Modeling the Peering and Routing Tussle between ISPs and P2P Applications

Jessie Hui Wang   Dah Ming Chiu
Department of Information Engineering
The Chinese University of Hong Kong
Email: {hwang3, dmchiu}@ie.cuhk.edu.hk

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

***Abstract*— The connectivity between millions of nodes on the Internet is provided by the interconnection of many ISPs' networks. These ISPs, in their decisions to peer with each other, define a set of *transit relationships*. These transit relationships are the primary factors that dictate how traffic flows through the Internet. BGP-based inter-domain routing that implements these transit relationships can be considered *economically efficient*. The advent of *peer-to-peer* (P2P) applications and overlay networks, however, changes the rules by providing traffic routing favoring the applications' needs. This can lead to reduced economic efficiency and upset the ISPs' business model. In this paper, we propose simple models to represent P2P traffic demand, peering and routing in a market place of two competing ISPs to illustrate this tussle of the Internet. Based on these models, we also propose and investigate alternative peering and provisioning strategies available to the ISPs and analyze their effectiveness.**

I. Introduction

The Internet is operated by many ISPs (or ASes) who decide to interconnect their networks together. Their peering relationships define a set of transit service agreements [1]. These transit agreements are in turn primary factors that determine how traffic navigates across ISPs' networks in the Internet.

There are primarily two kinds of peering relationships between Internet ISPs: (i) *provider to customer* relationship, and (ii) *free peering* (also called *peer to peer* relationship). Under the provider to customer relationship, the customer ISP pays the provider ISP for connectivity to the rest of the Internet. So the provider ISP provides unrestricted transit service to the customer ISP. The customer ISP also needs to provide some transit service to the provider, only to reach the customer ISP (or its customers), but not to any other destinations. The customer ISP is thus providing a form of *selective* transit service.

Under the free peering relationship, the traffic exchange on the peering link is nominally free of charge. Only local traffic, which is the traffic between the two free peering ISPs and their respective customer ISPs, can be exchanged on the free peering link [2]. Such traffic exchange helps both peering ISPs to reduce the dependence on their providers for transit service, and thus save money. Note that in this case, both peering ISPs provide *selective* transit service.

ISPs rely on BGP, a policy-based routing protocol to enforce these selective transit agreements. An ISP's routing policy includes *import routing policy* and *export routing policy*. Import policy, when applied to a particular neighbor X, determines what transit service the local ISP is to accept from the neighbor X. Export policy, applied as a filter to routes sent to neighbor Y, declares what transit service the local ISP is offering to the neighbor Y.

For example, in Figure 1, ISP_C is a transit provider, and ISP_A and ISP_B are customers of ISP_C. At the same time, ISP_A and ISP_B set up a free peering link to reduce the transmission costs on their links to ISP_C. In this network, if a subscriber i of ISP_A needs an object r from the Internet, the information flow can transit through ISP_C and arrive at ISP_A for subscriber i. Due to policy-based routing, the object cannot traverse the path (Internet $\rightarrow ISP_C \rightarrow ISP_B \rightarrow ISP_A$) since ISP_A does not receive the route to r from ISP_B under the free peering agreement.

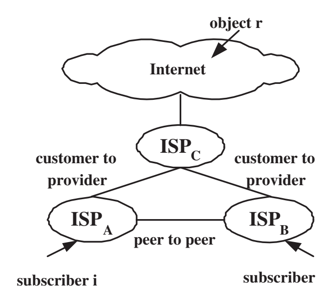

Fig. 1. Example to illustrate peering relationships.

It is reasonable to assume that the routing culminated from such decentralized peering agreements is *economically efficient*. If ISP_A drastically over-provisioned its provider link, another competing ISP (*i.e.* ISP_B) with a lower operating cost would be in the position to undercut the price ISP_A charges its subscribers thus grab customers away from ISP_A. Similarly, if ISP_A and ISP_B do not freely exchange their local traffic, other competing ISPs may do so to undercut their business.

Such economic efficiency, however, does not imply best possible service for individual subscribers and the applications they are running. This conflict may due to a variety of reasons. For example, one reason is that it would be too complicated for the network layer to learn about all the application requirements and tailor its routing accordingly; another reason is that applications would naturally choose *selfish optimal* routing, which does not always coincide with network-wide optimal routing [3][4].

An important class of applications that come into conflict with ISP controlled network routing is the peer-to-peer (P2P) applications. To provide efficient and speedy distribution of content to many receivers, peers play the role of information receiver as well as server at the same time. Imagine the case that the object r is a P2P object in the Internet and it is needed by both subscriber i and subscriber j. To improve performance and scalability, a P2P application makes both subscriber i and j to provide service to each other. So subscriber i may receive some pieces of r from subscriber j and vice versa. If we look at the routing of the pieces from subscriber j to subscriber i at the application layer, they traverse along the path (Internet $\rightarrow ISP_C \rightarrow ISP_B \rightarrow ISP_A$). This part of traffic on the peering link is beneficial only for ISP_A. In this sense, ISP_B is providing transit service for ISP_A without being paid by ISP_A, which is not the intention of the peering agreement between ISP_A and ISP_B. In this situation, policy-based routing (BGP) fails to implement the intended selective transit service because the path is implemented through two network layer connections and both connections are legitimate under the traditional free peering agreements. This example illustrates the routing tussle between ISPs and P2P applications and it shows that policy-based routing cannot always control the routing of P2P traffic effectively. Moreover, ISP_A may also provide transit service for ISP_B without being paid since subscriber j may also receive some pieces of r from subscriber i. So the question becomes which ISP receives more benefit from this free peering, and whether it is fair.

In many networks, the P2P traffic has already overtaken the traffic volume generated by traditional high-volume applications such as web and email. For this reason, the ISP and P2P application routing tussle has become a significant problem[1] for ISPs' business model. This problem also bears significance on Internet's service model, and possible future requirements for policy-based inter-domain routing protocols.

