Fault localization at the WDM layer

Carmen Mas (carmen.mas@epfl.ch), Patrick Thiran (patrick.thiran@epfl.ch)
and Jean-Yves Le Boudec (leboudec@epfl.ch)
Institute for computer Communications and Applications (ICA)
EPFL (Swiss Federal Institute of Technology)
CH-1015 Lausanne, Switzerland
Telephone:+41-21-6935260, Fax:+41-21-6936610

Abstract. A single failure in a communication network may trigger many alarms. When the communication network uses optical fibres as transmission medium and increases its capacity by using Wavelength Division Multiplexing (WDM), the number of alarms and the difficulty to locate the failure are considerably higher. In this case, a single failure may interrupt several channels which causes a large information loss.

We propose an Alarm Filtering Algorithm (AFA) for the fault management of an optical network that supports multiple failures and works in the presence of *non-alarming* elements, that is, network components which may fail but never generate an alarm (e.g. optical fibres). The algorithm provides a list of components whose failure explain the observed alarms. It avoids the use of failure probabilities, which are difficult to estimate, and does not need a global knowledge of the network topology. Moreover it also tolerates alarm losses and false alarms.

The algorithm is tailored to the specific behaviour of the hardware components of an optical network when a failure occurs. The classification of the network components according to the alarm signals they generate enables a formalisation of the alarm-filtering problem and results in an efficient algorithm for localising the failure(s). This algorithm is applied to the WDM rings of the COBNET network [1] (COBNET is a European ACTS project) and to a meshed optical network with the ARPA2 topology.

Keywords: Alarm filtering, optical communication networks, fault diagnosis

1. Introduction

Communication networks have been developing rapidly during the last few years. The evolution of the transmission technology and the rapid increase in user needs have influenced major changes in networks. One new technology is the use of optical fibre as transmission medium, which allows high transmission rates because of their low attenuation and high capacity. The capabilities of optical fibre are increased by Wavelength Division Multiplexing (WDM) that makes possible the transmission of several information channels through a single link so that the capacity of the link can become considerably higher. Nowadays, it is possible to find market equipment that supports from 80 to 128 wavelengths on a fibre, each fibre carrying up to 10 Gbit/s of information [9]. In a near future, this equipment may support up to one thousand wavelengths. But the capacity advantage comes however with some drawbacks if a failure occurs. This first one is that a failure can now result in the loss of a large amount of information. The second one is that the network management application is overwhelmed by a large number of alarms due to a single failure involving several channels. For these reasons a fast fault recognition process is needed and is critical in ensuring network availability. It is composed of two phases: alarm filtering and testing. The output of the first phase is the list of network components that might have failed and triggered the alarms received by the manager. The second phase confirms or

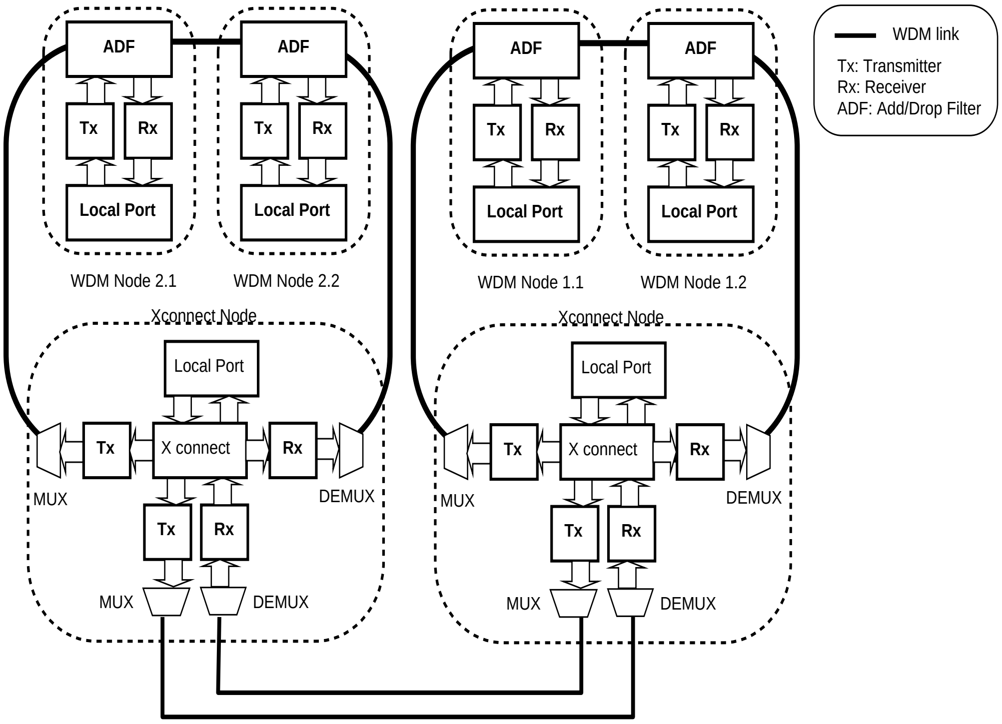

Figure 1. Multi-hop WDM network with ring topology

discards the previous result by checking these components. This testing part is outside the scope of this paper, which is devoted to the alarm processing phase only.

The problem of identifying failures is closely related to the physical layer. Upper layers react on physical failures in different ways in order to re-establish the communication (for example, by rerouting or changing switch positions). But even if the connection is successfully re-established, the physical layer management must localize the failure. This should be based on the alarms emitted by the hardware components and the knowledge of the components' behaviour. As the physical layer is migrating from electronics to optics new kinds of network components are appearing on the market. Each component behaves differently when a failure occurs, and the information that is sent to the manager or that the manager is able to retrieve is being standardised. This article considers the problem of localising hardware failures at the optical layer, in particular in WDM networks.

In this paper we consider a multi-hop optical network using WDM techniques where several channels are established. Multi-hop networks (in particular the one considered in this paper that is presented in Figure 1) use fixed-wavelength transmitters and receivers. When a channel is established between two nodes, it is the electro-optical switch at the WDM cross-connect node that converts the signal at the sender's wavelength into an electrical signal and pack it back into an optical signal at the receiver's wavelength. When a network component fails, all channels passing through this element are interrupted. Consequently, all the elements that were involved in the interrupted channels and that are able to send alarms to the manager will report him a problem. The messages from these elements will be different and the manager will have to determine where the failure is. This is more complex when multiple failures occur almost simultaneously. In this case, alarms due to different failures will reach the management application during the same period of time and they will intermingle. The Alarm Filtering Algorithm (AFA) can deal with the failure(s) of any component. As we will see later, to filter more alarms, we will sometimes use the (reasonable) assumption that a number n of failures is more probable that a number $n+1$ of failures.

The goal of the alarm filtering becomes more difficult to achieve under non ideal conditions. Two abnormal cases can arise: existence of (1) missing alarms (alarms that do not reach the manager) and (2) false or unexpected alarms. The first case occurs when there are alarms that are lost or that will arrive with such a delay that they cannot be considered during the ongoing computation of the algorithm. The second case occurs when, due to an abnormal situation, an element sends an alarm but there is no real failure. These alarms are not caused by a malfunction but by an exceptional problem, as for example, an incorrectly measured value. To handle these cases, a *mismatch threshold* will be introduced in Subsection 3.3.

Different methods have been proposed for alarm filtering in the literature. They differ on the network model that is used (e.g. model based on the physical topology [3] or on the description of the established channels [8]; on the information needed (e. g. probabilities [6]) as input for the algorithm; and on the information processing methodology to localise the failures (e. g. Finite State Machine (FSM) [5] or Artificial Neural Networks (ANN) [7]).

Our Alarm Filtering Algorithm (AFA) allows the location of multiple failures and the identification of failures in network components that are unable to communicate with the manager in the non-ideal scenario which allows the existence of lost and false alarms. The input of the algorithm is the list of established channels in the network and of the alarms received by the manager; the network topology is not needed. The output of the algorithm is a list of network components that may have failed.

The article is structured as follows. Section 2 defines the taxonomy of the network components considered in this work. Section 3 formalises mathematically the problem defined above and describes the AFA. Section 4 applies the AFA to different network topologies: a WDM meshed network and a WDM ring network.

2. A classification of optical network components in terms of their alarm processing properties

Walrand [2] defines communication networks as a set of nodes that are interconnected to allow the exchange of information. In this case, the term node can refer to bridges, switches and access ports. Such a definition is well suited for traffic control but is not very helpful for fault management because the view of the network is different. Fault management needs to know which are the network components, what are the alarms and their content and in which situation these alarms are sent. A model of optical communication networks used to demonstrate AFA is presented in this section. This model is based on channel relations between network components that may fail.

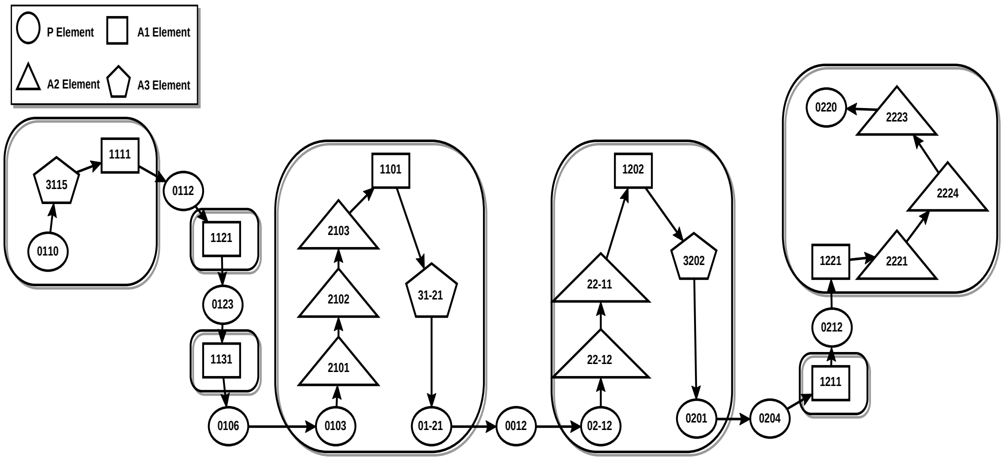

Figure 2. Channel between two WDM rings

2.1. Channel Relation

Channels are established between nodes that want to interchange information and can be unidirectional or bi-directional. In our model we will consider channels to be unidirectional and a bi-directional channel can always be considered as two unidirectional channels.

Unidirectional channels are ordered sets of network components. They correspond to the Telecommunication Management Network (TMN) M.3100 [10] *trail* concept when the *Trail termination Source point* is the sender access port and the *Trail termination Sink point* is the receiver access port. For example, the channel presented in Figure 2 is an ordered set of 30 network components where the first one is (0 1 1 0), the 14th one is (3 1 -2 1) and the last one is (0 2 2 0). The explanation about the component identifiers and their classification as P, A1, A2 and A3 elements will be given in Subsection 2.3.

