

A SPOT PRICING FRAMEWORK TO ENABLE PRICING AND RISK MANAGEMENT OF INTER-DOMAIN ASSURED BANDWIDTH SERVICES

Mehdi Aboufadi
Aparna Gupta

Decision Sciences & Engineering Systems
Rensselaer Polytechnic Institute
Troy, NY 12180, U.S.A.

Ritesh Pradhan
Shivkumar Kalyanaraman

Electrical, Computer, & Systems Engineering
Rensselaer Polytechnic Institute
Troy, NY 12180, U.S.A.

ABSTRACT

In the current bandwidth market, Internet Service Providers (ISPs) provide guaranteed Internet bandwidth within their domains. However, they are incapable of providing such assurances for data crossing their domain boundaries. In this paper, we present a spot pricing scheme for Internet bandwidth contracts within an ISP domain. These models when implemented at access or exchange points of different ISP domains would provide assured bandwidth for inter-domain traffic. Each contract will constitute a Quality of Service agreement between a customer and a provider within an ISP domain. By appropriately bundling derivative contracts defined on the intra-domain service contracts, a provider will not only be able to give inter-domain Quality of Service assurance, but will be able to add new services and manage its portfolio of services.

1 INTRODUCTION

Today, Internet mostly consists of ISPs providing a *best-effort* service, meaning that the network tries as much as it can to push data across the network from source to destination. But, in doing so, it does not give any guarantee to its customers regarding the data actually reaching the destination. Significant improvements in the technology over the last few years have enabled the providers to incorporate better assurances on *Quality Of Service (QoS)* of the network traffic. But, the assurances are restricted to the provider's own domain. Once the traffic crosses its boundaries, it is back in a *best-effort* service condition, with no assurance given to the customers.

One popular way of overcoming the difficulty in providing assured services is to provide more capacity than the average demand. This way the probability of data being lost in its transit from one end to the other decreases and assured service can be provided to the customers. But, the solution of commissioning more bandwidth to overcome

this problem is inefficient and has practical limitations as the costs of providing additional capacity are high. Additionally, due to high costs, the providers are sensitive to the poor utilization of the network.

In this paper, we lay the foundation for a model that will provide *inter-domain* (i.e. internet traffic that crosses multiple domains or administrative areas) bandwidth guarantees to the enterprise customers. The prices quoted in these contracts will be dynamic so that the providers can leverage on their network utilization and provide customers assured services for their *inter-domain* traffic.

Currently, as far as we know, no one (ISP or a Network Service Provider) provides assured bandwidth for inter-domain traffic. They only provide intra-domain bandwidth assurances for their customers. Effectively, the ISPs provide assurance to the data terminating in their respective domains. With our proposed model the ISPs will have an additional source of revenue, as now they could provide assurances to the traffic which terminates at other ISP's domain. An attractive feature of our model is that it is implementable on the differentiated services architecture (diff-serv) and can be overlaid on schemes which are capable of providing intra-domain assured services like Distributed Dynamic Capacity Contracting (Yuksel and Kalyanaraman 2002).

The paper is organized as follows. Section 2 provides a brief literature review done both in networking and finance areas. In section 3, we give an overview of current state of bandwidth market. Section 4 describes our Internet pricing model in detail. In section 5, we present simulation results of our spot pricing model. In section 6, we summarize our discussions and point out the future steps in our research.

2 RELATED WORK

Internet pricing is a growing area of research. Until recently, the providers had, in general, opted for flat rate or time-of-the-day pricing (Odlyzko 1998, Odlyzko 2000 and Paschalidis and Tsitsiklis 2000). These schemes don't

respond to the current state of the network. In other words, transmission of data through the network can cause congestion, which is not reflected in the flat rate pricing scheme. In the current market scenario, these static pricing schemes work satisfactorily, because of the factors such as replacing of copper with optical links, faster routers, better routing algorithms, and slower access bandwidths.

On the other hand, *dynamic pricing* schemes such as MacKie-Mason et al.'s *Smart Market* (MacKie-Mason and Varian 1995a, and MacKie-Mason and Varian 1995b), Kelly et al.'s *Proportional Fair Pricing Scheme* (Kelly, Maulloo and Tan 1998), and Gupta et al.'s *Priority Pricing* (Gupta, Stahl and Whinston 1997) take into account the state of the network. But due to the granularity of their pricing strategy (per-packet as opposed to contract time), there have been doubts about their implementations (Yuksel and Kalyanaraman 2002). Recently, however, Yuksel et al. have proposed an implementable *Pricing Over Congestion Control (POCC)* scheme for diff-serv architecture (Yuksel, Kalyanaraman and Goel 2002). Their scheme is overlaid on the congestion control framework proposed by Harrison, Kalyanaraman and Ramakrishnan (2001) POCC provides a range of fairness (e.g. max-min, proportional) in rate allocation by using pricing as a tool.

