Journal of High Speed Networks 10 (2001) 247–268
IOS Press

Node wrappers for QoS monitoring in transparent optical nodes[1]

Muriel Médard [a], Stephen R. Chinn [b] and Poompat Saengudomlert [a]

[a] *Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

[b] *Malachite Technologies, Methuen, MA 01844, USA*

Abstract. Transparent optical nodes (TONs), such as all-optical switches and erbium-doped fiber amplifiers, are an increasingly important part of wavelength-division multiplexed (WDM) networks. Our goal in this paper is to consider how quality of service (QoS) may be monitored at such TONs. The question is particularly important as access to WDM networks, and associated security concerns, increase. Our paper has four parts. First, we present an overview of the vulnerabilities of TONs to QoS degradation for two main classes of TONs, namely all-optical switching nodes and amplifiers in optical networks. Second, we discuss the applicability of traditional supervisory methods to such degradations. Third, we propose a novel approach to monitoring QoS degradations in TONs. Our approach works by comparing the input and output at a node and deciding whether unacceptable service degradation has occurred at that node. Finally, we analyze the performance, under simple attack scenarios, of our approach for jamming attacks at transparent optical switching nodes and amplifiers. We show that our method is several orders of magnitude faster than bit error rate testers in detecting QoS degradations.

1. Introduction

Two trends emerge in the development of wavelength-division multiplexed (WDM) networks. The first is that WDM networks, even if they are not all-optical, are increasingly enabled by transparent optical nodes (TONs), which do not require optical to electrical conversion. The second trend is towards providing access to WDM networks to a wider set of users, thus moving WDM networks beyond simple trunking or backbone functions. A wider set of users entails the risk of user misuse of the network, as evidenced by the denial of service attacks spread through the current Internet. The goal of our paper is to provide a possible answer to the question: how can we take advantage of the speed and flexibility of TONs and maintain quality of service (QoS) in WDM networks, even in the presence of nefarious users?

Transparent optical nodes (TONs) are nodes in which data does not undergo optical-to-electrical or electrical-to-optical conversion. The concept of transparency evokes the idea that the light emerging from a TON is basically the original optical signal which was sent into the TON, albeit possibly switched and amplified. Transparency enables the co-existence of a variety of formats and signaling systems on the same node and also frees us from electronic bottlenecks which currently limit the data rate. WDM all-optical network (AON) testbeds [2,17,19,23] use exclusively TONs. TONs are also used within electro-optic WDM networks in conjuction with transponders or other nodes with electrical-to-optical or optical-to-electrical conversion. A perspective on access methods in optical networks can be found in [25]. Transparency, from which stem many advantages of TONs, produces a new range of QoS issues. In a network where a single entity controls all access ports, QoS requirements are met by ensuring that all access ports behave cooperatively. However, as networks grow in span and functionality, this becomes harder.

[1]This work was performed mostly while the authors were at MIT Lincoln Laboratory. This work was supported by a DARPA program on optical network security.

The mode of transmission and the type of hardware used in WDM networks opens a whole set of vulnerabilities when users willfully or inadvertently cause denial of service attacks. One approach to mitigating this threat is to restrict significantly what users may do. Such an approach, however, would stunt the growth of new services in WDM networks. The vulnerabilities of the Internet to denial of service attacks, such as flooding attacks, has convincingly shown that maintaining QoS in networks is both crucial and difficult. Without adequate QoS monitoring procedures, however, no guarantees can be provided. In this paper, we consider the following issue: in a network where trusted and untrusted users share the same infrastructure and, in particular, the same TONs, how can TONs monitor their QoS effectively to guard against denial of service attacks? In order to answer this question, we consider what vulnerabilities concern us and to what extent current monitoring techniques for WDM networks, developed for networks overseen by a single entity, are applicable when we consider a wide range of natural degradations and denial of service attacks. Based on our discussion of the shortcomings of existing monitoring systems, we propose a new paradigm for QoS monitoring for TONs. We may summarize the purpose of our paper as follows:

– to present denial of service attacks as a security vulnerability in WDM networks using TONs;
– to consider how existing supervisory methods for WDM systems apply to such attacks;
– to present a new system, based on a secure wrapper for TONs, to detect denial of service attacks;
– to analyze the performance of our method in detecting denial of service attacks for simple constant amplitude attacks.

The organization of our paper is as follows. In the next section, we present a brief overview of TONs and their inherent vulnerabilities to denial of service attacks. Overviews of security issues in TONs and AONs are given in [47,49,50]. We concentrate on two main types of denial of service attacks – in-band and out-of band jamming attacks. We next motivate the need for detecting and identifying QoS degradations due to natural degradations or denial of service attacks at vulnerable TONs. In Section 3, we overview current monitoring of optical systems. These means of monitoring are designed to diagnose failures and we discuss their shortcomings with respect to detection of denial of service attacks. In Section 4, we propose a general approach to creating a monitoring device which can be fitted onto a TON for QoS degradation monitoring. We refer to such a device as a wrapper, since its fits around an existing TON. We also illustrate a scheme which can be used to implement our monitoring wrapper in hardware. In Section 5, we present a network monitoring system based upon our wrappers. The monitoring system we propose relies upon alarms generated by individual wrappers at TONs to evaluate the state of the network. Our monitoring system then appropriately generates system alarms which indicate, with very high reliability, that the network is undergoing a denial of service attack. We consider the speeds at which attacks can be detected and show that our system is several orders of magnitude faster than bit error rate testers (BERTs). While our scheme can be used for detecting failures at nodes, we concentrate our discussion on attacks, because attacks have been studied less often and are more difficult to detect than failures. We discuss our results and further areas of research in Section 6.

2. TONs and requirements for QoS monitoring

2.1. Overview of TONs and their vulnerabilities

While most of the traditional security issues pertaining to traditional networks are applicable to TONs, TONs in WDM networks also have certain intrinsic security issues which are particular to them [47,49,50]. Transparency entails that signals remain in the optical domain at all times, in contrast to current electronic or electro-optic networks. While there are many variations in the implementations of TONs, they generally fall into two main types of components needed to ensure network functionality. Switching nodes provide the switching necessary to route traffic. These switching TONs may include demultiplexers, multiplexers and wavelength-selective switching planes, with or without reconfigurability. Amplifiers are used to overcome the attenuation which occurs naturally in the network. These amplifiers are usually erbium-doped fiber amplifiers (EDFAs), which can simultaneously

amplify several wavelengths. Note that different types of doping, such as Erbium-Fluoride doping, exist. Also, Raman gain amplifiers are emerging as viable alternatives to EDFAs. While their specific behavior depends on their implementation and use, they have a certain set of common vulnerabilities.

Switching WDM TONs exhibit crosstalk, the process by which a portion of a signal, say s_1, at a certain wavelength is superimposed upon another signal, say s_2. The most pernicious type of crosstalk is coherent (in-band) crosstalk, in which two signals at the same wavelength, which traverse the same switch, interact with each other [4,9,14,28–30,34,37,43,58,61,63,66,67,74,76]. Crosstalk indicates the proportion of the signal s_1 that is superimposed upon s_2. The amount of crosstalk in common current demultiplexers is in the range of 10^{-2} to $10^{-3.5}$ (i.e., -20 dB to -35 dB). Crosstalk may be used by s_1 to jam s_2 if s_1 is sufficiently powerful with respect to s_1. Crosstalk may also be used by s_2 to eavesdrop on s_1. Note that some other technologies, such as micromachined optical cross-connects [38,44,45] have much lower crosstalk levels.

EDFAs exhibit out-of-band jamming or gain competition [10,20,22]. An EDFA can provide simultaneous gain for several different wavelengths over the bandwidth of the amplifier. This gain is caused by a population inversion between the ground and metastable manifolds of the Erbium ions created by an optical pump excitation. To the extent that homogeneous broadening applies to these manifolds (an excellent approximation for most EDFAs), all wavelengths are affected by the same population inversion, and their gains differ only through differences in the wavelength-dependent emission and absorption cross-sections. If a strong signal at one wavelength saturates the gain by reducing the population inversion, signals at other wavelengths will also experience a gain reduction. This cross-gain interference may allow an attacker to insert a strong signal, either in or out of the users' band, reduce the gain of the users' signals, and potentially cause a degradation in their signal to noise ratio. The gain of EDFAs will change only with signal variations that are slow with respect to the (saturated) recovery times of the Erbium inversion. This is what allows EDFAs to amplify high-frequency signals without cross-channel or inter-symbol interference effects. Cross-gain attacks need not be continuous, and may be harder to detect if the attacker uses pulses with a very low duty cycle (low average power) but high-to-normal peak power, having pulse durations long compared to the signal bit period and saturated recovery time (greater than 10's of microseconds). Simulated examples of such attacks are given [20], where a 100 mW pulse attack at 1530 nm can cause a 5–6 dB reduction in 1540–1555 nm signals within 0.5 ms. A common strategy to reduce such gain transients is to provide feedback in the pump excitation to compensate for gain reduction. However, this process, commonly termed automatic gain control, will be limited in speed by electronic detection circuits and the recovery dynamics of the Erbium ions. Another alternative is to provide optical stabilization by means of a continuous, out-of-band saturating signal that is generated by oscillation in the EDFA itself. In such gain-clamped EDFAs, cross-gain transients can be reduced by an order of magnitude. We consider gain-clamped EDFAs for our simulations in Section 5. Note that the jamming attack by in-band superposition of a user's signal onto another can also be realized by physically tampering with the fiber and introducing a jamming signal for some set of wavelengths (possibly with hopping in time among those wavelengths). EDFAs also cause natural degradation by introducing amplitude stimulated emission (ASE), which may be modeled as additive white Gaussian noise (AWGN).

We are interested in denial of service attacks. Thus, we do not consider here eavesdropping attacks, which affect privacy but do not disrupt service. Given our discussion of the main types of TONs and their vulnerabilities, we concentrate on two types of attacks: in-band jamming at switching nodes and out-of-band jamming at amplifiers.

2.2. Motivation for monitoring in WDM networks

There are many reasons why denial of service attacks which degrade QoS in WDM networks must be detected and identified at all points where attacks may occur. Moreover, the speed of attack detection must be commensurate with the data transmission rate of the WDM network. The four main reasons may be summarized as follows. First, the *high data rates* of WDM networks entail that large amounts of data can be compromized in short time. Second, the *large network latency* causes large amounts of data to be already in flight in the network by the time an attack is detected. Third, an *erroneous diagnostic* of an attack as a failure can cause widespread failure. Finally, if attacks are not identified at *all possible points* of attack in the network, inappropriate corrective action may be taken by the network management system (NMS). In the remainder of this section, we elaborate on each of these points.

