A Self-Organizing Control Plane for Failure Management in Transparent Optical Networks

Nina Skorin-Kapov[1,2] and Nicolas Puech[1]

[1]GET / Telecom Paris - LTCI - UMR 5141 CNRS, Networks and Computer Science Department, École Nationale Supérieure des Télécommunications, Paris, France
[2]Department of Telecommunications, Faculty of Electrical Engineering and Computing, University of Zagreb, Zagreb, Croatia
nina.skorin-kapov@fer.hr, npuech@enst.fr

Abstract. Self-organizing systems are present in many areas of nature and science, and have more recently been increasingly applied to telecommunications. These systems often exhibit common structural properties, such as the small-world property, and can react to changes in their environment with no centralized control. With ever-increasing capacity requirements, Transparent Optical Networks (TONs) have been established as the enabling technology for future long-haul high-speed backbone networks. Designing fast security mechanisms is critical, particularly due to the high speeds and transparency inherent in TONs. In this paper, we propose a self-organizing small-world control plane for failure management in TONs, which can improve scalability and adapt to changes in the network.

Key words: Self-organization, small-world phenomenon, transparent optical networks, control plane, failure management

1 Introduction

Self-organization is a phenomenon where low-level interactions between individual entities spontaneously emerge in certain global properties. These so-called 'emergent' properties, which are spontaneously achieved through the selfish actions of individuals, have certain functionality, i.e., fulfill a purpose beneficial for the system as a whole. Common structural properties have been observed in many such systems [1]. One of the most important is the 'small-world' property [2], a term coined to describe networks which are highly clustered with short average path lengths. Self-organizing systems and concepts have been observed in many areas of life and science, from fireflies flashing in perfect synchrony to

This work was supported by a Postdoctoral Research Fellowship from École Nationale Supérieure des Télécommunications, Paris, France. The authors are also grateful to the French and Croatian Governments who supported their work by funding their joint COGITO project HONeDT.

the interconnection of web pages on the World Wide Web [3]. Although self-organizing concepts have not yet been fully exploited in the design and functioning of telecommunication networks, applying these concepts to various areas in communications is currently being intensively researched. Examples include applications in peer-to-peer networks [4], as well as ad hoc and cellular wireless networks [5]. However, to the best of our knowledge, these concepts have not yet been systematically applied and explored in the context of transparent optical networks.

In Transparent Optical Networks (TONs), the physical network consists of an interconnection of optical fibers employing Wavelength Division Multiplexing (WDM). WDM is a technology which can exploit the large potential bandwidth of optical fibers by dividing it among different wavelengths. TONs are dynamically reconfigurable networks where a virtual topology is created over the physical optical network by establishing all-optical 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 making them transparent. In order to provision, establish, maintain, tear down, or reroute lightpaths due to new connection requests, changing traffic, and/or unexpected failures in the network, an optical control plane is maintained employing various routing and signalling protocols [6]. Control information is sent on a separate wavelength than data signals on each link and is electronically processed at each node.

Although the transparency of TONs offers many advantages, such as speed and insensitivity to data rate and protocol format, it makes monitoring much more difficult since it must be performed in the optical domain. Some of the optical monitoring (OPM) equipment and techniques available today include optical power monitors, optical spectrum analyzers, OTDRs (Optical Time Domain Reflectometer), eye monitors, BER (Bit-Error-Rate) estimation techniques, pilot tones, and others [7]. A survey of optical monitoring techniques can be found in [8]. Most OPM equipment generates alarms upon observing suspicious behavior. These alarms can be used to detect certain failures, but by no means all of them. Furthermore, due to the high cost of monitoring 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 necessary to ensure reliable network operation.

A failure management system is employed by the TON to deal with various failures, including both component malfunctions and deliberate attacks. Attacks can be particularly malicious since they can propagate through the network and appear sporadically. Attacks most often include jamming and/or tapping legitimate data signals by exploiting component weaknesses, such as gain competition in amplifiers and crosstalk in switches. Various failures have been described in [9], [10], and [11]. Failure management consists of preventing, detecting, and reacting to such failures. *Prevention* mechanisms, such as strengthening and/or alarming the fiber, are measures taken to prevent failures from occurring. *Detection* mechanisms are responsible for identifying and diagnosing failures according to

the alarms received from monitoring equipment (via the control plane), locating the source, and generating the appropriate notification messages to ensure successful reaction. Methods to locate and recover from various component faults are proposed in [12]. Localization algorithms to help locate the source of various attacks are given in [7], [13], and [11]. Finally, *reaction* mechanisms restore the proper functioning of the network by isolating the source of the failure, reconfiguring the connections, rerouting, and updating the security status of the network [10]. In the presence of attacks, reaction mechanisms should quickly isolate the source to preclude further attacks. Moreover, the source and destination nodes of failed lightpaths need to be notified quickly so they can launch their restoration mechanisms before triggering higher layer restoration. Additionally, efficiently restoring failed lightpaths is crucial due to the hight data rates involved which could potentially lead to huge data loss.

