

Chapter 7

A Multi-Market Approach to Multi-User Allocation

A primary limitation of the spot market approach (described in chapter 6) for multi-user allocation is the inability to provide resource guarantees. 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. These unpredictable price changes can cause the QoS of the user 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, which is the primary goal of the multi-market approach.

In this chapter, a multi-market approach to multi-user resource allocation is presented. Two types of markets exist in the economy, the spot market and the reservation market. The spot market has the advantage of immediate availability of link bandwidth; however, this market has the disadvantage of possible price fluctuations. In contrast, the reservation market has the advantage of resource guarantees (price stability), but incurs the overhead associated with traditional reservation methods. The multi-market approach provides users the unique ability to invest in either market type. In addition, users can dynamically switch from one market to another based on network (economic) conditions, which is a unique feature of this multi-market approach. Simulation results will demonstrate the benefits and risks associated with purchasing bandwidth from the various markets.

7.1 A Multi-Market Economy

This proposed distributed resource allocation method uses two types of independent markets to buy and sell link bandwidth: the reservation market and the spot market. Each market type is based on the competitive market model (described in chapter 5), where pricing is done to promote high utilization as well as Pareto-optimal and fair distributions [15]. As with the spot market approach, there are three entities in this multi-market network economy: users, Network Brokers and switches.

7.1.1 Switch

The network consists of several switches interconnected with links. For a unidirectional link between two switches, consider the sending switch as owner of the bandwidth of that link. For link i denote the total capacity of link bandwidth as s^i . The capacity is then divided into two types, *reserved* and *spot* bandwidth. Reserved bandwidth is sold as an amount for a duration of time, while spot bandwidth is sold as a non-storable resource. With this distinction, reserved and spot bandwidth are considered separate resources. Reserved bandwidth has the unique advantage of ownership over a period of time, while the advantage of spot bandwidth is its immediate availability. Each resource is sold in its own local (link) market, therefore the switch will associate two markets per output port (thus the “multi-market” designation). As with the spot market approach, these markets operate independently and asynchronously since there is no need for market communication (for example, price comparisons) or synchronization from switch to switch. Since the link capacity is divided into reserved and spot bandwidth, the switch must differentiate the traffic using either type. Assume that a bit will be set in the header of the packet, indicating the packet is using reserved bandwidth.

Reserved Bandwidth Market

Link i will sell a maximum percentage β^i of s^i as reserved bandwidth. The reserved bandwidth is divided into equal non-overlapping intervals of time called segments, where the length of each segment is denoted as T^i . Portions of the segment are then sold to users




Figure 7.1: Example segments and price calculation points for link i reserved bandwidth.

with an auction procedure. Users are only able to bid for an amount of the next segment; therefore, the reserved bandwidth of segment l is auctioned during segment $(l - 1)$. At the beginning of the auction for segment l of link i , users forward bids to the switch for an amount of reserved bandwidth. The sum of these bids, denoted as $h_m^{i,l}$, is recorded by the switch and is used to update the price. During the auction, the price of reserved bandwidth for segment l is adjusted at regular intervals G^i , as seen in figure 7.1. The price for reserved bandwidth of link i , segment l is denoted as $g_m^{i,l}$ and is adjusted using the following tâtonnement process,

$$g_{m+1}^{i,l} = g_m^{i,l} \cdot \frac{h_m^{i,l}}{\beta^i \cdot s^i} \quad (7.1)$$

After a new auction price is calculated, it is distributed to NB's, who may submit updated bids. It is important to note the switch does not need to store individual bids. Users can initially submit a bid amount, then send only changes (differences) to the switch. New bidders (users who have not yet participated in bidding for the segment) are not allowed to participate after $T^i \cdot \gamma$ has passed. This provides time for the auction to converge to the equilibrium price before the segment begins (convergence time is addressed in section 6.2). This process repeats until the end of segment $(l - 1)$, after which the switch notifies the users that a new segment has begun. Users are then able to use the amount of reserved bandwidth they defined in their last bid (explicit notification is not necessary). Since only aggregate information (not individual) is used and it is not necessary for the switch to store individual bids, the auction process for the reservation market can be considered a state-less implementation.

There are two important parameters associated with the reservation market; the supply of reserved bandwidth ($\beta^i \cdot s^i$) and the segment length (T^i). The percentage of capacity for the reservation market can be based upon various goals, such as maximizing utilization, QoS, or profits. Furthermore, β^i can be statically or dynamically set. For example, the switch could adjust β^i to maximize total revenue of that link by measuring the demand and the prices for each market. Similarly, segment length can be set to achieve several goals. Longer segment lengths provide more price stability, however this lengthens the time between auctions (and consequently the wait time to use reserved bandwidth for the next segment). Shorter segment lengths reduce the wait time, but reduce the stability of the reservation market as well. An alternative approach would partition the reserved bandwidth of a switch into smaller divisions, where each division would represent a reservation sub-market. These sub-markets could have various segment lengths and different auction start-times. This would offer users more choices of reserved bandwidth (different reservation times) and reduce the waiting time for the next segment. The effect of segment length on QoS is discussed further in section 7.3.2.

