Pricing Internet Services With Multiple Providers*

Linhai He and Jean Walrand
Dept. of Electrical Engineering and Computer Science
University of California at Berkeley
Berkeley, CA 94709
linhai, wlr@eecs.berkeley.edu

Abstract

One of the challenges facing the networking industry today is to increase the profitability of service providers. This calls for economic mechanisms that can enable providers to charge more for better services and collect a fair share of the resulting increased revenue. In this paper we present a generic model for pricing Internet services in a multi-provider network. We show that non-cooperative pricing is unfair and may discourage future upgrades of the network. As an alternative, we propose a simple revenue-sharing policy and show that it is more efficient and encourages providers to collaborate without cheating. We also suggest a scalable algorithm for providers to implement this policy in a distributed way and study its convergence properties.

1 Introduction

For historical reasons, the current architecture of the Internet lacks the support for implementing efficient market mechanisms. Consequently, service providers have limited economic incentives to invest in technology for new services. This situation limits the future evolution of the Internet. To correct this state of affairs, it is essential to implement economic mechanisms that would enable service providers to charge more for better services and collect a fair share of the resulting increased revenue. In this paper we investigate how to design pricing schemes that could meet these criteria.

The idea of using economic mechanisms in network design is not new. For example, [1] [2] [3] propose pricing mechanisms that can be used for congestion control in the Internet. However, in these schemes the network acts as a social-welfare maximizer with no self-interest. This assumption does not reflect the situation in today's Internet, as most network service providers are in the business for making profit and are primarily interested in maximizing their own benefits [4]. Our pricing schemes try to include these facts into the models. We believe that a good pricing scheme should provide the right incentives for providers to follow the protocol and not to cheat. In addition, it should be fair for all providers involved and encourage upgrades to the network. In other words, a provider should be able to collect more revenue by increasing the capacity of its network. Finally, a pricing scheme should be scalable, i.e., feasible for large-scale deployment.

*This research is supported in part by DARPA Grant No. BAA00-18.

Our paper is organized as follows. In Section 2, we describe the basic models for the providers and the services that they offer. In the following two sections, we first study the case in which providers adopt non-cooperative pricing strategies. Through simple examples, we show that such strategies would result in undesirable equilibria. We then suggest a revenue sharing policy as an alternative and show that it would lead to a better equilibrium. In addition, it could be reached through a distributed algorithm. We conclude the paper with discussions on future work.

2 Basic Model

We consider a group of providers offering services with a certain level of QoS guarantee. For simplicity, we assume that those QoS requirements could be translated into *local* capacity constraints. For instance, the maximum utilization on a link may be limited to be less than, say, 25% to ensure all packets experience only small delay going through that link.

We assume that there exists a set of routes across the network. On any of these routes each provider charges a price for its share of the service. The providers may adjust their prices dynamically and signal them to end users to control the demand for the services. There are many possible approaches for implementing such a pricing scheme, for applications with either fixed or elastic bandwidth requirements. However, for the purpose of modelling, we do not specify details of implementation in this paper. We simply model that when a price p is posted for a route r, the resulting traffic load on that route is given by a function $d_r(p)$, which is strictly decreasing and differentiable. Moreover, mechanisms exist for providers to collect revenues based on the amount of traffic that they have forwarded and the prices that they set.

We assume that when a provider sets its price, its objective is to maximize its own revenues, while maintaining the QoS for the service that it offers by respecting its local capacity constraint. Therefore, in the case of only one provider offering the service, the optimal price can be determined by solving the following constrained optimization program

$$\begin{array}{rl} \max_{p\geq 0} & J = p \cdot d(p) \\ \text{s.t.} & d(p) \leq C \end{array} \tag{1}$$

where C is the capacity constraint. The first-order condition for the solution is $p^* = \mu - d(p*)/d'(p*)$ where $\mu \geq 0$ is some constant that satisfies $\mu(d(p^*) - C) = 0$. It is easy to show that a unique solution exists if $d(p)/d'(p)$ is an increasing function of p. In that case, the solution is also a maximizer. So in the rest of the paper, we consider only demand functions that satisfy this property. For later use, we define $g(p) \triangleq -d(p)/d'(p)$. Notice that $g(p)$ indicates the elasticity of the demand function.

To simplify analysis, we assume that all providers have sufficient capacity on their internal links. Capacity may be limited only on the links between providers. Local QoS requirements by each provider are fixed and not affected by prices. All routes between sources and destinations are also fixed. We choose this assumption because in today's Internet routing between providers is often performed based on a set of provisioned policies instead of short-term costs or performance measures.

