

The Fairness and Utility of Pricing Network Resources Using Competitive Markets

Errin W. Fulp* and Douglas S. Reeves†

Department of Computer Science

North Carolina State University

1 Introduction

Current and future networks must accommodate a wide variety of network applications. These applications require a Quality of Service (QoS) which can be provided through proper network resource allocation. Examples include TCP congestion control, RSVP admission control, ATM ABR rate allocation, and others. Due to the finite amount of resources, contention frequently occurs in networks. For this reason, a method to manage these limited resources in a fair and efficient manner is needed. Recently, pricing has been promoted as a method for allocating network resources. Under these techniques, users are charged for their consumption and resources are priced to reflect supply and demand [1, 8, 9, 11, 16, 18, 19]. Benefits of pricing network resources include: flexible, precise, and distributed control of congestion; the ability to accomplish economic goals such as revenue generation; and provably fair resource allocations.

When pricing is used for network resource allocation, fairness definitions can be taken from traditional network theory [3] as well as microeconomics [20]. In addition, new fairness definitions have been introduced specifically for network pricing [16]. As a result of the many plausible definitions, it has often been difficult to determine which definition is achievable and appropriate. We believe the decision about which type of fairness to enforce is a policy, which should be independent of the mechanism used to allocate resources. The resource allocation method must also scale to a large number of users, and should be able to converge quickly to a new solution under network dynamics (users entering/exiting and variable bit rate sources). However, most pricing methods lack these attributes. For example, many pricing techniques are not intended to adapt to changing traffic demands that occur on various time-scales, nor have they been validated using realistic network topologies or actual traffic [1, 18]. In addition, the transient behavior and the method of distributing prices and allocations during convergence is generally ignored [9, 19]. Other limitations include the reliance on well-defined statistical models of source traffic and restrictions on the shape of the utility curve [8, 16, 18].

In this paper, we introduce a network resource pricing technique based on the competitive market model. Referred to as the “spot market approach,” the economic model consists of multiple dynamic competitive markets (spot markets) working asynchronously and independently. This allocation approach has all the advantages of other microeconomic-based methods; however, the spot market approach allows and encourages

*Post: N.C. State University, Box 7534, Raleigh, NC 27695 USA email: ewfulp@eos.ncsu.edu

†email: reeves@eos.ncsu.edu

network dynamics to occur. In addition, the allocation method is practical to implement, has no limiting assumptions, and is able to achieve multiple types of fairness.

The remainder of this paper is structured as follows. In the next section, the competitive market model is described, which serves as the basis for the proposed network economy. Section 3 reviews previous network pricing methods and introduces the spot market approach, a network implementation of the competitive market model. In section 4, experimental results will show that the spot market can quickly achieve fair allocations under realistic conditions. Finally in section 5, a summary of the pricing mechanism is provided and areas of future research are discussed.

2 Distributed Competitive Market Economy

The network economy described in this paper consists of multiple independent competitive markets. Each competitive market consists of one type of scarce resource and two types of agents: consumers and producers. The consumer has a wealth (budget) and acts independently (selfishly) purchasing resources to maximize utility (satisfaction). Producers own the resources sought by consumers and sell resources to maximize revenue. 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 can be determined in various manners; however, we will use a tâtonnement process [25]. First proposed by Léon Walras, the tâtonnement process iteratively adjusts the price with respect to supply and demand. Given a price, feedback from the consumers (bids, representing the demand) is collected and the price is adjusted. When a price is found that results in demand equaling supply, the market is in equilibrium. At this equilibrium price, consumers purchase resources and producers sell resources. Both agents are considered “price takers” and assume their individual actions will not significantly influence the market price determination. After the buying and selling of the resource is complete, the resulting allocation is provably optimal and fair.

2.1 Optimal and Fair Allocations

This section discusses the types of optimality and fairness that are possible when an economy, consisting of multiple competitive markets, is in equilibrium. First, the notation and assumptions about the economy are defined. Then, each type of optimality and fairness will be described, proven, and discussed.

Assume an economy consists of a set of independent competitive markets L . Each market i sells only a unique type of resource with supply s^i . This implies that an array of prices exists $\{p\}$ in the economy, where the price for the resource sold at market i is p^i . All consumers in the economy belong to set A , where consumer j desires resources belonging to the set $R^j \subset L$. Furthermore, let A^i denote the set of consumers who participate in market i . Consumer j has wealth w^j and desires a maximum amount b^j for any resource in R^j . Let $a^{j,i}$ be the allocation for consumer j in market i and, assume the consumer must purchase the

same amount in each market¹ ($a^{j,h} = a^{j,i}$, $\forall h, i \in R^j$). The second superscript of $a^{j,i}$ (i , indicating the market) will be dropped for brevity².

Depending on the price associated with each market, the consumer can afford different amounts; however, the consumer will prefer some amounts over others. These preferences are represented with a utility function, which maps a resource amount to a real number. Assuming $u^j(\cdot)$ is a utility function of consumer j , if the consumer prefers an amount x over y (this is represented using the notation $x \succ y$) then $u^j(x) > u^j(y)$. It is assumed that the utility function of each consumer is monotonically increasing [24], and a consumer becomes completely satiated with the desired resource amount b^j , above which the utility may decrease. As previously mentioned, assume the consumer will always purchase the same amount at each market $i \in R^j$. This amount a^j is equal to the minimum amount that is affordable at any market $i \in R^j$ (but no more than the maximum desired b^j),

$$a^j = \min \left\{ \min_{\forall i \in R^j} \left\{ \frac{w^j}{p^i} \right\}, b^j \right\} \quad (1)$$

The market in R^j with the highest price is considered *saturated* for consumer j (which ultimately determines the amount to purchase at the remaining markets in R^j). At the saturated market the consumer is non-satiated; however, for the remaining markets in R^j the consumer is considered satiated. In the case where the consumer can afford b^j at each market in R^j , then the consumer is considered completely satiated at each market in R^j .