Bandwidth Spot Market

The bandwidth spot market operates as described in chapter 6. Denote the amount of reserved bandwidth currently used as u^i . The supply of spot bandwidth for link i is $s^i - u^i$; therefore any reserved bandwidth that is not used can be sold in the spot market. As done in section 6.1, the spot market price is determined at regular intervals using the modified tâtonnement process,

$$p_{n+1}^i = p_n^i \cdot \frac{d_n^i}{\alpha^i \cdot s^i - u_n^i} \quad (7.2)$$

Once a new bandwidth spot price is determined, it is forwarded to users of this link. Users are also required to maintain a smooth transmission at their allowed rates (also required by [73]). Upon receiving a new price, users determine their allowable transmission rate.

Similarities of the Multi-Market and SENET

The goal of the multi-market economy is to provide guarantees and immediate availability of bandwidth in one resource allocation method. An earlier, non-economic based, bandwidth allocation method that sought a similar goal was the Slotted Envelope Network (SENET) [6, 26]. SENET is an integration of circuit and packet switching that transmits data in frames. Each frame consists of two compartments. The first compartment contains circuit switched traffic while the second compartment contains packet switched traffic. The boundary between the two compartments can be fixed or movable, depending on the implementation.

A circuit switched connection is assigned a slot in the first compartment large enough to handle the desired bit rate. Each circuit switched connection can hold its assigned slot for the duration of the session, which provides guaranteed service. The size of the first compartment is bounded by a value less than the frame size; therefore, some space is available for packet switched traffic. An arriving circuit switched session is blocked if the residual capacity of the first compartment cannot accommodate the requested bit rate. The second compartment is used for transmission of packets waiting in a queue. In the movable boundary version, packet switched traffic can occupy the second compartment bandwidth plus any unused circuit switched capacity. This provides high bandwidth utilization [64].


From this description, there are some similarities between the multi-market approach and SENET. The reserved bandwidth market is similar to the first compartment of SENET. Both provide a guarantee of bandwidth; however, the reserved market provides the guarantee only for the duration of the segment. This ensures users will eventually conform to market changes, preventing them from holding bandwidth indefinitely. The second compartment of SENET is similar to the spot market. No guarantees are provided for packet switched traffic in SENET or in the spot market, but high utilization of bandwidth is achieved in both methods. An advantage of the multi-market approach is the flexibility users have to move from one market type to the other based on QoS and/or prices. In addition, pricing in the multi-market approach (as well as other microeconomic-based methods) creates a disincentive for over-allocating bandwidth.

7.1.2 User

User j , executing a network application, requires link bandwidth for transmission. The amount of bandwidth desired is determined from the application and is denoted as b^j . Based on prices and wealth, W^j , user j can afford a range of bandwidth (less than or equal to b^j), where some amounts will be preferred over others. These preferences are represented with a QoS profile (described in section 6.1.2). Since there are two different types of resources in the economy (reserved and spot bandwidth), the user must identify how the reserved and spot bandwidth may be substituted for one another. In microeconomics, this is represented with an indifference curve.

An indifference curve indicates the combinations of resources that result in the same utility [15, 76]. For example in figure 7.2(a), the allocation (x_1, y_1) results in the same utility as the allocation (x_2, y_2) . The curve between these points is the indifference curve. The slope of the indifference curve is called the marginal rate of substitution and indicates the rate at which a user trades one resource for another. In figure 7.2(a), only one indifference curve was depicted; however, the graph actually contains an infinite number of these curves, where each curve indicates a level of utility. In figure 7.2(b), multiple indifference curves are drawn in the same graph. Typically the level of utility increases as you move towards the north-east (more is better). Given this assumption, the indifference curves given in figure 7.2(b) can be ranked based on utility as $I_3 > I_2 > I_1$.

For the multi-market economy, indifference curves indicate the combination of reserved and spot bandwidth that result in the same utility. These curves are normalized to the current desired amount of bandwidth b^j , as seen in figure 7.3. In this figure the indifference curve labeled “prefer-reserved” was generated from the equation $y^j = k \cdot (1 - x^j)^2$, where y^j is the amount of spot bandwidth and x^j is the amount of reserved bandwidth for user j . The preference for reserved bandwidth increases as k increases, which is reflected in the marginal rate of substitution. The other indifference curve given in figure 7.3 represents a user who has no preference for reserved or spot bandwidth. In this case the user will always prefer the *cheaper* (lower cost) bandwidth. In the case where the two types of bandwidth have the same price, assume the user will prefer equal amounts of both types



(a) A single indifference curve, where resource combinations (x_1, y_1) and (x_2, y_2) yield the same utility.

(b) An indifference map, where $I_3 > I_2 > I_1$.