3 Non-Cooperative Pricing Strategies

In this section we try to understand how providers would set their prices when they have to work together to offer a service. We assume that each provider acts in its own interest. In addition, each provider keeps its own capacity constraint as private information, but it may be possible for each provider to observe prices marked by others (depending on the implementation).

All these assumptions suggest a game-theoretic formulation of the problem in which each provider is a strategic player. Under different assumptions on what strategic information is available to the providers, different types of formulation, such as Nash, Stackelberg, etc., are possible. However, we argue that only a Nash game models closely how providers would interact in real situations. In a large-scale network with complex topology such as the Internet, not much information about the game is available. Providers generally do not know much about the global state of the network. A provider's best strategy probably is to optimize locally based on its observations on how its payoff changes as it changes its prices. This model fits naturally the "best-response" model of a Nash game. Therefore, we model the game between providers as follows:

$$\begin{aligned} \max_{p_{lr}\geq 0} \quad & J_i = \sum_{l\in E_i}\sum_{r\in R_l} p_{lr}\, d_r\left(\sum_{k\in L_r\setminus l} p_{kr} + p_{lr}\right) \\ \text{s.t.} \quad & \sum_{r\in R_l} d_r\left(\sum_{k\in L_r\setminus l} p_{kr} + p_{lr}\right) \leq C_l, \forall l \in E_i, \end{aligned} \tag{2}$$

where E_i is the set of egress links owned by provider i, L_r is the set of links that route r goes through, R_l is the set of routes going through link l, p_{lr} is the price charged for route r on link l, and C_l is the capacity constraint on link l.

It is known that the equilibria in Nash games are often inefficient and may have undesirable properties, due to the non-cooperative nature of the games. The game described in (2) is not an exception. In the following we use a simple example to show that non-cooperative pricing could lead to unfair distribution of revenues among providers, and that a bottleneck provider may not have an incentive to update its capacity.

Consider two providers connected in series with only one route going through them. The demand on that route is $d(p_1 + p_2)$, where p_i is the price charged by provider i for $i = 1, 2$. Without loss of generality, we assume $C_1 > C_2$, so that provider 2 is always the bottleneck. The resulting Nash game played by the two providers is as follows:

$$\begin{aligned} \text{Provider 1:} \quad & \max_{p_1\geq 0} \quad p_1\, d(p_1 + p_2) \\ \text{Provider 2:} \quad & \max_{p_2\geq 0} \quad p_2\, d(p_1 + p_2) \\ & \qquad \text{s.t.} \quad d(p_1 + p_2) \leq C_2 \end{aligned} \tag{3}$$

It is easy to show that this game has a unique Nash equilibrium. We are interested in the case in which the capacity constraint of provider 2 is active at equilibrium. We first show that in that case provider 2 always charges a higher price than provider 1, thus obtaining more revenues.

The capacity constraint C_2 is active if $C_2 \leq X^*$, where X^* is the traffic load at Nash equilibrium when there is no capacity constraint for either provider. By applying a symmetry argument to (3) with the capacity constraint removed, X^* can be found from the following equations:

$$\begin{cases} X^* = d(2p^*) \\ p^* = g(2p^*) \end{cases}, \text{ or equivalently, } \begin{cases} p^* = d^{-1}(X^*)/2 \triangleq K^*/2 \\ p^* = g(2p^*) = g(K^*) \end{cases}.$$

Here K^* is defined as the total price charged to the users at equilibrium. The above equations implies that $K^* = 2g(K^*)$.

When $C_2 < X^*$, define $K \triangleq d^{-1}(C_2)$. From the optimality condition for provider 1, $p_1 = g(p_1 + p_2)$, we get $p_1 = g(K)$. Since the constraint is active, $d(p_1 + p_2) = C_2$, or $p_1 + p_2 = K$. So $p_2 = K - p_1 = K - g(K)$. Since $d(\cdot)$ is a decreasing function, K is a decreasing function of C_2. Therefore, when $C_2 < X^*$, $K > K^*$. Moreover, since $g(\cdot)$ is a decreasing function, $2g(K) < 2g(K^*) = K^* < K$. As a result, $p_2 = K - g(K) > g(K) = p_1$. What this means is that when potential demand exceeds network capacity, bottleneck provider always charges higher price, thus obtain a larger share of the total revenue, than the unconstrained one. This is certainly very unfair. In addition, note that the ratio between the prices is $p_2/p_1 = K/g(K) - 1$. So the smaller C_2 is, the larger K is, and the higher the ratio is. For fixed C_2, the more elastic the demand is, the faster $g(K)$ decays with K, and the higher the ratio is.

