

Internet Research Task Force (IRTF)
Request for Comments: 6875
Category: Informational
ISSN: 2070-1721

S. Kamei
NTT Communications
T. Momose
Cisco Systems
T. Inoue
T. Nishitani
NTT Communications
February 2013

The P2P Network Experiment Council's Activities and Experiments with
Application-Layer Traffic Optimization (ALTO) in Japan

Abstract

This document describes experiments that clarify how an approach similar to Application-Layer Traffic Optimization (ALTO) was effective in reducing network traffic. These experiments were performed in Japan by the P2P Network Experiment Council in an attempt to harmonize peer-to-peer (P2P) technology with network infrastructure. Based on what was learned from these experiments, this document provides some suggestions that might be useful for the ALTO architecture and especially for application-independent ALTO-like server operation.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Research Task Force (IRTF). The IRTF publishes the results of Internet-related research and development activities. These results might not be suitable for deployment. This RFC represents the individual opinion(s) of one or more members of the Peer-to-Peer Research Group of the Internet Research Task Force (IRTF). Documents approved for publication by the IRSG are not a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6875>.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Introduction	3
2. Background in Japan	4
2.1. P2P Traffic	4
2.2. Impact on Network Infrastructure	4
2.3. Overview of the P2P Network Experiment Council	5
3. Objectives of the P2P Network Experiment Council	6
4. Details of the Experiment	7
4.1. Dummy Node	7
5. Hint Servers	9
6. High-Level Trial Results	13
6.1. Peer Selection with P2P	13
6.2. Peer Selection with the Hint Server	13
7. Considerations	14
7.1. Next Steps	14
7.2. Feedback to the ALTO WG	15
7.2.1. Hierarchical Architecture for ALTO Servers	15
7.2.2. Measurement Mechanism	15
8. Security Considerations	16
9. Acknowledgments	16
10. Informative References	16

1. Introduction

An overlay network, which is used by P2P and other applications, offers the advantage of allowing flexible provisioning of services while hiding the lower-layer network. The disadvantage is that inefficient routing without considering the lower-layer network may cause increasing the network load. Several proposals have been made to build an overlay network that takes into account the information about the lower-layer network [1] [2]. Since the management of the Internet is highly distributed, it is difficult to implement such proposals, and thus to optimize a network, without the cooperation of network providers.

Recently, the controversy between the overlay network and the network providers about network resource wastefulness has been rekindled. Under these circumstances, some researchers have studied overlay-network control technology that takes into account the network topology information obtained from network providers.

One research effort regarding this issue were experiments planned and performed by the P2P Network Experiment Council in Japan. This document reports on these experiments and the issues they addressed.

These experiments were performed from 2007 to 2008, because P2P traffic decreased after Japanese copyright law was revised. While more recently, the dominant traffic in Japan, the United States, and elsewhere has been HTTP-based flash streaming, a large amount of traffic in Asia (outside Japan) is still P2P traffic, like P2P streaming [3], and P2P technology is very useful in such a real-time streaming area.

Our experience in this experiment might be useful for ALTO architecture, especially for application-independent and multi-application ALTO-like server operations. We suggest that a generic measurement mechanism is important because each application has different mechanism, which makes it difficult to compare their effectiveness.

This document is a product of the P2P Research Group (RG). The views in this document were considered controversial by the P2P RG, but the RG reached a consensus that the document should still be published.

2. Background in Japan

2.1. P2P Traffic

As of 2008, the world's most popular P2P file-sharing application, BitTorrent, was not widely deployed in Japan. Instead, other file-sharing P2P applications specific to Japan, such as Winny [4] and Share [5], still account for 40% of the Internet traffic in Japan, even though many those P2P users were arrested for sharing illegal files with these P2P applications.

Each P2P file-sharing application has a unique protocol and none have a large market share, therefore making it hard to control them effectively.

2.2. Impact on Network Infrastructure

One advantage of using P2P technology for content delivery is that peers exchange content directly among themselves without server bottleneck. This reduces the load on servers. Also, P2P applications can reduce upstream traffic from an origin content server. This reduces server cost dramatically.