Figure 7.2: Example indifference curve and map.

of bandwidth. The only assumptions required for the indifference curve is that it must be continuously differentiable and convex to the origin (required for determining the amount of bandwidth to purchase).

Finally, the user is charged continuously for the duration of the session (analogous to a meter). To pay for the expenses, assume the user provides an equal amount of money over regular periods of time (as done in 6.1.2).

7.1.3 Network Broker

Users can only participate in 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. Representing the user in the economy the NB performs the following tasks: connection admission control, policing, packet marking, and purchasing/bidding decisions. Although the NB works as an agent for the user (making purchasing decisions), assume the NB operates honestly in regards to the switches and the user.

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, indifference curve, b^j

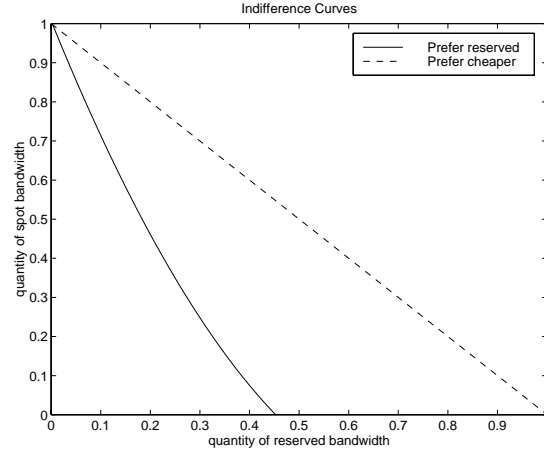


Figure 7.3: Example indifference curves, prefer-reserved and prefer-cheaper bandwidth.

and W^j . The NB also stores the route of user j , R^j , that connects source to the destination. For each link on R^j , the NB collects current reserved and spot bandwidth prices. The NB will divide W^j into separate budget rates, one for each link in the route (as described in section 6.1.3). Using this information 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. For example, the NB/user could require a minimum amount of reserved bandwidth be purchased before transmission begins. The NB also levies the user for their consumption. In addition, the NB polices the user, ensuring only the bandwidth purchased is used, and marks the packets (assigning which are to use reserved bandwidth). Finally, the NB determines the reserved bandwidth bid and the amount of spot bandwidth to purchase.

Bidding and Purchasing Bandwidth

Since reserved bandwidth is sold in an auction format, the NB must bid for reserved bandwidth at each link on the route. The bid (the amount the user will purchase) for link i is based on the budget w^j , a statistic of the spot market price, and the reserved bandwidth auction price for this link $g^{i,l}$. A statistic is required for the spot market bandwidth price due to its volatility. For this discussion (and simulation) the maximum spot price \hat{p}^i ,

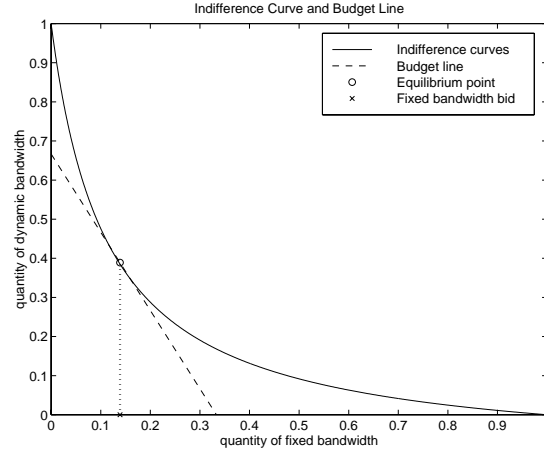


Figure 7.4: Example equilibrium point, where slopes of the the indifference curve and budget line are equal.

measured during the current segment, is used as the spot market price statistic. Using this information the NB maximizes the utility of the user $u^j(x^j, y^j)$,

$$\max \{u^j(x^j, y^j)\}, \quad g^{i,l} \cdot x^j + \hat{p}^i \cdot y^j \leq w^j \quad (7.3)$$

where x^j is the amount of reserved bandwidth and y^j is the amount of spot bandwidth of user j . As defined in [15], the first order condition of this constrained maximization problem is,

$$\frac{u^j(x^j)}{u^j(y^j)} = \frac{g^{i,l}}{\hat{p}^i} \quad (7.4)$$

The NB must spend the budget to equalize the ratio of marginal utility to the price of each resource. Assuming the indifference curve is convex to the origin, the derivative of the curve is,

$$\frac{dy}{dx} = -\frac{u^j(x^j)}{u^j(y^j)} \quad (7.5)$$

Plotting the budget line with the indifference curve, as seen in figure 7.4, the slope of the budget line is $g^{i,l}/\hat{p}^i$. Therefore, the point where the slope of the indifference curve and the budget line are equal is where the utility of the user is maximized. The second order condition is not required due to the convexity assumption of the indifference curve [15]. The reserved bandwidth bid is the x^j component of this point and is forwarded to link i .

