

Flow Aggregated, Traffic Driven Label Mapping in Label-Switching Networks

Ken-ichi Nagami, Hiroshi Esaki, *Member, IEEE*, Yasuhiro Katsube, *Member, IEEE*, and Osamu Nakamura, *Member, IEEE*

Abstract— Label-switching technology enables high performance and flexible layer-3 packet forwarding based on the fixed-length label information that is mapped to the layer-3 packet stream. A label-switching router (LSR) forwards layer-3 packets based on their layer-3 address information or their label information that is mapped to the layer-3 address information. Two label-mapping policies have been proposed. One is traffic-driven mapping, where the label is mapped for a layer-3 packet stream of each host-pair according to the actual packet arrival. The other is topology-driven mapping, where the label is mapped in advance for a layer-3 packet stream toward the same destination network, regardless of actual packet arrival to the LSR. This paper evaluates the required number of labels under each of these two label-mapping policies using real backbone traffic traces. The evaluation shows that both label-mapping policies require a large number of labels. In order to reduce the required number of labels, we propose a label-mapping policy that is a combination of the two label-mapping policies above. This is traffic-driven label mapping for the packet stream toward the same destination network. The evaluation shows that the proposed label-mapping policy requires only about one-tenth as many labels as the traffic-driven label mapping for the host-pair packet stream and the topology-driven label mapping for the destination-network packet stream.

Index Terms— ATM, communication systems, Internet.

I. INTRODUCTION

THE DATA networks using TCP/IP technology, i.e., the Internet and intranet, are growing both in geographical and numerical scale as well as in traffic volume. Also, recent and upcoming applications often require a specific quality-of-service (QoS) or a class of service (CoS), rather than the conventional best-effort service. Label-switching technology will satisfy these requirements, i.e., high speed and large processing capability while providing QoS and CoS, with reasonable implementation cost. A label-switching router (LSR) forwards layer-3 packets using either their layer-3 address information or their label information that is mapped to the layer-3 address information.

Several architectures for the LSR, including aggregate route-based IP switching (ARIS) [7], [8], cell switch router (CSR)

[3], [4], IP switch [1], [2], and tag switch router (TSR) [5], [6], have been proposed with their own protocol designs. The protocol to establish the mappings between a packet stream and a corresponding label is generally called a label-distribution protocol (LDP), discussed in the Internet engineering task force (IETF) multiprotocol label switching (MPLS) working group.

From the viewpoints of scalability and implementation feasibility of the LSR, it is important to evaluate the required number of labels in label-switching networks. The number of required labels depends on a label-mapping trigger (which is an event that invokes the label mapping), on the granularity of a packet stream, and on the scale of the network.

Two kinds of label-mapping triggers have been proposed. One is traffic-driven mapping, where the label is mapped to a packet stream according to the actual packet arrival. The other is topology-driven mapping, where the label is mapped in advance to a packet stream using the network-topology information, regardless of actual packet arrival to the LSR.

Several kinds of granularity for a packet stream can be defined. The typical definitions for the packet stream are a host-pair packet stream and a destination-network packet stream. The host-pair packet stream is defined as a set of packets having the same source and destination layer-3 addresses. The destination-network packet stream is defined as a set of packets having the same destination-network prefix. Other packet-stream definitions will be discussed in Section II.

Two typical label-mapping policies have been proposed. One is specified in the epsilon flow management protocol (IFMP) [2] and the flow attribute notification protocol (FANP) [4], where the LSR uses the host-pair packet stream as the granularity of the packet-stream definition and applies traffic-driven label mapping. Lin and Mckeown [10] have evaluated the number of required labels using actual traffic traces on the Internet and show that a large number of labels is required on the Internet backbone with this label-mapping policy. The other label-mapping policy is specified in the tag distribution protocol (TDP) [6], where LSR uses the destination-network packet stream as the granularity of the packet-stream definition and applies topology-driven label mapping. With this label-mapping policy, the number of required labels becomes as large as the number of routing-table entries maintained by the LSR.

This paper evaluates the number of labels required under each of the two label-mapping policies using real traffic traces on a wide area Internet backbone. The required number of

Manuscript received July 14, 1998; revised February 1, 1999.

K. Nagami and Y. Katsube are with the Information and Communication Systems Laboratory, Toshiba Corporation, Tokyo 191-8555 Japan (e-mail: ken.nagami@toshiba.co.jp; yasuhiro.katsube@toshiba.co.jp).

H. Esaki is with the Computer Center, University of Tokyo, Tokyo 113-8658 Japan (e-mail: hiroshi@wide.ad.jp).

O. Nakamura is with the Faculty of Environmental Information, Keio University, Kanagawa 252-8520 Japan (e-mail: osamu@wide.ad.jp).

Publisher Item Identifier S 0733-8716(99)04491-1.

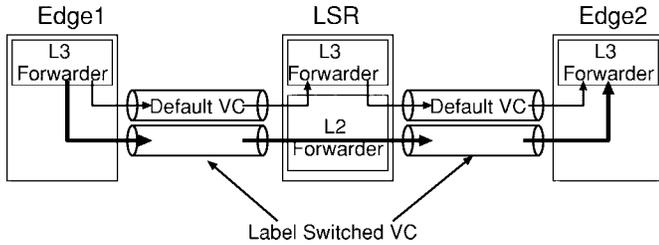


Fig. 1. The LSR.

labels with traffic-driven mapping for the host-pair packet stream is quite large, and the required number of labels with topology-driven mapping for the destination-network packet stream is also large.

