

Distributed Network Flow Control Based on Dynamic Competitive Markets

Errin W. Fulp^{*†} and Douglas S. Reeves[†]

Abstract

Network applications require a certain level of network performance for their proper operation. These individual guarantees can be provided if sufficient amounts of network resources are available; however, contention for the limited network resources may occur. For this reason, networks use flow control to manage network resources fairly and efficiently. This paper presents a distributed microeconomic flow control technique, that models the network as competitive markets. In these markets switches price their link bandwidth based on supply and demand, and users purchase bandwidth so as to maximize their individual Quality of Service (QoS). This decentralized flow control method provides a Pareto optimal bandwidth distribution, individual QoS, and high utilization. Simulation results using actual MPEG-compressed video traffic show utilization over 95% and better QoS control than max-min.

1 Introduction

Current and future networks must accommodate a wide variety of network applications. These applications range from programs that transmit simple text to complex multimedia applications that require voice and video transmission. All of these applications need network resources, such as bandwidth and buffer space, to obtain a certain Quality of Service (QoS). QoS may include bounds on, the delay of packets, variation in the delay or packet loss probability. Consequently, contention may occur for the finite amount of resources. For this reason networks need a method of flow control to manage these limited resources in a fair and efficient manner.

There are two goals associated with flow control, fairness among applications and the balance between throughput and QoS [1] [2]. Defining fairness is difficult because of the variations in application characteristics and requirements. The balance between throughput and QoS is the concept that the network should seek high resource utilization, but not at the expense of poor QoS (and vice versa). Hence, due to heterogeneous networks, diverse resource requirements and the goals associated with flow control, proper flow control is a challenging problem. Several different

methods of flow control have been proposed. We briefly discuss the general classes of flow control, as well as a new type based on economics theory.

Preventive flow control determines the transmission rate of each source that will avoid congestion. In this case congestion is prevented and some service guarantees can be provided. However, this type of flow control may lead to over allocation of resources and are not well suited for the dynamic changes (such as variable bit rate sources) that may occur in the network. Feedback flow control methods alter data transmission to adapt to changing network conditions. Window flow control is one example used in packet networks. In this strategy, network feedback is used to limit the number of packets transmitted; however this type of flow control is not well suited for large networks because of propagation delays and few (if any) QoS guarantees can be made [1]. In ATM networks, several feedback traffic management strategies have been proposed for Available Bit Rate (ABR) service. These traffic management techniques use network feedback to alter the rate of a source (instead of the number of packets). Examples of explicit rate techniques include EPRCA and ERICA [3]. These strategies rely on the circulation of a Resource Management (RM) cell per connection [3]. As the RM-cell travels along the path, a switch and/or the destination may alter its contents. Exactly how this is done depends on the strategy. Once the cell reaches the destination it is returned to the source, who must alter transmission based on the RM-cell information. When a switch becomes congested, many of these traffic management strategies seek to allocate the bandwidth in a *fair* max-min manner [3]. The result is a division of link bandwidth equally among bottlenecked users. However, these methods do not take into account the fact that some sources may be able to reduce their transmission rate (for example compressed video) more easily than others. Therefore when congestion occurs, this socialistic allocation may not be the best when considering the individual QoS expected by each user.

An economic flow control method models the network as an economy, then applies microeconomic principles for resource allocation. A simple network economy consists of two types of agents: consumers (network applications) and producers (switches). Consumers require resources to satisfy their QoS. Producers own the resources sought by consumers, and seek to maximize their satisfaction by selling or renting their resources. Using this framework, microeconomics can be used to

define how network resources are allocated. In this paper, we apply microeconomics theory only to the task of flow control (i.e., we do not suggest its use for revenue generation or usage-based billing).

One approach of applying microeconomics to computer networks uses maximization techniques to maximize utility [4] [5] [6] [7] [8]. A utility function maps a resource amount to a satisfaction value. Using this function, one can compare the satisfaction levels of different resource amounts. The maximization process determines the optimal resource allocation such that the utility of a group of users is maximized subject to budget and resource availability constraints. Since the computation required for the maximization process increases as the number of users increases, these methods are not scalable to networks with a large number of users. To provide scalability, some approaches group users and use a single utility curve to represent the group. The maximization process is then performed for the smaller number of groups instead of individual users. Groups can be created based on desired QoS [4] [5] or on traffic types (or service classes) [6]. Accurately grouping users together may be problematic due to the wide variety of applications and their diverse resource requirements. Another problem is that these approaches generally require a centralized entity to determine the optimal allocation amount. This is undesirable because the economy relies on one entity, which is not reliable or fault tolerant.

Another microeconomic approach, congestion pricing, charges users for their consumption of resources and resources are priced to reflect supply and demand [9] [10] [11] [12] [13] [7] [14]. Alternatively, prices can be set with respect to marginal costs [15]. With such a model, prices can be set to encourage high utilization of network resources as well as a fair distribution. Users act independently, attempting to maximize their own utility and prices are set based on local resource conditions. It has been shown that pricing based on supply and demand results in higher utilization than traditional flat (single) pricing [10] [7]. Ferguson, et al. is an example of virtual circuit flow control based on pricing network resources [11] [12]. Prices of links in the system were iteratively adjusted until an equilibrium of supply and demand was reached. They were able to prove that the system achieved a Pareto equilibrium, as long as demands remained constant. Limitations of these methods of congestion pricing include reliance on a well-defined statistical model of source

traffic, and restrictions on the shape of the utility curve. The transient behavior, and the method of distributing intermediate prices and allocations during convergence, is generally ignored. The methods are not intended to adapt to changing traffic demands, nor have they been validated in a detailed way using realistic network configurations and real traffic.

Our approach uses congestion pricing in a competitive market to provide high quality-of-service and high resource utilization, at a modest implementation cost. Similar to several microeconomic flow control methods [9] [11] [12] [14] our approach is decentralized, seeks an equilibrium price and achieves a Pareto optimal distribution. Our approach has the following unique features:

1. More realistic (measured) utility curves are incorporated.
2. There are no restrictions on the statistical behavior of user traffic.
3. Heterogeneous link and switch bandwidths are supported.
4. Control of individual QoS, rather than aggregate QoS, is provided.
5. The number of users who receive satisfactory QoS is maximized.
6. A technique for distributing prices during the convergence period is presented. The transient behavior of our approach is illustrated on a realistic network with realistic traffic.

Experimental results also demonstrate our approach adapts well to changing traffic demands, and controls QoS better than a well-known method of flow-control in ATM networks, with equivalent utilization.