Pricing being a critical issue, we reviewed the different pricing schemes described in Dolan and Simon (1996) that have been used in the literature. Some of these are uniform (linear) pricing schemes, all units quantity discount schemes, non-linear pricing schemes and mixed pricing schemes. All the pricing schemes are differentiated on the basis of their marginal price. In linear pricing scheme (refer Figure 1), the marginal price, or the *price schedule*, is a constant irrespective of the customer's demand or the supplier's capacity. In a non-linear pricing scheme, such as Ramsey Pricing Scheme ((Dolan and Simon 1996), (Wilson 1993)), the price schedule is based on the customers' demands (refer Figure 2). Higher the customer's demand lesser will its price be. Such a price schedule would attract customers with higher demand, thereby improving the capacity utilization of the network.

In all the congestion sensitive schemes we studied, the length for each contract is very small (in the range of milliseconds). Designing and delivering spot and forward contracts on such time-scales is difficult. So there is a growing need for pricing schemes for longer-term contract.

3 BANDWIDTH MARKET AT A GLANCE

Given the importance of the pricing of assured services, it is essential to understand the current bandwidth market to base our estimates of marginal costs. It is also essential to know the type of services that are available currently and position our model in the context of those services. In



Figure 1: Linear Pricing Scheme with Price as a Linear Function of Demand

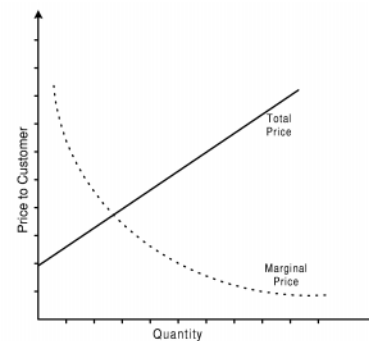


Figure 2: Non-Linear Pricing Scheme with Price as a Non-Linear Function of Demand

this section, we present a brief overview of the bandwidth market.

Due to the global growth of the Internet, today bandwidth is traded in the market as a commodity. In bandwidth market, companies sell different types of capacity (e.g. point-to-point fiber or coaxial lines, satellite frequency band) with various parameters: assured rate, burst rate, contract term. Several wholesaler companies such as, Qwest, Nortel, AT&T, Verizon, Sprint, etc., are players in this billion-dollar market. These companies have laid down physical links and sell bandwidth capacity over these links.

Some other companies, such as, Enron, RateXchange, Telco Exchange, have employed a newer economic model. These companies act as bandwidth market intermediaries where several functionalities can be provided to both buyers and sellers. A bandwidth intermediary can function as an exchange, broker or hybrid of the two. In an exchange, buyers and sellers meet and perform trades (for example, Bandwidth Market Ltd.). Just as the stock exchange, the Bandwidth Exchange earns revenue by charging fees to subscribers (i.e. both buyers and sellers) and/or by charging commissions for each trade made on their platform.

A bandwidth broker, on the other hand, buys bandwidth from wholesalers in large quantities and sells it to retailers and other customers. There are parties in the bandwidth markets that work both as a bandwidth exchange and a broker simultaneously. This is the hybrid model. Enron’s model was a hybrid model as it functioned both as a broker and an exchange. RateXchange, on the other hand, functions as an exchange and Telco Exchange functions as a broker.

Presence of bandwidth intermediaries has made it possible to introduce flexibility in effectiveness and efficiencies in service delivery. For example, given several wholesalers (bandwidth sellers), Enron could find the cheapest way of constructing a leased-line between any two given points in the world. Enron also provided risk management tools for buyers and sellers to handle price uncertainties. Similar to the stock market, buyers can buy a capacity at a pre-determined price before their required date in exchange for a premium, (longing a call option). Similarly, sellers can sell a capacity in advance for a particular period (shorting a put option).

More information on some of the specific prices and services available in the market has been included in the Appendix. We use this information for modeling marginal costs, competition (captured by Ramsey Number) as described in the following sections. It also validates our assumptions such as choice of fixed term contracts or non-linear price schedule.

