Telecommunication Systems Journal, Vol. 18, No. 4, pp. 315-29, Dec. 2001

State-Dependent Pricing and Its Economic Implications[1]

Qiong Wang[2] and Jon Peha[3]

Abstract:

In a packet-switched integrated-services network, the service provider can charge users a state-dependent price, which depends on the extent to which the network is congested. Alternatively, one can also charge users a long-term average price, which is set based on expected demand and capacity availability derived from prior experience at that time-of-day, but independent of instantaneous network conditions. In this paper, we study the economic implications of both state-dependent and long-term average pricing. Using dynamic programming and computer simulations, we compare benefits that the service provider, society, and consumers can derive under different pricing schemes. Our results suggest that adopting state-dependent pricing improves both profit to the service provider and total benefits to the society. Those improvements are achieved for two reasons. First, state-dependent pricing functions as a traffic management mechanism that leads to a better packet throughput. Second, because under state-dependent pricing, a higher price is charged during the congested periods, limited capacity can be used to send more valuable packets at those times. Meanwhile, the service provider can extract more wealth from each packet sent. While the first effect benefits the service provider, society, and consumers, the second effect can reduce consumer benefits. Therefore, consumers may or may not benefit, depending on which of these two effects is more significant. We demonstrate that with smaller buffer size or inelastic demand, the second effect is more likely to dominate the first, which will result in a lower consumer benefit.

---

[1] The authors gratefully acknowledge Brendan Dunham for his hard work in coding simulation programs, and Professor Marvin Sirbu for his helpful comments. Responsibility for errors are solely retained by the authors.

[2] Member of Technical Staff in Bell Laboratories and doctoral candidate in the Department of Engineering and Public Policy, Carnegie Mellon University. Address: 600 Mountain Avenue, Bell Labs, Murray Hill, NJ 07974. Tel: (908)-582-3798. Fax: (908)-582-3161. E-mail:chiwang@lucent.com

[3] Associate Professor at Carnegie Mellon University in the Department of Electrical and Computer Engineering, and the Department of Engineering and Public Policy. Address: Carnegie Mellon University, Dept. of ECE, Pittsburgh, PA, 15213-3890. Tel: (412)-268-7126. E-mail: peha@ece.cmu.edu. http://www.ece.cmu.edu/~peha

Section 1 Introduction

In integrated-services networks, a service is defined as *guaranteed service* if it is non-interruptable, meets specified quality of service, and requires some form of advance capacity reservation. A service is defined as *best-effort service* if it is interruptable, does not need to meet any specified quality of service, and can therefore be offered without capacity reservation. Because of packet arrival rate from guaranteed services fluctuates over time, and the network has guaranteed that these packets will be transmitted in a timely manner, capacity available for best-effort service varies over time. It has long been argued that the benefit from providing best-effort service can be enhanced if price is allowed to vary in very short time intervals such as microseconds ([MACK95]). In intervals when the network has more capacity than demand, the price should be lowered to encourage the use, and in intervals when the network is congested, the price should be raised to allocate limited capacity to the most valuable packets. Those approaches that vary price to follow network congestion status are defined as *state-dependent pricing*. Several state-dependent pricing schemes have been proposed ([DANI97], [GUPT97], [MACK97]).

An alternative to state-dependent pricing is *long-term average pricing* under which price is set according to expected packet arrival rate and capacity availability. A long-term average price can be changed to reflect time-of-day variation of these expected values, but does not vary with random fluctuations of network congestion status.

To decide whether state-dependent pricing or long-term average pricing should be adopted, one needs to examine benefits of each approach from the service provider, consumer, and society perspectives. Benefits to the service provider are measured by profit. For consumers, each user can derive some value from sending a packet, thus is willing to pay a price that equals that value to get the packet sent. The difference between the willingness to pay and the price paid is the net benefit that a user derives. Benefits to consumers as a whole can be measured by *consumer surplus*, which is the sum of net benefits of all users. For society, benefits are measured by the sum of profit and consumer surplus, which is defined as *social welfare*. In this paper, we will use profit, consumer surplus, and social welfare as criteria to evaluate different pricing schemes.

There are many factors that can affect evaluations of different pricing schemes. For example, state-dependent pricing requires more complicated billing and accounting systems in order to charge users based on network congestion status at time of sending their packets. As a result, the service provider's profit of adopting state-dependent pricing is affected by the cost of building those complicated systems. Moreover, networks are usually equipped with some kinds of congestion control mechanism, some of which are more effective than others. Therefore, even if packets arrival rate and capacity availability are exactly the same, the congestion status can be very different in networks that adopt different congestion control mechanisms. Consequently, the relative advantage of state-dependent pricing, which dynamically adjust prices to follow network congestion status, depends on what congestion control mechanisms are used in the network. In this paper, we will limit our analysis to the case in which neither billing and accounting costs nor congestion control mechanism is considered. Those limitations favors state-dependent pricing in the evaluation of pricing mechanisms. If the comparison shows even in that case, adopting state-dependent pricing does not improve profit, consumer surplus, or social welfare, then the mechanism does not deserve further consideration. Otherwise, more studies are needed to examine if the advantage is significant enough to justify the cost of building complicated billing and accounting systems, and whether state-dependent pricing can keep its superiority in case the network has an effective congestion control mechanism.

