Decreasing ISP transit cost in overlay routing based on multiple regression analysis

Kazuhito MATSUDA #1, Go HASEGAWA #, Masayuki MURATA #

#*Graduate School of Information Science and Technology, Osaka University*
1-5 Yamadaoka Suita, Osaka 560-0871, Japan
1k-matuda@ist.osaka-u.ac.jp

***Abstract*— Overlay routing, which is a routing mechanism that works at the application level, selects a route for overlay network traffic based on user-perceived metrics, such as end-to-end latency and available bandwidth. On the other hand, IP routing is configured based primarily on the commercial relationships with neighboring ISPs. This mismatch is one of the primary reasons why user-perceived performance is improved by overlay routing. However, overlay routing may be harmful to the monetary cost architecture of ISPs. In the present paper, we propose an overlay routing mechanism that can decrease the transit cost of ISPs as well as improve user-perceived performance. The proposed mechanism selects the overlay-level route using the number of transit links on the route, which is estimated by multiple regression analysis of measurable network performance. We confirm the effectiveness of the proposed mechanism in the PlanetLab environment and demonstrate that the proposed mechanism can maintain the level of user-perceived performance while significantly decreasing the number of traversing transit links by overlay routing.**

I. Introduction

Overlay routing, which is a type of service provided by overlay networks, does not assume a specific application, but instead concentrates only on routing for overlay network traffic. One early and typical example is the Resilient Overlay Network (RON) [1], in which each overlay node measures the end-to-end latency and packet loss ratio of the network path to other nodes and determines the route for the overlay network traffic, originating from the node, which can be either a direct route from the node to the destination node or a relay route that traverses other node(s) before reaching the destination node. In the present paper, we use the terms "overlay routing" and "IP routing" to refer to traffic routing at the application level and the IP level, respectively.

Recent studies have revealed that user-perceived performance can be improved by overlay routing [2-8]. Most of these studies focused on evaluations with only end-to-end latency as a routing metric. On the other hand, in [9], we revealed that a significant performance improvement can be provided by overlay routing based on bandwidth-related information (available bandwidth and TCP throughput). Such a performance improvement is caused primarily by the policy mismatch between IP routing and overlay routing. IP routing provided by an Internet Service Provider (ISP) is generally configured based on the monetary cost architecture of the individual ISP, which is based on transit and peering relationships with neighboring ISPs. On the other hand, overlay routing attempts to increase user-perceived performance. This policy mismatch also generates a problem for ISPs, whereby the traffic route determined by overlay routing would adversely affect the monetary cost structure of the ISPs. This problem is referred to as the *free-ride traffic problem* [3].

In the present paper, we propose a novel overlay routing mechanism that decreases the extent of the free-ride problem and prevents the monetary cost to ISPs from increasing, while maintaining the performance improvement provided by overlay routing. We first introduce a metric for evaluating the extent of the free-ride traffic problem by overlay routing, and select an overlay-level route that can keep the value of the metric small.

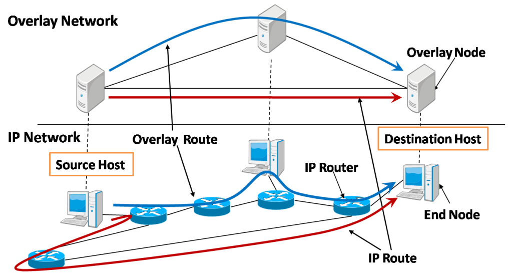

Fig. 1. Overlay routing

To calculate the proposed metric, we need to know the number of transit links traversed by the selected route. However, in general, the contract information between ASes (ISPs), which is in the from of transit or peering relationships, is not available directly. Therefore, we propose an estimation method of the number of transit links on the route from other network metrics, such as end-to-end latency and router-level hop count, which can be measured easily by overlay nodes. We use multiple regression analysis to evaluate the correlation between the number of transit links and such metrics and construct the regression equation with which to estimate the number of transit links on the route. Based on the equation we propose an overlay routing mechanism. We evaluate the effectiveness of the proposed mechanism by numerical evaluations, where the overlay network is constructed in the PlanetLab environment.

