

Survey on Dependable IP over Fiber Networks

Maciej Kurant, Hung X. Nguyen, and Patrick Thiran

LCA-School of Communications and Computer Sciences
EPFL, CH-1015 Lausanne, Switzerland
maciej.kurant, hung.nguyen, patrick.thiran@epfl.ch

Abstract. This paper gives a survey of the techniques for failure location, protection and restoration in IP over optical fiber networks.

The first part of the paper reviews failure location algorithms at the optical and the IP layers. We classify the failure location algorithms at the optical layer into two main categories: the model based approach, that builds an abstract model of the network and uses this model to diagnose failures, and the learning based approach, that views the network as a black box and diagnoses failures using a set of rules obtained either by learning or by the expertise of the human manager. At the IP layer, we focus on the location of one of the main sources of failure: lossy links. The lossy link location algorithms can also be classified into two categories: the correlation approach, that requires strong correlation between monitoring packets, and the simple tomography approach, that requires some knowledge of the distribution of lossy links.

The second part of the paper describes the main strategies that ensure survivability in IP-over-fiber networks. After a failure, traffic can be restored either at the optical layer or at the IP layer. Protection at the optical layer amounts to dedicate some lightpaths to reroute the traffic disrupted by the failure. Restoration at the IP layer eliminates the need to set up back-up optical paths, but requires to map the IP layer on the optical layer in a survivable way. We describe the most common approaches achieving this.

1 Introduction

Communication networks in general, and the Internet in particular, are overlays of multiple layers. Each layer has different functions and all the layers cooperate to deliver data from the source to the destination. The simplest layer stack is IP (Internet Protocol) over physical. The physical layer is the one where bits of data are sent. It can be wired or wireless. We consider the case where the physical layer is optical and where quick failure detection and restoration are crucial because a failure can result in the loss of tetra (10^9) bits of data per second. In today's backbone networks, to increase the capacity of the optical fibers, the optical layer uses the Wavelength Division Multiplexing (WDM) technique to send data simultaneously at different wavelengths over a single fiber. The upper layer in this simple stack is the IP layer, where packets of data are routed. Although there exist layers on top of IP (e.g., application layer), they are beyond the

scope of this paper and we do not consider them here. In reality, there may exist some other layers in between the IP and WDM layers; the most frequent layer in backbone networks is SONET/SDH. SONET and SDH are a set of network interface standards and multiplexing schemes developed to support the adoption of optical fiber as a transmission medium. They use Time Division Multiplexing (TDM). SDH is the European standard whereas SONET is the US counterpart. This means that IP packets are transported over optical fibers that multiplex several connections either in time (TDM) at the SONET or SDH layer or in frequency (WDM) at the optical layer.

Failures occur frequently in communication networks. For instance, an average inter-failure time for the Sprint backbone network is about 12 hours [1]. Every network needs therefore to have a failure management system that can detect failures and take measures to guarantee the successful and timely delivery of data. When a failure occurs, it first needs to be detected and located. Then the traffic needs to be rerouted around the failure and the failing component has to be replaced [2]. In communication networks, failures at a lower layer will affect the performance of its upper layers, but the latter also have their own failures, unrelated to the lower layers. For example, a high optical signal-to-noise ratio (SNR) caused by a bent fiber will cause heavy losses on the IP links traversing the optical link, but heavy losses on these IP links can also be caused by overloaded network traffic, which is not visible at the optical layer. For this reason, each network layer uses its own failure management system. Moreover the failure management mechanisms at different layers need to cooperate with each other to avoid task duplication and increase efficiency.

We begin the paper with the first failure management task, which is to detect and locate a failure. Section 2 explores the various methods that are used, first to locate a faulty link at optical layer and next to locate a lossy link at the IP layer. In Section 3, we move to the second step in failure management, which is to engineer the network so that traffic is restored after the occurrence of a failure. We review the main methods used at the optical and the IP layer of an IP-over-fiber network. We conclude the paper in Section 4.

2 Failure Location in Optical and IP Networks

All existing techniques performing failure diagnosis rely on the analysis of symptoms and events that are generated during the occurrence of the failure. Simple failure location mechanisms are often based on locally monitored variables, such as the temperature of a device. The irregular values reached by these variables are logged as errors. Critical errors are sent to the network manager as alarms. Based on them, a failure is located. This is not a trivial task because some particular sets of alarms can have multiple possible explanations. Moreover, the set of alarms is sometimes noisy making the problem even more difficult. The noise is introduced by corrupted alarms, which are those alarms that unexpectedly arrive at the management system when they should not (*false alarms*) or those that do not arrive at the management system when they should (*missing alarms*).

The nature of failures and the available monitoring information are significantly different for the optical and the IP layer. Therefore, each layer needs to have its own failure location method. We address them separately in the following two subsections.

2.1 Failure Location at the Optical Layer

We are interested in detecting and locating failures of equipments at the optical layer. Some of the most common optical equipments and their operations are shown in the simple network of Fig. 1. A detailed survey of the failure location algorithms at the optical layer of an IP/SDH/WDM network can be found in [3]. In this section, we summarize the the most important algorithms discussed in [3] and add new developments that were not covered in that paper. We begin with a discussion of the available monitoring information. We then describe the types of failures that can be found at the optical layer and the alarms they generate. Finally, we compare and contrast the various methods that have been proposed in the literature to solve the failure location problem at the optical layer.




Fig. 1. A simple WDM network with three transmitters and three receivers. The three wavelengths coming out of Tx1, Tx2 and Tx3 are multiplexed at the multiplexer MUX before being transmitted on fiber OF. At the destination, these three wavelengths are demultiplexed at the demultiplexer DEMUX and then forwarded to Rx1, Rx2 and Rx3, respectively.

Available Monitoring Information. A failure at the optical layer can generate a large number of alarms within the optical layer, as well as from all of the layers above, such as SONET/SDH and IP. A failure location algorithm at the optical layer needs therefore to correlate alarms from all of these layers.

- At the optical layer, the monitoring information is delivered by the micro-controllers that control the optical equipments. Not all optical equipments are controlled. In an optical network, the most common optical equipments able to provide monitoring information and generate alarms are transmitters, receivers, switches, 3Rs (Re-generators/Re-shaper/Re-timer), protection switches, and amplifiers (a more detailed description can be found

in [4]). Transmitters send alarms when either the temperature or the incoming power is beyond a prescribed range. Receivers send alarms when the input optical power is under an acceptable level. 3Rs send alarms when they cannot lock to the incoming signal. Protection switches send alarms when they change the switch position due to an unacceptable incoming optical power. Switches send alarms when the connection of a particular input and a particular output cannot be established. Amplifiers send alarms when the pump laser does not work properly or when incoming power is not sufficient. Furthermore, if adequate testing equipment is deployed, the management system can obtain information about the quality of the optical signal such as signal to noise ratio and crosstalk by measuring the Bit Error Rate (BER) [5]. Devices measuring direct optical signal can be divided into two categories. (i) Global testing equipment (GTE), such as spectrum analyzers, measures the quality of the overall optical signal in a fiber. GTEs are able to produce the measurements of frequency and time of all wavelengths in a fiber. Examples of the GTEs are the MS26665C and MS2667C of Anritsu [6]. (ii) Individual testing equipment (ITE) can measure only properties of a single wavelength and depends on the transmission technology (ATM, SONET, SDH, etc.). An example of the ITEs is the MP1552 of Anritsu [6].

- At the SDH/SONET layer, the important failure notifications handled by SDH are loss of signal (LOS), loss of frame (LOF), loss of pointer (LOP) [7], degraded signal, and excessive error [8]. The SDH/SONET interface has also a set of mechanisms that are used for sending alarms upstream and downstream of the optical path to guarantee fast failure detection and recovery [9].

All of the above monitoring information is obtained passively without introducing any additional traffic into the network. We call this approach *passive* monitoring. A complementary technique is *active* monitoring where additional end-to-end connections (called probes) are created to measure the optical signal quality, see e.g. [10]. A degradation of the probing signal indicates failures at some of the optical devices used by the connection. More details of the recent progress in monitoring the performance of optical networks can be found in [11].

Fig. 2 provides a simple illustration of the available monitoring information for failures at the optical layer in an IP/SDH/WDM optical network. The data format at the IP layer is in packets, at the SDH/SONET layer it is in frames multiplexed in several time channels, and at the WDM layer it is in connections multiplexed in several wavelengths. When there is a failure at the physical layer, alarms from several layers will be sent to their own management platforms and failure protection and restoration mechanisms will be triggered at each layer. In the example of Fig. 2, when Node 1 fails, the WDM layer could start a failure location mechanism based on the alarms generated at the physical layer (for example a *Loss of Optical Power* at the receiver of Node 2). Otherwise, the SDH/SONET layer will react by applying protection in order to restore the interrupted connection based on the SDH alarms *Loss of frame*, *Loss of Pointer*, etc. issued by the SDH equipments. If the SDH layer cannot restore the end-to-end connection, the IP routers will detect the failure and try to find an alternative IP path.

