PAPER

The Implications of Overlay Routing for ISPs' Peering Strategies*

Xun SHAO†a), *Student Member*, Go HASEGAWA††b), Yoshiaki TANIGUCHI††c), *and* Hirotaka NAKANO††d), *Members*

SUMMARY The Internet is composed of many distinct networks, operated by independent Internet Service Providers (ISPs). The traffic and economic relationships of ISPs are mainly decided by their routing policies. However, in today's Internet, overlay routing, which changes traffic routing at the application layer, is rapidly increasing and this challenges the validity of ISPs' existing agreements. We study here the economic implications of overlay routing for ISPs, using an ISP interconnection business model based on a simple network. We then study the overlay traffic patterns in the network under various conditions. Combining the business model and traffic patterns, we study the ISPs' cost reductions with Bill-and-Keep peering and paid peering. We also discuss the ISPs' incentive to upgrade the network under each peering strategy.

key words: *ISP, peering, overlay routing, routing game, Nash bargaining solution*

1. Introduction

The Internet is composed of many distinct networks, operated by independent ISPs. There are primarily two kinds of relationships among ISPs: transit and peering [2]. In a transit relationship, a traffic-originating provider pays a transit provider for the traffic destined for locations outside the originator's local network. On the other hand, in a peering relationship, only traffic between the two peering ISPs and their respective customer ISPs can be exchanged on the peering link. Such traffic exchange on a peering link helps both peering ISPs to reduce their dependence on transit providers and thus reduce monetary costs. In today's Internet, peering relationships are mostly "Bill-and-Keep (BK)" [3] due to ease of implementation. In this arrangement, the peering providers do not charge each other for the traffic on the peering links. There are other kinds of peering relationships in which ISPs make an agreement to charge for traffic [4], [5].

Various aspects of peering settlement have been analyzed in the literature [4], [6]–[8]. Laffont et al. [6] made the first in-depth analysis of ISP peering from an economic perspective. They looked at the impact of symmetric access charge on the strategies of providers and showed that operators set prices for their customers as if their customers' traffic were entirely off-net. Shakkottai et al. [7] extended the model of Laffont et al. [6] to include the geographical locations of ISPs, and analyzed local ISP interactions separately from distant and transit ISP interactions. Shrimali et al. [8] use a different model of symmetric ISP peering. They focus on the equilibrium of early exit routing and late exit routing, and gives the characteristics of the Nash equilibrium [9] and the corresponding conditions. Shrimali et al. [4] use a more general asymmetric peering network, and look at how ISPs could charge each other in response to the externalities caused by their traffic strategies.

All the above research is based on the routing policy with focus on business considerations and does not consider the performance of the networks and services for subscribers and their applications. However, in today's Internet, the use of overlay networks, which change traffic routing at the application layer to better satisfy the applications' demands, is rapidly increasing, and this challenges the ISPs' existing agreements and interconnections. Wang et al. [10] are the first to study the impact of the application layer routing of P2P applications on ISPs' peering and provisioning strategies. They propose simple models to represent P2P traffic demands, peering, and routing in a market place of two competing ISPs, and analyze the effectiveness of alternative peering and provisioning strategies available to ISPs. Wang et al. [11] then extend their original model [10] to include more general P2P traffic models and the subscribers' choice process. They build a multi-leader-follower game-theoretic model of subscribers choosing ISPs, and the ISPs making provisioning and peering decisions. However, Labovitz et al. in [12] find that the inter-domain P2P traffic keeps decreasing in recent years. At the same time, the total volume of inter-domain traffic has an average annual increase of 44.5%. In addition to traditional ISPs, content providers and consumer networks now also rival several global transit networks in inter-domain traffic contribution. In order to meet the demands posed by the new requirements such as heterogeneity and inter-domain QoS, we believe another important kind of overlay application, i.e. overlay routing [13]–[18] would become more and more important in the future. In addition to P2P, as another important kind of overlay applications, overlay routing also becomes more and more im-

Manuscript received August 7, 2012.
Manuscript revised December 28, 2012.

†The author is with the Graduate School of Information Science and Technology, Osaka University, Toyonaka-shi, 560–0043 Japan.
††The authors are with the Cybermedia Center, Osaka University, Toyonaka-shi, 560–0043 Japan.

*This paper is an extended version of the work originally to be presented at ITC 2011 [1].

a) E-mail: x-shao@ist.osaka-u.ac.jp
b) E-mail: hasegawa@cmc.osaka-u.ac.jp
c) E-mail: y-tanigu@cmc.osaka-u.ac.jp
d) E-mail: nakano@cmc.osaka-u.ac.jp

DOI: 10.1587/transinf.E96.D.1115

portant [15]–[18]. Hasegawa et al. [19] first study the tussle between overlay routing applications and ISPs' monetary profits, and discusses the guidelines for overlay routing applications to select paths that are more effective while having less of a negative effect on ISPs' profits.

In this paper we also focus on the interaction between ISPs' profits and overlay routing applications; however, we address this problem from the viewpoint of ISPs' interconnection strategies and economic issues of BK peering and paid peering. We assume a typical interconnection scenario with two ISPs, ISP_A and ISP_B, connecting via an abstracted transit service provider. The two ISPs have to decide whether to peer with each other and, if they decide to peer, what peering agreement to accept. We mainly study two peering arrangements: BK peering in which no money is exchanged between the two ISPs, and paid peering in which the peering agreement is determined by the Nash bargaining solution. We analyze some important properties of the two peering agreements, and compare them to the no-peering situation as well as to each other.

As the properties being studied are economic issues, we introduce an ISP cost model composed of the monetary cost and link latency cost. For the monetary cost, a linear pricing scheme is assumed in both transit service and paid peering agreements. For the latency cost, a general convex, increasing and continuous link latency function is assumed. The total cost of one ISP is taken as a weighted sum of monetary and link latency costs. Note that, as the peering link is shared between them, the two ISPs each pay a portion of the link latency cost. As the ISPs' costs are closely related to the inter-ISP traffic pattern, we study the inter-ISP traffic patterns in a Nash equilibrium with various peering link capacities. The traffic in our network is composed of non-overlay routing traffic and overlay routing traffic. Non-overlay routing traffic is transmitted in accordance with BGP routing, while all the overlay routing flows play a selfish routing game.

There are various overlay applications exist in the network, such as P2P, service overlay network and CDN. However, we find that most of the current overlay applications are latency-sensitive. In the early days of P2P file sharing networks, one peer selects neighbors randomly. But today, extensive research has been done on neighbor selection, and selecting neighbors with the lowest latency has become mainstream [20], [21]. Service overlay networks such as Detour [13] and RON [14] also take latency as the most important performance metric. One of the main objectives of the CDN networks such as Akamai is to minimize end-to-end latency by choosing alternative paths [22]. We can see that although various overlay applications exist, most of them choose paths according to latency. Therefore, in this paper, we assume the overlay routing applications are latency sensitive, and all tend to choose the paths with the least latency. We then find three Nash equilibrium traffic patterns exist, corresponding to different peering link capacity levels.

