

OBGP+: A Simple Approach to Drastically Improve OBGP

M. Yannuzzi, X. Masip-Bruin, Guillem Fabregó, Sergio Sánchez-López, and Josep Solé-Pareta

Department of Computer Architecture, Technical University of Catalonia (UPC)

Av. Víctor Balaguer s/n, 08800 Vilanova i la Geltrú, Barcelona, Spain.

Email: {yannuzzi, xmasip, gfabrego, sergio, pareta}@ac.upc.edu

Abstract—During the last few years, the research community has shown an increasing interest in the subject of inter-domain routing in optical networks, so new approaches like the optical extension of BGP, namely *OBGP*, have arisen. However, it is widely accepted now that a multi-domain routing model mostly centered on the exchange of reachability information – like the one we have today or the one provided by OBGP – is not going to be enough in the future. Routing domains must be able to exchange both, reachability, as well as aggregated *Path State Information* (PSI). Understanding that there is a missing piece in the routing models provided by BGP and OBGP is easy nowadays, but contributing with solutions capable of highly improving the performance of these routing protocols without increasing the number and frequency of the routing messages exchanged between domains is a challenging task.

This study makes the following contributions. First, we propose a straightforward strategy to compute highly aggregated PSI between routing domains. Second, we propose *OBGP+*, an extended version of OBGP that exploits this PSI to compute inter-domain lightpaths in a highly efficient way. Third, in order to avoid the typical increase in the number of routing messages associated with the update of PSI, we propose to piggy-back these updates in non-dummy Keepalive messages exchanged between *OBGP+* neighbors. Extensive simulations made with OPNET in the PAN European network topology reveal that: i) *OBGP+* is able to drastically reduce the blocking experienced with OBGP; ii) while even needing less number of routing messages than OBGP to achieve this performance.

I. INTRODUCTION

Future optical networks will require new protocols in order to route and support the provisioning of lightpaths on-demand between different domains. In light of this, some researchers have started to analyze the possibility of adopting an *Optical Border Gateway Protocol* (OBGP) as the future inter-domain routing protocol for optical networks [1]–[4]. The aim of these proposals is to extend BGP so that it can convey and signal optical path information between OBGP neighbors. The strength of this approach is that future optical networks will benefit from the well-known advantages of the BGP-based routing model, such as scalability, clear administrative limits of routing domains, fully-distributed network administration based on filtering and routing policies, etc. The weakness, on the other hand, is that the routing model of future optical networks will

This work was partially funded by the MCyT (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.

inheritate the well-known issues in BGP [5]. Indeed, a multi-domain routing model mostly centered on the exchange of reachability information, like the one offered by BGP, is not going to be enough. This is confirmed by a number of research initiatives recently started, like [6] and [7]. Accordingly, in this paper we consider that neighboring domains are able to exchange both *Network Reachability Information* (NRI), and enriched routing information consisting of aggregated *Path State Information* (PSI).

The subject of this study is to show that by integrating only plain and highly aggregated PSI in OBGP, it is possible to drastically improve its performance¹ (see Fig. 3), and this can be accomplished without increasing the number or the frequency of routing updates exchanged between domains (see Table I). We hope that the results presented here would encourage other researchers to devise novel ways of integrating PSI in the inter-domain routing model without impacting on its capability to scale.

The rest of the paper is organized as follows. Section II reviews related work. Section III presents the routing model, including the NRI and the PSI exchanged between neighboring domains. Then, in Section IV the Routing and Wavelength Assignment (RWA) algorithm proposed for *OBGP+* is detailed. The comparison between OBGP and *OBGP+* is shown in Section V. Finally, Section VI summarizes the main conclusions and describes our future work.

II. RELATED WORK

Most of the research and standardization efforts carried out so far in the area of routing in optical networks have been focused on intra-domain aspects. The discussions concerning multi-domain issues are in a very early stage yet, so despite some topics have started to be analyzed by the three relevant standardization bodies, namely, the International Telecommunications Union (ITU) [8], the Internet Engineering Task Force (IETF) [9], and the Optical Internetworking Forum (OIF) [10], the situation is that the majority of the multi-domain issues are largely open at present.

A few years ago, the OIF proposed the Domain-to-Domain Routing Protocol (DDRP) [11]. DDRP is basically a hierarchical extension of OSPF-TE, supported by a modified version

¹The performance metric considered here is the blocking ratio of inter-domain lightpath requests.

of Dijkstra’s algorithm. However, DDRP has mainly two drawbacks. First, it represents a major change in the routing system, since it proposes to move towards a fully hierarchical model. Second, the modified Dijkstra algorithm still offers limited functionalities. For instance, it returns a single path at a time, so complementary algorithms need to be adopted for path protection and diverse routing purposes.

On the other hand, some recent works have proposed to extend the reach of BGP so that it can be used in the framework of multi-domain optical networks [1]–[4]. However, neither OBGp nor DDRP will be able to provide the routing functionalities and expected performance for multi-domain optical networks. This has leveraged the proposal of different path state aggregation schemes and updating policies at the inter-domain level for wavelength-routed optical networks [12], [13].

With a more conservative approach, the standardization efforts being carried out at the IETF mention the need to work on new protocols, or extensions to the existing ones, in order to enable the advertisement of inter-domain Traffic Engineering (TE) information. In [14], the authors mention the possibility of adding TE extensions to BGP.

Given that inter-domain routing in optical networks is becoming an active research area, a sound approach is to start analyzing the issue of how routing protocols can be endowed with the ability to compute and efficiently convey aggregated PSI between domains. This is precisely the subject of this paper.

