Analysis of Traffic Engineering Information Dissemination Strategies in PCE-Based Multi-Domain Optical Networks

S. Spadaro, J. Perelló, G. Hernández-Sola, A. Moreno, F. Agraz, J. Comellas, G. Junyent
Optical Communications Group (GCO), Universitat Politècnica de Catalunya (UPC)
Jordi Girona, 1-3, 08034, Barcelona (Spain)
Tel: (+34) 93 4010966, Fax: (+34) 93 4017200, e-mail: spadaro@tsc.upc.edu

ABSTRACT

This paper presents and compares different TE information dissemination strategies between Path Computation Elements (PCEs) in multi-domain optical networks. In such network context, recent studies have found that path computation only with local domain visibility yields poor network performance. Accordingly, certain visibility between domains seems necessary. Aiming to fit the confidentiality requirements of the composing domains while improving the final network blocking probability, novel link aggregation techniques have been proposed, which summarize the state of network domains resources efficiently. This aggregated link information is afterwards disseminated to all the remainder domains in the network. To this end, we introduce different update triggering policies to make a good trade-off between routing information scalability and inaccuracy. The performance of all contributions has been supported by illustrative simulation results.

Keywords: Multi-domain, path computation element, topology abstraction.

1. INTRODUCTION

In the last decade, research efforts related to optical transport network infrastructures, have been mainly focused on single-domain scenarios, where scalability and confidentiality do not represent an issue. However, the future optical networks will include several domains, each controlled by a different service provider/network operator. From the routing perspective, a domain is a collection of network elements within a common address management or path computational responsibility, namely, an Interior Gateway Protocol (IGP) area or an Autonomous System (AS) [1]. In such scenario, the computation of end-to-end paths poses new challenges and traffic engineering solutions addressing the constraints imposed by scalability, domain information confidentiality, heterogeneous transmission technologies and physical layer impairments, must be provided. Due to the limited availability of effective solutions, TE is practically unavailable in multi-domain optical networks. The topological information exchange between domains is normally reduced to the minimum, only spanning the shared links and border nodes information. The lack of topological information related to neighbouring domains hinders the routing entities' capacity to compute inter-domain end-to-end paths efficiently. To address this problem, the Internet Engineering Task Force (IETF) and the Optical Internetworking Forum (OIF) standardization bodies have proposed both architectural and protocol specifications to facilitate the interoperability amid different network domains, technologies and vendor equipments.

Specifically, referring to IETF and Path Computation Element (PCE)-based solutions, path computation strategies such as per-domain path computation [3] and Backward Recursive PCE-based Computation (BRPC) [4] have been introduced. In the former, for end-to-end path computation, only border nodes information is used and the PCE of each traversed domain computes the local section of the path (i.e., intra-domain connection). In the latter, a domain sequence is previously provided to assure that the lower number of domains is traversed. However, it has to be noted that these two strategies do not rely on any kind of shared information among domains. Notwithstanding, to compute optimal or sub-optimal routes, a certain degree of information to be shared among the domains is required. In this context, topology abstraction schemes, which provide a simplified aggregated domain representation, have been presented as a viable solution aiming to attain both routing scalability requirements and confidentiality concerns among domains. Specifically, OIF suggests the application of topology abstraction schemes to disseminate routing information between different domains using the standard E-NNI interface [5]. Recent studies have analyzed three topology abstraction schemes for multi-domain optical networks, that is, *simple node*, *symmetric-star* and *full-mesh* [6], [7]. In this paper, we focus on the full-mesh abstraction scheme, since it provides better performance than the other abstraction schemes.

However, as the intra and inter-domain connections are established/released during network operation, the virtual topology information of each domain has to be periodically updated; then, the updated virtual topology must be disseminated, in order to assure that the responsible routing entity of each domain (i.e., the PCE in our network scenario) is able to compute inter-domain end-to-end paths efficiently.

In this paper, we propose two novel virtual topology update triggering policies aiming to keep the multi-domain topology abstraction synchronized. A trade-off between distributed routing information, scalability and security concerns across domains is pursued.

The rest of this paper is organized as follows. Section 2 describes our proposed mechanisms; the dissemination of the updated virtual topology information of each domain by using the GMPLS OSPF-TE

protocol is also discussed. The performance studies are highlighted in Section 3. Finally, some conclusions are discussed in Section 4.