II. Research background

A. Overlay routing and IP routing

Overlay routing is a technique to improve end-to-end network performance on overlay networks, using application-level routing based on user-perceived metrics, such as end-to-end latency, available bandwidth, and TCP throughput, as depicted in Figure 1. On the other hand, IP routing operated by ISPs is generally configured with metrics such as router-level and AS-level hop count, which do not always correlate with user-perceived performance.

In addition, ISPs have their own monetary cost architecture based on commercial contracts with neighboring ISPs, and the routing configurations are largely affected by the cost structure. There are two types of links between ISPs[1]. One is a transit link that connects the upper-level ISP and the lower-level ISP, and another is a peering link used for peering relationships. The monetary cost of a transit link is usually determined by the amount of traffic traversing the link. On the other hand, there is almost no monetary charge for the peering links, except for the cost paid to carrier companies for the physical links facilities, but the peering link is allowed to be

[1]We ignore sibling links because they connect ASes, which belong to the same organization.

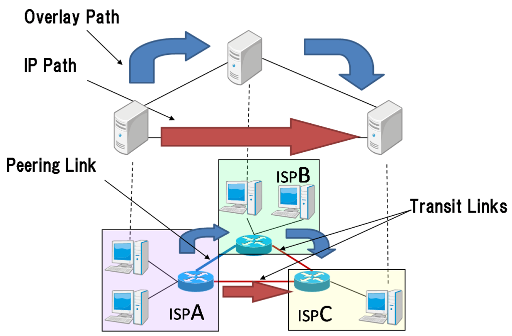

Fig. 2. Free-ride traffic problem

used only by the traffic between the interconnected ISPs. ISPs make routing decisions while considering these differences between transit and peering links.

The advantage of overlay routing is a result of the policy mismatch between IP routing and overlay routing. Figure 1 shows a typical example. When we compare the IP route and the overlay route from the source host to the destination host in this figure, the IP route has smaller router-level hops but longer latency. Therefore, the overlay routing provides better user-perceived performance (i.e., end-to-end latency).

B. Influence on the cost architecture of ISPs

Although overlay routing can improve user-perceived performance, as described above, overlay routing may also generate traffic that does not follow the monetary cost architecture of the ISPs, and the ISPs may incur additional monetary cost for such traffic.

Figure 2 shows a simple example. There are three ISPs, each of which has a router and some endhosts. ISP C is an upper-level ISP for ISPs A and B, and ISPs A and B have transit links to ISP C. In addition, ISPs A and B have a peering relationship and a peering link exists between the two ISPs.

We assume that the endhost in ISP A generates traffic to the endhost in ISP C. When using IP routing or overlay routing with a direct route (red arrow in the figure), the traffic traverses the transit link between ISP A and ISP C. In this case, under the normal monetary cost architecture, ISP A pays the transit cost to ISP C and collects the cost from its own customer, who generates the traffic.

On the other hand, when overlay routing uses a relay path via an endhost in ISP B (blue arrow in the figure), the traffic generated from ISP A traverses the peering link between ISP A and ISP B and the transit link between ISP B and ISP C. In this case, the transit cost is paid by ISP B to ISP C. However, ISP B cannot collect the cost from the customer of ISP A, who generates the traffic. Thus, by using overlay routing with relay paths, transit links, which have not been used in the corresponding direct path, may be used, and the ISPs on the relay path would pay the transit cost for the overlay-routed traffic, which is NOT generated from their customers. We refer to this problem as the free-ride traffic problem [3].

If the endhost in ISP B receives some explicit benefit from relaying the overlay traffic (e.g., content duplicating and caching), ISP B may be able to collect the cost from its own customers. However, in most cases of overlay routing, the relaying hosts are not aware of the relayed traffic. Another possible way to recoup the cost is to monitor the traffic from ISP A to ISP B on the peering link and differentiate this traffic as normal traffic or overlay-routed traffic. Then, ISP B can ask ISP A to pay the cost for the overlay-routed traffic. However, since overlay routing is operated by upper-layer protocols and applications, we cannot separate the overlay-routed traffic by simply checking the source or destination IP addresses of incoming packets.

