Pricing of Dialup Services: an Example of Congestion-Dependent Pricing in the Internet[1]

Stephen D. Patek and Enrique Campos-Náñez
Department of Systems Engineering
School of Engineering and Applied Science
University of Virginia

Abstract

Recent research on dynamic pricing of multiclass loss networks [18] has shown that the performance of optimal static pricing approaches that of optimal dynamic (congestion-dependent) pricing in the many small sources limit. In our own work with similar models, we have found it difficult to obtain large gains over static pricing in realistic settings, *even when the many small sources assumption is violated.* In this paper we give an example which is a stochastic control model for congestion-dependent pricing of Internet services. Our formulation captures the basic tradeoff in allocating bandwidth to two classes of users in maximizing average net revenue. Optimal pricing requires that the ISP anticipate and respond to changes in bandwidth consumption. Our goal is to quantify the gain that can be achieved through dynamic pricing over open loop pricing strategies which may or may not account for time-of-day effects. We frame the problem as a continuous-time Markov decision process for which we numerically compute optimal solutions.

1 Introduction

We consider a stochastic control model for congestion-dependent pricing of Internet services, involving a local Internet service provider (ISP) who provides service to two types of customers: large institutions and small dialup users. We assume that the institutions are small in number and that each one has paid in advance for a large amount of assured bandwidth. As for the dialup users, we assume that there can be many of these (subject to the number of dialup modems operated by the ISP) and that each one consumes only a small amount of bandwidth, when active. Our model for Internet traffic ignores the complex dynamics of TCP and actually assumes that all users transmit with fixed window sizes. The effect of congestion (excedence of link capacity) is packet loss which serves to modulate service times. Quality-of-Service (QoS) is defined only in terms of loss-rate (packet loss probability), for which we develop a very simple model. Loss rate is assured only to the institutional users, and this assurance is not explicit but rather takes the form of rebates proportional to the degree of excursion from nominal loss rates. We assume that the logged-in dialup users are charged each time they initiate an active session, specifically each time they click on an average-sized webpage.[1] (This would be in addition to a flat subscription fee.) We allow the "per-click" cost to be congestion-dependent, and we assume that the ISP can communicate this price to the dialup users in such a way that the instantaneous arrival rate of active sessions is price-dependent. The question we begin address in this paper is whether it's worth it to make the per-click price congestion (feedback) dependent, as opposed to just holding it fixed or perhaps varying it according to time-of-day effects. For this study, we make the simplifying, limiting-case assumption that price information can be communicated instantaneously to the dialup users; we will study the effect of delayed feedback in a subsequent paper.

The problem for the ISP in our model is to balance revenue from dialup usage with the cost of refunding the institutional users for occasional QoS violations. We employ a fluid model for load on the network, where the amounts of flow from various sources are modulated according to Markov processes that describe the number of active logged-in users, the state of the institutional users, prevailing loss rate, and the current price set for bandwidth consumption. Optimal congestion-dependent pricing requires that the ISP anticipate and respond to changes in institutional bandwidth consumption so that the highest possible average net revenue rate can be achieved. The optimal solution is a policy that maps network state to per-click price. In general the optimal policy is nonconstant (even with all of the parameters of the model time-invariant). This view of congestion-dependent pricing differs somewhat from other pricing models where distributed algorithms for usage-based prices are derived from the dual optimization of a *static* notion of network welfare, where changes in price are always toward the new underlying optimal solution that emerges when individual flows either commence or terminate. In our model, we explicitly capture the stochastic/dynamic nature of resource

[1]This work is supported by the National Science Foundation through grants ECS-9875688 (CAREER) and ANI-9903001.

[1]In reality, of course, web use does not account for all ISP traffic, but we make this assumption here for simplicity. The resulting optimal per-click prices could be translated into a price per unit bandwidth per second to account for other Internet applications. Also, for simplicity, we do not distinguish between incoming and outgoing flow on the ISP's link (even though the link is likely to be full duplex).

consumption.

