allocation of network resources among competing agents, and study it from a game-theoretic perspective. We prove the existence and the uniqueness of Nash equilibrium and present an update algorithm that allows users to converge to the Nash equilibrium in a decentralized manner using feedback of only the common information available to the resource. The necessary conditions for local stability of relaxed versions of the algorithm are derived and verified by simulations.

1 Introduction

The rapid expansion of the Internet can be attributed in part to the fact that the decisions regarding transmission rates are made at the source. This allows nodes to be added seamlessly without having to reconfigure the entire network. The detractor of this method is the management of resources by users, such as congestion control, evidenced by ever expanding delays. Similarly, mobile or transportable software agents allow users to port their code, accessing various network devices and processors in an autonomous manner [5]. However, there remains the concern that autonomous code may overrun open resources if agents are not properly managed.

Market-based control has been evolving as a method to incorporate control into systems with large autonomous agents. Economic paradigms allow for decentralized implementation while providing mechanisms to regulate the behaviors of users. In addition, markets enable systems to scale rather easily once the initial dynamics of the exchange between users and resources are established. There have already been several applications of market and auction mechanisms to computing resources, operating system memory allocation, factory scheduling and manufacturing systems [3]. There have also been extensive applications and support for market-based control in communication networks [2], [6], [10], [11], [12]. Many of these implementations advocate usage based pricing, and associate agents with individual utility functions. In this paper, we adopt this framework to create a stable bargaining structure where users can manage distributed resources in a decentralized manner.

The performance each user receives is often affected by the actions of other users attempting to access the same network resources. The level of service all the users obtain will be a result of the equilibrium reached after users negotiate among themselves to satisfy their utilities. The nature of this negotiation and the attempt to find a stable operating point between competing agents calls for game-theoretic analysis. Modeling communication network problems as dynamic games has produced Nash equilibria solutions in many settings such as capacity allocation in routing [7], congestion control in product form networks [8], flow control in Markovian queueing networks [9], and rate-based flow control [1]. To reach Nash equilibria in decentralized settings, necessitated by the fact that users do not have access to other agents' bidding strategies and utilities, iterative algorithms based on network feedback that converge to a stable operating point are often necessary. Game-theoretic perspectives have resulted in existence of Nash equilibria in multiclass traffic environments, and the convergence conditions of various algorithms have been investigated [1],[4].

In this paper, we attempt to integrate market-based modeling and game theory to introduce a new market mechanism to manage network resources. We show that a unique Nash equilibrium with certain profit maximization properties exist, and investigate decentralized methods to reach this equilibrium. The paper is organized as follows. In Section 2, we describe the basic model of the system including the market mechanism for distributing resources, and the behavior and util-

[1] Research supported in part by grants AFOSR MURI AF DC 5-36128, and ARMY OSP 35352-6086.

ity of our users. In Section 3, we prove the existence and uniqueness of a Nash equilibrium with respect to users' demand. In Section 4, we introduce a decentralized update algorithm to reach the Nash equilibrium. In Section 5, necessary conditions for the local stability of relaxed versions of the scheme are derived. In Section 6, we present the results of simulations of the algorithm that verify its convergence. Finally in Section 7, we make some concluding remarks and suggestions for future work.

2 System Model

In our model, a resource determines a capacity, $C(t)$, of its service that it is willing to open up to users at time t. This capacity can be a percentage of CPU cycles from a processor offered to software agents or a portion of bandwidth in a link allotted to elastic traffic. The competition in this market is among the users for the good being offered by the resource. We assume that there are $K(t)$ users competing for the offered resource at time t, where $K(t) \in Z^+$. The resource is allocated as follows. Each user bids an amount that he is willing to pay per unit time for a portion of the available capacity, i.e. the i-th user bids $u_i(t)$. After all users' bids are in, the i-th user will receive service at rate

$$v_i(t) = \left(u_i(t) / \sum_{k=1}^{K(t)} u_k(t) \right) C(t). \qquad (1)$$

This allocation is fair in the sense that each user pays the same amount per unit resource at a given instant in time. This allocation is also the *proportionally fair* allocation for a single resource [6]. The price at time t is $\sum_{k=1}^{K(t)} u_k(t)/C(t)$ dollars per unit of resource. This price serves as an indicator of the demand for this resource and can be used as feedback as a congestion indicator to users. This allocation mechanism induces a non-cooperative game between the $K(t)$ users attempting to partition the offered capacity of this resource. To ensure that the system is not raided by "scavenger" agents, the resource advertises that it will not provide service unless a minimum total bid of u_0 is supplied by all the agents at equilibrium. We assume $u_0 > 0$, is sufficiently small such that any equilibrium reached by competition between legitimate agents will exceed this value.

