EX-POST INTERNET CHARGING: AN EFFECTIVE BANDWIDTH MODEL

Joseph P. Bailey, Ioannis Gamvros, and S. Raghavan
The Robert H. Smith School of Business
University of Maryland
College Park, MD 20742-1815

Abstract Generally Internet Service Providers (ISPs) have charged their customers flat fees for their Internet connections. This has resulted in frequent congestion for many users. There are many different approaches to address this problem. Effective utilization of scarce resources is important to managers in the telecommunications industry, and thus usage-based pricing has become an important tool to address this problem—since it does not require large capital expenditures. In this paper we develop an ex-post charging mechanism based on the effective bandwidth concept. This model, effectively characterizes the utilization and burstiness of a user in a single metric. Further, we introduce a novel market for buffer size. In this market users purchase a specific buffer size from their ISP. Our model directs users with bursty traffic to purchase larger buffers, while users with well-behaved traffic are directed to purchase smaller buffers. From a resource usage standpoint, this is also the appropriate decision. We conduct computational experiments to show the viability of this approach, and also discuss real-world implementation issues.

1. Introduction

Over the past ten years there has been an ongoing debate over the issue of charging Internet traffic (see McKnight and Bailey, 1997). The growing numbers of Internet users coupled with the development of new applications that require large amounts of bandwidth has led to an explosive growth in Internet traffic resulting in frequent congestion that is widely perceived as poor service. More users are growing frustrated by slow connections and increasing packet delays (that result in slow applications like web browsing, ftp, e-mail etc.). Internet Service Providers (ISPs) are trying to solve this problem by over-provisioning (i.e., placing extra bandwidth) in the core of their backbone networks in order to alleviate the congestion experienced. However, there is a growing view amongst a group of researchers that this is a short-term (patch-up) solution that will not solve the problem. These researchers blame

instead the charging mechanisms that prevail in the Internet and insist that the Internet congestion problems can be alleviated to a large extent by using more sophisticated charging algorithms instead of investing heavily in faster routers and extra capacity. Investing in capacity is a significant capital expense which is difficult for telecommunications companies in the current market environment, and so it is critical for ISPs to develop pricing schemes to address this problem.

Although there has always been hope that increasing bandwidth availability will alleviate any need for bandwidth charging, problems of Internet congestion appear to be chronic. The supply and demand for Internet bandwidth appear to be in a virtual cycle whereby increasing supply of bandwidth allows for greater use of bandwidth-intensive applications. Increasing use of bandwidth-intensive applications leads to more demand of bandwidth. For example, the use of Internet Protocol networks to deliver video content is just the latest bandwidth-intensive application. As of 2006, there are a large number of users willing to watch compressed video that would be roughly equivalent to over-the-air broadcast quality. Already there are providers like www.movielink.com and www.cinemanow.com that deliver video movies over the Internet. It is reasonable to assume that demand for HDTV or HD-DVD quality video delivered over the Internet is not far behind! In this way, the evolution of the supply and demand of Internet bandwidth is similar to the evolution of memory and software. No matter how much memory is available today, one could imagine that future applications will inevitably demand even more. Recent scholarly articles continue to point to the fact that Internet congestion is chronic (see Srikant, 2004; Low et al., 2002) even though some argue it is not (Odlyzko, 2003). To help support the claim that Internet congestion is chronic, the most recent statistics that gauge Internet congestion (as measured by packet loss or packet delay, for example), continue to show problems of Internet congestion (see IHR, 2006; ITR, 2006).

In this paper we consider organizational users like small businesses, universities, government organizations etc., that lease a connection for their organization to communicate with the Internet. We believe that this is the right set of users to focus on, since the bulk of the traffic when the network is congested (i.e., during day time on weekdays) comes from these users. At present most of these Internet users are charged a price that is dependent solely on their connection bandwidth. In other words users pay a flat-fee every month to their ISP irrespective of the volume of traffic that they send over their connection. Some researchers like Odlyzko, 2001 are in favor of the status quo because they believe that the simplicity of the flat-fee model is essential and that over-provisioning can be more than a short-term solution. Another option is to charge users based on the actual traffic sent over their connection (i.e., usage-based charging). While another viewpoint by Shenker et al., 1996 is that flat-fee and usage-based charging can co-exist in the same market in the same

way as they do in telephony. Those that are against flat-fee pricing argue that it leads to the “tragedy of the commons” and should be replaced by a smarter charging mechanism. Right now all Internet users (obtaining service from the same ISP) pay the same price for the same connection speed even though their utilization rates of the connection can vary significantly. As a result low-end users end up subsidizing the high-end or heavy users who are possibly willing to pay more for their service.

Proponents of usage based pricing have already done a significant amount of work on the issue of usage based Internet charging (see MacKie-Mason and Varian, 1995; Falkner et al., 1999; Kelly, 1997; Courcoubetis et al., 1998; Courcoubetis et al., 2000). Most of this research has focused on devising optimal pricing strategies that aim at maximizing social welfare. This is achieved by estimating the marginal congestion created when a user sends a packet. The price the users are charged is proportional to the additional congestion generated by the packets they are sending. One of the most notable approaches that uses the marginal congestion paradigm is proposed in MacKie-Mason and Varian, 1995 where a so-called “smart market" mechanism is explained. In this charging scheme packets are assigned bids that are used to determine which packets are given priority. Usually these optimal approaches suffer from high complexity and difficult, if not impossible, implementations that make them unattractive for the real world. The idea of charging based on marginal congestion costs has also been criticized by Shenker et al., 1996. They claim that 1) marginal cost prices may not produce sufficient revenue to fully recover costs, 2) congestion costs are hard to compute, and 3) there are other structural goals of pricing that marginal congestion cost models do not address.

A different way to solve the congestion problem altogether is to make sure that any user that is given access to the network will under no circumstances slow down the traffic of other users. This is achieved by what is known as Call Admission Control (or simply Admission Control) and it involves policing who is connected to the network, what kind of traffic they are sending and either approving or rejecting more connections from other users. In CAC each user makes a request to the network specifying the traffic characteristics (i.e. peak rate, packet loss and acceptable delays) of the data flow he wishes to send. An admission control algorithm then checks the current status of the network to make sure that there are available resources to support the specific data flow with the required Quality of Service guarantees and either admits the user and assigns a charge to the connection or denies admission. If the connection is admitted then the network is required to monitor the traffic that the user is sending to make sure that it complies with the request that was made originally. One of the approaches that uses CAC is Falkner et al., 1999. In both the CAC and the “smart market" approach the charging mechanism is required to know in advance (in the case of CAC) or follow the entire path (in the case

of the smart market approach) that packets take from source to destination in order to assign charges. This may be quite difficult when traffic travels over multiple domains (service providers) to get from source to destination. Consequently, these requirements induce significant overhead and can cause scalability problems.

Unfortunately, proponents of new pricing models and complex admission control policies may never be able to adequately solve Internet congestion. One roadblock is the inability for the Internet to move away from the End-to-End design principles to a “Brave New World" (Blumenthal and Clark, 2001) where competing ISPs can coordinate their activities. Specifically, if ISPs wanted to develop a new “smart market" class of pricing, they would have to develop some settlement process whereby one ISP would reimburse other ISPs for carrying priority traffic. Alternatively, admissions control policies would also have to be closely coordinated in order to implement many QoS solutions. If the market were moving towards more industry concentration, then coordination across so many ISPs would not be a problem. However, ISP backbones have been unsuccessful in their attempts to integrate, in part because of merger guidelines that appear to be too stringent on defining market power (Besen et al., 2002). When there are multiple networks responsible for the transmission of packets, it is difficult to implement an End-to-End pricing scheme for a few reasons. First, it would require all of the involved networks to adhere to the same policy. This is very difficult because these networks are not only competing for the same customers, but they have an incentive to provide better service to their customers in preference over their competitor’s customers. Second, an End-to-End pricing scheme may be ripe for opportunism whereby an ISP can try to enhance its settlement money. For example, they may support larger routing tables to carry more traffic and it may even send the traffic over a greater number of hops within its network to increase its portion of the settlement. Therefore, some of the most promising solutions to Internet congestion are the ones that embrace, rather than abandon, the End-to-End design principles of the Internet.