Background Congestion dependent pricing of network resources has recently become an active area of research. Prices have been used in the context of market competition [12, 16, 10], through mechanisms such as auctioning of (bidding for) network resources [7, 23] or the fair pricing of pre-established priority classes for packet delivery [17, 6]. In such schemes, the user is typically an ISP competing for resources from a bandwidth reseller, and specific pricing mechanisms are developed to maximize overall welfare in the use of resources and achieve (Nash) equilibrium. Another approach, used in [9] is to reduce congestion during peak hours through discounts to promote a time shift of the demand. In [4], the authors use the effective bandwidths of flows as statistics for the characterization of admission regions in guaranteed-services networks and to establish fair charges for resource use that can be applied dynamically to optimize social welfare. Other approaches [20, 5, 14, 11, 14], also turn to price optimization to achieve high utilization, where to make optimization feasible, and eliminate centralization of processing and information, they show that the dual problem can be distributed to the network entities. Other studies identify price as means of congestion control and load balancing. Particularly, in [24, 15] prices do not have the usual economical significance, but are used as signals for communicating the consumption of resources and are continuously updated to convey such information to the policing entities in the network. Other researchers [3] have focused on comparisons between the performance of static and congestion-dependent schemes.

The model in this paper is closest to that in [18], where optimal pricing policies that maximize either average revenue or social welfare are characterized in terms of a continuous-time Markov decision process (MDP). Ours is a similar MDP model, but the probabilistic structure of our model is tuned to mimic Internet service in which congestion is a dominant feature. The model in [18] may be characterized as a controlled multiclass loss network, and the elegant structure of their model makes possible a clean theoretical analysis. One very compelling result they have obtained is the near optimality of (optimal) static pricing within the class of state-dependent policies. Specifically, in the case of "many small users" where all active flows represent a small fraction of the total bandwidth of a link, the difference between optimal dynamic and optimal static pricing (in terms of average revenue rate) can be made arbitrarily small. The Internet-inspired model of this paper does not have the same clean structure as in [18], and we do not attempt to establish an analogous result, at least for now. Rather, we have engaged in a computational study to gain insight into a case where the "many small users" assumption is violated. Our model mixes both small and large users, where pricing provides a means by which the ISP can adaptively shape dialup demand in response to changes in consumption by the large institutional users. Our model also differs in that we assume a finite and stochastically varying pool of logged-in users, modeled as two-state fluid sources, i.e. superposition of on/off sources, as opposed to a $G/M/\infty$-type model for call arrivals. Also, instead of enforcing a strict admission control mechanism based on the effective bandwidths of flows in various classes, we employ a very simple packet-loss model which modulates the services rates for individual sources during times of congestion. Finally, we use packet loss as a QoS measure for the large institutional users, somewhat akin to the assured service model proposed for differentiated services in the Internet.

In Section 2, we give a precise formulation of our model as a continuous time Markov decision process. In Section 3, we define an instance of the model involving two institutional users. By varying the average number of logged in users and the penalty experienced by the ISP for QoS violations, and by computing both the optimal static and congestion-dependent pricing policies, we are able to identify cases where feedback creates the potential for a substantial gain over static pricing. In Section 4, we discuss our results, make brief conclusions, and state opportunities for future research.

2 Pricing Model

We consider an ISP which has K institutional subscribers, a pool of D dialup modems, and a large base of customers who consume dialup resources, as shown in Figure 1. Let $z_k(t)$ be the load [in bits per second (bps)] submitted by institution k at time t. Let $z_0(t)$ be the load (in bps) due to dialup consumption of resources. We assume that packet loss on the link increases with total load $z(t) = \sum_{k=0}^{K} z_k(t)$. Let $\rho(z(t))$ denote the fraction of bits lost at time t due to congestion, where the function $\rho : \Re_+ \mapsto [0, 1]$ is continuous and nondecreasing. Implicitly, we assume that packet loss is a deterministic function of the total amount of traffic supported by the link. Here, we use a very simple loss model. Namely, if the ISP's link to it's peer network is C bps, then when the load submitted to the network is z the fraction of bits lost is $\rho(z) = \max\{0,\ (z - C)/z\}$. That is, the fraction of the traffic that exceeds the ISP's capacity is lost. Packet loss contributes to longer service times for both the institutional and dialup users as described below.

