Towards Interdomain Transit Traffic Reduction in Peer-assisted Content Delivery Networks

Hirochika Asai
The University of Tokyo
panda@hongo.wide.ad.jp

Hiroshi Esaki
The University of Tokyo
hiroshi@wide.ad.jp

***Abstract*—Transit traffic exchanged with transit providers costs more for network providers compared to intra-domain traffic or traffic exchanged over peering links. We have measured and analyzed the peer distribution in *BitTorrent*, which is one of peer-assisted content delivery networks (CDNs). From the peer distribution, we show the potential of the high-cost transit traffic reduction. We then propose a peer selection preference which takes into account the economical relationships among Autonomous Systems (ASes) in peer-assisted CDNs to reduce the high-cost transit traffic. Since most commercial Internet service providers cannot disclose the relationships due to their commercial contract, we employ degree-based heuristics for inferring the relationships; degree can be approximated from publicly available BGP routing tables. We show that the peer selection method utilizing the proposed preference can reduce interdomain transit traffic exchanged with provider ASes by trace-driven computer simulation. The significance of this paper are 1) we show the potential of the transit traffic reduction from peer distribution analysis, and 2) the peer selection method with the proposed preference appropriately reduces the high-cost transit traffic with degree-based AS relationships inference heuristics even though there is no public AS relationships information.**

***Index Terms*—content delivery network, transit traffic, Internet economics**

I. Introduction

Content Delivery Networks (CDNs) are widely employed to distribute large content files such as music files, movie files and software images. Various peer-assisted CDNs such as *BitTorrent* [1] have been developed and deployed to avoid excessive server load and to achieve effective and high-quality content delivery over the Internet. Since the topology of peer-assisted CDNs is generally different from the layer 3 network topology and it is hard to take into account the interdomain routing policies and economics, these CDNs frequently utilize a larger amount of network resources and cost more from the layer 3 operators' viewpoint [2], [3], [4].

The Internet consists of thousands of distinct administrative domains called Autonomous Systems (ASes). Customer ASes purchase Internet access over transit links from provider ASes by paying some amount of money according to their actual bandwidth usage [5]. Therefore, transit traffic exchanged with providers costs more for network providers compared to intra-domain traffic or traffic exchanged over peering links. In the interdomain routing, ASes control the transit traffic by a typical best path selection policy. However, in peer-assisted CDNs, peers are distributed into many distinct ASes, and ASes cannot control the peer selection because it is controlled by peers. So, peer-assisted CDNs often utilize high-cost transit links despite the existence of lower-cost delivery paths from the network providers' viewpoint. Hence, the peer selection control and the transit traffic reduction in peer-assisted CDNs is required.

We have measured and analyzed peer distribution in *BitTorrent* [1], which is one of peer-assisted CDNs. From the analysis, we show the potential of the transit traffic reduction in peer-assisted CDNs. We then propose an AS relationships-aware peer selection preference to control and reduce the high-cost transit traffic. Since most commercial Internet service providers (ISPs) cannot disclose the relationships due to their commercial contract, we employ degree-based heuristics for inferring the relationships. Here, degree can be approximated from publicly available BGP routing tables. We evaluate the proposed preference by trace-driven computer simulation, and show that the peer selection method utilizing the proposed preference reduces the transit traffic adequately.

The contributions of this paper are twofold:

1) We show the potential of the transit traffic reduction by peer distribution measurement and analysis.
2) We propose a simple AS relationships-aware peer selection preference and the peer selection method with the proposed preference adequately reduces the transit traffic without non-disclosure AS relationships information in trace-driven simulation.

II. Related Work

Round Trip Time (RTT) and router hop count have been used in peer selection algorithms to achieve high-quality (i.e., low delay and high throughput) content delivery [6], [7].

Xie et al. [3] have proposed a peer selection algorithm and traffic control architecture for the efficient intra-domain and interdomain network resource utilization. Their proposal, however, has focused mainly on the intra-domain traffic engineering for improving content delivery quality, and their interdomain traffic engineering approach pays little attention to the economical relationships among ASes (i.e., transit or peering). Furthermore, there are difficulties in disclosing the metric and integrating complex policies among ASes because the metric in their approach is converted from non-disclosure router configuration.

III. AS Relationships and Content Delivery

In peer-assisted CDNs, contents are delivered over the Internet utilizing peer-to-peer technologies. The Internet consists of thousands of ASes operated by many distinct administrative domains such as ISPs, companies and universities. The routing among ASes is determined by an interdomain routing protocol such as BGP [8].