In this paper, we are concerned with the control mechanisms enabling efficient detection and fast restoration in the presence of failures. Namely, when a failure occurs, optical monitoring equipment sends alarms via the control plane to be analyzed by the failure management system. Lightpaths affected by the failure are then restored as quickly as possible, while failure management works on locating, isolating, and repairing the failure. Here, we do not discuss the specific routing or signalling protocols involved, but present a general model for the optical control plane. Namely, we propose a self-organizing scheme to maintain an optical control plane whose structure implicitly enables fast monitoring information exchange for both detection and restoration purposes. The algorithm self-organizes the control plane into a 'small world'. The motivation for this is to reduce the average path length of the control plane to speed up the flow of control information, while maintaining high clustering to improve resiliency to false alarms and the resolution power of true alarms. Simulations show that the proposed scheme significantly reduces the average path length while maintaining fairly high clustering, and can adapt to changes in the network in a self-organized manner.

The rest of this paper is organized as follows. First, Sec. 2 gives an overview of the 'small-world' concept. Then in Sec. 3, we propose a self-organizing control plane which is supported by the simulation results presented in Sec. 4. Finally, Sec. 5 concludes the paper.

2 Small Worlds

Up until the 1990's, complex systems were generally modeled using regular and random graphs. However, many real-world self-organizing networks, from the collaboration of film actors to biological ecosystems, lie in between these extremes of order and randomness. Such complex networks have been successfully described using the small world [2] model. The term *small world* is used to describe networks that are highly clustered with short average path lengths. The average path length, L, is a global property describing the typical separation between any two nodes in the network. It is defined as the average hop distance

between all pairs of nodes. The clustering coefficient, C, is a local property describing 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[1]. These values, averaged over all the nodes in the network, define the clustering coefficient, C.

While regular lattices are highly clustered with long average path lengths, and random graphs exhibit low clustering with short average path lengths, small world structures are somewhere in between. Watts and Strogatz [2] proposed a 'rewiring' method, referred to as the WS algorithm, to generate such structures. The procedure initially starts with a ring lattice and then randomly replaces, or *rewires*, existing links with random ones with probability p. It has been shown that even for very small p, i.e., a tiny bit of rewiring, a small world is born. An example of a small world network generated in this manner is shown in Fig. 1.

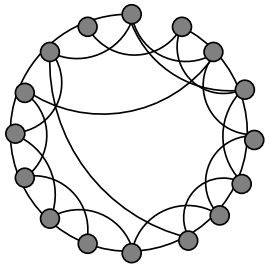

Fig. 1. A small world network generated by the WS procedure where $0 < p << 1$

Applying the small-world concept to communication networks could prove beneficial [14], helping to improve information flow and propagation speed in the Internet, ad hoc networks, and possibly transparent optical networks. 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.

3 The Proposed Self-Organizing Control Plane

The physical topology of the transparent optical network is far from being a random graph since geographic location and installation cost considerations play a major role. The physical topology of the mesh core network is usually more clustered and lattice-like. As already mentioned, a control plane is maintained in the network on a separate supervisory channel on each link. Thus, the control plane topology is equivalent to the physical topology, with point-to-point control lightpaths in each direction between every two physically neighboring nodes. Such a topology can have a fairly high average path length between distant parts of the network, making control information exchange relatively slow. Adding some 'shortcuts' to create a small world can help reduce the average path length.

It is not realistic to add physical long-range links between distant nodes due to the cost of installing fiber and the inherent need for optical regenerators. However, establishing some long-range control *lightpaths* between distant nodes

[1] It is assumed that there can be at most a single edge between a pair of nodes.

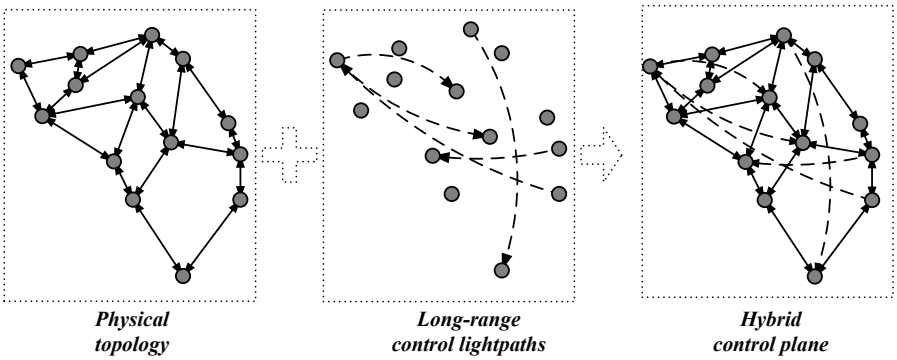

Fig. 2. An example of a hybrid control plane on a reference European core topology from [15].

over the existing physical topology is feasible. Basically, the control plane would be a hybrid control plane composed of point-to-point control lightpaths on each physical link and a set of directed long-range control lightpaths. An example of such a hybrid control plane for a reference European core topology from [15] is shown in Fig. 2. This idea was first introduced in [9] and further developed in [16]. In this paper, we propose a self-organizing scheme to create such a hybrid structure and maintain it in the presence of changes.