Internet congestion may be reduced by a class of charging mechanisms that assign prices based only on information collected at the ingress of the network, where the user’s packets enter. This paradigm is termed “edge pricing" (see Shenker et al., 1996) and it works by monitoring the packets that users send over their connection either constantly or at given intervals. While monitoring, the charging algorithms determine the traffic characteristics of different users and in return are able to estimate the network resources utilized by these users and the congestion they impose on others. Based on this information, charges that are proportional to the resource usage of each user are assigned. Edge pricing does not entail the risks and difficulties of the CAC or smart market approaches, but imposes the challenge of estimating resource consumption based on local information at the ingress point of the network. In many cases this challenge

is met with the use of effective bandwidth bounds (Kelly, 1997; Siris et al., 1999) that give good estimates of a users actual resource usage of the ingress connection. In these charging mechanisms users declare a utilization rate at which they will send data over their connections. If they respect this rate then they are charged according to the estimated effective bandwidth. However if their actual rate is different (even if it is lower) from the stated then they get penalized by paying more than what the effective bandwidth calculation indicates.

In this paper we develop a novel model for charging Internet connections based on effective bandwidth. This model falls under the class of the so-called "ex-post charging" models (see Bailey et al., 2006) where the pricing algorithm is determined ex-ante but the charges are determined after the traffic has been sent. Our effective bandwidth model is quite simple, and differs from other effective bandwidth models in the literature in several respects. First, we use the large buffer asymptotic method for calculating effective bandwidth (Guérin et al., 1991). As a consequence, unlike other effective bandwidth models used for pricing, we do not need to consider other sources of traffic to determine the charge for the traffic. This might seem a disadvantage of the model at first, since any possible multiplexing gains are not calculated, but in fact is an *extremely desirable property*. This is because (i) the charge is dependent solely on an individual user's traffic, and (ii) it can be calculated by the user without knowing any other users traffic (and thus can manage their traffic and charges without worrying about the affect of the behavior of other users on their charge). Second, one of the parameters that the effective bandwidth depends upon is a buffer size. We develop a market for buffers where ISPs charge users for buffer space, and based on this develop a coherent pricing model. Finally, our model satisfies a desirable feature that the ex-post charging mechanism has—namely Bayesian updating of parameters. This means terrabytes of traffic information need not be stored to determine the charge for the traffic. This is an important and critical issue, that seems to have been largely ignored in the literature. By ensuring that terrabytes of data need not be stored to implement the pricing mechanism (i) it is more likely to be accepted (ii) cheaply implemented, and (iii) removes a potential security risk associated with storing trace data.

In the rest of this paper we develop our effective bandwidth based pricing model. The remaining sections are organized as follows. In the rest of this section we review the ex-post charging model and philosophy. In Section 2 we will review the large buffer asymptotic model for effective bandwidth and the upper bound based on it that we use in our model. In Section 3 we develop our effective bandwidth charging model, introduce a market for buffers, and discuss issues concerning the fine tuning of the model to the needs of different ISPs. In Section 4 we present numerical results that showcase the performance of our pricing algorithm under different scenarios, and illustrate the behavior

of the pricing algorithm. Finally in Section 5 we present our conclusions and suggestions for future work on this area.

1.1 The Ex-Post Charging Approach

As we stated at the outset, ex-post charging may be most suitable between an ISP and organizational customers such as small businesses, universities, and government organizations. These users currently shape or manage their traffic, and are most concerned about their quality of service and of lowering their large Internet connectivity bills. These users are likely to benefit on both counts from the ex-post charging policy we propose and embrace it. On the other hand mass-market users (like residential customers) currently appear to prefer flat rate pricing (as evidenced by the shift in pricing schemes in the mobile and long-distance market). Interestingly, it is precisely for these customers with small individual revenues (in the $20-$50 range) the cost of calculating and metering usage makes usage-based pricing a costly proposition for the ISP.

The ex-post charging mechanism falls under the category of "edge-pricing" algorithms. In this model the charging algorithm is determined in advance while the actual charge is calculated after the fact. We note that the actual charging mechanism that might be used in practice will in fact consist of an ex-ante charge (i.e., a charge determined in advance) as well. So in essence the final price P can be viewed as:

$$P = P_{\text{ex-ante}} + P_{\text{ex-post}} \tag{1}$$

The ex-ante part of the price can be used as a mechanism that will prevent users from reserving connections that they don't really need. If there was no such component and the price depended only on resource usage a customer would be able to ask for multiple connections, not send any traffic over them, and pay nothing. Although, we do not study the ex-ante part of the price in this paper it plays an important role as well.[1] For example, it may affect the ability of an ISP to attract new customers in a competitive setting. However, when considering that an ISP has already made a contract with a customer, the ex-ante portion of the charge is sunk and will not affect a customer's incentive to manage its Internet traffic any differently. As we will not consider the ex-ante portion of the price within this paper, from now on we will use the term price to refer to the ex-post price.

For any new charging model to be accepted, and successfully implemented, in the current ISP marketplace, we believe there are two key desirable features—

[1]For example, the ex-ante portion of the price may typically cover the cost of maintaining a connection. Users with large bandwidths may thus have a higher ex-ante portion of the price.

simplicity, and Bayesian updating of parameters—as described above. Additionally, Bailey et al., 2006 identify the following desirable qualitative characteristics of an ex-post Internet charging model.

- The ex-post charge should be a monotonically increasing function of the total volume of traffic sent (and/or received). Utilization measures the volume of traffic divided by the speed of the connection times the duration (over which the volume of traffic is sent). Consequently, the ex-post charge should be a monotonically increasing function of utilization. Further, since it is likely that the provider would probably want to offer a service in which economies of scale are realized it is desirable for the relationship between the ex-post price as a function of the measured utilization to be concave.

- Burstiness expresses the notion of sudden, unpredictable and usually large transmissions of data from the customers to the provider. Bursty traffic can be problematic for an ISP because an ISP must either size their network large enough to accommodate peak periods or be willing to endure periods of congestion during peak periods. Consequently, bursty traffic should be charged a higher price than well-behaved traffic (i.e., if two traces have the same utilization but one is burstier than the other it should be charged more). Additionally, the relationship between the price and the measured burstiness should be a convex function. This corresponds to the notion that the effects of bursty traffic on a provider's network can have an additive effect resulting in prolonged network congestion as customers send more bursty traffic. We note however that burstiness is not a well defined metric like utilization. Consequently, it may not be easy to analytically verify whether a pricing model complies with this desired feature.

- Finally, the implementation of the charging mechanism should be transparent to the network. By transparent we mean that the algorithm should require very few or absolutely no network resources (e.g. bandwidth, CPU time, storage space) to complete its task.