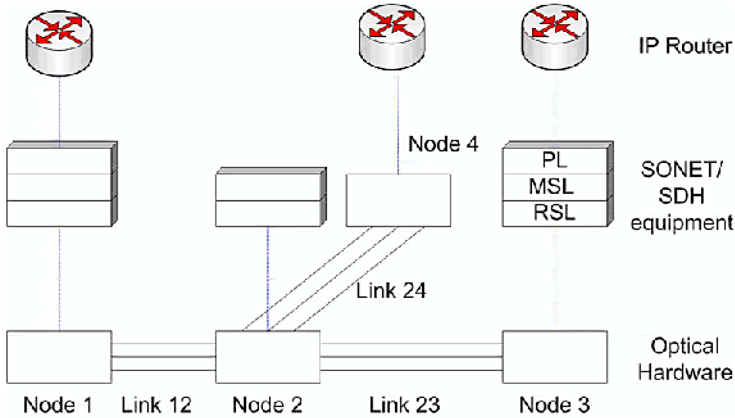


Fig. 2. Example of available monitoring information in an IP/SDH/WDM network. The notations PL, MSL, and RSL represent SDH Path, Multiplex Section and Regenerator Section layers, respectively.

Failures at the Optical Layer. We distinguish two types of failures at the optical layer: *hard* and *soft* failures.

- Hard failures are unexpected events that suddenly interrupt the optical channel. An example of a hard failure is a fiber cut. These failures can be detected at the optical layer from alarms sent by hardware devices.
- Soft failures are events that progressively degrade the quality of the signal transmission. An example of a soft failure is the variation of temperature of a laser: the output wavelength will drift as the laser heats up or cools down. In this case, the wavelength drift creates interferences with adjacent channels. The detection of soft failures often requires information from the upper layers, such as a SDH/SONET error frame rate. For example, when the wavelength is shifted, devices at the WDM layer will not detect any abnormality, but monitoring devices at the SDH and IP layer will observe increases in BER or SNR.

Failure Location Algorithms. In optical networks, a failure at an optical component not only results in faulty behavior at that component, but can also cause degradations in the signals sent from that component to other components. The other components may also forward the abnormal signals further. This manifestation is called *failure propagation*, and needs to be taken into account by the failure location mechanisms. Locating failures from the information provided by monitoring devices in systems with failure propagation such as optical networks is known to be NP-hard [12, 13, 14]. Many approaches have been proposed to solve this intractable problem. We divide them into two main categories: the *model based* methods and the *learning based* methods.

The model based methods [15, 16, 17, 18, 4] first construct an accurate and workable model for the networks on the basis of the functional and physical

properties of the network components, and then make a diagnosis by comparing actual observations with forecasts from the model. The advantages of the model based methods are that they are able to cope with incomplete information and unforeseen failures, and do not require learning. Their drawback is the difficulty of developing a good model for complex networks. We study three model based methods in this paper:

- (1a) the probabilistic reasoning system developed by Katzela and Schwartz [15],
- (1b) the FSM system developed by Li and Ramaswami [17],
- (1c) the deterministic system developed by Mas and Thiran [13].

The learning based methods view the system as a black box delivering outputs when a particular failure occurs. They learn the relationship between input events and output diagnosis, which can be done in different ways: by capturing the human expert knowledge and implementing it in an efficient way (expert systems) [19, 20, 21], by recording the history of previous cases that occurred in the past (cased-based systems) [22], by artificial neural networks [23, 24] or by any other algorithm with statistical learning capabilities [25]. The main advantage of the black box methods is that they do not require detailed model of the networks. However they need long learning processes. We consider the two learning based methods:

- (2a) the expert system presented by Jakobson et al. [26],
- and (2b) the case based system proposed by Lewis [22].

There are also hybrid methods [27,28,29] that combine the two aforementioned approaches and inherit both their advantages and disadvantages.

Comparison of Failure Location Algorithms. We now compare the failure location methods introduced in the previous section (1a, 1b, 1c, 2a, and 2b) by applying them to the example network in Fig. 2. These techniques are compared with respect to the *input data* they need and their *methodology*.

Input data is the information required by the failure location algorithms from the monitoring tools (timestamps, failure probabilities). It is different for the five studied algorithms.

- Method (1a) needs (i) network topology, (ii) the failure probabilities of each node and link and (iii) the probabilities that a failure at one component will propagate to the others (failure propagation probability). For the example network in Fig. 2, the graph representing the physical layer could look like the one shown in Fig. 3, where each element is either a network node or a fiber, and has an associated failure probability p_i . Every link between two nodes has a weight p_{ij} , which is the failure propagation probability.
- Method (1b) needs the network topology and the finite state machine (FSM) for that specific network. In order to design the FSM, the network manager




Fig. 3. The probabilistic graph model for the network in Fig. 2

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. In the simple example network in Fig. 2, considering up to two simultaneous failures, the FSM looks like the set of interconnected states shown in Fig. 4.

- Method (1c) does not need the topology but instead requires the set of established end-to-end connections in the network. Each connection is then viewed as a channel containing an ordered set of network elements. Fig. 5 shows the model for the example network with three channels.
- Method (2a) needs to have the manager experiences and translates these experiences into a set of “if/then/else” rules. In our example, the rules could be:

If loss of light 23 then
 If loss of light 24 then
 If loss of light 12 then link 12 fails
 else node 2 fails
 else link 23 fails

- Method (2b) takes as input the history of all previous solved failures: sets of alarms that are received and their diagnosis results. In our example, some solved scenarios would be:

Loss of light 23 is caused by failure of link 23.
Loss of light 23 and loss of light 24 are caused by failure of node 2.

Methodology is the actual algorithm used to locate the failure. The methodologies of the five studied techniques are:

- Method (1a) first designs the directed dependency graph as in Fig. 3 and then applies a divide and conquer algorithm [15], in two phases. In the first phase, called the *partitioning phase* that can be done off-line, it groups iteratively the nodes by taking the two nodes of the graph i, j , for which the failure propagation probability p_{ij} is largest, and by replacing them by a single

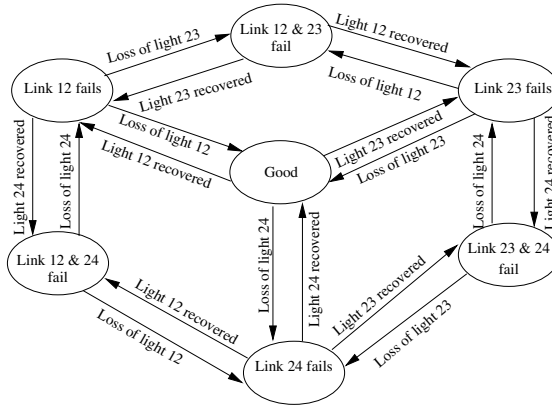


Fig. 4. The FSM model for the network in Fig. 2

new node k whose failure probability is $p_k = p_i + p_j \times p_{ji} + p_j + p_i \times p_{ij}$. The new propagation probabilities p_{kl} and p_{lk} , involving another node l and the new node k are the maximum of the previous propagation probabilities: $p_{kl} = \max\{p_{il}, p_{jl}\}$, $p_{lk} = \max\{p_{li}, p_{lj}\}$. The iterations stop when the all nodes in the dependency graph are merged into a single node that is also the root of the tree. The second phase, called the selection phase, is carried out on-line when alarms arrive at the management system. The algorithm starts from the root node of the tree obtained at the end of the partitioning phase and traverses the tree by choosing the branches that explain most alarms and that have a greater probability of containing the faulty element. It stops when it finds the smallest most likely set of elements explaining all the received alarms.

