Study on a Joint Multiple Layer Restoration Scheme for IP over WDM Networks

Yang Qin, Lorne Mason, and Ke Jia, Nanyang Technological University

Abstract

In this article we investigate the problem of restoration scheme for IP over WDM networks. Network reliability gains importance with the huge volume of traffic carried by such networks. Providing survivability at the optical layer is inherently attractive, but raises many questions and challenges, given the characteristic of optical aggregated lightpath and relatively coarse traffic granularity. The emergence of MPLS and its extension, MPλS, opens up new possibilities for developing simple integrated protection/restoration schemes that can be coordinated at both the IP and optical layers. This article first presents an overview of existing MPLS/MPλS recovery mechanisms. Then we propose a joint two-layer recovery scheme for IP-centric WDM-based optical networks where the optical layer will take the recovery actions first, and subsequently the upper IP layer initiates its own recovery mechanism, if the optical layer does not restore all affected services. A simulation-based analysis shows the benefits of the proposed two-layer recovery scheme over single-layer recovery schemes. We demonstrate the advantages of finer granularity in IP layer recovery and the effectiveness in speed on the optical layer. The impact of several network parameters on recovery performance is also studied in the article.

With the exponential growth of the Internet and the resulting traffic per host, an unprecedented shift has occurred in traffic pattern from fixed configured connection-oriented services (e.g., voice service) to dynamic connectionless IP services [1, 2]. Although there has been a slowdown in Internet growth from that of past recent years, it is expected that Internet traffic will continue to grow dramatically in the future. Wavelength-division multiplexing (WDM) technologies with substantially high bandwidth capacity are expected to play a dominant role in such networks. WDM technologies combine multiple signals, each at different carrier wavelengths, to increase capacity. Lightpaths are set up to provide end-to-end connections between optical cross-connects (OXCs).

IP over WDM is a simple example of a multilayer network, where the IP layer resides above an optical network. The Internet Engineering Task Force (IETF) has proposed multiprotocol label switching (MPLS) technology [3], where packets are forwarded based on appended labels. MPLS separates the routing decisions and forwarding of the data. Connection-oriented paths, called label switched paths (LSPs), are set up for connectionless IP packets. There are also current efforts to port MPLS to the photonic domain, resulting in multiprotocol lambda switching (MPλS) [4, 5]. The wavelength color is regarded as a label, and the label switching concept is used to provision wavelength-switched lightpaths.

Modern networks should be designed to be fault-tolerant. MPLS offers fast and efficient protection/restoration capability to provide network survivability. This fast protection/restoration capability is also a key feature of MPλS, inherited from MPLS. Thus, each layer in such a multilayer network provides its own recovery capability.

In this article we deal with the issue of providing survivability in such IP/WDM networks. A multilayer strategy is proposed, and simulations are performed to compare its performance with other single-layer survivability schemes, from both the capacity and restoration speed points of view.

The rest of this article is organized as follows. We provide a brief description of the existing protection and restoration schemes. We describe the proposed two-layer restoration scheme. We present and discuss simulation results. This is followed by a conclusion, and suggestions of areas where further work can be done in the last section.

Survivability of IP over WDM

Background on Network Survivability

Prior to performing recovery, the layer closest to the failure is responsible for detecting the failure. Network equipment communicates with each other to determine where failures have occurred and to notify other network equipment of failures.

Some spare capacity is needed for recovering traffic affected by failures. Depending on the different timescales in which the spare capacity is allocated, there are essentially two types of fault management techniques: protection and restoration. In protection, backup paths are established, and spare capacity is reserved for them at the time the working path is set up. In restoration, upon network failure, backup paths are estab-

0890-8044/03/$17.00 © 2003 IEEE

lished in real time while the spare capacity is allocated to them dynamically. Generally, protection may cost more resources, since it needs to pre-allocate spare capacity for pre-established backup paths to react to the failures. On the other hand, restoration may take longer to restore the connection, because real-time backup path establishment may involve dynamic route calculation and spare capacity allocation.

Protection and restoration have traditionally been addressed using two techniques: path switching and link switching [6].

Figure 1 shows path switching and link switching. In path switching, traffic is recovered along a new path between source and destination node pairs of each connection that traverses the failed links. In link switching, traffic is recovered around failed links.