2. Theoretical Background for Effective Bandwidth

In order to be able to charge customers for the use of a communications link we need to be able to identify scalars that will measure the resources they use when their packets are forwarded over the Internet. These scalars will then become the independent variables of a pricing function that will associate resource usage with a specific charge. Utilization of a traffic stream is a well defined metric and is easy to measure. However, burstiness is not so well defined. The effective bandwidth concept ties these two notions together, and summarizes

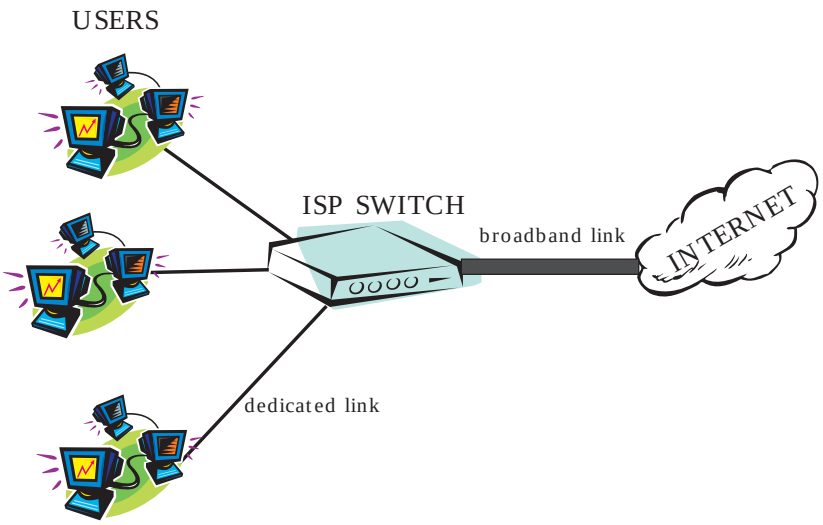

Figure 1. Multiplexing of many sources on an outgoing broadband link.

resource usage of a shared communications link by a specific source. We now review the effective bandwidth concept, and in particular the large buffer asymptotic method for computing effective bandwidth proposed by Guérin et al., 1991.

Effective Bandwidth is a scalar that summarizes resource usage on a communications link from a specific source in a packet-switched network. Specifically, at a given switch in the network where many traffic streams from different sources are multiplexed on a single outgoing link (Figure 1), the effective bandwidth of a specific source represents the capacity of the outgoing link used by that source.

It turns out that the effective bandwidth of a specific source depends not only on the statistical properties of the traffic source in question but also on the statistical properties of the other sources that it is multiplexed with, the characteristics of the switch (i.e., buffer size) and the characteristics of the link that the source utilizes (i.e., capacity). Moreover, effective bandwidth depends on the Quality of Service requirements (i.e., packet loss probability), which are imposed by the source.

A well known and widely accepted way to calculate the effective bandwidth of a specific traffic source is proposed by Kelly, 1996. The calculation proposed there takes into account all of the parameters that we mentioned previously and provides a very good estimate of network resources used by each source. However, Kelly's calculation requires the collection of large amounts of data that represent the traffic each source sends to the switch. This is somewhat

impractical from an individual users perspective, since to determine their charge for the connection they would need to know the traffic of other users. As a result instead of Kelly's model we will use an upper bound for effective bandwidth that was proposed by Guérin et al., 1991 and is easy to calculate.

Guérin et al. use what is called the large buffer asymptotic method to arrive at an *upper bound* for the effective bandwidth of a single source. This method is concerned with the overflow probability of the buffer at the switch as the buffer size increases. Additionally Guérin et al. do not take into account the traffic characteristics of other sources that send data to the switch. At first this seems to be a significant disadvantage of the calculation as it completely ignores possible gains in resource usage from multiplexing. However, for the purposes of pricing this model is *ideal* since the scalar that represents the resource usage for a specific user depends solely on the user's traffic and not the combined traffic characteristics of all the users that send traffic to the network. In simpler terms each customer is charged based solely on their individual traffic characteristics without considering traffic characteristics of other users. Since we are interested in using the effective bandwidth concept for pricing and not for traffic engineering the upper bound on the effective bandwidth is adequate as a measure of resource usage for our purposes.

2.1 Effective Bandwidth of a Single Source

In the following we review Guérin et al.'s large buffer asymptotic model for effective bandwidth (see Guérin et al., 1991). They assume that the traffic sources can be in one of two possible states at a time. Either the source is in a "Burst State" which means that it is sending data at the maximum rate of the connection or it is in an "Idle State" which means that there is no transmission. This assumption actually states what is happening on communication links that are utilized on the Internet. In order to be able to fully characterize the traffic source one needs to know the distributions of the "Burst State" periods and the "Idle State" periods. The second assumption that they make states that the length of the "Burst State" and "Idle State" periods are exponentially distributed. As a result they can be completely characterized by their means (i.e., mean length of the "Burst State" and "Idle State" periods).

Consequently, a traffic source can be fully identified if we know the peak rate (R_p) at which it can transmit, the mean of the "Burst State" periods (b) and the mean of the "Idle State" periods. Observe that given the mean of the "Burst State" periods and the mean of the "Idle State" periods one may calculate the utilization (ρ) as the mean of the "Burst State" periods divided by the sum of the means of the "Burst State" and "Idle State" periods. As a result given the source's utilization (ρ) and mean of the "Burst State" periods the mean of the "Idle State" periods can be computed. Therefore a traffic source can be fully

identified if we know the peak rate (R_p), the mean of the "Burst State" periods (b), and the source's utilization (ρ).

We are interested in calculating the capacity (C) in bits per second (bps) that for a specific buffer size (B) guarantees a buffer overflow probability less than ϵ. The capacity (C) is the effective bandwidth of the source and it means that the outgoing link shown in Figure 1 should have at least C bps reserved for the traffic source in question in order to be able to support the source's traffic for a specific overflow probability and the given buffer size. Guérin et al. show that an upper bound on C is given by the following equation:

$$C = \frac{\gamma b(1-\rho)R_p - B + \sqrt{(\gamma b(1-\rho)R_p - B)^2 + 4B\gamma b\rho(1-\rho)R_p}}{2\gamma b(1-\rho)} \quad (2)$$

where $\gamma = \ln(1/\epsilon)$.

This equation provides us with an estimate of the actual effective bandwidth of the source. Numerical experiments have shown that the value of C calculated by Equation 2 is very close to the exact value of effective bandwidth (see Guérin et al., 1991). With the help of Equation 2 we have a good approximation of the resource usage of a specific source and in turn of a specific customer.

2.2 Implementation Issues

We now show that the variables that are needed for the calculation of the effective bandwidth are readily available or can be easily measured. Specifically, the peak rate R_p of the connection is known in advance, the buffer size B that resides in the switch where the link that carries the customer's traffic is connected can be easily verified and the packet loss probability that the customer requests is agreed upon in the service level agreement. That leaves us with the actual measurement of the mean burst period b and the utilization ρ.

In order to calculate these values one needs to know the size of the packets that the source has transmitted and the time of their arrival at the port of the switch. Once the arrival times and the sizes of the packets are known, the mean burst period and the utilization can be calculated as follows. Utilization is given by the sum of the packet sizes (total volume of traffic) divided by the connection speed times the total period of measurement (i.e., divided by the maximum possible traffic that could have been sent on the connection over the duration of measurement). In order to calculate the mean burst period we have to determine consecutive packets that went over the connection in a burst (i.e., these packets were sent one after the other with no idle time between them). We define a series of sequential packets arriving over the span of a millisecond to be in the same burst. This assumption is supported by the fact that most measuring equipment cannot discriminate between arrival times under the millisecond

level (see Mills, 1989).[2] We set the size of the burst equal to the sum of the sizes of all the packets in the burst. Also we set the arrival time of the burst equal to the arrival time of the first packet in the sequence. The mean burst period of the entire data trace can then be calculated by first calculating the mean size of a burst as the sum of the sizes of the bursts (which is equal to the total volume of traffic) divided by the number of the bursts; and then dividing the mean size of a burst by the connection speed.

