IEICE TRANS. COMMUN., VOL.E86–B, NO.8 AUGUST 2003

PAPER *Special Issue on Photonic IP Network Technologies for Next Generation Broadband Access*

Traffic-Driven Optical IP Networking Architecture

Junichi MURAYAMA[†a)], **Takahiro TSUJIMOTO**[†], **Kenichi MATSUI**[†], **Kazuhiro MATSUDA**[†], ***and* Hiroshi ISHII**[†], ***Regular Members***

SUMMARY This paper proposes a traffic-driven optical IP networking architecture for service provider networks. Its design is derived from the optical GMPLS architecture, which provides high performance but is not scalable since both optical paths and IP routes need to be arranged in a mesh topology. To improve scalability, we first modified the configuration so that paths and routes can be arranged in a tree topology. However, this approach may degrade performance due to traffic concentration at each tree's root. To prevent such performance degradation, we further modified the architecture so that both cut-through optical paths and cut-through IP routes can be assigned reactively, according to traffic demand, and these can work together in cooperation. As a result, our architecture achieves both high performance and scalability, in that the whole network performance can be maintained without a massive increase in the number of optical paths and IP routes, even if the number of customer networks grows.

key words: *IP, GMPLS, optical path, overlay model, connectionless*

1. Introduction

As the Internet has spread widely, the ensuing traffic has increased explosively. To handle such traffic, the performance and scalability of the Internet needs to be improved. Here, performance means end-to-end throughput, while scalability reflects how many provider edge (PE) routers can be accommodated by a service provider (SP) network. A promising way to improve performance is to replace Internet Protocol (IP) forwarding by optical path cross-connection using generalized multi-protocol label switching (GMPLS) [1], [2]. In an optical GMPLS network, a GMPLS controller is attached to an optical path cross-connect (OXC). This controller performs IP routing and optical path signaling. Since IP packets are forwarded over broadband optical paths, high performance is achieved.

However, this solution does not improve scalability, because a 'connection-oriented' optical GMPLS network cannot accommodate a large number of PE routers when it simply emulates a 'connectionless' IP forwarding network. In the optical GMPLS network, in order to ensure full 'reachability' of all PE routers, optical paths need to be arranged in a mesh topology, regardless of the optical traffic demand on each. This may cause a massive increase in the number of optical paths. To reduce this number, it is effective to arrange them in a tree topology and to use electrical paths to maintain the reachability of PE routers [3], [4]. However, even in this solution, PE routers are forced to support a large number of IP routes. This means that IP networking scalability is not necessarily improved, even if optical networking scalability is improved.

To achieve both high performance and scalability in both the optical and IP layers, we propose a traffic-driven optical IP networking architecture. This architecture is based on the optical GMPLS architecture, but it is first modified so that optical paths and IP routes can be arranged in a tree topology. Then it is further modified so that both cut-through optical paths and cut-through IP routes can be assigned reactively, according to traffic demand, and in such a way that they can work together in cooperation.

In this paper, Sect. 2 describes the optical GMPLS architecture as a reference model and Sect. 3 summarizes its weak points, which need to be addressed in an improved design. Section 4 describes our proposed architecture and Sect. 5 evaluates its scalability. Section 6 provides a brief summary, by way of conclusion.

2. Reference Model

The optical GMPLS architecture is shown in Fig. 1, as a reference model. This model comprises customer networks, access networks, and an SP network.

A customer network is either an end-user network or another SP network. An SP network accommodates customer networks via access networks. An access network physically aggregates customer-side access lines into a provider-side access line. In addition, it logically provides point-to-point connectivity between customer edge (CE) routers and a PE router. In contrast, the SP network provides optical paths between PE routers. Here, an optical path means a tunnel created by optical cross-connections.

In the SP network, IP forwarding is replaced by optical path cross-connections according to a GMPLS-based signaling procedure. This means that optical paths can be routed using an IP address, and IP packets

---

Manuscript received November 25, 2002.
Manuscript revised February 25, 2003.

[†] The authors are with NTT Information Sharing Platform Laboratories, NTT Corporation, Musashino-shi, 180-8585 Japan.
a) E-mail: murayama.junichi@lab.ntt.co.jp