For simplicity, we will assume that time is partitioned into slots of unit length. We will first analyze the mechanism within a single time slot, and thus will drop the reference to time in our variables.

The i-th user has an associated utility which is a function of the cost or bid for that time slot and the service received. The service received is a function of the sum of the bids of all the other users attempting to gain access to that resource. The user's utility will then be a function of the total amount bid in that time slot by other users, denoted θ_{-i}. The i-th user can then optimize its utility over its bid u_i for any given value of θ_{-i}. This bid function can be mapped to an equivalent function whose independent variable is $\theta = \sum_{k=1}^{K} u_k$, the sum of all bids for the given time slot, by substituting $\theta_{-i} = \theta - u_i$ and solving for the u_i. Given that the sum of all bids at a given slot is θ and the optimal bid for the i-th user is $u_i(\theta)$, our resource will allocate $u_i(\theta)/\theta$ of the offered capacity to the i-th user. Because $u_i(\theta)$ was obtained by optimizing the user's utility, the quantity $d_i(\theta) := u_i(\theta)/\theta$ represents the optimal demand as a percentage of the resource available for the i-th user given that the sum of bids in that time slot is θ or equivalently, the price per unit capacity is θ/C. Thus, we know that if the allocation in the current time slot yields a point that lies on the curve $d_i(\theta)$, it will have optimized the user's associated given utility.

We assume that the functions $d_i(\theta) : R^+ \to [0, 1]$, $i = 1, \ldots, K$, have the following properties:

$$\begin{aligned} d_i(0) &= 1 \\ d_i(\theta) &= 0, \quad \forall \theta > \bar{\theta}_i \\ d_i(\theta_1) &> d_i(\theta_2) \quad \forall \theta_1 < \theta_2 < \bar{\theta}_i \\ d_i(\theta) & \text{ is continuous on } [0, \infty) \end{aligned}$$

The first property states that a user will want all of the available resource if the price is zero or equivalently if nobody else is bidding for it. Similarly, the second property states that there is some price beyond which the user will not desire any of the resource. The third property captures the fact that in the price region where a user will desire to purchase service, its optimal demand will decrease as the price increases. Finally the fourth property intuitively states that an infinitesimal change in price can cause only an infinitesimal change in the optimal demand. Even if the last property does not hold for some users, their optimal demands can be approximated with arbitrary precision if they have finite discontinuities. Sample optimal demand functions are shown in Figure 1.

If a user has linear benefit with respect to its allocation, and has the utility function, $U_i = \gamma_i v_i - u_i$, the resulting demand function is $d(\theta) = max\{0, 1 - \theta/\gamma_i\}$ which satisfies the conditions stated above. Another scenario where agents minimize the total time taken to complete a sequence of tasks leads to a demand function that also satisfies the aforementioned conditions on $d(\theta)$ [2].