It is important to point out here that there is no need to store large data files that contain the above information. For every new packet arrival the mean of the burst period and the utilization can be updated (since the mean of a time series can be updated in a Bayesian fashion) resulting in a new value for effective bandwidth. As a result the storage requirements for the calculation are minimal and the resource measurement can be done in a meter like fashion in a similar way as with utilities such as electricity and natural gas.

3. The Effective Bandwidth Ex-Post Charging Model

We now build an ex-post charging model based on effective bandwidth.

3.1 The Simple Effective Bandwidth Model

Since effective bandwidth is a good measure of resource usage it makes sense that a charging model could in fact consist of just the value of the effective bandwidth and not take any other variables into consideration. A possible model would look something like this:

$$P = a * C \tag{3}$$

where P is the ex-post price, C is the value of the effective bandwidth calculated in Equation 2 and a is a variable that can change during the billing period and its purpose is explained below.

The basic characteristic of this charging model is its simplicity. The price changes linearly with effective bandwidth and as a result it directly corresponds to a user's resource usage. The variable a is used to convert C into monetary units and its units are dollars per bps per unit of time. In addition a reflects congestion in the network and contention among the users for the scarce resources. As a result when the demand is high the ISP can increase the value of a whereas at times of the day when the demand is low the ISP can accordingly reduce a in order to attract more users. We note further that the model can be used to calculate the price on whatever time interval basis the ISP and/or user agree upon. For example, the effective bandwidth may be computed on an hourly

[2] It is possible to get measurements that are accurate up to a microsecond level (see Micheel et al., 2001) but that requires more sophisticated techniques and equipment.

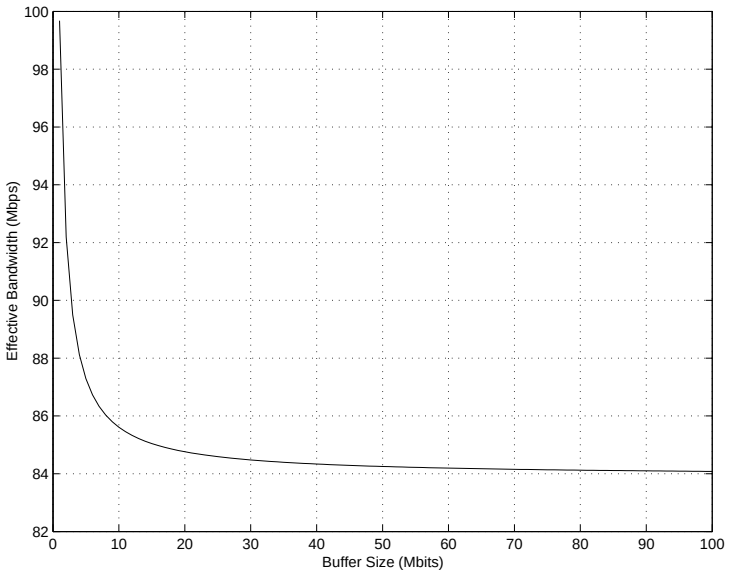

Figure 2. The effect of buffer size on effective bandwidth

basis (for the traffic sent to (and/or received) the ISP in the past hour) and the total charge for that hour may be determined using Equation 3. Although the value that a assumes is extremely important it is beyond the scope of this paper and we will not discuss it any further.

There is only one real disadvantage with the simple effective bandwidth model. The problem is that the value of effective bandwidth and consequently the price that customers pay for the service depends on the buffer size that is reserved for them. Everything else being equal (i.e., utilization, mean burst period, peak rate) the effective bandwidth value decreases with an increase in buffer size. This is shown in Figure 2 for a 90 second trace captured on an OC-3 link.

Looking at Figure 2 (and Equation 2) it should be clear that the choice of buffer can significantly affect the calculation of effective bandwidth and the price calculated by the model. As a result the customer will always want to have as large a buffer as possible while the ISP would prefer *exactly the opposite* since smaller buffers would mean higher prices and higher revenues. Additionally, buffer size is an important resource that can play a critical role in terms of the quality of service that the customer receives. Customers with large buffers can send large bursts of data to the ISP and be sure that their packets are not going to be dropped. As a result, we believe, buffer size should be taken into account in any charging mechanism for any packet switched network. Further,

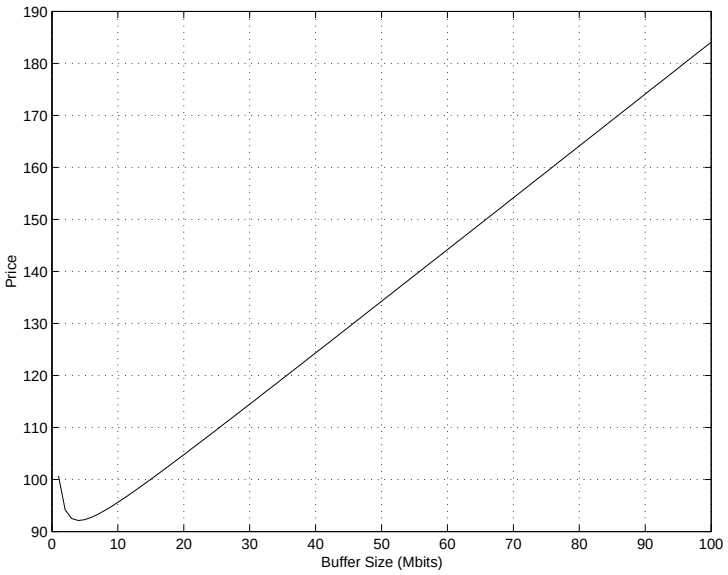

Figure 3. The calculated charge is overwhelmed by the choice of the buffer size

buffer is a resource provided by the ISP, and thus a cost to the ISP, and so we believe that it should play a role in the pricing model.

The next charging model that we propose resolves the shortcomings of the simple effective bandwidth model as it takes into account the buffer size that is reserved for the customer's traffic as well as the effective bandwidth calculated for that buffer size.

3.2 The Delta Model

In this model we assume that there is a market for buffer sizes and the customers are able to select buffer sizes that are consistent with their traffic demands. For example the model should direct customers with bursty traffic to choose larger buffers while customers with well-behaved traffic should be rewarded for selecting a smaller buffer. The Delta charging model that we propose is based on these assumptions and has the following form:

$$P = a * (\Delta * B + C) \tag{4}$$

where P is the ex-post price component in dollars per unit of time, C is the effective bandwidth in bps, B is the buffer size in bits and a is a variable that has the same role as in the previous model.

Delta (Δ) is a scaling constant that has as a primary purpose to balance the effect of the two resources B and C on the price. This is done because the buffer size can assume values that are comparable and some times even greater

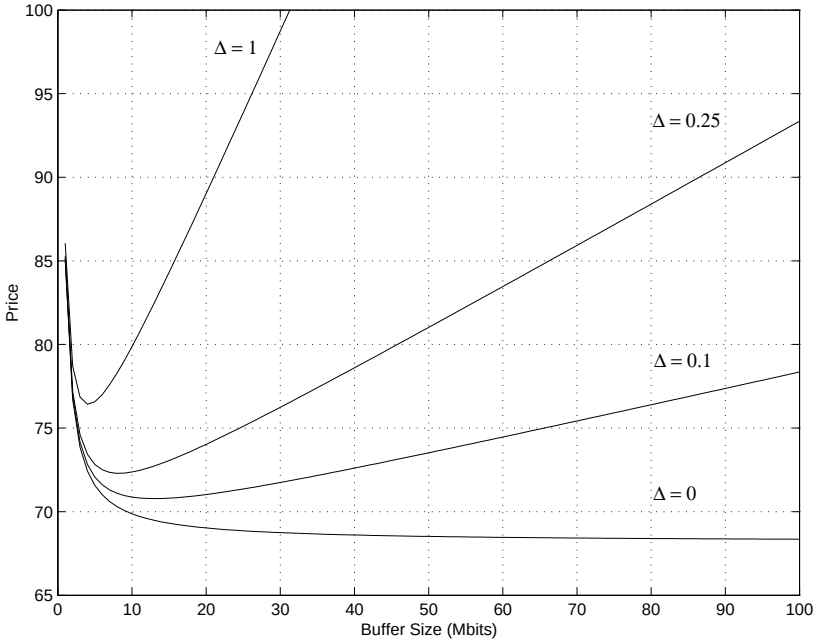

Figure 4. The effect of Δ on the calculated price

than the value calculated for C. If there was no scaling constant then the value of B would overwhelm the ex-post pricing component (as shown in Figure 3 for the same trace as in Figure 2) and lead to charges that would be based mostly on the selection of the buffer size and not the resource usage as it is represented by the value of C. Clearly, this was not our intention when we introduced B into the charging model and using a scaling constant Δ allows us to overcome this problem.