2. PROPOSED MECHANISMS

In this section, we propose two different mechanisms to enhance the domain virtual topology update and its dissemination across different domains. Intra-domain nodes are considered to be all-optical whereas border nodes perform full opto-electronic conversion. Each domain maintains at least one dedicated PCE to build a complete graph of virtual links between all its border nodes (i.e., a full-mesh abstraction). It is assumed that the PCEs have full resource visibility of their domains as provided by OSPF-TE link-state routing protocol.

Then, each PCE disseminates the abstracted topology of the domain to the remainder PCEs. This way, every PCE has a global graph of the multi-domain network, which will be used for the computation of inter-domain path upon an end-to-end connection provisioning request. Specifically, every connection request between two nodes is routed according to source-routing loose path selection based on the minimum cost path. The source PCE (which has a full view of its own domain and the aggregated graph from the rest of the network) calculates the complete path inside its domain and the domain sequence to reach the destination node. After that, PCEs from all the transit domains in the sequence select the proper shortest intra-domain route between their border nodes. At this point, the inter-domain end-to-end source-destination path is determined.

Nevertheless, as the number of allocated connections (both intra and inter-domain) increases, the cost of the virtual link of the corresponding full-mesh abstractions must be updated. Then, the updated topology abstraction must be disseminated. However, such dissemination must be performed in a controlled manner, in order to keep limited the network overhead and increase multi-domain routing scalability.

To address these problems, we introduce two different policies, each one based on different update criteria and different aggregated information to be shared among domains.

2.1 Inter-Domain Transit Path

Inter-Domain Transit Path (IDTP) is a triggering policy designed to maximize the efficiency of the distributed updates among PCEs. When a connection is allocated in a domain, the domain network resources are occupied and its full-mesh abstraction is modified. In the transit domains, if a huge number of inter-domain connections is allocated, the virtual link costs drastically change, driving to an outdated multi-domain network graph. Whether the number of established connections in a transit domain (M) exceeds a pre-defined threshold (N), the PCE of the corresponding domain disseminates the updated topology abstraction to all neighbouring PCEs. The updated topology abstraction consists of those virtual links whose costs have been changed with respect to the previous triggered update. The shortest path cost between border nodes is assigned according to the link administrative weights.

2.2 Inter-Domain Transit Path - Common Cost Balancing (IDTP-CCB)

In this policy, all virtual links in a certain domain have a unique common cost (CC). As the previous mechanism, the dissemination of the updated abstraction topology is triggered once the number of allocated connections, M, exceeds the threshold N. For each domain, the PCE computes the whole set of virtual links cost and assigns the same CC to all of them. This single cost value is then disseminated to all the remainder PCEs. This is to provide a lower amount of control overhead, and at the same time, allows avoiding confidentiality issue among domains (operators).

The disseminated CC is calculated (according to (1)) as the average cost (AC) between the set of virtual links in a domain (T) divided by the nodal degree (ND) of that domain.

$$ND = \frac{\text{Number of Links}}{\text{Number of Nodes}}; AC = \frac{\sum_{i=1}^{T}(\text{Virtual Link Cost})_i}{\text{Number of Virtual Links}}; CC = \frac{AC}{ND}. \qquad (1)$$

The (Virtual Link Cost)$_i$ is also calculated assuming the shortest path cost between border nodes in accordance with the resource occupancy in that specific domain. CC is different in each domain due to the physical intra-domain topology; therefore a dynamic network load balancing distribution of the number of allocated transit connections, in comparison to the physical available resources, is created.

From the implementation point of view, there are still several open issues to adapt the GMPLS routing protocol to the multi-domain requirements. However, some work has been already done on OSPF-TE by the IETF [8], [9].

Two scenarios have been considered as feasible candidates to deploy the aforementioned mechanisms, namely, OSPF-TE areas based scenario and an ASes based one. In the former a new summary opaque LSA with inter-area scope would be needed to disseminate the virtual topology information. In the latter, the already defined type 11 opaque LSAs which have AS external scope [10] could be used to deal with the virtual topology dissemination.

