Market-Based Local Wireless Bandwidth Allocation Using Customer Impact Valuations

Christopher J. Hazard
North Carolina State University
Raleigh, NC 27695
cjhazard@ncsu.edu

Peter R. Wurman
North Carolina State University
Raleigh, NC 27695
wurman@ncsu.edu

ABSTRACT

Finding the best methods of allocating transmission rights on the electromagnetic spectrum is a challenging problem. To better allocate the spectrum to increasing demand, the FCC has recently been deregulating spectrum trading, allowing for more open markets. Additionally, current findings indicate that customer loyalty to wireless service providers is most heavily influenced by service satisfaction. In this paper, we present a market-based bandwidth allocation model which performs short term leases between service providers. We also present a valuation model for bandwidth for use in these markets based on expectations of customer satisfaction. Customer satisfaction is modeled as the threshold of the number of blocked channel requests (BCR) that will quantifiably impact the customers' perception of the service. Our model is shown to scale well and increase service providers' utility as measured by lease revenue and customer satisfaction.

1. INTRODUCTION

Determining the best methods of allocating transmission rights on the electromagnetic spectrum is an old and difficult problem. Spectrum transmissions entail continuous ranges of frequency, transmission power, location, direction, and time. Different groups of users of the spectrum have different goals, including long range broadcasts, short range signaling, wireless Internet, cellular telephony, and emergency communications. These users' goals are met by integrating a wide variety of devices such as towers, satellites, and hand held devices. Proper allocation of the spectrum is important to prevent devices from interfering with each other in many applications such as long-range communication. On the other hand, regulating the spectrum as a public good has worked well for short-range household devices.

If the problem of spectrum allocation was not difficult enough, transmission rights have been marred by problematic historical decisions that still remain [4]. For example, until 1993 in the USA, the Federal Communications Commission (FCC) assigned spectrum based on a political process, which was not always optimal. Since 1993, the FCC has been auctioning unused portions of the spectrum. However, much of the current spectrum is underutilized due to spectrum allocation choices from as far back as the late 1920's. Economists have been arguing for market-based approaches to spectrum allocations. Most relevant to our work is the FCC's 2003 decision to permit long-term spectrum leasing [5]. While current regulations require sufficient paperwork and lead time for leases, we suggests further loosening of these restrictions. For comprehensive treatment of the arguments for market-based spectrum allocation and the relevant history in the USA, we refer the reader to [4]. Though mentioned in the context of the FCC, we feel that our work is applicable to other countries as well.

In this paper, we focus on efficiently allocating spectrum to service providers (carriers) in the same location and similar transmission power (a cell). The services all utilize multiples of a common quantum of bandwidth, which we will refer to a channel. This is most readily applicable to wireless telephony, where each voice call consumes a basic channel, and data calls consume a basic channel plus dynamically allocated supplementary channels. It is also applicable to Wi-Fi, where each connection is one channel. We use Poisson arrival and departure processes to approximate future unknown usage models from wireless Internet service providers.

The model we present here is intended to be a replacement for current roaming practices. Roaming practices in the USA involve agreements between providers to allow customers of one provider to use the network of the other provider. While these agreements often provide mutual benefit to both providers, large providers can use these agreements anti-competitively to prevent new entrants to the market. While an open market for trading bandwidth is not without its concerns, we argue that it is better than the current practices for several reasons. First, because usage of each slice of the spectrum (or code with respect to CDMA) is directly owned or leased, each owner or leaser can be held accountable for its transmissions. Second, the market forces incumbents to compete with new entrants on price rather than secretive reciprocal contracts. Incumbents will instead be forced to lease their bandwidth on an open market in order to realize the gains of sharing access to incumbents' networks. Third, very short term leases will maintain efficient utilization of the spectrum.

Current findings indicate that customer satisfaction is the primary influence in customer loyalty in the domain of wireless telephony [9]. Though quality of service (QoS) and other factors contribute to customer satisfaction, one of the primary factors is blocked channel requests (BCRs). With respect to QoS, BCRs occur when a level of service cannot be allocated. We derive the service provider bandwidth valuations from the negative impact of BCRs. While we use this valuation model within our market model, the two models are independent.

A further constraint is that, in order to have an open market between service providers, their underlying technology must be compatible. FDMA (frequency division multiple access) makes bandwidth trading simply trading slices of the spectrum. However, with modern technologies, the commodity being traded becomes more complicated. For example, TDMA (time division multiple access) and FDMA are both used by the GSM standard, where customers' transmissions are divided up between different time slots and frequencies. In this scenario, the service providers would trade in one channel's worth of time slices. In the spread-spectrum technology CDMA (code division multiple access), the traded commodity would be the codes representing channels. Though it may not be currently possible to trade bandwidth in small allotments with differing technologies in our model, future unification of transmission standards and greater demand for bandwidth will lessen this constraint. Due to the aforementioned complications in relating spectrum to actual communications channels and to keep our arguments generic, we will refer to what is being traded as bandwidth, rather than spectrum.

The rest of this paper is organized as follows. In Section 2, we compare and contrast other work related to our model. In Section 3, we present our model, first in terms of the bandwidth market, then in terms of bandwidth valuations. We simulated this model, and the results are covered in Section 4. Finally, in Section 5 we draw conclusions from our findings.