In the present paper, in an attempt to resolve the above-described problem, we evaluate the overlay routing mechanism focusing on the free-ride traffic problem and propose a novel mechanism that limits the problem while maintaining user-perceived performance.

III. METHODOLOGY

In this section, we first explain the dataset used for the evaluation of overlay routing and multiple regression analysis. In addition, we explain the metrics used for the overlay routing mechanism to choose overlay-level routes. Finally, we introduce a new metric for evaluating the extent of the free-ride traffic problem and introduce the overlay routing mechanism considering the metric.

A. Dataset

In the present paper, we assume that an overlay network is constructed of PlanetLab [10] nodes. To evaluate the overlay routing on the overlay network of PlanetLab nodes, we obtain the full-mesh measurement data of the end-to-end latency and the available bandwidth between each node from Scalable Sensing Service (S^3) [11]. In S^3, the measurement results for network paths between PlanetLab nodes, which are summarized every four hours, are available. In the present paper, we mainly used the dataset obtained on 11/12/2008, which has 476 nodes and 213,396 paths between nodes. In addition, we use the dataset on 10/25/2006, 09/02/2007, and 04/08/2009 to evaluate year-on-year changes in the effect of overlay routing in Subsection V-A.

To obtain the number of transit links on each path, we use the relationship information between ASes (peering or transit) available from CAIDA [12]. Since CAIDA does not provide the relationship information for all links between ASes, we infer unknown relationships based on the degree of each AS (the number of outgoing links to other ASes) using the method developed in a previous study [9]. We used the relationship information obtained on 08/18/2008 and 01/23/2009.

We also need to know the AS-level route of paths between PlanetLab nodes. We first obtained IP router-level routes using the `traceroute` command, and then convert these routes into AS-level routes. For this purpose, we used the database of the correspondence between IP address prefixes and AS numbers, which is available from the Route Views Project [13]. However, there are a number of IP addresses for routers that cannot obtain corresponding AS numbers by this method. Such routers do not contribute to the free-ride traffic problem because the database of the Route Views Project is based on BGP messages. Therefore, we assume that such links are peering links. In the present paper, we used `traceroute` results obtained on 11/12/2008.

B. Routing metrics for overlay routing

In this subsection, we explain the metrics for overlay routing to choose overlay-level paths. In the present paper, there are two candidate overlay paths, as described below.

- **direct path** - A direct path from the source node to the destination node, i.e., a one-hop path with overlay routing.
- **relay path** - A path from the source node to the destination node via another node. In the present paper, we consider only the two-hop path, because paths with more than two hops do not contribute to improving user-perceived performance. [14].

1) Latency: The overlay routing based on end-to-end latency is suitable for applications that require quick response rather than longer-term throughput. For the latency of a relay path, we use the sum of the latencies of the paths that make up the relay path. We denote the latency between overlay nodes i and j as δ_{ij}. Then, we determine the latency of the direct path between nodes i and j, which is denoted as D^1_{ij}, and the latency of the relay path via node k, D^2_{ikj}, respectively, as follows:

$$D^1_{ij} = \delta_{ij} \qquad D^2_{ikj} = \delta_{ik} + \delta_{kj} \tag{1}$$

We define the *best relay path* as the path that has the smallest latency among all possible relay paths. The latency of the best relay path can then be described as follows:

$$\hat{D}^2_{ij} \;=\; \min_{k \neq i,j} \left(D^2_{ikj} \right) \tag{2}$$

We also introduce the *improvement ratio*, which is the ratio of the performance of the best relay path to that of the direct path. The improvement ratio for latency is defined as follows:

$$I(D^2_{ij}) \;=\; \frac{D^1_{ij}}{\hat{D}^2_{ij}} \tag{3}$$

2) Available bandwidth: The overlay routing based on the available bandwidth is suitable for applications that generate a large amount of traffic, such as video streaming and file transmission. We denote the available bandwidth between overlay nodes i and j as β_{ij}. Then, we determine the available bandwidth of the direct path between nodes i and j, which is denoted as B^1_{ij}, and the available bandwidth of the relay path via node k, B^2_{ikj}, respectively, as follows:

$$B^1_{ij} = \beta_{ij} \qquad B^2_{ikj} = \min\left(\beta_{ik}, \beta_{kj}\right) \tag{4}$$