In order to reduce the required number of labels, we propose a label-mapping policy that is an integration of the previously mentioned label-mapping policies. The label is mapped to the packet stream toward a specific destination network, triggered by the arrival of its traffic. The required number of labels with the proposed label-mapping policy is evaluated using real wide area Internet backbone traffic.

Section II describes an overview of the LSR, the label-mapping triggers for the LSP, and the granularity of the packet stream. Section III describes the evaluation of the required number of labels with the two conventional label-mapping policies: traffic-driven mapping with host-pair packet stream and traffic-driven mapping with destination-network packet stream. Section IV describes the evaluation of the proposed label-mapping policy with various granularities of the packet stream. Section V gives a brief conclusion.

II. OVERVIEW OF THE LSR

A. Architectural Overview

The label-switching concept enables layer-3 packets to be forwarded using either the layer-2 label (e.g., VPI/VCI) or using the shim label [9] between the layer-2 and layer-3 headers. The label is used as information to forward the packets without analyzing the layer-3 address (e.g., IP address). This means that the label represents the destination address of the layer-3 packets. By using the label instead of the layer-3 address for packet forwarding, the LSR does not need to look up anything in the best-match policy-based routing table, whose search key is the layer-3 address.

The packet forwarding at the LSR is shown in Fig. 1. The LSR has a layer-3 forwarding engine and label forwarding engine. It is connected with two LSR-edge routers (Edge1 and Edge2). LSR's are connected with two types of VC's. One is the default VC, and the other is the label-switched VC.

The default VC is used for conventional packet forwarding. For example, when Edge1 sends a packet to Edge2 using conventional packet forwarding, Edge1 sends the packet through the default VC. The LSR receives it and transmits it to the layer-3 forwarding engine through the label-forwarding engine. The layer-3 forwarding engine of the LSR looks up the routing table according to the destination of the packet and

sends it through the default VC. Then Edge2 receives it. The LSR forwards the packets using the layer-3 engine.

The label-switched VC is used for cut-through packet forwarding [3]. For example, when Edge1 sends a packet to Edge2 using cut-through packet forwarding, Edge1 sends the packet through the label-switched VC. The LSR receives it and forwards it, using only the layer-2 forwarding engine, to send it through the label-switched VC. The LSR forwards these packets faster than with the conventional forwarding because the LSR does not need to look up the layer-3 packet. The conjunction of the label-switched VC is called the label-switched path (LSP).

The LSR has to establish the relationship between the packet stream and the corresponding label for label-switched packet forwarding. Label-distributed protocol (LDP) establishes the mapping between the packet stream and the label.

One of the major applications of an LSR is an ATM-LSR, which contains the cell-switch module as a label-forwarding engine. In an ATM-LSR, the LSR system has to allocate a packet-reassemble buffer space for every LSP. Since there is a hardware limitation regarding the available packet-reassemble buffer space, an ATM-LSR system has the maximum number of labels (i.e., the maximum number of LSP) to be able to provide for label-switching purposes. From the viewpoints of scalability and implementation feasibility of the LSR system, it is important to evaluate the required number of labels in a label-switching network. The required number of labels depends on the following three parameters:

- 1) label-mapping triggers;
- 2) granularity of packet stream;
- 3) scale of the network.

B. Label-Mapping Triggers

There are two types of triggers to control the LSP's, traffic-driven label mapping and topology-driven label mapping.

With traffic-driven mapping, the LSP is established on demand according to the appearance of data packets at a node. The LSP is maintained as long as packets are forwarded through the LSP. When the node recognizes that the LSP is not active anymore, it is released.

With traffic-driven mapping, the LSP is established in advance according to the routing-protocol information. For example, the LSP is established when a routing entry appears through the routing-protocol information. The LSR maintains the LSP as long as the corresponding routing entry exists. LSR tears down the LSP, when the corresponding routing entry is deleted.

The advantage of traffic-driven mapping would be that the required amount of label space is smaller than that for traffic-driven mapping. This is because, in the topology-driven system, the LSP is established and maintained, even though no data packet is forwarded through the LSP.

The advantage of traffic-driven mapping is that all data packets can be forwarded through the LSP. In the traffic-driven system, some portion of data packets are forwarded without label switching, i.e., with conventional layer-3 packet forwarding.

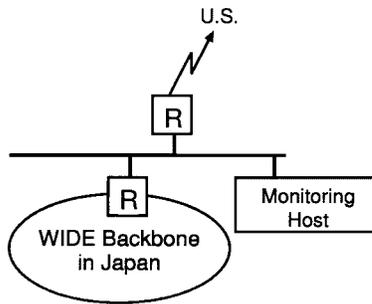


Fig. 2. Traffic-monitoring point.

