Practical Issues for the Implementation of Survivability and Recovery Techniques in Optical Networks

Georgios Ellinas♣, Dimitri Papadimitriou♦, Jacek Rak+, Dimitri Staessens‡
James P.G. Sterbenz†, and Krzysztof Walkowiak♠

♣ Department of Electrical and Computer Engineering, University of Cyprus, CY, email: *gellinas@ucy.ac.cy*

♦ Network Research Department, Alcatel-Lucent Bell Labs, BE, email: *dimitri.papadimitriou@alcatel-lucent.com*

+Department of Computer Communications, Gdansk University of Technology, PL, email: *jrak@{pg.gda.pl, ieee.org}*

‡ Internet Based Communication Networks and Services Department, Ghent University, BE, email: *dimitri.staessens@intec.ugent.be*

† Electrical Eng. and Computer Science Dept., The Univ. of Kansas, US / School of Computing and Communications, Lancaster Univ., UK, email: *jpgs@comp.lancs.ac.uk*

♠ Department of Systems and Computer Networks and ENGINE - European research centre of Network intelliGence for INnovation Enhancement, Wroclaw University of Technology, Wybrzeze Wyspianskiego 27, 50-370 Wroclaw, PL, email: *krzysztof.walkowiak@pwr.wroc.pl*

Abstract

Failures in optical networks are inevitable. They may occur during work being done for the maintenance of other infrastructures, or on a larger scale as the result of an attack or large-scale disaster. As a result, service availability, an important aspect of Quality of Service (QoS), is often degraded. Appropriate fault recovery techniques are thus crucial to meet the requirements set by the Service Level Agreements (SLAs) between the carriers and their customers.

In this paper, we focus on practical issues related to the deployment of fault recovery mechanisms in commercial optical networks. In particular, we outline the most important functionalities that, to our knowledge, need to be implemented, as well as discuss the related problems making deployment of fault recovery mechanisms difficult. Investigated topics include: fault recovery challenges (fault detection, location, and recovery), multiple failures recovery, as well as application of reliability mechanisms in elastic optical networks, and in multiprovider multilevel networks.

Keywords: dependability · reliability · availability · resilience · survivability · fault tolerance · fault recovery · multiple failures · multilevel multiprovider optical networks · implementation issues

1. Introduction

Network survivability, defined in [1] as the ability to provide continuous service in the presence of failures, is a critical issue for high-bandwidth backbone optical networks with arbitrary mesh topologies. Failures in fiber-optic networks occur often due to the fact that they are a cable-based technology and the infrastructure is co-located with networks for other utilities. Thus, damages usually happen during work being done for the maintenance of other infrastructures.

Furthermore, due to the use of wavelength division multiplexing (WDM) technology in these networks, each fiber can carry an extremely high volume of traffic, thus more traffic is concentrated on fewer routes, increasing the number of customers that can be potentially affected by a failure.

In this paper, we focus on deployment issues of fault recovery mechanisms in commercial optical networks. In particular, we discuss the current needs concerning implementation of failure recovery techniques as well as the related problems following e.g., from hardware constraints.

Over the past two decades, various approaches have been proposed for the recovery of the traffic when a failure event occurs. They are mainly based on utilization of alternate paths, called *backup paths* (*BPs*), used to redirect the traffic after a failure of a network element affecting the primary routes, called *working paths* (*WPs*) [2]. Specific schemes of backup routes include link-, path-, segment-, and cycle-based techniques.

The general requirement when providing protection against failures of nodes/links is that the respective backup paths should be node- (link-) disjoint from the working paths being protected [3]. Additionally, link capacities reserved for backup paths can be shared along certain links, if the considered backup paths protect mutually disjoint working paths [2]. Failures of single links/nodes are the most frequent types of failures. However, due to the often observed inter-failure correlation, a large set of solutions is dedicated to the case of multiple failures, i.e., a simultaneous failure of several network elements [4].

Survivability can be provided either in a proactive way - using a *protection* strategy, implying establishment of a backup path before the occurrence of a failure (i.e., at the time of working path establishment), or by means of a reactive *restoration* strategy. In the latter case, the network tries to establish a new connection using available resources only after a failure has occurred. Protection typically has faster recovery speed but lower resource-efficiency than restoration [5].

There are many types of service disruptions in optical networks, which can be classified in two major types: soft and hard failures [6]. *Hard failures*, such as fiber cuts and failure of a network linecard occur suddenly and have a severe impact on services, causing major loss of traffic. *Soft failures*, such as aging of an amplifier, cause subtle changes in performance, resulting in a wide spectrum of service degradations which are far more difficult to detect and localize.

Some failures, called *self-reported*, are very easily detected because they interfere with the correct functioning of the upstream device and are flagged‡ by internal control mechanisms. Most hard and a large number of soft failures are

self-reported [7]. Soft failures that are not self-reported can be very hard to detect and accurately localizing them is time-consuming and very costly.

Even though failures cannot be avoided, quick detection, identification, and recovery of faults are crucial aspects in the successful deployment of telecommunication networks. A network fault that goes unattended for a long period of time can cause both tangible and intangible losses for the company that provides the service, as well as for its clients. Therefore, the current trend is for more and more networks that are virtually uninterruptible.

Currently, carriers are bound to *service-level agreements* (*SLAs*) with their customers guaranteeing that the customer will be provided with services with a prescribed service availability (e.g., 99.999% availability - equivalent to less than 5 minutes of down time per year), with financial penalties if the SLA availability is not met. It is therefore clear that in optical backbone networks it is essential to have effective fault recovery mechanisms to prevent the loss of information due to fiber cuts or equipment failures, which may occur often enough to cause major service disruptions.

Furthermore, the constant growth of traffic in computer networks mainly due to popular services such as cloud computing and content-oriented networks, has triggered the need to develop an efficient and scalable optical transport platform for capacities beyond 100 Gb/s. One of the technologies, which enables improved use of flexible optical network is a scalable and efficient architecture called *Elastic Optical Networks* (*EONs*). The key innovation of EONs compared to currently used WDM (Wavelength Division Multiplexing) networks is the possibility of using sub-wavelength granularity with 6.25 GHz slices for low-rate transmission and super-channel connectivity for accommodating ultra-high capacity client signals within a common network [8]. Accordingly, the optical spectrum can be used much more flexibly compared to the fixed grid of 50 GHz channels in WDM.

One of the consequences of the flexible grid provided in EONs, is the possibility to provision asymmetric traffic where demands of the same bidirectional connection between a particular pair of nodes have different bandwidth requirements in each direction. Especially, in the context of network survivability based on path protection, this option seems very attractive, since significant savings of network resources in terms of optical spectrum should be obtained.

Although much research has been performed in the area of optical networks' reliability and survivability over the last two decades, there are still many practical issues that need to be addressed for the successful commercial implementation of fault recovery techniques in optical networks. In this paper, we identify the most important of them, with invited sections from panelists based on their presentations at IEE/IFIP RNDM (Reliable Networks Design and Modeling) 2013 [9]. In particular, in Section 2, we outline the fault recovery challenges related to fault detection, localization, and recovery related to physical layer impairment issues, shared protection, as well as switch design considerations, especially for the case of transparent or translucent optical networks, where the signal remains in the optical domain for the entire end-to-end path or for large parts of the path. In Section 3, we extend our investigation to the case of multiple failures. Next, in Section 4, we investigate applicability of reliability mechanisms in the context of Elastic Optical Networks. In Section 5, we address the resilience and recovery issues concerning multiprovider multilevel networks, while Section 6 concludes the paper.

2. Fault Recovery Challenges

A number of challenges can be identified for the practical implementation of recovery techniques in mesh optical networks including issues related to control and management (such as fault detection, isolation, and recovery), design of the optical switch architectures, as well as design of the protection/restoration algorithms and techniques.

An important aspect for fault management is the signaling capability in the network. Most current optical networks use an Optical Supervisory Channel (OSC) for remote node management, monitoring, and control [10]. The OSC is a low bandwidth (STM-1) out-of-band (usually at 1510 nm), full duplex point-to-point communication and control channel. It is a common practice to use the Digital Communication Channel (DCC) section of the STM-1 header or the General Communication Channel (GCC) of OTN for this purpose. In every managed node (e.g., amplifier, regenerator, cross-connect) the channel is dropped, the relevant data is inspected, instructions are performed, and possible replies are added. The OSC can efficiently detect link failures, node failures, and some soft failures if they are self-reported. Link failures are immediately detected at the upstream amplifier, so the span where the failure occurs is efficiently localized.

The cost of the OSC is quite low and does not currently warrant deployment of separate monitoring equipment for link failure localization. Moreover, its main function is to provide communication for configuring and maintaining amplifiers, equalizers, and other network equipment along the links. Some alternatives for the OSC have been proposed, such as work in [11], where a hybrid supervisory plane is envisioned, but the benefits, mainly speed, over the current OSC may not warrant the cost of deploying multiple transparent wavelengths, as time savings at each managed site would be on the order of 100s of microseconds.