The NB keeps a table storing the amount and price of reserved bandwidth purchased at each link in the route. When the segment for link i is sold, the amount user j purchased $c^{j,i}$ and the price f^i is updated in the table. The maximum amount of reserved bandwidth that can be used by user j is,

$$e^j = \min_{\forall i \in R^j} \{c^{j,i}\} \quad (7.6)$$

which is the minimum amount of reserved bandwidth purchased at any link. If the desired bandwidth is greater than the purchased reserved bandwidth, then spot bandwidth is used for transmitting the remaining portion $(b^j - e^j)$. The amount of spot bandwidth to use y^j by user j is,

$$y^j = \min \left\{ \min_{\forall i \in R^j} \left\{ \frac{w^j - e^j \cdot f^i}{p_n^i} \right\}, b^j - e^j \right\} \quad (7.7)$$

which is the maximum amount of spot bandwidth that is affordable, but no more than what is required $(b^j - e^j)$.

7.1.4 Equilibrium Surfaces

The bidding procedure described in section 7.1.3 uses the indifference curve and the budget line to determine the reservation bid amount. The bid is the x-component of the equilibrium point, where the slope of the budget line equals the slope of the indifference curve. At this point the user achieves the highest utility based on their budget constraint and indifference curve. However, it is difficult to visualize how the shape of the indifference curve impacts the bidding procedure.

To provide another perspective on the bidding procedure, the equilibrium points for all reservation and spot prices (normalized) can be graphed. This yields two equilibrium surfaces, one representing the amount of reserved bandwidth, the other surface representing the amount of spot bandwidth (both normalized). On these graphs, the x-axis measures the reservation price, the y-axis measures the spot price, and the z-axis indicates the amount (reserved or spot, depending on the surface).

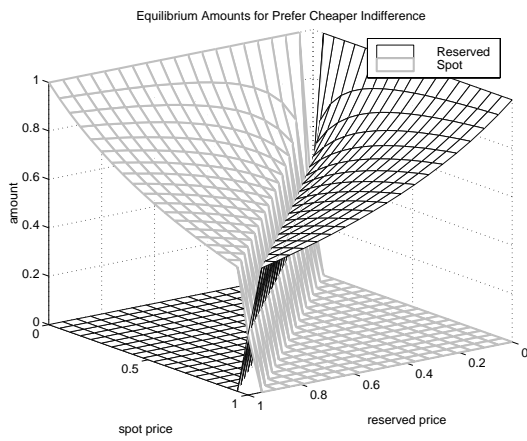
Given a constant amount of wealth, the graphs given in figure 7.5 display the equilibrium surfaces of three different indifference curves: prefer-reserved bandwidth, prefer-

cheaper and prefer-equal (equal amounts of reserved and spot bandwidth). The prefer-cheaper surface, figure 7.5(a), indicates that the user purchases only the cheaper resource. When the price for spot and reserved are equal, the user purchases equal amounts of both. In contrast, the surface for prefer-reserved, figure 7.5(b), indicates that the user continues to purchase only reserved bandwidth when the price for spot bandwidth is cheaper. The total amount of bandwidth (spot and reserved) for prefer-cheaper and prefer-reserved is given in figure 7.5(d). In this graph the prefer-reserved user purchases less (total bandwidth) than the prefer-cheaper user when the spot price is less. This is the “penalty” the prefer-reserved user faces. Therefore, the prefer-cheaper indifference curve always maximizes the amount of bandwidth purchased.

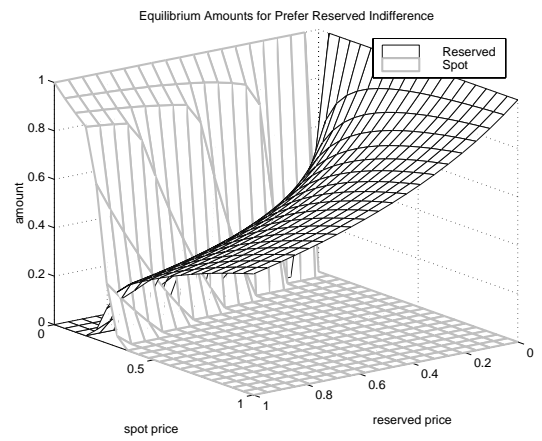
7.2 Network Dynamics and Optimality

As described in section 6.1, once the tâtonnement process reaches equilibrium the resulting allocation is Pareto-optimal and fair. Proofs that the resulting allocation of competitive markets (for the entire network) in equilibrium is Pareto-optimal and fair are given in chapter 5. Users are allocated bandwidth based on their budget as well as the competition they face in their markets. This results in a weighted max-min fair allocation.