Combining the business model and traffic patterns, we study the optimal peering link capacities of the two ISPs with BK peering and paid peering, and then determine the optimal agreement. With BK peering, the two ISPs announce their own preferred peering capacities simultaneously and, in general, the smaller one is accepted. In our research, we assume that ISP_A will free-ride by overlaying routing traffic when the peering link capacity is large enough. We show that ISP_A always prefers a larger peering link capacity than ISP_B, so the peering link capacity in an agreement is usually the optimal peering link capacity of ISP_B. With paid peering determined by Nash bargaining, the two ISPs coordinate their calculations of the optimal peering link capacity and corresponding price, so they can reach the optimal agreement together. With both BK peering and paid peering, the two ISPs costs may be lower than with no peering. Then we study the impact of network upgrading on the ISPs' peering decisions. The results show that with BK peering, transit link upgrading can improve the situation for ISP_B but not for ISP_A. In a special case in which ISP_A pays the complete cost of the peering link, it has an incentive to upgrade the transit link itself. With paid peering, ISPs have similar incentives for network upgrading: upgrading transit links can lead to smaller costs for both ISPs. Finally, we compare BK peering and paid peering from the aspect of total welfare.

Our work is different from [4], [6]–[8], because they focus on ISPs economic problems based on BGP policy routing, while we build a different ISP business model and introduce overlay routing traffic into it. It is different from [10] and [11], because their focus is on P2P applications, while ours is on overlay routing traffic. This research is also different from [19], because they are concerned with the overlay routing applications side, while our work considers ISPs' connection decisions.

The paper is organized as follows. In Sect. 2 we construct an ISP business model. Section 3 studies the overlay routing traffic patterns under various conditions. Section 4 analyzes the economic issues of the ISPs' peering connection. Section 5 provides our conclusions and offers an outlook on future work.

2. Network and Business Models

2.1 The Models

We consider a network as shown in Fig. 1. For clarity of presentation, Table 1 also lists most of the notations used in the model. ISP_A and ISP_B are two ISPs connecting with each other through a peering link of capacity c_{AB}. R represents the rest of the Internet, and both ISP_A and ISP_B have connections with R of capacities c_{AR} and c_{BR}, respectively. The network model is simple, but as it reflects some basic principles of transit and peering relationships, it is preferred to be used by researchers. The authors of [11] and [10] study the tussles between ISPs and P2P file-sharing applications with similar network models. Different from [11] and [10], we use this model to study overlay routing.

Subscribers of both ISPs may produce traffic demand

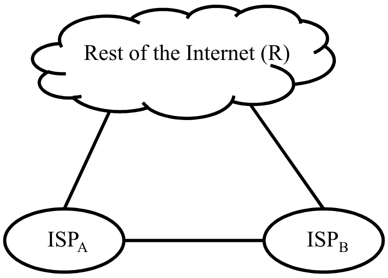

Fig. 1 Network model.

Table 1 Notations used in the mathematical model.

l_{ij}:	The communication link connecting ISP_i and ISP_j
c_{ij}:	The link capacity of l_{ij}
t_{ij}:	The total traffic demand between ISP_i and ISP_j
$\tilde{t}_{ij}$:	The actual traffic amount on link l_{ij}
t^o_{ijk}:	The multi-hop overlay traffic amount on path $ISP_i \leftrightarrow ISP_j \leftrightarrow ISP_k$
P_i:	Price per unit traffic that ISP_i pays for the transit service
p_{ij}:	Price per unit traffic for the peering link
D_{ij}:	The latency of link l_{ij}
J_i^{BK}:	The cost of ISP_i with BK peering
J_i^{PP}:	The cost of ISP_i with paid peering
ρ:	The ratio of overlay routing traffic
α:	The ratio of ISP_A's share of the latency cost on the peering link

between the ISPs as well as with the rest of the Internet. As we focus on inter-ISP traffic in the paper, the intra-ISP demand is not considered. The traffic is composed of non-overlay routing and overlay routing traffic. The overlay routing traffic is a proportion ρ of the total traffic amount, that is, if the traffic between ISP_A and ISP_B is t_{AB}, then the overlay routing traffic is ρt_{AB}. As the traffic amount between an ISP and the Internet is always much larger than the traffic amount between local ISPs, we assume $t_{AR} > t_{AB}$, and $t_{BR} > t_{AB}$. Non-overlay routing traffic is transmitted with policy routing, while overlay routing traffic is performance sensitive and user-directed, and chooses routes by itself. Overlay routing applications can get better performance by choosing multi-hop paths to avoid a bottleneck link. In this paper, the multi-hop overlay traffic is denoted by t^o_{ijk} $(i, j, k \in \{A, B, R\},\ i \neq j \neq k)$. We assume overlay routing traffic takes latency as the performance criterion, and the routes with the least latency are preferred. Because of the existence of overlay routing, the actual traffic on link l_{ij} may not be equal to t_{ij}, so we denote the actual traffic by $\tilde{t}_{ij}$. A general link latency function $D_{ij}(c_{ij}, \tilde{t}_{ij})$ is used to denote the latency of link l_{ij}, and c_{ij} is the link capacity. We assume $D_{ij}(c_{ij}, \tilde{t}_{ij})$ is continuous and twice differentiable with respect to both c_{ij} and $\tilde{t}_{ij}$, with the following properties: $\frac{\partial D_{ij}(c_{ij},\tilde{t}_{ij})}{\partial \tilde{t}_{ij}} > 0$, $\frac{\partial D_{ij}(c_{ij},\tilde{t}_{ij})}{\partial c_{ij}} < 0$, and $\frac{\partial^2 D_{ij}(c_{ij},\tilde{t}_{ij})}{\partial^2 \tilde{t}_{ij}} > 0$. The latency is assumed to be the same for both directions of traffic. Without loss of generality, we assume that

$$D_{AR}(c_{AR}, t_{AR}) > D_{BR}(c_{BR}, t_{BR}). \tag{1}$$

The results for the case $D_{AR}(c_{AR}, t_{AR}) < D_{BR}(c_{BR}, t_{BR})$ is similar, with the ISPs swapping roles. We note that the case $D_{AR}(c_{AR}, t_{AR}) = D_{BR}(c_{BR}, t_{BR})$ is not interesting, since in this case it can be verified that no free-riding happens.

An ISP's cost is composed of a monetary cost and a performance cost. In order to access the Internet, ISPs have to pay higher tier ISPs for transit service. We assume a linear pricing scheme is used, then ISP_A is charged P_A per unit of traffic transmitted, and ISP_B is charged P_B per unit of traffic transmitted. If a paid peering agreement is reached, a linear peering price p_{AB} is also assumed to be used. $p_{AB} > 0$ implies ISP_A pays ISP_B, while $p_{AB} < 0$ implies ISP_A charges ISP_B. The boundary case, $p_{AB} = 0$, implies that a BK peering agreement is in use, and no money is exchanged. Notice that in this work, we focus on the ISPs' long term average costs, and omit all the once only investment. Besides monetary cost, ISPs also suffer from link latency. We use as an ISP's latency cost the product of link latency and traffic on the link as in [23]. Then the ISPs' costs with paid peering can be written as

$$\begin{aligned} J_A^{PP} &= \lambda(\tilde{t}_{AR}D_{AR} + \alpha\tilde{t}_{AB}D_{AB}) + P_A\tilde{t}_{AR} + p_{AB}\tilde{t}_{AB} \\ J_B^{PP} &= \lambda(\tilde{t}_{BR}D_{BR} + (1-\alpha)\tilde{t}_{AB}D_{AB}) \\ &\quad + P_B\tilde{t}_{BR} - p_{AB}\tilde{t}_{AB}. \end{aligned} \tag{2}$$