In this paper, our main contribution is to capture the essence of this problem using relatively simple models. Since the P2P traffic is often routed at the application layer according to the availability and performance of paths created by the P2P application, oblivious of transit policies of the ISPs, we model it as the outcome of solving an optimization problem by individual peers in optimizing their file download performance. By applying this model to a simple market scenario of two competing ISPs, we analyze the ISPs' predicaments. We also investigate alternative peering and provisioning options and their effectiveness. Based on our P2P routing model, we conclude that the ISPs must control both their provider link capacity as well as free-peering capacity they provide to other peers to maintain certain parity between different incoming link capacity among all ISPs. Only when such parity is maintained, ISPs can sustain their relative economic positions hence economic equilibrium despite P2P routing.

The rest of the paper is organized as follows. In Section II, we propose a P2P traffic and routing model, and lay out the economic analysis methodology. In Section III, we analyze the situation when all traffic in the network is P2P routable. This assumption simplifies some parameters and gives us the insights for the asymptotic case. A numerical example is given in Section IV. The related work is discussed in Section V. And finally Section VI concludes the paper.

II. Models and Methodology

We focus our study on a simple network scenario with two ISPs competing for the same subscribers in a local market, as shown in Figure 2. The market share of ISP_i is assumed to be $\alpha_i \geq 0$, and we have $\alpha_1 + \alpha_2 = 1$. Without loss of generality, we assume $\alpha_1 \geq \frac{1}{2}$ in the rest of the paper.

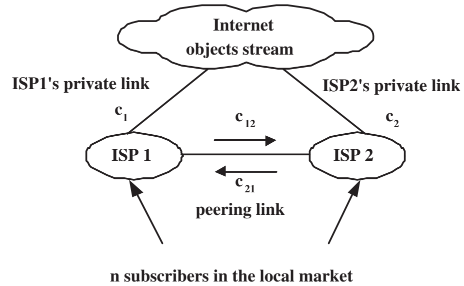

Fig. 2. Network Model.

Both ISPs buy transit service from their respective providers (which could be the same) for access to the Internet. In practice, many ISPs are multihomed, which means they have multiple provider links to reach the Internet. In our model, we use one virtual link between ISP_i and the Internet to denote these multiple physical links and this virtual link is referred to as ISP_i's *private link*. We define the capacity of ISP_i's private link as c_i.

These two local ISPs may also set up a free peering link between them to reduce traffic on their respective private links. We denote the capacity of the peering link in the direction from ISP_i to ISP_j as c_{ij}. When $c_{12} = c_{21} = 0$, it means that there is no peering link between them.

The traffic demand can be defined in terms of end-to-end flows. Specifically, *local traffic* is the traffic between two subscribers in the local market; whereas *remote traffic* is the traffic between local subscriber and non-local subscribers in the Internet. Conventionally, only local traffic is allowed to be exchanged through the peering link and this traffic exchange is beneficial to both ISPs. A previous paper [5] studied ISP peering in a network with only local traffic, providing a good understanding of peering strategies based on local traffic only. In this paper, we focus on the effects of remote traffic on local peering decisions. In particular, we focus on *incoming* remote traffic generated by local subscribers who want to download objects from the Internet. In local markets, incoming

[1]We have learned about the problem from talking to ISPs in our region.

remote traffic tends to dominate outgoing remote traffic and the normal settlement model for provider links is based on the maximum of the two directional traffic. So this assumption is reasonable to simplify the analysis.

The remote traffic can be further classified into two categories. The traffic that can not be routed across ISPs at application layer is defined as *private demand*. For this part of the traffic, inter-domain routing is determined by ISPs' policy configuration and the inter-domain routing protocol. Therefore private demand portion of the remote traffic does not traverse the peering link. For example, the remote web/email traffic is by and large private demand. Moreover, if a remote content object is only interested by subscribers from the same ISP, the traffic generated by subscribers downloading this object is also a kind of private demand. Although such traffic is also generated by P2P applications, it behaves the same as web/email traffic since the other ISP does not need it and thus the traffic cannot traverse through the peering link. The intensity of private demand is denoted by D_i for ISP_i.

The remaining traffic, which is generated by remote P2P contend objects needed by both ISPs, can be potentially routed at the application layer. A subscriber j from ISP_B may download a remote object r; then those subscribers from the other local ISP (*i.e.* ISP_A) may choose to download r either directly from the Internet, or from ISP_B, depending on the *peer selection algorithm* of the P2P application. In the latter case, the traffic will traverse the local peering link. We define this pubic P2P traffic demand as *application layer routed demand*, or simply *public demand*. The intensity of public demand is denoted by D_0.

In order to study the economic implications of these traffic on ISPs' peering strategies, we must *model the aggregate routing behavior* of the P2P applications so that we can predict the routing for this public demand. We also need to develop a methodology to analyze the economic consequences of the traffic demands and ISPs peering and provisioning decisions.

A. Routing Model of Public Demand

First, let us discuss all possible ways for two ISPs to download a content object that belongs to the public demand.

1) Both ISPs download the content object via their own private links.
2) One ISP downloads the content object via its own private link, and then the other ISP downloads the object via the peering link.

Different applications/peers may use different metrics and algorithms to optimize their routing decisions. The essence of the aggregate behavior, we believe, can be captured by an optimization model, as if the peers in the P2P application try to optimize the performance of transferring each byte of their traffic.

1) Routing Model: We define the following partitions of the public demand D_0: x_i $(i = 1, 2)$ is the fraction that goes to ISP_i only; and y_i $(i = 1, 2)$ is the fraction that goes to ISP_i first and is subsequently downloaded by peers in ISP_j $(j = 1, 2$ and $j \neq i)$ through the peering link. Obviously, x_i and y_i are parameters that reflect the routing of the public demand in this network.

Let ρ_i $(i = 1, 2)$ be the traffic intensity on ISP_i's private link, and ρ_{ij} be the traffic intensity on the peering link from ISP_i to ISP_j. We know that $\rho_i = D_i + D_0(x_i + y_i)$ and $\rho_{ij} = D_0 y_i$.