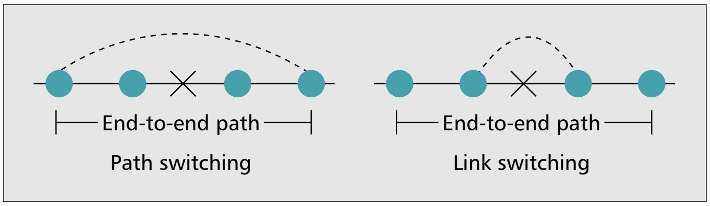

■ Figure 1. *Path switching and link switching.*

The MPLS Recovery Mechanism

The nature of IP implies that it reacts very slowly to network failures, and it is not desirable to employ IP layer recovery with MPLS. Therefore, some faster mechanisms for MPLS recovery have been proposed in [7–9].

In an end-to-end path protection scenario [7, 8], a pre-established backup LSP is set up from ingress label switched router (LSR) to egress LSR, which is physically disjoint from the working LSP. This pre-established backup LSP does not require any resources as long as the working LSP has not failed. When the working LSP fails due to the failure of a network component, the ingress LSR no longer forwards packets along this LSP, but switches over to the backup LSP.

A similar approach to path protection can be implemented on a link switching basis, which is called local protection. A backup LSP only spans a link (or a node) to protect this link (or node) [9]. The backup LSP originates in the protection switch LSR (PSL) and terminates in the protection merge LSR (PML) where the backup LSP is merged with the working LSP. If a working LSP spans several links, one backup LSP has to be set up for each link in the working LSP in order to protect the entire working LSP.

Path protection reacts to failures more slowly than local protection, since it takes a significant amount of time to notify the ingress LSR to switch over to the backup LSP. This leads to more packet loss. On the other hand, only one backup LSP is needed for one working LSP, and its global nature allows for less spare resource requirements.

A hybrid scheme named local loopback (also called fast reroute or alternative path) was also proposed. A backup LSP is provided in the opposite direction and concatenated to a physically disjoint LSP. It combines the best characteristics of both path and local protection schemes; see [9] for more detail.

Rerouting [7] is essentially a restoration mechanism, since it is based on real-time establishment of the backup LSP. The LSP is able to get the route of the backup LSP according to the automatically updated routing table after the failures.

The main disadvantage of rerouting is that the recovery time can be quite long, inheriting this feature from the IP routing protocol on which it relies. Therefore, some improvements are proposed to minimize the restoration time, such as explicit failure notification that accelerates the failure detection of the LSR and a set of precalculated reroutes used to reroute the time-critical traffic. Of course, rerouting has the advantage that it is able to deal with very complicated failure scenarios. In [9] the authors propose a novel rerouting mechanism named fast topology-driven constraint-based rerouting (FTCR) to mitigate some problems of the original rerouting scheme. The novelty of FTCR is that the first upstream LSR rather than the original LSR is responsible for rerouting.

Recovery in Optical Networks

The optical layer also can provide the resilience. References [6, 10] studied the resilience in optical networks within an ATM/SDH/WDM context. In the IP-over-WDM two-layer scenario, the MPLS concept has been extended to the optical domain via MPλS; thus, the MPLS recovery strategies can be adopted in the MPλS context in a straightforward manner [5, 11].

The protection mechanisms (i.e. path protection, local protection, and local loopback protection) are all applicable for an MPλS network, but restricted to the physical characteristics of optical networks. Four special considerations need to be included in this adaptation.

First of all, the backup lightpath cannot be established without allocating resources, which is different from MPLS protection. A wavelength is consumed by the backup lightpath once it is established. This leads to dedicated protection in the optical domain instead of shared protection in the MPLS domain. Since the number of wavelengths in a single fiber is limited, the problem becomes more critical.

The second issue is that there is no label stacking equivalent in MPλS. Thus, no statistical multiplexing is available, in contrast with MPLS, which allows statistical multiplexing between LSRs routed over the same link.

Third, we cannot merge the working and backup lightpaths into a single outgoing lightpath at the same bit rate. So, there are some problems for local protection.

Finally, since it is currently not economically feasible to place wavelength converters everywhere in the network, the utilization of spare wavelengths is low due to the wavelength continuity constraint.

Recently, Assi *et al.* proposed an approach with dynamic logical connection of lightpaths for IP over WDM networks in [12]. Gerstel *et al.* discussed a quantitive framework for best effort protection of the optical layer in [13]. Both were focus on the protection only on optical layer.