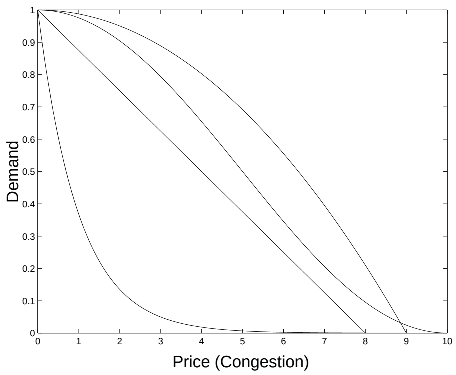

Figure 1: Sample Demand Functions

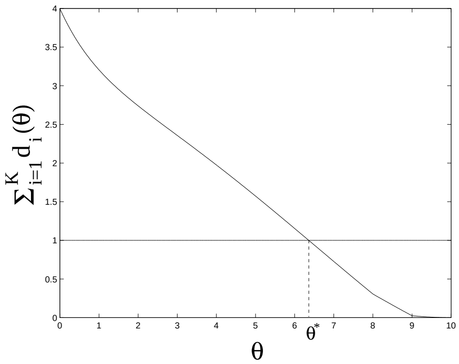

Figure 2: Depiction of Nash Equilibrium Value of $\theta = \theta^*$

3 Nash Equilibrium Allocation

We have introduced in the previous section a model where K users enter a bidding game to obtain a portion of an offered resource. If the resulting allocation does not lie on the user's optimal demand curve, the user will update its bid. An immediate question is whether there exists a set of bids $\{u_i\}_{i=1}^K$ that is a Nash equilibrium, i.e., a set of bids such that no single user wishes to deviate from its bid given that the bids of all the other users remain the same. To answer this question, we note that this is equivalent to asking whether there exists a value for the sum of total bids, θ, such that $\sum_{i=1}^K d_i(\theta) = 1$. This equivalence is valid because all of the offered capacity is partitioned among the users proportionally to their bids and the optimal demand function represents a percentage of the offered capacity.

Theorem 1 *Given any set of continuous functions $\{d_i(\theta)\}_{i=1}^K$ where $d_i(0) = 1$, $d_i(\theta) = 0\ \forall\theta > \bar{\theta}_i$, and $d_i(\theta_1) > d_i(\theta_2)\ \forall\theta_1, \theta_2$ such that $\theta_1 < \theta_2 < \bar{\theta}_i$ for $i = 1, \ldots, K$, there exists a unique value θ^* such that $\sum_{i=1}^K d_i(\theta^*) = 1$.*

Proof. Let $\bar{d}(\theta) = \sum_{i=1}^K d_i(\theta)$. Then $\bar{d}(\theta)$ is a continuously decreasing function whose maximum is $\bar{d}(0) = K > 1$. Let $\bar{\theta}_{\max} = \max_i \bar{\theta}_i$. We have $\bar{d}(\bar{\theta}_{\max}) = 0$. Applying the Intermediate Value Theorem for $\bar{d}(\theta)$ on $[0, \bar{\theta}_{\max}]$, we know that there exists at least one θ^* such that $\bar{d}(\theta^*) = \sum_{i=1}^K d_i(\theta^*) = 1$. 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 $d_i(\theta_1^*) > d_i(\theta_2^*)\ \forall i = 1, \ldots, K$, 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 depiction of the Nash equilibrium value of θ can be seen in Figure 2 where there are four users with optimal demand functions shown in Figure 1. The resulting set of bids that would form a Nash equilibrium allocation would be $\{u_i : u_i = d_i(\theta^*)\theta^*\}_{i=1}^K$. Given any set of competing agents with $K \geq 2$, we will have $\theta^* > 0$, thus if the i-th user is to receive any service from the resource, it is necessary that $d_i(\theta^*) > 0$. Thus, the value of θ^* will determine which users will get service in the equilibrium allocation.

4 A Decentralized Bidding Algorithm

If we had a way to communicate $\{d_i(\theta)\}_{i=1}^K$ to the resource, it could calculate the equilibrium allocation and enforce it immediately. This could be achieved by either the profit maximization procedure outlined in the previous section or a binary search algorithm. It would be desirable, however, if the users could reach the Nash equilibrium allocation in a decentralized manner. If, at time slot 0, each user bids $u_i^0,\ i = 1, \ldots, K$, then the total demand would be $\sum_{i=1}^K u_i^0$ and the i-th user would receive $u_i^0 / \sum_{i=1}^K u_i^0$ fraction of the resource. If the available common information, the sum of all bids and the number of users, is given as feedback to the users, then we seek a set of update policies $\{f_i\}_{i=1}^K$ such that if $u_i^{n+1} = f_i(u^n)$, where $u^n = [u_1^n u_2^n \cdots u_K^n]$, then $\lim_{n\to\infty} u_i^n = u_i^* = d_i(\theta^*)\theta^*,\ i = 1, \ldots, K$.

For the bid that the i-th user makes, there will be a pair

$$(x, y_i) = \left(\sum_{i=1}^K u_i, \frac{u_i}{\sum_{i=1}^K u_i} \right) \quad (2)$$

which denotes the congestion for that current time slot and the service rate received by the i-th user. We know that if this pair does not lie on the curve $(\theta, d_i(\theta))$, the current allocation cannot be a Nash equilibrium allocation. Thus, we propose an update algorithm where the user projects the current allocation pair to a point on its optimal demand curve as demonstrated in Figure 3.