In our model, we define the goal of application layer routing as to determine x_i and y_i to optimize the performance of transferring each byte of their demand. Let $f(\rho, c)$ be some performance metric that measures the transferring performance on a link with capacity c when the traffic intensity on this link is ρ. Since the flows on different links affect different number of subscribers and different amount of traffic, we need to give the transferring performance of each link a weighting factor $w(\alpha, \rho)$, where α represents the number of subscribers affected and ρ denotes the traffic intensity over this link. Thus, we add up the weighted performance of all links in the local network and get the following expression. It is used to measure the performance of transferring all traffic in the whole network.

$$\begin{aligned} & w(\alpha_1, \rho_1) f(\rho_1, c_1) + w(\alpha_2, \rho_2) f(\rho_2, c_2) \\ + \quad & w(\alpha_2, \rho_{12}) f(\rho_{12}, c_{12}) + w(\alpha_1, \rho_{21}) f(\rho_{21}, c_{21}) \end{aligned} \quad (1)$$

P2P applications optimize Equation (1) subject to the following demand constraints and capacity constraints:

$$\begin{aligned} \text{demand constraint} \quad & x_i + y_1 + y_2 = 1, (i = 1, 2) \\ \text{capacity constraints} \quad & \rho_i = D_i + D_0(x_i + y_i) \leq c_i, \\ & \rho_{ij} = D_0 y_i \leq c_{ij}. \end{aligned} \quad (2)$$

For this paper, we assume

$$f(\rho, c) = c - \rho, \quad (3)$$

which is the available bandwidth of the link, to measure the performance of transferring each byte of the traffic on the link. The weighting factor is defined as

$$w(\alpha, \rho) = \alpha\rho \quad (4)$$

where ρ is the traffic volume on the link and α is the percentage of subscribers involved. Roughly speaking, $\alpha\rho(c - \rho)$ can be interpreted as the sum of the performance of transferring each byte for all traffic and all related subscribers on the link. The overall routing objective of P2P applications is to optimize the transferring performance for the traffic on all different links in the network. It is a reasonable assumption although each peer in a P2P application does not solve this optimization problem directly.

Furthermore, from Equation (2), we know that $x_i = 1 - y_1 - y_2$, which implies $x_1 = x_2$. It is obviously true according to our definition. However, it does not mean that two ISPs download the same amount of the public demand since the portion of public demand ISP_i needs to download is $x_i + y_i$ and generally $y_1 \neq y_2$.

Let us use D_0 as the unit to measure the capacity or traffic intensity of different links, thus the intensity of the public demand can be considered as 1 and the intensity of

ISP_i's private demand can be defined as $\lambda_i = \frac{D_i}{D_0}$. The last optimization problem can be refined to the following form:

$$\begin{array}{ll} \mathbf{min}_{y_1,y_2} & \alpha_2(y_1 - y_1^{'})^2 + \alpha_1(y_2 - y_2^{'})^2, \\ \mathbf{subject\ to} & l_i \le y_i \le h_i, \\ & y_1 + y_2 \le 1, \end{array} \quad (5)$$

where

$$\begin{aligned} y_1^{'} &= \frac{c_{12} - c_2 + 2 + 2\lambda_2}{4}, \\ y_2^{'} &= \frac{c_{21} - c_1 + 2 + 2\lambda_1}{4}, \\ h_i &= \min\ (c_{ij},\ 1), \\ l_i &= \max\ (\lambda_j + 1 - c_j,\ 0). \end{aligned}$$

2) Analysis of the Routing Solution: From Equation (5), we can see that the objective function is a convex function in a three-dimensional space. As shown in Figure 3, the feasible region is confined by the rectangle ($y_i \in [l_i, h_i]$) and the line ($y_1 + y_2 = 1$), indicated by the shaded area. Although the solution cannot be expressed by a clean closed form equation, where the optimal solution lands in the feasible region has clear physical meanings. We define the solution of the above optimization problem in Equation (5) as (y_1^*, y_2^*), and discuss different possible scenarios as follows: [2]

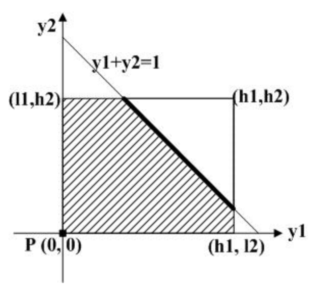

Fig. 3. Feasible Region.

1) If $(y_1^{'}, y_2^{'})$ satisfies all constraints, then $(y_1^*, y_2^*) = (y_1^{'}, y_2^{'})$. In this case, the optimal solution is also the global optimal solution of the unconstrained problem. If the ISPs are cooperative, they can collectively adjust their link capacities to achieve this global optimal result.
2) If $y_1+y_2 = 1$ is an active constraint, the solution lies on the bold line in the figure and it means all public demand would traverse only one of the private links. Whatever is downloaded by ISP_i would then be downloaded by ISP_j ($j \neq i$) over the peering link. This is achieved only when the capacity of the peering link is large enough for the two ISPs to exchange public demand, so we refer to this scenario as "*Unlimited Capacity Peering*".
3) If $y_1^* = y_2^* = 0$, which is indicated by the point P in the figure, it means there is no traffic exchange on the peering link. We refer to this scenario as "*No Peering*".

[2]Note the location of the rectangle and the line depends on the parameters, and this figure is not exact and it is only used to illustrate different scenarios.

4) If $y_1+y_2 = 1$ is not an active constraint and $y_1^* = y_2^* = 0$ is not true, the solution should be the point within the shaded rectangle and *closest* to $(y_1^{'}, y_2^{'})$. We also have $y_1^*+y_2^* \in (0, 1)$, which means the capacity of the peering link is not large enough so that two ISPs cannot share the public demand efficiently. We refer to this scenario as "*Limited Capacity Peering*".

The routing model predicts the aggregate behavior of the application layer routing for different traffic intensities (*i.e.* λ_1 and λ_2), and ISP peering (*i.e.* c_{12} and c_{21}) and provisioning (*i.e.* c_1 and c_2) agreements.