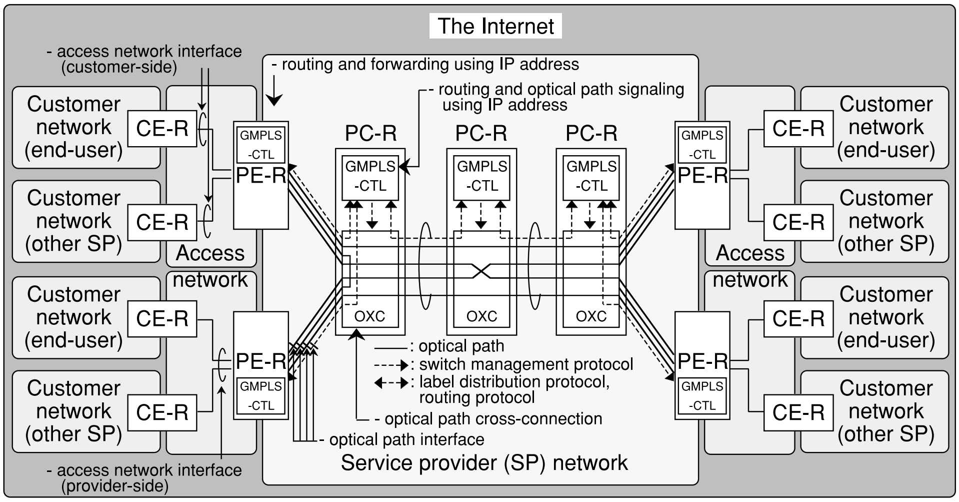

CE-R: customer edge router, PE-R: provider edge router, PC-R: provider core router, OXC: optical path cross-connect, GMPLS-CTL: generalized multi-protocol label switching controller.

Fig. 1 Reference model.

are forwarded via those paths.

Provider core (PC) routers located in the SP network are optical routers, each comprising an OXC and a GMPLS controller. The former handles the optical paths themselves, while the latter supports optical path signaling. PC routers are connected to each other by means of optical fiber links using wavelength-division multiplexing.

PE routers are composed of an IP forwarding device and a GMPLS controller. The former supports IP forwarding and the latter supports optical path signaling. PE routers are physically connected to a PC router by optical fiber links. In addition, they are logically connected with each other via optical paths. To maintain full reachability between PE routers, optical paths are arranged in a mesh topology.

Optical path signaling is performed using GMPLS controllers in a three-step procedure. First, IP routing information is exchanged between neighboring PC or PE routers using an IP routing protocol [5], [6] and IP routes are calculated within each router. Then, optical path cross-connection information is exchanged between neighboring routers using a label distribution protocol [7], [8], and optical path routes are calculated within each router. Finally, GMPLS controllers set up a cross-connection table within each router using a switch management protocol [9], [10].

After the signaling procedure has been implemented, PC routers simply cross-connect optical paths to forward IP packets. PE routers, on the other hand, in order to select optical paths or access connections, perform an IP forwarding procedure according to the destination IP address described in the IP packet header.

As the alternative conventional solution, electrical paths can replace some mesh-structured optical paths when the number of available optical paths is insufficient. Here, an electrical path means a tunnel created by electrical label switching. In order to establish electrical paths, some OXCs need to be replaced by electrical label switch routers, though the method of signaling can be almost the same between electrical path signaling and optical path signaling. With this approach, optical paths are also used to aggregate electrical paths.

3. Design Issues

As described in the following sections, scalability issues in the optical GMPLS architecture can be functionally classified into optical networking issues and IP networking issues.

3.1 Optical Networking Issues

There are three key optical networking issues, as shown in Fig. 2.

The first issue is the heavy processing load of the IP routing required for optical path signaling. Although IP packets are forwarded using optical paths, the IP address still needs to be used to determine an optical path route in the signaling procedure. In the Internet, a border router in an autonomous system (AS) should be able to handle more than a hundred thousand IP routes [11]. In an optical GMPLS network, a GMPLS controller, corresponding to an AS border router, should also handle the same number of IP routes, because optical path signaling is based on destination-based IP routing. Even if the number of PE and PC routers is very small, the number of IP routes cannot be reduced because the number of IP destinations is constant. To solve this problem, IP routing for the signaling should remain closed in the SP network because

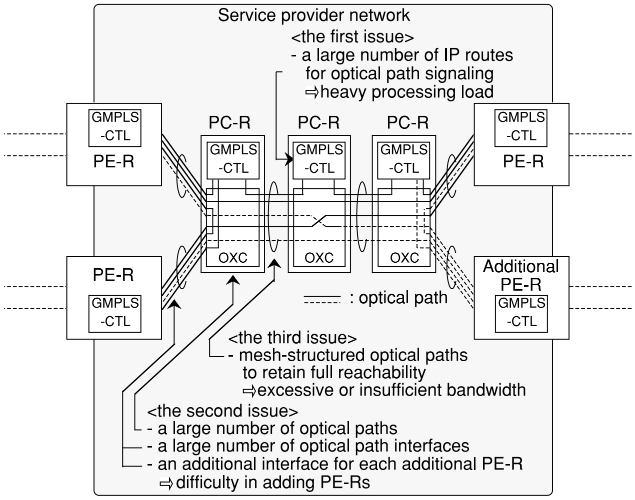

Fig. 2 Optical networking issues.

