Master's Thesis

Title

A study on free-riding traffic problem in overlay routing

Supervisor

Professor Hirotaka Nakano

Author

Yuichiro Hiraoka

February 13th, 2008

Graduate School of Information Science and Technology

Osaka University

Master's Thesis

A study on free-riding traffic problem in overlay routing

Yuichiro Hiraoka

Abstract

Recent research on overlay networks has revealed that user-perceived network performance could be improved by an overlay routing mechanism. However, most of these studies only consider end-to-end delay and packet loss ratio, and there are few works focusing on bandwidth-related information such as available bandwidth and TCP throughput that are important performance metrics especially for long-lived data transmission. In this thesis, we first investigate the effect of overlay routing using both delay and bandwidth information, using the public measurement results of delay and available bandwidth of network paths between PlanetLab nodes. We consider three metrics for selecting overlay route, end-to-end delay, available bandwidth, and TCP throughput. We then present that the available bandwidth-based overlay routing provided significant gain, compared with delay-based routing.

The effectiveness of overlay routing is caused mainly by the policy mismatch between the overlay routing and the underlay IP routing operated by ISPs. However, this policy mismatch causes "free-riding" traffic problem in routing overlay networks, which may become harmful for ISPs' cost structure. Therefore, we define the free-riding problem in routing overlay network, and evaluate the degree of free-riding traffic to reveal the effect of the problem on ISPs. We introduce the numerical metric to evaluate the degree of the free-riding problem, and confirm that most of relay paths that has better performance than direct path bring the free-riding problem. We also discuss the guideline of selecting paths that are more effective than direct path and that mitigate the free-riding problem.

Keywords

Overlay networks

Overlay routing

IP routing

Free-riding traffic

PlanetLab

Contents

List of Figures

List of Tables

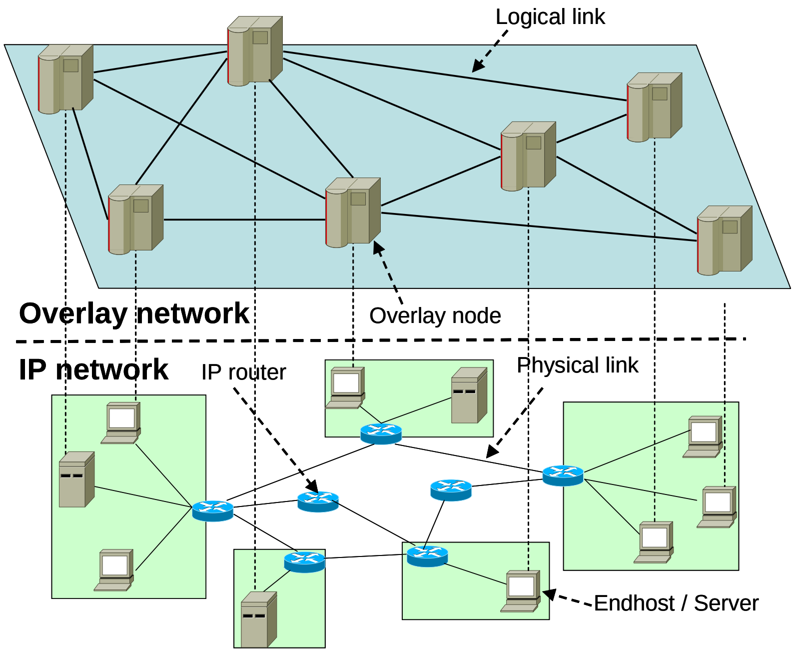

Figure 1: Overlay network

1 Introduction

As the Internet increasingly diversifies and the user population grows rapidly, new and varied types of service-oriented networks are emerging. Called service overlay networks [1] include P2P networks, anonymous file-sharing services, audio and video conferencing services, and Content Delivery/Distribution Networks (CDNs). Service overlay networks are defined as upper-layer networks providing special-purpose services that are built on the lower-layer IP network. Therefore, their performances depend primarily on how well they take advantage of the characteristics and resources of the underlying IP network.

In overlay networks, the endhosts and servers that run the applications become overlay nodes that form the upper-layer logical network with logical links between the nodes, as depicted in Figure 1. Some of the overlay networks select a route for data transmission according to network conditions such as link speed, delay, packet loss ratio, hop count, and TCP throughput between overlay nodes. In WinMX, an endhost can report the type of network link used to connect to the Internet when joining the network. CDNs such as NetLightning [2] and Akamai [3] distribute overlay nodes (content servers) over the entire Internet and select appropriate source and destina-

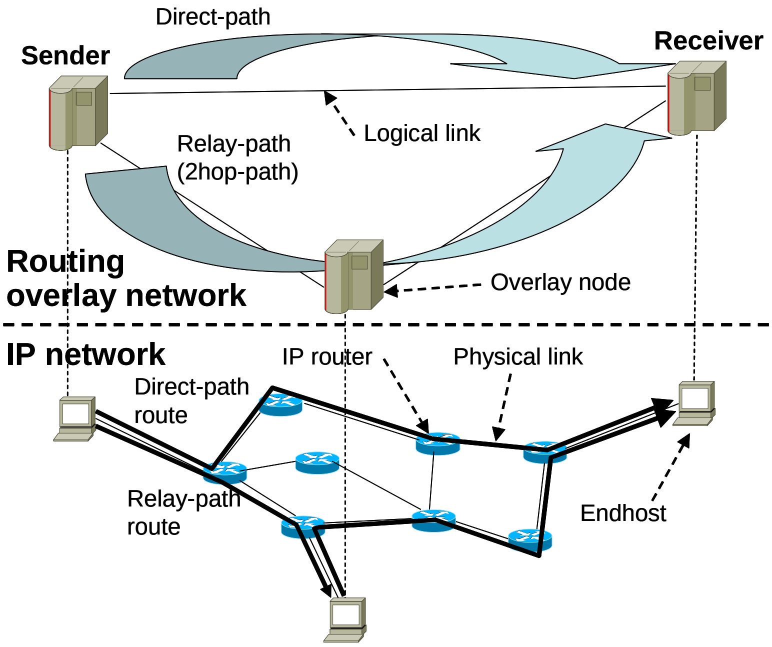

Figure 2: Overlay routing and routing overlay network

tion hosts according to the network condition when the contents are moved, duplicated, or cached.

Some overlay networks do not assume specific upper-layer applications and concentrate only on the routing of overlay network traffic. We call such application-level traffic routing *overlay routing*, and overlay networks for traffic routing are referred to as *routing overlay networks*, as depicted in Figure 2. In Resilient Overlay Networks (RON) [4], for example, 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 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 [4], the authors reported that RON can provide an effective traffic transmission path compared with lower-layer IP routing. Furthermore, RON can detect network failures (link and node failures, and mis-configured routing settings) and provide an alternate route faster than IP routing convergence.

Several studies have examined the effectiveness of overlay routing with respect to IP routing [5–11]. For example, in [8], the authors used actual measurement data of the transmission latency among several geographically-distributed hosts in two Internet Service Providers (ISPs) in Japan, and showed that the transmission latency of approximately 28% of end-to-end paths can be reduced by relaying another host, as compared to using the direct path. In [9], the authors investigated the effectiveness of the reactive overlay routing by using the measurement data on four ISPs in United States, and confirm its effectiveness compared than the IP routing and proactive overlay routing. However, most of these studies focus on end-to-end delay performance, and . bandwidth-related information, such as available bandwidth and TCP throughput, that are important performance metric especially for long-lived data transmission.

In this thesis, we first investigate the effectiveness of overlay routing, based on both of delay and bandwidth information. We assume that PlanetLab [12] nodes construct a routing overlay network, and utilize the measurement results obtained from Scalable Sensing Service (S^3) [13], which measures various properties of network paths between PlanetLab nodes. We utilize the following three metrics in selecting overlay route: we use end-to-end delay, available bandwidth, and TCP throughput. One of the important results of the present study is the investigation of the effectiveness of 3-hop relay overlay path, whereas almost all of the previous studies on overlay routing focused on the 2-hop relay overlay path. Another interesting result in this thesis is the correlation between transmission latency and available bandwidth of the end-to-end path. We revealed whether or not a network path with larger available bandwidth has smaller transmission latency, and vice versa.

The primary reason why overlay routing mechanisms can improve user-related performance metric is that the traditional IP routing operated by ISPs does not always determine the route according to user-perceived performance. In IP routing, the metrics determining the route are hop count and link loads, and the end-to-end delay and bandwidth-related information, which affect the data transmission throughput for short- and long-lived TCP connections, are not directly taken into account. In addition, inter-domain routing by Border Gateway Protocol (BGP) is based on autonomous system-level (AS-level) network topology, which is more abstracted than router-level IP network topology. Furthermore, most ISP-driven IP routings are configured by political and financial factors: the billing mechanism of transit links to upper-layer ISPs, the relationships between the ISP and other ISPs interconnected by public or private peering links, and the amount of traffic traversing transit and peering links. Therefore, the resulting IP routing policy cannot maximize network performance and user demand.

There are some previous works on problems of overlay network. In [14], the authors discuss the interaction between overlay routing and underlay IP routing, that causes routing and traffic oscillation. In [15], the authors discuss the effect of P2P-based content distribution on ISP's costs. We focus on other problem of overlay routing, where overlay routing can harm the profit of the ISPs which operate the lower-layer IP routing, and overlay routing mechanisms can generate network traffic which ignores an ISP's billing structure. In this thesis, we focus on this problem caused by overlay routing, which we call "free-riding" traffic.

Note that free-riding traffic problem in this thesis is a general problem for overlay networks, regardless of the kind of application and algorithms and metrics for overlay routing. Furthermore, it can occur even when the routing interaction between overlay routing and IP routing is stable.

We define the free-riding problem in routing overlay network and evaluate the degree of free-riding traffic problem. For this purpose, we introduce the numerical metric to assess the degree of the free-riding traffic problem. By using the evaluation results, we investigate the degree of free-riding problem of overlay routing mechanisms with three metrics in selecting overlay paths. We also show the results on the ratio of relay paths which has better performance than the direct path and which has smaller number of free-riding transit links than the optimized relay path. We finally present that we can obtain the reasonable performance gain by using relay path with limitation on the number of free-riding transit links.

The rest of this thesis is organized as follows. In Section 2, we explain the methodology and performance metrics for evaluating the effectiveness of overlay routing, and present extensive numerical results and discussions. In Section 3, we define the free-riding traffic problem and present evaluation results to assess the degree of free-riding problem. We also show the guideline in selecting overlay paths to mitigate free-riding traffic problem. Finally, in Section 4, we summarize the conclusions and discuss areas for future consideration.

2 Effectiveness of overlay routing based on delay and bandwidth information

2.1 Methodologies

2.1.1 Dataset used for evaluation

We investigate the effectiveness of overlay routing based on delay and bandwidth information under the assumption that PlanetLab nodes construct a routing overlay network. For that purpose, we utilize the measurement results obtained from S^3. S^3 measures various properties of end-to-end paths between PlanetLab nodes, including physical capacity, available bandwidth, end-to-end delay, and packet loss ratio. The measurement results are provided every four hours via a Web site. In this thesis, we use the data obtained on Oct. 25th, 2006.

There exist 588 PlanetLab nodes in the measurement data utilized herein. However, some nodes are located in the same subnetwork, as estimated from the IP address and host name of the nodes. In evaluating the effectiveness of overlay routing, we should avoid using the nodes in the same subnetwork as relay nodes for the following three reasons: (1) The measurement results of end-to-end delay and available bandwidth between nodes in the same subnetwork may be quite small for delay and quite large for available bandwidth, which may overestimate the effectiveness of overlay routing. (2) The measurement results between nodes in the same subnetwork may include large errors especially for available bandwidth. (3) There is almost no meaning in using a relay node in the same subnetwork as the source and destination nodes.

