ISP Uplink Pricing in a Competitive Market

Qian Wang Dah Ming Chiu
Department of Information Engineering
The Chinese University of Hong Kong
Email: {wq007, dmchiu}@ie.cuhk.edu.hk

John C.S. Lui
Department of Computer Science & Engineering
The Chinese University of Hong Kong
Email: cslui@cse.cuhk.edu.hk

Abstract—**In recent years, the upgrade of access networks to broadband networks together with the P2P technology has brought many new applications to the Internet. P2P applications have quickly become the biggest consumer of network resources. ISPs of access networks as well as backbone networks are all wondering how to better manage their network resources. We explore the idea of uplink pricing as a way to provide differential pricing to P2P and regular users. In particular, we formulate a simple economic model to analyze under what scenarios uplink pricing will be adopted by all ISPs in a competitive market.**

Index Terms—**ISP, P2P traffic management, network economics**

I. INTRODUCTION

It is no secret that the Internet is *filled* with Peer-to-Peer (P2P) traffic. P2P-based content distribution effectively distributes the load from a single server and its uplink to all the receivers of the content and the rest of the network. More large-scale Internet content providers (ICPs) are reportedly looking into this technology. While the P2P technology is going through improvements (in its network efficiency) via experimentation and research, the expectation is that it will continue to demand more and more network bandwidth.

The natural question to ask is whether Internet Service Providers (ISPs) will be able to estimate the P2P traffic growth and provision enough bandwidth for P2P users, and if not quite enough, how ISPs will be able to manage the resource contention by different types of users? Many ISPs, from backbone ISPs to small campus network administrators are all grappling with these questions.

In a non-competitive scenario, the network administrators can instrument a policy by fiat. For example, a campus network may let the P2P users consume the currently provisioned bandwidth to its capacity (the critical resource is often the link connecting to the uplink service provider), and then apply some measure to limit further growth of P2P traffic. In a competitive market, the situation is more interesting and complicated. Controlling P2P traffic may cause an ISP to lose market share, yet not controlling P2P traffic may alienate non-P2P users and increase the ISP's operation costs.

Leaving the practicality issues aside for a moment, it seems a perfectly reasonable approach is by properly applying pricing to reflect the utility of network resource usage. Internet users are used to flat-rate pricing. The reasons are mostly psychological - a consumer prefers not to repeatedly spend the energy making small decisions for incremental network usage [1]. Internet content providers, however, often have to negotiate private deals with their Internet service providers. The pricing is usually based on a combination of bandwidth usage (which costs the ICPs), as well as the value of the ICP content to the ISPs (which costs the ISPs). The use of the P2P technology shifts the bandwidth usage from the ICPs to the users, and at the same time makes all the users little ICPs (by offering content). Arguably, the negotiated ICP pricing must also be shifted to the P2P users.

In this paper, we consider a simple model of a competitive ISP market, to study if a new pricing scheme called *uplink pricing* will be adopted by ISPs to control P2P traffic. An ISP is said to adopt uplink pricing if it charges its users at a fraction of the original flat rate pricing for downloading services, and imposes a usage-based charge for uploading services. We defer the problem of how ISPs should set the uplink price, to (presumably) maximize their profits, as a topic for future studies. Instead, we assume the ISPs would take a profit-neutral stance when setting the uplink price. Under this setting, we analyze whether a single ISP adopting uplink pricing would lead to this form of pricing adopted by the market, or would it lead to the co-existence of both forms of pricing, or would it result in the market revert back to flat-rate pricing. The contribution of this paper is to determine the factors and conditions which lead to various equilibrium outcomes.

The rest of the paper is organized as follows. In Section II, we set up our model. In Section III, we apply a game-theoretic analysis to define all the possible outcomes, and give examples to illustrate how these outcomes can be realized. In Section IV, we derive the conditions for different outcomes to occur. In Section V, we discuss the implications of our analysis, and how our model can be extended to relax some of the assumptions. We also briefly discuss related works. In the last section, we give a conclusion for this paper.

