“Free-riding” Traffic Problem in Routing Overlay Networks

Go Hasegawa*, Masayoshi Kobayashi†, Masayuki Murata‡ and Tutomu Murase†
*Cybermedia Center, Osaka University
1-32, Machikaneyama-cho, Toyonaka, Osaka 560-0043, JAPAN
†System Platforms Research Laboratory, NEC Corporation
1753 Shimonumabe, Nakahara-ku, Kawasaki, Kanagawa, 211-866 JAPAN
‡Graduate School of Information Science and Technology, Osaka University
1-3, Yamadaoka, Suita, Osaka 560-0871, JAPAN

Email: hasegawa@cmc.osaka-u.ac.jp

Abstract—**In this paper, we discuss the “free-riding” traffic problem in routing overlay networks, which is mainly caused by a policy mismatch between the overlay routing and the underlay IP routing. We first define the free-riding problem in a routing overlay network, and construct mathematical models to calculate the amount of the free-riding traffic in routing overlay networks with various path-selection metrics. We also present the numerical examples to estimate the effects of the overlay routing. We conclude that there can be a significant amount of free-riding traffic in an actual Internet environment.**

I. 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], they include P2P networks, anonymous file-sharing services such as WinMX and Share, audio and video conferencing service such as Skype, and Content Delivery/Distribution Networks (CDNs). Service overlay networks are upper-layer networks providing special-purpose services built onto the lower-layer IP network. Therefore, their performance depends primarily on how well they take advantage of the characteristics and resources of the underlying network. To improve performance, service overlay networks need fast and accurate information concerning the resource availability in the IP network to realize adaptive control mechanisms. Some examples of these control mechanisms are as follows:

- P2P networks. When a resource discovery mechanism finds multiple peers having the same requested contents, this information is used to determine which peer should transmit the contents.
- Grid networks. When multiple sites contain the same data, this information is used to determine from which site data will be copied or read.
- CDNs. When backup data or cached data is transmitted, this information can be used to prevent other network traffic from being deprived of resources during the transmission.

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.

Some of the overlay networks construct a logical network and select a route for data transmission according to network conditions, such as the following: link speed, delay, packet loss ratio, hop count, and TCP throughput between overlay nodes. In WinMX, an endhost can report the kind 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 destination hosts according to the network condition when the contents would be moved, duplicated or cached. Some overlay networks do not assume specific upper-layer applications and concentrate only on the routing of overlay network traffic. In Resilient Overlay Networks (RON) [4], each overlay node measures the end-to-end delay and packet loss ratio of the network path between the node and 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 relayed route which passes through another node(s) before reaching the destination node. Thus, overlay routing can provide more effective traffic transmission compared to lower-layer IP routing. Furthermore, it can detect network failures (link and node failures, and mis-configured routing settings) and provide an alternate route in faster time than the IP routing convergence.

Previous papers such as in [5-7] show that overlay routing mechanisms can improve throughput and transmission time in data transfer. This is because the traditional IP routing operated by Internet Service Providers (ISPs) does not always determine the route according to user-perceived performance. In intradomain IP routing, the metrics determining the route are hop count and link loads, not end-to-end bandwidth-related information, which affects the data transmission throughput for long-lived flows. Furthermore, inter-domain routing is based on autonomous system-level (AS-level) topology and hop count, which are more abstract than router-level IP network topology. Furthermore, most ISP-driven 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. Therefore, the resulting IP routing policy cannot maximize the network performance and users' demands.

However, we believe that there are some situations where overlay routing can harm the profit of the ISPs which operate the lower-layer IP routing. These situations are caused by the difference in the billing mechanism of transit links and peering links. 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. The IP routing operated by ISPs takes the difference into account, and the peering link between two ISPs is used only for traffic which is either from or to the two ISPs. However, application-level overlay routing does not consider the billing structure of ISPs and determines the route according only to the user-perceived performance. As the result, overlay routing mechanisms can generate network traffic which ignores an ISP's billing structure. In this paper, we focus on one problem caused by overlay routing, which we call "free-riding" traffic.

