

# Interplay of ISPs: Distributed Resource Allocation and Revenue Maximization

Sam C.M. Lee<sup>†</sup>   Joe W.J. Jiang<sup>†</sup>   Dah-Ming Chiu\*   John C. S. Lui<sup>†</sup>

## Abstract

*The Internet is a hierarchical architecture comprising heterogeneous entities of privately owned infrastructures, where higher level Internet service providers (ISPs) supply connectivity to the local ISPs and charge the local ISPs for the transit services. One of the challenging problems facing service providers today is how to increase the profitability while maintaining good service qualities. In this work, we seek to understand the fundamental issues on the “interplay” (or interaction) between ISPs at different tiers. While the local ISPs (which we term peers) can communicate with each other by purchasing the connectivity from transit ISPs, there stands an opportunity for them to set up private peering relationships. Under this competitive framework, we explore the issues on (a) impact of peering relationship, (b) resource distribution and (c) revenue maximization. Firstly, a generalized model is presented to characterize the behaviors of peers and the transit ISP, in which their economic interests are reflected. We study how a peer can distributively determine its optimal peering strategy. Furthermore, we show how a transit ISP is able to utilize the available information to infer its optimal pricing strategy, under which a revenue maximization is achieved. A distributed algorithm is proposed to help ISPs to provide a fair and efficient bandwidth allocation to peers, avoiding a resource monopolization of the market. Extensive simulations are carried out to support our claims.*

**Keywords:** ISP peering, economic pricing, distributed resource allocation

## 1 Introduction

One of the challenging problems facing today’s Internet Service Providers (ISPs) is how to increase the profitability and at the same time, provide good performance

to users as we scale up the network. For the Internet, it is a hierarchical architecture comprising heterogeneous entities of privately owned infrastructures. Generally speaking, the networks can be categorized into two types of service providers: (1) local ISPs which consist of geographically close meshed networks. These local ISPs provide Internet access and connectivity services for consumers within their regions. And (2) large-scaled ISPs which traverse across large geographical distances, providing connectivity among the local ISPs.

For the local ISPs, in order to gain the Internet access, a common way is to purchase this service from higher level ISPs (or we called transit ISPs). These transit ISPs set charges for the service provisioning, which depend on the allocated transmission bandwidth as well as the amount of transferred traffic. One important issue is to come up with a good pricing model for the current Internet, especially to reflect the economic roles of different ISPs. Currently, most ISPs adopt a flat rate pricing scheme, i.e., end users pay a fixed amount of money to gain the Internet access in a certain period of time (usually on a monthly basis). Most broadband and ADSL services are examples of this type. Another approach is to charge users by the time they connect to the Internet, following the charging methods employed in the telephony industry. Still, there are ISPs who charge users based on the actual traffic volume transmitted. There are some existing work which investigate the pricing strategy for the service providers. In [4], authors discuss how a provider should price its services differentially based on their characteristics such that prices can match service qualities. Authors in [3] also discuss how to present a co-operative pricing strategy to provide a fair distribution of profits to ISPs.

Besides relying on the transit ISPs for Internet access, local ISPs can also inter-connect their networks together by signing up private peering agreements. For local ISPs which are geographically close to each other, there is an opportunity to exchange information between themselves and bypass the reliance on transit ISPs. One possible way to accomplish this is to establish a *private peering link* between two parties. In practice, these peering agreements can be quite complicated, involving many business considerations [15][5]. However, the basic nature of the *peering relation-*

\*Department of Information Engineering, The Chinese University of Hong Kong dmchiu@ie.cuhk.edu.hk.

<sup>†</sup>Department of Computer Science & Engineering, The Chinese University of Hong Kong, Shatin, N.T. Hong Kong; {cmlee, wjjiang, cslui}@cse.cuhk.edu.hk.

*ship* is to exchange local traffic between the two local ISPs through the peering link without paying for the traffic transfer. Note that *free* peering is only one special case of the peering relationship, having charges on peers are also considered in more generalized circumstances. Usually such peering relationship is beneficial to both ISPs since it can provide better performance and at the same time, reduce the operating cost since traffic does not need to go through the transit ISPs.

There are a number of existing work which explore the economics of network pricing with multiple ISPs on the Internet, recent work being [1, 2, 7, 13]. These authors all investigate a basic question: How to set prices for the Internet services, so as to fairly share revenues among providers, and at the same time encourage the network to grow? Unfortunately, these work underestimate the impact of local peering relationship on the traffic demand, since this will influence the proper pricing strategy to achieve a maximization of ISP's profitability. To bridge this gap, our work aims to seek a fundamental understanding of the interaction between ISPs with peering links. We explore how the peering relationship can affect the service purchasing strategies and pricing strategies played by ISPs. For the ease of presentation, in the rest of this paper, we term the local ISPs simply as *peers* since they tend to establish peering relationships with each other. Similarly, we refer to the transit ISPs as *ISPs*. In this work, we are interested to explore the interactions between the connecting peers, as well as the interactions between the peers and the ISP. We also investigate the implications of these interactions.

To communicate with another local ISP, a peer has two options: either to use the connection provided by the transit ISPs, or to use the peering link connecting the two peers. Even given a constant transmission demand, deciding on an appropriate proportion of traffic delivered via these two connections, however, is not a trivial matter. Another factor which makes the decision of traffic allocation difficult is that all peers are *rational*, i.e., they want to maximize their *happiness* by transmitting/receiving traffic, and at the same time, they also care about the quality of service they receive and the total payments for consuming the services. Also, one peer's optimal strategy may depend on strategies taken by other peers, as well as the pricing policy employed by the ISP. All these make it a challenging task to come up with an efficient resource allocation policy.