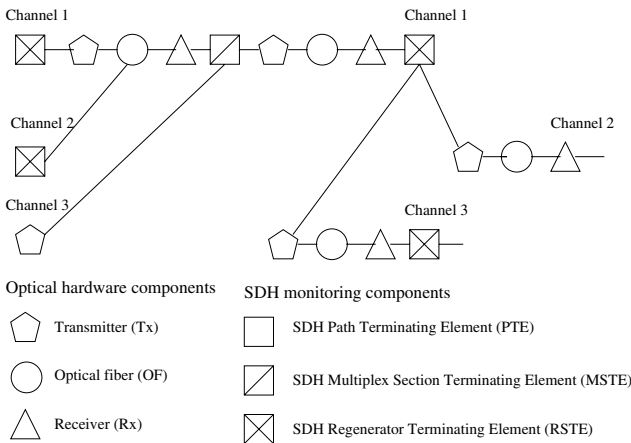


Fig. 5. The deterministic graph model for the network in Fig. 2

- Method (1b) needs to design the FSM that models the failure behavior 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 states is defined by the events that have to be received from the network. When events (alarms or repairs) arrive, the FSM changes its state. The output of the FSM is the actual state that the system is in.
- Method (1c) consists of two stages. In the first stage, which can be done off-line, the algorithm determines which alarms will ring when a network component fails. In the second stage, which is carried out when alarms arrive, the algorithm first corrects possible alarm errors by determining the most likely set of erroneous alarms. The algorithm then solves the resulting failure location problem with the cleaned alarms by iteratively picking the network component whose failure generates the largest number of alarms, until all alarms are covered by at least one failure.
- Method (2a) first defines all the rules given by the expert knowledge of the manager. When alarms are received, the corresponding rule will provide a diagnosis. If no rule has been established for the received alarms, either no result or a default result will be given.
- In method (2b), given the alarm input, the delivered outputs are the already solved failure scenarios corresponding to the closest match in the history database. When the result is confirmed by the manager, the new case is then added to the database.

Table 1 summarizes the comparisons of different failures location methods at the optical layer.

Table 1. Comparative table of the properties of the failure location algorithms at the optical layer described in this section: *HF*= Hard failure, *SF*= soft failure, *Memory*= Memory usage, *Diagn.*= On-line diagnostic phase (alarm processing) complexity, *Prepr.*= Off-line 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 database/rules/etc.

	<i>HF</i>	<i>SF</i>	<i>Memory</i>	<i>Diagn.</i>	<i>Prepr.</i>	<i>FP</i>
<i>1(a) Probabilistic model</i>	Yes	No	Medium	Medium	High	Yes
<i>1(b) FSM</i>	Yes	No	Low	Low	High	Yes
<i>1(c) Deterministic model</i>	Yes	Yes	High	Low	High	No
<i>2(a) Expert System</i>	Yes	<i>Ext.</i>	Low	Low	Low	No
<i>2(b) Case based System</i>	Yes	<i>Ext.</i>	High	Medium	Low	No
<i>2(c) Neural network</i>	Yes	<i>Ext.</i>	Low	Low	High	No
<i>2(d) Proactive system</i>	Yes	Yes	Medium	High	Low	No

2.2 Failure Location at the IP Layer

In this section we survey failure detection and location mechanisms at the IP layer. As in the previous section, we begin with a discussion of the available monitoring information. We then describe the types of failures that occur at this layer. Finally, we will compare and contrast the various methods that have been proposed in the literature.

Available Monitoring Information at the IP Layer. At the IP layer, performance information can be obtained in different ways, depending on the accessibility to the individual routers.

- **Direct measurements at routers:** Network managers can configure routers in their network to maintain information about their own performance. For example, routers can keep count of the numbers of packets dropped due to some reasons. These information can be collected and transmitted to the network manager using mechanisms such as SNMP [30] at regular time intervals. Collecting performance information requires significant memory and computing resources from the routers. Even though sampling techniques can be used to reduce these requirements, in practice, the direct measurements can only be made at intervals of minutes (typically five minutes). Therefore, despite their potential to give accurate information for the network manager, direct measurements are the least reliable and informative way to collect performance data.
- **Passive measurements using dedicated monitors:** Network managers can also deploy passive monitors in the network at multiple points to measure the performance of packets, such as the arrival time of a packet at a specific monitor [31]. From these measurements, performance metrics such as one way delay and packet losses on the segments between monitors can be inferred. These methods have the advantages of being non-intrusive and quite accurate, see e.g. [32]. The drawback is that they require monitors to be installed at multiple locations and can be deployed only by the network owner. Even though packet monitors are cheap, their deployment and maintenance costs are substantial.
- **End-to-end measurements:** In this approach, one infers the state of the network devices through the observed performance of end-to-end monitoring packets. A special feature of the probing approach is that it allows people without privilege rights to measure the networks. This approach is important in today's IP networks where traffic traverses different administrative domains and there is no incentive for the owners of each sub-network to collect and freely distribute vital statistics of their networks. There are many different types of probes one can use, namely ICMP response packets, TCP SYN/ACK, DNS, HTTP page downloads, as well as dedicated probe protocols. These factors have led the end-to-end measurement approach to be the most the widely deployed method [33, 34]. Note here that end-to-end information can be obtained either actively by injecting probing traffic into the network, or passively by listening to existing traffic in the network.

Since most of the measurement data available for the failure diagnosis of wide area IP networks is end-to-end, we only consider these measurements in this paper.

Failures at the IP Layer. The IP layer employs a sophisticated set of routing mechanisms to carry data between end points whenever possible. However, the IP layer can only provide a best effort service and does not guarantee the timely nor even the successful delivery of the data.

Many applications, such as voice or video, require strict loss and delay requirements for acceptable quality. For example, at loss rates of 4-6% or more, video conferencing becomes irritating and non-native language speakers are unable to communicate. The occurrence of long delays of 4 seconds or more at a frequency of 4-5% or more is also irritating for interactive activities, such as telnet or X windows. Paxson [33] reports that a loss of 5% has a significant adverse effect on TCP performance, because it will greatly limit the size of the congestion window and hence the transfer rate, whereas 3% is often substantially less serious. A loss rate of 2.5% makes conversations using Voice over Internet Protocol (VoIP) slightly annoying. A more realistic burst loss pattern results in VoIP distortion going from not annoying to slightly annoying when the loss rate goes from 0 to 1% [35]. Round trip times (RTTs) should be $RTT < 400\text{ms}$ for the interactive applications. VoIP requires a $RTT < 250\text{ms}$ or it is hard for the listener to know when to speak [35].

Failure management at the IP layer is mainly concerned with the ability of the network to deliver data within some bounds on loss rates and/or delays. Of the various metrics (loss, delay, throughput) that one can use to evaluate the performance of an IP network, loss is the most critical; this is because other metrics can be inferred from it. For example, the throughput of a TCP connection can be calculated using loss and delay information [36]. In this section, we concentrate only on the detection and location of IP links that have loss rates above 1% (*lossy links*). A lossy link can be caused either by failures at the optical layer or by congestions at the IP layer. We do not distinguish these two cases in this paper. Knowing the locations of lossy links, an application can significantly improve its performance by rerouting around them [37].

Lossy Link Location Problem Definition. We focus on the techniques that can be used to infer lossy links using end-to-end measurements. The inference of internal link properties given end-to-end observations is called network tomography. Most tomography works consider tree topologies like the one depicted in Fig. 6. Each node in the tree represents a router or an end-host. Each link represents a connection between two routers/hosts. Note here that the link can be a single physical link or a chain of physical links connected by intermediate routers. Probing packets are sent from the source at the root node to the receivers at the leaf nodes along paths that pass through several internal nodes. The goal of loss tomography is to estimate individual link loss rates based on the loss rate perceived at a few end nodes.

The network is modelled as a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where the set \mathcal{V} of nodes denote the network routers/hosts and the set \mathcal{E} of edges represent the communication links connecting them. The number of nodes and edges is denoted

by $n_v = |\mathcal{V}|$, and $n_e = |\mathcal{E}|$, respectively. For a known topology $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and a set of end-to-end paths \mathcal{P} , $n_p = |\mathcal{P}|$, we compute the routing matrix D of dimension $n_p \times n_e$ as follows. The entry $D_{ij} = 1$ if the path P_i contains the link e_j and $D_{ij} = 0$ otherwise. A row of D therefore corresponds to a path, whereas a column corresponds to a link.

Let ϕ_i denote the packet transmission probability on path P_i and ϕ_{e_j} the packet transmission probability on link e_j . Clearly, the loss rate of a link e_j equals to $1 - \phi_{e_j}$. Therefore, estimating the link loss rates amounts to estimating the variables ϕ_{e_j} from the measured path transmission rates ϕ_i . Assuming independence among loss events on links, the relation between the path-wise and link-wise transmission rates reads

$$y = Dx = \left[\sum_{j=1}^{n_p} x_j D_{ij} \right]_{1 \leq i \leq n_e} \tag{1}$$

where $y_i = \log(\phi_i)$ and $x_j = \log(\phi_{e_j})$: y is the vector of measurements, e.g. path packet transmission rates, and x is the vector of link transmission rates.

