Fault Detection and Localization Scheme for All-Optical Overlaid-Star TDM Networks

Jun Zheng, Cheng Peng, and Gregor v. Bochmann
School of Information Technology and Engineering
University of Ottawa
Ottawa, Ontario K1N 6N5, Canada

***Abstract*—Fault detection and localization is a crucial issue in all-optical networks. Since most commercially-available all-optical space switches are incapable of detecting the loss of optical signals along the data paths between its input ports and output ports, fault localization becomes a challenge for providing service survivability in such networks. This paper proposes a fault detection and localization scheme for an all-optical overlaid-star TDM network. The proposed scheme employs a fault localization technique that identifies the location of a failure by detecting the power loss of optical signals in data and control channels. Two alternatives are proposed. One requires a control channel on each wavelength of a fiber link while the other requires a small data block to be transmitted in each non-allocated data channel. Based on the proposed fault localization technique, a fault advertisement protocol is further presented, which can be incorporated into the signaling protocol used in the network to facilitate the provisioning of static protection or dynamic restoration. The data loss, fault detection time, and connection recovery time are analyzed for the different failure scenarios.**

I. INTRODUCTION

With recent advances in enabling technologies, all-optical switches have become commercially available, making it possible to deploy all-optical networks. An all-optical network has the characteristic that both signal transmission and switching are performed in the optical domain, which eliminates O/E/O conversions within the network and therefore improves network performance significantly. In an all-optical network, each fiber link carries a number of optical (wavelength) channels, each operating at a very high speed of several gigabits per second. A single network failure such as a fiber cut may cause a large amount of data loss, and thus greatly degrade and even disrupt network services. To guarantee network services, the network must incorporate effective protection and restoration mechanisms to provide a high level of service survivability against different types of network failures, such as a fiber cut or a node fault.

All-optical overlaid-star TDM networks are a class of all-optical networks that employ an overlaid-star topology and use time division multiplexing (TDM) for data transmission. This class of networks features the ability to dynamically allocate bandwidth on demand at a fine granularity, and the concentration of control and routing functionality at the electronic edge nodes that surround the optical core [1]. The architecture of such networks consists of a number of edge nodes interconnected via several core nodes in an overlaid-star topology, as shown in Figure 1. The overlaid-star topology provides robustness in the case of a network failure and at the same time relieves potential network congestion.

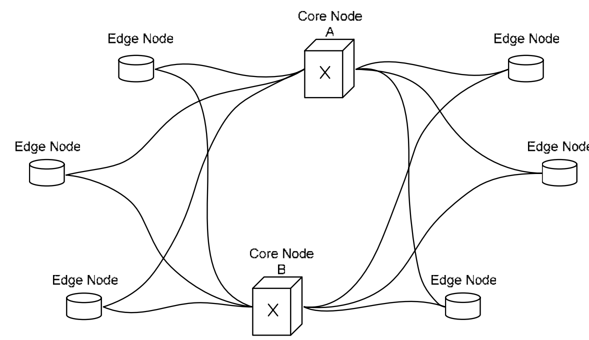

Figure 1 Overlaid star topology

In such networks, a connection between a pair of edge nodes is established on a primary path passing one of the core nodes. If the connection is disrupted because of a failure occurring on the path, the edge nodes should switch the connection over to a backup path passing through another core node to continue the transmission. The backup path can be provisioned using two different paradigms: static (preconfigured) protection and dynamic restoration [2]. In static protection, the backup path is established at the same time the primary path is established, but passing through a different core node. In the event of a failure, both the source and destination nodes will switch over to the backup path to continue the transmission. In dynamic restoration, no backup path is established before a failure occurs. In the event of a failure, the source and destination nodes dynamically establish a backup path passing through a different core node and then switch over to the backup path to continue the transmission. Obviously, static protection is faster in service recovery than dynamic restoration while dynamic restoration is more efficient in resource utilization than static protection.

Fault detection and localization is a crucial issue in all-optical networks. Since most commercially-available all-optical space switches are incapable of detecting the loss of

1-4244-0463-0/06/$20.00

optical signals along the data paths between its input ports and output ports, fault localization becomes a challenge for providing service survivability in such networks. In this paper, we propose a fault detection and localization scheme for an all-optical overlaid-star TDM network. The proposed scheme employs a fault localization technique that identifies the location of a failure by detecting the power loss in data and control channels. Two alternatives are proposed for the fault localization technique. One requires a control channel on each wavelength of a fiber link while the other requires a small data block transmitted in each unallocated data channel. Based on the proposed fault localization technique, a fault advertisement protocol is further presented, which can be incorporated into the signaling protocol used in the network to facilitate the provisioning of static protection or dynamic restoration.

