

Avoiding Path-Vectors in Multi-Domain Optical Networks

Marcelo Yannuzzi and Xavi Masip

*Technical University of Catalonia, Avda. Víctor Balaguer s/n, 08800 Vilanova i la Geltrú, Spain
{yannuzzi, xmasip}@ac.upc.edu*

Abstract: We show that a modified version of a path-vector protocol can drastically reduce the blocking and converge significantly faster, while exchanging less number of routing messages both during failure-free conditions and during a convergence.

© 2008 Optical Society of America

OCIS codes: (060.4250) Networks; (060.4251) Networks, assignment and routing algorithms

1. Introduction

One of the key building blocks of optical internetworking is the Routing and Wavelength Assignment (RWA) protocol. The consolidation of the Border Gateway Protocol (BGP) has motivated the proposal of its optical extension, namely, OBGp [1, 2]. We claim, however, that path-vector protocols like OBGp will not be sufficient for multi-domain optical networks, since the latter inherits the well-known weaknesses of BGP. These weaknesses are rooted not only in the intrinsic properties of path-vectors, but also, in some decisions made while implementing protocols like BGP or OBGp.

Among the intrinsic issues, the most important is that path-vectors were primarily designed for the exchange of reachability information. At present, there is wide consensus that optical domains must be able to exchange both reachability as well as aggregated Path-State Information (PSI). Understanding that the computation and advertisement of PSI is a missing piece in path-vector protocols is relatively easy nowadays, but contributing with solutions capable of highly improving the performance of a path-vector without impacting on key aspects of the protocol like its scalability, its convergence properties, and the number of routing messages exchanged between domains is a challenging task. On the other hand, among the most important issues related to the implementation of path-vectors is the path exploration phenomenon, which is responsible for the slow convergence of protocols like BGP and OBGp.

We have recently shown that even minor modifications to a path-vector protocol are sufficient to drastically improve its performance. In [3] we introduced OBGp+, which is a modified path-vector supporting the computation and advertisement of PSI between optical domains. In [4] we went one step further, by proposing a new route control paradigm together with a modified path-vector especially designed for the exchange of both reachability and PSI between optical domains. In those works we showed that, under failure-free conditions, a modified path-vector such as OBGp+ outperforms a pure path-vector like OBGp.

Differently from our previous works, we focus here on comparing the properties of OBGp and OBGp+ under highly stressful conditions. More precisely, we analyze the behavior of the protocols when a node can fail. The assessment of the impact of this kind of event on the inter-domain RWA protocol is particularly important, since applying protection and restoration techniques to all the lightpaths using a node might not be feasible in practice—e.g., due to budget constraints. As far as we know, this is the first study where the behavior and convergence properties of an optical path-vector are examined and improved under highly stressful conditions. In this sense, this paper complements the research started in [3] and [4]. For a detailed description of the algorithmic aspects of OBGp and OBGp+, the reader is referred to [3].

2. Performance Evaluation

The inter-domain setting chosen for the tests is the PAN European reference network illustrated in Fig. 1a, where we have conducted extensive simulations using OPNET Modeler. This multi-domain network is composed by 28 domains and 41 inter-domain links, and we used 5 fibers per-link, and 12 wavelengths per-fiber throughout the entire network. For the network topology inside each domain in the PAN, we have randomly chosen a minimum number of OXCs equal to the number of inter-domain links of that domain, up to a maximum of 6 OXCs. This approach guarantees that each inter-domain link is supported by a different border OXC, and at the same time offers a reasonable trade-off between the size of the network and the time needed to run the simulations. The OXCs inside each domain were connected by a full-mesh.

