Economically Efficient Interdomain Overlay Network Based on ISP Alliance

Xun Shao
Graduate School of Information Science and Technology
Osaka University
Osaka, Japan
Email: x-shao@ist.osaka-u.ac-jp

Go Hasegawa, Yoshiaki Taniguchi, Hirotaka Nakano
Cybermedia Center, Osaka University
Osaka, Japan
Email: (hasegawa, y-tanigu, nakano)@cmc.osaka-u.ac.jp

***Abstract*—As interdomain routing protocol, BGP is a 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 overcome 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 is still routed by BGP. As economic structure plays very important role in interdomain routing, we then propose an effective and fair charging and pricing scheme within the ISP alliance in correspondence with the overlay routing structure. At last, 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.**

Keywords-BGP; interdomain; overlay routing; charging; pricing

I. Introduction

The Internet is composed of thousands of networks owned by Internet Service Providers (ISPs), which are selfish, often competing economic entities. The task of establishing routes between ISPs is called interdomain routing. The standard interdomain routing protocol is the Border Gateway Protocol (BGP), which is a path-vector protocol. BGP allows routing policies to override distance-based metrics with policy-based metrics. ISPs often wish to control next hop selection so as to reflect agreements or relationships they have with their neighbors. Two common relationships ISPs have are: customer-provider, where one ISP pays another to forward its traffic, peer-peer, where two ISPs agree that connecting directly to each other would mutually benefit both. ISPs often prefer customer-learned routes over routes learned from peers and providers when both are available. This is because sending traffic through customers generates revenue for the ISP while sending traffic through providers costs the ISP money.

Although BGP is the sole interdomain routing protocol currently, the authors in [1] found that the default BGP paths are not often optimal with respect to end-to-end performance. At any time, for 30 to 80 percent of the paths we can find there are alternative paths with significantly improved measures of quality. There are both technical and economic reasons to expect that BGP routing is non-optimal. Theoretically, the BGP uses "shortest" path routing, where paths are chosen to minimize hop count. However, hop count correlates less well with performance than explicit measurements. Moreover, economic considerations can also limit routing options. Routing policies are driven by many concerns especially the contracts with neighbor ISPs and monetary prices.

In order to realize better end-to-end performance, overlay networks [2]–[6] have recently gained attention as a viable alternative to overcome functionality limitations of BGP. The basic idea of overlay networks is to form a virtual network on top of the physical networks so that overlay nodes can be customized to incorporate complex functionality without modifying the native IP network. Typically, these overlays route packets over paths made up of one or more overlay links to achieve a specific end-to-end objective.

Routing in overlay networks often violates BGP routing policies [7]–[11]. Consider, for example, a hypothetical ISP-level connectivity graph as shown in Fig. I. In that figure, overlay nodes exist in a, b, d and e. Overlay nodes are trying to obtain the best possible route to each other. Overlay node in b can route data to overlay node e using the overlay path bdce, which results in dfs ISP being used for transiting traffic. This is a violation of the ISP's transit policy at d. From an economic perspective, we see that the performance improvement comes at the expense of d, as d has to pay b and c for the overlay traffic through the illegitimate path. Because overlays operate at the application layer, the violations typically go undetected by the native layer. In order to take the advantage of overlay to improve the end-to-end performance, while maintaining the ISPs' benefit, we propose an economically efficient interdomain overlay structure operated by ISPs based on ISP alliance. The ISP alliance in this paper is formed by adjacent ISPs. Each ISP in the alliance operates one or more overlay nodes, and all the overlay nodes form an overlay network. The traffic between ISPs in the alliance can be routed by overlay routing

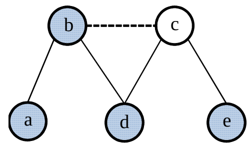

Figure 1. An example of routing policy violation. Solid circles represent ISPs with overlay nodes in their domains, empty circle represents ISP with no overlay node in its domain. Solid lines represent transit relation, and dashed line represents peering relation

for better end-to-end performance. The routes are chosen by traffic source ISPs, and multiple path routing is also employed. The alliance is limited to adjacent ISPs for three reasons: first, according to the results in [12] and [6], the alliance formed by regional ISPs can improve the end-to-end performance significantly, second, it is easy for regional ISP alliance to set up, manage and maintain compared with global one, third, the ISPs' loss of interest caused by policy violation can be avoided, which we go into detail in Section III. In the overlay network, in order to take full advantage of bandwidth resource, the business relationships such as provider-customer and peer-peer do not exist, neither BGP routing policies. In the overlay network, each ISP is responsible for transiting traffic across its network for its neighbors. As reward, it can receive money from the ISPs who send traffic. Within the ISP alliance, we avoid harming ISPs' interests caused by policy violation by introducing novel economic structure.