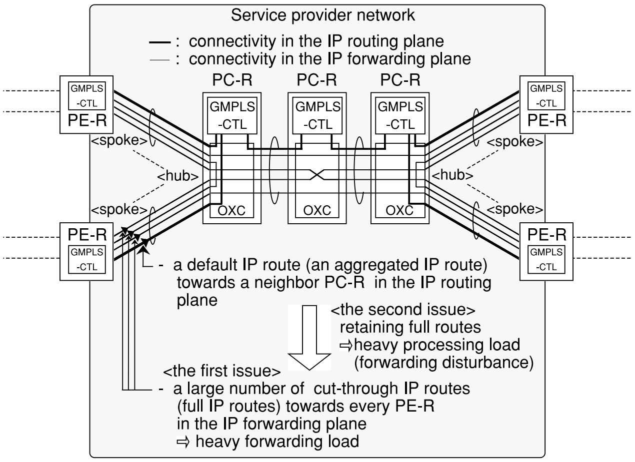

Fig. 3 IP networking issues.

optical paths are established only between PE routers.

The second issue is the large number of optical paths. To ensure full reachability among PE routers by means of optical paths alone, the paths would need to be arranged in a mesh topology, regardless of traffic demands. However, this arrangement may result in low utilization of each path. For example, each OXC would have to handle a large number of broadband optical paths even if the amount of traffic on some paths were small. In addition, PE routers equipped with only a few optical path interfaces could not be attached to a large-scale network composed of many PE routers, because the lack of interfaces would make it impossible to maintain full reachability. If some mesh-structured optical paths were replaced by electrical paths, which were then aggregated into tree-structured optical paths, the number of optical path interfaces required to retain full reachability could be reduced. However, PC routers might be still required to handle a large number of edge-to-edge electrical paths. To solve this problem, the reachability should be retained in a connectionless manner where edge-to-edge paths are not required.

The third issue is the degradation of effective performance in the SP network. Mesh-structured optical paths provide only a fixed bandwidth between PE routers regardless of performance requirements. On the other hand, tree-structured optical paths concentrate traffic at the tree's root, which may become a performance bottleneck. To solve this problem, optical paths should be distributed appropriately, according to IP traffic demands.

3.2 IP Networking Issues

There are two key IP networking issues, as shown in Fig. 3.

The first issue is the heavy processing load of IP forwarding. In the IP routing plane where IP routing table is created using an IP routing protocol, a PC router and several PE routers can be arranged in a hub-and-spoke topology. In this topology, the processing load of IP routing is concentrated in the hub PC router and this limits the number of spoke PE routers. On the other hand, each spoke PE router can simply route optical or electrical paths towards the hub PC router. Here, a simple IP route towards the hub is called a default IP route. This route is effective in reducing the processing load arising from IP routing in a PE router. However, in the IP forwarding plane where IP forwarding table is created using a label distribution protocol, in order to use every optical or electrical path arranged in a mesh topology, each PE router needs to divide a default IP route towards a PC router into multiple IP routes towards each PE router. This means that each PE router needs to be able to handle as many IP routes as are handled by AS border routers. In order to solve this problem, a default IP route should also be established in the IP forwarding plane.

The second issue is degradation of the whole network performance. When IP routes are arranged in a mesh topology to keep full reachability regardless of IP traffic demand, the processing load of managing those IP routes becomes heavy, and the forwarding performance may be degraded due to a shortage of shared processing resources. On the other hand, when a default IP route is established in the IP forwarding plane, each PE router uses only a single optical path and other paths are not used. This means that most optical paths are not used to carry traffic but are used only to retain full reachability. To solve this problem, IP routes should be appropriately distributed according to both IP traffic demands and optical path arrangements.

4. Architecture Design

The proposed architecture is composed of three elements: the optical GMPLS architecture, a traffic-

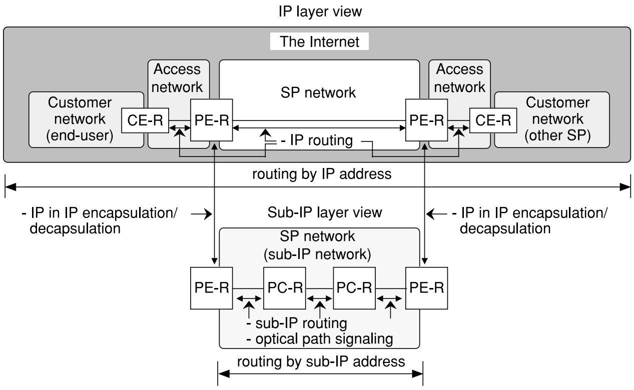

Fig. 4 Overlay networking.

