A review on Fault Location methods and their application to optical networks

Carmen Mas[*] and Patrick Thiran[†]

Abstract

Fault identification and location in optical networks must cope with a multitude of factors: (i) the redundancy and the lack of coordination (internetworking) of the managements at the different layers (WDM, SDH/SONET, ATM, IP); (ii) the large number of alarms a single failure can trigger; (iii) the difficulty in detecting some failures and the resulting need to cope with missing or false alarms.

This paper reviews some approaches used for fault location. We classify them in two main categories: the Model Based approaches and the Black box, Learning-based approaches. The first ones rely on an abstract model of the network, capturing the dependency relations between the elements and capable of pointing out the element(s) which is (are) most likely to be the cause of the received alarms. The second ones do not attempt to model the network in detail, but leave it as a black box. An abnormal situation is then diagnosed from a set of rules obtained by learning or thanks to expertise of the human manager.

1 Introduction

Today's optical WDM networks are able to transmit up to 164 wavelengths through a single optical fiber, each wavelength transporting 10 Gbps (OC-192/STM-64). Advanced cabling technologies group up to 432 fibers on a single ribbon [DDH+99]. As a result, a ribbon break yields the interruption of hundreds of thousands of flows, and the loss of thousands of megabits of data [Met00]. Survivability of optical networks as well as fault identification and location have thus become a

[*]NTUA, Electrical Engineering and Computer Science Dpt., Heroon Polytechniou 9, 15773 Zographou, Athens, Greece, cmas@ieee.org

[†]Sprint ATL, 1 Adrian Court, Burlingame, CA 94010, USA; and ICA-DSC, EPFL, CH-1015 Lausanne, Switzerland, Patrick.Thiran@epfl.ch

crucial problem.

When a failure occurs at the physical layer, the lightpaths that are affected or interrupted have to be restored as soon as possible so that higher layers do not see the the failure and do not start their own restoration mechanisms. In the meantime, the failure has to be located and repaired. Failures are located from the alarms received by the management system. When there are two or more simultaneous failures, the number of alarms considerably increases, the alarms arrive intermingled to the management system and the problem of locating the failures becomes even more difficult. The problem of locating multiple failures has been shown to be NP-hard by Rao [Rao93]. The location of the failure(s) must be fast and efficient (by efficient, we mean that the set of faulty candidates must be as small as possible) so that the faulty element(s) is (are) clearly identified before start digging or sending a technician to repair it(them). Such a task is expensive, yet it is not rare. For example, submarine cables, which are vulnerable to damage from submarines, anchors and fishing gears, have to be repaired once every five weeks [Sub99], whereas under-earth cables are cut even more frequently, because of agricultural engines, natural disasters animals or human interventions in construction and maintenance areas.

Network management is essential to ensure the good functioning of these networks. Every network management performs several functions, which have been classified into five different functional areas by OSI, and which are briefly recalled here.

- **Configuration Management** deals with initializing the network components, establishing relationships among them, maintaining, adding and updating these relationships, and keeping the manager informed about the status of the components. These relationships between components are based on the connections established and cleared down in the network. Configuration management must also reconfigure the network whenever necessary and include routines able to inform about any change at the configuration (for example, when a protection switch changes its position, the manager should be informed about the new path of the established channels).

- **Performance Management.** Data communication networks are made of a large variety of components. Performance management has to monitor and control them. *Monitoring* is the function that tracks activities in the network, whereas the *control* function enables adjustments in the components to improve network performance. The main performance issues are: network capacity utilization, existence of excessive traffic and of bottlenecks, and increase of response time. Performance management collects information from the network and analyzes it so that the management system can recognize situations of performance degradation.

- **Security Management** deals with the generation, distribution and storage of encryption keys. It also monitors and controls access to computer networks and to management information.

- **Accounting Management.** Many network services are charged to the users. Network management performs not only the internal accounting but also other tasks, such as checking allowed access privileges.

- **Fault Management** deals with [HHN98]:

  - fault detection, to know whether there is a failure or not in the network,
  - fault location, to know which is/are the component(s) that has/have failed and caused the received alarms,
  - fault isolation (so-called protection) so that the network can continue to operate, which is the fast and automated way to restore interrupted connections. In general, it is implemented with protection switches that change position when the optical power gets below a certain threshold. Recent work on fault isolation can be found in [Gop00].
  - network (re-)configuration (so-called restoration) that minimizes the impact of a fault by restoring the interrupted connections using spare equipment. This involves some processing to discover the best path to re-route the connections.
  - replacement of the failing component(s).

Several articles deal with protection and restoration mechanisms for optical networks [Dea99, Sea97, VJH97, DDH^{+}99]. This article focuses on the fault location problem. This problem is not specific to optical networks, but is encountered in many fields, such as electrical power plants monitoring, medical diagnosis, electric circuit analysis, nuclear power station maintenance and management of communication networks at large. Here we consider only its application to optical networks, although many methods are actually valid or have been developed in more general settings. Fault location is particularly important for optical networks, where the quantity of information carried within an optical network is much larger than within other networks, making the effects of a failure much more severe.

This paper is organized as follows: Section 2 describes the problem of fault location and the features that distinguish the existing approaches. We begin by describing them in a more general context than that of optical networks, together with their advantages and limitations, in Section 3. Section 4 focuses on the fault location problem in optical networks by providing the information that can be extracted from these networks, a possible classification of their failures and an example of a failure scenario, which is used in Section 5 to compare five different methods able to solve the fault location problem in optical networks. The article concludes with a summary of the most important features of the different methods.

2 Fault Location Problem definition

Optical communication networks, and all networks in general, need a fault management system able to perform fault diagnosis, that is, able to identify the faults that occur from the information given by the network components. A *fault* can be defined as an unpermitted deviation of at least one characteristic parameter or variable of a network element from acceptable/usual/standard values, whereas a *failure* can be defined as the manifestation of the fault. For example, if the ventilator in a laser stops, and if the temperature increases and overpasses an accepted temperature limit, the fault is the stopped ventilator and the failure is the temperature of the laser, which is too high. Because both terms are closely related, we use them indistinctly.

Fault detection relies on the monitoring of the state of the network components. Simple fault detection mechanisms are often based on locally monitored variables. The faulty values reached by these variables are logged as errors. Critical errors are sent to the network manager as alarms. However, it is not always possible to detect complex faults on the sole basis of locally monitored variables: it is then necessary to have a global knowledge of the network and to do some processing to diagnose the presence (or absence), the nature and the location of the fault. Also, and because the fault can propagate to components that depend on the failed component, the influences of faulty components on other components have to be taken in account to perform an efficient fault management.

Communication networks are built on several layers, performing each fault management functions independently. When a failure occurs, several symptoms or event indications are issued to the network manager from the different management layers, and fault management functions start in parallel. Research is carried out to allow interoperability between different layers, to avoid task duplication and to increase efficiency.

The fault location problem is solved by Fault Management Systems that take the events generated by the network elements as input (these events can be alarms, warnings or parameters of the network element itself), and produce an output, which is the set of network elements whose failure explains the input events (see Figure 1).