ISP, on the other hand, provides Internet access and connectivity between peers. Its goal is to maximize its own revenue by providing connectivity service. In order to maximize the total profit and attract more potential peers, a good pricing strategy is essential. In general, a transit ISP needs to address the following issues:

- *Resource Distribution*: how should the ISP sell and allocate its capacity resource to the competing peers, and

at the same time, avoid the monopolization of bandwidth resources by a small number of peers?

- *Maximization of Revenue*: does it exist a unique price by which the ISP's revenue can be maximized under a homogeneous pricing scheme (i.e., all peers are charged using the same pricing model)? If it exists, how can one find this optimal price?

While these questions have substantial impact and important implications, it is not straightforward to obtain an immediate answer. From the ISP's perspective, it is undesirable that its resource is *utilized* (or monopolized) by a small number of peers since the ISP wants to achieve *customer diversification*. To attract or retain a peer for the connectivity service, an ISP has to perform a "*fair*" resource distribution which *avoids* resource monopolization. To achieve this goal, the ISP and peers have to exchange traffic information. Note that one has to consider a minimal information exchange due to business confidentiality, as well as the necessity to perform resource allocation in a distributed manner. On the other hand, maximizing its own profit is also an important objective for the ISP. With a particular price offered by the ISP, every connecting peer *decides* the amount of traffic to send through the transit ISP. The aggregate traffic thus determines the total demand on the ISP link. Setting a lower price attracts more traffic from the peers, but this may lead to traffic congestion. Moreover, a low price does not guarantee the maximization of the ISP's revenue. Setting a higher price, on the other hand, may discourage peers to purchase the ISP service and the traffic demand will decrease, which does not ensure a maximal profit for the ISP as well. Therefore, finding an *optimal* unit price is an important issue.

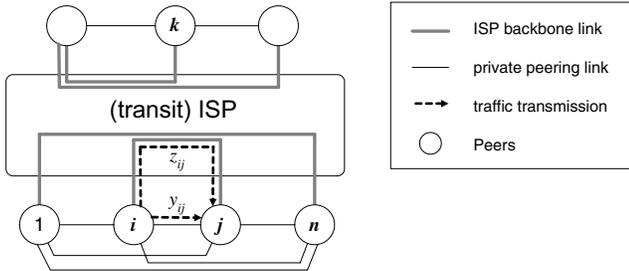
The contribution of our paper is to answer the questions listed above. We explore the *interplay* or *interaction* between ISPs at different tiers, discussing issues on (a) *resource distribution*, and (b) *revenue maximization*.

- We present a generalized model to capture a snapshot of the current Internet, a hierarchy consisting of ISPs of two tiers with peering relationships. We believe this two-tier-interaction represents a basic element of the complicated system, characterizing ISPs' behaviors beyond which their economic interests are reflected.
- We study how a peer can *distributively* determine its optimal peering strategy by solving a convex optimization problem.
- We propose a distributed algorithm, namely *Equal Share Algorithm* (ESA), to help ISP provide an efficient and fair bandwidth allocation to peers. We show that to avoid monopolization of the market is not a trivial issue.

- We further explore how a transit ISP is able to utilize its available information to *infer* an optimal pricing strategy under which its revenue maximization can be achieved.

The organization of the paper is as follows. In Section 2, we present our mathematical model and formulate the objective functions for a peer and the ISP. In Section 3, we show the conditions for a peer to obtain the maximum utility. In Section 4, we propose an algorithm for the ISP to distribute its resource among the peers. We carry out simulations to examine the performance of the algorithm. In Section 5, we propose a methodology on how an ISP can estimate its optimal pricing strategy to maximize its revenue. Section 6 presents the related work and Section 7 concludes.

## 2 The Mathematical Model



**Figure 1.** A model of  $n$  peers and one ISP. Each peer has one aggregate link to the ISP and possibly  $n - 1$  private links to other peers. Peer  $i$  can communicate with peer  $j$  in two possible ways: one through the peering link  $l_{ij}$  and one through the ISP link  $l_{ii}$ . The traffic rate on link  $l_{ij}$  is  $y_{ij}$  while the traffic rate on link  $l_{ii}$  is  $z_{ij}$ .

Consider a network which is depicted in Figure 1. The network consists of  $n$  peers and one Internet Service Provider (ISP), where a peer can be viewed as a local ISP<sup>1</sup> and the higher level ISP is to provide connection between these peers. Peers need to communicate with each other by sending data. They can communicate with each other either by sending traffic through the ISP, or by the private peering links between themselves. In order to provide connectivity, the ISP has a communication network (in which we abstract it as a link) that has a *total* capacity of  $n\mathcal{C}$  (in units of bps). For each peer  $i \in \{1, 2, \dots, n\}$ , it possesses an aggregate link to the ISP and possibly  $n - 1$  private peering links connecting to the other  $n - 1$  peers. Since the peering links

<sup>1</sup>Unless we state otherwise, we will use the term peer to denote a local ISP while the term ISP to denote higher level ISP, such as tier-1 ISP.

are privately owned infrastructures by two parties, we use the terms “*private links*” and “*peering links*” interchangeably in the rest of the paper. Let  $l_{ij}$  denote the peering link between peer  $i$  and peer  $j$  and this link has a capacity of  $c_{ij}$  (in unit of bps). Note that if we set  $c_{ij} = 0$ , it implies that there is no peering link between peer  $i$  and peer  $j$ . The link connecting peer  $i$  and the ISP is denoted as  $l_{ii}$ , and the ISP allocates  $\mathcal{C}_i$  amount of bandwidth (in units of bps) for this connection. Note that our model can be viewed as a generalization of the network model in [2], in which private peering links are not considered.