2.2. Alarms

Alarms are messages sent to the manager by the components of a network informing of an abnormal condition (e.g. some parameter of a component out of range, or a missing signal). As we will see in detail in the next subsection, only the so-called ’alarming’ components can send alarms to the manager. The content of these messages depend on the component [3] but always contains the identification of the component that sent the alarm, the nature of the alarm, the abnormal value of the parameter and/or the time when the alarm was sent (timestamp). The timestamp can be useful in localising the failure but it can also lead to errors when there are multiple failures and has therefore not been considered in our analysis. Our algorithm needs only the minimal amount of information present in all the alarms, that is, the alarm origin and its nature.

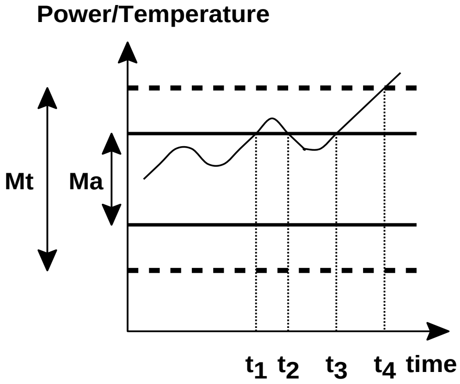

Figure 3. Transmitter Margins

2.3. Network Components

In this subsection we introduce the network components that belong to an optical network. Based on their behaviour when a failure occurs, we have classified them into four different categories which form the basis of the algorithm.

2.3.1. *Hardware components of an optical network*

The common hardware components that can be found in optical networks are:

– Optical fibres: They are the medium for transmitting optical signals between two points. The fibres have a huge transmission capacity, low attenuation and low cost. One of their advantages is that the same fibres are able to transmit several information channels, each channel being modulated at a different wavelength.

– Transmitters: In optical networks, transmitters are lasers or laser arrays that convert the electrical signal into an optical one 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. The new lasers used in advanced WDM networks are tunable and can change the emission wavelength within a prescribed range [4].

  Transmitters send alarms when either the temperature or the incoming power is beyond a prescribed range. In fact, for each variable (temperature or power) there are two ranges (see Figure 3): the first one, *Ma*, delimits the values for which the transmitter works correctly. Once the upper or lower margin is crossed by either the incoming power or the temperature, an alarm is sent informing that the corresponding variable is too high/too low, and hence that the emitted signal may be incorrect. When the larger range *Mt* is exceeded, not only a new alarm is sent but the transmitter is turned off to prevent it to cause any damage to the network. In the example of Figure 3, at time t_1 an alarm is sent informing that *Ma* is exceeded. Due to local temperature or power control of the board, the value is back to the expected range at t_2 and another alarm is sent to cancel the first alarm. Some time later, at t_3, *Ma* is again exceeded and a new alarm is sent. At t_4 the range *Mt* is also exceeded, a new alarm is sent and the transmitter is turned off

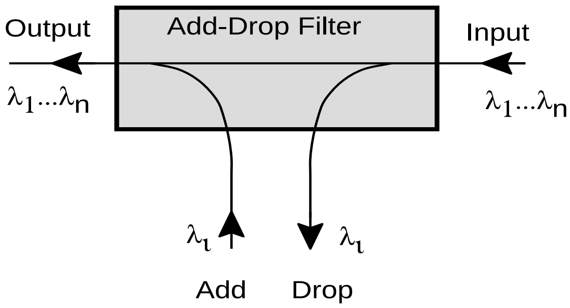

Figure 4. Add/Drop Filter

(because the local control mechanism of the board has not been able to restore the normal behaviour).

- Receivers: They convert an optical signal, which corresponds to a certain wavelength into an electrical signal. Some receivers can only convert a fixed wavelength, but tunable receivers are able to change the wavelength within a given range. The latter ones have been developed for advanced WDM networks, where wavelengths are reallocated.

- Add-Drop Filters (ADFs): They are able to drop and add a certain wavelength to an optical signal with several wavelengths without distorting the other wavelengths (see Figure 4). These components are used in networks where each node has to retrieve the data addressed to itself and modulated with its wavelength from a WDM signal. For example, in Figure 1, if there is a connection established between WDM Node 1.1 and Node 1.2, they will send and receive the information at respectively wavelengths λ_1 and λ_2 because their ADF will work at these wavelengths. The WDM cross-connect node will associate the input λ_1 to the output λ_2 and vice versa. In future optical networks, not only one but several wavelengths will be added/dropped, making channel routing possible. This component is still under research and it is called Add/Drop Multiplexer.

- 3Rs (Re-generator/Re-shaper/Re-timing): These components are able to amplify the electrical signal, give the original shape of the signal and readjust the time interval between pulses.

- Protection switches(PSs): These components receive more than one optical signal and select one of them having an acceptable power level. The criterion to choose which input to take depends on the switch implementation. For simplicity, the input to be chosen is given by default and in case this input does not reach an acceptable power level, the switch changes its position to take another input whose power is within the acceptable range.

- Multiplexers (MUXs) and Demultiplexers (DEMUXs): They are able to pass from several optical signals at different wavelengths into one optical signal that contains all the wavelengths and vice versa. The MUX composes the WDM signal from several optical signals and the DEMUX converts the WDM signal back into several single optical channels.

– Switches: These are the components that allow cross-connection, that is, the association of a particular input with a particular output. In single-hop networks, they transfer an optical input to an optical output without any electrical conversion. In multi-hop networks, the switches are opto-electrical or electrical.

2.3.2. *Alarming Properties of hardware components*

network components are composed of the components of which each may fail and each can have a different behaviour when there is a failure. Three features can be distinguished:

1. **Self-alarmed**: This property specifies whether a network component is able to send an alarm to the manager informing him about its own failure or not. This alarm corresponds to the *Equipment alarm* specified in X.721 [11]. The capacity of sending alarms is related to the existence of a micro-controller, i.e. of some software that controls the board and communicates with the manager via a management protocol such as SNMP (Simple Network Management Protocol) or CMIP (Common Management Information Protocol).
   An example of a self-alarmed component is a transmitter. This component usually has a micro-controller that controls its power and its temperature and sends alarms when the functioning is not the regular one, that is, when these values exceed the allowed ranges, as explained above.
   If there is no micro-controller, it is not possible for the manager to be informed about the hardware status. This is, for example, the case of an optical fibre.

2. **Out-alarmed:** This property applies for the components that communicate with the manager and send alarms about a problem external to them. These alarms corresponds to the *communication and environmental alarms* specified in X.721 [11]. For example, receivers are able to detect that there is no incoming power and send the correspondent alarm to the manager even if themselves are working correctly. On the contrary, multiplexers are not able to detect if some inputs are missing or not, therefore they are not out-alarmed components.

3. **Failure masking**: This property specifies whether the network component masks or not the failure to the components that follow it on the channel (remember that a channel is an ordered set of components as explained in Section 2.1). The laser of a transmitter sends power even if there is no incoming signal (due to a failure of some component located before the laser). Therefore, any out-alarmed component located after this transmitter on the channel will not send any alarm because it will keep receiving power, even if it does not receive data any more.

We can now check which one(s) of these three properties applies to each kind of optical component. This analysis is based on the capabilities of real optical components, in particular, on the components of the COBNET (Corporate Optical Backbone NETwork) optical network. The results of this analysis are summarised in Table I.

All the components have a micro-controller to communicate alarms to the manager, except optical fibres and (de)multiplexers, which are therefore unable to send any alarm in case of failure. Components sending alarms about their own failures are

Table I. Alarm properties of the Optical Components and the resulting classification

Network Component	Self-alarmed	Out-alarmed	Failure Masking	Category
Optical Fiber	No	No	No	P
Transmitters	Yes	No	Yes	A3
Receivers	No	Yes	No	A2
Add/Drop Filters	Yes	No	No	A1
3R	No	Yes	No	A2
Protection Switch	No	Yes	No	A2
MUX/DEMUX	No	No	No	P
Switch	Yes	No	No	A1

the transmitters (once either the power or the temperature is out of range), the ADFs (once they have an internal malfunctioning), the switches (they inform about a problem such as the impossibility to connect a certain input with a certain output). However, these components do not send any alarm if another component of the same channel has failed. These components are therefore self-alarmed but not out-alarmed. On the contrary, the remaining components listed in Subsection 2.3.1 are out-alarmed but not self-alarmed: receivers and 3Rs (they send alarm whenever the received incoming power is below a certain level) and protection switches (their alarm informs about a change of the switch position towards a better quality input due to the degradation of the default input). Finally, as we have seen above, a transmitter masks a failure that damaged a component located before it on the channel, to all the other components following the transmitter on the channel, but it is the only component to do so.

2.3.3. *Network Components Classification*

The properties of Table I enable us to classify all optical components in the following *categories*:

1. **'Non-alarming' components**: These are the components that do not give any information to the manager because they do not have any micro-controller. They are denoted my *P*.

2. **'Alarming' components**: These are the components that are able to communicate with the manager because they have a micro-controller. Within this group there are three kinds of *alarming* components:

   a) **A1 components**: these are the self-alarmed components that do not mask the failures.
   b) **A2 components**: these are the out-alarmed components.
   c) **A3 components**: these are the components that are self-alarmed and mask previous failures.

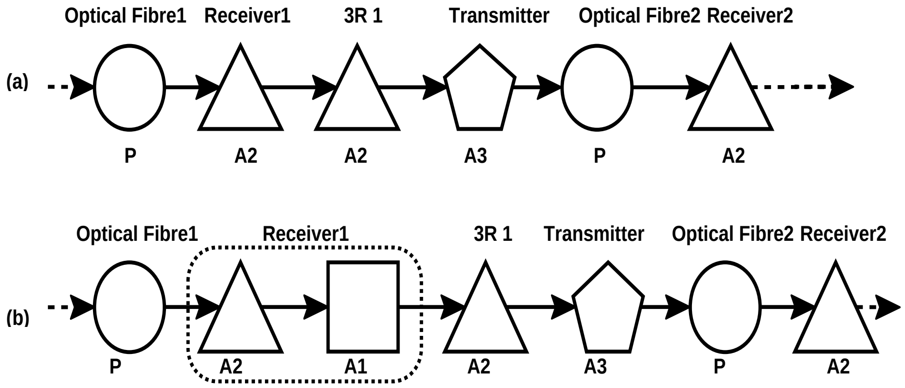

Figure 5. Modeling of a channel part

Comparing this classification with the TMN standard M.3100 [10], the *A3* elements corresponds to the *source* components. The *A2* elements may correspond to the *sink* components but not exclusively.
In this way, all the optical components have been classified as presented in the last column of Table I.

In Figures 2 and 5, we use different geometric forms to identify the different classes of components. We use circles for the *non-alarming* components, squares for *A1* components, triangles for *A2* elements and pentagons for *A3* components.