Given the preceding economic model, we can prove optimal and fair allocations are obtained when the markets are in equilibrium. However, before defining and proving the fairness of the allocations, we must define a feasible allocation and the competitive equilibrium.

Definition 2.1. Feasibility: For competitive market i , the price and an allocation array, $[p^i, \{a\}]$, are said to be feasible if and only if,

- (i) $s^i \geq \sum_{j \in A^i} a^j$
- (ii) $p^i \cdot a^j \leq w^j \quad \forall j \in A^i$

Definition 2.2. Competitive equilibrium: At price p_*^i and allocation array $\{a\}$, competitive market i is in equilibrium if and only if,

- (i) $[p_*^i, \{a\}]$ is feasible
- (ii) $u^j(a^j) \geq u^j(\hat{a}^j)$ for all \hat{a}^j , where $\hat{a}^j \leq b^j$ and $p_*^i \cdot \hat{a}^j \leq w^j$

¹This assumption becomes clear when the economy is a computer network and the resource is link bandwidth.

²If the requirement ($a^{j,h} = a^{j,i}$, $\forall h, i \in R^j$) is removed, then optimal and fair allocations can be proved for individual markets (instead of an entire economy).

A feasible allocation requires: the demand to be less than or equal to the supply, and users to adhere to their budget constraints. The competitive equilibrium definition builds on the feasible allocation definition by requiring each consumer to maximize utility.

Using the precursory definitions, we can identify the conditions required to achieve optimal and fair allocations. A common economic definition of optimality is Pareto-optimality. A Pareto-optimal allocation is one where no agent can increase their utility without decreasing the utility of another [20]. The following theorem indicates if the individual markets of the economy are in equilibrium, then the allocation is Pareto-optimal. This also referred to as the First Fundamental Theorem of Microeconomics [20].

Definition 2.3. Pareto Optimality: The feasible allocation array $\{a\}$ is said to be Pareto-optimal if there does not exist another feasible allocation array $\{\hat{a}\}$, such that $u^j(\hat{a}^j) \geq u^j(a^j) \quad \forall j \in A$ with a strict inequality for at least one j .

Theorem 2.1. *The allocation of an economy consisting of independent competitive markets in equilibrium is Pareto-optimal.*

Proof. A proof that the allocation of an economy consisting of independent competitive markets in equilibrium is Pareto-optimal is given in [24]. This proof can be used for the economy presented in this paper with a slight modification [10]; however for brevity, the modified proof is not presented here.

Many different Pareto-optimal allocations exist, and theorem 2.1 does not imply the best allocation will always result. For this reason, other fairness definitions are often used to reduce the set of desired allocations. A weighted max-min fair allocation is one where each agent receives their fair-share with respect to their weight (priority or wealth) and who they compete against. This definition was proposed for computer networks [14], and is the fairness sought by many bandwidth allocation mechanisms [2, 12]. The following indicates that the allocation of the economy in equilibrium is weighted max-min fair, where the wealth of each user is their weight.

Definition 2.4. Weighted max-min fair: An allocation of resources $\{a\}$ with weights $\{w\}$ is weighted max-min fair if it is feasible, and if, for any other feasible allocation $\{\hat{a}\}$,

$$\exists j : \hat{a}^j > a^j \implies \exists k : \frac{\hat{a}^k}{w^k} < \frac{a^k}{w^k} \leq \frac{a^j}{w^j} \quad (2)$$

Theorem 2.2. *The allocation of an economy consisting of independent competitive markets in equilibrium is weighted max-min fair, where the weight of each consumer is their wealth.*

Proof. Given in appendix A.1.

It is important to note that the weighted max-min fair definition is indifferent to the number (different types) of resources required. For example, if prices are applied to each link in a network, then this definition

is unfair (biased towards longer paths) since they are only constrained by the saturated market (bottleneck link). In contrast, an allocation that is proportionally fair per unit charge does consider the number of resources desired by the agent [16]. For this reason, this type of fairness is often cited when pricing network resources [7, 16]. Now the price faced by the consumer j is the sum of all the prices in R^j .

Definition 2.5. Proportional fairness per unit charge: Let $\{\phi\}$ be an array of weights, where ϕ^j is the weight for consumer j . An allocation of resources $\{a\}$ is proportionally fair per unit charge if it is feasible, and if, for any other feasible allocation $\{\hat{a}\}$

$$\sum_{j \in A} \phi^j \cdot \frac{\hat{a}^j - a^j}{a^j} \leq 0 \quad (3)$$

Theorem 2.3. *The allocation of an economy consisting of independent competitive markets in equilibrium is proportionally fair per unit charge, where*

- (i) $\phi^j = w^j, \quad \forall j \in A.$
- (ii) *The allocation of consumer j is,*

$$a^j = \frac{\phi^j}{\sum_{i \in R^j} p^i} \leq b^j \quad (4)$$

- (iii) *The price for any under-utilized market i ($\sum_{j \in A^i} a^j < s^i$) is zero.*

Proof. Given in appendix A.2.

This section has reviewed the wide variety of optimal and fair allocations the competitive market economy is able to achieve. In addition, the economy can achieve various social welfare goals, which are measured in terms of utility.