3.2.1 Effect of Delta on the Model. By comparing Figure 2 for which $\Delta = 0$, and Figure 3 for which $\Delta = 1$, it should be clear that Δ can significantly influence the shape of the charging model. In order to better demonstrate the role of Δ in our model we have calculated the prices that would be generated by our charging model for a specific trace but for different values of Δ for a wide ranging choice of buffer sizes.

Figure 4 shows that when Δ is small or zero the buffer size does not affect the price calculated by the model. (Note, this is for a different trace than used in the previous figures. All traces used in the figures are identified in the Appendix.) However as the value of Δ increases buffer becomes all the more important. Specifically for smaller values of Δ it seems that effective bandwidth is dominating the price while for larger values it is the buffer size

that has the major effect. Ideally Δ should have a value that balances the effect of the two parameters that control the price. In the following section we discuss extensively what we consider appropriate values for Δ and how they can be calculated.

3.2.2 Setting Delta. Since Δ was introduced in order to balance the effect of the two resources used by the customer we suggest setting it equal to the ratio of differences of these two resources as follows.

$$\Delta = \frac{|C_H - C_L|}{B_H - B_L} \tag{5}$$

where B_H and B_L are the highest and lowest values respectively that the buffer size can assume for a specific type of link; and C_H and C_L are the corresponding values that effective bandwidth assumes (for the trace) when B_H and B_L are used as the buffer sizes. In our computations (see Section 4) we assumed that the possible buffer size values will be in the range [$0.01R_p$, $0.9R_p$], where R_p is the peak rate of the incoming link (i.e., the link capacity in terms of bps). Actually, the exact bounds of this range are not that critical as long as the actual buffer sizes available to the customer fall into that range.

On further examination of Equation (5) it should be evident that to calculate Δ one needs not only the link capacity and the values selected for B_H and B_L, but also the specific trace for which C_H and C_L are calculated. If the ISP uses the trace that it is applying the ex-post charge to, in order to determine C_H and C_L, then the ISP will be able to calculate Δ only *after* the end of a billing period. However, this is contrary to the ex-post pricing concept where the charging model must be specified explicitly in advance so that the customers will be able to estimate their charges based on the traffic they send over the network (and thus be able to manage their traffic to potentially lower their charges). Also it would be fairer from the customer's point of view if the ISP offers to everyone using a link of a given capacity the same Δ. Moreover selecting a constant Δ will help the customers plan ahead and select an appropriate buffer size for their connection. If Δ were to change frequently (for example within every billing period) then it would be very difficult, if not impossible, for a user to determine the right choice of buffer size as well as to manage their transmissions in order to minimize their charges. Consequently, we impose the condition that Δ has to be constant for a given link type.

In Section 2 we mentioned that in order to calculate effective bandwidth one needs to measure the mean of the burst periods and the utilization (since all the other variables are known). Looking at Equation (5) we see that the only unknowns are the two values of effective bandwidth (i.e., C_h and C_L) in the numerator. As a result, since the value of effective bandwidth (for a given buffer size and packet loss probability) depends on only two variables utilization (ρ) and mean burst-period (b), we see that Δ for a specific link capacity actually

depends on only utilization (ρ) and the mean burst-period (b). Consequently, we suggest that they be set equal to the utilization and mean burst period that the provider considers an average “well-behaved" customer would have. This provides a uniform value of Δ for a given link type.

As an example we will now calculate the price that a “well-behaved" customer would have to pay with respect to the buffer size chosen, for an OC-3 link (capacity: 155Mbps). We calculate these prices using the Δ value set with Equation (5). Specifically we assume the “well-behaved” customer has a utilization of 35% and a mean burst period of 3.5E-04 sec. The choice of an acceptable utilization was based on conversations with our contacts in industry. For the mean burst period however things were somewhat more complicated. We mentioned earlier that we treat bursts as consecutive packets that were transmitted in the same millisecond. In the time frame of a millisecond an OC-3 link can transmit 155Kbits. However, well-behaved customers that transmit traffic only 35% of the time will probably send on average: 35%*155Kbits per millisecond. This amount of data (i.e., the burst) will be transmitted in: 35%*155Kbits/155Mbps = 0.35*0.001 sec = 3.5E-04 sec. Several experiments we have done on trace data indicate that the measured mean of the burst periods, that were calculated for different traces, are greatly dependent on utilization and could be in fact be approximated in the above way. So by using the utilization and mean burst period values mentioned we were able to determine Δ and also calculate the price that a “good" customer will be charged.

We plot the price of a well-behaved customer as a function of buffer size in Figure 5 (using the same trace as in Figure 4). One can see that neither effective bandwidth nor buffer size dominate the price. (We note that in the figure there are no specific monetary units associated with the price.) Specifically, we can see that for very small buffer sizes the price is high since the value that effective bandwidth assumes will be significantly large. However, as the customers choose larger buffers they are able to reduce the price they pay since effective bandwidth falls sharply. Nevertheless, choosing a very large buffer proves to be inefficient since buffer also affects the price and as a result there is a high charge associated with that choice. This behavior is in accordance with the points that we made earlier on the characteristics of a good charging model. Customers will want to select a buffer that will correspond to the minimum of the graph so that they can minimize their costs. In order to achieve this, customers will have to be aware of the type of traffic they are sending because as we will experimentally show in the next section the behavior of their traffic will shift the minimum of the curve.

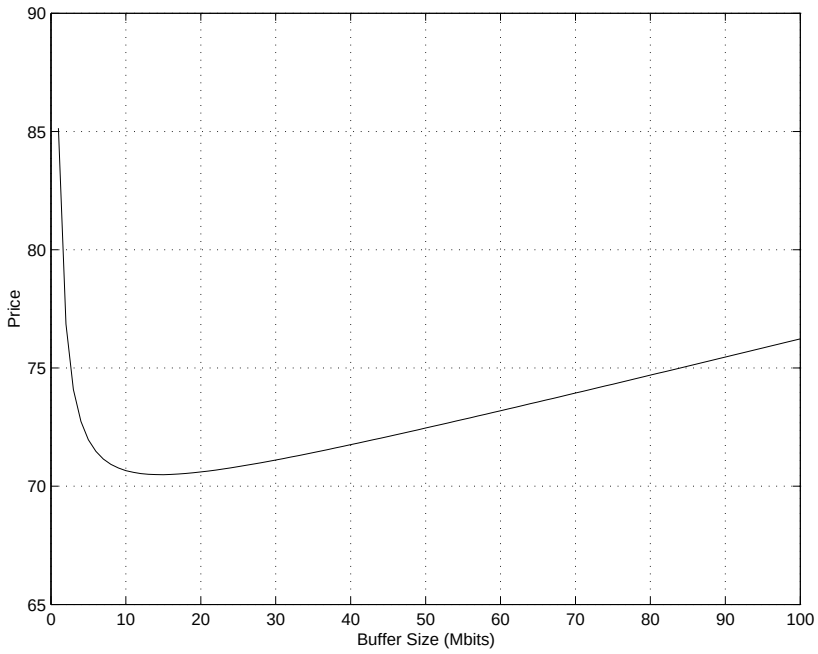

Figure 5. The price calculated for a well behaved customer

4. Numerical results

In order to test the proposed models we used real world traces captured on different networks. The traces consist of the IP headers of the packets captured on specific links. From the IP headers we were able to extract the byte count (i.e., number of bytes of the packet) and the timestamp (i.e., the time the packet arrived at the switch and was captured by the metering software) of the captured packets. These two variables, timestamp and byte count, are the only inputs required for the calculation of the effective bandwidth bound we use.

