

A novel end-to-end fault detection and localization protocol for wavelength-routed WDM networks

Hongqing Zeng^{*a,b}, Alex Vukovic^a, and Changcheng Huang^b

^a Communications Research Centre Canada, 3701 Carling Ave.,
P. O. Box 11490, Stn. H, Ottawa, ON, K2H 8S2, Canada

^b Dept. of Systems and Computer Engineering, Carleton University,
1125 Colonel By Drive, Ottawa, Ontario, Canada K1S 5B6

ABSTRACT

Recently the wavelength division multiplexing (WDM) networks are becoming prevalent for telecommunication networks. However, even a very short disruption of service caused by network faults may lead to high data loss in such networks due to the high data rates, increased wavelength numbers and density. Therefore, the network survivability is critical and has been intensively studied, where fault detection and localization is the vital part but has received disproportional attentions. In this paper we describe and analyze an end-to-end lightpath fault detection scheme in data plane with the fault notification in control plane. The endeavor is focused on reducing the fault detection time. In this protocol, the source node of each lightpath keeps sending *hello* packets to the destination node exactly following the path for data traffic. The destination node generates an alarm once a certain number of consecutive *hello* packets are missed within a given time period. Then the network management unit collects all alarms and locates the faulty source based on the network topology, as well as sends fault notification messages via control plane to either the source node or all upstream nodes along the lightpath. The performance evaluation shows such a protocol can achieve fast fault detection, and at the same time, the overhead brought to the user data by *hello* packets is negligible.

Key words: Optical network, fault detection and localization, wavelength-routed, wavelength-division multiplexing

1. INTRODUCTION

Recently, with the deployment of wavelength-division multiplexing (WDM) technology in optical networks, the number of wavelengths in a single link and the channel speed tend to become larger and larger. For example, it is reported even in 2001 that 432 wavelengths could be multiplexed into a single fibre [1]. In current commercial WDM systems, it is easy to boost the number of available wavelengths within a fibre to 192 or above [2-3]. In some real application systems, the channel speed has been updated to 10 Gbps from 2.5 Gbps. Experimental systems with 40 Gbps channel speed are also reported and they are emerging in market [4]. Furthermore it is widely believed that optical networks will eventually provide the function of dynamic wavelength switching and routing to end-users. Such high-speed, large-capacity and wavelength-routing flexibility enable WDM networks to be prevalent in the future telecommunication networks.

However, a very short disruption of service caused by a network fault may lead to a very high data loss in such networks. Thus it is essential to ensure continuous and reliable network operation. Numerous schemes have been proposed for such networks to improve the network survivability [5], where fault detection and localization is the vital part but has received disproportional attention.

Fault detection is relatively easier and faster than fault localization. Optical power detection, optical spectral analysis, pilot tone and optical time domain reflectometry can be deployed for fault detection in WDM networks [6]. Fault

* hongqing.zeng@crc.ca; phone 1 613 520-2600 ext. 2253; www.crc.ca

localization is to find the minimum set of potential failed network resources based on the alarms generated in the fault detection phase. Fault localization in general network is not a new problem and has been intensively studied for years in various areas, including power distribution systems, electrical circuits, industrial control systems, as well as communication networks [7]. However, existing fault localization approaches for conventional networks cannot be applied to WDM networks directly due to the lack of electrical terminations [8], or the unaffordable cost and complexity of implementation. To handle such particular issues for WDM networks, black-box based and network-model based methods have been proposed [8]. Black-box based methods consider the network as a black-box, which generates certain output data when a network fault occurs. They acquire the relationship between the network faults and output diagnosis through expert systems [9], artificial neural networks [10] or other proactive learning systems [11]. Such approaches don't need to know the exact model of the network and are suitable and efficient for large-scale networks. However, they usually have slow training or learning processes. Especially in WDM networks with the function of dynamic wavelength switching and routing to end-users, every new lightpath provisioning may initialize a new training or learning process. Thus the inaccuracy of fault detection and localization is inevitable during the training and learning processes. On the other hand, network-model based methods firstly construct an accurate model based on the network topology, components, transmission characterization, and so on. Then they compare the expected network behavior based on the model with the actual observation of the network, to detect and locate network faults. According to the applied network models, such approaches include probabilistic reasoning systems [12-13], finite state machine models [14-15], and deterministic fault-propagation models [16]. They are accurate for static networks but for WDM networks with dynamic lightpath provisioning, the network model has to be changed dynamically. This is extremely difficult and time-consuming, if not impossible, and thus makes such approaches unable to detect and locate the network faults within the strict time-constraint.