In the rest of the paper, we will demonstrate that introducing state-dependent pricing is always beneficial to the service provider, and to society as a whole, but can sometimes be harmful to consumers. We will also characterize situations in which consumer benefit is likely to decrease with the adoption of state-dependent pricing and explain why. We will first discuss different pricing schemes in Section 2. We will then compare profit, social welfare, and consumer surplus in section 3, and summarize the paper in Section 4.

Section 2 Different Pricing Schemes for Best-effort Service

In this section, we define different pricing schemes and derive procedures for calculating the profit-maximizing prices under each one.

2.1 Definition of Pricing Schemes

Long-term average pricing is to set the price for some specific time-of-day based on expected traffic volume at that time. The price does not change with instantaneous fluctuations of network congestion status.

State-dependent pricing associates the price with the current state of the network Packets of best-effort service are admitted only when the consumer's willingness-to-pay is no less than the current price, and the buffer has enough space to accept it.

Based on what is defined as the "state", a variety of pricing schemes can be considered as state-dependent pricing. Some schemes are more complex than others, but are more accurate in describing network congestion status. In this paper, we consider two schemes: response pricing and spot pricing.

In a network that offers both guaranteed and best-effort services, the number of guaranteed calls in progress determines how much shared capacity is used by the guaranteed service, thus reflects network congestion status. Consequently, one can define the number of guaranteed calls in progress as the state variable, and change the price whenever this variable changes, either because of the arrival of a new call and/or the departure of an existing call. We define this type of state-dependent pricing as response pricing, since under that scheme, the price change is invoked in response to a change in number of calls in progress.

Alternatively, one can use a more accurate but more complicated indicator of network congestion by defining state as a combination of buffer occupancy and number of guaranteed calls in progress. Under that scheme, the state-dependent price changes at a fixed interval based on the number of calls in progress and packets in the buffer at the beginning of that interval. We define this type of state-dependent pricing as spot pricing, since it is essentially a mechanism of allocating constrained network capacity in a spot market as discussed in previous literature [JORD95].

2.2 Assumptions

The above definitions of pricing schemes can lead to different procedures to calculate the optimal prices under each scheme, depending upon network architecture, service discipline, call/packets arrival process, etc. This subsection will specify a scenario, based on which our analysis is derived.

Assume one guaranteed service and one best-effort service are offered on an access link, of which the capacity is C_T fixed-length packets per second. Assume the call arrival process of the guaranteed service is Poisson, and call duration is exponentially distributed. Define λ_c and r_c as the call arrival and departure rates in calls/minute, respectively. Define M_c as the maximum number of guaranteed calls that can be carried simultaneously. For each guaranteed call, assume the packet arrival process is Poisson and define λ_g as the packet arrival rate in packets per second. Calls of guaranteed service are admitted on a first-come-first-serve basis.

Assume packet arrivals from the best-effort service also follow a Poisson process, and let the packet arrival rate be a function of price, p_b,$\lambda_b(p_b)$. $\lambda_b(p_b)$ is defined as the demand function for best-effort service.

There is a shared buffer for all packets. Packets of guaranteed service are always admitted into the buffer as long as there are empty spaces. Assume the buffer size is such that the probability of dropping a guaranteed-service packet due to buffer overflow is negligible. Nevertheless, whether an arriving packet of best-effort service will be admitted or not depends on the sum of the packet queue length of guaranteed service and that of best-effort service. Packets of best-effort service will always be admitted if that sum is below a given threshold B, and will be dropped once this threshold has been reached. B is defined as the maximum buffer size for best-effort service. Once in the buffer, all packets from guaranteed and best-effort services will be transmitted on first-come-first-serve basis.

2.3 Derivation of the Optimal Prices

In this section, we describe the calculation of optimal prices given the scenario defined in 2.2. The price is set to maximize expected profit from best-effort service.

Long-term Average Pricing

Long term average price, denoted as p_{bl}, is set to maximize the expected profit function:

$$\Phi(p_{bl}) = \sum_{i=0}^{M_c} \Pr_i [1 - d_i(\lambda_b)] * p_{bl} * \lambda_b(p_{bl}) \qquad (1)$$

where $\lambda_b(p_b)$ is the demand function. $\Pr_i$ is the probability that there are i guaranteed calls in progress, which equals $\dfrac{\frac{(\lambda_c / r_c)^i}{i!}}{\sum_{m=0}^{M_c} \frac{(\lambda_c / r_c)^m}{m!}}$. $d_i(\lambda_b)$ is the packet dropping rate of best-effort service when there are i guaranteed calls in progress and packet arrival rate is λ_b. $d_i(\lambda_b)$ can be derived from the steady-state analysis of the packet queue in the buffer. In that analysis, packet arrival rate is $\lambda_b + i\lambda_g$ if queue length is less than B, and $i\lambda_g$ otherwise. $d_i(\lambda_b)$ is the sum of probabilities that queue length is larger than or equal to B.

The value of p_{bl} that maximizes Equation 1 is the optimal long-term average prices, and can be found by exhaustive search.

Response Pricing

The optimal response price, denoted as p_{br}, is set to maximize the expected profit for the period when the number of calls in progress stays constant. If that number is i, the profit function is:

$$\Phi(p_{br}, i) = [1 - d_i(\lambda_b)] * p_{br} * \lambda_b(p_{br}) \qquad (2)$$

Given λ_b, C_T-i*λ_g, and B, $d_i(p_{br})$ can be determined in the same way as in the calculation of long-term average prices, and the optimal value of p_{br} can also be obtained by trial and error search.

