

Yale University
Department of Computer Science

**Optimal ISP Subscription for Internet Multihoming:
Algorithm Design and Implication Analysis**

Hao Wang Yale University	Haiyong Xie Yale University	Lili Qiu Microsoft Research
Avi Silberschatz Yale University	Yang Richard Yang Yale University	

YALEU/DCS/TR-1299

July 2004

(revised December 2004)

Hao Wang is supported by NSF grant ANI-0207399. Haiyong Xie is supported by NSF grant ANI-0238038. Avi Silberschatz is supported in part by NSF grants ITR-0331548 and CNS-0435201. Yang Richard Yang is supported in part by NSF grants ANI-0207399, ANI-0238038 and CNS-0435201.

Optimal ISP Subscription for Internet Multihoming: Algorithm Design and Implication Analysis

Hao Wang Haiyong Xie Lili Qiu
Yale University Yale University Microsoft Research

Avi Silberschatz Yang Richard Yang
Yale University Yale University

Abstract

Multihoming is a popular method used by large enterprises and stub ISPs to connect to the Internet to reduce cost and improve performance. Recently researchers have studied the potential benefits of multihoming and proposed protocols and algorithms to realize these benefits. They focus on how to dynamically select which ISPs to use for forwarding and receiving packets, and assume that the set of subscribed ISPs is given *a priori*. In practice, a user often has the freedom to choose which subset of ISPs among all available ISPs to subscribe to. We call the problem of how to choose the optimal set of ISPs the ISP subscription problem. In this paper, We design a dynamic programming algorithm to solve the ISP subscription problem optimally. We also design a more efficient algorithm for a large class of common pricing functions. Using real traffic traces and realistic pricing data, we show that our algorithm reduces users' cost. Next we study how ISPs respond to users' optimal ISP subscription by adjusting their pricing strategies. We call this

Hao Wang is supported by NSF grant ANI-0207399. Haiyong Xie is supported by NSF grant ANI-0238038. Avi Silberschatz is supported in part by NSF grants ITR-0331548 and CNS-0435201. Yang Richard Yang is supported in part by NSF grants ANI-0207399, ANI-0238038 and CNS-0435201.

problem the ISP pricing problem. Using a realistic charging model, we formulate the problem as a non-cooperative game. We first prove that if cost is the only criterion used by a user to determine which subset of ISPs to subscribe to, at any equilibrium all ISPs receive zero revenue. We then study a more practical formulation in which different ISPs provide different levels of reliability and users choose ISPs to both improve reliability and reduce cost. We analyze this problem and show that at any equilibrium an ISP's revenue is positive and determined by its reliability.

1 Introduction

Multihoming is a popular method used by large enterprises, stub ISPs, and even small businesses to connect to the Internet [33]. A user is said to be multihomed if it has multiple external links (either to a single provider, or to different providers). According to a study by CAIDA [9], as of June, 2004, 51% of stub ASes are multihomed. When a multihomed user actively controls how its traffic is distributed among its multiple links, we say that it implements *smart routing*. Smart routing is also referred to as route optimization, or intelligent route control.

In the past few years, there has been significant research on evaluating and realizing the benefits of multihoming. For example, in [1, 2], Akella *et al.* quantify the benefits of multihoming and show that selecting the right set of providers yields performance improvement. In [15], Goldenberg *et al.* propose smart routing algorithms to distribute traffic among multiple links to optimize both cost and performance. A recent economic analysis shows that smart routing has the potential to benefit not only the end users, but also the service providers [13]. Many companies are actively developing commercial products to realize the benefits of multihoming (*e.g.*, Internap, Profluent, Radware, RouteScience).

Although these previous studies have made much progress in realizing the potential benefits of multihoming, two important problems remain unaddressed. First, most of the previous studies focus on how to dynamically select which ISPs to use for forwarding and receiving packets, and do not consider the ISP subscription problem (*i.e.*, how to determine which ISPs among all available ISPs to subscribe to). Second, the freedom for users to choose ISPs introduces competitions among ISPs. ISPs will respond to users' selections by adjusting their pricing strategies. We call this problem the ISP pricing problem. While there is a large volume of literature on pricing and competition, most are based on abstract pricing

ing models. There is no previous study on this problem using realistic Internet pricing models.

To address the above issues, we first study the ISP subscription problem. We develop an optimal algorithm using dynamic programming to minimize a user's cost. Based on the observation that many pricing functions are concave due to diminishing marginal returns, we design a more efficient algorithm for this class of pricing functions. Using real traffic traces and realistic pricing data, we show that our algorithm reduces a user's cost by up to 24% compared with a greedy heuristic, and by up to 100% compared with random subscription.

Next we study the ISP pricing problem. Using the realistic percentile-based charging model, we formulate the problem as a non-cooperative game. We prove that if cost is the only criterion used by a user to determine which ISPs to subscribe to, all ISPs receive zero revenue at any equilibrium. We then study a more practical formulation of the ISP pricing problem in which different ISPs provide different levels of reliability and users choose ISPs to both improve reliability and reduce cost. We analyze this problem and show that an ISP's revenue is positive and determined by its reliability at any equilibrium. This result suggests that when users use multihoming to both improve reliability and reduce cost, the increasingly wide deployment of multihoming can be beneficial to the global Internet, since it provides incentives for the ISPs to improve their reliability and thus benefits users.

Our key contributions can be summarized as follows:

- We design a dynamic programming algorithm to solve the ISP subscription problem optimally. We also design a more efficient algorithm for concave pricing functions. We demonstrate the effectiveness of the general algorithm using real traffic traces and realistic pricing data.
- We study the effects of multihoming on ISPs by formulating the ISP pricing problem as a non-cooperative game using a realistic charging model. We prove that if cost is the only criterion used by a user to determine which ISPs to subscribe to, all ISPs receive zero revenue at any equilibrium.
- We also study a more general formulation in which different ISPs provide different levels of reliability and users choose ISPs to both improve reliability and reduce cost. We show that an ISP's revenue is positive and determined by its reliability at any equilibrium.

The rest of this paper is organized as follows. In Section 2, we describe the network and charging models. In Section 3, we propose dynamic programming

algorithms to solve the ISP subscription problem. In Section 4, we study the ISP pricing problem when cost is the only criterion. In Section 5, we investigate a more general formulation, in which different ISPs provide different levels of reliability. In Section 6, we review related work. Finally we conclude the paper in Section 7.

2 Network and Charging Models

We start with a description of our network and ISP charging models.

2.1 Network Model




Figure 1: An illustration of a user with K service providers.

A multihomed user has multiple links to the Internet for sending and receiving traffic, as shown in Fig 1. The implementation techniques of distributing traffic to the links are different for outgoing and incoming traffic. For outgoing traffic, a border router inside the user's network can actively control how traffic is distributed. For incoming traffic, a user can use NAT, BGP prepending, BGP selective announcement, and/or DNS to control the routes. For more detailed discussions about the implementations, we refer the readers to [1, 8, 11, 15, 16, 31]. In this paper, we consider only outgoing traffic.

2.2 Charging Models

Users pay ISPs for using their service. The cost incurred to a user is usually based on the amount of traffic a user generates, *i.e.*, $cost = c(p)$, where p is a variable determined by a user's traffic (which we will term the *charging volume*) and c is

a non-decreasing function which maps p to cost. Various charging models differ from one another in their choices of charging volume p and cost function c .

Usually, the cost function c is a piece-wise linear (non-decreasing) function, which we will use for our design and evaluation. There are several ways in which the charging volume p can be determined. Percentile-based charging and total-volume based charging are both in common use.

In this paper, we focus on percentile-based charging. This is a typical usage-based charging scheme currently in use by many ISPs [27]. Under this scheme, an ISP records the traffic volume a user generates during every 5-minute interval. At the end of a complete charging period, the q -th percentile of all 5-minute traffic volumes is used as the charging volume p for q -percentile charging. More specifically, the ISP sorts the 5-minute traffic volumes collected during the charging period in ascending order, and then computes the charging volume p as the traffic volume in the $(q\% \times I)$ -th sorted interval, where I is the total number of intervals in a charging period. For example, if 95th-percentile charging is in use and the charging period is 30 days, then the cost is based on the traffic volume sent during the 8208-th $(95\% \times 30 \times 24 \times 60/5 = 8208)$ sorted interval.

3 The ISP Subscription Problem

In this section, we first develop optimal algorithms to solve the ISP subscription problem. Then we demonstrate the effectiveness of our algorithms using real traffic traces and realistic pricing data.