Next we show that provider 2 may not have incentive to upgrade its link. In other words, increasing C_2 may not always increase provider 2's revenues! To estimate how the equilibrium changes with C_2, we are interested in the properties of $\partial J_2^*/\partial C_2$ (We designate the equilibrium value of all variables by the superscript *). If the solution to the equation

$$\frac{\partial J_2^*}{\partial C_2} = \frac{\partial\left(p_2^* C_2\right)}{\partial C_2} = \frac{\partial p_2^*}{\partial C_2} C_2 + p_2^* = 0 \tag{4}$$

exists, then it determines the capacity C_2 that yields the largest revenues for provider 2. By differentiating both sides of $d(p_1^* + p_2^*) = C_2$ w.r.t. C_2, we get

$$d'(p_1^* + p_2^*)\left(\frac{\partial p_1^*}{\partial p_2^*} + 1\right)\frac{\partial p_2^*}{\partial C_2} = 1. \tag{5}$$

Similarly, by differentiating both sides of $g(p_1^* + p_2^*) = p_1^*$ w.r.t p_2^*, we get

$$\frac{\partial p_1^*}{\partial p_2^*} = g'(p_1^* + p_2^*)\left(\frac{\partial p_1^*}{\partial p_2^*} + 1\right) \implies \frac{\partial p_1^*}{\partial p_2^*} + 1 = \frac{1}{1 - g'(p_1^* + p_2^*)} > 0,$$

by the fact that $g(p)$ is a decreasing function of p. Since the demand function is also a decreasing function of price, the first term in (5) is negative. This implies that $\partial p_2^*/\partial C_2$ must be negative as well. This result suggests that there might exist a solution for (4) and hence a possible maximum of J_2^*. If such a maximum does exist, then the bottleneck provider may stop upgrading its link after that capacity, even before the demand is fully met.

We find that it is not hard to find a demand function with which this maximum exists. For instance, it can be shown that $d(p) = A\exp(-Bp^\alpha), \alpha > 1$, is one class of such functions. Figure 1 shows J_2^* as a function of C_2, with the demand function chosen to be $d(p) = 10\exp(-p^2)$. It can be clearly seen that a maximum is achieved before C_2 moves into the unconstrained region.

4 Revenue-Sharing Policy

Given the undesirable properties of non-cooperative pricing strategy, it is natural to ask if a pricing scheme could be designed to overcome those drawbacks and yet be compatible with the providers' interest by giving them no incentive to cheat. Game theory itself provides many useful concepts and tools for such a design problem, such as mechanism

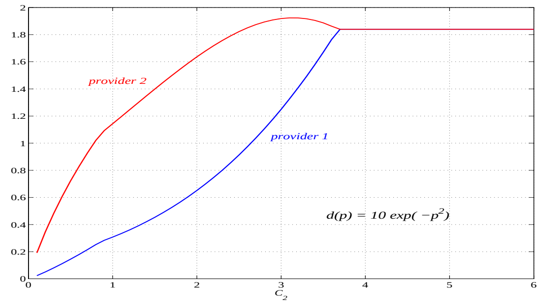

Figure 1: Revenues collected by two providers at the equilibrium, when demand function is $d(p) = 10\exp(-p^2)$.

design, cooperative game theory, and so on [5] [6]. However, we have found that these concepts are either too complex to compute and very difficult (if not impossible) to implement in a scalable way. As an alternative, we propose a straightforward revenue sharing policy that has the aforementioned properties. We first study how providers would behave under this policy, and then suggest a scalable algorithm for providers to reach that equilibrium.

