PAPER

An Interdomain Overlay Network Based on ISP Alliances for Economically Efficient Interdomain Traffic Routing*

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

SUMMARY As interdomain routing protocol, BGP is fairly simple, and allows plenty of policies based on ISPs' preferences. However, recent studies show that BGP routes are often non-optimal in end-to-end performance, due to technological and economic reasons. To obtain improved end-to-end performance, overlay routing, which can change traffic routing in application layer, has gained attention. However, overlay routing often violates BGP routing policies and harms ISPs' interest. In order to take the advantage of overlay to improve the end-to-end performance, while overcoming the disadvantages, we propose a novel interdomain overlay structure, in which overlay nodes are operated by ISPs within an ISP alliance. The traffic between ISPs within the alliance could be routed by overlay routing, and the other traffic would still be routed by BGP. As economic structure plays very important role in interdomain routing, so we propose an effective and fair charging and pricing scheme within the ISP alliance in correspondence with the overlay routing structure. Finally, we give a simple pricing algorithm, with which ISPs can find the optimal prices in the practice. By mathematical analysis and numerical experiments, we show the correctness and convergence of the pricing algorithm.
key words: *BGP, interdomain, overlay routing, charging, pricing*

1. Introduction

Today's Internet is composed of thousands of interconnected networks operated by independent Internet service providers (ISPs). Two common relationships between ISPs are customer–provider, where one ISP pays another to forward its traffic, and peer–peer, where two ISPs agree that connecting directly to each other would mutually benefit both. The standard interdomain routing protocol is the Border Gateway Protocol (BGP) which is a path-vector protocol. BGP is fairly simple, and allows a wide variety of routing policies to override distance-based metrics according to the preferences of the ISPs. For example, ISPs often prefer customer-learned routes over routes learned from peers and providers in cases where both are available. This is because sending traffic through customers generates revenue for the ISP while sending traffic through providers costs the ISP money.

Recent research has found that BGP routes are often not optimal in terms of end-to-end performance metrics. For example, the authors in [2] reported that for 30% to 80% of paths there are alternative paths with significantly improved measures of quality. This non-optimal BGP routing arises for both technical and economic reasons. In terms of technology, BGP uses "shortest" path routing, where paths are chosen to minimize hop count. However, hop count correlates less well with performance than explicit measurements. In terms of economics, routing policies are driven by many concerns, especially contracts with neighboring ISPs and monetary prices. The problem of how to improve interdomain traffic routing performance has therefore been gaining attention.

Overlay routing [3]–[6] is a potential candidate for overcoming the functional limitations of BGP. The basic idea of an overlay network is to form a virtual network on top of the physical network so that overlay nodes can be customized to incorporate complex functionality without modifying the native IP network. Overlay networks typically route packets over paths made up of one or more overlay links to achieve a specific end-to-end objective.

However, routing in overlay networks often violates BGP routing policies [7]–[11], which could cause ISPs suffer from loss in revenue. Consider, for example, a hypothetical ISP-level connectivity graph as shown in Fig. 1. In that figure, overlay nodes exist in a, b, d and e. According to the current BGP structure, the traffic between b and e should go through the peering link between b and c rather than bd and dc. Therefore, d does not have to pay any one for that traffic. However, overlay nodes in b, d and e may choose to use path $bdce$ (although no overlay node exists in c), which results in that d is used for transiting traffic. This is a violation of the ISP's transit policy at d, and b has to pay both b and c for the transit traffic. We can see that d suffers loss in revenue because of overlay traffic. However, as overlays operate at the application layer, the violations typically go undetected by the native layer.

We propose an economically efficient interdomain overlay structure operated by ISPs based on an ISP alliance for taking advantage of the benefits of overlay networks while remaining ISP friendly. The ISP alliance in this paper is formed between adjacent ISPs. Each ISP in the alliance operates one or more overlay nodes, with all of the overlay nodes forming an overlay network. The traffic be-

Manuscript received October 17, 2013.
Manuscript revised August 1, 2014.
[†]The author was with the Graduate School of Information Science and Technology, Osaka University, Osaka-shi, 565–0871 Japan.
[††]The authors are with the Cybermedia Center, Osaka University, Osaka-shi, 560–0043 Japan.
[†††]The author was with the Cybermedia Center, Osaka University, Osaka-shi, 560–0043 Japan.
*This paper is an extended version of the work originally to be presented at ICNS 2013 [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@kindai.ac.jp
d) E-mail: nakano@cmc.osaka-u.ac.jp
DOI: 10.1587/transinf.2013EDP7367

tween ISPs in the alliance can be routed by overlay routing for better end-to-end performance. Unlike in BGP routing, the routes in the overlay network are decided by traffic source ISPs and multiple path routing is also employed. To take full advantage of bandwidth resources, there are no BGP routing policies. In the overlay network, each ISP is responsible for transiting traffic across its own network on behalf of its neighbors. As a reward, it receives money from the ISPs who send the traffic. The alliance is limited to adjacent ISPs for three reasons. First, according to the results in [5], an alliance formed by regional ISPs can improve end-to-end performance significantly. Second, it is easier to set up, manage, and maintain a regional ISP alliance compared with a global alliance. Third, it is possible to avoid harm to the interests of the ISPs caused by policy violations, which we go into detail in Sect. 3.