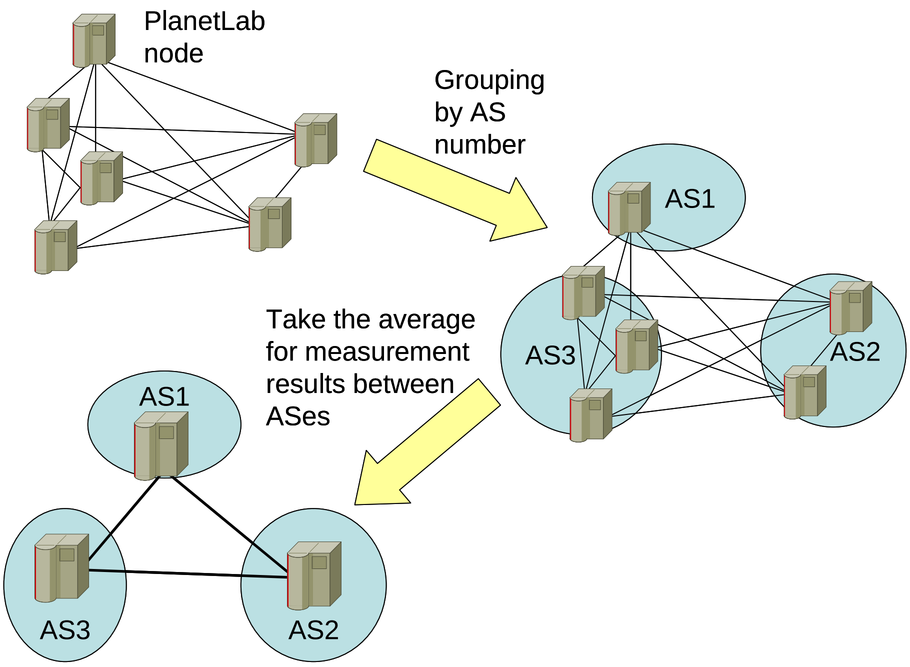

Figure 3: Grouping PlanetLab nodes

Therefore, we divide the PlanetLab nodes into groups according to their AS number and assume that there is only one overlay node in each AS. We obtain the AS number of PlanetLab nodes by tracerouting from a route server in traceroute.org [16] to the PlanetLab nodes. As a result, the number of overlay nodes decreases to 179, which is equal to the number of ASes of PlanetLab nodes. In grouping, we take the average for measurement results when we have more than one measurement result between the overlay nodes (ASes). Figure 3 depicts this process for node grouping.

2.1.2 Overlay path candidates

When one node (source node) selects the transmission path to another node (destination node), we compare the end-to-end latency and available bandwidth of the following three candidates (Figure 4):

- Direct path: the source node to the destination node
- 2-hop relay path: the source node to the destination node via a relay node

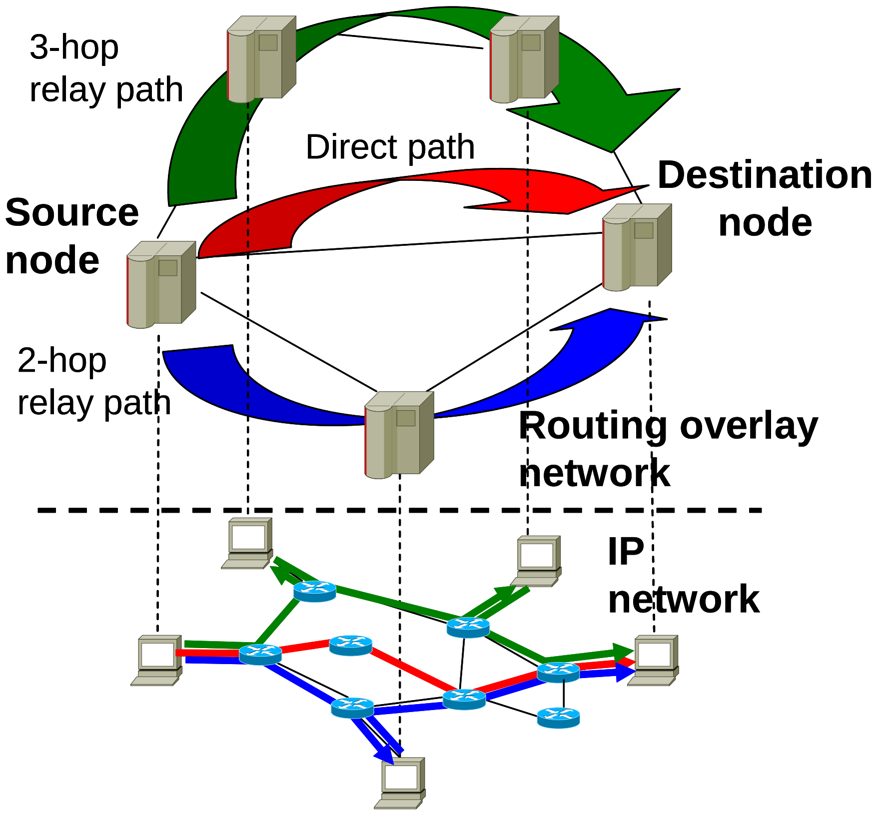

Figure 4: Definition of overlay path

- 3-hop relay path: the source node to the destination node via two relay nodes

2.1.3 Metrics

In this subsection, we explain the metrics utilized for selecting overlay paths.

End-to-end latency Overlay routing based on end-to-end latency would be adapted for applications, including voice chat applications such as Skype [17] that need quick response, rather than bandwidth-related resources. We utilize the measurement results from S^3 for the end-to-end latency of the direct path between nodes. We define the end-to-end latency of a relay path as the sum of the latencies of direct paths constructing the relay path. We assume the number of overlay nodes is M, and the measured results of the end-to-end delay of the network path between nodes N_i and N_j is τ_{ij} ($1 \leq i, j \leq M$). Then, we can describe the latencies of the direct path, 2-hop relay path, and 3-hop relay path, as follows:

$$\begin{aligned} D^1_{ij} &= \tau_{ij} && (1) \\ D^2_{ikj} &= \tau_{ik} + \tau_{kj} && (2) \\ D^3_{iklj} &= \tau_{ik} + \tau_{kl} + \tau_{lj} && (3) \end{aligned}$$

We denote that the relay node for the 2-hop relay path as N_k and the relay nodes for the 3-hop relay path as N_k and N_l ($1 \leq k, l \leq M,\ k \neq l,\ k, l \neq i, j$). Furthermore, we define the latency-optimized path as the relay path that has the smallest end-to-end latency. We can then obtain the respective latencies of the 2-hop and 3-hop latency-optimized paths as follows:

$$\begin{aligned} \hat{D}^2_{ij} &= \min_{k \neq i,j} (D^2_{ikj}) && (4) \\ \hat{D}^3_{ij} &= \min_{k \neq l,\ k,l \neq i,j} (D^3_{iklj}) && (5) \end{aligned}$$

In this thesis, we compare the performance of the direct path and relay path for each node pair. We therefore define the *improvement ratio* of the relay path with respect to the direct path as follows:

$$\begin{aligned} I(D^2_{ikj}) &= \frac{D^1_{ij}}{D^2_{ikj}} \\ I(D^3_{iklj}) &= \frac{D^1_{ij}}{D^3_{iklj}} \end{aligned}$$

When the above ratio is larger than 1, we can say that the relay path is effective compared with the direct path.

Available bandwidth Available bandwidth is important performance metric for audio video streaming services such as YouTube [18] and GyaO [19]. We simply use the measurement results of available bandwidth in S^3 for the available bandwidth of direct paths. We define the available bandwidth of a relay path as the minimum of the available bandwidth of direct paths constructing the relay path. By denoting, available bandwidth of the network path between node N_i and N_j as ρ_{ij}, we can describe the available bandwidths of the direct path, 2-hop relay path, and 3-hop relay path, as follows:

$$\begin{aligned} B^1_{ij} &= \rho_{ij} && (6) \\ B^2_{ikj} &= \min(\rho_{ik}, \rho_{kj}) && (7) \\ B^3_{iklj} &= \min(\rho_{ik}, \rho_{kl}, \rho_{lj}) && (8) \end{aligned}$$

We also define the bandwidth-optimized path as the relay path that has the largest available bandwidth among all possible relay paths. We can then obtain the respective the available bandwidths of the 2-hop and 3-hop bandwidth-optimized paths as follows:

$$\hat{B}^2_{ij} = \max_{k \neq i,j} (B^2_{ikj}) \tag{9}$$

$$\hat{B}^3_{ij} = \max_{k \neq l,\ k,l \neq i,j} (B^3_{iklj}) \tag{10}$$

Furthermore, we define the improvement ratio of the relay path with respect to the direct path as follows:

$$I(B^2_{ikj}) = \frac{B^2_{ikj}}{B^1_{ij}}$$

$$I(B^3_{iklj}) = \frac{B^3_{iklj}}{B^1_{ij}}$$

TCP throughput Overlay routing based on TCP throughput would be adapted for file sharing applications like Bittorrent [20] and WinMX. In [4], RON utilizes TCP throughput as one of performance metrics for overlay routing, which is calculated from end-to-end delay and packet loss ratio as follows:

$$T = \frac{\sqrt{1.5}}{RTT\sqrt{Loss}}\ (\text{packet/sec}) \tag{11}$$

Equation (11) is based on the formula for average throughput of long-lived TCP connection in [21]. This metric increases as packet loss ratio and RTT decrease, but never exceeds the available bandwidth of the path in the actual situation. In this thesis, we utilize Eq. (11) as a performance metric of TCP throughput. This equation includes the packet loss ratio of the path, and S^3 has the measurement results of packet loss ratios of the network path between the PlanetLab nodes. However, we do not use them, since they are obtained by sending only 100 probe packets. Instead, we utilize the following two models for packet loss ratio of the direct path:

(A) AS-hop-count-base loss ratio: the packet loss ratio of the path is determined proportionally to the AS-level hop count of the path.

(B) Overlay-hop-count-base loss ratio: the packet loss ratio of the direct path is constant value regardless of the path's other characteristics. Relay path has the packet loss ratio proportionally to its overlay-level hop count.

Note that from the measurement results of packet loss ratio obtained from S^3, we found that the packet loss ratio is not related to the AS-lelvel hop count of the path. However, we believe that the assumption where the packet loss ratio increases as the AS-level hop count increases is reasonable.

We define P^1_{ij} as TCP throughput of the direct path between node N_i and N_j, and we can describe P^1_{ij} as follows:

$$\begin{aligned} P^1_{ij} &= \min\left(\frac{(8 \cdot MSS)\sqrt{1.5}}{D^1_{ij}\sqrt{L_1}}\ ,\ B^1_{ij}\right)\ (\text{bps}) \qquad (12) \\ L_1 &= \begin{cases} n_{ij} \cdot L_{\text{A}} & (\text{case(A)}) \\ L_{\text{B}} & (\text{case(B)}) \end{cases} \end{aligned}$$

L_{A} and L_{B} are parameters which determines the packet loss ratio per one AS-level hop and that per one overlay-level hop, respectively. We also denote that maximum segment size as MSS, and AS-level hop count between node N_i and N_j as n_{ij}.

The definition of TCP throughput of the relay path is different whether or not TCP connection is terminated at each relay node, which means that we utilize the TCP proxy mechanism [22] at the relay node. When we do not use the TCP proxy mechanism, meaning that we utilize an end-to-end TCP connection, we calculate the TCP throughput of the relay path from the end-to-end latency and the available bandwidth of the relay path as follows:

$$\begin{aligned} P^2_{ikj}(e2e) &= \min\left(\frac{(8 \cdot MSS)\sqrt{1.5}}{D^2_{ikj}\sqrt{L_2}}\ ,\ B^2_{ikj}\right) \qquad (13) \\ P^3_{iklj}(e2e) &= \min\left(\frac{(8 \cdot MSS)\sqrt{1.5}}{D^3_{iklj}\sqrt{L_3}}\ ,\ B^3_{iklj}\right) \qquad (14) \\ L_2 &= \begin{cases} (n_{ik} + n_{kj}) \cdot L_{\text{A}} & (\text{case(A)}) \\ 2L_{\text{B}} & (\text{case(B)}) \end{cases} \end{aligned}$$

When the TCP proxy mechanism is deployed, on the other hand, we determine the TCP throughput of the relay path as the minimum of TCP throughput of direct paths constructing the relay path:

$$\begin{aligned} P^2_{ikj}(pxy) &= \min\left(P^1_{ik},\ P^1_{kj}\right) \qquad (15) \\ P^3_{iklj}(pxy) &= \min\left(P^1_{ik},\ P^1_{kl},\ P^1_{lj}\right) \qquad (16) \\ L_3 &= \begin{cases} (n_{ik} + n_{kl} + n_{lj}) \cdot L_{\text{A}} & (\text{case(A)}) \\ 3L_{\text{B}} & (\text{case(B)}) \end{cases} \end{aligned}$$

