

Performance Modeling on the interaction of ISPs

Sam C.M. Lee, Joe W.J. Jiang, John C.S. Lui

Department of Computer Science & Engineering,
The Chinese University of Hong Kong.
Email: {cmlee,wjjiang,cslui}@cse.cuhk.edu.hk

Abstract. In this paper, we consider interactions of Internet Service Providers (ISPs) and how these interactions can affect the overall traffic and resource allocation between ISPs. In particular, we consider a simplified two-level hierarchical model wherein there are a single tier-1 ISP and $N > 1$ tier-2 ISPs. Each tier-2 ISP needs to pay the tier-1 ISP for the connectivity service. At the same time, a tier-2 ISP can also arrange to have “private peering” links with other tier-2 ISPs. Therefore, each tier-2 ISP can optimize its utility by deciding on the proper traffic routing of transmitting traffic, either via the tier-1 ISP link, or via the private peering link with other tier-2 ISPs. The tier-1 ISP, on the other hand, needs to decide on the proper resource allocation for all its tier-2 peers so as to avoid monopolization of its resource by a single peer (i.e., to achieve customer diversity). We investigate a distributed framework wherein a tier-1 ISP can achieve customer diversity while tier-2 peers can perform their utility maximization in terms of traffic routing. We also explore other important issues such as sensitivity and convergency. Extensive simulations are carried out to quantify the merits of the proposed distributed framework.

Keywords: Interaction, Utility Maximization, Sensitivity and Convergency.

1 Introduction

Current Internet is basically hierarchical in nature: there are many tier-1 Internet Service Providers (ISPs) providing backbone connectivity service. Regional ISPs, which are usually called the tier-2 ISPs, are *customers* to these tier-1 ISPs and they pay for the connectivity service. There are some smaller ISPs, which are called the tier-3 ISPs, are then connected to tier-2 ISPs for connectivity and/or local access. Often times, ISPs of the same tier (or level) may negotiate with each other so as to provide mutual connectivity. Whenever two peers of the same tier are connected, we say that these two peers have established a *private peering* relationship. In essence, a private peering relationship allows two peers to transfer traffic for their local customers without going through their providers at the higher tier. Figure 1 illustrates the hierarchy and various connectivity relationships.

There are two common ways for a local ISP (which we call peer from now on) to gain the Internet access. The first way is to transmit traffic via the connectivity service provided by its providers (or higher tier ISPs). For this form of traffic transmission, cost will be incurred since the local ISP needs to pay for the amount of traffic transmitted via the *provider-customer* link. The other way to gain the Internet access is to transmit the

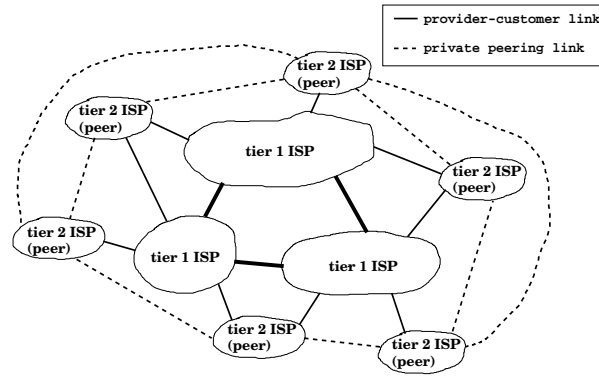


Fig. 1. Internet hierarchical relationship and various connectivity relationships.

traffic via the private peering link, and that the traffic is destined to that particular local ISP. Since there are two ways to transmit traffic, a local ISP often needs to decide on the proper traffic routing: the amount of traffic transmission via the provider-customer link and private peering links, so as to minimize its operating cost and at the same time, satisfy some bandwidth or quality-of-service constraints. It is important to note that for two peers that are geographically apart, it may not be economical or even possible to establish a private peering link, therefore, connectivity between these two peers will be provided by the higher tier ISPs.

The tier-1 ISP (which we call ISP from now on), needs to set a proper price for each of its customers so as to attract peers to transmit traffic via the provider-customer link. If the price is set too high, peers may opt to transmit via their peering links. On the other hand, if the price is set too low, then the ISP may not be able to maximize its profit. Another important issue that an ISP needs to consider is proper resource provisioning so that no peer can monopolize the connectivity resource. It is obvious that there exists certain level of *interaction* between peers and this interaction can affect the decision of the ISP. In particular, the level of pricing and resource allocation can affect the routing decisions of peers, and the routing decisions by peers also affect the level of resource allocation set by an ISP.

The contribution of our paper is to provide an *understanding* of this form of interaction, in particular, how it can affect the routing strategy of individual peers and resource allocation of an ISP. We propose a distributed algorithm wherein peers and ISP can communicate so that a peer can maximize its utility while the ISP can provide fair resource allocation. We also show that the distributed algorithm is stable and can converge to an equilibrium point quickly.