B. Methodology for ISP Economic Analysis

In this study we view each local ISP to be simply in the business of providing transit service, and compete with other ISPs in similar positions for subscriber share.

First, each ISP needs to pay its providers for the traffic on their private links. In real life, the cost structure for an ISP can be quite complicated. The dependency on traffic is often not exact, but based on some measured percentile of traffic intensity over a period of time. There is usually also a cost associated with the capacity and guaranteed traffic minimum. In order to simplify the analysis, we assume a constant cost q per unit intensity of traffic and this cost is the same for all ISPs in the same local market.

To cover its cost (and to make some money), ISP_i needs to charge its subscribers. As it is quite common with current ISP pricing, we assume each ISP uses flat rate charging and charges a fixed price p_i. It is possible that $p_1 \neq p_2$.

The net income of ISP_i is therefore

$$R_i = p_i n \alpha_i - q\rho_i$$

wherein $p_i n \alpha_i$ is the revenue from its subscribers and $q\rho_i$ is the transmission cost. If $R_i < 0$, ISP_i cannot survive economically. We define the minimum price that ISP_i needs to charge its subscribers in order to survive as ISP_i's *break-even price*, which is denoted as p_i^*. From the last equation, ISP_i's break-even price is given by

$$p_i^* = \frac{q}{n\alpha_i}\rho_i. \quad (6)$$

If one ISP's break-even price becomes lower after peering, we say it benefits from peering. We use the decrease in the break-even price to measure the benefit an ISP achieves from the peering. It is also an ISP's concern to solidify its market position by driving towards a lower break-even price than its competitors. In the two-ISP model, ISP_1 would like to minimize $p_1^* - p_2^*$, while ISP_2 would like to minimize $p_2^* - p_1^*$ in order to be more competitive in the market.

In the rest of paper, we use p_i^* to denote the break-even price of ISP_i after peering with limited capacity. In particular, $p_i^*(0)$ is the break-even price of ISP_i under no peering scenario; and $p_i^*(\infty)$ is the break-even price under unlimited capacity peering scenario.

III. Analysis of Network with Only Application Layer Routed Traffic

The methodology described in the last section can be readily applied to analyze the implication of the P2P traffic on ISPs. Since the routing of private demand is determined by policy-based routing, the routing of public demand in a network with private demand (λ_1, λ_2) and private capacity (c_1, c_2) is the same as that in a network without any private demand and with private capacity ($c_1 - \lambda_1$, $c_2 - \lambda_2$). Therefore, in this section we only analyze the scenario with $\lambda_1 = \lambda_2 = 0$. This simplifies the analysis by reducing the number of parameters and also lets us concentrate on the effect of the application layer routing. The optimization problem in Equation (5) thus becomes:

$$\begin{aligned}
\mathbf{min}_{y_1,y_2} \quad & \alpha_2(y_1 - \frac{c_{12} - c_2 + 2}{4})^2 \\
& + \alpha_1(y_2 - \frac{c_{21} - c_1 + 2}{4})^2, \\
\mathbf{subject\ to} \quad & max(0, 1 - c_j) \leq y_i \leq min(1, c_{ij}), \\
& y_1 + y_2 \leq 1.
\end{aligned} \tag{7}$$

Next, we analyze the economic implications in each of the following scenarios: (a) *no peering*, (b) *unlimited capacity peering*, and (c) *limited capacity peering*.

A. Scenario Analysis: No Peering

From Section II-A.2, we recall that the routing solution $y_1^* = y_2^* = 0$ means that there is no exchange of traffic over the peering link. For the no peering scenario to occur, the network must satisfy the following two conditions[3]:

$$c_i \geq 1 \qquad i = 1, 2 \tag{8}$$

$$c_{ij} = 0 \quad \text{or} \quad c_i - c_{ji} \geq 2 \qquad i, j = 1, 2 \quad i \neq j \tag{9}$$

Equation (8) guarantees feasibility, namely each ISP can independently satisfy its subscriber demand respectively. Equation (9) means either there is no peering link, or the peering link has significantly lower capacity than the private links such that there is no local traffic exchange under optimal routing. In the no peering scenario, we can make the following observation about the economic positions of the two ISPs:

Proposition 1: *In a network with only public demand, an ISP's break-even price is inversely correlated to its market share when there is no local peering. Specifically, the ISP with the highest market share has the lowest break-even price.*

In economics jargon, we have the *economy of scale*. Due to the caching effect of P2P applications, the ISP with more market share is more efficient in making use of its private link to satisfy its subscribers, and as a result has a more competitive market position.

Proof: From Equation (6), we have

$$p_i^*(0) = \frac{q}{n\alpha_i}\rho_i = \frac{q}{n\alpha_i}D_0, \tag{10}$$

and observe that $p_i^*(0)$ is a decreasing function of α_i. ∎

[3]The derivation of the conditions and the solutions for different scenarios can be found in the more complete technical report [6].

Note that the statement above does not necessarily mean ISPs with higher market share prefer no peering. As analyzed in [5], as a way to provide short circuits for local traffic, local ISPs may choose to set up free peering to reduce the cost of transiting local traffic over the private links, resulting in benefits for both ISPs. Now let us consider the possible effect of the local peering in a network with only application layer routed traffic.

