A GRASP-Based Heuristic to Design the GMPLS Control Plane Network Topology with Resilience Guarantees

Marc Ruiz, Luis Velasco, Gabriel Junyent, and Jaume Comellas
Advanced Broadband Communications Center (CCABA),
Universitat Politècnica de Catalunya (UPC), Barcelona, Spain
e-mail: mruiz@ac.upc.edu

ABSTRACT
The control plane topology has traditionally followed that of the data plane and therefore the design of the former, as well as its resilience, has been addressed in the latter's one. Owing to the fact that a clear separation between control and data planes is a main GMPLS architectural requirement, the control plane in GMPLS networks may describe a different topology than the data plane, even realized over a separated IP network. As a consequence of this, data and control network design become no more linked in such scenarios. In this paper, we design a heuristic algorithm to obtain the optimal GMPLS control plane design, minimizing the network Capital Expenditures (CAPEX) while matching specific resilience requirements. The performance of the heuristic algorithm is compared with exact solutions obtained by solving an Integer Linear Programming (ILP) model on several reference network scenarios. Furthermore, its benefits in terms of total execution time are also highlighted.
Keywords: Optical Networks, GMPLS Control Plane, Network Topology Design.

1. INTRODUCTION

Current transport network operation relies on three clearly separated planes, providing each one different functionalities. From the bottom to the top of the network architecture, we find the data, control, and management planes. The data plane represents those physical network resources that support the exchange of information between end-users (e.g., end-to-end connections). The control plane is devoted to automate the routing and signaling of the paths whereby the end-user data will flow. The goal of the control plane is the provisioning of advanced network services such as Bandwidth on Demand (BoD) to efficiently support short-term and long-term traffic fluctuations. On top, the management plane supervises the whole network operation.

The Automatically Switched Optical Network (ASON, [1]) has been adopted as the leading architecture towards flexible and easy-to-maintain optical transport networks. Basically, ASON describes the reference architecture for the control plane and its interfaces. Although all specifications in [1] are technology-independent, the Generalized Multi-Protocol Label Switching (GMPLS, [2]) technology has arisen as the most accepted solution for implementing the control plane functionalities in ASON.

A separation between control and data planes was introduced in GMPLS, letting the control information to be transmitted on a different wavelength of the same fiber (*in-fiber out-of-band*) or on a separated IP network (*out-of-fiber*). Moreover, control and data information in GMPLS do not have to be congruently routed. This opens the possibility of deploying asymmetrical control plane topologies, instead of the usual symmetrical topologies in the *in-band* configurations.

Some works have addressed the resilience of the GMPLS-enabled control plane. In particular, [3] highlighted the reasons of a decoupled control plane in all-optical networks and addressed the resilience requirements that this would impose. In our previous work [4] we presented an analytical expression to estimate the resilience of an asymmetrical *out-of-fiber* meshed control plane. This parameter (called P_d) was subsequently used in [5] to include resilience constraints to the first control plane design problem defined independently from the data plane. More specifically, in [5] we defined and formulated the problem of minimizing the control plane CAPEX subject to some resilience requirements (referred as ARCO problem). An iterative method based on an Integer Linear Programming (ILP) model was also presented. After solving several instances over real topologies, we concluded that the reduction of control plane links provides significant CAPEX reduction without affecting negatively to the data plane performance (i.e., blocking probability, failure recovery time).

Since the complexity of the ARCO problem is NP-hard, the computational effort increases exponentially with the size of the network, becoming thus impractical for large topologies. In light of that, in this paper we propose a meta-heuristic method for solving the ARCO problem. More specifically, we propose to apply the *Greedy Randomized Adaptative Search Procedure* (GRASP) meta-heuristic, whose effectiveness for solving hard combinatorial optimization problems has been abundantly proven [6].

This work has been partially founded by Spanish Science Ministry through TEC2008-02634 ENGINE project.

2. PROBLEM STATEMENT

Let G_{DP} (N_{DP}, E_{DP}) and G_{CP} (N_{CP}, E_{CP}) be the data and the control plane graphs of a GMPLS-enabled transport network, respectively, where N_{DP} and N_{CP} represent the sets of nodes and E_{DP} and E_{CP} the sets of arcs. For the ongoing definition, we assume that G_{DP} is two-connected and planar. Moreover, we also assume that G_{CP} is two-connected providing survivability to the control plane. Particularly, we restrict the control plane topology to be a subset (or the complete set) of the data plane one. More formally, we restrict G_{CP} to be a spanning subgraph of G_{DP}, i.e., $N_{DP} = N_{CP}$ and E_{CP} contained into E_{DP}. Thus, G_{CP} ranges from the *symmetrical* topology (equal to G_{DP}) to the *minimal* one (the minimum-size two-connected topology).