II. A SIMPLE MODEL OF A COMPETITIVE ISP MARKET

We consider the simplest representation of a competitive ISP market, consisting of two ISPs competing for a population of $2N$ subscribers. Initially, both ISPs adopt *flat-rate* pricing, with the same price p. Each ISP has a subscriber population of N.

The idea of *uplink pricing* is to divide the user subscription price into two parts: (a) a flat-rate charge as before; and (b) a usage-based[1] component which is designed to charge the user when it is behaving as a server (or ICP). Under uplink pricing, the subscription price, $\tilde{p}$, can be expressed as:

$$\tilde{p} = \frac{p}{2} + vq. \tag{1}$$

For simplicity of presentation (one fewer parameter), we assume the flat-rate part is half of the original flat-rate price. The parameter v represents the uplink traffic volume generated by a user, and q is the charge per volume of traffic.

We are interested in studying the user behavior due to uplink pricing, and in turn how the new market conditions affect the eventual adoption of uplink pricing by ISPs. We make a number of assumptions about types of users, traffic, ISP costs and ISP behavior (pricing decisions), and user behavior (bandwidth usage and ISP choices).

A. User Types

There are two types of users, *regular* users and *P2P* users. A regular user generates negligible uplink traffic compared to a P2P user, whereas a P2P user generates uplink traffic at a constant rate of $v = V$. Out of the whole subscriber population, a fraction β is of the P2P type.

B. Traffic

Regular users generate an insignificant amount of uplink traffic in comparison to P2P users. Out of the uplink traffic generated by a P2P user, V, a fraction α leaves the ISP's network, which contributes to the ISP's peering costs.

C. ISP Costs

Each ISP has sufficient funds for capital investments to support all users in the market if necessary, so we do not explicitly consider capital costs, but only an ISP's operating costs. We assume there are three components to the operating cost:

$$C(n, k) = C_f + C_m n + C_t v_r(k). \tag{2}$$

The first component C_f is a fixed cost; the second component depends on the number of users, where n is the subscriber population size and C_m is the marginal cost per additional subscriber; the third component depends on the outbound external traffic volume, v_r. Notice that v_r is a function of the number of P2P users, denoted by k, and C_t is the marginal cost per outbound external traffic volume. Normally, an ISP's payment for external traffic depends on both inbound and outbound traffic volume. Here, we assume the cost for inbound external traffic, if any, is part of the second component of Eq. (2). From the previous assumption on traffic, we get

$$v_r(k) = \alpha k V. \tag{3}$$

The fixed cost in Eq. (2) captures a simple form of *economies of scale*, namely, the more subscribers (or remote traffic), the less the per user (or per traffic volume) cost. As we will discuss in section V, in order to represent different levels of efficiency of the P2P technology, we will need to introduce more sophisticated models of economies of scale.

D. User Behavior

Users are assumed to select ISP based on price, to minimize what they pay. When both ISPs have the same pricing scheme, however, we assume the market is *symmetric*. In other words, each of the two ISPs will have half of the subscribers and the user types are also equally distributed in these two ISPs.

Another aspect of user behavior is how the P2P users react to uplink pricing. We assume there is some degree of elasticity represented by a parameter ρ. Given uplink price $\tilde{p} > p$, with probability ρ, a P2P user will lower its uplink usage v to maintain the same payment as under flat-rate pricing. If $\rho = 0$, it means the P2P traffic is *inelastic*, i.e. all users would stick to their original P2P traffic levels; if $\rho = 1$, however, it means that *all* users would decrease their P2P traffic to maintain their previous payment levels. On the other hand, we assume the uplink usage v is the maximum a P2P user would incur. Even if $\tilde{p} < p$, a P2P user would not increase its P2P traffic.