Thus, given a bid u_i and feedback $\theta = \sum_{i=1}^K u_i$, the user can calculate (x, y_i) and project it on the demand

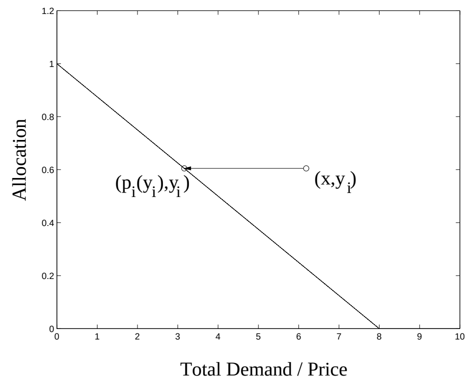

Figure 3: Horizontal Projection Algorithm

curve to get $(p_i(y_i), y_i)$ where $p_i(y_i)$ represents the level of demand or equivalently, the price at which the i-th user would want y_i as its allocation of the resource. The function $p_i(y_i) = d_i^{-1}(y_i)$ is a strictly decreasing function where $p_i(0) = \bar{\theta}_i$ and pair $p_i(1) = 0$. To achieve the price and allocation $(p_i(y_i), y_i)$, a user must bid $u_i = y_i p_i(y_i)$. This leads to the following decentralized update scheme:

$$u_i^{n+1} = (u_i^n/\bar{u}^n)p_i\left((u_i^n/\bar{u}^n)\right) \tag{3}$$

where n and $n+1$ denote the two consecutive iteration stages and $\bar{u}^n := \sum_{i=1}^{K} u_i^n$. We define:

$$q_i := (p_i(y_i^*) + y_i^* p_i'(y_i^*))/x^* \tag{4}$$

where y_i^* is the equilibrium allocation for the i-th user, $p_i'(\cdot)$ is the derivative of $p_i(\cdot)$, and x^* is the equilibrium bid total. Since $p_i(y_i^*) \leq x^*$, $q_i \leq 1$ for decreasing demand functions. We can show that the basic update scheme as described by (3) converges locally whenever $|q_i| < 1\ \forall i$. Because $p_i(\cdot)$ and $x^*$are finite, q_i can be interpreted as an indicator of the price sensitivity of equilibrium allocation, i.e. $p_i'(y_i^*)$. Thus, the basic algorithm will converge unless the user's sensitivity to his allocation is higher than a given value. We next investigate if we can extend local stability to a larger class of demand functions by relaxing the basic algorithm.

5 Local Stability of Relaxed Algorithm

We consider the following relaxed version of the basic algorithm:

$$u_i^{n+1} = \alpha_i(u_i^n/\bar{u}^n)p_i\left((u_i^n/\bar{u}^n)\right) + (1-\alpha_i)u_i^n$$

where $\alpha_i \in (0, 1]$, which also covers the *unrelaxed* case ($\alpha_i = 1$). This update scheme is decentralized as it only depends on common feedback, $\bar{u}^n$, and local information, u_i^n and α_i. We assume that the demand curves are restricted to those that yield the following conditions for $i = 1, \ldots, K$:

$$q_i \neq 0$$
$$q_i \neq -\infty$$

which ensure that the users are neither infinitely sensitive nor completely insensitive to the price at equilibrium. To investigate local stability, we linearize the update algorithm around the equilibrium bids $\{u_i^*\}$. Then, if we define $\epsilon_i^n := u_i^n - u_i^*$, the linearized system is $\epsilon^{n+1} = \tilde{J}\epsilon^n$, $\tilde{J} = AJ + (I - A)$, where

$$J = (I - Y)Q, \qquad Y = y^T 1_K$$
$$y = [y_1^*\ y_2^* \cdots y_K^*], \qquad 1_K = [1\ 1 \cdots 1]$$
$$Q = \mathrm{diag}(q_1, q_2, \cdots, q_K)$$
$$A = \mathrm{diag}(\alpha_1, \alpha_2, \cdots, \alpha_K)$$

We now investigate the conditions under which $\tilde{J}$ is stable. Let λ be an eigenvalue of $\tilde{J}$, and x a corresponding eigenvector. We have

$$\tilde{J}x = (A(I - Y)Q + (I - A))x = \lambda x.$$

Multiplying from the left by Q, we have

$$(QA(I - Y)Q + Q(I - A))x = \lambda Qx.$$

Since I, Q and A are diagonal matrices, we have the following relations

$$QA = AQ, \quad Q(I - A) = (I - A)Q$$

and thus,

$$\left[AQ(I - Y) + (I - A)\right]\ Qx = \lambda Qx$$
$$\left[AQ - AQY + I - A\right]r = \lambda r$$

where $r = Qx$. Let $\bar{r} := 1_K r$. Then,

$$AQr = \begin{bmatrix} \alpha_1 q_1 r_1 \\ \vdots \\ \alpha_K q_K r_K \end{bmatrix}$$
$$AQYr = \begin{bmatrix} \alpha_1 q_1 y_1^* 1_K \\ \vdots \\ \alpha_K q_K y_K^* 1_K \end{bmatrix} r = \begin{bmatrix} \alpha_1 q_1 y_1^* \\ \vdots \\ \alpha_K q_K y_K^* \end{bmatrix} \bar{r}.$$

