Toward Efficient Failure Management for Reliable Transparent Optical Networks

Nina Skorin-Kapov, University of Zagreb
Ozan K. Tonguz, Carnegie Mellon University
Nicolas Puech, Télécom ParisTech

ABSTRACT

Security and reliability issues are of utmost importance in transparent optical networks due to the extremely large fiber throughput. Fast and successful reaction and restoration mechanisms performed by failure management can prevent loss of large amounts of critical data, which can cause severe service disruption. In this article we discuss failure management issues in TONs, the mechanisms involved, and optical monitoring techniques. Furthermore, we propose applying structural properties of self-organizing systems to create a "small world" hybrid supervisory plane that can enable faster system-wide communication. We also investigate the possibility of a scale-free structure aimed at improving robustness in the network and propose various topology generation algorithms.

INTRODUCTION

The rapid growth of data traffic, primarily Internet traffic, in the past several years is driving the demand for high-speed communication networks. Optical networks based on wavelength-division multiplexing (WDM) have been established as the most promising solution for satisfying the ever increasing capacity requirements in telecommunication networks. WDM is a technology that can exploit the large potential bandwidth of optical fibers by dividing it among different wavelengths. Transparent optical networks (TONs) are dynamically reconfigurable WDM networks that establish and tear down all-optical data connections, called *lightpaths*, between pairs of nodes. These connections can traverse multiple links in the physical topology, and yet transmission via a lightpath is entirely in the optical domain. The reliability of such networks is critical since a single failure can cause tremendous data loss. Although transparency has many attractive features, such as speed and insensitivity to data rate and protocol format, it introduces several vulnerabilities to security. Optical performance monitoring is much more difficult since it must be performed in the optical domain.

A failure management system is used to deal with failures in the TON, which could be due to either component faults or deliberate attacks that aim to disrupt the proper functioning of the network. Due to the transparency inherent in TONs, nodes do not have access to service-bearing wavelengths except where data lightpaths terminate. Thus, management and control information is carried over a separate supervisory wavelength that is optoelectronically processed at each node [1]. We refer to this interconnection of supervisory channels as the supervisory plane. In case of failure, failure management receives alarms from the monitoring equipment available (via the supervisory plane), and then attempts to locate and isolate the source. Meanwhile, the source and destination nodes of failed lightpaths are notified of the failure, after which they launch their restoration mechanisms. In this article we propose creating a hybrid supervisory plane whose structure is such that it can speed up and improve critical security information exchange, and thus improve the network's ability to reconfigure and reestablish communication in the presence of failures. We propose adding a set of long-range supervisory lightpaths in addition to the point-to-point channels between physically neighboring nodes, aimed at creating a small-world scale-free topology.

The small-world and scale-free properties are structural properties that have been observed in many self-organizing complex systems. Self-organizing systems are those in which local low-level interactions and processes between individual entities spontaneously achieve global properties with certain functionality. Since structure affects function, these systems often self-organize into structures which enable efficient and successful operation. We propose applying these concepts to TONs in order improve the efficiency of failure management in them. This is a first step in applying self-organization to optical networks. Our ultimate goal, and vision for the future, is to develop a self-healing and self-management

0163-6804/09/$25.00 © 2009 IEEE

approach that will be able to supervise the functioning of TONs in the presence of increasing complexity and unforeseen attacks.

The rest of this article is organized as follows. In the next section we discuss failure management and optical monitoring in TONs. Complex network structures are then discussed. We propose a new hybrid optical supervisory plane and topology generation algorithms. We then discuss numerical results, and finish with some concluding remarks and ideas for future work.

Failure Management Issues in Transparent Optical Networks

Failure Management

Failure management in TONs deals with the countermeasures taken to compensate for vulnerabilities in the network and failures that can occur. Failures can be due to component faults and deliberate attacks on the proper functioning of the TON. Component *faults* include single or multiple component malfunctions that can be a consequence of natural fatigue, improperly installed or configured equipment, or external influence (e.g., power loss). *Attacks*, on the other hand, are deliberate attempts to interfere with the secure functioning of the network. Attacks differ from faults in that they can spread and propagate throughout the network and can appear sporadically. These characteristics make them much harder to locate and isolate. Various attacks have been described in [2, 3]. They most often include jamming and/or tapping legitimate data signals by exploiting component weaknesses such as gain competition in optical amplifiers and crosstalk in switches.

The countermeasures taken by failure management to ensure secure network operation include prevention, detection, and reaction mechanisms [2]. *Prevention* schemes can be realized through hardware (e.g., strengthening and/or alarming the fiber), transmission schemes (e.g., coding schemes), or network architecture and protocols. Detection mechanisms are responsible for identifying and diagnosing failures, locating the source, and generating the appropriate alarms or notification messages to ensure successful reaction. Due to attack propagation capabilities and the constraints inherent in optical performance monitoring, these tasks are more difficult than in electrical networks. Various alarms generated by monitoring equipment, changes in performance trends, and customer call-ins all help to detect failures.

The third aspect of failure management is reaction to failures. *Reaction* mechanisms restore the proper functioning of the network by isolating the failure source, reconfiguring the connections, rerouting, and updating the security status of the network. In order to establish, tear down, and reroute lightpaths in the presence of major traffic changes, new connection requests, and/or unexpected failures, a control plane employing various signaling and routing protocols is maintained in the TON [4]. In case of attacks it is crucial that reaction mechanisms quickly isolate the source to preclude further attacks. Survivability techniques, which are responsible for restoring failed lightpaths, utilize either preplanned backup paths or reactive rerouting schemes [5]. Both techniques require that the source and destination nodes of failed lightpaths be informed of the failure quickly to ensure high restoration speeds.