Let us summarize some related work in this topic. There is a rich literature on Internet pricing [6, 8, 4, 5], but they mainly concern about the pricing strategy of individual customers, i.e., access charging. Our work focuses on the interaction of a tier-1 ISP and its customers, namely, tier-2 ISPs. Authors in [1] present the revenue maximization and scalability of a service provider. It shows that using the suggested pricing scheme, there are sufficient incentives for an ISP to upgrade its network. However, they do not

consider the interaction between peers, namely, peers can exchange traffic via private peering links.

The organization of the paper is as follows. In Section 2, we present our mathematical model of representing the interaction of the ISP and its peers. We formulate the maximization function for a peer and show how each peer can perform routing so as to maximize its utility. Conditions of maximization are also presented. We also present the resource allocation algorithm of the ISP to achieve customer diversity. In Section 3, we investigate the convergency issue of the distributed algorithm. In Section 4, we investigate the sensitivity of the system behavior on various system parameters. Finally, conclusion is given in Section 5.

2 The Mathematical Model, Distributed Maximization and Resource Allocation

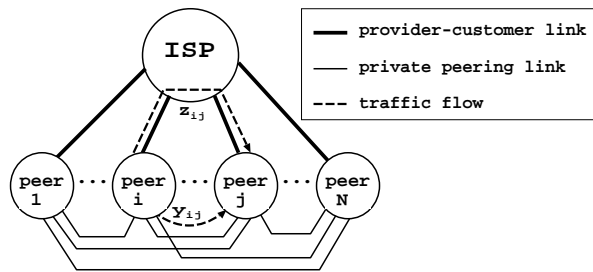


Fig. 2. A simplified two-tiers hierarchical model with one tier-1 ISP and N peers.

Let us consider a simplified two-tiers hierarchical model in which there is one tier-1 ISP and N tier-2 ISPs (or peers). Peers can communicate with each other either by sending traffic through the ISP, or by the private peering links between two peers. To provide connectivity, the ISP has a communication network which has a total capacity of \mathcal{R} (in units of Mbps). For a peer $i \in \{1, 2, \dots, N\}$, it has a link to the ISP, and possibly $N - 1$ private peering links connecting to the other $N - 1$ peers. Let l_{ij} denote the private peering link between peer i and peer j and this link has a capacity of r_{ij} (in unit of Mbps). Note that when $r_{ij} = 0$, it implies that there is no private 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{R}_i amount of bandwidth (in units of Mbps) for this connection. It is important to point out that our model can be viewed as a generalization of the network model presented in [1], in which private peering links are not considered and so there is no interaction between peers. Note that we only have one ISP in our model. The issues of multiple ISPs and multihoming are much more complicated and will appear in our future work. Lastly, Table 1 contains all notations used in our mathematical model.

Let x_{ij} denote the traffic demand (or transmission rate in unit of Mbps) from peer i to peer j . In essence, it is the traffic originated from peer i and destined to peer j . If

| | |
|-----------------|--|
| N | : Number of peers (the local ISPs) in the communication network. |
| l_{ii} | : An abstraction of the communication link between peer i and the ISP. |
| l_{ij} | : The private communication link connecting peer i to peer j . |
| \mathcal{R} | : Total capacity of the ISP link. |
| \mathcal{R}_i | : Allocation of ISP's link bandwidth to peer i . |
| r_{ij} | : Capacity of the private link l_{ij} connecting peer i to peer j . |
| x_{ij} | : Traffic demand or transmission rate from peer i to peer j , such that $x_{ij} = y_{ij} + z_{ij}$. |
| y_{ij} | : Traffic transmission rate from peer i to peer j going through the private link l_{ij} . |
| z_{ij} | : Traffic transmission rate from peer i to peer j going through the ISP link l_{ii} . |
| z_i | : Aggregate traffic rate that peer i sends through the ISP link. |
| \bar{z} | : Aggregate traffic rate through the ISP link from all peers. |
| \mathcal{P}_i | : Price per unit bandwidth on the ISP link for peer i . In this work we assume $\mathcal{P}_i = \mathcal{P}$ for all i . |
| p_{ij} | : Price per unit bandwidth of the private peering link l_{ij} . |
| \mathbf{y}_i | : $\mathbf{y}_i = (y_{i1}, y_{i2}, \dots, y_{in})$ denotes the traffic rate vector for peer i through its private links. |
| \mathbf{z}_i | : $\mathbf{z}_i = (z_{i1}, z_{i2}, \dots, z_{in})$ denotes the traffic rate vector for peer i through the ISP link. |

Table 1. Mathematical notations.

the peer i can sustain the transmission rate of x_{ij} , peer i receives a utility of $F_{ij}(x_{ij})$, where F_{ij} is a strictly concave function in x_{ij} . The utility $F_{ij}(x_{ij})$ represents the degree of happiness of peer i by sending data to peer j at the rate of x_{ij} . Noted that concave function is commonly used to represent elastic traffic[2], which is the dominant traffic in the current Internet.