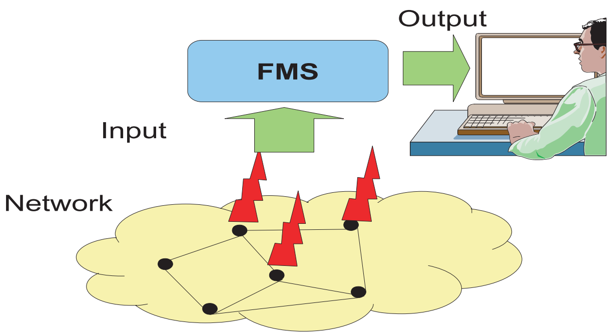

Figure 1: Fault Management System (FMS)

These fault management systems differ in:

1. the way they solve the problem: using neural networks [GH97, RRM$^+$96, ML97], Finite State Machines [LR95, WS92, BHS93], the history of previous cases [HE93], a (detailed) modeling of the network [MTB99], etc.

2. the information they need: failure propagation probabilities [DLW93], timestamps, set of established channels [MTB99, MT00], etc.

3. the assumptions on which they rely: existence of only single failures [Rao96], existence of multiple failures but assuming than x failures are more likely than $x + 1$ failures [MTB99, MT00, KW87, YS93], etc.

4. the quantity of memory they use: large memory requirements [HE93] to store failure history, or more modest memory requirements [LR95, BHS93, GH97] to store for example an FSM or neural network parameters.

5. intolerance [HE93, KS95] or tolerance [GH97, MTB99, MT00] to false and missing alarms.

3 Approaches to the fault location problem

All techniques performing fault diagnosis rely on the analysis of symptoms and events that are generated during the occurrence of the fault. We divide them in two main categories. The first ones are *Model based approaches*, which construct first an accurate and workable model of the network on the basis of the functional and physical properties of the network components, and perform diagnosis by comparing actual observations with forecasts from the model. The second ones are the *Black box, learning-based approaches*, which view the system as a black box delivering output data when a particular failure occurs. They "learn" the relationship between input events and output diagnosis, which can be done in different ways: by capturing and implementing in an efficient way the human expert knowledge (expert system), by recording the history of previous cases that occurred in the past (case-based systems), or by artificial neural networks or any other algorithm with statistical learning capabilities.

3.1 Model Based Systems

Model Based approaches to fault diagnosis are more appropriate than other approaches because they can cope with faults that did not occur before. These methods construct an abstract model of the network that has to be diagnosed. The model can be of any kind, from logic to differential equations, and represents the failure inter-dependency between the network elements. Depending on the kind of model, different approaches can be used, for example statistical approaches, artificial intelligence based approaches, or approaches using control theory [Nyb99]. The second step is the elaboration of an algorithm to search and deliver the elements that best explain the discrepancies between the observations and the predictions from the model (see Figure 2).

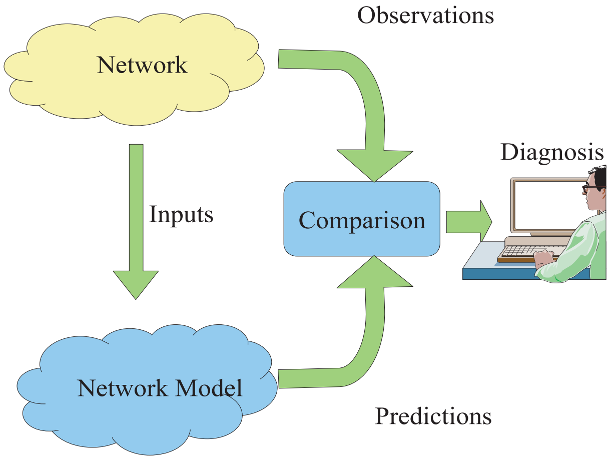

Figure 2: Model Based System

The advantages of these systems are that:

- they are able to cope with incomplete information,
- they do not require learning,
- they can cope with unforeseen failures

The drawback is the difficulty of developing good model for complex networks.

Some works belonging to this approach are:

1(a) probabilistic reasoning systems developed by Katzela and Schwartz [KS95] and Wang and Schwartz [WS93]. The input parameters of these systems are failure probabilities, which may be difficult to be determine for optical networks because of their time dependence. These systems also take as input the conditional probabilities of failure of an element given the failure of another one.

1(b) Finite State Machine (FSM) model proposed by Wang and Schwartz [WS92], Li and Ramaswami [LR95] and Bouloutas et al. [BHS93]. This approach models the network, and its behaviour when a failure occurs, as a Finite State Machine. It consists in a set of states, with transitions between states dictated by input events such as alarms coming from the network. Each of the states define either the correct functioning of the network or a failure scenario. Therefore, only the failure scenarios considered in the FSM can be identified. The FSM can become very complex and difficult to be designed.

1(c) Deterministic modeling algorithms proposed by Mas, Thiran and Le Boudec [MTB99, MT00]. These systems construct a deterministic model of the network, and efficient algorithms tailored to this model. The advantages are the small amount of information needed (no timestamps, nor failure probabilities), and the tolerance to false and missing alarms. The drawback is that the algorithms require a fairly large amount of pre-computation and storage off-line. Another disadvantage is that the alarms that will be sent for each failure scenario has to be known.

3.2 Black box, Learning-based approaches

This section presents different approaches to solve the fault location problem that do not model the network but consider it as a "black box" that outputs certain data when a failure occur.

2(a) Expert Systems use the knowledge of human experts acquired during their past experience with the network. They work in two phases (see Figure 3):

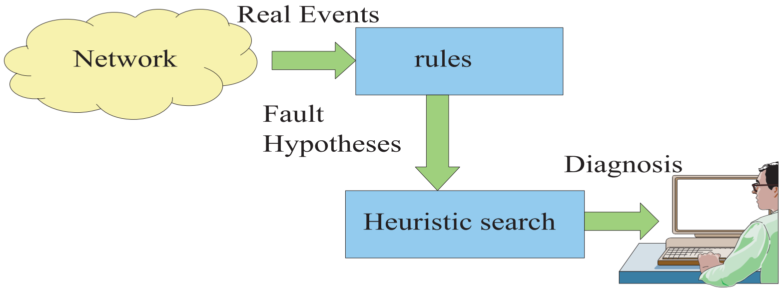

Figure 3: Expert System Approach

- The first phase is the *generation of rules*, based on the expert's knowledge. These rules associate a fault hypothesis with (network/alarm) events, and have the form:

  if `event` x **then** `fault hypothesis` y

- The second phase is the *operational or diagnosis phase*: whenever an event is prompted, the expert system will extract a fault hypothesis from the rules. When a large number of events are prompted, they will often result in different fault hypothesis. A heuristic search is performed to recover the most likely one.

Expert systems have the advantage of effectively taking human expertise into account, and coping with large scale systems. The operational phase is also rapid.

The drawbacks of expert systems are: the relationships between failure symptoms and the fault have to be completely pre-determined; the required translation of the human expertise into a set of rules, which cover all cases in an exhaustive manner is hard; the need to exhaustively search for all the possible fault hypothesis slows down the global functioning of the system; and, when the network components or connections change, the expert system has to be changed dramatically. Defining expert systems is a very difficult task, which may result in a large number of rules and therefore require a huge database.