A. Interdomain Economics and Routing Policies

The economical relationships between two neighboring ASes can be categorized into three types [9]; 1) *transit*, 2) *peering* and 3) *sibling*. There are certain routing policies for each type of the relationships. We describe these relationships, economics and routing policies below.

1) *Transit*: An AS purchases access to the Internet from another AS by paying some amount of money according to bandwidth usage. Hence, transit traffic costs more for customer ASes and they need to reduce this traffic. A transit link from a customer AS to a provider AS and a link with opposite orientation are called customer-to-provider (c2p) link and provider-to-customer (p2c) link, respectively. *A provider AS* exports its routes and its customer routes, as well as its provider or peer routes [10]. On the other hand, *a customer AS* exports its routes and its customer routes, but usually does not export its provider or peer routes [10].

2) *Peering*: A pair of neighboring ASes can exchange traffic directly and traffic exchanged between two peering ASes is free of charge. A peering link is called peer-to-peer (p2p) link. *A peer AS* exports its routes and its customer routes, but usually does not export its provider or peer routes [10].

3) *Sibling*: Multiple ASes can belong to same organization. Even though each AS may be managed separately from perspective of network administration, traffic can be exchanged among them *without* any payment. A sibling link is called sibling-to-sibling (s2s) link. *A sibling AS* exports its routes and routes of its customers, as well as its provider or peer routes [10].

Additionally, there is a typical best path selection policy when an AS receives multiple paths to an identical destination [11]; (a) highest priority for customer routes, (b) middle priority for peer routes, and (c) lowest priority for provider routes[1]. If the priority of each candidate link is equal, one of the shortest (i.e., minimum AS hop count) AS paths is selected. If AS path length is also equal, we refer to the AS number of its neighbors, then the route from the smallest AS number[2] is selected.

These routing policies result in *valley-free* paths [8] when we ignore s2s links. A valley-free path means that a path between any two ASes first traverses uphill (c2p) links, goes across at most one peering link, and then traverses downhill (p2c) links. We show an example of valid and invalid AS paths according to the valley-free path model in Fig. 1. From the economical viewpoint of the valley-free path model, the source and/or destination ASes possibly pay some amount of money for their communication, and intermediate ASes have no cost disadvantage.

[1]Sibling is a special relationship, and there is not a typical best path selection policy. In the AS path computation (Section IV-A), we set the lowest priority (i.e., lower than the priority for provider routes) for sibling routes.

[2]Most implementations refer to router ID instead of AS number [12], but we refer to AS number because the dataset we use in this paper does not include router IDs.

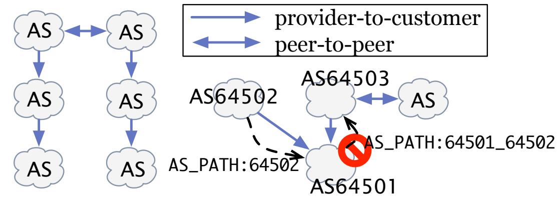

(a) valid path (b) invalid path

Fig. 1. An example of valid and invalid paths on the valley-free path model. The interdomain routing policy results in the valley-free paths. Path (b) is invalid (does not follow valley-free path model) because *AS64501* does not export routes received from a provider *AS64502* to another provider *AS64503*.

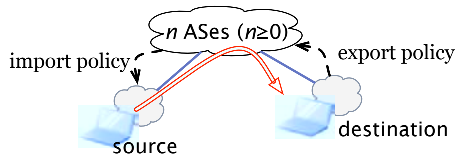

(a) Address-based routing

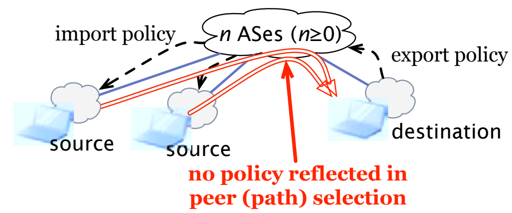

(b) Routing in peer-assisted CDNs

Fig. 2. Routing characteristics. (a) On the Internet, all paths between any two *ASes* follow interdomain routing policies because each AS operates its AS border routers by itself. (b) In peer-assisted CDNs, paths between any two *ASes* also follow interdomain routing policies, but the path from *a content* to *a peer* is selected by peers, and consequently the selection is not controlled by ASes.