4 SPOT PRICING MODEL

In this section, we will describe the spot pricing model assumptions. With these assumptions, we then describe Ramsey spot pricing model using the concepts of both demand profile and demand curve.

In commodities like energy and telecommunications, non-linear pricing models are highly popular. Compared to the erstwhile uniform pricing models, they are biased toward high demand customers. In a uniform pricing scheme, irrespective of the demand price per unit remains same. But in non-linear pricing schemes, the price per unit decreases as demand increases. Additional features like capturing the type of competition in the market, demand price elasticity of the buyers, and profit margin of the sellers make the Ramsey model an attractive choice.

4.1 Pricing Model: Definitions and Assumptions

First, we define some of the terms we will frequently use in our model. A *Domain* is defined as the *administrative reach* of an Internet Service Provider. In the context of the paper, a *buyer* is an enterprise customer consisting of multiple users, while a *seller* will be a provider which sells the different contracts to the buyers.

The basic intra-domain bandwidth pricing model is shown in Figures 3 and 4. This model will be implemented by a seller to interact with the buyers to sell its bandwidth. This feature of our model is essential for any seller-buyer relationship. Such models will be implemented at the access and/or exchange points of different domains to bring about an inter-domain service assurance to the buyers.




Figure 3: Basic Pricing Model Implemented at an ISP’s Access Point

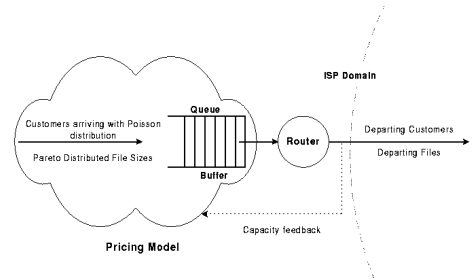


Figure 4: Figure Showing How the Pricing Model Works

To develop the spot pricing model, we make the following assumptions:

- We consider 3 alternatives for demand profiles of the buyers. Here, *demand profile* is a function of quantity and price such that it represents the number of buyers who will buy at least specific quantity at particular price (a more detailed explanation is given in §4.2).
- When buyers purchase a contract, they require bandwidth for a simple and immediate file transfer application and each buyer signs for a contract of a fixed duration.
- All buyers send the same number of files (denoted by N), and the size of each file has a Pareto distribution. The latter probability model has been studied thoroughly by Crovella, Taquq and Bestavros (1998) and Crovella and Lipsky (1998). Thus, the probability density of file-size, x , is given by

$$P(x) = \frac{ab^a}{x^{a+1}}. \quad (1)$$

- Buyers arrive by a Poisson process. This is also an observed feature of the Internet. Therefore, the inter-arrival times of the buyers is distributed as follows,

$$f_X(x) = \begin{cases} \lambda e^{-\lambda x} & \forall x \geq 0. \\ 0 & \forall x < 0. \end{cases} \quad (2)$$

- At each ISP, the arrival of the buyers depends on the *time of the day*. We assume that between 7 a.m. and 5 p.m. 70 % of the buyers arrive, between 5 p.m. and 11 p.m. 20 % arrive while the rest 10 % arrive between 11 pm and 7 am. These numbers are based on historical data (NLANR 2002).
- The seller is assumed to know the exact characteristics of the buyer like λ_t (arrival rate at the time of day), a and b .
- The capacity of the network is assumed to be constant, as is the marginal cost.

Figure 5 shows the flowchart to capture the essence of our model. A buyer arrives at *exponential* inter-arrival rate of $\lambda = 5$ customers/min averaged over a day and announces its volume requirement to the seller. Each buyer has a fixed number of files (N) to send, each file-size is Pareto distributed size with parameters: 0.35, 100. The total volume is then divided by the contract length to determine their *Asked_Capacity* (in Kbps). If the buyers have an *Asked_Capacity* that is lower than the *Available_Capacity* at the time of arrival, they are accepted into the system while the others leave the system and the *Available_Capacity* is updated for the next buyer. With the demand of the buyer known to the seller, the price schedule is calculated as explained in the next sub-section.

4.2 Ramsey Pricing Model

We know that a seller provides a buyer with a portion of the bandwidth for a price. Therefore, the most important part of the model is finding the right spot price for the buyer. The interesting feature of the Ramsey model is that it captures the characteristics of the market (oligopoly v/s monopoly) and buyer(demand elasticity) to calculate the optimal price schedule. Ramsey Pricing model has been widely popular in the Telecommunications (Dolan and Simon 1996), (Wilson 1993) and power sectors. In this section, we explore the Ramsey Pricing model and discuss its relevance in our model.