To deal with the challenge of the dynamic lightpath provisioning in WDM networks, we propose an end-to-end fault detection and localization protocol in this paper. In this protocol, the sender user keeps sending the "hello" packets in a certain pattern to the receiver user along the established lightpath during its lifetime. The receiver reports a fault once a given number of consecutive *hello* packets are missed within a time threshold. Then the receiver initializes a fault recovery process. At the same time the network management system (NMS) start to locate the network fault based on the alarm distribution in multiple lightpaths. Such a protocol requires neither the priori information about the network topology and characterization, nor the long-time training/learning process. It is especially suitable to be integrated into the destination-initiating path restoration protocols for wavelength-routed WDM networks [17]. The *hello* packet in this protocol could be borrowed from the MPLS echo request/reply packet [18], in order to reduce the cost of implementation of the proposed protocol in real world.

This paper is organized into the following sections. Section 2 describes the background of the work in this paper, e.g., the architecture of wavelength-routed WDM networks, the time estimation of fault recovery process, and the methods to insert *hello* packets to user data. The details of the proposed protocol are given in Section 3, including the *hello* packet format. Section 4 statistically analyzes the performance of the protocol, including the calculation of fault detection time and the overhead brought by the *hello* packets. Finally, Section 5 lists some concluding remarks and proposes some future work.

2. BACKGROUND

2.1 The model of wavelength-routed WDM networks

A typical architecture of wavelength-routed WDM networks is shown in Fig. 1. In such architecture a network node consists of an optical switch (either with or without opto-electro-opto conversion) and an electronic controller. The controller maintains either local or global network state information for wavelength routing, e.g., network topology, wavelength occupation, and port mappings of the optical switch. The controller also provides an interface for upper layer protocols to manipulate the switch through commercial or customer developed software. The optical switch connects its input ports with output ports optically to perform the wavelength switching, according to the control information stored in the electronic controller.

Those networks are interconnected by so-called WDM links. Each WDM link consists of a pair of uni-directional fibres where one fibre has the reverse transmission direction with the other one. A number of optical channels (or called

wavelengths) are multiplexed within each fibre by using the WDM technology. User data are transported along these optical channels. All the optical channels for user data constitute the data plane. All controllers in such a network communicate with each other either over extra electronic channels or over dedicated optical channels in WDM links. These channels for controller communications constitute the control plane of the network.




Fig. 1. The architecture of wavelength-routed WDM networks

At the input and output ports of optical switches, wavelength converters could be applied to improvement the global wavelength utilization. For the clarity and conciseness, we assume there is no wavelength converter in our WDM networks in this paper, although the research work is also applicable to WDM network with wavelength converters.

In a network with above architecture, before any data transportation between a pair of nodes, a lightpath must be established between them. Once a network fault occurs along any established lightpath, a backup path detour the affected network resources must be quickly set up to recover the disrupted network service. The techniques for provisioning the backup path fall into two categories, pre-designed protection and dynamic restoration. In pre-designed protection schemes, the protection routes are preplanned based on redundant resources (e.g., redundant fibres, transceivers, and switches) before the network faults occur. Pre-designed protection schemes include dedicated protection where every redundant resource is dedicated to a work path, and shared protection where the same redundant resources are shared by multiple work paths. Examples include $(1+1)$, $(1:1)$, $(n:m)$ and WDM Self Healing Ring protection [6]. Protection schemes can shift the affected traffic to the backup path/routes in a short time after the failure (less than 60 msec in SONET). On the other end, the restoration schemes dynamically look for backup routes of spare capacity after a network fault occurs. Such schemes are usually slower than protection schemes but they can provide higher utilization of network resources and may provide additional resilience against further network faults, e.g. optical attack and/or dual failures. The restoration approaches are based either on path or on link (segment) paradigm. Here the segment refers to a part of a path which consists of a group of consecutive links. In path restoration, a source-destination node pair provisions a backup route on an end-to-end basis, while in the link (segment) restoration, the end nodes of a failed link (segment) establish a detour route that bypass the failed link (segment). It is well known that the link (segment) restoration is faster than path restoration but, on the contrast, the later achieves the better capacity utilization.