E. ISP Behavior and Market Assumptions

We assume the ISP's profit is simply the sum of the payments from its subscribers minus its costs. Initially, under flat-rate pricing, based on the assumption of user behavior, $k = \beta N$, so the profit for each ISP can be expressed as:

$$P_o = Np - (C_f + NC_m + \alpha\beta VNC_t). \tag{4}$$

Without losing generality, suppose ISP_1 first converts to uplink pricing. The most critical question is how ISP_1 would set the incremental price q. After the conversion, the regular users will pay less ($p/2$ vs p). It is reasonable to assume the P2P users will pay more than what they pay under flat-rate pricing, hence:

$$q \geq \frac{p}{2V}, \tag{5}$$

but exactly how much the converting ISP will charge is a rather complicated question. On the one hand, ISPs want to maximize their profit by extracting as much out of the paying customers as possible depending on their utility functions; on the other hand, ISPs must also be concerned about market share, and growth of the business. For this paper, we assume the converting ISP takes a *profit-neutral* position. In other words, the converting ISP assumes that if it has the same customer base after the conversion, its profit would stay neutral. This also implies that if both ISPs convert to uplink pricing, they will both stay profit-neutral, since, in the case both convert they will become symmetric again based on the user behavior assumption, hence it is equivalent to each ISP retaining the same customer base.

---

[1] In general, the uplink price does not need to be usage-based. The ISP strategy analysis in this paper does not depend on the assumption of a usage-based uplink price either. However, we believe it is more reasonable to consider a usage-based uplink price because it is the form of pricing between the ISPs and ICPs.

Let P' denote the profit for each ISP after both convert. Based on the form of P2P traffic elasticity assumed above, we have

$$\begin{aligned} P' &= N\frac{p}{2} + \rho\beta N\frac{p}{2} + (1-\rho)\beta NVq \\ &\quad -(C_f + NC_m + (\rho\frac{p}{2q} + (1-\rho)V)\alpha\beta NC_t). \end{aligned} \quad (6)$$

The first term is the flat-rate pricing contribution by all subscribers; the second term is the usage-based payment made by the P2P users with elastic traffic; and the third term is the usage-based payment made by the inelastic P2P users. The forth term is the cost of the converted ISP. According to the profit-neutral assumption, we can derive q by setting $P' = P_o$.

Assume ISP_1 adopts the uplink pricing first, given P2P users will be paying more under uplink pricing than flat-rate pricing (Eq. (5)), and the user behavior assumed above, it follows that all P2P users will move to the flat-rate pricing ISP (ISP_2), and all the regular users will move to the uplink pricing ISP (ISP_1). As a result, the number of users in each network becomes:

$$n_1 = 2N(1-\beta), \quad (7)$$

$$n_2 = 2N\beta. \quad (8)$$

It follows that after ISP_1 converts to uplink pricing, the profits for the two ISPs become P_1 and P_2 respectively:

$$P_1 = n_1\frac{p}{2} - (C_f + n_1C_m), \quad (9)$$

$$P_2 = n_2p - (C_f + n_2C_m + \alpha Vn_2C_t). \quad (10)$$

To summarize, we list all the notations of our model in the following table.

Symbol	Explanation
$2N$	Total number of users in this market
n	Number of users in one ISP
k	Number of P2P users in one ISP
C_f	Fixed operation cost
C_m	Marginal cost per additional subscriber
C_t	Marginal cost per outbound external traffic volume
p	Original flat-rate price
$\tilde{p}$	Subscription price under uplink pricing
q	Price for per uplink traffic volume
v	Uplink traffic volume generated by one user
$v_r(k)$	Outbound external traffic volume generated by k P2P users
β	Fraction of P2P users
α	Fraction of traffic leaving the ISP's network
ρ	Degree of elasticity of P2P usage
P_o	Original profit of each ISP
P_i	Profit of ISP_i after ISP_1 converts to uplink pricing
P'	Profit of each ISP after both ISPs convert

TABLE I
NOTATIONS

III. GAME-THEORETIC ANALYSIS

From an ISP's perspective, whether to convert to uplink pricing depends on the possible outcomes of this conversion. An ISP would be the first one to adopt uplink pricing under two conditions: 1) if it will get a higher profit than its original profit; 2) if the other ISP would also convert.