The guiding principle of the Ramsey model is to develop prices to maximize buyer's benefits, subject to the constraint that the seller recovers its total costs (both fixed as well as variable). Additional constraint of the model is that the price schedule calculated from Ramsey Model must be lower than a uniform price schedule which provides same

net revenue to the seller. Thus, with respect to the above constraints an optimal price schedule can be obtained.

Thus if $p(q)$ is the optimal price schedule, $c(q)$ is the marginal cost of the q^{th} unit, α be the *Ramsey Number* and $\eta(p(q), q)$ be the elasticity of the demand profile, then according to the *Ramsey Rule*:

$$\frac{p(q) - c(q)}{p(q)} = \frac{\alpha}{\eta(p(q), q)}. \quad (3)$$

The demand profile $N(p(q), q)$ for a commodity is defined as the number of buyers who will buy at least q units at marginal price $p(q)$. This demand profile is, typically, a function of price and quantity desired by the buyers. It can be computed in two ways (Wilson 1993). The first method, which requires data on the distribution of customers' purchase sizes for several uniform prices, directly estimates the demand profile by obtaining a smooth estimate of it. The second method approximates the demand profile through the estimation of customers' demand functions.

Using the consumer's demand profile, we can compute their demand elasticity to price as follows:

$$\eta(p(q), q) = - \frac{\partial N(p(q), q) / N(p(q), q)}{\partial p / p}. \quad (4)$$

The *Ramsey Number*, denoted by α , is the fraction of the monopoly profit margin that is common to all the units of the customers' demand. The interesting aspect of Ramsey Number is that it indicates the nature of the provider as a firm. For instance, a *profit-maximizing monopolist* has $\alpha = 1$, while a regulated firm with no binding revenue requirement has an $\alpha = 0$. In case of a budget-constrained welfare maximization and an oligopolist competition, $0 < \alpha < 1$.

For simulation and analysis, we considered 3 sample demand profiles, all from Wilson (1993). The first demand profile represents a market in which the customers moderately react to the prices.

$$\begin{aligned} N(p, q) &= 1 - \frac{q}{1-p}. \\ \Rightarrow \eta(p(q), q) &= \frac{pq}{(1-p)(1-p-q)}. \end{aligned} \quad (5)$$

This demand profile is drawn in Figure 6. The price schedule, or simply variation of price with quantity, of such a demand profile is given as:

$$p(q) = 1 + \frac{q(1-\alpha)}{2\alpha} - \sqrt{\left(\frac{q(1-\alpha)}{2\alpha}\right)^2 + \frac{q(1-c)}{\alpha}}. \quad (6)$$