4.1 Equilibrium

In our revenue sharing policy, all providers agree to evenly split the total revenue collected on a route. Nevertheless, they are still allowed to choose prices based on their best interest. Mathematically, a provider finds its optimal prices by solving the following program:

$$\begin{aligned} \max_{p_{lr}\geq 0} \quad & J_i = \textstyle\sum_{l\in E_i}\sum_{r\in R_l}\frac{1}{N_r}\left(\sum_{k\in L_r\setminus l} p_{kr} + p_{lr}\right) d_r\left(\sum_{k\in L_r\setminus l} p_{kr} + p_{lr}\right) \\ \text{s.t.} \quad & \textstyle\sum_{r\in R_l} d_r\left(\sum_{k\in L_r\setminus l} p_{kr} + p_{lr}\right) \leq C_l, \forall l \in E_i, \end{aligned} \tag{6}$$

where N_r is the number of providers on route r. As we will see next, this change in the objective function completely alters providers' strategy. First we show the existence of an equilibrium.

Theorem 4.1. *A Nash equilibrium exists for the game described in (6).*

Proof. (To save space we only give an outline of the proof here.)

We first show that for any given strategy profile of other providers, a unique maximizer for (6) exists. Note that because there is no capacity constraint on internal links, routes existing through different egress links of a provider do not affect each other at all. So a provider can optimize over each egress link independently. By applying the first-order optimality condition, we get

$$p^*_{lr} = \max\{0, N_r\mu_l + g_r(\sum_{k\in L_r\setminus l} p_{kr} + p^*_{lr}) - \sum_{k\in L_r\setminus l} p_{kr}\}, \tag{7}$$

where $\mu_l \geq 0$ is the Lagrangian multiplier for link l. Define $t_{lr} \triangleq N_r\mu_l + g_r(t_{lr})$. Intuitively, t_{lr} is the locally optimal *total price* for route r, preferred by provider i on link l. Since $g(\cdot)$ is continuous and decreasing, t_{lr}, and thus p^*_{lr}, are non-decreasing, continuous functions of μ_l. Then by Intermediate Value theorem, we can conclude that either there exists a unique $\mu^*_l > 0$ which satisfies $\sum_{r\in R_l} d_r(\sum_{k\in L_r\setminus l} p_{kr} + p^*_{lr}) = C_l$, or $\mu^*_l = 0$. By duality theory [7], this set of (p^*_{lr}, μ^*_l) is the optimal solution to (6). This result also suggests that we only need to look at the dual variables when solving the equilibrium of the game. Following this idea, we first show the following Lemma.

Lemma 4.2. *On any route r, for any given set of $\mu_l \geq 0$, the solution to the following system of equations: $p_{lr} = \max\{0, t_{lr} - \sum_{k \in L_r \setminus l} p_{kr}\}, \forall l \in L_r$, is*

$$p_{lr} = \begin{cases} t_{\bar{l}r}, & \text{if } t_{\bar{l}r} = \max\{t_{lr} : l \in L_r\}; \\ 0, & \text{otherwise.} \end{cases}$$

In other words, only the link with the largest t_{lr} sets non-zero price.

Proof. (We only give a brief outline of the proof here. Also the subscript r is omitted for clarity.) Classify links into different sets according to the value of their t_l's. Denote those sets by $A_j, j = 1, \cdots, J$. Define the value of these sets, S_j, by the corresponding t_l of its members. It can be shown that for links belong to the same set A_j, either $\sum_{l_j \in A_j} p_{l_j}$ equals some positive number, or $p_{l_j} = 0, \forall l_j \in A_j$. Based on this result, define $y_j \triangleq \sum_{l \in A_j} p_l$, and consider the following system of equations: $y_j = \max\{0, S_j - \sum_{k \neq j} y_k\}, j = 1, \cdots, J$. By constructing a contradiction, one can show that y_1 must be zero. The same procedure is repeated until $j = J - 1$ to get $y_{J-1} = 0$ and $y_J = S_J$.

If there are more than one links in A_J, there are a set of prices that they could choose, as long as the sum of the prices equals S_J. To avoid possible ambiguity, we fix a rule that only the most upstream link in A_J sets its price to S_J, while the rest of the links all set zero price. □

Remark. Since $t_{lr} = N_r \mu_l + g_r(t_{lr})$, if $\mu_m > \mu_n$, then $t_{mr} > t_{nr}$. Roughly speaking, Lagrangian multiplier μ at the optimum indicates how congested a link is. So this lemma implies that on any route only the most congested link sets its total price. Moreover, since μ fully determines the optimal price, we may view μ as the actual strategy played in the game. We will use this argument to prove the existence of the equilibrium.

Consider the mapping $f_l : \underline{\mu} \triangleq \{\mu_k, \forall k\} \to \mu_l$, the best-response of link l to $\underline{\mu}$ on other links. It is easy to see that $f_l(\underline{\mu})$ is bounded in $[0, f_l(\underline{0})], \forall l$. So to show the existence of the equilibrium, we only need to show that the mapping f_l is continuous and then apply Brouwer's fixed-point theorem.