B. Content Delivery Path

As described in Section III-A, there are certain policies on the interdomain routing and paths between two any ASes follow these policies on the Internet. However, in peer-assisted CDNs, paths from a content to a peer in a certain AS is selected by peers, and consequently ASes cannot control the selection as shown in Fig. 2. Since transit traffic costs more for customer ASes, the transit traffic reduction in peer-assisted CDNs is essential for them.

We show three examples of possible content delivery paths in Fig. 3. From the viewpoint of economics, path (a) in Fig. 3 is obviously the worst case among these three examples because both of the edge ASes have to pay transit fee. Paths (b)

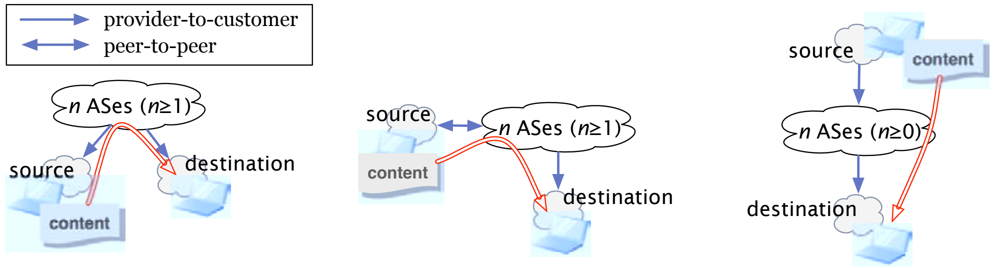

(a) src: to provider, dst: from provider (b) src: to peer, dst: from customer (c) src: to customer, dst: from provider

Fig. 3. Three examples of possible content delivery paths and the types of AS relationships between the source and its neighbor, and the destination and its neighbor. Obviously, path (a) is the worst case in these examples because both source and destination exchange traffic with their provider ASes; i.e., both the source and the destination ASes are customers and have a cost disadvantage. Path (b) is a better case compared to path (a) because one of the edges is not customer and it does not pay transit fee for the content delivery. Path (c) *may* be better than path (b) because the source can acquire transit fee for the content delivery.

TABLE I
THE TYPE OF RELATIONSHIPS OF EDGE ASES AND DELIVERY COST

$ToR^{dst} \setminus ToR^{src}$	p2c	p2p	c2p
p2c	N/A	N/A	$+, -$
p2p	N/A	$0, 0$	$0, -$
c2p	$-, +$	$-, 0$	$-, -$

Each element x, y ($x, y \in \{ -, 0, + \}$) denotes that the cost of destination AS and that of source AS are x and y. Symbol $-$ means the AS can pay transit fee, and symbol $+$ means the AS acquire transit fee. Symbol 0 means the AS delivers contents without any payment. N/A denotes the path does not follow valley-free path model.

TABLE II
THE NUMBER OF LINKS AND THE PROPORTION FOR EACH RELATIONSHIP

relationship	#links	proportion
sibling (s2s)	219	0.302%
peering (p2p)	6142	8.47%
transit (p2c/c2p)	66181	91.2%

and (c) are better than path (a) because one of the edges is not a customer and it does not pay transit fee. We cannot definitely determine which path is better between paths (b) and (c), but path (c) *may* be better than path (b) because one of the edges can acquire transit fee from its customer for the content delivery. From the viewpoint of content delivery quality, shorter AS paths are preferred because there is a strong correlation between AS hop count and Round Trip Time (RTT), which is one of the indicators of path quality [13], [14]. We discuss this point in detail as *AS path length vs. economics* in Section VI.

We list the type of relationships of edge ASes on possible delivery paths following the valley-free path model and the delivery cost in TABLE I. In this table, the symbols ToR^{src} and ToR^{dst} denote the type of relationships from source AS to its neighbor and that from destination AS to its neighbor on the delivery path, respectively; for example, if ToR^{dst} is p2c, the destination AS is a provider of its neighbor on the delivery path. TABLE I does not show quantitative transit fee because there exist some charging policies [15] and the charge fee is not flat to each inter-AS link. From TABLE I, we can qualitatively know that paths which both edge ASes are customers are the worst as we described with Fig. 3.