The remainder of this paper is structured as follows. Section 2 reviews the competitive market model. Section 3 describes the pricing technique in detail. Section 4 discusses how our pricing strategy achieves an equilibrium price and a fair Pareto optimal distribution. Section 5 discusses how the pricing policy contends with network dynamics such as, users entering/exiting and multimedia traffic. Section 6 describes the simulation results and comparison to max-min. Finally, section 7 reviews the pricing technique, summarizes the results and discusses some open questions.

2 Competitive Market Model

We will use a competitive market model for our network economy. The competitive market model consists of scarce resources and two types of agents, consumers and producers. A resource is an item (or service) which is valued by agents in the economy. Since it is scarce, there is never enough of the resource to satisfy all the agents all the time. For this reason, allocation decisions must be made. Consumers require resources to satisfy wants. Producers create or own the resources sought by consumers. These agents come together at a market, where they buy or sell resources. Usually these exchanges are intermediated with money and the exchange rate of a resource is called its price. Prices are set with respect to supply and demand. The price increases if the demand is greater than the supply and decreases when the demand is less than the supply. When they are equal, the market and price is in equilibrium. This moment is referred to as "clearing the market" and the resulting allocation is Pareto optimal [16]. Pareto optimality is the allocation of finite resources such that no sub-set of users can improve on their allocation without lowering the utility of another. This model was chosen for our computer network economy because of its ability to achieve certain desirable goals, such as Pareto optimal distribution and price stability. The competitive market also has a simple structure and a well founded mathematical basis for analysis. We again emphasize that our goal is flow control with QoS. Users are not billed, nor is there any element of cost recovery or profit generation.

3 A Proposed Pricing Policy

This proposed flow control method is based on a competitive market model, where pricing is done to promote high utilization and Pareto optimal distribution. As seen in figure 1 there are three entities in this network economy: users (those who execute network applications), Network Brokers (NB) and switches. Using the competitive market nomenclature, users are consumers, switches are producers and network brokers are used to assist the exchange of resources in the market. While there are many resources in a computer network, this paper focuses on the pricing of link bandwidth.

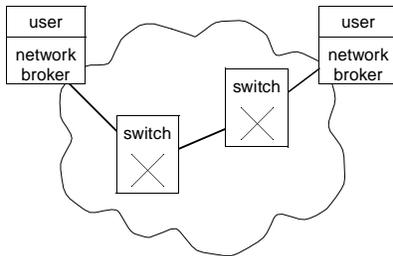


Figure 1: Example network consisting of users, network brokers and switches.

3.1 Switch

In our competitive market, the switch owns the link bandwidth that is sought by consumers. The network consists of several switches interconnected with links. For a unidirectional link between two switches, we consider the sending switch as owner of the bandwidth of that link. Each switch prices its link bandwidth based on local supply and demand for that link. Therefore a single switch, having multiple output links, will have one price associated with each output port. The entire network can be viewed as multiple competitive markets, one market per link (similar to the New York Stock Exchange). These markets operate independently and asynchronously since there is no need for market communication (for example, price comparisons) or synchronization from switch to switch. Consequently, this results in a decentralized economy, where the physical failure of one switch/link does not necessarily cause failure of the entire economy.

The price computation for link i is performed at the switch, at discrete intervals of time. We denote the n th calculation instant as t_n^i and the interval of time between the calculation points t_n^i and t_{n+1}^i as the n th price interval, P_n^i . The price during P_n^i is constant and is denoted as p_n^i . The demand for bandwidth at link i is measured as the total (aggregate) traffic received at its associated output port. During the n th price interval, P_n^i , the total demand is expected to change; even so, the calculation of p_{n+1}^i will only use the demand measured at the end of the interval. For this reason, let the demand for bandwidth at link i , at the end of the n th price interval, be denoted as d_n^i . The supply of bandwidth at link i is constant and denoted as S^i .

At the end of the price interval, P_n^i , the switch updates the price of link i using the following equation,

$$p_{n+1}^i = p_n^i + c \cdot \left(\frac{d_n^i - \alpha \cdot S^i}{\alpha \cdot S^i} \right) \quad (1)$$

The form of the price equation is referred to as a tâtonnement process and is used in a competitive market to set the price with respect to the current supply and demand [17]. In a tâtonnement process the new price is equal to the previous price plus a correction function. The correction function provides feedback based on the demand (received traffic) and the supply (bandwidth available). The bandwidth available is the total bandwidth times a constant α , where $0 < \alpha \leq 1$. This causes the price to increase after some percentage (α) of the total bandwidth has been reached. This is evident from the equation, since the price will only increase if the numerator is positive ($d_n^i > \alpha \cdot S^i$). The price will decrease as the demand decreases and will increase as the demand increases. An *equilibrium price* p_*^i is reached at link i when the supply equals the demand. At this point the market clears for link i and the allocation of bandwidth is Pareto optimal [16]. The positive constant c amplifies the feedback signal and its value ultimately controls how quickly the price will increase or decrease (speed of adjustment). Note that the equation can yield negative prices. We will assume that the price will not fall below a certain non-negative minimum price (set by the switch).

After the new price, p_{n+1}^i , is calculated, a new price quote is forwarded to each NB using this link. The price quote for link i , denoted as q_{n+1}^i , consists of; p_{n+1}^i , d_n^i , S^i , c and α . The NB will use all of the information in the price quote to determine the amount of bandwidth to purchase. The switch is only responsible for storing the current total (not individual, or even group) demand and price for each link, which requires a trivial amount of storage.

3.2 User

The user, executing a network application, requires bandwidth for transmission. The amount of bandwidth desired is determined from the application and is denoted as b_m . We assume b_m is constant for the duration of the application. In section 5 we will allow b_m to vary over time, which is desirable for multimedia transmission.

Based on prices and wealth, the user can afford a range of bandwidth (less than or equal to b_m),

and some amounts will be preferred over others. In economics these preferences are represented with a utility function. The utility function maps a resource amount to a real number, that corresponds to a satisfaction level. Assuming $U(\cdot)$ is a utility function, if the user prefers an amount x over y (this is represented using the notation $x \succ y$) then $U(x) > U(y)$. The utility curve can be used to compare resource amounts based on the satisfaction the user will receive. This provides an important link between resource amounts and user satisfaction. For this economy we will use *QoS profiles* [18] for the utility curves. Based on psycho-visual experiments, the QoS profile is a two dimensional graph, as seen in figure 2. The profile can be approximated by a piece-wise linear curve with three different slopes. The slope of each linear segment represents the rate at which the performance of the application degrades when the network allocates a percentage of the desired bandwidth (b_m). A steeper slope indicates the inability of the application to easily scale bandwidth (for example, high quality video), while a flatter slope signifies the application can more readily scale bandwidth requirements (for example, teleconferencing or data transmission). The horizontal axis measures the bandwidth ratio of allocated bandwidth to desired bandwidth (b_m). The vertical axis measures the satisfaction and is referred to as a QoS score. Our QoS scores range from one to five, with five representing an excellent perceived quality and one representing very poor quality. We will refer to an *acceptable* QoS score as any value greater than or equal to 3. As seen in the figure, if the allocated bandwidth is equal to the desired bandwidth (b_m), the ratio is one and the corresponding QoS score is 5 (excellent quality). As this ratio becomes smaller the QoS score reduces as well. Profiles can be created for a variety of applications and redefined as users gain more experience. New and updated profiles can be easily incorporated within the economy as they become available. More information about QoS profiles and the relationship between bit-rate and quality can be found in [18], [19], and [20].