Note that in this classification, based on the actual behaviour of hardware components in networks such as the COBNET WDM network, there is no element which is at the same time self-alarmed and out-alarmed. This is however not a restriction for the alarm filtering algorithm described in Section 3. Indeed, whenever an optical component has both, the *A1* (or *A3*) and *A2* properties, it can be represented by a pair of elements: one *A1* (or *A3*) and one *A2*. For example, if a receiver is able to send alarms when it does not receive optical power and/or when it has an internal problem, it is self-alarmed and out-alarmed. In this case, the receiver will be considered as a set of two elements, one belonging to *A1* category and the second belonging to *A2* category. Figure 5(a) presents the case when *Receiver1* is only out-alarmed and Figure 5(b) presents the case when *Receiver1* is self and out-alarmed.

3. Problem abstraction

The classification of Section 2 enables us to derive and implement the Alarm Filtering Algorithm (AFA). Before describing the AFA itself, we first need the following definitions.

3.1. Classes

Three different classes have been defined:

– *Component e* is a network component, which is either *non-alarming* or *alarming* (see Section 2). Its identifier is a set of integers that indicate the location and the component type (this identifier is the one used in COBNET project). In our scenarios we have considered 4 integers to identify all the components in a unique manner. The first integer always gives the class of component: 0 for *non-alarming* components, 1 for an *A1* class *alarming* component, 2 for an *A2* class *alarming* component and 3 for an *A3* class *alarming* component. The other integers give the location of the component within the network topology and are specific to this topology (see Section 4 for further explanation on particular networks). This addressing issue is different to the TMN standard but the one chosen for COBNET due to its simplicity.

- *Channel c* is defined as an ordered list of components $\{e_i\}$.
- *Alarm a* is a pair $(origin, information)$ where *a.origin* is an *alarming* component and *a.information* is an integer that expresses the nature of alarm.

We now define the sets that will be used throughout this article. The set of all network components is denoted by *NC*, the set of established channels in the network by *C* and the set of alarms that the manager receives by *R*. Moreover, let *P* be the set of *non-alarming* components, A_1 be the set of *alarming* self-alarmed components of *A1* class, A_2 be the set of *alarming* components that send alarms due to external failures, and A_3 be the set of *alarming* self-alarmed components of *A3* class.

3.2. Functions

The two following functions will be used in the different AFA modules:

- As channels are ordered list of network components, one can define the position of a component *e* in a channel *c*, *Pos(e,c)*, as the index of its occurrence in the ordered list. If *e* is not present in the channel *Pos(e,c)* returns 0. In other words,

$$Pos(e,c) = \begin{cases} 0 & \text{if } \forall e_i \in c, e_i \neq e \\ i & \text{if } \exists e_i \in c, e_i = e \end{cases} \quad (1)$$

- Component e_1 has a *PassivePath* with Component e_2 when there exists at least one channel that contains both components and such that all components between them are of class *P* or *A1*. Mathematically this can be expressed by the Boolean relation:

$$e_1 \beta e_2 = \begin{cases} 1 & \text{if} \exists c \in C \text{such that} \begin{cases} 1 \leq Pos(e_1,c) \leq Pos(e_2,c) \text{ and} \\ \forall e \in NC \text{ with} \\ Pos(e_1,c) < Pos(e,c) < Pos(e_2,c),\ e \in P \cup A_1 \end{cases} \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

3.3. Alarm Filtering Algorithm scheme

The AFA has as inputs the established channels in the network $C = \{c_j\}$ and the received alarms by the manager $R = \{a_i\}$. These inputs are updated whenever a new event occurs: either a channel event, which updates *C*, or an alarm event, which updates

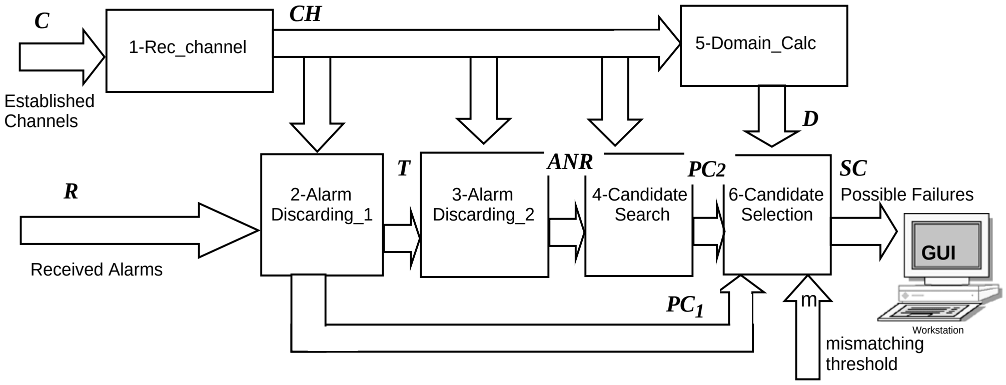

Figure 6. AFA scheme

R. By channel event we mean a modification of the set of established channels by the addition or deletion of a channel or by a change of some components involved in a channel. By alarm event we assume the of a new alarm arrival to the manager. The two kinds of events are independent.

The AFA has to deal with erroneous alarms, that is, false and missing alarms. The number of erroneous alarms is the *mismatching value*. There is therefore an additional input set a priori by the Human Manager: a *mismatch threshold*, denoted by m, which is an upper bound on the number of erroneous alarms (that is, false or missing alarms that are tolerated). m=0 corresponds to an ideal scenario where all the alarms are supposed to be correctly received. Depending on the status of the network, the manager can therefore choose a low value of m (which results in a small list of possible faulty elements but may miss a failure in case of corrupted alarms) or a large value of m (which yields a broader list of faulty elements but has more robustness against false and lost alarms).
The AFA is based on the discarding of some of the received alarms. This discarding phase is presented in two different modules: (2) *Alarm_Discarding_1* and (3) *Alarm_Discading_2*. The former filters, among others, the *A1* and *A3* alarms and the latter filters the *A2* alarms.
The output of the algorithm is a list of different sets of network components that may be faulty and that are an explanation of the observed symptoms with up to m erroneous alarms.

The AFA is structured in 6 modules, as shown in Figure 6, which are as follows:

1. ***Rec_Channel***: The input of this module is the set of channels presently established in the network whose set is denoted by *C*. This procedure delivers an updated set of ordered lists of network components (denoted by *CH*) that represent all the established channels after the latest channel event, and some of the channels before this event. This module corresponds to the configuration management functionality since it is called every time there is a channel event, that is, in any of these three cases: (i) when a new channel is established, (ii) when a channel is cleared down or (iii) when there is a change of the components that belong to the channel. An example of case (iii) occurs when due to a failure, the position of the protection switch is changed due to a failure. In this case, some network components stop participating in that channel and new ones start doing so.

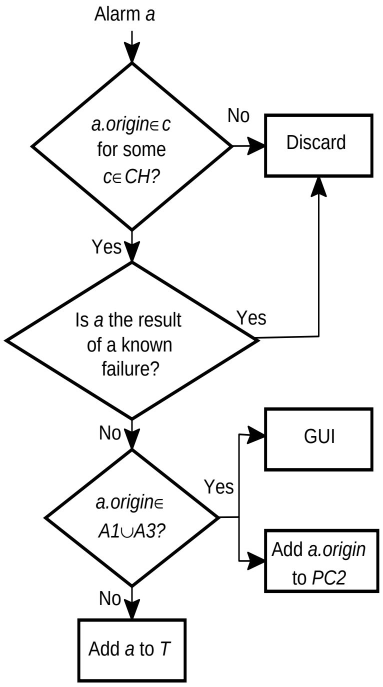

Figure 7. Alarm_Discarding_1 Module

This module performs differently for each of these cases. For case (i), it adds the ordered list of network components making the new channel to *CH*. In case (ii), it removes from *CH* the ordered list associated to the removed channel. For the last case (iii), it adds the ordered list related to the new path of the channel, as if it was a newly established channel, and it keeps the ordered list associated to the old path during a certain time. The reason of keeping the old path is to allow the AFA to localise failures that happened before the path adjustment (for example, the one that caused the change of the protection switch).

2. ***Alarm_Discarding_1***:This module is the first alarm discarding phase of the algorithm. It produces a subset T of alarms from the received alarms R, which contains the alarms having passed successfully the sequence of the three following tests, as shown in Figure 7:

   a) The alarm must be sent by a component belonging to one channel $c \in CH$. If this is not the case, the alarm is discarded because it has been generated by transient values reached by the network variables during the establishment of a channel.

   b) If the new alarm is consequence of an already known failure it is discarded because it is a delayed alarm. This can be performed by storing the missing alarms that correspond to the known failure and check whether the new alarm is one of them. In this case the use of timestamps can be useful.

   c) Is the alarm sent by a $A2$ component? Only in this case, the alarm will be added to T. The alarms sent by $A1$ and $A3$ components are forwarded to the Graphical

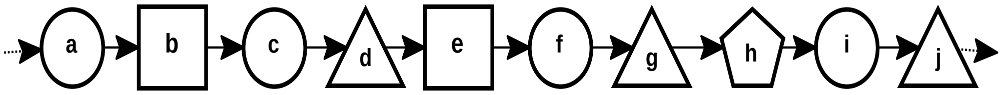

Figure 8. Channel example

User Interface (GUI) to be presented to the Human Manager and included in the set PC_1*(Possible_Candidates_1)*, which is input of the *Candidate Selection* module, so that the origin of the alarm is considered a likely faulty component:

$$PC_1 = \{e \in NC \,|\, \exists c \in CH \,\text{with}\, Pos(e,c) \neq 0 \,\text{and}\, \exists a \in R \,\text{with}\, e = a.origin\} \tag{3}$$

The *Alarm_Discarding_1* procedure is called each time a new event occurs. If the new event is a channel event, all the alarms *R* have to succeed the three sequential tests with the new set *CH* in case one of the previously discarded alarms do belong to the recently established channel. If the new event is an alarm event, the sequential conditions must be succeeded only by the new alarm in order for this one not to be filtered out as the other alarms are not changed.

3. ***Alarm_Discarding_2***: This module is the second discarding phase. It has as input *T*, the reduced set of received alarms obtained after *Alarm_Discarding_1*. Note that the set *T* contains only alarms originated by *A2* elements that belong to *CH* and are not the consequence of a known failure.
   Some of these alarms can be discarded thanks to the assumption that a number of *n* failures is always more likely than *n+1* multiple failures. Indeed, consider two contiguous out-alarmed elements e_1 and e_2 belonging to a channel *c* (i.e. $Pos(e_1,c)+1 = Pos(e_2,c)$). Suppose that an element *e*, with $Pos(e,c) < Pos(e_1,c)$ fails. If there is no *A3* element between *e* and e_1, both out-alarmed elements e_1 and e_2 send an alarm. We can disregard the alarm from e_2 because it does not give more information about the failure than the alarm from e_1. Indeed e_1 is closer to the failure than e_2. The same principle makes us discard the alarm sent by e_2 when these two out-alarmed elements are not contiguous but have *P* and *A1* between them. Consequently, when several alarms issued by out-alarmed elements having only *P* and *A1* elements between them are received, only the one whose origin has the smallest position in the channel will be retained.

