Exploiting Network Kriging for Fault Localization

K. Christodoulopoulos[1], N. Sambo[2], E. Varvarigos[1]

1: Computer Engineering and Informatics Department, University of Patras, and Computer Technology Institute and Press, Patras, Greece
2: Scuola Superiore Sant'Anna, Pisa, Italy

Abstract: We propose a correlation framework for fault localization, leveraging information from established lightpaths. This functionality is integrated in a reactive control plane employing a lightpath establishment algorithm to unambiguously localize failed link.
OCIS codes: 060.4257; 060.4256.

1. Introduction

Backbone and metro networks are evolving toward high flexibility. Transmission parameters, such as modulation format or forward error correction, are expected to be optimized for service setup and changed in case of re-routing (e.g., due to a fault). Moreover, operators are taking into consideration to install networks operating at their limits and reducing worst-case margins (e.g., for ageing or interference effects), so that optical reach can be increased and the number of regenerators – thus costs – reduced [1]. In a network working with low margins, soft failures (e.g., implying Quality of Transmission – QoT – degradations) may become more frequent [2]. So, in such a network, monitoring functionalities and the processing of monitoring information are key functions to re-act to both soft- and hard-failures. At the control plane layer, IETF Application-based Network Operations (ABNO) architecture [3] is emerging as an architecture providing control and management functionalities such as provisioning and Operation, Administration, and Maintenance (OAM) [4]. In particular, the ABNO OAM Handler is responsible to receive alarms of potential problems, to correlate the alarms, and to take actions to preserve services. In the literature, Network Kriging (NK) has been exploited as a mathematical framework to correlate linear parameters. In the field of optical networks, NK has been used to correlate physical-layer parameters with the aim of reducing the number of monitors [5] or estimating the Quality of Transmission (QoT) of lightpaths to be established [6].

Lightpaths established specifically for monitoring purposes have been used as a way to enable fast restoration performed purely at the optical layer [7]. So, taking a similar approach in this paper, we propose a correlation framework based on NK for fault localization, leveraging information from lightpaths established for data communication in the network. Since a fault can be localized with ambiguity with NK, the framework exploits *Monitoring as a service*. In particular, the control plane (i.e., the ABNO architecture) triggers the setup of new lightpaths with the scope of identifying unambiguously the failed elements. We also propose a heuristic Failure Localization-Aware Routing and Spectrum Allocation (FLA-RSA) algorithm that provisions lightpaths with the objective of reducing the failure localization ambiguity. A study on ambiguity, i.e. the capacity of the correlation algorithm and the proposed FLA-RSA algorithm to identify or not the failed links is also provided.

2. Control plane workflow

In this study we assume the control and management structure based on the ABNO architecture [3]. At the data plane, a wavelength routed optical network is considered to be equipped with monitors, e.g. installed in the digital signal processing (DSP) of the coherent receivers. In case a soft- or a hard-failure is detected by a monitor, an alarm is sent to the ABNO OAM Handler. In this study we focus on a single link hard- or soft-failure, which is the most common scenario. A centralized OAM Handler or a hierarchical OAM [4] as proposed within the European ORCHESTRA project is assumed. Fig. 1 shows the flow chart of the control plane actions upon a failure is detected by monitors. A single fault generates a certain number of alarms – received by the OAM Handler – which is proportional to number of affected lightpaths. By receiving the alarms, the OAM Handler triggers the re-routing of affected services on a link-disjoint path. The OAM Handler (or certain agents in case of a hierarchical plane) uses NK to localize the failed link. Based on the received alarms it may happen that the correlation algorithm identifies unambiguously the failed link or not. If more than one links are suspected to have failed, further actions are taken by the OAM Handler. In particular, we propose to use *monitoring as a service*, that consists in establishing probing lightpaths to provide more information for correlation with the purpose of solving unambiguously the fault localization problem. Moreover, the RSA algorithm can exploit NK to estimate the QoT of the lightpaths to be established for restoration, in a manner similar to [6]. This research direction is left outside of the current paper, being a future endeavor.

