Distributed Network Flow Control Based on Dynamic Competitive Markets*

Errin W. Fulp and Douglas S. Reeves
Departments of Computer and Electrical Engineering, and Computer Science
North Carolina State University
ewfulp|reeves@eos.ncsu.edu

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 and equitable (QoS-fair) bandwidth distribution. 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 [3] [7]. 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 economic 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 [3]. 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 [2]. These strategies rely on the circulation of a Resource Management (RM) cell per connection

*This work was supported by AFOSR grants F49620-96-1-0061 and F49620-97-1-0351. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the AFOSR or the U.S. Government.

[2]. 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 [2]. 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. A more detailed description is provided in [3]. However, max-min does not take into account the fact that some sources (compressed video) may be able to reduce their transmission rate more easily than others. Therefore when congestion occurs, this allocation criterion 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 microeconomic 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 [9, 10, 13, 14]. Utility is a measurement of satisfaction, and 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 [9, 10] or on traffic types (or service classes) [13]. 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 [1, 4, 5, 11, 20]. 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 [4]. Ferguson, et al. is an example of virtual circuit flow control based on pricing network resources [5]. 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 Nash 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 QoS and high resource utilization, at a modest implementation cost. We also introduce a new notion of allocation fairness, *equitable* allocation. Based on microeconomic social welfare theory, an equitable allocation is one in which all users receive approximately the same level of utility [17]. It is important to note that this does not necessarily correspond to equal amounts of a resource (the goal of max-min). An equitable allocation can also be referred to as "QoS-fair." Similar to several microeconomic flow control methods [1, 5, 20] 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. Can provide equitable (QoS-fair) allocations.
6. Allow network dynamics (VBR sources and users entering/exiting).

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 is in equilibrium. This moment is referred to as "clearing the market" and the resulting allocation is Pareto optimal [21]. 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. 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.

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 nth 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 nth 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 nth 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 nth 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,

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

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 [22]. 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 [21]. 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 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], [16], and [12].

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.

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. 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, $\{l^i, i = 1 \ldots v\}$. For each link on R, a price quote, q^i, is collected, where $\vec{q} = \{q^i, i = 1 \ldots v\}$ 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 $\vec{w} = \{w^i, i = 1 \ldots v\}$ 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.

[1]The 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.

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 rth 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.

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,

$$b^i_{max} = \frac{w^i}{p^i} \; , i = 1 \ldots v$$

therefore the maximum bandwidth the user can afford is,

$$\widehat{b}_{max} = \min_{i=1 \ldots v} \{b^i_{max}\} \; .$$

Note this equation maximizes the bandwidth at the current prices. The minimum bandwidth that can be used, b_{min}, is determined from the QoS profile, b_m and the value that corresponds to the lowest acceptable QoS score. It is possible that $\widehat{b}_{max} < b_{min}$ (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 $\widehat{b}_{max}$ 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,

$$u_{r+1} = \begin{cases} \min\{\widehat{b}_{max}, b_m\} & \text{if } \widehat{b}_{max} \geq b_{min} \\ \emptyset & \text{otherwise} \end{cases} \qquad (2)$$

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,

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

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

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

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,

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

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.

As described in section 2, 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. Several proofs have been developed to show that competitive markets reach a Pareto optimal distribution [21]. A proof that our computer network economy achieves a Pareto optimal distribution is given in [6].

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. This property is what is referred to as price stability. A proof that our proposed pricing technique has price stability is also given in [6].

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 either or both been 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 mth desired bandwidth change as t_m, and the interval of time between bandwidth changes t_m and t_{m+1} as the mth 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. 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 equilib-

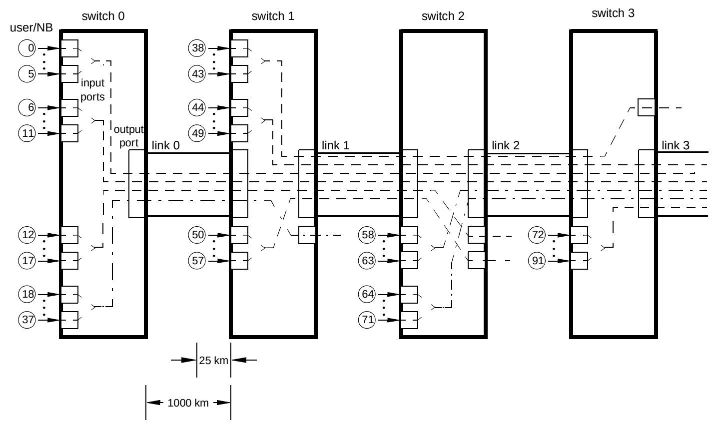

Figure 1. Network configuration used in simulations.

rium 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. A comparison with max-min is provided, since max-min fairness is a goal of many flow control techniques [3]. The max-min implementation was centralized and did not include communication overhead; therefore the max-min results presented here should be considered better than what is possible in practice. Experimental results will show that the proposed pricing technique achieves high network utilization and equitable (QoS-fair) allocations, as well as higher QoS scores than max-min.