As ISPs are individual economic entities, we cannot separate routing from economic. 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 [13] 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 II gives ISP alliance based overlay structure, routing and charging scheme. Section III goes into detail with ISPs' routing, charging and pricing. We conclude in Section IV.

II. ISP ALLIANCE BASED INTERDOMAIN OVERLAY NETWORK STRUCTURE

In order to take the advantages of overlay networks while overcome the disadvantages, we propose an interdomain overlay network, in which overlay nodes are operated by ISPs belonging to the same alliance. In this section, we elaborate the structure of the ISP alliance, and make a brief discussion with the routing and charging scheme within the alliance.

A. 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

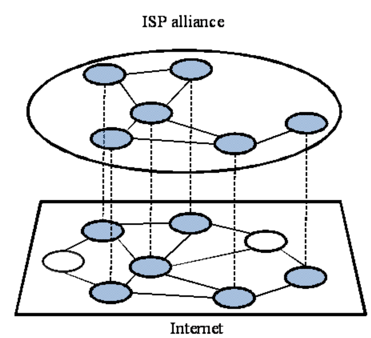

Figure 2. Overlay network based on an ISP alliance. The solid circles are ISPs within the alliance, and the empty circles are ISPs not in the alliance.

show the border routers of each ISP. ISP1, ISP2 and ISP3 form an alliance, while ISP4 and ISP5 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

 ISBN: 978-1-61208-256-1

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. By this limitation, ISPs' loss of interest can be avoided. For example, in Fig. I, suppose a, b and d form an ISP alliance, if e is accepted, then d may suffer a loss of interest as illustrated in Section I. Within the ISP alliance, we avoid harming ISPs' economic loss caused by policy violation with effective and fair charging scheme in correspondence with the overlay routing structure. We make a brief introduction of the charging scheme, and go into detail in Section III.

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

B. Comparison of routing and charging in and outside the ISP alliance

In order to make the intra-alliance routing and charging scheme more clearly, we make a brief comparison with the Internet. Fig. 3 shows the summery of the comparison. First, as the Internet is very huge and ISPs are located all across the world, hierarchical routing structure is adopted. Geographic distributed stub ISPs can connect to each other only with the transit service of local ISPs and the backbone. However, our ISP alliance is supposed to construct with tens of ISPs near each other geographically, so that a simple but effective flat routing structure is adopted. Second, the business relationships in the Internet include transit and peering. As known to us, customer ISPs do not transit traffic for their providers, and peering ISPs do not provide transit service for each other. It turns out that some routes are illegal because they may violate the routing polices even if they have better performance. As comparison, in our ISP alliance, every ISP provides transit service for all its neighbors in order to take advantage of all potential routes. As compensation, the ISP who provides transit service will be paid by the traffic sender. Third, we design a charging scheme intra-alliance, which is different from the charging scheme of the Internet. In the intra-alliance charging scheme, the traffic sender s pays the other ISPs along the route to t.

	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. 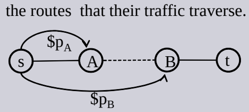	Customer ISP pays its providers. 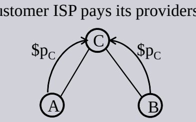

Figure 3. Routing structure and policies in and outside of ISP alliance

Note that with BGP routing and charging structure, a source ISP can only decide the next hop ISP, and has no control to the rest of the route. It is not necessary that the money the source ISP pays to the next hop ISP is positively correlated with the whole route performance. But with the routing and charging structure we propose, the correlation between source ISPs' routing decision, route performance and price is created. In the next section, we go into detail with the charging and pricing scheme.

III. Routing, charging and pricing within the ISP alliance

A. ISPs' routing decision and pricing strategies

The point to propose effective charging and pricing scheme is well capturing the properties of ISPs' routing decision. In the prominent work of [14], the authors introduce a model to capture 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. Then the relationship is 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 be decreasing with respect to p_r. With this restriction on $g_r(p_r)$, the demand is inelastic when price is low, which means the demand is dominated by ISPs' need to communicate; but when prices increases, the demand becomes elastic, which means price becomes a more important factor in ISPs' decisions once price passes a certain threshold. This model succeeds in grabbing the properties of Internet service, however, it can only be used in single path routing system. Moreover, in this model, price is the only factor to affect ISPs' routing decision. In the overlay network in our work, multi-path routing is supposed 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 to model

 ISBN: 978-1-61208-256-1

the relationships among ISPs' routing decision, price and performance of routes.