Spot Pricing

The optimal spot price is set at beginning of every time segment, of which the duration is k packet times. Let q_j be the number of guaranteed service calls in progress and n_j be the buffer occupancy when segment j starts. As an approximation, it is assumed that q_j remains constant for the period over which the spot price is optimized, and will be denoted as q in the description of algorithm. p_{js}, the spot price for segment j, is chosen to maximize the following profit function:

$$\Phi^*(n_j,q) = \underset{p_{js}}{Max}\,\Phi_j(p_{js},n_j,q)$$

$$= \phi_j(p_{js},n_j,q) + \sum_{l=0}^{B} P_{j+1}(n_{j+1} \mid p_{js},n_j,q) * \Phi^*_{j+1}(n_{j+1},q) \quad (3)$$

where $\Phi_j(p_{js},n_j,q)$ is the expected profit for all time after segment j by setting the spot price in segment j to p_{js}. $\Phi_j(p_{js},n_j,q)$ includes both $\phi_j(p_{js},n_j,q)$, the expected value from the current segment j, and $\sum_{l=0}^{B} P_{j+1}(n_{j+1} \mid p_{js},n_j,q) * \Phi^*_{j+1}(n_{j+1},q)$, the future profit from all time after segment j. The future profit is the weighted sum of $\Phi^*_{j+1}(n_{j+1},q)$, the maximum profit achievable from segment j+1 afterwards, given that buffer occupancy is n_{j+1} at the beginning of segment. The weight, $P_{j+1}(n_{j+1} \mid p_{js},n_j,q)$, is the conditional probability that buffer occupancy will be n_{j+1} under spot price p_{js}. Given n_j and q, both $\phi_j(p_{js},n_j,q)$ and $P_{j+1}(n_{j+1} \mid p_{js},n_j,q)$ can be uniquely determined by p_{js}.

Equation 3 can be solved by backward recursion, which initializes the calculation by setting $\Phi^*_{j+1}(n_{j+1},q)$ to zero for all n_{j+1}, and search for p_{js} that maximizes $\Phi_j(p_{js},n_j,q)$. Each iteration then starts, using $\Phi^*(n_j,q)$ calculated in the last round as $\Phi^*_{j+1}(n_{j+1},q)$ and searching for the value of p_{js} that maximizes $\Phi_j(p_{js},n_j,q)$. The optimal values of $p_{js}(n_j)$ has been reached when none of $p_{js}(n_j)$ (n_j=0,B) changes after the iteration.

Section 3 Comparison of Different Pricing Schemes

In this section, we compare state-dependent pricing (spot pricing and response pricing) with long-term average pricing based on the benefits of different schemes to the service provider, consumers, and society. The results demonstrate that service providers always obtain a greater profit under state-dependent pricing than under long-term average pricing. Adopting state-dependent pricing also increases social welfare. Between the two state dependent pricing schemes, spot pricing is more effective than response pricing, since the former allows more flexibility to vary prices. Nevertheless, consumers may not always benefit from state-dependent pricing. In some cases, spot pricing can result in lower consumer surplus than the other two schemes.

Section 3.1 presents a simplistic scenario to reveal the intuition behind our conclusions. The effects are demonstrated in a more complex scenario in Section 3.2.

3.1 A Simple Example

As a simple example, consider a case in which there are three users, each one generating one packet per unit of time. Their willingness to pay is 4, 3.2, and 2.4 per packet, respectively. Assume these three users share the network with others, and the network has enough capacity to admit one packet per unit of time 60% of the time when the network is “congested”, enough capacity to admit two packets 20% of the time when the network is “normal”, and enough capacity to admit three packets the remaining 20% of the time when the network is “off-peak”. Under long-term average pricing, the optimal price is 3.2, and the network admits one packet in the “congested” situation, and 2 packets in the two other situations. Under state-dependent pricing, service providers will set price to be 4, 3.2, and 2.4 , and sends 1, 2, and 3 packets at congested, normal, and off-peak times, respectively. Table 1 shows the resulting profit, social welfare, and consumer surplus, and their break-up in different time periods.

Table 1
Profit, Social Welfare, and Consumer Surplus (Example I)

	profit		social welfare		consumer surplus	
situation	long-term	state-dependent	long-term	state-dependent	long-term	state-dependent
congestion (60%)	3.2	4	3.6	4	0.4	0
normal (20%)	6.4	6.4	7.2	7.2	0.8	0.8
off-peak (20%)	6.4	7.2	7.2	9.6	0.8	2.4
expected values	4.48	5.12	5.04	5.76	0.56	0.64

In the normal situation, both schemes charge the same price. Therefore, there is no difference in profit, social welfare, or consumer surplus.

In the off-peak situation, even though the network can admit 3 packets, only 2 will be admitted under long-term average pricing. However, under state-dependent pricing, all capacity will be used instead of being idle because the scheme has the flexibility to reduce the price to increase packet inflow. In this case, the ability to vary price under state-dependent pricing can be viewed as a better traffic management function, which helps to achieve efficient use of network resources. As a result, profit, social welfare, and consumer surplus all increases.