Thus, for every λ that is an eigenvalue of $\tilde{J}$, we have:

$$\alpha_i q_i r_i - \alpha_i q_i y_i^* \bar{r} + (1 - \alpha_i) r_i = \lambda r_i \quad \forall i$$

If we assume that the candidate λ has an eigenvector such that $\bar{r} = 0$, we have

$$\lambda = \alpha_i q_i + (1 - \alpha_i) =: \tilde{q}_i$$

Thus, we will have a stable system if we choose $\alpha_i \in (0, 1]$ such that

$$\alpha_i < 2/(1 - q_i) \tag{5}$$

If, on the other hand, the candidate eigenvalue λ has an eigenvector such that $\bar{r} \neq 0$, we have

$$(\alpha_i q_i + (1-\alpha_i) - \lambda) r_i = \alpha_i q_i y_i^* \bar{r}$$
$$\Rightarrow \frac{\alpha_i q_i y_i^*}{\alpha_i q_i + (1-\alpha_i) - \lambda} = \frac{r_i}{\bar{r}}$$
$$\Rightarrow \sum_{i=1}^{K} y_i^* \frac{\alpha_i q_i}{\alpha_i q_i + (1-\alpha_i) - \lambda} = 1$$
$$\Rightarrow \sum_{i=1}^{K} y_i^* \frac{\alpha_i q_i}{\tilde{q}_i - \lambda} = 1$$
$$\Rightarrow \sum_{i=1}^{K} y_i^* z_i = 1$$

where $z_i = (\alpha_i q_i)/(\tilde{q}_i - \lambda) =: z_i^R + j z_i^I$ is complex.

If we can choose α_i such that $z_i^R < 1$ for all i, for some candidate eigenvalue λ, then we know that λ cannot be an eigenvalue for $\tilde{J}$, because $\sum_{i=1}^{N} y_i^* = 1$. We now investigate how to choose α_i in a way such that $z_i^R < 1$ for all candidate eigenvalues on or outside the unit circle. If we can do this, we know the resulting system is locally stable as the only valid eigenvalues for $\tilde{J}$ must lie inside the unit circle. Let $\lambda = \sigma + j\omega$. Then,

$$z_i = \alpha_i q_i/(\tilde{q}_i - \lambda) = \alpha_i q_i/(\tilde{q}_i - (\sigma + j\omega))$$
$$= \alpha_i q_i/((\tilde{q}_i - \sigma) - j\omega)$$
$$= \alpha_i q_i((\tilde{q}_i - \sigma) + j\omega)/((\tilde{q}_i - \sigma)^2 + \omega^2)$$
$$\Rightarrow z_i^R = \alpha_i q_i(\tilde{q}_i - \sigma)/((\tilde{q}_i - \sigma)^2 + \omega^2)$$

We assume that α_i has been chosen to satisfy (5), implying $-1 < \tilde{q}_i < 1$.

Case 1. Let us assume that $q_i > 0$. Then, $\sigma \geq q_i \Rightarrow z_i^R \leq 0$ so we assume $\sigma < q_i$. If $|\sigma| \geq 1$, $q_i > 0$, then

$$z_i^R < \alpha_i q_i(q_i - \sigma)/(q_i - \sigma)^2 < \alpha_i q_i/(q_i - \sigma) < 1$$

$\forall \alpha_i \in (0,1)$ since $q_i < 1$, $\sigma < q_i$, $|\sigma| \geq 1$ implies $\sigma < 0$. If $|\sigma| < 1$, $|\lambda| \geq 1 \Rightarrow \omega^2 \geq 1 - \sigma^2$. Then,

$$\begin{aligned} z_i^R &\leq \alpha_i \frac{q_i(q_i - \sigma)}{(q_i - \sigma)^2 + (1 - \sigma^2)} \\ &= \alpha_i \frac{q_i(q_i - \sigma)}{q_i^2 + 2\sigma q_i + \sigma^2 + 1 - \sigma^2} \\ &= \alpha_i \frac{q_i(q_i - \sigma)}{q_i(q_i - \sigma) + (1 - \sigma q_i)} < 1 \end{aligned}$$