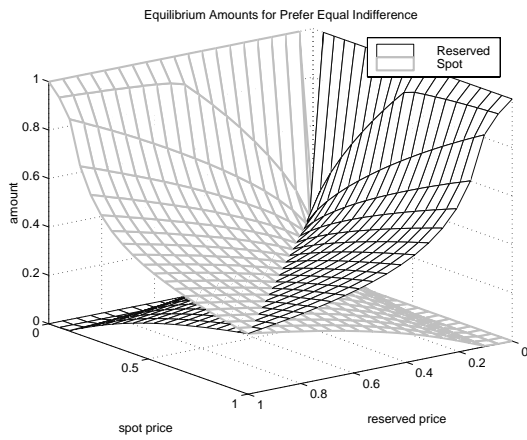
The multi-market model consists of two different types of competitive markets (spot and reservation) and can be view as two separate sub-economies. Collectively, the spot markets can be considered one sub-economy, while the reservation markets create the other. For either sub-economy the proofs of Pareto-optimal and fair allocations, given in chapter 5, are applicable. Therefore, when the spot markets are in equilibrium, their allocation is is Pareto-optimal and fair. Similarly, when the reservation markets are in equilibrium the allocation of reserved bandwidth is Pareto-optimal and fair. A user will receive a fair allocation of spot bandwidth based on the competition in the spot markets as well as the portion of their budget spent in this sub-economy. A user will also receive a fair allocation of reserved bandwidth based on the competition in the reservation markets as well as the portion of their budget spent in this sub-economy. The budget is divided between in the two market types based on the indifference curve of the user as well as the




(a) Prefer cheaper bandwidth.



(b) Prefer reserved bandwidth.



(c) Prefer equal amounts.



(d) Total bandwidth requested for prefer-reserved and prefer-cheaper.

Figure 7.5: Equilibrium allocation amounts for various types of indifference curves.

current spot and reservation prices, as described in section 7.1.3.

A single equilibrium price will not exist due to the changing source demands (VBR source and users entering/exiting the network), as described in section 6.3. However, the market can be viewed as having multiple equilibrium prices, each for some segment of time. During a segment the pricing technique will seek the equilibrium price as described in section 6.2. Once this price is found, the resulting distribution is Pareto-optimal and fair. When the aggregate demand changes, the stability of the price equation ensures that the price always moves towards p_* . Simulation results are provided in the next section to demonstrate the performance of the multi-market approach in this environment.

7.3 Multi-Market Experimental Results

Simulations are used in this section to demonstrate the performance of the multi-market economy under network dynamics (changing user demands and users entering/exiting the network). Experiments will consist of a realistic network configuration, allowing users to randomly enter the network and use actual MPEG-compressed traffic. The first experiment will provide an example of the benefits of purchasing different amounts in the reservation and spot markets. The second experiment will investigate the effect of the reservation market segment length on QoS.

7.3.1 A Demonstration of the Multi-Market Approach

This experiment will provide an example of the QoS obtained by users with different bandwidth preferences. The network simulated consisted of 160 users and their associated NB's, three switches and seven primary links, as seen in figure 7.6. Each output port carried traffic from 40 users and connected to a 45 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 consisting of one, two or three hops. 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 [46] for allocation comparisons since it provides competition among users with different routes and various

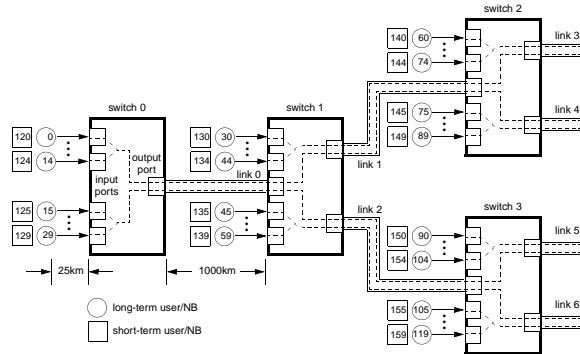


Figure 7.6: Network configuration used in multi-market simulations.

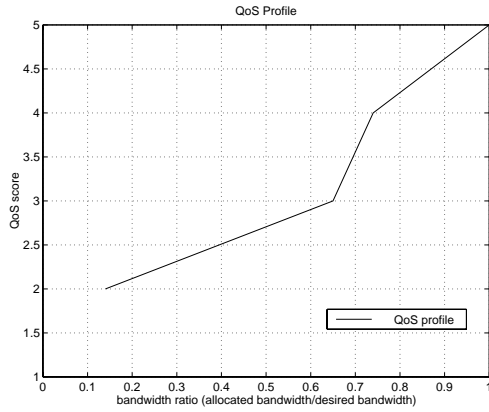
propagation delays. The multi-market economy had the following initial values. The spot market parameter α (targeted utilization) was 90%. Switches sold equal amounts of reserved and spot bandwidth; therefore β was 45%. Reserved bandwidth prices were initialized to 55 and segments were 15 minutes in duration. Longer segments could have been selected; however, transition effect from one segment to another is of interest. Spot market prices were initialized to 50. There was no propagation delay between the user and their NB.