C. Granularity of the Packet Stream

The LSR system can define a wide variety of packet-stream granularity. In this paper, the following packet-stream granularities are used for evaluation:

- a set of the packets that have the same source and destination addresses, i.e., the so-called microflow, denoted as (src, dst);
- a set of the packets that have the same source address and destination-network prefix, denoted as (src, dstnet);
- a set of the packets that have the same source and destination-network prefixes, denoted as (srcnet, dstnet);
- a set of the packets that have the same destination address, denoted as (*, dst);
- a set of the packets that have the same destination-network prefix, denoted as (*, dstnet).

III. EVALUATIONS OF CONVENTIONAL LABEL MAPPING POLICIES

A. Evaluated Traffic Conditions

This section briefly describes the traffic conditions used for the performance evaluations in this paper.

- *Traffic Monitoring Point:* Fig. 2 shows a traffic monitoring point. A traffic monitoring host is attached to the Ethernet segment that is located between the WIDE (widely integrated and distributed environment) project's Internet backbone network in Japan and the United States. The host records the traffic using the tcpdump application on the host. The measurement was performed for approximately 2 h (7437 s), from 14:08 JST November 4, 1997 to 16:11 JST November 4, 1997. The traffic traces are recorded for each direction separately.
- *Packet Forwarding Rate:* Table I shows the average packet transmission speed at the monitoring point for both directions, in megabits per second (Mbit/s) and in packets per second (p/s). The average packet transfer rate was 1.225 Mbit/s and 329 p/s for the traffic coming into the AS (WIDE project backbone network) and 0.926 Mbit/s and 564 p/s for the traffic going out from the AS.

B. Evaluation Results

1) *Evaluations for Topology-Driven Mapping:* In this section, we evaluate the number of labels for conventional topology-driven label mapping, where each label is mapped

TABLE I
AVERAGE TRANSFER RATE

Direction	Transfer Rate	Transfer Rate
	[Mbit/s]	[p/s]
to the AS	1.225	329
from the AS	0.926	564

TABLE II
NUMBER OF ROUTING ENTRIES

	In the AS	Full
	Route	Route
Total Entries	2865	50903
Entries directed outside of AS	5	48385
Entries directed inside of AS	2860	2518
Referred Entries		
directed outside of AS	5(100%)	7530(16%)
Referred Entries		
directed inside of AS	209(7%)	106(4%)

to the destination-network packet stream shown in the routing table entry.

The total number of routing table entries at the traffic-monitoring point was 2865. The number of routing entries for outside the AS was only five, and the routing entries for inside the AS was 2860. When the labels are established with conventional topology-driven mapping, the required number of labels is the same as the number of routing entries. This means that 2860 labels are required for the conventional topology-driven mapping policy toward the inside of the AS.

Next, we evaluate the required number of labels, where we use the full-route routing table. In this case, it could be assumed that the analyzed LSR is located at the Internet backbone. Table II shows the number of entries in the full-route routing table. The total number of entries in the full-route routing table was 50903. The number of routing entries for inside the AS was 2518, and the routing entries for outside the AS was 48385.

With a full-route routing table, if the labels are established with the conventional topology-driven mapping policy, we need 2518 labels for the traffic coming into the AS, and we need 48385 labels for the traffic leaving the AS. With the conventional topology-driven policy, the result shows that the number of required labels has to be large enough in an Internet backbone area.

2) *Evaluations for Conventional Traffic-Driven Mapping:* In this section, we evaluate the number of required labels for conventional traffic-driven mapping, where each label is mapped for each host-pair packet stream. The following are

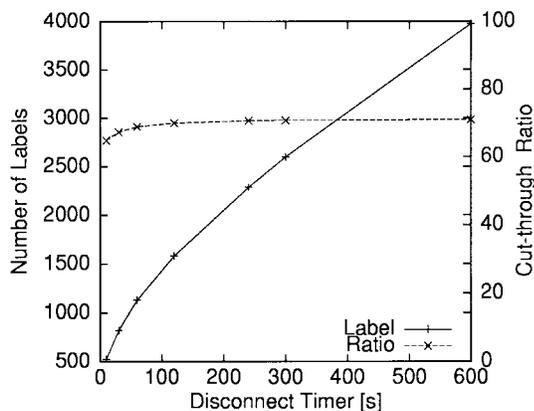


Fig. 3. Number of labels and cut-through ratio directed outside of the AS.

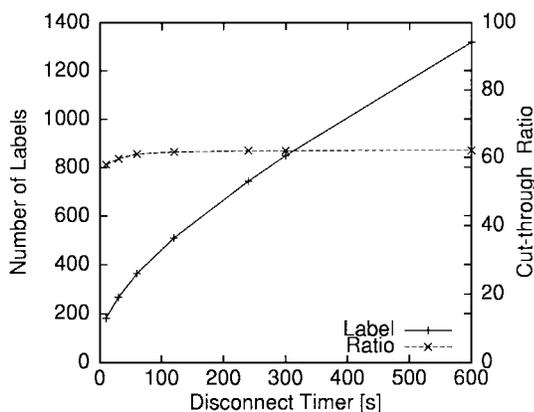


Fig. 4. Number of labels and cut-through ratio directed inside of the AS.

the parameters applied in the LSR system.