IV. Potential of Transit Traffic Reduction in Peer-Assisted CDNs

To show the potential of the transit traffic reduction in peer-assisted CDNs, we have measured peer distribution for a content in *BitTorrent* [1]. We then analyze the distribution and content delivery paths by using the AS-level Internet topology dataset [16].

A. Internet Topology Dataset

We employ "The CAIDA AS relationships dataset (10/08/2009) [16]" as AS-level Internet topology for the analysis. The relationships in this dataset are inferred by algorithms [17], [18] from collected AS paths. This dataset includes 32281 ASes and 72542 inter-AS links. We write up the number of inter-AS links and the proportion for each relationship included in the dataset in TABLE II; most of the inter-AS links (91.2%) are transit, and a few links (only 0.302%) are sibling.

We then compute AS paths for all ASes (i.e., paths between two any ASes) according to the routing policies and the best path selection policy described in Section III-A.

B. Peer Distribution Measurement

We had collected a list of peers from a tracker[3] every minute from 23/10/2009 to 19/12/2009 for the content: Debian Linux [19] DVD image; `debian-503-i386-DVD-1.iso` (4.4GB). The collected list includes sets of peer's IP address and port number.

We then aggregate one-minute peer lists into a five-minute peer list to complement incomplete[4] lists as shown in Fig. 4, and extract unique IP address list from the lists and annotate IP addresses with AS number by using a BGP routing table[5]

[3]http://bttracker.debian.org:6969/announce
[4]Not all active peers appear in a list because some peers are filtered out by the tracker.
[5]route-views2.oregon-ix.net archive at 01/08/2009.

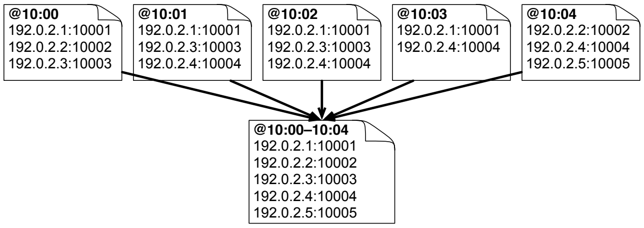

Fig. 4. Peer list aggregation procedure. One-minute peer lists in five minutes are aggregated into a five-minute peer list.

TABLE III
BREAKDOWN OF EDGE AS RELATIONSHIPS OF THE MESH PATHS AMONG ASES ACCOMMODATING BITTORRENT PEERS

ToR^{src}	ToR^{dst}	proportion
c2p	c2p	80.83%
p2p	c2p	11.99%
p2c	c2p	4.48%
others than the above		2.70%

collected by Route Views Project [20]. We observed 48844 unique IP addresses belonging to 2569 ASes in the lists.

C. Content Delivery Path Analysis

To show the potential of the transit traffic reduction in peer-assisted CDNs, we count the number of worst/better paths in case that all the ASes/peers connect each other (i.e., mesh connection).

First, we analyze the mesh paths among ASes which accommodate the peers observed during the measurement period (58 days). The number of the paths amounts 6597192 ($= 2569 \cdot (2569 - 1)$) because the number of ASes which accommodate observed peers is 2569. We show the breakdown of edge AS relationships of these mesh paths in TABLE III. This table shows that most of paths (80.83%) are the worst paths from the economical viewpoint; i.e., both edge ASes are customers. When we assume that the peers are randomly selected, the probability that the worst paths are selected is very high. Hence, there is a potential of the transit traffic reduction by selecting better paths such as p2p–c2p (11.99%) and p2c–c2p (4.48%) paths.

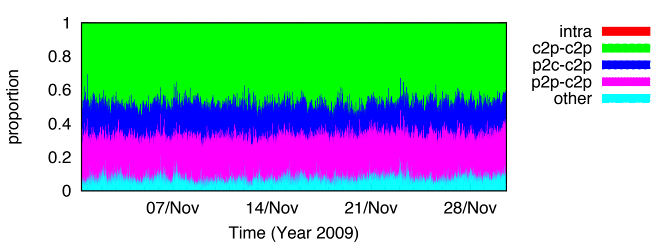

Fig. 5. Temporal breakdown of edge AS relationships of the mesh paths (11/2009). For legend x–y, x and y denote ToR^{src} and ToR^{dst}, respectively. This figure shows that the mean number of the worst paths (c2p–c2p) described in Section V reaches 52.7%.