We assume that the institutions have all entered into service contracts with the ISP with the following type of QoS agreement. For each institution $k = 1, \ldots, K$, the ISP must assure a packet loss probability of P_k, regardless of the bandwidth actually being consumed. The cost of failing to deliver this level of service is a refund to the institution at a rate proportional to the degree of the excursion, specifically: $\max\{0, R_k[\rho(z) - P_k]/P_k\}$, where R_k is a parameter. We assume that each institution may be modeled as a load-dependent Markov chain, with one of three states prevailing at any point in time: 0 (low use), 1 (medium use), 2 (high use). We will use $i_k(t)$ to denote the state of institution k at time t, and $S = \{0, 1, 2\}$ to denote the institutional users'

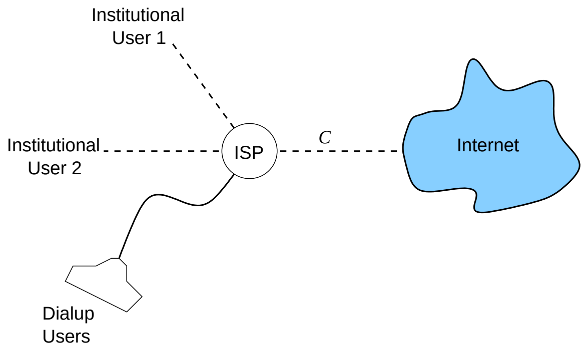

Figure 1: Model of an ISP provider. Institutional and dialup user share the single link with capacity C.

state space. The state of each institution determines its load (fluid flow) submitted to the link. Specifically, given that institution k is in state $i \in S$, it submits $c_{k,i}$ units of fluid flow to the link. Thus, if institution k is in state i at time t, then $z_k(t) = c_{k,i}$. We assume that the time an institution remains in a given state is an exponential random variable, and $\bar{q}_{k,ij}$ denotes the *loss-free* instantaneous rate at which institution k transitions from state i to $j = i \pm 1$. To capture the effects of congestion in the institutional traffic, we modulate the transition rates as

$$
\begin{aligned}
q_{k,i(i+1)}(z) &= \frac{\bar{q}_{k,i(i+1)}}{[1 - c_{k,i}\,\rho(z)/c_{k,2}]}, \quad i = 0, 1, \\
q_{k,i(i-1)}(z) &= \bar{q}_{k,i(i-1)}\,[1 - c_{k,i}\,\rho(z)/c_{k,2}], \quad i = 1, 2.
\end{aligned}
$$

As for the dialup users, we assume that the maximum number of customers who can be logged-in simultaneously is D, the number of modems available at the ISP's premises. The number of logged-in customers determines an upper bound on the number of "active sessions" contributing to the load on the ISP's link. We'll discuss this aspect of the model in the following paragraph. We assume that the number of customers logged-in is random, following a simple birth-death process. Let $\bar{n}(t) = \bar{n}_{lo} + \tilde{n}(t)$ denote the number of customers logged-in at time t, where $\bar{n}_{lo}$ is a predetermined minimum number of logged-in users at any time (could be zero) and $\tilde{n}(t)$ is the state of an auxiliary process defined as the superposition of $D - \bar{n}_{lo}$ on/off sources where off-times for individual sources are exponential with rate λ_d and on-times are exponential with rate γ_d.

The number of logged-in customers determines an upper bound on the number of "active sessions" contributing to the load on the ISP's link. In this way, each logged-in dialup user is modeled as a two-state fluid source. The total amount of dialup flow is proportional to the number of active application-layer sessions generating traffic, $z_0(t) = n(t)M$, where $n(t)$ is the number of active dialup sessions [with $0 \le n(t) \le \bar{n}(t)$] and M is the bandwidth associated with each active session. Our notation reflects the implicit assumption that all individual dialup users transmit at a common rate M (bps). We assume that when a source is active it stays active for a random amount of time, modeled by an exponential distribution with instantaneous rate $\gamma(z(t)) = M[1-\rho(z(t))]/B$. This formula reflects the time required to download a webpage of size B (bits), where M (as above) is the rate of the connection, and $(1 - \rho(z(t)))\mu$ accounts for degraded service when the link's capacity has been exceeded. Aside from the completion of a download, there is one other mechanism by which the number of active sessions may decrement. Namely, if all logged-in users are active (i.e. $n(t) = \bar{n}(t) > \bar{n}_{lo}$), then if one of the logged-in customers departs, both $n(t)$ and $\bar{n}(t)$ are reduced by one.