3. Network Kriging for fault localization, monitoring as a service, and failure localization aware RSA

3.1. Failure Localization

NK is a mathematical framework based on linear operations. Because of this intrinsically linearity we propose to associate to each link a linear parameter describing if it is active or failed. Let Active Parameter of link l, denoted by AP_l, be the parameter representing if the link l is active or failed. AP_l=1 if l is active, while AP_l=0 otherwise. Let AP_p be the parameter representing if the lightpath p is active or failed. AP_p is given by the sum of the APs of the links in lightpath p. Assuming that lightpath p traverses V links, if p is not involved in the failure (the path does include the failed link), AP_p=V. If p is involved in the failure, AP_p<V. Assuming a single link failure, if the OAM Handler receives an alarm of fault related to a lightpath p, the OAM Handler sets AP_p=V-1.

Consider a network with N nodes, L unidirectional links and P already established lightpaths in it. The routing matrix of established lightpaths is defined as $G \in \{0,1\}^{PxL}$ where $G_{p,l}$=1when a lightpath p contains link l. Consider the end-to-end parameters $\mathbf{y} \in R^P$, where y_p is a value for lightpath p. Vector $\mathbf{y}$ can be written as linear combination of link-level vector parameters $\mathbf{x} \in R^L$ so that $\mathbf{y}=G\mathbf{x}$. We denote by $\mathbf{y_m}$ (or $\mathbf{y_n}$) the parameters of the lightpaths for which monitoring data are available (or should be estimated), and set $\mathbf{y} = [\mathbf{y}_m^T, \mathbf{y}_n^T]^T$. Similarly, the routing matrix G is denoted as $G = [G_m^T, G_n^T]^T$ where G_m (or G_n) includes the rows that correspond to lightpaths for which monitoring information is available (or whose QoT parameters we want to estimate). Then, $[\mathbf{y}_m^T, \mathbf{y}_n^T]^T = [G_m^T, G_n^T]^T\mathbf{x}$. The objective is to determine the unknown parameters $\mathbf{y}_n$, where $\mathbf{y}_n = G_n\mathbf{x}$, which can be achieved using NK [6].

Based on the definition of the Active Parameter, presented above, we can formulate the failure localization problem as follows. We create the corresponding routing matrix G_m, with columns the links and rows the monitored established lightpaths. Then $\mathbf{y}_m$ includes the related Active Parameters of these lightpaths. The matrix G_n is the eye matrix, including all single links. By solving the problem with the NK method, we obtain the Active Parameter of the links and localize the failure, being the link with AP equal to zero. The localization of the failure might be unambiguous, meaning that we find a single link with AP=0, or might be ambiguous when more than one links have AP close to 0. The latter is the result of not having enough monitoring information, and can be caused by having few established lightpaths or the established ones have degenerate (similar and not diverse) routes. In this case monitoring as a service can be used to establish lightpaths crossing the ambiguous links and fill the G_m matrix so as to definitely solve the localization problem.


Figure 1 Flow chart of control plane actions upon failure.

3.2. Failure Localization Aware RSA

The proposed FLA RSA is a heuristic algorithm that establishes a single lightpath at a time. It is an extension of the heuristic proposed in [8], considering the current utilization state of the network to avoid spectrum overlapping. It considers a set of candidate lightpaths, formed by calculating the k-shortest paths between the given source and destination and combine that with the tuples that represent the transmission options of the available transponders. The novelty of the proposed algorithm is that instead of selecting the solution that optimizes the spectrum used, it examines whether adding the new path enriches the routing matrix with information that can improve failure localization. To be more specific, the lines of the routing matrix represent the links that the paths cross. A row that is linearly dependent on other rows gives no additional information for single link failure localization, that corresponds to a single column. Following this fact, before establishing a new lightpath we check the rank of the routing matrix including a candidate path, and preferably select a path that increases the rank. The proposed heuristic algorithm can be included in the PCE of the ABNO and used for serving new lightpath demands or iteratively in the network planning phase so as to serve all demands in the traffic matrix, in which case simulated annealing can be used to search among different orderings.