For efficient network reliability, we need to be able to detect, localize, and identify failures in the network [12]. This means that we need to have monitoring points in the network, either as part of the transmission system (e.g., Forward Error Correction (FEC) units in the receiver [13]), as part of the control loop of a device (photodiode in an amplifier), or a dedicated monitor. Protection mechanisms can restore the traffic immediately after a failure is detected on the working path (e.g., Loss-of-Signal Payload (LoS-P) in SDH [5] or Bi-directional Forwarding Detection (BFD) for IP/MPLS [14]-[15]), but for restoration mechanisms we also need to localize the failure in order to efficiently re-route traffic around it.

Detection should always be as fast and reliable as possible – false or missing alarms should be rare – but localization requirements can be a trade-off between accuracy, speed, and cost. Failure localization for recovery purposes should use simple, cost-effective devices, be as fast as possible, and be sufficiently accurate to allow routing the traffic around the failure. Localization for repair purposes could be allowed to take more time (hours or days), but should be as accurate as possible. The devices required for accurate localization (such as test devices, optical spectrum analyzers, and optical time-domain reflectometers - OTDRs) are prohibitively expensive to deploy ubiquitously in the network.

We can say that two (inter-dependent) aspects are important: where to place which type of monitor and how to efficiently correlate the alarms raised by these monitors. General probabilistic models for localizing network failures have been examined; however, they typically focus on the sub-problem of alarm correlation.

In [6], a probabilistic approach is examined where a hierarchical causality tree is used to identify the most likely problem. Work in [16] assumes that alarms only carry information about the emitting node, while work in [17] makes use of Alarm Reporting Functions in order to create classes of objects and [18] defines a hierarchical dependency graph consisting of services, protocols, and functions and defining multiple failure modes per element.

The placement of components is another important topic. Work in [19] showed the feasibility of a fault detection scheme for all-optical networks based on their decomposition into monitoring-cycles (*m*-cycles). To detect and localize network faults, it is not necessary to put monitors on all links, lightpaths, or nodes. *M*-cycles were extended to more general structures, such as *m*-trails [20] and *m*-trees.

While most solutions proposed in the literature are sound, their application domains may need revision. *M*-cycles/trails/trees types of solutions could prove useful for locating difficult to find soft failures that affect all channels (note that *m*-cycles/trails/trees operate on separate wavelengths than the actual traffic). Locating soft failures that affect only a single channel (necessary for efficient repairs) is still a difficult issue which remains, to our knowledge, largely unresearched.

After the failure has been detected, there are still challenges that need to be addressed in practical implementations of fault recovery in transparent/translucent optical networks. Some of these include:

- physical layer impairment considerations,
- shared protection considerations,
- switch design considerations for transparent nodes in opaque networks.

(a) Physical Layer Impairment Considerations

In backbone optical networks, the trend is for next-generation mesh optical networks that are evolving from opaque (with electrical components providing optical-electronic-optical (OEO) conversions at all network nodes), to translucent (where OEO conversions are sparsely provided at a few network locations, while some of the connections can stay in the optical domain throughout), and eventually to transparent (all-optical) networks where the nodes provide pure optical switching and the signal is never converted back to the electronic domain until it reaches the receiver at the destination node.

In the opaque approach, all switching and processing of the data at the nodes can be handled by electronics (opaque node/opaque network), or the node switch fabric can be transparent while still maintaining transponders at the WDM systems (transparent node/opaque network), thus again providing OEO conversions at all network nodes.

In terms of survivability, opaque architectures are flexible in the sense that they have access to the electrical signal overhead and they can readily provide all the necessary control and management functions, including fault detection, fault isolation, as well as fault recovery. Furthermore, they only require link-to-link engineering, thus simplifying the design of the fault recovery techniques.

Opaque switching nodes (with an electronic switch fabric and transponders present in the WDM systems) are the ones that are currently deployed by the network operators in core optical networks. However, although this is a well-established technology, the large number of optical-to-electrical-to-optical (OEO) conversions at each switching node greatly increases the network cost, the power consumption, as well as the footprint required to deploy these switches. Furthermore, these architectures cannot keep pace with the growth in capacity of optics in the near future and the rapidly growing customer demand for bandwidth [21].

It is envisioned that at some point in the future, the network operators will eventually move to all-optical architectures for future core mesh optical networks mainly driven by cost and bandwidth considerations [22]. These transparent networks are extremely desirable as they provide bit-rate, protocol, and modulation format transparency, and are more efficient in terms of cost, power, and footprint. In transparent networks, however, there are several challenges that need to be addressed before such architectures can be deployed. Some of these challenges relate to the implementation of efficient fault recovery techniques and include the following:

- The *physical layer impairments* (*PLIs*) incurred by the non-ideal optical transmission medium accumulate along an optical path limiting the transmission reach of optical signals [23]-[27]. As PLIs accumulate, end-to-end system engineering is now required in the network, something that creates several challenges when fault recovery techniques are designed.

- Transparent optical networking solutions fail to recover the full functionality of the opto-electronic versions they replace without byte-level access. Therefore, it is a challenge to provide the control and management functionalities associated with the survivability process (such as fault detection (especially for the case of degradation failures), fault isolation, and fault recovery) that are readily available when we have access to the electrical signal.

Even though, as mentioned earlier, there has been extensive work in the literature for the design of fault recovery techniques, only recently attempts were made to take the PLIs into consideration when designing fault protection/restoration algorithms.

While various works have been proposed for the impairment-aware routing and wavelength assignment (IA-RWA) problem while maintaining the acceptable level of quality of transmission (QoT) of signals (in [28] a survey presents and analyzes most of the works in the literature on this subject), there has not been considerable work on the problem of designing impairment-aware fault recovery techniques that will ensure that in case of a failure event the recovery signal will reach its destination while also having an acceptable level of QoT.

The importance of considering for the PLIs when designing fault recovery techniques is shown in Fig. 1 with the performance results obtained from [29].
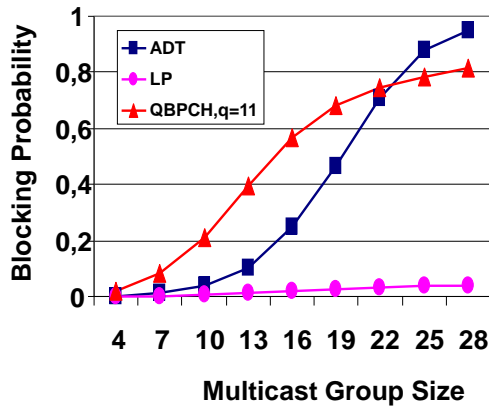

Fig. 1. Comparison of tree-, segment- and cycle-based heuristics, when the PLIs are considered [29].

In this case, multicast protection is the goal and as the figure clearly shows, the absence of protection techniques specifically designed so as to also account for the PLIs (the arc-disjoint tree (ADT) and the Q-Based *P*-Cycles Heuristic (QBPCH) approaches shown in Fig. 1) can result in the utilization of techniques with unacceptable network performance results. Thus, when designing fault recovery algorithms and techniques in transparent networks, clearly all approaches considered must take the QoT issues into consideration (such as the segment-based (LP) technique shown in Fig. 1).

The deployment of translucent network architectures that have been proposed as a compromise between opaque and all-optical networks is another possible solution to this problem. Translucent networks are networks where signal regeneration is performed only at some specific network locations. By performing opto-electronic signal regeneration at some of the intermediate nodes, it is possible to recover the signal degradation due to physical impairments [30], so long as these regenerators are strategically placed in the network.

(b) Shared Protection Considerations

In a number of protection techniques proposed for mesh optical networks, sharability plays an important role, since reducing the redundant capacity in the networks is (together with the recovery speed) one of the main considerations. When sharing of the redundant facilities by a number of different (disjoint) primary paths is used, this means that after the failure event, the affected primary path will capture the previously shared redundant capacity to be used for the failure recovery of its signal.

Running a protection protocol that includes sharing of the redundant network resources requires a certain number of information that is not readily available in core networks and may be difficult to acquire. Initially, the shared risk groups (SRGs) for a network need to be identified so that only primary paths that are not part of the same SRG may share redundant facilities. SRGs express the risk relationship that associates all the optical channels with a single failure as shown in Fig. 2 [23].
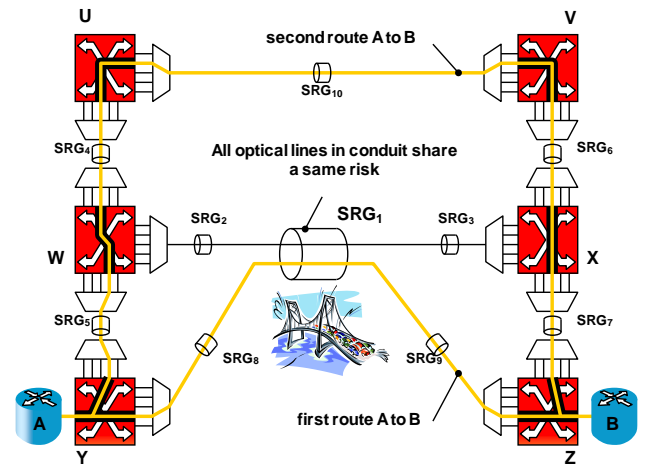

Fig. 2. Shared Risk Groups [23].