With regard to the arrival of active dialup sessions (webpage downloads by logged-in users), we assume that the ISP has a means of informing dialup users of the price for the next web click. This determines the random time that users stay inactive before initiating another download. We assume that this time is is exponentially distributed with rate $\lambda(u)$, where u is the ISP-determined price for each newly initiated active session. We assume that values of u are chosen from a compact (possibly finite) set U and that the demand (arrival rate) function $\lambda : U \mapsto \Re$ is continuous, nonincreasing, and is such that an arrival rate of zero can be achieved by setting u to a maximal price u_{max}.

Formulation as a Continuous-Time Markov Decision Process Given the current price u, we only need to know $\bar{n}(t)$, $n(t)$, $i_1(t)$, ..., $i_K(t)$ to fully characterize the uncertainty regarding: (1) when the next transition will occur and (2) what that transition will be. Thus, it serves as the state of a continuous-time Markov decision process, which we now proceed to formulate. It will be convenient to refer to $\mathbf{N}(t) = \left[\bar{n}(t), n(t), i_1(t), \ldots, i_K(t)\right]$ as the state of the process, and we'll use $\mathcal{N}$ to denote the (finite) set of states that the system can assume.

The ISP seeks to maximize profit through the judicious choice of u out of a compact (possibly finite) set of pricing options U. In doing this the ISP must balance expected dialup revenue with the costs associated with excursions from the nominal packet loss probabilities assigned to the institutions. We capture this as a continuous-time Markov decision process with the average reward criterion $J^*(\mathbf{N}) = \max_{\mu \in M} \lim_{T \to \infty}$

$$
\frac{1}{T}\mathbf{E}\left\{ \int_0^T [\alpha(\mathbf{N}(t), \mu(\mathbf{N}(t))) - \beta(\mathbf{N}(t))]\, dt \;\middle|\; \mathbf{N}(0) = \mathbf{N},\ \mu \right\}, \quad (1)
$$

where

1. M is the set of all functions (stationary policies) $\mu : \mathcal{N} \mapsto U$ that map states $\mathbf{N}$ to dialup per-click prices $\mu(\mathbf{N})$,

2. $\alpha(\mathbf{N}, u)$ is the average rate at which dialup revenue comes back to the ISP under the price u evaluated at the state $\mathbf{N} = (\bar{n}, n, i_1, \ldots, i_K)$:

$$\alpha(\mathbf{N}, u) = [\bar{n} - n]\lambda(u)u,$$

and

3. $\beta(\mathbf{N})$ is the rate at which the ISP must refund the institutional users for violation of the nominal packet loss probability:
$$\beta(\mathbf{N}) = \sum_{k=1}^{K} R_k \max\{0, (\rho(z) - P_k)/P_k\},$$
where z is the combined load on the ISP's link when the state of the system is $\mathbf{N}$.

Characterization of Optimality The probabilistic structure of our pricing model is such that for any two states $\mathbf{N}_1, \mathbf{N}_2 \in \mathcal{N}$ there is a stationary policy that makes $\mathbf{N}_2$ reachable from $\mathbf{N}_1$. As such, the theory of communicating, average reward, semi-Markov decision processes applies (cf. [22] and the references contained therein). We start by uniformizing the continuous-time decision process to obtain an equivalent discrete-time process with self-transitions, as in [2, 13, 21], with dynamic operator T (see [19] for a more detailed discussion).