Applications of expert systems to fault diagnosis have been proposed since 1985 by Wollenberg [Wol86]. Since then, many proposals using expert systems have presented as for example by Gu et al. [Gea93], by Brugnoni et al. [Bea93] and by Jakobson et al. [GJW94, JWB$^+$00].

2(b) Case Based Systems are on-line learning systems that use earlier experiences as the basis for making decisions. To solve a particular problem, a case based system remembers the occurrence of a similar problem that was solved in the past, and adopts the old solution to solve the newly submitted problem. In order to retrieve cases efficiently, it is crucial to decide what the attributes of a case are, and on which attributes the cases are indexed. The system can be divided into *four* steps: *retrieve*, which finds the best match of a previous case, *reuse*, which finds what can be reused from old cases, *revise*, which checks if the proposed solution is correct and *retain*, which learns from the problem solving experience (see Figure 4).

The advantages of these systems are their efficiency and speed when the submitted problem is a previously submitted and solved case, the on-line learning that allows storing newly solved cases and their implementation simplicity.

The main disadvantages are that they need a huge memory to store all the past cases and that they cannot cope with new problems that were not previously encountered.

Lewis [HE93] proposes a case base system called `CRITTER` to solve fault diagnosis with its

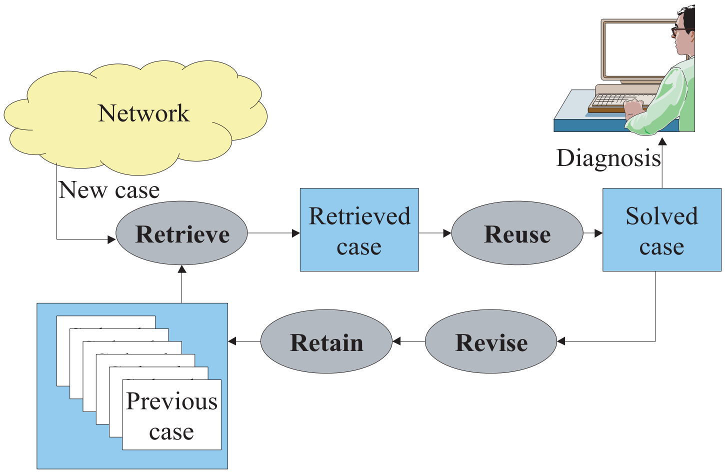

Figure 4: Case Based Systems Approach

techniques to retrieve, adapt and embed knowledge in the case library.

2(c) Artificial Neural Networks (ANNs) are a set of elements, the *neurons*, interconnected by parameters, the *weights*. Systems using ANNs realize a black box modeling of the network using statistical learning. The mapping between the input and the output of the neural network is realized in two phases:

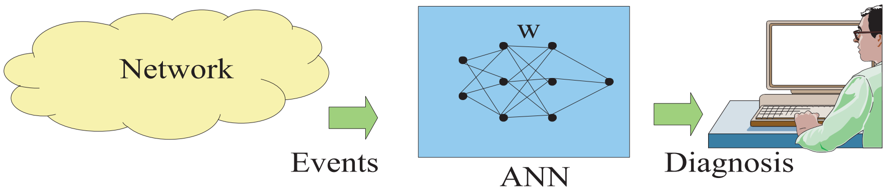

Figure 5: Artificial Neural Network Approach

- *Learning phase.* The initial values of the weights are randomly chosen. This first phase updates the weights by learning from a database containing sample pairs `(input events, desired output)`. In the case of a supervised learning algorithm, the weights are updated by some form of gradient descent of an error functional between the desired output and the actual output, as sketched in Figure 6. A unsupervised learning algorithm makes only use of the input network events. It is only at the end of the process, once the weights have reached their final values, that they are "labeled" and associated to a particular diagnosis output.

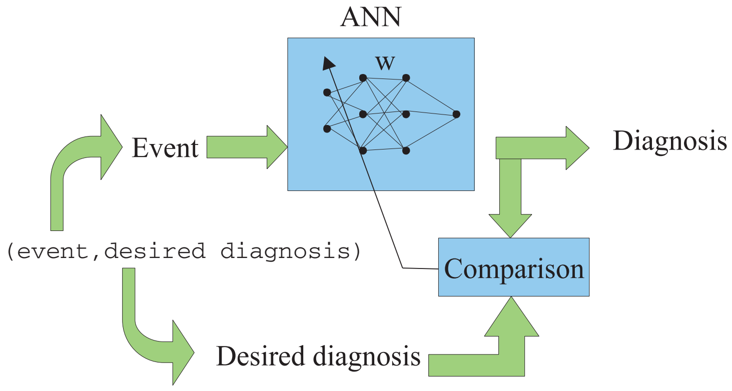

Figure 6: Learning phase of ANNs where w stands for *weight*

- *Operational phase.* When training is completed, the ANN should issue the correct diagnosis whenever an event is presented, provided it generalizes well from the samples in the training database to new samples not contained in the database.

The advantages of using ANNs are as follows:

- no explicit modeling is needed,
- learning is not only achieved from past experiences but from any set of measurements,
- if the system is properly trained, it can cope with new cases, with noisy and partial inputs,
- once the system is in the operational phase, the processing time for a new diagnosis is very short.

The drawbacks of these systems are mainly twofold: the training can be slow and the generalization to new cases can be poor. Moreover, neural networks are only useful as a tool for statistical classification for analog and noisy time series.

Examples of neural networks applied to the fault diagnosis problem are proposed by Maki and Loparo [ML97] and Rodriguez et al [RRM$^+$96].

2(d) (Other) Proactive learning systems. Ho et al. [HCPZ00] have proposed a proactive system to detect soft failures or performance anomalies by continuously analyzing performance data of the managed network. The performance anomalies are signatures of soft and hard failures in the network. The core of the method is to study the variation of the measured data under normal conditions to set dynamic thresholds that will closely follow the trend in the data. In this way, any anomaly that causes a deviation of the measured data with respect to the threshold, will be detected. The general idea of the method, as shown in Figure 7, is to receive performance data from the network, to sample it and to compare it with the dynamic threshold set on the basis of typical values of the data of the network under regular conditions, to detect anomalies and to send alarms to the human manager through a graphical interface.

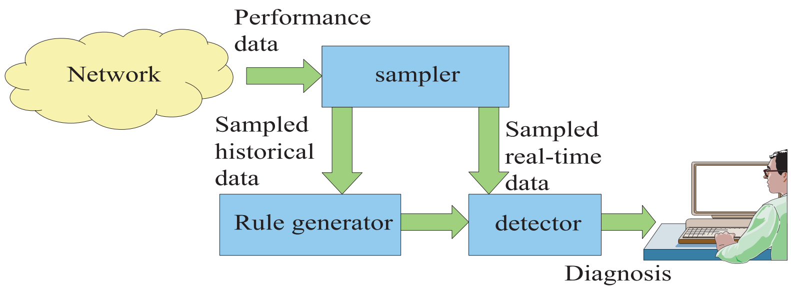

Figure 7: Proactive Learning System Approach

The advantage of this method is that it is sensitive to many network faults, and that it can cope with unforseen failures.