Users had a budget rate, W , of 3×10^8 /sec and used the QoS profile given in figure 7.7(a). Since all users have the same budget rate, they are considered equal (purchasing power) in the economy. The source for each user was one of 15 MPEG-compressed traces, as described in section 3.3.2. Although users have the same wealth and QoS profile, for this demonstration users are considered either *long-term* or *short-term*. Long-term users measure, over the duration of the simulation, the different QoS obtained from preferring different amounts of reserved and spot bandwidth. Short-term users are introduced in the economy to cause sudden demand shifts, which may occur in actual networks (peak load times). Together, how the various users impact the QoS achieved in the multi-market economy is of interest.

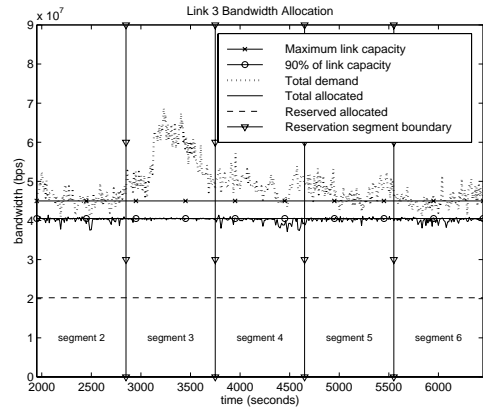
A total of 120 users were considered long-term and had sessions that lasted the duration of the simulation. Half of the long-term users preferred reserved bandwidth, the

remaining preferred cheaper bandwidth (indifference curves given in 7.3). The long-term users entered the network at random times uniformly distributed between 0 and 600 seconds. The remaining 40 users were considered short-term. These users transmitted a short segment of an MPEG video (under 3 minutes, randomly determined). Due to the relative shortness of their session, these users only purchased bandwidth from the spot market. Starting at 3000 seconds the short-term users entered the network with a Poisson distribution of mean 120 seconds. The link bandwidth utilization and the QoS provided to each type of long-term user is of interest. Allocation graphs are provided to measure the utilization of link bandwidth, while QoS graphs measure the average QoS observed by long-term users.

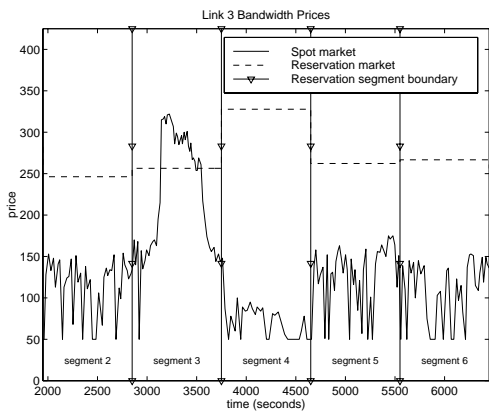
For this simulation, the example bandwidth allocations, prices and average QoS are given in figures 7.7(b) - 7.7(d). Only the results from 1800 to 6500 seconds are displayed, since the effect the short-term users have on the economy is of interest. As seen in figure 7.7(b), all of the available reserved bandwidth for link 3 was sold, while the total bandwidth allocated stayed within the vicinity of α (targeted utilization). Similar results were noted for the remaining links. As seen in figure 7.7(d), before the short-term users entered the network (time less than 3000 seconds) prefer-cheaper users enjoyed a higher QoS. During this time, prefer-cheaper users only purchased bandwidth from the spot markets, since prices were lower (as seen in figure 7.7(c)). In contrast, prefer-reserved users spent their entire budget in the reservation markets. Purchasing from the spot markets yielded higher QoS for the prefer-cheaper users, because these users only purchased what they needed at any time. This allows users to efficiently share the spot market bandwidth. In the case where the total demand of the prefer-cheaper users exceeded the spot market supply, each prefer-cheaper user received an equal-share (weighted max-min fair) of the spot market bandwidth. The prefer-reserved users were allocated an equal-share of the reserved bandwidth supply for the duration of the simulation (no more). For this reason, prefer-reserved users observed a lower QoS, until the short-term users arrived. When the short-term users arrived (during segment 3 of figure 7.7(b)), the spot market price increased in response to the increase in demand. During this time prefer-cheaper users received a lower QoS, since they had to compete with the new arrivals. In contrast, prefer-reserved users continued to receive approximately the



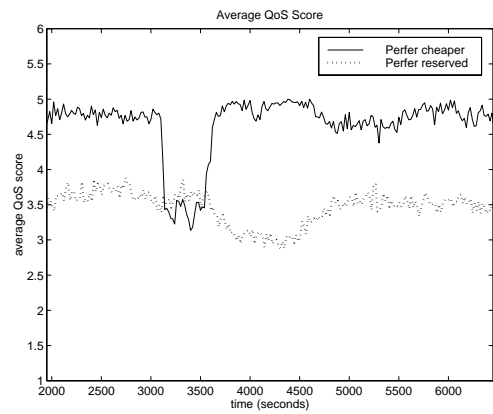
(a) QoS profile.



(b) Link 3 allocation.