This constant traffic demand x_{ij} 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 the ISP link l_{ii} . In other words, we have: $x_{ij} = y_{ij} + z_{ij}$, with $y_{ij}, z_{ij} \geq 0$, for $i, j \in \{1, \dots, N\}$.

Let x_{ii} be the traffic demand from peer i to destinations other than the $N - 1$ peers. This represents traffic to other part of the Internet wherein peer i has to send the data through the ISP. In this case, peer i can only use the provider-customer link l_{ii} for the traffic transmission. Therefore, we have the following relationship: $y_{ii} = 0$ and $x_{ii} = z_{ii}$ for $i \in \{1, \dots, N\}$. For the ease of presentation, let $z_i = \sum_{j=1}^N z_{ij}$ denote the aggregate traffic demand 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 network 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 issues on the cost of setting up peering links, since it is not part of the operating cost. We assume peers can utilize existing peering links with fixed capacities r_{ij} . For convenience, we denote $\mathbf{y}_i = (y_{i1}, y_{i2}, \dots, y_{iN})$ as the traffic rate vector for peer i , representing the traffic going through its private peering links and $\mathbf{z}_i = (z_{i1}, z_{i2}, \dots, z_{iN})$ as the traffic rate vector for peer i through the ISP link. We denote $\mathcal{P} = (\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_N)$ as the price vector set by the ISP for different peers.

In making the routing decision, each peer not only needs to consider the cost of transmitting the traffic, but also on the quality of service. In other words, each peer needs to take into consideration of the delay or congestion cost on the links. In this work, we assume each link is represented by an M/M/1 model as in [1], and one can take the average delay on the link as its congestion indicator.¹ Rather than informing all peers about the current transmission demands \bar{z} on the ISP link (this is considered as a confidential information by a peer), the ISP will compute and announce its bandwidth allocation to peer i as \mathcal{R}_i . Under this form of setting, the congestion cost \mathcal{D}_{ij} of a link l_{ij} can be represented as:

$$\mathcal{D}_{ij} = \begin{cases} \frac{1}{r_{ij} - y_{ij}} & \text{if } i \neq j, \\ \frac{1}{\mathcal{R}_i - z_i} & \text{if } i = j. \end{cases}$$

The traffic demand x_{ij} can be viewed as the long-term average aggregate request from the customers of peer i destined to peer j . For example, average of a monthly traffic from a peer. So the traffic demand is considered as a constant within a certain period. Consider the case when peer i can always obtain a *sufficient* bandwidth capacity to transmit all the aggregate requests, i.e. $\sum_j x_{ij} \leq \sum_j r_{ij} + \mathcal{R}_i$, then the peer will transmit all the requests, while maximizing its utility at the same time. With fixed traffic demands x_{ij} 's, the aggregate happiness $\sum_j F_{ij}(x_{ij})$ is therefore a constant.

Let us now formulate the objective function of a peer, say i . Peer i wants to maximize the following function:

$$\begin{aligned} \text{Max } U_i &= \sum_j F_{ij}(x_{ij}) - \mathbf{1}_{\{y_{ij} \neq 0\}} \left[\frac{1}{r_{ij} - y_{ij}} \right] - \sum_{j \neq i} p_{ij} y_{ij} - \mathbf{1}_{\{z_i \neq 0\}} \left[\frac{1}{\mathcal{R}_i - z_i} \right] - \mathcal{P}_i z_i \\ \text{Max } U_i &= K - \sum_{j \neq i} \mathbf{1}_{\{z_{ij} \neq x_{ij}\}} \left[\frac{1}{r_{ij} - x_{ij} + z_{ij}} \right] + \sum_{j \neq i} p_{ij} z_{ij} - \mathbf{1}_{\{z_i \neq 0\}} \left[\frac{1}{\mathcal{R}_i - z_i} \right] - \mathcal{P}_i z_i \\ \text{Min } C_i &= \sum_{j \neq i} \mathbf{1}_{\{z_{ij} \neq x_{ij}\}} \left[\frac{1}{r_{ij} - x_{ij} + z_{ij}} \right] - \sum_{j \neq i} p_{ij} z_{ij} + \mathbf{1}_{\{z_i \neq 0\}} \left[\frac{1}{\mathcal{R}_i - z_i} \right] + \mathcal{P}_i z_i \quad (1) \\ \text{s. t.} \quad & \max\{0, x_{ij} - r_{ij}\} \leq z_{ij} \leq x_{ij} \quad \text{for all } j \neq i, \quad z_{ii} = x_{ii}, \quad \sum_j z_{ij} \leq \mathcal{R}_i \quad (2) \end{aligned}$$

where $K = \sum_j F_{ij}(x_{ij}) - \sum_{j \neq i} p_{ij} x_{ij}$ is a constant.