It is also known that server cost could be reduced with P2P technology. However, the story is quite different for network providers. From the viewpoint of network providers, the traffic that content servers generate has shifted to the edge network and the amount of traffic has not necessarily been reduced by using P2P technology for reducing server cost. Another problem for network providers is that extremely inefficient routing may be selected because overlay network systems are configured without any regard to the structure of the lower-layer network or network geometry.

In some cases, the total amount of traffic on the Internet used to be limited by the capacity of servers. For those cases, P2P technology can improve the scalability of servers; however, it may exhaust network resources. Moreover, using P2P applications remarkably increases the volume of traffic per user.

Faced with an increase in the load on network infrastructure, network providers are compelled to take actions to overcome the sudden increase in facilities' costs. Representative actions include placing content in Internet Exchanges (IXs) or data centers, introducing bandwidth control, and raising access fees [6].

As mentioned above, the dominant traffic currently in Japan, the US, and elsewhere, is HTTP-based flash streaming. However, a large amount of traffic in Asia (outside Japan) is P2P traffic, like P2P

Streaming [3], and P2P technology is very useful in such real-time streaming. The increase in traffic arising from such a shift may be a great threat to the network.

2.3. Overview of the P2P Network Experiment Council

In order to reduce Internet traffic and encourage legitimate use of P2P technologies, in 2006 the Japanese government established a new council called the P2P Network Experiment Council, in conjunction with commercial P2P application vendors and ISPs.

The council developed regulations that include guidelines such as giving advance notice to heavy users before restricting their bandwidth. In accordance with the regulations, some ISPs introduced solutions that reduce traffic caused by P2P file-sharing applications.

In addition, the council, along with ISPs, carriers, contents providers, and P2P system vendors, looked for new ways to control traffic by commercial P2P applications. In this work, the council performed experiments that introduced an ALTO-like system and observed how the traffic was reduced when it was redirected to proper peers on the real Internet in Japan.

In our experiment, the council deployed hint servers, which are described in Section 5. Hint servers run a protocol that offers network distances to peers, and these distances are disclosed to P2P application vendors.

Using hint servers, P2P application vendors can introduce ALTO concepts easily into their P2P distribution systems. Because the protocol used by hint servers, as defined by the council, is independent of specific P2P application vendors like BitTorrent. The protocol needs to gather network information from ISPs so it can provide network distance to peers. However, many ISPs dislike disclosing such information to others. Therefore, hint servers are designed to offer little information about an ISP's network architecture to P2P application vendors.

To monitor the traffic of peers, the council also deployed a dummy node, which is described in Section 4.1.

The remainder of this memo provides an overview of the experiments.

3. Objectives of the P2P Network Experiment Council

The Japanese Ministry of Internal Affairs and Communications, which has jurisdiction over information and communication systems in Japan, held meetings of an advisory panel on network neutrality in 2006 and 2007 in order to study issues related to next-generation networks, such as how to ensure fairness in the use of networks and how to define fairness in the cost burden. The panel took an interest in P2P technology as a solution to the impending traffic saturation in the backbone network resulting from the rapid expansion of broadband access in Japan, and it formed a "Working Group on the P2P Network", which carried out an intensive study of P2P networks.

The working group reported that it would be necessary to undertake the following four activities, which are intended to encourage the government to adopt relevant policies [7]:

- o Formulate guidelines on P2P file-delivery applications to be self-imposed by the industry.
- o Promote feasibility tests of P2P networks.
- o Study the current state of traffic control and promote the sharing of information.
- o Hold working group meetings about traffic control.

The first two proposals led to the establishment of the P2P Network Experiment Council, supported by the Japanese Ministry of Internal Affairs and Communications [8] [9]. The Council, with membership from P2P delivery providers, content holders, and network providers, began a variety of delivery experiments, which were expected to strengthen cooperative control between different layers. In contrast to P4P (Proactive Network Provider Participation for P2P), which takes a relatively top-down approach of adopting an architecture based on a proposal from a university, the Council is characterized by its bottom-up approach. The aim of establishing the Council was described as follows (translated from [10]).

The rapid growth of broadband access enables content delivery systems to deliver high-quality and high-volume videos securely and efficiently. Although P2P technology is an effective technology for this requirement, it still has some issues to be coped with. Therefore, the "P2P Network Experiment Council" was established with the support of the Japanese Ministry of Internal Affairs and Communications, with its secretariat set up within the

Foundation for MultiMedia Communications (FMMC), in order to formulate guidelines for providers and conduct feasibility tests so that users can receive video delivery services safely.