In the congested situation, price will not change under long-term average pricing, and packets will be indiscriminately admitted or dropped. Therefore, service providers admit fewer packets and derive the same profit per packet as that in normal and off-peak periods. Under state-dependent pricing, the price will be raised so the limited capacity will be used to admit more valuable packets. Even though the number of packets transmitted in that period will be smaller, such a “selection” function results in a higher value and profit per packet sent. Consequently, the profit and social welfare are higher under state-dependent pricing than long-term average pricing. Nevertheless, raising price for high-willingness to pay consumers also means more wealth is transferred from users to the service provider. As a result, the table shows that consumer surplus will be lower under state-dependent pricing.

The last row of Table 1 gives expected values from all three situations. It shows state-dependent pricing enhances profit, social welfare, and consumer surplus. Nevertheless, the effect on consumer benefits may not always be positive. As discussed above, state-dependent pricing results in better traffic management, which benefits consumer surplus, and a higher price for some consumers, which harms it. In case the latter dominates the former, consumer surplus can actually decrease. Those cases can occur when congested periods happen more frequently, so state-dependent pricing will rely on extracting more wealth from high-willingness to pay users to improve profit and social welfare. In the example above, if frequency of congested period is increased to 70% of time, and that of off-peak period is reduced to 10%, the overall consumer surplus will become 0.48 under state-dependent pricing, which is less than 0.56 achieved under long-term average pricing (See Table 2 below).

Consumer benefits can also be reduced when demand become less elastic. In this case, the same price increase results in smaller decrease in packet arrival rate, therefore, the service provider can raise price even higher to keep packet arrival rate at the same level during the congested period. For example, if the three users' willingness- to-pay is changed from 4, 3.2 and 2.4 to 4.2, 3.1, and 2.4, then consumer surplus will also be less under state-dependent pricing than that under long-term average pricing (see Table 3 below).

Table 2
Profit, Social Welfare, and Consumer Surplus (Example II)

	profit		social welfare		consumer surplus	
situation	long-term	state-dependent	long-term	state-dependent	long-term	state-dependent
congestion (70%)	3.2	4	3.6	4	0.4	0
normal (20%)	6.4	6.4	7.2	7.2	0.8	0.8
off-peak (10%)	6.4	7.2	7.2	9.6	0.8	2.4
expected value	4.48	5.04	5.04	5.52	0.56	0.48

Table 3
Profit, Social Welfare, and Consumer Surplus (Example III)

	profit		social welfare		consumer surplus	
situation	long-term	state-dependent	long-term	state-dependent	long-term	state-dependent
congestion (60%)	3.1	4.2	3.65	4.2	0.55	0
normal (20%)	6.2	6.2	7.3	7.3	1.1	1.1
off-peak (20%)	6.2	7.2	7.3	9.7	1.1	2.5
expected value	4.34	5.2	5.11	5.92	0.77	0.72

3.2 Simulation Results

This section, we consider a network service model that is defined in Section 2.2 and use simulation results to demonstrate the effect described in Section 3.1. The change of packet arrival rate with respect to price is specified by the following demand function:

$$\lambda_{\mathrm{b}} = \lambda_{\mathrm{bmax}}[1-(\frac{p_{\mathrm{b}}}{p_{\mathrm{bmax}}})^{\alpha_{\mathrm{b}}}] \qquad (3)$$

λ_{bmax} is the maximum packet arrival rate (arrival rate when the price is 0), p_{bmax} is the maximum willingness to pay per packet of all consumers, and α_{b} is a parameter. When α_{b} is small, at the same price, there will be smaller number of packet arrivals. Moreover, a slight increase in price will result in a large decrease in packets arrival rate. Therefore, the value of α_{b} indicates both the strength of demand and demand elasticity.

We run the same simulation for multiple times, each time using different seeds to generate random numbers. Results from those runs are considered as independent samples of output variables. With probability 95%, mean values of revenue, social welfare are accurate within $\pm 3\%$, and that of consumer surplus are accurate within $\pm 8\%$. T-test is used to compare means of those sample values of different pricing schemes.

We first consider a base case of which input parameters are shown in Table 4. We will first show that in comparison with long-term average pricing, response pricing and spot pricing improve throughput. We will then demonstrate that for the same reason given in Section 3.1, both schemes result in higher profit and social welfare, but in some cases, can harm consumer benefits.

Table 4
Parameters Used in the Simulation (Base Case)

Symbol	Value	Interpretation
C_{T}	3640 packets/sec.	total capacity
guaranteed service		
λ_{g}	12/min.	call arrival rate
r_{c}	0.8/min.	call departure rate
λ_{c}	100/sec.	packet arrival rate
M_{c}	18	maximum number of calls can be carried
best-effort service		
λ_{bmax}	10000 packets/sec.	maximum packet arrival rate
p_{bmax}	$5*10^{-7}$/packet	maximum willingness to pay
a_{b}	0.5	demand elasticity parameter
B	20	maximum buffer size

In this case, the optimal long-term average price is $3.028*$10^{-7}$/packet. Figure 1 shows response price as a function of the number of guaranteed calls in progress, and figure 2 shows spot prices as a function of buffer occupancy given the number of guaranteed calls in progress is 0, 9, and 18. As those figure show, when the network is less congested, i.e. when the number of guaranteed calls is small and/or buffer occupancy is low, both response price and spot price stay constant because they are based only on demand and not constrained by capacity availability. As the network gets more congested, those prices increase, which means only more valuable packets can be transmitted during those periods.