In here, $\mathbf{1}_{\{p\}}$ is an indicator function. The objective of the optimization problem (1) is to minimize the aggregate congestion costs and payments under constant traffic demands. The variable transmission rate vector \mathbf{y}_i is absorbed and the remaining variable in the new optimization problem is \mathbf{z}_i . The constraints in Equation (2) represent the feasible region of the ISP link transmission rates. The first constraints give the lower and upper bounds for z_{ij} 's. When $r_{ij} \geq x_{ij}$, the bandwidth in the private peering link l_{ij} is larger than the demand x_{ij} , i.e. private peering link capacity is sufficient for the demand and so the minimum transmission rate in ISP link z_{ij} is zero. When $r_{ij} < x_{ij}$, the bandwidth in the private peering link is insufficient for the demand and so part of the traffic must go through the ISP link. It makes the minimum value of $z_{ij} = x_{ij} - r_{ij}$. The

¹ We have also investigated in using expected waiting time in the peers' utility functions. Please refer to our technical report[3].

second constraint again is due to the absence of private peering link to the “outsiders”. The third constraint is the ISP link capacity constraint.

It is important to point out that the optimization process is indeed a *coupled* optimization process. For each peer i , given the bandwidth allocation \mathcal{R}_i of the ISP link, it performs an optimization and determines its optimal rates z_i and will send a bid for this bandwidth allocated by the ISP. After collecting the bidding information from all peers, the ISP calculates the bandwidth allocation according to the biddings. Peers will offer their new biddings based on the allocated bandwidth by the ISP. We model this interaction process as a *non-cooperative game* wherein each peer offers a bid to the ISP so as to minimize its own cost.

For a given ISP price vector $\mathcal{P} = (\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_N)$, this defines a non-cooperative game between these N peers [7]. These peers interact with each other and determine their optimal transmission rates periodically and asynchronously. For each price vector $\mathcal{P} > 0$, a Nash equilibrium point for this N -peers game is defined as N -tuple $\mathbf{z}^* = (z_1^*, z_2^*, \dots, z_N^*)$, such that for *all* peers $i \in \{1, 2, \dots, N\}$:

$$\mathcal{C}_i(\mathbf{z}^*, \mathcal{P}) \leq \mathcal{C}_i(\mathbf{z}, \mathcal{P}) \quad (3)$$

for any other feasible traffic vector $\mathbf{z} = (z_1, z_2, \dots, z_N)$ that satisfies the constraints defined in Equation (2).

2.1 Distributed Solution of the Minimization Problem

In the following, we illustrate how a peer, say i , can determine its transmission rates, that is z_i , rates to other peers via the ISP’s link, as well as y_i , rates to other peers via private peering links, so as to minimize its cost when the bandwidth supply is *sufficient*. Assuming that the peer knows the price \mathcal{P}_i specified by the ISP and the associated bandwidth allocation \mathcal{R}_i , one can model an individual peer’s behavior as a convex optimization problem as defined in Equation (1). Let us first study the necessary and boundary conditions for a peer to minimize the cost.

Necessary conditions with positive transmission rate

Since the cost \mathcal{C}_i is discontinuous at $z_{ij} = x_{ij}$ (i.e., transmission rate through the private peering link l_{ij} is zero) and $z_i = 0$ (i.e., transmission rate through the ISP link is zero), we first investigate the necessary conditions when $z_{ij} \neq x_{ij}$ and $z_i \neq 0$. The optimization problem of Equation (1) has $N - 1$ variables (with $z_{ii} = x_{ii}$). The first and second order partial derivatives with respect to z_{ij} and z_{ik} for $k \neq j \neq i$ are:

$$\begin{aligned} \frac{\partial \mathcal{C}_i}{\partial z_{ij}} &= \frac{-1}{(r_{ij} - x_{ij} + z_{ij})^2} - p_{ij} + \frac{1}{(\mathcal{R}_i - z_i)^2} + \mathcal{P}_i, \\ \frac{\partial^2 \mathcal{C}_i}{\partial z_{ij}^2} &= \frac{2}{(r_{ij} - x_{ij} + z_{ij})^3} + \frac{2}{(\mathcal{R}_i - z_i)^3} > 0, \\ \frac{\partial^2 \mathcal{C}_i}{\partial z_{ij} \partial z_{ik}} &= \frac{2}{(\mathcal{R}_i - z_i)^3} > 0. \end{aligned}$$

This shows that the Hessian matrix of the objective function in Equation (1) is positive definite on the non-negative space bounded by the capacity constraints $x_{ij} - r_{ij} \leq z_{ij} \leq$

x_{ij} and $z_i \leq \mathcal{R}_i$. So the cost \mathcal{C}_i is strictly convex in z_{ij} for all $j \neq i$. The minimum cost and optimizer to this problem is unique and can be found by the Lagrangian method. The necessary conditions of z_{ij} for the minimization of \mathcal{C}_i are:

$$\frac{\partial \mathcal{C}_i}{\partial z_{ij}} \begin{cases} > 0 & \text{if } z_{ij} = 0, \\ = 0 & \text{if } z_{ij} > 0. \end{cases} \quad (4)$$