To create a small-world control plane topology in a self-organized manner, each node must choose to which distant nodes it wants to be connected to via lightpaths in such a way that their selfish behavior emerges in the desired global structure. Although the physical neighbors are fixed, each node is free to choose its distant neighbors, called 'informants', from which it obtains additional information about other parts of the network. These extra lightpaths are directed, originating at the informant and terminating at the node which chose it.

3.1 The Desired Global Structure

The motivation for creating a small world control plane is to be able to exchange monitoring and control information quickly, particularly in the context of failure management. It is desirable that the management system receive alarms generated from monitoring equipment (via the control plane) as quickly as possible to ensure fast failure detection and localization. In the meantime, it is of utmost importance that lightpaths affected by the failure be restored quickly due to the very high data rates inherent in TONs which can potentially lead to critical data loss causing severe service disruption. Additionally, fast restoration is necessary to ensure that lightpaths are restored before higher layers trigger their own restoration procedures creating a race condition. Failed lightpaths can be restored by utilizing preplanned back-up paths or reactive rerouting strategies. In both cases, the end nodes of the failed lightpath must efficiently be signalled

to handle the failure [17]. Since it is not realistic to assume that extensive optical monitoring is available at each node, failures along a particular lightpath trigger alarms only at a subset of optical monitoring nodes which the lightpath traverses. Thus, it is desirable that the source and destination nodes of lightpaths be well connected to the monitoring nodes they traverse.

Furthermore, clustering in the control 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 [18]. Recall that the physical topology is often highly clustered. By adding long-range control lightpaths, a trade-off is made by slightly decreasing the clustering coefficient in order to significantly lower the average path length. Our goal is to optimize this trade-off by minimizing the drop in clustering while maximizing the decrease in average path length.

In accordance with all of this, we deem the following properties of the control plane as the desired global structural properties.

Low L**,** where L is the average path length in the control plane in terms of hops. (A hop is considered to be a control lightpath.)

High C**,** where C is the clustering coefficient as described in Sec. 2. (Since the clustering coefficient is defined for an undirected graph, the directed long-range control lightpaths are considered undirected in the calculation of C.)

Low $L_{mon\_to\_s}$ **and** $L_{mon\_to\_d}$**,** where $L_{mon\_to\_s}$ and $L_{mon\_to\_d}$ are the average path lengths in hops from each monitoring node to the source and destination nodes, respectively, of all data lightpaths passing through it, averaged over all the monitoring nodes in the network.

3.2 Local Behavior Rules

Our goal is to create and maintain a control plane topology in a self-organized manner where the selfish behavior of individual nodes emerges in the desired global properties. In addition to its fixed physical neighbors, each node can choose distant 'informants' from which it obtains additional information about other parts of the network. Not all nodes are equally attractive to use as informants. Naturally, each node prefers to connect to nodes with access to more information relevant to it. For example, suppose node j has certain monitoring equipment available to monitor lightpaths passing through it. Furthermore, suppose node i happens to be the source node of a lightpath routed via node j. Node i would benefit from having j as an informant because if the monitoring equipment at node j detects a failure, node i could be informed very quickly (in a single hop) and could, thus, launch its restoration mechanism faster.

It is also important that the control plane self-maintains and self-organizes to adapt to changes in its environment. Namely, nodes can change over time causing a shift in the attractiveness of informants. In the presence of traffic changes and/or failures, several data lightpaths could be reconfigured. New monitoring equipment could also be acquired or existing equipment could fail. Furthermore,

informants could acquire a bad reputation after sending false information. Nodes in our control plane can choose new informants, in light of these changes, subject to certain constraints.

All nodes in the network have certain local information available. Each node is aware of all the lightpaths originating at it, terminating at it, and passing through it (called transient lightpaths). A node maintains the following information regarding each lightpath: its source node, its destination node, the wavelength it utilizes, the input port on which it arrives (unless it originates at the node), and the output port on which it is transmitted (unless it terminates at the node). Each node is also aware of the monitoring information available to it. It can have various optical monitoring equipment to monitor passing lightpaths, such as spectrum analyzers or power monitors.