Clearly, the network designers or the network operators need to identify and map these SRGs so that the recovery algorithms can correctly calculate SRG-diverse backup paths that will always provide a viable backup route in the event of a failure. However, a network topological view alone does not encompass the notion of SRGs and there is no obvious automated way to generate this information.

Some early work in the literature proposed location-based approaches for SRG auto-discovery and SRG management, to replace the (potentially) erroneous manually maintained databases [31]-[32]. Nevertheless, this is still a practical consideration that must be taken into account.

Another practical problem that needs to be addressed during the design of protection algorithms with sharability functionalities is the issue of the identification of the sharable channels. A channel is sharable, if all SRGs traversed by primary path are not already protected by the channel. This means that (with a deterministic approach) sharing information of every protection channel must be disseminated within the network. However, this information is not easily disseminated using a distributed protocol, thus limiting the feasibility of a deterministic approach to a centralized solution, something that can potentially create scalability problems in the future.

A probabilistic approach was proposed in [33] as a possible solution to this problem that only requires the dissemination of

summarized information (e.g., the limited information can comprise of the number of times each SRG is protected by a channel in the edge, rather than the exact information of SRGs protected by each channel as required by the deterministic approach). Even though performance results in [33] validated this approach, this issue requires further investigation for the practical implementation of shared mesh protection techniques.

(c) Switch Design Considerations for Transparent Nodes in Opaque Networks

Even if all algorithmic, protocol, and design considerations related to fault recovery are successfully addressed, there are still practical considerations for fault recovery that are inherent to the architecture of the switching nodes used in the network.

Consider, for example, an opaque network architecture with a transparent switch (OEO transponders are present at the WDM systems, while the switch is completely transparent). This is potentially the transition architecture on the way to all-optical networks, when the opaque switch approach eventually reaches scaling limitations in signal bit rate, switch matrix port count, and network element cost (even though opaque switches could still remain in the network architecture in order to provide network functions such as grooming and multiplexing). In such a network, architecture practical issues for effective fault recovery implementation are faced due to the lack of transmitters at the optical switch and the lack of direct access to the electrical signal and consequently to the overhead bytes [34].

For example, in opaque switches, an unequipped signal is generated at every idle transceiver on the switch's network-side to prevent alarms in other equipment connected to the switch. In the absence of a keep-alive signal the following problems arise:

- alarms will be generated at the WDM systems that have knowledge of provisioned channels but detect no light on those channels,
- there is lack of monitoring of the protection channels to ensure availability when or if a failure occurs,
- recovery time increases due to the additional time required to turn on the ITU grid WDM lasers.

Even though there has been some recent research activity on all-optical failure localization that does not need any optical layer signaling during recovery [35], lack of access to the overhead bytes still makes fault recovery features such as shared mesh recovery very difficult to achieve in a transparent switch without forfeiting the economies that the switch was designed to extract. Approaches to address these challenges principally consist of:

- using the WDM transponders and client equipment (where OEO is performed) as proxies for the opaque functionality,
- using out-of-band signaling between the WDM systems and the transparent switch, and between the switch and the client equipment (something that requires vendor interoperability and standardization),

- using a bank of a few lasers at each node to address the unequipped signal generation problem [34].

However, even though most of the issues can be addressed via clever innovation as well as standardization efforts, they are still practical considerations that must be taken into account during design and implementation of fault recovery techniques.

3. Assessment of Multiple Failures

Optical networks can be subject to temporally correlated failures. Such failures may affect simultaneously or consecutively multiple network components (being of the same type or not). In this context, analysis of lifetime data can provide key insight on the corresponding failure patterns involving multiple network components (the term component refers here to network and node resources).

The main difficulty in detecting and identifying these failure patterns and their underlying common (or root) cause arises from the limited amount of observation data available to the "decision" entity. Moreover, the spatial distribution of the network components that share common risks is often not directly observable or derivable from the individual failure occurrence that can be observed and corresponding rate that can be derived for each network component taken individually.

The temporal correlation between failures of network components (referred to as joint failure events) implies that these components are spatially or spectrally inter-dependent. The corresponding rate of failure occurrence can be characterized by a joint failure rate following a generalized multivariate distribution, for instance, the Generalized Multivariate Weibull distribution. In this context, the usual task consists in estimating the unknown parameters of such a distribution from observation data.

The applicability of parametric methods is nevertheless often limited (in non-formal terms: more specific/detailed distributions are not robust over time and simpler distributions lead to errors). Consequently, non-parametric alternatives are being thought including Recurrent Data Analysis by means of the Mean Cumulative Function (MCF) and Kaplan-Meier nonparametric maximum likelihood estimation. Nonparametric techniques, even though they provide interesting insight in terms of recurrence rate (slope of the MCF) and estimation of the survivor function, they provide little information for modeling the spatial and/or spectral relationships between the different components of optical networks.

The main question concerning the analysis of lifetime data in optical networks is thus as follows: which technique should be considered in order to model the hidden spatial inter-dependencies among its components? We first assume that information can be progressively inferred from the statistical nature of the failures experienced by the optical network (process referred to as learning from experience) as the input data, i.e., observations, is obtained online from a set of observers or monitoring points. Many classes of statistical learning techniques have been developed and each class

comprises multiple variants. Determining which technique would suit the problem at hand, first requires the examination of the fundamental relationship between the statistical learning technique and the input data properties on which it performs.

On the one hand, (most of) the commonly envisaged statistical learning techniques assume that:

- propositional data are identically and independently distributed ("i.i.d. assumption"), implying that an element in the sequence is independent of the random variables that came before it, and
- random samples of homogeneous data objects result from single relation.

These common assumptions have to be contrasted with the intrinsic properties of real-world data sets, in particular, those characterizing optical networks. These environments are indeed characterized by data that are not identically distributed (heterogeneous) and not independent (multi-relational structures). For instance, a logical link failure can result from spectrally different wavelength failures and each of them can induce spatially different forwarding path failures. In other terms, out-of-the-shelf statistical learning techniques cannot effectively account for the intrinsic properties of the data characterizing these environments.

Filling this gap is the main purpose of Statistical Relational Learning (SRL) [36] which combines i) relational logic learning to model complex relational structures and inter-dependency properties in data with ii) probabilistic graphical models (such as Bayesian networks or Markov networks) to model the uncertainty on the data. The resulting process can perform robust and accurate learning tasks out of multi-relational and inter-dependent data. In the context of optical networks, SRL is of particular interest when considering the objective of learning hidden dependencies between multi-relational, heterogeneous, and semi-structured but also noisy and uncertain data. The shared risk detection problem from link failure observation data provides a representative example of such learning problem.

Indeed, SRL is nowadays applied to social networks' data analysis, hypertext and web-graph mining, etc. It is thus reasonable to also consider its potential in the context of the optical networks' lifetime data analysis and mining. The motivations stem from the fact that the models learned from both intrinsic (propositional) and relational information perform better than those learned from intrinsic information alone. These models offer also better (predictive) accuracy, robustness, and understanding of the relational structures when processing heterogeneous and/or (inter-)dependent data sets. However, this learning technique induces a harder learning task and higher complexity.

In the context of optical networks, applicability to fault diagnosis/root cause analysis such as shared risk detection would provide substantial benefit in increasing reliability of optical routing decisions in particular when (predictive) reactive adaptation of these decisions would be possible in response to (future) environmental changes (or changes in its interacting parts).

(a) Statistical Relational Learning (SRL)

SRL combines probabilistic graphical models (probabilistic learning and inference) to model and reason about uncertainty with representation language to describe relational properties of the data and complex dependencies between them (logical learning and inference). Graphical models provide a principled approach to deal with uncertainty and relational data by means of probability theory. These models represent dependency structure between random variables by joint distributions.

Two types of graphical models are commonly considered: Markov(ian) networks and Bayes(ian) networks. On the one hand, Markov networks, described by undirected graphs, where edges do not carry arrows (no acyclic constraint) and have no directional significance, are useful for expressing symmetric relationships (soft constraints) between random variables. On the other hand, Bayesian networks, represented by Directed Acyclic Graphs (DAG), where edges have a particular directionality indicated by the arrows (acyclic constraint), are useful for expressing causal relationships between random variables.

In addition to the distinction between undirected and directed graphical models, the differentiation between main representation syntaxes, i.e., first-order logic vs. frame-based representation provides a complete categorization of the different SRL models. One distinguishes, as part of the directed models, between rule-based models Bayesian Logic Programs (BLP) [37] and frame-based models Probabilistic Relational Models (PRM) [38] and, as part of undirected models, between frame-based models Relational Markov Networks (RMN) [39] and rule-based models Markov Logic Networks (MLN) [40].

In Section 3(b), we extend the latter, i.e., the MLN model. Selection of this specific learning model stems from the following reasons: it suits control processes whose execution is causality-independent, it is more flexible when the data are made available sequentially (as it is the case when monitoring optical communication networks), and it enables exploiting data sparseness (condition met when only partial data is available upon occurrence of failures).