   Let us illustrate on the channel example of Figure 8. If *c* fails, the components that will send alarms are *d* and *g*. In this case, the alarm from *g* will be discarded by the alarm from *d*. However, if for example *f* also fails, it will not be detected, only the failure from *c* will be localized. It is only once *c* is repaired, that as explained above, a failure in *f* can be detected. This situation of a double failure is assumed much less likely than a single one, and hence are preferred to discard the alarm from *g* which most probably is redundant.
   This second discarding phase results in a set of non redundant alarms issued by out-alarmed elements denoted by *ANR (Alarms_Non_Redundant)* and defined by:

$$ANR = \{a \in T \,|\, a.origin \in A_2 \,\text{and for all}\, a_0 \in T,\, a_0.origin \,\beta\, a.origin = 0\} \tag{4}$$

4. ***Candidate Search***: After having obtained the set of non redundant alarms *ANR* from out-alarmed elements, one searches for the network components of *CH* whose failure may have prompted these alarms, called *fault candidates*. Their set, PC_2 *(Possible_Candidates_2)*, may contain elements of all the categories.
An element *e* is *fault candidate* of an alarm *a*, when *e* may have caused *a*, that is, when the origin of the alarm is an *A2* component and has a *PassivePath* with *e*. Mathematically it can be expressed by

$$PC_2 = \{e \in NC \,|\, \exists c \in CH \, with \, Pos(e,c) \neq 0 \; and \; \exists a \in ANR \, with \; e \beta a.origin = 1\} \tag{5}$$

At this point, we have obtained the set PC=$PC_1 \cup PC_2$ containing all the components whose likely failures account for the observed alarms (either *A1* and *A3* alarms, sent by elements in PC_1, or *A2* alarms, sent because a failure of an element in *PC*). The final result, presented to the Human Manager is not the whole set of elements *PC* but a refinement of this set (which is denoted by *SC (Subset Candidates)* and is a set of *PC* subsets), for two reasons:

(i) In the ideal scenario where there are neither false nor lost alarms (*m*=0), one looks for the smallest subsets of elements of *PC* that explain all the received alarms *R*. One begins therefore with singletons: all those which account for all the alarms in *R* are included in *SC*. If none of them qualifies, pairs of elements of *PC* are then examined, and so on, until a subset of sufficient size explaining all the received alarms has been found. Then, all the subsets of that size that explain all the received alarms *R* are included in *SC*.
Let us consider the example of the channel presented at Figure 8. If there are only two received alarms and if they are issued by *d* and *g*, the elements included in PC_2 after the two discarding and the Candidate Search modules will be *a*, *b* and *c*. Therefore, since there is no alarm issued by a self-alarmed element, $PC = PC_1 \cup PC_2 = \emptyset \cup \{a,b,c\} = \{a,b,c\}$. As the failure of only one of these elements is considered much more likely than the simultaneous failure of two or three of them, while yielding the same alarms from *d* and *g*, the singletons *{a}*, *{b}*, *{c}* have to be first examined. In this case, singletons *{a}* and *{c}* satisfy the condition because they explain alarms issued by *d* and *g*. On the contrary, singleton *{b}* does not fulfill this condition because if it had failed, one would have had PC_1=*{b}*. Therefore, $SC = \{\{a\},\{c\}\}$.
Let us now consider the channels of Figure 9. Because a double failure, the only received alarms are issued by *d* and *g*. This scenario will cause that not only *a*, *b* and *c* are included in PC_2 but also *e* and *f*. As no singleton can explain both received alarms, pairs of the elements of PC_2 are then considered, so that $PC = PC_2 = \{a,b,c,e,f\}$ and clearly, $SC = \{\{a,f\},\{c,f\}\}$.

(ii) If $m > 0$, another task will be performed by the two last modules, in addition to **(i)**. In this scenario, the set *SC* will contain all subsets of elements of *PC* which yield a mismatching value less than or equal to *m*.
In the example of Figure 8, if the only received alarms come from *d* and *g*, $PC = \{a,b,c\}$. Assume that *m*=1, that is, one mismatching between the received alarms and the expected alarms is tolerated. Apart from the singletons *{a}*, *{c}* that result

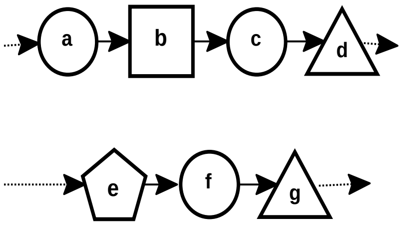

Figure 9. Double channel example

in perfect matching, as we have seen in **(i)**, singleton *{b}* is also included in *SC* because when it is faulty, not only alarms from *d* and *g* are expected but one from *b* itself. This scenario results in one mismatching that is within the mismatch threshold given. The resulting *SC* will be $SC = \{\{a\},\{b\},\{c\}\}$.

5. ***Domain_Calc***: This module is called every time *CH* changes. It computes, for each element of each channel belonging to *CH*, the set of *alarming* components that would send alarms if this element fails. This set is called *Domain(*e_0*)*, for a given e_0 that belongs to *CH*. The components that belong to *Domain* are e_0 itself if it has the self-alarmed property and the out-alarmed components located after e_0 in the channel but without any *A3* component between themselves and e_0. For example, consider the channel of Figure 9. If *Optical Fibre1* fails, *Receiver2* will not detect the failure because *Transmitter* is still sending optical power (even if it does not have any incoming data to modulate). In this case, *Domain(Optical Fibre1)={Receiver1 , 3R 1}* because *Receiver1* and *3R 1* are the two network components sending an alarm when *Optical Fibre1* fails.
   Mathematically, *Domain (*e_0) can be expressed by:
$$Domain(e_0) = \left\{ \begin{array}{ll} e \in A_2 \,| & \forall c \in CH,\ 1 \leq Pos(e_0,c) \leq Pos(e,c) \text{ and} \\ & \forall e' \in NC \text{ with } Pos(e_0,c) \leq Pos(e',c) \leq Pos(e,c),\ e' \notin A_3 \end{array} \right\} \quad (6)$$
   D stores the list of the *Domain* of all the components that belong to each channel of the set *CH*.

6. ***Candidate Selection***: This routine is called repeatedly as long as there is no updating of the entries. The inputs of this process are the mismatch threshold *m*, the output PC_2 of *Candidate Search* and the output PC_1 of *Alarm_Discarding_1*.
   Let us define R_{orig} as the set of network components (NC) that are the origins of the received alarms *R*:
$$R_{orig} = \{e \in NC \,|\, \exists a \in R \text{ with } a.origin = e\} \quad (7)$$
   (i) In the ideal scenario, there are no lost or false alarms, so the manager can choose *m=0*. The output *SC* is the set of subsets of elements whose failure will cause the same alarms as the ones received by the manager.
   Let $P = \{e_1, e_2, ..., e_n\}$ be any subset of *n* elements belonging to *PC* and whose domains form a union of R_{orig}, that is, whose domains verify equation (8).
$$R_{orig} = \bigcup_{i=1}^{n} Domain(e_i) \quad (8)$$

In this case, the output *SC* is the set of all subsets *P* for all possible values of n.

(ii) In the non ideal scenario, the output *SC* is the set of subsets of elements that when failing will cause a set of alarms which will differ from the received alarms by a mismatching value lower than the given threshold m.
Let $Q = \{e_1, e_2, ..., e_n\}$ be any subset of elements belonging to *PC* and such that:

$$\#\Big((UnionDom(Q) \setminus R_{orig}) \bigcup (R_{orig} \setminus UnionDom(Q))\Big) \leq m \qquad (9)$$

where # denotes cardinality and where $UnionDom(Q) = \bigcup_{i=1}^{n} Domain(e_i)$.
Then $UnionDom(Q) \setminus R_{orig}$ is the set of components which should have sent an alarm due to the failure of the *SC* components but whose alarms did not reach the manager (i.e. *lost alarms*) whereas $R_{orig} \setminus UnionDom(Q)$ is the set of components having sent an alarm but which cannot be explained by the failure of the *SC* components (i.e. *false alarms*).
In this case, the output *SC* is the set of all subsets Q for all possible values n.
Clearly, if m=0, (9) becomes (8).

3.4. AFA Result

The algorithm described in Section 3.3, has been implemented in Java. The algorithm inputs (channels and alarms) can be given either via file (from the management application) or via Input Frame (for testing purposes). The mismatch threshold m is set by the manager at the beginning of the management session. As explained above, the subsets of possible failures, which are presented to the Human Manager, form the set *SC* such that the mismatching value is lower than m. This parameter will depend on the network size and on how accurate the Human Manager wants the alarm diagnosis, where the lower m is, the less mismatching allowed. There is a compromise between the desired accuracy and the probability of losing an alarm or receiving false ones.

The output of the algorithm is presented in an Output Frame.It has three windows: the first one gives the number of received alarms, the second one presents *SC* (that is, each subset Q) and the third one gives, for each of the subsets, the mismatching value.

4. Examples

The AFA is applied on two different network topologies. The first network has a meshed topology, the ARPA2 topology and its nodes are either local nodes (the ones with ADF) or central nodes (the ones with switches). The second network has two protected WDM rings interconnected through the main node of each ring, which is the only node having a switch. The results of the AFA are shown in different scenarios: single failures, double failures and for some missing and false alarms. We will see the location of the failure is more exact when more channels are established due to the fact that the information that the manager receives is larger. In some cases, the

AFA presents several components as possible failures because none of them is a better candidate than the others. In this case, the possible existence of another parameter (as for example, the failure history of the component or how old is it) may help to refine the list of faulty candidates.
Two of the most usual failures at the physical layer are the following:

1. Optical fibre failures: The optical fibres are installed underground and they can suffer from cuts due, for example, to animals or earth movements. In case a fibre is broken, all the channels that were transmitted through it are interrupted and all the receivers, protection switches and 3Rs related to the interrupted channels and located after the failure will send alarms to the manager. In case a ribbon is broken, channels in both directions will be interrupted thus the number of alarms will be higher and more dispersed.

2. Wavelength instability: The stability of the emitted wavelength is an important parameter because it is the one that guarantees controlled levels of crosstalk and distortion. For example, if the temperature of the laser is too high, the emitted wavelength may be shifted. If the link is using optical amplifiers or if the signal is passing through add/drop filters, because they do not have a flat gain, this wavelength may get less amplified than the other wavelengths. Due to this problem, the optical power at the associated receiver may be under the threshold and an alarm informing about a missing optical signal will be send to the manager.