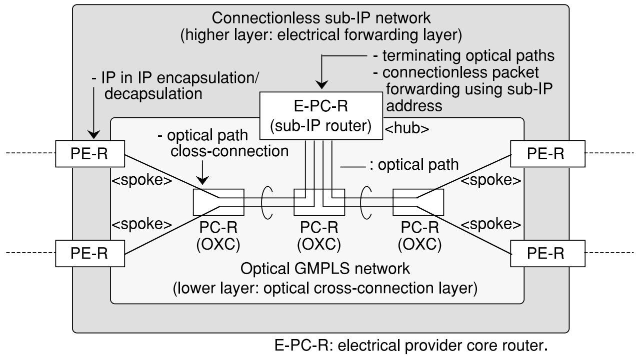

Fig. 5 Connectionless sub-IP forwarding.

driven optical networking architecture, and a traffic-driven IP networking architecture. The first, which is described in Sect. 2, is deployed to achieve high performance. The others, which are described in the following sections, are deployed to achieve scalability without performance degradation.

4.1 Traffic-Driven Optical Networking Architecture

The traffic-driven optical networking architecture uses three key technologies: overlay networking, connectionless sub-IP forwarding, and optical cut-through control.

As a first approach to reducing the processing load of IP routing required for optical path signaling, we deploy overlay networking. As shown in Fig. 4, the SP network is overlaid by the Internet, but its IP address space is separated from that of the Internet. Although Internet IP addresses are visible to customers, those of the SP network are concealed from customers and used only within the SP network. Here, the underlying SP network and concealed IP addresses are called the sub-IP network and sub-IP addresses, respectively. Thus, PE routers on the forwarding plane perform address mapping between the Internet and the sub-IP network so that a large number of visible IP addresses associated with destination hosts are aggregated into a concealed sub-IP address associated with the egress PE router accommodating these hosts. This procedure is performed within a PE router using IP routing information. However, optical path signaling is based on sub-IP routing and is independent of IP routing. In addition, the number of sub-IP routes within the SP network can be reduced to almost the same number as that of PE routers. Consequently, optical paths can be routed using a small number of sub-IP addresses in spite of a large number of Internet IP addresses.

As a solution to the second issue, the need to reduce the number of optical paths, we apply connectionless sub-IP forwarding to the optical sub-IP network. As shown in Fig. 5, the sub-IP network is logically divided into two layers: the higher layer is the electrical forwarding layer and the lower layer is the optical cross-connection layer. In the electrical forwarding layer, in order to retain full reachability among PE routers with a small number of optical paths, a sub-IP router is deployed as an electrical PC (E-PC) router. On the other hand, in the optical cross-connection layer, optical paths are established between an E-PC router and PE routers and these are arranged logically in a hub-and-spoke topology with an E-PC router at the hub. Here, the number of optical path interfaces in the hub E-PC router limits the number of spoke PE routers. When multiple E-PC routers are used to accommodate a large number of PE routers, they can be also distributed in a hub-and-spoke topology to limit the number of terminated optical paths at spoke E-PC routers.

Ingress PE routers perform IP-in-IP encapsulation [12]. The sub-IP address assigned to the egress PE router is resolved from the destination IP address in the packet header and the original IP packet is encapsulated into a sub-IP packet destined for the resolved sub-IP address. The output optical path is further resolved from the destination sub-IP address. The intermediate E-PC routers terminate optical paths and forward sub-IP packets along paths in a connectionless manner where the output optical path is resolved from the destination sub-IP address. Egress PE routers, which terminate optical paths, perform IP-in-IP decapsulation to terminate sub-IP routes. The extracted IP packets are forwarded towards CE routers according to the destination IP address.

Consequently, the number of optical paths in the SP network can be reduced to almost the same number as that of PE and PC routers. Furthermore, the number of optical path interfaces required for each PE router to retain full reachability can be reduced to only one. Although deploying a sub-IP router requires additional cost, it can reduce the required costs of both the required PE routers and the OXCs. As a result, the total required networking cost is reduced when a large number of PE routers and OXCs are deployed.

As a solution to the third issue, that is suppressing the performance degradation arising from traffic concentration at the hub, we apply optical cut-through

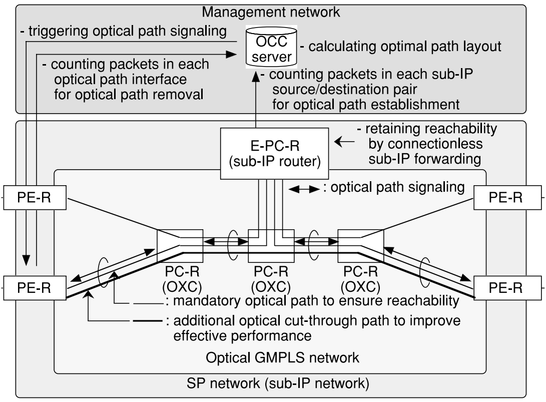

OCC server: optical cut-through control server.