Optical Performance Monitoring

Failure management mechanisms are highly dependent on alarms received from optical monitoring equipment. Optical monitoring devices that are currently available include optical power meters, optical spectrum analyzers, OTDRs, eye monitors, and others [6]. These devices help monitor passing signals and send alarms if they detect certain suspicious behavior. Various optical monitoring equipment can be used to detect certain failures, but by no means all of them. For example, optical power meters (which monitor changes in the power of an optical signal) can detect component faults or overt in-band jamming, but may not detect sporadic jamming. Some optical monitoring techniques can estimate the bit error rate (BER) without electronically processing the data payload. These methods include using subcarrier multiplexed pilot tones or evaluating histograms derived from eye diagrams. Additionally, some optical components can have monitoring capabilities themselves (e.g., transmitters may send an alarm if their temperature exceeds a given threshold). An excellent survey of optical monitoring techniques can be found in [7]. Due to the high cost of such equipment, it is not realistic to assume all nodes are equipped with full monitoring capabilities. Thus, obtaining monitoring information from nodes with high monitoring capabilities efficiently is critical for successful failure management.

Failure management mechanisms are highly dependent on alarms received from optical monitoring equipment. Optical monitoring devices which are currently available include optical power meters, optical spectrum analyzers, OTDRs, eye monitors, and others.

Structural Properties of Self-Organizing Complex Systems

Until the middle of the 20th century, complex systems were modeled using regular topologies and Euclidian lattices. After the pioneering work of Erdös and Rényi in the 1950s, random graphs became predominant. However, many real-world self-organizing networks, from the collaboration of film actors to biological ecosystems, lie somewhere between order and randomness. These complex networks have been successfully described using the small-world [8] and scale-free [9] models developed in the 1990s. In order to describe these models in more detail, we first define the basic parameters most often used to characterize complex network structures. They are:

- The average path length L: The average hop distance between all pairs of nodes.
- The clustering coefficient C: The typical cliquishness of a local neighborhood. For each node, we find the ratio of edges in its immediate one-hop neighborhood (including itself) to the total possible number of edges in this neighborhood. These values, averaged over all the nodes in the network, define the clustering coefficient C.

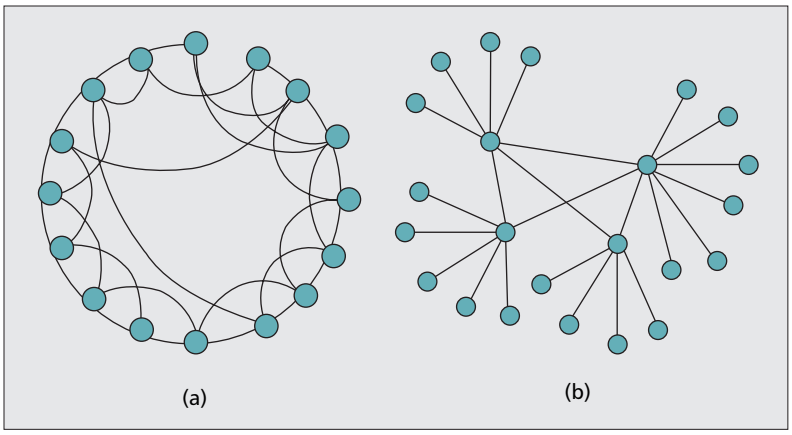

Figure 1. *An example of a) a small-world network generated using the WS small-world generation procedure from [8] where 0 < p << 1; b) a scale-free network.*

- The degree distribution $P(k)$: The probability that a randomly selected node has exactly k neighbors.

Small Worlds

Small-world networks are highly clustered (like lattices), and yet have low average path lengths (like random networks). Watts and Strogatz [8] proposed a rewiring method, which we refer to as the WS algorithm, to generate small-world graphs that can be tuned to lie at various points between regular and random graphs. The algorithm initially starts with a ring lattice and then randomly replaces, or *rewires*, existing links with random ones with probability p. If p is set to 0, the network remains regular. For a probability of $p = 1$, a random graph is created. It has been shown that even for very small p (i.e., a tiny bit of rewiring), the procedure dramatically lowers the average path length with respect to that of a regular lattice, and yet does not significantly affect the clustering coefficient. Thus, a small world is born. An example of a small-world network generated in this manner is shown in Fig. 1a. Such small worlds have Poisson degree distributions that peak at an average degree and then decay exponentially.

The realization that a small world can easily be created by introducing just a few shortcuts between cliques could prove advantageous in the context of communication networks [10]. Namely, applying these concepts has the potential to improve information flow and propagation speed in the Internet, ad hoc networks, and possibly TONs. Intuitively, high-speed shortcuts between distant parts of a network could enable faster system-wide communication, thus aiding dynamic processes such as synchronization, control, and management.

Scale-Free Networks

The characteristic property of scale-free networks is their power law degree distribution. This basically means that there are a few nodes with many neighbors and many nodes with just a few neighbors. An example of a scale-free topology is shown in Fig. 1b. The high-degree nodes are referred to as hubs and basically hold the network together. Barabasi, Albert, and Jeong [9] showed that such power law properties can emerge from stochastic growth and preferential attachment. Basically, as a network grows, new nodes tend to connect to already well connected nodes (the so-called *rich get richer* phenomenon) and thus self-organize into a scale-free state. They propose an algorithm to generate such a network, called the BA algorithm, which initially starts with just a small number of interconnected nodes (m_0). Each new node connects to $m < m_0$ existing nodes, where the probability of connecting to a node is proportional to its degree. Scale-free networks have been shown to be highly robust against accidental failures, but very sensitive to coordinated attacks. Hence, attacking only a few key hub nodes could devastate the entire system, while random failures rarely have a significant effect.