The rest of the paper is organized as follows. In Section II, we briefly introduce the network architecture, discuss the fault detection and localization issue, and review related work in the literature. In Section III, we describe the proposed fault localization technique and the fault advertisement protocol. In Section IV, we analyze the data loss, fault detection time, and connection recovery time with the proposed fault detection and localization scheme in different failure scenarios. In Section V, we present our conclusions.

II. BACKGROUND

In this section, we briefly introduce the network architecture, discuss the fault detection and localization issues, and review related work in the literature.

1. Network Architecture

An all-optical overlaid-star TDM network consists of a number of edge nodes interconnected via several central core nodes in an overlaid star topology, as shown in Figure 1. Each edge node is connected to a core node by a couple of fibers, one for transmission in each direction. An edge node is a hybrid electro-optical component that serves as an interface between the optical network and an electronic external network, based on IP, MPLS, or ATM. A core node employs an all-optical space switch that can switch an input wavelength on an input port to an output port, making data paths inside the core node purely optical and transparent.

The network uses time division multiplexing (TDM) for data transmission. Each fiber supports multiple wavelengths. Each wavelength is divided into a series of frames that consist of a fixed number of timeslots. The control of the network is performed in the electronic domain. Each core switch has an associated electronic controller that performs timeslot allocation, switch configuration, and other control functions. The control messages are exchanged between edge nodes and core nodes out-of-band over a dedicated control timeslot (or channel) on a particular wavelength of each fiber. There is one control timeslot per frame in either direction. Due to the use of TDM, a control channel from an edge node must go through a switch fabric to reach the controller and vice versa.

2. Fault Detection and Localization

As mentioned in Section 1, service survivability in an optical network can be provided in the form of either static protection or dynamic restoration. No matter what paradigm is employed, fault detection and localization is a prerequisite. In an optical network, a failure such as a fiber cut or a component failure would cause the loss of an optical signal and can thus be easily detected by measuring the power level of the signal at the end of the failed link. There are also other types of failures, such as a fault in a transmitter, which may result in no data block being received or received with transmission errors without causing the power loss of the optical signal. In these cases, the problem can be identified at the destination edge node by checking the validity of the received data in those time slots when a data block is expected.

In an all-optical overlaid-star network, a detector can theoretically be deployed in both edge nodes and core nodes. In this case, a fiber cut on an upstream link can be easily detected at the corresponding input port of a core node. A fiber cut on a downstream link or a fault in a component of a core node can be easily detected at the edge node of the downstream link. In the real world, however, most commercially-available all-optical space switches are not equipped with detectors along the data paths between its input ports and output ports. For this reason, a fiber cut on an upstream link cannot be detected at the corresponding port of a core switch. A fiber cut on a downstream link or a fault in a core switch on a physical path between a pair of edge nodes can only be detected at the receivers of the destination node, making it difficult to localize the failure. To solve this problem, an effective fault localization technique is needed.

3. Related Work

Network survivability has been extensively studied in optical networks. A variety of protection and restoration schemes have been proposed to provide service survivability against different types of network failures in various network scenarios [3]. As a prerequisite for protection and restoration, fault detection and localization has been widely studied in traditional optical networks. In SONET, fault detection is achieved by electronically monitoring a loss of data or a high bit error rate through digital cross-connects (DXCs) [3]. In opaque optical networks, fault detection can also be achieved by electronically monitoring the loss of signal power at the receivers in an electro-optical switch [4]. In [5], Mas and Thiran studied the fault localization problem in optical networks from the viewpoint of fault management and considered a multi-hop optical network using WDM technology and electro-optical switches. In such a network, a single failure may trigger a number of alarms generated by different elements in a switch, which are sent to the manager. This is more complex when multiple failures occur almost simultaneously. The manager has to identify where the failure is actually located based on all the alarms received. To address this problem, Mas and Thiran proposed an alarm filtering algorithm (AFA) for fault management in the network. In [6], Mas et al. presented a more comprehensive study on the fault management problem similar to that in [5]. This work assumed that different types of testing and monitoring equipment are available in different components and at different layers of a WDM network. Both hard failures and soft failures are considered. A fault location algorithm (FLA) is proposed to detect multiple soft and hard failures in a WDM network, which is an improved version of AFA in [5]. Despite the extensive studies for traditional optical networks, however, there is not much work conducted for all-optical networks. In fact, no effective technique has been available to address the localization problem described in Section 2.2 [7]. To the best of our knowledge, this work is the first attempt to solve the fault localization problem in an all-optical overlaid-star TDM network.