Suppose there is only one route R_1 from source ISP s to destination t, then s has no choice but to send the traffic through R_1. Denote the price of R_1 as p_1, then the traffic volume is $d(p_1)$, where d is the aggregate traffic demand function. We assume d is decreasing, differentiable, and $-d(p)/d'(p)$ is decreasing with respect to p as in [14]. Now, if a better route R_2 is added with price $p_2 > p_1$, then $d(p_2)$ traffic would change to R_2, $d(p_1) - d(p_2)$ traffic will remain on R_1, and the total traffic volume remains $d(p_1)$. Now suppose there are m routes $R_1, ..., R_m$ between source ISP s and one destination t. The performance indicator of R_i is Per_i and the price is p_i. The performance indicator is logical, and larger Per_i indicates better performance. Without loss of generality, we assume $Per_1 < Per_2 < ... < Per_m$, and $p_1 < p_2 < ... < p_m$ correspondingly. The traffic demand from s to t will be $d(p_1)$, because p_1 is the lowest price of all 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 routes is R_m, and the traffic volume $f_m = d(p_m)$.

Denote the revenue obtained from R_m as 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 $-d(p_m)/d'(p_m)$ 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 $-d(p_{m-1})/d'(p_{m-1})$ 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.

We can see that in this model, better route can decide optimal price with higher priority, and the optimal price of worse route always depends on the price of better route. The best route can decide optimal price independently to any other route. We believe that this model is more efficient than the models in which routing decision is not correlated with performance.

B. Analysis of route based pricing strategies

Although the charging scheme in Section III-A seems ideal, it is difficult to realize it in practice, because ISPs are selfish, and global cooperation cannot be expected. A very natural and easy way to realize the route based pricing scheme is non-cooperative pricing game, in which prices are determined for every individual route by the ISPs on those routes independently. We illustrate this scheme with a simple network example shown in Fig. 4. In the figure, s is an ISP

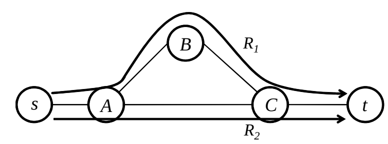

Figure 4. A simple network example

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 as R_1, and the other is ACt which is denoted as R_2. With the route based pricing, prices are determined based on routes. As the hierarchical structure does not exist, the commodity is specific route, the customer is the ISP who sends traffic through that route, and the provider being paid is every ISP on that routes. With non-cooperative pricing game, each AS could decide price 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 ASes is needed. But in fact, we find that this method is nether effective nor fair.

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 Section III-A, $f_1 = d(p_1)$, 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:

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

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:

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

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:

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

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 can find 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 route would not be acceptable. If the ISPs realize the undesirable properties of non-cooperative pricing game, they will look for some kind of cooperation. In the next section, we give our pricing scheme based on route bundle.

C. 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.

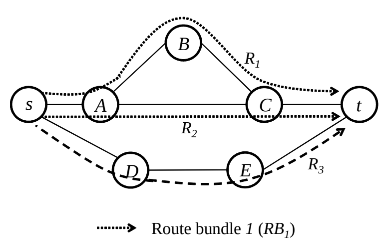

Figure 5. A network example with route bundles

D. Pricing algorithm

Section III-C 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 Section III-A. 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 $RB_1, ..., RB_n$ 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.

E. 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.

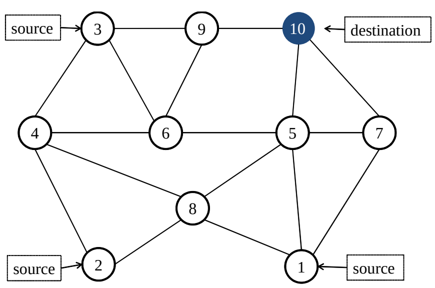

Figure 6. Network for experiment. Circles represent ISPs

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)

Figure 7. Route bundles and routes they contain

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 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.

IV. CONCLUSION

In this paper, we propose 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 functionality limitations of BGP. According to the definition of the ISP alliance and the economic structure within the alliance, the BGP policy violation problem can also be addressed.

As ISPs are individual economic entities, interdomain routing issues cannot be separated from economic factor. We study ISPs' routing decision facing multiple routes, and model the relationship between ISPs' routing decision and route properties – performance and price. Based on this model, we obtain the optimal price for each route to maximize the revenue.