To create and maintain our desired small world control plane, we propose the following self-organizing scheme. Initially, each node chooses one random informant and establishes a corresponding control lightpath. Periodically, each node i sends a *rating_request* message to a random node j in the network demanding its rating. The rating of node j, when requested by node i, represents its attractiveness as a potential informant to node i. We denote this as $Rating(j, i)$ and it depends on both i and j. Upon receiving a rating request, node j returns a *rating_reply* message, whose contents will be described later on. From the information provided in the *rating_reply* message, node i can calculate $Rating(j, i)$. It then compares j's rating to the rating of its current informant. If j's rating is better, it tears down the lightpath connecting it to its current informant and establishes a new lightpath from node j using the signaling protocol employed by the control plane. We set a limit on the maximum number of nodes for which a node can be an informant (i.e., each node has a maximum control plane out-degree) due to the limited resources available at each node. The pseudocode of the local node behavior protocol is shown in Fig. 3.

Node (*i*) Behavior Protocol
Initialization:
$currentInformant := NULL$;
$Rating(NULL, i) := 0$;
Begin:
Periodically, send $rating\_request$ to a random node j;
if received $rating\_request$ from a node k **then**
  Send $rating\_reply$ to k;
end if
if received $rating\_reply$ from node j **then**
  Compute $Rating(j, i)$;
  if $Rating(j, i) > Rating(currentInformant, i)$ **then**
    Tear down control lightpath $(currentInformant, i)$;
    Establish new control lightpath (j, i);
    $currentInformant := j$;
  end if
end if
End

Fig. 3. The pseudocode of the node behavior protocol.

To help describe function $Rating(j, i)$, we define the following parameters.

$Phy_{j,i}$ is a binary parameter indicating if nodes j and i are physical neighbors and, thus, already connected via one-hop lightpaths along the physical link connecting them.

$Free\_Port_j$ is a binary parameter indicating whether there are free resources at node j to handle becoming an informant for a new node.

Mon_j is an integer representing the level of optical monitoring equipment and techniques used at node j. If there is no optical monitoring available, $Mon_j = 0$. With increased equipment and better techniques, the level increases.

$TLPs_j^i$ is an integer which represents the number of transient data lightpaths passing through node j whose source node is node i.

$TLPd_j^i$ is an integer which represents the number of transient data lightpaths passing through node j whose destination node is node i.

$Hops_{j,i}$ represents the length of the shortest path in hops in the physical topology from node j to node i.

CP_j^{in} is an integer representing the in-degree of node j in the control plane topology.

The rating function is then defined as

$$Rating(j,i) = (1-Phy_{j,i}) \cdot Free\_Port_j \cdot [Hops_{j,i} \cdot Mon_j \cdot (TLPs_j^i + TLPd_j^i) + CP_j^{in}]. \quad (1)$$

If nodes j and i are already physical neighbors, i.e., $Phy_{j,i} = 1$, then there is no need for a new control lightpath between them since they are already one-hop away. Thus, rating $Rating(j,i) = 0$. The same is true if node j does not have any free resources (i.e., a free output port) to establish a new control lightpath originating at it. Otherwise, the rating depends on the information that can be obtained from node j which is relevant to node i.

Node j monitors all its transient lightpaths in accordance with the level of optical monitoring capabilities available to it, i.e., Mon_j. If node j detects a failure, it sends an alarm to failure management and the source and destination nodes of the corresponding lightpaths via the control plane. The more lightpaths that pass through node j that happen to have their source or destination at node i, and the better the optical monitoring performed at node j, the more attractive j is as an informant to i.

Furthermore, node i will receive alarm(s) from j in the presence of failure (provided j's monitoring equipment detects it) via the shortest path in the current control plane topology. Thus, the longer this path, the more desirable it is for node i to employ node j as an informant in order to reduce this path. In the $Rating(j,i)$ function, however, the parameter $Hops_{j,i}$ represents the shortest path in the physical topology and not the control plane. The motivation for this is as follows. As the control plane changes over time, the shortest paths between nodes in the control plane also change. Thus, if the shortest path between j and i in the control plane were included in function $Rating(j,i)$, the rating could change due to a shift in the control topology even if there are no significant changes in the network with respect to traffic flows, data lightpaths, monitoring equipment, etc. Since each change in the control plane requires certain signalling

overhead to tear down and establish a new informant, it is not desirable to have frequent modifications. We aim to optimize the trade-off between the stability of the control plane and its ability to adapt to changes in the network. By considering the shortest physical path between nodes j and i in the rating function, the protocol initiates fewer changes and yet often gives a good indication of the distance between the nodes. Essentially, it is a tradeoff between updated information and control overhead. Since the shortest path between two nodes in the physical topology is the longest possible shortest path in the control topology (i.e., adding informants can only lower this path), $Hops_{j,i}$ considers the worst case scenario for the node. Furthermore, preliminary testing indicated that considering the physical shortest path in the rating function, instead of the shortest path in the control plane, lowered L for most cases while performing the same with respect the remaining criteria.