There are some previous works on problems of overlay network. In [8], the authors discuss the interaction between overlay routing and underlay IP routing, that causes routing and traffic oscillation. In [9], the effect of P2P-based content distribution on ISP's costs. On the other hand, the "free-riding" traffic problem in this paper is a general problem for overlay networks, regardless of the kind of application. Furthermore, it can occur even when the routing interaction between overlay routing and IP routing is stable.

In this paper, we first introduce the network model and define the problem in this paper. We also introduce some overlay routing mechanisms which utilize the round-trip delay and available bandwidth of the network path between the overlay nodes as a performance metric. We then formulate the amount of relayed traffic, which defines as the traffic conveyed by relayed paths on the routing overlay networks. Next, we apply the formulation results to the PlanetLab [10] network and calculate the amount of relayed traffic when the PlanetLab nodes construct the routing overlay network. We finally estimate the amount of free-riding traffic from that of the relayed traffic and show that the problem of free-riding traffic cannot be ignored by ISPs in the current and future Internet environments.

II. "Free-riding" traffic problem caused by overlay routing networks

As we described in the previous section, overlay routing can improve user-perceived performance such as data transmission throughput and delay. However, there are some situations where overlay routing can have a negative effect on ISPs. These situations are caused by the difference between the billing structure of transit links and peering links. Generally, the lower ISP pays the monetary cost for usage of the transit links connected to the upper ISPs. On the other hand, two ISPs

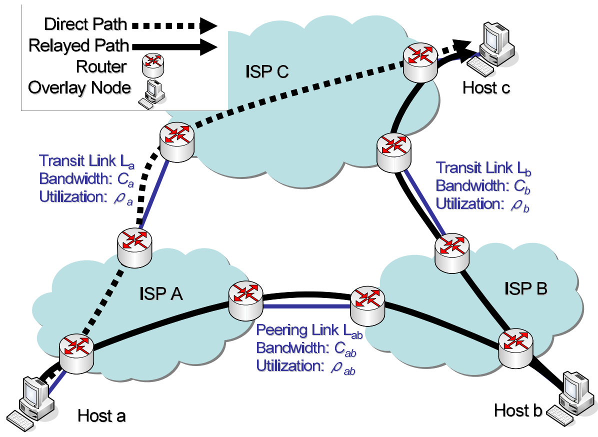

Fig. 1. Network model

interconnected by a peering link share the link's running cost. Although the IP routing operated by ISPs takes the difference into consideration, overlay routing does not cover the cost difference; instead, overlay routing just routes the overlay network traffic to improve the user-perceived performance in data transmission. Consequently, the network traffic routed by routing overlay mechanisms may violate the ISP's cost structure. We refer to this problem as "free-riding" traffic, which is caused by overlay routing.

We explain the free-riding traffic problem in this paper by using Figure 1. In this figure there are three ISPs (ISP A, B and C), where ISP C is the transit ISP for ISPs A and B. ISP A and B have transit links L_a and L_b to connect to ISP C. Furthermore, ISP A and B are interconnected by peering link L_{ab}. Hosts a, b, and c exist in ISP A, B, and C, respectively. These three hosts are the overlay node of the routing overlay network.

We consider a situation where Host a transmits overlay network traffic to Host c by using the routing overlay network. There are two paths for transmitting the traffic:

- Direct path: the path from Host a to Host c without any relaying host
- Relayed path: the path from Host a to Host c via Host b

If we use the direct path, the traffic is transmitted by transit link L_a. Therefore, the cost of conveying the traffic is charged to ISP A. However, if we use the relayed path, the traffic is transmitted by peering link L_{ab} from Host a to Host b, and the transit link L_b from Host b to Host c. In this case, ISP B pays the cost for using L_b to convey the traffic, although only the customers in ISPs A and C benefit from the transmission. We call this mis-match the "free-riding " traffic problem, and the relayed path as the free-riding path. Since more and more ISPs are making peering relationships in current ISP networks, we believe that this problem has a serious effect on an ISP's cost structure.

