Graph Partitioning for Survivability in Multi-Domain Optical Networks

Hamza Drid, Bernard Cousin, Miklós Molnár, and Nasir Ghani

***Abstract*—As optical networking deployments increase, multi-domain survivability is becoming major concern. To address this issue, in this paper we propose a novel graph partitioning technique for improved protection scalability in multi-domain optical networks. To demonstrate the efficiency of our method, we also extend the p-cycle concept to expanded multi-domain settings. Overall simulation results show the efficiency of our proposed solution in terms of resource utilization and running time.**

***Index Terms*—Survivability, multi-domain, optical networks.**

I. Introduction

SURVIVABILITY in optical networks is a key focus area and has been well-studied over the years. The term refers to the ability of a network to continue to provide services even in the presence of a failure. Now survivability in multi-domain optical networks entails the ability to recover end-to-end lightpaths crossing multiple domains. In general, this consists of two main tasks. The first task is to collect and maintain up-to-date network state (e.g., link resources, link diversity, etc). Meanwhile, the second task is to find and reserve working and backup resources for the connections based upon the above-collected information. Now the first task is generally implemented by running routing protocols to broadcast related link-state information between network nodes, thereby allowing each to have a complete view of the network (i.e., distributed settings). However, as network size grows, it is very unrealistic to broadcast complete topology information to every network node. This scalability challenge is further complicated when networks are delineated into multiple domains, each with its own administrative privacy constraints.

In order to address these scalability concerns, several multi-domain protection schemes have been studied and these can be classified into two types. Namely, the first type computes two ”end-to-end” link-disjoint paths for a source-destination request, i.e., working/backup paths can traverse different domains [1]. Meanwhile, the second type protects each domain-level segment of the working lightpath [2-5]. Hence this is more of a “per-domain” protection approach in which working/backup paths must traverse the same domain sequence. However, in order to compute inter-domain working/backup routes, both of the above protection schemes require some form of network aggregation to help reduce routing information exchange between domains. For example, a key strategy here is to use topology abstraction to transform domains into smaller “representative” graph models, e.g., mesh, star, tree, see [1],[6].

Manuscript received September 26, 2009. The associate editor coordinating the review of this letter and approving it for publication was X. Cao.

H. Drid, B. Cousin, and M. Molnár are with the University of Rennes I, IRISA, France (e-mail: hdrid@irisa.fr).

N. Ghani is woith the University of New Mexico, USA.

Digital Object Identifier 10.1109/LCOMM.2010.08.091923

Although topology abstraction schemes can hide the internal network topology of each domain, they still cannot scale for very large network sizes with sizeable domain counts. Take for example a multi-domain network with hundreds of domains. Here, even if each domain is aggregated into a single virtual node, i.e., highest level of state reduction, state dissemination overheads can still be very high. As a result, aggregation is not deemed as a sufficient means for achieving scalable multi-domain survivability provisioning, i.e., especially when the number of domains is large. In light of the above, in this paper we propose to use graph partitioning techniques to improve multi-domain optical network survivability. To understand how our method works and overcomes the problems described above, we will apply it towards the protection of multi-domain optical networks using p-cycles [7]. Namely, p-cycles are chosen as they are an effective means of deploying high-grade protection over flexible network routes. Nevertheless, our solution is generic and based upon a network partitioning strategy using spectral clustering [9]. As such, it can readily be combined with other multi-domain protection strategies, e.g., schemes in [1-5] can be applied within a given partition.

The overall paper is organized as follows. Section 2 presents a survey of the existing work in multi-domain optical network survivability. Also, Section 2 describes the challenges associated with graph partitioning and develops a novel solution to apply this technique in multi-domain networks. The objective here is to reduce the complexity of network state exchange. Finally, the performance of the proposed scheme is studied in Section 3 for the case of p-cycle protection and conclusions and directions for future work are presented in Section 4.