Because rerouting and FTCR do not need to pre-establish a dedicated lightpath, they do not suffer so much from the high-capacity cost problem as do the protection mechanisms. Also, we need to tear down the downstream part (from the failure) of the affected working lightpath, in order to release the capacity it occupied. This process may be rather complicated due to the failures in the network. Furthermore, the route of the rerouted lightpath is also constrained by the wavelength continuity problem, if no wavelength converter is available along this backup route. Compared to MPλS protection, MPλS restoration performs much better from a capacity point of view. On the other hand, it is much slower than protection.

Joint Two-Layer Restoration Scheme

In this kind of IP-over-WDM network, both the IP and optical layers have some recovery capability, and a single-layer recovery scheme may be deployed in either layer. An important question arises. In which layer should one provide net-

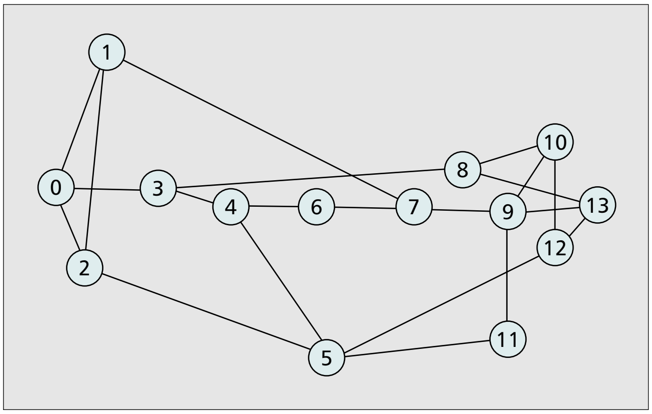

■ Figure 2. *The network under study: topology of NSFNET.*

work survivability? An intuitive possibility is to provide resilience in the optical layer, since the recovery actions are performed on the coarsest granularity. Although failures at the physical layer and optical transport equipment failures can be recovered at the IP/MPLS layer as well, in the optical layer a single element failure is treated and fewer recovery actions are taken. In addition, failures do not propagate through multiple layers before triggering any recovery actions. However, the optical layer is not always able to resolve problems caused by a failure that effects a higher layer. For example, when an OXC fails, the optical layer can only recover lightpaths transiting the failed OXC; hence, the LSR(s) residing on top of the failed OXC becomes isolated, and thus only the IP layer is able to restore all traffic transiting this isolated LSR.

Since the optical layer cannot recover, all kinds of failures in the network, the IP layer has no other choice than to take over the recovery job. Although providing resilience at the IP layer can deal with failures occurring at either the IP or optical layer, it suffers from the fact that many recovery actions are needed, due to the finer granularity of the LSPs at the IP layer. Furthermore, a single element failure in the optical layer will result in typically complex secondary failures in the virtual topology. However, the finer granularity also allows the differentiation between individual LSPs, based on their service class with different reliability requirements. That is to say, the IP layer may restore critical high-priority traffic before any action is taken on low-priority traffic, which is unachievable for the optical layer.

Recovery at either the optical or IP layer has its own pros and cons. A more advanced possibility is to provide recovery at both the IP and optical layers (i.e., a multilayer recovery strategy). This multilayer survivability strategy has been studied in previous work [14, 15]. The initial motivation of the multilayer strategies is to avoid duplicated survivability functionality at multiple layers that can yield reduced resource utilization and routing instabilities.

With the idea of multilayer recovery, we propose a novel joint two-layer restoration scheme for MPLS-based IP-over-WDM networks. Recovery work is done sequentially in a bottom-up fashion. The recovery starts in the optical layer, which is closest to the failures. If the optical layer is unable to restore all the affected traffic, the IP layer takes over the recovery actions.

The spectrum of recovery mechanisms, as mentioned earlier, can be deployed in either layer; one is free to deploy any of them in any layer of the network, and any combination of recovery technologies in different layers.

Some interworking mechanisms exist for handing over the responsibility for recovery from one layer to another [15]. One is called a holdoff timer. A timer is set at the moment the optical layer starts attempting to restore the traffic, and the IP layer takes over the recovery when this holdoff timer times out. Alternatively, in the recovery token method, the optical layer sends the explicit recovery token to trigger IP layer recovery.