The network loss tomography problem boils down to solving the linear system of equations (1) to find x given y and D .

Lossy Link Location Algorithms at the IP Layer. Equations (1) cannot be solved directly because most of the time the matrix D is rank deficient, that is, $\text{rank}(D) < \min(n_p, n_e)$ (the rank of D is the maximal number of columns (rows respectively) of D which are linearly independent). The non-uniqueness of link loss rates is illustrated in the example of Fig. 6 [38].




Fig. 6. In the figure, the nodes are the network routers/hosts and the directed links are the communication links connecting them. ϕ_{e_j} denotes the transmission rate of link e_j and c is a constant between $\max\{\phi_{e_2}, \phi_{e_3}\}$ and $1/\phi_{e_1}$. Both set of link transmission rates give the same end-to-end transmission rates. Link transmission rates therefore cannot be uniquely calculated from end-to-end transmission rates.

Two techniques can overcome the non-uniqueness solution problem to identify lossy links at the IP layer. The first approach, called the *correlation* approach, introduces additional constraints to (1) by creating a correlation between probing packets. The second approach, called the *simple tomography* approach, exploits the distribution of link loss rates on the Internet to solve (1).

The correlation approach can be realized by either using multicast [39] or unicast probing packets [40,41]. Multicast packets are sent to a group of subscribing receivers. At internal branching points, each multicast packet is replicated and sent along each branching path. In contrast, unicast packets are sent to only one receiver. To correlate them, the unicast packets to different receivers are sent almost at the same time such that they experience the same loss behavior on the common links shared by different receivers. Several challenges exist in bringing the multicast or unicast methods into widespread fruition. On one hand, multicast is not widely deployed. On the other hand, methods based on unicast probing incur costs to deploy appropriate data collection softwares. We study in this paper three correlation based methods:

- (1a) the multicast method developed by [39],
- (1b) the unicast packet pair method developed by [40],
- (1c) the unicast packet train method developed by [41].

The difficulties encountered in the previous methods motivate the simple tomography approach that does not require the correlations between probing packets [42,38]. The simple tomography approach is based on the assumption that network links are generally lossless and that only a few links are responsible for dropped packets. The major advantage of this approach is that applications already monitor packets from end-to-end. Simple tomography methods do not seek to calculate the exact loss rate for each link. Instead, they use a threshold t_l , called link threshold, to determine whether a link e_k is good ($\phi_{e_k} \geq t_l$) or bad (lossy) ($\phi_{e_k} < t_l$). The threshold t_l can be set either to meet a given transmission rate target, or on the basis of data history that shows a clear value separating well and badly performing links. The problem of identifying lossy links without finding exact link loss rates amounts to finding the most probable solution for the observed end-to-end data. Knowing that bad links are not frequent, the most probable solution is the one giving the least number of lossy links. Let us consider the example in Fig. 6. Assuming that the threshold separating good and bad links is 0.99, if the end-to-end transmission rates to the sink of both nodes 2 and 3 are below 0.98 ($\approx 0.99 \times 0.99$), the most probable explanation is that link 0-1 is lossy (having transmission rate less than 0.99). Other explanations require at least two links to be lossy, and are therefore much less likely. We consider in particular two simple tomography methods:

- (2a) the simple tomography using Monte-Carlo simulation method developed in [42],
- and (2b) the simple tomography using set-cover heuristic method in Duffield [38].

Comparison of Lossy Link Location Algorithms at the IP layer. We now compare the lossy link location methods introduced in the previous section (1a, 1b, 1c, 2a, and 2b) by applying them to the example network in Fig. 6. These techniques may be compared with respect to the *network support* they require and their *methodology*.

Network support is the additional support needed by the loss link location algorithms in addition to the network topology. It is different for the five studied algorithms.

- Method (1a) requires (i) multicast support from all the routers and (ii) specific software packages installed at the multicast sender and receivers to send, collect and analyze the multicast traffic.
- Methods (1b) and (1c) need to have some specific softwares installed at the unicast sender and receivers to send, collect and analyze the unicast traffic.
- Methods (2a) and (2b) do not need any additional support from the network.

Methodology is the actual algorithm used to locate the lossy links. The methodologies of the five studied techniques are:

- Method (1a) uses multicast packets to calculate the link loss rates. In the network of Fig. 6, if a multicast packet is sent by the sender at node 0 and received by the receiver at node 2 but not the one at node 3, then it can be immediately determined that the loss occurred on link 3 (successful reception at node 2 implies that the multicast packet reached the internal node 1). By performing such measurements repeatedly, the loss rates on the two links 2 and 3 can be estimated; these estimates and measurements enable to deduce the loss rate on link 1. To illustrate the method further, let $\phi_{2|3}$ be the ratio of the number of multicast packets simultaneously received at both nodes 2 and 3, relatively to the total number received at node 3. In other words, $\phi_{2|3}$ is the empirical probability of success on link 2 conditional upon success on link 3, which provides a simple estimate of ϕ_{e_2} . Similarly, we define $\phi_{3|2}$ as the probability of success on link 3 conditional upon success on link 2, and let ϕ_1 and ϕ_2 be the transmission rates of multicast packets for node 2 and 3. We can then write

$$\begin{pmatrix} \log \phi_2 \\ \log \phi_2 \\ \log \phi_{2|3} \\ \log \phi_{3|2} \end{pmatrix} = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \log \phi_{e_1} \\ \log \phi_{e_2} \\ \log \phi_{e_3} \end{pmatrix}. \quad (2)$$

A least square estimate of ϕ_{e_j} is easily computed for this over-determined system of equations. Sophisticated and effective algorithms have been derived for large scale network tomography in [39].

- Method (1b) was introduced in [40] to overcome the fact that most routers in the Internet today do not support multicast and that performance observed

by multicast packets differs significantly from that observed by unicast traffic. In method (1b), back-to-back packet pairs, each consisting of two packets sent one after the other by the sender, are sent to every pair of receivers. When two packets of a pair arrive back-to-back in a common queue, the successful transmission of one packet indicates, with high probability, the successful transmission of the other packet. If two back-to-back packets are sent to node j from the previous node $\rho(j)$ on a path, then let β_j be the conditional probability that the first packet of a pair arrives at node j from $\rho(j)$ given that the second packet sent by $\rho(j)$ has arrived successfully at node j . β_j is very close to 1. Denote the complete set of conditional success probability by $\beta = \{\beta_j\}_{j=1}^{n_e}$, where n_e is the number of links. In this case, the network tomography problem boils down to determining the values of ϕ_{e_j} and β_j that best explain the probing results. Maximizing the likelihood is not simple and, consequently, numerical optimization strategies are required. The most commonly used method is the expectation-maximization algorithm (EM) [40].

- Method (1c) also uses back-to-back unicast packets but in a different way. The main objective of method (1c) is to create multicast probing using unicast packets. In method (1c), the sender sends a sequence (a train) of many back-to-back packets to all receivers (instead of packet pairs to all pairs of receivers as in method (1b)). Contrary to method (1b) where the conditional probabilities β_j are treated as variables, method (1c) assumes that $\beta_j = 1$ for all j . By viewing each packet train as a multicast packet, method (1c) then uses techniques in method (1a) to infer the link loss rates.
- Method (2a) is a simple tomography method proposed by [42], which means that it does not seek to calculate the exact link loss rates, but that it determines whether a link is lossy or not. Method (2a) uses the sophisticated Monte Carlo Markov Chain Simulation (*MCMC* method) to determine the lossy link. It tries to determine the posterior distribution, $\mathbb{P}(x|y)$, of the link loss rates in logarithmic scale x given the observed path data, y . Knowing $\mathbb{P}(x|y)$, one can draw samples from this distribution where each sample is a vector containing the transmission rates for all links in the network that can explain the observed data. The method then collects the transmission rates of each link in all samples and compares them with the threshold t_l . If the majority of the sampled transmission rates of a link are bad ($< t_l$), then the link is declared as bad. Otherwise it is declared as good. In general, it is hard to compute $\mathbb{P}(x|y)$ directly because of the complex integrations, especially when x is a vector, as in the present case. It is also difficult to obtain samples of the distribution $\mathbb{P}(x|y)$. Hence method (2a) uses an indirect approach to collect them by constructing a Markov chain whose stationary distribution is exactly equal to $\mathbb{P}(x|y)$. When such a Markov chain is run for a sufficiently large number of steps, it converges to its stationary distribution. The method then gathers samples from this stationary distribution and views them as samples from the posterior distribution $\mathbb{P}(x|y)$. This way, it does not have to determine the distribution $\mathbb{P}(x|y)$ and then draw the

Table 2. Comparative table of the properties of the methods described in this section for lossy link location at the IP layer: the column *Loss Rates* indicates whether the method can infer the exact loss rates, and the column *Meas. Errors* indicates whether the method can handle measurements errors, e.g., errors in estimating the end-to-end loss rates.