The first terms are latency costs, and the second terms are monetary costs. The variable $\lambda > 0$ translates the latency cost into an appropriate monetary value. As ISP_A and ISP_B share the same peering link and the latency is experienced by users of both ISPs, the latency cost of the peering link is shared by the two ISPs. We use α $(0 < \alpha < 1)$ to measure the ratio of ISP_A's share, and $(1-\alpha)$ to measure ISP_B's share. If $p_{AB} = 0$, the situation is reduced to BK peering, and the ISPs' costs are

$$\begin{aligned} J_A^{BK} &= \lambda(\tilde{t}_{AR}D_{AR} + \alpha\tilde{t}_{AB}D_{AB}) + P_A\tilde{t}_{AR} \\ J_B^{BK} &= \lambda(\tilde{t}_{BR}D_{BR} + (1-\alpha)\tilde{t}_{AB}D_{AB}) + P_B\tilde{t}_{BR}. \end{aligned} \tag{3}$$

Analogously, we also define the cost functions without peering as

$$\begin{aligned} J_A^{NP} &= \lambda(t_{AR} + t_{AB})D_{AR} + P_A(t_{AR} + t_{AB}) \\ J_B^{NP} &= \lambda(t_{BR} + t_{AB})D_{BR} + P_B(t_{BR} + t_{AB}). \end{aligned} \tag{4}$$

Sometimes it is useful to consider the total cost of the two ISPs. We denote the total costs of paid peering, BK peering and no peering by

$$\begin{aligned} J_{total}^{PP} &= J_A^{PP} + J_B^{PP} \\ J_{total}^{BK} &= J_A^{BK} + J_B^{BK} \\ J_{total}^{NP} &= J_A^{NP} + J_B^{NP}. \end{aligned} \tag{5}$$

It is easy to find that

$$J_{total}^{PP} = J_{total}^{BK}, \tag{6}$$

which implies that paid peering cannot increase the warfare, but can just reallocate the warfare between the two ISPs. In order to reach a paid peering agreement, ISPs have to negotiate to decide p_{AB}, because the costs with paid peering

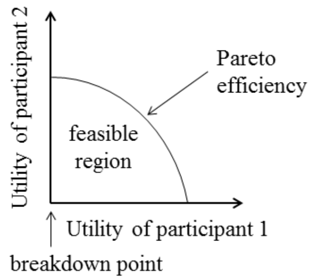

Fig. 2 The feasible region with Pareto-efficient frontier.

of them are interdependent, and the p_{AB} preferred by one ISP might not be preferable to the other. In order to come to an agreement, the gains from negotiation must be equitable (or fair) while operating at a Pareto efficient operation point. We introduce a Nash bargaining solution [24] for the negotiation and explain the idea next.

2.2 Paid Peering by Nash Bargaining Solution

In a general Nash bargaining game [24], the participants's utilities are interdependent and concave. Figure 2 shows the feasible region for the utilities of the two participants. The feasible region is defined as the region where both participants would obtain better utilities compared to no agreement. No agreement can be referred to as the breakdown point. If there is no other choice to improve the utility of one participant, without making that of the other participant worse, the solution found is Pareto-efficient. From [24], a fair and Pareto-efficient outcome can be obtained by optimizing the Nash product. In this work, we suppose the costs with no peering as the breakdown point. As the objective of ISPs is to minimize costs, the Nash product is transformed as the following minimal problem.

$$\min\left(J_A^{NP} - J_A^{PP}\right)^{\alpha}\left(J_B^{NP} - J_B^{PP}\right)^{1-\alpha}.$$

Note that the ISPs' should have different bargaining powers, which can be influenced by many factors. In this paper, we take α (the parameter controlling the sharing of latency costs) as the bargaining power of ISP_A, and $(1-\alpha)$ as the bargaining power of ISP_B because, for fairness, the ISP who pays more of the cost of the peering link should also benefit more from it. In order to obtain the solution, we can transform it into an equivalent problem by taking the logarithm of the objective function:

$$\min\left(\alpha\ln(J_A^{NP} - J_A^{PP}) + (1-\alpha)\ln(J_B^{NP} - J_B^{PP})\right).$$

In order to solve the minimum problem, we let the first order condition equal to 0,

$$\frac{\partial\left(\alpha\ln(J_A^{NP} - J_A^{PP}) + (1-\alpha)\ln(J_B^{NP} - J_B^{PP})\right)}{\partial p_{AB}} = 0,$$

and the Nash solution is

$$p_{AB}^{Nash} = \frac{(1-\alpha)(J_A^{NP} - J_A^{BK}) - \alpha(J_B^{NP} - J_B^{BK})}{\tilde{t}_{AB}}. \tag{7}$$

According to [24], the price p_{AB}^{Nash} is the only solution that is simultaneously Pareto-efficient and fair. Substituting (7) and (6) into (2), we get the ISPs' costs determined by Nash bargaining as

$$\begin{aligned} J_A^{PP} &= (1-\alpha)J_A^{NP} - \alpha J_B^{NP} + \alpha(J_{total}^{PP}) \\ J_B^{PP} &= \alpha(J_B^{NP} - (1-\alpha)J_A^{NP} + (1-\alpha)(J_{total}^{PP}). \end{aligned} \tag{8}$$

We can see that J_A^{PP} and J_B^{PP} are both proportional to J_{total}^{PP}.

3. Traffic Model with Overlay Routing

It was shown in the business model that ISPs' costs are closely related to the traffic model. In this section, we study the traffic patterns composed of overlay routing and non-overlay routing in the network shown in Fig. 1. Non-overlay routing traffic is routed with a policy routing strategy. In our model, traffic with source i and destination j ($i, j \in \{A, B, R\}$, $i \neq j$) is routed through the directed path l_{ij}. Overlay routing traffic generated by all overlay users are playing a non-atomic selfish routing game [25]. In this game, each unit of overlay routing traffic flow travels along the minimum-latency path available to it, where latency is measured with respect to the rest of the flows; otherwise, this flow would reroute itself on a path with smaller latency. In other words, all paths in use by an equilibrium flow have minimum-possible cost. In particular, all paths of a given commodity used by an equilibrium flow have equal latency. In our model, there are only three links, so we are able to analyze all the overlay routing traffic patterns explicitly.

First, suppose at certain time that the latencies of the three links are

$$D_{AR} + D_{BR} < D_{AB},$$

which happens when c_{AB} is so small that the latency of path $A \leftrightarrow R \leftrightarrow B$ is less than path $A \leftrightarrow B$. Then overlay routing traffic with source-destination pair in ISP_A and ISP_B would choose the multi-hop path for better performance. This process will continue until the latencies of the two paths become equal, or all the overlay routing traffic with source-destination pair in ISP_A and ISP_B has chosen path $A \leftrightarrow R \leftrightarrow B$. Then we have

$$D_{AR} + D_{BR} \leq D_{AB}. \tag{9}$$

Note that, given (9), we can also have

$$\begin{aligned} D_{AR} &< D_{BR} + D_{AB} \\ D_{BR} &< D_{AR} + D_{AB}, \end{aligned}$$

which suggests that $t_{ABR}^{o} = t_{BAR}^{o} = 0$, and that this is a Nash equilibrium traffic pattern. Note that a peering link with a very small capacity will cause serious congestion, so that this is not feasible in practice. In order to make the analysis mathematically tractable, we set a lower bound for c_{AB} as the value that makes (9) an equation, and denote it by c_1^l. Then the properties of this pattern can be summarized as follows.

$$
\begin{aligned}
&D_{AB} = D_{AR} + D_{BR}\\
&\tilde{t}_{AR} = t_{AR} + t^o_{ARB}\\
&\tilde{t}_{BR} = t_{BR} + t^o_{ARB}\\
&\tilde{t}_{AB} = t_{AB} - t^o_{ARB}\\
&t^o_{ABR} = t^o_{BAR} = 0, 0 \le t^o_{ARB} \le \rho t_{AB}.
\end{aligned} \tag{10}
$$

Equations (10) suggest that given transit link capacities and latency functions, t^o_{ARB} is determined only by c_{AB}. So t^o_{ARB} can be seen as a function of c_{AB}, and it is decreasing with respect to c_{AB}. We denote the upper bound of c_{AB} that allows (10) to hold as c^h_1.