We have some considerations in designing the scheme. Even though protection need hundreds of milliseconds and restoration time is on the order of seconds or even minutes in higher layers, restoration has high utilization and flexibility. We tend to use restoration instead of protection; hence, the speed of restoration is a very important factor. Regarding path switching and link switching, we tend to use path switching due to the limited wavelength resources around the failed link, especially since some networks have wavelength continuity constraints.

Hence, in our study, we are more interested in deploying rerouting (i.e., restoration) in both the IP and optical layers, since rerouting, especially optical rerouting, is very efficient from capacity and cost points of view. Some spare wavelengths are reserved that provision spare capacity to the optical layer used in rerouting. Upon receiving the failure notification message, the optical layer recovery actions will be carried out at the ingress OXCs of the affected lightpaths. A heuristic algorithm is implemented in the simulator for calculating the reroute path as follows.

First, an undirected graph G_l is constructed based on the physical topology and wavelength availability status where an edge connecting vertices $<s, d>$ in G_l denotes there are spare wavelength(s) in the link between corresponding physically adjacent node pair $<s, d>$.

Second, for rerouting an affected lightpath from node s to node d, an alternative path with least hops from s to d can be found in G_l using Dijkstra's algorithm, which implies that a minimal number of spare wavelengths are used for rerouting this affected lightpath.

Note that wavelength conversion is assumed in every node in the network here, and when the wavelength converter is not available everywhere, a similar procedure has to be duplicated for each wavelength.

One lightpath consists of many LSPs. If the failed lightpath(s) are successfully rerouted in the optical layer, those LSPs carried by the affected lightpath(s) are also successfully rerouted, but they are not aware of the rerouting actions taken in the optical layer. They only see the unchanged virtual topology, and the routes for those LSPs are also unchanged in the virtual topology.

But when the optical layer fails to reroute the affected lightpath(s) in the optical layer, the IP layer has to take charge of restoring the LSPs utilizing the affected lightpath(s). A lightpath has a fixed bandwidth, and once the lightpath is established the wavelengths are occupied, even if the fixed bandwidth is more than the LSP triggering the lightpath setup requires. Thus, some spare capacity usually exists in the working lightpaths, and it is possible to reroute the LSPs using this spare capacity in the virtual topology. Furthermore, new lightpaths could be established somewhere in the network for rerouting the LSPs, if it is necessary and free wavelengths are available. Thus, a problem similar to the virtual topology reconfiguration problem is involved here. Various heuristic algorithms have been proposed for this kind of virtual topology reconfiguration problem, which is NP-complete [16]. They can be ported into our scheme with some modification. Here we do not need global reconfiguration of virtual topology, that is, links in the existing virtual topology will not be removed, and only new lightpaths need to be added to form the new virtual topology. Likewise, a heuristic algorithm is implemented in the simulation in the following steps:

OXC connection matrix reconfiguration time	25 ms
Holdoff timer	150 ms
LSR forwarding table reconfiguration time	50 ms
Link propagation delay	0.3 km/µs

■ Table 1. *Timing parameter for restoration.*

Step 1: Construct a directed graph G_v, representing the current virtual topology. Each edge in the graph represents the existing lightpath that has enough spare bandwidth for the rerouted LSP, and they are tagged "exist."

Step 2: Then add an edge in each direction between the vertices where there are spare wavelengths in the link connecting the corresponding nodes in the physical topology, and tag those edges "new."

Step 3: So now the problem of minimizing used wavelengths is equivalent to the problem of finding a path with least "new" edges (i.e., the shortest path in G_v). The Dijkstra algorithm is used to find a shortest path on an unweighted graph.

Step 4: For the path found in step 3, if the consecutive edges are the "new" edges, combine them to be one "new" edge. The resulting path is the route for the rerouted LSP. Note that if the node is without wavelength conversion, only "new" edges with same the wavelength can be concatenated.

One lightpath usually carries many LSPs, which means a lightpath has coarser granularity than an LSP. The finer granularity of the LSP leads to better restoration ability, since each LSP is restored individually in the IP layer rather than all LSPs carried by the failed lightpath being restored together in the optical layer.