From [22, 8], optimality is characterized by solutions $\phi \in \Re$ and $H : \mathcal{N} \mapsto \Re$ to the functional equation (Bellman's equation)
$$H(\mathbf{N}) + \phi = TH(\mathbf{N}), \qquad \forall\, \mathbf{N} \in \mathcal{N}. \tag{2}$$

From [22] (which considers the case of compact constraint sets), a solution to Bellman's equation is sure to exist, and the solution is unique up to constant shifts in H. In fact, using $\phi^* \in \Re$ and $H^* : \mathcal{N} \mapsto \Re$ to denote the *unique* solution such that $H^*(\mathbf{0}) = 0$, we may interpret ϕ^* as the optimal average reward for the decision process [i.e. the optimal value in (1) for all $\mathbf{N} \in \mathcal{N}$] and H^* as the "value function" that describes the relative value associated with starting the process in different states $\mathbf{N} \in \mathcal{N}$. Moreover, any stationary policy $\mu^* : \mathcal{N} \mapsto \Re$ such that $\mu^*(\mathbf{N})$ achieves the minimum in $TH^*(\mathbf{N})$ for all states $\mathbf{N} \in \mathcal{N}$ is optimal (i.e. achieves the optimal average value ϕ^*). [8]

Note that to characterize the average reward for a fixed (static) price u the same analytical framework applies. We need only define a new dynamic programming operator $T_u : \mathcal{N} \mapsto \Re$ which is the same as T without the maximization over $u \in U$. Then, solving for the unique $\phi_u \in \Re$ and $H_u : \mathcal{N} \mapsto \Re$ such that $H_u(\mathbf{N}) + \phi_u = T_u H_u(\mathbf{N})$ for all $\mathbf{N} \in \mathcal{N}$ and $H_u(\mathbf{0}) = 0$, the average reward associated with u is ϕ_u.

Computing Solutions via Relative Value Iteration Since the state $\mathbf{0} = (0, \bar{n}_{lo}, 0, \ldots, 0)$ is recurrent under all policies it is possible to use relative value iteration [1, 21, 25] to compute solutions to Bellman's equation. Let $H_0 : \mathcal{N} \mapsto \Re$ be an initial guess for the optimal relative value function. The algorithm proceeds iteratively, as
$$H_{k+1}(\mathbf{N}) = T(H_k)(\mathbf{N}) - T(H_k)(\mathbf{0}), \qquad \forall\, \mathbf{N} \in \mathcal{N}. \tag{3}$$

Note that if this algorithm converges, say to $\bar{H}$, then the limit satisfies $\bar{H}(\mathbf{N}) + T(\bar{H})(\mathbf{0}) = T(\bar{H})(\mathbf{N})$ for all $\mathbf{N} \in \mathcal{N}$. In other words, if the method converges, then $T(H_k)(\mathbf{0})$ converges to the optimal average reward ϕ^*, and H_k converges to the optimal relative value function H^*. For the case of a finite set of pricing options U (i.e. a finite constraint set), convergence is assured from [21] because of the way we have uniformized the process. Note that relative value iteration also applies in evaluating the average reward associated with a fixed price u, where all we need do is replace T with T_u in the recursion.

3 Computational Results

Setup of the Experiments In this section we describe a set of experiments designed to gain insight into the model of Section 2. We consider the case of two institutions and up to 190 dialup users. The capacity of the link is set to $C = 10$ Mbps, approximately one quarter of a T3 link. The load that an institution pumps into the link depends on its state. For this experiment, institutional users have three possible states, corresponding to low (542.5 kbps), medium (2.17 Mbps), and high (4.34 Mbps) loads. The transition rates for the institutions were set as: $\bar{q}_{k,01} = .05$, $\bar{q}_{k,12} = .05$, $\bar{q}_{k,10} = .01$, and $\bar{q}_{k,21} = .1$. Notice that when both institutions reach the high load state, the combined institutional load is 86.8% of the link's capacity. We set the nominal QoS levels to be $P_1 = P_2 = .001$. Dialup users log in according to the birth death process for $\tilde{n}$ (described in Section 2), with parameters $\lambda_d = .025$ and $\gamma_d = .025$. Once logged in, users click on webpage links at rates that depend on the price set by the ISP. For this experiment, we set λ to decrease linearly from a free rate of $\lambda(0) = .1$ (clicks per second) to a rate of $\lambda(1) = 0$ when the price reaches its maximum of $u_{max} = 1$ monetary unit. We assume an average webpage size of $B = 640$ kb and that all users log in with $M = 56$ kbps modems. To get an idea of how these numbers add up, consider the case where the price per-click is set to $u = .6$: if one monetary unit equals \$.01, this yields a net usage-based revenue of roughly \$16000 per month assuming (1) an average of 100 logged-in users for 18 hours a day and (2) QoS penalty factors $R_1 = R_2 = .1$. In the following, one monetary unit = $\$.01$.