2. RELATED WORK

The topic of utilizing markets to control the allocation of network traffic has been explored in many ways, but most often in terms of QoS. [17] developed a QoS market to optimize allocation based on user preferences and services. [15] investigated similar QoS markets with a multi-level market, and [6] explored market bandwidth management in backbone Internet service providers.

QoS typically works on the level of individual packets to ensure data throughput rates for a communication session, whereas our model focuses on the level of communication channels to ensure service providers have the bandwidth their customers require. [7] also focused on communication channels in implementing an (M+1)st price auction between customers competing for connections with different QoS. However, Ibrahim's work differs from the scope of our work as it does not consider multiple service providers.

Though market-based resource management obviously works well in real-world monetary systems, it has been deployed in relatively few automated network environments. [12] deployed such a system for managing time on a wireless sensor network testbed. They concluded that the market yielded better results than traditional proportional allocation or batch scheduling, despite strategic behaviors demonstrated by its users.

[10] explore the broader economic interrelationships in bandwidth pricing with respect to common bandwidth pricing techniques, but do so on a broad timescale with a macroeconomic scope.

Technical issues of secondary usage of the electromagnetic spectrum, which would allow parties to lease or otherwise license unused blocks of the spectrum, are explored by [18]. [3] suggests that in order for a secondary market to succeed, it will need to favor market liquidity by reducing the barriers to trade, and utilize an electronic call market. With respect to Bykowsky's suggestions, our work assumes the former and implements the latter.

Peha and Panichpapiboon make the case for allowing secondary devices to directly negotiate their spectrum usage [14]. Spot markets allow participants to directly and instantly trade resources for money. While their method is efficient in utilizing spectrum resources, it complicates billing and payments, as well as makes accounting and tracking accountability more difficult.

[20] extend Peha and Panichpapiboon's work by presenting an architecture to manage authentication in secondary wireless markets. Their model, like that of [14], uses spot markets which allow secondary users to directly negotiate spectrum for their own use with the providers. This is in contrast to our model, which is based on a secondary market of primary users at regular intervals. Zhou et al. also derive an optimal clearing price for bandwidth trading in terms of the Erlang Loss Formula (also known as Erlang B). One of the primary limitations of their cost model is that it assumes the relationship between utility and blocking probability is linear. Additionally, the Erlang Loss Formula does not measure or account for multiple BCRs given to the same customer. Our model is not limited in these regards.

[1] model the blocking probability in terms of call arrival rate, and present two algorithms to allocate calls to nearby base-stations. Their model is limited, however, in the sense that it only considers call arrival rate, making it inaccurate for circuit-based switching, and the load balancing between base-stations does not consider different service providers.

With regard to spectrum allocation, [13] argues that the FCC should maintain ownership and continually lease the spectrum. Our proposal is similar to Noam's model in that we require the spectrum to be leased or owned in order to use it. However, our proposal differs in that spectrum ownership is maintained as it is currently, and on the clearing mechanisms. The customer valuation model and market details we present would be applicable to Noam's proposed spectrum management, but may need to be adapted depending on the market clearing mechanisms employed.

3. MODEL

3.1 Customer Model

The customers are situated in a set of tessellated hexagons, as is typically used to represent wireless telephony cells. Each customer belongs to exactly one service provider, and may only make connections using that provider's resources. Every cell has a fixed amount of bandwidth, with each provider owning a fixed portion of the bandwidth. Providers may lease bandwidth to any other provider, but the bandwidth cannot move from one cell to another.

Customers originate connections, release connections, and move from cell to cell, all at specified rates with exponential distributions. Every connection requires one channel of bandwidth. When a customer moves from one cell to another with a connection in progress, this is called a handoff. During a handoff, the channel is acquired in the destination cell and released in the source cell. Once the number of channel requests (handoffs and connection originations) is equal to the provider's available bandwidth, any subsequent channel requests will be blocked until channels are free again. An illustrative example of this model is depicted in Figure 1, where Provider C in the lower right cell will block 1 connection.

3.2 Bandwidth Market

Due to the competitive business of wireless services, providers may be unwilling to share their valuations of bandwidth with their competitors. To address this desire, we employ a sealed bid double auction with a uniform clearing price.

[11] proved that incentive compatibility (incentive to bid honestly) cannot be maintained for both buyers and sellers without subsidies. Additionally, [19] extended this work by showing that no incentive compatible mechanism exists for multi-unit allocations with

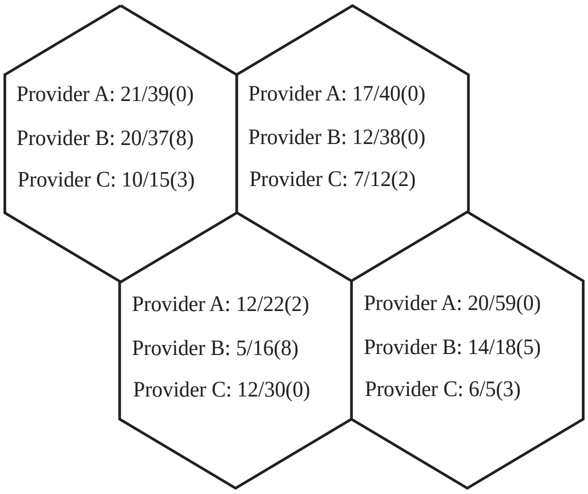

Figure 1: Each cell contains providers and customers for each provider. This figure is a random example, in the format of *Provider Name: requested bandwidth / available bandwidth (leased bandwidth)*.

uniform-price. Without an incentive-compatible mechanism, the bidders may have an incentive to behave strategically, reducing efficiency from the idealized case. Despite the strategic friction, it is widely accepted that markets provide an effective means of reallocating resources.

