Performance evaluation of a method to reduce inter-ISP transit cost caused by overlay routing

Kazuhito MATSUDA*, Go HASEGAWA*, Satoshi KAMEI† and Masayuki MURATA*
*Graduate School of Information Science and Technology, Osaka University
1-5 Yamadaoka Suita, Osaka 560-0871, Japan
Email: {k-matuda, hasegawa, murata}@ist.osaka-u.ac.jp
†NTT Service Integration Laboratories 3-9-11 Midori-cho Musashino, Tokyo 180-8585, Japan

***Abstract*—Overlay routing is an application-level routing mechanism and existing research has revealed that overlay routing can improve user-perceived performance. On the other hand, overlay routing may harm the ISPs' cost structure because of the policy mismatch between IP routing and overlay routing. In a previous study, we proposed a method to reduce inter-ISP transit cost caused by overlay routing while maintaining the effectiveness of the overlay routing. However, we evaluated the proposed method only in the PlanetLab environment, the node location of which is biased to North America. In addition, the previous evaluation focused only on limiting the degree of increase in the inter-ISP transit cost and did not explicitly consider the performance of the overlay routing. In the present paper, we evaluate the proposed method in more general network environments, which are the generalized PlanetLab environment and the Japanese commercial ISPs network environment. We also evaluate the performance of the proposed method in terms of the trade-off relationships between the overlay routing performance and the inter-ISP transit cost over the entire network. In addition, we discuss the differences in the network properties in both environments, which affect the performance of the proposed method.**

I. Introduction

Overlay routing is an application-level routing mechanism on overlay networks that provides application-level routes for network application traffic, as depicted in Fig. 1. In the present paper, the terms "overlay routing" and "IP routing" are used to refer to traffic routing at the application level and the IP level, respectively. One advantage of overlay routing is that user-perceived performance can be improved compared with IP routing [1]–[4]. Such performance improvement is caused primarily by the policy mismatch between IP routing and overlay routing.

However, this policy mismatch between overlay routing and IP routing also generates a problem for the ISPs' cost structure and increases the inter-ISP transit cost over the entire network [5], [6].

In a previous study [7], we proposed a method by which to reduce the inter-ISP transit cost caused by overlay routing. The proposed method uses the number of transit links on a overlay-routed path as a metric of inter-ISP transit cost and estimates the metrics from network performance values using multiple regression analysis. We limit the increase in the metric caused by overlay routing in order to reduce the inter-ISP transit cost. However, the proposed method [7] was evaluated only in the

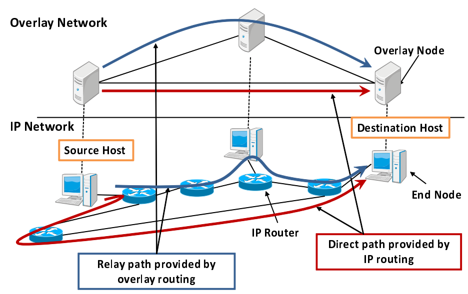

Fig. 1. Overlay routing

PlanetLab environment, the node location of which is biased to North America, which does not match the general host distribution in the Internet. In addition, the evaluation focused only on limiting the degree of increase in the inter-ISP transit cost and did not explicitly consider the relationship between the improvement in performance of the overlay routing and the reduction in inter-ISP transit cost.

In the present study, we evaluate the proposed method in various environments. For this purpose, we assume two network environments: 1) a PlanetLab environment, the node location of which is generalized to conform to the general Internet host distribution, and, in order to confirm the performance of the proposed method in another type environment, 2) a Japanese commercial network environment that is constructed of overlay nodes located at several commercial ISPs in Japan. In addition, based on a mathematical analysis, we discuss the differences of the network properties in the generalized PlanetLab environment and the Japanese commercial network environment that affect the performance of the proposed method. We also evaluate the performance of the proposed method in terms of the trade-off relationships between the overlay routing performance and the reduction in the inter-ISP transit cost over the entire network.