Proposition 2: *In a network with only application layer routed traffic, local peering improves the overall efficiency of the peering ISPs, and each ISP is always "better off" or "equal to" before.*

Proof: Let us consider two ISPs as one network. Before peering, the cost of this network is $q(D_0 + D_0) = 2qD_0$. After peering, the cost of this network is $qD_0(x_1 + y_1 + x_2 + y_2) = q(1 + x_1)D_0 = q(1 + x_2)D_0 \leq 2qD_0$. Although the local peering incurs some operating cost to the peering ISPs, it is usually not significant, and we can ignore the cost here. We see that the cost of this network is reduced, so the overall efficiency of this network is better than before. Furthermore, since

$$p_i^* - p_i^*(0) = \frac{q}{n\alpha_i}(x_i + y_i - 1)D_0 = \frac{q}{n\alpha_i}(-y_j)D_0 \leq 0,$$

So the break-even price would be lowered for both ISPs after peering if $y_i^* \neq 0$. ∎

This proposition shows the free peering is still beneficial for the whole network in a network with only application layer routed traffic. However, the proof also shows the break-even price reduction ($\frac{q}{n\alpha_i}y_j$) for the two ISPs might be different. Under certain conditions, we may have the routing solution $(y_1^*, y_2^*) = (1, 0)$. It means all the public demand would become traffic for ISP_1's private link, so subscribers of ISP_2 always download the remote objects from subscribers of ISP_1 with the same interests. Clearly, peering has not helped ISP_1 at all while ISP_2 takes full advantages of ISP_1's committed bandwidth without having to pay for it. In this extreme case, clearly there is no incentive at all for ISP_1 to peer with ISP_2.

This extreme case highlights some of the implications of application layered routed traffic on ISP peering agreements. It gives an intuitive explanation of why the recent trend of increasing P2P traffic in ISP network is causing some ISPs to revise their peering relationships. In the following subsections, we discuss the distribution of peering benefit under different scenarios, which can shed light on how an ISP should make peering decisions.

B. Scenario Analysis: Unlimited Capacity Peering

In this scenario, $y_1^* + y_2^* = 1$. In other words, each remote object downloaded by ISP_1 is shared with subscribers in ISP_2 and vice versa. The routing problem in Equation (7) can be transformed to the following equivalent problem:

$$\begin{aligned} \mathbf{min}_{y_1} \quad & (y_1 - \frac{\alpha_2(c_{12}-c_2) - \alpha_1(c_{21}-c_1) + 2}{4})^2, \\ \textbf{subject to} \quad & y_{1l} \le y_1 \le y_{1h}, \end{aligned} \tag{11}$$

where y_{1l} and y_{1h} are x-coordinates of two intersection points of the line $y_1 + y_2 = 1$ and the rectangle in Figure 3. In this simplified form, the optimal solution can be easily derived:

$$y_1^* = \begin{cases} \frac{1}{2} + t & : \text{ if } t \in [y_{1l}, y_{1h}] \\ y_{1l} & : \text{ if } t < y_{1l} \\ y_{1h} & : \text{ if } t > y_{1h} \end{cases} \tag{12}$$

where $t = \frac{\alpha_2(c_{12}-c_2) - \alpha_1(c_{21}-c_1)}{4}$. The break-even prices of two ISPs are:

$$\begin{aligned} p_1^*(\infty) &= \frac{q}{n\alpha_1}(1 - y_2^*) = \frac{q}{n\alpha_1}(\frac{1}{2} + t), \\ p_2^*(\infty) &= \frac{q}{n\alpha_2}(1 - y_1^*) = \frac{q}{n\alpha_2}(\frac{1}{2} - t). \end{aligned} \tag{13}$$

Based on the above routing solution, we can answer the following questions about the profitability and survivability of the ISPs.

Proposition 3: *In order for ISP_1, the ISP with the larger market share, to benefit more than ISP_2 from peering, the necessary condition in Equation (14) must be satisfied.*

$$\alpha_1(c_1 - c_{21}) < \alpha_2(c_2 - c_{12}) - 2(\alpha_1 - \alpha_2). \tag{14}$$

In order for ISP_1 to maintain a lower break-even price after peering, the necessary condition in Equation (15) must be satisfied.

$$\alpha_1(c_1 - c_{21}) < \alpha_2(c_2 - c_{12}) + 2(\alpha_1 - \alpha_2). \tag{15}$$

Both Equation (14) and Equation (15) are expressed in terms of $(c_i - c_{ji})$, the difference of incoming links capacity for a given ISP. Both conditions require a certain parity of incoming link capacity differences between different competing ISPs.

Before peering, ISP_1 has a lower break-even price than that of ISP_2, so that sometimes ISP_1 may still maintain a lower break-even price after peering even if ISP_2 achieves more benefit from peering than ISP_1. In this case, peering might still be a reasonable option for ISP_1 since it does not cause ISP_1 to lose its break-even pricing advantage over ISP_2. As expected, the condition of Equation (14) to extract more benefit from peering is more stringent than the condition of Equation (15) to maintain price advantage, since $2(\alpha_1 - \alpha_2) > 0$.

We enumerate the different possible outcomes after unlimited capacity peering as follows.

1) ISP_1 maintains a lower break-even price than ISP_2, and achieves more benefits than ISP_2 from peering,
2) ISP_1 maintains a lower break-even price than ISP_2, but ISP_2 achieves more benefits than ISP_1 from peering,
3) ISP_1 achieves less benefits from peering, and has a higher break-even price than ISP_2.

This order is also the preference for ISP_1. To visualize, let us assume $\alpha_1 = 3/5$ and depict the outcomes under different values of $(c_1 - c_{21})$ and $(c_2 - c_{12})$ in Figure 4.

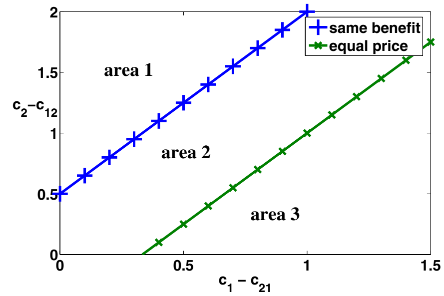

Fig. 4. Different Outcomes of Unlimited Capacity Peering.