The network simulated consisted of 92 users/NB, four switches and four primary links, as seen in figure 1. Each output port carried traffic from 38 users and connected to a 55 Mbps link. Links 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 [8] 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. We are interested in determining if the users achieve similar QoS scores (utility) regardless of their application type (equitable allocation). The QoS profiles associated with MoD and teleconferencing applications are given in figure 2. 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]. 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 has a larger acceptable bandwidth ration range of 0.4 to 1.0. This is primarily due to quantizer control as well as the ability to transmit below the standard 24 or 30 frames-per-second. Regardless of the type of application, the source for each user was one of 15 MPEG-compressed traces obtained from Oliver Rose at the University of Würzburg, Germany [19][2]. Identifying each trace with a unique number (0 - 14), user i transmitted video trace $mod(i, 15)$, where $i = 0 \ldots 91$. Each trace is a thirty minute segment of the original video and each was encoded with constant

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

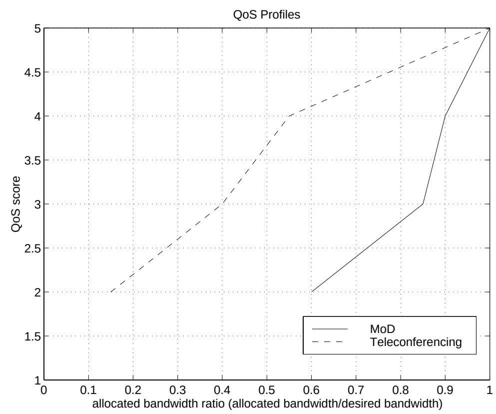

Figure 2. MoD and Teleconferencing QoS profiles.

quality using the same MPEG-1 encoder card.

The pricing strategy had the following initial values. MoD users had a budget rate, W, of 3×10^8/sec, while teleconferencing users had a budget rate of 1.5×10^8/sec (the denomination is based on bps; if based on Mbps, the budget would be 300/sec). Wealth was determined based on the bandwidth ratio required to provide a QoS score of 3. This was done to provide a more equitable allocation. 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 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.

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 3(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 3(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% 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 4(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 4(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; therefore the pricing method achieved a more equitable allocation.

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 achieves Pareto optimal and equitable resource distribution. There are fewer restrictions on the network

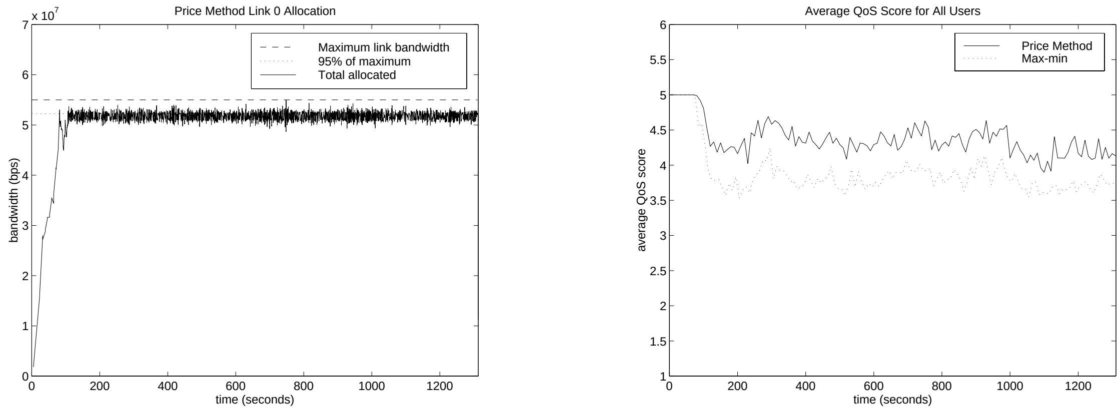

(a) Price method link 0 allocation. **(b) Average QoS score for all users.**

Figure 3. Allocation and average QoS score values.