Let  $x_{ij}$  denote the transmission rate (in unit of bps) from peer  $i$  to peer  $j$ . In short, it is the traffic originated from peer  $i$  destined to peer  $j$ . To sustain the transmission rate of  $x_{ij}$ , peer  $i$  obtains a utility of  $A_{ij}(x_{ij})$  where  $A_{ij}$  is a strictly concave function in  $x_{ij}$ . As noted in [6], concave function is commonly used to represent elastic traffic, which is the dominant traffic in the Internet. The utility  $A_{ij}(x_{ij})$  represents the happiness of peer  $i$  by sending data to peer  $j$  at a rate of  $x_{ij}$ . In this paper, we use a weighted log function as our utility function and  $A_{ij}(x_{ij}) = w_{ij} \log(1 + x_{ij})$ . The weighting  $w_{ij}$  can be interpreted as the *happiness weighting coefficient* of transmitting traffic between peer  $i$  and peer  $j$ . Therefore, it is possible for  $w_{ij} > w_{ik}$ , which represents that peer  $i$  prefers to communicate with peer  $j$  than peer  $k$ . Note that the log function is chosen as it leads to a proportionally fair resource allocation if proper congestion control is used. Additionally, this type of utility function is also commonly used for performing distributed admission control[2].

The traffic transmission rate  $x_{ij}$ , which has to be computed later, can either go through the ISP link  $l_{ii}$ , or the private link  $l_{ij}$ . We denote  $y_{ij}$  as the traffic rate that peer  $i$  decides to transmit through the private link  $l_{ij}$ , and  $z_{ij}$  as the traffic rate through ISP link  $l_{ii}$ . In other words, the traffic transmission rate  $x_{ij}$  is equal to

$$x_{ij} = y_{ij} + z_{ij} \quad \text{for } i, j \in \{1, \dots, n\}.$$

A particular case to note is the traffic rate  $x_{ii}$ , which denotes the traffic rate from peer  $i$  to destinations other than the  $n - 1$  peers. This type of traffic can represent data to other part of the Internet wherein peer  $i$  has to send the data through the ISP. Since there is no established private link to those outsiders, peer  $i$  can only rely on the ISP link for the traffic transmission. Therefore,

$$y_{ii} = 0 \quad \text{and} \quad z_{ii} = x_{ii} \quad \text{for } i \in \{1, \dots, n\}.$$

For the ease of presentation, let  $z_i = \sum_{j=1}^n z_{ij}$  denote the aggregate traffic rate that peer  $i$  sends through the ISP link, and let  $\bar{z} = \sum_{j=1}^n z_j$  denote the aggregate traffic on the ISP link from all  $n$  peers.

To transmit data across the ISP, peers need to pay the network operators for the transmission service. The price

per unit bandwidth through the ISP link  $l_{ii}$  is  $\mathcal{P}_i$ , which is determined by the ISP. Peer  $i$  can also send the traffic  $y_{ij}$  through the private link  $l_{ij}$ , and the price per unit bandwidth is  $p_{ij}$ , which can be mutually agreed upon between peers  $i$  and  $j$ . In this work, we do not consider the issues on the cost of setting up peering links, since it is not part of the operating cost. We assume peers can utilize existing peer links with fixed capacities  $c_{ij}$ 's. For convenience, we denote  $\vec{y}_i = (y_{i1}, y_{i2}, \dots, y_{in})$  as the traffic rate vector for peer  $i$  through its private links and  $\vec{z}_i = (z_{i1}, z_{i2}, \dots, z_{in})$  as the traffic rate vector for peer  $i$  through the ISP link. We denote  $\vec{\mathcal{P}} = (\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_n)$  as the vector of ISP prices set on different peers.

Besides paying the ISP for the transmission service, each peer also needs to take into consideration of the congestion costs on the links. If we assume that all links can be represented by an M/M/1 model as in [2], one can take the delay on the link as its congestion indication or cost. Rather than informing all peers about the current transmission rate  $\bar{z}$  on the ISP link (which can be considered as confidential information by a peer), ISP will do a pre-computation and announce its bandwidth allocation to peer  $i$  as  $\mathcal{C}_i$ . There is also a technical merit for this announcement which will be discussed in detail in later sections. Under this form of setting, the congestion cost  $D_{ij}$  of a link  $l_{ij}$ , is

$$D_{ij} = \begin{cases} \frac{1}{c_{ij} - y_{ij}} & \text{if } i \neq j, \\ \frac{1}{\mathcal{C}_i - z_i} & \text{if } i = j. \end{cases}$$