As in the case of the available bandwidth, we can define the throughput-optimized relay paths as follows:

$$\begin{aligned}
\hat{P}^2_{ij}(e2e) &= \max_{k \neq i,j} \left(P^2_{ikj}(e2e)\right) & (17)\\
\hat{P}^3_{ij}(e2e) &= \max_{k \neq l,\ k,l \neq i,j} \left(P^3_{iklj}(e2e)\right) & (18)\\
\hat{P}^2_{ij}(pxy) &= \max_{k \neq i,j} \left(P^2_{ikj}(pxy)\right) & (19)\\
\hat{P}^3_{ij}(pxy) &= \max_{k \neq l,\ k,l \neq i,j} \left(P^3_{iklj}(pxy)\right) & (20)
\end{aligned}$$

Furthermore, improvement ratio of the relay path with respect to the direct path can be described as follows:

$$\begin{aligned}
I(P^2_{ikj}(e2e)) &= \frac{P^2_{ikj}(e2e)}{P^1_{ij}}\\
I(P^3_{iklj}(e2e)) &= \frac{P^3_{iklj}(e2e)}{P^1_{ij}}\\
I(P^2_{ikj}(pxy)) &= \frac{P^2_{ikj}(pxy)}{P^1_{ij}}\\
I(P^3_{iklj}(pxy)) &= \frac{P^3_{iklj}(pxy)}{P^1_{ij}}
\end{aligned}$$

2.2 Evaluation results and discussions

2.2.1 Performance distribution of overlay path

In Figure 5, we show the distributions of end-to-end latency and available bandwidth of direct paths and relay paths for all node pairs. We can observe from Figure 5(a) that 80% of the direct paths have an available bandwidth of from 10 Mbps to 100 Mbps. However, using the relay path, the ratio increases to 90%. For end-to-end latency (Figure 5(b)), roughly half of the direct paths the end-to-end latency from 10 ms to 100 ms, and it increases to 80% by using relay paths. Furthermore, the degree of improvement is quite large, especially when the performance of the direct path is not good: less than 10 Mbps for available bandwidth and greater than 20 msec for end-to-end latency. From these results, we can expect to find a relay path that has better performance than that of the direct path in both terms of end-to-end latency and available bandwidth, especially when the performance of the direct path is not so good.

In Figure 6, we show the distributions of TCP throughput of direct paths, for four cases of

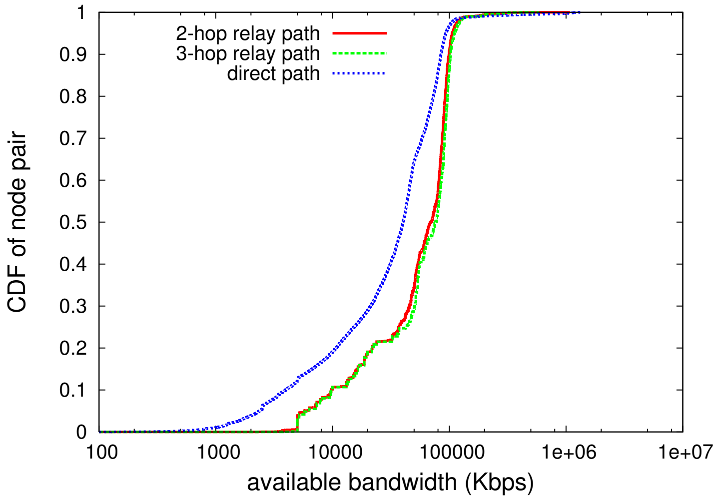

(a) Available bandwidth

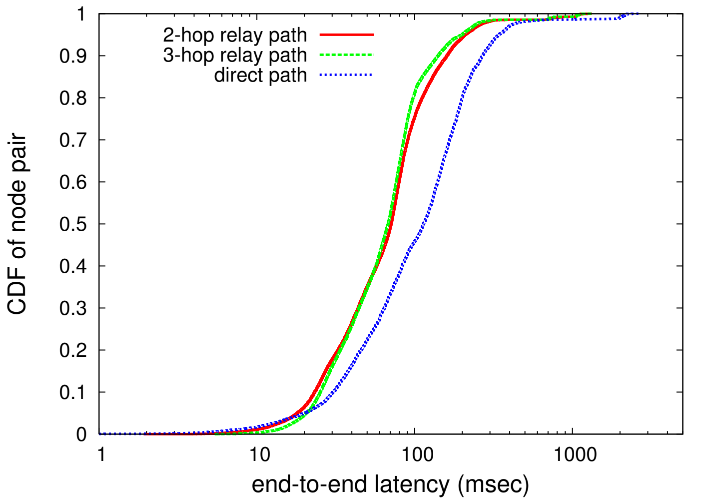

(b) End-to-end latency

Figure 5: Distribution of end-to-end latency and available bandwidth

combination of packet loss ratio and TCP connection setting in calculating TCP throughput. We denote the four cases as follows:

AS/e2e Packet loss ratio is proportional to AS-level hop count of the path (AS), and end-to-end TCP connection is utilized (e2e).

OL/e2e Packet loss ratio is proportional to overlay-level hop count of the path (OL), and end-to-end TCP connection is utilized (e2e).

AS/pxy Packet loss ratio is proportional to AS-level hop count of the path (AS), and TCP proxy mechanism is deployed (pxy).

OL/pxy Packet loss ratio is proportional to overlay-level hop count of the path (OL), and TCP proxy mechanism is deployed (pxy).

We can observe from Figure 6 that TCP throughput in OL cases (Figure 6 (b) and (d)) is better than that in AS cases (Figure 6 (a) and (c)), and that TCP throughput in pxy cases (Figure 6 (c) and (d)) is better than that in e2e cases (Figure 6 (a) and (b)). The reason of this is that packet loss ratio of the OL case is not depend on AS-hop-count, therefore it generally becomes smaller than that of AS case, and that TCP proxy mechanism can isolate the effect of packet loss, resulting the TCP throughput remains unaffected by the packet losses at other parts of the overlay path. We also observe that TCP throughput do not reach available bandwidth even when $L = 0.00001$ for all cases. This results may indicate that the bandwidth resource of PlanetLab nodes is enough large, and the available bandwidth is far larger than TCP throughput to be achieved.

2.2.2 Characteristics of relay path

In Figure 7, we present the distribution of the relationship between the available bandwidth of the direct path and that of the bandwidth-optimized relay path for each node pair, for 2-hop relay paths (Figure 7(a)) and 3-hop relay paths (Figure 7(b)), respectively. Figure 8 shows similar plots for end-to-end latency. For 96.6% of all node pairs, we can find a 2-hop relay path that has a larger available bandwidth than the direct path. When we compare the direct path and the 3-hop relay path, for 97.7% of all node pairs, we can find a 3-hop relay path that has a larger available bandwidth. For end-to-end latency, these percentages decrease to 87.5% and 85.4%, respectively.

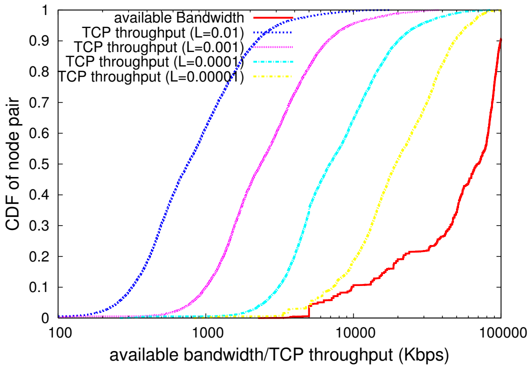

(a) AS/e2e

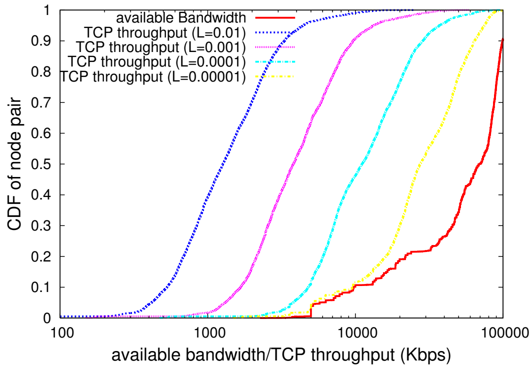

(b) OL/e2e

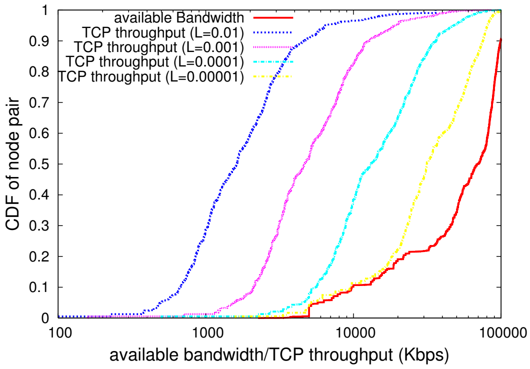

(c) AS/pxy

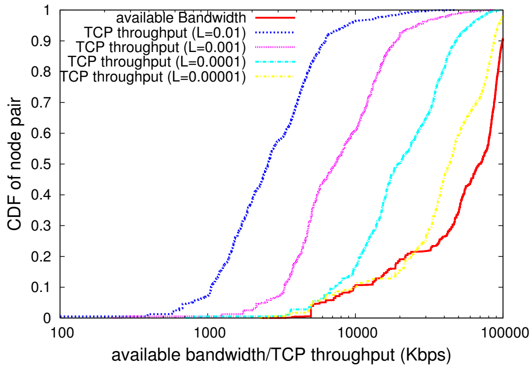

(d) OL/pxy

Figure 6: Distribution of TCP throughput

Furthermore, with respect to available bandwidth, 46.9% of node pairs for which a better 2-hop relay path cannot be found, a 3-hop relay path having a larger available bandwidth than the direct path can be found. In addition, for 51.6% of the node pairs that has a larger available bandwidth than the direct path, we can find a better 3-hop relay path than the bandwidth-optimized 2-hop relay path. With respect to end-to-end latency, these percentages decrease to 17.8% and 47.3%, respectively.

The above results indicate that the effectiveness of the latency-based relay path is smaller than that of the available bandwidth-based relay path. A reasonable explanation for this is that the underlying IP routing is configured based on router-level and AS-level hop count, which have some degree of correlation with the end-to-end delay.

In Figure 9, we present the distribution of the relationship between the TCP throughput of the direct path and that of the throughput-optimized 2-hop relay path for each node pair. We plot for four cases of combination of packet loss ratio and TCP connection setting in calculating TCP throughput. We set $L =$ 0.00001 in this figure.

From Figure 9(a), for 45.9% of all node pairs, we can find a 2-hop relay path that has a larger TCP throughput than the direct path, and 47.8% for Figure 9(b). for Figures 9 (c), and (d), these percentages increase 95.2%, and 95.8%, respectively. The reasons can be explained as follows.

When we use TCP proxy, the TCP throughput does not degrade significantly due to the effect of TCP proxy mechanism described in the previous subsection. Therefore, the effectiveness of relay path becomes almost similar to that of available bandwidth shown in Figure 7. On the other hand, by comparing Figures 9 (a) and (c), and Figures 9 (b) and (d), the effect of packet loss model does not affect the effectiveness of the relay path, compared with the that of using TCP proxy. From these results, if we utilize TCP throughput for the metric in the overlay routing, introducing the TCP proxy mechanism is a key issue to improve the performance.