(b) Incremental Markov Logic Networks (iMLN)

These models have been designed independently on the input data arrival process, i.e., the processing algorithm performs on complete data set (in "batch mode"). However, when performing online learning in optical communication networks, data arrives following different temporal patterns (in "sequential mode") and the model is to be updated as data arrives. As stated in the previous section, this is also one of the main reasons for selecting the Markov Logic Network (MLN) model as it supports sequential data.

For this purpose, we extend the MLN model which represents a probability distribution over possible worlds to cover incremental updates from arrival of input data. An MLN

is formally defined as a set of pairs of formulas F_i in first order logic and their corresponding weights w_i denoted $\{(F_i, w_i)\}$. In first-order logic, formulas are recursively built from atomic formulas (nodes of the Markov network). Each formula F_i has an associated weight w_i: the higher the weight, the greater the difference in probability between a world that satisfies the formula and one that does not, other things being equal. It is important to emphasize that an MLN becomes a Markov network only with respect to a specific grounding and interpretation. Indeed, atomic formulas do not have a truth value unless they are grounded and given an interpretation. Thus, one requires that each node represents a ground atom, i.e., an atomic formula all of whose argument terms contain no variables).

A possible world along with its interpretation assigns a truth value to each possible ground atom: when a world violates one formula, it is less probable although not impossible. The fewer formulas a world violates, the more probable it is.

Together with a set of constants in the domain of discourse, an MLN defines a (ground) Markov network with i) one binary node for each possible grounding of each atomic formula or atom appearing in the MLN (the value of the node is 1 if the ground atom is true and 0, otherwise) and ii) one feature f_i for each grounding of each first-order logic formula F_i in the MLN with the corresponding weight w_i (the value of this feature is 1, if the grounding of the formula is true; 0 otherwise). Each state of the ground Markov network, represented as a log-linear model,) presents a possible world x (i.e., assignment of truth values to all possible nodes or ground atoms. The probability distribution over possible worlds x specified by the ground Markov network probability is given by:

$$P(X = x) = \frac{1}{Z} \exp \left(\sum_{i=1}^{n} w_i n_i(x) \right) \qquad (1)$$

In equation (1), n is the number of formulas in the MLN, the denominator Z denotes the partition function used to make the summation of all possible groundings adding up to 1, w_i is the weight of the formula F_i, and $n_i(x)$ is the number of true (satisfied) groundings for the formula F_i in x. We also operate under the closed world assumption, i.e., if a ground atom is absent in the data, it is assumed to be false.

Besides MLN structure learning (not covered in this section), the main learning task consists in learning MLN weights. Assuming we have at our disposal a given set of formulas $\{F_1, F_2, \ldots, F_n\}$, the learning task consists in finding the respective weights $\{w_1, w_2, \ldots, w_n\}$ These weights can be learned generatively by maximizing the likelihood of one or more "possible worlds" that form training samples. To avoid requiring inference at each step, one can instead obtain the weights w_i from the pseudo-likelihood (PL) approximation of the joint probability distribution of a world x based on its Markov blanket. The Markov blanket (MB) of a node X is the minimal set of variables that must be observed to make this node independent of all other nodes in a model. In an undirected model, such as a Markov network, the Markov blanket includes the node's neighbors in the graph. If x is a possible world and x_k is the k^{th} ground truth value, the PL approximation of x given weights w is provided by:

$$PL_w(X = x) = \prod_{k=1}^{n} P_w(X_k = x_k | MB(X_k)) \qquad (2)$$

The use of the pseudo-likelihood approximation does not require inference at each step and avoids the use of the partition function Z. It is indeed impractical to perform exact inference on large Markov models because of the computations on the partition function Z.

On the other hand, a basic inference task consists in finding the most probable state of the world given the evidence x, i.e., the world in which the sum of the weight of all satisfied groundings is maximized[1]. For this purpose, given the evidence x, it suffices to compute the following [41]:

$$arg \max_{y} P(y|x) = arg \max_{y} \sum_{i} w_i n_i(x, y), \qquad (3)$$

where $n_i(x, y)$ is the number of true groundings of formula F_i involving atoms y.

Computation of equation (3) relies on a weighted SAT solver[2] as corresponding from equation (1) to the weighted MaxSAT problem. In order to find a truth assignment that maximizes the sum of the weights of satisfied formulas, one can use (to avoid local optima while searching) the MaxWalkSAT solver [42], a weighted variant of the WalkSAT stochastic local-search (SLS) satisfiability solver. In order to predict the occurrence of most likely patterns given the observation of certain events (predictive inference problem), it suffices to compute given evidence x, the equation (3) by means of the MaxWalkSAT algorithm.

Another key inference task consists in computing the probability $P(F_i)$ that a given formula F_i holds, given an MLN and possibly a set of one or more formulas as evidence. As the probability of a formula is the sum of the probabilities of the worlds x where it holds ($P(Fi) = \sum_x P(X = x)$), the conditional probability is given by:

$$P(F_i|F_j) = \frac{P(F_i \wedge F_j)}{P(F_j)} = \frac{\sum_{x \in \Xi_i \cap \Xi_j} P(X = x)}{\sum_{x \in \Xi_j} P(X = x)}, \qquad (4)$$

where Ξ_i (Ξ_j) is the set of worlds where F_i (F_j) holds and $P(X = x)$ is given by Eq.1. To avoid exponential time in the number of possible ground atoms, equation (4) can be approximated using probabilistic inference methods like Markov Chain Monte Carlo (MCMC). This algorithm samples a sequence of states according to their probabilities, counts the fraction of sampled states where the formula holds and rejects any state that violates one of them (i.e., it rejects all moves to states where F_j does not hold, and counts the number of samples in which F_i holds). However, as MCMC breaks down

---

[1] In Markov networks, this task is referred to as MAP estimates
[2] A SAT Solver is an algorithm that decides if a propositional logic formula is satisfiable ; if satisfiable it produces an example of a truth assignment that satisfies the formula.

when deterministic or near-deterministic dependencies are present, it is combined with satistifiability testing (by extending the WalkSAT solver) in the MC-SAT inference algorithm [43]. Using this procedure, it is possible to find for instance the probability that the formula F_i holds knowing that the formulas (evidence) F_j and F_k do.

(c) Application to the Shared Risk Detection Problem

Several applications of this learning technique include fault diagnosis and root cause analysis which covers shared risk detection. In the following, we aim at showing whether the MLN technique would be able (or not) to detect among the large possible set of (sometimes hidden) relationships between heterogeneous data which of these relationships characterize shared risks. Indeed, as stated earlier, MLN enables to compactly represent the dependencies between data and relations (compared to the approaches that process them independently) together with collective inference; hence, we could expect that MLN would deliver a more accurate predictive model about possible existence of risks shared by various optical network components. In the simplest instance of this problem, two links ℓ_1 and ℓ_2 share the same risk r (i.e., a physical resource shared by both links).

In case routing decisions would lead to the use of both link ℓ_1 (for the primary path) and link ℓ_2 (for the alternate path), if the resource underlying risk r fails then both paths would fail. For this purpose, assume we have at our disposal the following set of formulas $\{F_1, F_2, F_3\}$:

$F1 : \forall x \; Failure(x) \Rightarrow Alarm(x)$
$F2 : \forall x, y \; Shared\_risk(x, y) \Rightarrow \big(Failure(x) \Leftrightarrow Failure(y)\big)$
$F3 : \forall x, y \; \exists r \; Cross(x, r) \wedge Cross(y, r) \Rightarrow Shared\_risk(x, y)$

The learning task consists in finding the respective weights of these formulas considering we have the following constants in the domain of discourse $x = $ link ℓ_1, $y = $ link ℓ_2, and $r = $ shared risk. The ground Markov network corresponding to this MLN is depicted in Fig. 3, where A stands for Alarm, F for failure, and SR for Shared_risk.
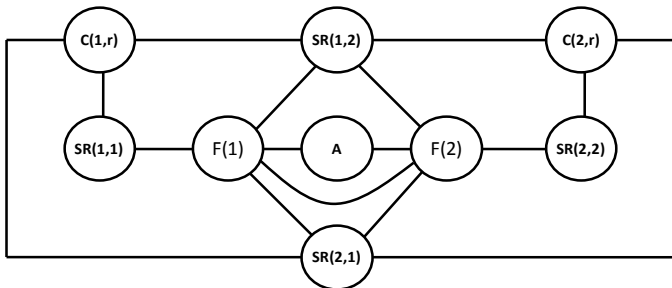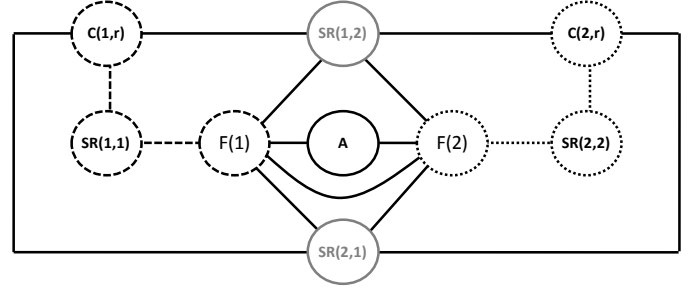

Fig. 3 Ground Markov network.