The activities of the P2P Network Experiment Council can be classified into two categories. The first is formulating guidelines for promoting the commercial use of P2P technology. These guidelines will enable users to use P2P technology safely and will give providers clear rules they must observe. The second is feasibility testing of P2P technology. Section 4 describes experiments conducted in 2007 and 2008.

4. Details of the Experiment

The Council investigated data offered by the members of the Council and learned that the server cost could be reduced by using P2P technology for content delivery. For example, the data from the vendors showed the following:

Traffic was reduced by 90% with UGLive by Utagoe, Inc. [11].

The cost of delivering to tens of thousands of subscribers was reduced by 80% with BBbroadcast with TV Bank Corp. [12]

On the other hand, these reduced server costs may have affected the network load. One of the goals of our experiments was to visualize the impact and propose an architecture to reduce network load caused by these new technologies.

In order to visualize the reduction of network cost, we modeled P2P applications and a multi-ISP environment. This model was also needed for visualizing the effectiveness of the ALTO-like approach.

4.1. Dummy Node

As mentioned above, while the effect of using P2P technology to reduce the traffic and the load on servers is well known; however, traffic behavior in the inter-ISP area is not known. In Japan, the ISPs and IXes cooperated to create a backbone traffic report [13]. However, the measurements gathered for that report required capturing packets on subscribers' lines in order to determine the end users' activities. It is not realistic to measure the behavior of P2P applications at user terminals connected to the Internet because that would require a large-scale arrangement for measurement, such as using deep packet inspection (DPI) on aggregated lines.

To solve these problems, we put several nodes called 'dummy nodes' in the ISP's networks. The dummy nodes emulate an end user's PC running P2P applications. Every P2P node provided by participating vendors in the experiment was configured so it always contacted the hint server.

By introducing dummy nodes and measuring the traffic on them, we were able to observe and evaluate how much the P2P applications affected the networks. Since this method can't measure every subscriber's traffic, the accuracy is less than other methods. However, using dummy nodes makes it possible to adapt to situations in which many different P2P applications coexist on a network. We decided that using dummy nodes was suitable for these experiments.

A dummy node consisted of an Intel PC server running Linux (CentOS), VMWare, and Windows XP on VMWare. With this configuration, all packets can be captured without any impact on the behavior of the network, nodes, or applications. Also, this configuration enabled us to use different P2P applications for Windows and evaluate them generally.

To see behaviors of the node, incoming and outgoing packets are captured on Linux because every packet is transmitted through it. To see flow information in these experiments, we captured the source and destination addresses, port number, amount of traffic, and start and end times.

We placed 60 dummy nodes on access networks of 40 different ISPs. They were placed as close as possible to the subscriber in each network.

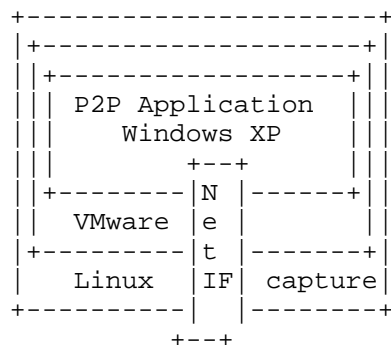


Figure 1: Dummy node

5. Hint Servers

Since fiber to the home (FTTH) has rapidly spread all over Japan, bottlenecks in IP networks have been shifting from access networks to backbone networks and equipment, such as bandwidth between ISPs and capacity in IXs. Under these circumstances, the Council proposed less restrictive and more flexible cooperation between ISPs than existent P4P experiments [14]. The proposed method consists of the following elements: (1) P2P clients, (2) P2P control servers, and (3) a hint server (specifically, a peer selection hint server). P2P clients and control servers are existing systems, but whether the P2P control servers exist is application dependent. The hint server is a server that provides a hint for peer selection and plays a role equivalent to that of the ALTO server. Note that this proposal was based on results of experiments using dummy nodes. The results showed that it was possible to reduce unnecessary traffic that flows across the boundaries of geographical districts or ISPs by providing information about the physical network to P2P applications.