2.2 Fault recovery process in WDM networks

The recovery process in WDM networks consists of 5 phases: fault detection, notification, localization, traffic switch and reversion (optional). The typical sequence of network events during the fault recovery process is shown in Fig. 2 [19]. For each lightpath in a WDM network, if we omit the optional step, the overall recovery time includes detection time,

notification time and traffic switchover time,

$$T_R = T_D + T_L + T_N + T_S \quad (1)$$

where T_R , T_D , T_L , T_N and T_S represents the overall fault recovery time, fault detection time, fault localization time, fault notification time, and traffic switchover time, respectively. Hereafter we define the overall fault recovery time as the time taken from the instant a network fault occurs to the instant the traffic is recovered successfully. Similarly, the fault detection time is defined as the time taken from the instant a network fault occurs to the instant the fault is detected, as shown in Fig. 2. The fault localization time is the period from the instant starting fault localization to the instant the fault is located to a network link. The fault notification time is the period from the instant the fault is located to the instant all corresponding nodes receive the fault notification. The traffic switchover time is calculated from the instant the corresponding node starts to set up the backup path to the instant the traffic is resumed to transfer along the backup path.

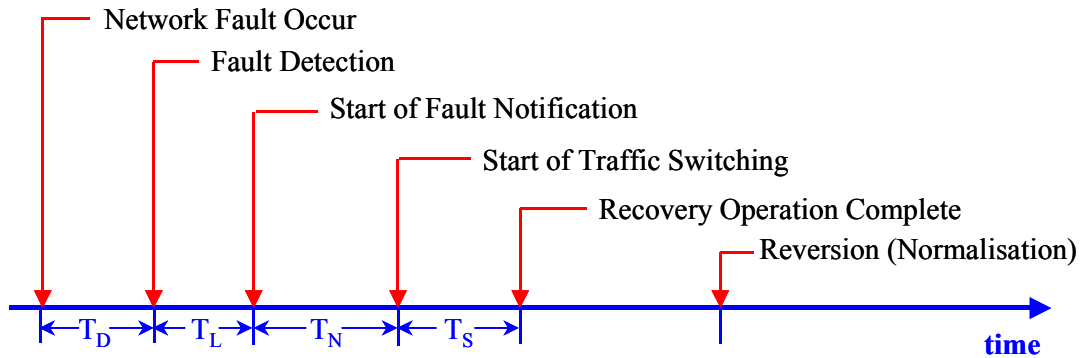


Fig. 2. The sequence of events in network fault recovery process

3. END-TO-END HELLO FAULT DETECTION AND LOCALIZATION PROTOCOL

In this section we introduce an end-to-end lightpath fault detection and localization protocol for wavelength-routed WDM networks. In this protocol, the communication source of an established lightpath keep inserting “hello” packets into user data and sending such packets to the communication destination of the lightpath, exactly along the corresponding user data route (in the data plane). At the lightpath destination, if a certain number of such *hello* packets are missed in a given time threshold, the controller of the destination node assumes that a network fault occurs along the lightpath. Then the destination node starts to send fault alarms to the network management system (NMS) and, at the same time, send notification to the affected nodes to invoke the traffic restoration scheme. If a path restoration scheme is deployed in the network, the notification will be only sent to the source node. Otherwise, if the segment (link) restoration is applied, the notification will be sent to all upstream nodes along the lightpath. In this protocol, the fault notification phase starts before the fault localization to accelerate the recovery process. The NMS collects alarms in real time and executes the fault localization algorithm based on the alarm distribution in lightpaths. For example, in Fig. 3 the common parts of the two lightpaths that report alarms would be located as the network fault source. In this protocol, the notification messages are transported in data plane, with the fault notification and localization messages in control plane. The procedure of this protocol is described as follows:

- 1) The source node periodically inserts “hello” packets to user data, the time interval between two consecutive *hello* packets is defined as the *hellointerval*;
- 2) The destination node reports a fault once a certain number of consecutive *hello* packets are missed within a given time threshold (defined as the *detectioninterval*);
- 3) The destination node sends fault notification messages via control plane to the source node;
- 4) Corresponding nodes trigger the traffic recovery scheme;
- 5) Fault localization in NMS based on the alarm distribution in multiple lightpaths, and
- 6) Traffic reversion after removing the fault (optional).