Fig. 4 Ground Markov network (decomposition).

The decomposition of the ground Markov network into "sub-networks" is represented in Fig. 4. The dashed entities represent the elements detected by one of the monitoring agents and the dotted entities those associated to the other. In the same figure, the grey vertices (SR(1,2) and SR(2,1)) represent the "shared atoms" between sub-networks.

The important characteristic underlying this (ground) Markov network is the following: assuming the ground Markov network is learned incrementally from the information detected by individual agents (each agent being associated to one sub-network), a set of relationships has to be established where the size of this set reflects the number of ground atoms interconnecting these sub-networks.

Further, we would like to predict the probability of occurrence of certain shared failure patterns between paths even though the detection that a given link "crosses" a given risk remains a local decision. To solve this predictive inference problem, it suffices to compute, using the a set of formulas as evidence x, the equation (4) by means of the MC-SAT algorithm. Using this procedure, the MLN model is able to determine for instance the probability that the formula:

$$cross(x, r) \wedge cross(y, r) \wedge cross(z, r) \Rightarrow shared\_risk(x, y, z)$$

holds, given that the formulas

$cross(x, r) \wedge cross(y, r) \Rightarrow shared\_risk(x, y)$ and
$cross(x, r) \wedge cross(z, r) \Rightarrow shared\_risk(x, z)$ do.

Simulation results obtained by running the MLN model for a predictive inference task such as the shared prediction task confirm the inherent problem of decomposing a learning method originally designed to perform on data presenting (hidden) correlations but without accounting for their spatial distribution and sequential arrival. This observation leads to consider the MLN model decomposition for which the corresponding "shared atoms" can themselves correspond to sub-networks to improve the relationship creation process; the latter is of major importance to support scaling with respect to the number of information sources/agents.

An important challenge thus consists in determining the best achievable tradeoff between learning performance (accuracy, sensitivity, specificity, etc.), relationship creation cost, and computational complexity of the SRL methods (in particular, when applied to predictive tasks).

4. Reliability in Elastic Optical Networks – Business Perspective versus Scientific Perspective

Elastic Optical Network (EON) is another interesting recent proposal of an optical technology. The main novelty of EONs compared to currently used Wavelength Division Multiplexing (WDM) technology is the provisioning of sub-wavelength granularity for low-rate transmission and super-channel connectivity for accommodating ultra-high capacity beyond 100 Gb/s. The EON provides a flexible optical grid; specifically, in EONs the frequency spectrum is divided into narrow frequency segments (we refer to them as *slices*) and the optical path (*lightpath*) is determined by its routing path and a *channel*, which consists of a flexibly (ad-hoc) assigned subset of slices around a nominal *central frequency*. EONs use slices of 12.5 GHz width, while WDM applies optical channels of 50 GHz. Therefore, EONs enable much more efficient usage of optical spectrum. Moreover, since in EONs lightpaths can be composed of many neighboring slices, high bit-rates (up to 400 Gb/s) can be achieved, while WDM is currently limited to 100 Gb/s.

There are two key elements in EON architectures, namely, bandwidth-variable transponders (BVTs) and bandwidth variable wavelength cross-connects (BV-WXCs). The role of BVTs is to adapt the client data signal to be sent to/received from the optical network with just enough frequency resources. Simultaneously, BV-WXCs are used to create an optical routing path through the network by switching transmitted signals within their frequency bandwidth to appropriate switch output ports [44]-[46].

In the context of EONs, a new optimization problem arises called Routing and Spectrum Allocation (RSA). The RSA problem comprises of the selection of a routing path for each demand and the assignment of a contiguous fraction of frequency spectrum (slices) to the demand on the selected routing path (spectrum contiguity constraint). The RSA optimization problem is *NP*-hard [47] and it is more difficult than the corresponding Routing and Wavelength Assignment (RWA) problem in WDM networks.

Two basic survivability approaches developed in the context of WDM networks can be also applied in EONs: Dedicated Path Protection (DPP) and Shared Backup Path Protection (SBPP). The former method assumes that each demand is assigned with two link- (or failure-) disjoint paths. Since both working and backup paths are allocated with spectrum resources and the signal is sent on both of them, the reaction to a potential failure is very fast and simple. The SBPP method also uses two disjoint paths. However, in order to reduce the utilization of resources, backup resources can be shared between the demands whose primary paths do not fail at the same time (contrary to DPP, for which each demand has its own backup resources). For more details on DPP and SBPP problems in EONs, refer to [48] and [49].

(a) Practicality from Business and Scientific Perspectives

According to the Macmillan dictionary, *practicality* is defined as "the quality of being effective, useful, or suitable for a particular purpose or situation" [50]. However, the evaluation of features like effectiveness, usefulness, or suitability of a particular method or an approach can be different and depends strongly on background and context.

Here, we want to analyze the practicality of protection methods from two perspectives: business and scientific. In the former case, the practicality is mostly perceived as the possibility to direct application of a particular method in a live production network. According to the business needs and competition on the market, the business parties (e.g., telecoms) prefer to use methods that are widely tested and with proven performance reliability.

Therefore, in some cases, business parties are conservative, that is, they are reluctant to new approaches and methods. An important issue that must be underlined is the cost related to the transition to a new solution. If the expected cost is high, in many cases business parties prefer to keep the old solutions.

Moreover, even if a new solution is finally introduced to the production network, many old procedures and habits are still in use, which can significantly limit the potential benefits of the new approach.

On the other hand, the scientific perspective provides a much more elastic approach, since the constraints that are very strong in the business world can be weakened or even totally removed from consideration. The scientific approach uses as the evaluation environment usually simulation software or some small testbeds. These tools provide much more flexibility in testing the network, compared to the live production network that is providing business services for numerous customers.

Moreover, possible disruptions of the research network do not have so severe consequences as the potential problems that could arise in a real production network. Therefore, some ideas that proved substantial benefits according to a wide range of experiments made in research labs are not always easily acceptable by the business world. However, since the researchers can have a broader view of the whole situation, compared to the business representatives, some innovative concepts can be developed earlier in scientific environments.

(b) Protection of Asymmetric Traffic

In recent years, we have observed a growing attention on services like cloud computing, content delivery networks (CDNs), grids, distributed storage, big data, video on demand (VoD), etc. As these new services are highly bandwidth-demanding, currently used WDM technology in the near future cannot be sufficient to provide the required capacity. Thus, the new concept of EONs must be considered as a possible future optical technology. Moreover, these new services mostly lead to asymmetric traffic patterns in a backbone network [51]. Therefore, in this section we address the protection of asymmetric traffic in the context of EONs. For better illustration of traffic asymmetry in new services, let us consider a CDN system.

A CDN system can be defined as a set of mechanisms that enable delivery of various types of content to end-users on behalf of origin Web servers. The original content is offloaded from source web sites to other content replica servers geographically spread over the network. For each content request, the CDN system selects the best server offering the requested Web page. As a consequence, the CDN delivers the content from a replica server that is much closer to the end user, compared to the original web server. Due to the architecture of the CDN system, the downstream connection from a content server to a client node, usually has much larger bandwidth (capacity), compared to the upstream connection in the opposite direction. Consequently, the traffic between a normal node and the server node is asymmetric [52].

On the other hand, a common practical approach concerning real networks is the assumption that the traffic matrix is symmetric, i.e., demands of the same bidirectional connection between a particular pair of nodes have the same bandwidth in both directions. The capacity is selected as the maximum value between the downstream and upstream connection.

This approach follows from practices from the past, when most of the traffic was symmetric according to the old point-to-point network services. As mentioned previosuly, business parties are in some cases conservative and not willing to introduce new concepts. Moreover, the fixed spectrum grid offered in WDM networks aggravated introduction of asymmetric demands (WDM provides connection of one fixed capacity (e.g., 10 Gb/s), therefore, on default the capacity of both the downstream and upstream connections between the same pair of nodes is the same). However, EON technology now enables quite easily the provisioning of asymmetric lightpaths. As a result, a common assumption regarding symmetric traffic demands in network planning and operation may be costly and not efficient in the new service context [51]-[53].

To verify this issue in the context of survivability of elastic optical networks, simulations were performed using the NSF15 network including 15 nodes and 46 links (Fig. 5).
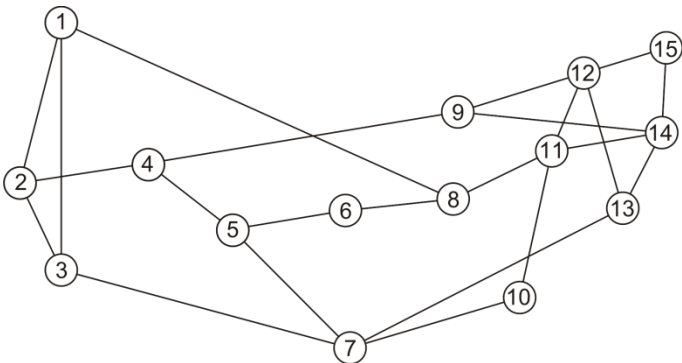

Fig. 5. Topology of NSF15 network.