To maintain confidentiality of the bids, two methods may be used to host the auction. The first would be to have an auction server hosted by a trusted, impartial third party, whether private or government owned. This has the benefit of accountability, but entails cost and management overhead. The second approach is to use a distributed (M+1)st price auction protocol, such as the one devised by [8]. The service providers could each maintain their share of servers as auctioneers without the need for third parties. This protocol keeps bids private and determines correct winners, as long as a sufficient number of auctioneer servers are reputable.

The double auction generalizes traditional first and second price auctions to M sell offers and N buy offers from any combination of buyers and sellers. All of the bids (buy and sell offers) are ranked by price. The clearing price, that is, the price at which all transactions for that round are made, is computed as $kp_{M+1} + (1-k)p_M$, where p_x is the xth highest price. The price range between p_{M+1} and p_M is where supply and demand are balanced. In this study, we used the value of $k = 0.5$.

Once the clearing price has been found, buyers with bids at or above the clearing price will purchase the commodity from the sellers with offers at or below the clearing price for the amount of the clearing price. As buyers and sellers may not be balanced, the buyers with the highest bidding prices and sellers with the lowest offering price trade.

In our model, service providers reevaluate their valuations at regular intervals on the scale of several minutes, though incorporating historical information from other scales such as days or weeks may also work well. Because our model assumes customer connection and movement rates will be the same until the next bandwidth market session, the market clearing frequency should maximize the accuracy of predicting and tracking customer connection and movement rates. Having auctions clearing every several minutes is long enough to allow for a clear distinctive ownership of bandwidth. At the same time, several minutes is short enough such that if a customer turns on a device in a cell where the provider has no other known customers, only a few minutes will lapse before the provider can lease the bandwidth and provide service. Common channels to track and report the location of customers to their respective providers would be required.

In each auction, each service provider places bids into the market to buy and sell leases of a unit of bandwidth. While companies may not trust their automated systems with large unmonitored purchases, the cost of a mistake of a 30 minute lease would be comparatively minor, and the bidding system could be taken offline if problematic.

To facilitate practical trading, bandwidth should be divided up into either individual usable communication channels, or small groups of communication channels. Because all blocks of bandwidth are considered equal, the valuations define how beneficial each additional unit of bandwidth would be for the next time interval. The valuation attributed to each incremental unit of bandwidth is the additional utility of having that unit of bandwidth compared to not having it. Further details of how the valuations are calculated are provided in section 3.3.

Providers calculate the valuation for each unit of bandwidth available in the cell before participating in the market. If a provider already owns t units of bandwidth, then the highest t marginal values are sell offers, and the remaining marginal values are buy offers.

As an example of pricing on the bandwidth market, consider a provider with very little traffic compared to the bandwidth it owns. The channels that the provider needs to accommodate its customers would be bid to sell at relatively high prices, as it would take a large incentive for the provider to sell this bandwidth. The provider would place channels that exceed its expected requirements on the market at lower prices. All the provider's buy bids would be very low prices, as it will likely not need the additional bandwidth. Conversely, a provider with heavy traffic would place high value on channels it owns, and some of its buy bids would have high prices, as it would be likely to run out of bandwidth from its customers' demand.

3.3 Bandwidth Valuation

The utility of bandwidth in our model is the cost of the detriment to the customers as modeled by the provider. We consider all situations for one customer, compute the expected value of the detriment to that customer, and then multiply it by the number of customers in the cell. Our model's input includes the current number of customers, k, channel request and release rates, μ_{req} and μ_{rel}, arrival and departure rates of customers to the cell, μ_{arr} and μ_{dep}, number of channels currently in use by customers, e, and the total bandwidth, t. In this section, we derive the expected cost of BCRs on an individual customer, and then compute the expected utility of the specified amount of bandwidth, based on the BCR's cost in terms of model inputs.

We first examine an individual customer's circumstances. When originating a new connection or moving to another cell with a connection in progress, a customer requests a channel from his/her provider. If all of the provider's channels in the cell where the customer is located are in use, then the customer's channel request is blocked.