Data rate and denial of service attack duration issues. For a given period of time when QoS is degraded owing to a denial of service attack or a failure, the number of bits affected is proportional to the data rate. For instance, a downtime of 1s may entail the loss of hundreds of Gbits of data in an AON. If we are to retransmit data affected by denial of service attacks, then a high data rate will entail large memory requirements for buffering data for retransmission. Such retransmissions will also burden the NMS. Therefore, we must have attack detection and identification which is commensurate with the data rate.

Latency issues. Latency refers to the number of bits in flight and is also related to the data rate. For a given delay, for instance for a given distance traversed by optical signals, latency increases proportionally to the bit rate. A delay of milliseconds may entail that several Gbits have already entered the network by the time a degradation was identified at a peripheral node. The data in flight may be beyond the reach of certain corrective measures, such as rerouting at the node where the failure was detected.

Differences between failures and attacks. Even if a network is set up to deal with failures, it may not respond properly to attacks which affect its QoS. A discussion of issues differentiating attacks from failures is given in [6]. In networks which perform automatic failure recovery, it is particularly important to identify an attack caused by the traffic itself from a failure which occurs because of natural fatigue of components or physical sabotage of the network. For instance, signals may be broadcast or multicast in the optical domain by using different splitting approaches [5,15,16,35,48]. Attacks may then spread to many users and many parts of the network. Rerouting away from a failure may be an appropriate response, but if several nodes appear to fail from a single attack, then rerouting away from all nodes under attack may lead to catastrophic network failure. Finally, attacks can be sporadic or intermittent in such a way as to attempt to avoid detection, whereas failures will not employ any means of avoiding detection.

Importance of diagnostic at all possible denial of service attack points. Identification of attacks should take place at all possible denial of service attack locations. Otherwise, an incorrect diagnostic may be given by the NMS. For instance, suppose an attack spreads through several nodes, say from node 1, to node 2, but the attack is only detected at node 2. Corrective measures aimed at node 2 may fail to overcome the problem, which actually originates at node 1. In particular, if multicasting or other splitting takes place, then a single attack originating at a node may affect several nodes, so localizing the first attacked node is important.

The type of QoS monitoring which is appropriate for TONs is an important issue, since such monitoring may differ significantly from the type of monitoring that would take place at traditional electronic nodes. In particular, since transparency enables the coexistence on the same node of several types of protocols, signaling schemes, rates and coding, QoS monitoring at TONs should be well suited to a variety of different data streams. Thus, QoS monitoring cannot be dedicated to a particular protocol, say IP packets or ATM cells. Moreover, coding protection and similar electronic monitoring mechanisms are not applicable to the very high data rates for which TONs are built. Since current WDM systems generally use on-off keying (OOK), we consider QoS monitoring based on OOK signaling. The criterion we select is worst-case bit error rate (BER) – i.e., what is the worst-case number of errors that will occur without an alarm being generated for the highest data rate supported by the network. From our discussion in this Section, we know that denial of service attack duration is critical for TONs. Thus we consider the amount of time that a certain BER may be degraded without an alarm being generated. Finally, since any monitoring system may be subject to errors, we consider FP, the probability of a false positive alarm, which is generated erroneously in the absence of an attack. We also consider FN, the probability of a false negative result, i.e., no alarm being generated when one should have been generated. Thus, we have four criteria which we consider for QoS monitoring: BER level, time during which that BER level was sustained, FP and FN. Note that TONs are subject to degradations of service due to naturally occurring effects in the network. In general, such degradations, including receiver noise, noise in the fiber and noise from network components, such as ASE, can be well modeled by additive white Gaussian noise (AWGN). In Section 5, we consider the effect of such natural degradations when evaluating the performance of a network monitoring system based upon wrappers.

3. Applicability of current testing methods

In this Section, we examine qualitatively the applicability of current diagnostic and supervisory techniques to certain broad classes of denial of service attacks. Following our previous discussion, we select two different general types of denial of service attacks: in-band jamming at switching devices and out-of-band jamming at amplifiers. The supervisory techniques we consider may be broadly arranged into two categories. The first category is of methods which perform statistical analysis of the communications payload: power detection, optical spectral analyzers (OSAs) and BERTs. The second category is of methods which measure a signal devoted to diagnostic purposes: pilot tones and optical time domain reflectometers (OTDRs). We first give a brief overview of these methods and then examine, first for in-band jamming and next for out-of-band jamming, their applicability. Note that an overview of the operation of supervisory techniques can be found in [21].

3.1. Overview of supervisory techniques

Power detection. Power detection generally describes the measurement of power over a wide band. Because we are comparing against an expected value, a slight decrease in power may take a long time to detect. If we use some law of large numbers for our statistical analysis, then a very long averaging time may be necessary to establish with reasonable certitude that a deviation of the sample mean from the statistical mean was statistically significant. As an example, consider that average power is obtained by integrating over 50 μs. Let us suppose that we monitor a single wavelength carrying 2.5 Gb/s. If we have a denial of service attack which significantly reduces power in one bit out of 10^5 bits, the BER will be brought down to 10^{-5}, which is significantly above the 10^{-9} requirement of SONET, for instance. However, the power reduction ($-100\,000$ dB) is far too small to warrant an alarm, since such a power reduction uniformly distributed over all bits would have no effect.

In the case of out-of-band jamming causing gain competition, the received power may be decreased. However, certain gain competition attacks may lead to a severe degradation in SNR without a degradation in total power. Suppose that a signal, s, must traverse two EDFAs, A_1 and A_2. If there is gain competition at A_1, the signal s may not be adequately amplified. At A_2, the signal received will have proportionally more ASE noise from A_1 than if there had been no gain competition. If the gain of A_2 were fixed, the signal output from A_2 would be lower in power than if there had been no gain competition at A_1. With some automatic gain control at A_2, the signal received after A_2 may have sufficient power but may consist mostly of amplified ASE from A_1 superimposed upon ASE from A_2. Thus, the power at the output of A_2 would be acceptable, while the SNR would not be.

OSAs. OSAs, as their name indicates, display the spectrum of an optical signal. There are many implementations of OSAs [21,64]. Note that OSAs are usually intended to be used by an operator. Therefore, unless there is significant programming to analyze the output of the OSA and map it to the generation of different types of alarms, it is not as convenient a diagnostic tool for the automatic generation of network alarms as the method discussed previously. Jamming attacks which significantly affect the spectrum will be detected by an OSA. However, OSAs suffer from the shortcomings of statistical comparisons between sample averages and statistical averages. An extreme example is that where every bit a random sequence with balanced 0 s and 1 s with Bernoulli distribution has every bit negated. The spectrum of the resulting sequence will still correspond to a Bernoulli distribution on the sequence, although all the bits will have been corrupted. Moreover, OSAs based on gratings have very slow responses with respect to data rates.

OSAs may be of use to determine the source of a gain competition attack, as long as the band that is analyzed by the OSA is sufficiently large to encompass the carrier frequency of an ut-of-band attack. An OSA may be able to show the presence of such an out-of-band attacker (say at 1530 nm), even though a power detection on the individual channels will not. Still, issues of difficulty of processing the output of OSAs and of slowness of OSAs remain.

BERTs. BERTs operate by comparing a received pattern with the pattern which was known to have been sent. Given the number of discrepancies which are found, the BER of the transmission is estimated. The time it

takes for a BERT to establish the BER will depend on the BER and the data rate. For instance, at 1 Gpbs, it takes several seconds for a BERT to establish with good statistical accuracy that the BER has been degraded from 10^{-9} to 10^{-8}. Thus, tens of errors, many of which might not be corrected if the error-correction codes are not designed to operate above 10^{-9} BER, may occur in the time it takes for the degradation to be detected. BERTs suffer from the same drawbacks as OSAs. Moreover, other characteristics of BERTs limit their use for attack detection at TONs. The first characteristic is the fact that BERTs work with a given test sequence. They do not examine the actual data communications. A sporadic jamming attack, for instance intermittent or wavelength-hopped, might therefore escape detection for a long time, until it coincided in time and wavelength with a test BERT sequence. Secondly, BERTs detect actual errors rather than degradations. This feature means that a degradation will not be noticed until after it has led to errors. This observation will be important when we develop our attack detection scheme. Finally, BERTs work on particular signaling schemes and thus do not allow the coexistence of various schemes on a TON.

Pilot tones. Pilot tones are signals which travel along the same links and nodes as the communication payload but which are distinguishable from the communication payload. Pilot tones in WDM networks are often at different carrier frequencies than the transmitted signal, but they might also be distinguished from the communications payload by certain time slots (in a TDM system) or certain codes (in a CDMA system). Discussion of pilot tones and their implementations, including sub-carrier multiplexed (SCM) pilot tones, can be found in [7,26,32,41,46, 53,70]. The purpose of the pilot tones is to detect transmission disruptions and possibly carry some signaling for the NMS.

Pilot tones will not be effective in detecting jamming attacks unless those attacks cover the wavelengths at which the pilot tones are carried. Even for SCM pilot tones, an attacker may be able to introduce a jamming signal which disrupts communications without significantly affecting the SCM pilot signal. Unless the pilot tones are hidden or dynamically hopped within the transmission band, the attacker may be able to avoid the pilot tones when jamming. Low rate amplitude modulation (using SCM) may be considered as a very slow (with respect to the data rate) averaging. Therefore, the communications signal may be significantly affected without impinging upon the detection of the pilot tone. Note that pilot signals may be subject to jamming themselves. Gain competition affects all wavelengths through an amplifier, although not all wavelengths are equally affected and there is dependence upon the saturating wavelength. If the pilot signals traverse the same amplifiers as the communication signals, then the pilot signals should be affected by gain competition when the communication signals are. If the pilot signals are amplified separately, then they will not enable detection of a gain competition attack.

The pilot tones may, under certain conditions, not be of use in detecting gain competition attacks. For a tone, detectability requires a much lower SNR than that required to obtain adequate BER on a communication.

OTDRs. OTDRs are a special application of pilot tones. Rather than analyze a pilot tone at the point where communication signal is received, the pilot tone's echo is analyzed. OTDRs are generally used to diagnose faults, bends and losses in fibers [3,5,64], but may be used as supervisory signals [18,42]. In branched networks, such as networks where wavelengths are demultiplexed onto different fibers, different branches may be individually probed [15,16,73].

If there is a wideband jamming attack, then some of the jamming signal will be returned in the reflections and should be observable. Such a diagnostic differs from that offered by a pilot tone in that the diagnostic may be done at the head-end. If there is some modulation on the OTDR probe signal [65], then detection of a jamming signal superimposed on the OTDR probe signal may be fairly sensitive.