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

One may think that the cost to a transit link may be balanced out if we consider bi-directional traffic. However, the monetary cost of the transit link of each ISP is usually different, and it is dependent on the total amount of traffic of each ISP. That is, even when the amount of free-ride traffic is equal, the additional monetary cost for the ISPs becomes different, that causes serious problem especially for the small ISPs.

In the current Internet, there is not a lot of overlay-routed traffic. For example, in [11] the authors ignore the amount of such traffic when calculating the traffic matrix for large-scale IP networks. In the future, however, when routing overlay networks such as RON become popular and more and more network applications utilize them by improving their service quality, we believe that the free-riding traffic problem will become a serious issue for ISPs. That is, when an ISP constructs a "better-quality" network in terms of link bandwidth, end-to-end delay, and packet loss ratio, it will induce a larger amount of overlay-routed traffic to pass through the ISP's network, but the ISP will not be able to collect the cost for conveying such traffic.

III. Formulation of conditions for overlay traffic

In this section, we formulate the condition in which overlay traffic is conveyed by the relayed paths and the amount of traffic on these relayed path. Note that the problem formulation in this section is quite simple, in order to emphasize the free-ride traffic problem defined in this paper. We believe that this simple model can make clear the problem.

We utilize the network model depicted in Figure 1. The link bandwidth of links L_a, L_b, and L_{ab} is C_a, C_b, and C_{ab}, and the current utilization is ρ_a, ρ_b, and ρ_{ab}, respectively. We also use the notation of the available bandwidth of each link, A_a, A_b, and A_{ab}, where $A_a = C_a(1-\rho_a)$, $A_b = C_b(1-\rho_b)$, and $A_{ab} = C_{ab}(1-\rho_{ab})$. We assume that the bandwidth inside each network of ISP A, B, and C is large enough to not limit the end-to-end bandwidth.

We consider the following two routes for conveying the overlay network traffic from Host a to Host b:

- Direct path: Host a → ISP A → L_a → ISP C → Host c
- Relayed path: Host a → ISP A → L_{ab} → ISP B → Host b → ISP B → L_b → ISP C → Host c

In the following, we consider that the total amount of overlay network traffic to be x (bps), and calculate the amount of overlay network traffic conveyed by the direct path and that by the relayed path, which we denoted as x_d (bps) and x_r (bps), respectively.

Note that all of the traffic on the relayed path is not the free-riding traffic defined in Section II. We discuss the relationship between relayed traffic and free-riding traffic in Subsection IV-C.

A. Selecting overlay paths according to the ratio of available bandwidth

We first consider the case when the overlay routing algorithm selects the route according to the ratio of available bandwidth of the direct path and the relayed path. We can easily derive x_d and x_r (bps) as follows:

$$x_d = \frac{A_a}{A_a + \min(A_{ab}, A_b)} x \quad (1)$$

$$x_r = \frac{\min(A_{ab}, A_b)}{A_a + \min(A_{ab}, A_b)} x \quad (2)$$

We assume that ISP B increase the bandwidth of its transit link L_b from C_b to C'_b. Then the link utilization changes from ρ_b to ρ'_b, where ρ'_b is the link utilization after the increase of the link bandwidth and where $C_b\rho_b = C'_b\rho'_b$ (we assume the amount of traffic on link L_b remains unchanged just after the increase of the link bandwidth).