Figure 1

Response Prices

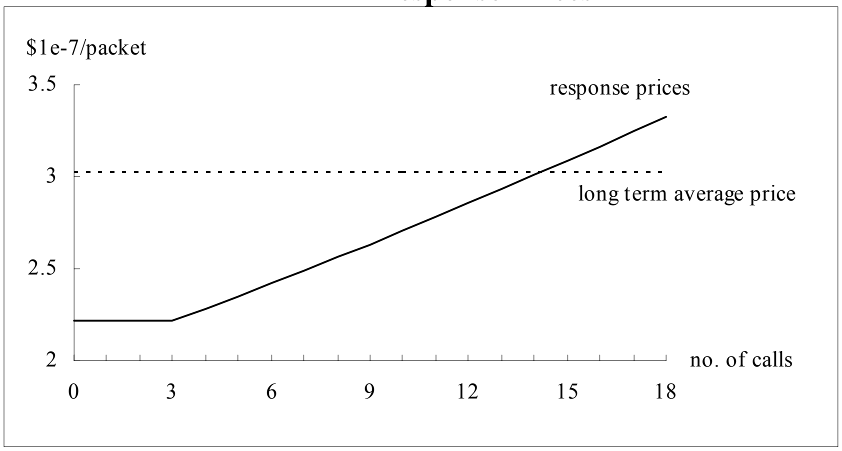

Figure 2
Spot Price Trajectories

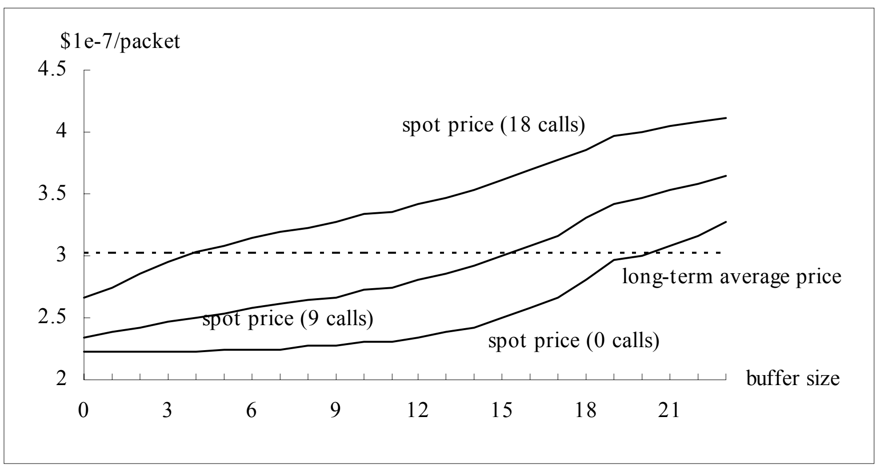

The role of state-dependent pricing in improving traffic management is also demonstrated by simulation results. Figure 3 shows distribution of buffer occupancy under different pricing schemes. Notice that zero buffer occupancy means capacity lays idle and full buffer occupancy means packets have to be dropped. Therefore, it is desirable to have buffer occupancy distribution take lower values at points where the buffer is full or empty. Based on this criteria, the distribution for spot pricing is best, and response pricing is better than long-term average pricing. This is because the price is directly related with buffer occupancy under spot pricing. Therefore, the service provider is most capable of adjusting prices under spot pricing scheme, therefore increase or decrease packets arrival rate when the buffer gets empty or full. Response pricing is also more capable of doing that than long-term average pricing since the latter fixes the price for all time while the former does not.

Figure 3
Distribution of Buffer Occupancy Under Different Pricing Schemes

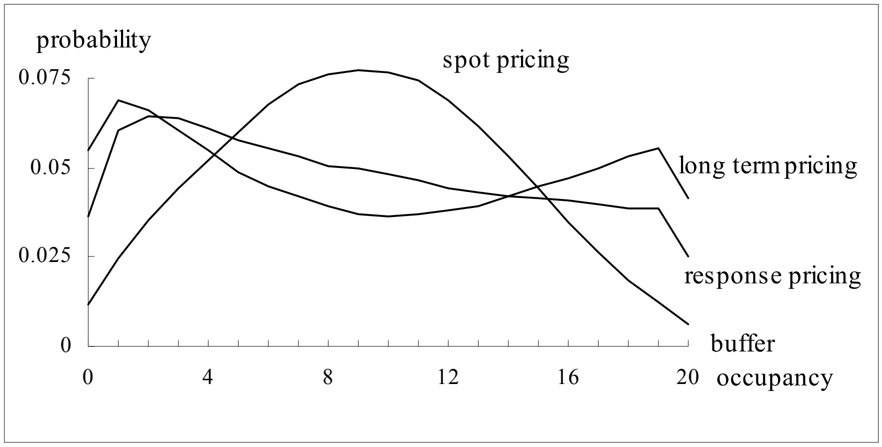

The fact that the buffer has a higher probability of being partially full under spot pricing leads to superior throughput under that scheme. The simulation shows that in a 60 minutes segment, mean throughput is 8.0 million packets under spot pricing, 7.7 million packets under response pricing, and 7.5 million packets under long-term average pricing (accuracy $\pm 4\%$, with probability 95%). T-test shows with 99.5% statistical significance that throughput is better under response pricing than that under long-term average pricing, and better under spot pricing than that under the other twos.

The same as the example in Section 3.1, adopting state-dependent pricing improves both profit and social welfare. Table 5 shows that profit and social welfare achieved in a 60 minutes segment under different pricing schemes. T-test shows with 99.5% confidence that profit and social welfare under spot pricing are higher than those under response pricing, which are higher than those under long-term average pricing.