From Equation (13) we see that there are at least two issues with unlimited capacity peering:

- It is in each ISP's selfish interest to lower the provisioning of its private link and uploading capacity of the peering link, in order to benefit more from the unlimited capacity peering. For a given c_2 and c_{21}, ISP_1 can successively lower c_1 or c_{12} to achieve better and better outcome for itself. From an overall network point of view, this tends to discourage network growth and service improvement. After a point, the subscriber demand cannot be satisfied.
- Each ISP's destiny may be affected by the other ISP's provisioning. For example, assuming c_1 is fixed by ISP_1 and two ISPs have to negotiate the capacity of the peering link. When ISP_2 lower c_2, the outcome for ISP_1 may still drop from area 1 to 2, to 3. The reason for this lost of control is because application layer routed traffic under unlimited peering makes the two ISP networks behave like a single network. Note that before P2P traffic becomes dominant, we did not have this problem with web traffic. With only web/email traffic, ISPs have more control over traffic routing through policy-based routing.

Limiting the peering link capacity c_{12} and c_{21} might be one of the remedies available to ISPs. When the capacity of the peering link is sufficiently reduced, we have the routing solution $y_1^* + y_2^* < 1$, thus limiting the usage of an ISP's private link by other ISPs. This scenario is studied in the next subsection.

C. Scenario Analysis: Limited Capacity Peering

In this case, the optimization problem in Equation (7) can be transformed to the following format:

$$\begin{aligned} \mathbf{min}_{y_1,y_2} \quad & \alpha_2(y_1 - \frac{c_{12} - c_2 + 2}{4})^2 \\ + \quad & \alpha_1(y_2 - \frac{c_{21} - c_1 + 2}{4})^2, \\ \textbf{subject to} \quad & l_i \le y_i \le h_i. \end{aligned} \tag{16}$$

The optimal solution is

$$y_1^* = \begin{cases} y_1^{'} & : \text{ if } y_1^{'} \in [l_1, h_1] \\ l_1 & : \text{ if } y_1^{'} < l_1 \\ h_1 & : \text{ if } y_1^{'} > h_1. \end{cases} \quad (17)$$

We can solve for y_2^* in the same way as y_1^*. And the break-even prices of two ISPs are:

$$p_1^*(\infty) = \frac{q}{n\alpha_1}(1 - y_2^*) = \frac{q}{n\alpha_1}(\frac{2 - c_{21} + c_1}{4}),$$
$$p_2^*(\infty) = \frac{q}{n\alpha_2}(1 - y_1^*) = \frac{q}{n\alpha_2}(\frac{2 - c_{12} + c_2}{4}). \quad (18)$$

Similar to the unlimited capacity peering scenario, the following proposition gives the ISPs economic predicaments in the limited capacity peering scenario:

Proposition 4: *In order for ISP_1, the ISP with the larger market share, to benefit more than ISP_2 from limited capacity peering, the necessary condition in Equation (19) must be satisfied.*

$$\alpha_2(c_1 - c_{21}) < \alpha_1(c_2 - c_{12}) - 2(\alpha_1 - \alpha_2). \quad (19)$$

In order for ISP_1 to maintain a lower break-even price after peering with limited capacity, the condition in Equation (20) must be satisfied.

$$\alpha_2(c_1 - c_{21}) < \alpha_1(c_2 - c_{12}) + 2(\alpha_1 - \alpha_2). \quad (20)$$

Again, in order for ISP_1 to maintain certain economic advantages, it still has the incentive to lower its provisioning and it must try to maintain certain parity of the difference in incoming link capacity with that of its competitor. However, in this case the conditions are further relaxed. For example, condition in Equation (19) is a more relaxed version of the condition in Equation (14) since $\alpha_1 \geq \alpha_2$. This reveals the fact that it is easier for ISP_1 to maintain economic advantages under limited capacity peering scenario than unlimited capacity peering scenario. It also explains part of the reason why large ISPs want to limit their peering links capacity.

Furthermore, from Equation (18), we can see that with limited peering, each ISP's break-even price can be determined by its own decision and would not be affected by the other ISP's private link provisioning. This advantage is also likely to make ISPs to favor limited capacity peering to gain more control of their own fate.

IV. Numerical Example

In this section, we illustrate and explain the results in Section III through a simple numerical example. We assume $\lambda_1 = \lambda_2 = 0$, $\alpha_1 = \frac{3}{5}$, $\alpha_2 = \frac{2}{5}$, $c_2 = 0.9$, and $c_{21} = 1.5$. All these parameter are fixed. We will study ISP_1's provisioning in this section.

A. Private Link Provisioning

Let us further assume $c_{12} = c_{21} = 1.5$. We will focus on ISP_1's private link provisioning and study how c_1 affects ISP_1's break-even price. The routing problem under these parameters is:

$$\begin{aligned} \textbf{min}_{y_1,y_2} \quad & 0.4(y_1 - 0.65)^2 + 0.6(y_2 - \frac{3.5 - c_1}{4})^2, \\ \textbf{subject to} \quad & 0.1 \leq y_1 \leq 1, \\ & \max(0, 1 - c_1) \leq y_2 \leq 1, \\ & y_1 + y_2 \leq 1. \end{aligned} \quad (21)$$

Given $c_2 = 0.9$, we must have $c_1 \geq 0.1$ in order to satisfy the total traffic demand. We solve this problem and plot p_1^* and p_2^* in Figure 5 for $c_1 \geq 0.1$. It shows ISP_1's beak-even price p_1^* continues to increase as c_1 increases. This is expected since ISP_1's operating cost increases. However, ISP_2's break-even price p_2^* decreases as c_1 increases since ISP_2 also takes advantages of ISP_1's increased private link capacity without any cost, due to the routing effect of public P2P demand. In real world, this phenomenon may not be that apparent because public demand is not the only traffic in the network. However, with the increasing popularity and intensity of P2P traffic, it would have more influence on ISPs' operation.

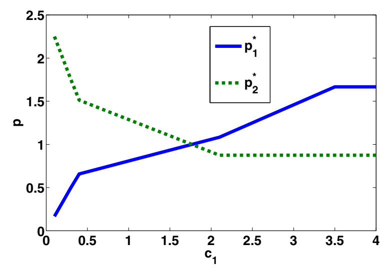

Fig. 5. Break-even Price of two ISPs as we Vary c_1.