In order to examine the influence of traffic asymmetry, a parameter called asymmetry ratio (AR) was introduced and applied in the experiments. AR shows the average value of asymmetry between upstream and downstream demands for a particular set of demands. The asymmetry of one pair of demands (downstream and upstream) between the same pair of nodes is defined as $\max(h^{Down}, h^{Up}) / \min(h^{Down}, h^{Up})$, where h^{Down} and h^{Up} denote the requested capacity (bandwidth) for downstream and upstream demands, respectively. The AR parameter is the average value over all demands. Two different lightpath provisioning models in terms of demand provisioning and traffic asymmetry were studied:

- Symmetric (Sym) − for both upstream and downstream demands, the requested bandwidth is equal to the largest value, i.e., $h = \max(h^{Down}, h^{Up})$.
- Asymmetric (Asym) − each demand (upstream and downstream) is established with its original values of requested bandwidth.

Full mesh connectivity demands were generated, with bandwidth requirements created at random but following the selected value of the AR parameter. AR values 1, 2, 3, 4, and 5 were tested and for each value of AR 10 demand sets were generated. The results presented below are average values over these 10 demand sets. To obtain the results, the AFA algorithm proposed and tested in [48] and [49] was applied.

Fig. 6 illustrates the corresponding results, showing the average link spectrum usage. In this case, three protection scenarios are compared: no protection (NP), SBPP, and DPP. For each of the protection approaches, the values obtained for the symmetric and asymmetric cases are shown. Obviously, for AR=1, both Sym and Asym approaches provide the same result.

With the increase of the AR value, it is observed that the symmetric scenario requires up to 30% more optical spectrum, compared to the asymmetric scenario. These results clearly demonstrate the potential gains of the flexible provisioning of bidirectional demands. In particular, if asymmetric demands are established between the same pair of nodes, the spectrum consumption in the network can be substantially reduced.
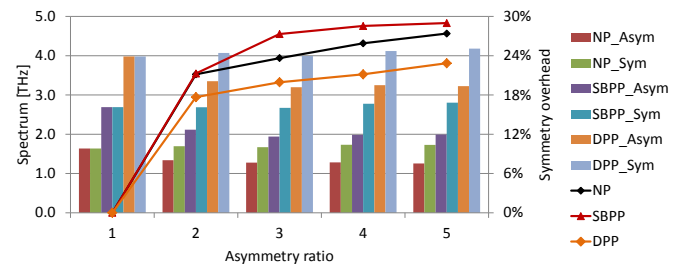

Fig. 6. Performance of various protection and asymmetry scenarios as a function of asymmetry ratio.

Clearly, the concept of Elastic Optical Networks is a very interesting proposal for a new scalable optical transport platform utilizing capacities beyond 100 Gb/s. However, business parties such as telecommunication companies are sometimes conservative and very often during the introduction of new technologies they are accustomed to previous concepts and methods. One of the examples illustrating this issue is the assumption of a symmetric traffic matrix, i.e., demands of the same bidirectional connection between a particular pair of nodes have the same bandwidth in both directions. However, the recent advent of many new services like cloud computing or content-oriented networking leads to a situation where the

traffic is strongly asymmetric. Therefore, in this section, the potential gains of using asymmetry in the context of survivability in EONs were discussed. Simulation results demonstrate that the asymmetry assumption can significantly reduce the optical spectrum usage (up to 30%), when compared to the classical approach with full symmetry of traffic.

4. Multiprovider Multilevel Survivability Challenges

Most research into the dependability of networks has concentrated on a single level (optical restoration or IP-level connectivity) and a single provider (or Autonomous System (AS) in Internet terminology). While these techniques are necessary, they are not sufficient to provide *end-to-end resilience* to users. *Resilience* is the ability of the network to provide and maintain an acceptable level of service in the face of various faults and challenges to normal operation [54], and subsumes the disciplines of survivability (including fault tolerance), disruption tolerance, dependability (including reliability and availability), and performability.

Users of network services are totally unaware of the mechanisms that individual optical or IP service providers use, nor are able to negotiate resilience contracts among service providers. Unless users are communicating across a single provider's network, there are currently no available or deployed resilience, survivability, protection, or dependability services that can provide service assurance to end-users.

As shown in Fig. 7, two users wish a resilient communications path between them. Assume that a single provider is not able to provide this service, and note that even if a single provider's optical and IP paths could be used, high-assurance communication requires multihoming in case a service provider fails. Therefore, each user is multihomed to two service providers, and multiple service providers are required on each path.
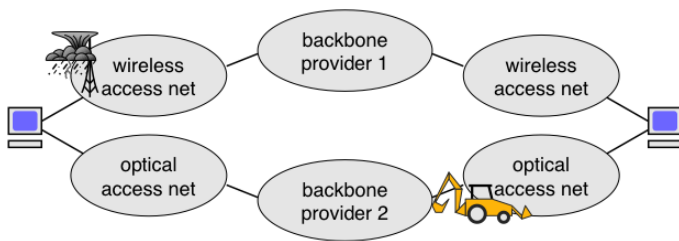

Fig. 7 Diverse communication.

While fault-tolerance requires only redundant components and paths, survivability to correlated failures requires diversity in multiple dimensions: geographic path and medium. Thus, not only should each path be provided by distinct service providers, these paths must be diverse, both in avoiding the use of shared components [55], [56] as well as geographically [56], [57].

While shared link risk groups (SLRGs) [58] prevent multiple paths from suffering failures due to shared fate, this does not work across service providers. For example, the 2001 Howard St. Tunnel fire [59] resulted in the melting of all service providers' fiber optic cables when a CSX Railroad train burned. Thus, we need new models, mechanisms, and multiprovider multilevel capabilities to coordinate dependability among service providers across geographic paths.

An abstraction of the multiprovider multilevel structure of the Internet is shown in Fig. 8. At the top level are users on end-systems (computers or mobile devices) attached to a service provider, as introduced in Fig. 7. The challenge is that the actual structure of the network is much more complex, and the practical aspects of a diversity service consist of getting sufficient *geographic* information at each level and across service providers.
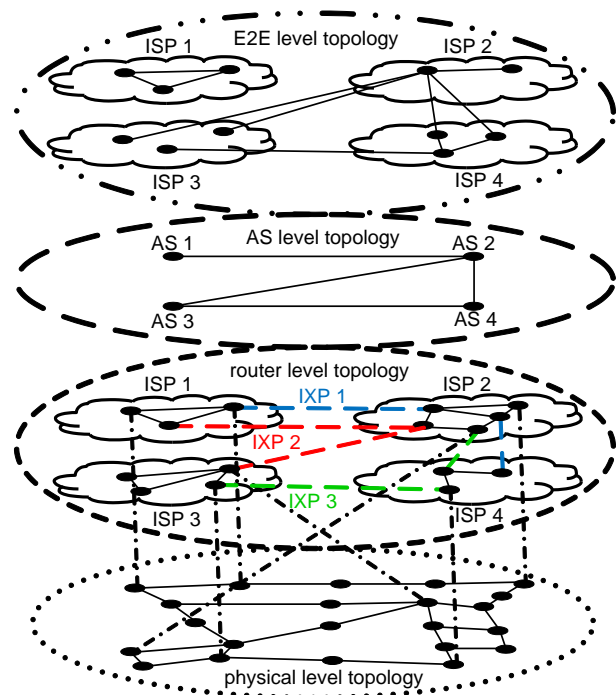

Fig. 8 Multiprovider multilevel Internet structure.

The end-to-end path transits a set of service providers forming an AS (autonomous system) graph; it is possible to get public information about AS connectivity, for example from RouteViews [60]. Each vertex in the AS graph expands to an entire IP-level PoP (point-of-presence) graph, as shown in the third level from the top of Fig. 8.

While the topology of each AS can be inferred using tools such as Rocketfuel [61], inferring the peering points that connect the ASs (dashed lines labeled IXP – interexchange provider) is more difficult. In some cases, likely interconnection can be inferred from membership in public IXPs such as MAE-East and MAE-West, but many peering points are privately linked between service providers [62].

Finally, each service provider is overlaid on its fiber physical graph, shown in aggregate for all service providers. While this information is also generally not made publicly available, third-parties that mine public data (such as from the FCC – Federal Communications Commission and PUCs – public utility commissions in the US) use them to create per-provider fiber maps [63].

Thus, what is needed is an agreed API among service providers to exchange the abstracted geographic information necessary for the cross-layering to provide an inter-provider resilience service to users.