3. PERFORMANCE STUDY AND DISCUSSION

This section evaluates the performance of IDTP and IDTP-CCB by running OMNeT++ simulations in a 9-domain transport network composed of 61 nodes and 95 links (19 inter-domain) carrying each one 8 wavelengths per link. The global network abstraction is composed by 78 links and each PCE manages in average 8.6 virtual links. For comparison purposes, we also evaluate the per-domain approach and the Static Balancing (SB) approach. In the per-domain approach, the end-to-end path is computed during the signalling phase; the responsible PCE of each domain performs the path computation for the segment of the end-to-end connection that crosses it, providing therefore connectivity to the next domain. The per-domain approach does not require any shared information among the domain. The SB policy is based on the computation of a unique cost CC for the whole set of virtual link costs in the domain full-mesh abstraction. More specifically, the same cost for the virtual links, as the one in IDTP-CCB is computed and then disseminated during the network boot-up phase; as a consequence, such costs are statically configured and the updated cost values are not disseminated during the network operation. In all simulations, 10^5 Poisson connection requests are generated following 70/30% intra/inter-domain ratio. For inter-domain connections, source and destination domains are uniformly selected and source/destination nodes are uniformly chosen in their respective domains. Source and destination nodes are randomly selected for intra-domain connections. All generated requests demand a whole wavelength capacity and mean holding time (HT) is set to 600 seconds following an exponential distribution. Request inter-arrival time is also exponentially distributed and varies with the network offered load. Figure 1 (left) depicts the overall network connection blocking probability (B_p) achieved by the per-domain, IDTP, IDTP-CCB and SB respectively, as a function of the global offered load. For the IDTP policy, the case $N = 1$ and $N = 5$ have been considered. Having only local domain information and visibility, per-domain approach leads to significantly higher blocking probability compared to the others. For instance, fixing the load at 150 Erlang, the B_p in per-domain is more than ten times higher than the mechanisms which rely on the dissemination of the domain virtual topology. Moreover, as expected, IDTP with $N = 1$ yields the best B_p figures, due to its immediate virtual link update policy. In fact, IDTP with $N = 1$ updates its virtual link cost each time that a connection is allocated in a domain. Nevertheless the amount of network control overhead due to the dissemination messages among the PCEs dramatically increases with the offered load. IDPT-CCB performs a closely identical B_p in relation to IDTP with $N = 1$ but it drastically reduces the network overhead. To make clearer this situation, fixing the network load at 250 Erlang the difference between IDTP $N = 1$ and IDTP-CCB is about 9%. In addition, if we the N parameter is increased, to reduce the network overhead, the difference in terms of B_p between IDTP with $N = 5$ and IDTP-CCB increases. As seen, for a fixed load of 200 Erlang, IDTP-CCB reduces the B_p a 6% in front of IDTP with $N = 5$. Additionally, Figure 1 (left) also depicts the global B_p performance by SB which is slightly higher than the other path computation methods having abstracted topological information, but it is still much better than the per-domain approach. As previously mentioned, SB only delivers a CC update message to other PCEs when the network is firstly configured, trying in this way, to reduce the overhead while increasing scalability and confidentiality in large multi-domain scenarios. Interestingly, SB is the most scalable solution but the loss of topology information reduces the capacity of computing inter-domain end-to-end paths efficiently. As depicted in Figure 1 (left), SB increases the B_p more than 32% for 300 Erlang in comparison to IDTP $N = 1$ as well as keeps lowest levels of network overhead.
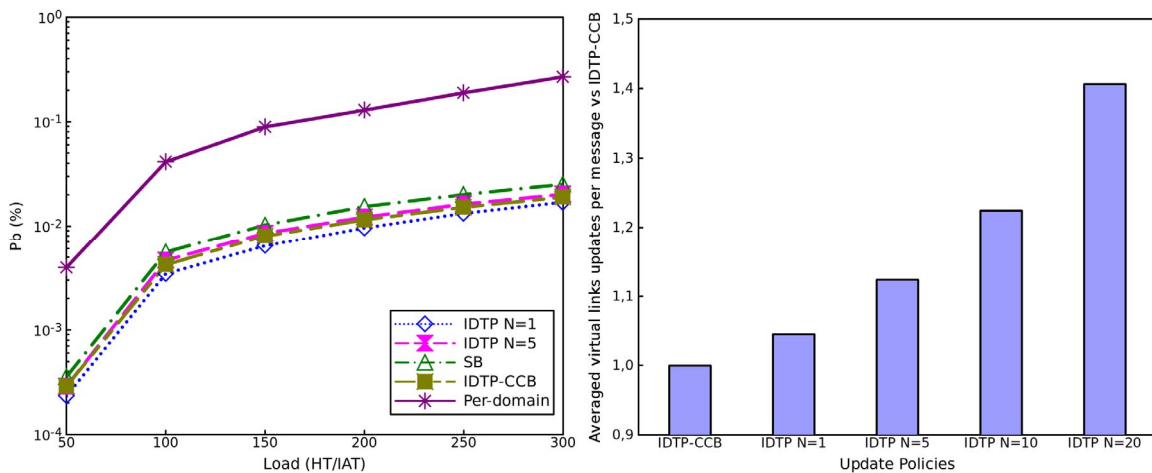