A Hybrid Supervisory Plane for Secure TONs

The Proposed Supervisory Plane

We would like to explore whether the scale-free and small-world models could help to design a more robust TON. We investigated certain small-world characteristics in [11] and further elaborate on them here. Unfortunately, applying these structural models to TONs is not straightforward. If we consider the physical interconnection of optical fibers, which is more lattice-like and clustered due to geographical considerations, utilizing the WS "rewiring" mechanism to achieve a small world is simply not realistic. Rewiring random edges would involve major cost concerns related to digging and laying down new fiber. Furthermore, physical optical networks do not grow continuously at a significant rate since most fiber plants have already laid down large amounts of extra fiber that has not yet been lit for use by Internet service providers and other users of bandwidth. When such networks do grow, fibers and/or nodes are added at locations that best suit the owner of the fiber plant, whose goal is to improve network performance as a whole and not the selfish needs of the newly added node. Thus, growth by preferential attachment to create scale-free topologies, which is the basis of the BA algorithm, may not be applicable.

However, recall that in TONs all-optical connections called lightpaths create a virtual topology over the underlying physical network. This topology is much more flexible and can be dynamically reconfigured, subject to certain constraints. Creating a small-world and/or scale-free topology of data lightpaths independent of the physical interconnection of fibers may be possible. However, in the context of failure management, ignoring the physical topology does not seem logical since optical monitoring information exchange between physical neighbors is crucial. For example, propagating attacks can trigger a large number of redundant alarms, which can often be resolved via communication between physically neighboring downstream and/or upstream nodes.

We propose creating a hybrid supervisory

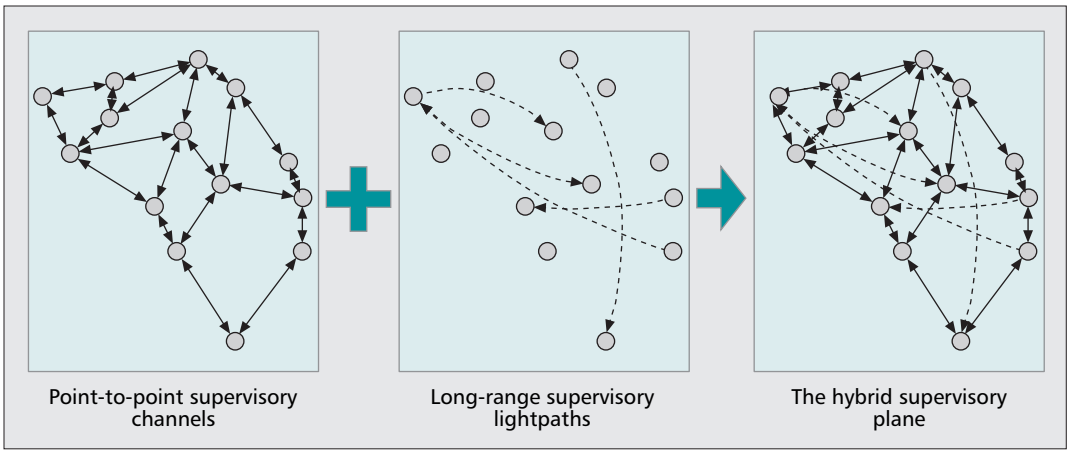

Figure 2. *An example of a hybrid supervisory plane on a reference European core topology.*

plane by maintaining the bidirectional point-to-point *supervisory channels* along each physical link, but also introducing a few long-range supervisory lightpaths between distant nodes. Thus, we could create a small-world supervisory plane, clustered as a result of the physical topology, but with a low average path length due to the small number of transparent shortcuts. An example of such a supervisory plane for a reference European core topology from [12] is shown in Fig. 2. Communication via these lightpaths would be somewhat slower than between physically neighboring nodes due to longer propagation delays, but would still be very fast as a result of their transparency. In addition to the small-world property, these shortcuts could be arranged to yield a scale-free topology that could possibly help create a more robust supervisory plane.

Motivation

The main motivation for creating a small-world supervisory plane is to speed up the exchange of monitoring and control information, particularly in the context of failure management. Our goal is to ensure that the management system receives monitoring alarms and messages as quickly as possible to ensure fast failure detection and localization. Furthermore, we aim to speed up the process of signaling the end nodes of failed lightpaths to start their restoration procedures quickly before triggering higher-level restoration and causing severe data loss and data contention.

Besides the faster exchange of monitoring information, long-range supervisory lightpaths could potentially be used to help nodes with access to local information obtain a better picture of the global network state. In the proposed supervisory plane, "local" information exchange would also include communication between distant parts of the network via virtual shortcuts. Important additional information could be exchanged and merged with local information obtained from physical neighbors to create a more robust network. This information could possibly be used to avoid suspicious parts of the network, help localize attacks, find routes for reconfiguration purposes more quickly, and/or share past experiences and preplanned responses.

Meanwhile, maintaining high clustering in the supervisory plane is desirable in the context of optical monitoring and security to help detect false alarms and resolve redundant ones. Clustered individuals in various self-organizing systems have been known to establish trust easier and communicate more frequently, and thus work together more efficiently [13].