II. Proposed Solution

As mentioned earlier, various strategies have been proposed to improve the multi-domain protection scalability [1-5]. For example [5] proposes a solution for protecting inter-domain links using p-cycles. Namely, this solution computes the set of p-cycles over a single-node virtual topology in which each domain is represented by a single virtual node (see Fig. 1b). However, this method suffers from some serious drawbacks. Foremost, the scheme uses the entire multi-domain topology to compute a set of p-cycles. As a result, calculated p-cycles can traverse all domains, and this will entail excessive resource management and lengthy recovery times. Next, optical impairment concerns will also rise for longer p-cycle routes. Additionally, p-cycle computation times will also be large, as they tend to increase exponentially with node counts. Finally, [5] uses a single virtual-node topology to compute inter-domain p-cycles, i.e., reduce a domain with multiple nodes/links into a single virtual node. As such, this high level of aggregation will remove critical domain-internal topological information, further impacting the overall quality (e.g., redundancy) of computed p-cycles. A novel partitioning

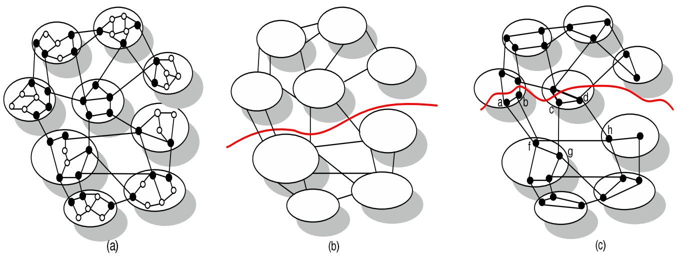

Fig. 1. Multi-domain networks.

strategy for multi-domain p-cycle protection is now presented to address the above concerns.

A. Detailed solution

In order to overcome the problems relating to protection scalability, we propose to divide a multi-domain topology into k independent, approximately equal "sub-multi-domain" networks. Next, we propose to separately solve the problem of survivability on each of these "sub-multi-domain" networks and thereby protect end-to-end primary lightpaths. Furthermore, since the physical topology of multi-domain networks is very important and graph partitioning is an NP-complete problem [9], the multi-domain topology that we propose to use in the partitioning process, again, is one that represents a domain as a single-virtual-node (see Fig. 1b).

Now after having built the k partitions, a set of p-cycles protecting each of the "sub-multi-domain" networks are computed in an independent manner. Here, clearly the size of the calculated p-cycles will be limited to the number of domains in each partition. Hence the amount of inter-domain routing information exchanged will now be notably reduced, i.e., between nodes in same partition. Moreover, associated computation times will also be lowered as the independent domain-level p-cycles can be computed in a parallel manner.

Note that when computing the partitions, inter-domain links connecting two separate partitions (i.e., links of the cut) will not be protected. To address this concern, we propose to further protect these links with the links of the partition that give the best redundancy after computing the set of p-cycles. Namely, the links of the cut to the first partition are joined and the redundancy of p-cycles which protect this new set of links is calculated. Specifically, these include the links of the partition, the links of the cut, and also the links connecting the nodes of the cut links in the other partition. This process is then repeated for the second partition. To illustrate this more clearly, consider the upper partition in Fig. 1c. Here, the links *a-f*, *b-f*, *c-g* and *d-h* are the links of the cut, and the links *a-b* and *c-d* connect the nodes of the cut in the other partition, i.e., nodes *a*, *b*, *c*, *d*. Overall, a lower redundancy value will allow us to decide with which partition the links of the cut are associated. Carefully note that this approach essentially implements (more scalable) link-level protection at the inter-partition level and not lightpath-level protection.

Finally, our proposed scheme also takes into account some internal topology information of each domain during p-cycles computation. This essentially addresses the earlier-cited problems related to imprecise resource state in p-cycle protection for multiple domains, i.e., in [5]. Now clearly, it is not feasible or practical to take into account the exact internal topology details of each domain here, i.e., owing to privacy concerns, computational overheads, etc. Hence we propose an intermediate solution that leverages existing topology aggregation models, e.g., mesh, to condense associated domain-level state, as shown in Fig. 1c. As a result, this yields a reduced "sub-multi-domain" topology on which the protection cycles can be computed [8]. Further details of the proposed graph partitioning process are now presented.