Table 5
Mean Profit and Social Welfare under Different Pricing Schemes
($/60 minutes)

	Spot Pricing	Response Pricing	Long-term Pricing
revenue	2.37	2.33	2.28
social welfare	2.99	3.06	3.15

The simulation also shows improvement in consumer surplus under state-dependent pricing. The means of consumer surplus are 0.71, 0.73, and 0.78 under spot, response, and long-term average pricing, respectively. Again, T-test shows differences among those means are at 99.5% significance level.

To test the robustness of results from the base case, we use different values of α_b, the demand elasticity parameter, and run the simulation. Figure 4 displays mean value of profit under each pricing scheme given different values of α_b. As α_b increases, the same price can induce more packets arrival. Consequently, mean profit increases under every pricing scheme. Moreover, the graph shows mean profit under spot pricing is higher than the mean profit under response pricing, and mean profit under each of them is higher than the mean profit under long-term average pricing. T-test shows in all cases, differences of mean profits between two pricing schemes are significant at 99.5% significance level.

Figure 5 shows changes of mean value of social welfare. Like profit, social welfare increases with α_b under each pricing scheme, and state-dependent pricing results in higher social welfare than long-term average pricing. T-test also shows those differences in mean values are statistically significant. Nevertheless, the same trend can not be observed in changes of consumer surplus. When α_b becomes larger, demand becomes less elastic, so the carrier can charge a higher price and extract more wealth from consumers. As Figure 6 shows, consumer surplus falls with α_b, and the complicated the pricing scheme, the faster the rate of decrease. As a result, consumer surplus under state-dependent pricing is higher when α_b is small and lower when α_b is large. For example, in comparison with long-term average pricing, consumer surplus is 9.1% higher under spot pricing and 2.6% higher under response pricing when α_b=0.5 (both differences are at 99.5% significance level). However, when α_b=1.0, consumer surplus is 5.1% lower under spot pricing than long-term average pricing (at 99.5% significance level), and the difference between response pricing and long-term average pricing is not statistically significant.