III. Fault Localization and Advertisement

In this section, we present a fault localization technique and a fault advertisement protocol to facilitate the provisioning of static protection or dynamic restoration.

1. Fault Localization Technique

We consider a network scenario with two core nodes, which is a reasonable deployment in practice. In such a network, there are two paths between each pair of edge nodes, passing different core nodes. A connection between a pair of edge nodes is established on one of the two paths. To establish a connection, the source node, destination node, and the core node participate in a signaling protocol to allocate a wavelength and timeslots on a selected path for the connection. When a connection request arrives at a source node, the source node first sends a REQUEST message to the core node it selects. The core node will allocate a wavelength and timeslots for the connection and then send an ACKNOWLEDGMENT message to the source node. At the same time, it will also send a NOTIFICATION message to the destination node, notifying the node of the allocated wavelength and timeslots for the connection. When the source node receives the ACKNOWLEDGEMENT message, it will start to transmit its data in the allocated timeslots on the allocated wavelength. Once the transmission is completed, the source node will send a RELEASE message to the core node. When the core node receives the RELEASE message, it will release the timeslots and wavelength allocated for the connection, and will forward the RELEASE message to the destination node. Therefore, the destination node is able to know the transmission duration of the connection through the NOTIFICATION and RELEASE messages. This information can be used in the detection and localization of a network failure.

We assume that both core switches are not equipped with detectors along the data paths between its input ports and output ports. Thus they have no ability to detect the power loss of an optical signal. On a physical path from one edge node to another edge node, there are three points that are easily exposed to failures: (1) an upstream link, (2) a downstream link, and (3) an optical switch fabric in the core switch, as shown in Figure 2. In the core switch, each switch fabric is an opto-electronic device that is more easily exposed to a failure. The multiplexers and demultiplexers are passive devices and are thus robust. For this reason, we consider three different failure scenarios: a link cut on an upstream link, a link cut on a downstream link, and a fault in a switch fabric.

We propose two alternative techniques for the fault localization. One requires a control channel on each wavelength of a fiber link and the other requires a small data block to be transmitted in each non-allocated data channel.

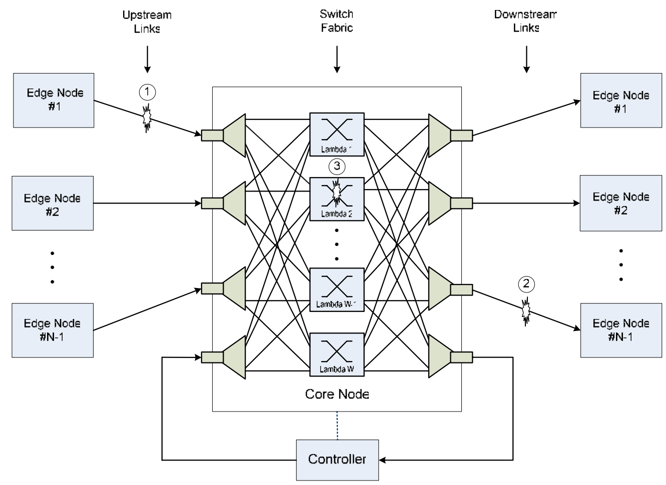

Figure 2 Illustration of failure scenarios

A. Alternative 1

We assume that a control channel is allocated on each wavelength of a fiber link (one time slot in each TDM frame) to facilitate fault localization. The fault localization process in different failure scenarios is described as follows.

(a) On an upstream link

In the network, each frame on a wavelength consists of a fixed number of timeslots, among which one is used as a control channel and the other are used as data channels. If a fiber cut occurs on an upstream link, the signal of all connections on the failed link will be lost. This loss can be detected at the destination edge node of each connection. However, it would take a considerable delay for the failure to propagate to the destination edge node. On the other hand, the failure would also result in the power loss in the control channel of the failed link. This loss can be detected at the receiver of the controller associated with the core node by checking the power loss in the control channel, which would significantly reduce the detection time. To avoid misjudgment in the case of no control messages, each edge node should continuously transmit a small data block in the control channel on each wavelength even if it has no control messages to send.

(b) On a downstream link

If a fiber cut occurs on a downstream link, the signals of all connections on the failed link will be lost, resulting in the power loss in all channels on each wavelength. Such power loss, including that in the control channel on each wavelength, can be detected at an edge node. Thus, the edge node can identify that the failure occurs on the downstream link. To avoid misjudgment in the case of no data channel allocated to any connection on the downstream link, the core node should continuously transmit a small data block in the control channel on each wavelength even if it has no control messages to send.