Boundary cases to Minimization problem

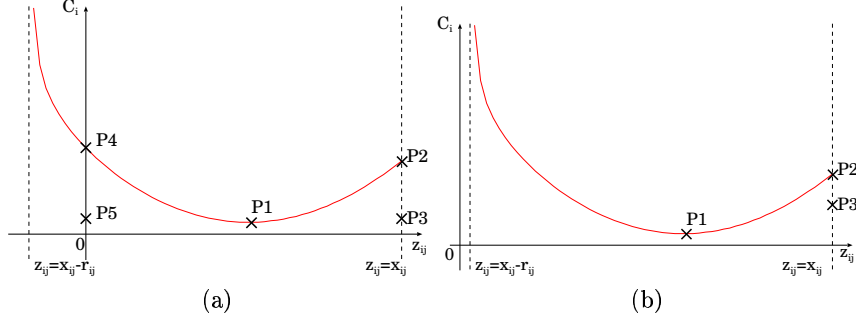


Fig. 3. Peer i 's cost against transmission rate z_{ij} when $z_{ij}^* = \arg\{\frac{\partial \mathcal{C}_i}{\partial z_{ij}} = 0\}$ is in the feasible range: (a) when $x_{ij} \leq r_{ij}$, (b) when $x_{ij} > r_{ij}$.

Due to the discontinuity of the objective function, the necessary conditions given above may not achieve the global minimum. In here we explore the boundary cases when the transmission rates are zero, i.e., $z_{ij} = x_{ij}$ or $z_i = 0$. Figures 3 and 4 show these cases. Figure 3 corresponds to the case when $z_{ij}^* = \arg\{\frac{\partial \mathcal{C}_i}{\partial z_{ij}} = 0\}$ is in the feasible range. The vertical axis shows the aggregate cost \mathcal{C}_i and the horizontal axis shows the transmission rate z_{ij} . Figure 3(a) considers if $x_{ij} \leq r_{ij}$, which implies the private peering link capacity is adequate for the transmission demand; and Figure 3(b) considers if $x_{ij} > r_{ij}$, which implies the private peering link capacity is inadequate for the transmission demand. In Figure 3, the minimum point of the curve is at P1 when $z_{ij} = z_{ij}^*$. We first consider the upper bound. When $z_{ij} = x_{ij}$, the transmission rate goes through the ISP link only. The congestion cost in peering link l_{ij} is not considered and is subtracted from \mathcal{C}_i , so P3 rather than P2 is the point of \mathcal{C}_i when $z_{ij} = x_{ij}$. We then consider the lower bound under two cases: case i) when $x_{ij} \leq r_{ij}$ (as in Figure 3(a)), the minimum value of $z_{ij} = 0$. If there is a $z_{ik} > 0$ for some $k \neq j$, P4 is the point when $z_{ij} = 0$. But if the aggregate traffic through the ISP link is zero ($z_i = 0$), P5 is the point when $z_{ij} = 0$. Note that the congestion cost in the ISP link is subtracted from \mathcal{C}_i in this case; case ii) when $x_{ij} > r_{ij}$ (as in Figure 3(b)), the minimum value of $z_{ij} = x_{ij} - r_{ij}$. This is because the maximum amount of traffic through the private peering link is r_{ij} , the remaining rate $x_{ij} - r_{ij}$ has to go through the ISP link and the congestion cost in the ISP link must be considered. In general, when z_{ij}^* is in the feasible range, the optimal transmission rate is either $z_{ij} = 0$, $z_{ij} = x_{ij}$ or $z_{ij} = z_{ij}^*$.

Figure 4 illustrates when $z_{ij}^* = \arg\{\frac{\partial \mathcal{C}_i}{\partial z_{ij}} = 0\}$ is not in the feasible range. Figure 4(a) considers when $z_{ij}^* \leq \min\{0, (x_{ij} - r_{ij})\}$. As \mathcal{C}_i is strictly convex in z_{ij} , the

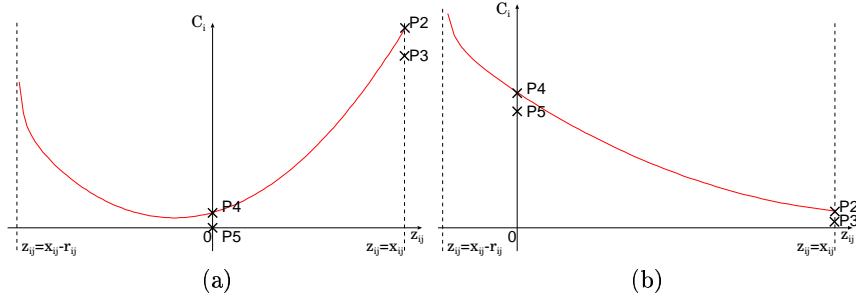


Fig. 4. Peer i 's cost against transmission rate z_{ij} when $z_{ij}^* = \arg\{\frac{\partial C_i}{\partial z_{ij}} = 0\}$ is not in the feasible range: (a) when $z_{ij}^* \leq \min\{0, (x_{ij} - r_{ij})\}$, (b) when $z_{ij}^* \geq x_{ij}$.