We assume that each customer has a threshold of BCRs at which the customer will either discontinue service (perhaps transferring to another provider) or incur some form of monetarily quantifiable dissatisfaction. We refer to the event at which the customer reaches a quantifiable dissatisfaction level as a *critical incident*. We further assume that the customers' distribution of dissatisfaction thresholds follows a normal distribution, though this may be easily changed. For use in real-world cases, this distribution may be empirically obtained by surveys, and would need to account for factors such as the distributions of customers' tolerances and usages of the system. Given our assumptions, the probability that a customer will incur a critical incident thus follows the cumulative density function of the normal distribution, Φ, with a mean of $\bar{b}$ channels and standard deviation of σ, given b BCRs, can be written as

$$\Phi(\frac{b-\bar{b}}{\sigma}). \tag{1}$$

Given a certain number of BCRs over a period of time, we assume that each BCR is given randomly to customers requesting a channel for a new connection or handoff. The number of different permutations with replacement that b BCRs can be allocated to c customers is c^b. We chose to use permutations with replacement because a customer could make and release several connections within the interval between market clearings. While permutations without replacements would not capture the possibility of a customer incurring multiple BCRs, it would be more appropriate for very short market cycles, where each customer would not have the opportunity to place a second channel request. If we isolate one customer, we can compute the probability of that customer having i of the b BCRs. The number of permutations of the $b-i$ remaining BCRs distributed among the $c-1$ customers is $(c-1)^{b-i}$. Additionally, we must account for the different ways of choosing the i BCRs, $\binom{b}{i}$. The probability, Q, of a customer having i BCRs is the number of permutations of the customer having i BCRs divided by the number of possible permutations of BCRs, expressed as

$$Q(b,c,i) = \frac{\binom{b}{i}(c-1)^{b-i}}{c^b}. \tag{2}$$

Using the probability of receiving i BCRs, Q, to weight the customer's probabilities of a critical incident, we can find the expected probability of a given customer incurring a critical incident given b BCRs among c customers. Multiplying this by the number of customers, c, gives us the number of expected critical incidents given b BCRs. This can also be multiplied by the utility loss of a critical incident, ν. The expected utility loss of critical incidents is thus

$$L(b,c) = \nu c\left(\sum_{i=0}^{b}\Phi(\frac{i-\bar{b}}{\sigma})Q(b,c,i)\right). \tag{3}$$

Because we want to find the probability of each number of simultaneous channels in use or current customers in the cell, we make some assumptions to simplify computations and increase scalability. We assume that customer channel requests and customer channel releases are Poisson processes. This assumption that channel releasing is a Poisson process follows from the Poisson channel requests process and an assumption of exponentially distributed channel hold times. While other distributions are more accurate, the exponential distribution is a decent approximation to call hold times [2], and greatly simplifies our computation. As wireless Internet connections are increasingly used for longer durations as primary Internet connections, we argue that Poisson arrival and departure processes provide a more generic framework than current data models for wireless telephony.

To find the probability that the number of connections has varied by each possible amount, we employ the probability distribution function of the difference between two Poisson processes as derived by [16]. The probability function is expressed as

$$P_k(\mu_1,\mu_2) = e^{-(\mu_1+\mu_2)}\left(\frac{\mu_1}{\mu_2}\right)^{k/2} I_k(2\sqrt{\mu_1\mu_2}), \tag{4}$$

where k is the integer net change in number of occurrences ($arrivals - departures$), μ_1 and μ_2 are the arrival and departure rates respectively, and $I_k(z)$ is the modified Bessel function of the first kind.

From these formulae, we construct a model representing the expected utility of a provider based on the number of critical incidences. $L(b,c)$ must be weighted to correspond with the probability of having the number of customers in a given cell as well as the probability of having the number of BCRs. To account for these dynamics, the model needs the current values of number of customers in the cell, c, number of channels currently in use, e, and total number of channels available, t. Each of these three parameters is with respect to a given service provider.

For a specified number of customers in the cell, k, channel request and release rates, μ_{req} and μ_{rel}, number of customers in the cell at the time when μ_{req} and μ_{rel} were measured, c, and specified number of current channels in use, j, the probability of exceeding the bandwidth is given by $P_{j-e}(\frac{k}{c}\mu_{req}, \frac{k}{c}\mu_{rel})$. This represents the probability of the number of channels in use changing by $j-e$. This probability can also be interpreted as the expected fraction of time that the system spends in that given state. The values of μ_{req} and μ_{rel} are assumed to correlate with the number of customers, thus are scaled by the change in customers, $\frac{k}{c}$. The number of BCRs is represented by $j-t$, and channel requests are only blocked when $j > t$. Weighting the utility of critical incidents for a given number of BCRs can therefore be written as

$$\begin{aligned}\hat{U}(t,k,e,c,\mu_{req},\mu_{rel}) = \\ \sum_{j=t+1}^{\infty} P_{j-e}(\frac{k}{c}\mu_{req}, \frac{k}{c}\mu_{rel})L(j-t,k).\end{aligned} \tag{5}$$

Similarly, the formula must be weighted with the probabilities of having different numbers of customers in each cell based on the customer movements. Each customer has a probability of whether it will move to a new cell or stay in its current one, which yields a sum of binomial distributions for the expected number of customers in a given cell. We approximate this binomial distribution with a normal distribution[1] in order to easily sum the expected arrivals and departures, as normal distributions may be easily summed. For every cell, n, each service provider knows the number of customers in that cell, c_n. To find the expected number of customers in a given cell, a provider can use the probabilities, p_n, that a customers would come from cell n (or stay in the current cell) to

[1]The normal distribution is an effective approximation of the binomial distribution for large populations with non-trivial probabilities.

find $\mu_{cust} = \sum_{\forall n} c_n p_n$. Similarly, the standard deviation may be found as $\sigma_{cust} = \sum_{\forall n} c_n p_n(1-p_n)$. Applying a continuity correction of $\pm\frac{1}{2}$ to the normal distribution on the customer movement process, we can arrive at an expected utility of a given amount of bandwidth as

$$U(t,e,c,\mu_{req},\mu_{rel},\mu_{cust},\sigma_{cust}) = \sum_{k=1}^{\infty}\left(\Phi(\frac{k-\mu_{cust}}{\sigma_{cust}}+\frac{1}{2})-\Phi(\frac{k-\mu_{cust}}{\sigma_{cust}}-\frac{1}{2})\right) \cdot\hat{U}(t,k,e,c,\mu_{req},\mu_{rel}). \quad (6)$$

The difference in utility computed by Formula 6 between two amounts of bandwidth yields the valuation of having that additional bandwidth.