3.1 Problem Formulation

The *ISP subscription problem* can be stated as follows: Given a set $\mathcal{K} = \{1, \dots, K\}$ of ISPs with cost functions c_k and charging percentiles q_k , where $k \in \mathcal{K}$, find a subset $S \subseteq \mathcal{K}$ of ISPs that minimizes the user's total cost $\sum_{k \in S} c_k(p_k)$, where p_k is the charging volume of ISP k . Formally,

$$\begin{aligned} \min_S \sum_{k \in S} c_k(p_k) \\ \text{subject to } S \subseteq \mathcal{K}. \end{aligned} \quad (1)$$

Compared with the cost optimization problem formulated in [15], the ISP subscription problem is different in that [15] assumes that the ISP subscription decision has already been made, so all ISPs can be used, while in our ISP subscription

problem the user has the freedom to select a subset of ISPs to use in order to minimize cost. A user can benefit from selecting a subset of the ISPs if the ISPs charge non-zero base prices.

3.2 A Dynamic Programming Algorithm

Table 1 introduces the notations we will use. We define *aggregated charging volume* and *total peak percentile* as follows. Suppose a user subscribes to a set S of ISPs, then the user's aggregated charging volume is defined as the sum of p_k , i.e., $\sum_{k \in S} p_k$, and the user's total peak percentile is defined as the sum of z_k , i.e., $\sum_{k \in S} z_k$, where $z_k = 1 - q_k$. Assume a user subscribes to a set of ISPs, denoted as S . Then aggregated charging volume and total peak percentile satisfy the following two properties [15]. First, if the cost functions c_k of all ISPs in S are non-decreasing, then the user's minimum total cost $\sum_{k \in S} c_k(p_k)$ is also a non-decreasing function of the user's aggregated charging volume. Second, the user's aggregated charging volume has a lower bound, which is $V_0(S) \stackrel{\text{def}}{=} \text{qt}(V, 1 - \sum_{k \in S} z_k)$, where qt is the quantile function, and V is the time series of the user's total traffic volume. The lower bound is achievable when each ISP has sufficient bandwidth to handle the user's traffic by itself. Below we will focus on this scenario, since multihoming is often used to provide high reliability – even when all other ISPs fail, a user can still use the single remaining ISP to carry out its traffic.

Based on the above properties, now we reformulate the ISP subscription problem as in (2).

$$\begin{aligned} & \min_S \sum_{k \in S} c_k(p_k) \\ & \text{subject to } S \subseteq \mathcal{K} \\ & \sum_{k \in S} p_k = V_0(S). \end{aligned} \tag{2}$$

The reformulation in (2) allows us to design efficient optimal algorithms.

Instead of solving the ISP subscription problem for a fixed K , we first generalize the problem. Let $\mathcal{K} = \{\infty, \dots, \mathcal{K}\}$ be the set of all ISPs. Let $C(n, k, p, z)$ denote the minimum cost when the user has aggregated charging volume p , total peak percentile z , and subscribes to no more than k out of the first n ISPs

Table 1: Notations

\mathcal{K}	The set of all ISPs, <i>i.e.</i> , $\mathcal{K} = \{1, \dots, K\}$, where K is the total number of ISPs.
c_k	The cost function of ISP k . We assume that c_k is a non-decreasing function.
I	The number of time intervals in a charging period.
$v^{[i]}$	The total traffic volume during interval i . Let time series $V = \{v^{[i]} \mid 1 \leq i \leq I\}$.
$t_k^{[i]}$	The volume of traffic distributed to ISP k during interval i . Let time series $T_k = \{t_k^{[i]} \mid 1 \leq i \leq I\}$. Note that $V = \sum_k T_k$ (with vector summation).
q_k	The charging percentile of ISP k , <i>e.g.</i> , $q_k = 0.95$ if an ISP charges at 95th-percentile.
z_k	$z_k \stackrel{\text{def}}{=} 1 - q_k$.
$\text{qt}(X, q)$	The $[q * X]$ -th value in X_{sorted} (or 0 if $q \leq 0$), where X_{sorted} is X sorted in non-decreasing order, and $ X $ is the number of elements in X .
p_k	The charging volume of ISP k , (<i>i.e.</i> , $p_k = \text{qt}(T_k, q_k)$). For example, if ISP k charges at 95th-percentile, then p_k is the 95th-percentile of the traffic assigned to ISP k .
$V_0(S)$	$V_0(S) \stackrel{\text{def}}{=} \text{qt}(V, 1 - \sum_{k \in S} z_k)$, where $S \subseteq \mathcal{K}$ is a subset of ISPs, and V is the time series of the total traffic volumes of a user.

$\{1, \dots, n\}$. Formally,

$$\begin{aligned}
C(n, k, p, z) &= \min_S \sum_{k \in S} c_k(p_k) \\
\text{subject to } & S \subseteq \{1, \dots, n\} \\
& |S| \leq k \\
& \sum_{k \in S} p_k = p \\
& \sum_{k \in S} z_k = z.
\end{aligned} \tag{3}$$

Note that for some combinations of n , k , p , and z , there may not be any $S \subseteq \mathcal{K}$ that satisfies all of the constraints. In such cases, we define $C(n, k, p, z) = +\infty$.

Given the definition of $C(n, k, p, z)$, we have that the solution to (2) is $\min_z C(K, K, \text{qt}(V, 1-z), z)$. Thus we can solve the ISP subscription problem (1) if we can compute $C(n, k, p, z)$ efficiently.

The generalization allows us to observe that $C(n, k, p, z)$ satisfies the recurrence relation shown in (4), assuming that the cut points on the cost functions are all integers. This recurrence relation leads naturally to a dynamic programming algorithm. The algorithm solves the ISP subscription problem optimally when there is no capacity constraint (i.e., each ISP can handle the user's traffic by itself).

$$C(n, k, p, z) = \min \begin{cases} C(n-1, k, p, z) \\ \min_{0 \leq y \leq p} (c_n(y) + C(n-1, k-1, p-y, z-z_n)) \end{cases} \tag{4}$$

$$C(n, k, p, z) = \min \begin{cases} C(n-1, k, p, z) \\ c_n(0) + C(n-1, k-1, p, z-z_n) \\ c_n(p) + C(n-1, k-1, 0, z-z_n) \end{cases} \tag{5}$$

Now, we analyze the complexity of the algorithm. Its time complexity is $O(K^2 Z P^2)$, and its space complexity is $O(K Z P)$, where Z is the total number of choices of z , and P is the total number of choices of p . The percentile z is of the form i/I , where I is the total number of intervals in a charging period, and i is an integer between 0 and I . So we have $Z = I$. Since the input specifies the user's traffic in each interval to decide the charging volumes, the input complexity is linear in I , instead of $\log I$. In the worst case, the dynamic programming algorithm

is exponential for general pricing functions. In practice, however, the cost functions are usually piece-wise linear or step functions with very coarse-grained cut points, so P is usually small. In addition, it is easy to use discretization to make tradeoffs between precision versus computational time and space complexity.

3.3 Polynomial-time Dynamic Programming Algorithm for Concave Functions

If the ISP's cost functions are concave (as is often the case), we can specialize the preceding dynamic programming algorithm to design a more efficient, polynomial-time algorithm. First, for concave cost functions, we have the following observation:

Lemma 1 *Let $S = \{1, \dots, n\}$ be a set of n ISPs. If the total cost function $c(p_1, \dots, p_n)$ is concave, then the following minimization problem*

$$\begin{aligned} & \min_{p_1, \dots, p_n} c(p_1, \dots, p_n) \\ & \text{subject to } \sum_{k=1}^n p_k = p > 0 \\ & \quad p_k \geq 0 \quad \forall k \in S \end{aligned} \quad (6)$$

has an optimal solution in which the charging volumes p_k are 0 for all but one ISP.

Proof: Denote by e_k the k -th unit vector. Suppose (p_1, \dots, p_n) is an optimal solution. Since $\sum_{k=1}^n \frac{p_k}{p} = 1$, we have $c(p_1, \dots, p_n) = c(\sum_{k=1}^n p_k e_k) = c(\sum_{k=1}^n \frac{p_k}{p} (p e_k)) \geq \sum_{k=1}^n \frac{p_k}{p} c(p e_k) = \sum_{k=1}^n p_k \frac{c(p e_k)}{p}$, where the inequality is due to the concavity of c .

Let $k^* = \operatorname{argmin}_k \frac{c(p e_k)}{p}$, we have $\sum_{k=1}^n p_k \frac{c(p e_k)}{p} \geq \sum_{k=1}^n p_k \frac{c(p e_{k^*})}{p} = c(p e_{k^*})$. In addition, $p e_{k^*}$ also satisfies the constraint in (6), so $p_{k^*} = p$, and $p_k = 0, \forall k \neq k^*$ is an optimal solution to (6).

Given the above lemma, we observe that if all cost functions c_k are concave, then for any subset $S = \{1, \dots, n\}$ of ISPs, the user's total cost $c(p_1, \dots, p_n) = \sum_{k=1}^n c_k(p_k)$ is also concave. Applying Lemma 1 to the second case of (4), we have that the minimum occurs either when $y = 0$ or $y = p$. Therefore, we do not need to search for y all the way from 0 to p . Instead, we only need to compare the user's total cost when $y = 0$ with that when $y = p$. This leads to a new recurrence

relation shown in (5). Notice that now, in order to compute $C(K, K, \text{qt}(V, 1 - z), z)$, instead of having to compute $C(n, k, p, z)$ for all p values as in (4), we only need to compute $C(n, k, p, z)$ values for $p = \text{qt}(V, 1 - z)$ and $p = 0$. Therefore, a dynamic programming algorithm based on the recurrence relation in (5) has time complexity $O(K^2Z)$ and space complexity $O(KZ)$, which are both polynomial.