		ISP_2	
		flat-rate pricing	uplink pricing
ISP_1	flat-rate pricing	(P_o,P_o)	(P_2,P_1)
	uplink pricing	(P_1,P_2)	(P_o,P_o)

TABLE II
PAYOFF MATRIX OF ISP PRICING GAME

In general, the situation can be viewed as a Stackelberg game with ISP_1 as the leader. The payoff matrix of this game is shown in Table II. The payoff when both ISPs adopt uplink pricing is the same as when both adopt flat-rate pricing, due to the profit-neutral price setting assumption.

In a Stackelberg game, the leader knows the other player's reaction to its own actions. With the knowledge of all the possible outcomes, the leader can then choose its own action to maximize its payoff.

In this case, if ISP_1 converts first, the resulting profit distribution would be (P_1, P_2). If $P_o > P_1 > P_2$, ISP_2 would predictably follow ISP_1 and also convert, which leads to a better outcome. From ISP_1's point of view, although it earns less money than before after the conversion, this is only temporary since it knows that after ISP_2 follows its decision to convert, they can both return to the original level of profitability. Therefore, we can conclude that if the parameters in our model lead to the condition $P_o > P_1 > P_2$, then ISP_1 would be willing to adopt uplink pricing first and would expect ISP_2 to follow.

There are seven possible relationships between P_o, P_1, P_2. Based on the same kind of logic, we conclude these distinct conditions may lead to three different outcomes, as shown in Table III.

Condition	Outcome
$P_o > P_1 > P_2$	Both choose uplink pricing
$P_1 > P_o > P_2$	
$P_o > P_2 > P_1$	Both choose flat-rate pricing
$P_2 > P_o > P_1$	
$P_1 > P_2 > P_o$	Co-existence of flat-rate and uplink pricing
$P_2 > P_1 > P_o$	
$P_o = P_1 = P_2$	

TABLE III
CONDITIONS ON PROFIT RELATIONS AND CORRESPONDING OUTCOMES

To illustrate some possible outcomes, let us consider some typical scenarios. The payoff matrix of each case is shown in Table IV (”UP” denotes ”uplink pricing” for short in the table):

Case 1 *Parameters:* $N = 100, p = 1, C_f = 10, C_m = 0.4, C_t = 0.25, \alpha = 0.3, V = 20, \beta = 0.2, \rho = 0.9$.

Under flat-rate pricing, each ISP earns a revenue of 100 (pN), incurring total costs of 80 (fixed cost 10, management cost 40, and bandwidth cost 30), so the profit is 20 (Eq. (4)). After ISP_1 converts to uplink pricing, ISP_1 would lose all P2P customers and earn a revenue of 80, with cost 74, which leaves a profit of 6 (Eq. (9)). Meanwhile, ISP_2 would get all the P2P customers to earn a revenue of 40, with cost 86, so ISP_2 would run a deficit of 46 (Eq. (10)). For ISP_2 to stay in business, a wise choice is to follow the ISP_1's decision to apply uplink pricing. Based on the profit-neutral uplink pricing assumption, both ISPs would set the usage-based part of uplink price to 0.39, which is reasonable.

Outcome: Both choose the uplink pricing strategy.

Case 2 *Parameters:* $N = 100, p = 1, C_f = 10, C_m = 0.55, C_t = 0.1, \alpha = 0.5, V = 6, \beta = 0.5, \rho = 0.9$.

Compared with Case 1, we now change the value of C_m, C_t, α, β, and V. The profit-neutral uplink price becomes 0.54. Applying similar calculations, we find that the original profit of each ISP is 20, ISP_1 gets a deficit of 15 after applying the uplink pricing, and ISP_2 gets a profit as 5. From ISP_1's perspective, this conversion to the uplink pricing does no good to itself, so it will revert to flat-rate pricing.