4.1. MESHED OPTICAL NETWORK WITH ARPA2 TOPOLOGY

ARPA2 is a well-known meshed topology that is going to be considered as the topology of the studied optical network (see Figure 10).

We will consider this network as a multi-hop network so that each node uses a different wavelength and that each link between two nodes consists on two separate optical fibers. Due to this bi-connectivity, the information addressed to a certain node can reach it from at least two different paths and the protection switch will choose the best input among all (for example, Node 12 can receive information from two different fibers and Node 13 from 3 different ones). There are two kinds of nodes: local nodes and central nodes. The former ones are the nodes that have an ADF and the latter ones are the nodes with a switch (contrary to local nodes, central nodes have more than two connected nodes). In the example of Figure 10, Nodes 1, 4, 5, 7, 9, 10, 11, 12, 14, 17, 18, 19, 20 and 21 are local nodes (marked with a single circle) and Nodes 2, 3, 6, 8, 13, 15 and 16 are central nodes (represented by a double circle). Each kind of node has a different internal structure hence its modeling results to be different. During the modeling, each element is associated to a four digits identifier that is unique in the network and helps to a fast localization within the network.
Figure 11(a) presents the internal structure of a local node and Figure 11(b) presents its modeling. Each element has its identifiers in accordance with Table II where the first digit gives the class of element, according to Section 2.3.3, the second digit x is the node identifier, the third digit is 0 and the fourth one determines the element within the node.
On the other hand, Figure 13(a) shows the components of a central node where k is

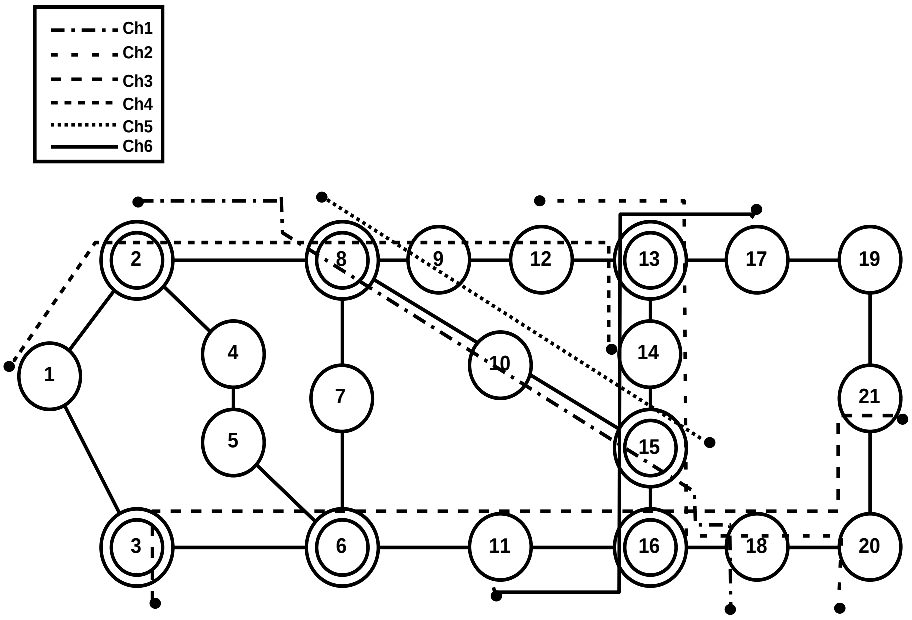

Figure 10. ARPA2 Network with optical links

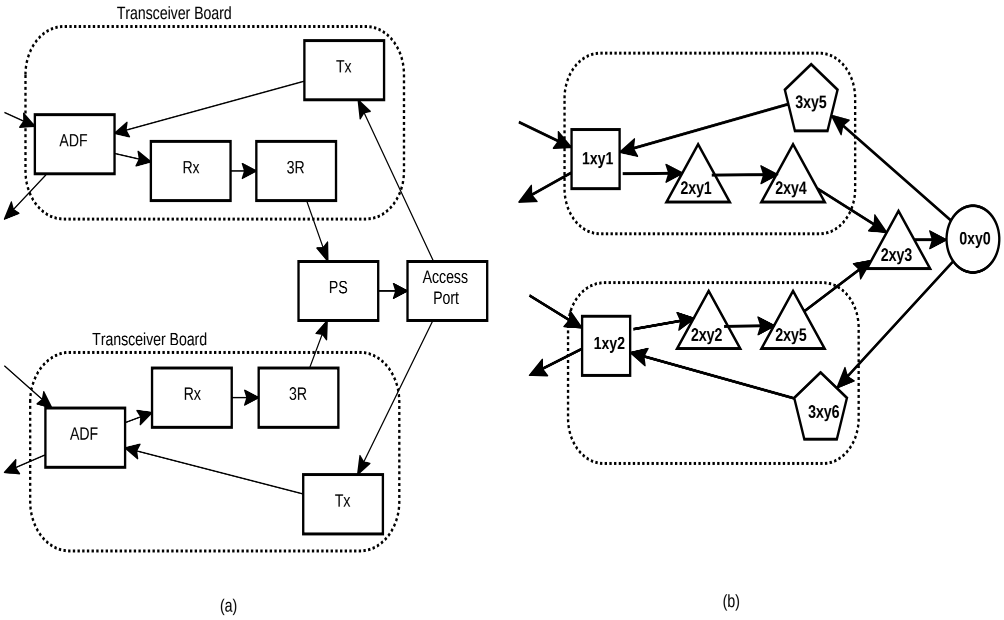

Figure 11. Local Node Internal structure (a) and its modeling (b)

the number of nodes connected to it. Figure 13(b) presents its modeling. The element identifiers have been assigned in accordance with Table III where the first digit gives the class of element, the second digit x is the node identifier and the combination of digits third and fourth is the identification of the component within the node.

For this example, six channels have been established (Table IV).

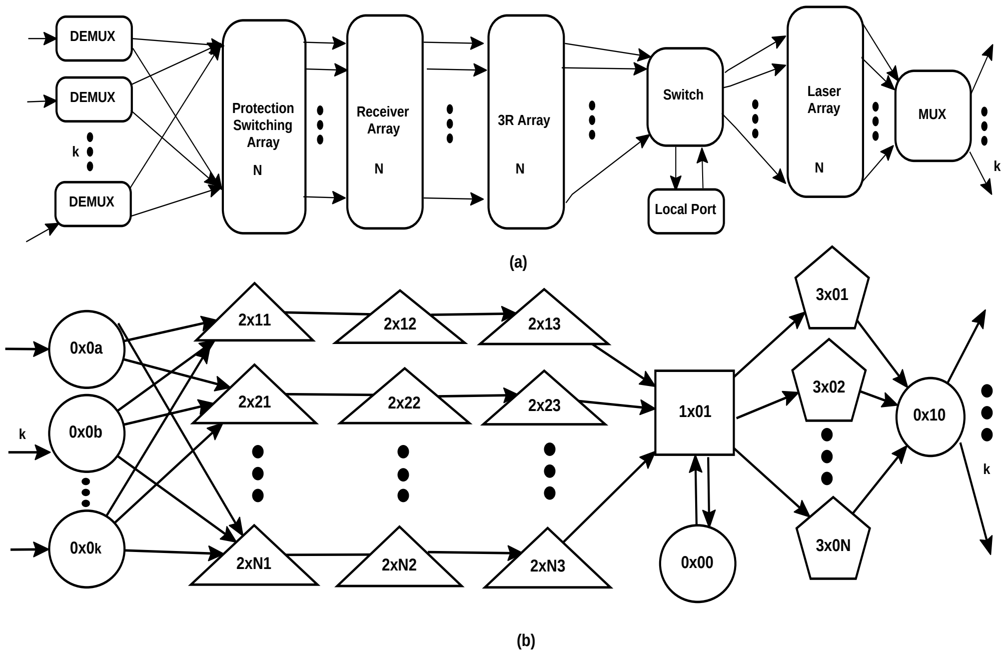

Figure 12. Central Node Internal structure (a) and its modeling (b)

Table II. Identifiers of a local node hardware components

Component	Identifier
ADF CW node *x*	1 *x* 0 1
ADF CCW node *x*	1 *x* 0 2
Rx CW node *x*	2 *x* 0 1
Rx CCW node *x*	2 *x* 0 2
PS node *x*	2 *x* 0 3
3R CW node *x*	2 *x* 0 4
3R CCW node *x*	2 *x* 0 5
Local Access Port node *x*	0 *x* 0 0
Laser CW node *x*	3 *x* 0 5
Laser CCW node *x*	3 *x* 0 6

Table III. Identifiers of the central node hardware components

Component	Identifier
MUX node x	$0\ x\ 1\ 0$
DEMUX node x from node y	$0\ x\ 0\ y+1$
PS node x at wavelength λ_z	$2\ x\ z\ 1$
3R node x at wavelength λ_z	$2\ x\ z\ 3$
Rx node x at wavelength λ_z	$2\ x\ z\ 2$
Switch node x	$1\ x\ 0\ 1$
Local Access Port node x	$0\ x\ 0\ 0$
Laser node x at wavelength λ_z	$3\ x\ 0\ z$

Table IV. Established channels

Channel	Input Node	Output Node	Intermediate Nodes
Ch #1	2	18	2-8-10-15-16-18
Ch #2	12	20	12-13-14-15-16-18-20
Ch #3	3	21	3-6-11-16-18-20-21
Ch #4	1	14	1-2-8-9-12-13-14
Ch #5	8	15	8-10-15
Ch #6	11	17	11-16-15-14-13-17

The routing of the channels has been done arbitrarily and does not enter within the scope of this paper. The input C to the AFA (presented in Appendix) is the ordered list of components of each channel. For example, the ordered list related to *Ch #1* is:

Ch #1: **(0 2 0 0)**(1 2 0 1)(3 2 0 18)(0 2 1 0)(0 0 2 8)(0 8 0 2)(2 8 18 1)(2 8 18 2)(2 8 18 3)(1 8 0 1)(3 8 0 18) (0 8 1 0)(0 0 8 10)(1 10 0 1)(0 0 10 15)(0 15 0 10)(2 15 18 1)(2 15 18 2)(2 15 18 3)(1 15 0 1)(3 15 0 18)(0 15 1 0) (0 0 15 16)(0 16 0 15)(2 16 18 1)(2 16 18 2)(2 16 18 3)(1 16 0 1)(3 16 0 18)(0 16 1 0)(0 0 16 18)(1 18 0 1) (2 18 0 1)(2 18 0 4)(2 18 0 3)**(0 18 0 0)**

Several failure scenarios have been tested with this configuration. The results are presented in Table V.