3.4 Greedy Subscription

The greedy algorithm chooses a set of k ISPs, denoted as S_k , as follows. In the first iteration, it examines all ISP sets with size no larger than r ($r \leq k$), and selects the one which yields the lowest cost. In the second iteration, it searches for a new ISP to add which in conjunction with the ISPs already picked yields the lowest cost. It iterates until k ISPs have been chosen. Here r is a tuning parameter of the algorithm, and all ISP sets with size no larger than r are exhaustively searched. If $r = n$, all subsets are searched, and hence the solution is optimal; however, in this case, its complexity is much higher than the dynamic programming algorithm. Using different values of r can trade off running time for solution quality. In our evaluation, we set $r = 1$.

3.5 Random Subscription

The random subscription algorithm randomly chooses a specified number of ISPs under the constraint that the total bandwidth of the subscribed ISPs is large enough to accommodate the user's traffic. In our evaluation, we run the random subscription algorithm 20 times and report the average.

3.6 Evaluations

In this subsection, we evaluate the performance of our ISP subscription algorithms using two sets of Abilene traffic traces. The traces contain netflow data from an institution (National Institutes of Health) and an enterprise (Red Hat Inc.) on the Internet-2 from October 8, 2003 to January 6, 2004. In our evaluations, We scale each set of traffic traces such that each ISP can handle the traffic by itself.

In each evaluation scenario, there are 10 ISPs and 1 subscriber. The 10 ISPs have 5 different pricing functions as shown in Fig. 2. Each pricing function has 2 ISPs associated with it. The shape of the pricing functions reflects the general pricing practice of decreasing unit cost as bandwidth increases; it is also consistent

with the pricing functions we are aware of (*e.g.*, [4, 24]). We refer readers to [15] for more details. The subscription cost is computed based on the 95-th percentile of the subscriber's traffic during each month.

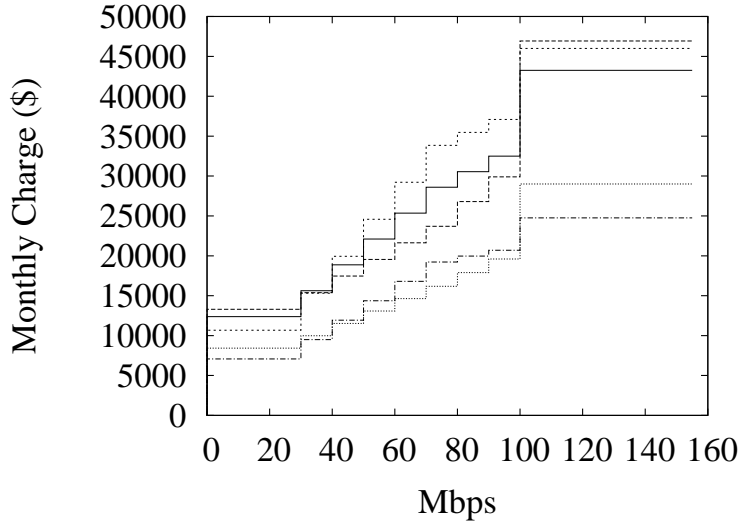


Figure 2: The complex OC3 pricing functions.

We compare our optimal subscription algorithm against the random subscription algorithm and the greedy subscription algorithm. In our first set of experiments, we assume that the user knows its traffic volume in advance. Fig. 3 compares the total cost incurred using the three subscription algorithms as we vary the number of ISPs the user subscribes to. We present here the results using traces obtained in December 2003. Results using other months' traces show the same relative ranking of the three algorithms. Random subscription continues to do much worse than the optimal, while the difference between the greedy and the optimal algorithms is much smaller.

We make the following observations. First, as expected, our optimal subscription algorithm yields the lowest cost in all cases. The random subscription algorithm incurs about 50% higher cost on average for both traces, and leads to over 100% higher cost in worst cases, especially when subscribing to a small number of ISPs. The greedy subscription yields similar cost to the optimal algorithm in most cases, but could result in up to 24% higher cost in worst cases. Second, we observe that adding ISPs initially helps reduce the total cost; as the number of ISPs increases further, the cost increases. To explain this, we note that an ISP's

cost involves two components: base charge and usage-based charge. Adding ISPs initially helps to accommodate burstiness of the traffic, thereby reducing usage-based charge. The initial reduction in usage-based charge is large enough to offset the additional ISPs' base charge. As the number of subscribed ISPs increases further, the reduction in usage-based charge becomes smaller than additional base charge. Therefore the total cost increases with additional ISPs.

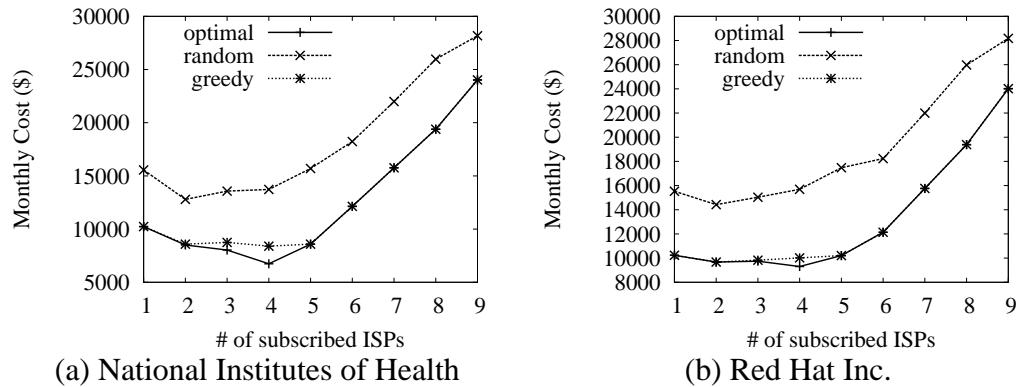


Figure 3: Comparison of the three subscription algorithms. Traces are obtained in December 2003.

In our second set of experiments, we study the case where the user does not know its traffic *a priori*, but predicts one month's traffic based on the previous month's traffic and applies the three subscription algorithms to the predicted traffic. We call this scheme predicted subscription. We compare the results with the optimal subscription that knows traffic in advance. We present the results using trace obtained in November to predict the traffic of December 2003, as shown in Fig. 4. Results using other months' traces are similar.

We observe that our optimal subscription algorithm using predicted traffic performs fairly well. It performs close to the optimal algorithm under perfect knowledge about traffic patterns, and much better than the random subscription algorithm. In most cases, uncertainty in traffic patterns yields less than 5% cost increase on average for the optimal subscription algorithm. The greedy algorithm performs close to the optimal in most cases, but could lead to 24% higher cost in the worst case.

Although our evaluation shows that the greedy subscription algorithm performs reasonably well in most cases, it is worth noting that its worst case approximation ratio is unbounded for $r < n - 1$, as shown below. Consider $r + 2$ ISPs

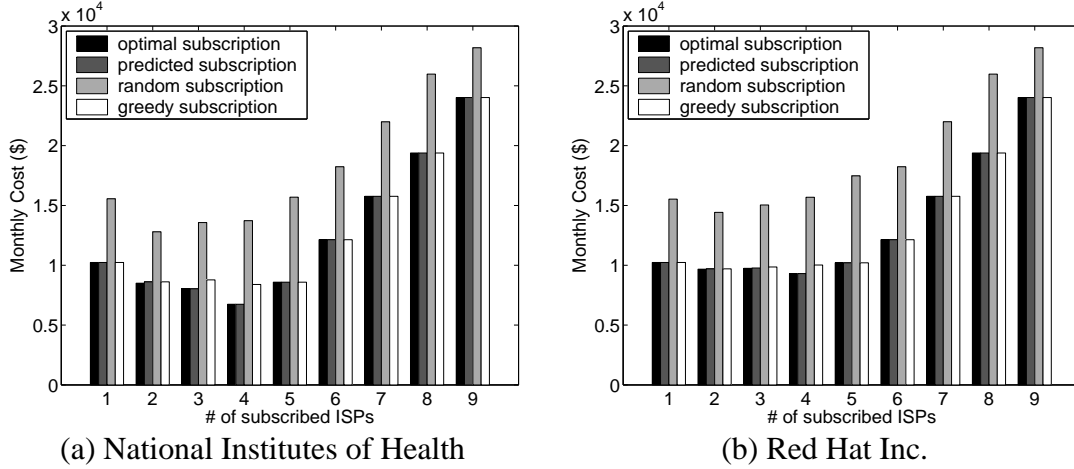


Figure 4: Impact of traffic fluctuation on subscription algorithms for December 2003. Prediction is based on traffic of November 2003.

that are available for subscription, and they all use 100-th percentile charging with the following pricing functions,

$$c_1(p) = A, \quad (7)$$

$$c_i(p) = \begin{cases} B & \text{if } p < V/r, \\ 2A & \text{if } p \geq V/r, \end{cases} \quad i = 2, \dots, r+2, \quad (8)$$

where $A \gg B > 0$, and V is the user's peak traffic volume. The greedy algorithm starts by exhaustively searching over all ISP sets within size r .