2.2 Social Welfare Theory

The First Fundamental Theorem of Microeconomics asserts that at equilibrium the resulting allocation is Pareto-optimal; however, it does not advance the optimistic claim that the allocation is the best of all possible allocations. The Pareto-optimal allocation is efficient in that any other allocation that would make an agent better off would be at the expense of others. Since there are many different possible wealth (or weight) distributions in the economy, there is a large number of possible fair allocations. Social welfare theory can be used to target a particular fair allocation that achieves a social goal, measured in terms of utility or QoS (instead of resource amounts). Two common social goals are the equality criterion and the utilitarian criterion [20]. These fairness goals are achieved in our economy only if the wealth is properly distributed (a condition in both theorems); therefore, a wealth distribution algorithm is provided for each.

The equality criterion represents an allocation (called an equitable or QoS-fair [4] allocation) where consumers enjoy the same level of utility. This criterion is appropriate for networks when equivalent QoS (utility) for a class of traffic is desired (e.g. ABR service class [12], Differential Services [5], and wireless transmission [4]). The conditions necessary for the economy to achieve an equitable allocation is specified in the following theorem.

Definition 2.6. Equitable allocation: An allocation of resources $\{a\}$ is equitable if it is feasible, and if, for any other feasible allocation $\{\hat{a}\}$,

$$\exists j : u^j(\hat{a}^j) > u^j(a^j) \implies \exists k : u^k(\hat{a}^k) < u^k(a^k) \leq u^j(a^j) \quad (5)$$

Theorem 2.4. *Allocating wealth using algorithm 1 yields an equitable allocation for an economy consisting of independent competitive markets in equilibrium.*

Proof: Given in appendix A.3.

The utilitarian criterion (also called utility maximization) describes an allocation where the sum of the individual utilities is the greatest. Network pricing methods that achieve this criterion under certain conditions include [7, 16]. The conditions necessary for the economy to achieve the utilitarian criterion is specified in the following theorem.

Definition 2.7. Utilitarian criterion: An allocation of resources $\{a\}$ meets the utilitarian criterion if it is feasible, and if, for any other feasible allocation $\{\hat{a}\}$

$$\sum_{j \in A} u^j(a^j) > \sum_{j \in A} u^j(\hat{a}^j) \quad (6)$$

Theorem 2.5. *Allocating wealth using a modified algorithm 1 (modification described in appendix A.4) achieves the utilitarian criterion for an economy consisting of independent competitive markets in equilibrium.*

Proof: Given in appendix A.4.

Both criteria require the distribution of wealth to achieve a certain social goal. Typically, wealth algorithms require complete information from consumers (wealth and utility function) to distribute (assign) wealth. Reliably collecting this information may not be possible; therefore, the application of social welfare theory to networks is limited. While not applicable to large integrated networks (WAN's); small networks (LAN's), Differentiated Service classes, and wireless networks may benefit from social welfare theory.

3 Pricing Network Bandwidth

In this section we discuss the application of pricing for network resource allocation. In addition, we present a pricing method based on the competitive market economy described in section 2. While many different types

of networks exist, we will apply pricing to the ATM Available Bit Rate (ABR) service class [2]. Specifically, ABR bandwidth will be priced to control the sending rate of users in a fair and efficient manner.

The ATM ABR service class is one in which network characteristics may change during the lifetime of a connection. For this reason, the ABR service class is suitable for traffic that can adapt to changing network conditions (elastic-traffic [23]). Due to the dynamic nature of the ABR service class, proper rate control is essential and becomes more difficult when video is transmitted. Explicit rate control relies on network feedback provided by Resource Management (RM) cells that are circulated for each connection [2]. A RM-cell traveling from the source to the destination will be referred to as moving *upstream*, while a RM-cell traveling from the destination to the source will be referred to as moving *downstream*. The RM-cell consists of several fields, one of which is the Expected Rate (ER). This field indicates the maximum rate the network can support for this user. As the RM-cell travels along the path, a switch and/or 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, which must alter transmission based on the RM-cell information.

3.1 Previous ABR Pricing Methods

Microeconomic-based techniques designed specifically for ABR rate control include [6, 7]. In [6], switches allocate ABR bandwidth in a proportionally fair manner based on the “willingness-to-pay” provided by each user. When conditions change, users determine a new willingness-to-pay via a curve fitting process which relies on a history of previously optimal decisions. In the ABR rate control method of [7], users bid for some amount of effective bandwidth. While effective bandwidth allocates over a longer time scale, these techniques are difficult to apply to sources with little or no a priori information (for example, live and interactive video) and can be considered too conservative [7].

3.2 Bandwidth Spot Market

In this section, we introduce a new method for pricing ATM ABR bandwidth. Instead of using the competitive market defined in section 2, a modified competitive market (dynamic competitive market or spot market [10]) will be used. It will be shown, that the spot market is well suited for pricing networks resources. It has the unique characteristic of quickly adapting to changing network demands, that are intrinsic to the ABR service class. Furthermore, the spot market has the advantages associated with the competitive market such as optimal and fair allocations and high utilization.

The network economy consists of multiple spot markets and three entities: 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 in the exchange of resources in the market.

3.2.1 Switch

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 ABR bandwidth based on local supply and demand. Therefore a single switch, having multiple output ports, will have one price associated with each output port, where i represents the i th link of the economy. These spot markets operate independently and asynchronously since there is no need for market communications (e.g. price comparisons) or synchronization from switch to switch. We consider ABR bandwidth a non-storable resource (similar to residential electricity); therefore, users are charged for their consumption at the current market price. Bandwidth is sold for immediate use, without reservation overhead.