The difficulty is that it needs a continuous (and non trivial) study of the normal behaviour of the network to set the threshold, and that the diagnosis information may be not accurate enough (only a warning when there is a problem, without providing an exact identification of the failure.)

3.3 Hybrid approaches

Some proposals combine different approaches among the ones described above. Lin [Lin98] for example proposes a hybrid approach that combines all of them. This system is able to cope with unknown faults thanks to model based reasoning, and because the complete and correct model is difficult to be achieve, expert systems and neural networks can diagnose faults based on the typical symptom based associations. The drawbacks are that it is time consuming and that it does not avoid all disadvantages of all approaches included in the model.

Hood and Ji [HJ97] also propose a hybrid approach combining a probabilistic method that provides a Bayesian model of the network, with a proactive learning system to set parameters.

Another hybrid approach has been proposed by Gardner and Harle [GH97]. They propose to combine a special type of ANN called Kohonen networks with a rule based system.

4 Fault Diagnosis on Optical Networks

Any method used for locating failures in other systems and networks than optical networks cannot be applied directly as such to optical networks, because of their special features such as available information of the physical layer, the kind of alarms prompted by a failure, etc.

Optical networks have several layers (typically IP/ATM/SONET/ WDM or IP/SONET/WDM) which all perform fault management independently from each other. Although today's networks cannot avoid task duplication, we believe that interoperability between layers will not only eliminate it, but will also improve the performance of the fault management tasks. More precisely, the fault location problem could be more efficient if the management system was able to correlate information from the different layers. At the physical layer, the management system receives information about the physical properties of the network, such as optical power and temperature of the equipment. At the WDM layer, if adequate testing equipment is deployed, the management system obtains information about the quality of the optical signals such as Signal to Noise Ratio (SNR) and crosstalk. The manager at the upper layers receives more detailed information about the quality of the signal such as Bit Error Rate (BER), which is specific to each transmission technology. Merging all this information may help to better locate failures and even to predict failures.

4.1 Available information in an optical network

The first important factor is the available alarms and monitoring information where the fault location should rely on.

4.1.1 Physical Layer

At this layer, the available information is delivered by the micro-controllers that control the hardware equipment. Only some hardware equipment is controlled. For example, an optical fiber does not have any software taking actions on it or retrieving information from it, whereas a transmitter has a micro-controller that is able to set it up, turn it off, control its temperature and power, etc. The information available to micro-controllers depends on the manufacturers. The most common components of an optical network able to provide information and alarms (a more detailed description can be found in [MTB99]) are:

- Transmitters (Txs), which are lasers or laser arrays converting an electrical signal into an optical signal at a certain wavelength. The resolution of the laser limits the spacing between the different wavelengths of the different channels, and hence the number of channels in WDM networks. New lasers used in advanced WDM networks are tunable and can change the emission wavelength within a prescribed range [Muk97]. Transmitters send alarms when either the temperature or the incoming power is beyond a prescribed range.

- Receivers (Rxs), which convert an optical signal corresponding to a certain wavelength, into an electrical signal. They send alarms when the input optical power is under an acceptable level.

- 3Rs (Re-generator/Re-shaper/Re-timer), which are framers able to amplify the electrical signal, give the original shape of the signal and readjust the time interval between pulses. They send alarms when they cannot lock to the incoming signal.

- Protection switches, which receive more than one optical signal, and select one among them, having an acceptable power level. They send an alarm when they change the switch position due to an unacceptable incoming optical power.

- Switches, which allow cross-connection, that is, the connection of a particular input with a particular output. An alarm is sent when the connection cannot be established.

- Amplifiers, which output a signal at a higher power level than the input signal. Optical amplifiers do not perform conversion to an electrical signal before amplifying it. Most amplifiers add distortion to the signal. They send alarms for example when the pump laser does not work properly, or when the incoming power is not sufficient.

4.1.2 WDM Layer

A typical meshed DWDM network comprises several lightpaths, having each their own optical characteristics (delay, jitter, wavelength, BER,...) [GNS00]. These quality parameters are difficult to obtain because they depend on the optical components along the path, which add defects to the optical signal. Typically, the optical path traverses transmitters, optical fiber segments, amplifiers, optical crossconnects, etc. [Mas00].

ITU [12199] has given some guidelines on how to measure the signal quality in optical networks. Bit Error Rate (BER) is the most significant signal and data quality parameter, because most of the above mentioned defects degrade it. BER has even been considered to classify lightpaths to negotiate different QoS [GNS00]. Bit errors are mainly caused by noise and distortion. Noise can be added, for example, by the laser or the receiver. Distortion appears, for example, when there is crosstalk, dispersion or non-linear effects in the fiber.

Devices measuring directly the quality of the optical signal can be divided into the two following categories:

- Global Testing Equipment (GTE) measures the quality of the *overall* optical signal. This is for example the case of *Spectrum Analyzers* and *DWDM Monitors*. GTEs are equipment able to analyze the frequency characteristics of the incoming signal and display them on a graphical interface. Spectrum analyzers are the most common GTEs.

  Their basic operation is the separation of an optical signal over several wavelengths into multiple signals over one single wavelength using filters, and the measurement of the amplitude of each of them using detectors. The information spectrum analyzers can provide is:

  - *Frequency measurement:* power at each wavelength, location of each wavelength on the frequency band, bandwidth at 3dB, the channel spacing, etc.
  - *Time measurement:* time-series of signal, signal bursts, transient response, etc.

  Spectrum analyzers can be electrical or optical. Examples of electrical spectrum analyzers are the new models MS2665C and MS2667C developed by Anritsu [Anr99] or the HP 70000 series of Hewlett-Packard [Pac99]. Examples of optical spectrum analyzers are MS9720A and MS9715A developed by Anritsu [Anr99], the Walics analyzer by Photonetics [Pho00], the HP86142 by Hewlett-Packard [Pac99] and the OSP102A by Wandel & Goltermann [Gol].

- Individual Testing Equipment (ITE) designates a device that can measure parameters relevant to a *single* channel, that is, to a single wavelength. Two examples are the network tester ANT-20 by Wandel & Goltermann [Gol], which can calculate the BER of a given channel, and the SDH/PDH/ATM Analyzer MP1552A manufactured by Anritsu [Anr99], which can evaluate PDH, SDH and ATM equipment, as well as provide performance functions by detecting management signals and monitoring the traffic. The drawback of this monitoring equipment is that they depend on the transmission technology (ATM, SONET, SDH, etc.). This detailed knowledge of the upper layer should be avoided since one of the advantages of WDM networks is their transparency with respect to the transmission technologies deployed.

4.1.3 SDH/SONET Layer

Synchronous Digital Hierarchy (SDH) and Synchronous Optical NETwork (SONET) are a set of network interface standards and multiplexing schemes developed to support the adoption of optical fiber as a transmission medium. SDH is the European standard and SONET is the US counterpart. For simplicity we refer only to SDH, but the concepts apply to SONET as well.

The most important failure notifications handled by SDH are: loss of signal (LOS), loss of frame (LOF) and loss of pointer (LOP).