(c) In a switch fabric

If a fault occurs to a switch fabric for a particular wavelength, the signals of all connections using that wavelength on each downstream link will be lost, resulting in the power loss in all channels on that particular wavelength, including the control channel. This can be detected at the edge node of each downstream link. At the same time, the signal power of all connections using other wavelengths remains a normal level. If such results are detected, the edge node can identify that a fault has occurred to a switch fabric corresponding to that particular wavelength. To avoid misjudgment in the case of no data channel allocated to any connection on a particular wavelength, a core node should also continuously transmit a small data block in the control channel on each wavelength even if it has no control messages to send.

B. Alternative 2

We assume that a source edge node transmits a small data block at the beginning of each non-allocated time slot. To reduce the detection time, the detection of an upstream link failure is still performed at the controller. With the aid of such small data blocks, a destination edge node can immediately detect and localize a downstream link failure or a switch fabric fault, significantly reducing the detection time. Moreover, there is no need for a control channel on each wavelength of a fiber link. Only one control channel is sufficient on each fiber link for signaling control, which is an advantage compared with alternative 1. However, this alternative introduces some power consumption for free (non-allocated) time slots. To minimize power consumption, this data block should be as short as possible, only including a few octets of control information. A performance comparison between alternative 1 and alternative 2 will be given in Section 4 in terms of data loss and fault detection time.

2. Fault Advertisement Protocol

Based on the fault localization technique proposed above, we now present a fault advertisement protocol, which can be incorporated into the signaling protocol used in the network to facilitate static protection or dynamic restoration.

For ease of exposition, one of the core nodes is called core node A and the other is called core node B, as shown in Figure 1. The subnet with core node A is called subnet A and the one with core node B is called subnet B. Without loss of generality, we consider a failure that occurs in subset A and assume that a static protection paradigm is employed in the network, in which a backup path is established at the same time a primary path is established for a connection request. The main procedures of the fault advertisement protocol for different failure scenarios are described as follows, as shown in Figure 3.

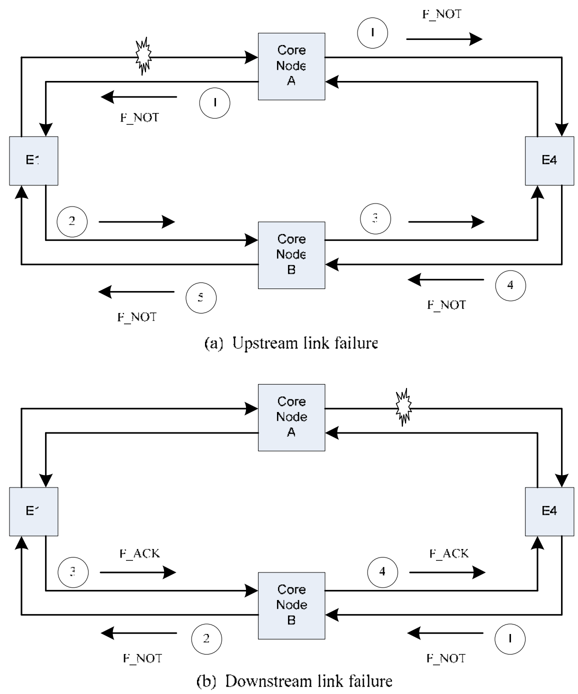

(a) Upstream link failure

(b) Downstream link failure

Figure 3 Signaling control process

(a) A failure on an upstream link

- Once a failure is detected by the controller associated with core node A, the controller immediately sends a FAILURE_NOTIFICATION (F_NOT) message to the source edge node of the failed link, say, E1, and the destination edge node of all disrupted connections (e.g., E4), as shown in Figure 3a (Step 1). At the same time, it releases all the timeslots allocated to the disrupted connections.
- Once node E1 receives the F_NOT message, it first sends a FAILURE_ACKNOWLEDGMENT (F_ACK) message to the destination edge nodes of all disrupted connections (e.g., E4) via core node B, as shown in Figure 3a (Step 2), and then modifies its link state table. Later, it switches over to the corresponding backup paths of the disrupted connections and continues to transmit data on the backup paths.
- When core node B receives the F_ACK message from

node E1, it will forward the message immediately to the destination edge nodes of all disrupted connections (e.g., node E4), as shown in Figure 3a (Step 3).

- Once node E4 receives the F_ACK message from core node B, it will switch over to the corresponding backup path of the connection and continue to receive data on the backup path.

Note that the destination edge node of each disrupted connection (e.g., E4) should also forward the F_NOT message it receives from core node A to node E1 via core node B in order to address the case in which both the upstream link and the downstream link are cut, as shown in Figure 3a (Step 4 and Step 5). In this case, an additional delay will be introduced, which will be analyzed in Sections 4.1 and 4.2.