When $A_{ab} < A_b$, that is, when the bandwidth of L_{ab} is not very large, the amount of traffic on the relayed path does not change by the bandwidth increase. When $A_{ab} > A_b$, on the other hand, the bandwidth enhancement increases the amount of traffic on the relayed path. We denote the amount of traffic on the direct path and that on the relayed path after the bandwidth increase as x'_d and x'_r, respectively. We then have

$$x'_d = \frac{A_a}{A_a + A'_b} x, x'_r = \frac{A'_b}{A_a + A'_b} x \quad (3)$$

where $A'_b = C'_b(1-\rho'_b)$. So, we can calculate the amount of increased traffic on the relayed path by the bandwidth increase, denoted as Δx_r, from Equations (1), (2) and (3) as follows:

$$\Delta x_r = x'_r - x_r = \frac{A_a(C'_b - C_b)}{(A_a + C'_b(1-\rho'_b))(A_a + A_b)} x$$

We also have the following results under the condition where the peering link bandwidth (C_{ab}) is enough large:

$$\lim_{C'_b \to \infty} \Delta x_r = \frac{A_a}{(A_a + A_b)} x = x_d$$

This equation means that when ISP B greatly increase the bandwidth of its transit link, ISP B absorbs almost all of the overlay network traffic in the free-riding path. In other words, the overlay network traffic will go to the ISPs with lower utilization of network links.

B. Selecting the overlay path with larger available bandwidth

We next consider the situation where the path with the larger available bandwidth is selected. We assume that the traffic is equally divided when both paths have the same available bandwidth.

When $x \leq |A_a - \min(A_{ab}, A_b))|$, that is, when the difference of the available bandwidth is larger than the amount of the overlay network traffic, all of the overlay network traffic is conveyed by the path with the larger available bandwidth. That is, x_r becomes

$$x_r = \begin{cases} 0 & (A_a > \min(A_{ab}, A_b)) \\ \frac{x}{2} & (A_a = \min(A_{ab}, A_b)) \\ x & (A_a < \min(A_{ab}, A_b)) \end{cases} \tag{4}$$

Note that the amount of traffic on the relayed path remains unchanged even when ISP B increase the transit link bandwidth.

On the other hand, when $x > |A_a - \min(A_{ab}, A_b))|$, we can calculate x_r as follows:

$$x_r = \frac{1}{2}(x - |A_a - \min(A_{ab}, A_b)|)$$

We assume that ISP B increases the bandwidth of its transit link L_b from C_b to C'_b. We use the same notation stated in Subsection III-A for ρ'_b and x'_r. When the peering link bandwidth, C_{ab} is not very large, the amount of traffic on the relayed path remains unchanged. So, we have

$$x'_r = x_r = \frac{1}{2}(x - (A_a - A_{ab})) \tag{5}$$

On the other hand, if C_{ab} is large enough, the amount of traffic on the relayed path increases when ISP B increase the transit link bandwidth. That is,

$$x_r = \frac{1}{2}(x - (A_a - A_b)), x'_r = \frac{1}{2}(x - (A_a - A'_b))$$

where we utilize the relation $A_a < A'_b = C'_b(1 - \rho'_b)$. Then we can obtain the amount of increased traffic on the relayed path, $\Delta x_r = x'_r - x_r$ as follows:

$$\Delta x_r = \frac{1}{2}(C'_b - C_b) \tag{6}$$

where we assume $C_b\rho_b = C'_b\rho'_b$ and $x > |\min(A_a - \min(A_{ab}, A'_b)|$. From Equation (6), we can conclude that when ISP B increase the transit link bandwidth, half of the increased bandwidth is "stolen" by the relayed traffic by overlay routing.

C. Selecting the overlay path with smaller round-trip delay

We last consider the overlay routing algorithm which selects the route with a smaller round-trip delay. We assume in the analysis here that the transit link between ISPs becomes the bottleneck causing queueing delay. The case when the bottleneck exists at user-side link is explained in the last part of this subsection.

We define the round-trip propagation delay of the direct path as τ_d, and that of the relayed path as τ_r. We assume that the processing delay at the routers on the path is included in τ_d and τ_r. We assume $\tau_d < \tau_r$, that is, the round-trip propagation delay of the relayed path is larger than that of the direct path. The processing delay at the relayed node (Host b) is assumed to be included in τ_r. We also denote the bottleneck link bandwidths of both paths as μ_d and μ_r, respectively. For simplicity, we use the M/M/1 queuing model for deriving the average queuing delay at the output buffer on the bottleneck link.