As ISPs are individual economic entities, we cannot separate routing from economic issues. ISPs always have dual roles: when sending traffic, they are customers, who pay the transit providers; when transiting traffic across their network, they are providers who charge the traffic sender. As customer, ISPs prefer the paths with better performance and lower price; while as provider, ISPs make pricing decision to maximize revenue. In this paper, we deal with the two roles in a unified effective economic structure. By the word “effective”, we mean that the ISP who is willing to pay more money can enjoy better routes. On the other hand, if specific route has better performance, the ISPs along it should gain more revenue by making optimal pricing decision. Besides effectiveness, fairness among ISPs along identical route is also important.

In the above routing and charging structure, we model the relationship between ISPs’ routing decision and properties of routes – performance and price. As customer, ISPs’ routing decision is decided by route performances, prices and ISP’s own property. The decision includes which path to choose, and how much traffic to send. Based on this model, we study ISPs’ pricing scheme as provider, and obtain the optimal price to maximize the revenue. In order to realize the optimal price in the practice, we study the non-cooperative pricing game [12] played by individual ISPs, and find that it is neither effective nor fair. We believe that if ISPs realize the undesired properties of the non-cooperative pricing game, they would seek cooperation. We then propose a pricing scheme based on route bundle – a bundle of routes having the same entrance ISP with each other – and prove that it is a better pricing scheme than non-cooperative pricing game. At last, we give a simple algorithm for route bundles to find the optimal prices, which can maximize the revenue. According to mathematical analysis and numerical experiments, we show that our pricing algorithm is correct, and can always converge to the optimal price.

The remainder of this paper proceeds as follows. Section 2 gives ISP alliance based overlay structure, routing and charging scheme. Section 3 goes into detail with ISPs’ routing, charging and pricing. We conclude in Sect. 5.

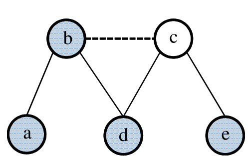

Fig. 1 An example of routing policy violation. Solid circles represent ISPs with overlay nodes in their domains, empty circle represents ISP with no overlay node in its domain. Solid lines represent transit relation, and dashed line represents peering relation.

2. ISP Alliance Based Interdomain Overlay Network Structure

To realize the advantages of overlay networks while overcoming the disadvantages, we propose an interdomain overlay network in which the overlay nodes are operated by ISPs who belong to the same alliance. In this section, we elaborate the structure of the ISP alliance, and give a brief discussion of the routing and charging scheme within the alliance.

2.1 Overlay Network Structure

An ISP alliance is formed by adjacent ISPs by bilateral contract. An example of interdomain overlay network based on ISP alliance is shown in Fig. 2. In this figure, we only show the border routers of each ISP. 6 shaded ISPs form an alliance, while the other two do not belong to the alliance. The three ISPs in the alliance construct an overlay network by setting virtual links between border routers. If the traffic demand is between two ISPs in the alliance, then it could be routed by overlay network with overlay routing. Otherwise the traffic demand is routed by the origin BGP routing. The two routing schemes co-exist, and can be applied for different kinds of traffic. That is, our approach does not preclude the Internet as it is today neither does it exclude BGP policies. Instead of competing with BGP, our architecture can be seen as a complementary tool for ISPs.

Note that the ISP alliance can only be formed by adjacent ISPs. An ISP with no direct connection to any ISP in a specific alliance cannot be accepted. This limitation makes it possible to avoid harming the interests of an ISP. For example, in Fig. 1, suppose a, b, and d form an ISP alliance.

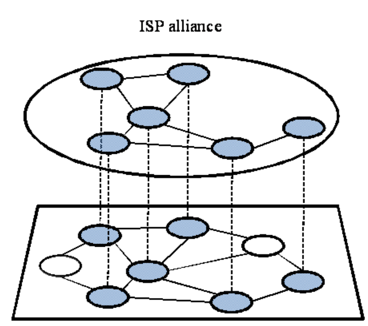

Fig. 2 Overlay network based on an ISP alliance. The solid circles are ISPs within the alliance, and the empty circles are ISPs not in the alliance.

	ISP alliance	Internet
Routing structure	Flat routing structure is adopted because an ISP alliance is assumed to consist of tens of ISPs. Source routing is employed, and multiple paths are allowed.	Hierarchical routing structure is adopted because the Internet is very large. Single path routing is generally used.
Business relations	Every ISP in the alliance provides transit service to all its neighbors.	Transit and peering. Routes violating the policies are not permitted.
Charging	Traffic users pay every ISP along the routes that their traffic traverse. 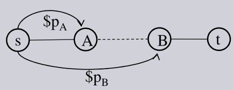	Customer ISP pays its providers. 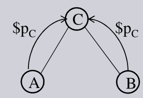

Fig. 3 Routing structure and policies in and outside of ISP alliance.