The available bandwidth of the best relay path, which has the largest available bandwidth among all possible relay paths, can be described as follows:

$$\hat{B}^2_{ij} \;=\; \max_{k \neq i,j} \left(B^2_{ikj} \right) \tag{5}$$

The improvement ratio for the available bandwidth, which is the ratio of the performance of the best relay path to that of the direct path, is defined as follows:

$$I(B^2_{ij}) \;=\; \frac{\hat{B}^2_{ij}}{B^1_{ij}} \tag{6}$$

C. Metric for the free-ride traffic problem

In Section IV, we propose a novel overlay routing mechanism that decreases the extent of the free-ride traffic problem. For this purpose, we generalize the concept of the free-ride traffic problem described in Section II and set our goal to decreased monetary cost for the entire network, rather than for an individual ISP.

When we consider the free-ride traffic problem in the entire network, we should decrease the number of transit links used in relay paths selected by overlay routing. Therefore, we use the number of transit links on the path as the metric of the free-ride traffic problem and select the relay path by overlay routing based on the metrics explained in Section III-B, with the constraint on the increased number of transit links.

In what follows, τ_{ij} represents the number of transit links of the path between nodes i and j. Then, the number of transit links of the direct path between nodes i and j and that of the relay path via k are given, respectively, as follows:

$$T^1_{ij} = \tau_{ij} \qquad T^2_{ikj} = \tau_{ik} + \tau_{kj} \tag{7}$$

TABLE I

CORRELATION COEFFICIENTS

Router-level hop count	0.420
End-to-end latency	0.300
Available bandwidth	-0.027

When we limit the increase in the number of transit links through the use of relay paths, we add the following constraint to Equation (2) for end-to-end latency and to Equation (5) for available bandwidth:

$$T_{ikj} \leq T_{ij} + \alpha \qquad \alpha = 0, 1, \ldots, n \tag{8}$$

where α is the upper limit of the increase in the number of transit links. This constraint means that the overlay routing mechanism uses relay paths that increase the number of transit links by α or fewer, as compared with the direct path.

IV. ESTIMATION OF THE NUMBER OF TRANSIT LINKS BY MULTIPLE REGRESSION ANALYSIS

As described above, we use the number of transit links on the path between the source and destination overlay nodes as a metric for evaluating the free-ride traffic problem. However, the number of transit links cannot be determined by overlay nodes because the contract information between ISPs is not disclosed in general. Therefore, we propose a method for estimating the number of transit links on a path from other network performance metrics, which are easily obtained by overlay nodes. We use multiple regression analysis to derive an equation for estimating the number of transit links on a path.

To select parameters for multiple regression analysis, we first calculate the correlation coefficients between the true number of transit links on a path and three types of metrics, which can be measured easily. We then derive the regression equation from the selected parameters and evaluate the accuracy of the equation.

A. Correlation between metrics

We first evaluate the correlations between the true number of transit links and network performance metrics, such as end-to-end latency, available bandwidth, and router-level hop count. We use *Pearson's correlation coefficient* C to calculate the correlation coefficient.

$$C \;=\; \frac{\sum (x_{ij} - \bar{x})(y_{ij} - \bar{y})}{\sqrt{\sum (x_{ij} - \bar{x})^2}\sqrt{\sum (y_{ij} - \bar{y})^2}} \tag{9}$$

Table I lists the correlation coefficients and their values between the true number of transit links and each metric. Unlike the correlation between the router-level hop count and the end-to-end latency, the correlation between the number of transit links and the available bandwidth is close to zero. From the viewpoints of calculation complexity and accuracy of regression analysis, we exclude available bandwidth from the multiple regression analysis.

B. Regression equation by multiple regression

For the reasons described above, we use router-level hop count and end-to-end latency for multiple regression analysis to estimate the number of transit links on a path. We deploy the linear least squares method to derive the regression equation.

