

Evaluation of free-riding traffic problem in overlay routing and its mitigation method

Go Hasegawa
Cybermedia Center
Osaka University

1-32, Machikaneyama-cho, Toyonaka, Osaka 560-0043, JAPAN
Email: hasegawa@cmc.osaka-u.ac.jp

Yuichiro Hiraoka and Masayuki Murata
Graduate School of Information Science and Technology
Osaka University

1-3, Yamadaoka, Suita, Osaka 560-0871, JAPAN
Email: {y-hiraok, murata}@ist.osaka-u.ac.jp

Abstract—Recent research on overlay networks has revealed that user-perceived network performance could be improved by an overlay routing mechanism. The effectiveness of overlay routing is mainly a result of the policy mismatch between the overlay routing and the underlay IP routing operated by ISPs. However, this policy mismatch causes a “free-riding” traffic problem, which may become harmful to the cost structure of Internet Service Providers. In the present paper, we define the free-riding problem in the overlay routing and evaluate the degree of free-riding traffic to reveal the effect of the problem on ISPs. We introduce a numerical metric to evaluate the degree of the free-riding problem and confirm that most multihop overlay paths that have better performance than the direct path brings the free-riding problem. We also discuss the guidelines for selecting paths that are more effective than the direct path and that mitigate the free-riding problem.

Index Terms—Overlay networks, Overlay routing, Free-riding traffic, Available bandwidth, End-to-end delay, TCP throughput

I. INTRODUCTION

Overlay networks are the upper-layer logical networks where the endhosts and servers that run the applications become overlay nodes that form the logical network with logical links between the nodes. 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. For example, 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 [1] and Akamai [2] distribute overlay nodes (content servers) over the entire Internet and select appropriate source and destination 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 1. In Resilient Overlay Networks (RON) [3], 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 [3], the authors reported that RON can provide an effective traffic transmission path compared with lower-layer IP routing. Furthermore, RON




Fig. 1. Overlay routing and routing overlay network

can detect network failures (link and node failures, and mis-configured routing settings) and can provide an alternate route faster than IP routing convergence.

Several studies have examined the effectiveness of overlay routing with respect to IP routing [4-11]. For example, in [7], 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 [8], 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 to the IP routing and proactive overlay routing. In [11], the authors focused on bandwidth-related information, such as available bandwidth and TCP throughput, for metric in selecting overlay routes, and revealed that the bandwidth-based overlay routing provided significant gain, compared with latency-based routing.

The primary reason why overlay routing mechanisms can improve the 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 abstract than router-level IP network topology. Furthermore, most ISP-driven IP routings are configured by political and financial factors, namely, 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 studies on the problems of overlay network. In [12], the authors discuss the interaction between overlay routing and underlay IP routing, which causes routing and traffic oscillation. In [13], the authors discuss the effect of P2P-based content distribution on ISP's costs. We focus on other problem of overlay routing, whereby overlay routing can reduce the profits of ISPs that operate the lower-layer IP routing, and overlay routing mechanisms can generate network traffic that ignores an ISP's billing structure. In the present paper, we focus on this problem caused by overlay routing, which we call "free-riding" traffic. Note that the free-riding traffic problem in this paper is a general problem for overlay networks, regardless of the type of application and the algorithms and metrics for overlay routing. Furthermore, the free-riding problem can occur even when the routing interaction between overlay routing and IP routing is stable.

We define the free-riding problem in the routing overlay network and evaluate the degree of the 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 the 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 a smaller number of free-riding transit links than the optimized relay path. We finally report that we can obtain a reasonable performance gain by using the relay path with a limitation on the number of free-riding transit links.

The remainder of this paper is organized as follows. In Section II, we define the free-riding traffic problem and explain the dataset and metrics to evaluate problem in Section III. In Section IV we present the evaluation results to assess the degree of the free-riding problem. We also present guidelines for selecting overlay paths to mitigate the free-riding traffic problem. Finally, in Section V, we summarize the conclusions and discuss areas for future consideration.

II. FREE-RIDING TRAFFIC PROBLEM

The free-riding traffic problem is caused by 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 hop count, which may not be directly related to user-perceived performance. Furthermore, most IP routing is configured by political and financial factors, such as 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 that ignores an ISP's billing structure. We refer to such traffic as "free-riding" traffic, and this problem is called the free-riding traffic problem.