This would permit a geodiverse resilient end-to-end service, as shown in Fig. 9. The application specifies its service (e.g., "real-time critical") and threat model (e.g., "300 km diameter disaster") via cross-layer control knobs to the resilient transport protocol ResTP. ResTP can then determine that the best way to deliver this service to the application is to establish a k=3 multipath transport connection, and request that the geodiverse routing protocol GeoDivRP establishes k=3 distinct paths such that no components are as close as d=300 [km] apart, and may also specify other parameters such as the stretch h (max additional hop count of paths with respect to the shortest) or maximum skew t delay between the paths.
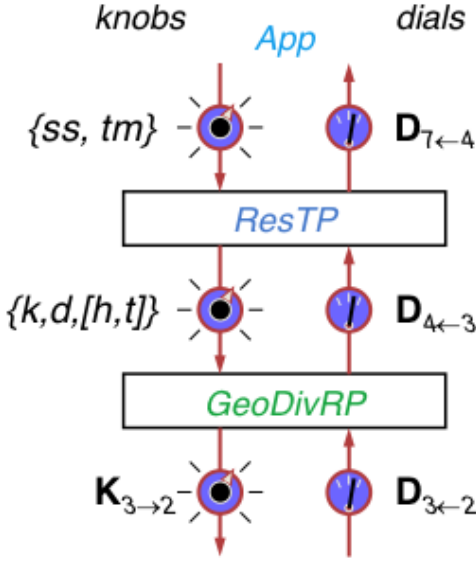
Fig. 9 Cross-layering for geographic resilience.

While we can design and deploy ResTP without the support of service providers, a geodiverse routing protocol has two components. An *intra*domain routing protocol such as GeoDivRP can be deployed unilaterally by a service provider, by adding geolocation of routers to link state routing and a heuristic for creating geographically diverse paths [57], [64] using the cross-layering shown in Fig. 9. However, a geodiverse *inter*domain routing protocol must also present an API between the service providers that exchanges sufficient information to ensure that paths across providers are at least d apart. This is future research, but also provides a greater practical challenge for deployment.

Finally, we need new graph-theoretic complex-system models to analyze the resilience of existing topologies and architectures to attacks based on the structure of the network, as well as to area-based challenges that result from large-scale disasters as given in the previous example.

For this we need a multilevel and multirealm (domain) graph model, as shown in Fig. 10a (only a single provider is shown in this example) [65]. The bottom graph is the physical level; each higher level is an overlay with a subset of corresponding vertices and an arbitrary set of edge connections.
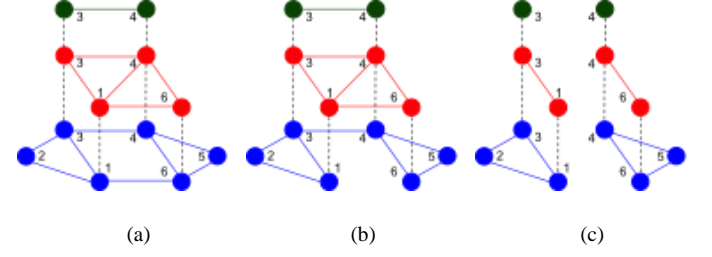

Fig. 10 Multilevel graph model.

The overlay can survive edge deletion, and node deletions (except at a given location when all equipment simultaneously fail or are destroyed) as shown in Fig. 10b, unless the lower level graph is partitioned. In this case, the partition must propagate up the graph levels, as shown in Fig. 10c. By removing the nodes and links *at the correct level for a given challenge* (e.g., physical level for destruction of infrastructure), and propagating up to the user level, we can determine the resilience of the entire network.

These deletions are either based on the structural properties of the network based on an attacker going after the most important nodes (e.g., based on degree or betweenness – the number of shortest paths passing through a node or link), or within a given area affected by a large-scale disaster [66]. Alternatives can then be tested that add links and nodes under cost constraints to increase the resilience of the network with the least cost, and protocol mechanisms that increase resilience [67].

6. Conclusions

In this paper, we focused on deployment issues of survivability and dependability mechanisms in optical networks. In particular, we discussed the current challenges concerning implementation of failure recovery techniques, as well as the related problems following e.g., from hardware constraints, and across multiple providers and levels.

Although failures are inevitable, resilience, survivability, and fault recovery mechanisms can provide efficient means to provide uninterrupted service. However, as we discussed in this paper, the problem of fault detection, localization, and recovery becomes more difficult for the scenario of all-optical forwarding due to lack of electronic processing of a signal by the transit nodes of the lightpath.

In the latter part of the paper, we concentrated on fault recovery issues for the new concept of Elastic Optical Networks, being in our opinion an important direction of optical networks evolution. Finally, we discussed issues related to fault recovery in multilevel multiprovider environments, where e.g., due to ownership-related problems, fault recovery strategies can be confined only to certain network areas.

References

1. A. Haider, R. Harris, "Recovery techniques in next generation networks," *IEEE Communications Surveys and Tutorials*, vol. 9, no. 3, pp. 2-17, 2007.

2. J. Rak, "Fast service recovery under shared protection in WDM networks," *IEEE/OSA Journal of Lightwave Technology*, vol. 30, no. 1, pp. 84-95, 2012.

3. J. Rak, "*k*-Penalty: A novel approach to find *k*-disjoint paths with differentiated path costs," *IEEE Communications Letters*, vol. 14, no. 4, pp. 354-356, 2010.

4. J. Rak, W. Molisz, "A new approach to provide the differentiated levels of network survivability under a double node failure," *Proc. 11th International Conference on Transparent Optical Networks (ICTON)*, pp. 1-4, 2009.

5. J.P. Vasseur, M. Pickavet and P. Demeester, *Network recovery, protection and restoration of optical, SONET-SDH, IP and MPLS*, Morgan Kaufmann Series in Networking, Elsevier, 2004.

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

7. D. Staessens, K. Manousakis, D. Colle, U. Mahlab, M. Pickavet, E. Varvarigos, and P. Demeester, "Failure localization in transparent optical networks," *Proc. RNDM*, pp.589-594, Oct. 2010.

8. O. Gerstel, M. Jinno, A. Lord, and S. J. B. Yoo, "Elastic optical networking: A new dawn for the optical layer?," *IEEE Communications Magazine*, vol. 50, no. 2, pp. 12-20, 2012.

9. IEEE/IFIP 5th International Workshop on Reliable Networks Design and Modeling, Almaty Kazakhstan, Sep. 10-12, 2013, *http://www.rndm.pl/2013/*

10. R. Ramaswami and K. N. Sivarajan, *Optical Networks, A Practical Perspective*, 2nd ed., Morgan Kaufmann, New York, 2001.

11. N. Skorin-Kapov, O. Tonguz, and N. Puech, "Towards efficient failure management for reliable transparent optical networks," *IEEE Communications Magazine*, vol. 47, no. 5, pp. 72-79, May 2009.

12. A.T. Bouloutas, S. Calo, and A. Finkel, "Alarm correlation and fault identification in communication networks," *IEEE Transactions on Communications*, vol. 42, no. 2/3/4, pp. 523-533, February/March/April 2004.

13. K. Azadet, E. F. Haratsch, H. Kim, F. Saibi, J. H. Saunders, M. Shaffer, L. Song and M.-L. Yu, "Equalization and FEC techniques for optical transceivers," *IEEE Journal of Solid-State Circuits*, vol. 37, no. 3, pp. 317-327, March 2002.

14. D. Katz and D. Ward, "Bidirectional forwarding detection", *IETF RFC5880*, 2010.

15. R. Aggarwal, K. Kompella, T. Nadeau, and G. Swallow, "Bidirectional forwarding detection (BFD) for MPLS label switched paths (LSPs)," *IETF RFC5884*, 2010.

16. I. Katzela and M. Schwartz, "Schemes for fault identification in communication networks," *IEEE/ACM Transactions on Networking*, vol. 3, no. 6, pp. 753-764, 1995.

17. J. Choi, M. Choi and S.H. Lee, "An alarm correlation and fault identification scheme based on OSI managed object classes," *Proc. IEEE Int'l Conference on Communications (ICC)*, pp.1547-1551, 1999.

18. M. Steinder and A.S. Sethi, "Non-deterministic diagnosis of end-to-end service failures in a multi-layer communication system," *Proc. IEEE Int'l Conference on Computer Communications and Networks*, pp. 374-379, 2001.

19. T. Wu and A.K. Somani, "Necessary and sufficient condition for *k* crosstalk attacks localization in all-optical networks," *Proc. IEEE Globecom*, pp. 2541-2546, 2003.

20. J. Tapolcai, B. Wu and P.H. Ho, "On monitoring and failure localization in mesh all-optical networks," *Proc. IEEE Infocom*, 2009.

21. M. Bourouha, *et al.*, "Advances in optical switching and networking: Past, present, and future," *Proc. IEEE SoutheastCon*, Columbia, SC, Apr. 2002.

22. P. Soproni, T. Cinkler, and J. Rak, "Methods for physical impairment constrained routing with selected protection in all-optical networks," *Telecommunication Systems (Springer)*, DOI: 10.1007/s11235-013-9827-6, pp. 1-12, 2013.

23. E. Bouillet, *et al.*, *Path routing in mesh optical networks*, Wiley, 2007.

24. B. Mukherjee, *Optical WDM networks*, Springer, 2006.

25. C. Saradhi and S. Subramaniam, "Physical layer impairment aware routing (PLIAR) in WDM optical networks: Issues and challenges," *IEEE Communications Surveys Tutorials*, vol. 11, no. 4, pp. 109–130, Dec. 2009.