	<i>Monitoring Costs</i>	<i>Processing Time</i>	<i>Loss Rates</i>	<i>Meas. Errors</i>
1(a) <i>Multicast tomography</i>	<i>High</i>	<i>Low</i>	<i>Yes</i>	<i>Yes</i>
1(b) <i>Unicast packet pair</i>	<i>Medium</i>	<i>High</i>	<i>Yes</i>	<i>Yes</i>
1(c) <i>Unicast packet train</i>	<i>Medium</i>	<i>Low</i>	<i>Yes</i>	<i>Yes</i>
2(a) <i>MCMC</i>	<i>Low</i>	<i>High</i>	<i>No</i>	<i>Yes</i>
2(b) <i>SCFC</i>	<i>Low</i>	<i>Low</i>	<i>No</i>	<i>No</i>

samples from it. For the detailed construction of the Markov chain, we refer to [42].

- Method (2b) is a simple tomography method proposed by Duffield [38] (Duffield called this approach the *SCFC* algorithm). It first determines a threshold $t_p = t_l^{n_{link}}$, with n_{link} is the number of links in the path, for all paths. It then determines all end-to-end paths that have bad transmission rates, that is, whose transmission rate is below t_p . By observing that a path is bad if and only if one of its links is bad and that bad links are rare, it tries to find the smallest number of links whose badness can explain the badness of all bad end-to-end paths. The SCFC method adopts a greedy heuristic that iteratively chooses at each step the link that can explain the largest number of bad paths and infers that the link is lossy.

Table 2 gives a summary of the comparisons of different lossy link location methods at the IP layer.

3 Failure Protection and Restoration in IP/WDM Networks

Failures must not only be identified and located. The network must be designed so that the traffic is protected against them, which implies rerouting rapidly the traffic when a failure occurs, until it is repaired.

So far, we distinguished between a first generation IP/SONET/WDM network and a second generation IP/WDM network. In the context of failure protection, however, speaking at a functional level, there are no big differences between IP/SONET/WDM networks protecting traffic at the SONET layer, and IP/WDM networks protecting traffic at the optical layer, since the optical layer should take over all protection and restoration functionalities of the SONET layer. Therefore in this section we do not make a distinction between both architectures and we use the SONET/SDH and optical layer indifferently. This brings us to the analysis of two layers: the *physical layer* (optical) and the *logical layer* (IP).

The physical layer topology is a set of optical switches (nodes) and fibers (links) interconnecting them. Each logical link is mapped on the physical topol-

ogy as a *lightpath*. A lightpath may span multiple fiber links. A set of all lightpaths defines a *mapping* of the logical graph on the physical graph. Given the physical and logical topologies, the problem of finding a mapping and assigning wavelengths to all logical links is called the Routing and Wavelength Assignment (RWA) problem. In its general form, the RWA problem does not take failure resilience into account - its objective is to minimize the network resources. A survey on RWA algorithms can be found in [43]. The difficulty of the RWA problem depends partially on the type of physical nodes used in the network. Perhaps the simplest kind of physical node is an optical crossconnect (OXC). It switches the optical signal from an input port to an output port without wavelength conversion. In this case a lightpath must occupy the same wavelength on all fiber links it traverses, which is called a *wavelength continuity constraint*. Physical nodes can be equipped with *wavelength converters* to alleviate this constraint: some can offer a full conversion capability (that is, any wavelength can be converted to any other wavelength), others only offer limited conversion capability (that is, a wavelength can only be converted to the neighboring wavelengths on the spectrum). When the full wavelength conversion is available at every node, the RWA problem is in fact reduced to the routing problem only.

Generally, there are two approaches for providing survivability of IP-over-WDM networks: protection and restoration [44]. *Protection* is the mechanism in which the traffic is switched to available resources when a failure occurs. It needs to be very fast, the commonly accepted standard for physical layer is 50 ms. Protection routes must therefore be pre-computed, and wavelengths must be reserved in advance at the time of connection setup. For speed requirements, protection may require fairly simple topologies (rings rather than complex meshes) and may be performed in a distributed way, without relying on a central management entity to coordinate actions. *Restoration* is the mechanism in which new routes are established on the fly, after a failure has occurred. This is much slower than protection and requires enough free resources available at the moment of the failure.

The failure protection and restoration tasks can be carried out at different layers. Often, the logical layer uses restoration (IP restoration) and the physical layer uses protection (optical protection). We illustrate these approaches in the toy example in Fig. 7, where three IP connections are mapped on a six-node physical topology. Assume that each fiber can carry two wavelengths λ_1 and λ_2 , each of the capacity of one unit of traffic. Fig. 7a shows an example of protection at the physical layer. This is achieved by setting up three primary lightpaths (set in bold), all on wavelength λ_1 , which are used to carry the traffic in absence of failures. For each primary lightpath we also prepare a backup lightpath (dashed) on wavelength λ_2 . If the fiber (5,6) fails then all the traffic on the primary lightpath (1,6,5,4) on λ_1 , is routed over the backup lightpath (1,3,4) on λ_2 . Note that due to very small reaction times, these mechanisms are transparent to the logical layer.

The Internet Protocol (IP) is also capable of restoring the traffic around a failed facility. It is illustrated in Fig. 7b. Here each logical (IP) link has only

one corresponding lightpath. Routers periodically exchange keep-alive or hello messages to check the health of neighboring links and nodes. A failure of the fiber (5,6) does not trigger any action at the physical layer. Instead, after the loss of a few successive hello messages between the routers 1 and 4, the logical link (1,4) is deduced to have failed. Now, the traffic between the nodes 1 and 4 is rerouted in the logical topology via node 3. In order to enable this, two requirements must be met; first, a single physical failure cannot cut the connectivity at the logical layer, and second, the links at the logical layer must be overprovisioned, in order to be able to absorb the additional traffic rerouted after a failure. Let us assume that in the example in Fig. 7b all logical links initially carry (before a failure) the same amount t of traffic. In order to enable the IP restoration, t can be, at most, half of one traffic unit (i.e., half of the optical channel capacity). Note, however, that overprovisioning has a positive side effect of keeping the links under-utilized during regular operations and therefore of maintaining all delays short in the network.

Despite the requirements it imposes, the IP restoration approach turns out to be more resource efficient than optical protection. This is partially due to the

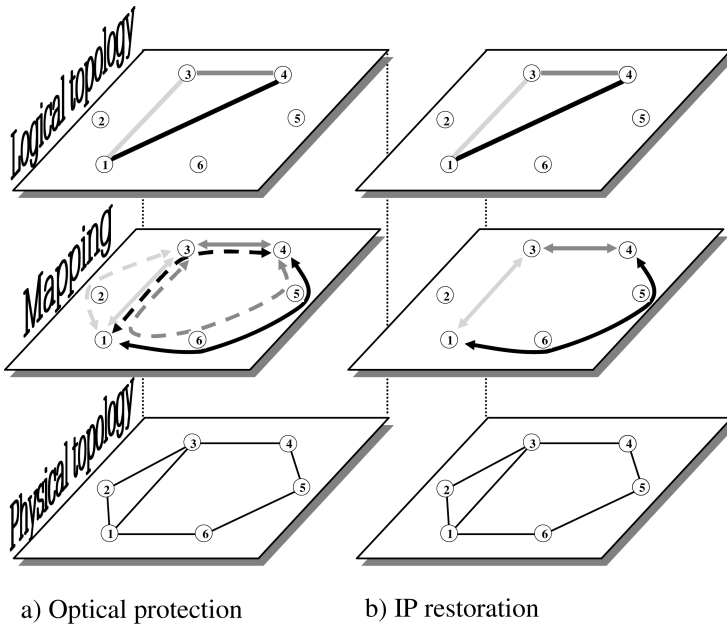


Fig. 7. An illustration of the basic concept of optical protection (left) and IP restoration (right). The logical topology consisting of three IP links is mapped on the physical topology with six nodes and seven optical fibers. Each fiber can carry two wavelengths λ_1 and λ_2 , each of the capacity of one unit of traffic. The lightpaths are represented by the arrows in the mapping. For example, the logical link (1,4) uses the lightpath (1,6,5,4). In the WDM protection scheme, the primary lightpaths use λ_1 and are set in solid, whereas the backup lightpaths use λ_2 and are dashed.

different granularity of the approaches: at a packet level in IP restoration vs. at a wavelength level in optical protection. Assume, for instance, that initially every logical link is loaded with at most half a unit of traffic. Under this assumption, both failure protection schemes in Fig. 7 can deal with any single fiber failure. Note that the primary lightpaths in the optical protection scheme are exactly the same as the lightpaths set up in IP restoration, hence the resources used by them are the same. The optical protection approach, however, commits additional resources by setting up the backup lightpaths, which makes it more resource-consuming. This effect is even stronger in denser and bigger topologies.