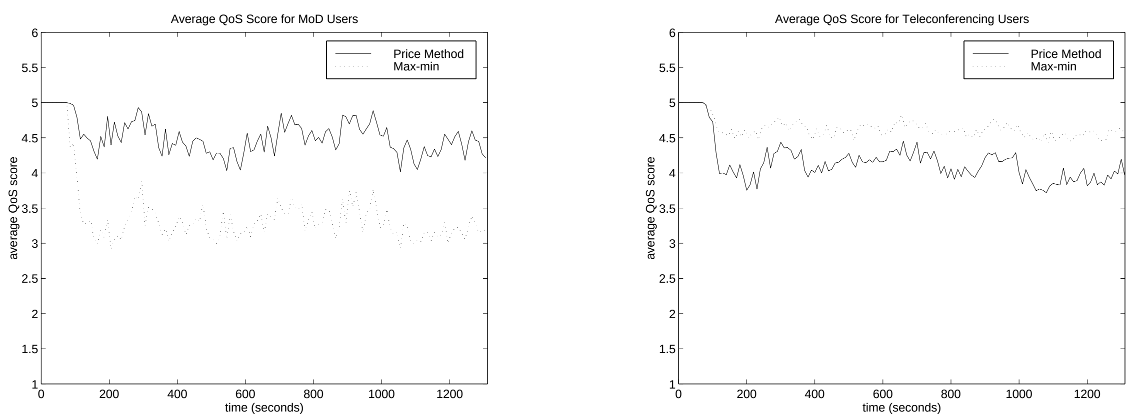

(a) Average QoS score for MoD users. **(b) Average QoS score for teleconferencing users.**

Figure 4. Average QoS score values for MoD and 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.

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 [6]. 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.

Acknowledgements The authors wish to thank Maximilian Ott and Daniel Reininger of C & C Research Laboratories, NEC USA for their significant contributions to this research.

References

[1] N. Anerousis and A. A. Lazar. A framework for pricing virtual circuit and virtual path services in atm networks. *ITC-15*, pages 791 – 802, 1997.

[2] ATM Forum Technical Committee. Traffic management specification. Available through `ftp://ftp.atmforum.com/pub/approved-specs/af-tm-0056.000.ps`, 1996.

[3] D. Bertsekas and R. Gallager. *Data Networks.* Prentice Hall, second edition, 1992.

[4] R. Cocchi, S. Shenker, D. Estrin, and L. Zhang. Pricing in computer networks: Motivation, formulation, and example. *IEEE/ACM Transactions on Networking*, 1(6):614 – 627, Dec 1993.

[5] D. F. Ferguson, C. Nikolaou, J. Sairamesh, and Y. Yemini. Economic models for allocating resources in computer systems. In S. Clearwater, editor, *Market Based Control of Distributed Systems.* World Scientific Press, 1996.

[6] E. W. Fulp, M. Ott, D. Reininger, and D. S. Reeves. Congestion pricing flow control for computer networks. Technical report, Center for Advanced Computing and Communication, April 1998.

[7] M. Gerla and L. Kleinrock. Flow control: A comparative survey. *IEEE Transactions on Communications*, 28(4):553 – 574, April 1980.

[8] R. Jain. Congestion control and traffic management in atm networks: Recent advances and a survey. *Computer Networks and ISDN Systems*, Feb. 1995.

[9] H. Ji, J. Y. Hui, and E. Karasan. Gos-based pricing and resource allocation for multimedia broadband networks. In *Proceedings of the IEEE INFOCOM*, pages 1020 – 1027, 1996.

[10] H. Jiang and S. Jordan. A pricing model for high speed networks with guaranteed quality of service. In *Proceedings of the IEEE INFOCOM*, pages 888 – 895, 1996.

[11] J. F. Kurose and R. Simha. A microeconomic approach to optimal resource allocation in distributed computer systems. *IEEE Transactions on Computers*, 38(5):705 – 717, May 1989.

[12] J. G. Lourens, H. H. Malleson, and C. C. Theron. Optimization of bit-rates for digitally compressed television services as a function of acceptable picture quality and picture complexity. In *Proceedings of the IEE Colloquium on Digitally Compressed Television by Satellite*, 1995.

[13] S. Low and P. Varaiya. An algorithm for optimal service provisioning using resource pricing. In *Proceedings of the IEEE INFOCOM*, pages 368 – 373, 1994.

[14] J. K. MacKie-Mason and H. R. Varian. Pricing congestible network resources. *IEEE Journal on Selected Areas in Communications*, 13(7):1141 – 1149, Sept 1995.

[15] J. Murphy, L. Murphy, and E. C. Posner. Distributed pricing for atm networks. *ITC-14*, pages 1053 – 1063, 1994.

[16] E. Nakasu, K. Aoi, R. Yajima, K. Kanatsugu, and K. Kubota. Statistical analysis of mpeg-2 picture quality for television broadcasting. In *Proceedings of the 7th International Workshop on Packet Video*, volume 11, pages 702 – 711, Nov. 1996.

[17] W. Nicholson. *Microeconomic Theory, Basic Principles and Extensions.* The Dryden Press, 1989.

[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] O. Rose. Statistical properties of mpeg video traffic and their impact on traffic modeling in atm systems. Technical Report 101, University of Würzburg Institute of Computer Science, Feb. 1995.

[20] 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 *Proceedings of the IEEE INFOCOM*, pages 1111 – 1119, 1995.

[21] A. Takayama. *Mathematical Economics.* Cambridge University Press, 1985.

[22] L. Walras. *Elements of Pure Economics.* Richard D. Irwin, 1954. trans. W. Jaffé.