Maintaining high clustering in the supervisory plane is desirable in the context of optical monitoring and security to help detect false alarms and resolve redundant ones. Clustered individuals in various self-organizing systems have been known to establish trust easier and communicate more frequently.

Supervisory Plane Generation Algorithms

In order to generate a supervisory plane with the desired structural properties, we investigated the possibilities of applying various *rewiring*, *preferential attachment*, and *growth techniques*. Topology generation algorithms for wireless networks were proposed in [14].

Preliminaries — We refer to the source nodes of supervisory lightpaths as *informants* since they provide the destination nodes with additional information. It is important to note that not all nodes are equally attractive to use as informants. Nodes with access to more information, better monitoring equipment, and perhaps a good reputation for providing trustworthy and quick responses may provide more reliable information. We define the attractiveness of a node i as an informant to be based on a combination of the following factors:

- The number of data lightpaths that traverse the node, called *transient lightpaths*, denoted DP_i^{tr} after the data plane. Nodes that have more lightpaths passing through them will be able to monitor and analyze more data connections.
- The node's optical monitoring capabilities, denoted Mon_i.
- The number of data lightpaths terminating at the node, denoted DP_i^{dest}. Here the optical data signal is converted into the electrical domain; hence, extensive BER monitoring can be performed.
- The number of data lightpaths originating at the node, denoted DP_i^{source}. Here, the node can obtain detailed information regarding the traffic being sent along these lightpaths.
- The node's in-degree in the supervisory plane, SP_i^{in} (i.e., how well informed it is)
- The node's out-degree in the control plane,

SP_i^{out}, multiplied by a factor α. This is a measure of the node's reputation and desirability among other nodes, which is crucial in growth procedures to enable the rich get richer phenomenon. Varying parameter α allows us to tune the effect of this phenomenon to the desired level.

The overall attractiveness of node i, $A(i)$, is calculated as

$$A(i) = Mon_i DP_i^{tr} + DP_i^{in} + DP_i^{out} + SP_i^{in} + \alpha SP_i^{out} \quad (1)$$

Note that the first element ensures that a node can only provide information regarding transient lightpaths if it employs optical monitoring. Otherwise, data lightpaths simply pass transparently through the node. Herein, we propose four supervisory plane topology generation algorithms:

- *The Random Attachment algorithm*: The Random Attachment (RA) algorithm, inspired by the WS rewiring procedure, considers nodes in random order and chooses for each a random informant. The algorithm terminates after a desired number of shortcuts have been established.
- *The Preferential Attachment algorithm*: Instead of randomly choosing informants to which to attach, considering their attractiveness could prove beneficial. The Preferential Attachment (PA) algorithm selects nodes at random and chooses for each an informant with a probability proportional to its attractiveness. Potential informants are all nodes in the network, except for those that are physically neighboring the node choosing the informant since they are already connected in the supervisory plane via point-to-point supervisory channels. This process is repeated until a desired number of long-range shortcuts are established.
- *The Randomized Preferential Attachment via Growth algorithm*: Since most of the supervisory plane is fixed (i.e., the links corresponding to the physical topology), all the nodes are already included in the topology and thus cannot be grown as in the BA algorithm. However, we can grow an informant web of supervisory lightpaths and then superimpose it onto the physical topology to get our hybrid supervisory plane. The Randomized Preferential Attachment via Growth (R-PAG) algorithm runs as follows. It first chooses a set, m_0, of the most attractive nodes that are not physical neighbors, and interconnects them in an informant web. The algorithm then randomly selects nodes not yet included and assigns to each of them m informants from the existing informant web (provided they are not physical neighbors) with a probability proportional to their attractiveness. This differs from the PA algorithm in that potential informants are only those nodes already included in the informant web. After a desired number of long-range shortcuts are assigned, the algorithm terminates, and the directed informant web is merged with the physical topology to form the supervisory plane.
- *The Ordered Preferential Attachment via Growth algorithm*: Since the informant web grown by the R-PAG algorithm may not include all nodes (depending on the desired number of shortcuts), it may prove beneficial to not only select informants according to their attractiveness, but also select the nodes that choose informants according to their attractiveness. The Ordered Preferential Attachment via Growth (O-PAG) algorithm, like R-PAG, begins by interconnecting m_0 of the most attractive nodes. The algorithm then iteratively selects the most attractive node not included in the informant