A major difficulty in optical networks that support various upper layers is that each layer performs its own protection and mechanisms independently from the others. This can lead to undesirable races between layers to protect traffic. For example, if the optical layer is protection enabled and if it did not recover from the failure very rapidly, the logical layer might happen to detect the failure. It will start rerouting the IP traffic around the failed link(s) or router(s). The lack of coordination between layers can therefore create a quite intricate situation. The problem of inter-layer coordination is addressed for example in [45, 46, 47]. Having protection at only one layer might simplify the problem, but one still needs then to choose the layer at which it should be done. It is not obvious which layer is more suitable for failure protection/restoration; each has pros and cons [48, 44]. First of all, some failures, such as a failed line card in a router, cannot be detected at the lower layer, but only at the IP layer (i.e., by IP restoration). Another advantage of IP restoration, as we have seen above, is its resource efficiency. Unfortunately, it is also inherently slow - failure detection at the logical layer takes tens of seconds at least, and time scales at which restoration occurs are typically at least three orders of magnitude larger than the protection processes at the physical layer. However, many real network operators deploy IP restoration only, and find it an effective and cost-efficient solution (see e.g., [49]). Some multi-layer protection/restoration schemes can adequately combine the advantages of each layer and still avoid most of their disadvantages [46], but do not eliminate the complexity of coordinating the different restoration schemes at the various layers (some solutions are proposed in [46]). One way to bypass some of this complexity race is to allocate the restoration task to a different layer for different traffic classes¹, which also brings benefits in resource usage [50].

We discuss now in more details the techniques used to protect and restore traffic, first at the physical layer, and next at the logical layer.

3.1 Protection and Restoration at the Physical Layer Only

All protection techniques involve some redundant capacity with the network to reroute traffic in case of a failure. There are essentially two basic protection mechanisms used in point-to-point links: 1+1 protection, and 1:1 protection (and its generalization to 1:n protection).

¹ A traffic class might be defined for instance by its origin/destination, bandwidth or maximal delay and jitter.

In 1+1 protection, traffic is transmitted simultaneously on two separate fibers on disjoint routes. The receiver selects the signal at the destination that has the largest incoming power. If that fiber is cut, it will automatically switch to the other fiber. This is the fastest and simplest protection, because no signalling is needed. It is however very inefficient in terms of resources, as every unit of traffic requires twice as much capacity. As a result, it is used in some ring networks (Unidirectional path-switched rings, see [51]), but not in large, meshed WDM networks.

In 1:1 and, more generally speaking, 1: n protection, traffic is transmitted only on one fiber (called working or primary fiber). If this fiber is cut, the sender and receiver both switch to the other fiber (called protection or back-up fiber). This is not as fast nor as simple as 1+1 protection, because the destination must detect the failure first and then signal it to the source, which will then switch over to the protection fiber. The advantage of 1:1 protection is that the capacity on the back-up fiber can be spared for unprotected traffic, which will be preempted in case of a failure, or can be shared between n multiple, physically disjoint working paths, in which case one speaks of 1: n protection rather than 1:1 protection (the latter applies only to a back-up path which is not shared among multiple primary paths). Sharing a back-up path among n disjoint working paths can spare a large amount of bandwidth, but at the cost of an increased amount of signalling. On the contrary, having a dedicated path requires the reservation of many more resources, but requires less signalling. The gain in spatial reuse of 1:1 protection schemes over 1+1 is already important for rings [51], but gets even much larger in meshed WDM networks.

Protection around the failed facility can be done at different points in the network: either around the two end-points of the the failed link, by line protection; or between the source and destination of each connection traversing the failed link, by path protection. Protection at the line layer is simpler, but path protection requires less bandwidth and can better handle node failures.

Routing and assigning wavelength in an optical network to guarantee its survivability by either 1+1 or 1:1 protection can be formulated as an Integer Linear Programming (ILP) problem. Ramamurthy and Mukherjee [52] use the ILP formulation to compare quantitatively the two schemes, together with the variants of link and path protection. The 1:1 path protection leads to significant savings in capacity utilization over the 1:1 link and 1+1 protection schemes. Since for large topologies the ILP formulation approach becomes computationally difficult, a number of heuristics have been proposed [53].

Protection is the most common mechanism deployed at the optical layer, because WDM or SONET/SDH connections are usually long-lived, and rarely set up on demand. Some authors advocate the possibility of restoration at the optical layer, which would spare more bandwidth than protection, but can also introduce significant delays to restore the traffic [54]. The complexity of restoring traffic at the optical layer (compared to protection at same layer, or restoration at the IP layer) makes it unlikely that operators rely primarily on restoration at the optical layer in the near future.

3.2 Restoration at the Logical Layer by Survivable Mapping

Recall from the beginning of Section 3, that in order to make the IP restoration work, the logical topology must remain connected after a failure. This requirement can be met by an appropriate mapping of the logical topology on the physical topology.² More specifically, if the logical topology remains connected after any single physical link failure, then the underlying mapping is called a *survivable mapping*.

Although the survivable mapping problem can be viewed as a specific version of the Routing and Wavelength Assignment (RWA) problem, it is often defined relaxing some basic assumptions of RWA, such as the wavelength continuity or even the capacity constraints. This results in a survivable mapping problem that is independent of RWA and can be addressed separately.

The problem of finding survivable mapping is NP-complete [55] and has drawn recently a lot of attention. It was first identified by Crochat et.al. [56], and named “design protection”. Some authors focused on simplified versions of the survivable mapping problem, assuming a cycle (ring) topology at the physical layer [57, 58] or the logical layer [55, 59]. The others addressed the general case, with arbitrary topologies at both layers. In general, the existing approaches can be divided into three groups: (i) exact algorithms based on Integer Linear Programming (ILP), (ii) pure heuristics and (iii) heuristics with provable properties. Below we describe each of them in more details, and compare in Table 3.

ILP. The ILP solutions can be found for example in [55, 44, 60]. In [55] it was observed that a mapping is survivable if and only if no physical link is shared by all logical links belonging to a cut-set of the logical graph.³ This observation is used in [55] to formulate an ILP model for the survivable mapping problem: for each logical link and for each cut-set of the logical graph, a constraint is added to the ILP. This leads to exact solutions, but also to excessive run-times [61] for networks of a non-trivially small size (few tens of nodes). To overcome this difficulty two relaxations to ILP are proposed in [55], by including only cut-sets of small sizes. This considerably accelerates the algorithm, but can easily lead to suboptimal solutions. Facing the same time-complexity problem of ILP, the authors of [44] and [60] decided to try a heuristic approach.

Heuristics. Despite many differences, the heuristics used to solve the survivable mapping problem share the same general methodology. They start with some initial mapping (e.g., shortest path) and try to improve it at subsequent iterations. Probably the most often used heuristic is *Tabu Search*. It is a version of a steepest descent search algorithm that stores a list (called a Tabu List) of recent moves to avoid them. This allows Tabu Search escape the local minima.

² We assume that the logical and the physical topologies are given and cannot be changed.

³ A *cut-set* of a network is defined by a cut of the network: a cut is a partition of the set of nodes V into two sets S and $V - S$, and the cut-set defined by this cut is the set of edges which have one endpoint in S and one in $V - S$.

Table 3. Comparison of efficiency and functionalities of four approaches to search for a survivable mapping. The question mark “?” means that the option might be possible to realize, but, to the best of our knowledge, nobody did it to date.