The remainder of the present paper is organized as follows. In Section II, research background on the overlay routing and the associated problem of increased inter-ISP transit cost are described. In Section III, we describe the dataset used in the present paper. In Section IV, we briefly introduce the previously proposed method [7] to reduce the inter-ISP transit cost. In Section V, we present the results of the evaluation of the proposed method in various environments and discuss the

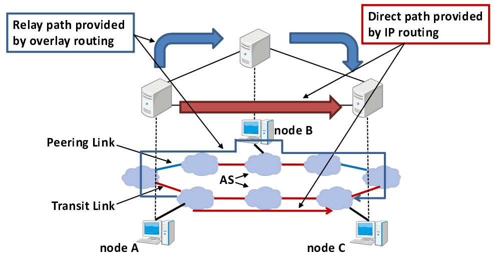

Fig. 2. Increase in the number of transit links by overlay routing

differences of the network properties in these environments. Finally, in Section VI, we summarize the conclusions and discuss future research.

II. Overlay routing and its impact on the ISPs' cost structure

Overlay routing is a technique for network application to improve end-to-end network performance by choosing the paths based on application-level network performance metrics, such as end-to-end latency, available bandwidth, and TCP throughput. On the other hand, IP routing is based primarily on metrics such as router-level and AS-level hop count, which do not always correlate with user-perceived performance. In addition, ISPs that perform IP routing have their own cost structure 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]: transit links, which connect the upper-level ISP and the lower-level ISP, and peering links, which are used for peering relationships. The monetary cost of the transit link is usually determined by the amount of traffic traversing the link, and transit links can be used by customers belonging to the interconnected ISPs. On the other hand, there is almost no monetary charge for peering links, except for the cost paid to carrier companies for the physical link facilities. Therefore, peering links are allowed to be used only by the traffic for which the origin and destination are interconnected ISPs. Internet service providers make routing decisions by considering such differences between transit and peering links.

The advantage of overlay routing is mainly a result of the policy mismatch between IP routing and overlay routing. Figure 1 shows a typical example of this advantage. We assume that IP routing uses the direct path and that overlay routing chooses the relay path. The length of the arrow in the IP network represents the value of the end-to-end latency. When we compare the IP routing and the overlay routing from the source host to the destination host in this figure, the direct path has smaller router-level hop counts, but longer end-to-end latency, as compared to the relay path. Therefore, the overlay routing provides better user-perceived performance (i.e., end-to-end latency) than the IP routing.

Although overlay routing can improve user-perceived performance, overlay routing may also generate traffic that does not follow the cost structure of the ISPs (i.e., the policy of IP routing provided by the ISPs), and the ISPs may incur additional monetary costs caused by such traffic. It can increase the inter-ISP transit cost over the entire network.

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

TABLE I
Number of ASes in each RIR and number of nodes for evaluation

RIR (region name)	Number of ASes	Number of nodes
ARIN (North America)	24,422	50
RIPE NCC (Europe)	21,065	43
APNIC (Asia)	5,782	12
LACNIC (South America)	2,815	6

Figure 2 shows a simple example of this problem. There are three endhosts, all of which work as overlay nodes. There are three overlay links: between Node A and Node B, Node B and Node C, Node C and Node A. Each overlay link includes underlay (i.e., IP) links, each of which is either a transit link or a peering link. In Fig. 2, red links between ASes represent transit links, and blue links between ASes represent peering links. It is assumed that Node A generates traffic to Node C. When using the IP routing or the overlay routing that chooses the direct path, the traffic traverses two transit links. On the other hand, when the overlay routing uses the relay path via Node B, the traffic traverses the transit links between Node A and Node B, and the transit links between Node B and Node C. Therefore, the sum of the transit links traversed by the relay path increases by two compared with the direct path. As a consequence, the inter-ISP transit cost over the entire network increases.

In an attempt to resolve the above-mentioned problem, in a previous study, we evaluated the overlay routing mechanism while focusing on this problem and proposed a novel method by which to alleviate the problem while maintaining user perceived performance [7]. In the present paper, we demonstrate the effectiveness of the proposed method in various environments and discuss the network properties that affect the performance of the proposed method.