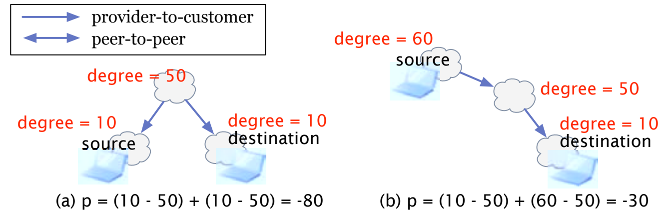

Fig. 6. Example of peer selection preference

Second, we analyze the mesh paths among peers which observed in each five-minute peer list. We show the temporal breakdown of edge AS relationships of these mesh paths in Fig. 5. This figure shows that the worst paths from the economical viewpoint are dominant on every list, and the mean number of the worst paths reaches 52.7%. Furthermore, since few intra-domain paths exist on every list, interdomain traffic control is effective for peer-assisted CDNs. Here, we note that intra-domain traffic control is also important because the traffic goes through each intra-domain network.

V. AS RELATIONSHIPS-AWARE PEER SELECTION

We showed the potential of the transit traffic reduction in Section IV. We then propose AS relationships-aware peer selection preference to reduce transit traffic exchanged with provider ASes in peer-assisted CDNs. The relationships among ASes, however, are generally non-disclosure information due to their commercial contract. Consequently, we employ degree-based heuristics [10] for inferring the relationships in the proposed method. The heuristics is very simple; higher degree AS is a provider and lower degree AS is a customer. We do not use path analysis also proposed in [10] because this requires lots of AS paths and it is difficult to collect them in the peer selection procedure. On the contrary, degree cam be approximated from publicly available BGP routing tables, which is spanning subgraphs, such as Route Views Project [20].

A. Peer Selection Preference

We define peer selection preference p of the path from source peer s belonging to AS S to destination peer d belonging to AS D by Equation (1) according to TABLE I.

$$p(s, d) := (\rho_D - \rho_{N_D}) + (\rho_S - \rho_{N_S}) \tag{1}$$

Here, $\rho_\bullet$ denotes degree of AS $\bullet$, and N_S and N_D are the neighbor of AS S and that of AS D on the delivery path, respectively. The first and second terms of Equation (1) mean transit cost for destination and source edges, respectively; i.e., if the term is negative number, the AS pays transit fee. Consequently, a path with larger preference costs less for ASes accommodating peers and is better.

We show an example of peer selection preference computation in Fig. 6. As shown in Fig. 6, larger number of preference means better paths.

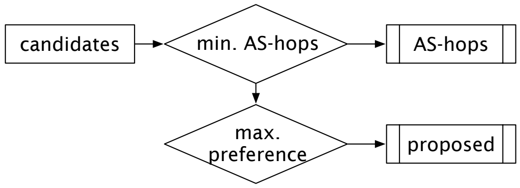

Fig. 7. Peer selection methods for the evaluation

B. Evaluation

We evaluate the proposed preference by comparing with a traditional peer selection method, which is minimum AS hop count selection (AS-hops), through a trace-driven computer simulation. Minimum AS hop count selection is the alternative to the minimum router hop count selection since we do not deal with intra-domain topology in the simulation. The minimum router hop count selection is one of the common approaches for traffic localization. To evaluate the influence of the proposed preference appropriately, we refer to the preference after minimum AS hop count selection in the method with the proposed preference (Fig. 7); i.e., a peer selects one which maximize the preference from candidate peers minimizing AS hop count.

1) Trace-driven Simulation: We use the five-minute peer lists described in Section IV-B for the simulation; the number of lists is 8928 (i.e., 31 days from 31/10/2009 to 30/11/2009). In this simulation, we use a communication model described below.

- Peers exchange 4457 pieces of the 4.4GB content (i.e., the size of a piece is 1MB)[6] over the topology described in Section IV-A.
- The pieces are delivered in series. For example, a peer starts downloading piece 2 after it completes the download of piece 1.
- A peer starts downloading the content (i.e., the first piece) when the peer first appears in the list, and continue downloading the content while the peer is in the list or until the peer completes downloading whole the content. For example, a peer first appears in the list of *01/11/2009 10:05*, and then the peer starts downloading the first piece at *01/11/2009 10:05*. This means every peer downloads the content once, and peers never delete pieces.
- A piece is provided by peers which have already downloaded the piece.
- A peer selects one of the peers which provide the piece in the list according to the peer selection method. When the selected peer disappears from the next list, the downloading peer re-selects another peer.
- The peers measured on the first hour (31/10/2009 00:00–01:00) have all pieces.
- Every peer has 1Mbps upload and download bandwidth. Backbone networks have unlimited bandwidth.