Figure 4
Changes of Profit with α_b

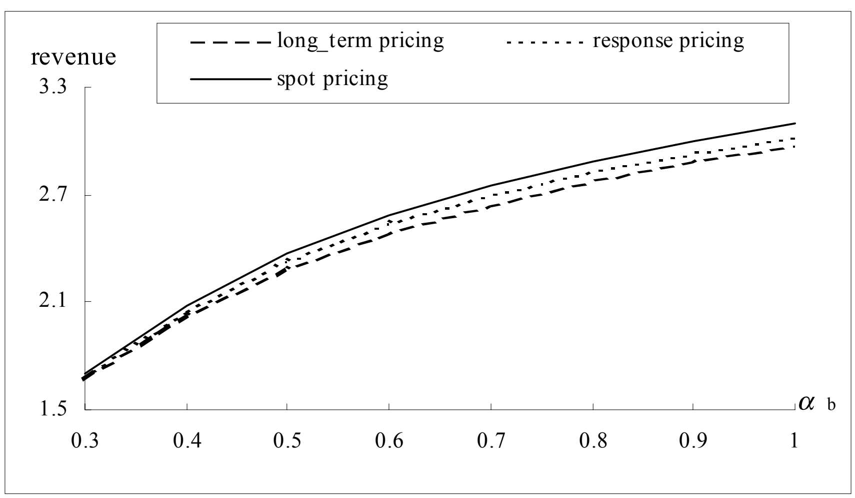

Figure 5
Changes of Social Welfare with α_b

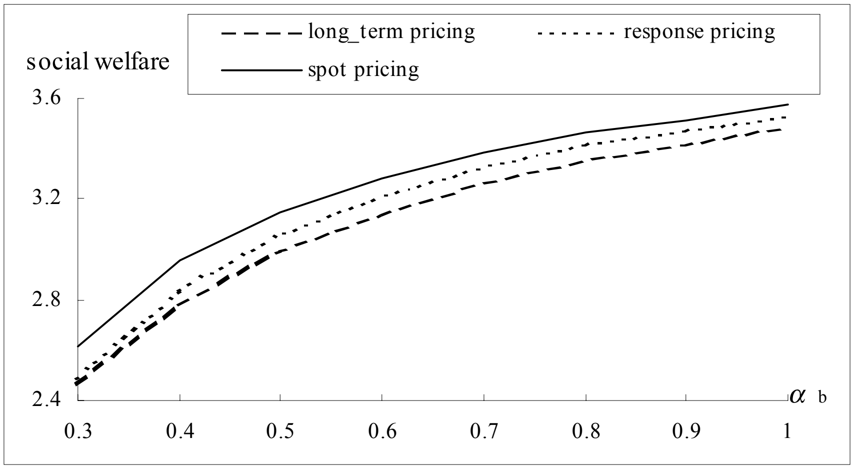

Figure 6
Changes of Consumer Surplus with α_b

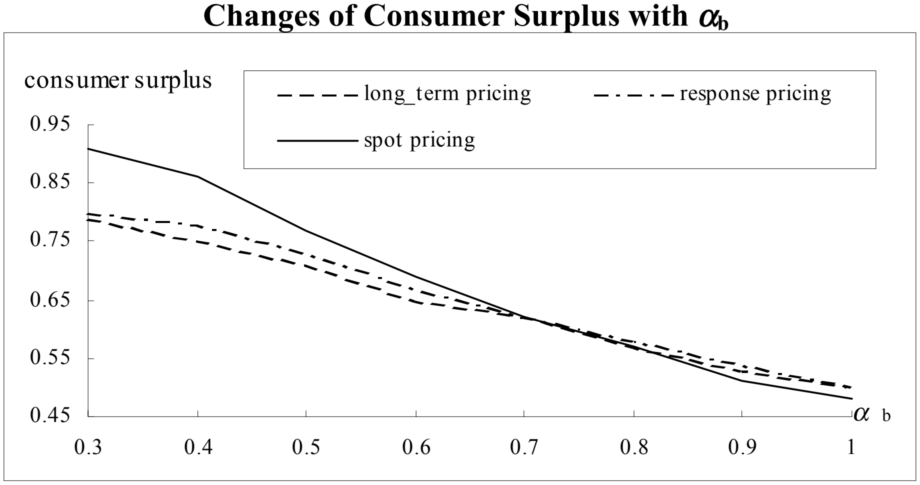

The faster fall of consumer surplus under state-dependent pricing is consistent with the intuition we derive from the example in 3.1. State-dependent pricing achieves higher profit and social welfare by better traffic management, which benefit consumer surplus, and by raising price during congested periods, which can be detrimental to consumer surplus. Small α_b means elastic demand, so a slight change in price can cause large variations in packet arrivals. In this situation, to keep sufficient number of packet arrivals, the service provider can only change state-dependent prices by small increments, so the effect of throughput improvement is more significant than the effect of price increase. As a result, consumer surplus grows with the adoption of state-dependent pricing. Large α_b means inelastic demand, so there will be smaller changes in traffic flow in response to the price variation. Therefore, varying the price has more do to with increasing value of packets to be transmitted and deriving more profit from those packets, as opposed to improving throughput. As a result, consumer benefits will suffer.

We also conducted other simulations in which buffer size takes different values. Figures 7,8 and 9 show profit, social welfare, and consumer surplus under each scheme. As T-test shows, with a small buffer size, profit and social welfare are higher under state-dependent pricing than those under long-term average pricing, but there are no difference in consumer surplus between response pricing and long-term average pricing. Consumer surplus under spot pricing are significantly lower than both response pricing and long-term average pricing. This phenomenon is consistent with the insights we developed in the example of Section 3.1. The discussion demonstrates as congested periods occur more frequently, the service provider relies more on raising price to send more valuable packets as opposed to increasing throughput to improve profit. As a result, the benefits of adopting more complicated state-dependent pricing will mainly be captured by the service provider. Since a smaller buffer size means more stringent capacity constraint, thus a higher likelihood of occurrence of congestion, there should be no surprise that adopting state-dependent pricing won't benefit consumers.

Figure 7
***t*-Test of Profit Comparisons (buffer size change)**

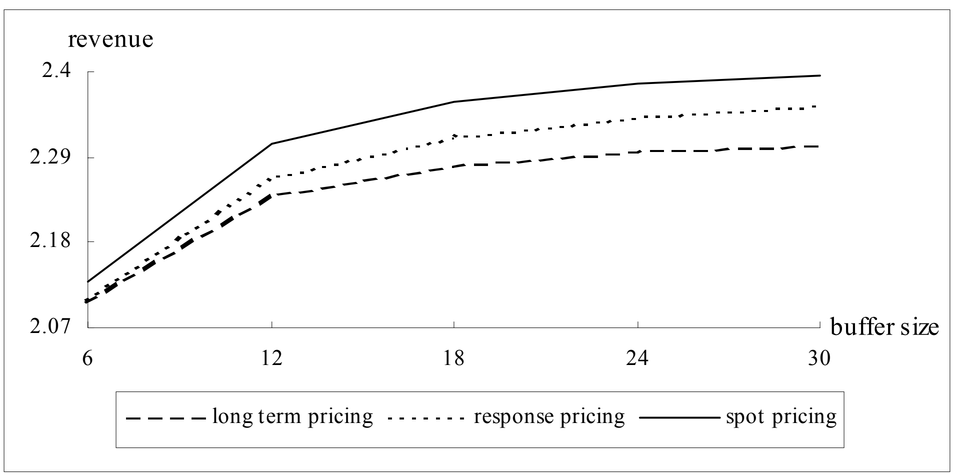

Figure 8
Changes of Social Welfare with Buffer Size

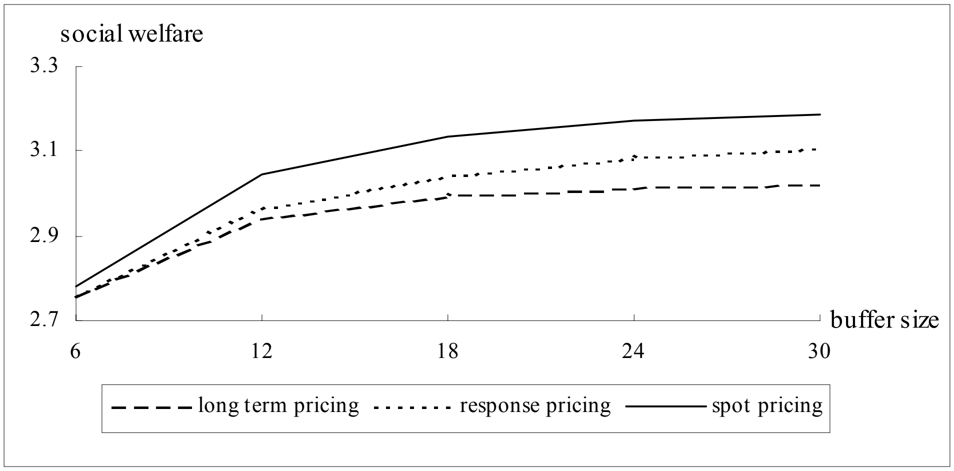
social welfare
3.3
3.1
2.9
2.7
6
12
18
24
30
buffer size
long term pricing
response pricing
spot pricing