In this setting, we have randomly placed 18 sources and 10 destinations, covering the entire PAN European network with a source or a destination OXC inside each of its 28 domains. The overall set up process was repeated 10 times,

OWL5.pdf

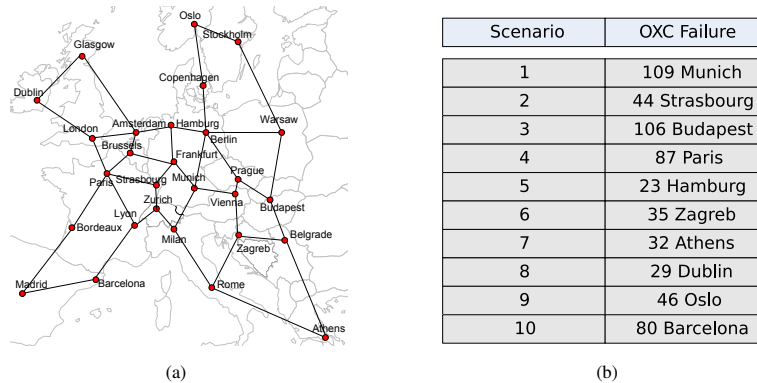


Fig. 1. a) PAN European reference network; b) Different scenarios for assessing the convergence properties of OBGP and OBGP+.

and for each of these settings we have randomly torn down a node during the simulation run-time. The resulting set of failing nodes is listed in Fig. 1b. The table shows the node number together with the domain where the failure occurred. All the results shown in Fig. 2 are the averages over those 10 settings. Our interest is to compare four different performance metrics, namely:

- (i) The Blocking Ratio (BR) of inter-domain lightpath requests (cf. Fig. 2a).
- (ii) The overall number of routing messages exchanged to achieve this blocking (cf. Fig. 2b).
- (iii) The convergence time after a node failure (cf. Fig. 2c).
- (iv) The number of routing messages exchanged during a convergence (cf. Fig. 2d).

As shown in Fig. 2, the trials were performed for different traffic loads, varying from 100 Erlangs up to 300 Erlangs, where traffic was modeled according to a Poisson distribution with exponentially distributed arrival and departure rates. Figure 2a shows the overall blocking of inter-domain lightpath requests for the different traffic loads assessed. It is worth highlighting that the results shown in Fig. 2a take into account the blocking experienced during the whole simulation run-time, i.e., before, during, and after the node failure. Clearly, OBGP+ outperforms OBGP. Whereas OBGP experiences blocking for all the traffic loads tested, OBGP+ starts to show some negligible blocking only after reaching 200 Erlangs.

In [3] and [4] we have shown that, under failure-free conditions, our modified path-vectors always need less number of routing messages than OBGP. Figure 2b confirms that, even under node failures, OBGP+ still exchanges less overall number of routing messages than OBGP. Once again, the results shown in Fig. 2b correspond to the total number of routing advertisements during the simulation run-time, i.e., those exchanged before, during, and after the node failure.

In order to assess the convergence properties of both protocols, we measured the time that elapses between a node failure, and the instant when the last message originated by this event is processed. Figure 2c shows the convergence time of OBGP and OBGP+. The results obtained confirm that OBGP+ needs to explore less paths, hence converges considerably faster than OBGP. It is important to notice that even for a small multi-domain optical network (see Fig. 1a), the time need by a path-vector to converge upon a node failure is in the order of seconds. OBGP+ is able to reduce considerably the OBGP convergence time, but still, its convergence remains in the order of seconds. This is because the main difference between OBGP and OBGP+ is that the latter effortlessly incorporates PSI in the RWA process (see [3] for the details), hence OBGP+ still suffers from the path exploration issues present in both BGP and OBGP.

Following the same approach described above, we assess the number of routing messages exchanged by both OBGP and OBGP+ during a convergence. The results are shown in Fig. 2d. We observe that, in general, OBGP+ needs to process less number of routing messages, which is consistent with the previous finding that OBGP+ shows shorter convergence times than OBGP.

3. Conclusion

This paper complements the research that we started in [3] and [4]. In those works we showed that even minor variations to a path-vector protocol like OBGP, are enough to drastically improve its performance under failure-free conditions. In this paper, we have extended our previous work by assessing the performance of OBGP and OBGP+ under highly stressful conditions, like the event of a node failure.