The optimal subscription cost is $(r+1)B$. In comparison, the greedy algorithm first selects ISP 1, since all other ISP sets of size within r have higher cost. Its final subscription is no less than A , since ISP 1 is included in the final selection. So the ratio between the greedy solution and optimal solution is no less than $A/(r+1)B$, which is unbounded. The above analysis can easily be generalized to the case of more than $r+2$ available ISPs by having $c_j(p) = 3A$ for $j > r+2$.

To summarize, in this section we develop a dynamic programming algorithm for solving the ISP subscription problem, and demonstrate its effectiveness using real traffic traces.

4 The ISP Pricing Problem

Our ISP subscription algorithm allows users to choose a subset of ISPs to subscribe to and minimize their costs. In response, ISPs may adjust their prices to maximize their revenue. How ISPs will adjust their prices is an interesting question because it helps us understand the evolution of Internet multihoming. In this section, we formulate the problem as a non-cooperative game and prove that, if cost is the only criterion used by a user for ISP subscription, all ISPs receive zero revenue at any equilibrium.

4.1 Problem Formulation

To make our game theoretical analysis more realistic, we use the realistic percentile-based charging model in our formulation. Using this model makes our analysis more involved, but we believe the results can be more relevant. In our formulation, we focus on the case where multiple ISPs compete for a single subscriber. Hereafter, we use subscriber and user interchangeably. We assume a special structure of pricing functions: ISP k receives revenue by charging the subscriber $c_k = a_k p_k + b_k$ if it is selected by the subscriber, and 0 otherwise. Here a_k is the unit price; p_k is the charging volume determined by the charging percentile q_k and time series of the subscriber's traffic assigned to ISP k ; and b_k is the base price.

We now define the game formally. The players of the game are a set $\mathcal{K} = \{\infty, \epsilon, \dots, \mathcal{K}\}$ of ISPs. The action space of player k is $\mathcal{R}^+ \times \mathcal{R}^+ \times [t, \infty]$. Specifically, a player adjusts its charging parameters $\{a_k, b_k, q_k\}$, where $a_k, b_k \in \mathcal{R}^+$, and $0 \leq q_k \leq 1$, such that its revenue is maximized. When an ISP changes its charging parameters, it should consider how the subscriber and the other ISPs will respond. Specifically, there exists competition among ISPs, and the subscriber takes advantage of our subscription algorithm to select a set of ISPs to minimize cost. It is worth noting that, since all pricing functions are concave, when subscribing to a set S of ISPs, the subscriber always allocates traffic in such a way that only ISPs in S with minimum unit cost can have non-zero charging volume.

To state our assumptions clearly, we first introduce some more terms. We call the set of ISPs computed by our subscription algorithm as a feasible set. An ISP in a feasible set is called a feasible ISP. There may exist multiple feasible sets. Let \mathcal{F} denote the set of all feasible sets. Note that a subscriber has equal cost on any feasible set of ISPs. Let $a_{min}(S)$ denote the minimum unit price of all ISPs in a set S , and $n_a(S)$ the number of ISPs having the same unit price as a in set S . Note

that S may have multiple ISPs with the same minimum unit price. For a complete list of notations we use hereafter, please refer to the Appendix.

Finally, we explicitly make the following assumptions in the analysis below:

- We assume that each feasible set has equal probability of being selected by the subscriber. We also assume that when the subscriber is multihomed to a feasible ISP set S , the aggregated charging volume traffic $c p(S)$ is distributed evenly across those ISPs with minimum unit price.
- We assume that each of the ISPs has enough capacity to accommodate all of the subscriber's traffic c , and that the total amount of traffic c that the subscriber generates is bounded. We also assume that each ISP only charges a finite price and can adjust its unit price and base price in an infinitesimal amount.
- We assume that there is perfect information sharing among the subscriber and ISPs; that is, each of them has perfect information about the others when making decisions.

4.2 Summary of Results

Our analysis based on the percentile-based charging model is quite involved. In the interest of clarity, we first summarize our results and the structure of our analysis.

The main result of this section is that an action profile of the ISP pricing problem is an equilibrium if and only if all ISPs receive zero revenue in the outcome. It is obvious that any action profile with an outcome in which all ISPs receive zero revenue is an equilibrium of the game, since if an ISP unilaterally increases its price, the subscriber can always switch to other ISPs that charge zero; thus the revenue of the ISP is not increased.

The remaining challenge then is to prove that all ISPs receive zero revenue at any equilibrium. We first show that at any equilibrium, either all ISPs receive zero revenue or all of them receive positive revenue. There does not exist an equilibrium in which some ISPs receive zero revenue while others receive positive revenue. Therefore, we only need to show that there does not exist an equilibrium with positive revenue for all ISPs, which we call a positive-revenue equilibrium. To do so, we derive the following properties that a positive-revenue equilibrium should have. We first show that a subscriber is not able to free-ride all providers (pay only the base price), and that any feasible ISP k must have a unit price equal

to the maximum of all positive minimum unit prices of all of the feasible sets containing ISP k , at any positive-revenue equilibrium. We then show that a feasible ISP can reduce its unit price by a small amount without introducing any new feasible set with the same minimum cost. We then prove that any ISP k in a feasible set must have the unique minimum unit price in that set. Using these properties, we prove that there exists no positive-revenue equilibrium.

4.3 Equilibrium Analysis

We consider an arbitrary ISP k . Without loss of generality, we assume k appears in feasible sets $Z_{i \in \{1, \dots, N_k\}}$. Note that there must exist at least one such set thus $N_k \geq 1$. Also, there are $N_{-k} \geq 0$ feasible sets that do not contain k , and $N = N_k + N_{-k}$ is the total number of feasible sets. We sort Z_i in non-increasing order of $a_{min}(Z_i)$. Let $S_{i \in \{1, \dots, N_k\}}$ denote the sorted sets. Without loss of generality, we assume that S_1, \dots, S_{n_k} have the same minimum unit price, and denote this minimum unit price by a . Apparently, $a_{min}(S_j) = a, \forall j \in \{1, \dots, n_k\}$, and $a_k \geq a$. In addition, a subscriber has equal cost on all feasible sets of ISPs; that is, $c(S_i), \forall i \in \{1, \dots, N_k\}$, are equal and we denote it by c_{min} . We denote by R_k the total revenue of ISP k .

We first notice that at any equilibrium, either all ISPs or none of them receives positive revenue:

Theorem 1 *At any equilibrium, either $R_k = 0, \forall k \in \mathcal{K}$, or $R_k > 0, \forall k \in \mathcal{K}$.*

Proof: Proof by contradiction. Assume $R_k = 0, R_{k'} > 0, k \neq k'$. Then ISP k can increase its revenue from 0 to some positive value by reducing its charge to the minimum of $R_{k'}$ (e.g., by setting $a_k = 0, q_k = 0$, and $b_k = \min_{k'} R_{k'}$), which leads to a contradiction.

The above theorem tells us that there does not exist an equilibrium in which some ISPs receive zero revenue while others receive positive revenue. Now we only need to show that there does not exist an equilibrium with positive revenue for all ISPs. Therefore, in the remaining part of this subsection, we consider only these positive-revenue equilibria.

Next we show the first property of a positive-revenue equilibrium: the subscriber is not able to free-ride all providers at a positive-revenue equilibrium.

Lemma 2 *At any positive-revenue equilibrium, there exists at least one $S_{k'} \in \mathcal{F}$ such that $p(S_{k'}) > 0$, for some $k' \in \{1, \dots, n_k\}$.*

Proof: Proof by contradiction. Assume that $p(S_{k'}) = 0, \forall k' \in \{1, \dots, n_k\}$. Then we have $p_k = 0$. Therefore, ISP k 's total revenue is

$$R_k = \frac{N_k}{N_k + N_{-k}} b_k.$$

We first prove that $c_{min} > b_k$. Suppose that $c_{min} = b_k$, then $b_j = 0$ for an arbitrary ISP $j \in S_{k'}$, for some $k' \in \{1, \dots, n_k\}$. We consider the following two cases:

1. all feasible sets contain ISP k . Then any other feasible ISP j receives zero revenue since $b_j = 0$ and $p_j = 0$ (because $p(S_{k'}) = 0, \forall k' \in \{1, \dots, n_k\}$). This contradicts with Theorem 1.
2. some feasible sets do not contain ISP k . Then ISP k can set $a_k = 0$ and choose a small positive value

$$\epsilon < \frac{N_{-k}}{N_k + N_{-k}} b_k$$

so that it can attract all of the user's traffic and receive revenue $R'_k = b_k - \epsilon > R_k$.