Outcome: Both choose the flat-rate pricing strategy.

	Case 1		Case 2	
	flat-rate	UP	flat-rate	UP
flat-rate	(20,20)	(-46,6)	(20,20)	(5,-15)
UP	(6,-46)	(20,20)	(-15,5)	(20,20)

TABLE IV
PAYOFF MATRICES OF THE TWO EXAMPLES

Given the current assumptions in our model, there is no example for the *co-existence* outcome. After ISP_1 converts to uplink pricing, the total P2P traffic in this market does not decrease since all the P2P users can switch to ISP_2 (flat-rate pricing hence unrestrained P2P traffic). Therefore, the total costs of these two ISPs remain the same as before. On the other hand, the total revenue of the two ISPs has reduced since regular users are charged only half of the original price and P2P users are charged the same price as before. That means the total profit in the market has reduced:

$$P_1 + P_2 < 2P_o. \tag{11}$$

Refer to Table III, we can conclude that the outcome of co-existence of both pricing strategies cannot be realized under the current problem formulation. In general, however, if usage-based pricing leads to increased efficiency hence increased profit for both ISPs, then co-existence is still possible. This will be further discussed in Section V.

IV. WHEN WILL UPLINK PRICING BE ADOPTED

In the previous section, we point out three different outcomes in deploying a new pricing scheme in a competitive ISP market, namely (1) both adopt new pricing, (2) both stay with old pricing, and (3) co-existence of different pricing schemes. Simple game-theoretic analysis illustrates how to determine the outcome of a particular market scenario, based on computing P_o, P_1, P_2 and hence the payoff matrix. In this section, we further analyze the parameter space and characterize all the conditions that lead to each outcome.

Since the incremental uplink price q is determined by other parameters under the profit-neutral assumption, the key parameters of the model are: flat price (p), and the ratio of P2P users (β). Using Eq. (4), (9) and (10) and the rules in Table III, we can summarize all the conditions that lead to each outcome in terms of different values of p and β and their relations to various ISP cost parameters, as shown in Fig. 1, 2 and 3.

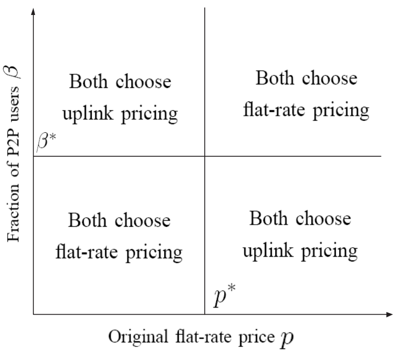

Fig. 1. Pricing outcomes resulting from different model parameters when $2C_m = C_m + \alpha V C_t$

In the parameter space of p and β, there is a *watershed* for each value, as shown in Fig. 1. The watershed for β is defined by β^*:

$$\beta^* = \frac{\frac{p}{2} - C_m}{(\frac{p}{2} - C_m) + (p - C_m - \alpha V C_t)}. \tag{12}$$

The watershed for p is defined by two values:

$$p^* = 2C_m, \tag{13}$$

$$p^{**} = C_m + \alpha V C_t. \tag{14}$$

For ease of presentation, we temporarily assume $p^* = p^{**}$, as shown in Fig. 1.