III. Dataset

In the present paper, we assume two network environments for performance evaluation of the proposed method. The first is an overlay network constructed of PlanetLab [8] nodes, and the second is an overlay network constructed of nodes of Japanese commercial ISPs. In order to evaluate the proposed method in each environment, we need to know the following properties with respect to the end-to-end path between overlay nodes: end-to-end latency, available bandwidth, router-level path and hop count, and AS-level path and hop count. In addition, the information of the transit/peering relationships between ISPs is necessary in order to evaluate the transit cost of the overlay routing. In the remainder of this section, we describe these environments and explain how to obtain these values for each environment.

A. Generalized PlanetLab environment

The performance of the overlay routing in the PlanetLab environment with all nodes may be biased by the distribution of the locations of the PlanetLab nodes because most of the PlanetLab nodes are located in North America. In order to confirm the performance of the proposed method in the general host distribution on the Internet, we investigated the performance of the proposed method in the *generalized* PlanetLab environment, in which the overlay node distribution is generalized to conform to the Internet host distribution. For this purpose, we referred to the number of ASes in Regional Internet Registries (RIRs) in the current Internet. TABLE I summarizes the distribution of the number of ASes obtained from [9] and the number of PlanetLab nodes used in each region, which is determined in proportion to the number of ASes. We randomly selected the PlanetLab nodes from each region and evaluated the performance of the overlay routing performed by the selected nodes.

We obtained a dataset between all node pairs in PlanetLab using the methods described below. We use the data obtained for 64,077 end-to-end paths between nodes.

End-to-end latencies and available bandwidth: We obtained the end-to-end latency and available bandwidth between PlanetLab nodes from a Scalable Sensing Service (S^3) [10]. In S^3, the measurement results for all network paths between PlanetLab nodes, which are summarized every four hours, are available. In order to avoid the effect of day-to-day fluctuation of the measurement results, we used the median of the datasets obtained over two weeks from November 12, 2008 to November 25, 2008.

IP-level paths and router-level hop counts: We conducted `traceroute` commands between all node pairs in PlanetLab. In the present paper, the `traceroute` results obtained on November 12, 2008 are used.

AS-level paths and AS-level hop counts: We converted the IP-level path to the AS-level path using the relationships between the IP address prefix and the AS number, which was available at the Route Views Project [11].

Transit/peering information: In order to obtain the number of transit links on each path, we used the transit/peering relationship information between ASes that was available from CAIDA [12]. However, CAIDA does not provide the relationship information for all links between ASes and there are several IP addresses for routers for which the corresponding AS numbers cannot be obtained by the method described above. Therefore, we stochastically inferred the relationship information using the ratio at which the relationships were peering for each pair of degrees of ASes (the number of outgoing links to other ASes) in the CAIDA dataset. In the present paper, the number of transit links obtained using the above method is assumed to be the "true" number of transit links, which is used in the proposed method in Section IV and in the performance evaluation in Section V.

B. Japanese commercial network environment

The dataset on the Japanese commercial network environment was obtained from a colleague. The environment is composed of 18 nodes of 13 Japanese commercial ISPs. We used 289 data of the end-to-end path between nodes. This dataset includes the full-mesh `traceroute` results and end-to-end latency measured using `ping` commands, so that end-to-end latency, IP-level paths, and router-level hop counts could be obtained from the dataset. Since this dataset does not include data on the available bandwidth, the evaluation on the available bandwidth-based overlay routing is excluded in Section V. The dataset obtained on March 22, 2009 is used in the present paper.

The AS-level paths and hop counts and transit/peering information were obtained in a manner identical to the information in the generalized PlanetLab environment.

IV. Proposed method

The proposed method consists of two components. We first explain the definition of the metric of the inter-ISP transit cost and the path selection methods for overlay routing. We then introduce a method of estimating the proposed metric, because the metric cannot be obtained directly.

A. Limited overlay routing

Since the inter-ISP transit cost is generated by the traffic that traverses inter-ISP transit links, we can reduce the inter-ISP transit cost in overlay routing by decreasing the number of transit links on a path used by overlay networks. Therefore, we consider the number of transit links on a path chosen by the overlay routing as a metric of inter-ISP transit cost.