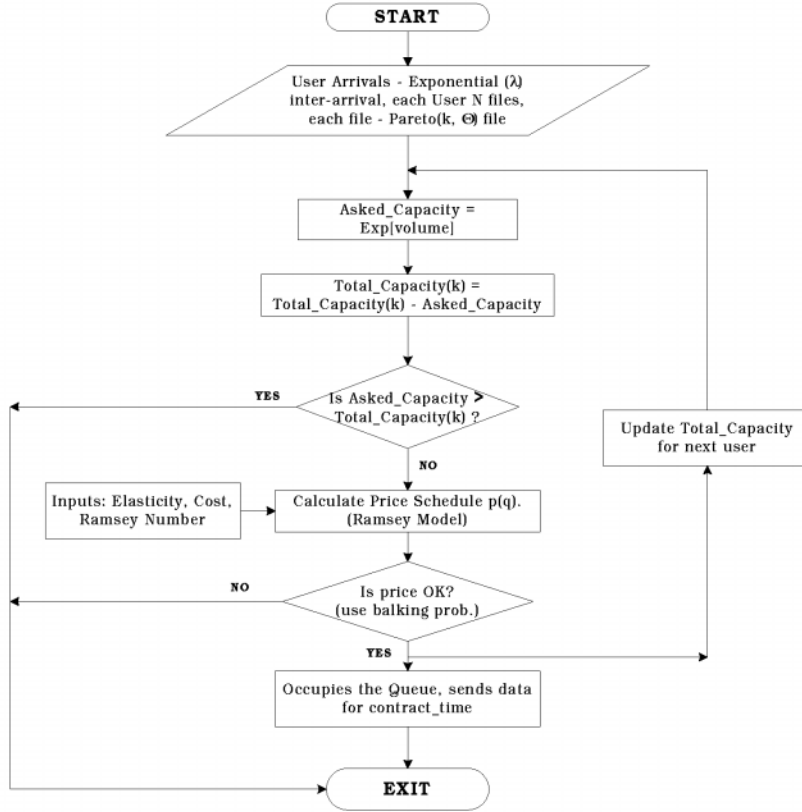


Figure 5: Flow Chart Indicating the Working of Basic Pricing Model




Figure 6: Graphical Representation of Demand Profile Given by Eqn. 5

The second demand profile that we chose, reflected a market in which the customers were very sensitive to prices.

$$N(p, q) = 1 - p - q. \quad (7)$$

This implies that

$$p(q) = \frac{1 - q}{p}. \quad (8)$$

To analyze yet another extreme case of the market, we chose the third demand profile that represented a market which was not sensitive to prices and the customers are willing to pay whatever the price is demanded.

$$N(p, q) = 2 + \frac{\ln(1 + p)}{\ln(q)}. \quad (9)$$

Solving this equation, we get:

$$p(q) = \exp \left(\frac{f \left(\frac{(c+1) \exp \left(\frac{2\alpha \ln(q)+1}{\alpha} \right)}{\alpha} \right) \alpha - 1 - 2\alpha \ln(q)}{\alpha} \right) - 1. \quad (10)$$

where, Lambert's **W** function (a series expansion of Lambert's function was used in the simulation) is defined as

$$f(W) = We^W. \quad (11)$$

The demand profile approach of Wilson 1993 for Ramsey Pricing model is a relatively new approach (Dolan and Simon 1996). A more traditional approach for building Ramsey models has been the use of demand functions instead of demand profile. To understand the demand function approach, we use a simple linear demand equation (12), where parameters m and n may be estimated by the provider, as:

$$\begin{aligned}
 q &= m + np & (12) \\
 \text{or, } p(q) &= \frac{q - m}{n}. \\
 \text{Now, } \eta(p(q), q) &= \frac{\partial q/q}{\partial p/p} \\
 \Rightarrow \eta &= np/q.
 \end{aligned}$$

Thus, according to the Ramsey rule (3), the optimal price schedule for the given demand function (12), the type of firm α , and marginal cost $c(q)$, can be given by,

$$p(q) = c(q) + \frac{\alpha q}{n}. \quad (13)$$

All the price schedules obtained above give us the spot prices that can be advertised to the customers. Depending on their demand and the available network capacity, the contract will be set for a certain duration. Once the contract is accepted, the customer will be provided with a guaranteed bandwidth in the network for the contract duration. It will then be the provider's responsibility to transfer the customer's data with assured characteristics (such as, latency, packet loss, etc.).

5 SIMULATION RESULTS AND CONCLUSIONS

Based on the Ramsey model discussed in the earlier section, we will now present our simulation model itself and discuss the results. For simulating our model, we used ProModel version 4.2.

Simulation of the model discussed in the previous sections was done to validate our expectations from the model. The following were the insights we intended to gain by performing the simulation experiments:

- Generation of a price schedule based on the demand profile of the customers.
- Observe the relationship of the price schedule with the type of the firm as indicated by the Ramsey number.
- See the relationship between price schedule and duration of the contract length.
- Comment on exploiting hedging options based on the price schedules.

Given the buyer arrival process discussed in detail in section 4, a price schedule, $p(q)$, was generated using Ramsey Pricing Model (equation 6). To simulate the different types of providers (from a monopolist to welfare maximization firm), we used two different values of α (0.2, and 0.8). The marginal cost was assumed to be uniform, at the value of \$0.5/Kbps capacity. The total capacity of the network was 12 Mbps. The simulation time of the model was 240 hours with 5 replications for each run.

The first table (Table 1) represents the variations in the Asked_Capacity and Available_Capacity with respect to the contracts lengths and the volume of demand, expressed in terms of number of files. As we may notice, an increase in contracts length results in a decrease in the average Available_Capacity, thereby stabilizing the system. This was due to the demand-supply mismatch at the seller; the longer the contract length, the lower is the possibility of mismatch.