For computational purposes, both infinite summations in this model are approximated by summing over the range of ± 3 standard deviations from the mean for each of the bivariate Poisson distributions. Accordingly, the standard deviation is computed as $\pm\sqrt{\mu_1+\mu_2}$.

4. SIMULATION

4.1 Implementation

To evaluate our model, we implemented a simulation that runs with steady-state customer behavior. The model is geographically consistent; the number of customers in the entire system is fixed, and customers may move from their current cell to any adjacent cell. The customers can either stay in the same cell or move into one of the six neighboring cells, each choice has its own probability. When a customer is at the edge of the system and moves off the board, it will wrap around to the other side in a toroidal fashion. We also ran our simulation with customers staying in place rather than wrapping around the edge, but it did not make a significant difference in any of the results. Customers may move while connections are in session.

The discrete-time Markov chain used to model customer connections has three states: idle, singe-channel connection (such as voice and low bandwidth data calls), and multi-channel connection (large downloads, intensive data calls). When making a connection, if bandwidth resources are insufficient, the connection is considered blocked and the customer remains in the idle state. Single-channel and multi-channel connections each use a primary channel, but multi-channel connections may also use some number of supplemental channels. These supplemental channels increase the bandwidth of each multi-channel connection, bettering the customers' experiences. However, supplemental channels are allocated only after all the single-channels and primary channels.

The simulation iterates over both the movement and connection Markov chains a specified number of times between each round of bandwidth trading. Each service provider calculates its bandwidth valuations from the measurements of movement, connection origination, and connection release rates, collected between each round of trading. The greater the number of iterations, the more closely the behavior approximates the exponential distribution of connections and movements.

It is also a simultaneous strength and weakness of our model that it is based on the assumption that the next period between bandwidth trading is statistically the same as the previous one. This effect is useful in that it can rapidly adapt to a situation, but does not work as well if the system's rates change too fast between markets. One of the necessary precautions of implementing this model in a real system would be to make sure that the markets are frequent enough such that the system's connection and movement rates are not drastically changing between every round of bandwidth trading.

We ran each simulation for 50 market periods, which was enough for the results to converge. Between otherwise identical runs, the randomness caused only about a 1% difference. For measuring the simulation, we gathered the global average number of simultaneous successful connections, average number of BCRs per provider, each provider's utility values, each provider's leases, and average data rate (accumulated per cell as the number of the provider's unused channels divided by the number of multi-channel connections, or 0 if no multi-channel connections were in progress).

4.2 Results

To evaluate our model, we ran the simulation with different configurations of number of providers, number of customers, and cell layouts, as detailed in Table 4.2. For each of the 8 simulation environments, we ran 32 tests varying the connection rates, customer movement rates, and the number of customer iterations between each market. These customer behavior parameters are detailed in Table 2. We ran each of these 256 configurations with both the bandwidth trading market enabled and disabled. The customer dissatisfaction threshold was set to $\bar{b} = 3$ and $\sigma = 1$. Because the values of $\bar{b}$ and σ define the valuation curve with low numbers of customers, they would need to be calibrated to customer expectations for a real implementation.

The primary effects of our simulation runs are shown in Figure 2. Each point represents the results of one provider in a particular simulation run. All of the dimensions in this graph are scaled by the inverse of the average number of customers per provider. This scaling makes the values comparable between large and small scale simulations. The vertical axis is the increase in the provider's utility with the horizontal plane representing 0. The utility is computed by adding any profit from leasing bandwidth to the provider's valuation of its bandwidth, minus the provider's valuation of the bandwidth it owns if the market were not used. While utility may technically be negative, a rational provider would not participate in trading if it expected to lose money. Each provider's utility was non-negative in all of our trials.

The axis labeled "bandwidth difference/customer" represents the difference between the provider's average bandwidth in a given cell, and the average bandwidth of all the other providers in the system, scaled by the number of the provider's customers. A positive bandwidth difference means the provider has more bandwidth per customer than average. The axis labeled "BCR reduction/customer" is an output of the system, and represents the average decrease in the number of BCRs due to enabling the market. Both horizontal axes have a dark line to indicate the value of 0. Each data point is represented as a diamond shape. Lines are drawn from each data point to the corresponding position on the horizontal plane. The shadow at the base of each line is for clarity, with darker shades indicating more points in the same area. From the diamond shape, a line protrudes representing the amount of bandwidth leased per customer. If the line is to the right, it means the provider leased bandwidth from other providers, whereas the left means it leased bandwidth to other providers. The length of the line from the edge of the diamond shape indicates the amount of bandwidth leased per customer.