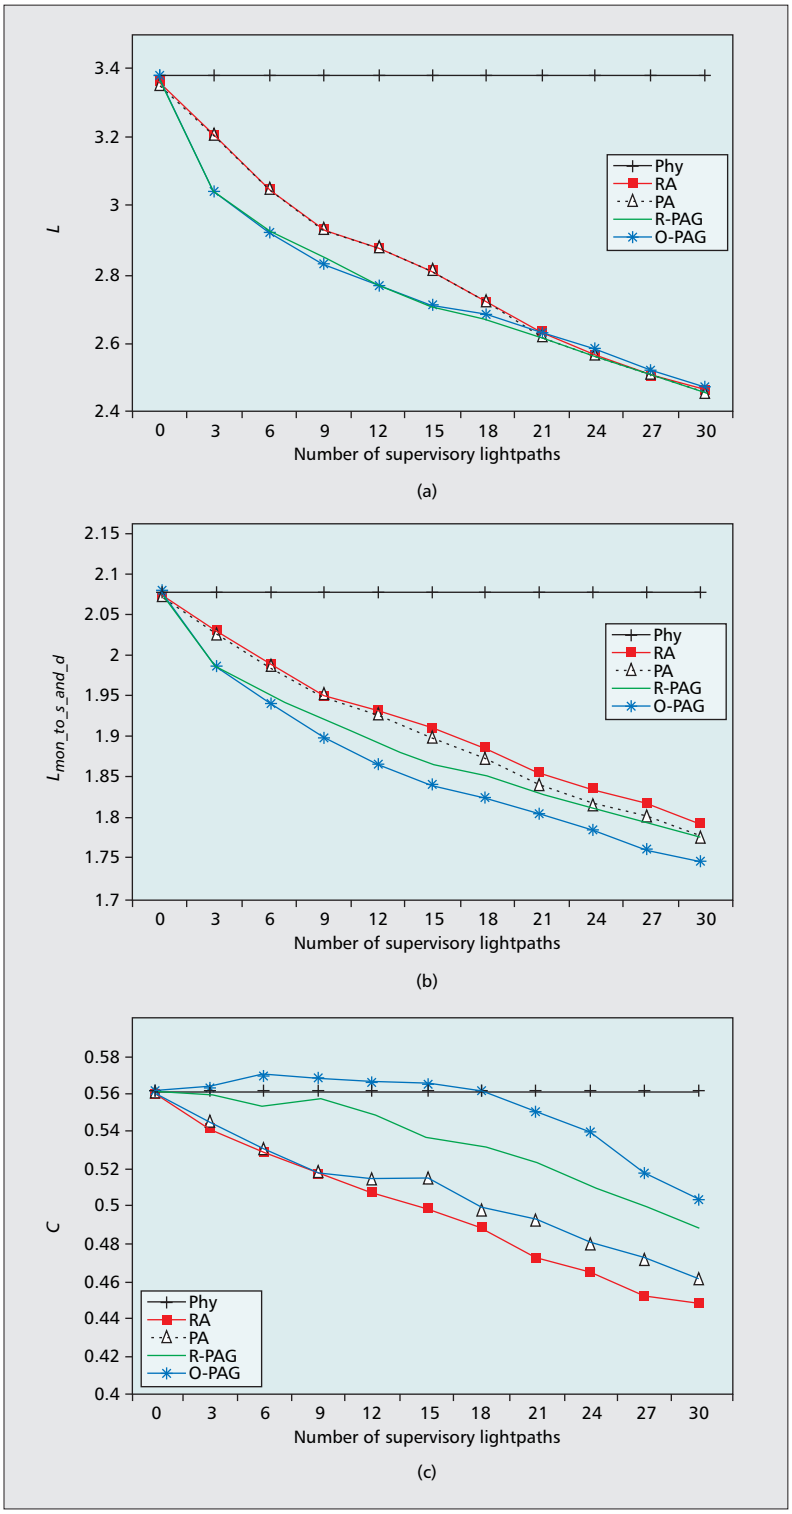

Figure 3. *The average path lengths: a) L; b) $L_{mon\_to\_s\_and\_d}$; c) clustering of the supervisory planes generated by the proposed algorithms and the physical topology for traffic type 1.*