Finally, the user is charged continuously for the duration of the session (analogous to a meter). To pay for the expenses, we will assume the user provides an equal amount of money over regular periods of time. We will refer to this as the budget rate of the user, W (\$/sec). A single initial endowment could have been used, but would necessitate defining how it is spent during the session. To simplify simulation and analysis, budget rates are used.

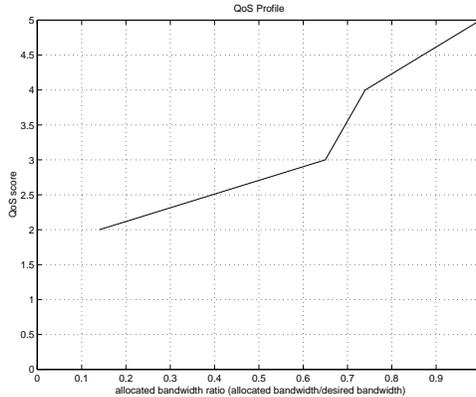


Figure 2: Example profile relating bandwidth allocation to perception of quality.

3.3 Network Broker

Users can only enter the network economy through a network broker (NB). This entity is an agent for the user and is located between the user and the edge of the network. The functions of the NB can be part of the protocol stack that executes on a host system, just as current protocol stacks provide flow control to individual users. Representing the user in the economy the NB performs the following tasks: connection admission control, policing, and purchase decisions. Although the NB works as an agent for the user (making purchasing decisions), we assume that the NB operates honestly in regards to both the switches and the user.

The NB controls network admission by initially requiring the user to have enough wealth to afford at least an *acceptable* QoS; otherwise, the user is denied access. The purpose of this requirement is to be certain all users are viable consumers in the market and to prevent overloading the economy. We believe the social welfare of the economy is better when it consists of fewer users each receiving a good QoS, instead of many users each receiving a poor QoS. Hence, we are attempting to maximize the number of users in the economy, where each user can afford an acceptable QoS. If the desired bandwidth is constant, then the test is relatively simple. However, for sources where the desired bandwidth will change over time, a more complex admission test is required.

The NB monitors the user and the prices by gathering and storing information about each, as shown in figure 3. From the user, the NB collects and stores; the QoS profile, b_m and W . The NB also stores the route, R , that connects source to the destination, where R consists of v links,

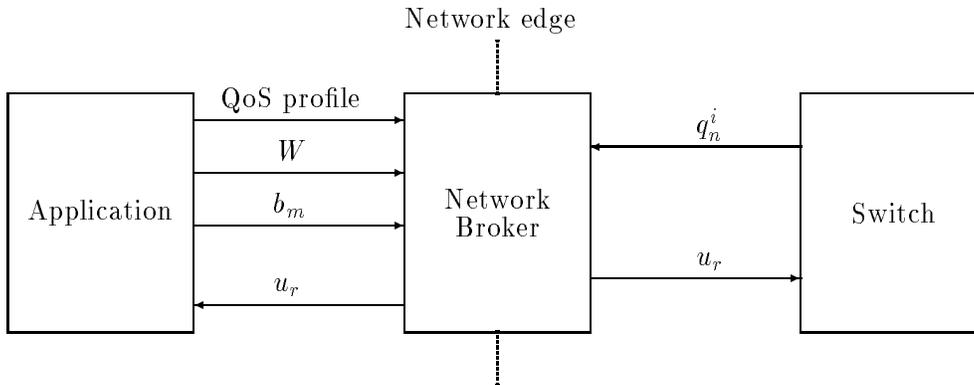


Figure 3: Information exchanged between the application, NB and switch.

$\{l^i, i = 1 \dots v\}$. For each link on R , a price quote, q^i , is collected, where $\vec{q} = \{q^i, i = 1 \dots v\}$ is the vector of price quotes for the route.¹ Price quotes will change over time, since they represent link supply and demand. The NB will only store the most recent price quote from each link in the route. The NB will divide the budget rate, W , into a vector of v budget rates \vec{w} , where $\vec{w} = \{w^i, i = 1 \dots v\}$ and w^i corresponds to link i . Separate budgets are used to localize the effect of prices to each link. This prevents spending the entire budget on one expensive link. Of course depositing and withdrawing to and from these individual budgets is possible and perhaps advantageous. Using this information the NB levies the user for their consumption. Users will be charged based on usage (similar to electricity), since bandwidth is a non-storable item. Using this information the NB polices the user, ensuring only the bandwidth purchased is used.

Finally, the NB determines the amount of bandwidth to purchase. This value is based on the budget, current prices and QoS profile of the user. Denote the r th amount of bandwidth to purchase (use) as, u_r . Once the NB determines u_r , the user will start sending at this rate immediately. There is no need for direct confirmation/feedback from the switches. A new amount of bandwidth to purchase, u_{r+1} , will be determined in response to a new price (or change in demand, as will be described in section 5). Exactly how the NB determines u_{r+1} is described next.

¹The requirement that the NB must know the entire route, and store a distinct price per link, can be relaxed. The NB can periodically circulate a RM cell or packet per connection, as in ATM flow control. This cell delivers demand information on the forward trip and collects price information on the return trip. Details and experimental results are omitted due to space restrictions.

3.3.1 Determining the Bandwidth to Use

When determining u_{r+1} , the NB will first calculate the maximum and minimum bandwidth that can be used. The maximum bandwidth that can be used at link i is,

$$b_{max}^i = \frac{w^i}{p^i}, i = 1 \dots v$$

therefore the maximum bandwidth the user can afford is,

$$\hat{b}_{max} = \min_{i=1 \dots v} \{b_{max}^i\}.$$

Note this equation maximizes the bandwidth at the current prices. The minimum bandwidth that can be used is determined from the QoS profile, b_m and the value that corresponds to the lowest acceptable QoS score. It is possible that $\hat{b}_{max} < b_{min}$ (the minimum is not affordable), due to the QoS constraint, prices and budgets. If this case arises, the user must either; increase the budget rate, accept a lower QoS, or drop the connection. Properly managing such a situation is an area for future work.