Fig. 6 Optical cut-through control.

control in the sub-IP network. As shown in Fig. 6, an optical cut-through control (OCC) server is deployed to arrange appropriate cut-through optical paths according to sub-IP traffic demands. Here, cut-through optical paths are not paths which are essential in order to retain full reachability but additional paths, added to improve performance. The OCC server is connected to PC and PE routers via the management network in the management plane.

An E-PC router in the management plane classifies transit packets according to the source-destination pair in the sub-IP address and counts the number of occurrences of each pair. The OCC server collects these numbers from E-PC routers and identifies any source-destination pair whose frequency of occurrence exceeds the decision threshold for establishing a path. Then it triggers the ingress PE router associated with the source-destination pair to establish a cut-through optical path using GMPLS-based signaling.

To avoid running short of cut-through path resources, paths that are not being effectively used need to be released. Thus, a PE router in the management plane also counts the number of transit packets at each interface which terminating a cut-through optical path. The OCC server also collects these numbers from PE routers and identifies any path for which the number is smaller than the path release decision threshold.

When cut-through paths cannot be established due to a shortage of path resources, the OCC server calculates the optimal path arrangement for minimizing the traffic load on the E-PC routers and may replace some paths with new ones depending on the result. In order to perform frequent path reassignment, overlay networking technology is necessary to reduce the processing load of optical path signaling. As the result, cut-through optical paths are optimally arranged according to sub-IP traffic demands, so limiting any performance degradation of the sub-IP network.

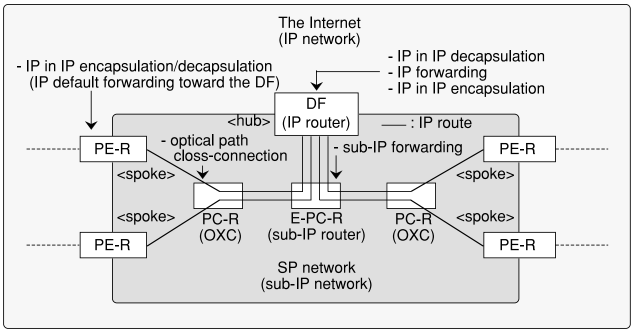

DF: default forwarder.

Fig. 7 Hub-and-spoke IP routing.

4.2 Traffic-Driven IP Networking Architecture

The traffic-driven IP networking architecture deploys two key technologies: hub-and-spoke IP forwarding and IP cut-through control.

As a first solution, we use hub-and-spoke IP forwarding to reduce the processing load of IP forwarding. As shown in Fig. 7, PE routers are logically arranged in a hub-and-spoke topology. Here, the hub PE router is called the default forwarder (DF) and a route from a PE router to the DF is called a default IP route. On the other hand, a route from a spoke PE router to the hub PC router in the sub-IP layer is called a default sub-IP route. In hub-and-spoke IP forwarding, the processing load of IP routing at the hub DF limits the number of spoke PE routers. When multiple hub-and-spoke topologies are established using multiple DFs, they can also be distributed in a hub-and-spoke topology to reduce the processing load of IP routing at the spoke DFs. IP routes within the SP network can be controlled using an IP routing protocol between PE routers including DFs. IP routes towards other SP networks can be also controlled using an IP routing protocol between PE routers and CE routers. As a result, IP routing information towards other SP networks is also concentrated from PE routers to the DF. In a hub-and-spoke topology, only a single default IP route is required for each spoke PE router, while, in a mesh topology, full IP routes are required for each PE router. Consequently, hub-and-spoke IP forwarding can limit the number of IP routes required to retain full reachability and thus can limit the processing load of PE routers.

As a solution to the second issue, the need to avoid performance degradation resulting from traffic concentration at a DF, we use IP cut-through control, which is composed of redirection control and purge control.