In the following sub-sections we will present results generated by using traces obtained from the National Laboratory for Applied Network Research (NLANR) repository (http://www.nlanr.net/Traces/). In our experiments we used many of the short 90 second traces provided by the Passive Measurement Project team of NLANR (http://pma.nlanr.net/PMA/) and some of the day-long traces that were captured on the New Zealand to US link. The short traces were originally used to explore the reaction of our models in different contexts while the long traces were used to verify the consistency of our approach in real world settings. Below we present some results obtained with the use of a few selected short traces, in order to demonstrate the behavior of our charging model (the particular trace(s) used in each figure are identified in the Appendix).

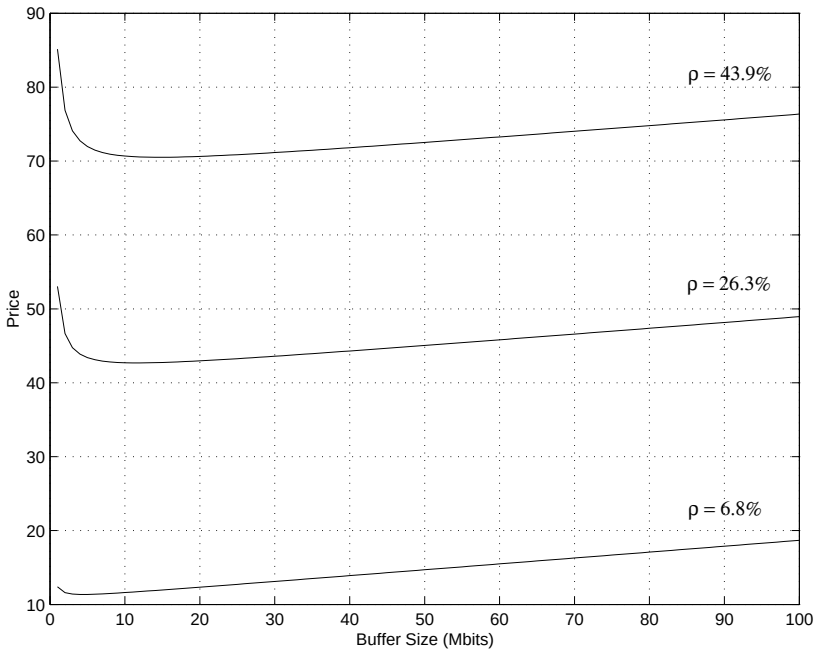

Figure 6. The effect of a customer's utilization on the calculated price

Utilization. Utilization (ρ) represents the percentage of time that the customer is sending traffic over the connection. It is probably the most important of the customer's traffic characteristics. A higher utilization value means that the customers are sending more traffic and consequently they should pay more. Figure 6 presents the results produced by the Delta model for three different traces captured on the same link at different time intervals. The graph shows the price calculated with the same Δ for different values of buffer size for the three different traces. From the graph it is evident that the model performs consistently in the sense that increased utilization is penalized by higher prices.

Packet Loss Probability. Packet loss probability determines the average number of packets that are lost over the customer connection during a billing period. Lost packets can occur because of bit errors that might corrupt the header of the packet. A switch that sees a corrupted header drops the packet because it cannot trust the information in the header. In addition to this, packets might get dropped when they reach a switch and the buffer of the switch is full so there is no room to store the packet. When we discuss packet loss probability in this paper we refer only to the latter case where a packet is dropped because of a buffer overflow. Depending on the size of the link or a customer's specific needs different packet loss probabilities might be requested from the provider.

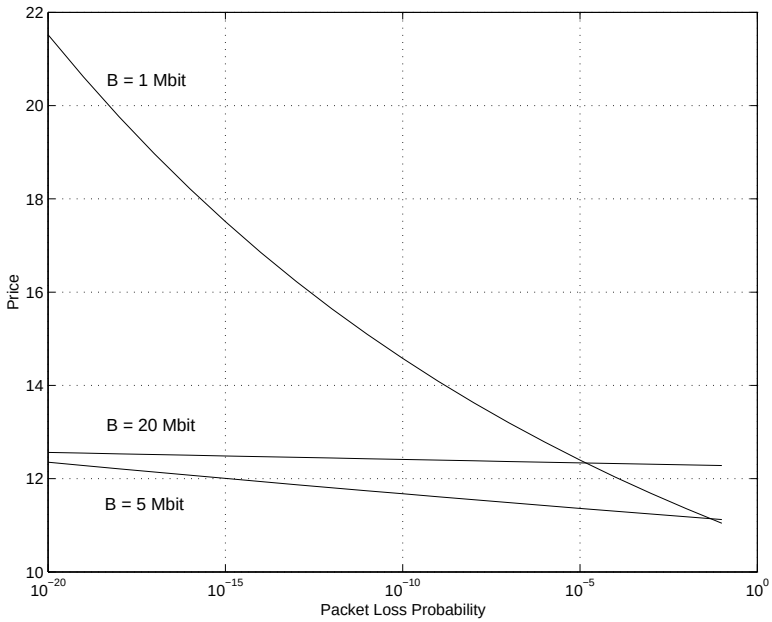

Figure 7. The effect of packet loss probability on the calculated price for a low utilization customer ($\rho = 6.8\%$).

The following graphs provide insight as to how different customers with varying packet loss probability requirements will be charged by our model.

Figure 7 shows the effect of packet loss probability on the calculated price for three different choices of buffer size. We can see that for smaller buffer sizes customers that are requesting lower average packet loss have to pay a higher price. However if the customer has already opted for a higher buffer then the increase is significantly smaller. Figure 8 provides the same information for a different trace with significantly higher utilization, 43.6% as opposed to the utilization of the trace used in Figure 7, 6.8%. For the higher utilization trace the slopes of the lines remain roughly the same. The only difference is the lower price of the 20 Mbit buffer choice with respect to the other two choices. This occurs because for higher utilization the minimum price with respect to buffer size occurs for larger values of buffer (see Figure 6). As a result small values of buffer will generate higher prices.

Price vs. Burstiness. As we have already pointed out customers with bursty traffic should be charged more than customers whose traffic is well behaved. Moreover the model should motivate customers with bursty traffic to select larger buffer sizes since these buffers will make sure that when large bursts occur there will be little or no packet loss.