After \hat{b}_{max} and b_{min} have been calculated, u_{r+1} can be determined. The following procedure will attempt to find the maximum bandwidth at the current prices and budgets. It also calculates the price impact of the change in consumption on itself. In microeconomics this is similar to *internalizing externality*. The initial u_{r+1} is,

$$u_{r+1} = \begin{cases} b_m & \text{if } \hat{b}_{max} \geq b_m \\ \hat{b}_{max} & \text{if } \hat{b}_{max} < b_m \text{ AND } \hat{b}_{max} \geq b_{min} \\ \emptyset & \text{otherwise, } b_{min} \text{ was not affordable} \end{cases} \quad (2)$$

Using the price quotes, the NB must determine if the u_{r+1} will cause a price change that the user cannot afford, minimizing the externality of the bandwidth used. The highest price that the user can afford at link i is,

$$\frac{w^i}{u_{r+1}}. \quad (3)$$

The new price caused by u_{r+1} at link i is,

$$p^i + c \cdot \left(\frac{u_{r+1} + d^i - \alpha \cdot S^i}{\alpha \cdot S^i} \right) . \quad (4)$$

where d^i is the aggregate bandwidth demand of all users on link i . The new price given in equation 4 can not exceed the maximum price affordable, given in equation 3. Using these equations the following inequality provides a bound on feasible u values,

$$w^i \geq u_{r+1} \cdot \left[p^i + c \cdot \left(\frac{u_{r+1} - u_r + d_n^i - \alpha \cdot S^i}{\alpha \cdot S^i} \right) \right] . \quad (5)$$

Solving (5) for u_{r+1} yields the bandwidth at link i whose price change the user can afford. The inequality (5) has a closed form or it can be solved iteratively.

As described earlier, once the NB has determined its u_{r+1} it will start sending immediately at this rate. No signaling is performed. This technique provides a significant reduction in overhead; however an over allocation of resources may occur. Consider the following scenario. Assume many users are using one link and the price has reached an equilibrium value. Now assume one user ends their session and this reduction of bandwidth results in a lower price. If the remaining users react to this lower price, over-allocation of bandwidth may occur. One simple approach to prevent this situation is to have the switch adjust c so the price decreases at a slower rate. An over-allocation may still occur if many users using a link start sending at a higher rate simultaneously due to their application (not price); however this would require a correlation of these events. In general, adjusting the price based on $\alpha \cdot S^i$ and the high capacity of most links diminish the significance of this problem.

4 Optimality

As with any allocation strategy there are certain optimal allocation goals. Since pricing is used, optimality will be described in microeconomics terms. There are two important goals this technique strives for; Pareto optimal allocation and price stability.

4.1 Pareto Optimality

In this section we define the conditions, that are necessary for the competitive markets to reach a Pareto optimal distribution of bandwidth. Pareto optimality is the allocation of finite resources such that no sub-set of users can improve on their allocation without lowering the utility of another, given that supply equals demand. This is a standard goal in microeconomics for social benefit of resource distribution. The proof provided is based on one by Akira Takayama [16] and was adapted for our competitive market model. This proof was chosen because it does not require strict assumptions on the utility functions of users, as other Pareto proofs require.

The notation for this section is as follows. The network economy consists of several individual markets. A single market is composed of n consumers (user and NB pair) and a single producer (link). Let d^i be the demand of consumer i where $i = 1 \dots n$ and $d = \sum_{i=1}^n d^i$. Denote the demand set of i as D^i , where D^i is a subset of R^n . An initial allocation of resources to consumer i is denoted as \bar{d}^i and $\bar{d} = \sum_{i=1}^n \bar{d}^i$. Let S be the resource supply of the producer. Given a price p , the profit of the producer is $p \cdot S$.

Definition 4.1 *Feasibility:* An array of demand vectors $\{d^i\}$ is said to be feasible if $d = S + \bar{d}$.

Definition 4.2 *Pareto Optimal:* A feasible $\{\hat{d}^i\}$ is said to be Pareto optimal if there does not exist a feasible $\{d^i\}$ such that $d^i \succeq \hat{d}^i$ for all $i = 1 \dots n$ with \succ for at least one i .

Definition 4.3 *Competitive Equilibrium:* An array of vectors $[p, \{\hat{d}^i\}, S]$ is called a competitive equilibrium, if $\hat{d}^i \in D^i$, $i = 1 \dots n$ and

$$(i) \hat{d}^i \succeq d^i \text{ for all } d^i \in D^i \text{ such that } \hat{p} \cdot d^i \leq \hat{p} \cdot \hat{d}^i, i = 1 \dots n$$

$$(ii) \hat{d} = S + \bar{d}$$

Definition 4.4 *Local Nonsatiation Point:* A point $d^i \in D^i$ is called a local nonsatiation point, if there exists a $\delta > 0$ with $B_\delta(d^i) \cap (D^i \setminus d^i) \neq \emptyset$ such that for any $\epsilon, 0 < \epsilon < \delta$, with $B_\epsilon(d^i) \cap (D^i \setminus d^i) \neq \emptyset$, we have $d_0^i \succ d^i$ for some $d_0^i \in B_\epsilon(d^i) \cap D^i$, where $B_\delta(d^i)$ and $B_\epsilon(d^i)$ are open balls with center d^i and radii δ and ϵ , respectively.

Definition 4.5 *Locally Nonsaturating \succeq* : The preference ordering \succeq is called locally nonsaturating if given any local nonsatiation point d^i , $d_0^i = d^i$ implies that d_0^i is also a local nonsatiation point.

Nonsaturation assumption. The preference \succeq ordering is locally nonsaturating for every consumer.

Lemma 4.1 Let \hat{d}^i be a locally nonsatiation point for the i th consumer when price \hat{p} prevails. Then under the nonsaturation assumption, $d^i \succ \hat{d}^i$ implies $\hat{p} \cdot d^i > \hat{p} \cdot \hat{d}^i$.

Proof. Suppose this is not true; therefore $\hat{p} \cdot d^i \leq \hat{p} \cdot \hat{d}^i$. Since \hat{d}^i is the chosen point at price \hat{p} , $\hat{d}^i \succeq d^i$, which is a contradiction. \square

Theorem 4.2 Let $[\hat{p}, \{\hat{d}^i\}, S]$ be a competitive equilibrium such that \hat{d}^i is a local nonsatiation point for all $i = 1 \dots n$. Suppose the nonsaturation assumption holds true for all i . Then $[\{\hat{d}^i\}, S]$ is a Pareto optimum.