The result can be summarized as follows, when $p < p^*$ and $\beta > \beta^*$, or when $p > p^*$ and $\beta < \beta^*$, a single ISP adopting uplink pricing will lead to the adoption of uplink pricing by the market (both ISPs). Otherwise, both ISPs will stay with flat-rate pricing. Intuitively, we can explain the reasons as follows:

1) If $p < p^*$, meaning that $p < 2C_m$ and $p < C_m + \alpha V C_t$, it is easy to see that both ISPs would be running a deficit for each additional user or additional volume of traffic. Under this condition, the larger the fraction of P2P users (larger β), the more losses for ISP_2, and therefore ISP_2 will follow ISP_1 to adopt uplink pricing. This is true unless when the fraction of P2P users is so small, smaller than a threshold as β^*, such that ISP_1 would get too many regular users to cause it to lose more money than ISP_2. In this latter case, ISP_1 will revert to flat-rate pricing.
2) If $p > p^*$, meaning that $p > 2C_m$ and $p > C_m + \alpha V C_t$, then both ISPs can expect a positive profit from each additional user, so that both of them prefer larger market share. Larger β would decrease ISP_1's market share, so ISP_1 would revert. On the other hand, smaller β would decrease ISP_2's market share, and therefore encourage ISP_2 to convert.

To make it easier to understand Fig. 1, we temporarily assumed $p^* = p^{**}$, or in other words $C_m + \alpha V C_t = 2C_m$. The cases when $p^* < p^{**}$ and $p^* > p^{**}$ are depicted in Fig. 2 and 3. In these cases, when $p < \min(p^*, p^{**})$ or when $p > \max(p^*, p^{**})$, the result is the same as shown in Fig. 1. When $\min(p^*, p^{**}) < p < \max(p^*, p^{**})$, the outcome is independent of β, but depends on whether p^* or p^{**} is larger:

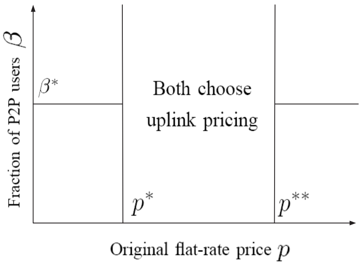

Fig. 2. Pricing outcomes resulting from different model parameters when $2C_m < C_m + \alpha V C_t$

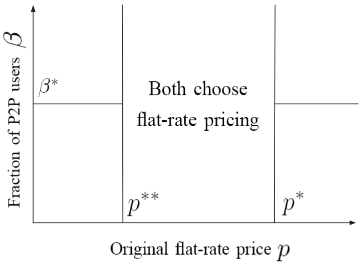

Fig. 3. Pricing outcomes resulting from different model parameters when $2C_m > C_m + \alpha V C_t$

1) If $2C_m < p < C_m + \alpha V C_t$, the situation is ISP_1 will make a profit whereas ISP_2 will lose money. Therefore, whatever is the fraction of P2P users (β), ISP_2 will definitely follow ISP_1 and adopt uplink pricing.
2) If $C_m + \alpha V C_t < p < 2C_m$, however, the situation is just the opposite. Irrespective of β, ISP_1 will lose money whereas ISP_2 will make a profit; so ISP_1 will revert to flat-rate pricing.

Although the analysis assumes ISP_1 makes a move first, and then possibly reverts to flat-rate pricing, such an exercise does not necessarily need to be carried out. ISPs can estimate some of the parameters through network measurement and other parameters through market studies. The analysis in this paper can then shed some light on the likely outcomes.

If we assume the current ISP market competition is fierce, p is possibly lower than p^* and p^{**}. In that case, the market would convert to uplink pricing only if β is sufficiently large.

V. FURTHER DISCUSSIONS OF THE MODEL

A. Distribution of P2P Users

In the ISP model, we assume that there are only two types of users in the market, regular users with negligible upload traffic and P2P users with constant upload volume V. A more realistic model would be to represent the users' upload volume as a random variable, following either an exponential distribution or some heavy tail type of distribution. In this case, the expected ISP profits before and after (one ISP) adopting uplink pricing can still be derived. Based on $E[P_o]$, $E[P_1]$ and $E[P_2]$, the same game-theoretic analysis can be carried out to characterize what market conditions will lead to uplink pricing adoption. Due to space limitations, we have not included the analysis in this paper.