OWL5.pdf

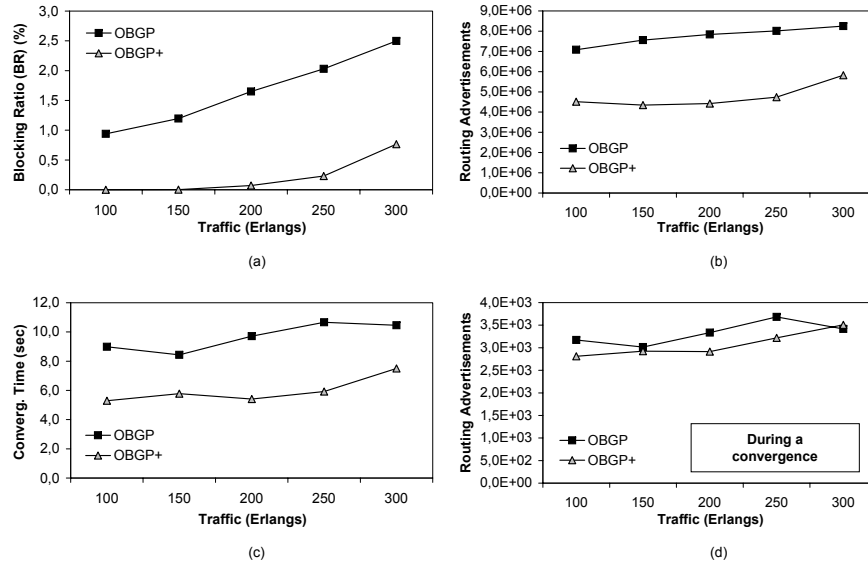


Fig. 2. a) Overall blocking ratio (%); b) Total number of routing advertisements exchanged, i.e., considering the messages before, during, and after the node failure; c) Convergence time (in seconds) during a node failure for the scenarios considered in Fig 1b; and d) Number of routing advertisements exchanged during a convergence.

We have shown that significant improvements can be achieved without needing to exchange more routing messages than with OBGP. In particular, our work shows that:

- (i) OBGP+ drastically reduces the blocking experienced with OBGP, both under failure-free [3] and under node failure conditions.
- (ii) OBGP+ reduces the overall number of routing messages exchanged to achieve this blocking, and this claim stands also under failure-free [3] and under node failure conditions.
- (iii) OBGP+ converges considerably faster than OBGP under the highly stressful event of a node failure.
- (iv) In general, OBGP+ also exchanges less number of routing messages than OBGP during a convergence.

When considering the results found in [3], [4], and in this paper, jointly, it becomes clear that we should avoid relying on plain path-vectors to support the inter-domain RWA protocol. Our research demonstrates that even minor variations (e.g., OBGP+) are sufficient to outperform a plain path-vector like OBGP. Although these are promising findings, the research in this direction has only started. Our main results and conclusions apply to a small multi-domain optical network (the PAN European topology shown in Fig. 1a), so further studies are needed to analyze the performance of the proposals made here—or more advanced ones—in a large-scale environment composed by thousands of domains. This is particularly challenging in terms of computational power, especially when running event-driven simulations (see the discussion in [4]). Moreover, we have not yet considered the more general problem of RWA and wavelength converters placement, which introduces many open problems given the distributed nature of routing policies between optical domains. We plan to explore these issues in the near future.

Acknowledgments This work was partially funded by the Spanish Ministry of Science and Technology under contract TEC2005-08051-C03-01. The authors also acknowledge the support received from OPNET Technologies Inc., in using OPNET Modeler.

References

1. M. Blanchet, F. Parent, and B. St-Arnaud, "Optical BGP (OBGP): InterAS lightpath provisioning," IETF draft, ietf-draft-parent-obgp-01, March 2001.
2. L. Wang, H. Zhang, and L. Zheng, "Reducing the OBGP protection switching time in WDM mesh networks," in Proc. of OFC, Anaheim, CA, USA, March 2006.
3. M. Yannuzzi, X. Masip-Bruin, G. Fabrego, S. Sanchez-Lopez, and J. Sole-Pareta, "OBGP+: A simple approach to drastically improve OBGP," in Proc. of IFIP/IEEE ONDM 2008, Vilanova i la Geltrú, Spain, March 2008.
4. M. Yannuzzi, X. Masip-Bruin, G. Fabrego, S. Sanchez-Lopez, A. Sprintson, and A. Orda "Toward a new route control model for multi-domain optical networks," in IEEE Communications Magazine, 46(6):104-111, June 2008.