Next, we present the distribution of the improvement ratio of the bandwidth-optimized 2-hop and 3-hop relay paths with respect to the direct path in Figure 10(a). In the figure, we also plot the improvement ratio of the bandwidth-optimized 3-hop relay path with respect to the bandwidth-optimized 2-hop relay path. In Figure 10(b), we present similar results for end-to-end latency. These figures indicate that by using the relay path, we can obtain a significant improvement in terms of both available bandwidth and end-to-end latency. However, the effectiveness of 3-hop relay path is quite limited when compared to 2-hop relay path. Thus, seeking 3-hop relay path

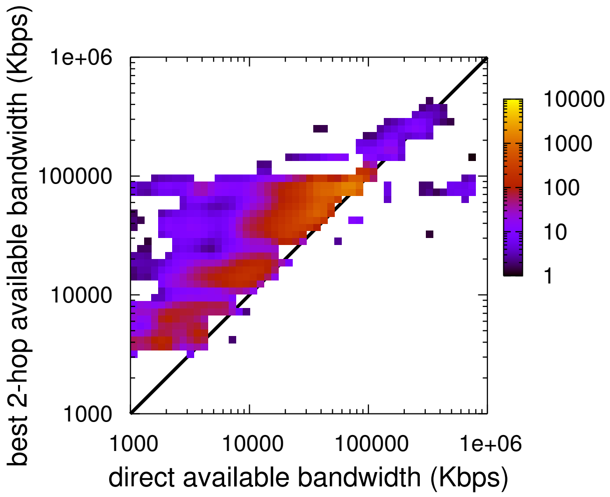

(a) 2-hop relay path

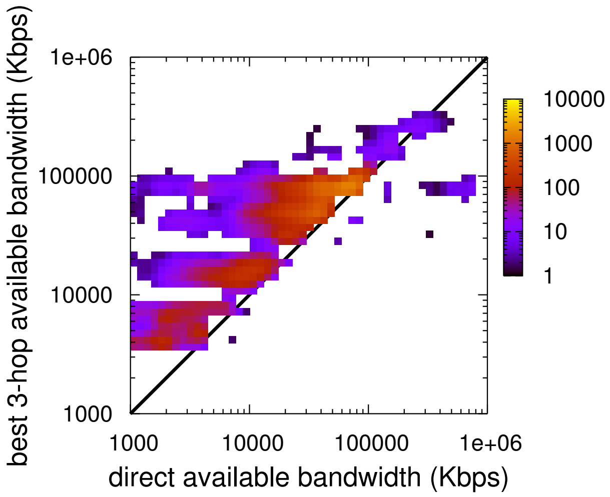

(b) 3-hop relay path

Figure 7: Available bandwidths for the direct path and the bandwidth-optimized relay path

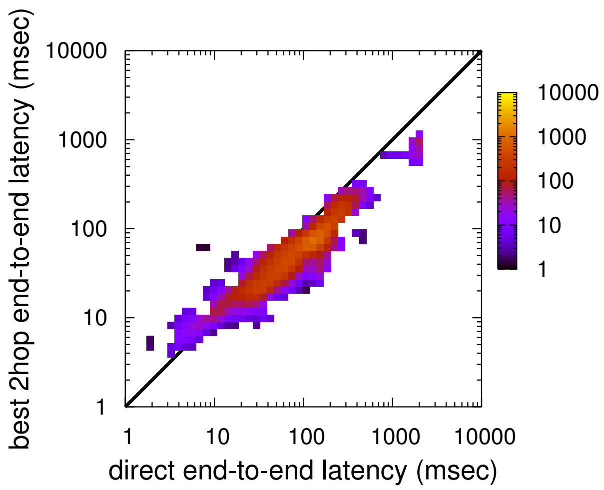

(a) 2-hop relay path

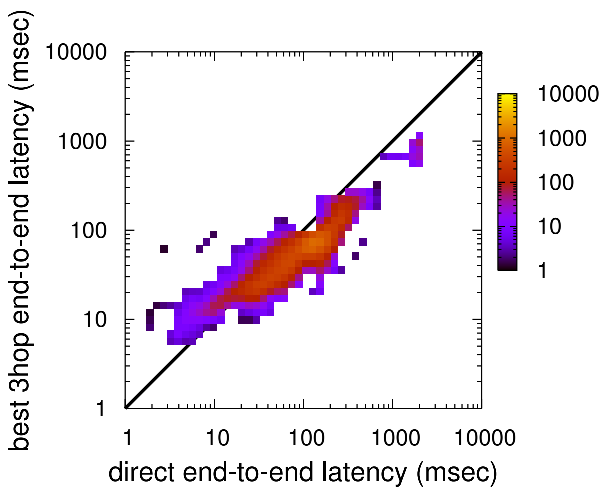

(b) 3-hop relay path

Figure 8: End-to-end latencies for the direct path and the latency-optimized relay path

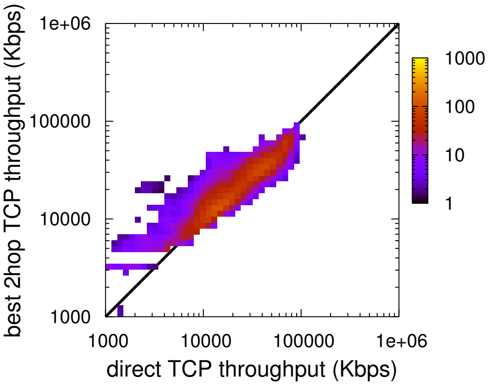

(a) AS/e2e

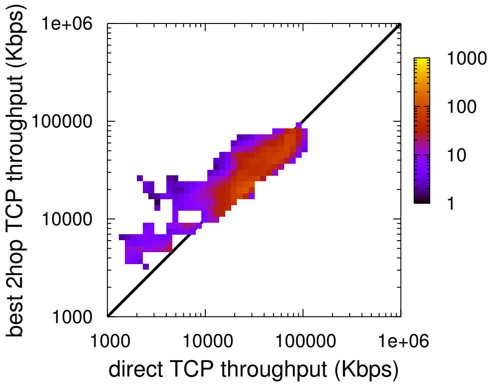

(b) OL/e2e

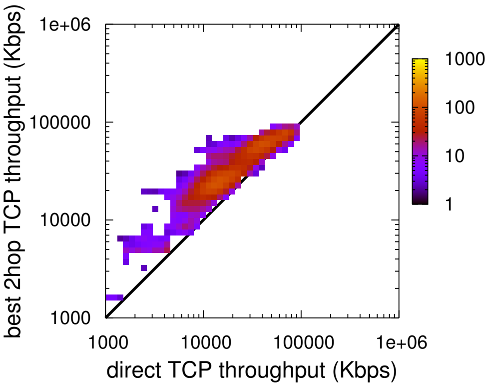

(c) AS/pxy

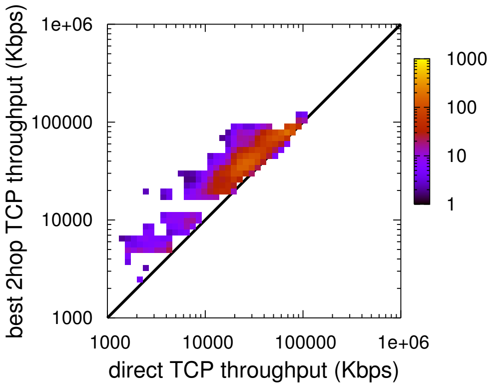

(d) OL/pxy

Figure 9: TCP throughputs for the direct path and the throughput-optimized relay path

has a limited effect for overlay routing when we consider normal data transmission using a single path. However, when we consider multipath data transmission, 3-hop relay paths may become possible candidates for path selection. The effectiveness of the 3-hop relay path for multipath data transmission is discussed in Subsection 2.2.3.

We present the distribution of the improvement ratio of the throughput-optimized 2-hop path with respect to the corresponding direct path in Figure 11. The results of the cases of end-to-end latency and available bandwidth are also plotted in these graphs. We can observe from Figure 11 that, for $50 - 70\%$ of all node pairs, we cannot find any 2-hop relay path which has better performance than the corresponding direct path in AS/e2e case. However, the performance of 2-hop relay path slightly increses in OL/e2e case. Furthermore, by using TCP proxy, we can obtain almost similar performance gain to available bandwidth, when the improvement ratio falls between 1 and 2. However, in the region where the improvement ratio is larger than 2, the effectiveness of using TCP throughput does not reach that of available bandwidth. This is the same reason as in Figure 6: since TCP throughput is affected by packet loss ratio of the network, the performance gain of the relay path degrades.

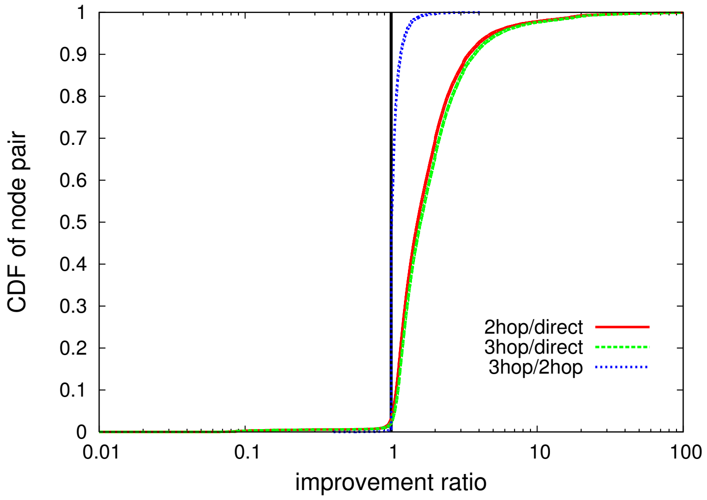

(a) Available bandwidth

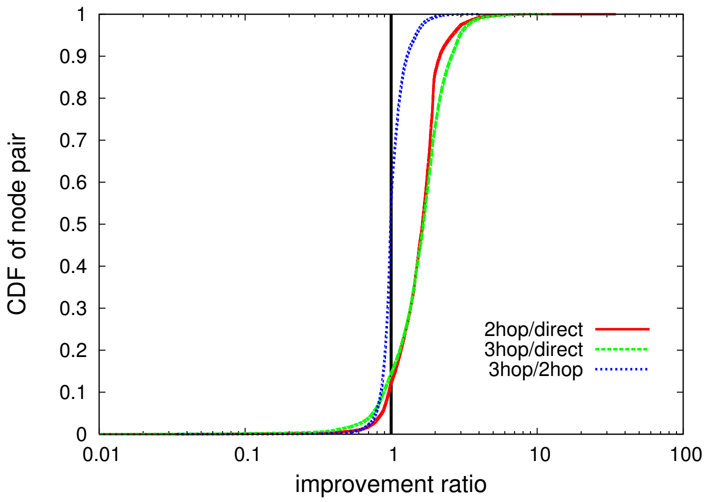

(b) End-to-end latency

Figure 10: Distribution of improvement ratio for available bandwidth and end-to-end latency

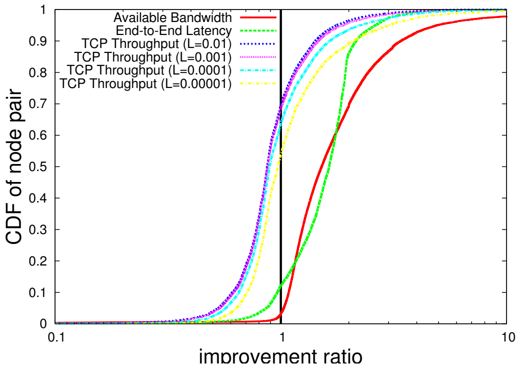

(a) AS/e2e

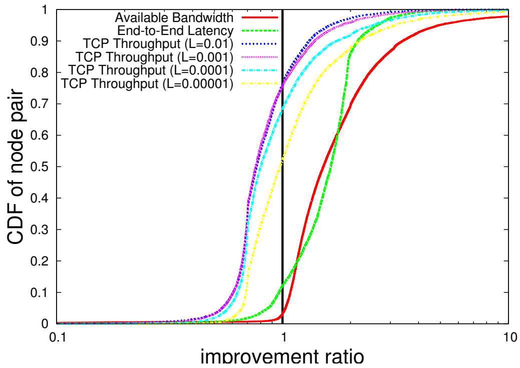

(b) OL/e2e

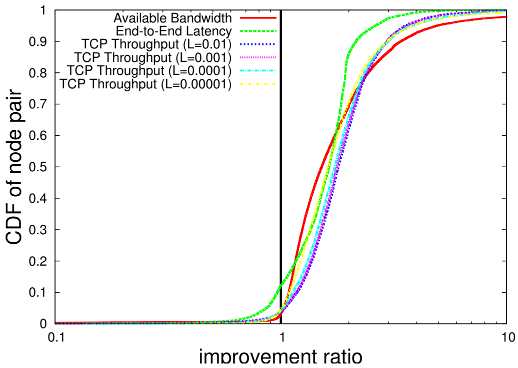

(c) AS/pxy

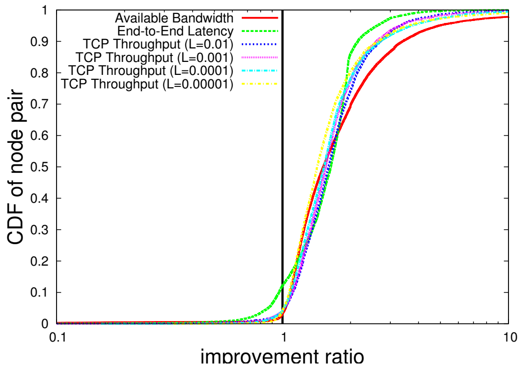

(d) OL/pxy

Figure 11: Distribution of improvement ratio for TCP throughput

2.2.3 Effectiveness in multipath transmission

We next investigate the effectiveness of seeking the 3-hop relay path in multipath transmission. Here, we define multipath transmission as data transmission using multiple paths for one data transmission between source and destination nodes. We assume that we choose the multiple paths in the best order of available bandwidth or end-to-end latency from all of the direct, 2-hop, and 3-hop paths with considering the path disjointness of selected paths. Figure 12 shows the average ratio of the number of direct, 2-hop, and 3-hop paths in the selected multiple paths, as a function of the total number of using paths in multipath transmission, when we use end-to-end latency and available bandwidth for performance metric. This figure shows that seeking 3-hop relay paths is meaningful in multipath transmission with a few paths, but its effectiveness gradually decreases as the number of total using paths in multipath transmission increases.