minimum feasible $z_{ij} = \min\{0, (x_{ij} - r_{ij})\}$, is either at $P4$ (when $z_i > 0$) or at $P5$ (when $z_i = 0$). For the upper bound of z_{ij} , when $z_{ij} = x_{ij}$, the congestion cost in the private peering link is subtracted from the cost. This concludes that the minimum point of C_i is either when $z_{ij} = \min\{0, (x_{ij} - r_{ij})\}$ (optimizer is either $P4$ or $P5$) or when $z_{ij} = x_{ij}$ (optimizer is $P3$). Figure 4(b) shows the case when $z_{ij}^* > x_{ij}$. The maximum feasible $z_{ij} = x_{ij}$ due to the convexity of C_i . For the lower bound of z_{ij} , when $z_i = 0$ (which implies $z_{ij} = 0$), the value of C_i at $P5$ may be smaller than that at $P3$. This concludes that the minimum point of C_i is either $P3$ when $z_{ij} = x_{ij}$ or $P5$ when $z_{ij} = 0$. Lastly, after the ISP link transmission rates z_{ij} 's are computed, the private peering link transmission rates y_{ij} 's can be found by $y_{ij} = x_{ij} - z_{ij}$.

2.2 Distributed Resource Allocation by ISP

Let us present the algorithm in which the ISP can determine the appropriate capacity \mathcal{R}_i for every peer i , for $i \in \{1, \dots, N\}$. At the beginning, the ISP distributes its capacity equally among all peers, so $\mathcal{R}_i = \frac{\mathcal{R}}{N}$ for all i , and sends the distribution \mathcal{R}_i to every peer i . Upon receiving the information, each peer i applies the procedure described in the previous sub-section to compute its own optimal transmission rates (y_i and z_i) and sends the information ($z_i = \sum_j z_{ij}$) to the ISP as its bidding for the ISP capacity. The ISP gathers the biddings from the peers. Then it allocates the resource represented by the following formula:

$$\mathcal{R}_i = z_i + \frac{(\mathcal{R} - \bar{z})}{N}.$$

3 Convergency of Traffic Rates

With the traffic distribution algorithm by peers and resource allocation by the ISP we described in the previous section, one important issue that we need to address is whether these traffic rates and biddings will converge. In this section, we investigate the convergency of rates and biddings of peers when the number of peers is large.

The experiment considers the case when *all* peers are of similar sizes and thus have similar traffic demands. We show that the traffic rates distributions and biddings of all

peers converge rapidly. We have also investigated in the case when *some* of the peers are of *larger and smaller* sizes. For more details, please refer to our technical report [3].

The environment of the experiment is constructed as follows. There is a network of $N = 50$ peers and one ISP. Each peer has an aggregate link connecting to the ISP and peering links connecting to the other peers. A peering link l_{ij} exists with a probability of 0.5. If the link l_{ij} exists, its capacity r_{ij} is chosen from a random variable uniformly distributed in $[0, 10]$, and its unit price p_{ij} is another random variable uniformly distributed in $[1, 2]$. The traffic demand x_{ij} is a random variable equal to 0 with probability 0.3 and uniformly distributed in $[0, 10]$ with probability 0.7. The ISP provides a bandwidth of $\mathcal{R} = 7500$ units with unit price $\mathcal{P} = 2.5$. After each period of one second, the ISP applies the resource allocation algorithm we described in the previous section, and sends the distribution \mathcal{R}_i to every peer i . Upon receiving the signal from the ISP, peer i performs the traffic rate distribution algorithm with the parameters \mathcal{R}_i and $\mathcal{P}_i = \mathcal{P}$. To introduce noise in the information exchange, a bidding is successfully transmitted to the ISP with a probability of 0.8.

Experiment I: Homogeneous Peers with similar traffic demands : In this experiment, we have a network of 50 peers and one ISP and the parameters are generated as described above. The peers are of similar sizes and have similar traffic demands. Figure 5(a) shows the biddings of peers 1, 2, 3 and 4 throughout the experiment. The vertical axis shows the biddings of the peers and the horizontal axis shows the time. Note that the biddings converge rapidly. Figure 5(b) shows the transmission rates of peer 8. The vertical axis shows the transmission rates and the horizontal axis shows the time. Again, we observe that the traffic rates converge to the equilibrium quickly.