Next we discuss the situation when c_{AB} increases so that

$$|D_{AR} - D_{BR}| < D_{AB} < D_{AR} + D_{BR}. \tag{11}$$

In this case no multi-hop overlay routing traffic exists. With the assumption $D_{AR}(t_{AR}) > D_{BR}(t_{BR})$, we can summarize the properties of this pattern as

$$
\begin{aligned}
&D_{AR} - D_{BR} < D_{AB} < D_{AR} + D_{BR}\\
&\tilde{t}_{AR} = t_{AR}\\
&\tilde{t}_{BR} = t_{BR}\\
&\tilde{t}_{AB} = t_{AB}\\
&t^o_{ABR} = t^o_{BAR} = t^o_{ARB} = 0.
\end{aligned} \tag{12}
$$

We denote the lower and upper bounds of c_{AB} for which (12) hold by c^l_2 and c^h_2, and note that $c^l_2 = c^h_1$.

If c_{AB} continues increasing and exceeds c^h_2, then

$$D_{AB}(c^h_2, t_{AB}) \le D_{AR}(t_{AR}) - D_{BR}(t_{BR}). \tag{13}$$

If so, a portion of overlay routing traffic with source-destination pair in ISP_A and R will move to the path $A \leftrightarrow B \leftrightarrow R$ until the latencies of the two paths become equal, or all overlay routing traffic with source-destination pair in ISP_A and R have chosen the alternative multi-hop path. Note that c_{AB} can be increased indefinitely in the theory, but this is not feasible in practice. Also, as c_{AB} exceeds some very large value, the traffic pattern will become very complicated. In order to make the analysis practical and mathematically tractable, we choose as an upper bound for c_{AB} the relatively large value that makes (13) an equation, and denote it by c^h_3. In this case, t^o_{ABR} is an increasing function of c_{AB}, and $t^o_{ABR}(c^h_3) \le \rho t_{AR}$. We also denote the lower bound of c_{AB} by c^l_3, and $c^l_3 = c^h_2$. We can summarize the properties of this pattern as follows.

$$
\begin{aligned}
&D_{AB} = D_{AR} - D_{BR}\\
&\tilde{t}_{AR} = t_{AR} - t^o_{ABR}\\
&\tilde{t}_{BR} = t_{BR} + t^o_{ABR}\\
&\tilde{t}_{AB} = t_{AB} + t^o_{ABR}\\
&t^o_{ARB} = t^o_{BAR} = 0, 0 \le t^o_{ABR} \le \bar{t}^o_{ABR}.
\end{aligned} \tag{14}
$$

Note that ρt_{AR} may not be a tight upper bound of t^o_{ABR}. We denote the maximum of t^o_{ABR} as $\bar{t}^o_{ABR}$. The traffic traveling through the path $A \leftrightarrow B \leftrightarrow R$ is called free-riding traffic. ISP_A does not need to pay for the transit service of the free-riding traffic, and ISP_B pays instead.

In the three patterns discussed above, if the properties of the two ISPs are given, the traffic patterns are completely determined by the peering link capacity. So peering capacity is an important factor for ISPs to consider when making connection decisions.

4. Economic Issues for ISP Connections

In this section, we combine the business model and the traffic patterns considered in the previous sections to obtain the ISPs' costs under different peering agreements and with various peering capacities. First, we summarize the costs functions by combining (3), (10), (12) and (14) as follows.

$$
J^{BK}_A = \begin{cases}
\lambda((\tilde{t}_{AR} + \alpha\tilde{t}_{AB})D_{AR} + \alpha\tilde{t}_{AB}D_{BR})\\
+P_A\tilde{t}_{AR}, \text{if } c^l_1 \le c_{AB} \le c^h_1;\\
\lambda(t_{AR}D_{AR} + \alpha t_{AB}D_{AB}) + P_A t_{AR}\\
\text{if } c^l_2 < c_{AB} \le c^h_2;\\
\lambda((\tilde{t}_{AR} + \alpha\tilde{t}_{AB})D_{AR} - \alpha\tilde{t}_{AB}D_{BR})\\
+P_A\tilde{t}_{AR}, \text{if } c^l_3 < c_{AB} \le c^h_3;
\end{cases} \tag{15}
$$

$$
J^{BK}_B = \begin{cases}
\lambda((\tilde{t}_{BR} + (1-\alpha)\tilde{t}_{AB})D_{BR}\\
+(1-\alpha)\tilde{t}_{AB}D_{AR}) + P_B\tilde{t}_{BR},\\
\text{if } c^l_1 \le c_{AB} \le c^h_1;\\
\lambda(t_{BR}D_{BR} + (1-\alpha)t_{AB}D_{AB}) + P_B\tilde{t}_{BR},\\
\text{if } c^l_2 < c_{AB} \le c^h_2;\\
\lambda((\tilde{t}_{BR} - (1-\alpha)\tilde{t}_{AB})D_{BR}\\
+(1-\alpha)\tilde{t}_{AB}D_{AR}) + P_B\tilde{t}_{BR},\\
\text{if } c^l_3 < c_{AB} \le c^h_3.
\end{cases} \tag{16}
$$

Costs of paid peering can be calculated from (8), (15) and (16). As J^{PP}_A and J^{PP}_B are both proportional to J^{PP}_{total} from (8), we only show J^{PP}_{total} here.

$$
J^{PP}_{total} = \begin{cases}
\lambda((\tilde{t}_{AR} + t_{AB})D_{AR} + (\tilde{t}_{BR} + t_{AB})D_{BR})\\
+P_A\tilde{t}_{AR} + P_B\tilde{t}_{BR}, \text{if } c^l_1 \le c_{AB} \le c^h_1;\\
\lambda(t_{AR}D_{AR} + t_{AB}D_{AB} + t_{BR}D_{BR})\\
+P_A t_{AR} + P_B t_{BR}, \text{if } c^l_2 < c_{AB} \le c^h_2;\\
\lambda((\tilde{t}_{AR} + t_{AB})D_{AR} + (\tilde{t}_{BR} - t_{AB})D_{BR})\\
+P_A\tilde{t}_{AR} + P_B\tilde{t}_{BR}, \text{if } c^l_3 < c_{AB} \le c^h_3;
\end{cases} \tag{17}
$$

Figure 3 shows two examples of ISPs' costs vs. peering link capacity with and without overlay routing traffic. We can see that all the cost functions with overlay routing traffic are in three pieces. In the examples, we assume that the two ISPs use the same M/M/1 latency model. For an M/M/1 queue, the latency can be expressed as $l(x) = \frac{1}{\mu - x} + prop$, where x is the traffic load, μ is the link capacity, and $prop$ is the propagation delay. This model satisfies all the assumptions of our latency model. The values of variables are set to $P_A = P_B = 0.0001$, $t_{AR} = 200$, $c_{AR} = 600$, $t_{BR} = 150$,

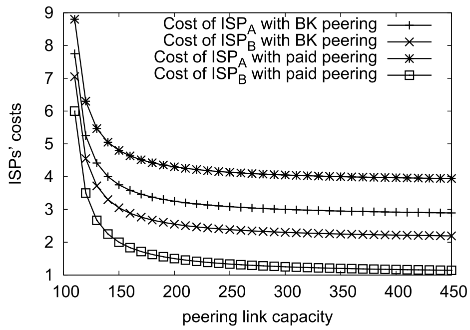

(a) $\rho = 0.0$

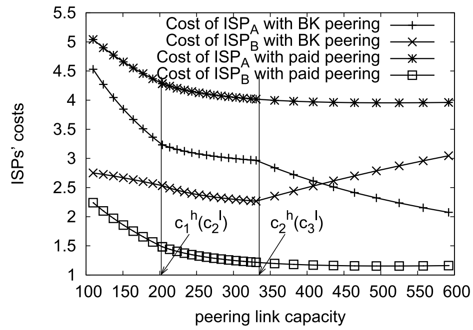

(b) $\rho = 0.5$. The boundaries of different peering levels are indicated by the vertical lines.