(b) A failure on a downstream link or in a switch fabric

- Once a failure is detected and localized at an edge node, say E4, the node first sends a F_NOT message to core node B via a control channel on its upstream link, reporting the failure to the core node (and indirectly to all other edge nodes), as shown in Figure 3b (Step 1). The F_NOT message contains related information on the failure, such as failure location.
- Once core node B receives the F_NOT message, it will modify its link state table and at the same time broadcast the F_NOT message to each of the other edge nodes via a control channel on the corresponding downstream link, as shown in Figure 3b (Step 2).
- Once an edge node, say E1, receives the F_NOT message, it will immediately send an F_ACK message to node E4 via core node B, as shown in Figure 3b (Step 3), and then modify its link state table. After that, an edge node may need to switch over to a backup path, depending on the failure location.
  - If the failure is located on the downstream link to node E4, the edge node of each disrupted connection should switch over to the corresponding backup path.
  - If the failure is located in a switch fabric for a particular wavelength in core node A, each edge node that uses that particular wavelength should switch over to the corresponding backup path.
- If core node B receives an F_ACK message, it will forward the message immediately to node E4, as shown in Figure 3 (Step 4).
- Once node E4 receives an F_ACK message from the source edge node of a disrupted connection, it will switch over to the corresponding backup path and continue to receive data on the backup path.

Note that if the network employs a dynamic restoration paradigm, each edge node of a disrupted connection should first initiate a process to establish a backup path for the disrupted connection before it performs a switchover. Obviously, a dynamic restoration scheme would take longer time to restore the transmission than a static protection scheme.

IV. Performance Analysis

In this section, we analyze the data loss, detection time, and recovery time with the proposed fault detection and localization scheme in different failure scenarios. For a disrupted connection, we define the fault detection time as the time from the instant a failure occurs to the instant the destination node detects and localizes the failure, and the connection recovery time as the time from the instant the destination node detects and localizes the failure to the instant the source node switches over to the corresponding backup path of the disrupted connection. Without loss of generality, we still consider static protection and assume that the point of an upstream link failure or a downstream link failure is uniformly distributed along the corresponding link. The other notations used in the analyses are defined as follows.

- ρ : the upstream link load
- γ : the link transmission rate
- D_L : the total amount of data lost over a disrupted connection caused by a failure
- T_R : the average connection recovery time, i.e., the average time from the instant the destination node detects and localizes the failure to the instant the source node switches over to the corresponding backup path of the disrupted connection.
- F : the number of timeslots in each frame
- T_s : the length of a timeslot
- T_d : the average fault detection time, i.e., the average time from the instant a failure occurs to the instant the destination node detects and localizes the failure
- T_w : the average time it takes an edge or core node to wait for a control timeslot to send out an F_NOT message.
- T_p : the propagation delay of each fiber link. Here we assume for simplicity that the propagation delay is the same for all fiber links.

1. Data Loss

In the event of a failure, the source edge node of a disrupted connection will continue its transmission on the primary path until it receives an F_NOT message and switches over to the corresponding backup path. Accordingly, the data transmitted during the service recovery time will be lost. Moreover, the data that was already transmitted onto the link before the failure occurred, but did not yet pass the failure point, will also be lost.

Figure 4 illustrates a failure case in which the failure occurs somewhere on the primary path of a connection between edge node E1 and edge node E4 at time t_0. Obviously, the data that was already transmitted onto the link for a period T_2 but did not pass through the failure point at t_0 cannot be transferred to the destination node. The time at which the controller or the destination node detects the failure is $(t_0+T_1+\Delta_d)$, where T_1 is the propagation delay between the failure point and the controller or the destination node. There is also a propagation delay T_N for sending the F_NOT message from the detection point to the core node and the source edge node and some additional time Δ_d required to recover the disrupted service (in addition to the propagation delay). Therefore, the total amount data loss of the connection caused by the failure is

$$D_L = \begin{cases} \gamma \cdot \rho \cdot (T_1 + T_2 + \Delta_d + T_N) = \gamma \cdot \rho \cdot (2T_p + \Delta_d) & \\ \quad \textit{on an upstream link} & \\ \gamma \cdot \rho \cdot (T_1 + T_2 + \Delta_d + T_N) = \gamma \cdot \rho \cdot (4T_p + \Delta_d) & \\ \quad \textit{on a downstream link} & \\ \gamma \cdot \rho \cdot (T_1 + T_2 + \Delta_d + T_N) = \gamma \cdot \rho \cdot (4T_p + \Delta_d) & \\ \quad \textit{in a switch fabric} & \end{cases} \tag{1}$$

where $$T_N = \begin{cases} T_p & \textit{on an upstream lin} \\ 2T_p & \textit{on a downstream link} \\ 2T_p & \textit{in a switch fabric} \end{cases}$$

Note that in the case of a failure occurring on both the upstream and downstream links, an additional $2T_p$ propagation delay is added to T_N. The additional time Δ_d in the different failure scenarios is analyzed as follows.