To model the economic incentives and behaviors of all peers, we consider the following optimization. The objective of peer  $i$  is to maximize the following function:

$$\begin{aligned} \text{Max } U_i &= \sum_j w_{ij} \log(1 + y_{ij} + z_{ij}) - \mathbf{1}_{\{z_i \neq 0\}} \left[ \frac{1}{\mathcal{C}_i - z_i} \right] - \mathcal{P}_i z_i \\ &\quad - \sum_{j \neq i} \mathbf{1}_{\{y_{ij} \neq 0\}} \left[ \frac{1}{c_{ij} - y_{ij}} \right] - \sum_{j \neq i} p_{ij} y_{ij} \quad (1) \\ \text{s. t.} \quad &0 \leq y_{ij} \leq c_{ij} \text{ for all } j \neq i, \\ &y_{ii} = 0, \quad \sum_j z_{ij} \leq \mathcal{C}_i, \quad z_{ij} \geq 0 \text{ for all } j. \quad (2) \end{aligned}$$

where  $\mathbf{1}_{\{p\}}$  is an indicator function which equals to 1 if the condition  $p$  is true, or 0 otherwise. The objective function of Equation (1) represents the economic incentive for peer  $i$  to perform traffic transmission. In here,  $w_{ij} \log(1 + y_{ij} + z_{ij})$  is the happiness of peer  $i$  by sending traffic to peer  $j$ . The term  $\frac{1}{\mathcal{C}_i - z_i}$  is the congestion cost of peer  $i$  on the ISP link. However, if peer  $i$  does not transmit through the ISP link, it does not bear the delay load and the congestion cost will be zero. The term  $\mathcal{P}_i z_i$  is the total payment of peer  $i$  to the ISP. Similarly,  $\frac{1}{c_{ij} - y_{ij}}$  is the congestion cost on the peering link connecting peer  $i$  to peer  $j$ , when the transmission rate on

$l_{ij}$  is non-zero. Lastly, peer  $i$  has a payment of  $p_{ij} y_{ij}$  to peer  $j$  for using the private link<sup>2</sup>. Note that if these parameters, e.g.,  $w_{ij}, \mathcal{C}_i, \mathcal{P}_i, c_{ij}, p_{ij}$  are set properly, peer  $i$ 's happiness, congestion cost and payment can be measured in monetary units.

Meanwhile, constraints represented in Equation (2) define the feasible region of this optimization problem. The first are the non-negative and capacity constraints of the peering links. The second constraint is due to the absence of peering links established to the "outsiders". The third and fourth are the capacity and non-negative constraints of the ISP link respectively. In summary, each peer  $i$  needs to determine the proper traffic rates vectors  $\vec{z}_i$  and  $\vec{y}_i$  so as to maximize its aggregate utility in Equation (1).

Note that the optimization processes of different peers are *not* independent. For each peer  $i$ , given the bandwidth allocation  $\mathcal{C}_i$  of the ISP link, it performs an optimization and determines its optimal transmission rate  $z_i$  and bids to the ISP. After collecting the bidding information from all peers, the ISP will calculate the new bandwidth allocation according to the new biddings. Therefore, the interaction process between peers can be modelled as a *non-cooperative game* such that each peer offers a bid to maximize its own utility.

Under this framework, for a given ISP price vector  $\vec{\mathcal{P}} = (\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_n)$ , this defines a non-cooperative game between these  $n$  peers [12]. They interact with each other and determine their optimal transmission rates periodically and asynchronously. Given the existence of an equilibrium point, the operating point for  $n$  peers is the solution to the Nash equilibrium of this game. For each price vector  $\vec{\mathcal{P}} > 0$ , a Nash equilibrium point for this  $n$ -peers game is defined as two  $n$ -tuples  $y^* = (\vec{y}_1^*, \vec{y}_2^*, \dots, \vec{y}_n^*)$  and  $z^* = (\vec{z}_1^*, \vec{z}_2^*, \dots, \vec{z}_n^*)$ , such that for all peers  $i \in \{1, 2, \dots, n\}$ :

$$U_i(y^*, z^*, \vec{\mathcal{P}}) \geq U_i(y, z, \vec{\mathcal{P}}) \quad (3)$$

for any other feasible traffic vector  $y = (\vec{y}_1, \vec{y}_2, \dots, \vec{y}_n)$  and  $z = (\vec{z}_1, \vec{z}_2, \dots, \vec{z}_n)$  that satisfies the constraints defined in Equation (2).

On the other hand, the ISP is associated with a revenue maximization problem:

$$\text{Maximize } \mathcal{P} \cdot \bar{z}^*(\mathcal{P}) \quad \text{over } \mathcal{P} \geq 0 \quad (4)$$

where  $\bar{z}^*(\mathcal{P}) = \sum_j z_j^*(\mathcal{P})$  is the aggregate traffic on ISP link at the Nash equilibrium. In here, note that we assume the ISP charges the same price for all peers and there is no price discriminate. Therefore  $\mathcal{P}_i = \mathcal{P}$  for all  $i$ . This equivalently defines a Stackelberg game [12] with one leader (ISP) and the non-cooperative Nash followers ( $n$  peers). The ISP has a first move advantage to determine the optimal price such that its own revenue can be maximized.

<sup>2</sup>It is also possible for us to model the case that peer  $i$  and  $j$  do not charge each other for sending peering traffic, i.e.,  $p_{ij} = p_{ji} = 0$ .

### 3 Solution to Maximization Problem by Individual Peers

In this section, we illustrate how a peer, say  $i$ , can determine its transmission rates, which is  $z_i$ , to other peers via the ISP, as well as the transmission rate  $\bar{y}_i$ , rates to other peers via peering links, so as to maximize its utility. Assuming that the peer knows the price  $\mathcal{P}$  specified by the ISP and the associated bandwidth allocation  $\mathcal{C}_i$ , one can model an individual peer's behavior as a convex optimization problem as defined in Equation (1). In this section, we investigate the necessary and boundary conditions for a peer to maximize its utility.