The probing of EDFAs by OTDRs is not similar to the probing of fiber lines. If the EDFAs are unidirectional, then they are not useful for amplifying reflected signals and a bi-directional amplifier is required [42]. Therefore, we cannot expect OTDRs to be useful in determining gain competition among signals over a cascade of EDFAs. If an EDFA is used as a pre-amplifier for the OTDR [71] as well as a power amplifier for the communications system, then gain competition at that EDFA should be detectable over the reflected OTDR probe. The EDFA probe signal, for the purpose of gain competition detection, then fulfills the same purpose as a pilot tone.

4. Monitoring denial of service attacks at TONs

Our discussion in the previous section has indicated that traditional supervisory and diagnostic methods, geared towards detecting failures, is to generate alarms in case of a detected failure. Our goal in creating our monitoring method is rapid, reliable detection of service denial attacks at each vulnerable TON in a WDM network. Moreover, we wish to provide wrappers, i.e., monitoring devices that can be placed around different TONs. There are many benefits to offering security plug-ins such as wrappers rather than securing the TONs internally. Secure wrappers can be inserted in legacy systems. They can also be adapted to different types of devices. As security needs change, wrappers can be added, upgraded or replaced. We do not address the issue of how to correlate the alarms due to attacks. The problem of identifying failures using alarms has been examined in [6,39]. There are several issues concerning the interpretation of and the reaction to alarms in very high-speed networks.

In this section and the following section, we present and analyze the performance of a novel approach for QoS monitoring to guard against denial of service attacks. We often refer to denial of service attacks simply as 'attacks'. In this section, we present a general scheme to construct wrappers for TONs. These wrappers generate alarms when their output exceeds a certain threshold. We discuss the main features needed to implement such wrappers and propose a sample design for hardware implementation [51,52]. In the next section, we present a network QoS monitoring system based on alarms generated by wrappers. We investigate the performance of this decision system in the presence of in-band and out-of-band jamming attacks. The QoS monitoring criterion we consider is in terms of the following parameters: FP probability, FN probability, BER degradation detected, time until detection. In Section 6, we give conclusions and present directions for further research.

4.1. General scheme to build monitoring wrappers

Our QoS monitoring device is designed to wrap around vulnerable TONs. Each wrapper contains two taps, one at the input and the other at the output of a network node. Figure 1 shows the schematic diagram of the approach we propose. Denote the signal at the first tap by $s(t)$, which after passing through a device delay T_D, becomes $s(t - T_D)$. Denote the sum of the signal and noise at the second tap by $s'(t - T_D) + n(t - T_D)$, or $\gamma s'(t - T_D) + n(t - T_D)$ for a node with an amplifier. Note that we describe the signal component at the second tap by $s'(t - T_D)$ instead of $s(t - T_D)$ to take into account a possible shift in phase and polarization due to normal operation of the network node. Our approach is to take these two signals and compensate for the phase shift and polarization due to normal operation of the network node, i.e., the wrapper makes the signals from the two taps coherent with each other. The wrapper then performs a signal subtraction and a square-law magnitude detection (i.e., direct detection) of the result.

Because our approach requires vector subtraction of the two optical inputs to the comparator, difficulties are caused by the need to maintain a steady phase relationship and equal polarization states in the input fields. Maintaining a constant phase relation is a general problem encountered in using optical interferometers. Operation of the optical comparator requires two separate functions, phase stabilization and polarization alignment. We give below one example of how each of these functions might be implemented, but many others are possible. For instance, the use of optical phase locked loops or interferometer stabilization, of which many examples exist in the literature, perform the function of optical phase stabilization. Active polarization alignment is a less common function, but at least one commercial instrument exists that dynamically measures the Stokes polarization state, which is a large part of the polarization vector control process. The major factors involved in choosing how to build the comparator will depend on desired frequency response and allowable cost and complexity.

We must point out that coherent subtraction is difficult to implement and that the wrapper must be carefully calibrated to the TON on which it resides. In the case of amplifiers, the amplification may need to be tracked to take into account amplification changes due to benign variations. Such variations may arise, for instance, as wavelengths are legitimately added or removed in WDM systems. In reconfigurable WDM switches, the state of the switch must be known so that the correct output/input comparison is carried out. We do not address such implementation issues but merely present a general set up for QoS monitoring.

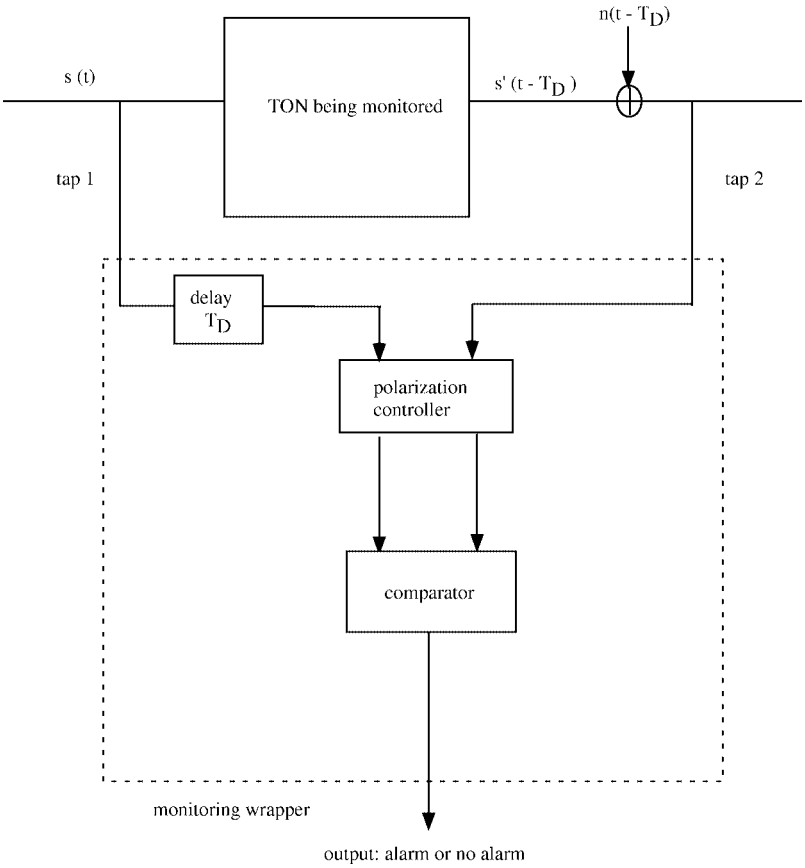

Fig. 1. Attack detection wrapper around a node.

4.2. Example implementation design

We discuss one possible method using planar optical waveguide technology to implement a device that can provide optical sum, difference, and quadrature fields for maintaining a stabilized difference output. Figure 2 illustrates the device which we discuss below. The optical comparator can be a 90° optical hybrid made from a monolithic multi-mode interferometric coupler [54,57]. In this case, two single-mode waveguides carrying the optical signals enter a wide, multi-mode region terminated with four single-mode output waveguides. To a first approximation, the multi-mode region is nW^2/λ long, where n is the effective index of refraction of the planar guided mode, W is the width of the multimode region, and λ is the free-space wavelength of the optical fields. With $\pi/4$ phase shift introduced between equal-intensity input fields, the output intensities from ports 1 through 4 are $\cos^2(\phi/2)$, $2[1 - \sin(\phi)]$, $2[1 + \sin(\phi)]$, and $\sin^2(\phi/2)$, respectively, with $\phi =$ optical phase difference between inputs. With $\phi = 0$, port 4 provides the difference intensity. Port 1 provides a sum intensity (which can be used for normalization purposes), and ports 2 and 3 provide the important quadrature outputs, whose detected intensities are used to provide feedback control to either of the input guides to maintain $\phi = 0$. If the input intensities are not equal, offset components are added to all the detected signals, but the difference between the detected quadrature signals still provides a phase-controlling feedback signal proportional to $\rho \sin(\phi)$ where ρ is the amplitude ratio. For example, in a silica waveguide system, these detected outputs can control the electrical input to a small heating element in proximity to a waveguide, to alter thermally the phase of light traveling through the guide. With a small enough heating element, the response of such a feedback loop should be adequate to track slow environmental changes in relative phase. In the nulled state ($\phi = 0$), the difference output has an intensity proportional to $(1 - \rho)^2$, which

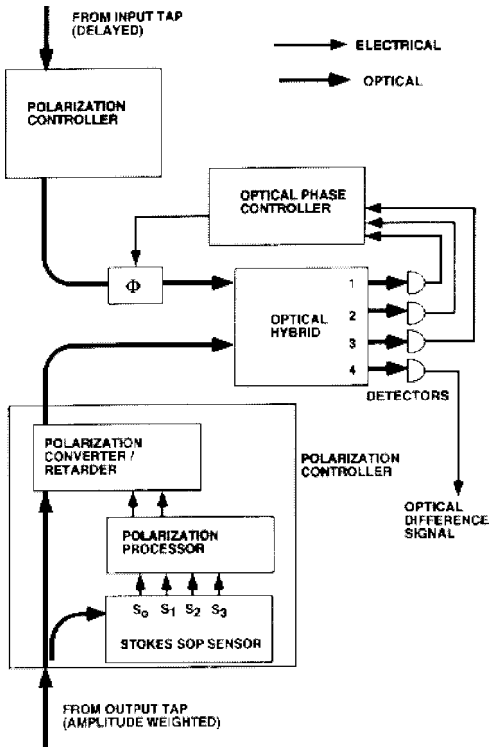

Fig. 2. Example implementation of an attack detection wrapper around a node.

can be used to sense attacks that perturb the amplitude. For rapid phase attacks, a difference signal proportional to $4\rho \sin^2(\delta\phi/2)$ can be used to detect phase perturbations outside the bandwidth of the slow phase-tracking loop.