The round trip time (RTT) of the direct path and that of the relayed path can be calculated as follows:

$$RTT_{d,r} = \tau_{d,r} + \frac{1}{\mu_{d,r}(1 - \rho_{d,r})} \tag{7}$$

The condition in which the relayed path is selected is $RTT_d > RTT_r$. From Equation (7), we have

$$\tau_r - \tau_d \quad < \quad \frac{1}{\mu_d(1 - \rho_d)} - \frac{1}{\mu_r(1 - \rho_r)}$$

That is, the overlay routing algorithm in this subsection uses the relayed path until the difference in the queuing delay becomes equal to the difference in the round-trip propagation delay, $\Delta\tau$ $(= \tau_r - \tau_d)$. Here we denote the difference in the utilization of both paths when the RTTs of both paths are equal as $\Delta\rho$. By assuming $\mu_d = \mu_r = \mu$, we have

$$\Delta\rho = \frac{\Delta\tau \cdot \mu(1 - \rho_d)^2}{1 - \Delta\tau \cdot \mu(1 - \rho_d)} \tag{8}$$

From Equation (8), we can observe that when ρ_d approaches 1, $\Delta\rho$ approaches 0. This means that, when the overall network load becomes high, the overlay routing algorithm utilizes both paths so that the utilization of both paths becomes equal. In other words, the larger the amount of network traffic becomes, the larger amount of overlay traffic is conveyed by the relayed path.

We also note that if μ_r becomes large as the results of the increase of the transit link bandwidth by ISP B, RTT_r becomes small due to the decrease of the queuing delay of the relayed path. This means that some of the overlay traffic on the direct path would move to the relayed path. Therefore, the performance of the normal traffic from the endhosts in ISP B does not improve much because of the movement of the overlay network traffic.

Next, we consider the case when the bottleneck of both paths are located at the link near Host c, meaning that both paths share the same bottleneck link. So we have $\rho_d = \rho_r$ regardless of the link bandwidth and the utilization of other part of the paths. When $\tau_d < \tau_r$, we always have $RTT_d < RTT_r$, so all of the overlay network traffic is conveyed on the direct path. However, when $\tau_d > \tau_r$, that is, when we find the relayed path with smaller propagation delay than the direct path, $RTT_d > RTT_r$ is always satisfied and all of the overlay network traffic is transmitted on the relayed path. This phenomenon cannot be found in the path selection based on the available bandwidth, discussed in Subsections III-A and III-B. From Equations (1), (2), (4) and (5), we calculate that 50% of the overlay network traffic would use the relayed path when both paths have the same available bandwidth.

IV. NUMERICAL EXAMPLES WITH PLANETLAB DATA

In Scalable Sensing Service [12], which is one of the research projects on PlanetLab, full-mesh measurement results such as RTT, available bandwidth, physical capacity and packet loss ratio between PlanetLab nodes are provided at the Web site. In this section, we apply the measurement data to the formulation results in the previous section and calculate the ratio of traffic on the relayed paths when the PlanetLab nodes construct the routing overlay network. We discuss the relationships between the relayed traffic and free-riding traffic defined in Section II, and show that most of relayed traffic is not welcomed by the ISPs.

The results here are not surprising, as previous work has already demonstrated that relay nodes can offer viable paths. By presenting this results in this paper, we provide some clear data to justify our claim that ISPs with relay nodes may bear an unfair burden for the extra traffic they will carry.