#### 3.1 Necessary conditions with positive transmission rate

Since  $U_i$  is discontinuous at  $y_{ij} = 0$  (i.e., the traffic rate through the peering link  $l_{ij}$  is zero) and  $z_i = 0$  (i.e., the traffic rate through the ISP's link  $l_{ii}$  is zero), we first investigate the necessary conditions when  $y_{ij} \neq 0$  and  $z_i \neq 0$ . The optimization problem of Equation (1) has  $2n - 1$  variables (with  $y_{ii} = 0$ ). We first write down the second order partial derivatives with respect to  $y_{ij}$  and  $z_{ij}$ :

$$\begin{aligned}\frac{\partial^2 U_i}{\partial y_{ij}^2} &= \frac{-w_{ij}}{(1 + y_{ij} + z_{ij})^2} - \frac{2}{(c_{ij} - y_{ij})^3} < 0, \\ \frac{\partial^2 U_i}{\partial z_{ij}^2} &= \frac{-w_{ij}}{(1 + y_{ij} + z_{ij})^2} - \frac{2}{(\mathcal{C}_i - z_i)^3} < 0, \\ \frac{\partial^2 U_i}{\partial y_{ij} \partial z_{ij}} &= \frac{-w_{ij}}{(1 + y_{ij} + z_{ij})^2} < 0.\end{aligned}$$

And for  $k \neq i \neq j$ , the second order partial derivatives of Equation (1) with respect to  $y_{ik}$  and  $z_{ik}$  are:

$$\frac{\partial^2 U_i}{\partial y_{ij} \partial y_{ik}} = 0, \quad \frac{\partial^2 U_i}{\partial y_{ij} \partial z_{ik}} = 0, \quad \frac{\partial^2 U_i}{\partial z_{ij} \partial z_{ik}} = \frac{-2}{(\mathcal{C}_i - z_i)^3} < 0. \quad (5)$$

Therefore, the Hessian matrix of the objective function in Equation (1) is negative definite on the non-negative orthant bounded by  $y_{ij} \leq c_{ij}$  and  $z_i \leq \mathcal{C}_i$ . So  $U_i$  is strictly concave in  $y_{ij}$  and  $z_{ij}$  for all  $j$ . The maximum utility and optimizer to this problem is unique and can be found by the Lagrangian method. The necessary conditions of  $y_{ij}$  and  $z_{ij}$  for the maximization of  $U_i$  are:

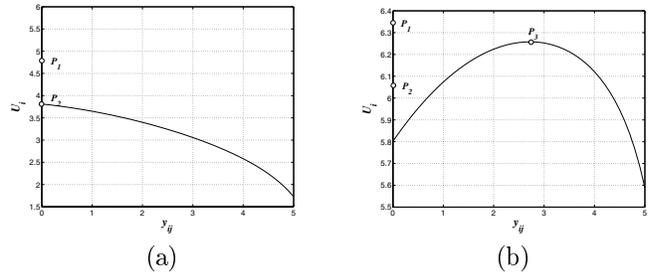
$$\frac{\partial U_i}{\partial y_{ij}} \begin{cases} < 0 & \text{if } y_{ij} = 0 \\ = 0 & \text{if } y_{ij} > 0 \end{cases}, \quad \frac{\partial U_i}{\partial z_{ij}} \begin{cases} < 0 & \text{if } z_{ij} = 0 \\ = 0 & \text{if } z_{ij} > 0. \end{cases} \quad (6)$$

#### 3.2 Boundary cases

Due to the discontinuity of the objective function, the necessary conditions given above may not achieve the global maximum. In here, we are going to explore the

boundary cases when the transmission rate tends to be zero, i.e.,  $y_{ij} = 0$  or  $z_i = 0$ . Figure 2 shows the illustration of an example. We plot the utility of peer  $i$  against one particular variable  $y_{ij}$  ( $z_{ij}$  is similar). Figure 2(a) corresponds to the case when  $\frac{\partial U_i}{\partial y_{ij}}|_{y_{ij}=0} < 0$ . The optimizer is  $y_{ij} = 0$ , but the maximum utility is at point  $P_1$  rather than point  $P_2$  since there is no congestion cost at the private link when  $y_{ij} = 0$ . Figure 2(b) corresponds to the case when  $y_{ij}^* = \arg\{\frac{\partial U_i}{\partial y_{ij}} = 0\}$  is positive ( $P_3$  in the figure). If the utility  $U_i$  at the boundary point  $P_2$  which is less than  $P_3$ ,  $P_3$  is the maximum point and  $y_{ij}^*$  is the optimizer. However, there exists a case when the utility  $U_i$  at the boundary point  $P_1$  is greater than that of point  $P_3$ . Therefore,  $P_1$  should be the maximum point and  $y_{ij} = 0$  is the optimizer.