In the following, c_{ij} represents the number of transit links of the path between nodes i and j. The number of transit links of the direct path between nodes i and j and that of the relay path via node k are then given, respectively, as follows.

$$C_{ij} = c_{ij} \tag{1}$$

$$C_{ikj} = c_{ik} + c_{kj} \tag{2}$$

In order to reduce the inter-ISP transit cost caused by overlay routing, we propose the following constraint.

$$C_{ikj} \le C_{ij} + \alpha \tag{3}$$

where α is the upper limit of the increase in the number of transit links used by a relay path instead of a direct path. The overlay routing selects the path that has the best performance under the condition in Eq. (3). When the performance of the direct path between nodes i and j is denoted as P_{ij}^{dct}, and that of the path between nodes i and j via node k is denoted as P_{ikj}, the *improvement ratio* of the user-perceived performance, which is denoted as $\hat{I}_{ij}$, can be described as follows.

$$\hat{I}_{ij} = P_{ij}^{dct} / \min_{k \neq i,j} (P_{ikj}) \tag{4}$$

$$\hat{I}_{ij} = \max_{k \neq i,j} (P_{ikj}) / P_{ij}^{dct} \tag{5}$$

Note that Eq. (4) is for performance metrics, such as end-to-end latency, for which smaller values indicate better performance, and Eq. (5) is for performance metrics, such as available bandwidth, for which larger values indicate better performance. This means that the overlay routing mechanism provides performance improvement $\hat{I}_{ij}$, under the limitation on the degree of increase in the number of transit links.

As another method by which to reduce the transit cost, we use the following path selection method to explicitly evaluate the performance of the proposed method from the viewpoint of the trade-off relationships between the overlay routing performance and the degree of the reduction in the inter-ISP transit cost. When the best performance provided by the overlay routing between nodes i and j without considering the inter-ISP transit cost is denoted as P_{ij}^{opt}, we select the overlay-routed path between nodes i and j from the possible paths that satisfy the following conditions.

$$P_{ikj} \leq P_{ij}^{opt} \times (1+\beta) \tag{6}$$

$$P_{ikj} \geq P_{ij}^{opt} \times (1-\beta) \tag{7}$$

where β determines the lower limit of the overlay routing performance as compared to the best one. Note that Eq. (6) is for performance metrics for which smaller values indicate better performance, and Eq. (7) is for performance metrics for which larger values indicate better performance. The overlay routing selects the path that has the smallest number of transit links while satisfying Eqs. (6) and (7). When the number of transit links on the path that has the best performance is denoted as C_{ij}^{opt}, the *reduction in the number of transit links*, which is denoted as $\hat{C}_{ij}$, can be described as follows.

$$\hat{C}_{ij} = C_{ij}^{opt} - \min_{k \neq i,j} (C_{ikj}) \tag{8}$$

In other words, this path selection method can reduce the number of transit links $\hat{C}_{ij}$ when we allow a decrease in the user-perceived performance as compared with the user-perceived performance of the best path.

B. Estimation of the number of transit links on a network path

Although the limited overlay routing uses the number of transit links on a path as the metric of the inter-ISP transit cost, it cannot be explicitly determined by overlay nodes because the contract information between ISPs is generally not disclosed. Furthermore, there is no effective method of measuring the number of transit links in an end-to-end manner. Therefore, we propose a method for estimating the number of transit links on a path using other network performance values, which can be measured easily by overlay nodes.

For this estimation method, we first investigate the correlations between the true number of transit links on the path between overlay nodes and network performance values that can be obtained easily by end-to-end measurement, such as router-level hop count, end-to-end latency, and available bandwidth. For the estimation, we then select a number parameters having strong correlations with respect to the number of transit links. We apply a multiple regression analysis to the selected parameters and derive the regression equation from the results to estimate the number of transit links.