As defined by the ABR service class, a switch will periodically receive RM-cells. The RM-cell provides the user feedback about the links (or destination) in their route. We propose using the link price as feedback, since users must scale bandwidth consumption due to budget constraints. The price in the RM-cell is initialized to 0 by the source node. How the price in the RM-cell is updated by a switch depends on the targeted fairness. If a weighted max-min allocation (definition 2.4) is desired, then the switch would insert the *current* price for link i into the downstream RM-cell traversing link i if it is greater than the price already stored in the RM-cell. In contrast, if a proportionally fair per unit charge allocation (definition 2.5) is sought, then the switch would add the current price for link i to the price in the downstream RM-cell. Regardless of the desired fairness, we assume that the price is stored in the ER field of the RM-cell; therefore, no additional field is required and no other information is placed/alterd in the RM-cell.

The price for link i is calculated at the switch, at discrete intervals. At the end of the n th interval, the switch updates the price of link i using a modified tâtonnement process. A limitation of the tâtonnement process, in its original form, is the inability to dynamically adapt to changing demands (which are prevalent in networks). To handle such dynamics the following *modified* tâtonnement process [11] is used,

$$p_{n+1}^i = p_n^i \cdot \frac{d_n^i}{\alpha \cdot s^i} \quad (7)$$

where p_{n+1}^i is the new price, p_n^i is the current price, s^i is the link capacity and d_n^i is the aggregate demand for bandwidth. The modified tâtonnement process adjusts the price at regular intervals, based on the demand (received traffic) and the supply. The bandwidth supply is the total bandwidth times a constant α , where $0 < \alpha \leq 1$. This modification causes the price to increase after some percentage (α) of the total bandwidth has been sold. An *equilibrium price* p_*^i is reached at link i when the supply equals the demand. At equilibrium, the resulting allocation is optimal and fair as described in section 2.1. Refer to the prices calculated before the equilibrium price is reached as *intermediate prices*. Buying and selling normally do

not occur with intermediate prices. However, this constraint does not apply to the bandwidth spot market, since we consider ABR bandwidth a non-storable resource [24]. This provides users with immediate access of link bandwidth, without waiting for the equilibrium price to be determined. A single equilibrium price does not exist for all time; however, the process will move towards the new equilibrium price when demand changes [10].

3.2.2 User

The user, executing a network application, requires ABR bandwidth for transmission. In the economy, the user must rely on a Network Broker (NB) for purchasing bandwidth. For this reason, the user provides the NB information regarding: the application bandwidth demands, allocation and QoS correlation, and wealth. The amount of bandwidth desired by user j is determined from the application and is denoted as b^j . Since different amounts of bandwidth are affordable, the user must define a utility function that identifies the QoS obtained from certain bandwidth amounts (for example a QoS profile [21]). 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.

3.2.3 Network Broker

As previously mentioned, a user can only gain access to the network economy through a Network Broker (NB). The NB serves as an agent for the user and is located on the edge of the network. The functions of the NB can be part of the protocol stack that executes on the 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 purchasing decisions. Although the NB works as an agent for the user, assume the NB operates honestly in regards to the switches and the user.

The NB monitors the user and link prices by gathering and storing information about each. From the user, the NB collects and stores: the utility curve, b^j and w^j . Link prices along the route are gathered using the ABR RM-cell. As described in the previous section, the returned RM-cell contains either the highest price or the sum of the prices in the route (depending on the targeted fairness definition). Using this information, the NB can make connection admission control and purchasing decisions based on the current market conditions.

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, which also prevents overloading the economy. Hence, the goal

is 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.

Finally, the allowable transmission rate, a^j , is determined in response to a new price or a change in application demand. If weighted max-min is sought, then the amount of ABR link bandwidth is given from equation 1. Similarly, if proportionally fair per unit charge is desired, then the amount is determined from equation 4. Since the fairness definition is for the entire network (a network-wide objective), all NB's must use the same allocation formula.

4 Experimental Results

Due to the complexity of the changing source demands, simulations are used to demonstrate the performance of the spot market approach under network dynamics. Two sets of experiments were performed. The first experiment investigated the impact of the initial price on the number of iterations required to find the competitive equilibrium. Results will show the modified tâtonnement process can quickly determine the equilibrium price given various initial prices. The second set of experiments investigated the fairness achieved when ABR bandwidth is allocated using the spot market under realistic network dynamics. Simulations will demonstrate ability of the spot market to provide high bandwidth utilization as well as fair allocations under realistic conditions.

4.1 Price Convergence

As previously described, the spot market price is determined using a modified tâtonnement process. Given the demand and supply at a price, the tâtonnement process iteratively adjusts the price until demand equals supply. It has been proven that the equation will reach equilibrium [10], but it is equally important to provide some insight into the speed of convergence. For this reason, multiple simulations were performed to measure convergence (number of iterations required) as the initial price varied. Each experiment simulated a 20 Mbps link (one spot market) and 20 users/NB's. Each user had equal wealth and a random demand (static demand curve), uniformly distributed between 1 and 2 Mbps. The initial price ranged from two orders of magnitude above and below the equilibrium price. Given an initial price, the bandwidth market was allowed to determine the equilibrium price, and the number of iterations required was recorded.

The results are summarized in figure 1, where each point represents the average of 10000 simulations and a 95% confidence interval. The x-axis measures the initial price ratio (the initial price divided by the equilibrium price), while the y-axis indicates the number of iterations required. On average less than six iterations were required when the initial price was two orders of magnitude above (or below) the equilibrium price. As the initial price ratio moved closer to one, the number of iterations required dropped. We have




Figure 1: Impact of the initial price on the number of iterations required to find equilibrium.

noted in other simulations that the price typically remains within the vicinity of one order of magnitude of the equilibrium price, as seen in figures 3 and 4. Therefore, under these conditions the modified tâtonnement process can quickly adjust to market changes.