Proper operation of the hybrid combiner requires that both input electric fields have the same state of polarization (SOP). Design of the hybrid is easier if this state is either TE (linear, in plane) or TM (linear, perpendicular). Since uncontrollable environmental factors will likely cause changes in the SOP of both inputs to the hybrid, these SOPs must be controllable. This requires the ability to sense the SOP of an input, and then to transform it to the desired state. One technique for measuring the SOP without the use of variable retarders and polarizers, is to measure the Stokes polarization parameters, normally used for display of the SOP on the Poincaré sphere. S_1 gives the difference between horizontal and vertical linear polarization powers, S_2 gives the difference between $+45^\circ$ and -45° linear polarization powers, and S_3 gives the difference between right- and left-hand-circular polarization powers (these are usually normalized by the total power, S_0). By using waveguide splitters, polarization selectors, and waveplate retarders, these Stokes components can be measured directly using fixed waveguide components [36,60] and photodiode detection. External processing of the signals can be used for normalization, automatic gain control, etc. Having obtained parameters describing the SOP from a sampled portion of the field, we must also have the means of altering the SOP of the field entering the hybrid. One means of doing this in a planar waveguide component has been described in [31] using a $LiNbO_3$ TE $\leftrightarrow$ TM converter/retarder. Using only two control voltages, any input SOP can be converted to any output SOP. The feedback loop can be closed using an external processor to convert the Stokes parameters to the control voltages necessary to obtain the desired TE or TM SOP at the hybrid input. This feedback would be applied to both inputs tapped off the optical device being probed. Drifts in the SOP are expected to be much slower than the feedback loop response (primarily limited by the external signal processing, whether digital or analog).

In order for the optical subtraction process to work, the two interfering signals must be phase coherent. Since the two signals in this case have a common source (one of them passing through the device under test), this means that

interference can occur if the optical path lengths to the comparator are matched to less than a coherence length of the signals. Another more intuitive way of understanding this requirement is as follows. For signals whose spectral width is determined by the data rate [fast modulation rates (e.g., 10 Gb/s) that exceed a typical source linewidth (e.g., 20 MHz)], the optical path delays must be matched to less than a bit interval. At 10 Gb/s, a one-bit path difference in optical fiber is 2 cm, so the path lengths must be matched to about 1 mm. Careful adjustment of fiber lengths is required, but an adjustable air-gap or fiber stretcher can provide fine tuning. Also, for high data rates and long devices (such as some EDFAs) or long return tap lengths, it may be necessary to match the optical path dispersions as well. This latter constraint is less severe, since the communication system itself generally requires the tested device to have low enough dispersion so that significant pulse spreading through it does not occur. The path matching or coherence length restriction also means that practical optical comparison will be limited to one WDM data channel (one wavelength) at a time.

5. A network QoS monitoring system based on alarms at wrappers

We would like to find out how our wrappers behave and construct a scheme for detecting jamming denial of service attacks. The gist of our scheme is to use individual alarms generated by our monitoring wrappers in order to generate system-level alarms which indicate, very rapidly and with great accuracy, whether a denial of service attack is taking place. In this section, we set up a scenario to base our analysis on and we analyze the performance of our scheme for both in-band and out-of-band attacks.

5.1. Scenario for analysis

Assume the transmission of data across M successive TONs equipped with monitoring wrappers. We shall consider the cases when M is equal to 1 and 10, respectively. For the analysis in this section, we assume that all M network nodes are similarly affected by a jamming attack, and that the corresponding effects are constant throughout an observation period. At each network node at any given bit time, we have two independent additive white Gaussian noise components denoted by N_R (real component) and N_I(imaginary component). Assume the noise variance $\sigma^2_{N_R}$ and $\sigma^2_{N_I}$ to be equal to $(1/M)\sigma^2_N$, so that the total noise variance across the light path at any given bit time is $2\sigma^2_N$. To obtain later numerical results, we assume σ^2_N to be 0.5.

For the sake of illustration, we assume an end-to-end SNR of 16 dB (approximated from $10\log_{10} 40$). In addition, we assume that the SNR degradation at TONs is much more significant than the SNR degradation along the fiber. Furthermore, we assume that all network nodes cause the same level of SNR degradation. Based on these assumptions, when M = 10, we have an SNR of 16 dB across the light path and an SNR of 26 dB across each TON. We transmit, using non-return to zero (NRZ) OOK, data bits whose values are equally likely to be 0 or 1. The transmission rate is 1 Gb/s and the bandwidth is 1 GHz. In addition, we assume direct detection at the receiver. Denoting the square magnitude of the ON level by P, we have that SNR $= (1/2)(P/2\sigma^2_N) + (1/2)(0/2\sigma^2_N)$. For the analysis in this section, we set (s_R, s_I) (representing the real and imaginary components) to be $(\sqrt{160}\sigma_N,0)$.

Knowing the value of (s_R, s_I), we want to find the cumulative probability distribution (c.d.f.) of $Y = |s + N|^2$, i.e., $Pr\{Y \leqslant A\}$ for a real $A > 0$. Y is a random variable corresponding to the output of a square-law detector. Note that $|s + N|^2 = |s_R + N_R|^2 + |s_I + N_I|^2$. We can consider Y as the sum of two random variables $Y_R = |s_R + N_R|^2$ and $Y_I = |s_I + N_I|^2$. Let us denote the characteristic function of Y_R by Φ_{Y_R}. We may write that

$$\Phi_{Y_R}(\omega) = \frac{1}{\sigma_N\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{j\omega(n+s_R)^2} e^{-\frac{n^2}{2\sigma^2_N}}\, \mathrm{d}n. \tag{1}$$

By performing a change of variable $y = (n + s_R)^2$, we can write Eq. (1) as

$$\Phi_{Y_R}(\omega) = \frac{1}{2\sigma_N\sqrt{2\pi}} \int_0^{\infty} e^{j\omega y} \left(e^{-\frac{(\sqrt{y}-s_R)^2}{2\sigma^2_N}} + e^{-\frac{(\sqrt{y}+s_R)^2}{2\sigma^2_N}} \right) \frac{\mathrm{d}y}{\sqrt{y}}. \tag{2}$$

From Eq. (2), the probability distribution function (p.d.f.) of Y_R is

$$p_{Y_R}(y) = \begin{cases} \frac{1}{2\sigma_N\sqrt{2\pi}} \left(e^{-\frac{(\sqrt{y}-s_R)^2}{2\sigma_N^2}} + e^{-\frac{(\sqrt{y}+s_R)^2}{2\sigma_N^2}} \right) \frac{1}{\sqrt{y}}, & \text{if } y \geqslant 0 \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

A similar expression holds for Y_I. Thus, using a change of variable, we can express the c.d.f. $Pr\{Y \leqslant A\}$ as

$$Pr\{Y \leqslant A\} = \frac{1}{2\pi\sigma_N^2} \int_0^{\sqrt{A}} \int_0^{\sqrt{A-z_R^2}} \left(e^{\frac{-(z_R-s_R)^2}{2\sigma_N^2}} + e^{\frac{-(z_R+s_R)^2}{2\sigma_N^2}} \right) \times \left(e^{\frac{-(z_I-s_I)^2}{2\sigma_N^2}} + e^{\frac{-(z_I+s_I)^2}{2\sigma_N^2}} \right) \mathrm{d}z_I \, \mathrm{d}z_R. \tag{4}$$

Note that we may simplify the expression on the right hand side of Eq. (4) by considering the integral as the sum of four terms. Using symmetry arguments and polar coordinates, we may write that

$$Pr\{Y \leqslant A\} = \int_0^{2\pi} \int_0^{\sqrt{A}} \frac{1}{2\pi\sigma_N^2} r e^{\left(-\frac{(r\cos\theta - s_R)^2}{2\sigma_N^2} - \frac{(r\sin\theta - s_I)^2}{2\sigma_N^2} \right)} \mathrm{d}r \, \mathrm{d}\theta. \tag{5}$$

We denote this cumulative distribution function (c.d.f.) as a function of σ_N^2, A, s_R and s_I by $F_d(\sigma_N^2, A, s_R, s_I)$. The dependence of $F_d(\sigma_N^2, A, s_R, s_I)$ on s is only through its norm $\|s\|$. We can reexpress $F_d(\sigma_N^2, A, s_R, s_I)$ as

$$F_d(\sigma_N^2, A, \|s\|) = \int_0^{2\pi} \int_0^{\sqrt{A}} \frac{1}{2\pi\sigma_N^2} r e^{\left(-\frac{(r\cos\theta - \|s\|)^2}{2\sigma_N^2} - \frac{(r\sin\theta)^2}{2\sigma_N^2} \right)} \mathrm{d}r \, \mathrm{d}\theta. \tag{6}$$

In the absence of an attack, we can compute the end-to-end BER. For analytical purposes, we shall occasionally express relevant quantities in terms of the function F_d defined in (6) before carrying out numerical computation. Let A denote the optimal threshold for the decision rule at the receiver, it follows that

$$BER_{\text{no attack}} = \frac{1}{2}(1 - F_d(\sigma_N^2, A, 0)) + \frac{1}{2} F_d(\sigma_N^2, A, \sqrt{160}\sigma_N). \tag{7}$$

The optimal decision threshold A is approximately 22 (or $44\sigma_N^2$). We then have $BER_{\text{no attack}} \approx 4.6 \times 10^{-10}$. From here on, we shall refer to this value as BER_{baseline}. We would like to construct a monitoring scheme that generates a system-level alarm whenever the end-to-end BER (denoted by $\text{BER}_{\text{end-to-end}}$) is greater than 10^{-8} (in comparison to BER_{baseline} of 4.6×10^{-10}). We shall compare the performance of our monitoring scheme to that of a BERT in terms of the required observation period.

If there is an in-band jamming signal denoted by J, the output of the square law detector will be $|n + J|^2$ (or $|(1/\gamma)(n + J)|^2$), which is likely to be much greater than $|n|^2$ (or $|(1/\gamma)n|^2$) for an attack with a sufficiently large in-band jamming signal. Based on the output of the square law detector, we set a decision threshold for an alarm generation. As we shall see later on, the value of such a threshold is a design parameter for our attack detection system. In analyzing gain competition, our analysis assumes that each EDFA at a network node makes up precisely for signal attenuation in the previous link of the light path. Consider an out-of-band jamming attack causing gain competition at an EDFA in a network node. An out-of-band jamming attack which robs the EDFA gain for the legitimate signal by a% will cause the output of the square law detector to be $|(a/100)s + (n/\gamma)|^2$.

5.2. Network QoS monitoring system for in-band jamming attacks

We investigate the effects of in-band jamming attacks. First, we find the $BER_{\text{end-to-end}}$ given the degradation of $a\%$ of the ON level (square magnitude) at each TON, or equivalently a possible total degradation of $Ma\%$ across M nodes. Note that we take a rather pessimistic view. In general, it is very hard for an attacker to jam multiple network nodes coherently so that the effects of a jamming attack add up constructively across M network nodes. In addition, we consider the worst case degradation, i.e., signals are always degraded in such a way that errors are more likely. Thus, the results of our analysis constitute worst case guarantees against service denial attacks. The following expression provides the upper bound on $BER_{\text{end-to-end}}$ with $a\%$ degradation allowed at each of the M network nodes

$$\overline{BER}_{\text{end-to-end}} = \frac{1}{2}\big(1 - F_d(\sigma_N^2, A, \sqrt{(.01Ma)160}\sigma_N)\big) + \frac{1}{2}F_d\big(\sigma_N^2, A, \sqrt{(1-.01Ma)160}\sigma_N\big), \tag{8}$$

Note that we use the notation $\overline{BER}_{\text{end-to-end}}$ to emphasize that the value is an upperbound on $BER_{\text{end-to-end}}$. Figure 3 shows the curve of $\overline{BER}_{\text{end-to-end}}$ versus the degradation at each of the M network nodes in percentage of the ON level ($160\sigma_N^2$).