Proof. Suppose $[\{\hat{d}^i\}, S]$ is not a Pareto optimum. Then there exists $\{d^i\}, S$ such that $d^i \in D^i$, $i = 1 \dots n$ and

- (i) $d = S + \bar{d}$
- (ii) $d^i \succeq \hat{d}^i$ for all $i = 1 \dots n$
- (iii) $d^i \succ \hat{d}^i$ for some i

Hence from lemma 4.1 we have

$$\sum_{i=1}^n \hat{p} \cdot d^i > \sum_{j=1}^n \hat{p} \cdot \hat{d}^i \quad \text{or} \quad \hat{p} \cdot d > \hat{p} \cdot \hat{d}$$

But definition 4.3, condition (ii) requires

$$\hat{p} \cdot \hat{d} = \hat{p} \cdot S + \hat{p} \cdot \bar{d}$$

Therefore we have

$$\hat{p} \cdot d > \hat{p} \cdot S + \hat{p} \cdot \bar{d}$$

Which contradicts the feasibility of $\{d^i\}$. \square

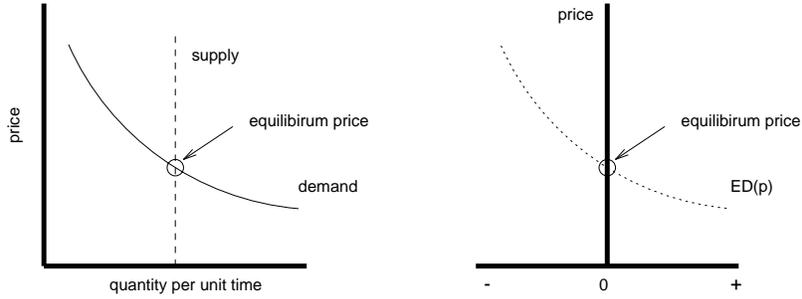


Figure 4: Example supply, demand and excess demand curves.

4.2 Price Stability

The equilibrium price (p_*) occurs when a price is reached such that the demand equals the supply. At this point, the resources are fully utilized. If the demand changes, pricing mechanism should alter the price to return to equilibrium. A proof that the proposed pricing technique achieves this is given next.

Walrasian price stability states that prices will adjust in response to supply and demand. For that reason, adjustments in the price are driven by knowledge from the market concerning the *excess demand* at a specific price. Denote the demand for bandwidth at price p as $d(p)$. For a link in the network the excess demand at price p is,

$$x(p) = (d(p) - \alpha \cdot S) . \quad (6)$$

Example supply, demand and excess demand curves for the system are given in figure 4. As seen in this figure, the demand curve has a negative slope. This represents that an increase in price will reduce demand. The supply curve is a vertical line, because the supply of bandwidth is constant (the link does not produce bandwidth). From the supply and demand curves the *excess demand* curve can be derived.

Using these graphs we can predict the behavior of the price rule (1). We will define stability as,

$$\lim_{t \rightarrow \infty} p \rightarrow p_* .$$

The price rule will increase the price p when it is lower than equilibrium price p_* . This is done because a positive excess demand exists. When p is greater than p_* , it is lowered towards p_* because

the excess demand is negative. Therefore the price rule always moves the price towards p_* , resulting in a stable equilibrium price. It should be noted that the slope of the supply curve must be positive for this to be true.

The equilibrium price can be proven stable mathematically as well. Using the excess demand equation, the price adjustment from the price equation (1) over time can be written as,

$$\frac{dp}{dt} = c \cdot \frac{x(p)}{\alpha \cdot S} . \quad (7)$$

Since $\alpha \cdot S$ is constant (the switch does not produce bandwidth) define the constant a as,

$$a = \frac{c}{\alpha \cdot S} .$$

Using the previous definitions, the price adjustment can be rewritten as,

$$\frac{dp}{dt} = a \cdot x(p) . \quad (8)$$

The price adjustment can be viewed as a first-order differential equation. The local response of the equation can be analyzed in the region of an equilibrium price using the Taylor approximation,

$$\frac{dp}{dt} = a \cdot x(p_*) + a \cdot x'(p_*) \cdot (p - p_*)$$

$$\frac{dp}{dt} = a \cdot x'(p_*) \cdot (p - p_*) . \quad (9)$$

The general solution to this equation is,

$$p_n = (p_0 - p_*) e^{a \cdot x'(p_*) \cdot t} + p_* \quad (10)$$

where p_0 is the initial price. As seen from the solution, system is stable as time increases. However it must be the case that $x'(p_*)$ is negative, as illustrated in figure 4 . The constant a is immaterial for the stability property [16]. However, the number of iterations required to reach the equilibrium price depends on the traffic, the budgets and the constant c . The switch has no control over first two items, yet some basic information can help in the selection of c . As an example, the impact

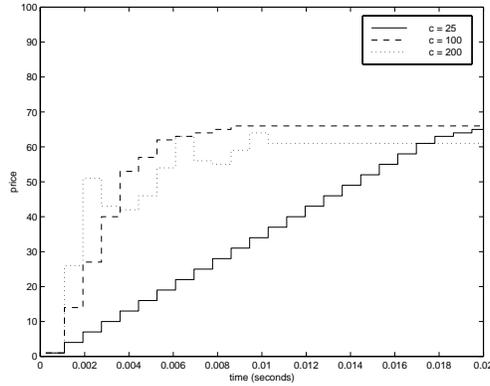


Figure 5: Price change over time using different c values.

of c on the number of iterations is given in figure 5. Values too low will cause the price to change slowly, resulting in more iterations. In the graph a c value of 25 requires 22 iterations, while c values of 100 to 350 require approximately 12 iterations. However values too large may raise the price too quickly, causing some users to exit before a price correction can be made.

5 Network Dynamics

Thus far, the description and analysis of the network economy has not considered the dynamic nature of an actual computer network. The dynamics we are interested in include; users entering/exiting the network, and allowing Variable Bit Rate (VBR) sources. Although prevalent in actual networks, these dynamics have been either or both excluded in other microeconomic flow control methods.

As described in the introduction, multimedia applications will constitute a large portion of the applications in current computer networks. The traffic generated by these applications can be described as VBR, which means the bandwidth required will change often and unexpectedly. Restricting the user to a constant desired bandwidth, as described in section 3.2, requires the user to purchase the highest amount of bandwidth expected (peak rate). For VBR sources, this approach is both difficult to implement and inefficient. Implementation is difficult since the peak rate may not be known in advance (consider live or interactive video). Purchasing only the peak rate is inefficient since the application may only require the peak rate for a short period of time. For

these reasons it is advantageous to allow the user to change the desired bandwidth over time. For a particular application, denote the m th desired bandwidth change as t_m , and the interval of time between bandwidth changes t_m and t_{m+1} as the m th application interval, A_m . The bandwidth desired during A_m is constant and is denoted as b_m . It is important to note the length of A_m depends on the application and will vary over time. At the end of A_m the new desired bandwidth b_{m+1} is sent to the NB. Now the NB determines a new amount of bandwidth to use, u_{r+1} , when either a new price or new desired bandwidth is received. The procedure for determining u_{r+1} is described in section 3.3.1. Once u_{r+1} has been determined the user starts sending at this rate immediately.