Aiming at measuring the resilience of the control plane, P_d is defined as the probability that any connection request or teardown message becomes affected by a control plane link failure during the failure recovery time Δt (i.e., it is forwarded onto the failed link). For the sake of easiness, equation (1) reproduces from [4] the analytical expression of P_d, where λ and μ represent the inter-arrival and service rate for the traffic distribution, respectively. The parameter C identifies the number of active connections in the network and can be accurately computed as ceil(λ/μ).

$$P_d = 1 - e^{-\lambda\Delta t(1+P_L)} \cdot \sum_{k=0}^{C} \binom{C}{k} \left[\left(e^{\mu\Delta t} - 1\right)\cdot\left(1 - P_L\right)\right]^k = 1 - e^{-\lambda\Delta t(1+P_L)} \cdot \left[1 + \left(e^{\mu\Delta t} - 1\right)\cdot\left(1 - P_L\right)\right]^C . \quad (1)$$

In the equation, P_L provides information regarding the control plane topology. More specifically, P_L is defined as the probability that the signaling of a setup/teardown connection is supported on a certain control plane link and can be computed as follows:

$$P_L = \frac{\langle H_{CP} \rangle}{|E_{CP}|} \quad (2)$$

being $\langle H_{CP} \rangle$ the average hop length of control plane paths.

Then, the ARCO problem can be defined as follows:

- *Given*: a) The data plane topology represented by the graph G_{DP} (N_{DP}, E_{DP}); b) the set of traffic intensities represented by a set Λ of inter-arrival rates and a service rate μ; c) a P_d threshold value (P_d^{max}); and d) a Δt control plane failure recovery time.
- *Output*: A control plane topology G_{CP} (N_{CP}, E_{CP}), and its corresponding P_d.
- *Objective*: Minimize the number of control plane links (i.e: $|E_{CP}|$)
- *Subject to*: a) G_{CP} is two-connected and b) $P_d \leq P_d^{max}$ for every inter-arrival rate in Λ.

The P_d constraint formulated using equation (1) converts the problem into non-linear. To deal with this non-linearity, in our previous work [5] we proposed an iterative method where an ILP was solved at each iteration. Fixing the number of control plane links as parameter, the ILP minimized $\langle H_{CP} \rangle$ in order to minimize P_L. P_d was computed outside of the ILP, as soon as the solution was found and the decision of adding or removing one link was taken in base of the obtained P_d value. We proved the validity of this iterative procedure to find the global optimal solution defining some mathematical propositions. Besides ensuring the minimization of $|E_{CP}|$, that method also selects, among the alternative optimal topologies, the one with the lowest P_d. The importance of this additional minimization underlies in the possibility of increasing Δt without violating P_d^{max}, thus leading to savings as a consequence of the deployment of less stringent recovery mechanisms.

In spite of the quality of the solutions obtained with this method, execution times became impractical for real-sized networks as a consequence of the size of the problem. In view of this, in the next section we present a GRASP-based meta-heuristic to solve ARCO with the aim of finding near optimal solutions with significant lower computational effort.

3. GRASP-BASED HEURISTIC

The GRASP meta-heuristic is an iterative procedure consisting of a two-phase main algorithm which finds a good-quality solution at each iteration. Within the first phase of the algorithm (*constructive phase*) one feasible solution is built by means of an *ad-hoc* randomized greedy algorithm. The degree of randomness is determined by the parameter α. Next, the *local search phase,* designed to explore the neighborhood of the solution, is applied aiming at improving the current solution. This iterative method is executed until a stop criterion is met. In this work, we consider as stop criterion a given maximum number of iterations without improving the best solution (hereafter *itemax*).

Since the *symmetrical* topology is the one with lowest P_d, the *constructive phase* starts from this topology and iteratively removes one link at each step. The reduction concludes when the *minimal* topology is reached or when the P_d^{max} threshold cannot be ensured. The selection of the link is based on its associated greedy cost (*GrC*). Being l a link of the control plane topology G_{CP}, $GrC(l)$ is defined as the resultant $\langle H_{CP} \rangle$ associated to the topology where the set of links is E_{CP}-$\{l\}$. Thus, the lower $GrC(l)$ the lower increase of P_L causes the candidate link l and, consequently, of P_d. At each *constructive phase* step, a candidate list (*CL*) is built adding every link which elimination results into a two-connected topology. Then, the restricted candidate list (*RCL*) is defined as follows:

$$RCL = \left\{ l \in CL : GrC(l) \leq GrC^{\min} + \alpha \cdot \left(GrC^{\max} - GrC^{\min} \right) \right\} \tag{3}$$

where GrC^{min} and GrC^{max} are the minimum and maximum greedy costs in *CL* respectively. The link to be removed is selected at random from the *RCL*.