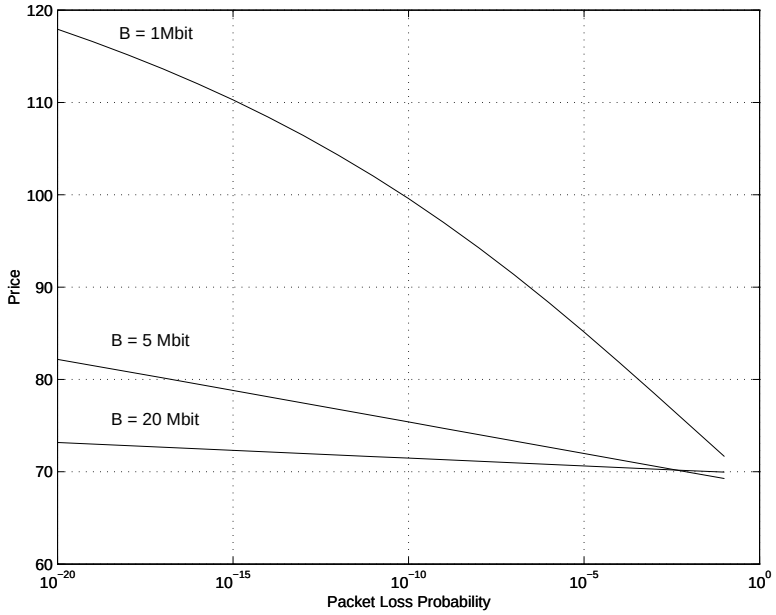

Figure 8. The effect of packet loss probability on the calculated price for a high utilization customer ($\rho = 43.9\%$.)

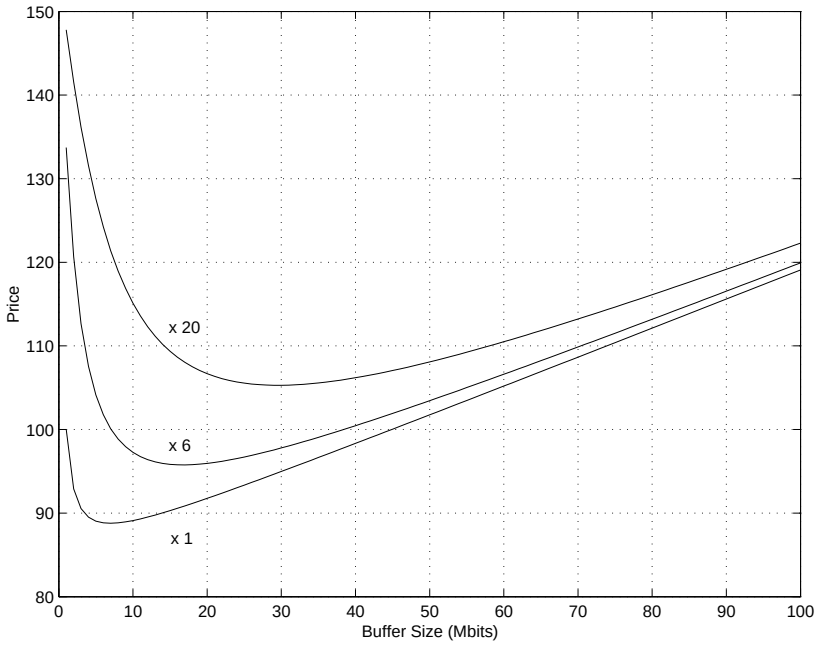

Figure 9. The effect of burstiness on the price curve

Trace	Buffer Size (Mbits)
x 1	7
x 6	17
x 20	30

Table 1. Optimal buffer sizes for different traffic behaviors

Figure 9 shows how the price changes for a specific trace that has undergone special manipulation so that although it's utilization remains constant it's burstiness is increased. In order to achieve this we aggregated consecutive packets so that they appear as if they were one. This way we keep utilization constant since we are not inserting packets to the trace while at the same time we force the packets to arrive at bursts. The more we aggregate the more bursty the traffic becomes. The multipliers in Figure 9 indicate the number of sequential packets aggregated (i.e. "x 1" corresponds to the original trace,"x 6" corresponds to a trace created from the original by aggregating every 6 packets into one and so on).

From Figure 9 one observes that apart from the overall increase in price each curve reaches the minimum point for a different buffer size. As the traffic becomes burstier customers will be directed towards larger buffers in order to reduce their costs. In Table 1 we can see the buffer sizes that correspond to the minimum point for each of the manipulated traces.

5. Final Remarks and Conclusions

By looking at the various experiments in the previous section it is evident that the charging model that we proposed behaves in accordance with the desirable properties of a pricing model that we specified earlier in this paper. It consistently penalizes customers with high utilization and/or bursty traffic and charges higher prices to those who seek better packet loss guarantees for their traffic. Although, for brevity, we presented limited results in the previous section to demonstrate the behavior of our charging algorithm we have actually conducted an extensive set of experiments of a similar nature on a wide range of trace data and for links that varied from 10Mbit Ethernet buses to OC-12 fiber optic carriers. The model has behaved consistently in all of these cases.

To get a better assessment of our charging model we would have liked to test this charging model in practice at an ISP. This will enable us to understand better many of the real-world aspects of implementing this charging model, as well as to observe user (and ISP) behavior in response to such a charging mechanism. We think it likely that an ISP is better able to shape the "long tail" of customer preferences through an ex-post charging approach. Rather than an ISP having to build its network for peak congested periods, the ex-

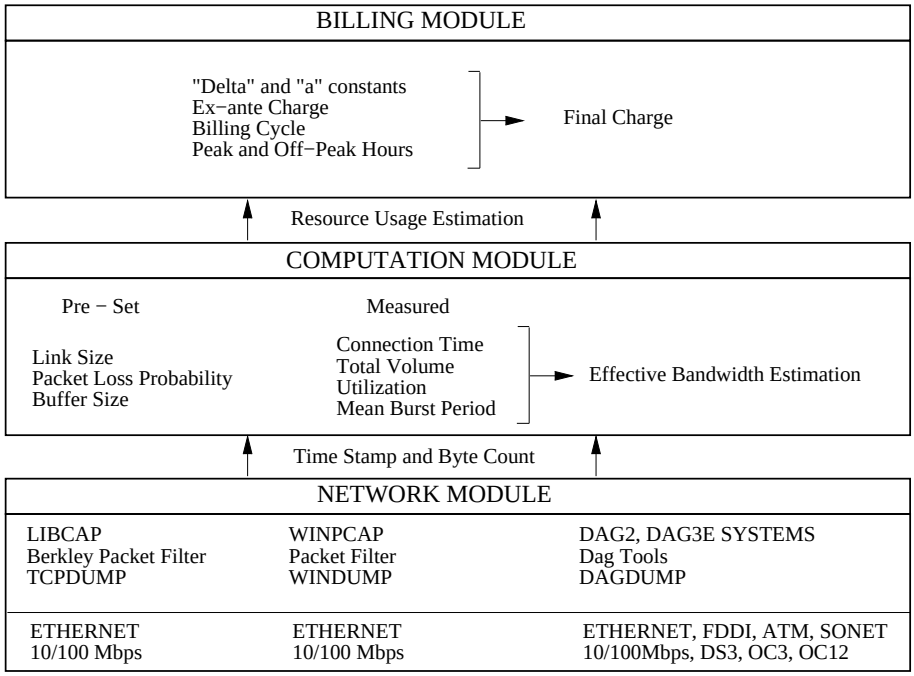

Figure 10. Architecture for implementing ex-post charging at an ISP.

post charge provides an incentive for customers to help shape and manage their network traffic before it is received by the ISP. Although the market has not yet achieved a level of sophistication to immediately implement our ex-post charging mechanism, there is some evidence that service level agreements are more complex and address some of the incomplete contracting problems that had previously existed. To further understand how the ex-post charging model works in practice, we have been discussing our charging model with service providers. As a consequence of those discussions, we elaborate briefly as to how our charging model can be easily implemented by ISPs with no significant cost.