When we denote x_{ij}^{l} as the value of the l-th parameter between nodes i and j, the regression equation used to estimate the number of transit links on the path between nodes i and j is as follows.

$$C_{ij}^{e} = b_0 + b_1 x_{ij}^{1} + b_2 x_{ij}^{2} + \ldots + b_n x_{ij}^{n} \tag{9}$$

TABLE II
PARTIAL COEFFICIENT VALUES OF THE REGRESSION EQUATION

	b_r	b_d	b_y
Generalized PlanetLab	0.145 (7.56×10^{-4})	0.00120 (1.08×10^{-6})	0.846 (0.20)
Japanese commercial network	0.240	-0.000889	-1.48

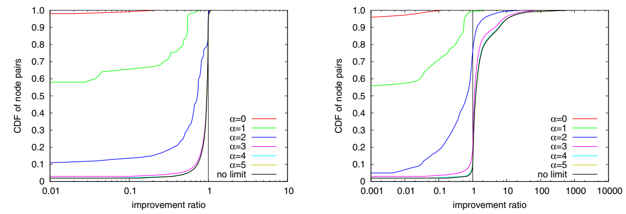

(a) End-to-end latency (b) Available bandwidth

Fig. 3. Distribution of the improvement ratio of limited overlay routing (generalized PlanetLab environment)

where b_0 is the intercept of the equation, b_l is the partial coefficient of the l-th parameter calculated by the multiple regression analysis, and n is the number of parameters.

V. NUMERICAL EVALUATION

A. Evaluation of limited overlay routing

For the evaluation, we first derived the regression equation (Eq. (9)) for the generalized PlanetLab environment and the Japanese commercial network environment. We calculated the correlation between the number of transit links and the following three metrics: router-level hop count, end-to-end latency, and available bandwidth. Based on the calculation results, we excluded the available bandwidth from the regression equation because the correlation was quite weak compared with the other two metrics. For the Japanese commercial environment, since there was no data on the available bandwidth, as mentioned above, the same parameters (i.e., router-level hop count and end-to-end latency) were selected. When the router-level hop count and the end-to-end latency between nodes i and j are denoted as h_{ij} and δ_{ij}, respectively, Eq. (9) can be rewritten as follows.

$$C_{ij}^{e} = b_0 + b_r h_{ij} + b_d \delta_{ij} \tag{10}$$

where b_0 is the intercept of the equation, and b_r and b_d are the partial coefficients of the router-level hop count and the end-to-end latency, respectively. The partial coefficients (b_r, b_d, and b_0), which are the results of the multiple regression analysis, for both environments are listed in TABLE II. For the generalized PlanetLab environment, we conducted twenty node selections and calculated the partial coefficients. The average values and the variance values (in parentheses) are listed in TABLE II. Since the variances are significantly smaller than the average values, it is thought that the node selection in the generalized PlanetLab environment does not affect the performance of the proposed method.

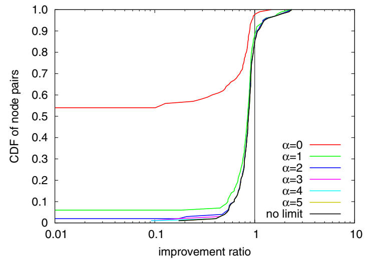

Fig. 4. Distribution of the improvement ratio of limited overlay routing (Japanese commercial network environment)

1) Distribution of improvement ratio: Figure 3 plots the results of the distribution of the improvement ratio of the limited overlay routing, where the overlay path is selected with the limitation in Eq. (3), in the generalized PlanetLab environment. Figures 3(a) and 3(b) show the results obtained using the end-to-end latency and the available bandwidth, respectively, as the routing metric. This figure indicates that when α is smaller than three, there are several node pairs for which the improvement ratio is zero, because such node pairs cannot find any relay paths that satisfy the limitation (Eq. (3)). 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 without the limitation. These results are almost identical to those reported in [7], where we used all of the nodes in the PlanetLab environment.

Figure 4 shows the results in the same manner as Fig. 3 for the Japanese commercial network environment. The tendency of the results is similar to that in the generalized PlanetLab environment. That is, when α is greater than or equal to one, the overlay routing performance is approximately the same as the case without the limitation. From the above results, the proposed method can achieve roughly the same overlay routing performance compared with that of the no limited overlay routing. We also concluded that the proposed method is appropriate to various network environments.