<i>Functionality</i>	<i>ILP</i>	<i>Tabu Search</i>	<i>FastSurv</i>	<i>SMART</i>
<i>Scalability</i>	<i>Low</i>	<i>Average</i>	<i>High</i>	<i>Very high</i>
<i>Capacity and other constraints</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>No</i>
<i>Verification of a solution existence</i>	<i>Yes</i>	<i>No</i>	<i>No</i>	<i>Yes</i>
<i>Node/span/multiple failures</i>	<i>?</i>	<i>?</i>	<i>Yes</i>	<i>Yes</i>
<i>Tracing and repairing the vulnerable areas</i>	<i>No</i>	<i>No</i>	<i>No</i>	<i>Yes</i>

For more details refer to [62]. Tabu Search was used with success to solve the survivability problem in many settings, e.g., without capacity constraints [56], with capacity constraints [63, 64] or additionally meeting maximum delay requirements [60]. Another general heuristic applied to solve the survivable mapping problem is Simulated Annealing in [48]. There is also a number of heuristics developed specifically to solve this problem, e.g., in [44] and [65]. The FastSurv algorithm introduced recently in [65], exploits the observation already mentioned in the ILP paragraph, which takes use of cut-sets in the logical topology. However, unlike in [55], the FastSurv algorithm systematically and indirectly learns about the importance of particular cut-sets and focuses only on the most relevant ones. This approach results in much better efficiency and scalability than those of other heuristics.

Heuristics with Provable Properties. The SMART algorithm proposed in [66, 67] does not fall in either group above. It is based on a breakdown of the problem into a set of independent smaller problems, which are easy to solve. Each of them is solved separately, and then the solutions are combined to obtain a survivable mapping of the entire topology. This makes SMART the fastest and most scalable heuristic to date. Moreover, the formal analysis in [67] revealed that SMART can also serve as a scalable method of verification of the existence of a survivable mapping and a tool tracing and repairing the vulnerable areas of the network. These two features are completely novel in the field.⁴ It should be noted, however, that one of the main assumptions of the analysis in [67] is relaxing the capacity constraints. In the presence of some additional real-life constraints such as limited fiber capacity or maximum delay, the SMART approach loses its efficiency and properties. Therefore SMART is more used to getting some topological insight into the problem than to finding an engineering solution, which makes this approach in a sense complementary to others.

3.3 Other Types of Failures

So far we have only considered single physical link failures. They may result from a fiber cut, a fault of a single interface card in the optical switch, or a fault of an

⁴ The ILP can also verify the existence of a survivable mapping, but as we argued before, it is not scalable.

optical amplifier. They are the most common type of failures in optical networks, but not the only one. If we allow for the physical location of the fibers, we extend single link failures to single *span failures*. A *span* is a bundle of fibers partially placed together for cost reasons (e.g., along railway and electricity lines). A single cut can break all of these fibers at once, in which case we speak of a span failure. We can also encounter *node failures*; they are the consequence of a failure of equipment at nodes, such as switches. In our context a node failure is equivalent to a failure of all physical links neighboring to the node. Finally, we consider *double-link failures*, i.e., independent failures of any two physical links. Usually such a situation takes place when the second failure occurs before the first one is repaired. This is not very common, but possible. For example, in the Sprint network, the time between two successive optical failures ranges from 5.5 sec to 7.5 days with a mean of 12 hours [1]. Most of them are repaired automatically within several minutes, but those requiring human intervention (e.g., after a fiber cut) may last hours or days. It is quite probable that during that period another physical failure occurs.

These failure scenarios were addressed mainly by physical layer protection: the span failures in [68,69], the node failures in [70], and the double-link failures in [71, 72, 73]. The IP restoration mechanisms considered these failures in [66] (all types of failures) and [67] (link and node failures).

4 Conclusion

We have addressed the failure management problem in IP/WDM optical networks. This issue can be decomposed with respect to two criteria. First, we distinguish the failure location from the failure restoration. The former aims at identifying the failing component based on the feedback from the network, whereas the latter consists in rerouting the traffic affected by the failure. These two tasks have different objectives and require different approaches. The second line of division is defined by the existence of at least two layers in the network: the IP layer and the optical layer. Each layer applies its own specific mechanisms to transport traffic, which significantly affects the way a failure is handled.

Following this view, we have discussed and made a detailed comparison of numerous failure management techniques, separately for failure location and restoration, and distinguishing between the IP and the optical layer. In contrast to previous surveys that have focused only on some particular aspects, our approach results in a global overview of failure management possibilities in IP/WDM networks.

References

1. Markopoulou, A., Iannaccone, G., Bhattacharyya, S., Chuah, C.N., Diot, C.: Characterization of Failures in an IP Backbone. In: Proceedings of the IEEE INFOCOM'04. (2004)
2. Abek, F., Hegerin, H., Neumair, B.: Integrated Management of Networked Systems. Morgan Kaufmann Publishers (1998)

3. Mas, C., Thiran, P.: An efficient algorithm for locating soft and hard failures in WDM network. JSAC special issue on Protocols and Architectures for next generation optical WDM networks **18** (2000) 1900–1911
4. Mas, C., Nguyen, H.X., Thiran, P.: Failure location in WDM networks. In: Optical WDM Networks: Past Lessons and Path Ahead. Kluwer Academic Publishers (2004)
5. ITU-T COM 15 121: Signal Quality Monitoring in Optical networks. (1999)
6. Anritsu: Catalog of measuring instrument (1993)
7. ITU-T Rec. G.872. Architecture of Optical Transport Networks (1998)
8. ITU-T Rec. G.806. Characteristics of Transport Equipment - Description Methodology and Generic Functionality (2000)
9. Wautersa, N., Ocachoglu, G., Struyve, K., Falcao, P.: Survivability in a new pan-european carrier's network based on WDM and SDH technology: Current implementations and future requirements. IEEE Communication Magazine **37(8)** (1999) 63–69
10. Tao, W., Somani, A.K.: Attack monitoring and monitor placement in all-optical networks. In: Proceedings of IEEE GBN 2001. (2001)
11. Kilper, D., Bach, R., Blumenthal, D.J., Einstein, D., Landolsi, T., Ostar, L., Preiss, M., Willner, A.E.: Optical performance monitoring. Journal of Lightwave Technology **22** (Jan 2004) 294–304
12. N.S.V.Rao: Computational complexity issues in operative dianosis of graph based systems. IEEE Transactions on Computers **42** (1993) 447–457
13. Nguyen, H.X., Thiran, P.: Failure location in all optical networks: the assymetry between false and missing alarms. In: Proceedings of ITC 19. (2005)
14. Ducatelle, F., Gambardella, L.M., Kurant, M., Nguyen, H.X., Thiran, P.: Algorithms for Failure Protection in Large IP-over-Fiber and Wireless Ad Hoc Networks. In Dependable Systems: Software, Computing, Networks, eds. J. Kohlas, B. Meyer, A. Schiper, Lecture Notes in Computer Science 4028, Springer, 2006 (this volume)
15. Katzela, I., Schwartz, M.: Scheme for fault identification in communication networks. IEEE/ACM Transaction on Networking **3** (1995)
16. Wang, C., Schwart, M.: Identification of faulty links in dynamics-routed networks. IEEE Journal on selected Areas in Communications (1993) 1449–1460
17. Li, C.S., Ramaswami, R.: Fault Detection and Isolation in transparent All-Optical Networks. In: IBM Research Report. Volume RC-20028. (1995)
18. Bouloutas, A., Hart, G., Schwartz, M.: Fault identification using a fsm model with unreliable partially observed data sequences. IEEE Transactions on Communications **41** (1993) 1074–1083
19. Gu, K., et al.: Realization of an expert system for an online fault diagnosis and restoration in a bulk power system. In: Proc. 4th International Symposium expert Systems Application Power Systems. (1993)
20. Brugnoli, S., et al.: An expert system for rel time fault diagnosis of the italian communications network. In: Proceedings of Integrated network management. Volume 3. (1993) 617–628
21. Jakobson, G., Weissman, M.E., Brenner, L., Lafond, C., Matheus, C.: Grace: Building next generation event correlation services. In: IEEE/IFIP: Network Operations and Management Symposium NOMS, 2000. (2000)
22. Lewis, L.: A case-based reasoning approach to the resolution of faults in communications networks. In Integrated network management III (1993) 671–682