4. Performance results

To evaluate the performance of the proposed failure localization framework and the failure localization aware RSA algorithm we performed simulation experiments. In particular we compare the performance of a failure localization unaware RSA (referred to as RSA) to the propose FLA-RSA, assuming k=3, 6 and 10 alternative paths. We assumed

the DT network topology with 12 nodes and randomly created traffic matrices for specific traffic loads (50 matrices for each reported load). The matrix defines 100 Gbps PM-QPSK lightpaths required to be established, each assuming to consume 37.5 GHz and having 1500 km of maximum reach, which poses constraints on the length of the established lightpaths reducing the number k of paths considered for specific communicating pairs. The traffic matrices were created for specific loads defined as a percentage, with load equal to 1 being the all-ones traffic matrix, that is, corresponding to an all-to-all communication pattern. Note that at load equal to 1, the proposed correlation framework finds always unambiguously the failure, since all links are representing by a single lightpaths and the routing graph has full rank. The performance metrics we used in our study is (a) the number of monitors as service (probing lightpaths) that need to be established to solve the ambiguity, and this is for all single link failures in the network, and (b) the maximum number of used slots (i.e., a slot is a standard portion of 12.5 GHz) and the graphs report on the average values of these metrics over the examined traffic matrices.

Fig. 2a shows the average number of monitors as a service required to obtain unambiguous failure localization. As expected, as the traffic grows the number of monitors as a service required decreases (i.e. at low load, some links are unused by any lightpath and require extra monitors). For example, for load 0.2 we need about monitors in the order of ten, but typically, the network is not supposed to be used at that low load. The simple (failure localization unaware) RSA is shown to require much more extra monitors compared to the proposed FLA-RSA solution that chooses paths taking into account the failure localization process. The price paid by the FLA-RSA is the use of some longer paths (still calculated to be QoT feasible) that yields in higher spectrum utilization (Fig. 2b). However, from Fig. 2b it seems that the increase in spectrum utilization is quite low, that in the worst case was lower than 5 slots. Using a higher number of k paths in the FLA-RSA algorithm reduces the extra monitors required, with a relatively small increase in the spectrum used. From the results and for the network under study we found that for loads higher than 0.4, the FLA-RSA achieves unambiguous localization without any extra monitor and even with k=3 paths considered (so with limited number of extra spectrum required). This finding supports the key idea of this paper: leveraging information of data lightpaths can be used for failure localization. The need for monitors as a service is low and can be used as a safeguard, while the control workflow can be quite fast and responsive to failures.


Figure 2: (a) average number of monitors as a service required for achieving unambiguous failure localization, (b) average number of slots required for serving traffic (used for failure localization) as a function of the traffic load.

5. Conclusions

We proposed a correlation framework for fault localization, leveraging information from established lightpaths. Since a fault can be localized with ambiguity, the control plane triggers the setup of new lightpaths (monitors as a service) with the scope of identifying the failed element. The ambiguity can be reduced using the proposed Failure Localization-Aware Routing and Spectrum Allocation (FLA-RSA) algorithm. The study on ambiguity that we performed showed that at typical traffic loads, the framework can achieve unambiguous failure localization, the need for monitors as a service is low and the proposed control workflow is quite fast and responsive to failures.

6. References

[1] J.-L. Auge, "Can we use flexible transponders to reduce margins?" in Proc. of OFC/NFOEC, 2013, March 2013.
[2] K. Christodoulopoulos, et al. "ORCHESTRA - Optical performance monitoring enabling flexible networking," Proc. of ICTON, 2015.
[3] D. King and A. Farrel, "A PCE-based architecture for application-based network operations," IETF RFC 7491.
[4] N. Sambo, F. Cugini, A. Sgambelluri, P. Castoldi, "Hierarchical monitoring architecture and OAM Handler", Proc. of ECOC 2015.
[5] Y. Pointurier, M. Coates, and M. Rabbat, "Active monitoring of all-optical networks" in Proc. ICTON, 2008.
[6] N. Sambo, et al. "Lightpath Establishment Assisted by Offline QoT Estimation in Transparent Optical Networks," JOCN, vol.2, no.11, pp.928-937, November 2010.
[7] J. Tapolcai, P.-H. Ho, P. Babarczi, L. Ronyai, "Neighborhood Failure Localization in All-Optical Networks via Monitoring Trails," in Networking, IEEE/ACM Transactions on , vol.PP, no.99.
[8] K. Christodoulopoulos, P. Soumplis, E. Varvarigos, "Planning flexible optical networks under physical layer constraints," in Optical Communications and Networking, IEEE/OSA Journal of , vol.5, no.11, pp.1296-1312, Nov. 2013