Since the number of users and demands for bandwidth change over time, the aggregate demand, d_n , for a link will vary as well. As a result there is not a single equilibrium price, p_* , for all time. However, the market can be viewed as having multiple equilibrium prices, each for some segment of time. During a segment the pricing technique will seek the equilibrium price as described in section 4. Once this price is found, the resulting distribution is Pareto optimal. When the aggregate demand changes, the stability of the price equation ensures that the price of bandwidth always moves towards p_* .

6 Experimental Results

In this section the performance of the network economy is investigated via simulation. Previous microeconomic flow control techniques either do not provide experimental results or simulate limited networks (network size and/or traffic source types). Experiments performed will consist of a realistic network configuration, allow users to enter/exit the network, have different application types and use actual MPEG-compressed traffic. Since max-min fairness is a goal of many flow control techniques [1], a comparison with max-min is provided. Experimental results will show that the proposed pricing technique achieves a fair Pareto distribution, provides higher QoS scores than max-min, and high network utilization.

The network simulated consisted of 92 users/NB, four switches and four primary links, as seen in figure 6. Each output port carried traffic from 38 users and connected to a 55 Mbps link. Links

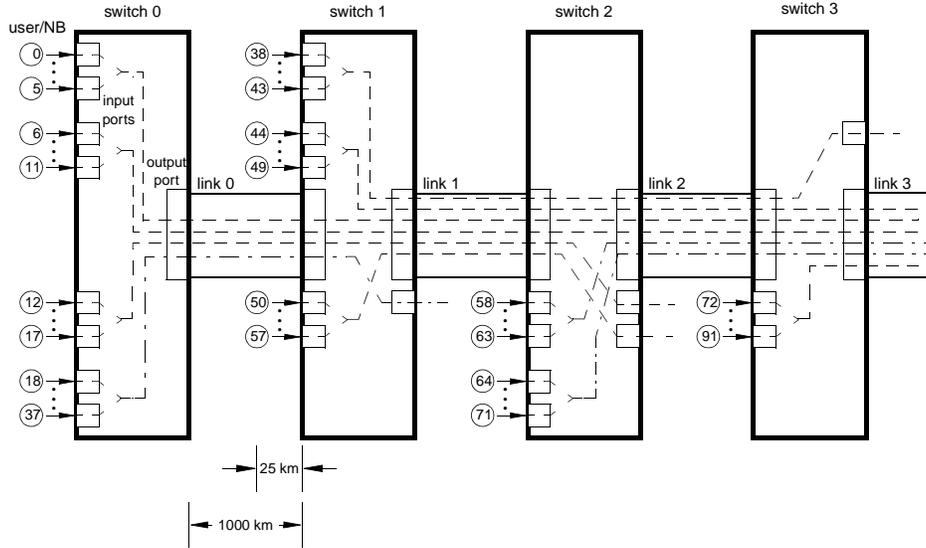


Figure 6: Network configuration used in simulations.

interconnecting switches were 1000 km in length, while links connecting sources to their first switch were 25 km in length. Users had routes ranging from one to four hops, and entered the network at random times uniformly distributed between 0 and 120 seconds. The network can be described as a "parking lot" configuration, where multiple sources use one primary path. This configuration was agreed upon by members of the ATM Forum [21] as a suitable benchmark for allocation methods; it models substantial competition between users with differing routes and widely-varying propagation delays.

For this simulation applications were one of two types, Multimedia on Demand (MoD) or teleconferencing. MoD applications require the transmission of high quality voice and video. These applications can scale bandwidth requirements only within a limited range, since bandwidth control is achieved through quantizer control [18]. The QoS profile associated with MoD applications is given in figure 7(a). As seen in the profile, the acceptable bandwidth ratio range (i.e., resulting in a QoS score greater than or equal to 3) is relatively small, 0.85 to 1.0. Teleconferencing applications, in contrast, transmit a lower quality voice and video and can scale bandwidth requirements within a larger range. This is primarily due to quantizer control as well as the ability to transmit below the standard 24 or 30 frames-per-second. The QoS profile associated with teleconferencing applications is given in figure 7(b); the acceptable bandwidth ratio range is 0.4 to 1.0. Regardless of the type

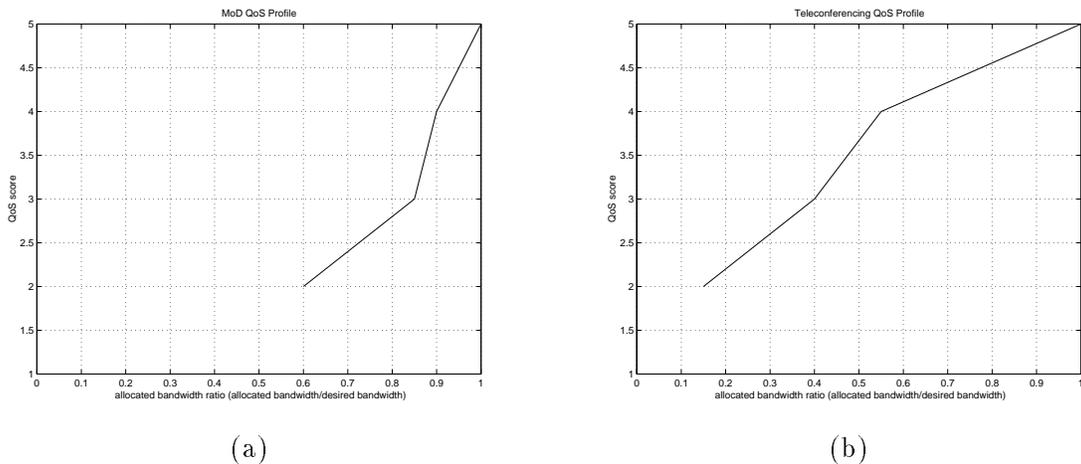


Figure 7: (a) MoD QoS profile. (b) Teleconferencing QoS profile.

of application, the source for each user was one of 15 MPEG-compressed traces obtained from Oliver Rose at the University of Würzburg, Germany [22]². Identifying each trace with a unique number (0 - 14), user i transmitted video trace $mod(i, 15)$, where $i = 0 \dots 91$. Each trace is a thirty minute segment of the original video and each was encoded with constant quality using the same MPEG-1 encoder card. Relevant statistics of each video are presented in [23] and [22]. As reported in [22], the Hurst parameters indicate all videos exhibit long-range dependency, and significant peak-to-mean ratios ranging from 18.4 to 4.63 based on frame size. Therefore, it is evident that these are very difficult sources to regulate. To date no other microeconomic flow control method has provided experimental results with actual MPEG sources or diverse application types.