Define $\bar{\mu}_r \triangleq \max\{\mu_l : l \in L_r\}$, and the corresponding total price for that route by $\bar{p}_r \triangleq h_r(\bar{\mu}_r)$, where h_r is the implicit function defined through (7). So now the demand on route r can be expressed in terms of $\bar{\mu}_r$, i.e. $d_r(\bar{p}_r) = d_r(h_r(\bar{\mu}_r)) = d_r \circ h_r(\max\{\mu_l : l \in L_r\})$. For use later, define $\bar{x}_r(\bar{\mu}_r) \triangleq d_r \circ h_r(\bar{\mu}_r)$. Clearly, $\bar{x}_r$ is strictly decreasing and continuous.

Consider $\underline{\mu}^{(i)}, i = 1, 2$, and suppose that $\| \underline{\mu}^{(1)} - \underline{\mu}^{(2)} \| < \epsilon$ in some norm. Without loss of generality, we may assume that $| \mu_l^{(1)} - \mu_l^{(2)} | < \epsilon, \forall l$. It can be shown that $| \bar{\mu}_r^{(1)} - \bar{\mu}_r^{(2)} | < \epsilon, \forall r$, as well. Denote $\tilde{\mu}_l^{(i)}$ as the best-response of link l under $\underline{\mu}^{(i)}$. Notice that $\bar{\mu}_r^{(1)}$ and $\bar{\mu}_r^{(2)}$ may be associated with different links, so we have to study $| \tilde{\mu}_l^{(1)} - \tilde{\mu}_l^{(2)} |$ under different possible scenarios. Before we proceed, we introduce a small lemma (we omit the proof).

Lemma 4.3. *Suppose $f(x)$ is a strictly decreasing, continuous function over a finite interval. Then $\forall \epsilon > 0$, if $|f(x_1) - f(x_2)| < \epsilon$, then $\exists \delta_\epsilon > 0$ s.t. $|x_1 - x_2| < \delta_\epsilon$.*

Case 1. The constraint on link l is active under both $\underline{\mu}^{(1)}$ and $\underline{\mu}^{(2)}$. Define $R_{l,1} \triangleq \{r \in R_l \mid \bar{\mu}_r = \tilde{\mu}_l\}$, and $R_{l,2} \triangleq R \setminus R_{l,1}$. In other words, $R_{l,1}$ is the set of routes whose total price is set by link l.

Case 1.1 Suppose $R_{l,1}$ and $R_{l,2}$ do not change under $\underline{\mu}^{(1)}$ and $\underline{\mu}^{(2)}$. It can be shown that $\forall \epsilon > 0$, if $\mid \bar{\mu}_r^{(1)} - \bar{\mu}_r^{(2)} \mid < \epsilon$, $\exists \delta_\epsilon > 0$, s.t. $\mid \sum_{r \in R_{l,1}} \bar{x}_r(\tilde{\mu}_l^{(1)}) - \sum_{r \in R_{l,1}} \bar{x}_r(\tilde{\mu}_l^{(2)}) \mid < \delta_\epsilon$. Then by Lemma 4.3, $\exists \gamma_\epsilon > 0$, s.t. $|\tilde{\mu}_l^{(1)} - \tilde{\mu}_l^{(2)}| < \gamma_\epsilon$.

Case 1.2. Suppose $R_{l,1}$ and $R_{l,2}$ change with $\underline{\mu}^{(1)}$ and $\underline{\mu}^{(2)}$. Define $R_{l,i}^{\triangle}$ as the set of routes that belong to $R_{l,i}^{(1)}$ under $\underline{\mu}^{(1)}$ but switch to $R_{l,j}^{(2)}$ under $\underline{\mu}^{(2)}$, for $i, j = 1, 2$, and $i \neq j$. There are two possibilities:
Case 1.2.a. Suppose neither $R_{l,1}^{\triangle}$ nor $R_{l,2}^{\triangle}$ is empty. One can show that $\mid \tilde{\mu}_l^{(1)} - \tilde{\mu}_l^{(2)} \mid < \max\{\bar{\mu}_2^{(1)}, \bar{\mu}_1^{(2)}\} - \max\{\bar{\mu}_1^{(1)}, \bar{\mu}_2^{(2)}\} \leq 2\epsilon$.
Case 1.2.b. Suppose $R_{l,1}^{\triangle} \neq \phi$ but $R_{l,2}^{\triangle} = \phi$. If $\tilde{\mu}_l^{(1)} \leq \tilde{\mu}_l^{(2)}$, then we may choose any $r_1 \in R_{l,1}^{\triangle}$, and by the fact that