Redirection control is shown in Fig. 8, and is used to establish a cut-through IP route in which any DFs between the ingress and egress PE routers are bypassed. The DF sends a redirection message to the ingress PE router when it forwards a packet from the ingress PE to the egress PE. This message contains IP-layer for-

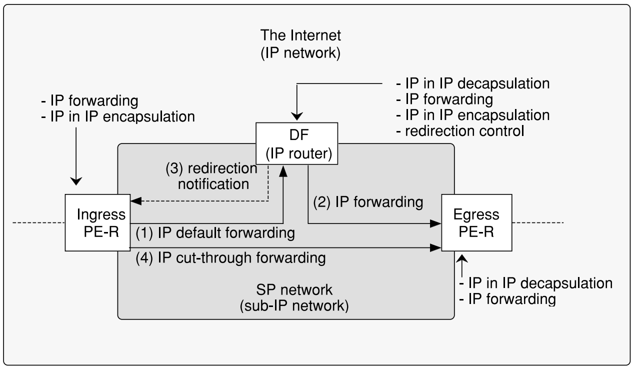

Fig. 8 IP redirection control.

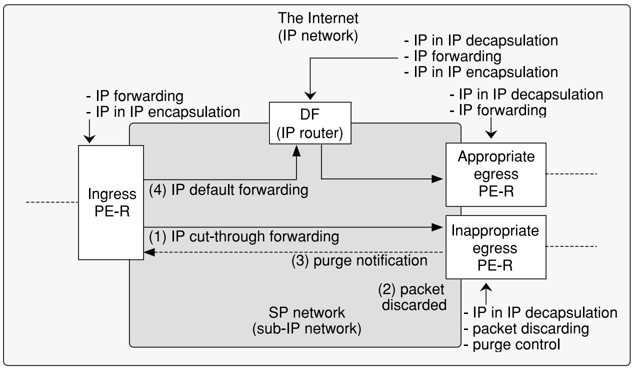

Fig. 9 IP purge control.

warding information consisting of the IP address prefix and sub-IP address associated with the destination IP subnet and egress PE router, respectively. An ingress PE router that receives a redirection message adds the contained information to its IP forwarding table as a redirection entry. Thus, the succeeding packets traveling towards the same destination are forwarded via a cut-through IP route in accordance with this entry. Since this control requires certain conditions in which IP forwarding information should be shared between the DF and the ingress PE routers, it can be deployed only when connectionless sub-IP forwarding is available in the SP network.

Purge control, as shown in Fig. 9, is used for an egress PE router to release a cut-through IP route. The egress PE router sends a purge message to the ingress one when an egress PE router cannot resolve an access connection towards the CE router associated with the destination IP address. This message contains only the destination IP address. The ingress PE router that receives a purge message removes the redirection entry associated with the contained information from its IP forwarding table. Thus, succeeding IP packets traveling towards the same destination are forwarded via the default IP route towards the DF. Finally, they are forwarded via the appropriate cut-through IP routes towards the egress PE router because the DF sends a redirection message again. The ingress PE router can also release a cut-through IP route that has been used rarely by simply removing the corresponding redirection entry.

Since a default IP route retains full reachability for any destination, cut-through IP routes can be established according to IP forwarding demands. In addition, when a significant number of IP packets are forwarded via a cut-through IP route, optical cut-through control is also triggered. As a result, cut-through IP routes and cut-through optical paths are established cooperatively according to IP traffic demands. This allows a satisfactory performance to be retained for the whole network, even if the number of IP destinations and PE routers grows.

5. Scalability Evaluation

We evaluated the scalability of the proposed architecture from the viewpoint of optical and IP networking scalability.

5.1 Optical Networking Scalability

We calculated the number of optical path interfaces required for each PE router and the number of optical paths required in the SP network. An optical networking model shown in Fig. 10 was used for the calculation. We assumed that the number of access network interfaces at each PE router (Ia) is constant regardless of the number of PE routers in the SP network (Ne) and that the interface bandwidths are the same among optical path interfaces and access network interfaces. We also assumed that the maximum bandwidth demand between any pair of PE routers is less than the bandwidth of the optical path interfaces.

First, we calculated the number of optical path interfaces required for each PE router. In the optical GMPLS architecture, optical paths are simply arranged in a mesh topology. Thus, the number of optical path interfaces required for each PE router (Ic) is given by

$$Ic = Ne - 1 \qquad (1)$$

This means that Ic should be increased linearly as Ne increases. In practice, however, Ic cannot be increased in an unlimited manner, so Ne should be kept within the limits of Ic. On the other hand, in our architecture, optical paths are first arranged in a tree topology and cut-through paths are arranged additionally according to sub-IP traffic demands. Here, in each PE router, the number of optical path interfaces can be less than or equal to Ia and Ic in order to increase a utilization ratio of each optical path. Thus, the maximum number of optical path interfaces (Ip_max) is given by

$$Ip\_max = \min\{Ia, (Ne - 1)\} \qquad (2)$$

This means that Ip_max can be practically constant

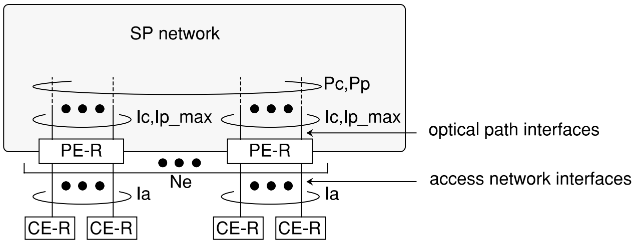

Fig. 10 Optical networking model.