B. Peering Link Provisioning

Now let us fix $c_1 = 1.1$ and study the influence of different c_{12} on two ISPs' break-even price. The routing problem under these parameters is:

$$\begin{aligned} \textbf{min}_{y_1,y_2} \quad & 0.4(y_1 - \frac{c_{12} + 1.1}{4})^2 + 0.6(y_2 - 0.6)^2, \\ \textbf{subject to} \quad & 0 \leq y_2 \leq 1, \\ & 0.1 \leq y_1 \leq \min(1, c_{12}), \\ & y_1 + y_2 \leq 1. \end{aligned} \quad (22)$$

After solve this optimization problem, we plot p_1^* and p_2^* in Figure 6. From it we can see ISP_1's break-even price increases while ISP_2's break-even price decreases as ISP_1 increases the capacity of the peering link from ISP_1 to ISP_2. Because of the "free-ride" phenomenon, ISP_2 can take more advantages from the peering link when c_{12} increases.

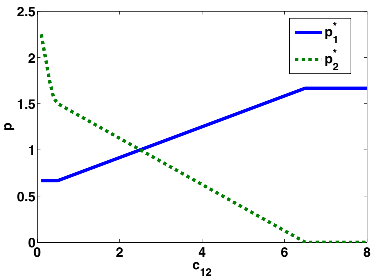

Fig. 6. Break-even Price of two ISPs as we Vary c_{12}.

This example shows that an ISP must be very careful when provisioning its private links and peering links. Otherwise, other ISPs will be able to *free-ride* and gain an economic advantage over it.

V. Related Work

Due to the inter-dependency of the economic principles and the technical architecture of the Internet operations, there has been increasing interest in applying economic and game-theoretic analysis to study Internet protocol and operational issues. In [7], the authors gave an interesting analysis of pricing strategies by local ISPs competing for the same customers. In [8], the problem of where a pair of ISPs will make their peering connections so as to minimize their own transit traffic is formulated and studied based on a game-theoretical model. Furthermore, authors in [4] explored the interactions between the routing behaviors of P2P networks and the underlay network, revealing the properties and implications of inefficient routing equilibria. All of these studies apply economic and game-theoretic analysis to networking problems and generate nice insights to the design of networking protocols and operational principles.

Over the last few years, overlay networks have become increasingly popular as a promising platform to provide customizable and reliable services at the application layer [9] [10]. In [11], the author pointed out that overlay networks attempt to take control over routing in the hope that they might achieve better performance and illustrated how an uncoordinated effort of the two layers to recover from failures may cause performance degradation for both overlay and non-overlay traffic. It showed how current traffic engineering techniques are inadequate to deal with emerging overlay network services. As an important type of overlay networks, the rise of P2P applications in broadband networks also draws a lot of attentions. It has been already reported and discussed academically as well as in popular press [12][13][14][15]. There are also various efforts in trying to characterize such P2P traffic and studying the caching behavior [16][17].

However, the study on economic implications of these overlay traffic, especially P2P traffic, on ISPs' peering decisions has so far received less attention. In [5], the authors propose some important free peering conditions based on their traffic model and economic model. But their traffic model only considers local web traffic. Although it also applies if P2P objects are located in the local network, their analysis does not include the influence of remote P2P traffic. As far as we know, we are the first to study the impact of P2P traffic's application layer routing on ISPs' peering and provisioning strategies.

VI. Conclusion and Future Work

The shift from primarily web and email traffic to a predominantly P2P workload is a relatively abrupt and new phenomenon. Since there is a growing trend that users are adopting better technologies to gain broadband access, P2P traffic will be more prominent. The implications of this shift on the traffic engineering of Internet operation, as presented in this paper, have not been investigated before. Therefore, this research not only provides the fundamental understanding on why there is such a shift on Internet operation, but also open doors for potential research on P2P traffic management, peering relationship establishment policy, as well as routing decision between autonomous systems.

The contributions of this paper include (a) reveal the phenomenon that inter-domain routing protocol cannot control the traffic routed at application layer, (b) show that peering relationship with unlimited capacity (or very high capacity) discourages network growth and service improvement if P2P traffic dominates, (c) show that in network with dominant P2P traffic, peering relationship with unlimited capacity may make two ISP networks behave like a single network, so that each ISP looses complete control of transit traffic routing through its own part of the network, (d) show that while peering with limited capacity may reduce ISPs' benefit from economy of scale, but it can remove or minimize the above "fairness" problems in distributing the peering benefits.

In this work, we focus on application layer multicast or streaming type of P2P traffic. In fact, there are many other kinds of traffic that may influence the traffic engineering policy and peering decision. For example, in order to optimize the performance or save money, some people set up "proxies" to redirect their inter-domain traffic to go through one ISP that is different from the default BGP routing. Such traffic is also a kind of "application layer routed" traffic. Another example of application layer routing is routing overlay networks. We plan to study the implication of other kinds of application layer routed traffic and potential traffic engineering solutions to the above problems in our future work.

Acknowledgments

The work described in this paper was partially funded by RGC Project no. 411505. We also thank Mr Echo Chan of Hutchison Global Crossing for helpful discussions that lead to the study of this problem.

References

[1] G. Huston, "Interconnection, peering and settlements," *inet*, 1999, available at http://www.isoc.org/inet99/proceedings/1e/1e_1.htm.

[2] W. B. Norton, "Internet service provider and peering," in *NANOG*.

[3] T. Roughgarden and E. Tardos, "How bad is selfish routing?" in *IEEE Symposium on Foundations of Computer Science*, 2000, pp. 93–102.

[4] W. Jiang, D. M. Chiu, and J. C. Lui, "On the interaction of multiple overlay routing," in *Proceeding of PERFORMANCE 2005*, Juan-les-Pins, France, Oct. 2005.

[5] M. B. Weiss and S. J. Shin, "Internet interconnection economic model and its analysis: Peering and settlement," *Netnomics*, vol. 6, no. 1, pp. 43–57, 2004, available at http://ideas.repec.org/a/kap/netnom/v6y2004i1p43-57.html.