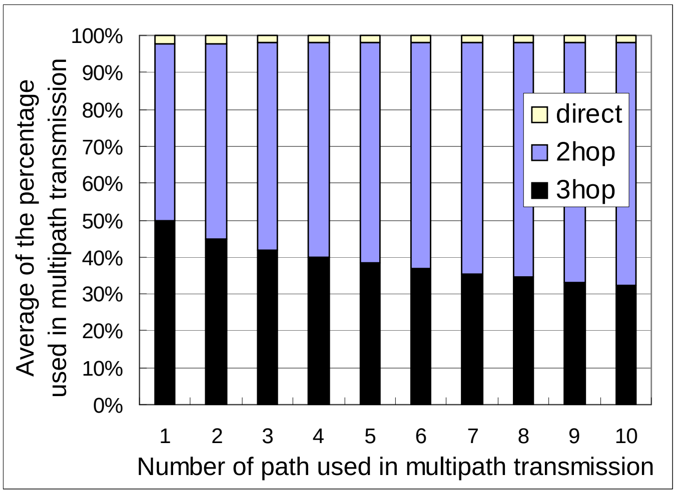

(a) Available bandwidth

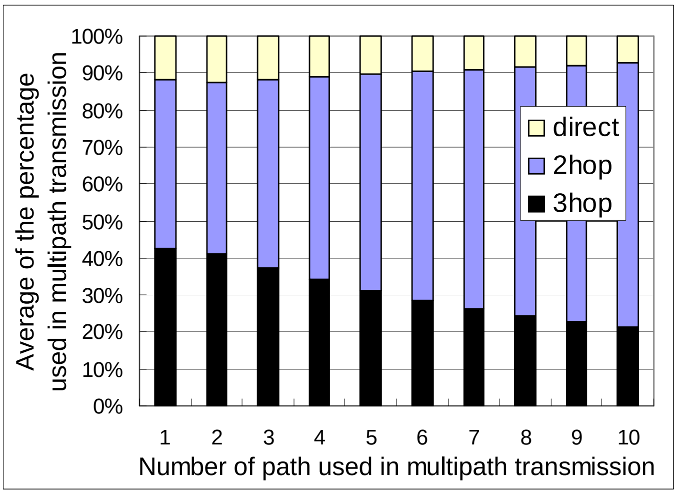

(b) End-to-end latency

Figure 12: Breakdown of paths used in multipath transmission

2.2.4 Correlation between available bandwidth and end-to-end latency

Finally, we investigate the correlation between improvement ratio in available bandwidth and end-to-end latency, in order to clarify whether or not a "good" relay overlay path for available bandwidth is also good for end-to-end latency, and vise versa. Note that we do not consider TCP throughput, since this metric is calculated from end-to-end latency and available bandwidth. In Figure 13(a), we plot the relationship between the improvement ratio of the bandwidth-optimized 2-hop relay path and the improvement ratio of the path in end-to-end latency. Figure 13(b) shows a similar graph for the bandwidth-optimized 3-hop relay path.

From these figures, we can provide the following observations: when we can find a multi-hop relay path that has a larger available bandwidth than the direct path, such path has a larger end-to-end latency than the direct path. That is, when we select the overlay path based on the available bandwidth, the selected path generally has a large end-to-end latency. Therefore, we should carefully choose the metric in selecting overlay paths according to the characteristics of upper-layer applications. We also note that, when we cannot find a relay path that has a larger available bandwidth than the direct path ($x < 1.0$ in Figures 13 (a) and b)), such relay paths have a significantly larger end-to-end latency. In such cases, simply choosing the direct path is reasonable, regardless of the type of upper-layer applications.

Figure 14 is the opposite graph to Figure 13: plots the relationships between the improvement ratio of the latency-optimized 2-hop relay path and the improvement ratio of the path in available bandwidth. Figure 14(b) is a similar graph for the latency-optimized 3-hop relay path. In contrast to the previous results in Figure 13, these figures indicate that when we choose the latency-optimized relay path, it is likely that the path also has a larger available bandwidth than the direct path. This means that when we choose the path based on end-to-end latency, the path generally has a larger available bandwidth than the direct path.

One can imagine from these results that it is sufficient to select the overlay path based only on end-to-end latency and that it is meaningless to observe the available bandwidth. However, Figure 15, which plots the distribution of the ratio of the available bandwidth of latency-optimized relay path with respect to the available bandwidth of bandwidth-optimized relay path for all node pairs, clearly shows that the available bandwidth of the latency-optimized relay path is significantly smaller than that of the bandwidth-optimized relay path. That is, when we want to find a

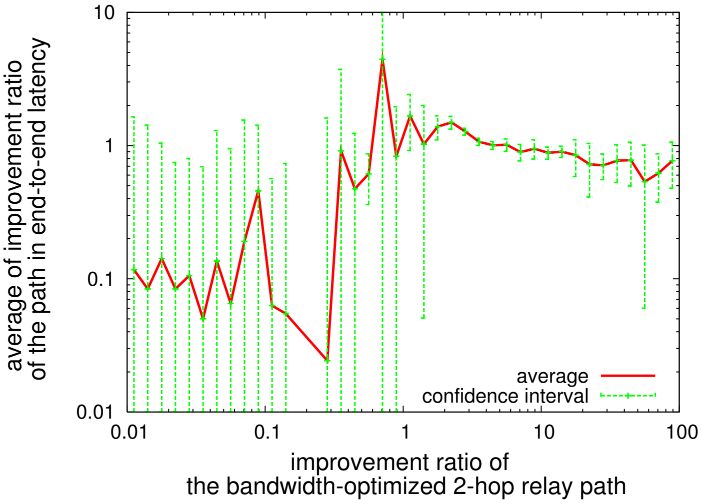

(a) 2-hop relay path

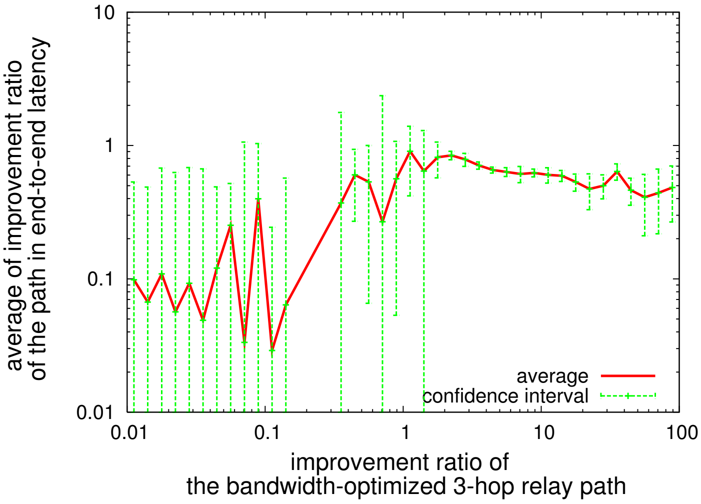

(b) 3-hop relay path

Figure 13: Correlation between end-to-end latency and available bandwidth of overlay paths (1)

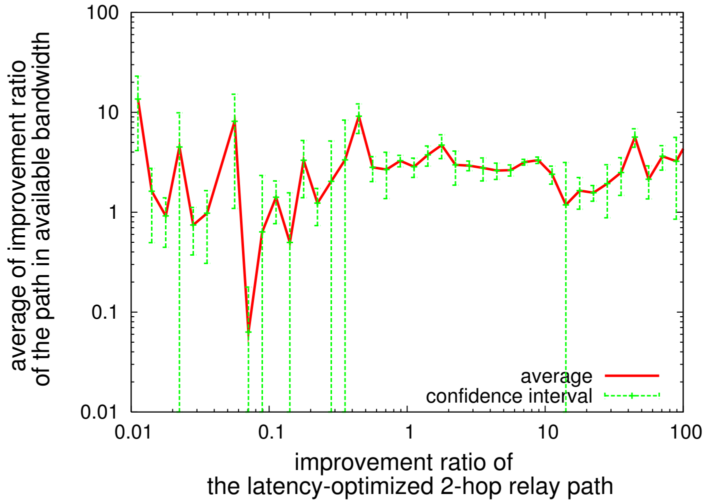

(a) 2-hop relay path

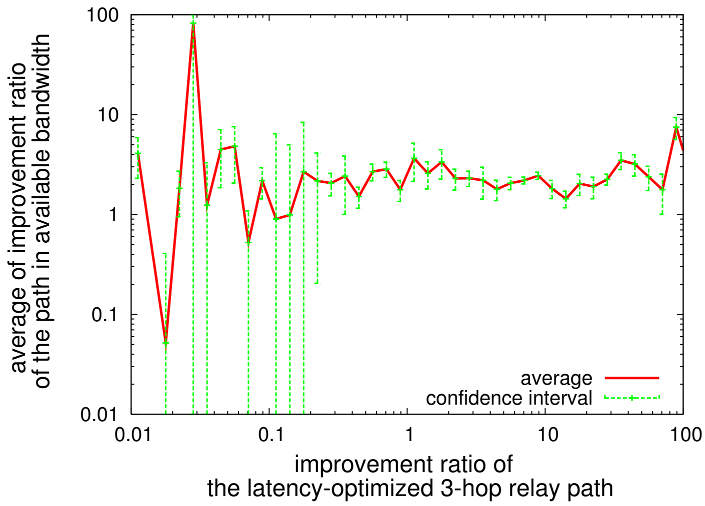

(b) 3-hop relay path

Figure 14: Correlation between end-to-end latency and available bandwidth of overlay paths (2)

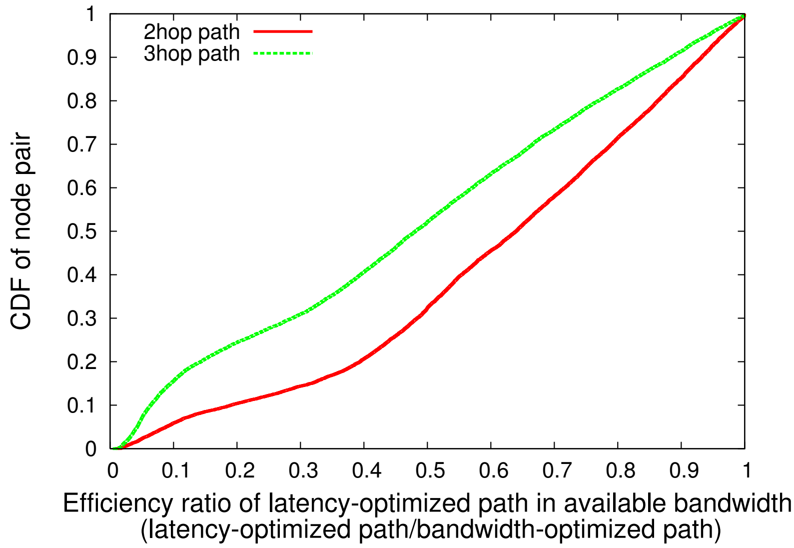

Figure 15: Distribution of ratio of available bandwidth of latency-optimized relay path to that of bandwidth-optimized relay path

data transmission path with sufficiently large available bandwidth, we should directly measure the available bandwidths of the overlay network paths.

However, since a large number of packets are required for measuring the available bandwidth than for measuring end-to-end latency, we propose one possible guideline for selecting the data transmission path in routing overlay networks for the bandwidth-centric applications as follows: When we transmit the data to a destination where we do not have sufficient information on the available bandwidth, we select the path based on end-to-end latency. When we have sufficient and accurate information on the available bandwidth, we choose the path based on available bandwidth.

3 Free-riding traffic problem

In this section, the free-riding traffic problem, which causes by policy mismatch between overlay routing and IP routing, is defined in detail, and define the numerical metric for assessing the degree of the problem. We then evaluate the free-riding problem possibly occurred in routing overlay networks and show some guidelines for selecting overlay paths to mitigate the problem.