Therefore, we have $c_{min} > b_k$.

We next show that ISP k can increase its revenue by reducing its base price b_k . Because all original feasible sets of ISPs have equal costs and $c_{min} > b_k$ we can find a small positive value δ satisfying

$$\delta < c_{min} - \frac{N_k}{N_k + N_{-k}} b_k$$

such that ISP k can set $b_k = c_{min} - \delta$ to increase its revenue to $R'_k = c_{min} - \delta > R_k$. Therefore we derive a contradiction.

Next we study the second property of a positive-revenue equilibrium: any feasible ISP k must have the same unit price as all other ISPs in feasible sets containing k .

Lemma 3 *At a positive-revenue equilibrium, $a_k = a$ if $a > 0$.*

Proof: Proof by contradiction.

Assume $a_k > a$. We compare the expected revenue of ISP k before and after letting $a_k = a$.

The expected revenue of ISP k when $a_k > a$ is

$$R_k = \frac{N_k}{N_k + N_{-k}} b_k,$$

since each of the N feasible sets is chosen by the subscriber with equal probability, and ISP k receives b_k revenue when any set $S_{i \in \{1, \dots, N_k\}}$ is chosen (recall that the subscriber runs smart routing algorithm to optimize cost such that all ISPs in S_i with unit price higher than $a_{\min}(S_i)$ have zero charging volume; therefore, the charging volume of ISP k is 0).

We next consider the expected revenue of ISP k , R'_k , after letting $a_k = a$, in the following three cases:

1. There are no new feasible sets introduced by ISP k 's action. Then we have

$$R'_k = \frac{1}{N_k + N_{-k}} \left(a \sum_{k'=1}^{n_k} \frac{p(S_{k'})}{n_a(S_{k'})} + b_k N_k \right).$$

Therefore,

$$R'_k - R_k = \frac{a}{N_k + N_{-k}} \sum_{k'=1}^{n_k} \frac{p(S_{k'})}{n_a(S_{k'})}.$$

Since $a > 0$, we only need to show that $p(S_{k'}) > 0$ for some $k' \in \{1, \dots, n_k\}$. Applying Lemma 2, we know that $\exists k' \in \{1, \dots, n_k\}$ such that $p(S_{k'}) > 0$. Therefore, $R'_k > R_k$.

2. There are $N'_k \geq 1$ new feasible sets with the same total cost as c_{\min} introduced by setting $a_k = a$.

Denote these new sets by $S'_i, \forall i \in \{1, \dots, N'_k\}$. Note that ISP k 's unit price must be the minimum unit price of all ISPs in S'_i . Then the expected revenue

of ISP k is

$$\begin{aligned}
R'_k &= \frac{1}{N_k + N'_k + N_{-k}} \left(a \sum_{k'=1}^{n_k} \frac{p(S_{k'})}{n_a(S_{k'})} + b_k N_k \right) \\
&\quad + \frac{1}{N_k + N'_k + N_{-k}} \left(a \sum_{k'=1}^{N'_k} p(S'_{k'}) + b_k N'_k \right) \\
&> \frac{N_k + N'_k}{N_k + N'_k + N_{-k}} b_k \\
&> \frac{N_k}{N_k + N_{-k}} b_k.
\end{aligned}$$

Note that the first inequality is derived by applying Lemma 2 and the fact that $a > 0$.

3. There are $N'_k \geq 1$ new feasible sets with total cost $c'_{min} < c_{min}$ introduced by setting $a_k = a$. Denote these new sets by $S'_i, i \in \{1, \dots, N'_k\}$. Note that none of the old feasible sets in \mathcal{F} is feasible now since they have higher cost. Note also that $k \in S'_i, \forall i \in \{1, \dots, N'_k\}$.

Therefore, the expected revenue of ISP k is

$$R'_k = \frac{a}{N'_k} \sum_{k'=1}^{N'_k} \frac{p(S'_{k'})}{n_a(S'_{k'})} + b_k.$$

Then by applying Lemma 2 and the fact that $a > 0$, we have

$$R'_k - R_k > \frac{N_{-k}}{N_k + N_{-k}} b_k \geq 0.$$

Apparently, we have contradiction $R'_k > R_k$ in all cases. Therefore, $a_k = a$ if $a > 0$ at a positive-revenue equilibrium.

The above lemma shows that a feasible ISP k is not able to increase its revenue by increasing its unit price. However, it is still unknown if it is possible for ISP k to increase revenue by reducing its unit price. We show below (Lemma 4) that a feasible ISP can reduce its unit price by a small amount such that all feasible sets remain unchanged; based on this lemma, we then prove by contradiction that a feasible ISP can reduce its unit price by a small amount to increase its revenue, if that unit price is not the unique minimum (Lemma 5).

Lemma 4 *At a positive-revenue equilibrium, if $a_k > 0$, there exists a small number $\epsilon > 0$ such that ISP k can reduce its unit price to $a_k - \epsilon$ without introducing any new feasible sets.*

Proof: Consider the sets in $\mathcal{U}_{||} - \mathcal{F}$, where $\mathcal{U}_{||}$ is the set of all subsets of \mathcal{K} containing ISP k . Note that we can safely drop those sets that do not contain ISP k since they are not affected by ISP k 's action of reducing a_k .

For any set $Z \in \mathcal{U}_{||} - \mathcal{F}$, we show that we can find a small value ϵ such that $c(Z)$ is still larger than c_{min} after ISP k reduces its unit price to $a_k - \epsilon$. Note that when $\mathcal{U}_{||} - \mathcal{F} = \phi$, no new feasible set is introduced by ISP k 's reducing its unit price.

Suppose $Z = \{k, u_1, \dots, u_l\} \in \mathcal{U}_{||} - \mathcal{F}$. Therefore, we have $c(Z) > c_{min}$. Let $c(Z)$ and $c'(Z)$ denote the expected total cost of the subscriber before and after ISP k reduces its unit price, respectively, if the subscriber uses ISPs in Z as the providers. Now we derive the condition for ϵ such that $c'(Z) > c_{min}$.

Note that

$$c(Z) = a_k p_k + b_k + \sum_{k' \in \{u_1, \dots, u_l\}} (a_{k'} p_{k'} + b_{k'})$$

and

$$c'(Z) = a_k \hat{p}_k - \epsilon \hat{p}_k + b_k + \sum_{k' \in \{u_1, \dots, u_l\}} (a_{k'} \hat{p}_{k'} + b_{k'}),$$

where $p_{k'}$ and $\hat{p}_{k'}$ are the charging volumes of ISP k' before and after ISP k reduces its unit price, respectively.

Consider three cases as follows.

- case 1: $a_k - \epsilon$ is not the minimum unit price in set Z . Then we have $c(Z) = c'(Z) > c_{min}$ since $\hat{p}_k = p_k = 0$.
- case 2: $a_k - \epsilon$ is the minimum unit price while a_k is not. Then we can reduce ϵ such that $a_k - \epsilon$ is no longer the minimum unit price and this case degenerates to case 1.
- case 3: both $a_k - \epsilon$ and a_k are the minimum unit price. Then we have $\hat{p}_k = p_k \geq 0$ (we have equality here because $p(Z) \geq 0$); therefore, $c(Z) = c'(Z) > c_{min}$.

Let ϵ_Z denote the appropriate ϵ value for a set Z satisfying the above conditions. Therefore, we can always find an ϵ by taking the minimum of all ϵ_Z .

Lemma 5 *At a positive-revenue equilibrium, if $a_k > 0$ and $a > 0$, then $a_k = a$ is the unique minimum unit price in $S_i, \forall i \in \{1, \dots, N_k\}$.*

Proof: Proof by contradiction. Assume at a positive-revenue equilibrium, a_k is not the unique minimum unit price in $S_i, \forall i \in \{1, \dots, n_k\}$.

By applying Lemma 3, we have

$$a_k = a_{\min}(S_1) = a_{\min}(S_2) = \dots = a_{\min}(S_{n_k}) = a.$$

By applying Lemma 4, we can find $\epsilon > 0$ such that no new feasible set is introduced by reducing ISP k 's unit price to $a_k - \epsilon$.

Now we compare the revenue received by ISP k before and after ISP k reduces its unit price. The revenue received before the reduction of a_k is

$$R_k = \frac{1}{N_k + N_{-k}} \left(a_k \sum_{k'=1}^{n_k} \frac{p(S_{k'})}{n_a(S_{k'})} + N_k b_k \right).$$

The revenue after a_k is reduced to $a_k - \epsilon$ is

$$R'_k = \frac{a_k - \epsilon}{n_k} \sum_{k'=1}^{n_k} p(S_{k'}) + b_k,$$

because

1. the sets $S_{n_k+1}, \dots, S_{N_k}$ are no longer feasible because the user has lower cost on S_1, \dots, S_{n_k} and the user will subscribe to these sets; therefore, ISP k receives no revenue from the sets $S_{n_k+1}, \dots, S_{N_k}$;
2. all of the sets S_{N_k+1}, \dots, S_N that do not contain k are no longer feasible because the user has smaller cost on sets S_1, \dots, S_{n_k} ;
3. we have chosen ϵ carefully such that no new set is introduced by reducing a_k ; therefore, ISP k receives no revenue from other sets in $\mathcal{U} - \mathcal{F}$.