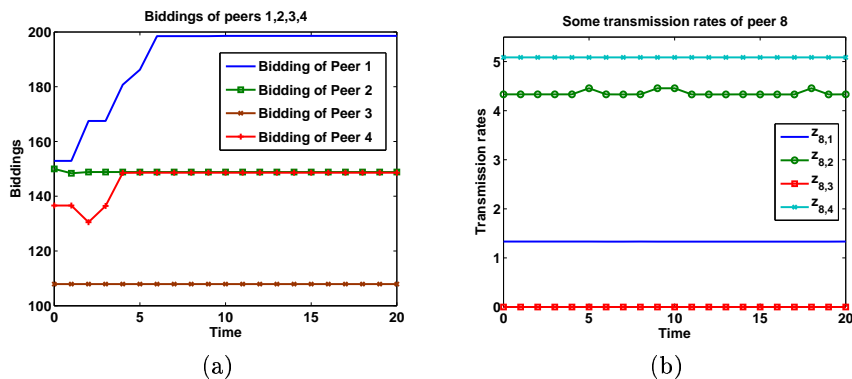


Fig. 5. Exp. I: (a) biddings of peers 1, 2, 3 and 4, (b) samples of transmission rates of peer 8.

4 Sensitivity to System Parameters

In this section, we investigate in the sensitivity of the equilibrium point as we vary some of the system parameters. The observation is made to the variation of the transmission

rates and the biddings from peers. We have three experiments, each corresponds to only one variation in the system parameters.

- Experiment A: variation in the unit price of the private link (p_{ij})
- Experiment B: variation in the traffic demand (x_{ij})
- Experiment C: variation in the unit price of the ISP link (\mathcal{P})

The environment of the experiments is constructed as follows. We have a network of 10 peers and one ISP. Each peer has an aggregate link connecting to the ISP and 9 private links connecting to the other peers with capacity $r_{ij} = 10$ and unit price $p_{ij} = 1.0$. All peers have the same traffic demands (i.e., $x_{ij} = 10$ for all i, j). The ISP provides a bandwidth of $\mathcal{R} = 300$ units with unit price $\mathcal{P} = 1.2$. After every period of one second, the ISP applies the resource allocation algorithm, and sends the distribution \mathcal{R}_i to every peer i . Upon receiving the signal from the ISP, peer i performs the rate distribution algorithm with the parameters \mathcal{R}_i and $\mathcal{P}_i = \mathcal{P}$. To introduce noise in the information exchange, a bidding is successfully transmitted to the ISP with a rate of 80%.

In the following experiments, we only change the parameters of peer 10 while keeping the parameters of the other peers unchanged. Our observation is made to the changes of traffic rate distributions and biddings of peer 10 and other peers (eg. peer 8).

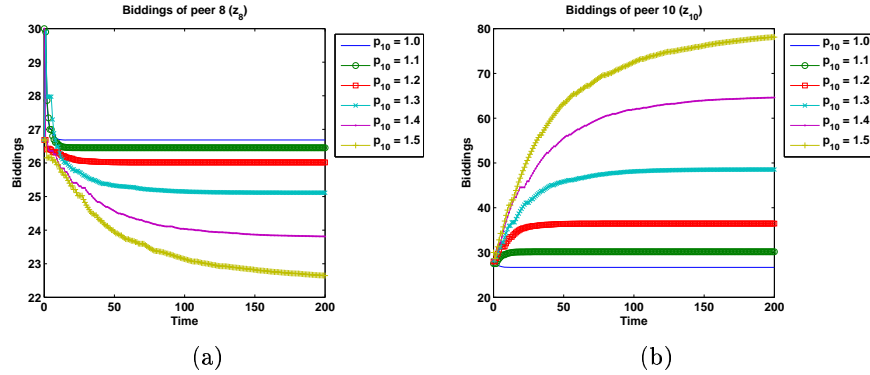


Fig. 6. Experiment A: (a) peer 8's bidding (z_8), (b) peer 10's bidding (z_{10}).

Experiment A: Change of unit price in private link (p_{ij}): In here, we show how the change in the unit price of private link affects the biddings of all peers. The parameters are constructed as described above. We vary the unit prices of the private links of peer 10 from 1.0 to 1.5, and investigate in the effects in the biddings of peers. Figure 6 illustrates the biddings in the ISP link bandwidth of peers 8 and 10. The vertical axis shows the biddings and the horizontal axis shows the time. When the unit price of the private links of peer 10 increases, peer 10 shifts its traffic from the private links to the ISP link in order to reduce the payment. So peer 10 offers a larger bid. This increases the congestion cost in the ISP link. Other peers (e.g., peer 8) detect this increase in

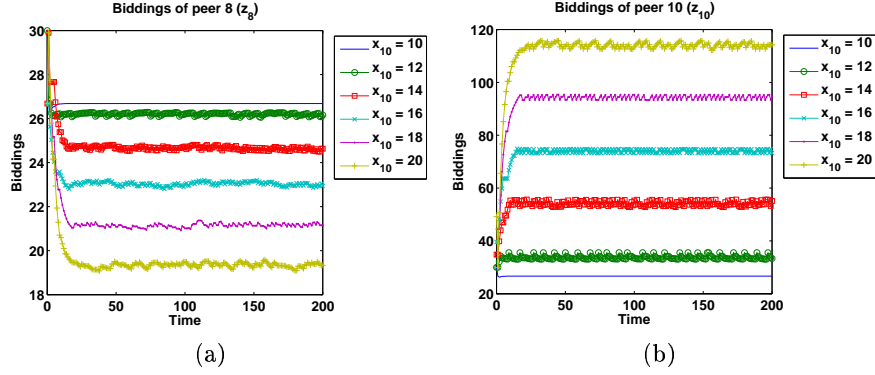


Fig. 7. Experiment B: (a) peer 8’s bidding (z_8), (b) peer 10’s bidding (z_{10}).

congestion cost and so they shift their traffics from the ISP link to their private links, and give smaller bids.