3.1 Problem definition

Free-riding traffic problem is occured from policy mismatch between overlay routing and IP routing. In overlay routing, the metrics determining the route are user-perceived performance such as data transmission throughput and end-to-end delay. On the other hand, in IP routing, the metric determining the route is mainly a hop count, which may not directly related to user-perceived performance. Furthermore, most of IP routing is configured by political and financial factors: the billing mechanism of transit links to upper-layer ISPs, relationships between the ISP and other ISPs interconnected by public or private peering links, and the amount of traffic to each ISP. In ordinary cases, the monetary cost for the usage of transit links, which lower-layer ISPs must pay to upper-layer ISPs, is determined by the amount of traffic passing through the transit links. 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. However, the application-level overlay routing does not consider such billing structure of ISPs and determines the route according only to the user-perceived performance. As a result, overlay routing mechanisms can generate network traffic which ignores an ISP's billing structure. We call such traffic as "free-riding" traffic, and this problem as free-riding traffic problem.

We explain the free-riding traffic problem by using Figure 16. In this figure, there are three ISPs (ISP I, J and K), where ISP J is the transit ISP for ISPs I and K. ISP I and K have transit links C_{ij} and C_{kj} to connect to ISP J. Furthermore, ISP I and K are interconnected by a peering link C_{ik}. Hosts $N_i,\ N_j$, and N_k exist in ISP I, J and K, respectively. These three hosts are the overlay nodes of the routing overlay network.

We consider a situation where Host N_i transmits the overlay network traffic to Host N_j by using the routing overlay network. When we use the direct path from N_i to N_j, the traffic is transmitted by transit link C_{ij} as shown by an orange arrow in Figure 16. Therefore, the cost of

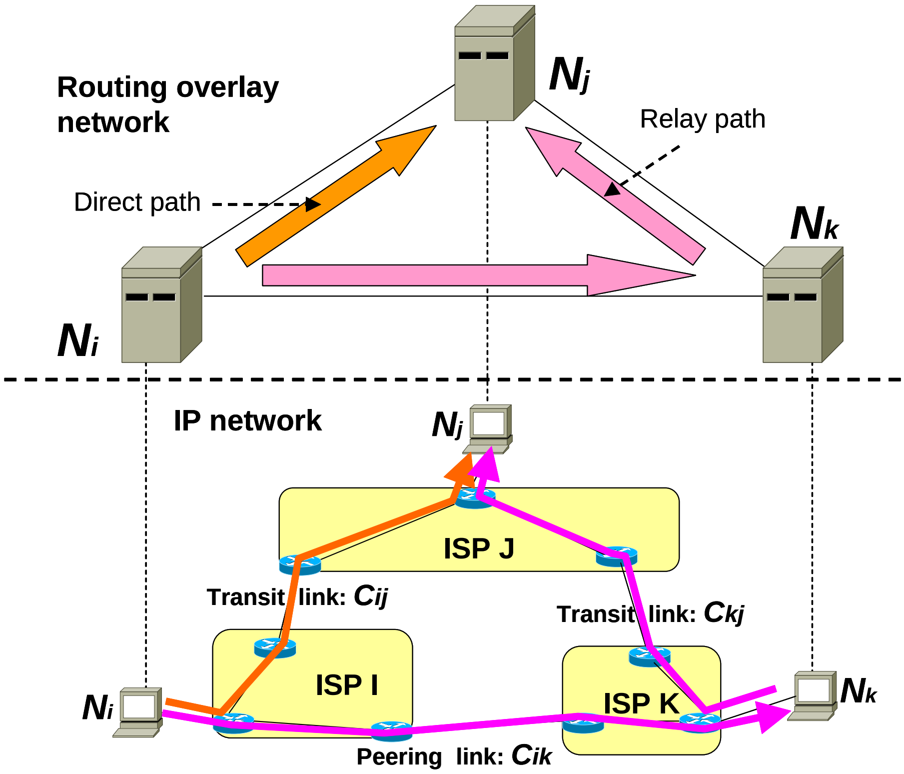

Figure 16: Free-riding traffic problem

conveying the traffic is charged to ISP I. This is the normal billing architecture, and ISP I can take the transit cost for this traffic from Host N_i. However, when we use the relayed path via Host N_k (pink arrows in Figure 16), the traffic is transmitted by peering link C_{ik} from Host N_i to Host N_k, and the transit link C_{kj} from Host N_k to Host N_j. In this case, ISP K pays the cost for using transit link C_{kj} to convey the traffic, although only the customers in ISPs I and J benefit from the transmission. We call this mismatch the "free-riding" traffic problem.

If ISP K monitors the traffic coming from the peering link C_{ik} and differentiates the free-riding traffic from the normal traffic, we can resolve this problem by restricting the free-riding traffic from coming into ISP K, or charging ISP I for the cost of the free-riding traffic. However, since overlay routing is operated by upper-layer protocols and applications, we cannot recognize the free-riding traffic by simply checking source and destination IP addresses of incoming packets. Therefore, we believe that this problem will become harmful for ISPs especially when the amount of traffic conveyed by the overlay network increases.

3.2 Evaluation methodology

3.2.1 Dataset used for evaluation

For the measurement results of end-to-end delay and available bandwidth which are necessary to select overlay path in routing overlay network, we utilize the same dataset explained in the previous section. We also utilize the information on the relationships between ASes, which is obtained from CAIDA. In [23], the investigated results on the relationship of links interconnected two ASes (inter-AS links) are provided, which is esitmated from BGP table information and a degree of each AS (the number of links of an AS to other ASes). The methods of AS relationship is refered in [24]. We obtain the AS-level route between PlanetLab nodes with traceroute command between the nodes, and AS-level traceroute results from a route server in traceroute.org to each node. However, 36.1% of inter-AS links utilized by routes between PlanetLab nodes are not found in [23]. The possible reason is that such inter-AS links do not exist in the BGP tables utilized in [24].

For determining the relationships of those unknown inter-AS links, we consider the following two cases: (i) Assuming that all of unknown inter-AS links are peering links, and they do not have any influence on free-riding problem. (ii) Determining the relationships of unknown inter-AS

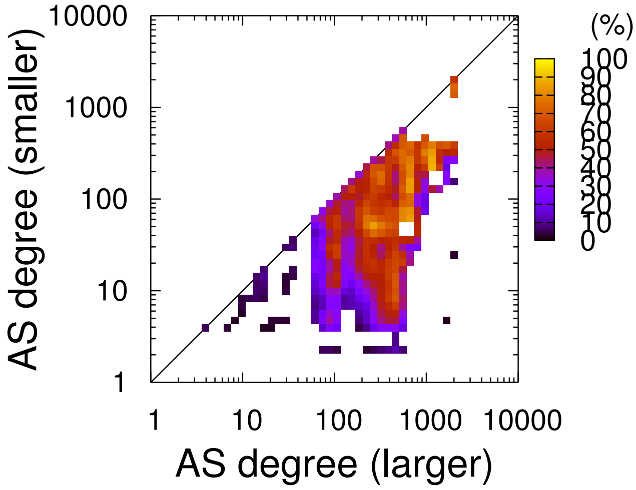

Figure 17: Distribution of ratio of peering links

links by the degree distribution of ASes revealed in [24]. In Figure 3.2.1, we show the distribution of the ratio of peering links as functions of degrees of interconnected ASes, which is obtained from [23]. In case (ii), we use this distribution to determine the relationship of unknown inter-AS links stochastically.

3.2.2 Evaluation metric

According to the problem definition in Subsection 3.1, we define the numerical metric of the degree of the free-riding problem as the number of transit links increased by utilizing relay paths. In detail, we introduce the metric called *the number of free-riding transit links* of a relay path as follows. We denote T_{ij} and T_{ikj} as sets of transit links constructing a direct path between nodes N_i and N_j, and that of a relay path via node N_k. Furthermore, we set F_{ikj} as the set of transit links that exists in T_{ikj} and that does not appear in T_{ij}. Then we obtain:

$$F_{ikj} = \{x|(x \in T_{ikj})\&(x \notin T_{ij})\} \quad (21)$$

In this thesis, we define $|F_{ikj}|$, the number of members in F_{ikj}, as the number of free-riding transit links used the relay path via node N_k.

Table 1: Average of $|F_{ikj}|$ in case (i)

	Available bandwidth	End-to-end latency	TCP throughput			
			$L = 0.01$	$L = 0.001$	$L = 0.0001$	$L = 0.00001$
AS/e2e	2.55	2.56	1.44	1.38	1.24	1.31
OL/e2e	2.55	2.56	2.58	2.41	2.20	2.29
AS/pxy	2.55	2.56	1.56	1.39	1.35	1.66
OL/pxy	2.55	2.56	2.65	2.32	2.34	2.47

3.3 Evaluation results

In this subsection, we evaluate the free-riding traffic problem, and discuss about guideline to mitigate problem.

In the following evaluation, MSS is 1460 bytes, and L is set to 0.00001, 0.0001, 0.001, and 0.01.

3.3.1 Number of free-riding transit links

In Figure 18, we show the distribution of $|F_{ikj}|$ of optimized 2-hop relay path for each metric in case (i), meaning that we ignore unknown inter-AS links. The average of $|F_{ikj}|$ are summarized in Table 1 for all metrics.

From these results, when we use either available bandwidth or end-to-end latency, the free-riding traffic problem occurs on roughly 90% of node pairs. In the case in using TCP throughput, the results depend on the packet loss model: when we use the AS-hop-count-base packet loss model, the degree of free-riding problem decreases significantly. This is because throughput-optimized relay path tends to have small AS-level hop count, resulting in that the number of transit links utilized by such relay path also decreases. On the other hand, when we use the overlay-hop-count-based packet loss model, the degree of the free-riding problem is similar to that in using available bandwidth. This is because in this case the relay path is selected regardless of AS-level hop count, which is the same situation as available bandwidth.

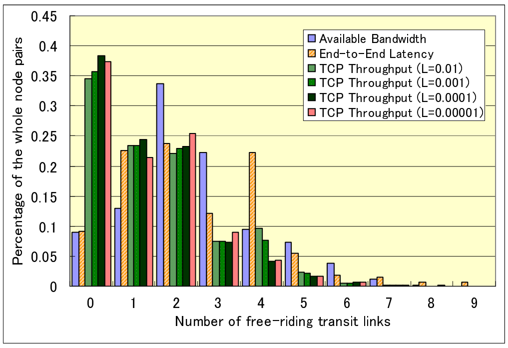

(a) AS/e2e

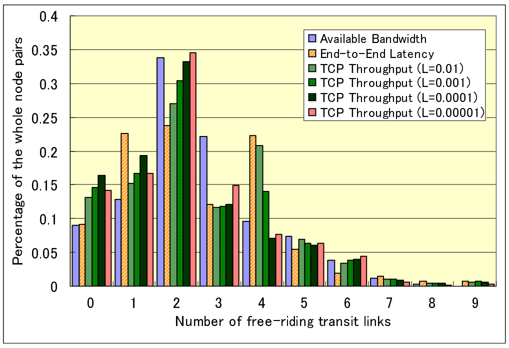

(b) OL/e2e

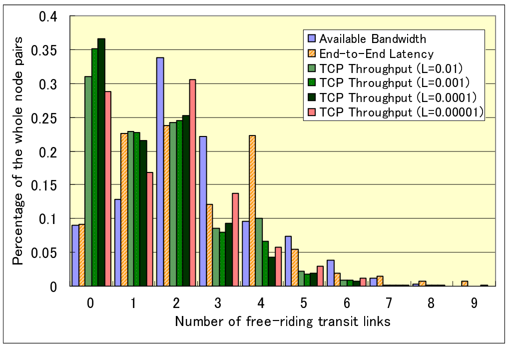

(c) AS/pxy

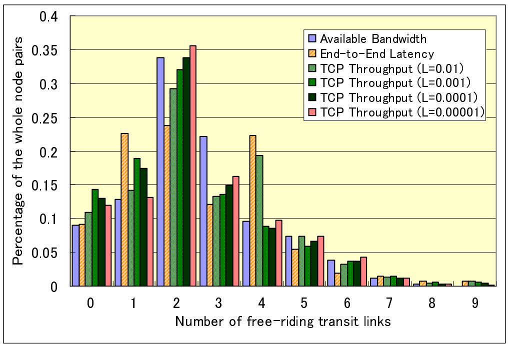

(d) OL/pxy

Figure 18: Distribution of the number of free-riding transit links in case (i)

Table 2: Average of $|F_{ikj}|$ in case (ii)

	Available bandwidth	End-to-end latency	TCP throughput			
			$L = 0.01$	$L = 0.001$	$L = 0.0001$	$L = 0.00001$
AS/e2e	4.04	4.02	4.00	3.94	3.68	3.68
OL/e2e	4.04	4.02	3.96	3.84	3.62	3.89
AS/pxy	4.04	4.02	2.93	2.77	2.76	3.15
OL/pxy	4.04	4.02	4.18	3.81	3.98	4.17

Figure 19 and Table 2 are the results of case (ii). We observe that for all routing metrics there is almost no optimized relay path without free-riding. This means that we cannot avoid the free-riding traffic problem in actual situation, when we want to maximize performance gain by using overlay routing.