Figure 9
Changes of Consumer Surplus with Buffer Size

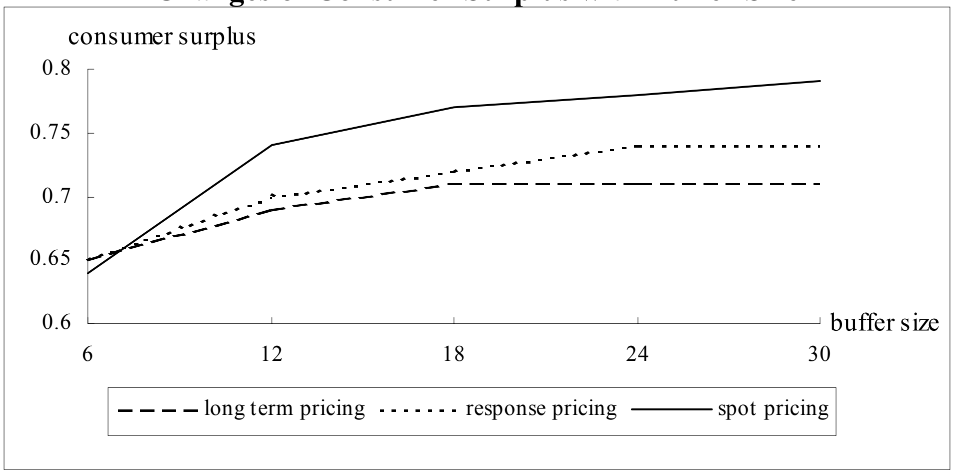

Section 4 Conclusions on Policy Implications

In this paper, we discussed the implementation of state-dependent pricing, and use simulations to compare two such schemes, spot pricing and response pricing, with long-term average pricing. Simulation results show that state-dependent pricing can achieve a higher profit and social welfare than long-term average pricing. This is largely due to two effects: 1) state-dependent pricing serves as a traffic management mechanism that results in a higher throughput of packets; and 2) state-dependent pricing has the flexibility of raising price based on network status. Therefore, more valuable traffic will be carried during congested periods. While improving traffic management benefits consumers as well, raising price during congested periods enables the service provider to extract more wealth from consumes. Therefore, in some cases, while profit and social welfare will be enhanced by adopting state-dependent pricing, consumer surplus will be reduced.

Those analysis are carried out without considering the impact of congestion control mechanisms. As demonstrated in literature [PEHA97], some congestion control mechanisms are more efficient than state-dependent pricing in maximizing throughput. Therefore, in networks equipped with those congestion control mechanisms, improvements of profit and social welfare achieved by state-dependent pricing would be less significant since those schemes no longer have the advantage of increasing throughput.

State-dependent pricing can also improve profit by raising price during congested periods so limited capacity can be used to send more valuable packets. This feature can not be implemented through technical means. Therefore, the service provider may have incentives to adopt state-dependent pricing even in networks equipped with a good congestion control mechanism, as long as the additional revenue exceeds the cost of building complicated billing and accounting systems. Sending more valuable packets during congested periods improves the efficiency of the service, so social welfare increases as well. However, because price increase, consumer benefits of transmitting more valuable packets is not guaranteed. It is possible that the state-dependent prices increase faster than the value of packets currently being transmitted, so users will be worse-off. The paper demonstrated cases in which consumer surplus either does not change or decreases when state-dependent pricing is adopted, even though both profit and social welfare improve. In those cases, whether to adopt state-dependent pricing presents an interesting tradeoff between maximizing social efficiency and protecting consumer benefits.

REFERENCES:

[DANI97] Danielsen, K. and Weiss, M. (1997) "User Control Modes and IP Allocation," *Internet Economics*, Lee McKnight and Joseph Baileys eds., The MIT Press, pp. 305-322.

[GUPT97] Gupta, A., Stahl, D. O., and Whinston, A. B. (1997) "Priority Pricing of Integrated Services Networks," *Internet Economics*, Lee McKnight and Joseph Bailey eds., The MIT Press, pp. 323-352.

[MACK95] Mackie-Mason, J. F. and Varian, H. (1995) "Pricing Congested Network Resources," *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 7, pp. 1141-1149.

[MACK97] Mackie-Mason, J. F., Murphy, L. and Murphy, J. (1997) "Responsive Pricing In The Internet," *Internet Economics*, Lee McKnight and Joseph Bailey eds., The MIT Press, pp. 279-304.

[PEHA97] Peha, J. M. (1997) "Dynamic Pricing as Congestion Control in ATM Networks," *IEEE Globecom-97*.

[WANG97] Wang, Q., Peha, J. M. and Sirbu M. A. (1997) "Optimal Pricing for Integrated-Services Networks," *Internet Economics*, Lee McKnight and Joseph Bailey eds., The MIT Press, pp. 353-378.