Since nodes in the network maintain only local connectivity information, they do not have knowledge of the shortest paths to all other nodes in either the physical or the control plane topology. Thus, a counter is included in the *rating_reply* message which counts the number of hops for the message to get from node j to node i. In this message, node j provides all the elements required to calculate $Rating(i, j)$, except for $Hops_{j,i}$. Once the message arrives at node i, the final $Rating(i, j)$ is calculated by node i using the information held in the counter and the *rating_reply* message. Since it is not crucial that the periodic updates performed at each node be extremely fast, we send *rating_request* and *rating_reply* messages using only the point-to-point lightpaths in the control plane (and not via informants). The 'shortcuts' in the control plane are reserved only for crucial monitoring information when a failure occurs and are not used up by other less-important signalling and control overhead. This way the counter would calculate the shortest path in the physical topology $Hops_{j,i}$. If we were to define $Hops_{j,i}$ as the shortest path in the current *control plane* topology, then the *rating_request* and *rating_reply* messages could be sent over any link in the control plane.

The last element in the rating function is simply the control plane in-degree of node j. For the case when j has a high monitoring level and many transient lightpaths relevant to node i, this parameter will not significantly affect the rating. However, if two nodes have similar ratings with respect to monitoring transient lightpaths, the node with a higher control in-degree is considered more attractive since it has access to more one-hop control information.

In the approach, we suppose that every node has exactly one 'informant'. This assumption is made for simplicity but need not be so for the general case. Furthermore, we assume nodes have global knowledge of the existence of all other nodes in order to send random *rating_request* messages. Since the physical topology is for the most part fixed[2], this is feasible but limits scalability. We are currently investigating various modifications of the model to deal with these issues.

[2] Changes in the physical topology do not occur very frequently due to the difficulties involved in laying down fiber

4 Numerical Results

In order to evaluate the proposed self-organizing scheme, we developed an event-driven simulator in C++. For simplicity, we assumed that the periodic updates of nodes are performed synchronously. We tested the algorithm on a reference topology of a pan-European basic network from the COST Action 266 project [15] with 30 nodes and 48 bidirectional edges, shown in Fig. 4. We assumed two levels of monitoring, differentiating between non-monitoring nodes ($Mon_j = 0$) and nodes which are equipped with at least some optical monitoring equipment ($Mon_j = 1$). To decide which nodes have optical monitoring equipment, we used the monitoring placement policy described in [13]. According to this policy, 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.

Before running the simulation, an initial virtual topology was created for the data plane as follows. First, a traffic matrix was generated using the method suggested in [19] where a fraction F of the traffic is uniformly distributed over $[0, C/a]$ while the remaining traffic is uniformly distributed over $[0, C * \Upsilon/a]$. The values were set to $C = 1250$, $a = 20$, $\Upsilon = 10$, and $F = 0.7$ as in [19]. To establish the initial virtual topology, we set up lightpaths between pairs of nodes in decreasing order of their corresponding traffic, with at most 5 lightpaths originating and 5 lightpaths terminating at each node, i.e., we assumed 5 transmitters and receivers were utilized per node[3]. Lightpaths were routed on their shortest physical paths, in terms of hops, and we assumed that there were enough available wavelengths on all links.

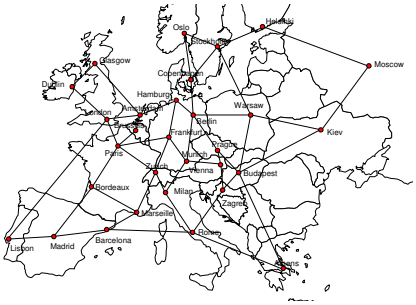

Fig. 4. The European basic network topology, from [15].

In the first simulation scenario, referred to as Scenario 1, requests to tear down the lightpaths comprising the initial virtual topology described above, arrived according to a Poisson process with rate $\lambda = 5$. New lightpath requests also arrived according to a Poisson process with rate $\lambda = 5$, with exponentially distributed holding times with mean $b = 10$. We assumed that the monitoring equipment at nodes was fixed. In this scenario, the values of $TLPs_j^i$ and $TLPd_j^i$ in the informant *Rating* function can change over time while the remaining parameters remain constant.

Simulations were run for 3 cases. In the first case, the control plane topology was kept equivalent to the physical topology with no long-range shortcuts. This is denoted as Phy_CP. In the second case, a hybrid control plane was created at

[3] At most one lightpath was established between the same pair of nodes.