When a peer joins the network, it registers its location information (IP address) and supplementary information (line speed, etc.) with the hint server. The hint server calculates the network distance between peers (P2P clients) based on network topology information obtained from the ISP and generates a priority table for peer selection. The hint server returns the table to the peer.

If all information is public, the above procedure can produce results that are nearly optimal. However, some information held by ISPs is often confidential. Also, in some cases, the volume of calculation required to process all information can be excessive. To avoid these problems, the plan is to conduct experiments with a limited set of functions, analyze the results, and gradually expand the scope of optimization.

A control mechanism that makes use of all possible information is difficult not only technically but also because it requires coordination among providers. In light of these difficulties, the council has been limiting the implementation and experiments to the technical scope.

Figure 2 shows an outline of the hint server.

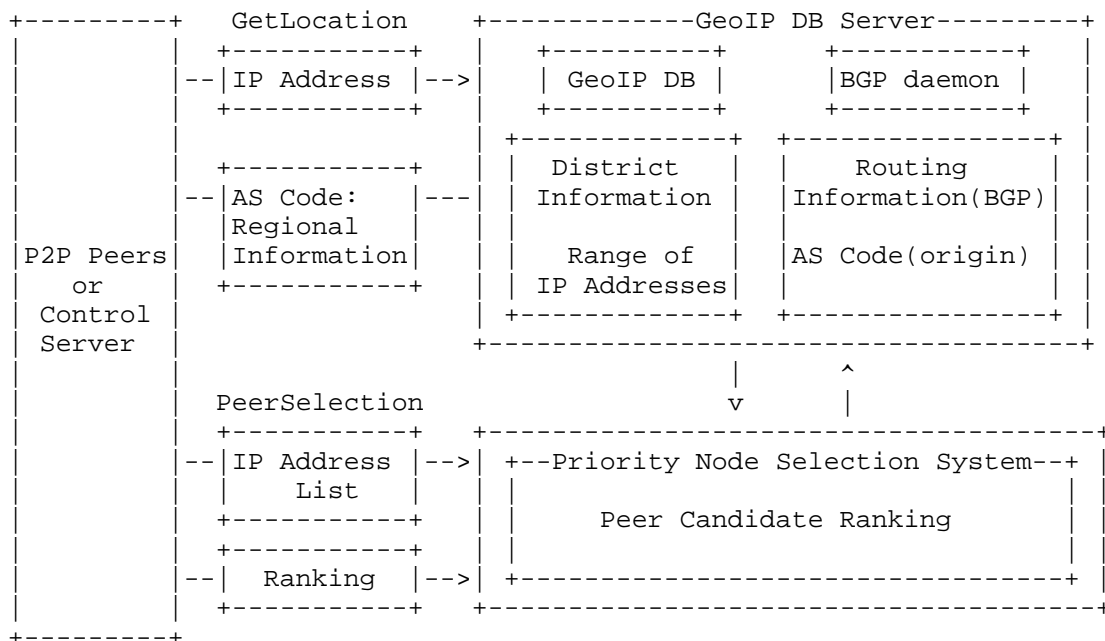


Figure 2: Hint server for peer selection

The network information used by the hint server is not information solicited from individual ISPs but is the Autonomous System (AS) number and district information, which are more or less public already. Routing tables are not generated. Instead, peers within the same ISP or the same district are selected with higher priority in order to confine traffic to within the same ISP or the same district.

When the hint server receives an IP address, it returns its attribute information, in order to confine the traffic to within the nearer ISP or district. A peer can select another based on the returned information. This operation is called `GetLocation`. However, in preparation for the time when it becomes necessary to hide topology information, an interface is provided through which a priority order is returned in response to an input of a list of candidate peers. This operation is called `PeerSelection`.

Although the target node is selected based on the criterion that it is within the same ISP or the same district, this type of selection is not very effective if the number of participating peers is small. Table 1 shows the percentage of peers within the same AS or the same prefecture calculated from the distribution of ASes and prefectures in the IP address space from one-day data on a Winny network.

Conditions	Percentage
AS matches	6.70%
Prefecture matches	12.76%
Both match	2.09%
Neither match	78.45%

Table 1: AS and prefecture distributions

Because, in addition to the above, the presence or absence of content affects the results, controlling peer selection within the same district may be inadequate. Therefore, it is necessary to introduce the weight of a continuous quantity that reflects the physical distance or the AS path length as an indicator of the proximity of the areas involved.