2) Trade-off evaluation between overlay routing performance and the reduction in the number of transit links: We used two methods to evaluate the performance of the proposed method from the viewpoint of the trade-off relationships between the overlay routing performance and the degree of reduction in the inter-ISP transit cost. We first investigated the relationships between the actual number of transit links on the selected path and the improvement ratio of the user-perceived performance under the limitation in Eq. (3). Figure 5 is the distribution of the relationships between the performance improvement ratio defined in Eqs. (4) and (5) and the number of transit links on the overlay routed path in the generalized PlanetLab environment with different values of α in Eq. (3) when the available bandwidth is used as the routing metric. The lines represent the density contours of 50, 100, 150, and 200 nodes pairs. The marked point in the figure represents the mode value for each condition. Figure 5 indicates that the reduction in the number of transit links is greatest when α is equal to two, but the improvement ratio becomes less than

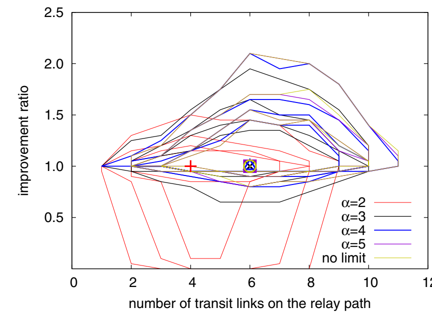

Fig. 5. Overlay routing performance vs. the number of transit links under the limitation on the increase in the number of transit links (generalized PlanetLab environment)

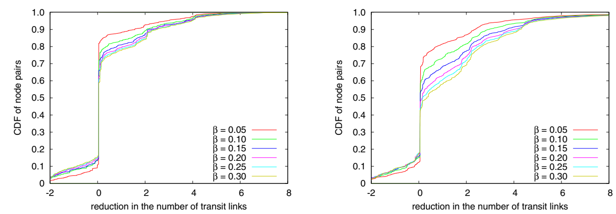

(a) End-to-end latency

(b) Available bandwidth

Fig. 6. Distribution of the reduction in the number of transit links of limited overlay routing with respect to the decrease in the overlay routing performance (generalized PlanetLab environment)

one in a number of node pairs. On the other hand, when α is greater than or equal to four, there is almost no reduction in the number of transit. We conclude that $\alpha = 3$ is a better value from the viewpoint of the trade-off.

Figure 6 shows the results of limited overlay routing with the limitations given in Eqs. (6) and (7). Figures 6(a) and 6(b) are for the end-to-end latency and the available bandwidth, respectively, as the routing metric. The figure plots the distribution of the degree of the reduction in the number of transit links in the generalized PlanetLab environment. Note that the degree of the reduction in the number of transit links is calculated using the true number of transit links. This figure indicates that the proposed method can reduce the number of transit links by at least one in 11% of node pairs when the end-to-end latency is used as the routing metric and by at least one in 18% of node pairs when the available bandwidth is used as the routing metric, allowing only a 5% decrease in the overlay routing performance. Although we can achieve a greater degree of reduction in the number of transit links by allowing a greater decrease in user-perceived performance, the decrease in the user-perceived performance becomes significant. For example, when allowing a 30% decrease in the overlay routing performance, we can reduce the number of transit links by at least one in 22% and 42% node pairs for the end-to-end latency and the available bandwidth, respectively.

B. Network properties affecting the regression equation

TABLE II reveals the difference in the regression equations for the two environments. The router-level hop count is weighed more heavily in the Japanese commercial network environment than in the generalized PlanetLab environment.

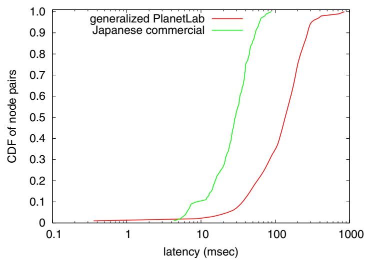

Fig. 7. End-to-end latency distribution

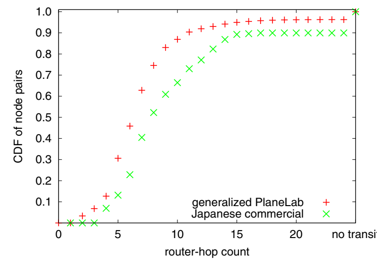

Fig. 8. Router-level hop count of the first traversing transit link

Furthermore, the intercept value in the generalized PlanetLab environment is positive, whereas that in the Japanese commercial network is negative.