$$\bar{\mu}_{r_1}^{(1)} \triangleq \max\{\mu_k^{(1)} : k \in L_{r_1} \setminus l\} < \tilde{\mu}_l^{(1)} \leq \tilde{\mu}_l^{(2)} < \bar{\mu}_{r_1}^{(2)} \triangleq \max\{\mu_k^{(2)} : k \in L_{r_1} \setminus l\}$$

to conclude $\mid \tilde{\mu}_l^{(1)} - \tilde{\mu}_l^{(2)} \mid < \epsilon$. For the case that $\tilde{\mu}_l^{(1)} > \tilde{\mu}_l^{(2)}$, it can be shown that $\exists \delta_\epsilon > 0$,

$$0 < \sum_{r \in R_{l,1}^{(2)} \setminus R_{l,1}^{\triangle}} \bar{x}_r(\tilde{\mu}_l^{(2)}) - \sum_{r \in R_{l,1}^{(1)} \setminus R_{l,1}^{\triangle}} \bar{x}_r(\tilde{\mu}_l^{(1)}) < \delta_\epsilon.$$

Then by Lemma 4.3, $\exists \gamma_\epsilon > 0$, s.t. $\left|\tilde{\mu}_l^{(2)} - \tilde{\mu}_l^{(1)}\right| < \gamma_\epsilon$.

The opposite case, i.e. $R_{l,1}^{\triangle} = \phi$ but $R_{l,2}^{\triangle} \neq \phi$, can be proved using the same idea.

Case 2. $\tilde{\mu}_l^{(1)} = 0$, but $\tilde{\mu}_l^{(2)} > 0$. Define $x_r(\mu_l : l \in L_r) \triangleq d_r \circ h_r(\max\{\mu_l : l \in L_r\})$. By continuity of x_r, $\exists \delta_\epsilon > 0$, s.t.

$$\begin{aligned} \delta_\epsilon &> \textstyle\sum_{r \in R_l} x_r(0, \mu_{-l,r}^{(2)}) - \sum_{r \in R_l} x_r(0, \mu_{-l,r}^{(1)}) > \sum_{r \in R_l} x_r(0, \mu_{-l,r}^{(2)}) - C_l \\ &= \textstyle\sum_{r \in R_l} x_r(0, \mu_{-l,r}^{(2)}) - \sum_{r \in R_l} x_r(\tilde{\mu}_l^{(2)}, \mu_{-l,r}^{(2)}), \end{aligned}$$

because $\sum_{r \in R_l} x_r(0, \mu_{-l,r}^{(1)}) < C_l < \sum_{r \in R_l} x_r(0, \mu_{-l,r}^{(2)})$. By Lemma (4.3), $\exists \gamma_\epsilon > 0$ s.t. $\mid 0 - \tilde{\mu}_l^{(2)} \mid = \mid \tilde{\mu}_l^{(1)} - \tilde{\mu}_l^{(2)} \mid < \gamma_\epsilon$.

The opposite case, i.e. $\tilde{\mu}_l^{(1)} > 0$, but $\tilde{\mu}_l^{(2)} = 0$, can be proved in the same way. □

Remark. If the decision on the distribution of revenues is made by a central agent, the optimal total price for route r is $p_r = \sum_{l \in L_r} \mu_l + g_r(p_r)$, where μ_l's are the corresponding Lagrangian multipliers which satisfy capacity constraints on *all* links. In our sharing policy, the total price at equilibrium for the same route is $p_r = N_r \max\{\mu_l : l \in L_r\} + g_r(p_r)$. Therefore, one may consider this as a tradeoff between system efficiency and fairness for individual providers.

This proof can also be used to show that providers under revenue-sharing policy always have incentive to upgrade their links. Consider any provider with a constrained link. Denote $R_{l,1}$ as the set of routes whose total prices are set by link l and $R_{l,2} = R_l \setminus R_{l,1} = \{r \in R_l \mid \bar{\mu}_r \geq \mu_l\}$. Because the total price for a route is determined by the maximum of μ_k, over all $k \in L_r$, routes in $R_{l,2}$ are not affected by increase in C_l at all, because $\partial \mu_l / \partial C_l < 0$. Moreover, on any route $r \in R_{l,1}$, since $\mu_l > \mu_k$ for any $k \in L_r \setminus l$, any infinitesimal increase in C_l will not change $R_{l,1}$. Therefore, $\partial J_l^* / \partial C_l = \mu_l > 0$.