Although the optimal price exists, it is difficult to realize it in practice. We show that a non-cooperative pricing game by selfish ISPs would lead to ineffective and unfair result. We believe that if ISPs realize the above fact, 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 show that it is better than the non-cooperative pricing game. At last, we give a simple pricing algorithm with which ISPs can find the optimal prices without precise knowledge of traffic source ISPs. With

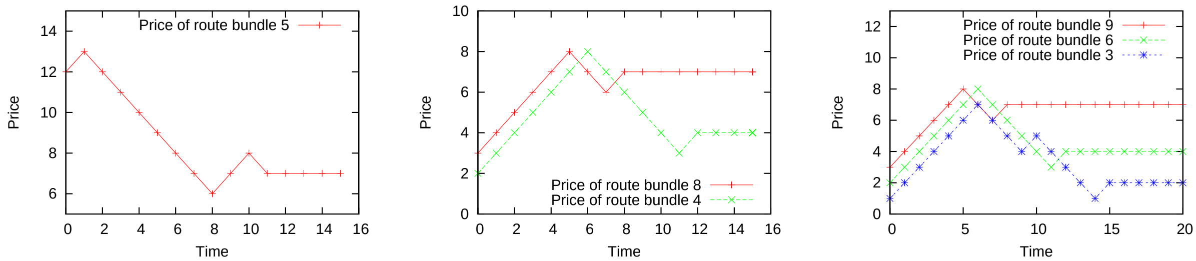

(a) Price of route bundle 5 (b) Price of route bundles 8 and 4 (c) Price of route bundles 9, 6 and 3

Figure 8. Price of route bundles

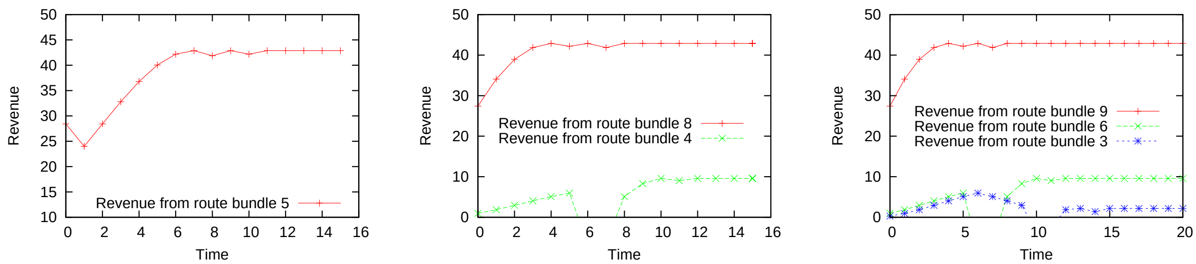

(a) Revenue of route bundle 5 (b) Revenue of route bundles 8 and 4 (c) Route bundle 3

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

mathematical analysis and numerical experiments, we show the correctness and convergence of the pricing algorithm.

REFERENCES

[1] 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, January 1999.

[2] D. Anderson, H. Balakrishnan, M. F. Kaashoek, and R. Morris, "Resillient Overlay Networks," in *Proceeding of the ACM SIGCOMM*, November 2001.

[3] Z. Duan, Z. L. Zhang, and Y. T. Hou, "Service Overlay Networks: SLAs, QoS, and Bandwidth Provisioning," *IEEE/ACM Transactions on Networking*, vol. 11, pp. 870–883, December 2003.

[4] H. T. Tran and T. Ziegler, "A Design Framework towards the Profitable Operation of Service Overlay Networks," *Computer Networks*, vol. 51, pp. 94–113, January 2007.

[5] A. Capone, J. Elias, and F. Martignon, "Routing and Resource Optimization in Service Overlay Networks," *Computer Networks*, vol. 53, pp. 180–190, October 2009.

[6] Y. Hei, A. Nakao, T. Ogishi, T. Hasegawa, and S. Yamamoto, "AS Alliance for Resilient Communication over the Internet," *IEICE Transactions on Communication*, vol. E93-B, no. 10, pp. 2706–2714, October 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," in *Proceedings of IEEE IWQoS 2006*, June 2006, pp. 51–59.

[8] S. Seetharaman and M. Ammar, "Characterizing and Mitigating Inter-Domain Policy Violations in Overlay Routes," in *Proceedings of IEEE ICEP 2006*, November 2006, pp. 259–268.

[9] G. Hasegawa, Y. Hiraoka, and M. Murata, "Evaluation of Free-Riding Traffic Problem in Overlay Routing and Its Mitigation Method," *IEICE Transactions on Communication*, vol. E92-B, no. 12, pp. 3774–3783, December 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," *Computer Networks*, vol. 52, pp. 2961–2974, October 2008.

[11] X. Shao, G. Hasegawa, Y. Taniguchi, and H. Nakano, "The Implication of Overlay Routing on ISPs' Connecting Strategies," in *Proceedings of ITC 2011*, no. September, 2011, pp. 286–293.

[12] R. Winter, "Modeling the Internet Routing Topology with a Known Degree of Accuracy," in *Proceedings of ACM/IEEE/SCS PADS 2009*, June 2009.

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

[14] L. He and J. Walrand, "Pricing and Revenue Sharing Strategies for Internet Service Providers," *IEEE Journal on Selected Areas in Communications*, vol. 24, pp. 942–951, May 2006.