4.2 Allocation Fairness

In this section the fairness of the network economy is investigated under realistic conditions. Previous microeconomic flow control techniques either do not provide experimental results or simulate limited networks (network size and/or traffic source types). The experiments were conducted on a network configuration consisting of seven links and with 152 users, as seen in figure 2. The network can be described as a “parking lot” configuration, where multiple sources use a primary path. This configuration was agreed upon by members of the ATM Forum [17] as a suitable benchmark for allocation methods; it models substantial competition between users with differing routes and widely-varying propagation delays. User applications were considered 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. Teleconferencing applications, in contrast, can transmit lower voice and video quality. This is primarily due to quantizer control as well as the ability to transmit below the standard 24 or 30 frames per second. MoD users had a budget rate³ of 3×10^8 /sec, while teleconferencing users had a budget rate of 1.5×10^8 /sec. Teleconferencing users are given a lower budget since they are able to scale bandwidth requirements more readily. 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 [22]⁴. Users entered the network at random times

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

⁴Traces can be obtained from the ftp site <ftp-info3.informatik.uni-wuerzburg.de> in the directory /pub/MPEG

uniformly distributed between 0 and 120 seconds. Switches initialized their prices to 50 and α (the target utilization) to 95%. Switches updated their link prices at 10 msec intervals, a compromise between the desire for responsiveness, and the need for stability.

We are interested in the link bandwidth utilization and fairness of the allocation. Allocation graphs are provided to measure the utilization of link bandwidth. The fairness of the allocation is given in the *fairness index* graph, which indicates how far the allocation is from fair [15]. Suppose the allocation among n users is $\{\hat{a}^1, \hat{a}^2, \dots, \hat{a}^n\}$ and the fair allocation (weighted max-min or proportionally fair per unit charge) is $\{a^1, a^2, \dots, a^n\}$. Define the normalized allocation as $\tilde{a}^i = \hat{a}^i/a^i$ for each source, then the fairness index is computed as,

$$\text{fairness index} = \frac{(\sum \tilde{a}^i)^2}{n \sum (\tilde{a}^i)^2}$$

and will be plotted as a function of time. A fairness index of 1.0 indicates a perfectly fair allocation while 0 indicates an unfair distribution. A measurement equal to or greater than 0.99 will be considered fair [15].

Two experiments were performed using either weighted max-min or proportionally fair per unit charge (described in section 2.1) as the targeted fairness. The results for the weighted max-min experiment are given in figure 3. As seen in figure 3(a), the utilization of link 3 (also representative of other links) stayed in the vicinity of 95%, which was the target utilization. The fluctuation around this value indicates users entering and exiting, as well as changing individual demands. This causes the variability of the bandwidth price. The fairness index graph given in figure 3(b), shows the allocation was considered fair for the duration of the simulation. The average fairness index was 0.9991, with a standard deviation of 7.757×10^{-4} . The variation in the fairness index is due to changing demands and the corresponding price response. Similar results were observed for the proportionally fair per unit charge simulation, as seen in figure 4. However, the fairness index (average of 0.9987 with a standard deviation of 8.912×10^{-4}) was lower than the weighted max-min experiment. This difference is due to the proportionally fair allocation equation 4, where the allocation of a user depends on all the link prices instead of just the saturated market (bottleneck link). For this reason, the equilibrium price of the markets are interdependent, which increases the convergence time.

5 Conclusions

Resource management is necessary when contention occurs for limited network resources. In many instances a fair and efficient resource allocation is desired. Pricing is a promising mechanism to control the demand for resources under these conditions. When pricing is applied to network resource allocation, fairness definitions can be taken from traditional network theory (max-min fairness and proportionally fair per unit charge) and microeconomics (Pareto-optimal and social welfare criteria). However, many pricing models and methods






Figure 2: Network topology used in fairness simulations.




(a) Link 3 allocation and prices.




(b) Weighted max-min fairness index.

Figure 3: Results for the weighted max-min fairness simulation.



(a) Link 3 allocation and prices.



(b) Proportionally fair per unit charge fairness index.

Figure 4: Results for the proportionally fair per unit charge simulation.

are not able to achieve all of these fairness goals.

In this paper, we proved an economy consisting of multiple competitive markets is able to achieve all of the network and microeconomic fairness goals. The economy was then applied to the ATM ABR service class, resulting in distributed allocation method that has all of the competitive market advantages (high utilization and a variety of fair allocations) and can quickly adapt to changing network demands (inherent to the ABR service class). There are fewer restrictions on the network than required by other methods based on microeconomics, and behavior during the convergence period is described. Simulation results indicate the ability to achieve fair allocations to a large number of diverse users under dynamic conditions (users entering/exiting and the transmission of actual MPEG-compressed video). This method could be used for other types of congestion control or resource management (for example, TCP flow control and wireless access control).

One disadvantage of the spot market approach is the inability to provide users with guarantees of bandwidth availability. It is possible that a user may enter the network during a period when the prices are low, only to find at a later time prices have increased significantly. This may be acceptable for elastic-traffic that can gracefully adapt to network conditions. For other applications (such as high definition video) these unpredictable price changes can cause QoS to suffer or even force the user to exit the network prematurely. In either case, the economy should provide some protection from these possible market changes. One solution is a second market (reservation market) that offers guaranteed bandwidth over a period of time [13]; however, this does increase the complexity of the pricing method.