Failure of an optical fibre : In Scenario 1 we have considered the failure of an optical fibre, as described in the beginning of this Section. In this scenario the problem is the failure of the optical fibre that connects Node 16 with Node 18 which is identified by (0 0 16 18). In this case, three channels are interrupted: Ch #1, #2 and #3. Therefore, the components that will send an alarm are the *A2* components located after the failure and without any *A3* component between the failure and themselves. These components are: (2 18 0 1)(2 18 0 4)(2 18 0 3)(2 20 0 1)(2 20 0 4)(2 20 0 3)(2 21 0 1)(2 21 0 4)(2 21 0 3), that is, the three Protection Switches, 3Rs and Receivers at respectively Nodes 18, 20 and 21. The solution of the AFA having *m=0*, is either MUX (0 16 1 0)or Optical Fibre (0 0 16 18). The AFA is not able to distinguish which of these adjoint *non-alarming* elements has failed because any of them provide information to the manager. In this case, as well as in all the next scenarios that have alarms from Protection Switches, there are channel events that are considered in the *Rec_Channel module*. The agent that controls the hardware components, informs to the manager that the Protection Switch has changed its position so that the channel path should be updated. In this case, *CH* changes. For example, for Ch #1, the channel path will pass by Nodes 2, 8, 9, 12, 13, 17, 19, 21, 20 and 18. This is one possibility that will depend on the channel routing applied and is outside the scope of this paper. When the Protection Switch changes its position, both paths are stored so that the failure that provoked this switch can be located given the previous path. After localizing the failure, the path is removed from *CH*.

Laser wavelength instability In Scenario 2, the failure case of Subsection 2.4.2 has been considered. The laser at Node 6 that emits at wavelength λ_{16} has a problem. If the micro-controller detects any anomaly at the laser, as for example the temperature out of the *Ma* range, it will send an alarm. Because this alarm is sent by an *A3* component, this alarm will be processed by the *Alarm_Discarding_1* module and the origin of the problem will be perfectly localised. But if the micro-controller of the laser does not detect any anomaly, no alarm is sent. In this case, the resulting alarms from this problem are processed in the *Alarm_Discarding_2* module. These alarms are sent by the *A2* components that after demultiplexing the WDM signal detect the levels of each of the demultiplexed signals. In this case the alarms are coming from (2 16 21 1)(2 16 21 2)(2 16 21 3). The result of the algorithm in this case will be also the correct one but, with the mismatching value to 1 because the alarm from the laser was expected and was lost.

Two simultaneous failures In Scenario 3, two simultaneous failures have been considered: the two optical fibres between Node 13 and 14 fail: (0 0 13 14) and (0 0 14 13). This is the case when an optical ribbon is broken. In this case 3 channels are interrupted. Ch #2 and #4 use the fibre from Node 13 to Node 14 and Ch #6 uses the optical fibre from Node 14 to Node 13. The algorithm localises the failure perfectly.

4.2. A WDM RING NETWORK: COBNET

This Alarm Filtering Algorithm is applied in the COBNET Project [1]. The scope of the project is to interconnect, through the Public Network (PN), Corporate Premises Networks (CPNs) using new photonic technologies with multichannel transmission

Table V. Testing results

Scenario	no. of alarms	Involved channels	Alarm reduction	Results
Scenario 1	8	Ch #1,Ch #2,Ch #3	62%	MUX(0 16 1 0)(m=0) O. F.(0 0 16 18)(m=0)
Scenario 2	4	Ch #3	50%	Laser(3 6 0 21)(m=0)
Scenario 2’	3	Ch #3	66%	Laser(3 6 0 21)(m=1)
Scenario 3	9	Ch #2,Ch #4,Ch #6	66%	O. F.s (0 0 13 14)& (0 0 14 13)(m=0)

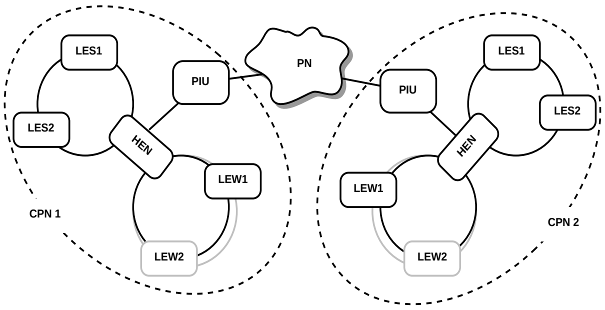

Figure 13. The COBNET network with two nodes at each protected ring

that are realized with space- and wavelength-division multiplexing (SDM and WDM) techniques (see Figure 4.2 where two CPNs are connected through the PN).

CPNs are composed of:

- an interface to the PN called Public Interface Unit (PIU),
- a central node called High End Node (HEN) that contains the main switch. The internal structure of this node is presented in Figure 15(a). This switch interconnects inputs from/outputs to any ring, its own local ports and the ports to/from the PN that are located at the PIU,
- 2 rings: a WDM add/drop ring for the network nodes separated by more than about 2km and a SDM add/drop ring for nodes that are closer to each other. Both rings are duplicated in hardware, so that the CPNs are less vulnerable to failures. Data is sent in both rings: the clockwise (CW ring) and the counter-clockwise ring (CCW ring). A WDM ring is composed of only one optical fibre whereas an SDM ring has as many optical fibres as nodes in the ring. The Low End Nodes (LENs) of the rings are denoted by LEW for the WDM ring and by LES for the SDM ring. Each LEW node, whose structure is shown at Figure 12(a), has two transmitters and two receivers because of the double ring.

Here we consider only the WDM part of the CPNs for two reasons:

(i) it is the technology that will be used in future optical networks and

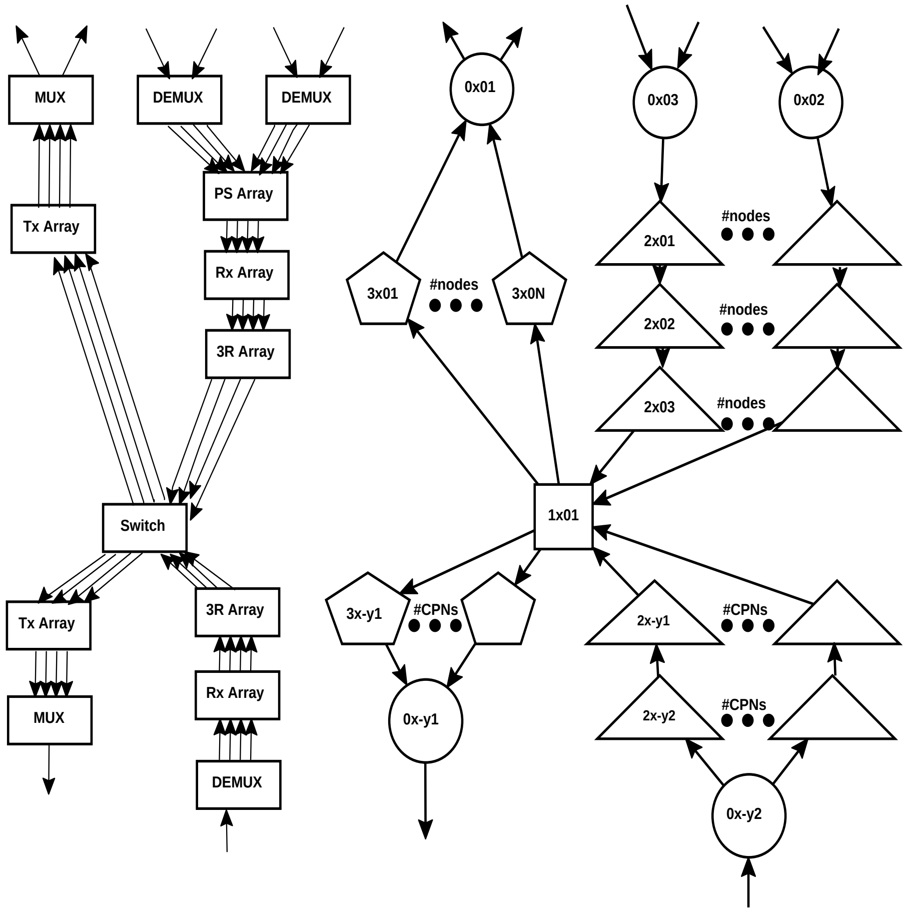

Figure 14. HEN internal structure (a) and its modeling (b)

(ii) these are the rings where the failures are the most difficult to be localized because of the higher number of channels interrupted by a single failure and the resulting increase of the number of alarms.

The management of the COBNET network is done hierarchically: a global manager controls all CPNs and interacts with the CPN entities. The CPN entities act as agents when they receive commands from the global manager, and as managers when they interact with the node agents. The node agents interact with the micro-controllers at each board. The fault management can be done either at the CPN level or at the global manager level. The manager receives alarms from all the elements of its management area and has to identify the failure(s). Once the global manager has this result, it forwards it to the Human Manager so that the testing phase can begin.

In order to apply our algorithm to the COBNET network, the CPN has been modeled as the interconnection of network components through a protected WDM ring. The structure and modeling of the LENs are the same as the local nodes of the previous example (Figure 12(a) and (b)) where x is now the ring identifier and the third digit is the node identifier (instead of 0 as in the local nodes of the meshed network). The internal structure of the HEN and its modeling as interconnection of components are

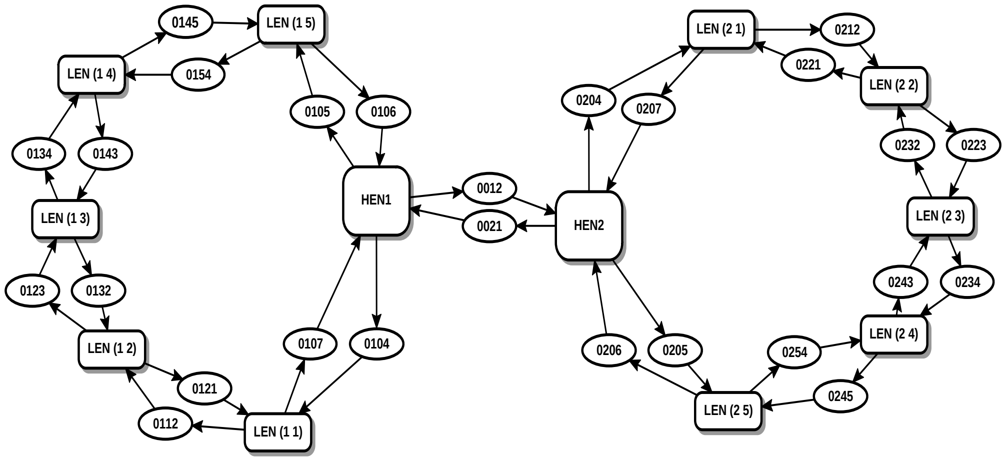

Figure 15. Optical network with 5-node protected rings

shown in Figure 15. The four digits identifiers are: the first digit gives the class of element, the second digit *x* is the ring identifier and the combination of digits third and fourth is the identification of the component within the node where *y* is the connected node identifier to this HEN.