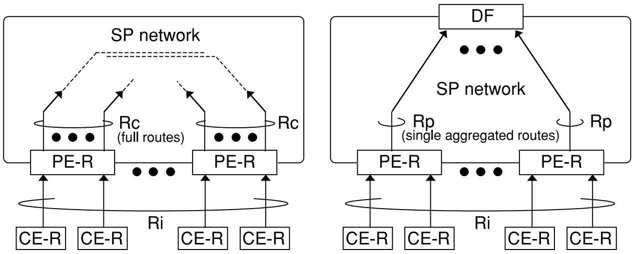

(A) Conventional network model (B) Proposed network model

Fig. 11 IP networking models.

even as Ne increases. Thus, the number of PE routers in the SP network (Ne) can be increased regardless of the maximum number of optical path interfaces in each PE router (Ip_max).

Second, we calculated the number of optical paths in the SP network. From the results of the first evaluation, the number of paths in the optical GMPLS architecture (Pc) and our architecture (Pp) are given by

$$Pc = \frac{(Ne-1)Ne}{2} \tag{3}$$

$$Pp = \frac{min\{Ia, (Ne-1)\}Ne}{2} \tag{4}$$

This means that Pp is smaller than or equal to Pc regardless of Ne. In addition, the difference between them increases as Ne increases. However, the performances achieved are almost the same because, in our architecture, cut-through optical paths are assigned according to sub-IP traffic demands, while, in the optical GMPLS architecture, an excessive number of paths are used with low utilization. As a result, our architecture is superior to the optical GMPLS architecture in that the whole network performance can be maintained without a massive increase in the number of optical paths, even if the number of PE routers grows.

5.2 IP Networking Scalability

To evaluate IP networking scalability, we calculated the number of IP routes required for each PE router to retain full reachability for any IP destination in the IP forwarding plane. An IP networking model shown in Fig. 11 was used for the calculation.

In the optical GMPLS architecture, a default IP route in the IP routing plane is divided into a number of IP routes in the IP forwarding plane. Since divided IP routes are arranged in a mesh topology, the number of IP routes handled by each PE router (Rc) is the same as that handled by the Internet backbone routers (Ri); typically this is more than a hundred thousand routes [11].

$$Rc = Ri \quad (Ri > 100,000) \tag{5}$$

On the other hand, in our architecture, each PE router is required to handle only a single default route to achieve full reachability. Thus, the number of IP routes handled by each PE router (Rp) is only one.

$$Rp = 1 \tag{6}$$

Consequently, Rp is always less than Rc regardless of the number of PE routers. In addition, their performances are almost the same when IP routes are stabilized. This is because IP forwarding is performed using forwarding caches according to IP traffic demands in order to improve practical performance. In the optical GMPLS architecture, forwarding caches are internally controlled in each PE router, while in the proposed architecture, they are controlled network-wide by the IP cut-through control. Thus, our architecture is superior to the optical GMPLS architecture in that the whole network performance can be maintained without an increase in the number of IP routes, even if the number of IP destinations grows.

6. Conclusion

Our architecture is composed of three elements: the optical GMPLS architecture, a traffic-driven optical networking architecture, and a traffic-driven IP networking architecture. The first is promising for achieving high performance, but is not scalable. To improve its scalability without performance degradation, we applied the second and third elements. In our architecture, sub-IP and IP layer reachability are retained using tree-structured optical paths and IP routes, respectively. This is effective in achieving scalability. In addition, sub-IP and IP layer performances are improved by using cut-through optical paths and cut-through IP routes, respectively. For both these, cut-through control is performed reactively, according to traffic demands, and the two work closely with each other. This is effective in maintaining high performance economically. Our scalability evaluation showed that this architecture is superior to the original optical GMPLS architecture in that the whole network performance can be maintained without a massive increase in the number of optical paths and IP routes, even if the number of customer networks grows. As the result, our architecture achieves both high performance and scalability.

Acknowledgments

This research was supported by a grant from the Telecommunications Advancement Organization of Japan (TAO).

References

[1] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol label switching architecture," IETF RFC3031, Jan. 2001.

[2] E. Mannie, "Generalized multi-protocol label switching (GMPLS) architecture," IETF Internet Draft, draft-ietf-ccamp-gmpls-architecture-03.txt, Aug. 2002.