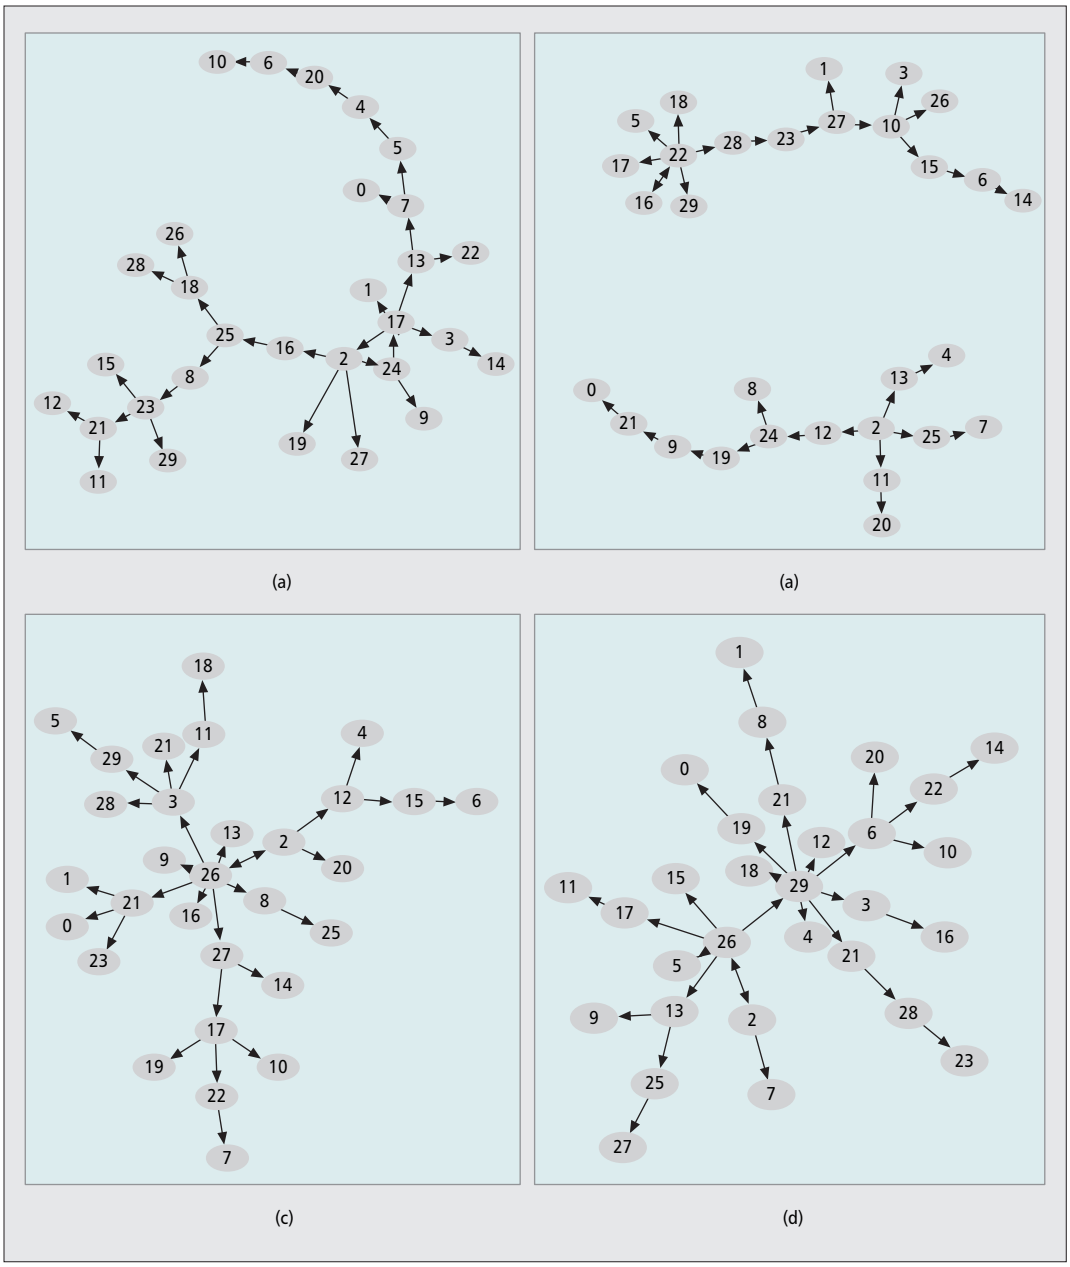

Figure 4. *Sample informant web topologies with 30 shortcuts for traffic type 1 generated by the a)* RA*; b)* PA*; c)* R – PAG*; (d)* O – PAG *algorithms.*

The Random Attachment algorithm, inspired by the WS rewiring procedure, considers nodes in random order and chooses for each a random informant. The algorithm terminates after a desired number of shortcuts have been established.

plane and assigns to it m informants from the existing informant web (provided they are not physical neighbors) with a probability proportional to their attractiveness. The informant web is then superposed onto the physical topology.

Numerical Results

In order to assess the potential benefit of the proposed failure management model and topology generation algorithms, we implemented these four algorithms in C++ and tested them on a reference pan-European topology from the COST Action 266 project [12] with 30 nodes and 48 bidirectional edges. To create a set of data lightpaths, we generated traffic matrices where a fraction F of the traffic is uniformly distributed over $[0, C/a]$, while the remaining traffic is uniformly distributed over $[0, C * \Upsilon/a]$ as in [15]. Here, C represents the lightpath channel capacity, a is an arbitrary integer greater than or equal to 1, and Υ represents the average ratio of traffic intensities between node pairs with high and low traffic values. We ran 25 test cases for three different types of traffic: Traffic type 1 had the values set to $C = 1250$, $a = 20$, $\Upsilon = 10$, and $F = 0.7$, as in [15]. Traffic type 2 considered all traffic to be uniformly distributed over the same value ($\Upsilon = 1$ and $F = 1$), while traffic type 3 had mostly uniformly distributed traffic, but with a few very long bursts ($\Upsilon = 100$ and $F = 0.95$). Lightpaths were then established on the shortest paths between pairs of nodes in decreasing order of their corresponding traffic, with at most five lightpaths originating and terminating at each node.[1] Monitoring capabilities were assigned to