A monitoring wrapper generates an alarm when the output of a square-law detector exceeds a certain threshold. We shall refer to an alarm generated by an attack detection device at a network node as a 'wrapper alarm'. FP depends solely on the threshold value, while FN depends on both the threshold value and the magnitude of a jamming signal. It is necessary at this point to distinguish between FP and FN at a single network node and FP and FN across multiple network nodes when $M > 1$. We will refer to the former case with FP_{node} and FN_{node}, and to the latter case with $FP_{\text{end-to-end}}$ and $FN_{\text{end-to-end}}$ respectively.

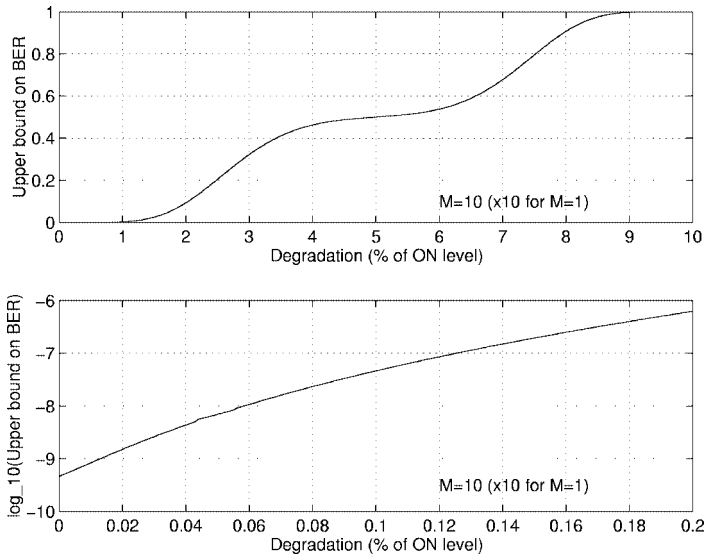

Fig. 3. Upper bound on BER versus degradation at each of the M (equal to 10) network nodes in percentage of the ON level ($160\sigma_N^2$). For $M = 1$, multiply the degradation level (horizontal axis) by 10. Note that the bottom plot is the zoomed version of the top one.

Suppose the threshold is set at t% of the ON level ($160\sigma_N^2$) and the magnitude of the jamming tone is a% of the ON level. Then

$$FP_{\text{node}} = 1 - F_d\left(\frac{1}{M}\sigma_N^2, (.01t)160\sigma_N^2, 0\right), \tag{9}$$

$$FN_{\text{node}} = F_d\left(\frac{1}{M}\sigma_N^2, (.01t)160\sigma_N^2, \sqrt{(.01a)160}\sigma_N\right). \tag{10}$$

The noise variance is adjusted to $(1/M)\sigma_N^2$ since we are looking at each individual node with the SNR of $16 + 10\log_{10} M$ dB instead of the end-to-end SNR of 16 dB. In the case of multiple network nodes ($M > 1$), we consider that a monitoring wrapper alarm at any bit time has occurred if at least one wrapper at any TON generates a wrapper alarm. This assumption leads to the computation of $FN_{\text{end-to-end}}$ and $\text{FP}_{\text{end-to-end}}$ given next.

We can think of $FN_{\text{end-to-end}}$ as the probability that none of the M TONs generates a wrapper alarm. Given a jamming signal, an event that an alarm is not generated depends on AWGN introduced at a network node. Since our model assumes that noise components introduced at different network nodes are independent, we have that the generation of wrapper alarms at different network nodes are independent. Therefore,

$$FN_{\text{end-to-end}} = (FN_{\text{node}})^M. \tag{11}$$

$FP_{\text{end-to-end}}$ is the probability that at least one of the M TONs generates a wrapper alarm. From the independence of wrapper alarms at TONs,

$$FP_{\text{end-to-end}} = 1 - (1 - FP_{\text{node}})^M. \tag{12}$$

Figures 4 and 5 show example curves of $FP_{\text{end-to-end}}$ and $FN_{\text{end-to-end}}$. The expected wrapper alarm rate is equal to $FP_{\text{end-to-end}}$ when no attack is present, and is equal to $1 - FN_{\text{end-to-end}}$ when there is an attack. Figures 6 and 7 show the curves of the expected wrapper alarm rate versus $\overline{BER}_{\text{end-to-end}}$ for different values of detection wrapper threshold. For each curve in Figs 6 and 7, the starting point on the left has the expected wrapper alarm rate equal to $FP_{\text{end-to-end}}$ while the corresponding BER is BER_{baseline} (equal to 4.6×10^{-10}). This is the situation with no

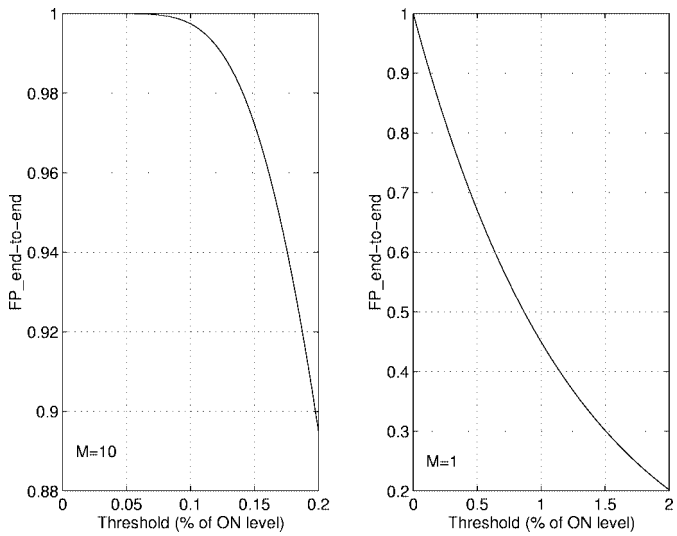

Fig. 4. $FP_{\text{end-to-end}}$ versus wrapper threshold (allowed degradation) at each network node.

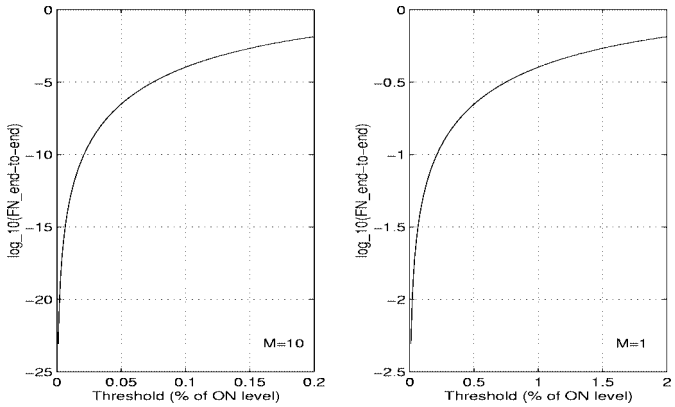

Fig. 5. $FN_{\text{end-to-end}}$ versus wrapper threshold (allowed degradation) at each network node given that the jamming signal magnitude is 0.06% of the ON level for $M = 10$ (0.6% for $M = 1$). With these jamming signal magnitudes, $BER_{\text{end-to-end}}$ can be as high as 10^{-8}.

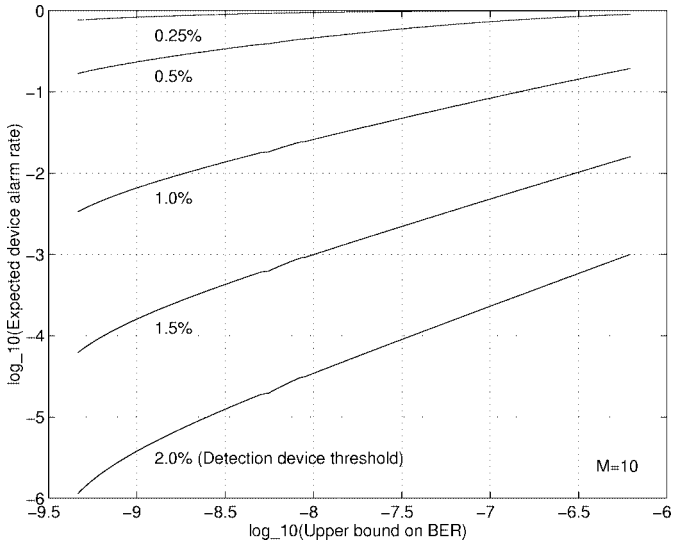

Fig. 6. Expected wrapper alarm rate versus $\overline{BER}_{\text{end-to-end}}$ for different wrapper thresholds ($M = 10$, in-band jamming).

attack. The points along the curve in the positive direction correspond to in-band jamming attacks with increasing jamming signal magnitudes. We would like to have a significant difference between the expected wrapper alarm rate with no attack and the rate with an attack causing $\overline{BER}_{\text{end-to-end}}$ to exceed 10^{-8}. In addition, we would like to be able to detect an attack in a small amount of time.

The presence of an alarm at a wrapper does not mean a transmitted bit is corrupted. An alarm simply notifies the user that the bit in transmission has a higher probability of decision error at the receiver. If an attack does not significantly corrupt transmitted bits and does not affect our communication more severely than naturally occurring noise, we do not worry about detecting such an attack. Based on these arguments, we propose two levels of alarms which will be referred to as a 'wrapper alarm' (first-level) and an 'system alarm' (second-level). We consider that a wrapper alarm at any bit time has occurred if at least one monitoring wrapper at any TON generates a wrapper alarm. A system alarm will turn on when the number of wrapper alarms in a given observation period (in bits) exceeds a certain threshold. Let FP_{attack} and FN_{attack} denote the FP and FN, respectively, of the system alarm at a single TON.