A Competitive Market Fairness Proofs

In section 2.1, theorems were introduced regarding the types of optimal and fair allocations achieved by an economy consisting of multiple competitive markets in equilibrium. In this appendix, the associated proofs for these theorems are provided. The notation and assumptions presented in section 2.1 will be used. Supporting definitions and a lemma are presented in this appendix, followed by the proofs for each type of fairness.

As described in section 2.1, consumer j is either “completely satiated” or “non-satiated” with their allocation a^j at market i .

Definition A.1. Completely satiated: At market i with price p^i , consumer j is completely satiated with a^j if the amount of resources affordable is greater than what is desired, b^j .

$$\text{if } \frac{w^j}{p^i} \geq b^j \quad \text{then} \quad a^j = b^j \tag{8}$$

Definition A.2. Non-satiated: At market i with price p^i , consumer j is non-satiated with a^j if the amount of resources affordable is less than or equal to what is desired, b^j .

$$\text{if } \frac{w^j}{p^i} < b^j \quad \text{then } a^j = \frac{w^j}{p^i} \quad (9)$$

Let C^i be the set of completely satiated consumers and N^i be the set of non-satiated consumers at market i ; therefore, $C^i \cup N^i = A^i$ and $\sum_{j \in A^i} a^j \leq s^i$ must always be true for all markets in the economy.

To prove the fairness of the allocation, we must understand the relationship between the allocation of satiated and non-satiated consumers. The following lemma identifies this relationship.

Lemma A.1. *If $[p_*^i, \{a\}]$ is the allocation of competitive market i in equilibrium, then the following is true*

$$\frac{a^j}{w^j} \geq \frac{a^k}{w^k}, \quad \forall j \in N^i, \quad \forall k \in A^i \quad (10)$$

Proof. Assume $[p_*^i, \{a\}]$ is the allocation of competitive market i in equilibrium, and $j, k \in A^i$. Consider two cases, (i) consumer k is non-satiated and, (ii) consumer k is satiated.

Case (i), consumer $k \in N^i$.

From definition A.2, the allocation of non-satiated consumers is,

$$a^j = \frac{w^j}{p_*^i}, \quad a^k = \frac{w^k}{p_*^i} \quad (11)$$

Substituting into equation 10,

$$\frac{\frac{w^j}{p_*^i}}{w^j} \geq \frac{\frac{w^k}{p_*^i}}{w^k} \quad \rightarrow \quad \frac{1}{p_*^i} = \frac{1}{p_*^i} \quad (12)$$

Case (ii), consumer $k \in C^i$.

Assume $[p_*^i, \{a\}]$ is the allocation of competitive market i in equilibrium. Denote $a^j = \max_{j \in C^i} \{a^j/w^j\}$ and $k \in N^i$. Suppose contrary to lemma A.1 that, $\frac{a^j}{w^j} < \frac{a^k}{w^k}$. Substituting for a^j and a^k ,

$$\frac{\frac{w^j}{p_*^i}}{w^j} < \frac{b^k}{w^k} \quad \rightarrow \quad \frac{1}{p_*^i} < \frac{b^k}{w^k} \quad (13)$$

From definition A.1,

$$b^k \leq \frac{w^k}{p_*^i} \quad \rightarrow \quad \frac{b^k}{w^k} \leq \frac{\frac{w^k}{p_*^i}}{w^k} \quad \rightarrow \quad \frac{b^k}{w^k} \leq \frac{1}{p_*^i} \quad (14)$$

Combining equations 13 and 14,

$$\frac{1}{p_*^i} < \frac{b^k}{w^k} \leq \frac{1}{p_*^i} \quad (15)$$

which is not feasible. □

A.1 Weighted Max-Min Fairness Proof

Using the definitions of section 2.1 and appendix A, we can prove theorem 2.2, which states that the resulting allocation of an economy consisting of multiple competitive markets in equilibrium is weighted max-min fair.

Proof. Assume $[\{p_*\}, \{a\}]$ is the allocation of an economy consisting of independent competitive markets in equilibrium. Let $\{\hat{a}\}$ be any other feasible allocation, where $\hat{a}^j = a^j + \delta^j \geq 0$ and $\sum \delta^j = 0$. Only non-satiated consumers may increase their allocation, requiring other consumer(s) to decrease their allocation. Let two consumers j and k participate in market i ($i \in R^j, R^k$). Assume consumer j is a non-satiated and gains resources under $\{\hat{a}\}$ implying $\delta^j > 0$. Denote consumer k as a consumer that loses resources under $\{\hat{a}\}$ implying $\delta^k < 0$. Combining the assumptions above with lemma A.1

$$\hat{a}^j > a^j \quad \text{and} \quad \frac{\hat{a}^k}{w^k} < \frac{a^k}{w^k} \leq \frac{a^j}{w^j} \quad (16)$$

which satisfies the requirement for weighted max-min fairness. □

A.2 Proportionally Fair Per Unit Charge Proof

Using the definitions of section 2.1 and appendix A, we can prove theorem 2.2, which states that the resulting allocation of an economy consisting of multiple competitive markets in equilibrium is proportionally fair per unit charge.