Fig. 3 ISPs' costs.

$c_{BR} = 300$, $t_{AB} = 100$, $prop_{AB} = prop_{BR} = 0.001$, $prop_{AR} = 0.008$, $\alpha = 0.65$ and $\lambda = 1$. In the examples, we let l_{AR} much longer than l_{AB} and l_{BR}, so that the costs of ISP_A can be separated widely from ISP_B's costs in Fig. 3. Note that as the results and conclusions in this paper are obtained from mathematical analysis, changing parameter values will not affect our results and conclusions. We decide the parameter values only to make the experiment results more clear to understand. In Fig. 3 (a), the ratio of overlay routing traffic is 0, and the ISPs' cost curves are smooth. In Fig. 3 (b), the ratio of overlay routing is 0.5, and the cost curves become three-piece curves corresponding to different peering levels, with breaking points of $c_{AB} = c_1^h$ and $c_{AB} = c_2^h$. From these examples, we can see the significant impact of overlay routing on ISPs' cost functions.

In the following, we first study the best peering agreement with BK peering and paid peering respectively. Then we analyze the impact of network upgrading. Finally we compare the total welfare with BK peering and paid peering.

Lemma 1. *With BK peering, ISP_A can reach the minimum cost with a peering capacity of c_3^h, while ISP_B can reach the minimum cost with c_{BK}^* in $[c_3^l, c_3^h]$.*

In fact, from the proof in Appendix, the peering link capacity is the larger the better for ISP_A, because it can obtain additional profit from free-riding traffic, and with BK peering agreement, it does not need to pay ISP_B for the free-riding traffic. In contrast, when the peering link capacity c_{AB} exceeds c_3^l, ISP_B has to pay for the free-riding traffic instead of ISP_A. At the same time, ISP_B can also enjoy the profit from the performance enhancement of the peering link. According to (16), we can see that the optimal peering decision for ISP_B is to minimize J_B^{BK} with the link capacity constraints. From the proof in Appendix, we show that ISP_B can reach the minimal cost with a certain peering link capacity c_{BK}^* in $[c_3^l, c_3^h]$. Based on Lemma 1, we can have the following conclusion.

Proposition 1. *With BK peering, ISPs can reach the optimal agreement with $c_{AB} = c_{BK}^*$. With this agreement, both ISPs can also do better than with no peering.*

As comparison, then we analyze the optimal agreement with paid peering. We have the following result.

Lemma 2. *With paid peering, both J_A^{PP} and J_B^{PP} reach the minimum value with the same peering capacity $c_{PP}^* \in [c_3^l, c_3^h]$.*

As in paid peering, the peering capacity and peering price are decided by the two ISPs cooperatively with Nash bargaining solution, so that the peering link capacity c_{PP}^* is preferred by the two at the same time. Based on Lemma 2, we have the following conclusion.

Proposition 2. *With paid peering, ISPs can reach the optimal agreement with $c_{AB} = c_{PP}^*$. With this agreement, both ISPs do better than with no peering.*

As both ISPs prefer $c_{AB} = c_{PP}^*$, obviously they can reach optimal agreement with $c_{AB} = c_{PP}^*$. And Proposition 2 also tells paid peering is more efficient than no peering.

In today's Internet, good policies not only minimize ISPs' costs, but can also promote the upgrading of Internet, while bad policies do the opposite. So we study the promotion effects of BK and paid peering agreements.

Proposition 3. *With the optimal BK peering in Proposition 1, ISP_B will make further cost savings through the upgrading of transit links.*

Proof. The problem we are interested here is the effect of c_{BR} and c_{AB} to the minimum value of J_B^{BK}. Essentially, in non-linear constraint optimization problem, it is to study the effect of the parameters in the constraint conditions to the extreme value of the objective function. In economics, it is well known that such problem can be solved by the envelope theorem [26] as follows.

From (14), we know that t_{ABR}^o is a strictly increasing function of c_{AB}. Therefore, we can also take J_B^{BK} as the function of t_{ABR}^o. Denote x as t_{ABR}^o, and x_{BK}^* as $x(c_{BK}^*)$ for convenience, we can have the Lagrangian equation of J_B^{BK} in (16) as

$$L_B^{BK}(x, \gamma_1, \gamma_2) = J_B^{BK}(x) - \gamma_1(x) - \gamma_2(-x + \bar{x}),$$

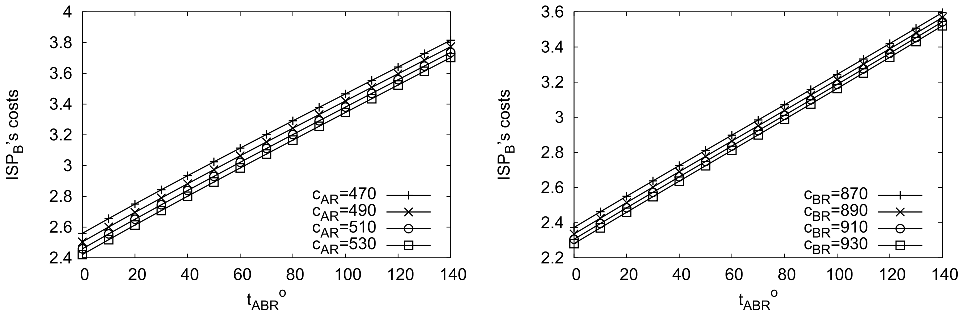

(a) Increase c_{AR} (b) Increase c_{BR}

Fig. 4 ISP_B's optimal cost and network upgrading.

where γ_1 and γ_2 are Lagrangian multipliers, and $\bar{x}$ is the upper bound of x as in (14). Write γ_1^* and γ_2^* for the Lagrangian multipliers when J_B^{BK} reaches its minimum value, and define a function of c_{AR} and c_{BR} as

$$\begin{aligned}&V_B^{BK}(c_{AR}, c_{BR})\\&\quad = L_B^{BK}(x_{BK}^*(c_{AR}, c_{BR}), \gamma_1^*(c_{AR}, c_{BR}), \gamma_2^*(c_{AR}, c_{BR})).\end{aligned}$$

According to the envelope theorem, the effects of increasing c_{AR} and c_{BR} can then be obtained as

$$\frac{\partial V_B^{BK}}{\partial c_{AR}} = \lambda(1-\alpha)(x_{BK}^* + t_{AB})\frac{\partial D_{AR}}{\partial c_{AR}} < 0, \tag{18}$$

and as we assume $t_{BR} > t_{AB}$,

$$\frac{\partial V_B^{BK}}{\partial c_{BR}} = \lambda(t_{BR} + \alpha x_{BK}^* - t_{AB} + \alpha t_{AB})\frac{\partial D_{BR}}{\partial c_{BR}} < 0. \tag{19}$$

□

In order to indicate the size of these savings, we show the relationship between ISP_B's cost and transit link capacities in Figs. 4 (a) and 4 (b). In Fig. 4 (a), we let c_{AR} increase from 470 to 530 while keeping c_{BR} constant; and in Fig. 4 (b), we let c_{BR} increase from 870 to 930 while keeping c_{AR} constant. The values of other variables are the same as in the example depicted in Fig. 3 (b). We can see that ISP_B's optimal cost decreases with both types of network upgrade.