[3] A. Banerjee, J. Drake, J.P. Lang, B. Turner, K. Kompella, and Y. Rekhter, "Generalized multiprotocol label switching: An overview of routing and management enhancements," IEEE Commun. Mag., vol.39, no.1, pp.144–150, Jan. 2001.

[4] K. Sato, N. Yamanaka, Y. Takigawa, M. Koga, S. Okamoto, K. Shiomoto, E. Oki, and W. Imajuku, "GMPLS-based photonic multilayer router (Hikari router) architecture: An overview of traffic engineering and signaling technology," IEEE Commun. Mag., vol.40, no.3, pp.96–101, March 2002.

[5] J. Moy, "OSPF version 2," IETF RFC2328, April 1998.

[6] K. Kompella and Y. Rekhter, "OSPF extensions in support of generalized MPLS," IETF Internet Draft, draft-ietf-ccamp-ospf-gmpls-extensions-08.txt, Aug. 2002.

[7] L. Andersson, P. Doolan, N. Feldman, and B. Thomas, "LDP specification," IETF RFC3036, Jan. 2001.

[8] O.A. Magd, S. Ballare, E. Tempest, R. Jain, L. Jia, B. Rajagopalan, R. Rennison, Y. Xu, and Z. Zhang, "LDP extensions for optical user network interface (O-UNI) signaling," IETF Internet Draft, draft-ietf-mpls-ldp-optical-uni-01.txt, July 2001.

[9] P. Newman, W. Edwards, R. Hinden, E. Hoffman, F. Ching Liaw, T. Lyon, and G. Minshall, "Ipsilon's general switch management protocol specification version 1.1," RFC 1987, Aug. 1996.

[10] J.K. Choi, M.H. Kang, J.Y. Choi, G.M. Lee, and J.U. Um, "General switch management protocol (GSMP) v3 for optical support," IETF Internet Draft, draft-ietf-gsmp-optical-spec-00.txt, Oct. 2002.

[11] http://www.potaroo.net/bgp/

[12] A. Conta and S. Deering, "Generic packet tunneling in IPv6 specification," IETF RFC 2473, Dec. 1998.

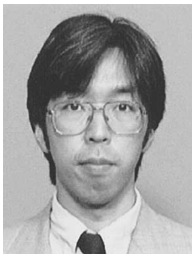

Junichi Murayama received the B.E. and M.E. degrees from Waseda University, Tokyo in 1989 and 1991, respectively. Since joining Nippon Telegraph and Telephone Corporation (NTT) in 1991, he has been engaged in research and development of ATM networks, large-scale IP networks, and IP VPN service platforms. He is now a senior research engineer in the Secure Communication project of NTT Information Sharing Platform Laboratories.

Takahiro Tsujimoto received his B.S. degree in computer science from Tokyo Institute of Technology in 1998. He earned his M.S. degree in computer science from the University of Tokyo in 2000 and joined Nippon Telegraph and Telephone Corporation (NTT) the same year. In 2003, he joined Nokia Japan. His research interests cover the fields of advanced networking technology, especially issues pertaining to provider provisioned virtual private networks.

Kenichi Matsui received the B.E. degree in information engineering and the M.S. degree in information sciences from Tohoku University, Japan in 1995 and 1997, respectively. He joined Nippon Telegraph and Telephone Corporation (NTT) in 1997. He works in NTT Information Sharing Platform Laboratories, where his study focuses on IP networking technology. His research interests include traffic engineering for optical IP networks, on-demand QoS management, and managed IP multicast platforms. He is a member of the IPSJ and the IEEE Computer Society.

Kazuhiro Matsuda received B.E. and M.E. degrees in electronic engineering from Hokkaido University, Japan in 1983 and 1985, respectively. Since joining Nippon Telegraph and Telephone Corporation (NTT) in 1985, he has been involved in LSI CAD systems, the design of high-speed protocol processing LSIs, and managed L2/L3 VPNs. He is now a senior research engineer, and supervisor in the Secure Communication project of NTT Information Sharing Platform Laboratories. He is a member of the IEEE Computer Society.

Hiroshi Ishii received B.E. and M.E. degrees in communication engineering from Osaka University in 1977 and 1979, respectively, and a Ph.D. degree in engineering from University of Tsukuba in 2001. At Nippon Telegraph and Telephone Corporation (NTT) Laboratories from 1979 to 1993, he was engaged in CCITT (ITU-T) standardization of ISDN protocols, and R&D on ATM switching systems, TINA systems, and managed IP networking systems. He is now a professor at Department of Communications Engineering, Tokai University. His research interest encompasses telecommunication and information networking systems and protocols. Dr. Ishii is a member of IPSJ and a Senior member of IEEE.