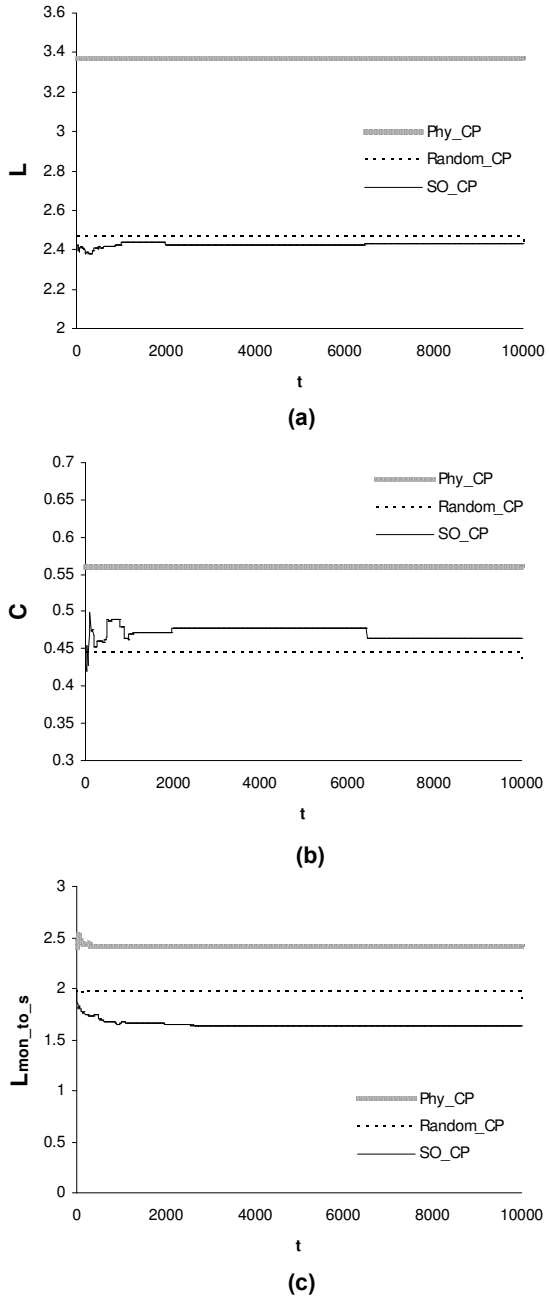

Fig. 5. The average path length in the control topology (a), the clustering coefficient (b), and the average path lengths from monitoring nodes to the source nodes of their transient lightpaths (c) for Scenario 1.

simulation start time by choosing a random informant for each node in the network, establishing the corresponding directed lightpath, and super-positioning it onto the physical topology. This control plane, denoted as $Random\_CP$, was then kept constant throughout the simulation. The third case ran the self-organizing scheme proposed, starting initially with the random control plane topology $Random\_CP$ and then self-organizing to adapt to the network state. This way we could analyze the benefits of the proposed scheme in comparison with the random case employing the same number of 'shortcuts' but self-organizing itself in the presence of changes. The self-organizing control plane for the third test case is denoted as SO_CP.

Each simulation was run for 10000 time units. For the SO_CP algorithm, nodes sent *rating_request* messages to random nodes periodically every 10 time units. Furthermore, every 10 time units we recorded the structural properties of the control plane and the data plane, and calculated the values for L, C, $L_{mon\_to\_s}$, and $L_{mon\_to\_d}$. The results for L, C, and $L_{mon\_to\_s}$ are shown in Fig. 5 in plots (a), (b), and (c), respectively. The results for $L_{mon\_to\_d}$ are analogous to those of $L_{mon\_to\_s}$ and are, thus, omitted for lack of space. We can see from plots (a) and (c) that the average path length of the control plane (L) and the average path lengths from monitoring equipment to the source nodes of transient lightpaths ($L_{mon\_to\_s}$) of the Phy_CP control plane are significantly decreased with the addition of extra long-range control lightpaths (SO_CP and $Random\_CP$). This makes sense since there are an increased number if links in the control plane topology. Naturally, the more lightpaths we add, the lower the average path length. However, it is not desirable to establish too many control lightpaths due to extra overhead and resource consumption. The Self-Organizing Control Plane SO_CP, obtained lower values for L, and even more so for $L_{mon\_to\_s}$ (and $L_{mon\_to\_d}$), than the $Random\_CP$ even though they use the same number of extra long-range lightpaths. With respect to the clustering coefficient C, adding random edges to the control plane naturally decreases clustering to some extent. However, applying the self-organizing scheme caused a smaller drop in clustering than the random case.