Next, we compare R'_k and R_k by considering the following two cases:

1. $N > 1$: since $N \geq N_k$, we note that if

$$\epsilon < \frac{N - N_k}{N} \frac{b_k}{\sum_{k'=1}^{n_k} p(S_{k'})} + a_k \sum_{k'=1}^{n_k} \left(1 - \frac{n_k}{N}\right),$$

then $R'_k > R_k$ (the right-hand side of the above inequality is always positive). Therefore, we could choose ϵ such that the above condition is satisfied, which leads to a contradiction.

2. $N = 1$: Then $N = N_k = 1$, and we have one single feasible set S_1 ($k \in S_1$). If $|S_1| = 1$, then we already have $a_k = a_{\min}(S_1)$ and a_k is the unique minimum unit price; otherwise, we can choose a small positive value ϵ satisfying the preceding inequality and derive a contradiction.

Finally, we prove the property of any equilibrium based on the preceding lemmas.

Theorem 2 *At any equilibrium, every ISP k has zero revenue.*

Proof: Proof by contradiction.

By applying Theorem 1, we know that there can be only two possible cases: (1) $R_k = 0, \forall k \in \mathcal{K}$, or (2) $R_k > 0, \forall k \in \mathcal{K}$. Therefore, we only need to prove that the second case does not show up in any equilibrium.

For an arbitrary ISP k , $R_k > 0$, ISP k must be in some feasible set. We next consider the case where ISP k is in some feasible set. Specifically, we examine the following two cases:

1. $a > 0$.

By applying Lemma 5, $a_k = a$ is the unique minimum unit price in all of the sets $S_{i \in \{1, \dots, N_k\}}$. Therefore, all the sets with non-zero minimum unit price can have only one single ISP. Then we only need to consider $N > 1$ since when $N = 1$ it is trivial to show that each ISP k has zero revenue.

Assume ISP k has positive revenue R_k ; therefore, $c_{\min} > 0$. Then, ISP k can lower its price to increase its revenue. Let c_{\min} denote the revenue of each feasible set with a single ISP. Then we have

$$R_k = \frac{c_{\min}}{N}.$$

ISP k can lower its charge to $c_{\min} - \delta$, where $\delta > 0$ is a small number, and receives revenue

$$R'_k = c_{\min} - \delta.$$

Since $c_{\min} > 0$, we are always able to find a $\delta < \frac{N-1}{N}c_{\min}$ such that $R'_k > R_k$. This leads to a contradiction.

2. $a = 0$.

Note that in this case $b_k = 0$, because otherwise the subscriber can dump all traffic to the ISP with zero unit price in S_i without using ISP k , which means that $R_k = 0$ and leads to contradiction. Therefore, ISP k is not in $S_i, \forall i \in \{1, \dots, n_k\}$. By the assumption $R_k > 0$, ISP k should be in any feasible set since the subscriber can assign free-riding traffic to ISP k (by making charging volume zero) without incurring any extra cost. Therefore, all feasible sets containing k should have zero minimum unit price and k is in all of these sets; otherwise, it contradicts with our assumption that $a = 0$.

Therefore, we prove that all ISPs receive zero revenue at any equilibrium if cost is the only criterion used by a user to determine which subset of ISPs to subscribe to.

5 Reliability and the ISP Pricing Problem

We have shown in the previous section that if the only difference among ISPs is pricing, then all ISPs receive zero revenue at any equilibrium. However, in reality, pricing is not the only difference among ISPs, and cost is not the only concern of subscribers, either. Subscribers also consider many other factors, *e.g.*, reliability, ease of management, and security. In particular, reliability is a major motivation for the deployment of multihoming.

Given the importance of both cost and reliability, we investigate a more realistic formulation of the ISP pricing problem: how ISPs respond to multihomed subscribers when the subscribers optimize both cost and reliability.

5.1 Problem Formulation

Similar to the previous section, we formulate the problem as a non-cooperative game. We consider the percentile-based charging model and focus on the case where multiple ISPs compete for a single subscriber. The players, action spaces, and ISPs' revenue-maximization objectives are the same as those of the previous section.

The major difference between this formulation and the preceding formulation is that we consider both cost and reliability. Specifically, the subscriber takes advantage of our subscription algorithm and smart routing algorithms to minimize

cost and maximize reliability. To characterize the objective of the subscriber, we define a utility function of the subscriber on a subset S of ISPs as follows:

$$U(S) = w \sum_{k \in S} \log m_k - \sum_{k \in S} c_k,$$

where λ_k is the instantaneous failure rate of ISP k , $m_k = \frac{1}{\lambda_k}$ is the mean time between failures (MTBF) of ISP k , and $w > 0$ is the weight of the subscriber's preference of reliability over cost. We consider finite constant mean time between failures in this paper. The weight w reflects how much the subscriber is concerned with reliability. The higher the weight is, the more the subscriber prefers reliability over cost. The subscriber's objective is to choose a subset of ISPs such that its utility is maximized: $\max_{S \in \mathcal{U}} w \sum_{k \in S} \log m_k - \sum_{k \in S} c_k$. We assume that the subscriber always chooses as many ISPs as possible when maximizing its utility, *e.g.*, when the subscriber has equal utility over multiple feasible sets of ISPs, the subscriber prefers to multihome to ISPs in the largest set in order to improve reliability.

Note that our formulation and approach can be easily extended to consider other metrics that subscribers are concerned with.

5.2 Analysis of Existence and Non-uniqueness of Equilibrium

Given our non-cooperative game-theoretic formulation, we will prove the existence and non-uniqueness of equilibrium of the game in this section. The intuition of our proof is that no matter how ISPs change their charging parameters, the subscriber excludes a particular ISP k from subscription if that ISP charges more than $w \log m_k$, because the subscriber's utility becomes less if ISP k is included in subscription.

Formally, we have the following theorem stating the existence and non-uniqueness of equilibrium:

Theorem 3 *There exist multiple equilibria in the ISP pricing game.*

Proof: We first show that $\{a_k = 0, q_k = 0, b_k = w \log m_k\}, \forall k \in \mathcal{K}$, is an equilibrium. Note that the subscriber has zero utility and uses all ISPs as providers. Let S denote the set of ISPs used (*i.e.*, all ISPs) in this scenario. We look at all of the possible actions $\{a'_k, q'_k, b'_k\}$ taken by an arbitrary ISP k :

1. $b'_k < b_k$. We first consider the case where ISP k changes its base price only. In this case, $w \log m_k - b'_k > 0$ and the subscriber's utility is maximized by including all ISPs as providers. However, ISP k 's revenue decreases.

We next consider the cases where ISP k changes its unit price and/or charging percentile simultaneously. Note that no matter how a_k and q_k change, the subscriber can always distribute the traffic in such a way that $d_k p'_k = 0$. For instance, if $p_k > 0$ and $a'_k > 0$, then the subscriber re-distributes traffic such that $p'_k = 0$ and incurs no extra cost. This is achievable because all of the other ISPs have zero unit price. Therefore, no matter how a_k and q_k change, the subscriber's utility is always maximized by using all ISPs as providers, while ISP k 's revenue decreases.

2. $b'_k = b_k$. There are three cases to be considered: $a'_k > a_k$ and $q'_k = q_k$; $a'_k = a_k$ and $q'_k > q_k$; $a'_k > a_k$ and $q'_k > q_k$. We present the proof for the last case here and the proof for the first and second cases can be constructed similarly. The intuition of our proof is that the subscriber can always dump all aggregated charging volume traffic $p(S)$ to the ISP with zero unit price such that no extra cost is incurred. Therefore, no matter how a_k and q_k change, the subscriber can always distribute the traffic in such a way that $a'_k p'_k = 0$. For instance, if $a'_k > a_k$ and $q'_k > q_k$, the subscriber assigns $p'_k = 0$ amount of traffic to ISP k . The difference between utilities of including and excluding ISP k is:

$$\begin{aligned}
& U(S) - U(S - \{k\}) \\
&= \left(w \sum_{k'=1}^K \log m_{k'} - \sum_{k'=1}^K b_{k'} \right) \\
&\quad - \left(w \sum_{k'=1, k' \neq k}^K \log m_{k'} - \sum_{k'=1, k' \neq k}^K b'_{k'} \right) \\
&= w \log m_k - b'_k \\
&= 0.
\end{aligned}$$

Therefore, no matter how a_k and q_k change, the subscriber's utility is always maximized by using all ISPs as providers, while ISP k cannot increase its revenue by setting $b'_k = b_k$.