We explain the free-riding traffic problem by using Figure 2. In this figure, there are three ISPs (ISP I, J and K), where ISP




Fig. 2. Free-riding traffic problem

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 2. Therefore, the cost of 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 2), 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 of ISPs I and J benefit from the transmission. We refer to this mismatch as 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 to 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 the 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.

III. METHODOLOGIES

A. Dataset used for evaluation

We investigate the free-riding traffic problem under the assumption that PlanetLab[14] nodes construct a routing overlay network. For this 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 paper, we use the data obtained on Oct. 25th, 2006. Note that we have investigated with the datasets on other dates and obtained the similar results to those in this paper.

There exist 588 PlanetLab nodes in the measurement data utilized herein. However, a number of nodes are located in the same subnetwork, as estimated from the IP address and the 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.

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 [15] 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).

In addition, we use the information on the relationships between ASes, which is obtained from CAIDA [16]. In [17], the results on the relationship of links connecting two ASes (inter-AS links) are provided, as estimated from BGP table information and the degree of each AS (the number of links of one AS to other ASes). Methods for investigating AS relationship is reported in [18]. We obtain the AS-level route between PlanetLab nodes with the traceroute command between the nodes, and the AS-level traceroute results from a route server in traceroute.org to each node. However, 36.1% of the inter-AS links used by routes between PlanetLab nodes are not found in [17]. The reason for this may be that such inter-AS links do not exist in the BGP tables used in [18].

In determining the relationships of these unknown inter-AS links, we consider the following two cases: (i) All of unknown inter-AS links are assumed to be peering links, that they do not have any influence on the free-riding problem. (ii) The relationship of unknown inter-AS links are determined by the degree distribution of ASes revealed in [18]. In this paper, we only show the evaluation results of Case (i) due to space limitation. Note that the results of Case (ii) are similar to those in Case (i).

B. 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 1):

- Direct path: the source node to the destination node
- Two-hop relay path: the source node to the destination node via a relay node
- Three-hop relay path: the source node to the destination node via two relay nodes

C. Routing metrics

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

1) *End-to-end latency*: Overlay routing based on end-to-end latency would be adapted for applications, including voice chat applications such as Skype [19] 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 that the number of overlay nodes is M and that 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, the two-hop relay path, and the three-hop relay path, as follows:

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

We denote the relay node for the two-hop relay path as N_k and the relay nodes for the three-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 two-hop and three-hop latency-optimized paths as follows:

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

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

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

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

2) *Available bandwidth*: Available bandwidth is an important performance metric for audio video streaming services such as YouTube [20] and GyaO [21]. 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 available bandwidth of direct paths constructing the relay path. Denoting the 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, the two-hop relay path, and the three-hop relay path, as follows:

$$\begin{aligned} B_{ij}^1 &= \rho_{ij}, B_{ikj}^2 = \min(\rho_{ik}, \rho_{kj}) \\ B_{iklj}^3 &= \min(\rho_{ik}, \rho_{kl}, \rho_{lj}) \end{aligned} \quad (4)$$

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 available bandwidths of the two-hop and three-hop bandwidth-optimized paths as follows:

$$\begin{aligned} \hat{B}_{ij}^2 &= \max_{k \neq i, j} (B_{ikj}^2) \\ \hat{B}_{ij}^3 &= \max_{k \neq l, k, l \neq i, j} (B_{iklj}^3) \end{aligned} \quad (5)$$

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

$$I(B_{ikj}^2) = \frac{B_{ikj}^2}{B_{ij}^1}, I(B_{iklj}^3) = \frac{B_{iklj}^3}{B_{ij}^1}$$

3) *TCP throughput*: Overlay routing based on TCP throughput would be adapted for file sharing applications like Bittorrent [22] and WinMX. In [3], RON utilizes TCP throughput as a performance metric for overlay routing, which is calculated from the end-to-end delay and the packet loss ratio as follows:

$$T = \frac{\sqrt{1.5}}{RTT \sqrt{Loss}} \text{ (packet/sec)} \quad (7)$$

Equation (7) is based on the formula for the average throughput of the long-lived TCP connection in [23]. This metric increases as the packet loss ratio and RTT decrease, but never exceeds the available bandwidth of the path in the actual situation. In the present study, we use Eq. (7) 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 (AS): the packet loss ratio of the path is determined in proportion to the AS-level hop count of the path.
- (B) Overlay-hop-count-base loss ratio (OL): the packet loss ratio of the direct path is constant value regardless of the other characteristics of the path. The relay path has a packet loss ratio proportionally that is proportional to its overlay-level hop count.