Proof. Let $\{\hat{a}\}$ be any other feasible allocation, where $\hat{a}^j = a^j + \delta^j \geq 0$. Under $\{a\}$ market i must be either fully-utilized $\sum_{\forall j \in A^i} a^j = s^i$ or under-utilized $\sum_{\forall j \in A^i} a^j < s^i$. If market i is fully-utilized, then the following is true

$$\sum_{\forall j \in A^i} \delta^j \leq 0 \quad (17)$$

We know for any under-utilized market i its price p_*^i is zero. Definition 2.5 can be re-written as

$$\sum_{\forall j \in A} \phi^j \cdot \frac{\hat{a}^j - a^j}{a^j} = \sum_{\forall j \in A} \phi^j \cdot \frac{(a^j + \delta^j) - a^j}{a^j} = \sum_{\forall j \in A} \phi^j \cdot \frac{\delta^j}{a^j} \quad (18)$$

Substituting equation 4 for a^j in the denominator, then separating based on markets

$$\sum_{\forall j \in A} \left(\delta^j \cdot \sum_{\forall i \in R^j} p_*^i \right) = \sum_{\forall i \in L} \left(p_*^i \cdot \sum_{\forall j \in A^i} \delta^j \right) \leq 0 \quad (19)$$

This is always true given the conditions for fully-utilized and under-utilized markets. \square

A.3 Proof of an Equitable Allocation

As described in section 2.2, proper wealth distribution is required to achieve an equitable allocation in an economy consisting of multiple competitive markets. For this reason the wealth distribution algorithm is presented, followed by the fairness proof.

Consumer j has utility function $q^j = u^j(a^j)$, that indicates the utility gained q^j from an allocation a^j . The inverse of the utility function, denoted as $\bar{u}^j(q^j)$, indicates an allocation amount a^j that achieves a utility value of q^j . Define the aggregate inverse utility function for all consumers who participate in and consider market i saturated as,

$$\bar{v}^i(\cdot) = \sum_{j \in A^i} \bar{u}^j(\cdot) \quad (20)$$

Since $\bar{u}^j(\cdot)$ is monotonic, $\bar{v}^i(\cdot)$ is monotonic. At equilibrium the supply equals the demand; let q_*^i be the utility value for all consumers at which this occurs, i.e.,

$$s^i = \bar{v}^i(q_*^i) = \sum_{j \in A^i} \bar{u}^j(q_*^i) \quad (21)$$

q_*^i can be found quite easily, since $\bar{u}(\cdot)$ is monotonic. To provide each consumer the same utility level q_*^i when the market is in equilibrium, the wealth of consumer j is set as follows:

$$w^j = \bar{u}^j(q_*^i) \quad (22)$$

The units of w^j are not relevant, just the fact that the wealth of each user is in the proper ratio.

The previous description determined the wealth distribution that achieves an equitable allocation for a single competitive market. Using this as a basis, algorithm 1 determines the wealth distribution that achieves an equitable allocation for an entire economy consisting of multiple independent markets.

Using the definitions of section 2.1 and appendix A, we can prove theorem 2.4, which states that the resulting allocation of an economy consisting of multiple competitive markets in equilibrium is equitable using algorithm 1 for distributing wealth.

Proof. Assume $[\{p_*\}, \{a\}]$ is the allocation of an economy consisting of independent competitive markets in equilibrium, where the wealth of consumers $\{w\}$ was determined from algorithm 1. Let $\{\hat{a}\}$ be any other feasible allocation, where $\hat{a}^j = a^j + \delta^j \geq 0$ and $\sum \delta^j = 0$. Only non-satiated consumers may increase their allocation, requiring other consumers(s) to decrease their allocation. Let two consumers j and k participate in market i ($i \in R^j, R^k$). Assume consumer j is a non-satiated (considers market i saturated) and gains resources under $\{\hat{a}\}$ implying, $\delta^j > 0$ and $u^j(\hat{a}^j) > u^j(a^j)$. Denote consumer k as a consumer that loses resources under $\{\hat{a}\}$ implying, $\delta^k < 0$ and $u^k(\hat{a}^k) < u^k(a^k)$. Combining the assumptions above with lemma A.1

$$u^j(\hat{a}^j) > u^j(a^j) \quad \text{and} \quad u^k(\hat{a}^k) < u^k(a^k) \leq u^j(a^j) \quad (23)$$

which satisfies the requirement for an equitable allocation. □

Algorithm 1 Wealth calculation algorithm for an equitable allocation.

```

1:  /**** variable initialization ****/
2:   $D \leftarrow L$  /* set of markets */
3:  for all  $i \in L$  do
4:     $C^i \leftarrow \emptyset$ 
5:    for all  $j : i \in R^j$  do
6:       $N^i = N^i \cup j$  /* assume all consumers of market  $i$  are non-satiated */
7:    end for
8:  end for
9:  /**** start wealth calculation algorithm *****/
10: while  $D \neq \emptyset$  do
11:    $q_{min} = \infty$ 
12:   for all  $i : i \in D$  do
13:     calculate  $q_*^i$  using consumers in  $N^i$ 
14:     /* determine market with smallest  $q_*^i$  */
15:     if  $q_*^i \leq q_{min}$  then
16:        $q_{min} = q_*^i$ 
17:        $h = i$ 
18:     end if
19:   end for
20:   /* assign wealth to all consumers participating in and non-satiated with market  $h$  */
21:   for all  $j : h \in R^j$  and  $j \in N^h$  do
22:      $w^j = \bar{u}^j(q_*^h)$ 
23:     /* consumer is satiated w.r.t. remaining markets in  $R^j$  */
24:     for all  $i : i \in R^j$  and  $i \neq h$  do
25:        $C^i \leftarrow C^i \cup j$ 
26:        $N^i \leftarrow N^i \setminus j$ 
27:     end for
28:   end for
29:    $D \leftarrow D \setminus h$  /* market  $h$  has been processed, remove from set */
30: end while

```

A.4 Proof of Utility Maximization

As described in section 2.2, proper wealth distribution is required to achieve utilitarian criterion in an economy consisting of multiple competitive markets. For this reason the wealth distribution algorithm is presented, followed by the fairness proof.