A. Effect of overlay routing in the PlanetLab network

We first show the calculation results on the effect of overlay routing in the PlanetLab network. For each pair of PlanetLab nodes, we find the relayed node which gives the largest available bandwidth, and another node which gives the lowest RTT, of the relayed path. We denote the relayed path giving the largest available bandwidth or lowest RTT as the "best relayed path". Figure 2(a) shows the relationships between the available bandwidth of the direct path and that of the best relayed path. Figure 2(b) shows the relationships between the RTT of the direct path and that of the best relayed path. From this figure, we observe that for almost all of the node pairs (96.2% in the available bandwidth case and 96.3% in the RTT case) we can find the "better" relayed path. We find that for each node pair, when we select a relayed node randomly, the average available bandwidth of the relayed paths is larger than that of the direct path for 27.4% of the node pairs, but only 0.05% for the RTT case. This is because IP routing in the current Internet is generally based on hop count, which correlates in some degree with the end-to-end propagation delay.

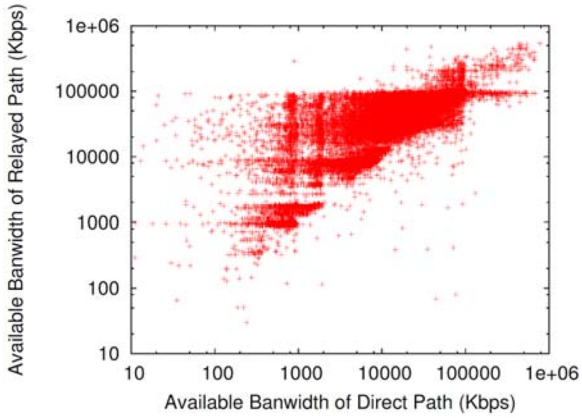

(a) Metric: Available bandwidth

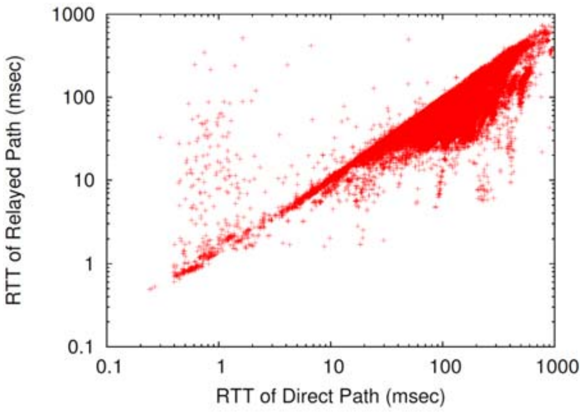

(b) Metric: Round Trip Time

Fig. 2. Relationships between direct path and best relayed path

B. Calculation of the amount of traffic on relayed paths

We next consider that the nodes in PlanetLab construct the routing overlay network and calculate the ratio of overlay network traffic conveyed on relayed paths. We use the available bandwidth as the metric, as explained in Subsections III-A and III-B, and use the following assumptions and settings:

- Overlay network traffic is generated equally for each node pair.
- We do not consider the effect of network topology. That is, the available bandwidth of the path between a certain node pair is not affected by the overlay network traffic generated on the paths between other node pairs.
- We use the following three methods to select a relayed path:
  - Best relayed path: Uses the relayed path with the largest available bandwidth
  - Good relayed path: Selects the relayed path randomly from the relayed paths with larger available bandwidth
  - All relayed path: Selects the relayed path randomly from all of the relayed paths
- We use the following two methods for distributing the overlay network traffic to the direct path and the relayed path selected by the above methods:
  - Ratio of available bandwidth of the direct path and the relayed path (Subsection III-A)
  - Selection of the path with the larger available bandwidth (Subsection III-B)

Table I presents the ratio of overlay network traffic on relayed paths in each case. From this table, we can observe that, regardless of the path selection methods and traffic distribution methods, a significant amount of overlay network traffic is conveyed on relayed paths.