If e is accepted, then d may suffer a loss as illustrated in Sect. 1. Within the ISP alliance, we avoid harming the economic interests of ISPs caused by policy violation by using an effective and fair charging scheme that corresponds with the overlay routing structure. We give a brief introduction to the charging scheme here, and go into detail in Sect. 3.

Two charging schemes co-exist in the ISP alliance. One is the original Internet charging scheme in which ISPs make either provider–customer (transit) or peer–peer (peering) contracts with other ISPs. In a transit contract, the customer pays the provider for both up-streaming and down-streaming traffic, while in a peering contract, the traffic transport is free of charge in both directions. The BGP charging scheme applies to traffic with a source ISP or destination ISP outside the alliance. The other pricing scheme applies to the overlay network. In the overlay network, since every ISP provides transit services, the ISPs act as providers when transiting the traffic for their neighbors, and charge the traffic sender. When sending traffic to other ISPs in the alliance, they act as customers, and pay the ISPs along the routes they use.

2.2 Comparison of Routing and Charging in and Outside the ISP Alliance

In order to make the intra-alliance routing and charging scheme more clear, we give a brief comparison with the Internet. Figure 3 gives a summary of the comparison.

First, as the Internet is very large and ISPs are located all across the world, the Internet employs a hierarchical routing structure. Geographically distributed stub ISPs can connect to each other via only the transit services between local ISPs and the backbone. However, our ISP alliance is supposed to be constructed of tens of ISPs that are near each other geographically, and so a simple but effective flat routing structure is adopted. Second, the business relationships in the Internet include transit and peering. As we know, customer ISPs do not transit traffic for their providers, and peering ISPs do not provide transit services for each other. It turns out that some routes are illegal because they may violate routing polices even if they offer better performance. By comparison, in our ISP alliance, every ISP provides transit services for all neighbors in order to take advantage of all potential routes. As compensation, ISPs who provide transit services are paid by the traffic sender. Third, we have designed an intra-alliance charging scheme which is different from the charging scheme used in the Internet. In the intra-alliance charging scheme, the traffic sender s pays the other ISPs along the route to t.

Note that in the BGP routing and charging structure, a source ISP can only decide the next hop ISP and has no control over the rest of the route. The money that the source ISP pays to the next hop ISP is not necessarily positively correlated with whole route performance. The routing and charging structure we propose creates a correlation between the routing decision of the source ISP, and the route performance and price. In the next section, we go into detail on the charging and pricing scheme.

3. Routing, Charging and Pricing within the ISP Alliance

3.1 ISPs' Routing Decision and Pricing Strategies

The point of proposing an effective charging and pricing scheme is to capture the properties of ISP routing decisions. In the prominent work of [13], the authors introduced a model for capturing the relationship between traffic demand and prices of routes. Suppose the price of a route r is p_r, which is the sum of prices determined by every ISP along r. The relationship is then abstractly modeled by a demand function $d_r(p_r)$, which is strictly decreasing and differentiable. Moreover, if a function $g_r(p_r)$ is defined as

$$g_r(p_r) = -d_r(p_r)/d_r'(p_r),$$

then $g_r(p_r)$ must decrease with respect to p_r. Under this restriction on $g_r(p_r)$, the demand is inelastic when the price is low, which means the demand is dominated by the need of the ISPs to communicate. However, as the price increases, the demand becomes elastic, which means that price becomes a more important factor in the decisions of the ISP once the price passes a certain threshold. Although this model succeeds in capturing the properties of Internet services, it can only be used in single path routing systems. Moreover, in this model, price is the only factor affecting the routing decisions of an ISP. The overlay network in our work assumes multi-path routing in order to make full use of network resources. When making routing decisions, ISPs do not only consider the prices, but also the performance. In the rest of this section, we introduce our method for modeling the relationships among ISPs' routing decisions, price, and route performance.

If we suppose that there is only one route R_1 from a source ISP s to a destination t, then s has no choice but to send the traffic through R_1. Denoting the price of R_1 by p_1, the traffic volume is $d(p_1)$, where d is the aggregate traffic demand function. We assume that d is decreasing, differentiable, and that $-d(p)/d'(p)$ decreases with respect to p as in [13]. Now, if a better route R_2 is added with price $p_2 > p_1$, then $d(p_2)$ traffic changes to R_2, $d(p_1)-d(p_2)$ traffic remains on R_1, and the total traffic volume remains $d(p_1)$. Now suppose there are m routes $R_1, \ldots, R_m$ between a source ISP s and a destination t. The performance indicator of R_i is Per_i and the price is p_i. The performance indicator is logical, with larger Per_i indicating better performance. Without loss of generality, we assume $Per_1 < Per_2 < \ldots < Per_m$, and $p_1 < p_2 < \ldots < p_m$ correspondingly. The traffic demand from s to t is $d(p_1)$, because p_1 is the lowest price among all the routes. The traffic volume through R_i is $d(p_i) - d(p_{i+1})$. We can see that the traffic volume on R_i is dependent on the traffic volume on R_{i+1}. The only route on which the traffic volume does not depend on any other route is R_m, over which the traffic volume $f_m = d(p_m)$. Note that when we refer to the route performance, we do not assume some rigorous performance metrics. Performance is logical which is just used to reflect ISPs routing preferences. We keep the flexibility in the practice for ISPs to use any metric they prefer.