- The label-established trigger is invoked whenever the HTTP/FTP/TELNET/NNTP packets are sent out from the LSR.
- The label-released trigger is invoked when no packet is sent during the disconnect-timer interval. The values of the disconnect-timer interval were 10, 30, 60, 120, 240, 300, and 600 s.
- The LDP establishment is completed with a 100 ms delay.¹ The packets are transmitted using the conventional layer-3 forwarding function for 100 ms, after the label-established trigger is invoked.

We evaluate the number of required labels and the cut-through ratio. Here, the cut-through ratio is defined as the number of packets transferred through label switching divided by the number of all packets transferred both through label switching and the conventional layer-3 forwarding function.

Fig. 3 shows the number of the required labels and cut-through ratios for the outgoing traffic from the AS. Fig. 4 shows those for the incoming traffic to the AS.

In both Figs. 3 and 4, the cut-through ratio is almost the same when the disconnect-timer interval is more than 60 ms.

¹The label-mapping establishment time depends on several parameters, e.g., LSR's software and hardware implementation or system configuration. An actual required label-mapping establishment time is 100 ms in a CSR system using an SVC, which is one of LSR's implementations. Here, using a PVC, the label-mapping establishment time is far less than that of the SVC.

This result indicates that the value of the disconnect-timer interval should be 60 ms. Therefore, the number of required labels and the cut-through ratio, when the disconnect-timer interval is 60 ms, are shown below.

The number of labels with the host-pair packet stream going out from the AS is 1131, and the cut-through ratio is 68%. The required number of labels coming into the AS is 365, and the cut-through ratio is 61%.

When we implement an LSR system using ATM technology, there is a limitation regarding the number of labels (VC's). This is because, in the ATM-LSR, the LSR system has to allocate a packet-reassemble buffer space for every LSP. The existing ATM hardware could provide a few hundred labels for the LSR system. The result shows that, in the Internet backbone area, a fairly large number of labels is required with conventional traffic-driven mapping for the host-pair packet stream.

IV. FLOW AGGREGATED, TRAFFIC DRIVEN, LABEL-MAPPING POLICY

A. Flow Aggregated, Traffic Driven, Label-Mapping Policy

In Section III, we show that the conventional label-mapping policies require a fairly large number of labels. In order to reduce the required number of labels, we propose a new label-mapping policy. The label is mapped to the packet stream toward a specific destination network triggered by the actual packet arrival. In this paper, this policy is denoted as a traffic-driven label mapping with the flow-aggregated packet stream. In this section, the required number of labels with the proposed flow-aggregated traffic-driven policy is evaluated using the real Internet backbone traffic traces.

B. Evaluation of Flow Aggregated, Traffic Driven Policy

The flow-aggregated traffic-driven label-mapping policy is evaluated using the same Internet backbone traffic traces used in Section III. The performance of the proposed policy is evaluated with different granularities of the packet stream, including packet-stream granularity with the conventional topology-driven mapping policy.

1) *Evaluation of Actually Referred Routing Entries:* In this section, we evaluate the number of actually referred routing entries using real Internet backbone traffic. We evaluated both regarding the incoming packet flow to the AS and regarding the outgoing packet flow from the AS.

First, we evaluated the actually referred routing entries, where we use a default routing with a small number of individual routing entries for packets going out from the AS. We analyzed which routing entries were actually referred to when the individual packets were transmitted from the LSR toward each AS. The results are shown in Figs. 5 and 6 and Table II.

Figs. 5 and 6 show the referred frequency of each individual routing entry during the monitoring period. The horizontal axis shows the destination prefix of the packets. For example, when the LSR forwards the packet whose destination IP address is "192.69.251.56," and the LSR has a correspond-

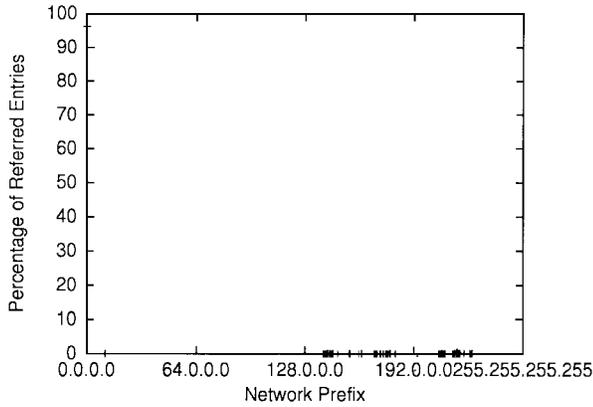


Fig. 5. Percentage of referred entries directed outside the AS.

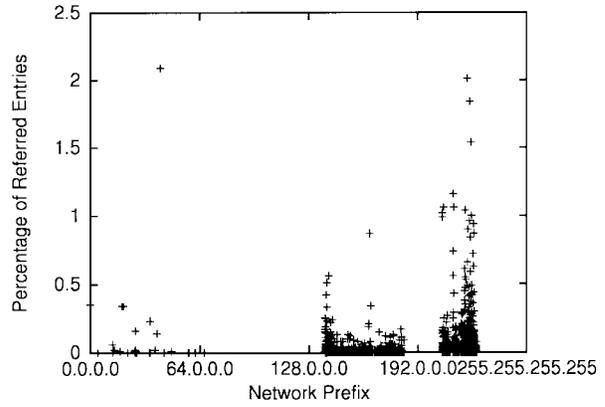


Fig. 7. Percentage of referred entries directed outside of the AS.