Simulation Results and Discussion

Simulation Details

For evaluation of the proposed multilayer recovery strategy and interworking mechanisms, a simulation platform was developed based on OMNeT++ [17]. The results presented in this article are based on the topology (Fig. 2) extracted from NSFNET, containing 14 nodes and 21 links, leading to an average nodal degree of 3.0.

Each node in the topology consists of an OXC and an attached LSR. They are connected with a single bidirectional fiber carrying $C = 10$ wavelengths, and each wavelength has OC-12 capacity (i.e., 622 Mb/s). We assume that wavelength converters are deployed at all network nodes in most of the simulations. We will discuss the impact of the equipment of wavelength converters in one case. A single link failure scenario is assumed, which is the most probable type of failure in networks.

We assume that a network design has planned optimal bandwidth allocation and traffic loading on the network. The spare wavelength coefficient θ denotes the ratio of spare wavelengths to all wavelengths in a fiber link. For example, by default, $\theta = 0.2$ means eight wavelengths are used as working wavelengths and two as spare wavelengths in each fiber link.

Data traffic is flow-based, and one LSP will be established for one traffic flow (i.e., no traffic aggregation at the MPLS level). Packets are generated according to a Poisson process based on the traffic demand matrix with rate

$$\lambda = \sum_{i=1}^{\eta} \lambda_i,$$

where λ_i is the packet arrival rate for flow i. The size of packets follows a exponential distribution with mean 2 kbytes.

The granularity coefficient is η, which denotes the number of LSPs carried by one wavelength. And the traffic load density ρ depicts how much wavelength capacity is occupied by data traffic. Reference values are $\eta = 5$ and $\rho = 0.75$.

A recovery token is usually used as the mechanism to trigger MPLS layer recovery. The timing parameters used in the simulations are shown in Table 1.

Comparison of Recovery Schemes

The performance of two-layer restoration with recovery token interworking is compared with a single IP/MPLS layer restoration approach in Fig. 3. The graphs show the average recovery ratio of affected traffic volume.

The most noticeable characteristic of the graphs is the two-step curve in the two-layer recovery. The optical layer rerouting is rather fast. With the given network and parameters, optical recovery takes up to 40 ms and reaches an approximate mean recovery ratio of about 65 percent . Then the affected LSPs that cannot be rerouted in the optical layer will be rerouted in the IP/MPLS layer, triggered by the explicit recovery token without any delay. The MPLS rerouting finished in less than 350 ms. Totally, 91 percent affected traffic is recovered at the end of the MPLS restoration interval. The single IP/MPLS layer restoration has a generally smooth curve, which means the recovery speed is slower than in the first case. Around 93 percent of affected traffic is recovered after 800 ms. This is because a much higher number of LSPs have to be individually rerouted in the IP/MPLS layer. Thus, two-layer restoration generally has a better recovery performance, since a majority of failed LSPs are recovered with a coarse granularity at high speed, which leads to less traffic loss and high throughput.

Another interesting feature is that the maximum restoration ratio is almost the same for both approaches. The main reason here is that the restoration (rerouting) scheme is used for both layers. On the other hand, we can expect that if the protection scheme is deployed in the optical layer, a significant decrease in mean recovery ratio will be noticed due to the dedicated backup lightpath of MPλS protection.

Influence of Granularity Coefficient η

The granularity coefficient η represents the number of LSPs in each lightpath. Larger η implies finer granularity of the