Table 1: Variation of the Average Asked Capacity and Average Available Capacity for the Contract Lengths of 1 Hr., 4 Hr. and Mix (1, 2, and 4 Hrs). This Variation Is Independent of the Type of Seller (i.e. Ramsey Number) but Depends on the Number of Files That a Buyer Customer Sends and the Contract Length

Contract Length T (hrs.)	N = 1000 files		N = 3500 files	
	Avg.Avbl. Capacity	Avg.Ask. Capacity	Avg.Avbl. Capacity	Avg.Ask. Capacity
1	4142.95	344.32	4867.01	1211.65
4	1309.26	86.11	1295.86	304.43
Mix	2195.46	203.77	2059.94	720.33

The time series plots of the spot prices for the first two demand profiles (Figures 7 and 8) were the most interesting ones. As we expected, the price levels between a competitive market and a profit-maximizing monopoly were quite different - higher for the latter, keeping the demand profile and volume the same. The average prices for the two types of markets are also shown in Table 2.

The third demand profile (Figure 9) was used to validate our solution. Using this extreme case, we obtained extremely high marginal price for low quantities and vice-versa. In such cases, Wilson (1993) recommends a cap on the marginal prices which is equal to the revenue equivalent uniform price.

Regarding the price levels between the first two demand profiles for each type of market, we notice that: first, prices are higher with more fluctuations for customers with demand profile - 1 (eqn. 5) than for customers with demand profile - 2 (eqn. 7), which can be explained by the sensitivity of the types of customers to prices (introduced earlier, and reflected by the demand profile); second, the average price (\$ /Kbps) is higher for low volumes than for high volumes of data (the volume expressed in terms of number of files),

Table 2: Variation of the Average Asked Capacity and Average Available Capacity for the Contract Lengths of 1 Hr., 4 Hr., and Mix (1, 2, and 4 Hrs). This Variation Is Independent of the Type of Seller (i.e. Ramsey Number) but Depends on the Number of Files That a Buyer Customer Sends and the Contract Length. * in the Table Indicates That the Quantity Was Negative

Demand profiles N(p,q)	Contract Length T (hrs.)	$\alpha = 0.2$		$\alpha = 0.8$	
		N=1000 files	N=3500 files	N=1000 files	N=3500 files
$N(p, q) = 1 - \frac{q}{1-p}$ Average price (Average revenue)	1	0.780 (42750)	0.661 (123231)	0.868 (47507.8)	0.761 (141982)
	4	0.882 (3789.36)	0.794 (12790.4)	0.935 (4020.29)	0.877 (14134.4)
	Mix	0.858 (10436.3)	0.751 (38298.9)	0.920 (11382.1)	0.842 (43594)
$N(p, q) = 1 - p - q$ Average price (Average revenue)	1	0.578 (31688.1)	0.566 (105626)	0.709 (3837.5)	0.677 (126306)
	4	0.582 (2503.87)	0.579 (9332.57)	0.719 (3092.83)	0.711 (11456.9)
	Mix	0.581 (7380.92)	0.575 (31166.4)	0.717 (9077.45)	0.701 (37617)
$N(p, q) = 2 + \frac{\ln(1+p)}{\ln(q)}$ Average price (Average revenue)	1	3.98 (178059)	* (*)	350.83 (1.85 e7)	29.52 (5.49 e6)
	4	* (*)	10.14 (157821)	6435.85 (2.65 e7)	456.01 (7.25 e6)
	Mix	* (*)	* (*)	4199.56 (2.38 e7)	265.38 (6.19 e6)

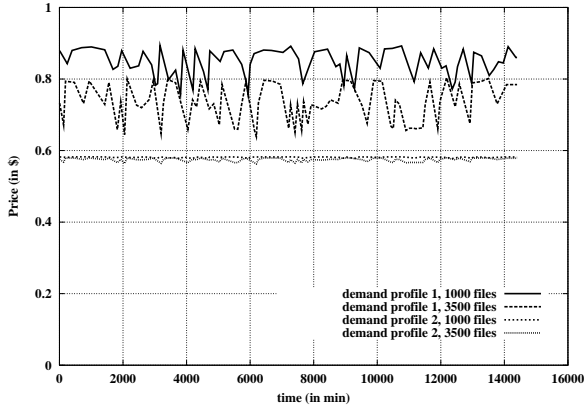


Figure 7: Variation of Price with Respect to Time for Demand Profiles 1 and 2 in a Competitive Market (Ramsey Number = 0.2)




Figure 8: Variation of Price with Respect to Time for Demand Profiles 1, 2, and 3 in a Monopolistic Market (Ramsey Number = 0.8)

a finding that validates our results (logically, a lower price is quoted for high volumes).