23. Maki, Y., Loparo, K.A.: Neural network approach to fault detetin and diagnosis in industrial processes. *IEEE Transactions on Control Systems Technology* **5(6)** (2001) 529–541
24. Rodriguez, C., Rementeria, S., Martin, J., Lafuente, A., Perez, J.: A modular neural network approach to fault diagnosis. *IEEE Transactions on Neural Networks* (March 1996)
25. Ho, L., Cavuto, D., Papavassilou, S., Zawadzki, A.: Adaptive and automated detection of service anomalies in transaction-oriented wans. *IEEE Journal on Selected Areas Communications* **18(5)** (May 2000) 744–757
26. Jakobson, G., Weissman, M.E.: Alarm correlation. *IEEE Network* (1993) 52–59
27. Hood, C., Ji, C.: Proactive network-fault detection. *IEEE Transactions on reliability* **46(3)** (Sep 2000)
28. Lin, A.: A hybrid approach to fault diagnosis in network and system management. HP Technical Report (1998)
29. Gardner, R., Harle, D.: Alarm correlation and nerwork fault resolution using kohonen self-organising map. In: In proceedings of Globecom 97. (1997) 1398–1402
30. Stallings, W.: *SNMP, SNMPv2, SNMPv3 and RMON 1 and 2*. Addison-Wesley Longman Inc (1999)
31. Zhang, Y., Breslau, L., Paxson, V., Shenker, S.: On the characteristics and origins of internet flow rates. In: Proceedings of the ACM SIGCOMM Conference. (2002)
32. Choi, B.Y., Moon, S., Zhang, Z.L., Papagiannaki, K., Diot, C.: Analysis of point-to-point packet delay in an operatinal network. In: Proceedings of the INFOCOM. (2004)
33. Paxson, V.: *Measurement and Analysis of End-to-End Internet Dynamics*. PhD thesis, Univ. of Cal., Berkeley (1997)
34. Almes, G., Kalidini, S., Zekauskas, M.: A one-way delay metric for IPPM. IETF, IP Performance metrics, request for comments:2680 (1999)
35. ITU-T Rec. G.113. [G.113 Appendix I (05/02)] Provisional planning values for the equipment impairment factor I_e and packet-loss robustness factor B_{pl} (2002)
36. Mathis, M., Semke, J., Mahdavi, J., Ott, T.: The macroscopic behaviour of the TCP congestion avoidance algorithm. *Computer Communication Review* **27** (1997)
37. Tao, V., Xu, K., Estepa, A., Fei, T., Gao, L., Guerin, R., Kurose, J., Towsley, D., Zhang, Z.L.: Improving voip quality through path switching. In: Proceedings of IEEE Infocom. (March 2005)
38. Duffield, N.: Simple network perormance tomography. In: Proceedings of the IMC'03, Miami Beach, Florida (2003)
39. Caceres, R., Duffield, N.G., Horowitz, J., Towsley, D.: Multicast-based inference of network-internal loss characteristics. *IEEE Transactions on Information Theory* **45** (1999) 2462–2480
40. Coates, M., Nowak, R.: Network loss inference using unicast end-to-end measurement. In: Proceedings of the ITC Seminar on IP Traffic, Measurements and Modelling, Monterey (2000)
41. Duffield, N., Presti, F.L., Paxson, V., Towsley, D.: Inferring link loss using striped unicast probes. In: Proceedings of the IEEE Infocom 2001, Alaska (2001)
42. Padmanabhan, V.N., Qiu, L., Wang, H.J.: Server-based inference of internet performance. In: Proceedings of the IEEE INFOCOM'03, San Francisco, CA (2003)
43. Zang, H., Jue, J.P., Mukherjee, B.: A review of routing and wavelength assignment approaches for wavelength-routed optical wdm networks. *SPIE Optical Networks Magazine* (1) 47–60

44. Sahasrabudde, L., Ramamurthy, S., Mukherjee, B.: Fault management in IP-Over-WDM Networks: WDM Protection vs. IP Restoration. *IEEE Journal on Selected Areas in Communications* **20** (2002) 21–33
45. Demeester, P., et al.: Resilience in multilayer networks. *IEEE Communications Magazine* (August 1999) 70–75
46. Colle, D., et al.: Data-centric optical networks and their survivability. *IEEE Journal on Selected Areas in Communications* **20** (2002) 6–20
47. Zhang, H., Durresi, A.: Differentiated Multi-Layer Survivability in IP/WDM Networks. *Proceeding of Network Operations and Management Symposium* (2002)
48. Fumagalli, A., Valcarenghi, L.: IP Restoration vs. WDM Protection: Is There an Optimal Choice? *IEEE Network* (2000)
49. Iannaccone, G., Chuah, C.N., Bhattacharyya, S., Diot, C.: Feasibility of IP restoration in a tier-1 backbone. (Sprint ATL Research Report Nr. RR03-ATL-030666)
50. Nucci, A., Taft, N., Barakat, C., Thiran, P.: Controlled use of excess backbone bandwidth for providing new services in IP-over-WDM networks. *IEEE Journal on Selected Areas in Communications* **JSAC-22** (2004) 1692–1707
51. Gerstel, O., Ramaswami, R.: Optical Layer Survivability-An Implementation Perspective. *IEEE Journal on Selected Areas in Communications* **18** (2000) 1885–1923
52. Ramamurthy, S., Mukherjee, B.: Survivable WDM mesh networks, Part I - Protection. *Proc. of IEEE INFOCOM'99* (1999)
53. Mohan, G., Somani, A.K.: Routing dependable connections with specified failure restoration guarantess in WDM networks. *Proc. of IEEE INFOCOM'02* (2002)
54. Ramamurthy, S., Mukherjee, B.: Survivable WDM mesh networks, Part II - Restoration. *Proc. of IEEE ICC'99* (1999)
55. Modiano, E., Narula-Tam, A.: Survivable lightpath routing: a new approach to the design of WDM-based networks. *IEEE Journal on Selected Areas in Communications* **20** (2002) 800–809
56. Armitage, J., Crochat, O., Boudec, J.Y.L.: Design of a Survivable WDM Photonic Network. *Proceedings of IEEE INFOCOM 97* (1997)
57. Lee, H., Choi, H., Subramaniam, S., Choi, H.A.: Survival Embedding of Logical Topology in WDM Ring Networks. *Information Sciences : An International Journal, Special Issue on Photonics, Networking and Computing* (2002)
58. Lee, H., Choi, H., Choi, H.A.: Restoration in IP over WDM optical networks. In *Proceedings of the 30th ICPP Workshop on Optical Networks* (2001)
59. Sen, A., Hao, B., Shen, B., Lin, G.: Survivable routing in WDM networks logical ring in arbitrary physical topology. *Proceedings of the IEEE International Communication Conference ICC02* (2002)
60. Giroire, F., Nucci, A., Taft, N., Diot, C.: Increasing the Robustness of IP Backbones in the Absence of Optical Level Protection. *Proc. of IEEE INFOCOM 2003* (2003)
61. Leonardi, E., Mellia, M., Marsan, M.A.: Algorithms for the Logical Topology Design in WDM All-Optical Networks. *Optical Networks Magazine* (2000)
62. Glover, F., Taillard, E., Werra, D.: A user's guide for tabu search. *Annals of Operations Research* (1993) 3–28
63. Crochat, O., Boudec, J.Y.L.: Design Protection for WDM Optical Networks. *IEEE Journal of Selected Areas in Communication* **16** (1998) 1158–1165
64. Nucci, A., Sansò, B., Crainic, T., Leonardi, E., Marsan, M.A.: Design of Fault-Tolerant Logical Topologies in Wavelength-Routed Optical IP Networks. *Proc. of IEEE Globecom 2001* (2001)
65. Ducatelle, F., Gambardella, L.: Survivable routing in ip-over-wdm networks: An efficient and scalable local search algorithm. *Optical Switching and Networking* (2005) To appear.

66. Kurant, M., Thiran, P.: Survivable Mapping Algorithm by Ring Trimming (SMART) for large IP-over-WDM networks. Proc. of BroadNets 2004 (2004)
67. Kurant, M., Thiran, P.: On survivable routing of mesh topologies in IP-over-WDM networks. Proc. of IEEE INFOCOM'05 (2005)
68. Li, G., Doverspike, B., Kalmanek, C.: Fiber Span Failure Protection in Mesh Optical Networks. Optical Networks Magazine **3** (2002) 21–31
69. Zang, H., Ou, C., Mukherjee, B.: Path-protection routing and wavelength-assignment (rwa) in wdm mesh networks under duct-layer constraints. IEEE/ACM Transactions on Networking **11** (2003) 248–258
70. Kim, S., Lumetta, S.: Addressing node failures in all-optical networks. Journal of Optical Networking **1** (2002) 154–163
71. Choi, H., Subramaniam, S., Choi, H.A.: On Double-Link Failure Recovery in WDM Optical Networks. Proc. of IEEE INFOCOM'02 (2002)
72. He, W., Sridharan, M., Somani, A.K.: Capacity Optimization for Surviving Double-Link Failures in Mesh-Restorable Optical Networks. Proc. of OptiComm'02 (2002)
73. Clouqueur, M., Grover, W.D.: Mesh-restorable Networks with Complete Dual-failure Restorability and with Selectively Enhanced Dual-failure Restorability Properties,. Proc. of OptiComm'02 (2002)