Here we provide the physical interpretation of the two cases illustrated in Figure 2(b). If the utility  $U_i$  at the boundary point is  $P_2$ , it indicates that when the transmission rate  $y_{ij}$  increases, the increase in happiness outweighs the increases in congestion cost and its total payment, thus achieving the maximum utility at point  $P_3$ . However, if the utility  $U_i$  at the boundary point is  $P_1$ , it means that when the transmission rate increases, the increase in happiness cannot compensate for the increases in congestion cost and its payments. That is, although peer  $i$  achieves the maximum utility at  $P_3$ , the utility is negative. So the best strategy for peer  $i$  is not to transmit data through  $l_{ij}$  at all. Note that when a peer  $i$  does not send through any links, it gets a zero happiness, zero congestion cost and zero payment, and thus a zero utility. Therefore, a peer will always achieve a non-negative utility, since in the worse case, it can opt not to transmit and leave the network (or market).



**Figure 2.** Utility of peer  $i$  against one particular  $y_{ij}$  (a) negative optimizer (b) positive optimizer.

### 4 Distributed Resource Allocation by ISP

From an ISP's point of view, a monopolized use of its link bandwidth surely reduces its customer size and so increases the risk of the business. Moreover, in order to max-

imize its revenue, an ISP has to know approximately the demand of its link bandwidth. Therefore, an ISP wants to have an efficient resource allocation algorithm. Now, given the total amount of resource  $n\mathcal{C}$  (ISP's link bandwidth), the ISP needs to determine how to distribute this common resource to all the  $n$  peers. In this section, we propose a resource allocation algorithm that can be adopted by the ISP.

Before we proceed to the formal presentation of the algorithm, let us illustrate the general framework under which the ISP can interact with peers so that the ISP is able to discover the actual resource demands from peers, and also peers are informed about the pricing information and the available resources. Initially, the ISP equally distributes its capacity to all peers at time  $t = 0$ . Each peer  $i$  calculates its optimal traffic transmission rates based on the currently allocated ISP capacity  $C_i$ , as well as the ISP link price  $\mathcal{P}$ . Then the peer reports its transmission rate (resource usage),  $z_i$ , back to the ISP. We refer to the feedback information  $z_i$  as the *bidding* of peer  $i$ . The ISP receives the biddings from peers within a period of time  $T$ . At the end of each period, ISP recomputes the link resource distribution and sends the new bandwidth allocation  $C_i$  to peer  $i$ , where  $i \in \{1, \dots, n\}$ . Based on the new bandwidth allocation, peers calculate their optimal transmission rates again and the process repeats.

There are two advantages for this framework. First, all the information that a peer  $i$  requires are the unit prices  $p_{ij}$ 's and capacities  $c_{ij}$ 's of its private links and the allocated link capacity  $C_i$ , as well as the informed price  $\mathcal{P}$ . These can be seen as the private information of peer  $i$ . Peer  $i$  does not have to know the bandwidth allocation  $\{C_1, \dots, C_n\}$  and transmission rates  $\{z_1, \dots, z_n\}$  of all other peers, which is considered as confidential information. On the other hand, when the ISP makes the bandwidth allocation, what it requires to know are the biddings  $(z_1, z_2, \dots, z_n)$  from all peers. The ISP may not know the utility functions and the information about private links of these peers (i.e.,  $p_{ij}$  and  $c_{ij}$  for all  $i, j$ ). Secondly, the overhead of information exchange in this framework is small. ISP only needs to inform each peer its allocated capacity, while each peer only needs to reply to the ISP its bidding.

In the following, we present a resource allocation algorithm by which the ISP can determine the appropriate capacity  $C_i$  for all peers  $i, i = 1, \dots, n$ .

#### 4.1 Equal share algorithm (ESA)

Under the Equal Share Algorithm, ISP distributes its *remaining* capacity equally among all peers after satisfying their bandwidth consumption demands indicated by their biddings. Initially, ISP distributes its capacity equally to every peer, i.e.,  $C_i = \mathcal{C}$  for all  $i$ , and sends the capacity distribution  $C_i$  to every peer  $i$ . Upon receiving the informa-

tion from the ISP, each peer uses the algorithm proposed in Section 3 to find its optimal transfer rates (i.e.,  $z_i$  for peer  $i$ ) and sends the information back to the ISP as its resource bidding. Within the following period, ISP gathers all the feedbacks from peers. ISP first allocates to each peer the capacity it bids, and then the ISP distributes the remaining resource equally to the peers. Formally, we have:

$$C_i = z_i + \frac{(n\mathcal{C} - \bar{z})}{n}.$$

The algorithm of the ESA is described as follows:

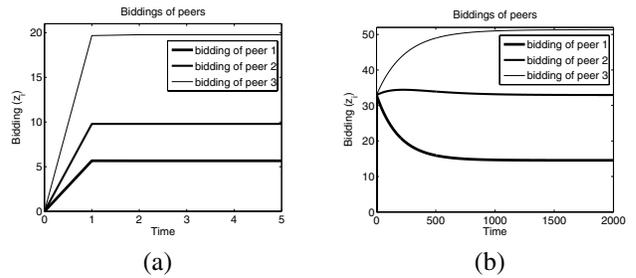
---

#### Equal Share Algorithm (ESA):

1. ISP initiates  $C_i^{(0)} := \frac{n\mathcal{C}}{n} := \mathcal{C}$  to each peer  $i$ .
  - Set counter  $k := 0$ .
  2. **while** (TRUE) {
  3.   ISP passes  $C_i^{(k)}$  to each peer  $i$ ;
  4.   **for** ( $i=1$  to  $n$ ) {
  5.     Peer  $i$  computes  $y_i^{(k)}$  and  $z_i^{(k)}$  ;
  - and sends  $z_i^{(k)} = \sum_j z_{ij}^{(k)}$  back to ISP;
  6.   } /\* termination of for-loop \*/
  7.   ISP updates  $C_i^{(k+1)} = z_i^{(k)} + \frac{n\mathcal{C} - \bar{z}^{(k)}}{n}$  for every peer  $i$ ;
  - update counter  $k:=k+1$ ;
  8.   } /\* termination of while-loop \*/
- 

#### 4.2 Illustration of ISP Resource Allocation

To illustrate the effectiveness and performance of the above algorithm, we carry out two experiments and illustrate the resource distribution under two different scenarios, namely, (a) the ISP has sufficient capacity, and (b) the ISP has insufficient capacity.