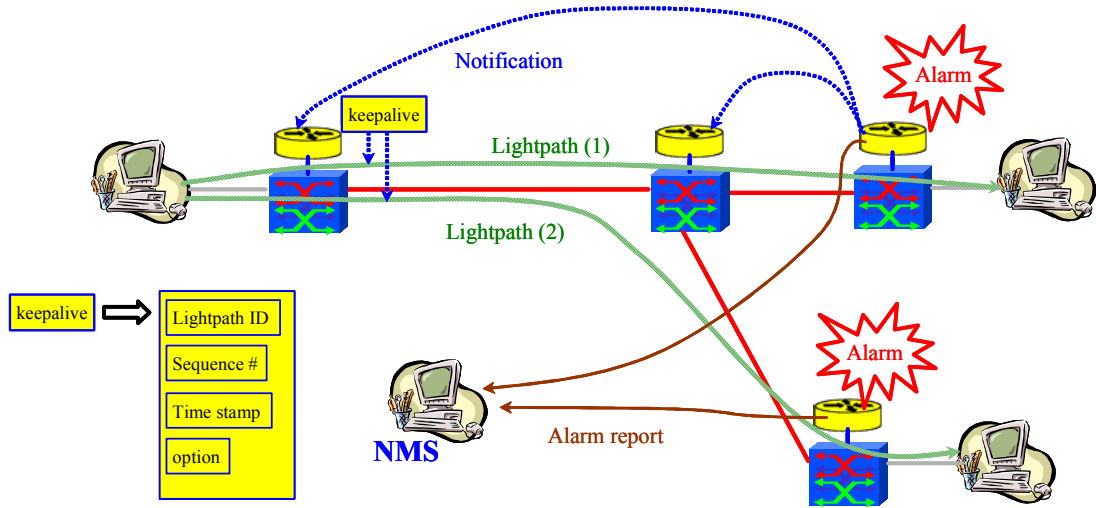


Fig. 3. End-to-end “Hello” packet based fault detection and localization

The format of the *hello* packet is shown in Fig. 3, which includes at least the lightpath ID, the packet sequence number, and the time stamp. Some optional fields are also reserved for future use. Furthermore, we can also obtain some end-to-end statistics of the lightpaths, such as latency, loss and rout availability, by statistically analyzing such packets. Thus, this scheme can be also considered as an active measurement [20] of the network quality of service (QoS) parameters and be helpful for QoS-sensitive network fault recovery [21].

4. PERFORMANCE ANALYSIS

4.1 Fault detection time

As described in Section 2.2, the overall recovery time for a network fault includes the fault detection time, the fault notification time, and the traffic switchover time (see equation (1)). The endeavour of our proposal is focused on reducing the fault detection time, which is the main component of the overall recovery time in WDM networks. In the proposal, the destination node of a path periodically receives *hello* packets to maintain the health of the path. If a certain number, say n , of *hello* packets are missed at the destination, i.e., the destination nodes do not receive any *hello* packet from the source within a *detectioninterval*, e.g., 4 *hellointervals* as in OSPF, it assumes that a fault occurs along the path. To estimate the fault detection time of the proposal, we define the following notations:

T_R : overall fault recovery time

T_D : fault detection time

T_N : fault notification time

T_S : traffic switchover time

T_{proc} : time that the source or destination node takes to process a *hello* packet

T_{prop} : propagation delay on each link (a hop along the path)

τ : *hellointerval*, the time between two consecutive *hello* packet at the source node

T_{int} : *detectioninterval*, the time threshold for the destination node to report a fault without receiving any *hello* packet

n : the expected number of consecutive *hello* packets within a *detectioninterval*

The *hello* packets transportation between the source and destination nodes is depicted in Fig. 4. According to the description in the figure, the fault detection time, T_D , can be calculated as,




Fig. 4. Transportation of *hello* packets between a source-destination node pair

$$(n-1) \times \tau \leq T_D \leq n \times \tau + T_{prop} \quad (2)$$

Formula (2) shows that the fault detection time can be reduced by decreasing τ but it is not effective in reducing the fault detection time beyond a limit, due to the fact that a smaller τ also increases the risk of false fault report. Because when the *hello* packet is inserted, if there are already a number of traffic data packets are waiting in the queue (at the source node), then several consecutive *hello* packet will be lost even no fault occur along the lightpath. If we consider some optical switching schemes, e.g., burst switching, then such risk also exists in middle nodes. More important, with a smaller τ , the fault detection protocol will bring more overhead to the user data and therefore decrease the network utilization. The number of lost *hello* packets, n , before report a failure is another factor that affects the fault detection time. Similar to the effect of τ , a smaller n speeds up the fault detection but may increases the risk of false fault report. Therefore a tradeoff regarding τ and n is expected for achieving fast fault detection and localization under various logical network topologies and traffic loads.