Note that it is desirable that there be a minimal number of changes in the control plane due to high control overhead, and yet we want it to achieve the desired global structure even in the presence of changes. To analyze our model, we recorded all changes made to the SO_CP topology during the simulation. Initially, there were 80 changes in the first 2000 time units. However, once the control plane stabilized, it only performed 2 changes from time 2000 until 10000, even though there were 55103 changes in the virtual topology. This shows that learning the location of the monitoring equipment and physical distances between nodes has a more significant impact on the control plane topology than changes in the virtual topology, i.e. the node protocol is more sensitive to variations in $Hops(j,i)$ and Mon_j than the remaining parameters in the $Rating(j,i)$ function. Thus, intense rearrangement of the control plane would more likely occur in the presence of drastic changes in monitoring equipment or the physical topology, rather than the virtual topology. This is very fortunate since monitor-

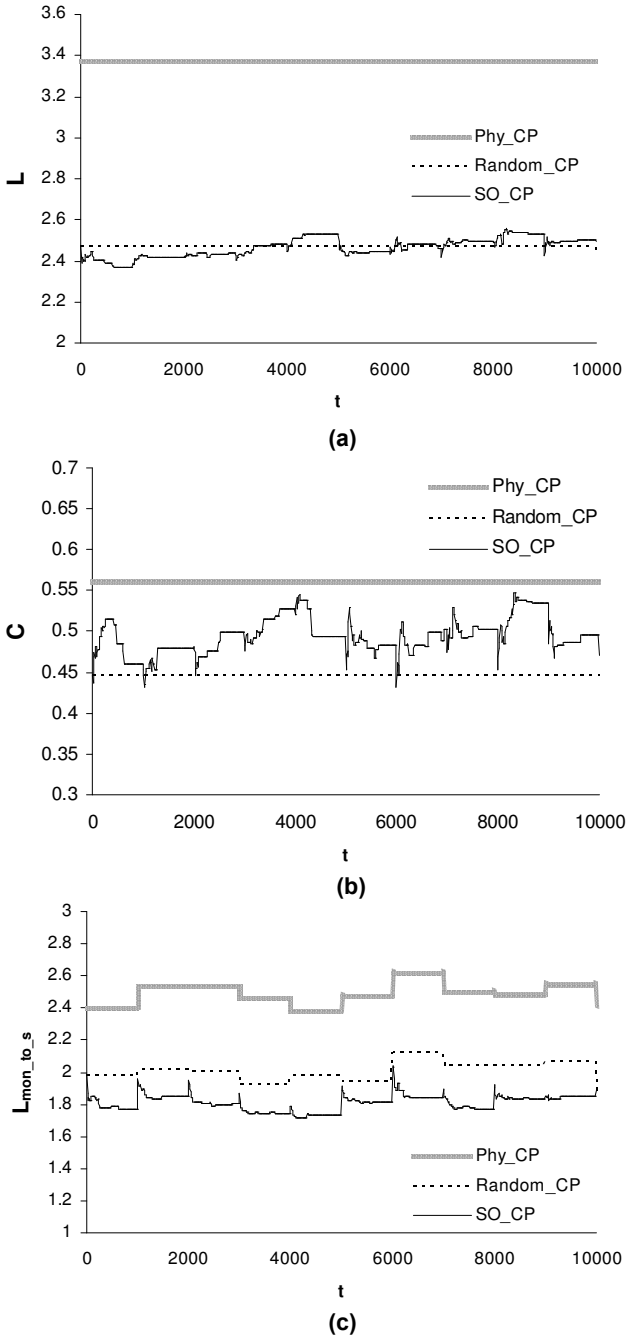

Fig. 6. The average path length in the control topology (a), the clustering coefficient (b), and the average path lengths from monitoring nodes to the source nodes of their transient lightpaths (c) for Scenario 2.

ing equipment at nodes and the physical topology are generally fairly constant and change slowly over time. Thus, the control plane would be quite stable.

To see how SO_CP would adapt to more drastic changes in the network, we created a second simulation scenario, referred to as Scenario 2. Here we ran simulations for 10000 time units where every 1000 time units there were major changes in both the virtual topology and the monitoring equipment. The virtual topology would be completely torn down and the same number of new random lightpaths would be established. Furthermore, the optical monitoring available at each node would fail with probability $P_{mon} = 0.5$, while non-monitoring nodes would gain new optical monitoring equipment with the same probability. This is, of course, a much hyperbolized situation but can help us see how SO_CP can adapt and self-organize into a stable state with the desired global structural properties in the presence of drastic changes. The results of the simulations[4] are shown in Fig. 6.

For the average path length L shown in Fig. 6.(a), SO_CP oscillates around $Random\_CP$ but both remain close and significantly lower than Phy_CP. We can see from Fig. 6.(b) that the clustering coefficient C for the control planes with long-range edges is lower than the physical topology. However, the self-organizing control plane performs better than the random constant one. With respect to the number of hops from optical monitoring equipment to the source of the lightpaths they monitor (Fig. 6.(c)), SO_CP outperformed $Random\_CP$ and Phy_CP in all cases. The situation is analogous for the number of hops from monitoring nodes to the destination nodes of transient lightpaths. When drastic changes occur, SO_CP performs a series of changes to adapt in a self-organizing manner and then stabilizes after achieving the desired properties. We are currently investigating the behavior of the control plane in the presence of node failure and growth of the network with the addition of new nodes or links.