When a failure occurs at the optical physical layer, three phases can be identified [WOSF99]:

- At the optical layer, a hard failure is detected by, for example, sensing a loss of optical power. Maintenance information upward and downward the lightpath has been specified by ITU-T G.872 [G.898].

- At the Regenerator Section layer, a notification is received from the optical layer informing about the failure.

- Indications are then sent upstream and downstream along the optical path, informing the other SDH layers about the failure.

4.2 Types of failures

Fault diagnosis in optical networks has to cope with two kinds of failures: *hard* and *soft* failures.

- Hard failures are unexpected events that suddenly interrupt the established channels. An example of a hard failure is a fiber cut, which interrupts all connections going through this fiber.

  Hard failures can be detected with the information at the optical layer, i.e. binary alarms sent by the hardware equipment, such as `Loss of Optical Power` or `Incoming Signal not locked`. Information from monitoring equipment such as GTE or ITE, as well as information from higher layers such as SDH equipment, may help to better locate hard failures thanks to their sensitivity to all failures.

- Soft failures are events that progressively degrade the quality of transmission. An example of a soft failure is the aging or the variation of temperature characteristics of a laser: the output wavelength will drift as the laser ages or as the laser heats up or cools down. In this case, the wavelength drift creates interference among adjacent channels.

  Soft failures can sometimes be detected at the optical layer if proper testing equipment is deployed, but often require performance monitoring at a higher layer (SDH, ATM or IP). For example, if the wavelength emitted by a laser is shifted, the hardware equipment will not notice any problem, but monitoring equipment will detect a BER increase or a degraded SNR. This information helps detecting soft failures and issuing early warnings against some hard failures.

4.3 Example of a failure scenario in an optical network

This section introduces a simple example of a hard failure scenario to illustrate the methods described earlier. Figure 8 shows a part of an IP/SONET/WDM optical network. The data format at the IP layer are packets, at the SDH/SONET layer are frames multiplexing in time several channels, and at the WDM layer are connections multiplexing in frequency several connections.

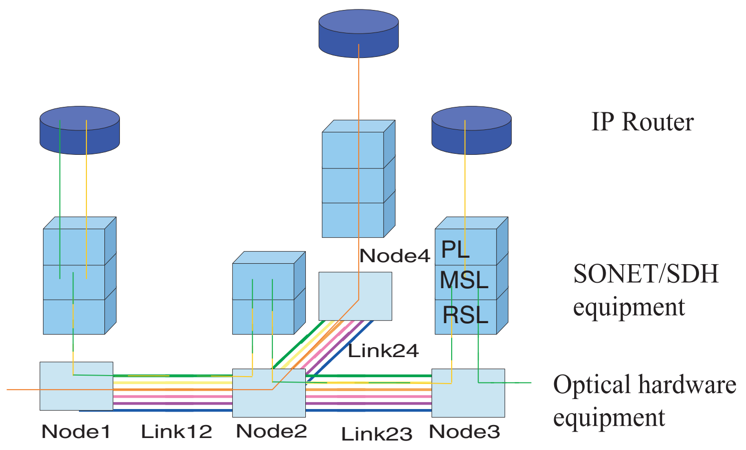

Figure 8: Optical fiber broken in an IP/SDH/WDM optical network

When there is a failure at the physical layer, alarms from several layers will be sent to their own management platforms and fault protection and restoration mechanisms will be triggered at each layer. In the example shown in Figure 8, the WDM layer could start a fault location mechanism based on the alarms generated at the physical layer (for example a `Loss of Optical Power` at the receiver of Node2). Also, the SDH/SONET layer will react by applying protection in order to restore the interrupted connections based on the SDH alarms `Loss of Frame (LOF)`, `Loss of Pointer (LOP)`, etc. issued by the SDH equipment. When SDH is unable to restore the connection (for instance if the 1+1 protection is not working because of a problem at the protection link), IP would also react to the failure by re-routing the packets through alternative paths. The failure will be detected at the IP layer when neighboring routers will no longer be able to talk to each other.

5 Comparison of the different approaches locating failures on an optical network

This section compares the different approaches, by applying them to the fault location problem introduced in Section 4.3. In particular, we will give the most important features of each approach and see how they could be applied to the particular case of an optical network. Five methods are studied in this section: three Model based systems: (1a) a probabilistic reasoning system developed by Katzela and Schwartz [KS95], (1b) an FSM system developed by Li and Ramaswami [LR95], (1c) a deterministic system developed by Mas and Thiran [MT00], and two Black box, Learning-based methods: (2a) the expert system presented by Jakobson et al. [GJW94], and (2b) the case based system proposed by Lewis [HE93].

- **Input data** needed in addition to the alarms from the network:

  (1a) Method (1a) needs (i) the network topology, to build a directed graph with the network nodes and the links that interconnect them, (ii) the failure probabilities of each node and link and (iii) the failure propagation probabilities. In our example, the directed graph representing the physical layer could look like the one shown in Figure 11(a), where each element is either a network node or a fiber, and has a failure probability p_i. Every link between two network components has a weight p_{ij}, which is the failure propagation probability and which is the conditional probability that Network Component j fails if Network Component i has already failed. These probabilities, can be estimated by network management, system specification information, failures history, etc.

  (1b) Method (1b) needs the system topology and the FSM for that specific system. In order to design the FSM, the network manager has to define the failures that have to be located, and the events that determine the status of the network, which can be either alarms informing about a problem or notifications informing about the resolution of a problem.

  The goal of the method proposed in [LR95] is to locate link failures in an all-optical network. In our network example, considering up to two possible simultaneous failures to keep the FSM simple, the FSM would look like the set of interconnected states shown in Figure 9.

  Events are the alarms `Loss of light` at Linkij informing that there is not optical power and the notifications `Light recovered` from Linkij when the optical power is re-established. The state transition depends on the expected event. For example, the FSM could pass from the "Good Functioning" state to the "Link24 fail" state when `Loss of light 24` is sent from Node4 and received by the management platform. The FSM machine becomes more complex with a larger number of failure scenarios and possible

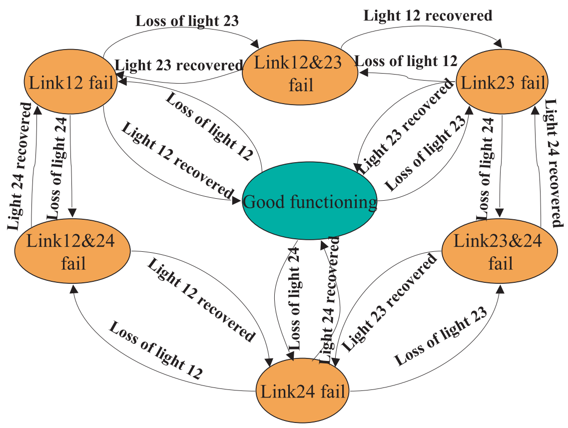

Figure 9: FSM modeling the system behaviour when link failures occur.

events.