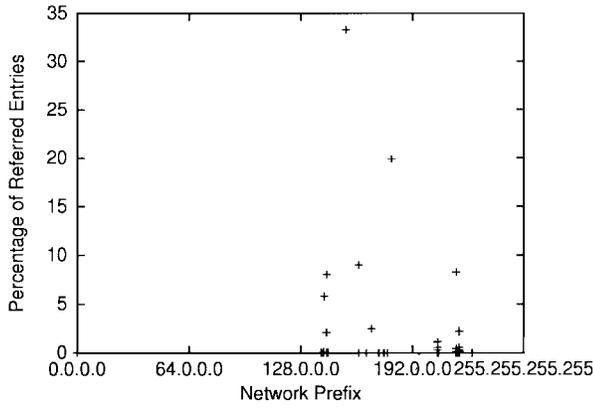


Fig. 6. Percentage of referred entries directed inside of the AS.

ing best-matched routing entry of “192.69.251.0/24” for the forwarded packet, this packet-forwarding event is included for the referred routing entry at the position of “192.69.251.0” on the horizontal axis. Here, “0.0.0.0” on the horizontal axis corresponds to where the packets are forwarded with the default routing.

Fig. 5 shows that 96% of packets directed outside the AS are forwarded using the default route (at “0.0.0.0” on the horizontal axis). Table II shows the number of routing entries at the end of the monitoring period and the actually referred routing entries during the monitoring period. The LSR has five routing entries for outside the AS, i.e., one entry is for the default routing and four entries are for individual destination networks outside the AS. All five external routing entries are referred to during the monitoring period. Here, in Fig. 5, more than five routing entries are referred to, even though there are only five routing entries for outside the AS. Some routing entries temporally disappeared from the routing entry due to some routing instability. In this case, the packets directed toward the temporally disappearing network(s) inside the AS are forwarded to outside the AS using the default routing. In the figure, these kinds of packet forwarding are included as the sampled packet forwarding. Therefore, the number of referred routing entries during the monitoring period is more than five, which exists in the routing entries directed to outside the AS.

Regarding the incoming packet flow toward the AS, 209 routing entries are referred to during the monitored period.

In this evaluation, the network, corresponding to the most frequently referred routing entry, receives about 33% of incoming packets through the LSR. When a label is established with the proposed mapping policy, it is sufficient to provide only 209 labels. The actual number of labels required during the analyzed period is less than 209. This is because 209 is not the maximum, but the accumulated number of routing entries referred to during the analyzed period. On the contrary, as evaluated in Section III, when we apply conventional topology-driven mapping, the number of required labels for traffic toward the AS has to be 2860. As a result, the number of required labels with the proposed mapping policy is less than 7% of that with the conventional topology-driven mapping policy.

The second evaluation is where we use a full-route routing table for outgoing packets from the AS. The total number of the full-route routing table entries was 50 903. As shown in Table II, the number of routing entries for the AS was 2518, and the number of routing entries for outside the AS was 48 385. The results of the evaluation are shown in Fig. 7 and Table II. Again, the horizontal axis shows the destination prefix of the packets, and “0.0.0.0” on the horizontal axis corresponds to where the packets are forwarded with the default routing. The default routing in Fig. 7 is the case where the packets are forwarded using the default route, since those packets do not belong to any routing entries in the LSR.

As shown in Table II, 106 labels are required for the packets toward the AS, and 7530 labels are required for the packets toward outside the AS. The number of required labels with the conventional topology-driven policy is 50 903. On the contrary, the total number of required labels with the proposed policy is only 7636 ($7376 = 7530 + 106$). The means that the number of required labels with the proposed mapping policy is only 15% of that with the conventional topology-driven mapping policy.

2) *Evaluations with Granularity of Packet Streams:* In this section, we evaluate the required number of labels and the cut-through ratio for the proposed label-mapping policy with several packet-stream granularities. The parameters of traffic-driven mapping in the LSR are the same as in Section III. Five granularities of the packet stream, (src, dst), (src, dstnet), (srcnet, dstnet), (*, dst), and (*, dstnet), are evaluated.

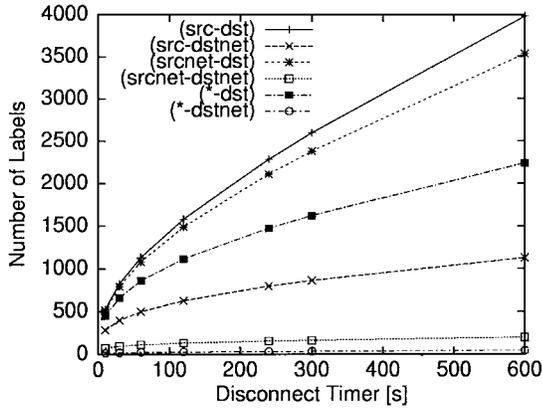


Fig. 8. Number of labels directed outside the AS.