The packet capturing architecture behind our charging model that the ISPs are required to implement can be seen in Figure 10. At the lower level of this diagram the “Network Module" is responsible for capturing the packets and reporting the time they arrived and their size. It’s essential components are a network adapter (e.g. a simple Ethernet Card in the case of Ethernet connections or an ATM card in the case ATM connections) that receives the user’s packets, an appropriate driver (e.g. tcpdump, windump) that can process all the packets on the physical link and a set of libraries (e.g. winpcap, dag2) that will be used

as an interface with the module. Fortunately all these components are readily available and well documented so there is no developing work to be done on the ISP's part. The output that is required from the "Network Module" is the packet size and arrival time of the different packets. However instead of just capturing packets it is possible to filter packets as well. Filtering might be a desirable feature if the ISP wants to charge different prices for different kinds of traffic such as TCP and UDP. In that case the "Network Module" would have to explore a packet's header, determine the protocol used for its transmission and report the appropriate values to the higher modules.

At the middle level the "Computation Module" receives the packet information and is responsible for processing that information to determine a user's connection time, utilization, mean burst period and total bytes send. These are the values that are required for the calculation of the effective bandwidth bound that is required by our charging mechanism. This model can easily be implemented in software and calls upon the network module libraries that we mentioned previously so that it can retrieve all the required information. It is important to note that the values calculated here can be updated continuously with the arrival of every new packet in a Bayesian fashion. As a result there are no extraordinary space requirements. Moreover the actual code that can be used to make these calculations can be only a few lines (depending on the platform and coding language). The output of this module is the estimate (based on the bound) of the effective bandwidth of the captured data trace.

Finally at the higher level of this architecture we find what we call the "Billing Module". This will probably be part of an ISP's usual billing system where necessary information for the billing of all customers is gathered. Mainly this information will be comprised of billing cycles, peak and off-peak periods of the day and/or week and the resource usage scalars (i.e., effective bandwidth and buffer size) for every customer. Once all this information is available the ISP will be in a position to calculate the actual charges of different users.

The elements of the "Network Module" and "Computation Module" can reside on a PC that can be connected to the user's access point (e.g., a switch) with a simple ethernet link. In this configuration the switch that is receiving all of the user's packets can be configured to send a copy of every packet to the SPAN (Switched Port Analyzer) port where the monitoring PC will be connected.[3] In this configuration the monitored link should have a capacity equal or smaller to the capacity of the ethernet link that connects the monitoring station to the switch (if this is not the case then the utilization of the monitored link should be low in order to avoid packet losses). A different configuration is to use an optical splitter to make a copy of everything the user is sending and direct it to

[3]This configuration was used for the monitoring of a 5-day trace at the New Zealand Internet Exchange (http://pma.nlanr.net/Traces/long/nzix2.html).

the monitoring station. Both of these configurations are passive in the sense that they don't intervene with the operation of the network and they don't require additional resources.

We hope to be able to convince an ISP to test out our charging model. If we are indeed successful, then our future research will focus on enhancing our charging models in response to our empirical observations of user and ISP behavior to our charging mechanism. We envision a testbed similar to the Internet Demand Experiment (INDEX) Project (Rupp et al., 1998) which studied the behavior of individual (micro) users of dialup connections in response to usage based pricing. Since INDEX used a university setting for its research, we are hopeful to also use a university network to examine the effect of an ex-post charging model. University networks are a likely candidate because the network administrators may be more open to supporting the research objectives of such a test of the ex-post charging model. A similar testbed for organizational users of high-speed bandwidth connections will go a long way in understanding better many of the practical issues related to usage based pricing, as well as in validating our pricing models.

Acknowledgement:. Support for this research was provided in part by the DoD, Laboratory for Telecommunications Sciences, through a contract with the University of Maryland Institute for Advanced Computer Studies. We are thankful to two anonymous referees for their helpful comments.

References

Bailey, J. P., Nagel, J., and Raghavan, S. (2006). Ex-post internet charging. In McKnight, L. W. and Wroclawski, J., editors, *Internet Services: The Economics of Quality of Service in Networked Markets*. MIT Press. to appear.

Besen, S. M., Spigel, J. S., and Srinagesh, P. (2002). Evaluating the competitive effects of mergers of internet backbone providers. *ACM Transactions on Internet Technology (TOIT)*, 2(3):187–204.

Blumenthal, M. S. and Clark, D. D. (2001). Rethinking the design of the internet: the end-to-end arguments vs. the brave new world. *ACM Transactions on Internet Technology (TOIT)*, 1(1):70–109.

Courcoubetis, C., Kelly, F. P., and Weber, R. (2000). Measurement-based usage charges in communications networks. *Operations Research*, 48(4):535–548.

Courcoubetis, C., Siris, V. A., and Stamoulis, G. D. (1998). Integration of pricing and flow control for available bit rate services in atm networks. In *IEEE Globecom '96*, pages 644–648. London, UK.

Falkner, M., Devetsikiotis, M., and Lambadaris, I. (1999). Cost based traffic shaping: A user's perspective on connection admission control. In *IEEE ICC*.

Guérin, H., Ahmadi, H., and Naghshineh, M. (1991). Equivalent capacity and its application to bandwidth allocation in high speed networks. *IEEE J. Selected Areas Communications*, 9(7):968–981.

IHR (2006). Internet health report http://www.internetpulse.net/. Technical report, Keynote Systems.

ITR (2006). Internet traffic report. Technical report. http://www.internettrafficreport.com/.

Kelly, F. P. (1996). *Notes on effective bandwidths in Stochastic Networks: Theory and Applications Telecommunication Networks*, volume 4, pages 141–168. Oxford University Press. Oxford, UK.

Kelly, F. P. (1997). Charging and accounting for bursty connections, in internet economics. In McKnight, Lee W. and Bailey, Joseph P., editors, *Internet Economics*, pages 253–278. MIT Press.

Low, S. H., Paganini, F., and Doyle, J. C. (2002). Internet congestion control. *IEEE Control Systems Magazine*, 22(1):28–43.

MacKie-Mason, K. and Varian, H. R. (1995). Pricing congestible network resources. *IEEE Journal of Selected Areas in Communications*, 13(7):1141–149.

McKnight, L. W. and Bailey, J. P., editors (1997). *Internet Economics*. MIT Press.

Micheel, J., Graham, I., and Brownlee, N. (2001). The Auckland data set: an access link observed. In *14th ITC Specialists Seminar on Access Networks and Systems*. Catalonia, Spain.

Mills, D. L. (1989). Measured performance of the network time protocol in the internet system. Technical Report RFC-1128. http://www.faqs.org/rfcs/rfc1128.html.

Odlyzko, A. (2001). Internet pricing and the history of communications. *Computer Networks*, 517:493–517.

Odlyzko, A. M. (2003). Internet traffic growth: Sources and implications. Technical report, University of Minnesota. Available at http://www.dtc.umn.edu/publications/publications.php.

Rupp, B., Edell, R., Chand, H., and Varaiya, P. (1998). Index: A platform for determining how people value the quality of their internet access. In *6th IEEE/IFIP International Workshop on Quality of Service*, pages 85–90.

Shenker, S., Clark, D., Estrin, D., and Herzog, S. (1996). Pricing in computer networks: Reshaping the research agenda. *ACM Computational Comm. Review*, pages 19–43.

Siris, V. A., Songhurst, D. J., Stamoulis, G. D., and Stoer, M. (1999). Usage-based charging using effective bandwidths: studies and reality. In *16th International Teletraffic Congress (ITC-16)*.

Srikant, R. (2004). *The Mathematics of Internet Congestion Control (Systems and Control: Foundations and Applications)*. SpringerVerlag.

Appendix

Trace identifier	Figure number
BWY-976126448-1	2,3,9
BWY-20000916	6,7
BWY-20010214	4,5,6,7,8
BWY-20001203	6

Identification of traces from NLANR repository used in this paper.