If we denote the revenue obtained from R_m by Re_m, then

$$Re_m = p_m d(p_m).$$

The ISPs on R_m can set price p_m to maximize Re_m independent to the other routes. The first order condition of Re_m with respect to p_m is

$$Re'_m(p_m) = d(p_m) + p_m d'(p_m).$$

Let $Re'_m(p_m) = 0$, then we have

$$p_m = -d(p_m)/d'(p_m).$$

As p_m is decreasing, the unique solution exists for the optimization problem. Denote the optimal price of R_m is p^*_m, then revenue of R_{m-1} is

$$Re_{m-1} = p_{m-1}(d(p_{m-1}) - d(p^*_m)).$$

The first order condition of Re_{m-1} with respect to p_{m-1} is

$$Re'_{m-1}(p_{m-1}) = d(p_{m-1}) + p_{m-1}d'(p_{m-1}) - d(p^*_m).$$

Let $Re'_{m-1}(p_{m-1}) = 0$, then we have

$$p_{m-1} = -d(p_{m-1})/d'(p_{m-1}) + d(p^*_m).$$

As p_{m-1} is decreasing with respect to p_{m-1}, the unique solution exists to the optimization problem. The optimal prices of the other routes can be obtained in the same way as above.

In this model, better routes have higher precedence for deciding the optimal price, and the optimal prices of worse routes always depend on the prices of better routes. The best route can decide the optimal price independently of any other route. We believe that this model is more efficient than models in which routing decisions are not correlated with performance.

3.2 Analysis of Route Based Pricing Strategies

Although the charging scheme in Sect. 3.1 seems ideal, it is difficult to realize in practice because ISPs are selfish and global cooperation cannot be expected. A natural and easy way to realize a route-based pricing scheme is the non-cooperative pricing game, in which prices are determined for each individual route independently by the ISPs on those routes. We illustrate this scheme with a simple network example in Fig. 4. In the figure, s is an ISP who sends traffic to t. A, B, and C are intermediate ISPs. There are two routes for s to reach t. One is $ABCt$, which is denoted by R_1, and the other is ACt which is denoted by R_2. Under route-based pricing, prices are determined based on the routes. As a hierarchical structure does not exist, the commodity is specific route, the customer is the ISP who sends traffic through that route, and the providers being paid are every ISP on that route. In the non-cooperative pricing game, each autonomous system could decide the prices for each route in a non-cooperative way to maximize the revenue obtained from that route. It seems natural and easy to realize because no cooperation among autonomous systems is needed. In actuality, however, we find that this method is neither effective nor fair.

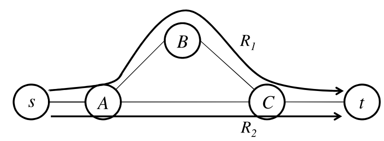

Fig. 4 A simple network example.

In Fig. 4, suppose route R_1 is better than R_2. Denote p_{A1} as A's price on R_1, p_{A2} as A's price on R_2, p_{B1} as B's price on R_1, p_{C1} as C's price on R_1, and p_{C2} as C's price on R_2. p_1 is the price of R_1, and $p_1 = p_{A1} + p_{B1} + p_{C1}$. p_2 is the price of R_2, and $p_2 = p_{A2} + p_{C2}$. f_1 is the traffic volume through $R1$, and f_2 is the traffic volume through R_2. The demand function is $d(p) = exp(-p^2)$, which is continuous, deceasing, and $-d(p)/d'(p)$ is also decreasing. According to the model in Sect. 3.1, $f_1 = d(p_1)$, and $f_2 = d(p_2) - d(p_1)$. If the ISPs on R_1 and R_2 play a non-cooperative pricing game fairly, the prices can be obtained as follows:

For ISP A:

$$\begin{aligned} \max Re_{A1} &= p_{A1}d(p_{A1} + p_{B1} + p_{C1}) \\ \max Re_{A2} &= p_{A2}(d(p_{A2} + p_{C2}) - d(p_{A1} + p_{B1} + p_{C1})), \end{aligned} \tag{1}$$

where Re_{A1} is A's revenue obtained from R_1, and Re_{A2} is A's revenue obtained from R_2.

For ISP B:

$$\max Re_{B1} = p_{B1}d(p_{A1} + p_{B1} + p_{C1}), \tag{2}$$

where Re_{B1} is B's revenue obtained from R_1.