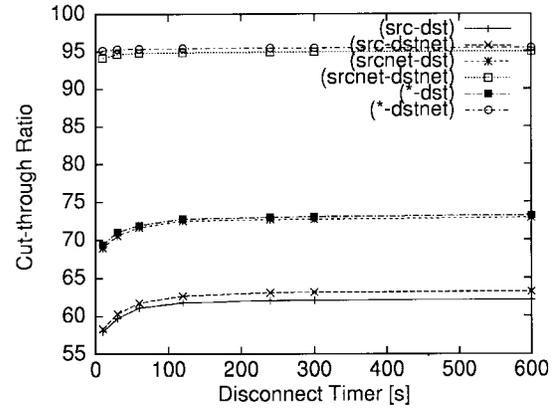


Fig. 11. Cut-through ratio directed inside the AS.

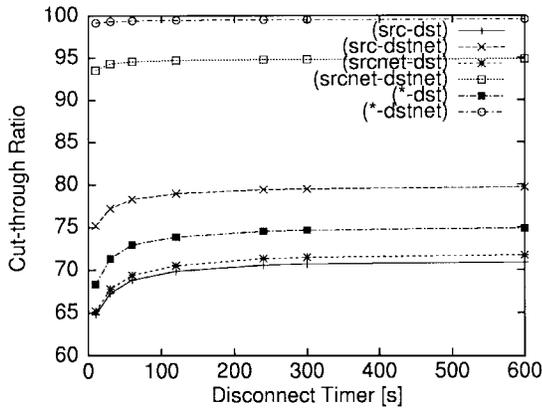


Fig. 9. Cut-through ratio directed outside the AS.

TABLE III
NUMBER OF LABELS AND CUT-THROUGH RATIO DIRECTED OUTSIDE THE AS

Granularity	Number of Labels	Cut-through Ratio
(src, dst)	1131	68
(* , dst)	859	73
(src, dstnet)	498	78
(srcnet, dstnet)	113	94
(* , dstnet)	99	99

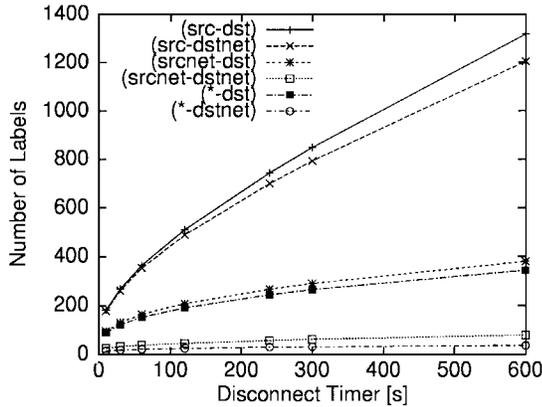


Fig. 10. Number of labels directed inside the AS.

TABLE IV
NUMBER OF LABELS AND CUT-THROUGH RATIO DIRECTED INSIDE THE AS

Granularity	Number of Labels	Cut-through Ratio
(src, dst)	365	61
(src, dstnet)	353	62
(* , dst)	152	72
(srcnet, dstnet)	38	95
(* , dstnet)	21	95

labels and the cut-through ratio when the disconnect timer is 60 ms is shown in Tables III and IV.

Figs. 8 and 9 show the number of the required labels and the cut-through ratio for the traffic going out from the AS. Figs. 10 and 11 show the number of the required labels and the cut-through ratio for the traffic coming into the AS. The horizontal axis shows the value of the disconnect-timer interval. In the evaluation, we use the default routing with several individual routing entries for packet forwarding toward outside the AS.

As shown in Table III, the number of required labels with the host-pair packet stream is 1131 for the packets going out from the AS. This corresponds to the case where the conventional topology-driven policy (evaluated in Section III) is applied. On the contrary, the number of required labels mapped to the same destination network (*, dstnet) granularity is only 99. It is more than ten times less, compared to the case using the host-pair packet-stream granularity. Also, the cut-through ratio increases from 68 to 99%.

As shown in Figs. 9 and 11, the cut-through ratios for both incoming and outgoing packets are almost the same when the disconnect timer is more than 60 ms. Therefore, the disconnect-timer value should be 60 ms. The number of

Table IV shows that the number of required labels with (*, dstnet) granularity, for the packets coming into the AS, is

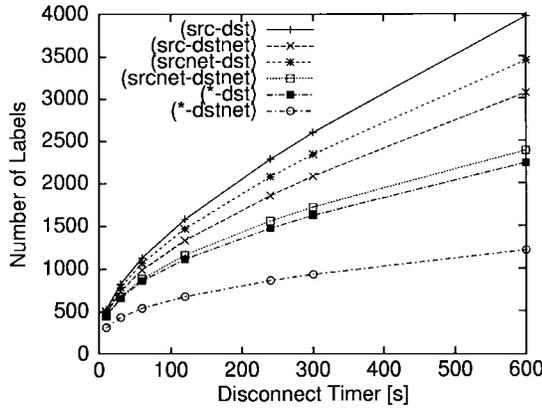


Fig. 12. Number of labels directed outside the AS using full route.