We omit the detailed process of the regression analysis due to space limitation. Here the router-level hop count and the end-to-end latency (ms) between nodes i and j are denoted as h_{ij} and δ_{ij}, respectively. Then, we obtain the regression equation for estimating the number of transit links on a path, denoted as T^e_{ij}, as follows:

$$T^e_{ij} \;=\; 0.1419 h_{ij} + 0.002482 \delta_{ij} + 1.136 \tag{10}$$

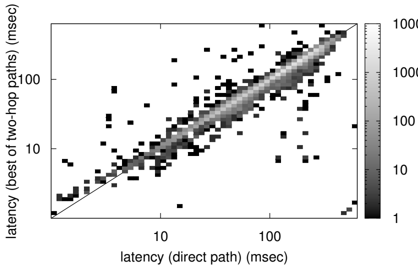

(a) end-to-end latency

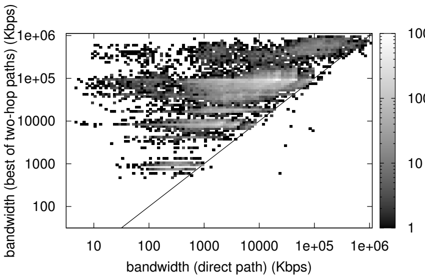

(b) available bandwidth

Fig. 3. Comparison between the direct path and the best relay path (08/11/12)

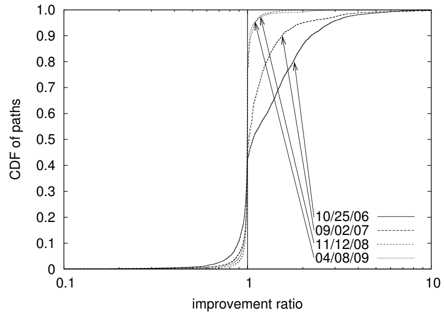

(a) end-to-end latency

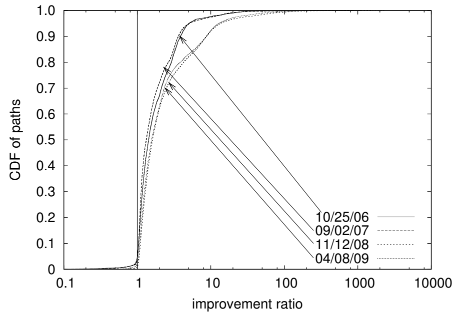

(b) available bandwidth

Fig. 4. Year-on-year comparison of improvement ratio

We assessed the estimation accuracy of the above equation, and confirmed that the maximum estimation error of the regression equation in Equation (10) is smaller than four and that the estimation error is smaller than one for 80% of the paths.

V. Evaluation results

In this section, we evaluate the performance of the overlay routing with the metrics described in Subsection III-B, while considering the decreasing transit cost of the ISPs using the metric proposed in Subsection III-C. We first present the results of the overlay routing without considering the free-ride traffic problem as a baseline for the discussion. We then present the results obtained while limiting the number of transit links on a overlay-routed path and evaluate the effect of the proposed mechanism.

A. Overlay routing without limitation of traversing transit links

Figure 3 shows the distribution of the relationships between the performance of the direct path and that of the best relay path for all node pairs from 11/12/2008 data. The x-axis represents the performance of the direct path, and the y-axis represents the performance of the best relay path corresponding to the direct path. Figure 3(a) plots the results using latency as an overlay routing metric, and Figure 3(b) plots the results using available bandwidth as an overlay routing metric. These figures show that there is little difference between the direct path and the best relay path in the case of latency (Figure 3(a)). This means that latency-based overlay routing may improve user-perceived performance only slightly. In contrast, using available bandwidth as an overlay routing metric reveals a significant improvement in user-perceived performance.