B. Graph partitioning

Graph partitioning is a fundamental optimization problem and been widely studied both in theory and in practice. Overall, this technique is very germane and powerful means of dissecting a larger problem into a set of semi-independent sub-problems that are easier to solve. Now the main objective of graph partitioning is to divide a given graph, $G = (\mathbf{V}, \mathbf{E})$, into k separate parts $(P_1, P_2, ..., P_k)$, i.e., where $\mathbf{V}$ is the set of vertices and $\mathbf{E}$ is the set of links.Generally we try to select a partition where all parts have approximately the same number of nodes (i.e., $|V|/k$, where $|\mathbf{V}|$ is the number of vertices) and where the number of edges connecting different parts is minimized. Hence in our case, we divide the topology of the multi-domain network into k parts (to be detailed shortly using spectral clustering techniques [9]),with each part contains approximately n_d/k domains, where n_d is the number of domains in the multi-domain network. For example, Fig. 1b shows a partition for $k = 2$. Overall, the k graph partitioning problem can be described as follows:

$$V = P_1 \cup P_1 \cup ... \cup P_k \tag{1}$$

where,

$$P_1 \cap P_1 \cap ... \cap P_k = \emptyset \tag{2}$$

Here, the subsets are assumed to be collectively exhaustive and also mutually exclusive. Hence the total cost of the partition is defined as:

$$C = \sum_{i,j \in V, P(i) \neq P(j)} c(i, j) \tag{3}$$

where $c(i, j)$ is the cost of the link between node i and node j in the graph G. The objective of the above k graph partitioning problem is to minimize the cost C of the edges connecting the different parts

Overall, graph partitioning is an NP-complete problem [9], and over the years a large number of heuristic solutions have been proposed to solve this problem efficiently.Now for the purposes of multi-domain network partitioning herein, any of these heuristics can be applied. However, without lack of generality, we choose to apply the spectral clustering method to divide a graph into disjoint sub-graphs, as this method has many advantages. Foremost, this algorithm has been shown to yield very satisfactory results and simple to implement. Specifically, the spectral clustering method computes an eigen-vector corresponding to the smallest nonzero eigen-value of the Laplacian matrix **L**, where **L=D-A**, and **A** is the adjacency matrix and **D** is the degree matrix [1].The graph is then divided

[1]Degree matrix is a diagonal matrix which contains the degree of each vertex.

TABLE I
REDUNDANCY

# of partitions	No partitioning	2 partitions	4 partitions
Partition Redundancy	No ILP solution	89.17%	93.84%
			91.24%
		88.48%	91.67%
			91.45%
Average Redundancy	No ILP solution	88.82%	92.05%

into two parts according to the sign of the eigen-vector components. As a result, the graph is divided into two equal sized sub-graphs. In turn, these sub-graphs can be recursively divided to achieve higher partitioning.

III. SIMULATION RESULTS

The performance of the proposed multi-domain survivability solution is evaluated using simulation analysis. Meanwhile, comparative *integer linear programming* (ILP) analysis from [7] is done for some cases using the MATLAB toolkit. Furthermore, all tests are carried out using a commonly used dense national topology as used in [10]. Namely, this topology contains 13 different domains (where each region is considered as a domain) and has a total of 50 nodes and 88 links. Traffic demands are uniformly distributed among all source-destination pairs. For each source-destination pair, an integer number between zero and the maximum allowable demand (which is 5 in our simulation) is randomly generated, and it is assumed that each request is for one wavelength capacity. The working capacities on the network links are also obtained by routing each demand over the shortest path. Now in order to properly evaluate the proposed solution, several criteria are defined:

Redundancy: This is defined as the ratio of protection capacity to working capacity. Here, low redundancy is deemed more efficient than high redundancy [7].

Computational complexity: This is also an important factor when evaluating any algorithm. Here it is simply defined as the time it takes for an algorithm to compute a solution.

Number of p-cycle structures: This is an important factor for evaluating any algorithm and is simply defined as the time it takes for an algorithm to compute a solution.

Performance analysis results are now presented for various scenarios. In the first scenario, the complete multi-domain network topology is considered when calculating the set of p-cycles. Namely, protection cycles are computed using the approach in [7], which uses an ILP model to find the optimal set of p-cycles in terms of redundancy. This scenario is used to provide a baseline case for comparing the performance of the proposed graph partitioning strategies. Meanwhile, in the second scenario the partitioning scheme of Section 3 is used to divide the multi-domain network topology into two parts. The optimal set of p-cycles are then computed for each of these sub-topologies using the same algorithm used in the previous scenario, i.e., from [7]. Finally, for the third scenario, the above process is repeated by further dividing the multi-domain topology into four parts.