Case 2. Let us now consider, $q_i < 0$. If $\sigma \leq \tilde{q}_i$, then $z_i^R \leq 0$, so we assume $\sigma > \tilde{q}_i$. Consider $\sigma \geq 1 > \tilde{q}_i$. Then, we have

$$\begin{aligned} z_i^R &> \frac{\alpha_i q_i(\tilde{q}_i - \sigma)}{(\tilde{q}_i - \sigma)^2} = \frac{\alpha_i q_i}{(\tilde{q}_i - \sigma)} \\ &= \frac{\alpha_i q_i}{\alpha_i q_i + 1 - \sigma - \alpha_i} < 1. \end{aligned}$$

If $\tilde{q}_i < \sigma < 1, |\lambda| \geq 1 \Rightarrow \omega^2 \geq 1 - \sigma^2$. Then, we have

$$z_i^R \leq \alpha_i q_i \frac{\tilde{q}_i - \sigma}{(\tilde{q}_i - \sigma)^2 + (1 - \sigma^2)} =: \alpha_i q_i f(\sigma)$$

where

$$\frac{\partial f(\sigma)}{\partial \sigma} = \frac{\tilde{q}_i^2 - 1}{[(\tilde{q}_i - \sigma)^2 + (1 - \sigma)^2]^2} < 0$$

if $|\tilde{q}_i| < 1$, and $f(\tilde{q}_i) = 0$. Thus $f(\sigma)$ is minimized and hence the bound on z_i^R is maximized as $\sigma \to 1$. This implies

$$\begin{aligned} z_i^R &< \frac{\alpha_i q_i(\tilde{q}_i - 1)}{(\tilde{q}_i - 1)^2 + (1 - 1^2)} = \frac{\alpha_i q_i}{\tilde{q}_i - 1} \\ &= \frac{\alpha_i q_i}{\alpha_i q_i + 1 - \alpha_i - 1} = \frac{\alpha_i q_i}{\alpha_1(q_i - 1)} = \frac{q_i}{q_i - 1} < 1. \end{aligned}$$

Thus, we have shown that if $\alpha_i < 2/(1 - q_i)$, the resulting system will be locally stable for all values of $q_i \in (-\infty, 1)$. However, q_i is determined from the allocation at equilibrium which is not known *a priori* to the user. Each agent must take into account all possible equilibrium allocations when choosing its relaxation parameter. If the user minimizes the function

$$g(y) = (p_i(y) + y p_i'(y))/p_i(y)$$

over the domain $y \in (0, d_i(u_0))$ where $\hat{q}_i$ denotes the minimum, then $\hat{q}_i < q_i$ (unless $g(y) > 0 \;\; \forall y$ in which case relaxation is not necessary). Since $p_i(y)$ and $d_i(u_0)$ are known to each agent *a priori*, α_i can be chosen in a decentralized manner assuring local stability. This analysis can be applied to all decreasing demand functions that do not have an infinite or zero derivative (sensitivity) with respect to price. The only exception is the case where a user receives an allocation of zero and his upper limit is equal to the equilibrium price, i.e. $\bar{\theta} = p_i(0) = x^*$. This is a topic for further investigation.

6 Simulation Results

The algorithm was tested over the sample demand functions

$$\begin{aligned} d(x) &= \max\{0, 1 - x/\bar{\theta}\}, \\ d(x) &= \max\{0, 1 - (x/\bar{\theta}\}, \\ d(x) &= \max\{0, 1 - \sqrt{x/\bar{\theta}}\}, \end{aligned}$$

assigned with equal probability to each agent. For each test, the quantities $K \in \{2, \cdots, 100\}$ and $\bar{\theta} \in (0,1)$ for each agent were chosen uniformly from the given sets. Simulations were run for 200 iterations, and in all cases the allocations matched the Nash equilibrium allocations within 100th of a decimal place. A sample evolution for 100 agents is shown in Figure 4. The tests showed that relaxation was only necessary in scenarios where only two users received non-zero allocations. Sample evolutions displaying the effect of relaxation can be seen in Figure 5.