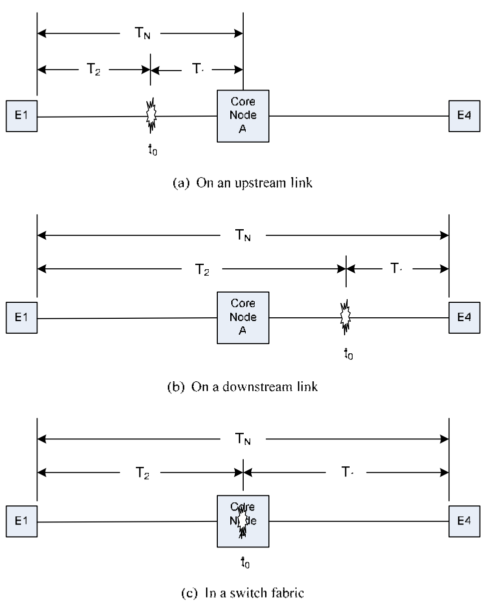

Figure 4 Illustration of a failure case

A. With Alternative 1

(a) A fault on an upstream link:

If a failure occurs on an upstream link, the average time it takes the controller to detect the upstream link failure based on the power loss in the control channel is $F \cdot T_s / 2$. The average time it takes the controller to wait for a downstream control timeslot to send out an F_NOT message T_w is also $F \cdot T_s / 2$. Hence, we have

$$\Delta_d = \frac{F \cdot T_s}{2} + T_w = F \cdot T_s \tag{2}$$

Note that in the case of a failure occurring on both the upstream and downstream links, an additional $2T_w$ propagation delay is added to Δ_d.

(b) A fault on a downstream link:

If a failure occurs on a downstream link, the destination edge node identifies the failure by checking the control channels. The average time it takes the destination edge node to detect the power loss in the control channels is $F \cdot T_s / 2$. The average time it takes the destination edge node to wait for an upstream control timeslot and the average time it takes core node B to wait for a downstream control timeslot to send out an F_NOT message are T_w equal to $F \cdot T_s / 2$, respectively. Hence, we have

$$\Delta_d = \frac{F \cdot T_s}{2} + 2T_w = \frac{3}{2} F \cdot T_s \tag{3}$$

(c) A fault in a switch fabric:

If a failure occurs in a switch fabric for a particular wavelength, the destination edge node identifies the failure by checking the control channel on that wavelength. The average time it takes the destination edge node to detect the power loss in the control channel on that wavelength is $F \cdot T_s / 2$. Hence, we have

$$\Delta_d = \frac{F \cdot T_s}{2} + 2T_w = \frac{3}{2} F \cdot T_s \tag{4}$$

B. With Alternative 2

With Alternative 2, the additional time Δ_d is the same as that with Alternative1 in the case of an upstream link failure. If a failure occurs on a downstream link or in a switch fabric, a destination edge node can immediately detect and identify a failure once it detects a power loss. Therefore, we have

$$\Delta_d = 2T_w = F \cdot T_s \tag{5}$$

Note that if the network uses a dedicated out-of-band control channel that is available at any time, a destination edge node or a core node can send an F-NOT message without any delay. In this case, T_w is approximately equal to zero.

2. Fault Detection Time

The average fault detection time T_d in different failure scenarios with alternatives 1 and 2 is analyzed as follows.

A. With Alternative 1

(a) On an upstream link

If a failure occurs on an upstream link, the average propagation delay between the failure point and the controller is $T_p/2$. The average time it takes the controller to detect the upstream link failure based on the power loss in the control channel is $F \cdot T_s/2$. Hence, we have

$$T_d = \frac{T_p}{2} + \frac{F \cdot T_s}{2} = \frac{1}{2}(T_p + F \cdot T_s) \quad (6)$$

(b) On a downstream link

If a failure occurs on a downstream link, the destination edge node identifies the failure based on the signals in the control channels. The average time it takes the destination edge node to detect the problem in the control channels is $F \cdot T_s/2$. Hence, we have

$$T_d = \frac{T_p}{2} + \frac{F \cdot T_s}{2} = \frac{1}{2}(T_p + F \cdot T_s) \quad (7)$$

(c) In a switch fabric

If a failure occurs in a switch fabric for a particular wavelength, the destination edge node identifies the failure based on the signal in the control channel on that wavelength. The average time it takes the destination edge node detects the power loss in the control channel on that wavelength is $F \cdot T_s/2$. Hence, we have

$$T_d = T_p + \frac{F \cdot T_s}{2} \quad (8)$$

Note that in the above analyses the propagation delay inside a switch, the time it takes for a node to process a control message, the transmission time of a control message, and the time it takes a source node to perform a switchover are ignored.