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

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

$$P_{ij}^1 = \min \left(\frac{(8 \cdot MSS) \sqrt{1.5}}{D_{ij}^1 \sqrt{L_1}}, B_{ij}^1 \right) \text{ (bps)} \quad (8)$$

$$L_1 = \begin{cases} n_{ij} \cdot L_A & \text{(case(A))} \\ L_B & \text{(case(B))} \end{cases}$$

L_A and L_B are parameters that determine the packet loss ratio per AS-level hop and the packet loss ratio per overlay-level hop, respectively. In addition, we denote that the maximum segment size as MSS and the AS-level hop count between node N_i and node N_j as n_{ij} .

The definition of the TCP throughput of the relay path is different depending on whether the TCP connection is terminated at each relay node, which means that we utilize the TCP proxy mechanism [24] 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:

$$P_{ikj}^2(e2e) = \min \left(\frac{(8 \cdot MSS) \sqrt{1.5}}{D_{ikj}^2 \sqrt{L_2}}, B_{ikj}^2 \right) \quad (9)$$

$$P_{iklj}^3(e2e) = \min \left(\frac{(8 \cdot MSS) \sqrt{1.5}}{D_{iklj}^3 \sqrt{L_3}}, B_{iklj}^3 \right) \quad (10)$$

$$L_2 = \begin{cases} (n_{ik} + n_{kj}) \cdot L_A & \text{(case(A))} \\ 2L_B & \text{(case(B))} \end{cases}$$

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

$$P_{ikj}^2(pxy) = \min(P_{ik}^1, P_{kj}^1) \quad (11)$$

$$P_{iklj}^3(pxy) = \min(P_{ik}^1, P_{kl}^1, P_{lj}^1) \quad (12)$$

$$L_3 = \begin{cases} (n_{ik} + n_{kl} + n_{lj}) \cdot L_A & \text{(case(A))} \\ 3L_B & \text{(case(B))} \end{cases}$$

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

$$\hat{P}_{ij}^2(e2e) = \max_{k \neq i, j} (P_{ikj}^2(e2e)) \quad (13)$$

$$\hat{P}_{ij}^3(e2e) = \max_{k \neq l, k, l \neq i, j} (P_{iklj}^3(e2e)) \quad (14)$$

$$\hat{P}_{ij}^2(pxy) = \max_{k \neq i, j} (P_{ikj}^2(pxy)) \quad (15)$$

$$\hat{P}_{ij}^3(pxy) = \max_{k \neq l, k, l \neq i, j} (P_{iklj}^3(pxy)) \quad (16)$$

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

$$I(P_{ikj}^2(e2e)) = \frac{P_{ikj}^2(e2e)}{P_{ij}^1}, I(P_{iklj}^3(e2e)) = \frac{P_{iklj}^3(e2e)}{P_{ij}^1}$$

$$I(P_{ikj}^2(pxy)) = \frac{P_{ikj}^2(pxy)}{P_{ij}^1}, I(P_{iklj}^3(pxy)) = \frac{P_{iklj}^3(pxy)}{P_{ij}^1}$$

D. Evaluation metric

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

$$F_{ikj} = \{x | (x \in T_{ikj}) \& (x \notin T_{ij})\} \quad (17)$$

In the present study 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 .

IV. EVALUATION RESULTS

In this section, we evaluate the free-riding traffic problem, and discuss guidelines to mitigate this problem. In the following evaluation, MSS is 1460 bytes, and L is set to 0.00001, 0.0001, 0.001, and 0.01. We only show the evaluation results of Case (i) due to space limitation. Note that the results of Case (ii) are similar to those in Case (i).


A. Number of free-riding transit links

In Figure 3, we show the distribution of $|F_{ikj}|$ of the optimized two-hop relay path for each metric in Case (i), which means that we ignore unknown inter-AS links. Due to space limitation, we only show the AS/e2e and OL/e2e case. The average of $|F_{ikj}|$ is shown in Table I for all metrics.