In consideration of this, the following two measures are used to evaluate the proximity of peers in a hint server.

- o AS path length (distance between ISPs)

AS path length is calculated from BGP full routes. Since a full routing table retrieved at an ISP can show only a best path, it may not get an accurate length if the AS hop count of both ISPs is too large. To avoid this, we use BGP information received from different ISPs and combine them. Based on this concept, we used BGP routing information offered by three ISPs operated by big telecommunication couriers and made a topology tree. Then, we were able to calculate the shortest path between two given ASes.

- o Geographical distance

Distances between peers are measured using the physical distance between the capitals of the prefectures to which the peers belong. Distances between prefectural capitals are sorted into ascending order, and then into bands, with weights 1 to 15 assigned to them so that each band contains roughly the same number of "capital pairs". If either of the peer's locations is indefinite, the distance is equal to 15; if they are in the same prefecture, the distance is equal to 0.

Evaluation of distances between peers showed that the distribution of distances was almost uniform when distances between peers are normalized. This result suggests that using normalized distances

expands the area where the control by a hint server is effective. The geographical distance is used only when the AS path length is the same between some candidates.

An example of the request and the response follows.

o Request

```
POST /PeerSelection HTTP/1.1
Host: ServerName
User-Agent: ClientName
Content-Type: text/plain; charset=utf-8

v=Version number
[application=Application identifier]
ip=IP address of physical interface
port=Port number of physical interface
[nat={no|upnp|unknown}]
[nat_ip=Global IP address using UPnP]
[nat_port= Global port number using UPnP]
[trans_id=transaction ID]
[pt=Flag of port type]
[ub=upload bandwidth]
[db=download bandwidth]
```

o Response

```
HTTP/1.1 200 OK
Date: Timestamp
Content-Type: text/plain; charset=utf-8
Cache-control: max-age=max age
Connection: close

v=Version number
ttl=ttl
server=hint server name
...
trans_id=transaction ID
pt=Flag of port type
client_ip=Peer IP address observed from server
client_port=Peer port number observed from server
numpeers=number of responding peers
n=[src address] dst address / cost / option
```

6. High-Level Trial Results

6.1. Peer Selection with P2P

Table 2 shows the result of the analysis of communication in a node of an ISP in Tokyo, as an example of measurement results.

In these two experiments, we evaluated different P2P applications. In the first experiment, the P2P topology was generated by a tree algorithm; in the second experiment, it was generated by a mesh algorithm. Both resulted in similar performance.

Conditions	Experiment 1	Experiment 2
Peers selected within the same ISP	22%	29%
Peers selected within the same district	19%	23%
Peers selected within the same district and the same ISP	5%	7%

Table 2: Percentage of communication within the same ISP

Table 2 shows that the probability of communication with peers in the same ISP is proportional to the population size and the share of the ISP in each district. The data show that peers were selected at random. Note that the vendor of a P2P application used in these experiments demonstrated that the mechanism for selecting a peer using network information can be implemented. However, peer selection is normally based on past information because users often cannot actually perceive the effect of using network information.

6.2. Peer Selection with the Hint Server

The main objective of these experiments was to verify the operation of the hint server and P2P applications. The distances between a dummy node and a peer were obtained from data on the dummy nodes. An examination of the distances between a dummy node and a peer revealed that the mean value of distance after the hint server was introduced was reduced by 10% and that the 95th percentile was reduced by 5%. The results show that introducing a hint server can reduce the network loads that result from P2P applications.

7. Considerations

We clarified the following during our experiments.

1. Dispersed dummy nodes can determine the behavior of peers and traffic between inter-ISP networks and can determine the peer that each peer selects. Therefore, this result proves the importance of the peer-selection control mechanism that is proposed by ALTO.
2. Using our peer-selection control mechanism, called hint servers, can result in significant differences. Hint servers can lead each peer to select a closer peer.
3. The 10% reduction of network cost is not satisfactory for ISPs, but the controllability of P2P applications is the most important point. When ISPs apply this mechanism to their real networks, they will set a very large cost for the most expensive network link.