For ISP C:

$$\begin{aligned} \max Re_{C1} &= p_{C1}d(p_{A1} + p_{B1} + p_{C1}) \\ \max Re_{C2} &= p_{C2}(d(p_{A2} + p_{C2}) - d(p_{A1} + p_{B1} + p_{C1})), \end{aligned} \tag{3}$$

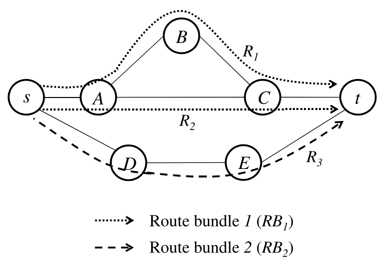

Fig. 5 A network example with route bundles.

where Re_{C1} is C's revenue obtained from R_1, and Re_{C2} is C's revenue obtained from R_2. Then the only Nash equilibrium is achieved when $p_{A1} = p_{B1} = p_{C1} = 0.24$, and $p_{A2} = p_{C2} = 0.15$. The traffic through R_1 is $f_1 = 0.61$, the traffic through R_2 is $f_2 = 0.31$. A's revenue is 0.19, B's revenue is 0.15, and C's revenue is 0.19.

In the above example, each ISP plays the game by considering R_1 and R_2 separately, and the result is efficient and fair for ISPs on the same route. But if, for example, A, realizes that it is disjoint point of R_1 and R_2, it would change to an alternative behavior as follows:

$$\begin{aligned}\max R_A =& R_{A2} + R_{A1}\\ =& p_{A1}d(p_{A1} + p_{B1} + p_{C1}) + p_{A2}(d(p_{A2} + p_{C2})\\ &- d(p_{A1} + p_{B1} + p_{C1})).\end{aligned} \tag{4}$$

When Nash equilibrium is achieved, $p_{A1} = 0.82$, $p_{A2} = 0.34$, $p_{B1} = 0.12$, $p_{C1} = 0.12$, and $p_{C2} = 0.34$. The traffic though R_1 is $f_1 = 0.11$, and the traffic through R_2 is $f_2 = 0.40$. The revenue of A is 0.23, the revenue of B is 0.01, and C's revenue is 0.06. From the above results, we see that the traffic through the better route R_1 decreases dramatically, which reduces the efficiency of the traffic routing. Moreover, on both R_1 and R_2, A obtains more revenue than the other ISPs on the identical route, which is unfair to the other ISPs. As above, the non-cooperative pricing game based on routes is not acceptable. If ISPs realize the undesirable properties of the non-cooperative pricing game, they will look for some kind of cooperation. In the next section, we present our pricing scheme based on route bundles.

Note that, the utility of players in Game Theory has various expressions in different problems. In this paper, stub ISPs and transit ISPs have different utilities. Stub ISPs care about both monetary cost paid to transit ISPs and quality of routes, while transit ISPs care only revenue. Therefore, when modeling transit ISPs' utility, only revenue is considered. In the next section, we give our pricing scheme based on route bundles.

3.3 Pricing Based on Route Bundle

In this paper, route bundle is defined as a set of routes having the same entrance ISP with each other. For example, in Fig. 5, R_1 and R_2 have the same entrance A, so that they are in the same route bundle RB_1. R_3 has different entrance from routes in RB_1, so that R_3 itself is route bundle RB_2. In fact, the inefficiency and unfairness in the non-cooperative route based pricing only happens at the disjoint point of multiple routes within identical route bundle. With pricing based on route bundle, the price is determined for route bundle, rather than individual route, so that the undesirable properties with route based pricing do not exist. In order to realize bundle based pricing scheme, cooperation with ISPs in the same bundle is required. Source ISP s would be noticed by the entrance A and D the price for RB_1 and RB_2 respectively, and decides how to route traffic. The traffic sent to RB_1 also has two options R_1 and R_2, and ISPs can choose a better one freely. The accounting can be done as follows. As source routing is employed, the route information can be found in the head of the packet. When a packet with entrance A and destination t enters A, A could write the price in the head of the packet, and forward it. Thus, every ISP on the route can keep record of the price and the packet amount. In the end of the contract cycle, the ISPs can share the revenue obtained from routes in identical route bundle. The share of each ISP can be calculated with bilateral negotiation. Although in the overlay network, the hierarchical structure does not exist, in fact, neighboring ISPs do not really have equal position. In practice, the two ISPs have either customer-provider contract or peering contract, so that ISPs may not be satisfied to share the revenue equally. One possible negotiation is, neighboring ISPs bargain with each other to decide the relative sharing. After every pair of ISPs finish the bargaining, the share of every ISP can be calculated.

Note that although the route bundle based pricing scheme is closely related with routing performance, it is independent on any detailed QoS metric. The routing and pricing scheme can work well only if ISPs rank parallel route bundles with similar long term performance metrics, which gives flexibility to ISPs to decide the QoS measurements.