Based on these results, when we use either the available bandwidth or end-to-end latency, the free-riding traffic problem occurs on roughly 90% of node pairs. In the case of using the TCP throughput, the results depend on the packet

TABLE I
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



(a) AS/e2e



(b) OL/e2e

Fig. 3. Distribution of the number of free-riding transit links in case (i)

loss model. When we use the AS-hop-count-base packet loss model, the degree of the free-riding problem decreases significantly because the throughput-optimized relay path tends to have a small AS-level hop count, resulting in a decrease in the number of transit links used by such a relay path. 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 when using the available bandwidth because, in this case, the relay path is selected regardless of the AS-level hop count, which is the same situation as for the available bandwidth.

Figures 3 reveals that most of the optimized relay paths have a number of free-riding transit links.

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

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

$$O(D_{ij}^2) = \{k | (|F_{ikj}| < |F_{i\hat{k}j}|) \& (I(D_{ikj}^2) > 1)\} \quad (18)$$

$$O(B_{ij}^2) = \{k | (|F_{ikj}| < |F_{i\hat{k}j}|) \& (I(B_{ikj}^2) > 1)\} \quad (19)$$




Fig. 4. Distribution of improvement ratios of the latency-optimized, latency-second-best, and latency-third-best paths

$$O(P_{ij}^2) = \{k | (|F_{ikj}| < |F_{i\hat{k}j}|) \& (I(P_{ikj}^2) > 1)\} \quad (20)$$

Figure 5 shows 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 5, in Case (i), there are several nodes in $O(B_{ij}^2)$ when we use the available bandwidth. When we use the TCP throughput in the AS/e2e case, we cannot find any relay node in $O(P_{ij}^2)$ for approximately 90% of node pairs. However, this percentage decreased when we use the TCP throughput in OL or pxy case. However, when we use end-to-end latency, over 90% of node pairs, we cannot find the relay node in $O(D_{ij}^2)$, even if the two-hop latency-optimized relay path improves the performance in 87.5% of node pair as compared to the corresponding direct path. The reason for this result can be explained using the distribution in Figure 4.

Figure 4 shows the distribution of the improvement ratio of latency-optimized, latency-second-best, and latency-third-best paths. This figure shows that the improvement ratio of the second- and third-best paths decreases greatly. Therefore, when we search the relay node in $O(D_{ij}^2)$, we cannot find a relay path with better performance than the corresponding direct path.

C. Mitigation of number of free-riding transit links

Based on the above results, we conclude that most of the relay paths that have better performance than the corresponding direct path have a significant 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 6 we show the relationships between the maximum number of free-riding transit links permitted in selecting overlay paths and the improvement ratio of the optimized two-hop relay path for each node pair. Figure 6 shows that the improvement ratio sharply increases as the limit of $|F_{ikj}|$ increases to 4, and remains approximately unchanged when the limit of $|F_{ikj}|$ exceeds 4. When we focus on available bandwidth, the improvement ratio exceeds 80% when the limit of $|F_{ikj}|$ is 2 and exceeds 90% when the limit of $|F_{ikj}|$ is 3. When we consider end-to-end latency,

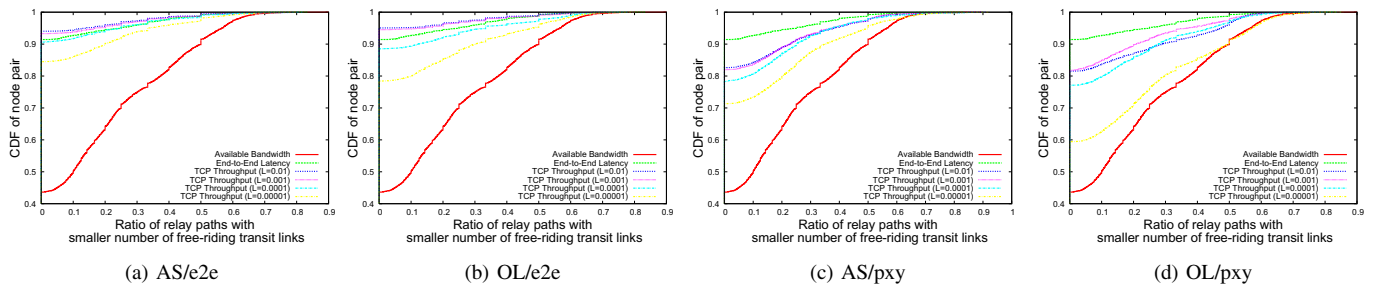


Fig. 5. Ratio of relay path with smaller number of free-riding transit links in case (i)