For ISP_A, as it just accepts ISP_B's decision of c_{BK}^* passively as shown in the proof of Proposition (1), the effect of transit link upgrading also depends on ISP_B's decision. For example, we can compute ISP_A's incentive to upgrade l_{AR} as

$$\begin{aligned}&\frac{\partial J_A^{BK}(c_{AR}, c_{BR}, x_{BK}^*(c_{AR}, c_{BR}))}{\partial c_{AR}}\\&\quad = \frac{\partial J_A^{BK}}{\partial c_{AR}} + \frac{\partial J_A^{BK}}{\partial x_{BK}^*}\frac{\partial x_{BK}^*}{\partial c_{AR}}.\end{aligned} \tag{20}$$

For the model in this paper, whether ISP_A can make further cost savings through the upgrading of transit links depends on external factors. It is not difficult to construct examples of both better case and worse case for ISP_A. Due to limitations of space, we neglect the examples.

In the case of paid peering, the two ISPs have similar incentives for transit link upgrading. We summarize these in the following proposition.

Proposition 4. *With the optimal paid peering, both ISPs can reduce costs further by upgrading transit links.*

Proof. As J_A^{PP} and J_B^{PP} are both proportional to J_{total}^{PP} from (8), we only have to obtain the impact of increasing transit links on J_{total}^{PP} instead of J_A^{PP} and J_B^{PP}. It is analogous to Proposition (3), so that we can prove it with similar method. From (14), we know that t_{ABR}^o is a strictly increasing function of c_{AB}. Therefore, we can also take J_{total}^{PP} as the function of t_{ABR}^o. Denote x as t_{ABR}^o, and x_{PP}^* as $x(c_{PP}^*)$ for convenience, we can have the Lagrangian equation of J_{total}^{PP} in (17) as

$$L_{total}^{PP}(x, \mu_1, \mu_2) = J_{total}^{PP}(x) - \mu_1(x) - \mu_2(-x + \bar{x}),$$

where μ_1 and μ_2 are Lagrangian multipliers, and $\bar{x}$ is the upper bound of x as in (14). As in Proposition 2, J_{total}^{PP} reaches the global minimum value when $c_{AB} = c_{PP}^*$. Denote μ_1^* and μ_2^* as the Lagrangian multipliers when J_{total}^{PP} reaches the minimum value, we can define a function of c_{AR} and c_{BR} as

$$\begin{aligned}&V_{total}^{PP}(c_{AR}, c_{BR})\\&\quad = L_{total}^{PP}(x_{PP}^*(c_{AR}, c_{BR}), \mu_1^*(c_{AR}, c_{BR}), \mu_2^*(c_{AR}.c_{BR})),\end{aligned}$$

According to the envelope theorem, the effects of increasing c_{AR} and c_{BR} can then be obtained as

$$\begin{aligned}\frac{\partial V_{total}^{PP}}{\partial c_{AR}} &= \lambda(t_{AR} + t_{AB})\frac{\partial D_{AR}}{\partial c_{AR}} < 0\\ \frac{\partial V_{total}^{PP}}{\partial c_{BR}} &= \lambda(t_{BR} + t_{AB})\frac{\partial D_{BR}}{\partial c_{BR}} < 0.\end{aligned}$$

These imply that, with paid peering, both ISPs can reduce costs further by upgrading transit links. □

In Figs. 5 (a) and 5 (b), we depict the relationship between ISPs' total optimal costs and transit link capacities with a paid peering agreement. In Fig. 5 (a), we let c_{AR} increase from 470 to 530 while keeping c_{BR} constant; and in Fig. 5 (b), we let c_{BR} increase from 870 to 930 while keeping c_{AR} constant. The values of other variables are the same

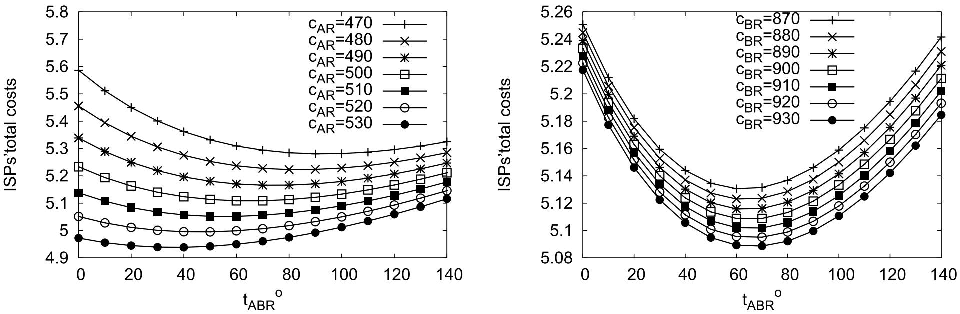

(a) Increase c_{AR} (b) Increase c_{BR}

Fig. 5 Total optimal cost and network upgrading.

as in the example depicted in Fig. 3 (b). We can see that the total optimal cost decreases with network upgrading in both cases. From (8), we know that the relation of optimal cost to transit link capacities follow a similar pattern for both ISPs.

Besides ISPs' individual incentives for network upgrading, the social welfare is also an important factor for policy makers. If we take ISP_A, ISP_B and the rest of Internet as a system, the cost of the system is equal to the total latency cost of ISP_A and ISP_B. The monetary cost is not accounted because it is the cost for the payer, but also the profit for the recipient. In this work, the social welfare can be evaluated with the total latency cost of the two ISPs. Smaller total latency cost implies bigger social welfare. Then we have the following result.

Proposition 5. *The social welfare with paid peering is equal to or bigger than that with BK peering.*

Proof. Denote x as t^o_{ABR}, and take J^{PP}_{total} and J^{BK}_B as functions of x. As $J^{PP}_{total}{}''(x) > 0$, it is a convex function with respect to x. Suppose J^{PP}_{total} reaches its minimum value at x^*_{PP}, and J^{BK}_B reaches its minimum value at x^*_{BK}. Then, given any $x_0 \geq x^*_{PP}$, we have

$$\begin{aligned} J^{PP}_{total}{}'(x_0) &= \lambda(t_{BR} - t_{AB})D'_{BR}(t_{BR} + x_0) + P_B \\ &\quad - \lambda(t_{AR} + t_{AB})D'_{AR}(t_{AR} - x_0) - P_A \\ &\geq 0, \end{aligned}$$

and

$$\begin{aligned} J^{BK}_B{}'(x_0) = &- \lambda(1-\alpha)(t_{AB} + x_0)D'_{AR}(t_{AR} - x_0) \\ &+ \lambda(t_{BR} + \alpha x_0 - t_{AB} + \alpha t_{AB})D'_{BR}(t_{BR} + x_0) \\ &+ \lambda(1-\alpha)D_{AR}(t_{AR} - x_0) + P_B \\ > &\; J^{PP}_{total}{}'(x_0), \end{aligned}$$

which implies $J^{BK}_B(x)$ increases in $(x^*_{PP}, \bar{x})$, where $\bar{x}$ is the upper bound of x as in (14). We can also claim that $x^*_{BK} \leq x^*_{PP}$. Specifically, if $x^*_{PP} = 0$, then $x^*_{BK} = 0$; if $x^*_{PP} = \bar{x}$, then $x^*_{BK} \leq \bar{x}$; and if $x^*_{PP} \in (0, \bar{x})$, then $0 \leq x^*_{BK} < x^*_{PP}$. Note that according to (15), (16) and (17), the total latency cost is $\lambda((t_{AR} + t_{AB} - t^o_{ABR})D_{AR} + (t_{BR} - t_{AB} + t^o_{ABR})D_{BR})$. As in this case, $D_{AR} > D_{BR}$, the total latency cost is decreasing with respect to t^o_{ABR}, and the proposition can be confirmed. □