Table 1: Simulation Environments and Results

Environment Number	1	2	3	4	5	6	7	8
Number of Providers	3	3	6	6	6	6	6	9
Cell Grid Size	4x6	2x1	2x1	4x6	4x6	4x6	4x6	10x10
Number of Customers	1400	120	225	2350	2350	2350	2350	16820
Average Bandwidth/Customer	1.11	0.62	0.60	0.56	0.55	0.69	0.45	0.50
Standard Deviation Provider-Cell's Bandwidth/Customer	0.81	0.30	0.30	0.36	0.31	0.50	0.13	0.30
Standard Deviation Provider's Customers/Cell	0.80	1.67	1.11	0.77	0.77	0.77	0.77	0.83
Reduction of Blocked Channel Requests	43.6%	3.9%	2.7%	16.5%	18.9%	34.6%	-0.2%	9.7%
Change in Utility Per Customer	8.5	0.007	0.5	6.5	5.9	14.0	0.5	4.2

Table 2: Simulation Customer Behavior

Probability of Customer Moving	.1, .2, .5
Probability of Placing Single-Channel Connection	.1, .2, .5
Probability of Placing Multi-Channel Connection	0, .1, .4
Probability of Releasing a Connection	.2, .5, .9
Iterations Between Markets	1, 10, 20, 30

From Figure 2, it is clear that the increase in utility gained from the market correlates with the reduction of BCRs (towards the upper right of the graph). The correlation of 0.70 between utility and BCR reduction confirms the effectiveness of the primary market function to reallocate bandwidth to reduce the number of BCRs. Another notable effect from Figure 2 is that providers with the largest deviation from the average bandwidth per customer had the largest utility gains from the market. This effect can be seen by the points being higher on both ends of the bandwidth difference axis. The highest utility gains were those by providers with the least bandwidth per customer relative to other providers, and is supported by a 0.78 correlation between the utility change per customer and the amount of bandwidth leased per customer (ignoring lease direction). These highest utility gains were made by leasing bandwidth to reduce the number of BCRs, as depicted by the other horizontal axis. The high utility gain of the providers on the positive side of the bandwidth difference is primarily due to the profit made by leasing excess bandwidth to other providers, incurring small increases in BCRs.

Across all simulation environments, the market reduced the average number of BCRs by 22%. The market also reduced the variance in BCRs by 19%, which reduced the probability that a customer would incur multiple BCRs. However, we note that the tests were designed to test a range of behaviors of the model rather than resemble a specific real-world distribution of test cases. The range of reduction in BCRs from these 8 simulation environments had a wide range, from 0.2% to 74%. From Table 1, the only environments that did not see a noticeable reduction in BCRs were environments 2, 3, and 7. Environment 7 could not have improved much, because it had the lowest standard deviation of provider-cell bandwidth per customer, meaning that all providers having very similar amounts of bandwidth per customer. Environments 2 and 3 were very small, allowing fewer fluctuations in the system. Environments 1 and 6 had the greatest reduction in BCRs due to the opposite conditions: many resources and a large difference between large and small providers. Environment 8, with 16820 customers over 100 cells, had results that were in line with the rest of the simulations, showing that our model scales to real-world sizes.

Some providers incurred a small number of additional BCRs when the market was added (negative BCR reduction/customer). This increase is primarily caused by the provider leasing bandwidth to another provider that has a higher valuation. This higher valuation stems from having few customers to absorb BCRs, since our valuation model attempts to prevent customers from receiving multiple BCRs channel to prevent critical incidents. It is also notable that the number of iterations between markets had no effect, in that its correlation to all outputs was trivially close to 0. This supports the case for our model being resistant to random perturbations of user behavior with a stationary user model, as well as supports our case for our model having flexibility in the inter-market time as long as the customer behavior can be assumed to be reasonably consistent in that interval. We note that the perturbations from customer movement had little effect as well, despite though the providers' tracking. This is largely due to the unbiased customer movement distribution. The only correlation with customer movement that was not trivially close to 0 was the average data rate, with a correlation of -0.14 without the market and -.16 with the market. A negative correlation between these values indicates that when customers move around rapidly, our model slightly reduces the average data transfer rate, though our measurement of customer dissatisfaction thresholds does not take this into account.

Figure 3 shows three time sequences of a single provider's available bandwidth after the market. To view the dynamics of a market for this figure, we changed the simulator to only record the rates of the previous iteration, instead of continually accumulating customer rates. Without this tweak the bandwidth trading converges. The convergence is more useful to measure the steady state, but causes the bandwidth line on the graphs to slowly converge to an equilibrium. In region A of Figure 3, bandwidth is somewhat scarce, influ-