From the above figures (Figure 18, and 19), we can clearly find that almost of the optimized relay path have a number of free-riding transit links.

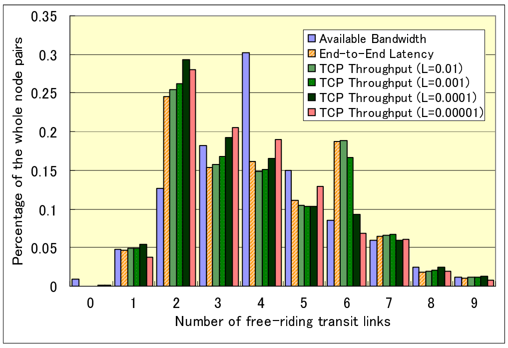

(a) AS/e2e

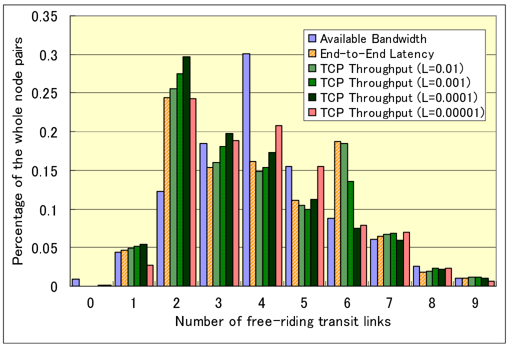

(b) OL/e2e

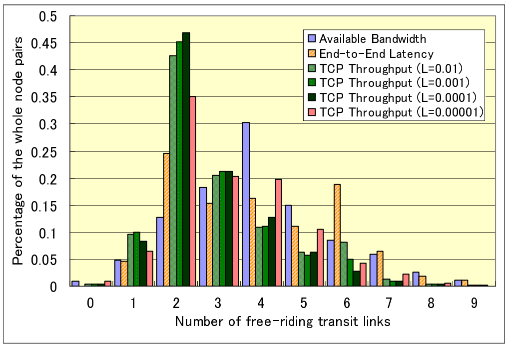

(c) AS/pxy

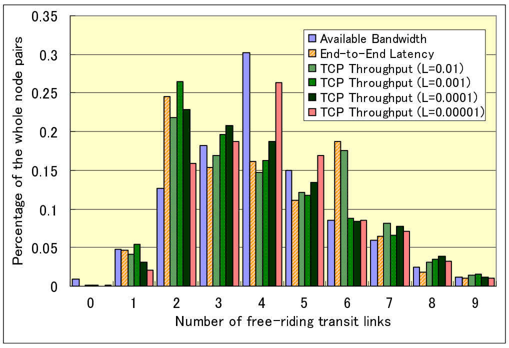

(d) OL/pxy

Figure 19: Distribution of the number of free-riding transit links in case (ii)

3.3.2 Correlation between improvement ratio and number of free-riding transit links

Here we define the set of relay nodes used in the 2-hop relay path that has smaller $|F_{ikj}|$ than optimized 2-hop relay path and higher performance than the corresponding direct path in terms of end-to-end latency, available bandwidth, and TCP throughput as $O(D_{ij}^2)$, $O(B_{ij}^2)$, $O(P_{ij}^2)$, and can be described as follows:

$$O(D_{ij}^2) = \{k|(|F_{ikj}| < |F_{i\hat{k}j}|)\&(I(D_{ikj}^2) > 1)\} \tag{22}$$

$$O(B_{ij}^2) = \{k|(|F_{ikj}| < |F_{i\hat{k}j}|)\&(I(B_{ikj}^2) > 1)\} \tag{23}$$

$$O(P_{ij}^2) = \{k|(|F_{ikj}| < |F_{i\hat{k}j}|)\&(I(P_{ikj}^2) > 1)\} \tag{24}$$

Figures 20 and 21 are the distribution of $O(D_{ij}^2)$, $O(B_{ij}^2)$, and $O(P_{ij}^2)$ of the relay path with smaller $|F_{ikj}|$ than the optimized relay path, and higher performance than the corresponding direct path. From Figure 20, in case (i), there are many node in $O(B_{ij}^2)$, when we select available bandwidth, and when we select TCP throughput in AS/e2e case, about 90% of node pair, we cannot find any relay node in $O(P_{ij}^2)$, but this percentage decreased when we select that in OL or pxy case. From Figure 21, in case (ii), we can obtain the similar results as Figure 20. However, when we select end-to-end latency, over 90% of node pair, we cannot find the relay node in $O(D_{ij}^2)$, even if in 87.5% of node pair the 2-hop latency-optimized relay path improves performance than the corresponding direct path evaluated at Figure 10 (b). We can explain the reason of this result from Figure 22.

Figure 22 is the distribution of improvement ratio of latency-optimized, latency-secondary, and latency-third paths. We can obtain from this figure that improvement ratio of the second and third path is declined greatly. Therefore, when we search relay node in $O(D_{ij}^2)$, cannot find relay path that improve performance than corresponding direct path. This result is influenced to Figure 20 and 21.

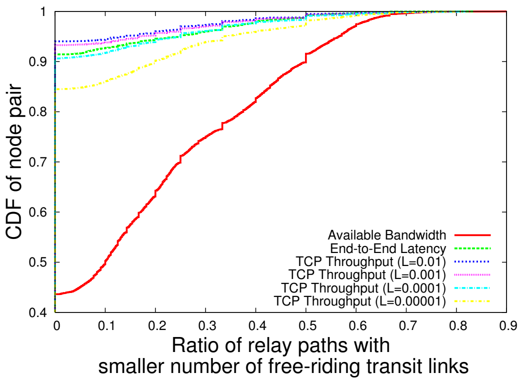

(a) AS/e2e

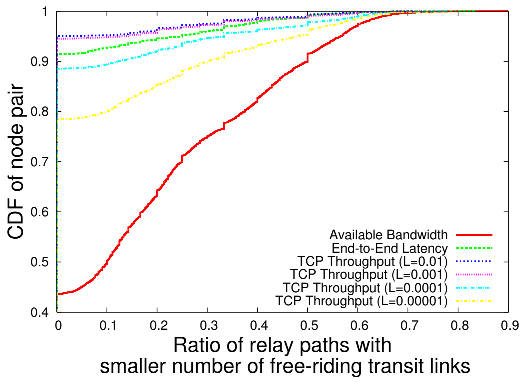

(b) OL/e2e

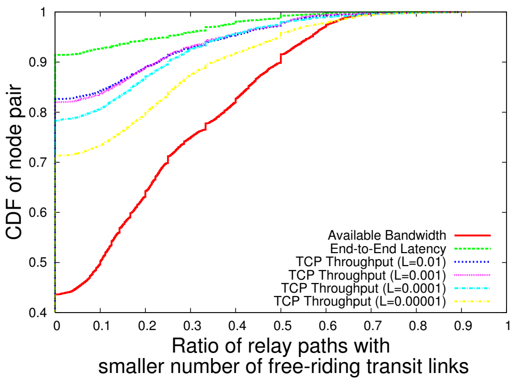

(c) AS/pxy

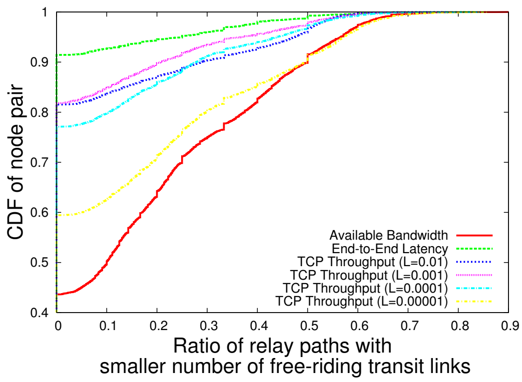

(d) OL/pxy

Figure 20: Ratio of relay path with smaller number of free-riding transit links in case (i)

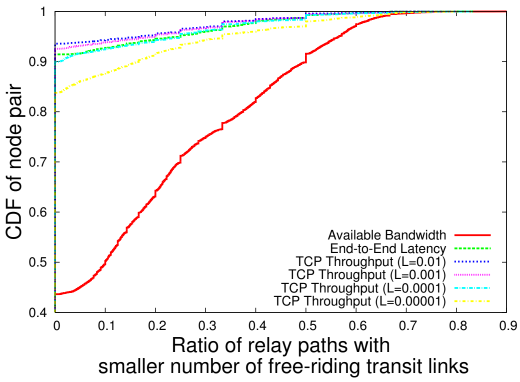

(a) AS/e2e

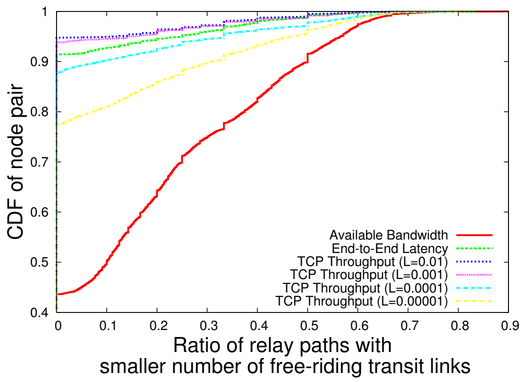

(b) OL/e2e

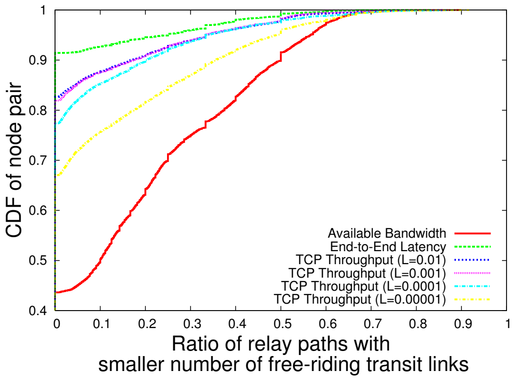

(c) AS/pxy

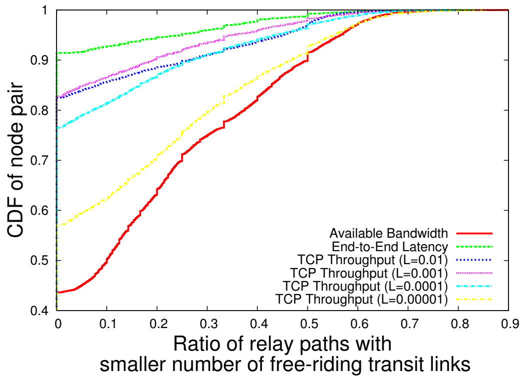

(d) OL/pxy

Figure 21: Ratio of relay path with smaller number of free-riding transit links in case (ii)

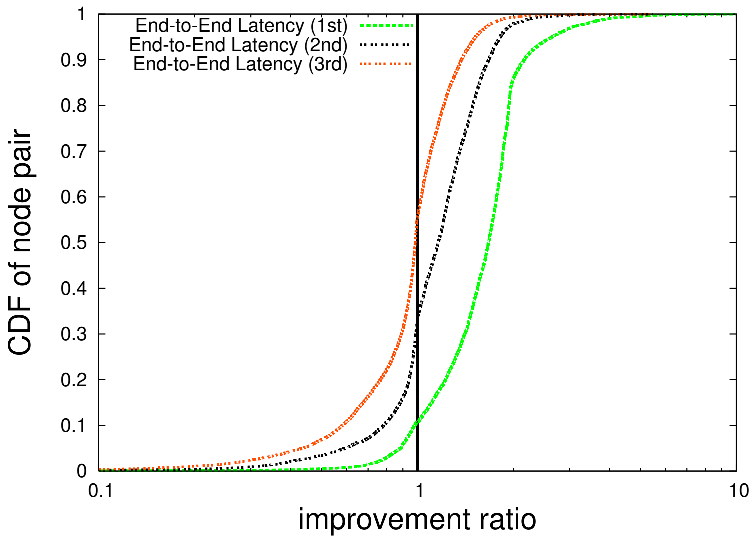

Figure 22: Distribution of improvement ratio in top three of end-to-end latency

3.3.3 Mitigation of number of free-riding transit links

From above results, we conclude that most of relay paths that have better performance than the corresponding direct path bring a sighificant degree of free-riding problem. Therefore, we should permit some degree of free-riding in order to improve the user-perceived performance in routing overlay networks. In Figure 23 and 24, we show the relationships between the maximum number of free-riding transit links permitted in selecting overlay paths and the improvement ratio to optimized 2-hop relay path for each node pair. From Figure 23, the rate of improve ratio increases as the limit of $|F_{ikj}|$ increases to 4, and remain almost unchanged as the limit of $|F_{ikj}|$ is over 4. When we focus on available bandwidth, such rate is over 80% when limit of $|F_{ikj}|$ is 2, and increased to over 90% when limit of $|F_{ikj}|$ is 3. When we focus on end-to-end latency, these percentages decrease to about 65% and 70%, respectively. Furthermore, when we focus on TCP throughput, TCP throughput in AS case is better than that in OL case, and e2e case is better than pxy case. From Figure 24, we can see the same tendency. However, when we compare Figure 23 and 24 about TCP throughput, the maximum improvement ratio in case (ii) is quite worse than that in case (i).