To investigate the impact of dialup user load and the effect of the QoS penalty factors R_1 and R_2, we have constructed a matrix of scenaria as shown in Table 1. In the experiments we have set $R_1 = R_2 = R \in \{.01, .1, 1\}$, and for each value of R there are different values for the minimum and maximum numbers of logged-in users, $\bar{n}_{lo}$ and D respectively. For each case, we evaluate each possible static price in the set $U_s = \{.5, .51, .52, \ldots, 1.0\}$ by relative value iteration. (Prices below $u = .5$ never turn out to be optimal.) After identifying the optimal state price u^* in this way, we then compute the dynamic optimal policy over the set $U = \{u^*, .5, .6, .7, .8, .9, 1.0\}$, again using rel-

Experiment	Light ($R = .01$)	Moderate ($R = .1$)	Heavy ($R = 1$)
1	90-110	50-70	10-30
2	130-150	90-110	50-70
3	170-190	130-150	90-110

Table 1: Definition of the experiments in which different load ranges and penalty factors are tested.

ative value iteration. In a given instance of the problem, the number of states depends on the numbers of logins allowable. Specifically, the number of states is lower bounded by $9\bar{n}_{lo}(D - \bar{n}_{lo})$, and each relative value iteration involves a maximization over U for each state. Generally, between 3000 and 10000 iterations are required for convergence to within 10^{-6} in the sup-norm of the difference between successive iterates. Thus, the computational burden, even for this very simple formulation, can be quite large.

Computational Results Table 2 summarizes the results of our study. The table lists for each value of the QoS penalty factor and associated range of logged-in users the static optimal price in U_s, the associated average net revenue, and the dynamic optimal average net revenue for the associated constraint set U along with the percentage improvement that the dynamic optimal policy affords. Note that in some cases, particularly when the number of logged-in users is small (relative to the penalty factor R), the percentage improvement with the dynamic optimal policy is quite small. However, as the number of logged-in users is increased (and the demands on the network become excessive), the value of the dynamic optimal policy becomes apparent, with up to 22.5% improvement over the optimal static price. The question remains, however, whether real-time congestion-dependent pricing is worth the effort and expense for its implementation. It is interesting to note that the optimal *static* price is generally strongly dependent on the number of logged-in users, suggesting that it is still reasonable to vary price in accordance with time-of-day variations in traffic. As an alternative presentation of our results, Figure 2 shows average instantaneous revenue for each possible static price compared to the optimal dynamic revenue (shown in each case as a horizontal line). The optimal policies (not shown) are nonconstant, generally increasing monotonically from some low price to the cutoff price $u_{max} = 1$ (monetary units) as a function of the number of active dialup sessions for each possible combination of instiution states.

4 Discussion

In the computational study of the preceding section we were able to identify scenaria where congestion-dependent pricing has the potential to substantially increase average net revenue for the hypothetical ISP. On a more philosophical level, we have shown through an example (admittedly somewhat contrived) that real-time feedback through pricing *can*

Low Penalty ($R = .01$)				
# Logged-in	Opt. Static	Revenue	Dynamic Rev.	% improv.
90-110	.63	1.5782	1.5938	0.99
130-150	.64	2.1087	2.1424	1.59
170-190	.65	2.6226	2.6728	1.91
Medium Penalty ($R = .1$)				
# Logged-in	Opt. Static	Revenue	Dynamic Rev.	% improv.
50-70	.62	.9772	.9863	0.93
90-110	.75	1.3555	1.4953	10.31
130-150	.82	1.5267	1.8405	20.55
Heavy Penalty ($R = 1.0$)				
# Logged-in	Opt. Static	Revenue	Dynamic Rev.	% improv.
10-30	.59	.3305	.3305	0
50-70	.69	.9343	.9797	4.86
90-110	.82	1.1485	1.4074	22.54

Table 2: Summary of the optimal performance for static and dynamic pricing schemes for the three sets of experiments.