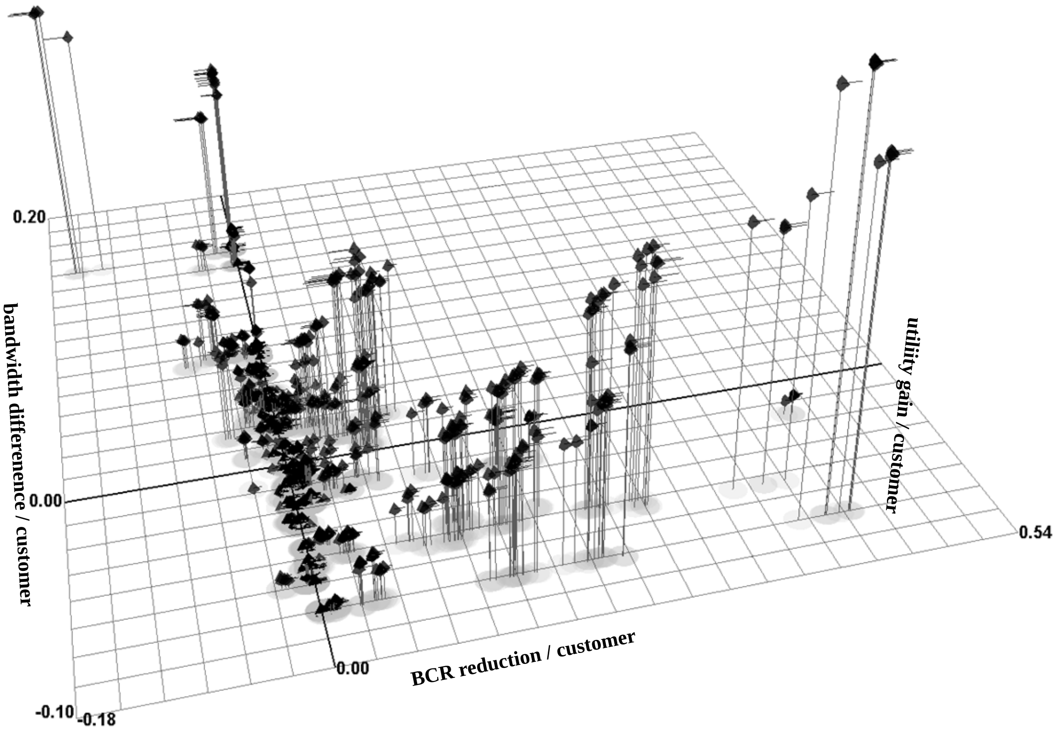

Figure 2: Provider utility (vertical axis) plotted against several simulation features, depicting the model's primary effects.

encing the provider to only lease it when necessary. This environment exhibits tight tracking, where the available bandwidth closely mimics the customers. Because the available bandwidth is reactive to the number of customers and the number of channel requests, this delay emphasizes the necessity of a sufficiently small time between market to ensure that traffic is reasonably steady. Region B shows the case where the provider has insufficient bandwidth, and the bandwidth is very scarce and expensive. In this case, the provider leases only a small (yet still insufficient) amount of bandwidth during the peak demands, representing a worst-case scenario when a cell is overloaded. A more typical example is region C, where the provider generally maintains sufficient bandwidth, and leases out bandwidth accordingly. Region C is more typical of real-world situations because the cell's bandwidth is reasonably close to its utilization. The case where bandwidth is plentiful is not shown; when bandwidth is plentiful, little trading commences. The little trading leaves flat lines for available bandwidth above the current number of channels requested, and occasionally resembles the characteristics of region C during high volume intervals of channel requests.

Even though our model does not account for supplemental data channel requests in its valuation, Figure 4 indicates that the rate is generally not greatly affected (the data points are close to the line $y = x$). If a provider deems that a lack of supplemental data channels will impose a critical negative incident for customers, it may be accounted for and calibrated in the customer threshold. However, if the provider deems it to be less critical to customers, Figure 4 shows that omitting supplemental data channels to simplify the customer model is not problematic.

5. CONCLUSIONS

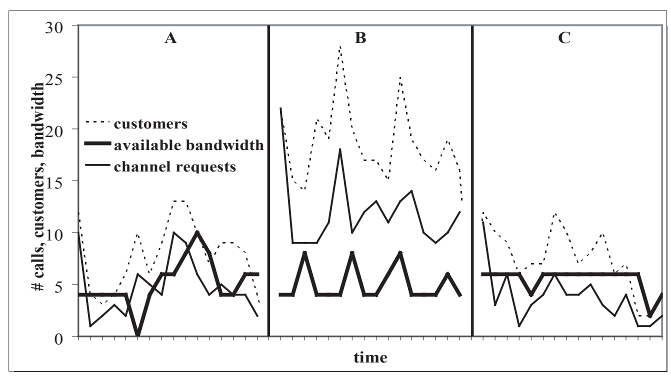

Figure 3: Three samples (regions following A, B, C) of an individual provider's available bandwidth after trading, plotted along with the number of customers and channel requests.

In general, our simulation results confirm that using a bandwidth trading market with our customer valuation model is viable for providers. This is not surprising, as markets are often very efficient at allocating resources according to valuations, but our model has the added benefit of more accurate market prices and fewer timing-related issues than spot markets. Our results show that our model is scalable and adaptable to many different situations.

We believe our work is aligned with the future goals of the FCC and builds off other technologies. We describe a model based on customer thresholds of critical incidences, which accommodates customer connection rates and movements. We implemented a simulation of this model and evaluate it on a variety of inputs.