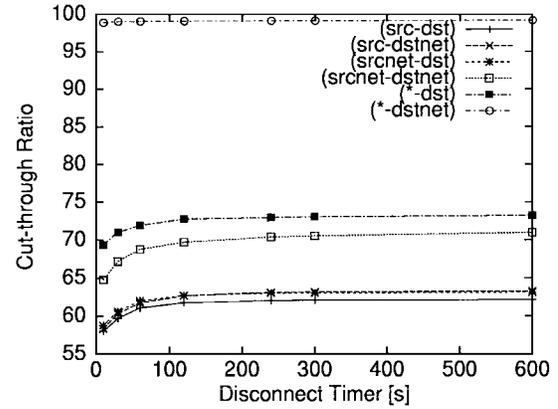


Fig. 15. Cut-through ratio directed inside the AS using full route.

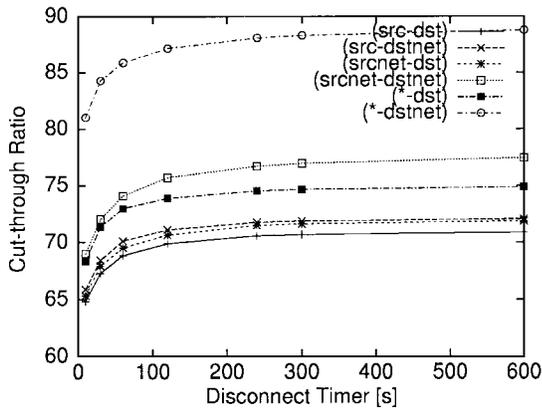


Fig. 13. Cut-through ratio directed outside the AS using full route.

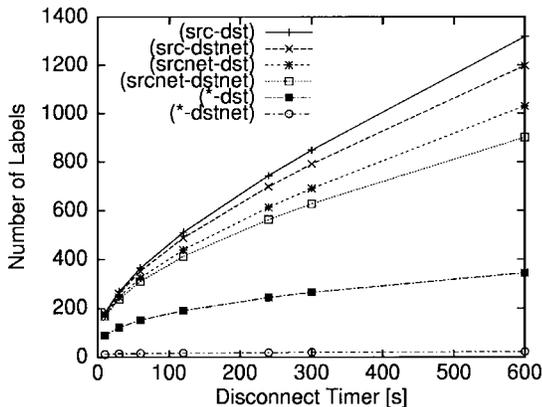


Fig. 14. Number of labels directed inside the AS using full route.

20 times less than that with (src, dst) granularity. Again, the cut-through ratio increases from 61 to 95%.

Furthermore, we evaluated the number of the required labels using the full-route routing table for packet forwarding toward outside the AS. Figs. 12 and 13 show the number of the required labels and cut-through ratio for the traffic directed to outside the AS. Figs. 14 and 15 show those for the traffic directed toward the AS. Identical to the previous evaluation, the disconnect-timer value is selected as 60 ms. The number of the required labels is shown in Tables V and VI.

Table VI shows the evaluation results for the incoming packets toward the AS. The number of required labels with

TABLE V
NUMBER OF LABELS AND CUT-THROUGH RATIO
DIRECTED OUTSIDE THE AS WITH FULL ROUTE

Granularity	Number of Labels	Cut-through Ratio
(src, dst)	1131	68
(src, dstnet)	990	70
(* , dst)	859	73
(srcnet, dstnet)	883	74
(* , dstnet)	542	86

TABLE VI
NUMBER OF LABELS AND CUT-THROUGH RATIO
DIRECTED INSIDE THE AS WITH FULL ROUTE

Granularity	Number of Labels	Cut-through Ratio
(src, dst)	365	61
(src, dstnet)	352	62
(srcnet, dstnet)	310	69
(* , dst)	152	72
(* , dstnet)	16	99

(src, dst) granularity that corresponds to the conventional traffic-driven policy is 1131. On the contrary, the number of required labels with (*, dstnet) granularity is 542. The number of required labels with (*, dstnet) granularity is about a half of that with (src, dst) granularity.

Table V shows the evaluation result for the outgoing packet flows from the AS. The number of labels for (src, dst) granularity that corresponds to the conventional traffic-driven policy is 365. On the contrary, the number of required labels with (*, dstnet) granularity is 16. The number of required labels with (*, dstnet) granularity is about 22 times less than that with (src, dst) granularity.

The results show that traffic-driven mapping with the flow aggregation (e.g., aggregated-packet flow with a destination network) significantly decreases the number of required labels and also increases the cut-through ratio.

V. CONCLUSIONS

This paper evaluates the number of required labels and the cut-through ratio in the LSR using actual Internet backbone traffic traces. The number of required labels depends on the label-mapping triggers, the scale of the network, and on the granularity of packet stream. With the conventional label-mapping policies, which are traffic-driven mapping with a destination-network packet stream or traffic-driven mapping with a host-pair packet stream, the required number of labels becomes large.