The reason for these differences can be found in Figures 7 and 8. Figure 7 shows the distributions of the end-to-end latencies of the direct paths of all node pairs in each environment. Figure 7 indicates that the latencies on the Japanese commercial network are significantly smaller than those in the generalized PlanetLab environment. Therefore, the transit-link estimation by end-to-end latency is difficult in the Japanese commercial network. This is one reason why the router-level hop count is heavily weighted in the Japanese commercial network environment.

Figure 8 represents the cumulative distribution of the router-level hop count from the source node at which the path traverses the first transit link for the direct paths of all node pairs. The paths that have no transit links are counted at the right endpoint of the x-axis. In the generalized PlanetLab environment, 6% of all paths traverse the first transit links within the first three hops and 30% of all paths traverse the first transit links within the first five hops. On the other hand, in the Japanese commercial network environment, these values are 0% and 13%, respectively. This means that the paths in the generalized PlanetLab environment have the first transit links within the first few hops, so that the equation has a positive intercept. On the other hand, since the paths in the Japanese commercial network environment have no transit links in the beginning part of the path, the equation has a negative intercept.

These differences in network properties may be caused by the differences between the generalized PlanetLab environment and the Japanese commercial network environment. PlanetLab is a global research network that is constructed of nodes that are located at universities and enterprises, whereas the Japanese commercial network is constructed of nodes located at Japanese commercial ISPs. The proposed method can obtain the appropriate regression equations for each network environment.

VI. Conclusion

In the present study, we evaluated the proposed method of reducing inter-ISP transit cost using the estimation of the number of transit links on network paths in the generalized PlanetLab environment and the Japanese commercial network environment. The evaluation results revealed the advantages of the proposed method whereby we can control the number of transit links in overlay routing in various network environments, while maintaining the performance improvement provided by the overlay routing. We also confirmed that the limitation whereby the increase in transit links compared with direct path was set to three was better from the viewpoint of trade-off between the overlay routing performance and the degree of the reduction in the inter-ISP transit cost. We also confirmed that we could achieve a certain degree of reduction in the number of transit links allowing only a 5% decrease in the overlay routing performance. In addition, we discussed the network properties that affected the proposed method and confirmed that the proposed method could obtain the regression equation appropriately for both network environments.

In the future, we intend to consider different mechanisms by which to decrease the inter-ISP transit cost by explicit cooperation between ISPs and overlay network applications, such as [13].

References

[1] 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.

[2] D. G. Andersen, A. C. Snoeren, and H. Balakrishnan, "Best-path vs. multi-path overlay routing," in *Proceedings of IMC 2003*, Oct. 2003.

[3] 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.

[4] 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.

[5] P. Rodriguez, S.-M. Tan, and C. Gkantsidis, "On the feasibility of commercial, legal P2P content distribution," *SIGCOMM Computer Communication Review*, vol. 36, no. 1, pp. 75–78, 2006.

[6] S. Seetharaman and M. Ammar, "Characterizing and mitigating inter-domain policy violations in overlay routes," in *Proceedings of ICNP 2006*, Nov. 2006, pp. 259–268.

[7] K. Matsuda, G. Hasegawa, and M. Murata, "Decreasing ISP transit cost in overlay routing based on multiple regression analysis," in *Proceedings of ICOIN 2010*, Jan. 2010.

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

[9] IANA, "IANA AS Numbers assignment data page," available at http://www.iana.org/assignments/as-numbers/.

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

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

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

[13] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. G. Liu, and A. Silberschatz, "P4P: Provider portal for applications," *SIGCOMM Computer Communication Review*, vol. 38, no. 4, pp. 351–362, 2008.