3.4 Pricing Algorithm

Section 3.3 showed that the price of a specific route bundle is decided by the entrance ISP of the bundle. In fact, what the entrance ISP faces is simple optimization problem with just a single variable. Although the objective function may be neither convex nor concave, we have shown that it has a unique optimal point in Sect. 3.1. Therefore, it can be solved by a one-dimensional search method. The entrance ISP could set a starting price from the empirical value p_0, and then update it periodically. Supposing prices are updated in steps of u, the ISP can update the price as follows:

1. Set the price p to the empirical value p_0
2. Loop step 3 to step 5 periodically until the optimal price being found
3. Increase p by one unit. If the revenue decreases, go to step 5. Else, go to step 4
4. Keep increasing p, until revenue begins decreasing
5. Keep decreasing p, until revenue begins decreasing

This method is valid for the following reason. Suppose a set of route bundles $RB_1, \ldots, RB_n$ are competing for traffic with each other. Without loss of generality, we assume the route bundles are in ascending order with respect to performance. The revenue of a specific route bundle RB_i can be represented by

$$Re_i = p_i(d(p_i) - d(p^*_{i+1})).$$

The first order condition is

$$Re'_i(p_i) = (p_i + \frac{d(p_i)}{d'(p_i)} - \frac{d(p^*_{i+1})}{d'(p_i)})d'(p_i), \tag{5}$$

where p^*_{i+1} is the optimal price of RB_{i+1}. As $-\frac{d(p_i)-d(p_{i+1})^*}{d'(p_i)}$ is decreasing, a unique solution to maximize Re_i exists, which is denoted by p^*_i. If $p_i \leq p^*_i$, then $Re'_i(p_i) \geq 0$, which means that Re_i increases with respect to p_i in $(0, p^*_i]$. If $p_i > p^*_i$, then $Re'_i(p_i) < 0$, which implies that Re_i decreases with respect to p_i. The validity of the pricing method can then be proved straightforwardly. We also find that, with this method, entrance ISPs can determine the optimal prices without knowing the exact formula for the demand function d.

Note that, if multiple route bundles have the same performance, we need to make a tie–breaking rule. In this work, the traffic source ISP should choose any one of the route bundles to transmit traffic. Also note that, the performance of routes we talk about means the average performance in long period. Therefore, it is not necessary for ISPs to update the price too frequently. ISPs can update prices weekly or monthly, so that the inter-domain routing stability will not be affected significantly.

3.5 Numerical Experiments

In this section, we describe numerical experiments for showing the validity and convergence of our pricing method. We conduct experiments based on a network with as shown in Fig. 6.

In the figure, ISP 1, 2 and 3 are source ISPs transmitting traffic to ISP 10. We assume links have the same propagation delay, and queuing delay is not considered. Therefore, the hop count can represent the latency, and latency is taken as the performance indicator in the experiments. The route bundles and routes they contain in Fig. 6 can be summarized as Fig. 7. At the beginning of the experiments, entrance ISPs set prices based on values from previous experience, and then adjust the prices periodically and independently. To make the experiments more clear, we assume that competing route bundles adjust prices in turn. Prices are assumed to be adjusted in steps of 1.0. Changes in price and revenue with respect to time are shown in Figs. 8 (a), 9 (a), 8 (b), 9 (b), 8 (c) and 9 (c).

Note that between ISP 1 and 10, there are two route bundles with entrance ISP 5 and 7, which have the same latency. According to our tie–breaking rule, 1 can choose any route bundle to transmit traffic. We assume route bundle with entrance ISP 5 (route bundle 5) is chosen. The initial price is set as 12.0 which is higher than the optimal price. After some steps of adjusting, the optimal price 7.0 is found (Fig. 8 (a)), and the revenue achieves the highest (Fig. 9 (a)). Between ISP 2 and 10, there are also two route bundles 8 and 4. The route in route bundle 8 has less hop count than the routes in route bundle 4, which indicates route bundle 8

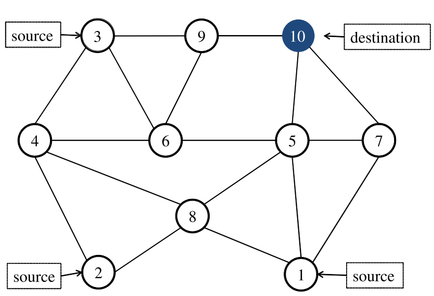

Fig. 6 Network for experiment. Circles represent ISPs.

Source ISP	Route bundle (distinguished with entrance ISP)	The best routes in the bundle
1	5	(1,5,10)
	7	(1,7,10)
2	8	(2,8,5,10)
	4	(2,4,8,5,10), (2,4,6,9,10), (2,4,3,9,10) and (2,4,6,5,10)
3	9	(3,9,10)
	6	(3,6,9,10) and (3,6,5,10)
	4	(3,4,6,9,10), (3,4,6,9,10) and (3,4,8,5,10)

Fig. 7 Route bundles and routes they contain.