B. With Alternative 2

With Alternative 2, the detection time is the same as that with Alternative 1 in the case of an upstream link failure. In the case of a downstream link failure or a switch fabric failure, a destination edge node can immediately detect and identify the failure once it detects the problem. Therefore, we have

$$T_d = \begin{cases} \frac{1}{2}(T_p + F \cdot T_s) & \textit{on an upstream link} \\ T_p & \textit{in a switch fabric} \\ \frac{T_p}{2} & \textit{on a downstream link} \end{cases} \quad (9)$$

3. Connection Recovery Time

The average connection recovery time in different failure scenarios is given as

$$T_R = \begin{cases} T_w + T_p = T_p + \frac{1}{2}F \cdot T_s & \textit{on an upstream link} \\ 2T_w + 2T_p = 2T_p + F \cdot T_s & \textit{on a downstream link} \\ 2T_w + 2T_p = 2T_p + F \cdot T_s & \textit{in a switch fabric} \end{cases} \quad (10)$$

Note that in the case of a failure occurring on both the upstream and the downstream links, an additional $2T_w$ and $2T_p$ delay is added in T_R.

4. Numerical Results

We investigate the data loss and fault detection time in the different failure scenarios as well as the impact of different frame sizes on the data loss and fault detection time. We consider two network scenarios with different link distances: one has a link distance of 20 km, such as a metropolitan area network, and the other has a link distance of 2000 km, such as a wide area network (WAN). In addition, we assume that the length of a timeslot T_s is 10 μs and the link rate γ is 10 *Gbps*.

Figure 5 shows the normalized data loss (divided by γ×ρ) in the different failure scenarios with different alternatives and different network sizes. In a smaller network of 20 km link distance, one sees that an upstream link failure causes less data loss than a downstream link failure and a fabric fault, while a down stream link failure causes the same data loss as a fabric fault. In the case of an upstream link, Alternative 2 causes the same data loss as Alternative 1. In the other two cases, Alternative 2 causes less data loss than Alternative 1. Moreover, the frame size has a remarkable impact on the data loss. The larger the frame size, the larger the data loss. In a larger network of 2000 km link distance, however, the frame size has only a small impact on the data loss because this loss is dominated by the propagation delay.

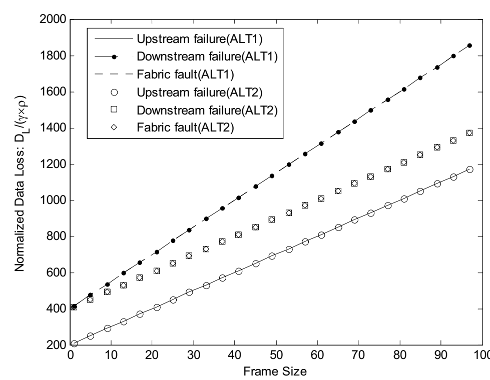

(a) Link distance=20 km

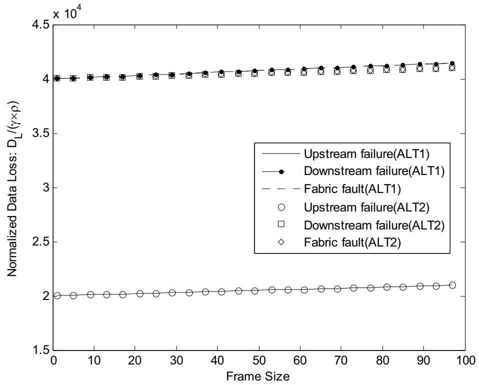

(b) Link distance= 2000 km

Figure 5 Data loss versus frame size

Figure 6 shows the fault detection time in the different failure scenarios with both alternatives. In a smaller network of 20 km link distance, the detection time in the case of a fabric fault is larger than that in the case of a link failure with Alternative 1. The frame size has a remarkable impact on the detection time in all the three failure cases. With Alternative 2, the detection time is the same as that with Alternative 1 in the case of an upstream link failure. In the other two cases, the detection time is smaller than that with Alternative 1 and is independent of the frame size. In a larger network of 2000 link distance, the impact of the frame size and the different alternatives on the detection time become much smaller. This is because of the domination of the propagation delay in the network. The detection time with a fabric fault is much larger than that with a link failure.