Experiment B: Change of traffic demand of a peer (x_{ij}): In here, we show how the change in the traffic demand affects the biddings of all peers. The parameters are constructed as described above. We increase all the traffic demands of peer 10 from 10 to 20, and investigate the effects in the biddings of peers. Figure 7 illustrates the biddings in the ISP link bandwidth of peers 8 and 10. When peer 10 increases the traffic demands, it has to increase the transmission rates both in the private links and the ISP link. So peer 10 gives a larger bid to ISP asking for more bandwidth. This increases the congestion cost in the ISP link. Other peers (eg. peer 8) detect this increase and shift their traffics from the ISP link to their private links, and give smaller bids.

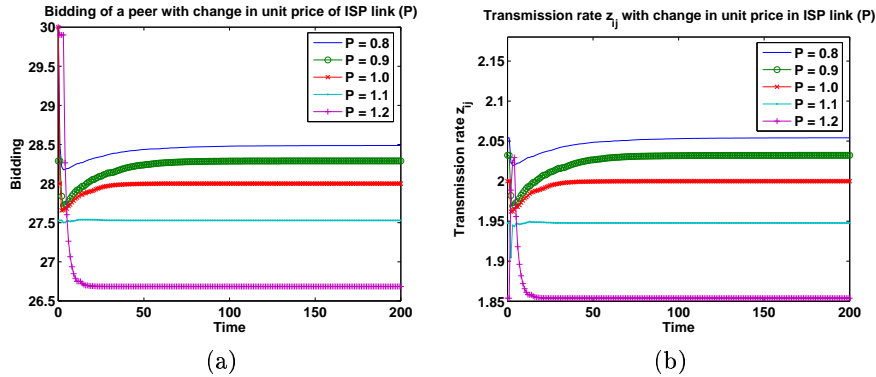


Fig. 8. Experiment C: (a) biddings in ISP link (z_i), (b) transmission rates (z_{ij}).

Experiment C: Change of unit price in ISP link (\mathcal{P}): In here, all peers have the same traffic demands, and the unit prices and capacities of all private links are identical. The only changing parameter is the unit price of the ISP link. The unit price decreases

from 1.2 to 0.8. Figure 8(a) shows the biddings in the ISP link of a peer throughout the experiment. We observe that a peer increases the bidding in the ISP link with decreasing price in the link. Figure 8(b) shows the transmission rate z_{ij} for $i \neq j$ throughout the experiment. We see that a peer increases the transmission rate in the ISP link with decreasing price in ISP link. The increases in both the biddings and transmission rates in ISP link are due to the decrease in payment to the ISP link. As a result, peers shift some of their traffics from the private link to the ISP link.

We have two extra experiments in the sensitivity test of (i) private link capacity (r_{ij}) and (ii) ISP link capacity (\mathcal{R}). For more details, please refer to our technical report [3].

5 Conclusion

In this paper, we investigate the interplay between a tier-1 ISP and N tier-2 ISPs (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 minimize its cost. The ISP, on the other hand, needs to perform proper resource allocation to distribute its resource properly. We show the necessary and boundary conditions for the transmission rate vectors of a peer to obtain the minimum cost. We present an algorithm for the ISP to do the resource allocation. We then show the optimal rates and biddings of peers converge with the resource allocation algorithm of ISP even when the number of peers is large. Finally, we show and explain how the change in a single parameter can affect the optimal rates and biddings of all peers and that peers can adapt to these changes and quickly converge to an equilibrium solution. The complicated issues of multiple ISPs and multihoming will appear in our future work.

Acknowledgement: This work is supported in part by the RGC grant.

References

1. 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.
2. F. Kelly. Charging and rate control for elastic traffic. In *European Transactions on Telecommunications*, volume 8, 1997.
3. S. C. M. Lee, J. W. J. Jiang, and J. C. S. Lui. Performance modeling on the interaction of isps technical report.
4. P. Marbach. Pricing differentiated services networks: Bursty traffic. In *Proceedings of the IEEE Infocom*, Alaska, March 2001.
5. P. Marbach. Priority service and max-min fairness. In *Proceedings of the IEEE Infocom*, New York, March 2002.
6. A. M. Odlyzko. The economics of the internet: Utility, utilization, pricing, and quality of service. In *AT&T Research Lab, Tech Report*, 1998.
7. M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, 1997.
8. 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.