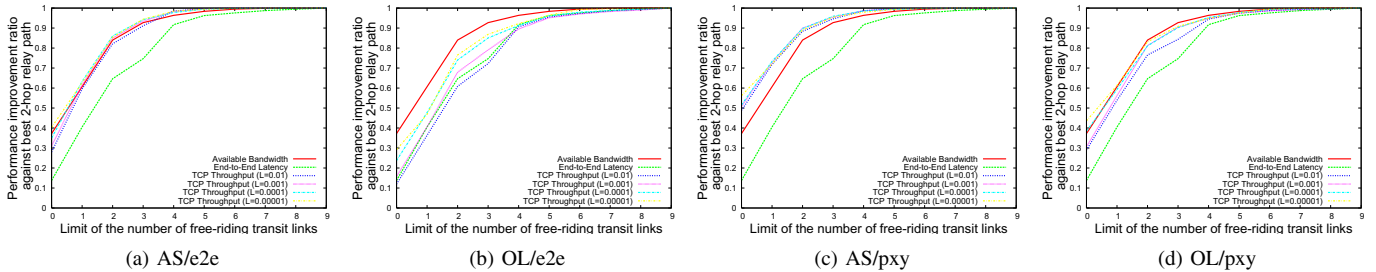


Fig. 6. Correlation between improvement ratio and number of free-riding transit links in case (i)

these percentages decrease to approximately 65% and 70%, respectively. Furthermore, when we consider TCP throughput, the TCP throughput in the AS case is better than that in the OL case, and the TCP throughput in the e2e case is better than that in the pxy case.

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

V. CONCLUSION

In this paper, we 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 additional transit links when using relay paths. We have revealed that most relay paths that have better performance than the corresponding direct path have a significant degree of free-riding. It was also shown that available bandwidth was the best metric for mitigation of the free-riding problem. We also found that when the number of free-riding transit links is restricted to 2, roughly 80% of the performance gain, as compared with using the optimized relay path was obtained.

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

REFERENCES

- [1] NetLightning Web Page. available at <http://www.netli.com/services/netlightning/>.
- [2] Akamai Home Page. available at <http://www.akamai.com/>.
- [3] 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.
- [4] 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.
- [5] G. Hasegawa, M. Kobayashi, M. Murata, and T. Murase, "Free-riding traffic problem in routing overlay network," in *Proceedings of ICON 2007*, Nov. 2007.

- [6] M. Uchida, S. Kamei, and R. Kawahara, "Performance evaluation of QoS-aware routing in overlay network," in *Proceedings of ICOIN 2006*, Jan. 2006.
- [7] S. Kamei, "Applicability of overlay routing in Japan using inter-domain measurement data," *Overlay Network Workshop*, Dec. 2006.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] Y. Hiraoka, G. Hasegawa, and M. Murata, "Effectiveness of overlay routing based on delay and bandwidth information," in *Proceedings of ATNAC 2007*, Dec. 2007.
- [12] R. keralapura, N. taft, C. Chuah, and G. Iannacone, "Can ISPs take the heat from overlay networks?," in *Proceedings of ACM HotNets Workshop 2004*, Nov. 2004.
- [13] T. Karagiannis, P. Rodriguez, and K. Papagiannaki, "Should internet service providers fear peer-assisted content distribution?," in *Proceedings of ACM Internet Measurement Conference 2005*, Oct. 2005.
- [14] PlanetLab Web Page. available at <http://www.planet-lab.org/>.
- [15] traceroute.org. available at <http://www.traceroute.org/>.
- [16] CAIDA Web Page. available at <http://www.caida.org/home/>.
- [17] AS-relationship data Page. available at <http://www.caida.org/data/active/as-relationships/index.xml>.
- [18] 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.
- [19] Skype Home Page. available at <http://www.skype.com/>.
- [20] YouTube Home Page. available at <http://www.youtube.com/>.
- [21] GyaO Home Page. available at <http://www.gyao.jp/>.
- [22] Bittorrent Home Page. available at <http://www.bittorrent.com/>.
- [23] 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.
- [24] 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.