(1c) Method (1c) does not need the network topology but needs the set of established channels. Each channel is defined as an ordered set of network elements. Consider the three following channels established in the example of Figure 8: the yellow one from Node1 to Node3 (labeled with Channel1 in Figure 10), the green one from Node1 to somewhere after Node3 (labeled with Channel2) and the red one somewhere before Node1 to Node4 (labeled with Channel3). The input to the fault location algorithm will be the set of elements, at any layer, located between the input and output ports, and belonging to each channel. Figure 10 shows how these three channels are modeled as ordered sets of network elements, representing each element by a box colored according to their alarming properties. For example, an optical fiber and a multiplexer can be put in the same class of so-called passive elements (represented by green, round boxes and labeled with p), because both can fail and none of them has a micro-controller to communicate with the manager.

(2a) Here the expertise of the network manager is the needed input, and is translated in a set of “if/then/else” rules. In our example, some rules could be:

if `Loss of light 23` **then**

  if `Loss of light 24` **then**

    if `Loss of light 12` **then** `Link12 fail`

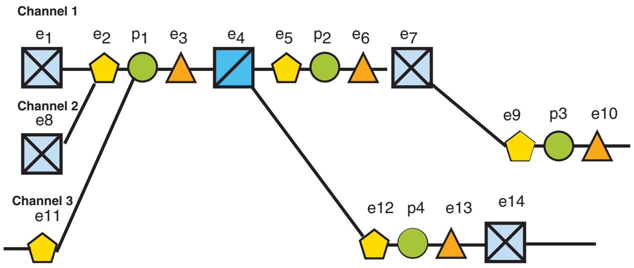

Figure 10: Channels modeled as ordered sets of network elements, classified according to their alarm and failure propagation properties.

Case	Events	Result
1	`Loss of light 23`	Link23 fail
2	`Loss of light 23 + Loss of light 24`	Node2 fail
3	`Loss of light 12 + Loss of light 23 + Loss of light 24`	Link12 fail

Table 1: Example of stored cases in a Case-based system

else `Node2 fail`

else `Link23 fail`

(2b) The input is the history of all previously solved failure scenarios: the set of alarms that are received and the failure associated to them. In the proposed example, some solved failure scenarios could be the ones shown in Table 1.

- **Pre-processing** needed to prepare the settings for an efficient fault diagnosis:

(1a) Method (1a) designs first the directed graph and then applies a *Divide and conquer* algorithm, in two phases. The first one, called *partitioning*, can be done off-line. It groups iteratively the nodes by taking the two nodes of the graph i, j (remember that these nodes represent links or nodes in the network diagram of Figure 8), for which the failure propagation probability p_{ij} is the largest, and by replacing them by a single new node k whose failure probability is $p_k = p_i + p_j * p_{ji} + p_j + p_i * p_{ij}$. The new propagation probabilities p_{kl} and p_{lk}, involving another node l and the new node k are obtained by taking the maximum of the previous propagation probabilities: $p_{kl} = \max\{p_{il}, p_{jl}\}$,

$p_{lk} = \max\{p_{li}, p_{lj}\}$.

This step is repeated until a single node is retrieved. Figure 11 shows the two first steps and the resulting hierarchical structure of the partitioning phase of the network, with failure probabilities of each node equal to 0.01 and with failure propagation probabilities given in Figure 11(a). Because Node6 and Node7 have the largest propagation probability ($p67$ =0.1), they are the first nodes to be grouped into a new one (Node with a failure probability of 0.022 and marked in green in Figure 11(b)). At the same time, the first level of the hierarchical structure displayed in Figure 11(c) is obtained.

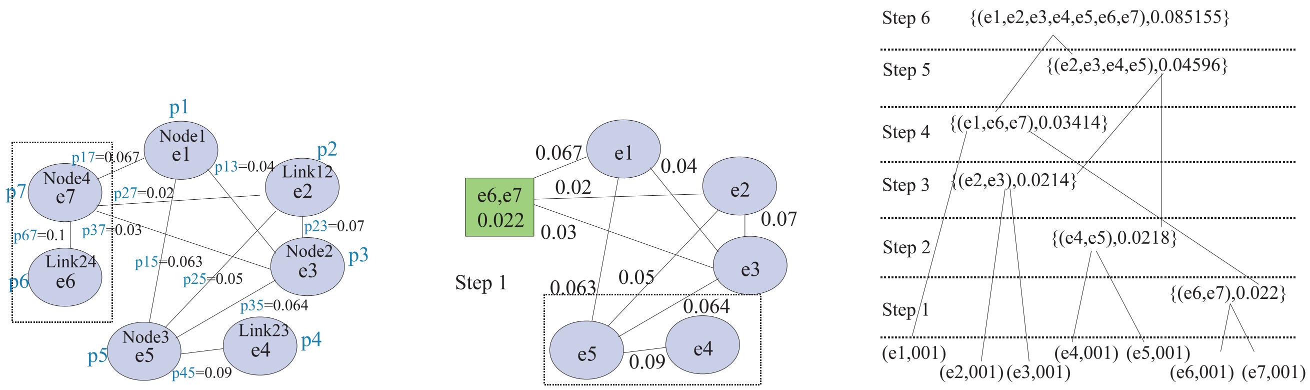

(a) Dependency graph with network elements with $p_i = 0.01$ and propagation probabilities marked over the links.

(b)After the first step Node6 and Node7 have been grouped into a new node, which has a failure probability of 0.022.

(c)Resulting hierarchy from the partitioning phase after grouping all the nodes into a single one.

Figure 11: Steps of the partitioning phase.

(1b) Method (1b) designs the FSM that models the failure behaviour of the network. It first defines a set of states, each state being associated to a failure scenario that may occur in the network. A last state is added to this set to represent the normal operation of the network.

Given this set of states, the transition between one state to another is defined by the event(s) (alarm(s) or notification(s)) that have to be received from the network.

(1c) The fault location algorithm, tailored to the model of Figure 10, has two parts, as Method (1a). The first part, which is called here Pre-Computation Part (PCP), involves all computations that are independent from the received alarms. The computation left to the second part, executed upon reception of the alarms, is kept minimal and the diagnosis is as fast as possible.

Once the established channels have been modeled, the PCP part consists in building a binary tree whose depth is the number of network elements able to send alarms, and whose leaves point to (potentially) faulty elements. Figure 12 shows the tree computed by the PCP part for identifying soft failures in the network of Figure 8, modeled as in Figure 10. The occupied leaves point to (a set of) pairs { network identifier, type of failure }. The type of failure is H for hard failures and S for soft failures.

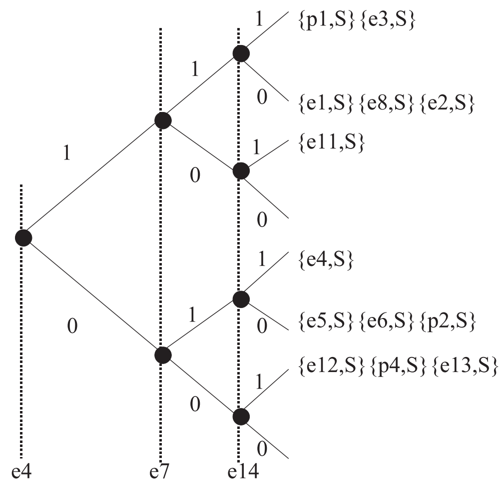

Figure 12: The binary tree obtained in the PCP part of Method(1c). For the sake of exposure, only a part of the tree is shown here, with a length limited to three active elements: $e4$, $e7$ and $e14$.