Aiming at reducing the P_d of a feasible solution, the *local search* algorithm explores the neighborhood of the current solution looking for feasible interchanges between one link in the solution (in E_{CP}) and one link not in the solution (in E_{DP}-E_{CP}). When an interchange reduces the current P_d, the solution is modified. In that case, the *local search* algorithm continues exploring the new neighborhood. On the contrary, when the neighborhood is completely explored without any improvement, it finishes.

4. ILLUSTRATIVE NUMERICAL RESULTS

The performance of the proposed GRASP-based heuristic has been evaluated over the 14-node DT and the 37-node EON network topologies (Fig. 1). For each topology, five different traffic intensities were tested. The *inter-arrival* time ($iat = 1/\lambda$) ranges from the lowest value (normalized to 1) to the highest one (4 times the lowest). The *holding* time ($ht = 1/\mu$) is proportional to the *iat*, being the intensity $I = ht/iat$ equal to the one that provides a blocking probability near to 1%. Moreover, two different resilience requirements configurations are compared (P_d= 5% / Δt = 1.5 s and P_d= 10% / Δt = 3.5 s), thus considering 10 different instances per topology. The heuristic has been implemented in Matlab and the optimal solutions were obtained using CPLEX v.11.0 optimizer. Both implementations were executed on 2.4 GHz Quad-Core machines with 8 Gb RAM.

Aiming at determining the best value of α, we solved each instance 100 times with a low value of *itemax* given by the number of links of the data plane. Figure 2 shows the average gap obtained for each α. As illustrated, the best performance is obtained with $\alpha = 0.5$ for the DT and $\alpha = 0.1$ for the EON topologies respectively. In general, a low degree of randomness (i.e.: $\alpha \leq 0.3$) provides the best performance for large networks. However, it is necessary to increase α when the size of the network is reduced (e.g.: DT topology), in order to allow enough amount of links in the *RCL*. The previous experiment has been repeated for different *itemax* values fixing α to the best configuration. Figure 3 shows the average gap of the best 95% of repetitions for each tested *itemax*. Note that for *itemax* equal to 100 for the DT and equal to 520 for the EON topologies respectively, the gap is equal to 0. Then, we conclude that our GRASP heuristic provides the optimal solution in every case.

Figure 4 shows the running times needed to reach the optimal solution and compares them with the obtained using CPLEX. As illustrated, the heuristic provides affordable and scalable computational times in contrast with the obtained solving the ILP model. In fact, the computational time of the heuristic presents a polynomial increase with respect to the network size, in contrast with the exponential increase of the ILP model.

Finally, Figure 5 shows the P_d values of each optimal topology, obtained solving the ILP model and with the heuristic. In just few cases the P_d value obtained with the heuristic is higher than the optimal value. Although the topologies contain the same number of links, this P_d difference entails that the topologies are slightly different. Nevertheless, the post-optimization process proposed by [5] to increase Δt can be also successfully applied as extension of the heuristic optimization.

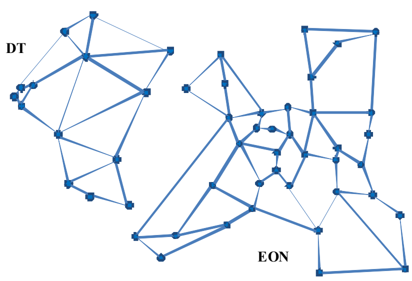

	$\lvert N\rvert$	$\lvert E_{DP}\rvert$	d_{DP}
DT	14	23	3.29
EON	37	57	3.08

Figure 1. Network topologies used for evaluation with their most relevant characteristics

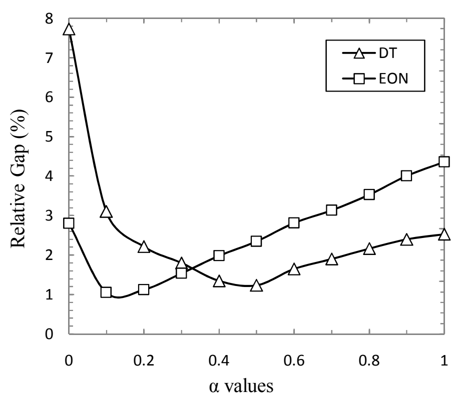

Figure 2. Relative gap against α values, with itemax= $\lvert E_{DP}\rvert$.