Since providers always have incentive to upgrade their links under sharing policy, eventually the network will move into the capacity region in which none of the links is constrained. For that case, we show next that revenue collected by each provider under sharing policy strictly dominates that with non-cooperative pricing.

Because there is no capacity constraint, we only need to prove the result for the case of a single route. Consider a route r connected by N providers and has associated demand d.

- Under revenue-sharing policy, by Lemma 4.2 in the proof, providers would agree on a single total price, $p_s = \arg\max\{p\, d(p)\}$, and then equally split the total revenue $p_s\, d(p_s)$. So in this case revenue-sharing policy is equivalent to centralized allocation, which results in a revenue of $J_s = p_s\, d(p_s)/N$ for each provider. Note that p_s is solved from the optimality condition $p_s = g(p_s)$.

- Under non-cooperative pricing, suppose that the price set by individual provider i is $p_{n,i}$. By symmetry, $p_{n,i}$ must be equal for all providers at the equilibrium. So the revenue collected by each provider is $J_n = p_n\, d(Np_n) = Np_n\, d(Np_n)/N$, where p_n is the solution to the local optimality condition $p_n = g(Np_n)$.

Because $g(\cdot)$ is a decreasing function, it is straightforward to show that $Np_n > p_s$ for any $N > 1$. Since p_s is the only maximizer to the unconstrained optimization program $\max_{p \geq 0} p\, d(p)$, we can conclude that $p_s\, d(p_s) > Np_n\, d(Np_n)$, i.e. $J_s > J_n$.

4.2 Implementation

The proof of the existence of equilibrium also suggests an algorithm for providers to compute it. Lemma 4.2 suggests that the optimal total price for a route is determined by the link with the largest Lagrangian multiplier. Because the duality gap for the local optimization program is zero, these Lagrangian multipliers can be computed iteratively based on the local traffic load. Theorem 4.1 then guarantees that an equilibrium exists even if each provider performs updates using local information only. Based on these results, we suggest the following algorithm:

- Each provider maintains a state variable μ_l for each link l, which is updated periodically according to the following rule: $\mu_l := \max\{0, \mu_l + \omega_l(X_l - C_l)\}$, where $\omega_l > 0$ is a small constant and X_l is the total traffic load on link l.

- Each packet has two dedicated fields in its header, denoted by $\bar{\mu}$ and N. Both fields are initialized to zero when a packet enters the network. As a packet passes through a link on its route to destination, the router on that link increments N by one and updates $\bar{\mu}$ by the rule: $\bar{\mu} := \max\{\bar{\mu}, \mu_l\}$, i.e. the router updates $\bar{\mu}$ only if its own μ_l is larger.

- After the packet reaches its destination, the values recorded in μ_l and N are returned to the sending host via either ACK or some special control packets.

- We assume that a provider is able to keep some estimates of demand on each route that initiates from its network. When it receives an ACK or control packet returned from a destination, it updates the price for the corresponding route with the solution from $p = N\bar{\mu} + g(p)$.

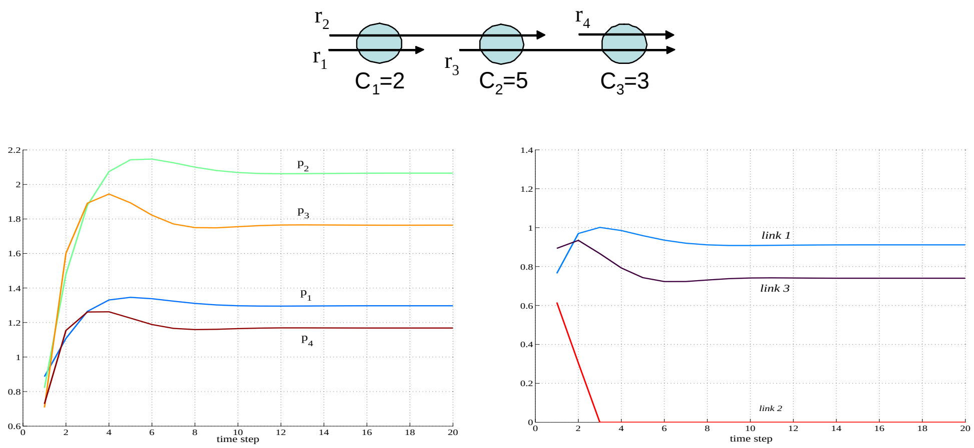

Figure 2: Adaptation of p and μ_l over time in numerical simulation.