If we assume the network faults are uniformly distributed over the time, then the average fault detection time can be estimated as the following according to formula (2),

$$E[T_D] = n\tau - \frac{\tau}{2} + \frac{1}{2}T_{prop} \quad (3)$$

Table 1 lists the average fault detection times for some τ and n combinations with $T_{prop} = 0.50 \text{ ms}$ and $T_{proc} = 0.01 \text{ ms}$ [17].

It is worth to point out that the detection time is also affected by the *hello* packet insertion scheme. Besides the uniformly packet insertion, some other insertion schemes, e.g., adaptive sampling [20], are also applicable and might be more efficient. Also note that the fault notification phase is moved before the fault localization phase in the proposal to reduce the overall fault recovery time. Accordingly, if we replace (3) in (1), the overall fault recovery time based on the proposal is,

$$T_R = (n - 1/2) \times \tau + T_{prop} + T_N + T_S \quad (4)$$

Table 1. Average fault detection times (*ms*) for given τ and n combinations

$\tau \backslash n$	1	2	3	4	5	6
0.25 ms	0.375	0.625	0.875	1.125	1.375	1.625
0.50 ms	0.500	1.000	1.500	2.000	2.500	3.000
0.75 ms	0.625	1.375	2.215	2.875	3.625	4.375
1.00 ms	0.750	1.750	2.750	3.750	4.750	5.750

4.2 Overhead due to *hello* packets

Clearly, due to the inserted *hello* packets, the utilization of the user data channel is decreased, i.e., the *hello* packets bring some overhead to the user data. To quantitatively evaluate such channel utilization degradation, we define the *hello* overhead, *HOH*, as the ratio of *hello* packets over the sum of user data and *hello* packets within a time unit. Let the user data load be ρ packets/second, the average user data packet length be L_u , and the *hello* packet length be L_h , then the *hello* overhead can be written as,

$$HOH(\%) = \frac{\frac{1}{\tau} \times L_h}{\frac{1}{\tau} \times L_h + \rho \times L_u} = \frac{L_h}{L_h + \rho \times L_u \times \tau} = \frac{1}{1 + \frac{\rho \times L_u}{L_h / \tau}} \quad (5)$$

In equation (5), $\rho \times L_u$ is the user data rate and L_h / τ is the data rate of *hello* packets. Usually the former is very large relative to the later. For example, the typical data rate is 155Mbps or 625Mbps in classic SDH/SONET. In the meanwhile, the length of the basic MPLS echo request packet is 44 bytes or 68 bytes [18]. If we borrow it as the *hello* packet, L_h is no more than 544 bits. Even with the smallest τ in Table 1, the data rate of *hello* packets, L_h / τ , is about 2Kbps, which leads the overhead brought to the data channel by *hello* packets to be negligible.

5. CONCLUSIONS AND FUTURE WORK

In this paper, we presented an end-to-end lightpath fault detection scheme in data plane with the fault localization and notification in control plane. The endeavor was focused on reducing the fault detection time. In this protocol, the source node of each lightpath keeps sending *hello* packets to the destination node exactly following the path for data traffic. The destination node generates an alarm once a certain number of consecutive *hello* packets are missed within a given time period. Then the network management unit collects all alarms and locates the faulty source based on the network topology, as well as sends fault notification messages via control plane to corresponding traffic recovery nodes. The fault detection time depends on the setting of the number of the lost *hello* packets before generating an alarm, and the time interval between two *hello* packets. We statistically analyzed the fault detection time and estimated the overhead to user data brought by the *hello* packets. The estimation showed that such overhead is negligible.

The future work will be focused on the validation of the proposal upon a WDM network testbed. The integration with some existing recovery protocols, e.g., the destination-initiating path restoration protocol, is another consideration. Furthermore, some fault localization algorithms for the network management are expected. Such algorithms are based on the alarm information generated by destination nodes of established lightpaths, as well as the information of lightpath distribution in the WDM network. Some schemes to suppress the redundant alarms in the control plane are strongly expected.