The pricing strategy had the following initial values. MoD users had a budget rate, W , of $3 \times 10^7/\text{sec}^3$, while teleconferencing users had a budget rate of $1.5 \times 10^7/\text{sec}$. Teleconferencing users are given a lower budget because they are able to scale bandwidth requirements more readily. Switches initialized their prices to 1, their price equation c constant to 50 and α (the target utilization) to 95%. This utilization target is extremely aggressive when coupled with QoS requirements. We assumed there was no propagation delay between the user application and its NB, since they are expected to run on the same host system. Switches updated their link prices at an interval equal

²Traces can be obtained from the ftp site [ftp-info3.informatik.uni-wuerzburg.de](ftp://ftp-info3.informatik.uni-wuerzburg.de) in the directory `/pub/MPEG`

³The denomination is based on bps; if based on Mbps, the budget would be 300/sec.

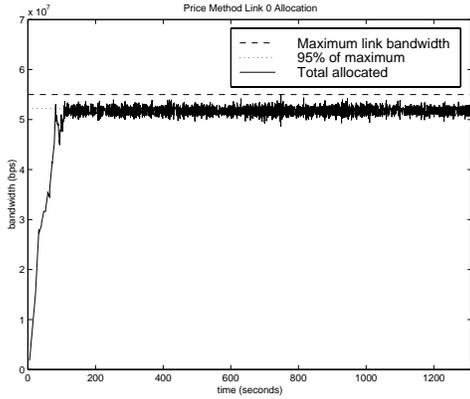
to 20 times the shortest propagation delay to any user to which it is connected. This interval is a compromise between the desire for responsiveness, and the need for stability.

As described in the introduction, the max-min fairness criterion states that any user is entitled to as much bandwidth as any other. When a link is bottlenecked, the bandwidth is divided equally among the users of the link. If a user requires less than this amount, the difference is divided equally among the remaining users. This process is repeated until all users of the link have been allocated a maximum amount of bandwidth. There is no distinction between application types. A more detailed description for networks is provided in [1]. In this simulation, the max-min allocation for the entire network was calculated after each source renegotiated and the resulting QoS scores were then recorded. The exact max-min algorithm, rather than an approximation, was implemented. No implementation overhead or propagation delays were included. We are therefore comparing to the best performance max-min can provide, which is probably better than the performance that can be achieved in practice.

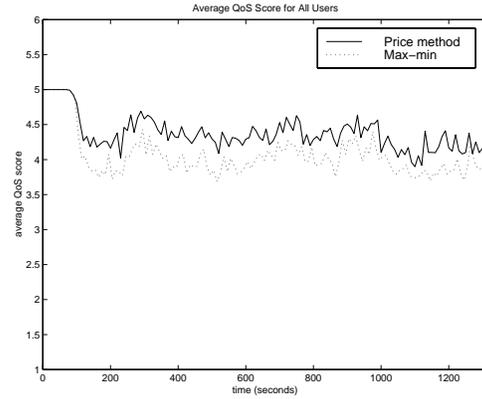
For comparisons, we are interested in the link bandwidth utilization and the QoS provided to each user. Allocation graphs are provided to measure the utilization of link bandwidth. To measure the QoS observed, average QoS graphs, percent Good or Better (GoB) measurements and average QoS scores are provided. Average QoS graphs measure the average QoS score observed over time and are based on all users or on individual type. The percent Good or Better (GoB) measurement is the average percentage of time a user had a quality score of at least 3.

For this simulation, the price method bandwidth allocation for link 0 is given in figure 8(a). The results for other links are very similar. The allocation graph indicates that the total allocation of bandwidth stayed in the vicinity of 95% (α , the target utilization), yet never crossed 100%. The fluctuation around 95% is the result of the changing demands created by the variable-bit-rate sources. Note that the time required for convergence, and the number of bandwidth changes, is no greater than for max-min.

The average QoS score graph, figure 8(b), shows that the price method always provided a higher average QoS score. This is also indicated in table 1, where the price method average QoS score was 4.37 as compared to 3.95 for max-min. The percent GoB for the price method was also 20%

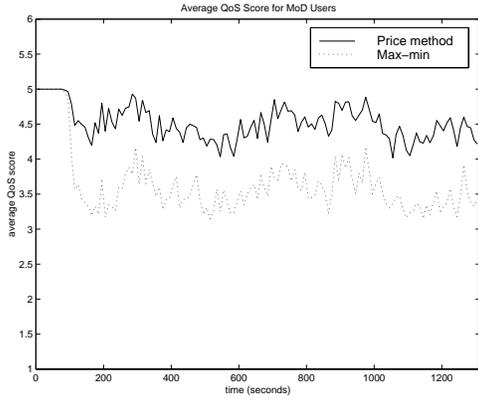


(a)

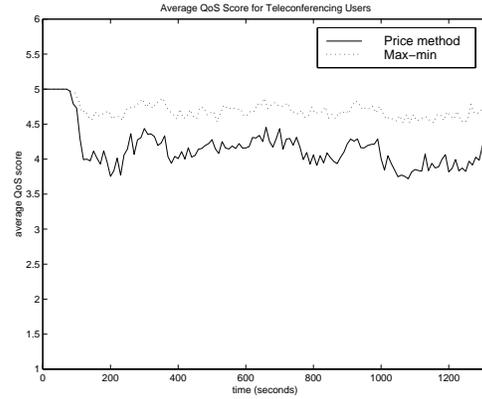


(b)

Figure 8: (a) Price method link 0 allocation. (b) Average QoS score for all users.



(a)



(b)

Figure 9: (a) Average QoS score for MoD users. (b) Average QoS score for teleconferencing users.