3. $b'_k > b_k$. The subscriber has the option to choose from two possible feasible sets S and $S - \{k\}$ in this case. Note that we implicitly apply our previous assumption that the subscriber chooses as many ISP as possible when maximizing its utility. Specifically, the subscriber has the same utility on any subset of $S - \{k\}$.

Similarly, we know that the subscriber can always distribute the traffic in such a way that $a'_k p'_k = 0$ no matter how ISP k changes a_k and q_k . We compute the difference of utilities on S and $S - \{k\}$ as follows:

$$\begin{aligned}
& U(S) - U(S - \{k\}) \\
&= \left(w \sum_{k'=1}^K \log m_{k'} - \sum_{k'=1}^K b_{k'} \right) \\
&\quad - \left(w \sum_{k'=1, k' \neq k}^K \log m_{k'} - \sum_{k'=1, k' \neq k}^K b'_{k'} \right) \\
&= w \log m_k - b'_k \\
&< 0.
\end{aligned}$$

Therefore, the subscriber chooses $S - \{k\}$ as the feasible set, and ISP k receives zero revenue by increasing b_k .

Therefore, we have found an equilibrium. Furthermore, it is obvious to see that $\{a_k = 0, q_k = 0.95, b_k = w \log m_k\}$ is another equilibrium. Therefore, there exist multiple equilibria.

5.3 Properties of Equilibria

Given the results of existence and non-uniqueness of equilibria, we consider a more challenging and important problem in this subsection: what properties does an equilibrium have? In particular, we are interested in understanding how the revenue is distributed across ISPs at an equilibrium.

We first show that every ISP has positive revenue at any equilibrium by proving Theorem 4. We summarize our intuition of the proof as follows. Any ISP k can attract the subscriber's subscription by charging the subscriber with some amount smaller than $w \log m_k$. By doing so, an ISP k can always receive positive revenue because the subscriber will receive a higher utility if ISP k is subscribed to.

Theorem 4 *At any equilibrium, $R_k > 0, \forall k \in \mathcal{K}$. In other words, all ISPs receive positive revenue at any equilibrium.*

Proof: Proof by contradiction. Suppose $R_k = 0, \exists k \in \mathcal{K}$. For ISP k , we consider the following two cases:

1. $R_k = 0$ because ISP k is not in any of the subscriber's feasible sets.

ISP k can set its charging parameters to $\{a'_k = 0, q'_k = 0, b'_k = w \log m_k - \delta\}$, where $w \log m_k > \delta > 0$ to increase its revenue. Note that the subscriber's utility is maximized if ISP k is included in the feasible sets since $w \log m_k - b'_k = \delta > 0$.

2. $R_k = 0$ and ISP k is in some of the subscriber's feasible sets.

Similarly, ISP k can set its charging parameters to $\{a'_k = 0, q'_k = 0, b'_k = w \log m_k - \delta\}$, where $w \log m_k > \delta > 0$, and the subscriber's utility is maximized if ISP k is still included in the feasible sets. Then ISP k is able to increase its revenue from 0 to $w \log m_k - \delta$.

Therefore, ISP k can increase its revenue by taking the above actions. This contradicts with our equilibrium assumption.

Theorem 4 above shows that every ISP has positive revenue when we consider the competition among all ISPs. It also indicates that new ISPs have incentives to join the competition and obtain a share of the total revenue. However, it is still not clear how the revenue is distributed across ISPs, or, an equivalent question is that what an ISP should do in order to increase its revenue in the game.

Next, we show that the revenue an ISP receives at any equilibrium is determined by both its own reliability of services and the weight of the subscriber's preference of reliability. Specifically, we prove the following theorem:

Theorem 5 *At any equilibrium, $R_k = w \log m_k, \forall k \in \mathcal{K}$.*

Proof: Theorem 3 shows that there exist (non-unique) equilibria, where $R_k = w \log m_k, \forall k \in \mathcal{K}$. We now prove that this property indeed holds for every equilibrium.

The proof follows from the fact that if an ISP k charges the subscriber $R_k < w \log m_k$, then it must be included in all of the feasible sets of the subscriber, because by including ISP k in the feasible set, the subscriber always increases its utility.

Suppose at an equilibrium, a particular ISP k charges the subscriber $R_k < w \log m_k$, then ISP k can increase its charge by a small positive amount δ . As long as $R_k + \delta < w \log m_k$, ISP k can be sure that it will be included in all of the feasible sets of the subscriber. Thus ISP k can increase its revenue to $R_k + \delta$, which contradicts with our equilibrium assumption.

On the other hand, suppose that at an equilibrium, a particular ISP k charges the subscriber $R_k > w \log m_k$ if the subscriber uses ISP k as a provider. This leads to a contradiction: ISP k receives zero revenue because the subscriber has higher utility by excluding ISP k from the feasible set; however, ISP k can increase its revenue by charging the subscriber $w \log m_k$ if the subscriber takes it as a provider. This contradicts with our equilibrium assumption. Note that here we assume that the subscriber uses as many ISPs as possible when there are multiple utility-maximizing feasible sets.

A couple of comments follow.

First, by considering both reliability and cost, we show that ISPs receive positive revenue in the competition; therefore, new providers have incentives to join the competition and share the total revenue.

Second, at any equilibrium, an ISP's revenue is jointly determined by that ISP's reliability and the subscriber's weight of preference. Therefore, ISPs have incentives to improve their reliability by upgrading their networks. On the other hand, the subscriber also benefits from the competition among ISPs. By adjusting relative preference between reliability and cost, a subscriber can trade reliability for cost, or vice versa. Our results indicate that the wide deployment of multihoming can be beneficial to the global Internet, since it provides incentives to ISPs to improve their reliability.

6 Related Work

We classify the related work into four areas: analysis of multihoming benefits, algorithm design for smart routing, implementation techniques for smart routing, and Internet pricing.

There are several papers that evaluate the potential benefits of smart routing, including [1, 13, 28, 29]. For example, in [1], Akella *et al.* quantify the potential performance and reliability benefits of multihoming using real Internet traces, and conclude that a careful choice of upstream providers is crucial. Dai *et al.* quantify the potential economic benefits to both subscribers and ISPs [13]. Our

work differs from the above in that we use both cost and performance as metrics of interest.

The potential performance and economic benefits of smart routing motivate research studies on designing algorithms for smart routing (*e.g.*, [1, 3, 15, 21, 23]). For example, Akella *et al.* [3] propose and evaluate a series of schemes to optimize the performance of multihomed users. In [15] the authors design smart routing schemes to dynamically distribute traffic among different external links to optimize cost and performance. They also study the interactions between multiple smart routing users, and between smart routing and single-homed users. Our work is complementary to the above work in that both of the above work consider the case where users have already decided which ISPs to subscribe to, whereas in this paper we study the ISP subscription problem. Moreover we analyze the implications of users' cost minimization on ISP pricing strategies.

On the implementation side, [8, 11, 18, 26, 31] propose implementing smart routing using BGP peering, whereas F5 Networks [14] and Radware [25] implement smart routing using DNS and NAT.

Finally, there is a large body of literature on Internet pricing strategies and competition (*e.g.*, [5, 6, 7, 10, 12, 17, 19, 20, 22, 30, 32]). These papers consider abstract charging models, while our work studies the percentile-based charging model which is widely used by today's ISPs. We believe analysis using a realistic charging model can provide much needed insight in understanding the implications of multihoming.

7 Conclusion

In this paper, we study two related problems — which subset of ISPs a user subscribes to to minimize cost, and how ISPs respond to the user's selection by changing their pricing strategies. Our results show that a user can apply the dynamic programming algorithm to effectively reduce its cost. In response to users' cost optimization, ISPs will adapt their pricing strategies. Using the percentile-based charging model which is widely used by today's ISPs, we formulate the pricing problem as a non-cooperative game. Our results show that if cost is the only criterion used by a user to determine which ISPs to subscribe to, at any equilibrium all ISPs receive zero revenue. To be more practical, we consider the case where different ISPs provide different levels of reliability, and users choose ISPs to both improve reliability and reduce cost. In this case, at any equilibrium an ISP's rev-

enue is positive and determined by its reliability.

Acknowledgments

Rick Summerhill and Mark Fullmer provided us research access to Abilene's traffic data. We would also like to thank Yin Zhang and the anonymous reviewers for giving us valuable comments.