The redundancy results for the above four scenarios are first presented in Table 1. Meanwhile, Table 2 also summarizes the respective running times for these different solutions.

TABLE II
RUNNING TIME

# of partitions	No partitioning	2 partitions	4 partitions
Max running time	864000 (s)	96703 (s)	20.9 (s)
Average running time	864000 (s)	48401 (s)	8.4 (s)

TABLE III
NUMBER OF P-CYCLES

# of partitions	No partitioning	2 partitions	4 partitions
# max of p-cycles	-	10	6
Average # of p-cycles	-	8,5	4

Foremost, the results clearly indicate that network redundancy increases with the number of partitions chosen. Although the relative percentage is about 4% (between 2 and 4 partitions) this value effectively represents a differential in network overbuild costs and hence can translates into a large monetary ”cap-ex” savings. These gains can be explained by the fact that larger partition sizes result in longer, augmented p-cycles. In turn, this lowers the spare capacity required for the set of p-cycles protecting the network and hence boosts overall redundancy. In other words, longer p-cycles tend to include more straddling links and therefore result in higher levels of protection resource efficiency.

Meanwhile, run time results in Table 2 clearly indicate that larger partition sizes are very problematic. For example, in the first scenario, i.e., non-partitioned, the ILP does not return a result due to the large network size. Conversely, the graph partitioning strategies quickly curtail run times by distributing the p-cycle computation overheads over semi-independent sub-problems that are easier to solve. For example, results show that increasing the number of partitions from 2 to 4 results in almost 2 orders magnitude lower runtimes, a very notable finding.

Finally, the number of distinct p-cycles generated by a design is also evaluated as has a direct impact upon network management costs/complexities. These findings are shown in Table 3 and indicate that the number of distinct structures actually decrease with the number of partitions chosen in the design, e.g., by almost 50% from 2 to 4 partitions. Consequently, network partitioning can be a very attractive option for network operators, allowing them to streamline and facilitate overall network survivability provisioning.

IV. CONCLUSION

This paper addresses survivability in multi-domain optical networks. Owing to the scale of these infrastructures, protection scalability is indeed a key challenge here. In order to address this concern, a novel graph partitioning approach is proposed based upon spectral clustering techniques to help reduce the overall multi-domain protection problem in to series of ”sub-multi-domain” problems. P-cycle protection strategies are then applied to these individual entities. Detailed simulation results show that the proposed partitioning approach can effectively resolve multi-domain protection scalability concerns. These results are very promising and future efforts will look at extending the work to multi-layer/multi-domain scenarios, e.g., IP-DWDM networks

REFERENCES

[1] D. Truon and B. Thiongane, "Dynamic routing for shared path protection in multidomain optical mesh networks," *J. Optical Networking*, 2006.

[2] L. Guo, "LSSP: a novel local segment-shared protection for multi-domain optical mesh networks," *Computer Commun.*, pp. 1794–1801, 2007.

[3] X. Zhang et al., "On segment shared protection for dynamic connections in multi-domain optical mesh networks," *Int. J. Elect. Commun.*, pp. 1–6, 2009.

[4] L. Guo et al., "A novel domain-by-domain survivable mechanism in multi-domain wavelength division-multiplexing optical networks," *Int. J. Optical Fiber Technol.*, pp. 192–196, 2009.

[5] J. Szigeti et al., "p-cycle protection in multi-domain optical networks," *Springer J. Photon. Network Commun.*, pp. 35–47, 2009.

[6] K. Liu et al., "Routing with topology abstraction in delay bandwidth sensitive networks," *IEEE/ACM Trans. Networking*, 2004.

[7] W. D. Grover and D. Stamatelakis, "Cycle oriented distributed preconfiguration: ring-like speed with mesh-like capacity for self-planning network restoration," in *Proc. IEEE Int. Conf. Commun.*, 1998.

[8] H. Drid et al., "A topology aggregation model for survivability in multi-domain optical networks using p-cycles," *6th IFIP Inter. Conf. Network and Parallel Computing*, 2009.

[9] U. Luxburg, "A tutorial on spectral clustering," *Statistics and Computing*, vol. 17, no. 4, pp. 395-416, Dec. 2007.

[10] A. Zymolka, "Design of survivable optical networks by mathematical optimization," Ph.D. thesis, Technische University at Berlin, 2006.