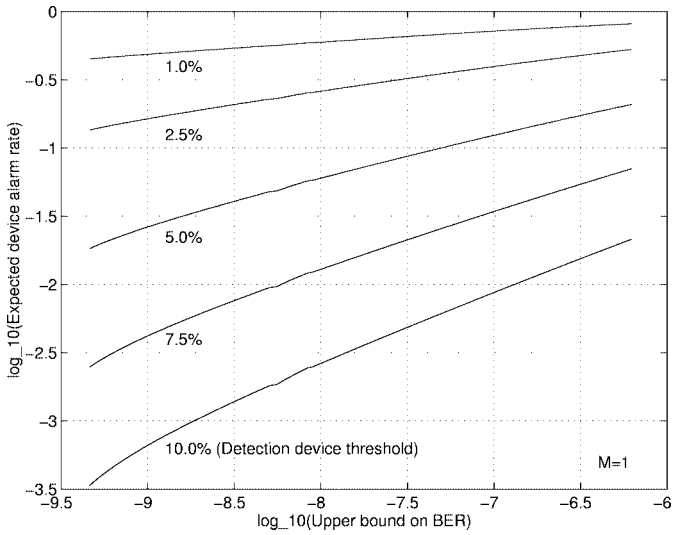

Fig. 7. Expected wrapper alarm rate versus $\overline{BER}_{\text{end-to-end}}$ for different wrapper thresholds ($M = 1$, in-band jamming).

Table 1

Expected number of wrapper alarms in an observation period. The top row contains the detection wrapper threshold values in percentage of the ON level ($160\sigma_N^2$). The first column contains the length of an observation period. The first number of each entry in the Table is the expected number of wrapper alarms when there is no attack, while the second number is the expected number of wrapper alarms when there is an attack causing $\overline{BER}_{\text{end-to-end}}$ to be approximately 10^{-8}. In this case, $M = 10$

	0.25%	0.5%	1.0%	1.5%	2.0%
10 μs	7,700/9,500	1,700/4,600	34/260	0.7/10	0.01/0.3
1 μs	770/950	170/460	3.4/26	0.07/1	0.001/0.03
0.1 μs	77/95	17/46	0.34/2.6	0.007/0.1	0.0001/0.003

Table 2

Expected number of wrapper alarms in an observation period (M = 1). All the entries have the same meanings as in Table 1

	1.0%	2.5%	5.0%	7.5%	10.0%
10 μs	4,500/6,000	1,400/2,600	270/600	25/130	6.5/26
1 μs	450/600	140/260	27/60	2.5/13	0.65/2.6
0.1 μs	45/60	14/26	2.7/6	0.25/1.3	0.065/0.26

We now have enough information to find an appropriate threshold value for the wrapper alarm as a percentage of the ON level ($160\sigma_N^2$) and an appropriate threshold value for the system alarm in terms of the number of wrapper alarms in an observation period. Note that our transmission rate is 1 Gb/s. Using the information in Figs 6 and 7, we can calculate the estimated numbers of wrapper alarms in an observation period. Tables 1 and 2 show the expected number of wrapper alarms for different detection wrapper thresholds and for different lengths of an observation period. In general, if the detection wrapper threshold is set too high, we need a long observation period to wait for a sufficient amount of wrapper alarms to occur before we can recognize an attack. On the other hand, if the threshold is set too low, we may not be able to distinguish between events associated with the presence and the absence of a denial of service attack. In what follows, we present two examples to demonstrate how we can assign threshold values for both levels of alarms. Our goal is to have both FP_{attack} and FN_{attack} approximately no greater than 10^{-10}.

Example 1. Assume $M = 10$ and the detection wrapper threshold of 0.25%. If an observation period is 1 μs, we have from Table 1 that the expected number of wrapper alarms in an observation period is 950 and 770 with and without an attack, respectively. At the transmission rate of 1 Gb/s, we transmit 1,000 bits in 1 μs. Hence, the number of wrapper alarms in an observation period is at most 1,000.

Denote the system alarm threshold by β, i.e., a system alarm is generated if β wrapper alarms occur during an observation period. Let X denote the number of wrapper alarms in an observation period. The values of FP_{attack} and FN_{attack} are

$$\begin{aligned} FP_{\text{attack}} &= Pr\{X \geqslant \beta \mid \text{no attack}\} \\ &= \sum_{i=\beta}^{1{,}000} \binom{1,000}{i} FP_{\text{end-to-end}}^{i} (1 - FP_{\text{end-to-end}})^{1{,}000-i}, \end{aligned} \tag{13}$$

$$\begin{aligned} FN_{\text{attack}} &= Pr\{X < \beta \mid \text{attack}\} \\ &= \sum_{i=0}^{\beta-1} \binom{1,000}{i} (1 - FN_{\text{end-to-end}})^{i} FN_{\text{end-to-end}}^{1{,}000-i}. \end{aligned} \tag{14}$$

We shall set β to be such that $FP_{\text{attack}} \approx 10^{-10}$ and find the corresponding value of FN_{attack}. We find the value of β to be ≈851. The corresponding FN_{attack} is $< 10^{-16}$. We conclude that, with an observation period of 1 μs, a detection wrapper threshold of 0.25%, and a system alarm threshold of 851, our detection scheme can detect a constant jamming attack (yielding the BER above 10^{-8}) that lasts for longer than 1 μs.

Example 2. Assume $M = 1$ and the detection wrapper threshold of 1.0%. For an observation period of 1 μs, we have $\beta \approx 550$. The corresponding FN_{attack} is 7.41×10^{-4}, which is too high for our purpose. If we change the observation period to 10 μs, we can choose β to be 4,960. The corresponding FP_{attack} is $< 10^{-10}$ and the corresponding FN_{attack} is $< 10^{-16}$. Thus, with an observation period of 10 μs, a detection wrapper threshold of 1.0%, and a system alarm threshold of 4,960, our detection scheme can detect a constant jamming attack (yielding the BER above 10^{-8}) that lasts for longer than 10 μs.

In both examples, the required observation period is much shorter than the time required if we were to use a BERT to detect the same jamming attacks. In particular, we detect simple in-band jamming attack scenarios in less than 10 μs in our examples. A similar task of distinguishing between the BERs of 4.6×10^{-10} (the value of BER_{baseline}) and 10^{-8} would require several seconds if we were to use a BERT. For example, using an observation period of 10 s, a BERT can expect 4.6 bit errors when there is no attack and 100 bit errors when there is an attack yielding the BER of 10^{-8}. A shorter observation period will not work since the expected number of bit errors in the presence of an attack do not differ significantly from the expected number of bit errors in the absence of an attack. Therefore, our QoS monitoring scheme is *6 orders of magnitude faster than a BERT*. In conclusion, we have demonstrated with examples how one can construct a jamming attack detection system based on wrappers at TONs and identify the corresponding observation period.

5.3. Network QoS monitoring system for gain competition at EDFAs

In this section, we investigate the effects of out-of-band jamming attacks causing gain competition at EDFAs. We then demonstrate with examples that our proposed monitoring scheme using alarms at wrappers can also be applied to detect out-of-band jamming attacks. We assume that the gain value γ is fixed to be that of a gain clamped EDFA without attack. In the case of in-band jamming, we do not specify the source of noise. Note that an important property is the dependence of the ASE noise variance on the EDFA gain. We obtain the data of noise variances together with EDFA gains from [20]. This set of data corresponds to the transmission at 1540 nm where the gain fluctuation is the largest during an attack. Table 3 shows the values of noise variances associated with EDFA

Table 3
Noise variances (σ_N^2 and $\sigma_{\tilde{N}}^2$) associated with EDFA square magnitude gains (γ^2) for the transmission at 1540 nm and a bit rate of 1 Gb/s using 1 GHz bandwidth

EDFA gain	Degradation	Noise variance	EDFA gain	Degradation	Noise variance
86.160	0%	1.7590×10^{-7}	85.585	0.6674%	1.7489×10^{-7}
86.052	0.1253%	1.7571×10^{-7}	85.424	0.8542%	1.7461×10^{-7}
85.908	0.2925%	1.7546×10^{-7}	85.135	1.1896%	1.7410×10^{-7}
85.786	0.4341%	1.7524×10^{-7}	84.580	1.8338%	1.7313×10^{-7}
85.679	0.5583%	1.7506×10^{-7}	84.234	2.2354%	1.7252×10^{-7}

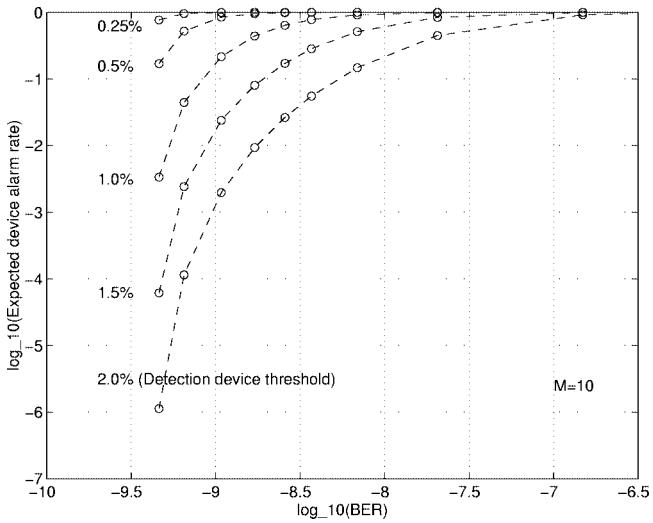

Fig. 8. Expected wrapper alarm rate versus $BER_{\text{end-to-end}}$ for different wrapper thresholds ($M = 10$, out-of-band jamming).

gains. We shall base our analysis on this set of data. We let σ_N^2 denote the noise variance (corresponding to the gain γ) when there is no attack, and $\sigma_{\tilde{N}}^2$ denote the noise variance when the EDFA gain is degraded owing to gain competition. Given that an attack attenuates the EDFA gain at each of the M network nodes by a%, we have the following expression for $BER_{\text{end-to-end}}$,

$$BER_{\text{end-to-end}} = \frac{1}{2}(1 - F_d(\sigma_{\tilde{N}}^2, A, 0)) + \frac{1}{2}F_d(\sigma_{\tilde{N}}^2, A, \sqrt{(1-0.01a)^M 160}\sigma_N). \tag{15}$$

As a reminder, A is the optimal decision threshold at the receiver, which is approximately $44\sigma_N^2$. Figures 8 and 9 show the curves of the expected wrapper alarm rate versus $BER_{\text{end-to-end}}$ for different values of detection wrapper thresholds. Tables 4 and 5 present the corresponding expected numbers of wrapper alarms in an observation period with and without an attack. Note that we consider the case when $M = 5$ instead of $M = 1$ since the maximal gain degradation in Table 3 does not yield a degraded level of BER above 10^{-8} in the case of $M = 1$. We can construct the detection scheme based on the procedures given in example 1. Our goal is to have both FP_{attack} and FN_{attack} approximately no greater than 10^{-10}. We end this section with two examples.