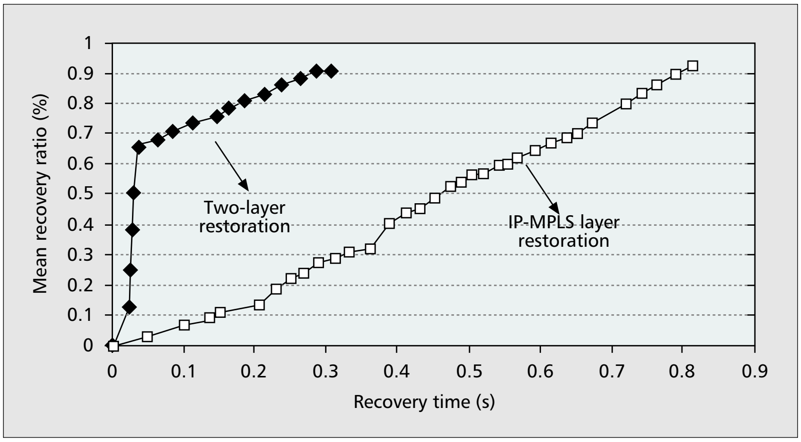

■ Figure 3. *Performance comparison for different recovery strategies.*

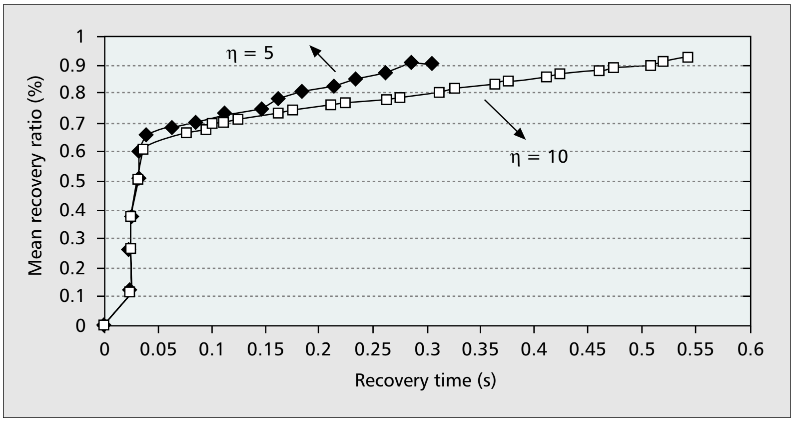

■ Figure 4. *Performance comparison for varying η value.*

LSPs. The results presented in Fig. 4 give recovery performance for η = 5 and η= 10, respectively.

With the same traffic demand, when η = 10, each LSP carries less traffic than it does when η = 5 on average. Hence, the MPLS layer restoration speed for η = 10 is slower than that for η = 5, while for optical layer restoration it is almost the same. It is clearly shown in Fig. 4 that the second phase of the curves have different slopes, while the first phase of the curves overlap. Another noticeable feature is that the curve with η = 10 has a slightly higher maximum recovery ratio. This result shows that finer granularity leads to better recovery ability but slower recovery speed, as we expected.

Handoff Timer vs. Recovery Token

The holdoff timer mechanism aims to simplify the implementation in the real network. The IP/MPLS layer recovery should be activated after the holdoff time elapses if optical layer recovery fails.

In Fig. 5 the performance of the recovery token strategy is compared to a holdoff timer interworking strategy. As can be seen from the graphs, recovery in the IP/MPLS layer using the holdoff timer starts indeed about 150 ms later than recovery using a recovery token. This holdoff delay worsens the recovery performance in situations when optical layer recovery fails. The network suffers more packet loss and lower throughput.

Influence of Wavelength Convertibility

In this section we study the impact of wavelength convertibility on system performance. Since wavelength converters are rather expensive, this study is important and necessary.

Figure 6 shows that recovery performance of the proposed two-layer restoration scheme degrades significantly in a non-wavelength-convertible network. With the same amount of spare capacity in both layers, optical layer rerouting is only able to recover up to 27.5 percent of affected traffic in around 25 ms, while the IP/MPLS layer can reach a maximum recovery ratio of about 92 percent in 650 ms. We can see that the wavelength continuity constraint has a major effect on optical layer rerouting. Without any wavelength conversion, a (backup) lightpath has to be set up between two nodes on the same wavelength. This greatly reduces the recovery speed . As we can see from Fig. 6, the network can reach 50 percent recovery ratio in 50 ms with wavelength conversion, while it needs more than 250 ms to reach 50 percent recovery ratio without it.

Conclusions and Future Work

The main goal of this article is to investigate survivability in IP/MPLS directly over a MPλS-enabled WDM-based multilayer network. In a multilayer environment, an important question is "which layer is responsible for network recovery?" It was found that single-layer recovery had its own drawbacks. Recovery at the IP/MPLS layer results in slower recovery and recovery at optical layer suffers from the fact that it cannot resolve any problem due to failure in or affecting a higher layer. A multilayer recovery scheme that combines the advantages of single-layer recovery would be appropriate to provide network survivability. Therefore, the question of how to coordinate single-layer recovery actions arises. In this article, an escalation recovery scheme is proposed. Recovery starts from the optical layer, and the IP/MPLS layer is activated if the optical layer cannot restore all affected traffic. Both the recovery mechanisms at each layer and interworking strategies are studied in the article.