[6] J. H. Wang, D. M. Chiu, and J. C. Lui, "Economics implications of p2p traffic on isp peering," http://personal.ie.cuhk.edu.hk/∼hwang3/outside/tech.pdf, Tech. Rep., 2006.

[7] S. Shakkottai and R. Srikant, "Economics of network pricing with multiple isps," in *IEEE Infocom*, Miami, 2005.

[8] R. Johari and J. Tsitsiklis, "Routing and peering in a competitive internet," MIT, Tech., 2003. [Online]. Available: http://citeseer.ist.psu.edu/johari03routing.html

[9] B.-G. Chun, R. Fonseca, I. Stoica, and J. Kubiatowicz, "Characterizing selfishly constructed overlay networks," in *Proceedings of the 23rd IEEE International Conference on Computer Communications (INFOCOM 2004)*, Hong Kong, 2004.

[10] Y. Zhu and B. Li, "Overlay networks with linear capacity constraints." in *IWQoS*, 2005, pp. 21–36.

[11] R. Keralapura, N. Taft, C.-N. Chuah, and G. Iannaccone, "Can isps take the heat from overlay networks?" in *ACM HotNets-III*, 2004.

[12] J. Glasner, *P2p fuels global bandwidth binge*. Wired News, Apr. 2005.

[13] B. Godfrey, K. Lakshminarayanan, S. Surana, R. Karp, and I. Stoica, "Load balancing in dynamic structured P2P systems," in *Proc. IEEE INFOCOM*, Hong Kong, 2004.

[14] K. Tutschku and P. Tran-Gia, *Peer-to-Peer-Systems and Applications*. Springer, 2005, traffic Characteristics and Performance Evaluation of Peer-to-Peer Systems.

[15] D. Xu, M. Hefeeda, S. Hambrusch, and B. Bhargava, "On peer-to-peer media streaming," in *Proc. of IEEE International Conference on Distributed Computing Systems (ICDCS'02)*, Vienna, Austria, 2002, pp. 363–371.

[16] K. Gummadi, R. Dunn, S. Saroiu, S. Gribble, H. Levy, and J. Zahorjan, "Measurement, modeling, and analysis of a peer-to-peer file-sharing workload," 2003. [Online]. Available: citeseer.ist.psu.edu/article/gummadi03measurement.html

[17] S. Sen and J. Wang, "Analyzing peer-to-peer traffic across large networks," *IEEE/ACM Trans. Netw.*, vol. 12, no. 2, pp. 219–232, 2004.

APPENDIX

APPENDIX: PROOF OF PROPOSITIONS

Proof of Proposition 3:

From Equation 13, we know the break-even price after unlimited capacity peering is:

$$p_1^*(\infty) = \frac{q}{n\alpha_1}(\frac{1}{2}+t), \qquad p_2^*(\infty) = \frac{q}{n\alpha_2}(\frac{1}{2}-t).$$

This allows us to write down the benefits for the ISPs as:

$$p_1^*(0) - p_1^*(\infty) = \frac{q}{n\alpha_1}(1-(x_1+y_1)) = \frac{q}{n\alpha_1}(\frac{1}{2}-t),$$
$$p_2^*(0) - p_2^*(\infty) = \frac{q}{n\alpha_2}(1-(x_2+y_2)) = \frac{q}{n\alpha_2}(\frac{1}{2}+t).$$

To let ISP_1 benefit more than ISP_2, we must have:

$$\frac{q}{n\alpha_1}(\frac{1}{2}-t) > \frac{q}{n\alpha_2}(\frac{1}{2}+t) \quad \Rightarrow \quad t < \frac{\alpha_2-\alpha_1}{2}$$
$$\Rightarrow \quad \alpha_1(c_1-c_{21}) < \alpha_2(c_2-c_{12}) - 2(\alpha_1-\alpha_2).$$

In order for ISP_1 to have a lower break-even price, we must have:

$$\frac{q}{n\alpha_1}(\frac{1}{2}+t) < \frac{q}{n\alpha_2}(\frac{1}{2}-t) \quad \Rightarrow \quad t < \frac{\alpha_1-\alpha_2}{2}$$
$$\Rightarrow \quad \alpha_1(c_1-c_{21}) < \alpha_2(c_2-c_{12}) + 2(\alpha_1-\alpha_2).$$

■

Proof of Proposition 4:

From Equation 18, we know the break-even price after limited capacity peering is

$$p_1^*(\infty) = \frac{q}{n\alpha_1}(\frac{2-c_{21}+c_1}{4}), \quad p_2^*(\infty) = \frac{q}{n\alpha_2}(\frac{2-c_{12}+c_2}{4}).$$

This allows us to write down the benefits for the ISPs as:

$$p_1^*(0) - p_1^* = \frac{q}{n\alpha_1}(1-(x_1+y_1)) = \frac{q}{n\alpha_1}\frac{c_{21}-c_1+2}{4},$$
$$p_2^*(0) - p_2^* = \frac{q}{n\alpha_2}(1-(x_2+y_2)) = \frac{q}{n\alpha_2}\frac{c_{12}-c_2+2}{4}.$$

To let ISP_1 benefit more than ISP_2, we must have:

$$\frac{q}{n\alpha_1}\frac{c_{21}-c_1+2}{4} > \frac{q}{n\alpha_2}\frac{c_{12}-c_2+2}{4}$$
$$\Rightarrow \quad \alpha_2(c_1-c_{21}) < \alpha_1(c_2-c_{12}) - 2(\alpha_1-\alpha_2).$$

In order for ISP_1 to have a lower break-even price, we must have:

$$\frac{q}{n\alpha_1}\frac{2-c_{21}+c_1}{4} < \frac{q}{n\alpha_2}\frac{2-c_{12}+c_2}{4}$$
$$\Rightarrow \quad \alpha_2(c_1-c_{21}) < \alpha_1(c_2-c_{12}) + 2(\alpha_1-\alpha_2).$$

■