**Figure 3.** Biddings of peers 1,2 and 3 (a) under sufficient bandwidth (b) under insufficient bandwidth.

The first experiment illustrates the case when ISP has sufficient resource. There are three peers in the network. Each peer has two private links to other peers in the system with capacity  $c_{ij} = 10$  and a unit price  $p_{ij} = 1$ . Peers 1,

2 and 3 have different values of happiness weighting coefficients,  $w_{1j} = 10$ ,  $w_{2j} = 15$  and  $w_{3j} = 20$  for  $j = 1, 2, 3$ . ISP provides a link with capacity  $n\mathcal{C} = 100$  and charges a unit price of  $\mathcal{P} = 1.5$ . The ISP updates the distribution with ESA and sends signals to peers every one second. Figure 3(a) shows the bidding of each peer during the experiment. The vertical axis shows the bidding of each peer and the horizontal axis shows the time. This experiment shows that peers offer the biddings efficiently under ESA. The biddings converge within only one single period.

The second experiment illustrates the case when ISP has insufficient resource. There are three peers in the network. Each peer has two private links to each other with capacity  $c_{ij} = 10$  and a unit price  $p_{ij} = 1$ . Peers 1, 2 and 3 have different values of happiness weighting coefficients,  $w_{1j} = 100$ ,  $w_{2j} = 150$  and  $w_{3j} = 200$  for  $j = 1, 2, 3$ . Note that in here the happiness weighting coefficients are much larger than the previous experiment, meaning that peers do have stronger desires to transmit traffic. Thus keeping the ISP link capacity at the same level leads to an insufficient resource supply. ISP provides a link with capacity  $n\mathcal{C} = 100$  and charges a unit price of  $\mathcal{P} = 1.5$ . The ISP updates the distribution under ESA and sends a signal to peers every one second. Figure 3(b) shows the bidding of each peer during the experiment. The vertical axis shows the bidding of each peer and the horizontal axis is the time axis. The experiment runs for 2000 seconds. We observe that the three peers share the limited ISP link and there is no monopolization.

**Remarks:** In this section, we show an algorithm for the ISP to distribute its capacity resource among the peers. ESA is efficient and prevents monopolized utilization of resource by a small number of peers.

## 5 ISP: Maximization of its Revenue

In this section, we investigate various approaches for an ISP to maximize its revenue. The revenue of an ISP is the aggregate payments received from  $n$  peers for consuming the ISP's link bandwidth. Formally, the ISP's revenue can be expressed as

$$\mathcal{R}(\mathcal{P}) = \sum_i \mathcal{P} \cdot z_i(\mathcal{P}) \quad (7)$$

In here,  $z_i(\mathcal{P})$  indicates that one peer's bandwidth consumption on the ISP link is a function of the price  $\mathcal{P}$  set by the ISP. It is obvious that, if the price is set too high, peers may switch their traffic to the private links where the service is cheaper, thus the ISP's revenue reduces. On the other hand, a lower price may attract peers to send more traffic via the ISP, however, too low a price may not ensure an increase in the total revenue. These characteristics leave the door open for ISP to search for an optimal price to ensure the maximization of its revenue. Normally, it only

makes sense for the ISP to obtain the optimal price in a "blind search" manner. Namely, the ISP randomly proposes a price to see the aggregate bandwidth consumption  $\bar{z}$  at the equilibrium point. After finding its revenue at this pricing level, the ISP may adjust its price a little bit to see how it affects its total revenue. With the feedback information the ISP can readjust its price. However, this type of "local search" method may not ensure the global optimality and it can be very time-consuming. So a natural question arises: is there any effective approach for an ISP to find its optimal price, assuming the ISP can estimate some necessary information about the system, e.g., the happiness weighting coefficients of peers ( $w_{ij}, \forall i, j$ ), capacities and unit prices of the private links?

Instead of doing a blind search on the unit price  $\mathcal{P}$  to maximize the revenue, we have an analysis on the estimation of the aggregate traffic  $\bar{z}$  given a fixed ISP link price  $\mathcal{P}$ , rather than implementing the price and waiting for an equilibrium point to reach. With the estimation of  $\bar{z}$ , we can easily calculate ISP's revenue  $\mathcal{R}(\mathcal{P})$  for a given  $\mathcal{P}$ .

This estimation has the following three assumptions. First, the ISP applies the Equal Share Algorithm (ESA) in resource distribution. Second, the ISP takes an indiscriminate pricing approach and charges the same unit price of  $\mathcal{P}$  to all peers. Third, when the ISP maximizes its revenue, it only considers the case of  $z_{ij} > 0$  for all  $i, j$ , which means all peers want to transfer data via the ISP's links.