The outcome of the current work shows that the proposed two-layer recovery scheme was superior to the traditional single layer recovery scheme, and was able to provide network survivability in MPLS-based IP-over-WDM networks. The simulation results show that traffic demand was found to have a significant influence on the performance of two-layer restoration. Simulation results also confirmed our prediction that finer granularity leads to higher recovery ratio but slower recovery speed. Two frequently used interworking strategies are studied here. Recovery token is found to be more efficient than holdoff timer, although it is more complicated in real implementations. Finally, we find that wavelength conversion has a major effect on our two-layer recovery scheme, especially for optical layer restoration.

As mentioned before, in the proposed two-layer recovery scheme, various existing recovery mechanisms can be deployed at either layer. In the current study, only path-based rerouting (restoration) was used in both the IP and optical layers. Since protection is essentially different from restoration, there is interest in investigating the performance of the proposed two-layer recovery scheme with protection deployed at each layer or even combining protection and restoration at either layer.

The cases of a network with and without wavelength conversion were studied in this article. Studies about the network with partial wavelength conversion and the wavelength converters allocation problem will be included in future work.

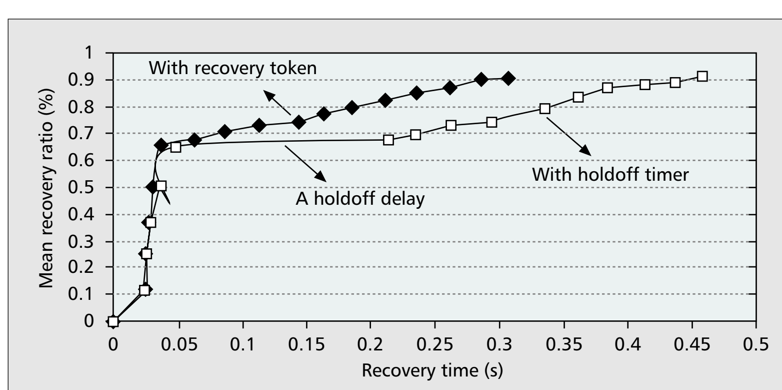

■ Figure 5. *Comparison of recovery token and holdoff timer interworking strategies.*

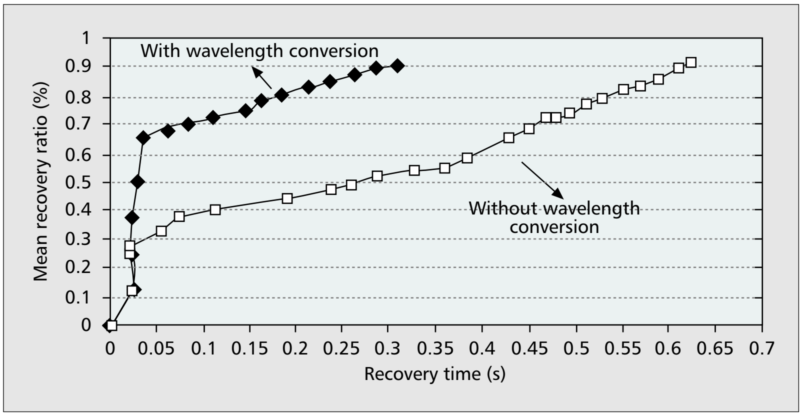

Figure 6. *Performance comparison with/without wavelength conversion.*

Recently, the emerging generalized multiprotocol label switching (GMPLS) has extended MPLS to encompass time, wavelength, and space domains, which opens an opportunity to have a common control plane operate across dissimilar network types. It would be interesting to study survivability in multilayer transport networks equipped with GMPLS. This also will be left for further study.

References

[1] Telcordia Netsizer Website, http://www.netsizer.com

[2] M. Noll, "Voice vs. Data: Estimates of Media Usage and Network Traffic," *Proc. 29th TPRC Conf.*, Alexandria, VA, 2001.

[3] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol Label Switching Architecture," RFC 3031, Jan. 2001, http://www.ietf.org/rfc/rfc3031.txt

[4] D. Awduche and Y. Rekhter, "Multiprotocol Lambda Switching: Combining MPLS Traffic Engineering Control with Optical Cross Connects," *IEEE Commun. Mag.*, vol. 39, no. 3, Mar. 2001.