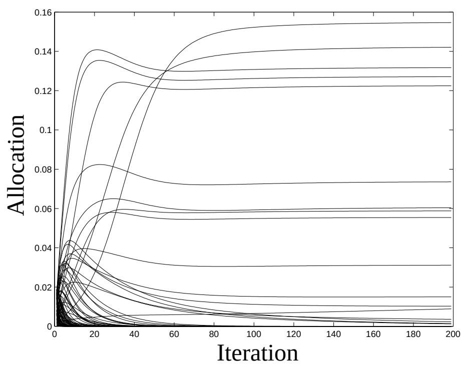

Figure 4: Sample Evolution of 100 Agents

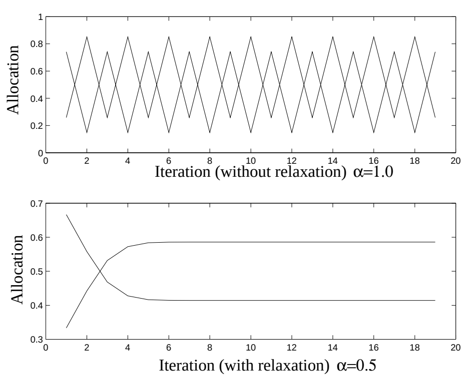

Figure 5: Sample Effect of Relaxation

7 Conclusions

We have introduced a market mechanism for allocating network resources in a proportionally fair manner where users pay the same price per unit of service. Users are characterized by decreasing optimal demand functions derived from an original associated utility function. A unique Nash equilibrium has been shown to exist for users with heterogeneous demand functions. We have proposed a decentralized update scheme that uses a projection onto the optimal demand curve. We investigated relaxed versions of the algorithm and derived the necessary conditions under which it is locally stable. Possibilities for future work include extensions to analysis of global stability, exploration of the relationships between utility and demand, and studying multiple resource competition.

References

[1] E. Altman and T. Başar. "Multi-User Rate-Based Flow Control," in *IEEE Transactions on Communications*, vol. 46, pp. 940-949, July 1998.

[2] J. Bredin, R. T. Maheswaran, Ç. Imer, T. Başar, D. Kotz and D. Rus, "A Game-Theoretic Formulation of Multi-Agent Resource Allocation," in *Proceedings of the Fourth International Conference on Autonomous Agents*, June 2000.

[3] S. H. Clearwater, *Market-Based Control: A Paradigm for Distributed Resource Allocation.* Singapore: World Scientific Publishing Co. Pte. Ltd., 1996.

[4] C. Douligeris and R. Mazumdar, "A Game Theoretic Perspective to Flow Control in Telecommunication Networks," in *Journal of the Franklin Institute*, vol. 329, pp. 383-402, 1992.

[5] R. S. Gray, "Agent Tcl: A transportable agent system," in *Proceedings of the CIKM Workshop on Intelligent Information Agents, Fourth International Conference on Information and Knowledge Management (CIKM 95)*, December 1995.

[6] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate Control for Communication Networks: Shadow Prices, Proportional Fairness and Stability," in *J. Oper. Res. Soc.*, vol. 48, pp.237-252, 1998.

[7] Y. A. Korilis and A. A. Lazar, "Architecting Noncooperative Networks," in *IEEE Journal on Selected Areas in Communication*, vol. 13, pp. 1241-1251, September 1995.

[8] Y. A. Korilis and A. A. Lazar, "On the Existence of Equilibria in Noncooperative Optimal Flow Control," in *Journal of the Association for Computing Machinery*, vol. 42, pp. 584-613, May 1995.

[9] M. T. Hsiao and A. A. Lazar, "Optimal Decentralized Flow Control of Markovian Queueing Networks with Multiple Controller," in *Performance Evaluation*, vol. 13, pp. 181-204, November 1991.

[10] J. K. Mackie-Mason and H. R. Varian, "Pricing Congestible Network Resources," in *IEEE Journal on Selected Areas on Communications*, vol. 13, pp. 1141-1148, September 1995.

[11] R. T. Maheswaran, Ç. Imer, and T. Başar, "Agent Mobility Under Price Incentives," in *Proceedings of the 38th IEEE Conference on Decision and Control*, December 1999.

[12] I. C. Paschalidis and J. N. Tsitsiklis, "Congestion-Dependent Pricing of Network Resources," in *IEEE/ACM Transactions on Networking*, vol. 8, pp. 171-184, April 2000.