Example 3. Assume that $M = 10$, the detection wrapper threshold is 0.5%, and the observation period is 0.1 μs. In this case, we find β to be 44. The corresponding FP_{attack} is $< 10^{-10}$, and the corresponding FN_{attack} is $< 10^{-16}$.

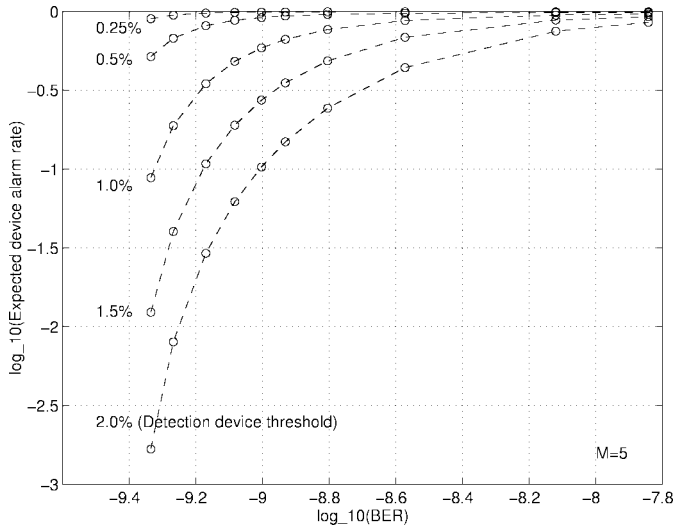

Fig. 9. Expected wrapper alarm rate versus $BER_{\text{end-to-end}}$ for different wrapper thresholds ($M = 5$, out-of-band jamming).

Table 4

Expected number of wrapper alarms in an observation period ($M = 10$, out-of-band jamming). All the entries have the same meaning as in Table 1

	0.25%	0.5%	1.0%	1.5%	2.0%
1 μs	770/1000	170/997	3.4/937	0.07/623	0.001/250
0.1 μs	77/100	17/99.7	0.34/93.7	0.007/62.3	0.0001/25
0.01 μs	7.7/10	1.7/9.97	0.034/9.37	0.0007/6.23	0.00001/2.5

Table 5

Expected number of wrapper alarms in an observation period ($M = 5$, out-of-band jamming). All the entries have the same meanings as in Table 1

	0.25%	0.5%	1.0%	1.5%	2.0%
1 μs	900/996	520/983	89/954	13/900	1.7/794
0.1 μs	90/99.6	52/98.3	8.9/95.4	1.3/90	0.17/79.4
0.01 μs	9/9.96	5.2/9.83	0.89/9.54	0.13/9	0.017/7.94

Therefore, we can detect a constant out-of-band jamming attack (yielding the BER above 10^{-8}) which lasts longer than 0.1 μs.

Example 4. Assume that $M = 5$, the detection wrapper threshold is 1.0%, and the observation period is 0.1 μs. In this case, we find β to be 31. The corresponding FP_{attack} is $< 10^{-10}$, and the corresponding FN_{attack} is $< 10^{-16}$. Therefore, we can detect a constant out-of-band jamming attack (yielding the BER above 10^{-8}) which lasts longer than 0.1 μs.

Examples 3 and 4 show, for out-of-band jamming, that the required observation time for our monitoring system is *8 orders of magnitude smaller* than the time required by a BERT.

6. Conclusions and directions for further research

We have presented some vulnerabilities of TONs to denial of service attacks. These attacks can occur through in-band crosstalk at switching nodes and through out-of-band crosstalk at amplifier nodes. While these attacks may lead to degradation of service rather than outright communications outage, they are difficult to detect rapidly and effectively. An overview of traditional fault detection mechanisms shows that they are ill suited for detecting the types of attacks we consider. We have presented a new network monitoring system, based on monitoring wrappers, for denial of service attack detection at TONs. We have analyzed the behavior of our wrappers and considered the performance of our network monitoring system for detecting denial of service attacks at switching nodes and amplifiers. The performance of our scheme yields several orders of magnitude improvement in speed of detection with respect to BERTs.

While we have considered NRZ OOK, our scheme can be applied to any type of modulation scheme and is thus applicable to TONs. The application of our scheme to other modulation schemes, such as return-to-zero (RZ) OOK, may yield different performance. Obtaining results for common types of signaling would provide bounds on the applicability of our scheme to TONs supporting different types of modulation. Furthermore, while we have considered FP and FN probabilities, we have not discussed the interaction of our attack detection scheme with the NMS. A distributed algorithm which uses attack alarms for the localization of the first attacked node in a spreading attack in an AON is given in [6]. The alarms there are considered to have virtually 0 FP and FN, which is reasonable given the analysis we have presented here. The effect of different FP and FN values on the reaction to alarms is an interesting issue, as is the use of alarms for long-term diagnostic purposes, such as post-mortem analysis of attack incidents. Finally, a natural extension of this work and an important question is the issue of what policies may be adopted by an attacker to thwart our detection scheme. A possible framing of the problem is: how can the jammer reduce the BER to some level while generating the fewest possible alarms in our detection wrapper? In the analysis we have presented here, the policy of the jammer was constant during an attack, but such a policy may not be the most advantageous for the attacker. Results in this area show that a sporadic jammer, distributed over all possible TONs, is the most difficult to detect for a given QoS degradation in a WDM network.

References

[1] S.B. Alexander, R.S. Bondurant, D. Byrne, V.W.S. Chan, S.G. Finn, R.G. Gallager, B.S. Glance, H.A. Haus, P. Humblet, R. Jain, I. Kaminow, M. Karol, R. S. Kennedy, A. Kirby, H.Q. Le, A.A.M. Saleh, B.A. Schofield, J,H, Shapiro, N.K. Shankaranarayanan, R.E. Thomas, R.C. Williamson and R.M. Wilson, A precompetitive consortium on wideband optical networks, *Journal of Lightwave Technology* **11**(May/June) (1993), 714–735.

[2] R.C. Alferness, J.E. Berthold, D. Pompey and R. Tkach, MONET: New Jersey demonstration network results, in: *Optical Fiber Communication Conference 97*, WI1.

[3] *Guide to OTDR Measurements*, Anritsu Wiltron, 1995.

[4] A.A. Al-Orainy, Analysis of crosstalk in WDM-ring networks, *IEEE Photonics Technology Letters* **5**(12) (1993), 1445–1447.

[5] N. Araki, Y. Enomoto and N. Tomita, Improvement of fault identification performance using neural networks in passive double star optical networks, in: *OFC 98*, WM38, 1998.

[6] R. Bergman, M. Médard and S. Chan, Distributed algorithms for attack localization in all-optical networks, in: *1998 Network and Distributed System Security Symposium*, sponsored by the Internet Society, session 3, paper 2.

[7] S. Betti, E. Bravi and M. Giaconi, Analysis of distortion effects in subcarrier-multiplexed (SCM) externally modulated systems: a generalized approach, *IEEE Photonics Technology Letters* **9**(1) (1997), 118–120.

[8] M. Bischoff, M.N. Huber, O. Jahreis and F. Derr, Operation and maintenance for an all-optical transport network, *IEEE Communications Magazine*, November, 1996, 136–142.

[9] D.J. Blumenthal, P. Granestrand and L. Thylen, BER floors due to heterodyne coherent crosstalk in space photonic switches for WDM networks, *IEEE Photonics Technology Letters* **8**(2) (1996), 184–286.

[10] A. Bononi. L. Tancevski and L.A. Rusch, Fast dynamics and power swings in doped-fiber amplifiers fed by highly variable multimedia traffic, in: *OFC 98*, paper WM31, 1998.

[11] J.H. Bowen, D.L. Baldwin and P.R. Couch, Secure fiber optic data transmission system, United States Patent No. 4,435,850, March, 1984.

[12] C.A. Brackett, Dense wavelength division mulriplexing networks: principles and applications, *IEEE Journal on Selected Areas in Communications* **8**(6) (1990).

[13] C.A. Brackett, A.S. Acampora, J. Sweitzer, G. Tangonan, M.T. Smith, W. Lennon, K.C. Wang and R.H. Hobbs, A scalable multiwavelentgh multihop optical network: a proposal for research on all-optical networks, *Journal of Lightwave Technology* **11**(May/June) (1993), 736–753.

[14] L.A. Buckman, L.P. Chen and K.Y. Lau, Crosstalk penalty in all-optical distributed switching networks, *IEEE Photonics Technology Letters* **9**(2) (1997), 250–252.

[15] C.-K. Chan, F. Tong, L.-K. Chen, J. Song and D. Lam, A practical passive surveillance scheme for optically amplified aassive branched optical networks, *IEEE Photonics Technology Letters* **9**(4) (1997), 526–528.

[16] C.K. Chan, F. Tong, L.K. Chen, J. Song and D. Lam, A passive surveillance scheme for passive branched optical networks, in: *OFC 97*, TuK1, 1997.

[17] G.-K. Chang, G. Ellinas, J.K. Gamelin, M.Z. Iqbal and C.A. Brackett, Multiwavelentgh reconfigurable WDM/ATM/SONET network testbed, *Journal of Lightwave Technology* **14**(6) (1996), 1320–1340.

[18] Y.-K. Chen and S. Chi, Fault-Locating and Supervisory Technique for multistaged branched optical networks, *IEEE Photonics Technology Letters* **6**(7) (1994).

[19] M.W. Chbat, Recent progress in the optical Pan-European network preoject (ACTS/RACE), in: *Optical Fiber Communication Conference 98*, ThP1.

[20] S.R. Chinn, Simplified modeling of transients in gain-clamped erbium-doped fiber amplifiers, *Journal of Lightwave Technology* **16**(6) (1998).

[21] D. Derickson, editor, *Fiber Optic Test and Measurement*, Hewlett–Packard Professional Books, 1998.

[22] E. Desurvire, *Erbium-Doped Fiber Amplifiers*, John Wiley & Sons, New York, 1994.

[23] H. Edwards, Real world experience in the deployment and operation of NTON consortium testbed, in: *Optical Fiber Communication Conference 97*, WI2.

[24] R. Erkander, Optical fibre security system ZAT 4, *Ericsson Review* **67**(1) (1986), 35–41.

[25] S.G. Finn and R.A. Barry, Optical services in future broadband networks, *IEEE Network* **10**(6) (1996), 7–13.