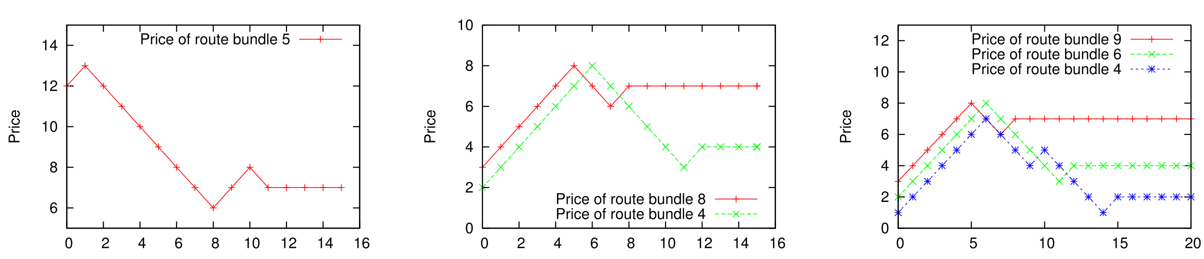

(a) Price of route bundle 5
(b) Price of route bundles 8 and 4
(c) Price of route bundles 9, 6 and 3

Fig. 8 Price of route bundles.

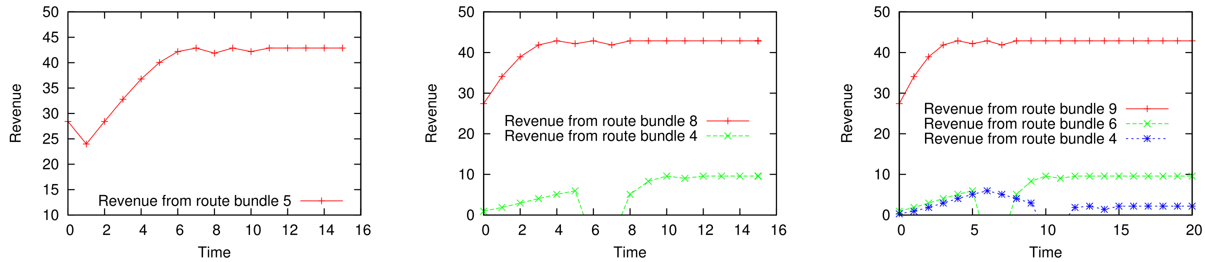

(a) Revenue of route bundle 5 (b) Revenue of route bundles 8 and 4 (c) Route bundle 3

Fig. 9 Revenue of route bundles 9, 6 and 3.

is better than 4. At the beginning, route bundle 8 initializes p_0 as 2.0 and route bundle 4 initializes p_0 as 1.0. Both of the prices are lower than the optimal prices. The price adjusting process is shown in Fig. 8 (b). In Figs. 9 (b), we can find that route bundle 4 receives 0 revenue in a period of time. This is because during that period, route bundle 4 sets higher price than route bundle 8, so that ISP 2 transmits all the traffic through route bundle 8. From Figs. 8 (b) and 8 (c), we can also find that the convergence of route bundles depends on the converge of better route bundles. The price adjusting of a route bundle can not converge before all the better route bundles finish adjusting prices.

4. Further Discussion on the Practice Issues

In this section, we discuss some practice issues about source routing technologies and the network scale.

In recent years, some important steps supporting source routing have been made. The authors of [14] introduced a distributed inter-ISP service plane, which is coupled with a Path Computation Element (PCE) based control plane. In their framework, routing is source-based at the AS-level and distributed at the router-level. The authors of [6] further extended the source routing framework, and made explicit approaches on composing ISP alliance, calculating inter-ISP routes and inter-domain QoS issues technically. Based on the above work, we believe that inter-ISP source routing in ISP alliance is promising and realistic in the future.

In the experiments in Sect. 3.5, we present some basic properties of the pricing method with a simple network model. However, in a large scale network, there are two significant differences. First, longer routes would appear. Second, some ISPs would connect with a large number of neighboring ISPs. In this section, we explain that our pricing method is still feasible and efficient for large scale network.

As the price of a route bundle is only determined by the entrance ISP, the price convergence time is not relevant to the length of routes. In fact, the convergence time for a route bundle is, if it is not the best route bundle between specific source–destination pair, dependent on the convergence times of all the better route bundles between that source–destination pair. It implies that if there are plenty of route bundles between one source–destination pair, the price adjusting process for the route bundles with poorer performance might be very slow. According to the definition of route bundle, theoretically, the maximum number of route bundles for an ISP to a specific destination is equal to the number of its neighbors. Thus in a large scale network, this number for some ISPs could be very large. However, we think in the practice, the slow price converge process caused by too many competing route bundles could hardly happen, because it is unlikely that the traffic source ISP would like to split the traffic to too many route bundles. As we introduced in Sect. 3.1, each ISP has an upper limit in the traffic volume to send. If the stable price of a route bundle is no more than a certain threshold γ_i, then ISP i will send all the traffic through that route bundle and all the better route bundles. The route bundles with poorer performance have to withdraw from the pricing game with 0 traffic from ISP i. As a result, the number of route bundles used by ISP i is decided by γ_i, which is completely determined by ISP i. We have reason to believe that ISP i will adjust γ_i in order to keep the number of route bundles it uses in reasonable range, otherwise, it will suffer from the extra overhead of splitting the traffic too much and slow price convergence.