TABLE I
RATIO OF OVERLAY NETWORK TRAFFIC ON RELAYED PATHS

	Ratio of available bandwidth	Path with larger available bandwidth
Best relayed path	72.8%	96.2%
Good relayed path	58.5%	96.2%
All relayed path	49.2%	22.6%

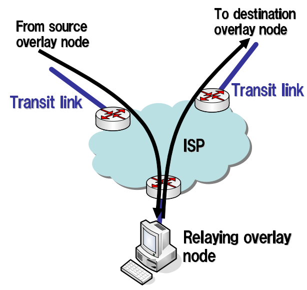

(a) From: transit link, To: transit link

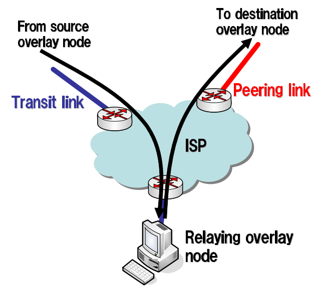

(b) From: transit link, To: peering link

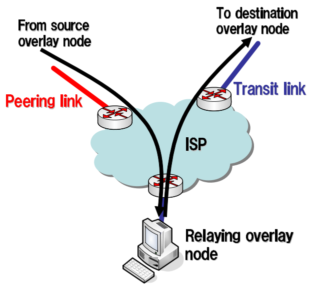

(c) From: peering link, To: transit link

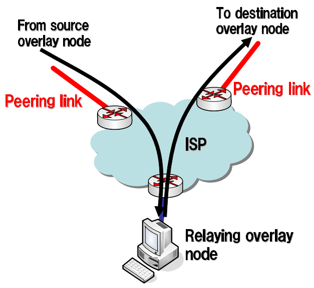

(d) From: peering link, To: peering link

Fig. 3. Four types of relayed path

C. Relationships between relayed traffic and free-riding traffic

In the previous subsection, when the PlanetLab nodes construct the routing overlay network, we showed that a significant amount of overlay network traffic is conveyed on relayed paths. However, we note that all of the traffic on the relayed paths does not correspond to the free-riding traffic defined in this paper. This is because we cannot identify the link type (transit link or peering link) between the PlanetLab nodes.

In general, the ISP to which the relayed host belongs has multiple transit and peering links for connecting to other ISPs. Figure 3 shows the four types of relayed overlay path, categorized by the type of ISP's links incoming from and outgoing to the relayed overlay node. Since the ISP must take the traffic cost passing through transit links, cases in Figures 3(a)–(c) generates free-ride traffic. In the case in Figures 3(d), the relayed path includes only peering links. So the path is not a free-riding path. However, the relayed traffic consumes some amount of resources in the ISP network, so the ISP does not welcome such relayed traffic passing through its network.

V. CONCLUSION

In this paper, we discussed the free-riding traffic problem, which can be caused by routing overlay networks because of the mismatch of policies between overlay routing and IP routing. We then formulated the amount of overlay network traffic on relayed paths by various overlay routing metrics. We applied PlanetLab measurement data to the formulation results and verified that we can have a significant amount of overlay network traffic conveyed on relayed paths, most of which corresponds to the free-riding path.

For future work, we plan to extend the formulation in Section III to consider the effect of network topologies. We also plan to propose methods to estimate the amount of free-riding traffic based on the passive measurements in an ISP network. We are also interested in building a new cost structure for ISPs to accommodate the effect of overlay-routed network traffic.

ACKNOWLEDGEMENT

This work is partly supported by the Strategic Information and Communications R&D Promotion Programme (SCOPE) from the Ministry of Internal Affairs and Communications of Japan.

REFERENCES

[1] T. H. Zhenhai Duan, Zhi-Li Zhang, "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 Web 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] 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.
[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] 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, pp. 91–100.
[8] 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.
[9] 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.
[10] PlanetLab Web Page, available at http://www.planet-lab.org/.
[11] Y. Zhang, M. Roughan, N. Duffield, and A. Greenberg, "Fast, accurate computation of largescale IP traffic matrices from link loads," in *Proceedings of SIGMETRICS 2003*, June 2003.
[12] Scalable Sensing Service, available at http://networking.hpl.hp.com/s-cube/.