It implies paid peering may be a better strategy than BK peering from the viewpoint of social welfare. In some cases, ISPs with paid peering can obtain bigger total welfare than with BK peering. Even in the worst case of paid peering, they can still obtain no less welfare than BK peering.

5. Conclusion

In this paper, we studied the economic issues surrounding ISP peering with the advent of overlay routing, using a simple network model. We focused on BK peering and paid peering with a Nash bargaining solution and studied the optimal agreements and related properties.

First, we introduced a cost model for ISPs that is a weighted sum of monetary cost and latency cost. Then we studied the overlay routing traffic patterns under various link conditions and found three Nash equilibrium patterns. Combining the cost model and traffic patterns, we found the ISPs' costs as functions of peering link capacity. Based on this, we determined the agreements which optimize ISPs' costs and analyzed the properties of such agreements. We got the optimal peering link capacities for different ISPs with BK peering and paid peering and proposed optimal agreements that can be accepted by both ISPs. We proved that, with BK peering, network upgrading is welcomed by the ISP that is not the free-rider; in some special cases, the free-riding ISP also has an incentive to upgrade the network. In comparison, with paid peering, we proved that network upgrading is welcomed by both ISPs, in other words, both of them have an incentive to upgrade the network. We also compared the total welfare under BK and paid peering, and showed that the total welfare with paid peering is equal to or bigger than that with BK peering.

Although the network model in the paper is simple, it is quite informative, and makes the traffic and economic models mathematically tractable. We also think the method can be applied to a more complicated network environment by iteration. In fact, even in a complicated network, we believe that ISPs would make bilateral contract with the others individually. However, when making peering contract with a potential neighbor, the ISP has to take into consider the effect of the new connection to the existing connections. As the ISP does not have explicit knowledge of the whole network, one way is to estimate the effect of the potential connection by experience, and then improve it iteratively. In the future, we would also like to extend the model to include more realistic features and implementation issues.

References

[1] X. Shao, G. Hasegawa, Y. Taniguchi, and H. Nakano, "The implication of overlay routing on ISPs' connecting strategies," Proc. ITC 2011, pp.286–293, Sept. 2011.
[2] W.B. Norton, "Internet service providers and peering," Equinix White Papers, 2001.
[3] G. Huston, ISP Survival Guide: Strategies for Running a Competitive ISP, John Wiley and Sons, 1999.
[4] G. Shrimali and S. Kumar, "Paid peering among internet service providers," Proc. GameNets Workshop on Game Theory for Communications and Networks, 2006.
[5] E. Jahn and J. Prüfer, "Interconnection and competition among asymmetric networks in the internet backbone market," Information Economics and Policy, vol.20, no.3, pp.243–256, Sept. 2008.
[6] J.J. Laffont, S. Marcus, P. Rey, and J. Tirole, "Internet interconnection and the off-net-cost pricing principle," RAND Journal of Economics, vol.34, no.2, pp.370–390, 2003.
[7] S. Shakkottai and R. Srikant, "Economics of network pricing with multiple ISPs," IEEE/ACM Trans. Netw., vol.16, pp.1233–1245, Dec. 2006.
[8] G. Shrimali and S. Kumar, "Can bill-and-keep peering be mutually beneficial?," Internet and Network Economics, vol.3828, pp.738–747, 2005.
[9] G. Fudenberg and J. Tirole, Game Theory, MIT Press, 2000.
[10] J.H. Wang, D.M.C. Chiu, and J.C.S. Lui, "Modeling the peering and routing tussle between ISPs and P2P applications," Proc. IEEE IWQoS'06, pp.51–59, June 2006.
[11] J.H. Wang, D.M. Chiu, and J.C.S. Liu, "A game-theoretic analysis of the implications of overlay network traffic on ISP peering," Computer Networks, vol.52, no.15, pp.2961–2974, Oct. 2008.
[12] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, "Internet inter-domain traffic," Proc. ACM SIGCOMM'10, pp.75–86, Sept. 2010.
[13] S. Savage, T. Anderson, A. Aggarwal, D. Becker, N. Cardwell, A. Collins, E. Hoffman, J. Snell, A. Vahdat, G. Voelker, and J. Zahorjan, "Detour: A case for informed internet routing and transport," IEEE Micro, vol.19, no.1, pp.50–59, Jan. 1999.
[14] D. Anderson, H. Balakrishnan, M.F. Kaashoek, and R. Morris, "Resillient overlay networks," ACM SIGOPS Operating Systems Review, vol.35, no.5, pp.131–145, Dec. 2001.
[15] A. Capone, J. Elias, and F. Martignon, "Routing and resource optimization in service overlay networks," Computer Networks, vol.53, no.2, pp.180–190, Feb. 2009.
[16] D. Adami, C. Callegari, S. Giordano, M. Pagano, and T. Pepe, "Optimal design of service overlay networks with economic and performance constraints," International Journal of Communication Systems, vol.23, no.3, pp.369–389, March 2010.
[17] F. Gu, C. Xie, M. Peng, C. Cavdar, S. Khan, and N. Ghani, "Virtual overlay network scheduling," IEEE Commun. Lett., vol.15, no.8, pp.893–895, Aug. 2011.
[18] G. Smaragdakis, N. Laoutaris, V. Lekakis, and A. Bestavros, "Selfish overlay network creation and maintenance," IEEE/ACM Trans. Netw., vol.19, no.6, pp.1624–1637, Dec. 2011.
[19] G. Hasegawa, M. Kobayashi, M. Murata, and T. Murase, "Free-riding traffic problem in routing overlay network," Proc. IEEE ICON'07, pp.118–123, April 2007.
[20] Y. Jiang and J. Wu, "Employing QoS driven neighbor selection for heterogeneous peer-to-peer streaming," Proc. IEEE ICC'11, pp.1–6, July 2011.
[21] R.L. Pereira, "On the impact of P2P file sharing traffic restrictions on user perceived performance," Proc. ICOIN'08, pp.1–5, Jan. 2008.
[22] Akamai, "http://www.akamai.com."
[23] T. Basar and R. Srikant, "Revenue-maximizing pricing and capacity expansion in a many-user regime," Proc. IEEE INFOCOM'02, pp.294–301, April 2002.
[24] J.F. Nash, "Two-person cooperative games," Econometrica, vol.21, pp.128–140, 1953.
[25] T. Roughgarden and E. Tardos, "How bad is selfish routing?," J. ACM, vol.49, no.2, pp.236–259, March 2002.
[26] A. Mas-Colell, M.D. Whinston, and J.R. Green, Microeconomic Theory, Oxford University Press, 1995.