Therefore, this paper proposes and evaluates a new label-mapping policy (i.e., flow-aggregated traffic-driven mapping policy) to reduce the required number of labels and to increase the cut-through ratio. The proposed policy uses label mapping with the aggregated packet stream toward a specific destination network, triggered by the actual packet arrival belonging to the defined aggregated packet stream. By evaluating the use of actual Internet backbone traffic traces, it is shown that the proposed label-mapping policy will work well. For example, the required number of labels with the proposed label-mapping policy is reduced to only one tenth of that with the conventional policy; that is, 1131 labels with the conventional policy is reduced to only 99 labels with the proposed policy. Also, the proposed policy increases the cut-through ratio from 68 to 99%.

These evaluation results indicate that, for LSR's in the Internet backbone area, the proposed flow-aggregated traffic-driven label-mapping policy will be practically useful. Future work should include the evaluation of the required number of labels using other traffic traces (e.g., at Internet NAP) and other flow aggregation policies because the results depend on traffic traces. We think the proposed policy is useful in some scales of the network. In other scales of the network, the simple topology-driven policy or the simple traffic-driven policy may be practical and appropriate. Also, in some scales of networks, as well as the proposed policy, a hybrid scheme (e.g., the traffic-driven operation across AS's over the topology-driven operation within the AS) would be appropriate. Which label-mapping policy is appropriate for various scales of the networks should be further studied.

REFERENCES

- [1] P. Newman, T. Lyan, and G. Minshall, "Flow labeled: Connectionless ATM under IP," in *Eng. Conf., Network+Interop '96*, Las Vegas, NV.
- [2] P. Newman, W. Edwards, R. Hinden, E. Hoffman, F. C. Liaw, T. Lyon, and G. Minshall, "Epsilon flow management protocol specification for IPv4 version 1.0," May 1996.
- [3] Y. Katsube, K. Nagami, and H. Esaki, "Toshiba's router architecture extensions for ATM: Overview," Toshiba Corp., Kawasaki, Japan, Rep. IETF RFC2098, Feb. 1997.
- [4] K. Nagami, Y. Katsube, Y. Shobatake, A. Mogi, S. Matsuzawa, T. Jinmei, and H. Esaki, "Toshiba's flow attribute notification protocol (FANP)," Toshiba Corp., Kawasaki, Japan, Apr. 1997.
- [5] Y. Rekhter, B. Davie, D. Katz, E. Rosen, and G. Swallow, "Cisco system's tag switching architecture overview," Rep. IETF RFC2105, Feb. 1997.
- [6] P. Doolan, B. Davie, D. Katz, Y. Rekhter, and E. Rosen, "Tag distribution protocol," IETF Internet-Draft, May 1997.
- [7] A. Viswanathan, N. Feldman, R. Boivie, and R. Woundy, "ARIS: Aggregate route-based IP switching," IETF Internet-Draft, Mar. 1997.
- [8] N. Feldman and A. Viswanathan, "ARIS specification," IETF Internet-Draft, Mar. 1997.
- [9] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol label-switching architecture," IETF Internet-Draft, Mar. 1998.
- [10] S. Lin and N. McKeown, "A simulation study of IP switching," in *ACM SIGCOMM '97*.



Ken-ichi Nagami received the B.S. degree from Chiba University, Chiba, Japan, in 1990 and the M.S. degree from the Tokyo Institute of Technology, Japan, in 1992.

In 1992, he joined the Research and Development Center, Toshiba Corp., Kawasaki, Japan, where he has been working on communication systems. He is currently at the Information and Communications Systems Laboratory, Toshiba Corp. He is interested in Internet technology for high-speed networks.



Hiroshi Esaki (M'90) received the B.E. and M.E. degrees from Kyushu University, Fukuoka, Japan, in 1985 and 1987, respectively, and the Ph.D. degree from the University of Tokyo, Japan, in 1998.

In 1987, he joined the Research and Development Center, Toshiba Corp., Kawasaki, Japan, where he researched ATM systems. Since 1998, he has been an Associate Professor at the University of Tokyo, Japan and works for the WIDE project as a Board Member. From 1990 to 1991, he was a Residential Researcher at Bellcore, NJ and has researched high-speed computer communications. From 1994 to 1996, he was a Visiting Scholar at the Center for Telecommunications Research, Columbia University, New York. He is currently interested in a high-speed Internet architecture, including MPLS technology and mobile computing.



Yasuhiro Katsube (M'92) received the B.E. and M.E. degrees in electrical engineering from Keio University, Yokohama, Japan, in 1984 and in 1986, respectively.

In 1986, he joined the Toshiba Research and Development Center, Kawasaki, Japan, where he has been engaged in the research and development of ATM network systems. From 1992 to 1994, he was a Visiting Researcher at Bellcore, NJ. He currently works at the Information and Communications Systems Laboratory, Toshiba Corp., Tokyo, Japan.



Osamu Nakamura (M'82) received the S.B., M.S., and Ph.D. degrees in mathematics from Keio University, Yokohama, Japan, in 1983, 1985, and 1990, respectively.

From 1990 to 1993, he was an Assistant Professor at the University of Tokyo, Japan. Since 1993, he has been an Assistant Professor on the Faculty of Environmental Information, Keio University. He has been a Board Member of the WIDE Project since 1988. His research interests include Internet technologies, operating systems, high-speed networks, and network management.