The PCP part has been implemented on the basis of the algorithm devised by Rao [Rao96], to locate single failures in a network with two kinds of network components (referred in Figure 10 as *passive (p)* components, which can fail and cannot communicate with the manager; and *active (e)* components, which can fail and send alarm to the manager informing that they detected an error).

The location was done in this example for an ideal scenario, that is, assuming that there are no lost and false alarms. For an extension to non-ideal scenarios, and to multiple hard and soft failures, see [MT00].

(2a) The only preprocessing needed here is the formalization of the human expertise into a correct set of rules.

(2b) The preprocessing in case based systems is limited to the storage of known failure scenarios in the database. Lewis' system `CRITTER` [HE93], applies case base reasoning to enhance a standard trouble ticketing system. A solved trouble ticket is a case that records a former problem and a trouble ticket database is the case library.

- (On-line) **processing and output** delivered by each method.

(1a) Method (1a) performs the diagnosis on the graph model of the network obtained at the end of the partitioning phase. When the alarms reach the manager, the second part of the algorithm, called *selection*, begins. This phase returns the smallest subset of nodes that explains all the alarms and has the maximum probability of containing the faulty element(s). It starts from the single node obtained at the end of the partitioning phase and which is on the top of the hierarchical structure shown in Figure 11(c). It goes down this structure choosing the branches that explain more alarms and has greater probability to contain the faulty element. It stops when it finds the smallest, most likely set of elements explaining all the received alarms.

(1b) Events (alarms and notifications) from the network cause transitions from one state to another. The output delivered by the FSM is the actual state of the system reached with the received events. For example in Figure 8, if the event is the alarm `Loss of light 24`, the system will pass from "Good Functioning" to "Link24 fail" and this will be the diagnosis delivered to the Human Manager.

(1c) As already mentioned, the first phase of the proposed algorithm produces the binary tree shown in Figure 12. Once the manager starts receiving alarms from the network, the algorithm simply has to traverse the binary tree. The binary vector representing the received alarms is used to traverse the tree and to retrieve the set of elements pointed by the corresponding leaf.

For example in Figure 12, if Elements $e7$ and $e14$ send alarms, the associated vector is (0 1 1) and the result returned to the manager is therefore Element $e4$, with a soft failure.

(2a) When alarms are received, the corresponding rule will provide a diagnostic. If no rule has been established for the received alarms, either no result, or a default result will be given. For example, if the received alarms are `Loss of light 12` and `Loss of light 23`, a pure Expert System would not give any solution.

When multiple alarms are received, and more than one rule applies, the proprietary system of Jakobson and Weissman [JW93] has the advantage of combining rules in such a way that the joint information on the failure is larger than the sum of individual information carried by each individual rule.

The system is used in a commercial shell called Intelligent Management Platform for Alarm Correlation Tasks (`IMPACT`), which is a GTE proprietary system, and has been extended into an event correlation system called Global Real-Time Advanced Correlation Environment (`GRACE`) [JWB$^+$00], which includes an IP/ATM/SONET fault management among other utilities.

(2b) Given alarm inputs, the delivered output are the already solved failure scenarios corresponding to the same input, or to the closest match in the history database. In the proposed example, and looking at Table 1, if the received alarms are `Loss of light 12` and `Loss of light 23` but not `Loss of light 24`, the case that could be delivered as an output is either Case1 or Case3 because they have one difference in the input (the first one corresponds to the existence of a lost alarm and the second one to the existence of a false alarm). When the result is confirmed by the manager, the new case is then stored in the database.

An enhancement in `CRITTER` is to adapt, if necessary, the "ticket" retrieved when a new problem is submitted to the system, and embed this new experience in the trouble ticket database.

	HF	SF	Corr. al.	Memory	Diagn.	Prepr.	FP
1(a) Probab. reas. model	Yes	No	No	Medium	Medium	High	Yes
1(b) FSM	Yes	No	No	Low	Low	High	No
1(c) Determinstic model	Yes	Yes	Yes	High	Low	High	No
2(a) Expert System	Yes	Ext.	No	Low	Low	Low	No
2(b) Case based system	Yes	Ext.	Yes	High	Medium	Low	No
2(c) Neural network	Yes	Ext.	Yes	Low	Low	High	No
2(d) Proactive system	Yes	Yes	Yes	Medium	High	Low	No

Table 2: Comparative table of the properties of the methods described in this paper: HF = Hard Failure Detection, SF = Soft Failure Detection, Corr. al. = Tolerance to false and missing alarm, Memory usage, Diag. = Diagnostic phase (alarm processing) complexity, Prepr. = Pre-processing phase complexity, FP = Knowledge of failure probabilities required. Ext. means that the method could have this property, but at the expense of a quite important extension of the databases/rules/etc.

6 Conclusion

The huge amount of information that can be lost in an optical network when a failure occurs, makes the fault location problem a crucial issue to solve. This article gave an overview on different approaches for fault location and their application to optical networks. Several features of these approaches allow to distinguish them. Table 2 shows a comparison of the methods based on the following features: possibility of detect hard and soft failures, tolerance to corrupted alarms, memory they need, computational complexity of the pre-computing and diagnosis phase and knowledge of probabilities. An important feature is the possibility to cope not only with hard but with soft failures, which are result of misalignement and equipment aging, quite frequent in optical networks. Another important feature to be considered in optical networks is the possibility of coping with lost and false alarms, that is, with corrupted alarms. False and lost alarms may occur when an analog signal is translated to a binary alarm, depending on whether it overpasses a threshold or not. Some approaches are able to cope with the existence of these alarms, such as the methods presented by Ho et al. [HCPZ00] and Mas and Thiran [MT00], or methods relying directly on the analog values such as the Neural Networks [GH97].

We can conclude that the ideal method to use in optical networks has to be sensitive to hard and soft failures, should rely on the analog signals and cope with lost and false alarms, should avoid the use of probabilities since they are time-varying parameters difficult to estimate, should be aware of all possible failure scenarios and how they can propagate by having correlation rules [BCF94] or an appropriate taxonomy of network elements [MT00]. The accurate location of the failure will also depend on the available information from the network. Optical networks provide, as shown in Section 4.1.1, quite few informations about the optical signal. A possible – but expensive – solution is to install more testing equipment in the optical physical layer. A second, cheaper solution is to exploit interoperability between layers with a single fault management system.

References

[12199] ITU-T COM-15 121. Signal Quality Monitoring in Optical networks, 1999.

[Anr99] Anritsu. Catalog of Electronic Measuring Instruments, 1999.

[BCF94] A.T. Bouloutas, S. Calo, and A. Finkel. Alarm Correlation and Fault Identification in Communication Networks. *IEEE Transactions on Communications*, 42(2/3/4):523–533, Feb/March/April 1994.

[Bea93] S. Brugnoni and et al. An expert system for real time fault diagnosis of the italian communications network. *Integrated Network Management*, III:617–628, 1993.