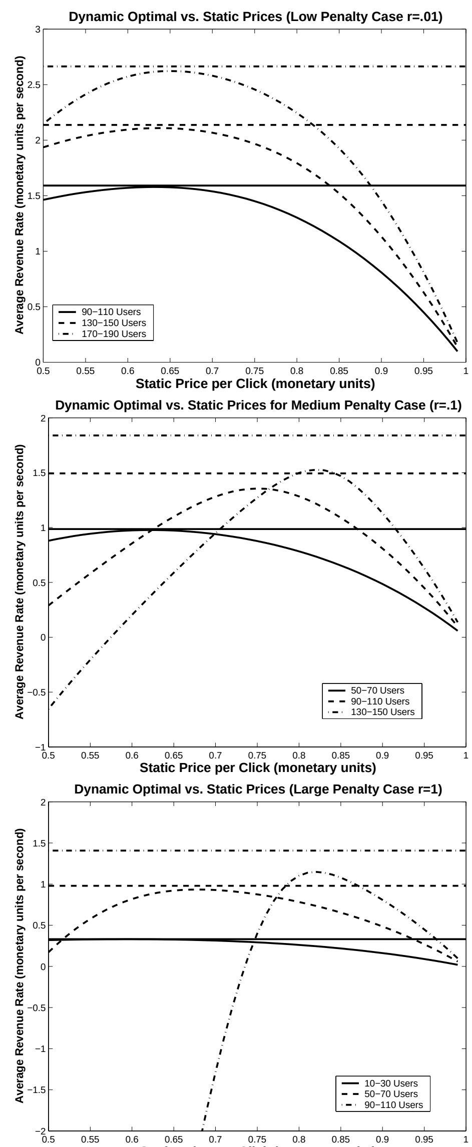

Figure 2: Comparison of the performance of all possible static prices for the low penalty case ($R = .01$), medium penalty ($R = .1$), and heavy penalty case ($R = 1.0$). The performance levels associated with the dynamic optimal policies are shown as horizon lines.

have enough “control authority” to significantly impact the utilization of network resources. This contrasts to some degree the finding in [18] of the near optimality of static pricing in the case of “many small users.” However, a gain of roughly 20% seems low, given the degree to which we have rigged the model, through mixing small and large users.

Further investigation is required to determine whether the gain associated with congestion dependent pricing is substantial enough to warrant an implementation of the model. Many details of the model are set completely arbitrarily. For example, the loss model (the equation that defines ρ) from Section 2 and the linear model for demand in Section 3 are highly suspect and should be refined. The way in which QoS violations are penalized through the definition of β also bears further scrutiny. Of course the assumption of exponential switching times in the on/off traffic model requires validation, and perhaps other distributions, perhaps with heavy tails, may be more descriptive of real Internet traffic. Finally, our continuous-time pricing model assumes that price information can be communicated instantaneously to the set of logged-in users. In a real network, there would be a significant delay in conveying price information, and it would be essential to resolve the impact of this delay before continuing with an implementation. Without question, the delay will result in some reduction of achievable performance, but how much remains to be determined.

References

[1] D. P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 2. Athena Scientific, Belmont, MA, 1995.

[2] F. J. Beutler and K. W. Ross. Uniformization for Semi-Markov Decision Processes under Stationary Policies. *J. Appl. Prob.*, 24:644–656, 1987.

[3] Ll. Cerdà and O. Casals. Dynamic and Static Charging of the ABR Service in ATM Networks. *Proc. of the GLOBECOM’98*, 1998. http://www.ac.upc.es/recerca/XARXES/expert/hotpapers.html.