In the experimental results for peer-selection control, the selection is smaller in intra-ISP traffic than in other experiments [15]. We think this is because there are fewer peers in each area of traffic control. When there are many peers in one ISP, it is easy to select peers in the same ISP. However, when there are fewer peers in one ISP, it is difficult to select peers in the same ISP. In our experiments, most of the ISPs had many peers in their networks, i.e., there were a small number of ISPs that had few peers in their networks.

Moreover, we didn't force P2P vendors to limit their implementation policy; therefore, we observed differences in how each implementation weighs the information from the hint servers. Specifically, in P2P applications when a tree topology is used, the hint-server mechanism is very effective; on the other hand, when a mesh topology is used, it is less effective.

7.1. Next Steps

In recent research, we've changed to an ALTO-based communication protocol on hint servers because the requirements of ALTO are documented in RFC 6708 [16] and the ALTO protocol is a work in progress [17]. In our implementation, protocol identifiers (PIDs) and the cost value are mapped to ISP subnets and to ISP distance, respectively. We also implement services for compatibility required by ALTO such as Map Services and Endpoint Cost Service. The Endpoint Cost Service (defined in [17]) is mainly used because of backward compatibility with our experiments.

We are also studying a hierarchical structure of hint servers, in order to control traffic at a coarse level (in inter-ISP areas) and at a finer level (in intra-ISP areas). It is also effective for limiting the areas where information is disclosed.

7.2. Feedback to the ALTO WG

This section describes what the authors learned from these experiments that might be useful to the ALTO WG.

7.2.1. Hierarchical Architecture for ALTO Servers

In our experiments, we present the possibility of traffic control among multiple ISPs and multiple P2P applications using an ALTO mechanism. We found several problems when ISPs try to adopt the mechanism. One is the granularity of network information from Council members. Among inter-ISP areas, it is relatively easy to handle information for public purposes by using BGP full routes. On the other hand, among the intra-ISP areas, it may be difficult to disclose the private information of each ISP. Kiesel [18] proposes some modifications for the ALTO protocol in order to hide ISP information. We propose hierarchical structures. From the viewpoint of cooperation between ISPs, fine-grained information is not necessarily required. Moreover, it is difficult to exchange the fine-grained information between ISPs. Considering this situation, we used only coarse-grained information to control backbone traffic in these experiments; however, in the future, there may be a demand for controlling traffic within an ISP using fine-grained information. Therefore, we decided to introduce hierarchical structures into ALTO in order to cope with both situations. Actually, adopting a hierarchical control mechanism that includes the following two steps will be useful.

- o First, use coarse-grained information about whole the network to select ISPs.
- o Second, use fine-grained information within the ISP to select a peer.

7.2.2. Measurement Mechanisms

In these experiments, there were two difficulties as follows.

- o Evaluating the effect of introducing a hint server was difficult because the P2P applications had their own measurement mechanisms.
- o How to treat the priority order of peers suggested by a hint server could not be predetermined for P2P applications.

From these experiences, the authors consider that clarifying the requirements about measurement mechanisms for P2P applications is necessary in ALTO.

8. Security Considerations

This document does not propose any kind of protocol, practice, or standard.

9. Acknowledgments

The P2P Network Experiment Council was established thanks to strong support by the Japanese Ministry of Internal Affairs and Communications. These experiments were performed with cooperation among the P2P Network Experiment Council members. DREAMBOAT Co., Ltd., Bitmedia, Inc., Utagoe, Inc., and Toyama IX have especially supported the analyses of the experiments. The authors appreciate Tohru Asami, Hiroshi Esaki, and Tatsuya Yamashita for their constructive comments.

The authors would also like to thank Martin Stiernerling, Stefano Previdi, and Vijay K. Gurbani for their comments on this document.