[5] N. Ghani, "Lambda-Labeling: A Framework for IP-over-WDM Using MPLS," *Opt. Net. Mag.*, vol. 1, no. 2, Apr. 2000.

[6] S. Ramamurthy and B. Mukherjee, "Survivable WDM Mesh Networks, Part I-Protection," *Proc. IEEE Infocom '99*, vol. 2, New York, NY, Mar. 1999.

[7] V. Sharm *et al.*, "Framework for MPLS-based Recovery draft-ietf-mpls-recovery-frmwrk-04.txt," May 2002, http://www.ietf.org/internet-drafts/draft-ietf-mpls-recovery-frmwrk-04.txt

[8] C. Huang *et al.*, "A Path Protection/Restoration Mechanism for MPLS Networks," July 2001, http://www.watersprings.org/links/mlr/id/draft-chang-mpls-path-protection-03.txt

[9] D. Colle *et al.*, "MPLS Recovery Mechanisms for IP-over-WDM Networks," Special Issue on IP over WDM and Optical Packet Switching, *Photonic Net. Commun.*, vol. 3, no. 1, Jan. 2001.

[10] S. Ramamurthy and B. Mukherjee, "Survivable WDM Mesh Networks, Part II-Restoration," *Proc. IEEE INFOCOM '99*, vol. 2, New York, NY, Mar. 1999.

[11] D. Colle *et al.*, "Porting MPLS-recovery Techniques to the MPS Paradigm," *Opt. Net. Mag.*, vol. 2, no. 4, July/Aug. 2001.

[12] C. Assi *et al.*, "Design a Survivable IP over WDM Networks," *OptiComm 2001*, Denver, CO, Aug. 21–22, 2001.

[13] O. Gerstel and G. Sasaki, "Quality of Protection (QoP): A Quantitative Unifying Paradigm to Protection Service Grades," *OptiComm 2001*, Denver, CO, Aug. 21–22, 2001.

[14] P. Demeester *et al.*, "Resilience in Multilayer Networks," *IEEE Commun. Mag.*, vol. 37, no. 8, Aug. 1999.

[15] D. Colle *et al.*, "Data-Centric Optical Networks and Their Survivability," *IEEE JSAC*, vol. 20, no. 1, Jan. 2002.

[16] C. Murthy and M. Gurusamy, *WDM Optical Networks: Concepts, Design, and Algorithms*, Prentice Hall PTR, Nov. 2001, pp. 225–55.

[17] OMNeT++ Website, http://www.hit.bme.hu/phd/vargaa/omnetpp.htm

Biographies

YANG QIN [M] (eyqin@ntu.edu.sg.) received her B.Sc. (with first class honors) in computer science at Southwest Jiaotong University, China, in 1989, her M.S. in computer science at Huazhong University of Science & Technology, Wuhan, Hubei, in 1992, and her Ph.D. in computer science from Hong Kong University of Science & Technology, Kowloon, in November 1999. She worked as a software engineer at Elite Business Machine Co., Shenzhen, Guangdong, China from 1992 to1994. From November 1999 to November 2000 she worked as a research associate at the School of Electrical Engineering and Computer Science, Washington State University. Since November 2000, she has joined the School of Electrical and Electronic Engineering at Nanyang Technological University as an assistant professor. Her research interests include optical high-speed networks, WDM networks, QoS, performance analysis, multimedia communication, and wireless mobile ad hoc networks.

LORNE G. MASON [SM] is a professor emeritus at INRS Telecommunications, University of Quebec. He has over 35 years experience in the telecommunication networks area and has held research positions with Nortel (formerly BNR) and B.C. Telephone before joining INRS-Telecommunications in 1977. Since that time he has taught at the post graduate level and conducted research in broadband networks with more than 80 publications in leading journals and conferences. His current interests include Internet traffic engineering, and traffic management and control in MPLS and related networks. He has provided consulting services for a number of industrial and academic research centers including, B.C. Telephone, Yale University, INRS-Telecom, TRLabs, the European Commission, and MediaRing. He is the Canadian representative to the Telecommunications Committee of IFIP.

KE JIA received his B.Sc. from the University of Electronic Science and Technology, Chengdu, China in 2001. He studied as a postgraduate student at the School of Electrical and Electronics Engineering, Nanyang Technological University from July 2001 to September 2002. Currently, he works for Motorola Company in Singapore.