	%GoB			Average QoS Score		
	All	MoD	Teleconf.	All	MoD	Teleconf.
Price method	90	88	91	4.37	4.51	4.15
Max-min	75	60	99	3.95	3.64	4.69

Table 1: Percent GoB and average QoS scores.

higher than max-min. This indicates that users, under the price method, enjoyed an acceptable QoS for a longer duration. The difference between the price method and max-min is more distinct when considering the QoS provided to the two types of applications individually. In figure 9(a), the price method provides a higher QoS score for MoD applications than max-min. This is also indicated in the MoD values in table 1, where the average QoS score is 24% higher and the percent GoB was 47% greater. This is due to the inability of max-min to differentiate between MoD users and teleconferencing users. When a link becomes congested, the max-min distributes bandwidth equally among bottlenecked users. However, a reduction in bandwidth reduces the QoS for MoD users more quickly than teleconferencing users (as defined by their profiles). This is also evident in the average QoS graph for teleconferencing users, figure 9(b) and the average QoS scores in table 1. In contrast, the pricing method provides more bandwidth to MoD users than teleconferencing users. As a result the average QoS score for either type is almost equal. We believe it is more desirable to allocate so as to provide comparable QoS rather than equal bandwidth amounts. In table 1 all percent GoB and the average QoS scores differ by no more than 9% for the price method. In contrast, the MoD and teleconferencing percent GoB values differ by more than 65% for max-min. For this simulation, the price method was able to price link bandwidth in such a manner that lead to high utilization and better QoS performance than tradition max-min. Users were able to purchase link bandwidth, maximizing their QoS score individually and yielding a high percent GoB.

7 Conclusions

This paper introduced a decentralized flow control method based on microeconomics. A computer network was viewed as an economy consisting of three entities; users, Network Brokers (NB) and switches. Switches own the resources sought by users, and price their resources based on local supply and demand. A user requires these resources to maximize their individual QoS. Representing the user in the economy, the NB makes the resource purchasing decisions based on current needs of the user and prices. Users and switches act independently, which yields a decentralized flow control method. This competitive market structure encourages high utilization, with equilibrium

pricing and Pareto optimal resource distribution. There are fewer restrictions on the network than required by other methods based on microeconomics, and behavior during the convergence period is described, as well as illustrated experimentally. This paper also discussed how this economy properly handles network dynamics, such as users entering/exiting, and VBR traffic sources. Simulation results demonstrate the ability of the economy to successfully allocate bandwidth of a network to a large number of users, each transmitting an actual MPEG-compressed video trace. Utilization for this network was over 95% and the allocation of link bandwidth provided substantially better control of QoS than max-min. The price method has also been shown to perform better than other standard flow control schemes [24]. Finally, we believe the implementation cost will be very reasonable, since most of the functionality is in the host systems (NB) rather than in the switches or routers.

Some areas for future work include application to ABR traffic in ATM networks, wealth distribution (an issue for any economy), and appropriate parameter selection. While this paper has advocated microeconomics theory solely for flow control, our approach can potentially be applied to usage-based billing and cost recovery.

References

- [1] D. Bertsekas and R. Gallager, *Data Networks*. Prentice Hall, second ed., 1992.
- [2] M. Gerla and L. Kleinrock, "Flow control: A comparative survey," *IEEE Transactions on Communications*, vol. 28, pp. 553 – 574, April 1980.
- [3] ATM Forum Technical Committee, "Traffic management specification." Available through <ftp://ftp.atmforum.com/pub/approved-specs/af-tm-0056.000.ps>, 1996.
- [4] H. Ji, J. Y. Hui, and E. Karasan, "Gos-based pricing and resource allocation for multimedia broadband networks," in *IEEE INFOCOM'96*, pp. 1020 – 1027, 1996.
- [5] H. Jiang and S. Jordan, "A pricing model for high speed networks with guaranteed quality of service," in *IEEE INFOCOM'96*, pp. 888 – 895, 1996.
- [6] S. Low and P. Varaiya, "An algorithm for optimal service provisioning using resource pricing," in *IEEE INFOCOM'94*, pp. 368 – 373, 1994.
- [7] J. K. MacKie-Mason and H. R. Varian, "Pricing congestible network resources," *IEEE Journal on Selected Areas in Communications*, vol. 13, pp. 1141 – 1149, Sept 1995.
- [8] J. Murphy and L. Murphy, "Bandwidth allocation by pricing in atm networks," in *ITC*, June 1995.

- [9] N. Anerousis and A. A. Lazar, "A framework for pricing virtual circuit and virtual path services in atm networks," *ITC-15*, pp. 791 – 802, 1997.
- [10] R. Cocchi, S. Shenker, D. Estrin, and L. Zhang, "Pricing in computer networks: Motivation, formulation, and example," *IEEE/ACM Transactions on Networking*, vol. 1, pp. 614 – 627, Dec 1993.
- [11] D. F. Ferguson, C. Nikolaou, and Y. Yemini, "An economy for flow control in computer networks," in *IEEE INFOCOM'89*, pp. 110 – 118, 1989.
- [12] D. F. Ferguson, C. Nikolaou, J. Sairamesh, and Y. Yemini, "Economic models for allocating resources in computer systems," in *Market Based Control of Distributed Systems* (S. Clearwater, ed.), World Scientific Press, 1996.
- [13] J. F. Kurose and R. Simha, "A microeconomic approach to optimal resource allocation in distributed computer systems," *IEEE Transactions on Computers*, vol. 38, pp. 705 – 717, May 1989.
- [14] J. Sairamesh, D. F. Ferguson, and Y. Yemini, "An approach to pricing, optimal allocation and quality of service provisioning in high-speed packet networks," in *IEEE INFOCOM'95*, pp. 1111 – 1119, 1995.
- [15] J. Murphy, L. Murphy, and E. C. Posner, "Distributed pricing for atm networks," *ITC-14*, 1994.
- [16] A. Takayama, *Mathematical Economics*. Cambridge University Press, 1985.
- [17] L. Walras, *Elements of Pure Economics*. Richard D. Irwin, 1954. trans. W. Jaffé.
- [18] D. Reininger and R. Izmailov, "Soft quality-of-service for vbr+ video," in *Proceedings of the International Workshop on Audio-Visual Services over Packet Networks, AVSPN'97*, Sept. 1997.
- [19] E. Nakasu *et al.*, "A statistical analysis of mpeg-2 picture quality for television broadcasting," *SMPTE Journal*, pp. 702–711, November 1996.
- [20] J. Lourens, H. Malleson, and C. Theron, "Optimization of bit-rates for digitally-compressed television services as a function of acceptable picture quality and picture complexity," *IEE Colloquium on Digitally-Compressed TV by Satellite*, 1995.
- [21] R. Jain, "Congestion control and traffic management in atm networks: Recent advances and a survey," *Computer Networks and ISDN Systems*, Feb. 1995.
- [22] O. Rose, "Statistical properties of mpeg video traffic and their impact on traffic modeling in atm systems," Tech. Rep. 101, University of Würzburg Institute of Computer Science, Feb. 1995.
- [23] E. W. Fulp and D. S. Reeves, "Dynamic Bandwidth Allocation Techniques," Tech. Rep. TR-97/08, Center for Advanced Computing and Communication, Aug. 1997.
- [24] E. W. Fulp, M. Ott, D. Reininger, and D. S. Reeves, "Congestion pricing flow control for computer networks," tech. rep., Center for Advanced Computing and Communication, April 1998.