Finally, considering the revenues, we notice a net increase by switching from four-hour contracts to one-hour contracts; this finding may be explained by the fact that customers with shorter-term contract lengths free up system capacity rapidly, enabling other customers to buy service contracts. Clearly further detailed investigation of these phenomena is warranted.

All these results point out to the fact that prices are highly volatile from one scenario to the other. Thus, if both the seller and the buyer have some knowledge about the characteristics of the market they are evolving in, better decisions can be made. For instance, a seller would know which type of contracts (short/medium/long) to sell and at what price depending on the buyer. On the other hand, a buyer would know what type of contracts (short/medium/long) to buy and how much volume of data to send.

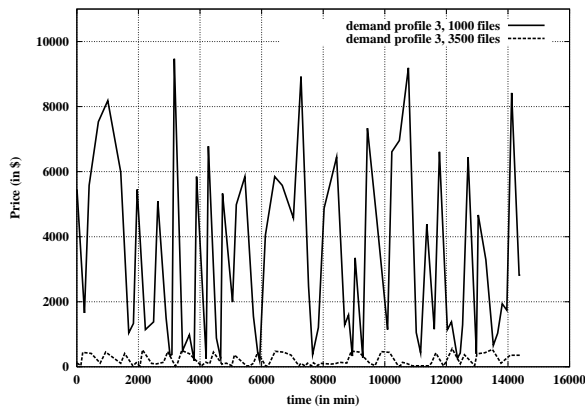


Figure 9: Variation of Price with Respect to Time for Demand Profiles 1,2 and 3 in a Monopolistic Market (Ramsey Number = 0.8)

6 FUTURE WORK

The model presented in this paper can be implemented at the access points of different domains. The price schedules generated at these points, given that the seller collected the required historical data to estimate the parameters used in the simulation, would then be combined together in order to structure specialized services across ISP's. Once creation and pricing of such specialized inter-domain services is worked out, risk management tools such as forward contract, and other derivative contracts, can be explored and designed for data transfer across domains.

The relevance of our model is further reinforced by the fact that the actual technology limits the use of a deterministic approach of QoS guarantees for deriving spot prices as presented in Cheliotis (2001); contracts for Internet bandwidth can at best be of probabilistic type in terms of the QoS specifications of the service contracts (they could be viewed as reliability guarantees).

Finally, the pricing model we introduced requires a complete characterization of the market. For instance, the Ramsey pricing model is primarily developed for a monopoly. In market settings such as an oligopoly, a game theoretic approach will be more relevant.

In summary, we recognize that this paper represents only the first step in designing an implementable inter-domain assured bandwidth provisioning and risk management model and that there is a plenty of work that could be done in this area.

ACKNOWLEDGMENTS

We wish to thank the National Science Foundation which supported this work through contract ANI9819112 and Intel Corporation. We also wish to thank our colleagues, Satish

Raghunath, and Murat Yuksel, for the suggestions they gave while writing this paper.

APPENDIX

Highlights of the current bandwidth market:

Currently, services provided by an ISP in the area of telecommunications are private lines, Internet access, co-location, dark optical fiber and long distance minutes. Below are some of the prices for these services.

- Internet **Access** rates in 200 U.S. cities at the monthly prices of: \$150/Mbps for DS3 (equivalent to 44 Mbps). \$125/Mbps for OC3 (equivalent to 155 Mbps). \$ 99/Mbps for OC12 (equivalent to 622 Mbps).
- T1 (equivalent to 1.54 Mbps) Internet Access in 132 U.S. cities for \$347/month.
- Point of Presence (POP) To POP Private Line Monthly Prices DS3 = \$6,044 OC3 = \$11,216 OC12 = \$30,678 Any of these City pairs: Atlanta, Boston, Chicago, Dallas, Denver, Houston, Kansas City, Los Angeles, Miami, New York, Santa Clara, Seattle.
- OC192 (equivalent to 9.95 Gbps) link from NY to Washington, DC costs \$43,000/month.
- DS-3 link from Atlanta to Seattle costs \$4,845/month.
- OC-3 link from New York City to Los Angeles costs \$9,290/month.