[6]This piece size follows the `.torrent` file.

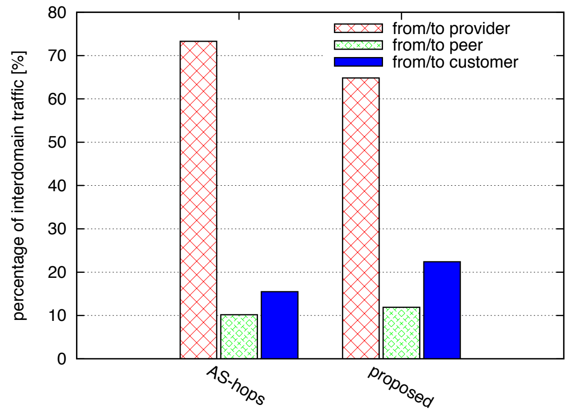

Fig. 8. Breakdown of total exchanged interdomain traffic volume of the ASes accommodating peers by types of AS relationships.

TABLE IV
MEAN THROUGHPUT

method	throughput [Kbps]
AS-hops	22.6
proposed	17.2

2) Results: We show the breakdown of total transferred interdomain traffic volume of the ASes which accommodate peers by types of AS relationships in Fig. 8. The method with the proposed preference reduces transit traffic exchanged with provider ASes by 8.46 percentage point compared to minimum AS hop count selection. This result shows that even the simple preference using degree-based heuristics achieves the transit reduction.

However, the method with the proposed preference deteriorates the mean throughput as shown in TABLE IV. This is because the number of candidate peers decreases and the distribution of selected peers concentrates. We discuss this point in Section VI as *quality concerns and impact on users*.

VI. Discussion

AS path length vs. economics: In Section III, we described AS relationships and the economics. According to the best path selection policy, the economical policy takes higher priority than AS path length. Naturally, ASes can adapt the policy to cases because the routing is controlled by themselves. Then, how about content delivery paths? In Section V, we use the proposed preference after filtering out by AS hop count to compare with minimum AS hop count selection. This is not consistent with the interdomain routing policy. We consider that AS path length, which is one of the quality parameters, should take higher priority than economics in the peer selection by following two reasons: 1) The degree-based heuristics possibly makes inaccurate inference, so quality deterioration by this inaccurate inference should be avoided. 2) The peer selection is not controlled by ASes but by peers, so quality parameter should take higher priority. We will quantitatively evaluate the balance between quality and delivery cost in future, though we naively filtered out by AS hop count to compare the proposed preference with minimum AS hop count selection in this paper.

How to provide degree database and look up AS path: We proposed only peer selection preference and we do not mention the peer selection procedure. The proposed preference requires two types of information in the peer selection procedure; 1) degree of ASes and 2) AS path. Currently, we consider that the degree is provided by degree database server. We will design the protocol to look up the degree in future. AS path is easily resolved by network management tools such as `traceroute` command in Unix/Linux system. Since these tools are not applicable to some networks due to firewall etc., AS path look-up architecture is also problem to be solved.

Quality concerns and impact on users: This mechanism as well as other localization techniques may ameliorate or deteriorate content delivery quality. The localization seems to ameliorate the quality because the distance for content delivery becomes much shorter, but the distribution of selected peers may converge more. It is easily to imagine that some other parameters such as the number of simultaneous uploads would be effective for avoiding this convergence. Moreover, in our evaluation, the quality deterioration occurs because backbone networks have unlimited bandwidth and bandwidth for paths between any two peers is flat. In the real environment, the bandwidth on backbone networks is also limited and throughput decreases according to the distance between peers. We will evaluate these environments in future.

VII. Conclusion

We had measured and analyzed the peer distribution in *BitTorrent* [1], which is one of peer-assisted CDNs. We showed that most of the mesh paths among ASes were the worst paths from the economical viewpoint, and also showed the potential of the transit traffic reduction from the peer distribution analysis.

We then proposed an AS relationships-aware peer selection preference to reduce transit traffic exchanged with provider ASes in peer-assisted CDNs. We employed degree-based heuristics for inferring non-disclosure AS relationships. We evaluated the proposed preference by trace-driven computer simulation and showed that the peer selection method with the proposed preference adequately reduced transit traffic compared to minimum AS hop count selection.