The overall network with the ring topology that we have considered is presented in Figure 15. Thanks to the double ring structure, the established connections in the network are not interrupted if there is a failure. The protection mechanism works as follows. Let us consider a bidirectional communication between LENs Nodes (1 4) and (2 2). The protection switch (PS) at each node is set at the clock position by default. In this case, the two paths for the two unidirectional channels (Ch #1 and CH #1') are:

Ch #1 from LEN (1 4) to LEN(2 2): *LEN(1 4)-LEN(1 5)-HEN1-HEN2-LEN(2 1)-LEN(2 2)*

Ch #1' from LEN(2 2) to LEN(1 4): *LEN(2 2)-LEN(2 3)-LEN(2 4)-LEN(2 5)-HEN2-HEN1-LEN(1 1)-LEN(1 2)-LEN(1 3)-LEN(1 4)*

If the LEN stops receiving signal from the CW ring, its PS switches and receives the signal from the CCW ring so that the channel is not interrupted. Following the example, if the fibre (0 1 1 2) between LEN(1 1) and LEN(1 2) breaks, the PS at LEN(1 4) will switch to the CCW ring so that the return path of Ch #1 will become:

From LEN(2 2) to LEN(1 4): LEN(2 2)-LEN(2 3)-LEN(2 4)-LEN(2 5)-HEN2-HEN1-LEN(1 5)-LEN(1 4)

Even if the connection is restored and no data is lost, the failure has to be localised and repaired. We have studied different failure scenarios considering the unidirectional channels presented at Table VI.

Each bi-directional connection (*node a-node b*) is considered as two unidirectional channels Ch #x and Ch #x' where the former is the direct path (from *node a* to *node b*) and the latter is the reverse path (from *node b* to *node a*). There are 10 established channels that are the *C* input of the AFA. The channel paths are the following are presented in Appendix 2.

Several scenarios has been considered:

Table VI. Established channels

Channel	Input Node	Output Node	Channel	Input Node	Output Node
Ch #1	(1 4)	(2 2)	Ch #1’	(2 2)	(1 4)
Ch #2	(2 1)	(1 3)	Ch #2’	(1 3)	(2 1)
Ch #3	(1 2)	(2 3)	Ch #3’	(2 3)	(1 2)
Ch #4	(1 1)	(1 5)	Ch #4’	(1 5)	(1 1)
Ch #5	(2 4)	(2 5)	Ch #5’	(2 5)	(2 4)

Failure of an optical fibre : Scenario 1 considers the hardware failure of the Optical Fibre (0 1 2 3), as presented in Subsection 2.4.1. In this case, four channels are interrupted Ch #1’, Ch #2, Ch #3, Ch #4 and restored thanks to the PS. Twelve alarms are sent to the manager: (2 1 4 1)(2 1 5 4)(2 1 4 4)(2 1 3 1)(2 1 4 3)(2 1 0 6)(2 1 3 3)(2 1 3 4)(2 1 0 4)(2 1 5 1)(2 1 0 5) and (2 1 5 3). The solution of the algorithm presents as perfect matching (*m=0*) the optical fibre that connects LEN(1 2) with LEN(1 3).

Two simultaneous failures : Scenario 2 considers two simultaneous failures in two different optical fibres: (0 1 2 3) and (0 1 0 6). In this case, the alarms related to both failures reach the manager simultaneously and intermingled. The manager receives a total of 21 alarms. The failures are well identified, with *SC* being the two sets: ((0 1 2 3) and (0 1 0 6)) or ((0 1 2 3) and (0 1 0 3)) for a perfect matching *m=0*. The algorithm can not distinguish between (0 1 0 6) -Optical Fibre- and (0 1 0 3) -Demultiplexer- because they are two consecutive *non-alarming* components with no *alarming* component between them to give more information to the manager (same case than Scenario 1 of ARPA2).

Two simultaneous failures with lost alarm : Scenario 3 studies the two previous simultaneous failures but considering the existence of a lost alarm. In this scenario the alarms related to both failures also reach the manager simultaneously and mixed but there is one that get lost. Giving *m=2*, the AFA result is also ((0 1 2 3) and (0 1 0 6)) or ((0 1 2 3) and (0 1 0 3)), which are the best candidates since they achieve the lowest mismatching value.

Two simultaneous failures with false alarm In this fourth scenario we studied the behaviour of the AFA when a false alarm is generated. Having the same double failure as in Scenario 2, the received alarms at the manager level are the same 21 alarms and the alarm from (2 1 5 4). In this scenario, the AFA result is also ((0 1 2 3) and (0 1 0 6)) or ((0 1 2 3) and (0 1 0 3)), which are the best candidates since they achieve the lowest mismatching value 1 which corresponds to the unexpected alarm.

Vague result : In this scenario 5 we considered that they were established only Ch #1, Ch #2 and Ch #3. After a simulated failure, the manager received 9 alarms coming

Table VII. Testing results

Scenario	no. of alarms	Involved channels	Results
Scenario 1	12	Ch #1', Ch #2, Ch #3, Ch #4	O. F.(0 1 2 3)(m=0)
Scenario 2	21	Ch #1, Ch #1', Ch #2, Ch #2', Ch #3, Ch #4, Ch #4'	O. F. (0 1 2 3) and (0 1 0 6)(m=0)
Scenario 3	20	Ch #1, Ch #1', Ch #2, Ch #2', Ch #3, Ch #4, Ch #4'	O. F. (0 1 2 3) and (0 1 0 6)(m=1)
Scenario 4	22	Ch #1, Ch #1', Ch #2, Ch #2', Ch #3, Ch #4, Ch #4'	O. F. (0 1 2 3) and (0 1 0 6)(m=1)

from: (2 1 4 1)(2 1 3 4)(2 1 2 3)(2 1 4 4)(2 1 3 3)(2 1 2 1)(2 1 4 3)(2 1 3 1)(2 1 2 4). In this case the result of the AFA is weak because it gives three possible components: (0 1 0 1)(0 1 0 4)(0 1 1 2). A possible solution is to establish a ghost channel such that involves the node in the failure domain so that the manager will have extra alarms with their information and the AFA will be able to locate the failure better. In this case, considering as ghost channel the one going from LEN(1 1) to LEN(1 1) the new alarms will be (2 1 1 1)(2 1 1 4)(2 1 1 3) and the new AFA result with perfect matching is: (0 1 1 2)

5. Conclusion

We have studied the components of a real optical network and their behaviour when a failure occurs. We have introduced a model as a first step for the Alarm Filtering Algorithm (AFA), an optical network as relations between classes. We have developed and implemented an AFA to localise failures in an optical WDM network. This algorithm allows to locate multiple failures, of both '*alarming*' and '*non-alarming*' elements. It requires a minimal amount of information as input, namely the established channels in the network, and the origin and type of alarms. No knowledge of the network topology or of failure probabilities is needed. The output of the algorithm is the set of subsets containing the most likely elements that may have failed and have prompted the alarm messages received by the manager. The more channels are established in the network, the easier it gets to locate the failure(s) because the amount of information that the manager receives due to the failure(s) is larger. Finally, the application of the AFA has been implemented in Java and its application to two different network topologies has been described: an ARPA2 meshed network and the European ACTS project COBNET network.

Authors' Vitae

Carmen Mas
Carmen Mas, born in Huesca (Spain) in 1970, earned her Telecommunication Engineering Degree from the Polytechnic University of Catalonia ETSETB, Barcelona (Spain) in 1995. She completed her Graduate Thesis at the Applied Physics Department of the Free University of Brussels VUB (Belgium) on the design of diffractive optical elements for optical interconnections in computing and sensing systems. In 1995 she completed her research at the VUB and she moved to the Electrical and Engineering Department of King's College London (UK) to work in fabrication of optical devices. In May 1996 she began as a research assistant in the Institute for computer Communications and Applications in EPFL (Switzerland) where she pursues her Ph.D. Among other activities, she has been an assistant for the Computer Networking and Traffic Control graduate course and has participated in the COBNET European project where her contribution was the implementation of the management platform.

Patrick Thiran
Patrick Thiran received the electrical engineering degree from the Université Catholique de Louvain,Louvain-la-Neuve, Belgium, in 1989, the M.S. degree in electrical engineering from the University of California at Berkeley in 1990 and the PhD degree from the Swiss Federal Institute of Technology at Lausanne (EPFL) in 1996. He then joined the Institute of Computer Communications and their Applications of EPFL, where he is was given the title of professor in 1998. His research interests are in fields of traffic control in communication networks, vulnerability of optical networks, system theory and neural networks. He received the 1996 Doctoral Prize of EPFL, and he serves as an associate editor of the IEEE Transactions on Circuits and Systems, Part II.

Jean-Yves LeBoudec
Jean-Yves Le Boudec is full professor at EPFL. He was born in Brittany, France in 1958. He graduated from Ecole Normale Superieure de Saint-Cloud, Paris, where he obtained the Agregation in Mathematics in 1980. He received his doctorate in 1984 from the University of Rennes, France, and became an assistant Professor at INSA/IRISA, Rennes. In 1987 and 1988 he was with Bell Northern Research, Ottawa, Canada, as a member of scientific staff in the Network and Product Traffic Design Department. In 1988, he joined the IBM Zurich Research Laboratory at Rushchlikon, Switzerland, where he was manager of the Customer Premises Network Departement. In 1994 he formed the Laboratoire de Reseaux de Communication at EPFL, which in autumn 1997 became part of ICA. His interests are in the architecture and performance of communication systems.

6. Acknowledgements

This work is supported by the ACTS project COBNET and the 'Office Federal pour l'Education et la Science (OFES) of the Swiss government.