References

- [1] A. Akella, B. Maggs, S. Seshan, A. Shaikh, and R. Sitaraman. A measurement-based analysis of multihoming. In *Proceedings of ACM SIGCOMM '03*, Karlsruhe, Germany, Aug. 2003.
- [2] A. Akella, J. Pang, A. Shaikh, S. Seshan, and B. Maggs. A comparison of overlay routing and multihoming route control. In *Proceedings of ACM SIGCOMM '04*, Portland, Oregon, Aug. 2004.
- [3] A. Akella, S. Seshan, and A. Shaikh. Multihoming performance benefits: An experimental evaluation of practical enterprise strategies. In *Proceedings of USENIX Technical Annual Conference*, Boston, MA, June 2004.
- [4] Amextel. <http://www.amextel.com/dedicated.htm>.
- [5] L. Anania and R. Solomon. *Internet Economics*, chapter Flat: The Minimalists Price, pages 91–118. MIT Press, 1997.
- [6] R. E. Azouzi, E. Altman, and L. Wynter. Telecommunications network equilibrium with price and quality-of-service characteristics. In *Proceedings of ITC*, Berlin, Germany, Sept. 2003.
- [7] T. Basar and R. Srikant. Revenue-maximizing pricing and capacity expansion in a many-users regime. In *Proceedings of IEEE INFOCOM '02*, New York, NY, June 2002.
- [8] T. Bates and Y. Rekhter. *Scalable Support for Multi-homed Multi-provider Connectivity*, RFC 2260, Jan. 1998.

- [9] CAIDA.
<http://www.caida.org/analysis/routing/atypes/>.
- [10] X.-R. Cao, H.-X. Shen, R. Milito, and P. Wirth. Internet pricing with a game theoretical approach: Concepts and examples. *IEEE/ACM Transactions on Networking*, 10(2):208–216, 2003.
- [11] Cisco Inc. Sample configurations for load sharing with BGP in single and multihomed environments. Available at <http://www.cisco.com/warp/public/459/40.html>.
- [12] D. Clark. *Internet Economics*, chapter Internet Cost Allocation and Pricing, pages 216–252. MIT Press, 1997.
- [13] R. Dai, D. O. Stahl, and A. B. Whinston. The economics of smart routing and QoS. In *Proceedings of the Fifth International Workshop on Networked Group Communications (NGC'03)*, 2003.
- [14] F5 Networks, Inc. <http://www.f5networks.com/>.
- [15] D. Goldenberg, L. Qiu, H. Xie, Y. R. Yang, and Y. Zhang. Optimizing cost and performance for multihoming. In *Proceedings of ACM SIGCOMM '04*, Portland, Oregon, Aug. 2004.
- [16] F. Guo, J. Chen, W. Li, and T. Chiueh. Experiences in building a multihoming load balancing system. In *Proceedings of IEEE INFOCOM '04*, Hong Kong, China, Apr. 2004.
- [17] A. Gupta, D. Stahl, and A. Whinston. *Internet Economics*, chapter Priority Pricing of Integrated Services Networks, pages 323–352. MIT Press, 1997.
- [18] Internap Networks, Inc. <http://www.internap.com>.
- [19] F. Kelly. *Internet Economics*, chapter Charging and Accounting for Bursty Connections, pages 253–278. MIT Press, 1997.
- [20] P. Marbach. Pricing differentiated services networks: Bursty traffic. In *Proceedings of IEEE INFOCOM '01*, Anchorage, AK, Apr. 2001.
- [21] A. Mihailovic, G. Leijonhufvud, and T. Suihko. Providing multi-homing support in IP access networks. In *PIMRC 2002*, 2002.

- [22] A. Odlyzko. A modest proposal for preventing Internet congestion. <http://www.research.att.com/~amo>, 1997.
- [23] A. Orda and R. Rom. Multihoming in computer networks: A topology-design approach. *Computer Networks and ISDN Systems*, 18(2):133–141, 1989.
- [24] Pacific Bell. <https://ebiznet.sbc.com/calnetinfo/RiderC>.
- [25] Radware, Inc. <http://www.radware.com/content/products/pd/default.asp>.
- [26] RouteScience Technologies, Inc. <http://www.routescience.com>, June 2003.
- [27] RouteScience Technologies, Inc. Route optimization for ebusiness applications. White Paper. Available at <http://www.routescience.com>, 2003.
- [28] G. Schreck, C. Rustein, and M. Porth. The end of the private WAN. Forrester Brief, Mar. 2002.
- [29] P. Sevcik and J. Bartlett. Improving user experience with route control. Technical Report NetForecast Report 5062, NetForecast, Inc., 2002.
- [30] S. Shakkottai and R. Srikant. Economics of network pricing with multiple ISPs. In *Proceedings of IEEE INFOCOM '05*, Miami, FL, Mar. 2005.
- [31] S. Sharma, J. Chen, W. Li, K. Gopalan, and T. Chiueh. Duplex: A reusable fault tolerance extension framework for network access devices. In *Proceedings of 2003 International Conference on Dependable Systems and Networks (DSN 2003)*, June 2003.
- [32] S. Shenker, D. Clark, D. Estrin, and S. Herzog. Pricing in computer networks: Reshaping the research agenda. *Communications Policy*, 20(1), 1996.
- [33] P. Smith. BGP multihoming techniques. NANOG 23. <http://www.nanog.org/mtg-0110/smith.html>, Oct. 2001.

- a_k : the unit price of ISP k .

- $a_{min}(Z_i)$: the minimum unit price of all ISPs in set Z_i , i.e., $a_{min}(Z_i) = \min\{a_{k'} | k' \in Z_i\}$.
- b_k : the base price of ISP k .
- c_k : the cost function of ISP k . We assume that c_k is a linear non-decreasing function.
- $c(S)$: the minimum total cost of subscribing to all ISPs in set S (single-homing when $|S| = 1$).
- \mathcal{F} : the set of all feasible sets of ISPs which the subscriber obtains by applying our subscription algorithm. A feasible set is a set of ISPs that the subscriber pays minimum cost to deploy multihoming over all possible sets in \mathcal{U} . There may exist multiple feasible sets.

- I : the total number of intervals in a charging period.
- \mathcal{K} : the set of all ISPs, *i.e.*, $\mathcal{K} = \{\infty, \dots, \mathcal{K}\}$, where K is the total number of ISPs. We use k as the index.
- m_k : $m_k = \frac{1}{\lambda_k}$ is the mean time to failure (MTTF) of ISP k .
- N : $N = |\mathcal{F}|$ is the number of feasible sets.
- N_k : the number of feasible sets in \mathcal{F} that contain ISP k .
- N_{-k} : the number of sets in \mathcal{F} that do not contain ISP k . $N_{-k} = N - N_k$.
- n_k : the number of feasible sets containing ISP k and having maximum unit price over all N_k sets which contain ISP k . In other words, $n_k = |\{Z_i | a_{\min}(Z_i) = \max\{a_{\min}(Z_j) | j = 1, \dots, N\}\}|$.
- $n_a(Z_i)$: the number of ISPs in set Z_i that have the same unit price as a . $n_a(Z_i) = |\{a_k | a_k = a, k \in Z_i\}|$.
- p_k : the charging volume of ISP k , (*i.e.*, $p_k = \text{qt}(T_k, q_k)$). For example, if ISP k charges at 95th-percentile, then p_k is the 95th-percentile of the traffic assigned to ISP k .
- $p_k(S)$: the charging volume of ISP $k \in S$ when the subscriber subscribes to all ISPs in the feasible set S .
- $p(S)$: the aggregated charging volume of ISPs in set S . $p(S) = V_0(S)$.
- q_k : the charging percentile of ISP k , *e.g.*, $q_k = 0.95$ if an ISP charges at 95th-percentile.
- $\text{qt}(X, q)$: the $\lceil q * |X| \rceil$ -th value in X_{sorted} (or 0 if $q \leq 0$), where X_{sorted} is X sorted in non-decreasing order, and $|X|$ is the number of elements in X .
- R_k : the total expected revenue of ISP k .
- $t_k^{[i]}$: the volume of traffic distributed to ISP k during interval i . Let time series $T_k = \{t_k^{[i]} | 1 \leq i \leq I\}$. Note that $V = \sum_k T_k$ (with vector summation).

- \mathcal{U} : the set of all subsets of \mathcal{K} . \mathcal{U}_k denotes the set of all subsets of \mathcal{K} containing k .
- $U(Z)$: the subscriber's utility function on a set Z of ISPs. $U(Z) = w \sum_{k \in Z} \log m_k - \sum_{k \in Z} c_k$. The subscriber's objective in ISP pricing problem is $\max_{Z \in \mathcal{U}} U(Z)$.
- $v^{[i]}$: the total traffic volume during interval i .
- V : time series of traffic volumes $V = \{v^{[i]} \mid 1 \leq i \leq I\}$.
- $V_0(S)$: $V_0(S) \stackrel{\text{def}}{=} \text{qt}(V, 1 - \sum_{k \in S} z_k)$, where $S \subseteq \mathcal{K}$ is a subset of ISPs,
- w : the weight of the subscriber's preference of reliability over cost.
- Z_i : the enumeration of all the feasible sets in \mathcal{F} . Here $i \in \{1, \dots, N\}$.
- z_k : $z_k \stackrel{\text{def}}{=} 1 - q_k$.
- λ_k : the instantaneous failure rate of ISP k . We consider constant failure rate in this paper.