(c) Link 3 bandwidth prices.



(d) Average QoS scores for all long-term users.

Figure 7.7: Multi-market simulation QoS profile and results.

same level of QoS since they purchased reserved bandwidth for segment 3. The spot market price increase, during segment 3, did cause a higher reservation market price for segment 4, as seen in figure 7.7(b). The prefer-reserved users could afford less reserved bandwidth during this segment, resulting in a slightly lower QoS. Afterwards, prices and QoS observed by the users returns to the previous values.

This simulation provides some insight to the rewards and risks of purchasing various amounts of bandwidth in the spot and reservation markets. In the example, prefer-cheaper users enjoyed the reward of a higher QoS, but were susceptible to the risk of spot market price fluctuations, that can cause large changes in their QoS. In contrast, prefer-reserved users opted for the lower risk reservation market, but generally received a lower QoS.


7.3.2 The Effect of Varying the Segment Length on QoS

In this experiment the effect of varying the segment length (of the reservation market) on QoS is considered. Similar to the experimental setup of section 7.3.1, the QoS of long-term users was measured for different segment lengths. The simulation consisted of a single 12 Mbps link and 12 users. Each user transmitted a MPEG-compressed video and had a wealth of 3×10^8 /sec. A total of 10 users were considered long-term. Half of the long-term users were prefer-cheaper, while the remaining long-term users were prefer-reserved (section 7.3.1 describes users and their preferences). Long-term users started their sessions at random times, uniformly distributed between 0 and 10 seconds. The remaining 2 users were short-term and started their sessions at 2020 seconds. The multi-market economy had the following initial values. The spot market parameter α (targeted utilization) was 90%. Switches sold equal amounts of reserved and spot bandwidth; therefore β was 45%. Reserved bandwidth prices were initialized to 55. For the first experiment, segments were 50 seconds in duration. This was increased by 50 seconds for each subsequent experiment, until segments were 400 seconds in duration. For each experiment, the link bandwidth utilization and the QoS provided to each type of long-term user is of interest. Allocation graphs are provided to measure the utilization of link bandwidth, while QoS graphs measure the average QoS observed by long-term users.

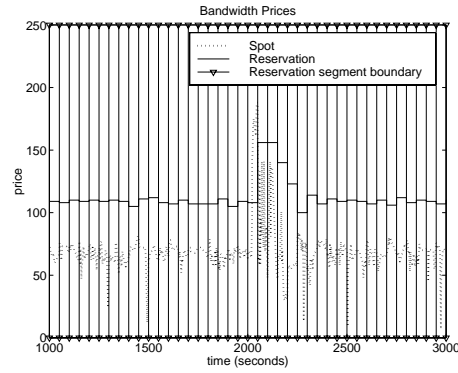
Figures 7.8 - 7.11 show the allocation, bandwidth prices and QoS scores for segment lengths of 50, 100, 300 and 400 seconds. The presence of short-term users (starting at 2020 seconds and ending at 2180 seconds) increases the demand as well as the spot market price, which is evident in each of the experiments. The spot market price increase also impacts the reservation market price; however, this change occurs in the next segment. At this time prefer-reserved users have a slightly lower QoS score due to the higher reservation market prices. In contrast, prefer-cheaper users have a slightly higher QoS due to the lower spot market prices. This typically lasted for one segment; however, the worst case (for these experiments) is given in figure 7.10. In this experiment, short-term users enter the network at the end of one segment and leave at the start of the next segment. This causes the reservation market price to remain “high” for two segments, as seen in figure 7.10(b). This indicates the need for a *better* spot market price statistic (other than the maximum spot market price) for determining the reservation market bid. For example, a better spot market price statistic would account for the duration of the segment as well as the duration of any spot market price change.

7.4 Chapter Summary

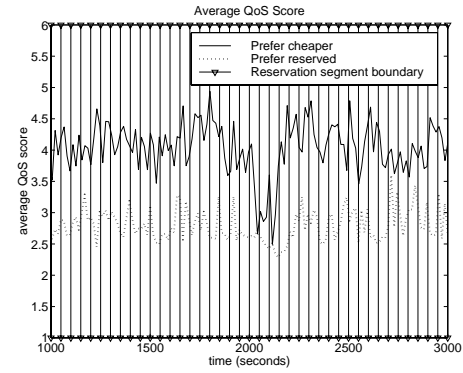
This chapter introduced a distributed, multi-user bandwidth allocation method based on a multi-market economy. A computer network can be viewed as an economy consisting of three entities (users, Network Brokers and switches) and two different markets/resources (reserved and spot bandwidth). Switches own the bandwidth, which is sold in the reservation and spot markets. Reserved bandwidth has the advantage of ownership over a period of time, providing the user with some predictability of their expected QoS. In contrast, spot bandwidth has the advantage of immediate availability without the reservation overhead. Both market types are modeled as competitive markets; therefore efficient as well as Pareto-optimal and fair allocations are possible. Users require link bandwidth for their applications and are represented in the economy by a Network Broker (NB). The NB buys bandwidth to maximize the utility (QoS) of the user and considers the risks and benefits associated with the two bandwidth types (demonstrated in the simulation



(a) Bandwidth allocation.