5. Conclusion

In this paper, we proposed an interdomain overlay network in which nodes are operated by ISPs within an ISP alliance. The traffic between ISPs within the alliance could be routed by overlay routing to overcome the functional limitations of BGP. The problem of BGP policy violations can also be addressed through the definition of the ISP alliance and the economic framework within the alliance.

As ISPs are individual economic entities, interdomain routing issues cannot be separated from economic factors. We studied the routing decision of ISPs facing multiple routes, and modeled the relationship between the routing decisions of the ISPs and the route properties of performance and price. Based on this model, we obtained the optimal price for each route for maximizing revenue.

Although optimal prices exist, these are difficult to realize in practice. We show that a non-cooperative pricing game by selfish ISPs would lead to ineffective and unfair results. We believe that if ISPs realize this, they would seek cooperation. We then proposed a pricing scheme based on route bundles, which are bundles of routes having the same entrance ISP as each other, and show that it is better than the non-cooperative pricing game. Finally, we presented a simple pricing method with which ISPs can find the optimal prices without precise knowledge of traffic source ISPs. We showed the validity and convergence of the pricing method through mathematical analysis and numerical experiments.

In order to make the research more impractical, in the further, we are planning to address some related issues in both technical and economic aspects. First, we would like to make the performance metrics explicitly which can be accepted by most of the ISPs. Second, as our proposal is based on source routing, the corresponding routing discovery mechanism is needed. Third, in our proposal, revenue from a route bundle should be distributed among all the participating ISPs within it. We are planning to explore the negotiation mechanism in the future.

References

[1] X. Shao, G. Hasegawa, Y. Taniguchi, and H. Nakano, "Economically efficient interdomain overlay network based on ISP alliance," Proc. ICNS 2013, pp.153–159, March 2013.

[2] 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, pp.50–59, Jan. 1999.

[3] D. Anderson, H. Balakrishnan, M.F. Kaashoek, and R. Morris, "Resillient overlay networks," Proc. SOSP 2001, pp.131–145, Nov. 2001.

[4] Z. Duan, Z.L. Zhang, and Y.T. Hou, "Service overlay networks: SLAs, QoS, and bandwidth provisioning," IEEE/ACM Trans. Netw., vol.11, no.6, pp.870–883, Dec. 2003.

[5] Y. Hei, A. Nakao, T. Ogishi, T. Hasegawa, and S. Yamamoto, "AS alliance for resilient communication over the Internet," IEICE Trans. Commun., vol.E93-B, no.10, pp.2706–2714, Oct. 2010.

[6] J.R.S. Secci and A. Pattavina, "AS-level source routing for multi-provider connection-oriented services," Comput. Netw., vol.54, no.14, pp.2453–2467, Oct. 2010.

[7] 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 2006, pp.51–59, June 2006.

[8] S. Seetharaman and M. Ammar, "Characterizing and mitigating inter-domain policy violations in overlay routes," Proc. IEEE ICEP 2006, pp.259–268, Nov. 2006.

[9] G. Hasegawa, Y. Hiraoka, and M. Murata, "Evaluation of free-riding traffic problem in overlay routing and its mitigation method," IEICE Trans. Commun., vol.E92-B, no.12, pp.3774–3783, Dec. 2009.

[10] 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," Comput. Netw., vol.52, no.15, pp.2961–2974, Oct. 2008.

[11] X. Shao, G. Hasegawa, Y. Taniguchi, and H. Nakano, "The implications of overlay routing for ISPs' peering strategies," IEICE Trans. Inf. & Syst., vol.E96-D, no.5, pp.1115–1124, May 2013.

[12] G. Fudenberg and J. Tirole, Game theory, MIT Press, 2000.

[13] L. He and J. Walrand, "Pricing and revenue sharing strategies for Internet service providers," IEEE J. Sel. Areas Commun., vol.24, pp.942–951, May 2006.

[14] R. Douville, J.L. Roux, J. Rougier, and S. Secci, "A service plane over the PCE architecture for automatic multidomain connection-oriented services," IEEE Commun. Mag., vol.46, no.6, pp.94–102, June 2008.

Xun Shao received his B.E. degree from Civil Aviation University of China, in 2005, M.E. degree from Beijing Jiaotong University, China, in 2008, and Ph.D. degree from Osaka University, Japan, in 2013. He is currently an researcher at the National Institute of Information and Communications Technology (NICT), Japan. His research interests include Internet structure, overlay technologies and sensor network. He is a member of IEICE.

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 and IEICE.

Yoshiaki Taniguchi received his B.E., M.E. and Ph.D. degrees from Osaka University, Japan, in 2002, 2004 and 2008, respectively. He was an Assistant Professor at the Cybermedia Center, Osaka University from 2008 to 2014. Since 2014, he has been a Lecturer at the Faculty of Science and Engineering, Kindai University, Japan. His research interests include wireless networks and energy management systems. He is a member of IEEE, IEICE, IPSJ and IIEEJ.

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. He was a Professor at the Cybermedia Center, Osaka University until from 2004 to 2013. His research is in the area of ubiquitous networks. He is a member of IEEE, IEICE and ITE.