Appendix: Proofs of Results

Proof of Lemma 1. For ISP_A, from (15) and (10), the cost of ISP_A with any $c_{AB} \in [c_1^l, c_1^h]$ is

$$\begin{aligned} J_A^{BK} &= \lambda((\tilde{t}_{AR} + \alpha\tilde{t}_{AB})D_{AR}(\tilde{t}_{AR}) \\ &\quad + \alpha\tilde{t}_{AB}D_{BR}(\tilde{t}_{BR})) + P_A\tilde{t}_{AR} \\ &> \lambda((t_{AR} + \alpha t_{AB})D_{AR}(t_{AR}) \\ &\quad - \alpha t_{AB}D_{BR}(t_{BR})) + P_A t_{AR} \\ &= J_A^{BK}(c_2^h). \end{aligned}$$

If $c_{AB} \in (c_2^l, c_2^h]$, from (15) and (12), we have

$$J_A^{BK\prime}(c_{AB}) = \lambda\alpha t_{AR} D_{AB}{}'(c_{AB}) < 0.$$

If $c_{AB} \in (c_3^l, c_3^h]$, from (15) and (14), J_A^{BK} can be seen as a function of t_{ABR}^o. Then we have

$$\begin{aligned} J_A^{BK\prime}(x) &= \lambda((t_{AR} - x + \alpha(t_{AB} + x))D'_{AR}(x) \\ &\quad - \alpha(t_{AB} + x)D'_{BR}(x)) \\ &\quad + (\alpha - 1)D_{AR} - \alpha D_{BR} - P_A \\ &< 0, \end{aligned}$$

where x is t_{ABR}^o. From (14), $x'(c_{AB}) > 0$, so that $J_A^{BK\prime}(c_{AB}) = J_A^{BK\prime}(x)x'(c_{AB}) < 0$. Thus we can conclude that ISP_A reaches the minimum value with BK peering for $c_{AB} = c_3^h$.

For ISP_B with any $c_{AB} \in [c_1^l, c_1^h]$, from (16) and (10), we can prove that $J_B^{BK}(c_{AB}) > J_B^{BK}(c_2^h)$ in similar way to that used for J_A^{BK}. If $c_{AB} \in (c_2^l, c_2^h]$, from (16) and (12), we have

$$J_B^{BK\prime}(c_{AB}) = \lambda(1 - \alpha)t_{AR}D'_{AB}(c_{AB}) < 0,$$

so the minimum value of J_B^{BK} in $[c_1^l, c_2^h]$ is $J_B^{BK}(c_2^h)$. Also, as J_B^{BK} is continuous in $[c_3^l, c_3^h]$, J_B^{BK} can reach the minimum value for some c_{BK}^* in $[c_3^l, c_3^h]$, so we claim that $J_B^{BK}(c_{BK}^*)$ is the global minimum value. □

Proof of Proposition 1. With BK peering, the general process is that ISPs announce their preferred peering link capacities simultaneously, and the smaller one is accepted. As we have shown in Lemma 1, in $[c_3^l, c_3^h]$, c_{AB} should be as large as possible for ISP_A, while c_{BK}^* is preferred by ISP_B. So the agreement with $c_{AB} = c_{BK}^*$ is the best agreement that can be reached.

Also, from (4), (12), (14) and (15),

$$\begin{aligned} J_A^{BK}(c_{BK}^*) &\leq J_A^{BK}(c_2^h) \\ &= \lambda(t_{AR} + \alpha t_{AB})D_{AR}(t_{AR}) + P_A(t_{AR}) \\ &\quad - \lambda\alpha t_{AB} D_{BR}(t_{BR}) \\ &< \lambda(t_{AR} + t_{AB})D_{AR}(t_{AR} + t_{AB}) + P_A(t_{AR} + t_{AB}) \\ &= J_A^{NP}. \end{aligned}$$

We can also prove $J_B^{BK}(c_{BK}^*) < J_B^{NP}$ in a similar way. To summarize the above arguments, we claim that ISPs will reach an agreement with $c_{AB} = c_{BK}^*$. □

Proof of Lemma 2. As J_A^{PP} and J_B^{PP} are both proportional to J_{total}^{PP} from (8), we can study the properties of total cost instead of J_A^{PP} and J_B^{PP}. When $c_{AB} \in [c_1^l, c_1^h]$, from (15), (16), (5) and (10), we have

$$\begin{aligned} {J_{total}^{PP}}'(c_{AB}) &= \lambda((t_{AR} + t_{AB})D'_{AR}(t_{ARB}^o) \\ &\quad + (t_{BR} + t_{AB})D'_{BR}(t_{ARB}^o)){t_{ARB}^o}'(c_{AB}) \\ &\quad + (P_A + P_B){t_{ARB}^o}'(c_{AB}) \\ &< 0. \end{aligned}$$

When $c_{AB} \in (c_2^l, c_2^h]$, from (15), (16), (5) and (10), we have

$${J_{total}^{PP}}'(c_{AB}) = \lambda t_{AB} D'_{AB}(c_{AB}) < 0.$$

So we can see that J_{total}^{PP} is decreasing in $[c_1^l, c_2^h]$. When $c_{AB} \in [c_3^l, c_3^h]$, as J_{total}^{PP} is a continuous and bounded function, it has at least one minimum value. Suppose the minimum value is reached when $c_{AB} = c_{PP}^*$. Then, for $c_{AB} = c_{PP}^*$, both J_A^{PP} and J_B^{PP} reach their minimum value. □

Proof of Proposition 2. With paid peering, the general process is that ISPs calculate the optimal peering link capacities and price in coordination. As we have shown in Lemma 2, the agreement with $c_{AB} = c_{PP}^*$ is the best agreement that can be reached.

Also, from (4), (5), (8), (15) and (16), we find

$$J_{total}^{PP}(c_{PP}^*) < J_{total}^{PP}(c_2^h) < J_{total}^{NP}.$$

So, from the formulation of the Nash bargaining solution, both ISPs can be better off than with no peering. □

Xun Shao received his B.E. degree from Civil Aviation University of China, in 2005, and his M.E. degree from Beijing Jiaotong University, China, in 2008. His research interests include Internet structure and overlay routing.

Go Hasegawa received his M.E. and Ph.D. degrees from Osaka University, Japan, in 1997 and 2000, respectively. From 1997 to 2000, he was a Research Assistant at the Graduate School of Economics, Osaka University. He is currently an Associate Professor at the Cybermedia Center, Osaka University. His research is in the area of transport architecture for future high-speed networks. He is a member of IEEE.

Yoshiaki Taniguchi received his B.E., M.E. and Ph.D. degrees from Osaka University, Japan, in 2002, 2004 and 2008, respectively. Since 2008, he has been an Assistant Professor at the Cybermedia Center, Osaka University. His research interests include wireless sensor networks, wireless mesh networks, and object tracking systems. He is a member of IEEE and IPSJ.

Hirotaka Nakano received his B.E., M.E. and D.E. degrees in Electrical Engineering from The University of Tokyo, Japan, in 1972, 1974 and 1977, respectively. He joined NTT Laboratories in 1977 and has been engaged in research and development of videotex systems and multimedia-on-demand systems. He was an executive manager of the Multimedia Systems Laboratory at the NTT Human Interface Laboratories from 1995 to 1999. He was the head scientist of the Multimedia Laboratory at NTT DOCOMO until 2004 and is now a Professor at the Cybermedia Center, Osaka University. His research is in the area of ubiquitous networks. He is a member of IEEE and ITE.