REFERENCES

1. S. V. Kartalopoulos, *Fault Detectability in DWDM – Toward Higher Signal Quality & System Reliability*, Piscataway: IEEE Press, 2001
2. W. D. Grover, *Mesh-based survivable networks*, Upper Saddle River, NJ: Prentice Hall, 2004
3. R. Elliott, Dark fibre pricing analysis Europe 1998-2002, [Online document], Nov. 2002, available at http://www.band-x.com/information/Dark_Fibre_Report_prices_98-02webversion.pdf
4. G. Held, "On the road to OC-768," *IT Professional*, **Vol. 3, No. 2**, pp. 46 – 48, March-April 2001
5. S. Ramamurthy, L. Sahasrabudhe, and B. Mukherjee, "Survivable WDM mesh networks," *IEEE Journal of Lightwave Technology*, **Vol. 21, Issue 4**, pp. 870 – 883, 2003
6. R. Ramaswami and K. N. Sivarajan, *Optical networks: a practical perspective*, Morgan Kaufman, 1998
7. R.J. Patton, P.M. Frank, and R. N. Clark, *Issues of Fault Diagnosis for Dynamic Systems*, Springer, 2000
8. C. Mas and P. Thiran, "A review on fault location methods and their applications in optical networks," *Optical Network Magazine*, **Vol. 2, No. 4**, 2001
9. R. Weihmayer, G. Jakobson and M. Weissman, "A domain oriented expert system for telecommunication network alarm correlation," *Network Management and Control*, **Vol. 2**, 1994
10. C. Rodriguez, S. Rementeria, J.I. Martin, A. Lafuente, J. Muguerza and J. Perez, "A modular neural network approach to fault diagnosis," *IEEE Trans on Neural Networks*, **Vol. 7, No. 2**, pp. 326-340, 1996
11. L. L. Ho, D. J. Cavuto, S. Papavassiliou, and A. G. Zawadzki, "Adaptive and automated detection of service anomalies in transaction-oriented WANs," *IEEE Journal on Selected Area in Communications*, **Vol. 18, No. 5**, pp. 744-757, May 2000
12. R.H. Deng, A.A. Lazar, and W. Wang, "A probabilistic approach to fault diagnosis in linear lightwave networks," *IEEE Journal on Selected Areas in Communications*, **Vol. 11**, pp. 1438-1448, 1993
13. N.S.V. Rao, "On parallel algorithms for single-fault diagnosis in fault propagation graph systems," *IEEE Transactions on Parallel and Distributed Systems*, **Vol. 7, No. 12**, pp. 1217-1223, 1996
14. A.T. Bouloutas, G.W. Hart, and M. Schwartz, "Fault identification using a finite state machine model with unreliable partially observed data sequences," *IEEE Transactions on Communications*, **Vol. 41, No. 7**, pp. 1074-1083, 1993
15. C.-S. Li and R. Ramaswami, "Automatic fault detection, isolation, and recovery in transparent all-optical networks," *IEEE Journal of Lightwave Technology*, **Vol. 15, No. 10**, pp. 1784-1793, 1997
16. Carmen Mas and Patrick Thiran, "An Efficient Algorithm for Locating Soft and Hard Failures in WDM Networks," *IEEE J on Selected Areas in communications*, **Vol. 18, No. 10**, pp 1900-1911, Oct. 2000
17. J. Zheng and H. T. Mouftah, "Destination-initiating path restoration protocol for wavelength-routed WDM networks," *IEE Proceedings Communications*, **Vol. 149, No. 1**, pp. 18-22, Feb. 2002
18. K. Kompella and G. Swallow, "Detecting MPLS data plane failure," *IETF Draft: draft-ietf-mpls-lsp-ping-09.txt*, May 2005
19. P. Czezowski and T. Soumiya, "Optical Network Failure Recovery Requirements," *IETF Draft: draft-czezowski-optical-recovery-reqs-01.txt*, Feb. 2003
20. W. Ma, J. Yan, and C. Huang, "Adaptive Sampling for Network Performance Measurement under Voice Traffic," *Proceedings of IEEE ICC'04*, Paris, June, 2004
21. M. Li, C. Huang, and A. Srinivasan, "Distributed Backup Path Bandwidth Management for Network Reliability," *Proceedings of 2004 International Conference on Dependable Systems and Networks (DSN'04)*, Florence, Italy, June 28-July 1, 2004
22. M. Médard, D. Marquis, and S. R. Chinn, "Attack detection methods for all-optical networks," *Network and Distributed System Security Symposium*, 1998