[26] R. Gaudino, M. Shell, M. Len, G. Desa, C. Juckett and D.J. Blumenthal, Experimental demonstration of MOSAIC: a multiwavelength optical subcarrier multiplexed controlled network, in: *OFC 98*, paper ThP4.

[27] A.H. Gnauck et al., One Terabit/s Transmission Experiment, in: *OFC 96*, Post deadline paper PD 20, 1996.

[28] E.L. Goldstein and L. Eskildsen, Scaling limitations in transparent optical networks due to low-level crosstalk, *IEEE Photonics Technology Letters* **7**(1) (1995), 93–94.

[29] E.L. Goldstein, L. Eskildsen and A.F. Elrefaie, Performance implications of component crosstalk in transparent lightwave networks, *IEEE Photonics Technology Letters* **6**(5) (1994), 657–660.

[30] E.L. Goldstein, L. Eskildsen, C. Lin and Y. Silberbberg, Polarization statistics of crosstalk-induced noise in transparent lightwave nertworks, *IEEE Photonics Technology Letters* **7**(11) (1995), 1345–1347.

[31] F. Heismann and R. C. Alferness, Wavelength-tunable electrooptic polarization conversion in birefringent waveguides, *IEEE Jour. Quantum Electron.* **24** (1988), 83–93.

[32] G.R. Hill, P.J. Chidgey, F. Kaufold, T. Lynch, O. Sahlen, M. Gustavsson, M. Janson, B. Lagerstron, G. Grasso, F. Meli, S. Johansson, J. Ingers, L. Fernandez, S. Rotolo, A. Antonielli, S. Tebaldini, E. Vezzoni, R. Caddedu, N. Caponio, F. Testa, A. Scavennec, M.J. O'Mahony, J. Zhou, A. Yu, W. Sohler, U. Rust and H. Herrmann, A transport network layer based on optical network elements, *Journal of Lightwave Technology* **11**(5/6) (1993), 667–679.

[33] P.M. Hill, R. Olshansky and M. Abdollahian, Novel carrier and clock-recovery circuit for multigigabit/second lightwave systems, *IEEE Photonics Technology Letters* **5**(1) (1993), 96–98.

[34] K.-P. Ho and J.M. Kahn, Methods for crosstalk measurement and reduction in dense WDM systems, *Journal of Lightwave Technology* **14**(6) (1996), 1127–1135.

[35] P.P. Iannone, K.C. Reichmann and N.J. Frigo, Broadcast digital video delivered over WDM passive optical networks, *IEEE Photonic Technology Letters* **8**(7) (1996), 930–932.

[36] Y. Inoue, H. Takahashi, S. Ando, T. Sawada, A. Himeno and M. Kawachi, Elimination of polarization sensitivity in silica-based wavelength division multiplexer using a polyimide half waveplate, *Jour. Lightwave Technol.* **15** (1997), 1947–1957.

[37] Y.D. Jin, Q. Jiang and M. Kavehrad, Performance degradation due to crosstalk in multiwavelength optical networks usnig dynamic wavelength routing, *IEEE Photonics Technology Letters* **7**(10) (1995), 1210–1212.

[38] *Journal of Lightwave Technology*, Special Issue on MEMS, **17**(1) (1999).

[39] I. Katzela and M. Schwartz, Schemes for fault identification in communication networks, *IEEE/ACM Transactions on Networking* **3**(6) (1995), 753–764.

[40] I. Katzela, G. Ellinas and T.E. Stern, Fault diagnosis in the linear lightwave network, in: *LEOS Summer Topical Meetings*, 1995, pp. 41–42.

[41] A. Kloch, B. Mikkelsen and K.E. Stubkjaer, Pilot tones in WDM networks with wavelength converters, in: *OFC*, TuE6, 1997.

[42] Y.W. Lai, Y.K. Chen and W.I. Way, Novel supervisory technique using wavelength-division-multiplexed OTDR in EDFA repeatered transmission systems, *IEEE Photonics Technology Letters* **6**(3) (1994), 446–451.

[43] C.-S. Li and F. Tong, Crosstalk and interference penalty in all-optical networks using static wavelength routers, *Journal of Lightwave Technology* **14**(6) (1996), 1120–1126.

[44] L.Y. Lin, Micromachined free-space matric switches with submillisecond switching time for large-scale optical crossconnect, in: *OFC 98*, WH5, 1998.

[45] L.Y. Lin, E.L. Goldstein, J.M. Simmons and R.W. Tkach, High-density connection-symmetric free-space micromachined polygon optical crossconnects with low loss for WDM networks, in: *OFC 98*, Post-deadline paper PD-24-1, 1998.

[46] C.-L. Lu , D.J.M. Sabido IX, Perluigi Pogglioni, R.T. Hofmeister and L.G. Kazovsky, CORD – A WDMA optical network: subcarrier-based signaling and control scheme, *IEEE Photonics Technology Letters* **7**(5) (1995), 555–557.

[47] D. Marquis, M. Médard, R.A. Barry and S.G. Finn, Physical security considerations in all-optical networks, *SPIE* **3228** (1997), 260–271.

[48] T. Maekawa, Y. Suzuki, K. Kumozaki and R. Watanabe, Ultrahigh-splitting-ratio optical subscriber system for small-capacity services, in: *OFC 95*, TuK4.

[49] M. Médard, Secure optical communications, Invited Paper, FE3, *LEOS '98*, pp. 323–324.

[50] M. Médard, D. Marquis, R.A. Barry and S.G. Finn, Security issues in all-optical networks, *IEEE Network* **11**(3) (1997), 42–48.

[51] M. Médard, D. Marquis and S.R. Chinn, Attack detection methods for all-optical networks, in: *1998 Network and Distributed System Security Symposium*, sponsored by the Internet Society, session 3, paper 1.

[52] M. Médard, S.R. Chinn and P. Saengudomlert, Attack detection in all-optical networks, in: *OFC 98*, ThD4, 1998.

[53] M.J. Minardi and M.A. Ingram, Adaptive crosstalk cancellation and laser frequency drift compensation in dense WDM networks, *Journal of Lightwave Technology* **13**(8) (1995), pp. 1624–1635.

[54] T. Niemeier and R. Ulrich, Quadrature outputs from fiber interferometer with 4×4 coupler, *Opt. Lett.* **11** (1986), 677–679.

[55] H. Onaka et al., 1.1 Tb/s WDM transmission over a 150 km 1.3 μm zero-dispersion single-mode fiber, in: *OFC 96*, Post deadline paper PD 19, 1996.

[56] Y. Peng and J.A. Reggia, A probabilistic causal model for diagnostic problem solving – Part I: Integrating symbolic causal interference with numeric probabilistic inference, *IEEE Transactions on Systems, Man, and Cybernetics* **SMC-17**(2) (1987), 146–162.

[57] E.C.M. Pennings, R.J. Deri, R. Bhat, T.R. Hayes and N.C. Andreakis, Ultracompact, all-passive optical 90-hybrid on InP using self-imaging, *IEEE Photon. Technol. Lett.* **5** (1993), 701–703.

[58] C. Saxtoft and P. Chidgey, Error rate degradation due to switch crosstalk in large modular switched optical networks, *IEEE Photonics Technology Letters* **5**(7) (1993), 828–831.

[59] N. Schroff and M. Schwartz, *Fault Detection/Identification in the Linear Lightwave Network*, CU/CTR/TR 243-91-24, Columbia University, 1991.

[60] Y. Shani, C.H. Henry, R.C. Kistler, R.F. Kazarinov and K.J. Orlowsky, Integrated optic adiabatic devices on silicon, *IEEE Jour. Quantum Electron.* **27** (1991), 556–566.

[61] Y. Shen, K. Lu and W. Gu, Coherent and incoherent crosstalk in WDM optical networks, *IEEE Journal of LIghtwave Technology* **17**(5) (1999), 759–764.

[62] M.H. Slonecker, Method and apparatus for securing information communicated through optical fibers, United States Patent No. 4,973,169, November, 1990.

[63] R.A. Soref, Secure optical matrix switch, reprt RADC-TR-79-51, April 1979, Rome Air Development Center.

[64] L.F. Stokes and D. Derickson, Lightwave component ans system measurements, Short Course Notes, in: *OFC 97*.

[65] M. Sumida, OTDR performance enhancement using a quaternary FSK modulated probe and coherent detection, *IEEE Photonics Technology Letters* **7**(3) (1995), 336–338.

[66] K. Takada, H. Yamada and Y. Inoue, Origin of channel crosstalk in 100 GHz-spaced silica-based arrayed-waveguide grating multiplexer, *Electronics Letters* **31**(14) (1995), 1176–1177.

[67] H. Takahashi, K. Oda and H. Toba, Impact of crosstalk in an arrayed-waveguide multuplexer on NxN interconnection, *Journal of Lightwave Technology* **14**(6) (1996), 1097–1105.

[68] H. Tsushima, M. Shabeer, P. Barnsley and D. Pitcher, Demonstration of an optical packet add/drop with wavelength-coded header, *IEEE Photonics Technology Letters* **7**(2) (1995).

[69] R.E. Wagner, R.C. Alferness, A.A.M. Saleh and M.S. Goodman, MONET: multiwavelentgh optical networking, *Journal of Lightwave Technology* **14**(6) (1996), 1349–1355.

[70] B.H. Wang, K.Y. Yen and W.I Way, Demonstration of gigabit WDMA systems using parallel processed subcarrier pilot-tone signaling technique, in: *OFC '96*, paper TuE1.

[71] W.I. Way, Y.W. Lai and Y.K. Chen, The effect of transient gain compression in a saturated EDFA on optical time domain reflectrometry testing, *IEEE Photonics Technology Letters* **6**(10) (1994), 1200–1202.

[72] A.V. Yakovlev, An optical-fiber system for transmitting confidential information, *Telecommunications and Radio Engineering* **49**(4) (1995), 1–6.

[73] F. Yamamoto, I. Sankawa, S. Furukawa, Y. Koyamada and N. Takato, In-service remote access and measurement methods for passive double star networks, in: *Conference on Optical Hybrid Access Networks*, pp. 5.02.01–06.

[74] C.X. Yu, W.K. Wang and S.D. Brorson, System degradation due to coherent cross talk in WDM network nodes, in: *OFC 98*, WM30, 1998.

[75] R.O. Yudkin, On testing communication networks, *IEEE Journal on Selected Areas in Communications* **6**(5) (1988), 805–812.

[76] J. Zhou, M.J. O'Mahony and S.D. Walker, Analysis of optical crosstalk effects in multi-wavelength switched networks, *IEEE Photonics Technology Letters* **6**(2) (1994), 302–307.