Figure 4 plots the cumulative distribution of the improvement ratio for all paths, as defined in Subsection III-B, when using latency (Figure 4(a)) and available bandwidth (Figure 4(a)). In order to investigate year-on-year changes, the figure includes the data obtained on 10/25/2006, 09/02/2007, 11/12/2008, and 04/08/2009. In the case of latency (Figure 4(a)), the overlay routing shows the best performance on 10/25/2006, and the performance decreases on 09/02/2007 and 11/12/2008. On 04/08/2009 the performance similar to the performance on 11/12/2008. The ratios of paths that have at least one relay path that is better than the direct path are 67%, 63%, 22%, and 22%, respectively, for these data sets. We are currently investigating the reason for these findings. One possible reason is the decrease in the degree of "distorted" routing configurations due to commercial inter-ISP relationships. On the other hand, in the case of available bandwidth, although the performances of the overlay routing on 11/12/2008 and 04/08/2009 are better than the performances on 10/25/2006 and 09/02/2007, there are significant improvements for all years. The ratio of paths that have at least one relay path that is better than the direct path is larger than 95% for all years.

B. Overlay routing with a limitation on the true number of transit links

Next, we present the results for the case in which we limit the increase in the number of transit links in selecting an overlay-routed path. The detailed algorithm can be found in Subsection III-C. Here, we use the true number of transit links of each path with the data obtained on 11/12/2008.

Figure 5 exhibits the cumulative distribution of the improvement ratio of all paths when limiting the increase in the number of transit links, where α is the upper limit of the increase in the number of the transit links, as described in Subsection III-C. Note that when α is small, we cannot find any relay paths that satisfy the limitation for some node pairs. Figure 5 indicates that, for the cases of using latency and available bandwidth as routing metrics, as α increases, the performance approaches that for the case without limitation, and when α is greater than or equal to three, the performances become approximately equal. From this result, we conclude that the overlay routing with a limitation on the number of transit links can provide performance similar to the case without the limitation, when the limitation on the increase in the number of transit links is

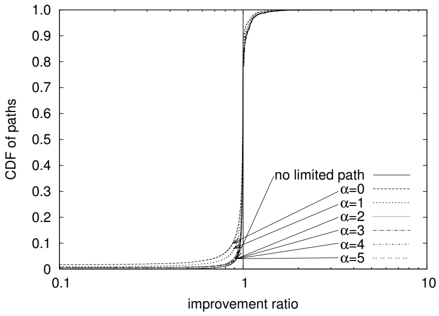

(a) end-to-end latency

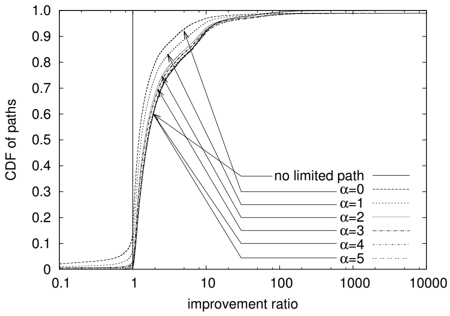

(b) available bandwidth

Fig. 5. Improvement ratio distribution with a limitation on the true number of transit links

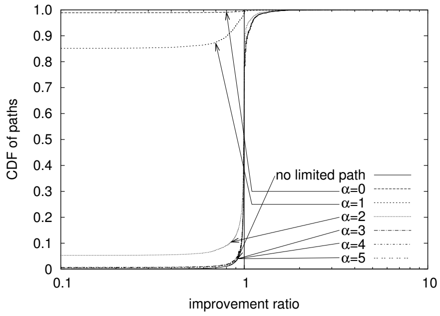

(a) end-to-end latency

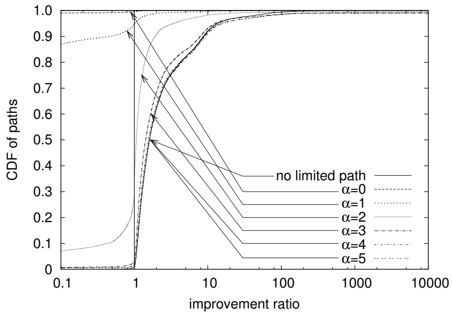

(b) available bandwidth

Fig. 6. Improvement ratio distribution with a limitation on the estimated number of transit links

greater than two.

C. Overlay routing with a limitation on the estimated number of transit links

Finally, we present the results for the case in which we use the estimated value of the number of transit links obtained by the method described in Section IV.