Figure 1. Network connection blocking probability improvement compared to per-domain mechanism (left); Average number of virtual links cost fields between different triggering policies (right).

Figure 1 (right) reflects the average number of virtual links cost fields performed by IDTP-CCB (using only one cost field for each virtual link) and IDTP $N = 1, 5, 10, 20$ (using as much as virtual link cost fields has each

domain full-mesh abstraction). This has been quantified to serve the same set of offered intra- and inter-domain connections under a fixed offered load of 200 Erlang ($B_p \approx 1\%$). As expected, for the IDTP with $N = 1$ policy, an increment of about 5% in terms of the present fields in its distributed opaque LSAs is experimented in front of IDTP-CCB policy, that only sends a cost field including the same CC for the whole set of virtual links. Assuming an increment of the parameter N, IDTP based solutions require on average a higher number of fields (virtual link costs) in their opaque LSAs while IDTP-CCB presents a flat performance maintaining just one field (since the disseminate cost is unique) in each opaque LSA. As seen in Figure 1 (right), for IDTP with $N = 5$, the size of the virtual link cost fields is about 15% higher that the IDTP-CCB policy.

4. CONCLUSIONS AND FUTURE WORK

In this paper, we compared different updated domain topological information triggering policies which improve the routing scalability in multi-domain optical networks. To compute optimal end-to-end paths crossing different domains, the exchange of some aggregated topological information of each domain is required, at expenses to increase the network control overhead among the PCEs. The policies discussed in this paper allow reducing the control messages exchanges while keeping the confidentiality requirements among the domain. On the basis of the simulation results, IDTP-CCB is presented as the best solution to reach a trade-off between network overhead (and thus routing scalability) and global connection blocking probability. At the same time, it allows to face in an efficient way the confidentiality concerns among domain operators. Future efforts will be devoted to design more advanced triggering update policies, able to provide even better solutions in terms of blocking probability and routing scalability. The extensions to current standard GMPLS protocols required to disseminate the information to be shared will be also defined.

ACKNOWLEDGEMENTS

REFERENCES

[1]   A. Farrel, JP. Vasseur, J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, Aug. 2006.
[2]   E. Mannie, "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, Oct. 2004.
[3]   JP. Vasseur, A. Ayyangar, R. Zhang,"A per-domain Path Computation Method for Establishing Inter-Domain Traffic Engineering (TE) Label Switched Paths (LSPs)",RFC 5152, Feb. 2008.
[4]   JP. Vasseur, R. Zhang, N. Bitar, JL. Le Roux, "A Backward-Recursive PCE-Based Computation (BRPC) Procedure to Compute Shortest Constrained Inter-Domain Traffic Engineering Label Switched Paths", RFC 5441, Apr. 2009.
[5]   The OIF Forum, 'External network-network interface (E-NNI) OSPF-based routing - 1.0 (intra-carrier) implementation agreement,' OIF-ENNI-OSPF-01.0 (OIF, 2007), http://www.oiforum.com/public/documents/OIF-ENNI-OSPF-01.0.pdf .
[6]   Q. Liu, N. Ghani, "Topology abstraction schemes in multi-domain full wavelength conversion DWDM networks," in *IEEE Symposium on Advanced Networks & Telecommunications Systems (ANTS) 2007*, Bombay, India, Dec. 2007.
[7]   G. Maier, C. Busca, A. Pattavina, "Multi-domain routing techniques with topology aggregation in ASON networks," in *Optical Network Design and Modelling (ONDM 2008)*.
[8]   D. Katz, K. Kompella, D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, Sep. 2003.
[9]   K. Kompella, Ed., and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, Oct. 2005.
[10] L. Berger, I. Bryskin, A. Zinin, and R. Coltun, "The OSPF Opaque LSA Option", RFC 5250, July 2008.