[4] C. Courcoubetis and V. A. Siris. An Approach to Pricing and Resource Sharing for Available Bit Rate (ABR) Services. Technical report, Institute for Computer Science, P.O.Box 1385,GR 711 10 Heraklion, Crete, Greece, 1997.

[5] R. J. Gibbens and F. P. Kelly. Resource pricing and the evolution of congestion control. *Automatica*, 35, 1999.

[6] A. Gupta, D. O. Stahl, and A. B. Whinston. Priority Pricing of Integrated Services Networks. *Internet Economics,MIT Press*, pages 323–378, 1998.

[7] L. Murphy J. K. MacKie-Mason and J. Murphy. Responsive Pricing in the Internet. *Internet Economics,MIT Press*, pages 279–304, 1998.

[8] W. S. Jewell. Markov-Renewal Programming. I and II. *Operations Research*, 11:939–971, 1963.

[9] N. Keon and G. Anandalingam. Adaptive Flow Control for VoD Using Price Discounts. Technical report, Department of Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104-6315, February 1999.

[10] Y. A. Korilis, T. A. Varvarigou, and S. R. Ahuja. Incentive-compatible pricing strategies in noncooperative networks. *(Conference Paper) IEEE INFOCOM ’98, Conference on Computer Communications*, pages 439–446, 1998.

[11] J. F. Kurose and R. Simha. A Microeconomic Approach to Optimal Resource Allocation in Distributed Computer Systems. *IEEE Transactions on Computers*, 38(5):705–717, May 1989.

[12] R. J. La and A. Venkat. Network pricing using a game theoretic approach. *Proceedings of Conference on Decision and Control*, 1999.

[13] S. A. Lippman. Applying a New Device in the Optimization of Exponential Queuing Systems. *Operations Research*, 23(4):687–710, 1975.

[14] S. Low and D. Lapsley. Optimization Flow Control, I: Basic Algorithm and Convergence. *IEEE/ACM Transactions on Networking*, 7(6), 1999.

[15] S. H. Low and P. P. Varaiya. A New Approach to Service Provisioning in ATM Networks. *IEEE/ACM Transactions on Networking*, I(5), May 1993.

[16] R. T. Maheswaran, O. C. Imer, and T. Basar. Agent mobility under price incentives. *Proceedings of Conference on Decision and Control*, 1999.

[17] P. Marbach. Pricing Priority Classes in a Differentiated Services Network. *Allerton Conference*, 1999. http://www.ccsr.cam.ac.uk/ pm259/publ.html.

[18] I. Ch. Paschalidis and J. N. Tsitsiklis. Congestion-Dependent Pricing of Network Services. Technical report, Systems Group, Department of Manufacturing Engineering, Boston University, October 1998. http://ionia.bu.edu/papers/pricing.pdf.

[19] S. D. Patek and E. Campos-Nánez. Congestion-dependent pricing of dialup services. *Technical Report, Department of Systems Engineering*, 2000.

[20] J. R. Perkins and R. Srikant. The role of queque length in congestion control and resource pricing. *Proceedings of Conference on Decision and Control*, 1999.

[21] P. J. Schweitzer. Iterative Solution of the Functional Equations of Undiscounted Markov Renewal Programming. *J. Math. Anal. Apple.*, 34:495–501, 1971.

[22] P. J. Schweitzer. A Brouwer Fixed-Point Mapping Approach to Communicating Markov Decision Processes. *Journal of Mathematical Analysis and Applications*, 123:117–130, 1987.

[23] N. Semret, R. R. F. Liaro, A. T. Campbell, and A. A. Lazar. Market Pricing of Differentiated Internet Services. Technical report, CTR/Department of Electrical Engineering, Columbia University, New York, NY, 10027-6699, USA., 1998.

[24] I. Stoica and H. Zhang. LIRA: An Approach for Service Differentiation in the Internet. *Proceedings of NOSSDAV’98*, pages 115–128, 1998. http://www.cs.cmu.edu/ istoica/pubs.html.

[25] D. J. White. Dynamic Programming, Markov Chains, and the Method of Successive Approximations. *J. Math. Anal. Appl.*, 6:373–376, 1963.