7. Appendix 1: Established channels at the meshed network

Ch #1: **(0 2 0 0)**(1 2 0 1)(3 2 0 18)(0 2 1 0)(0 0 2 8)(0 8 0 2)(2 8 18 1)(2 8 18 2)(2 8 18 3)(1 8 0 1)(3 8 0 18) (0 8 1 0)(0 0 8 10)(1 10 0 1)(0 0 10 15)(0 15 0 10)(2 15 18 1)(2 15 18 2)(2 15 18 3)(1 15 0 1)(3 15 0 18)(0 15 1 0) (0 0 15 16)(0 16 0 15)(2 16 18 1)(2 16 18 2)(2 16 18 3)(1 16 0 1)(3 16 0 18)(0 16 1 0)(0 0 16 18)(1 18 0 1) (2 18 0 1)(2 18 0 4)(2 18 0 3)**(0 18 0 0)**

Ch #2: **(0 12 0 0)**(3 12 0 5)(1 12 0 1)(0 0 12 13)(0 13 0 12)(2 13 12 1)(2 13 12 2)(2 13 12 3)(1 13 0 1) (3 13 0 20)(0 13 1 0)(0 0 13 14)(1 14 0 1)(0 0 14 15)(0 15 0 14)(2 15 20 1)(2 15 20 2)(2 15 20 3)(1 15 0 1) (3 15 0 20)(0 15 1 0)(0 0 15 16)(0 16 0 15)(2 16 20 1)(2 16 20 2)(2 16 20 3)(1 16 0 1)(3 16 0 20)(0 16 1 0) (0 0 16 18)(1 18 0 1)(0 0 18 20)(1 20 0 1)(2 20 0 1)(2 20 0 4)(2 20 0 3)**(0 20 0 0)**

Ch #3: **(0 3 0 0)**(1 3 0 1)(3 3 0 21)(0 3 1 0)(0 0 3 6)(0 6 0 3)(2 6 21 1)(2 6 21 2)(2 6 21 3)(1 6 0 1)(3 6 0 21) (0 6 1 0)(0 0 6 11)(1 11 0 1)(0 0 11 16)(0 16 0 11)(2 16 21 1)(2 16 21 2)(2 16 21 3)(1 16 0 1)(3 16 0 21)(0 16 1 0) (0 0 16 18)(1 18 0 1)(0 0 18 20)(1 21 0 1)(0 0 20 21)(1 21 0 1)(2 21 0 1)(2 21 0 4)(2 21 0 3)**(0 21 0 0)**

Ch #4: **(0 1 0 0)**(3 1 0 5)(1 1 0 1)(0 0 1 2)(0 2 0 1)(2 2 1 1)(2 2 1 2)(2 2 1 3)(1 2 0 1)(3 2 0 14)(0 2 1 0) (0 0 2 8)(0 8 0 2)(2 8 14 1)(2 8 14 2)(2 8 14 3)(1 8 0 1)(3 8 0 14)(0 8 1 0)(0 0 8 9)(1 9 0 1)(0 0 9 12)(1 12 0 1) (0 0 12 13)(0 13 0 12)(2 13 14 1)(2 13 14 2)(2 13 14 3)(1 13 0 1)(3 13 0 14)(0 13 1 0)(0 0 13 14)(1 14 0 1) (2 14 0 1)(2 14 0 4)(2 14 0 3)**(0 14 0 0)**

Ch #5: **(0 8 0 0)**(1 8 0 1)(3 8 0 15)(0 8 1 0)(0 0 8 10)(1 10 0 1)(0 0 10 15)(0 15 0 10)(2 15 15 1)(2 15 15 2) (2 15 15 3)(1 15 0 1)**(0 15 0 0)**

Ch #6: **(0 11 0 0)**(3 11 0 5)(1 11 0 1)(0 0 11 16)(0 16 0 11)(2 16 11 1)(2 16 11 2)(2 16 11 3)(1 16 0 1)(3 16 0 17) (0 16 1 0)(0 0 16 15)(0 15 0 16)(2 15 17 1)(2 15 17 2)(2 15 17 3)(1 15 0 1)(3 15 0 17)(0 15 1 0)(0 0 15 14) (1 14 0 1)(0 0 14 13)(0 13 0 14)(2 13 17 1)(2 13 17 2)(2 13 17 3)(1 13 0 1)(3 13 0 17)(0 13 1 0)(0 0 13 17) (1 17 0 1)(2 17 0 1)(2 17 0 4)(2 17 0 3)**(0 17 0 0)**

8. Appendix 2: Established channels at the COBNET network

Channel1: **(0 1 4 0)**(3 1 4 5)(1 1 4 1)(0 1 4 5)(1 1 5 1)(0 1 0 6)(0 1 0 3)(2 1 0 1)(2 1 0 2)(2 1 0 3)(1 1 0 1)(3 1 -2 4)(0 1 -2 1)(0 0 1 2)(0 2 -1 2)(2 2 -1 2) (2 2 -1 1)(1 2 0 1)(3 2 0 2)(0 2 0 1)(0 2 0 4)(1 2 1 1)(0 2 1 2)(1 2 2 1)(2 2 2 1)(2 2 2 4)(2 2 2 3)**(0 2 2 0)**

Channel1': **(0 2 2 0)**(3 2 2 5)(1 2 2 1)(0 2 2 3)(1 2 3 1)(0 2 3 4)(1 2 4 1)(0 2 4 5)(1 2 5 1)(0 2 0 6)(0 2 0 3)(2 2 0 4)(2 2 0 5)(2 2 0 6)(1 2 0 1)(3 2 -1 2) (0 2 -1 1)(0 0 2 1)(0 1 -2 2)(2 1 -2 4)(2 1 -2 3)(1 1 0 1)(3 1 0 4)(0 1 0 1)(0 1 0 4)(1 1 1 1)(0 1 1 2)(1 1 2 1)(0 1 2 3)(1 1 3 1)(0 1 3 4)(1 1 4 1) (2 1 4 1)(2 1 4 4)(2 1 4 3)**(0 1 4 0)**

Channel2: **(0 2 1 0)**(3 2 1 5)(1 2 11)(0 2 1 2)(1 2 2 1)(0 2 2 3)(1 2 3 1)(0 2 3 4)(1 2 4 1)(0 2 4 5)(1 2 5 1)(0 2 0 6)(0 2 0 3)(2 2 0 1)(2 2 0 2)(2 2 0 3)(1 2 0 1) (3 2 -1 1)(0 2

-1 1)(0 0 2 1)(0 1 -2 2)(2 1 -2 2)(2 1 -2 1)(1 1 0 1)(3 1 0 3)(0 1 0 1)(0 1 0 4)(1 1 1 1)(0 1 1 2)(1 1 2 1)(0 1 2 3)(1 1 3 1)(2 1 3 1) (2 1 3 4)(2 1 3 3)**(0 1 3 0)**

Channel2': **(0 1 3 0)**(3 1 3 5)(1 1 3 1)(0 1 3 4)(1 1 4 1)(0 1 3 5)(11 5 1)(0 1 0 6)(0 1 0 3)(2 1 0 7)(2 1 0 8)(2 1 0 9)(1 1 0 1)(3 1 -2 3)(0 1 -2 1)(0 0 1 2) (0 2 -1 2)(2 2 -1 6)(2 2 -1 5)(1 2 0 1)(3 2 0 1)(0 2 0 1)(0 2 0 4)(1 2 1 1)(2 2 1 1)(2 2 1 4)(2 2 1 3)**(0 2 1 0)**

Channel 3: **(0 1 2 0)**(3 1 2 5)(1 1 2 1)(0 1 2 3)(1 1 3 1)(0 1 3 4)(1 1 4 1)(0 1 4 5)(1 1 5 1)(0 1 0 6)(0 1 0 3)(2 10 4)(2 1 0 5)(2 1 0 6)(1 1 0 1)(3 1 -2 2) (0 1 -2 1)(0 0 1 2)(0 2 -1 2)(2 2 -1 4)(2 2 -1 3)(1 2 0 1)(3 2 0 3)(0 2 0 1)(0 2 0 4)(1 2 1 1)(0 2 1 2)(1 2 2 1)(0 2 2 3)(1 2 3 1)(2 2 3 1) (2 2 3 4)(2 2 3 3)**(0 2 3 0)**

Channel 3': **(0 2 3 0)**(3 2 3 5)(1 2 3 1)(0 2 3 4)(1 2 4 1)(0 2 4 5)(1 2 5 1)(0 2 0 6)(0 2 0 3)(2 2 0 7)(2 2 0 8)(2 2 0 9)(1 2 0 1)(3 2 -1 3)(0 2 -1 1)(0 0 2 1)(0 1 -2 2)(2 1 -2 6)(2 1 -2 5)(1 1 0 1)(3 1 0 2)(0 1 0 1)(0 1 0 4)(1 1 1 1)(0 1 1 2)(1 1 2 1)(2 1 2 1)(2 1 2 4)(2 1 2 3)**(0 1 2 0)**

Channel 4: **(0 1 1 0)**(3 1 1 5)(1 1 1 1)(0 1 1 2)(1 1 2 1)(0 1 2 3)(1 1 3 1)(0 1 3 4)(1 4 1)(0 1 4 5)(1 1 5 1)(2 1 5 1)(2 1 5 4)(2 1 5 3)**(0 1 5 0)**

Channel 4': **(0 1 5 0)**(3 1 5 5)(1 1 5 1)(0 1 0 6)(0 1 0 3)(2 1 0 10)(2 1 0 11)(2 1 0 12)(1 1 0 1)(3 1 0 1)(0 1 0 1)(0 1 0 4)(1 1 1 1)(2 1 1 1)(2 1 1 4)(2 1 1 3)**(0 1 1 0)**

Channel 5: **(0 2 4 0)**(3 2 4 5)(1 2 4 1)(0 2 4 5)(1 2 5 1)(2 2 5 1)(2 2 5 4)(2 2 5 3)**(0 2 5 0)**

Channel 5': **(0 2 5 0)**(3 2 5 5)(1 2 5 1)(0 2 0 6)(0 2 0 3)(2 2 0 10)(2 2 0 11)(2 2 0 12)(1 2 0 1)(3 2 0 4)(0 2 0 1)(0 2 0 4)(1 2 1 1)(0 2 1 2)(1 2 2 1)(0 2 2 3)(1 2 3 1)(0 2 3 4)(1 2 4 1)(2 2 4 1)(2 2 4 4)(2 2 4 3)**(0 2 4 0)**

References

1. COBNET Consortium, 'COBNET: Corporate Optical Backbone NETwork', URL: http://lrcwww.epfl.ch/COBNET/index.html
2. J. Walrand: Communication Networks, Aksen Associates.
3. A. T. Bouloutas et al.: Alarm Correlation and Fault Identification in Communication Networks, IEEE Transactions on Communications, Vol. 42, no. 2/3/4 (Feb/March/April 1994) pp.523-533.
4. B. Mukherjee: Optical Communication Networks, McGraw-Hill Series, 1997.
5. T. Bouloutas et al.: Fault Identification Using a FSM model with Unreliable Partially Observed Data Sequences, IEEE Transactions on Communications, Vol. 41, no. 7 (July 1993), pp. 1074-1083.
6. R. H. Deng et al.: A probabilistic Approach to Fault Diagnosis in Linear Lightwave Networks, IEEE Journal on Selected Areas in Communications, Vol. 11, no. 9 (December 1993), pp.1438-1448.
7. R. Gardner et al.: Alarm Correlation and Network Fault Resolution using Kohonen Self-Organising Map, Globecom 97 proceedings, 1997, pp. 1398-1402.
8. C. Mas et al.: Fault Localization for Optical Networks, All-optical networking: Architecture, Control and Management issues (SPIE'98), Boston, Vol. 3531, pp. 408-419, 1998.
9. G. Stix: Nothing but Light, Scientific America, In Focus, December 1998.
10. ITU-T Recommendation M.3100: Generic Network Information Model, 1995.
11. ITU-T Recommendation X.721: Information Technology-Open Systems Interconnection-Structure of management Information: Definition of management information, 1992.

List of Figures

List of Tables