10. Informative References

- [1] Kawahara, R., Lua, E., Uchida, M., Kamei, S., and H. Yoshino, "On the Quality of Triangle Inequality Violation Aware Routing Overlay Architecture", INFOCOM 2009, pages 2761-2765.
- [2] Li, Z. and P. Mohapatra, "QRON: QoS-aware routing in overlay networks", IEEE Journal on Selected Areas in Communications, Vol. 22, No. 1, January 2004.
- [3] Sandvine, Inc., "Global Internet Phenomena Report: 2H 2012", September 2012, <http://www.sandvine.com/news/global_broadband_trends.asp>.
- [4] Wikipedia, "Winny", July 2012, <<http://en.wikipedia.org/w/index.php?title=Winny&oldid=500744660>>.
- [5] Wikipedia, "Share (P2P)", January 2013, <[http://en.wikipedia.org/w/index.php?title=Share_\(P2P\)&oldid=532999898](http://en.wikipedia.org/w/index.php?title=Share_(P2P)&oldid=532999898)>.
- [6] Taniwaki, Y., "Broadband Competition Policy in Japan", March 2008, <<http://unpan1.un.org/intradoc/groups/public/documents/apcity/unpan040329.pdf>>.

- [7] Ministry of Internal Affairs and Communications, "Disclosure of the Report 'Working Group on P2P Networks'" (in Japanese), 2007, <http://www.soumu.go.jp/menu_news/s-news/2007/070629_11.html>.
- [8] The Foundation for MultiMedia Communications, "The P2P Network Experiment Council" (in Japanese), 2007, <<http://www.fmmc.or.jp/P2P/about.htm>>.
- [9] Ministry of Internal Affairs and Communications, "P2P Network Experiment Council Symposium to Be Held", February 2008, <http://www.soumu.go.jp/main_sosiki/joho_tsusin/eng/Releases/Telecommunications/news080201_1.html>.
- [10] The Foundation for MultiMedia Communications, "The Aim of P2P Network Experiment Council" (in Japanese), 2007, <http://www.fmmc.or.jp/p2p_web/aim.html>.
- [11] Shudo, K., "A Review of ALM Software in Practical Use", IRTF SAMRG (Scalable Adaptive Multicast Research Group) meeting, Proceedings of IETF 76, November 2009, <<http://www.ietf.org/proceedings/76/slides/SAMRG-6.pdf>>.
- [12] TV Bank Corp., "Live Delivery Using 'BB Broadcast' Achieved a 96% Saving in Traffic!" (in Japanese), October 2008, <<http://www.tv-bank.com/jp/20081031.html>>.
- [13] Cho, K., Fukuda, K., Esaki, H., and A. Kato, "The Impact and Implications of the Growth in Residential User-to-User Traffic", SIGCOMM '06, pages 207-218, September 2006.
- [14] Xie, H., Yang, R., Krishnamurthy, A., Liu, Y., and A. Silberschatz, "P4P: Provider Portal for Applications", SIGCOMM '08, pages 351-362, 2008, <<http://www.cs.yale.edu/homes/yry/projects/p4p/p4p-sigcomm08.pdf>>.
- [15] Griffiths, C., Livingood, J., Popkin, L., Woundy, R., and Y. Yang, "Comcast's ISP Experiences in a Proactive Network Provider Participation for P2P (P4P) Technical Trial", RFC 5632, September 2009.
- [16] Kiesel, S., Previdi, S., Stiernerling, M., Woundy, R., and Y. Yang, "Application-Layer Traffic Optimization (ALTO) Requirements", RFC 6708, September 2012.
- [17] Alimi, R., Ed., Penno, R., Ed., and Y. Yang, Ed., "ALTO Protocol", Work in Progress, September 2012.

- [18] Kiesel, S. and M. Stiemerling, "ALTO H12", Work in Progress, March 2010.

Authors' Addresses

Satoshi Kamei
NTT Communications Corporation
Granpark Tower 16F, 3-4-1 Shibaura
Minato-ku, Tokyo 108-8118
Japan

Phone: +81-50-3812-4697
EMail: skame@nttv6.jp

Tsuyoshi Momose
Cisco Systems G.K.
9-7-1 Akasaka
Minato-ku, Tokyo 107-6227
Japan

Phone: +81-3-6738-5154
EMail: tmomose@cisco.com

Takeshi Inoue
NTT Communications Corporation
Kuredo Hakushima Building 3F, 14-15 Higashihakushimacho
Chuo-ku, Hiroshima-City, Hiroshima 730-0004
Japan

Phone: +81-82-563-5030
EMail: inoue@jp.ntt.net

Tomohiro Nishitani
NTT Communications Corporation
1-1-6, Uchisaiwaicho
Chiyodaku, Tokyo 100-8019
Japan

Phone: +81-50-3812-4742
EMail: tomohiro.nishitani@ntt.com