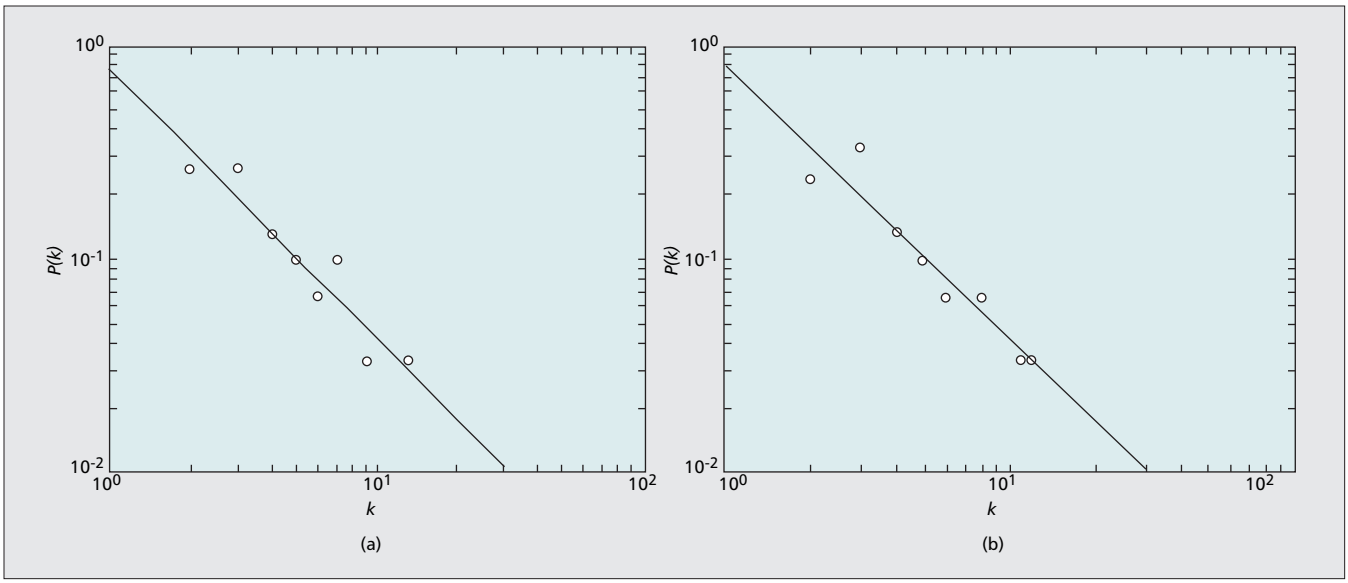

Figure 5. *The out-degree distributions of the supervisory planes generated by the a)* R-PAG*; (b)* O-PAG *algorithms by superposing the informant webs from Fig. 4 onto the physical topology.*

nodes according to the monitoring placement policy described in [16]: if a node is non-monitoring, all its neighbors must be monitoring nodes. Furthermore, if a node is of degree one, its neighboring node must be a monitoring node.

We ran the proposed algorithms for all test cases with the desired number of shortcuts ranging from 0 to 30, in increments of three, assuming that each node could be assigned a maximum of one informant. In the growth algorithms, R-PAG and O-PAG, m_0 was set to 2 and m was set to 1. Various values for α in the attractiveness function were tested. The results shown in Figs. 3, 4, and 5 are those with $a = 10$.

For each test case, we recorded the average path length, L, and the clustering coefficient, C. Since the clustering coefficient is defined on an undirected graph, the supervisory lightpaths were considered undirected in the calculation of C. Furthermore, we found the average path length in hops from each monitoring node to the source and destination nodes of all data lightpaths passing through it averaged over all the monitoring nodes in the network. We refer to this as $L_{mon\_to\_s\_and\_d}$. This is a measure of how fast an alarm can get from a monitoring node to the corresponding end nodes of failed lightpaths to signal that they are to launch their restoration mechanisms. The results, averaged over the 25 test cases for traffic type 1, are shown in Fig. 3. The results are compared with the standard supervisory plane composed of only point-to-point supervisory channels along all physical links, denoted *Phy*. Results for traffic types 2 and 3 are analogous, and are thus omitted for lack of space.

We can see from Figs. 3a and 3b that a significant decrease in the average path lengths L and $L_{mon\_to\_s\_and\_d}$ are already achieved by assigning informants to only 10–30 percent of the nodes, adding only 3–9 long range lightpaths to the fixed 48 bidirectional physical links (i.e., 96 directed point-to-point supervisory channels). Further increasing the number of informants seems inefficient due to the increase in overhead and resources used, as well as the decrease in clustering. When comparing the clustering coefficients in Fig. 3c, we can see that for a small number of informants, a high level of clustering is maintained. In fact, the ordered growth procedure O-PAG actually increased the clustering coefficient for cases with up to 18 extra lightpaths.

In order to determine the kind of patterns generated by the proposed algorithms, we plotted the interconnection of supervisory lightpaths for a large number of informants. An example with 30 lightpaths is shown in Fig. 4. It is evident that the growth algorithms (Figs. 4c and 4d) generate topologies more hierarchical in nature, centered around certain hub nodes. Figure 5 shows the corresponding degree distribution of the supervisory planes generated by the growth algorithms. We can see from the graphs that they are fairly close to following a power law. Although this property may not be very pronounced for a small number of informants, supervisory lightpaths are still centered around a small number of the most attractive nodes. A potential advantage of having such hub nodes in the supervisory plane is robustness to random failure, although it may increase vulnerability to attacks on hubs. Fortunately, hub nodes in our hybrid supervisory plane generated via R-PAG or O-PAG are mainly those with the best monitoring equipment due to the attractiveness function and thus are inherently better protected.

Conclusions and Future Work

As a result of the increasing complexity of transparent optical networks and the tremendous amount of information they carry, efficient failure management is crucial. While transparency offers many advantages, it also imposes various vulnerabilities in optical network security. Self-organizing concepts could possibly be applied to

[1] *We assumed that there were enough available wavelengths on all links.*

develop a highly scalable and robust failure management scheme. Commonly observed structural properties in many self- organizing networks can be described by the small-world and scale-free models. In this article we propose using these models to develop a more efficient supervisory plane to deal with failure management in transparent optical networks. A small-world scale-free supervisory plane could significantly speed up monitoring information exchange and potentially improve reliability. We propose various topology generation algorithms and show how they can achieve the desired structure. After establishing such a supervisory plane, several things will need to be considered in order to design an efficient self-organizing failure management architecture, which is our ultimate goal. Future work will include embedding individual nodes with sufficient intelligence aimed at migrating failure management from its currently centralized form to a more distributed self-organizing approach. This will include developing individual node behavior protocols, defining the content of local information exchange , and introducing mechanisms to establish trust between nodes.

ACKNOWLEDGMENTS

The work described in this article was carried out with the support of the BONE project (Building the Future Optical Network in Europe), a Network of Excellence funded by the European Commission through the 7th ICT-Framework Programme, research project 036-0362027-1641, funded by the Ministry of Science, Education and Sports of the Republic of Croatia, and the HONeDT Cogito Project supported by the Croatian and French governments.