B. Modeling External P2P Traffic

In this paper, our analysis is based on a crude model of P2P external traffic. The external traffic generated by a peer is assumed to be a fraction of the peer's uplink traffic. This assumption implies that there is no economies of scale; in other words, as the number of peers in an ISP increases, the external traffic per peer does not decrease. It is not hard to show that in this case, the condition for co-existence cannot arise.

Nonetheless, there ought to be some economies of scale by intuition, although this may be difficult to model. One way to approach is to assume that there are a total of $\hat{N}$ P2P peers (engaging in some common P2P activity) in the entire Internet, and each peer in the local market randomly selects other peers to communicate with. Recall that k represents the number of P2P users in one ISP, so the volume of outbound external traffic of one ISP can be expressed as

$$v_r(k) = kV(1 - \frac{k}{\hat{N}}). \tag{15}$$

Comparing Eq. (15) with Eq. (3), we can see that the fixed percentage of external traffic is replaced with a function of k. This is one way to characterize the efficiency (in terms of external traffic) of the P2P technology. Under this new model, we can find scenarios for two ISPs adopting different pricing strategies to co-exist. One such example is for the following parameters: $\hat{N} = 1000, N = 100, C_m = 0.1, C_t = 0.1, C_f = 10, V = 19, p = 1.8, \beta = 0.86, \rho = 0.9$. According to the profit-neutral assumption, $q = 0.18$, and it follows that the resulting profits are: $P_o = 10.65, P_1 = 12.4, P_2 = 11.8$, which satisfies the co-existence condition $P_1 > P_2 > P_o$. In this example, after ISP_1 adopts uplink pricing, ISP_2 would prefer to stay with its flat-rate pricing because otherwise its profit would be lowered; meanwhile ISP_1 has no incentive to revert, as uplink pricing has increased its profit. In conclusion, there will be different ISP pricing schemes in this market, serving different types of users. In the above example, the fraction of P2P users, β, needed to produce the co-existence scenario is rather high; but with improved peer-selection strategy hence improved economies of scale of serving P2P users, the required fraction of P2P users for co-existence will be lower.

In Eq. (15), we note that when $\hat{N}$ is much greater than k, the fraction of external traffic approaches 1, meaning that nearly all the traffic generated by P2P peers are going outside. Then this new model degenerates to a special case of the original model (with outbound external traffic linear in k), with $\alpha = 1$.

C. The Effect of Complete Adoption of Uplink Pricing

In this paper, we made two assumptions regarding the complete adoption of uplink pricing and its effect on P2P users. First, we assumed the ISP would take a profit-neutral stance towards pricing change. This is but one view of the situation. How ISP should provide network services and price them is partly a public policy issue, and has been under considerable debate [2], [3]. From a business growth viewpoint, ISPs must price their services at a rate conmensurate to user perceived utility. From this perspective, the profit-neutral assumption is probably a safe bet. Under this assumption, if there is economies of scale to the P2P external traffic, then P2P users will be rewarded with all the efficiency savings from the economies of scale.

Secondly, we tried to model the user reaction to uplink pricing by assuming a percentage of elastic users, and assuming the elastic users will restrain their P2P traffic to maintain the same subscription payment as flat-rate pricing. Obviously, there is much room for alternative views about this treatment of P2P traffic elasticity. The best approach would be to conduct some market research in this area. Note, the economies of scale of P2P external traffic also affect this assumption. Under profit-neutrality, the more economies of scale, the less elastic the users would be. A more realistic treatment would be to independently model these factors, which is an item for future studies.

D. Accounting Cost

One of the concerns with any usage-based pricing is the accounting cost associated with implementing such schemes. While a full-fledged user traffic accounting system may indeed be a significant undertaking, the required accounting effort for uplink pricing can be at a very coarse level. For example, in our analysis, we actually characterized users into two types only. One possible implementation is to put any user exceeding a certain threshold of uplink traffic volume into one class, and treat the rest as regular users. Such coarse classification is no more difficult than most traffic rate limiting mechanisms deployed by some ISPs for managing P2P traffic.