In this implementation, only first-hop providers need to keep states for each of their routes, and possibly the on-going price charged for each flow. This is feasible because at the edge of the network the numbers of active flows and routes are relatively small, and the providers have to maintain those information for charging purpose anyway. Transit providers do not need to keep any per-flow nor per-route state, and they do not even need to estimate any demand function in order to maximize its revenue. So this algorithm is quite scalable for implementation. Next we show the following convergence result.

Theorem 4.4. *The algorithm described above converges to the equilibrium described in Theorem 4.1.*

Proof. We take a continuous-time approximation to the adaptation process of μ_l, i.e. $d\mu_l/dt = \max\{0, \omega_l(X_l - C_l)\}$. Consider the following function

$$V(\underline{\mu}) = \sum_l \mu_l(C_l - \sum_{r\in R_{l,2}} \bar{x}_r(\bar{\mu}_r)) - \sum_r \int_0^{\bar{\mu}_r} \bar{x}_r(t)dt.$$

It is easy to show that V is continuous and achieves its minimum at the equilibrium. Then we have

$$\begin{aligned} \frac{\partial V}{\partial \mu_l} &= C_l - \textstyle\sum_{r\in R_{l,2}} \bar{x}_r(\bar{\mu}_r) - \sum_r \bar{x}_r(\bar{\mu}_r) I\{\mu_l = \bar{\mu}_r\} \\ &= C_l - \textstyle\sum_{r\in R_{l,2}} \bar{x}_r(\bar{\mu}_r) - \sum_{r\in R_{l,1}} \bar{x}_r(\mu_l) = C_l - X_l \\ \Rightarrow \frac{\partial V}{\partial t} &= \textstyle\sum_l \frac{\partial V}{\partial \mu_l}\frac{d\mu_l}{dt} = \sum_l \max\{0, \omega_l(X_l - C_l)\}(C_l - X_l) \le 0 \end{aligned}$$

i.e. V is a Lyapunov function for the algorithm. □

Figure 2 shows the results of a numerical simulation of the above algorithm. In the plots p and μ_l adapt over time and converge to their equilibrium values.

5 Conclusion

In this paper we have presented a generic model for pricing Internet services in a multi-provider network. We have showed that non-cooperative pricing is unfair and may discourage future upgrades of the networks. As an alternative, we have proposed a simple revenue sharing policy and have shown that it is fair, more efficient, and encourages providers to collaborate without cheating. We also have suggested a scalable algorithm for providers to implement this policy in a distributed way and studied its convergence property.

The model that we have presented in this paper is probably more suited for users who prefer predictable service quality and are not very sensitive to fluctuations in price. Certainly there are users whose preference is the other way around. So it would be useful to extend our pricing schemes to support both service types in a flexible way. Also it would be interesting to understand how optimal pricing strategies would change accordingly to reflect different natures of these two services.

References

[1] F. P. Kelly, A. K. Maulloo, and D. H. K. Tan, "The Rate Control for Communication Networks: Shadow Prices, Proportional Fairness and Stability," *Journal of the Operational Research Society*, pp. 237-252, vol.49, 1998.

[2] S. Kunniyur and R. Srikant, "Analysis and Design of an Adaptive Virtual Queue Algorithm for Active Queue Management," *Proc. ACM Sigcomm*, 2001.

[3] S. H. Low and D. E. Lapsley, "Optimization Flow Control, I: Basic Algorithm and Convergence,", *IEEE/ACM Transactions on Networking*, 7(6):861-75, Dec. 1999.

[4] G. Huston, "Interconnection, peering, and settlements," *Proc. INET*, June 1999.

[5] D. Fudenberg and J. Tirole, "Game Theory," *MIT Press*, Cambridge, Mass., 1991.

[6] G. Owen, "Game Theory," *Academic Press*, San Diego, 1995.

[7] D. G. Luenberger, "Linear and Nonlinear Programming," *Addison-Wesley*, Reading, Mass., 1984.