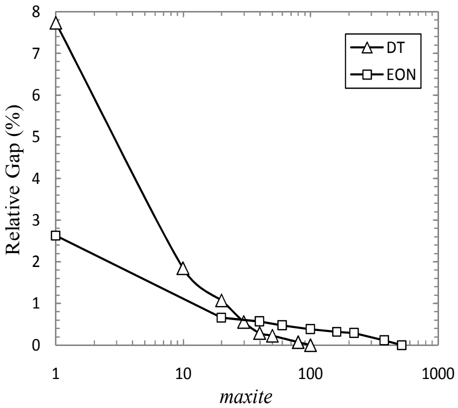

Figure 3. Relative gap against itemax values, with α=0.5 for DT and α=0.1 for EON.

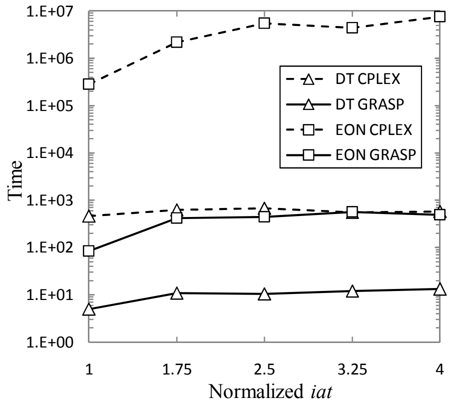

Figure 4. Comparison between CPLEX and the proposed heuristic in terms of running time.

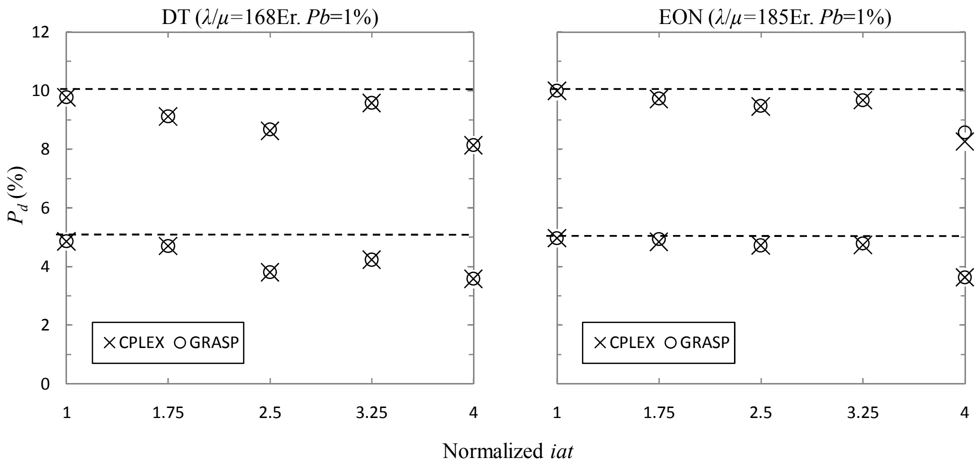

Figure 5. Comparison between CPLEX and the proposed heuristic in terms of Pd minimization.

5. CONCLUSIONS

We have presented a GRASP-based meta-heuristic to solve the control plane topology design problem ensuring some resilience requirements (ARCO). After tuning some heuristic parameters, the best configuration has been used to solve a set of instances using two well-known backbone optical networks. The performance of the proposed heuristic has been compared with that of an exact method solved using CPLEX. Our heuristic algorithm provided the optimal solution for every instance and non-significant differences in terms of P_d were detected. Moreover, regarding computational time, our heuristic needed just some seconds to few minutes to provide the optimal solution, in contrast with the time, several orders of magnitude higher, needed by CPLEX.

In light of these results, we conclude that the proposed heuristic is a highly efficient method to solve ARCO providing an excellent trade-off between solution quality and computational effort.

REFERENCES

[1] ITU. Rec. G.8080/Y.1304: "Architecture for the automatically switched optical network (ASON)," ITU-T, Nov. 2001.
[2] E. Mannie, "Generalized multi-protocol label switching architecture," IETF RFC 3945, Oct. 2004.
[3] A. Jajszczyk, P. Różycki, "Recovery of the control plane after failures in ASON/GMPLS networks", *IEEE Network*, vol. 20, pp. 4-10, 2006.
[4] M. Ruiz, J. Perelló, L. Velasco, S. Spadaro, and J. Comellas, "An analytical model for GMPLS control plane resilience quantification," *IEEE Commun. Letters*, vol. 13, pp. 977-979, 2009.
[5] M. Ruiz, J.Perelló, L. Velasco, S. Spadaro, J. Comellas and G. Junyent, "GMPLS control plane network design with resilience guarantees", *IEEE/OSA Journal of Lightwave Technology*, vol. 29, pp. 37-47, 2011.
[6] M. G. C. Resende, C. Ribeiro, "*Greedy Randomized Adaptive Search Procedures: Advances and Applications*," Handbook of Metaheuristics. 2nd ed., Springer Science+Business Media, 2010.