From the above results, we can conclude that the most efficient metric that save the number of free-riding transit links and improve performance than the corresponding direct path is available bandwidth.

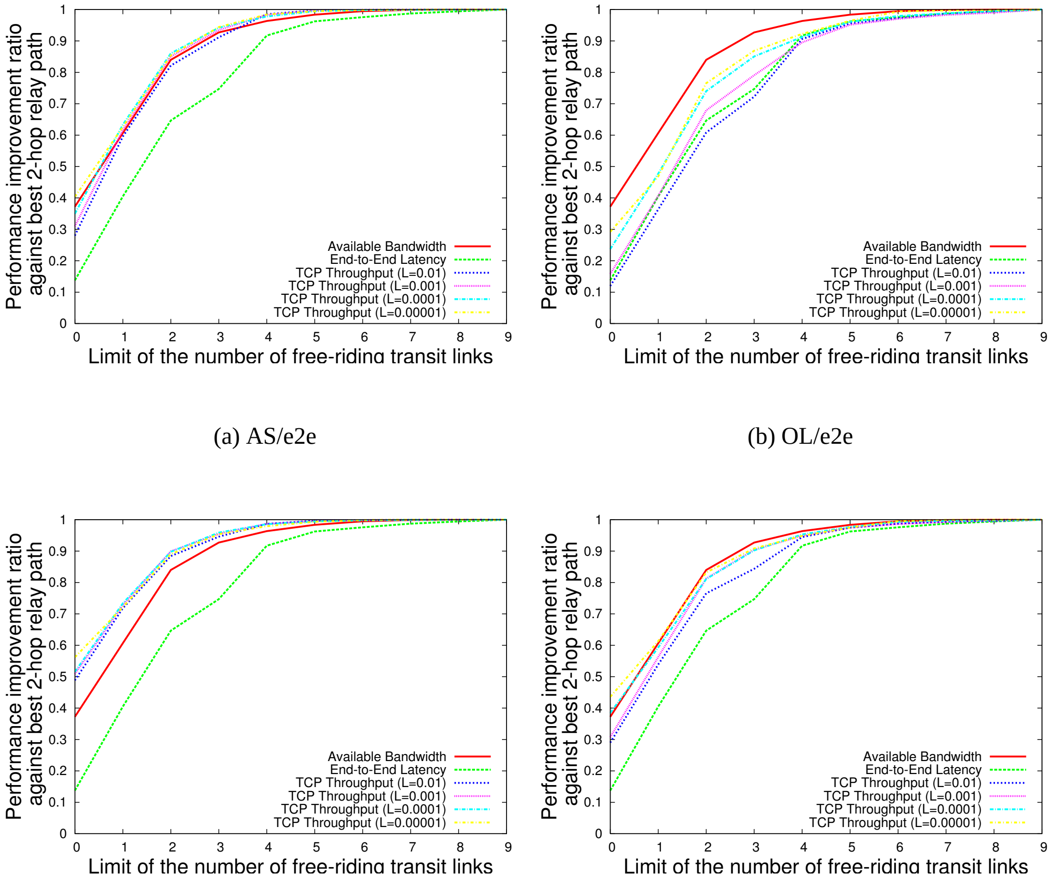

(a) AS/e2e

(b) OL/e2e

(c) AS/pxy

(d) OL/pxy

Figure 23: Correlation between improvement ratio and number of free-riding transit links in case (i)

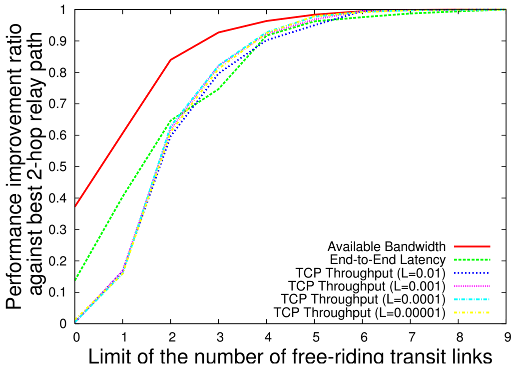

(a) AS/e2e

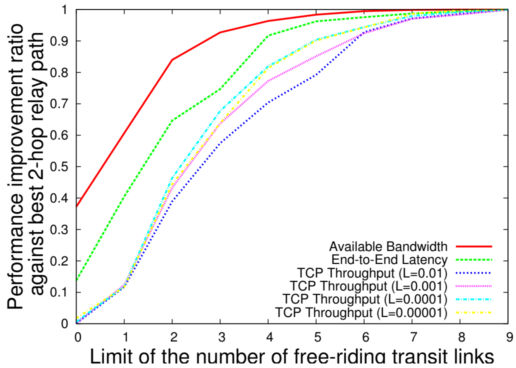

(b) OL/e2e

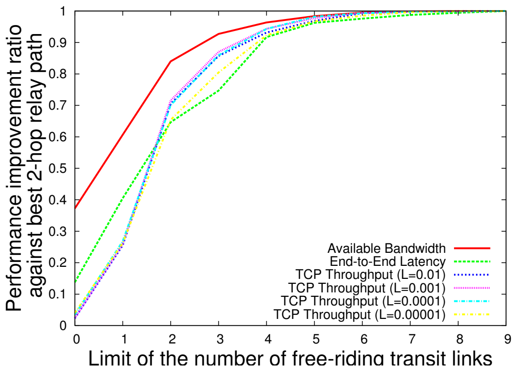

(c) AS/pxy

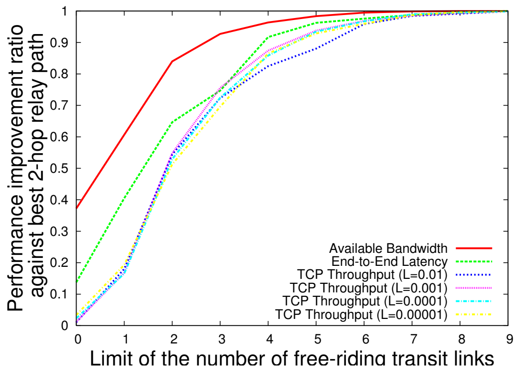

(d) OL/pxy

Figure 24: Correlation between improvement ratio and number of free-riding transit links in case (ii)

4 Conclusion

In this thesis, we first focused on overlay routing based on delay and bandwidth information. We considered three metrics in selecting overlay route: available bandwidth, end-to-end latency, and TCP throughput. By investigating the effectiveness of overlay routing on the assumption that the PlanetLab nodes construct the routing overlay network, the following results was presented: when we select bandwidth-optimized relay path, for 96.6% of node pair we could find a 2-hop relay path that has higher available bandwidth than the corresponding direct path. When we select TCP's throughput-optimized relay path, introducing TCP proxy mechanism at relay nodes is a key issue for obtaining performance gain by overlay routing. We also find that the 3-hop relay path becomes effective especially when we deploy the multipath data transmission. Furthermore, latency-optimized relay path is likely to have larger available bandwidth than the direct path.

We next focused on the free-riding traffic problem caused by overlay routing.

The numerical metric of the degree of the free-riding problem was defined as the number of transit links increased by utilizing relay paths. For the degree of free-riding problem, the following result was confirmed: most of relay paths that have better performance than the corresponding direct path bring a sighificant degree of free-riding problem. It was also shown that available bandwidth was the best metric for mitigation of free-riding problem. We also found that when the number of free-riding transit link is restricted to 2, roughly 80% of performance gain compared with using optimized relay path was obtained.

For future work, we plan to investigate the control policy by ISPs for free-riding traffic problem, and propose new cost structure for ISPs in which ISPs can co-exist with routing overlay networks.

Acknowledgements

I would like to express my sincere appreciation to Professor Hirotaka Nakano of Osaka University for his supervision of this thesis. His guidance and inspiration have provided an invaluable experience that has helped me.

All works of this thesis would not have been possible without the support of Associate Professor Go Hasegawa of Osaka University. It gives me great pleasure to acknowledge his assistance. He has been constant sources of encouragement and advice through my studies and preparation of this manuscript. I can't express how grateful I am.

I would like to express my great thankfulness to Professor Masayuki Murata of Osaka University. All works of this thesis also would not have been possible without his appropriate guidance and invaluable firsthand advice.

I am most grateful to Professors Koso Murakami, Makoto Imase, Teruo Higashino of Osaka University, for their technical guidance and helpful advice.

I am also indebted to Associate Professor Naoki Wakamiya, and Research Assistants Shin'ichi Arakawa, Masahiro Sasabe, and Yuichi Oshita of Osaka University, who gave me useful comments and feedbacks.

Finally, I want to say thanks to all my friends and colleagues in the Department of Information Networking of the Graduate School of Information Science and Technology of Osaka University for their support. Our conversations and work together have greatly influenced this thesis.

References

[1] Z. Duan, Z.-L. Zhang, and Y. T. Hou, "Service overlay networks: SLAs, QoS, and bandwidth provisioning," in *Proceedings of IEEE ICNP 2002*, Nov. 2002.

[2] NetLightning Web Page. available at `http://www.netli.com/services/netlightning/`.

[3] Akamai Home Page. available at `http://www.akamai.com/`.

[4] 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.

[5] C. L. T. Man, G. Hasegawa, and M. Murata, "Inferring available bandwidth of overlay network paths based on inline network measurement," in *Proceedings of ICIMP 2007*, July 2007.

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

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

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

[9] 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.

[10] 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.

[11] 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 IEEE INFOCOM 2003*, Apr. 2003.

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

[13] Scalable Sensing Service. available at http://networking.hpl.hp.com/s-cube/.

[14] R. keralapura, N. taft, C. Chuah, and G. Iannaconne, “Can ISPs take the heat from overlay networks?,” in *Proceedings of ACM HotNets Workshop 2004*, Nov. 2004.

[15] T. Karagiannis, P. Rodriguez, and K. Papagiannaki, “Should internet service providers fear peer-assisted content distribution?,” in *Proceedings of ACM Internet Measuremet Conference 2005*, Oct. 2005.

[16] traceroute.org. available at http://www.traceroute.org/.

[17] Skype Home Page. available at http://www.skype.com/.

[18] YouTube Home Page. available at http://www.youtube.com/.

[19] GyaO Home Page. available at http://www.gyao.jp/.

[20] Bittorrent Home Page. available at http://www.bittorrent.com/.

[21] M. Mathis, J. Semke, and J. Mahdavi, “The macroscopic behavior of the TCP congestion avoidance algorithm,” *ACM SIGCOMM Computer Communication Review*, vol. 27, No.3, July 1997.

[22] I. Maki, G. Hasegawa, M. Murata, and T. Murase, “Performance analysis and improvement of TCP proxy mechanism in TCP overlay networks,” in *Proceedings of ICC 2005*, May 2005.

[23] AS-relationship data Page. available at http://www.caida.org/data/active/as-relationships/index.xml.

[24] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, K. claffy, and G. Riley, “AS relationships: Inference and validation,” *ACM SIGCOMM Computer Communication Review*, vol. 37, No.1, Jan. 2007.