5 Conclusions

In this paper, we propose a self-organizing scheme to create and maintain a hybrid small world control plane for more efficient failure management in transparent optical networks. The motivation for such a control plane lies in the fact that fast detection, localization and restoration in the presence of failures are particularly important in TONs due to very high data rates and their inherent transparency. A small world control plane could significantly speed-up monitoring information exchange and potentially improve reliability. Furthermore, maintaining such a topology in a self-organized manner makes it more scalable and robust to changes in the network. Simulations performed on a reference European topology indicate the benefits of this model. We are currently investigating the possibilities of extending this model with feedback loops to minimize the control overhead incurred by periodic node updates. Furthermore, develop-

[4] The results for $L_{mon\_to\_d}$ are again omitted since they are analogous to those of $L_{mon\_to\_s}$.

ing trust models to establish trust between nodes and the exchange of reputation information could prove beneficial.

References

1. Strogatz, S.H.: Exploring Complex Networks. Nature **410** (2001) 268–276
2. Watts, D.J., Strogatz, S.H.: Collective Dynamics of 'Small-World' Networks. Nature **393** (1998) 440–442
3. Flake, G.W., Pennock, D.M., Fain, D.C.: The Self-Organized Web: The Yin to the Semantic Web's Yang. IEEE Intelligent Systems **18**(4) (2003) 75–77
4. Hales, D., Arteconi, S.: SLACER: A Self-Organizing Protocol for Coordination in Peer-to-Peer Networks. IEE Intelligent Systems **21**(2) (2006) 29–35
5. Dixit, S., Yanmaz, E., Tonguz, O.K.: On the Design of Self-Organized Cellular Wireless Networks. IEEE Communications Magazine **43**(7) (2005) 86–93
6. Li, G., Yates, J., Kalmanek, C.R., Wang, D.: Control Plane Design for Reliable Optical Networks, IEEE Communications Magazine **40**(2) (2002) 90–96
7. Mas, C., Tomkos, I., Tonguz, O.: Failure Location Algorithm for Transparent Optical Networks, IEEE Journal on Selected Areas in Communications **23**(8) (2005) 1508–151
8. Kilper, D.C., *et al.*: Optical Perfprmance Monitoring. Journal of Lightwave Technology **22**(1) (2004) 294–304
9. Skorin-Kapov, N., Tonguz, O., Puech, N.: Self-Organization in Transparent Optical Networks: A New Approach to Security. In: The 9th International Conference on Telecommunications (Contel 2007), Zagreb, Croatia (2007), invited paper, 7–14
10. Médard, M., Marquis, D., Barry, R., Finn, S.: Security Issues in All-Optical Networks. IEEE Network **11**(3) (1997) 42–48
11. Bergman, R., Médard, M., Chan, S.: Distributed Algorithms for Attack Localization in All-Optical Networks. In: Network and Distributed System Security Symposium (NDSS'98), San Diego, Cal., USA (1998) session 3, paper 2
12. Li, C.-S., Ramaswami, R.: Automatic Fault detection, isolation, and Recovery in Transparent All-Optical Networks. Journal of Lightwave Technology **15**(10) (1997) 1784–1793
13. Wu, T., Somani, A.: Cross-talk Attack Monitoring and Localization in All- Optical Networks. IEEE/ACM Transactions on Networking **13**(6) (2005) 1390–1401
14. Collins, J.J., Chow, C.C.: It's a Small World. Nature **393** (1998) 409–410
15. Inkret, R., Kuchar, A., Mikac, B.: Advanced Infrastructure for Photonic Networks: Extended Final Report of COST Action 266, Faculty of Electrical Engineering and Computing, University of Zagreb, Zagreb (2003) 19–21
16. Skorin-Kapov, N., Tonguz, O., Puech, N.: A 'Small World' Hybrid Control Plane for Reliable Transparent Optical Networks, (submitted to IEEE Journal of Selected Areas in Communications).
17. Sivakumar,M., Shenai, R.K., Sivalingam, K.M.: A Survey of Survivabilty Techniques for Optical WDM Networks. In: Sivalingam, AK.M., Subramaniam, S. (eds.): Emerging Optical Network Technologies: Architectures, Protocols and Performance. Springer Science+Media, Inc., New York, USA, Chapter 3 (2005) 297–332
18. Buchanan, M.: Nexus: Small Worlds and the Groundbreaking Theory of Networks, W. W. Norton & Company, Inc., New York (2002) 199–204
19. Banerjee, D., Mukherjee, B.: Wavelength-Routed Optical Networks: Linear Formulation, Resource Budgeting Tradeoffs, and a Reconfiguration Study. IEEE/ACM Transactions on Networking **8**(5) (2000) 598–607