This paper delivers the following contributions; 1) we showed the potential of the high-cost transit traffic reduction and 2) we proposed an AS relationships-aware peer selection preference without non-disclosure AS relationships information, and we showed that the peer selection method with the proposed preference reduced transit traffic.

Acknowledgment

We especially thank Sho Fujita and Yosuke Himura for their support and valuable advice.

References

[1] "BitTorrent," http://www.bittorrent.com/.
[2] V. Aggarwal, A. Feldmann, and C. Scheideler, "Can ISPs and P2P users cooperate for improved performance?" *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 3, pp. 29–40, 2007.
[3] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. G. Liu, and A. Silberschatz, "P4P: provider portal for applications," in *SIGCOMM '08: Proceedings of the ACM SIGCOMM 2008 conference on Data communication*. New York, NY, USA: ACM, 2008, pp. 351–362.
[4] J. Seedorf and E. Burger, "Application-Layer Traffic Optimization (ALTO) Problem Statement," RFC 5693 (Informational), Internet Engineering Task Force, Oct. 2009. [Online]. Available: http://www.ietf.org/rfc/rfc5693.txt
[5] M. B. Weiss and S. J. Shin, "Internet interconnection economic model and its analysis: Peering and settlement," *Netnomics*, vol. 6, no. 1, pp. 43–57, 2004.
[6] R. Carter and M. Crovella, "Server selection using dynamic path characterization in wide-area networks," *INFOCOM '97. Sixteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, vol. 3, pp. 1014–1021 vol.3, Apr 1997.
[7] W. Li, S. Chen, and T. Yu, "Utaps: An underlying topology-aware peer selection algorithm in bittorrent," *Advanced Information Networking and Applications, 2008. AINA 2008. 22nd International Conference on*, pp. 539–545, March 2008.
[8] T. Erlebach, A. Hall, A. Panconesi, and D. Vukadinovi, "Cuts and disjoint paths in the valley-free path model of Internet BGP routing," in *Combinatorial and Algorithmic Aspects of Networking*, ser. Lecture Notes in Computer Science, vol. 3405/2005. Springer Berlin / Heidelberg, 2005, pp. 49–62.
[9] S. Shakkottai and R. Srikant, "Economics of network pricing with multiple ISPs," *IEEE/ACM Trans. Netw.*, vol. 14, no. 6, pp. 1233–1245, 2006.
[10] L. Gao, "On inferring autonomous system relationships in the Internet," *Networking, IEEE/ACM Transactions on*, vol. 9, no. 6, pp. 733–745, Dec 2001.
[11] F. Wang and L. Gao, "On inferring and characterizing Internet routing policies," in *IMC '03: Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*. New York, NY, USA: ACM, 2003, pp. 15–26.
[12] M. Caesar and J. Rexford, "BGP routing policies in ISP networks," *Network, IEEE*, vol. 19, no. 6, pp. 5 – 11, nov.-dec. 2005.
[13] K. Obraczka and F. Silva, "Network latency metrics for server proximity," in *Global Telecommunications Conference, 2000. GLOBECOM '00. IEEE*, vol. 1, 2000, pp. 421–427 vol.1.
[14] A. Nakao, L. Peterson, and A. Bavier, "A routing underlay for overlay networks," in *SIGCOMM '03: Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*. New York, NY, USA: ACM, 2003, pp. 11–18.
[15] X. Dimitropoulos, P. Hurley, A. Kind, and M. P. Stoecklin, "On the 95-percentile billing method," in *PAM '09: Proceedings of the 10th International Conference on Passive and Active Network Measurement*. Berlin, Heidelberg: Springer-Verlag, 2009, pp. 207–216.
[16] Cooperative Association for Internet Data Analysis, "The CAIDA AS Relationships Dataset (10/08/2009)," http://www.caida.org/data/active/as-relationships/.
[17] B. Huber, S. Leinen, R. O'Dell, and R. Wattenhofer, "Inferring AS relationships beyond counting edges," D-INFK, Tech. Rep. Nr.446, 2004.
[18] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, kc claffy, and G. Riley, "As relationships: inference and validation," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 1, pp. 29–40, 2007.
[19] "Debian Linux," http://www.debian.org/.
[20] University of Oregon, "Route Views Project," http://www.routeviews.org/.