E. Capacity Growth

In this study, we have generally ignored the need for further capacity investments by ISPs (ISPs can take all users in the market). P2P applications are still in a growth phase; more broadband users may become P2P users in the future. These factors can also affect our models, and can be an interesting topic for further studies.

VI. RELATED WORKS

The extent of P2P traffic in ISP networks has been studied and mentioned in many papers [4],[5]. A number of papers attempted to model the tussle between ISP and P2P users; for example, [6] reveals some of the ISP peering and routing issues assuming P2P as the dominant traffic in the network. [7] provides an understanding of the tussle between ISP and P2P with a system performance model. Several papers discussed approaches to management P2P traffic for ISPs. These works all help to motivate our problem.

In this paper, we study a form of pricing that can be used to manage P2P traffic. There are many papers on Internet pricing, for example, flat-rate pricing [8], usage-based pricing [9], ”Smart” Market pricing [10]. Our work is a new look at applying pricing to manage different types of users in a network.

VII. CONCLUSION

When you talk to people who work for ISPs, large or small, you quickly realize that all ISPs are grappling with the P2P problem, which ironically is also an opportunity to ISPs as P2P can be viewed as a renewed content provider technology. ISPs have tried different methods, such as P2P traffic blocking, and rate limiting, which sometimes have led to negative sentiments by their customers. In this paper, we revisit pricing as a possible mechanism to manage ISP networks in the P2P era.

In proposing uplink pricing, we are advocating the uplink and downlink of subscribe connection should be treated separately, to some extent. A user’s uplink, when used for a high volume of traffic, is serving the same purpose as the uplink of an Internet Content Provider’s uplink. Therefore, similar charging plans as that for an ICP’s link should be applied to P2P users as well.

In this preliminary study, we focus on the issue of whether and under what market and user behavior conditions uplink pricing will be adopted. We use a game-theoretic analysis to suggest three possible outcomes: both ISPs adopting uplink pricing; both ISPs keeping flat-rate pricing, and uplink pricing co-existing with flat-rate pricing. We completely characterize all the market conditions for each outcome. In Section V, we discuss various assumptions and extensions, and how the framework can be applied to further studies.

REFERENCES

[1] A Odlyzko, ”Should flat-rate Internet pricing continue”, in *IT Professional, 2000*

[2] S Shakkottai, R Srikant, ”Economics of network pricing with multiple ISPs”, in *IEEE/ACM Transactions on Networking (TON), 2006*

[3] SCM Lee, JWJ Jiang, DM Chiu, JCS Lui, ”Interaction of ISPs: Distributed Resource Allocation and Revenue Maximization ”, in *IEEE Transactions on Parallel and Distributed Systems, 2007*

[4] http://www.cachelogic.com/

[5] L Plissonneau, JL Costeux, P Brown, ”Analysis of Peer-to-Peer Traffic on ADSL”, in *Passive And Active Network Measurement: 6th International Workshop, PAM 2005.*

[6] JH Wang, DM Chiu, JCS Lui, ”Modeling the peering and routing tussle between isps and p2p applications”, in *Proc. of IEEE International Workshop on Quality of Service (IWQoS)*, (Yale University, New Haven, CT), pp. 51C59, June 2006.

[7] M Garetto, DR Figueiredo, R Gaeta, M Sereno,”A modeling framework to understand the tussle between ISPs and peer-to-peer file-sharing users”, in *Performance Evaluation, 2007*

[8] M Falkner, M Devetsikiotis, ”An Overview of Pricing Concepts for Broadband IP Networks”, in *IEEE Communications Surveys, 2000*

[9] Altmann, J., and K. Chu, ”How to Charge for Network Services - Flat-Rate or Usage-Based,” in *Computer Networks 36 (2001), pp.519-531*

[10] MacKie-Mason, J. K. and Varian, H. R. 1993. ”Pricing the Internet.” *Public Access to the Internet, Kahin, Brian and Keller, James.*