### 5.1 Estimation of aggregate traffic $\bar{z}$ on ISP link

The estimation of aggregate traffic is divided into two cases:  $y_{ij} > 0$  and  $y_{ij} = 0$ . For the first case when  $y_{ij} > 0$  for all  $i \neq j$ , we obtain

$$\bar{\mathcal{C}}_0 + n^2 \approx \frac{\bar{W}}{k} + \frac{n}{\sqrt{k - \mathcal{P}}} + \frac{n(n-1)}{\sqrt{k - p_{av}}} \quad (8)$$

where  $\bar{\mathcal{C}}_0$  is the aggregate capacities in the system,  $p_{av}$  is the mean of  $p_{ij}$ ,  $\bar{W}$  is the sum of happiness weighting coefficients of all peers, and  $k$  is an artificial variable introduced.

For the second case when  $y_{ij} = 0$  for all  $i \neq j$ ,

$$n\mathcal{C} + n^2 = \frac{\bar{W}}{k} + \frac{n}{\sqrt{k - \mathcal{P}}} \quad (9)$$

Having some information of the happiness weighting coefficients and capacities and unit prices of private links, the ISP can estimate the value of  $k$  using Equations (8) or (9). Then ISP can estimate the aggregate biddings  $\bar{z}$  and its revenue  $\mathcal{R}(\mathcal{P})$  with Equation (7).

For more details of the derivation of Equations (8) and (9), please refer to our technical report [8].

## 6 Related Work

Let us present a brief review of some related work. There are a large model of work about Internet pricing [11, 14, 9, 10], but they are mostly about customer pricing strategy, or to provide differentiated service. On the other hand, our work focus on the “interaction” between the major ISP and local ISPs. In [2], authors investigated the revenue maximization and scalability of a service provider. Their work showed that there is rationale for the service provider to upgrade its capacity. But their model is different from ours in two ways. There is only one common link in the ISP and each peer considers the congestion cost of that common link. Our model differs from them in a sense that we allow a more realistic representation of today’s Internet, that is, we allow peers to have private links so as to reduce their cost. In [13] proposed a model consisting of local and transit ISPs. They showed that optimal strategy of local ISPs is to play “cooperatively” by threat. Our work found the conditions for every peer to achieve its maximum utility.

## 7 Conclusion

In this paper, we investigate the interplay between a higher tier ISP and  $n$  local ISPs (which we terms as peers). A peer has a connection to the ISP, and possibly connected to other peers with some private links. Each peer needs to determine the appropriate amount of traffic via the ISP’s link and the private links so as to maximize its utility. The ISP, on the other hand, needs to perform proper resource allocation so as to avoid resource monopoly and to maximize its revenue. We show the necessary and boundary conditions for the traffic rate vectors of a peer to obtain the maximum utility. We present a distributed algorithm for the ISP to do the resource allocation. The distributed algorithm converges quickly in case the ISP has sufficient resource. We show how the ISP can estimate its revenue with a unit price so as to maximize its revenue. The proposed methodology provides us a systematic way to determine pricing and resource allocation even when the ISP and peers interact with each other.

**Acknowledgement:** The work of D.M. Chiu is supported in part by the RGC grant 411505. This work is affiliated with the Microsoft-CUHK Joint Laboratory for Human- centric Computing and Interface Technologies. John C.S. Lui was supported in part by the RGC grant.

## References

- [1] A.M.Odlyzko. Pricing and architecture of the internet: Historical perspectives from telecommunications and transportation. Available at:

- <http://www.dtc.umn.edu/~odlyzko/doc/networks.html>, March 2004.
- [2] T. Basar and R. Srikant. Revenue-maximizing pricing and capacity expansion in a many-user regime. In *Proceedings of the IEEE Infocom*, New York, June 2002.
- [3] L. He and J. Walrand. Pricing and revenue sharing strategies for internet service providers. In *Proceedings of the IEEE Infocom*, Miami, March 2005.
- [4] L. He and J. Walrand. Pricing differentiated internet services. In *Proceedings of the IEEE Infocom*, Miami, March 2005.
- [5] G. Huston. Interconnection, peering and settlements.
- [6] F. Kelly. Charging and rate control for elastic traffic. In *European Transactions on Telecommunications*, volume 8, 1997.
- [7] J.-J. Laffont, S. Marcus, P. Rey, and J. Tirole. Internet interconnection and the off-net cost pricing principle. In *Institut d’economic Industrielle (IEDI)*, Mimeo, France, 2001.
- [8] S. C. M. Lee, J. W. J. Jiang, D. M. Chiu, and J. C. S. Lui. Interplay of isps: Distributed resource allocation and revenue maximization, technical report.
- [9] P. Marbach. Pricing differentiated services networks: Bursty traffic. In *Proceedings of the IEEE Infocom*, Alaska, March 2001.
- [10] P. Marbach. Priority service and max-min fairness. In *Proceedings of the IEEE Infocom*, New York, March 2002.
- [11] A. M. Odlyzko. The economics of the internet: Utility, utilization, pricing, and quality of service. In *AT&T Research Lab, Tech Report*, 1998.
- [12] M. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, 1957.
- [13] S. Shakkottai and R. Srikant. Economics of network pricing with multiple isps. In *Proceedings of the IEEE Infocom*, Miami, March 2005.
- [14] S. Shenker, D. Clark, D. Estrin, and S. Herzog. Pricing in computer networks: Reshaping the research agenda. In *ACM Computer Communication Review*, Vol 26, 1996.
- [15] W.B.Norton. A business case for isp peering. In *NANOG*.