[BHS93] A.T. Bouloutas, G. W. Hart, and M. Schwartz. Fault Identification Using a FSM model with Unreliable Partially Observed Data Sequences. *IEEE Transactions on Communications*, 41(7):1074–1083, July 1993.

[DDH+99] B. T. Doshi, S. Dravida, P. Harshavardhana, O. Hauser, and Y. Wang. Optical Network Design and Restoration. *Bell Labs Technical Journal*, pages 58–84, January-March 1999.

[Dea99] P. Demeester and et al. Resilience in multilayer networks. *IEEE Communications Magazine*, pages 70–75, August 1999.

[DLW93] R. H. Deng, A. A. Lazar, and W. Wang. A probabilistic Approach to Fault Diagnosis in Linear Lightwave Networks. *IEEE Journal on Selected Areas in Communications*, 11(9):1438–1448, December 1993.

[G.898] ITU-T Rec. G.872. Architecture of Optical Transport Networks, 1998.

[Gea93] K. P. Gu and et al. Realization of an expert system for an on-line fault diagnosis and restoration in a bulk power system. *Proc. 4th International Symposium Expert Systems Applications Power Systems*, pages 327–331, January 1993.

[GH97] R. Gardner and D. Harle. Alarm Correlation and Network Fault Resolution using Kohonen Self-Organising Map. *Globecom 97 proceedings*, pages 1398–1402, 1997.

[GJW94] R. Weihmayer G. Jakobson and M. Weissman. A domain-oriented Expert System for Telecommunication Network Alarm Correlation. *Network Management and Control*, 2, 1994.

[GNS00] N. Golmie, T. Ndousse, and D. Su. A differentiated optical services model for WDM networks. *IEEE Communications Magazine*, pages 68–73, February 2000.

[Gol] Wandel & Goltermann. Understanding Dense WDM.

[Gop00] R. Gopal. Layered model for supporting fault isolation and recovery. *IEEE/IFIP: Network Operations and Management Symposium NOMS*, 2000.

[HCPZ00] L. L. Ho, D. J. Cavuto, S. Papavassiliou, and A. G. Zawadzki. Adaptive and automated detection of service anomalies in transaction-oriented WANs. *IEEE Journal on Selected Areas in Communications*, 18(5):744–757, May 2000.

[HE93] H.-G. Hegering and Y. Yemini (Editors). *Integrated Network Management.* Elsevier Science Publishers B.V., 1993.

[HHN98] S. Abek H. Hegerin and B. Neumair. *Integrated Management of Networked Systems.* Morgan Kaufmann Publishers, 1998.

[HJ97] C.S. Hood and C. Ji. Proactive Network-Fault Detection. *IEEE Transactions on Reliability*, 46(3), September 1997.

[JW93] G. Jakobson and M. Weissman. Alarm Correlation. *IEEE Network*, pages 52–59, November 1993.

[JWB+00] G. Jakobson, M. Weissman, L. Brenner, C. Lafond, and C. Matheus. GRACE: Building next generation event correlation services. *IEEE/IFIP: Network Operations and Management Symposium NOMS*, 2000.

[KS95] I. Katzela and M. Schwartz. Schemes for Fault Identification in Communication Networks. *IEEE/ACM Transactions on Networking*, 3(6), December 1995.

[KW87] J. Kleer and B.C. Williams. *Diagnosing multiple faults-Artificial Intelligence*, volume 32. Elsevier Science Publishers, 1987.

[Lin98] Along Lin. A hybrid approach to Fault Diagnosis in Network and System Management. *HP Technical Report*, HPL-98-20, 1998.

[LR95] C. S. Li and R. Ramaswami. Fault Detection and Isolation in transparent All-Optical Networks. *IBM Research Report*, RC-20028, April 1995.

[Mas00] C. Mas. *Fault Location Algorithms for Optical Networks.* These No. 2164, EPFL, 2000.

[Met00] C. Metz. IP Protection and Restoration. *IEEE Internet Computing*, pages 97–102, March-April 2000.

[ML97] Y. Maki and K. A. Loparo. Neural Network approach to Fault Detection and Diagnosis in industrial processes. *IEEE Transactions on Control Systems Technology*, 5(6):529–5541, November 1997.

[MT00] C. Mas and P. Thiran. An efficient algorithm for locating soft and hard failures in WDM network. *IEEE JSAC special issue on Protocols and Architectures for next generation optical WDM networks*, October 2000.

[MTB99] C. Mas, P. Thiran, and J.-Y. Le Boudec. Fault location at the WDM Layer. *Photonic Network Comm.*, 1(3):235–255, November 1999.

[Muk97] B. Mukherjee. *Optical Communication Networks.* Mc Graw-Hill Series, 1997.

[Nyb99] M. Nyberg. *Model based Fault Diagnosis: Methods, Theory and Automotive Engine Applications.* Linköping University, 1999.

[Pac99] Hewlett Packard. Test & Measurement Catalog. http://www.hp.com/go/tmc99, 1999.

[Pho00] Photonetics. Multi-wavelength analyzer. http://www.photonetics.com/products/pages/ifm/walics.html, 2000.

[Rao93] N. S. V. Rao. Computational Complexity Issues in Operative Diagnosis of graph-based Systems. *IEEE Transactions on Computers*, 42(4):447–457, April 1993.

[Rao96] N. S. V. Rao. On Parallel Algorithms for Single-Fault Diagnosis in Fault Propagation Graph Systems. *IEEE Transactions on Parallel and Distributed Systems*, 7(12):1217–1223, December 1996.

[RRM+96] C. Rodriguez, S. Rementeria, J. I. Martin, A. Lafuente, J. Muguerza, and J.Perez. A Modular Neural Network approach to Fault Diagnosis. *IEEE Transactions on Neural Networks*, 7(2):326–340, March 1996.

[Sea97] K. Struyve and et al. Design and evaluation of multiplayer survivability for SDH-based ATM networks. *Globecom 97 proceedings*, 1997.

[Sub99] Submarine Networks. *FibreSytems*, 3(5):26, June 1999.

[VJH97] P. A. Veitch, D. Johnson, and I. Hawker. Design of resilient core ATM networks. *Globecom 97 proceedings*, 1997.

[Wol86] B. Wollenberg. Feasibility study for an energy management system intelligent alarm processor. *IEEE Transaction on Power Systems*, PWRS-1:241–247, May 1986.

[WOSF99] N. Wauters, G. Ocahoglu, K. Struyve, and P. Falcao. Survivability in a New Pan-European Carriers' Carrier Network Based on WDM and SDH Technology: Current implementation and future requirements. *IEEE Communications Magazine*, 37(8):63–69, August 1999.

[WS92] C. Wang and M. Schwartz. Fault Detection with Multiple Observers. *IEEE INFOCOM Proc.*, pages 2187–2196, 1992.

[WS93] C. Wang and M. Schwartz. Identification of faulty links in dynamic-routed networks. *IEEE Journal on Selected Areas in Communications*, 11(9):1449–1460, December 1993.

[YS93] Y. Y. Yang and R. Sankar. Automatic failure isolation and reconfiguration. *IEEE Network*, pages 44–53, September 1993.