26. L. Gillner, "Transmission limitations in the all-optical network," *Proc. 22nd European Conf. on Optical Commun.*, Sep. 1996.

27. J. Simmons, "On determining the optimal optical reach for a long-haul network," *IEEE/OSA Journal of Lightwave Technology*, vol. 23, no. 3, pp. 1039–1048, March 2005.

28. S. Azodolmolky, *et al.*, "Impairment-aware optical networking: A survey", chapter in book titled *WDM systems and networks: modeling, simulation, design and engineering*, N. Antoniades, G. Ellinas, and I. Roudas (Eds), Springer, November 2011.

29. T. Panayiotou, *et al.*, "Segment-based protection of multicast connections in metropolitan area optical networks with quality-of-transmission considerations," *IEEE/OSA Journal of Optical Commun. and Networking*, vol. 4, no. 9, pp. 692-702, 2012.

30. N. Sambo, *et al.*, "GMPLS-controlled dynamic translucent optical networks," *IEEE Network*, vol. 23, no. 3, pp. 34–40, May 2009.

31. P. Sebos, *et al.*, "Auto-discovery of shared risk link groups", *Proc. IEEE/OSA Optical Fiber Commun. Conf.* Anaheim, CA, 2001.

32. P. Sebos, *et al.*, "Effectiveness of shared risk link group auto-discovery in optical networks," *Proc. IEEE/OSA Optical Fiber Commun. Conf. (OFC)*, Anaheim, CA, March 2002.

33. E. Bouillet, *et al.*, "Stochastic approaches to compute shared mesh restored lightpaths in optical network architectures," *Proc. IEEE Infocom*, New York, NY, June 2002.

34. G. Ellinas, *et al.*, "Network control and management challenges in opaque networks utilizing transparent optical switches," *IEEE Communications Mag.*, vol. 42, no. 2, pp. S16–S24, 2004.

35. J. Tapolcai, P-H. Ho, P. Babarczi, and L. Ronyai, "On signaling-free failure dependent restoration in all-optical mesh networks", *IEEE/ACM Transactions on Networking*, pp. 1-12, 2014.

36. L. Getoor, D. Jensen, (Eds.). *Proc. of AAAI-2000 Workshop on Learning Statistical Models from Relational Data*, Austin, TX, AAAI Press, 2000.

37. M. des Jardins and M.E. Gaston, "Speaking of relations: Connecting statistical relational learning and multi-agent systems", *Proc. ICML Workshop on Open Problems in Statistical Relational Learning*, Pittsburgh, PA, 2006.

38. K. Kersting and L. De Raedt, "Adaptive Bayesian logic programs," *Proc. Eleventh Conference on Inductive Logic Programming (ILP)*, LNCS, vol. 2157, 2001.

39. D. Koller and A. Pfeffer, "Probabilistic frame-based systems," *Proc. of 15th National Conference on Artificial Intelligence (AAAI)*, Madison, WI, pp.580-587, 1998.

40. M. Richardson and P. Domingos, "Markov logic networks," *Machine Learning*, vol. 62, no. 1-2, pp. 107–136, 2006.

41. P. Taskar, P. Abbeel, and D. Koller, "Discriminative probabilistic models for relational data," *Proc. of 18th Conference on Uncertainty in Artificial Intelligence (UAI)*, pp.485–494, Edmonton (AB), Canada, 2002.

42. H. Kautz, B. Selman, Y. Jiang. "A general stochastic approach to solving problems with hard and soft constraints", In The Satisfiability Problem: Theory and Applications, pp.573-586, American Mathematical Society, New York (NY), 1997.

43. H. Poon and P. Domingos, "Sound and efficient inference with probabilistic and deterministic dependencies," *Proc. Twenty-First National Conference on Artificial Intelligence (AAAI)*, 2006.

44. O. Gerstel, M. Jinno, and A. Lord, "Elastic optical networking: A new dawn for the optical layer?", *IEEE Communications Magazine*, vol. 50, no. 2, pp. 12-20, 2012.

45. M. Jinno, H. Takara, B. Kozicki, Y. Tsukishima, Y. Sone, and S. Matsuoka, "Spectrum-efficient and scalable elastic optical path network: Architecture, benefits, and enabling technologies," *IEEE Communications Magazine*, vol. 47, no. 11, pp. 66-73, 2009.

46. M. Klinkowski and K. Walkowiak, "On advantages of elastic optical networks for provisioning of cloud computing traffic," *IEEE Network Magazine*, vol. 27, no. 6, pp. 44-51, 2013.

47. M. Klinkowski and K. Walkowiak, "Routing and spectrum assignment in spectrum sliced elastic optical path network," *IEEE Commun. Lett.*, vol. 15, no. 8, pp. 884-886, 2011.

48. M. Klinkowski and K. Walkowiak, "Offline RSA algorithms for elastic optical networks with dedicated path protection consideration," Proc. of RNDM, pp. 1-7, St. Petersburg, Russia, 2012.

49. K. Walkowiak and M. Klinkowski, "Shared backup path protection in elastic optical networks: Modeling and optimization," Proc. of RNDM, Budapest, Hungary, 2011.

50. *http://www.macmillandictionary.com/dictionary/british/practicality*

51. E. Palkopoulou *et al.*, "Traffic models for future backbone networks - A service-oriented approach," *Eur. Trans. on Telecomm.*, vol. 22, no. 4, pp. 137-150, 2011.

52. K. Walkowiak, "Anycasting in connection-oriented computer networks: models, algorithms and results," *Int. J. of Appl. Math. and Comp. Sci.*, vol. 20, no. 1, pp. 207-220, 2010.

53. L.M. Contreras *et al.*, "Toward cloud-ready transport networks," *IEEE Comm. Mag.*, vol. 50, no. 9, pp. 48-55, 2012.

54. J.P.G. Sterbenz, D. Hutchison, E.K. Çetinkaya, A. Jabbar, J.P. Rohrer, M. Schöller, and P. Smith, "Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines," *Computer Networks*, Elsevier, vol. 54, no. 8, pp. 1245–1265, 2010.

55. J.P.G. Sterbenz, D. Hutchison, E.K. Çetinkaya, A. Jabbar, J.P. Rohrer, M. Schöller, and P. Smith, "Redundancy, diversity, and connectivity to achieve multilevel network

resilience, survivability, and disruption tolerance," *Telecommunication Systems Journal*, Springer, pp. 1-15, DOI 10.1007/s11235-013-9816-9, 2014.

56. J.P. Rohrer, A. Jabbar, and J.P.G. Sterbenz, "Path diversification for future Internet end-to-end resilience and survivability," *Telecommunication Systems Journal*, Springer, pp. 1-19, DOI 10.1007/s11235-013-9818-7, August 2013.

57. Y. Cheng, J. Li, and J.P.G. Sterbenz, "Path geo-diversification: Design and analysis," *Proc. IFIP/IEEE Fifth International Workshop on Reliable Networks Design and Modeling (RNDM)*, Almaty, Kazakhstan, pp. 11–18, 2013.

58. S. De Cnodder *et al.*, "Exclude routes – extension to Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", *IETF RFC 4874*, April 2007.

59. H.C. Styron, "CSX Tunnel Fire: Baltimore, MD*"*, *US Fire Administration Technical Report USFA-TR-140*, Federal Emergency Management Administration, Emmitsburg, MD, 2001, available from *http://www.usfa.dhs.gov/downloads/pdf/publications/tr-140.pdf*

60. Route Views project, *www.routeviews.org*

61. N. Spring, R. Mahajan, D. Wetherall, and T. Anderson, "Measuring ISP topologies with Rocketfuel", *IEEE/ACM Transactions on Networking*, vol. 12, no. 1, pp. 2–16, 2004.

62. *https://www.peeringdb.com*

63. KMI Corporation, "North American Fiberoptic Long-haul Routes Planned and in Place," 1999.

64. Y. Cheng, M. Todd Gardner, J. Li, R. May, D. Medhi, and J.P.G. Sterbenz, "Optimised heuristics for a geodiverse routing protocol," *Proc. 10th IEEE/IFIP International Conference on Design of Reliable Communication Networks (DRCN)*, Ghent, Belgium, April 2014.

65. E.K. Çetinkaya, A.M. Peck, and J.P.G. Sterbenz, "Flow robustness of multilevel networks," *Proc. 9th IEEE/IFIP International Conference on Design of Reliable Communication Networks (DRCN)*, Budapest, Hungary, pp. 274–281, 2013.

66. E.K. Çetinkaya, D. Broyles, A. Dandekar, S. Srinivasan, and J.P.G. Sterbenz, "Modelling communication network challenges for future Internet resilience, survivability, and disruption tolerance: A simulation-based approach", *Telecommunication Systems Journal*, Springer, vol. 52, no. 2, pp. 751–766, 2013.

67. M.J.F. Alenazi, Egemen K. Çetinkaya, and James P.G. Sterbenz, "Network design and optimisation based on cost and algebraic connectivity," *Proc. IFIP/IEEE Fifth International Workshop on Reliable Networks Design and Modeling (RNDM)*, Almaty, Kazakhstan, September 2013.