Figure 6 plots the results of the proposed method in the same manner as Figure 5. When α is smaller than three, a significant portion of the paths cannot find any relay paths that satisfy the limitation, and this portion increases significantly compared to that shown in Figure 5. This occurs as a result of the estimation error described in Subsection IV-B. On the other hand, when α is greater than or equal to three, the overlay routing performance is approximately the same as in the case with true limitation (Figure 5) and the case without limitation (Figure 4). This result reveals the advantage of the proposed mechanism, whereby we can control the number of transit links in overlay routing using measurable network performance metrics, while preserving the performance improvement by overlay routing.

VI. Conclusion

In the present paper, we proposed an overlay routing mechanism that decreases the transit cost of ISPs in selecting relay paths. The proposed mechanism estimates the number of transit links on a path using multiple regression analysis of metrics that can be obtained easily by overlay nodes. Based on this estimation, the proposed mechanism limits the increase in the number of transit links on a relay path selected by overlay routing. We evaluated the proposed method assuming the PlanetLab environment and found that when the limitation on the increase in the number of transit links was greater than or equal to three, the proposed mechanism achieved a performance equivalent to that achieved using the true number of transit links or that achieved without a limitation.

In the future, we intend to improve the accuracy of the regression equation considering outlier values. We also intend to consider different mechanisms by which to decrease ISP cost by explicit cooperation between ISPs and overlay network applications, such as P4P [15].

References

[1] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris, "Resilient overlay networks," in *Proceedings of 18th ACM Symposium on Operating Systems Principles*, Oct. 2001.

[2] C. L. T. Man, G. Hasegawa, and M. Murata, "Monitoring overlay path bandwidth using an inline measurement technique," *IARIA International Journal on Advances in Systems and Measurements*, vol. 1, no. 1, pp. 50–60, 2008.

[3] G. Hasegawa, M. Kobayashi, M. Murata, and T. Murase, "Free-riding traffic problem in routing overlay network," in *Proceedings of ICON 2007*, Nov. 2007.

[4] M. Uchida, S. Kamei, and R. Kawahara, "Performance evaluation of QoS-aware routing in overlay network," in *Proceedings of ICOIN 2006*, Jan. 2006.

[5] S. Kamei, "Applicability of overlay routing in Japan using inter-domain measurement data," *Overlay Network Workshop*, Dec. 2006.

[6] Y. Zhu, C. Dovrolis, and M. Ammar, "Dynamic overlay routing based on available bandwidth estimation: A simulation study," *Computer Networks Journal*, vol. 50, pp. 739–876, Apr. 2006.

[7] D. G. Andersen, A. C. Snoeren, and H. Balakrishnan, "Best-path vs. multi-path overlay routing," in *Proceedings of ACM SIGCOMM conference on Internet measurement*, Oct. 2003.

[8] S. Banerjee, C. Kommareddy, K. Kar, B. Bhattacharjee, and S. Khuller, "Construction of an efficient overlay multicast infrastructure for real-time applications," in *Proceedings of INFOCOM 2003*, Apr. 2003.

[9] G. Hasegawa, Y. Hiraoka, and M. Murata, "Evaluation of free-riding traffic problem in overlay routing and its mitigation method," in *Proceedings of 5th International Conference on Networking and Services (ICNS 2009)*, Apr. 2009.

[10] PlanetLab Web Page, available at http://www.planet-lab.org/.

[11] Hewlett-Packard Laboratories, "Scalable Sensing Service," available at http://networking.hpl.hp.com/s-cube/.

[12] University of California, "CAIDA," available at http://www.caida.org/home/.

[13] University of Oregon, "Route Views Project," available at http://www.routeviews.org/.

[14] G. Hasegawa, Y. Hiraoka, and M. Murata, "Effectiveness of overlay routing based on delay and bandwidth information," *IEICE Transactions on Communications*, vol. E92-B, no. 4, pp. 1222–1232, Apr. 2009.

[15] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. G. Liu, and A. Silberschatz, "P4P: Provider portal for applications," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 4, pp. 351–362, Oct. 2008.