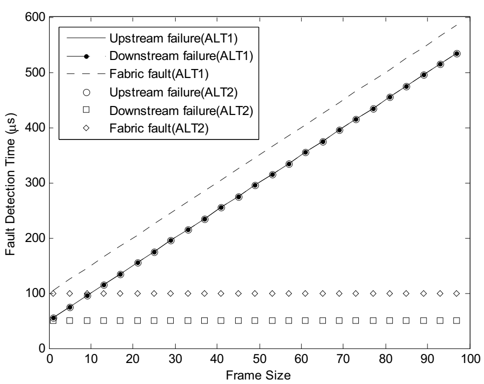

(a) Link distance=20 km

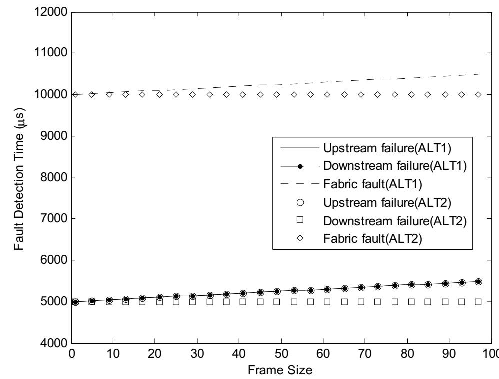

(b) Link distance=2000 km

Figure 6 Fault detection time frame size

V. CONCLUSIONS

Fault detection and localization is a crucial issue in all-optical networks. In this paper, we proposed a fault detection and localization scheme to facilitate the provisioning of service survivability in all-optical overlaid-star TDM networks. This scheme includes a fault localization technique and a fault advertisement protocol. For fault localization, two alternatives are proposed. Alternative 1 requires a control channel on each wavelength of a fiber link while Alternative 2 requires a small data block to be transmitted in each non-allocated data channel. The performance analyses show that in a smaller network Alternative 2 causes the same data loss as Alternative 1 in the case of an upstream link. In the other two cases, Alternative 2 causes less data than Alternative 1. In terms of detection time, a fabric fault needs a larger detection time than a link failure with Alternative 1. With Alternative 2, the detection time is the same as that with Alternative 1 in the case of an upstream link failure. In the other two cases, the detection time is smaller than that with Alternative 1 and is independent of the frame size. Moreover, the frame size has a remarkable impact on both the data loss and the detection time with Alternative 1. In a larger network, the impact of the frame size and the different alternatives on the data loss and the detection time becomes very small because of the domination of the propagation delay. Therefore, for smaller networks Alternative 2 is better than Alternative 1. But this is at the cost of power consumption for transmitting small data blocks in non-allocated data channels. For larger networks, Alternative 1 and Alternative 2 do not make much difference in terms of data loss and detection time. We believe that the results obtained provide useful guidelines for the design of all-optical overlaid-star TDM networks.

ACKNOWLEDGMENT

This work was supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada and industrial and government partners, through the Agile All-Photonic Networks (AAPN) Research Network.

The authors acknowledge some preliminary work on this topic performed by Danjiang Li (from the Design Institute of Radio, Film & TV (DIRFT), Beijing) during his visit at the University of Ottawa in 2004.

References

[1] G. v. Bochmann et al., "The agile all-Photonic Network: an architecture outline," Proc. of Queen's University Biennial Symposium on Communications, 2004.

[2] J. Zheng and H. T. Mouftah, *Optical WDM Networks: Concepts and Design Principles*, Wiley-IEEE Press, New Jersey, Jul. 2004.

[3] W. D. Grover, *Mesh-based Survivable Transport Networks: Options and Strategies for Optical, MPLS, SONET and ATM Networking*, Prentice Hall PTR, Upper Saddle River, New Jersey, Aug. 2003.

[4] C. Mas and P. Thiran, "A review on fault location methods and their application to optical networks," *SPIE Optical Networks Magazine*, Vol. 2, No. 4, Jul./Aug. 2001, pp. 73-87.

[5] C. Mas, P. Thiran, and J. L. Boudec, "Fault localization at the WDM layer," *Kluwer Photonic Network Communications*, 1:3, 1999, pp. 235-255.

[6] C. Mas and P. Thiran, "An efficient algorithm for locating soft and hard failures in WDM networks," *IEEE Journal on Selected Areas in Communications*, Vol. 18, No. 10, Oct. 2000, pp. 1900-1911.

[7] N. Ghani, S. Dixit, and T. Wang, "On IP-over-WDM integration," *IEEE Communications Magazine*, vol. 38, no. 3, Mar. 2000, pp. 72-84.