REFERENCES

- Cheliotis, G. 2001. Structure and Dynamics of Bandwidth Markets. PhD dissertation, N.T.U. Athens.
- Crovella, M. E., M. S. Taqqu and Azer Bestavros. 1998. Heavy-Tailed Probability Distributions in the World Wide Web. In *A Practical Guide to Heavy Tails: Statistical Techniques and Applications* with eds. Adler, Feldman and Taqqu. Birkhäuser, Boston, pp. 3-25.
- Crovella, M. E. and L. Lipsky. 1998. Simulations with Heavy-Tailed Workloads. In *Self-Similar Network Traffic and Performance Evaluation* with eds. Park and Willinger. Wiley Inter-Science Publications, pp. 89-100.
- Dolan, R. J. and H. Simon. 1996. In *Power Pricing - How Managing Price Transforms the Bottom Line*. Free Press, NY.
- Gupta, A., D. O. Stahl and A. B. Whinston. 1997. Priority pricing of Integrated Services Networks. In *Internet Economics* with eds. McKnight and Bailey, MIT Press, pp. 323-352.
- Harrison, D., S. Kalyanaraman, and S. Ramakrishnan. 2001. Overlay band-width services: Basic framework and

- edge-to-edge closed-loop building block. Poster in SIGCOMM.
- Kelly, F. P., A. K. Maulloo and D. K. H. Tan. 1998. Rate control in communication networks: shadow prices, proportional fairness and stability. In *Journal of the Operational Research Society*, vol. 49, pp. 237-252.
- MacKie-Mason, J. K. and H. R. Varian. 1995a. Pricing the Internet. In *Public Access to the Internet* with eds. Brian Kahin and James Keller, MIT Press.
- MacKie-Mason, J. K. and H. R. Varian. 1995b. Pricing the congestible network resources. In *IEEE Journal on Selected Areas of Communications*, vol. 13, pp. 1141-1149.
- NLANR. 2002. National Laboratory for Applied Network Research. <<http://www.nlanr.net/NA/Learn/daily.html>>, 2002.
- Odlyzko, A. M. 1997. A modest proposal for preventing Internet congestion. In Technical Report - AT&T Labs.
- Odlyzko, A. M. 1998. The economics of the Internet: Utility, utilization, pricing and quality of service. In Technical Report - AT&T Labs.
- Odlyzko, A. M. 2000. Internet pricing and history of communications. In Technical Report - AT&T Labs.
- Paschalidis, I. and J. N. Tsitsiklis. 2000. Congestion-Dependent Pricing of Network Services. *IEEE/ACM Transactions on Networking*.
- Wilson, R. B. 1993. In *Nonlinear Pricing*. Oxford University Press, Inc., NY.
- Yuksel, Murat, Shivkumar Kalyanaraman. 2002. Distributed Dynamic Capacity Contracting: An overlay congestion pricing framework for diff-serv. Submitted to GlobeCom.
- Yuksel, Murat, Shivkumar Kalyanaraman, Anuj Goel. 2002. Congestion Pricing Overlaid on Edge-to-Edge Congestion Control. Submitted to GlobeCom.

AUTHOR BIOGRAPHIES

MEHDIABOULFADL is a graduate student in the Decision Sciences and Engineering Systems Department at Rensselaer Polytechnic Institute. Currently working with Dr. A. Gupta for his doctoral research, he has a background in General Engineering and Operations Research and Statistics. His areas of interest include derivatives pricing and emerging stock markets. Contact: <aboulm2@rpi.edu>.

RITESH PRADHAN is a graduate student in the Department of Electrical, Computer and Systems Engineering. He is currently working with Prof. Shivkumar Kalyanaraman for his doctoral studies. He earned his BE (EE) from Nagpur University in India and MS (EE) from Rensselaer Polytechnic Institute. His interests are in the area of networking, specifically in ad-hoc networks and Internet pricing. Contact: <pradhr@rpi.edu>.

APARNA GUPTA is an Assistant Professor in the Department of Decision Sciences and Engineering Systems at Rensselaer Polytechnic Institute. Her research interests are in modeling and optimization of stochastic systems, risk management, and financial engineering. Contact: <guptaa@rpi.edu>.

SHIVKUMAR KALYANARAMAN is an Associate Professor at the Department of Electrical, Computer and Systems Engineering at Rensselaer Polytechnic Institute in Troy, NY. His research interests are in network traffic management, traffic engineering, automated network management, multicast, and multimedia networking. His special interest lies in the interdisciplinary areas between traffic and network management, control theory, economics, scalable simulation technologies, wireless technologies, and video compression. Contact: <shivkuma@ecse.rpi.edu>.