Assume the utility function $u^j(a^j)$ for each user j is monotonic and concave. To achieve the utilitarian criterion, competitive market i seeks to maximize the sum of the utilities, given the market supply constraint.

$$q_*^i = \max \left\{ \sum_{\forall j \in A^i} u^j(a^j) \right\}, \quad \sum_{\forall j \in A^i} a^j \leq s^i \quad (24)$$

This maximization problem can be solved using the Lagrangian-multiplier method, since the utility functions are monotonic and concave [20]. Denote $\{a_*\}$ as the solution to the maximization problem. Using equation 9, the wealth of each consumer is $w^j = a_*^j \cdot p_*^i$; however, p_*^i can be dropped since it is a scalar applied to each consumer in the market.

The previous description determined the wealth distribution that achieves the utilitarian criterion for a single competitive market. Using this method as a basis, algorithm 1, can be modified to maximize the sum of the utilities for an entire economy consisting of multiple independent competitive markets. Line 22 of the algorithm changes to $w^j = a_*^j$, while the remaining algorithm operates as defined.

Using the definitions of section 2.1 and appendices A and A.4, we can prove theorem 2.5, which states that the resulting allocation of an economy consisting of multiple competitive markets in equilibrium achieves the utilitarian criterion using a modified algorithm 1 for distributing wealth.

Proof. Assume $[\{p_*\}, \{a_*\}]$ is the allocation of an economy consisting of independent competitive markets in equilibrium, where the wealth of consumers $\{w\}$ was determined from a modified algorithm 1 (modifications described in appendix A.4). If market i is in equilibrium, then the supply equals demand. Therefore,

$$s^i = \sum_{\forall j \in A^i} a^j = \frac{\sum_{\forall j \in A^i} w^j}{p_*^i} \implies p_*^i = \frac{\sum_{\forall j \in A^i} w^j}{s^i} \quad (25)$$

The allocation of user j is

$$a^j = \frac{w^j}{p_*^i} = s^i \cdot \frac{w^j}{\sum_{\forall j \in A^i} w^j} \quad (26)$$

From the algorithm, we know $w^j = a^j$. Substituting into the previous allocation equation

$$a^j = s^i \cdot \frac{a^j}{\sum_{\forall j \in A^i} a^j} = a^j \quad (27)$$

□

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] G. Bianchi, A. T. Campbell, and R. R.-F. Liao. On Utility-Fair Adaptive Services in Wireless Networks. In *Proceedings of the IEEE Sixth International Workshop on Quality of Service*, pages 256 – 267, 1998.
- [5] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss. An Architecture for Differentiated Services. IETF RFC 2475, December 1998.
- [6] C. Courcoubetis, C. Manolakis, and G. D. Stamoulis. An Intelligent Agent for Negotiating QoS in Priced ABR Connections. In *Proceeding of the International Conference on Telecommunications*, 1998.
- [7] C. Courcoubetis and V. A. Siris. An Approach to Pricing and Resource Sharing for Available Bit Rate (ABR) Services. In *Proceedings of the IEEE GLOBECOM*, November 1998.
- [8] C. Courcoubetis, V. A. Siris, and G. D. Stamoulis. Integration of Pricing and Flow Control for Available Bit Rate Services in ATM Networks. In *Proceedings of the IEEE GLOBECOM*, pages 644 – 648, 1996.
- [9] 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.
- [10] E. W. Fulp. *Resource Allocation and Pricing for QoS Management in Computer Networks*. PhD thesis, North Carolina State University, 1999.
- [11] E. W. Fulp, M. Ott, D. Reininger, and D. S. Reeves. Paying for QoS: An Optimal Distributed Algorithm for Pricing Network Resources. In *Proceedings of the IEEE Sixth International Workshop on Quality of Service*, pages 75 – 84, 1998.
- [12] E. W. Fulp and D. S. Reeves. ABR Rate Control for Multimedia Traffic Using Microeconomics. In *Proceedings of the International Conference on ATM*, pages 78 – 84, 1999.
- [13] E. W. Fulp and D. S. Reeves. QoS Rewards and Risks: A Multi-Market Approach to Resource Allocation. In *Proceedings of the IFIP Networking 2000*, 2000.
- [14] J. M. Jaffe. Bottleneck Flow Control. *IEEE Transactions on Communications*, COM-29:954 – 962, July 1981.
- [15] R. Jain. Congestion Control and Traffic Management in ATM Networks: Recent Advances and A Survey. *Computer Networks and ISDN Systems*, Feb. 1995.
- [16] F. Kelly, A. K. Maulloo, and D. K. H. Tan. Rate Control for Communication Networks: Shadow Prices, Proportional Fairness and Stability. *Journal of the Operational Research Society*, 49:237 – 252, 1998.
- [17] A. Kolarov and G. Ramamurthy. Comparison of Congestion Control Schemes for ABR Service in ATM Local Area Networks. In *Proceedings of the IEEE GLOBECOM*, pages 913 – 918, 1994.
- [18] 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.
- [19] J. Murphy, L. Murphy, and E. C. Posner. Distributed Pricing for ATM Networks. *ITC-14*, pages 1053 – 1063, 1994.
- [20] W. Nicholson. *Microeconomic Theory, Basic Principles and Extensions*. The Dryden Press, 1989.
- [21] 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.
- [22] 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.
- [23] S. Shenker. Fundamental Design Issues for the Future Internet. *IEEE Journal on Selected Areas in Communications*, 13(7):1176 – 1188, 1995.
- [24] A. Takayama. *Mathematical Economics*. Cambridge University Press, 1985.
- [25] L. Walras. *Elements of Pure Economics*. Richard D. Irwin, 1954. trans. W. Jaffé.