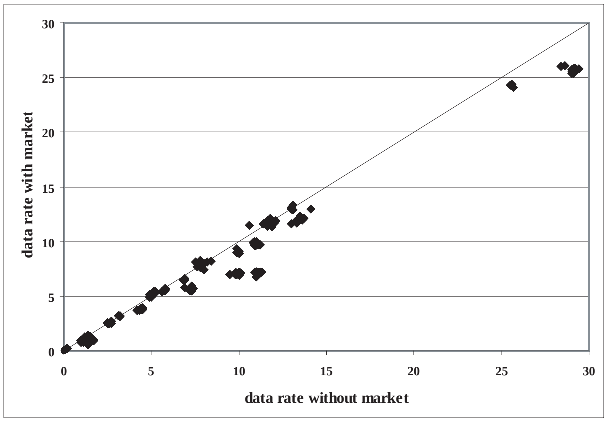

Figure 4: Average data rate for each simulation run with and without the market enabled, measured by the number of the provider's unused channels when at least one multi-channel connection is active divided by the total number of multi-channel connections.

While a case-study would be required to determine the benefits in a particular real-world situation, we feel the simulation results are indicative that our approach would be beneficial in practice. Though the bandwidth leasing we require for our model is not yet permitted by the FCC, we feel that it fits with the spirit of the FCC regulations in maintaining accountability, and provides the added benefit of maintaining competitive markets.

6. REFERENCES

[1] B. Aazhang, J. Lilleberg, and G. Middleton. Spectrum sharing in a cellular system. In *IEEE Eighth International Symposium on Spread Spectrum Techniques and Applications*, pages 355 – 359, 2004.

[2] F. Barceló and J. Jordán. Channel holding time distribution in public telephony systems (pamr and pcs). *IEEE Transactions On Vehicular Technology*, 49(5):1615–1625, September 2000.

[3] M. Bykowsky. A secondary market for the trading of spectrum: Promoting market liquidity. *Telecommunications Policy*, 27(7):533–541, August 2003.

[4] G. R. Faulhaber and D. J. Farber. Spectrum managment: Property rights, markets, and the commons. In *Proceedings of the 30th Telecommunications Policy Research Conference*, 2002.

[5] Federal Communications Commission (FCC). FCC adopts spectrum leasing rules and streamlined processing for license transfer and assignment applications, and proposes further steps to increase access to spectrum through secondary markets. *FCC News Release*, 2003.

[6] J. Hwang, M. B. H. Weiss, and S.-J. Shin. Dynamic bandwidth provisioning economy of a market-based ip qos interconnection: Intserv-diffserv. In *Proceedings of the 28th Telecommunications Policy Research Conference*, 2000.

[7] W. Ibrahim. Customer's satisfaction based call admission control scheme for next generation wireless networks. In *Proceedings of the 2004 International Research Conference on Innovations in Information Technology*, 2004.

[8] H. Kikuchi. (m+1)st-price auction protocol. In *Proceedings of The Fifth International Conference on Financial Cryptography*, pages 291–298, February 2001.

[9] M.-K. Kim, M.-C. Park, and D.-H. Jeong. The effects of customer satisfaction and switching barrier on customer loyalty in korean mobile telecommunication services. *Telecommunications Policy*, 28(2):145–159, 2004.

[10] S. Mandal and D. Saha. A bandwidth pricing model that is attractive to regulators, users, and providers. In *Proceedings 12th IEEE International Conference on Networks*, volume 2, pages 738–742, 2004.

[11] R. B. Myerson and M. A. Satterthwaite. Efficient mechanisms for bilateral trade. *Journal of Economic Theory*, 29:265–281, 1983.

[12] C. Ng, P. Buonadonna, B. N. Chun, A. C. Snoeren, and A. Vahdat. Addressing strategic behavior in a deployed microeconomic resource allocator. In *Proceeding of the 2005 ACM SIGCOMM Workshop on Economics of Peer-to-Peer Systems*, pages 99 – 104, 2005.

[13] E. Noam. Spectrum auctions: Yesterday's heresy, today's orthodoxy, tomorrow's anachronism. taking the next step to open spectrum access. *Journal of Law and Economics*, 41(2):765–790, October 1998.

[14] J. M. Peha and S. Panichpapiboon. Real-time secondary markets for spectrum. *Telecommunications Policy*, 28:603618, 2004.

[15] D. Reininger, D. Raychaudhuri, and M. Ott. Market based bandwidth allocation policies for qos control in broadband networks. In *Proceedings of the First International Conference on Information and Computation Economies*, pages 101 – 110, 1998.

[16] J. G. Skellam. The frequency distribution of the difference between two poisson variates belonging to different populations. *Journal of the Royal Statistical Society*, 109(3):296, 1946.

[17] S. Tanaka, H. Yamaki, and T. Ishida. Mobile-agents for distributed market computing. In *Proceedings of the 1999 International Conference on Parallel Processing*, pages 472 – 479, September 1999.

[18] A. Tonmukayakul and M. B. Weiss. Secondary use of radio spectrum: A feasibility analysis. In *Proceedings of the 31st Research Conference on Communication, Information and Internet Policy*, 2004.

[19] P. R. Wurman, W. E. Walsh, and M. P. Wellman. Flexible double auctions for electronic commerce: Theory and implementation. *Decision Support Systems*, 24(1):17–27, 1998.

[20] Y. Zhou, D. Wu, and S. M. Nettles. On the architecture of authentication, authorization, and accounting for real-time secondary market services. *International Journal of Wireless and Mobile Computing*, To Appear.