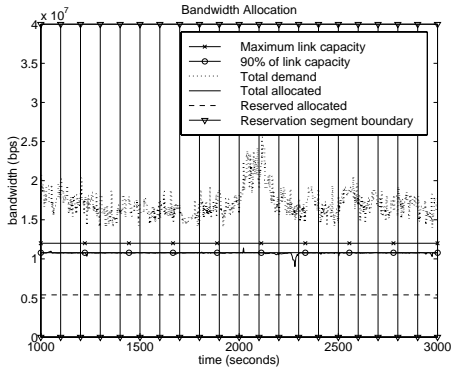


(b) Bandwidth prices.

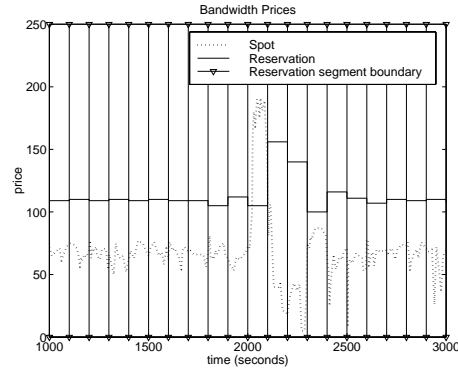


(c) Average QoS score.


Figure 7.8: Allocation, prices and QoS scores for segment length of 50 seconds.



(a) Bandwidth allocation.

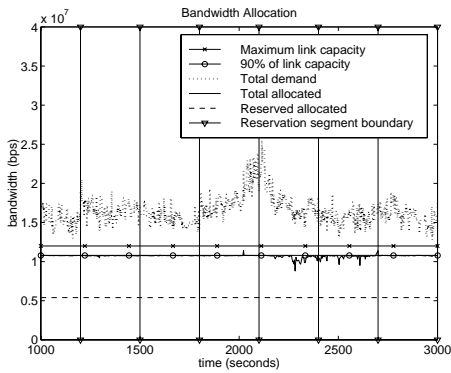


(b) Bandwidth prices.

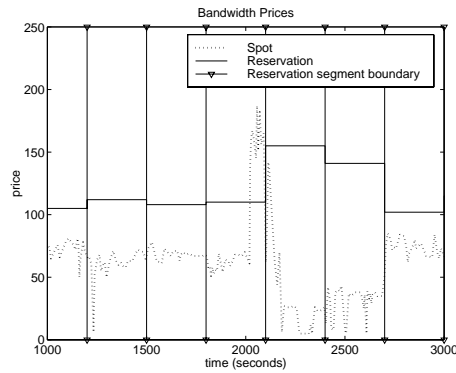


(c) Average QoS scores.


Figure 7.9: Allocation, prices and QoS scores for segment length of 100 seconds.



(a) Bandwidth allocation.




(b) Bandwidth prices.

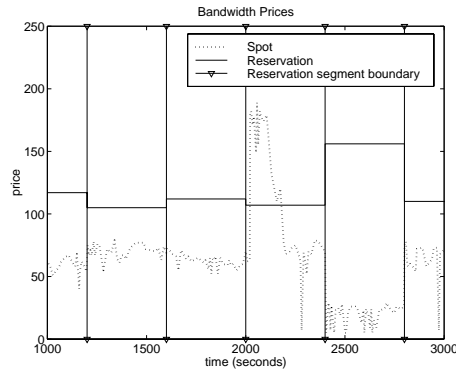


(c) Average QoS score.

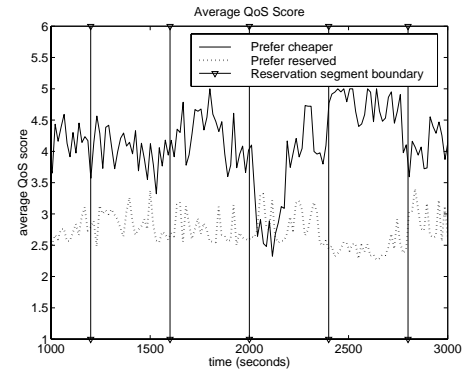
Figure 7.10: Allocation, prices and QoS scores for segment length of 300 seconds.



(a) Bandwidth allocation.



(b) Bandwidth prices.



(c) Average QoS scores.

Figure 7.11: Allocation, prices and QoS scores for segment length of 400 seconds.

results). This multi-market approach uniquely integrates the benefits of the spot market (such as Pareto-optimal and equitable allocations) with the price stability offered with the reservation market. This is done in a distributed and state-less manner. Unique to the multi-market approach is the ability of users to dynamically change bandwidth amounts in response to market and source requirements. This provides the user greater flexibility when maximizing their QoS.