REFERENCES

[1] M. W. Maeda, "Management and Control of Transparent Optical Networks," *IEEE JSAC*, vol. 16, no. 7, 1998, pp. 1008–23.

[2] M. Médard *et al.*, "Security Issues in All-Optical Networks," *IEEE Network*, vol. 11, no. 3, 1997, pp. 42–48.

[3] N. Skorin-Kapov, O. Tonguz, and N. Puech, "Self-Organization in Transparent Optical Networks: A New Approach to Security," *9th Int'l. Conf. Telecommun.*, invited paper, Zagreb, Croatia, 2007, pp. 7–14.

[4] G. Li *et al.*, "Control Plane Design for Reliable Optical Networks," *IEEE Commun. Mag.*, vol. 40, no. 2, 2002, pp. 90–96.

[5] M. Sivakumar, R. K. Shenai, and K. M. Sivalingam, "A Survey of Survivabilty Techniques for Optical WDM Networks," Ch. 3, *Emerging Optical Network Technologies: Architectures, Protocols and Performance*, K. M. Sivalingam and S. Subramaniam, Eds., Springer Science+Media, Inc., 2005, pp. 297–332.

[6] C. Mas, I. Tomkos, and O. Tonguz, "Failure Location Algorithm for Transparent Optical Networks," *IEEE JSAC*, vol. 23, no. 8, 2005, pp. 1508–19.

[7] D. C. Kilper *et al.*, "Optical Performance Monitoring," *J. Lightwave Tech.*, vol. 22, no. 1, 2004, pp. 294–304.

[8] D. J. Watts and S. H. Strogatz, "Collective Dynamics of 'Small-World' Networks," *Nature*, vol. 393, 1998, pp. 440–42.

[9] A.-L. Barabasi and R. Albert, "Emergence of Scaling in Random Networks," *Science*, vol. 286, 1999, pp. 509–12.

[10] J. J. Collins and C. C. Chow, "It's a Small World," *Nature*, vol. 393, 1998, pp. 409–10.

[11] N. Skorin-Kapov and N. Puech, "A Self-Organizing Control Plane for Failure Management in Transparent Optical Networks," *Proc. IWSOS '07*, LNCS 4725, 2007, pp. 131–45.

[12] R. Inkret, A. Kuchar, and B. Mikac, "Advanced Infrastructure for Photonic Networks: Extended Final Report of COST Action 266," Faculty Elec. Eng. and Comp., Univ. of Zagreb, 2003, pp. 19–21.

[13] M. Buchanan, *Nexus: Small Worlds and the Groundbreaking Theory of Networks*, W. W. Norton, 2002, pp. 199–204.

[14] S. Dixit, E. Yanmaz, and O.K. Tonguz, "On the Design of Self-Organized Cellular Wireless Networks," *IEEE Commun. Mag.*, vol. 43, no. 7, July 2005, pp. 76–83.

[15] D. Banerjee and B. Mukherjee, "Wavelength-Routed Optical Networks: Linear Formulation, Resource Budgeting Tradeoffs, and a Reconfiguration Study," *IEEE/ACM Trans. Net.*, vol. 8, no. 5, 2000, pp. 598–607.

[16] T. Wu and A. Somani, "Cross-Talk Attack Monitoring and Localization in All-Optical Networks," *IEEE/ACM Trans. Net.*, vol. 13, no. 6, 2005, pp.1390–1401.

BIOGRAPHIES

NINA SKORIN-KAPOVIS (nina.skorin-kapov@fer.hr) is an assistant professor at the University of Zagreb, Faculty of Electrical Engineering and Computing, Croatia. She received her Dipl.-Ing. (2003) and Ph.D. (2006) degrees in electrical eengineering from the same university, and completed a post-doctoral fellowship at Télécom Paris — École Nationale Supérieure des Télécommunications, France from September 2006 to September 2007. Her main research interests include optimization in telecommunications (particularly in WDM wide-area optical networks), optical networks planning, and security.

OZAN K. TONGUZ (tonguz@ece.cmu.edu) is a tenured full professor in the Electrical and Computer Engineering Department of Carnegie Mellon University (CMU). He currently leads substantial research efforts at CMU in the broad areas of telecommunications and networking. He has published about 300 papers in IEEE journals and conference proceedings in the areas of wireless networking, optical communications, and computer networks. He is the author (with G. Ferrari) of *Ad Hoc Wireless Networks: A Communication-Theoretic Perspective*(Wiley, 2006). His current research interests include vehicular ad hoc networks, wireless ad hoc and sensor networks, self-organizing networks, bioinformatics, and security. He currently serves or has served as a consultant or expert for several companies, major law firms, and government agencies in the United States, Europe, and Asia.

NICOLAS PUECH (npuech@enst.fr) graduated from the École Nationale Supérieure des Télécommunications (Télécom ParisTech), Paris, France, in 1987 as a telecommunications engineer. He received a Ph.D. degree in computer science (honors) in 1992 from the University of Paris 11 and the Habilitation in computer science and networks from the University of Paris 6 (2007). He joined TELECOM ParisTech as an associate professor in 2002. His research interests include network planning and modeling, optimization, and computer algebra. He is a co-author of over 40 papers in journals and international conferences. He has published several books as an author or a translator. He is editor of the IRIS book series published by Springer Verlag.

Future work will include embedding individual nodes with sufficient intelligence aimed at migrating failure management from its currently centralized form to a more distributed self-organizing approach.