III. THE ROUTING MODEL IN OBGp+

The routing information exchanged by means of OBGp+ fulfills the following requirements:

- (i) PSI must be advertised between domains in addition to the usual reachability information.
- (ii) The PSI received from downstream domains must be assembled and aggregated together with local PSI, and advertised to upstream domains.
- (iii) This PSI flow must supply a standardized coupling between the different segments along a lightpath. This will support the computation of end-to-end optical paths in an efficient way.
- (iv) The PSI exchanged must be completely independent of the intra-domain routing and signaling protocols. In this sense, enhancements or even a complete replacement of any of the protocols used inside a domain must not affect the routing information exchange model between domains.
- (v) Special care must be taken while developing aggregated PSI schemes, and while deciding the frequency of the updates associated with the routing information sent across domain boundaries.

These issues are the subject of the study in this section. A detailed description of the RWA algorithm provided by OBGp+ shall be introduced in Section IV.

OBGP+ is responsible for distributing inter-domain routing information, and deciding within each domain the best path to reach a destination. To this end, the routing advertisements in OBGp+ contain the usual NRI in addition to PSI.

The role of the PSI is to capture the “state” of resources along an inter-domain path. During the composition of the advertisements, OBGp+ nodes aggregate the PSI along a path taking into account the state of both the intra- and the inter-domain segments of the path. The advertised PSI is rich enough so that upstream domains can drastically reduce the number of blocked lightpath requests, and at the same time is sufficiently aggregated so that administrative limits and business protection considerations of domains are respected.

The flow of routing advertisements between OBGp+ nodes from a destination domain AS_D towards a source domain AS_S can be summarized as follows. A border node in AS_i assembles the PSI received from AS_{i+1} with its local PSI, and advertises AS_{i-1} the aggregate: $PSI^{(i-1)} = PSI^{(i)} \oplus PSI^{(i+1)}$, where the operator \oplus denotes an appropriate PSI assembling function. The data conveyed in the PSI as well as the strategy to update them are detailed later in Section III-B.

A. Network Reachability Information (NRI)

In this paper we assume that the optical nodes, namely, the *Optical Cross-Connects* (OXC) do not perform wavelength conversion, so each lightpath computed by OBGp+ is subject to the wavelength continuity constraint. We proceed now to describe the NRI and the aggregated PSI conveyed by OBGp+.

Let L , F , and Ω denote the number of links, the number of fibers per-link, and the number of wavelengths per-fiber, respectively, at each destination OXC. For the sake of simplicity we assume that all destination OXC are identical, and that each network sinking traffic is connected to only one OXC². Thus, $LF\Omega$ is an upper bound of the number of available wavelengths to reach any destination within a domain. Each AS may select – according to its local TE and routing policies – the particular subset of wavelengths that can be used by an upstream domain to reach the local networks. Consequently, the reachability information contained in the NRI messages conveyed by OBGp+ consists of:

- (i) The set of destination networks $\{d\}$ and their associated AS-path.
- (ii) The Next-Hop (NH) to reach those destinations, i.e., the address of the ingress OXC in the neighboring domain from which the advertisement was sent. It is worth noticing that the NH concept is basically the same as in the case of BGP.
- (iii) A set of pairs $(\Lambda_1, M_{\Lambda_1}), \dots, (\Lambda_N, M_{\Lambda_N})$ available for each destination d , where Λ_i , $i \in \{1, \dots, N\}$ denotes a particular wavelength, and M_{Λ_i} denotes the maximum multiplicity of Λ_i . Clearly, $N \leq \Omega$, and $M_{\Lambda_i} \leq LF \forall i$.

In sum, the NRI distributed between OBGp+ nodes is composed by:

²The routing information exchange model described here can be easily generalized if these assumptions are not met.

$$\Phi_{NRI}(d) = \left[\text{AS-path, NH, } (\Lambda_i, M_{\Lambda_i}) \right]_d \quad (1)$$

For each destination network, a transit AS may filter and advertise a subset of Φ_{NRI} to its upstream domains, or simply retransmit the NRI messages received. When a new destination network becomes available, or an already known one becomes unavailable, the NRI messages are triggered immediately by OBGp+. In any other case, the NRI should only change over large timescales compared to the PSI, according to the local optimizations and TE actions performed by the different routing domains.

Henceforth, and especially during the performance evaluation in Section V, we assume that both OBGp and OBGp+ handle exactly the same NRI and treat it exactly in the same way.

B. Aggregated Path State Information (PSI)

The PSI is composed by aggregated wavelength availability information. OBGp+ advertises PSI messages by aggregating and assembling the following three pieces of information:

- (i) Intra-domain PSI.
- (ii) PSI related to the inter-domain links towards its downstream domains.
- (iii) The already aggregated PSI contained in the inter-domain advertisements received from downstream domains.

The aggregation process operates as follows. Let r and q be a pair of OXCs inside an AS, $P(r, q)$ be a candidate path between r and q , and l be a link within the path $P(r, q)$. OBGp+ nodes compute the *Effective Number of Available Wavelengths* (ENAW) of type Λ_i between the OXCs r and q as follows:

$$W_{r,q}(\Lambda_i) = \max_{P(r,q)} \left\{ \min_{l \in P(r,q)} [W_l(\Lambda_i)] \right\} \quad (2)$$

The rationale in (2) can be easily interpreted by means of Fig. 1. For instance, in AS1 the ENAW of type Λ_1 between the nodes $OXC15$ and $OXC12$ is $W_{15,12}(\Lambda_1) = 3$. This is because from the two possible paths between these nodes, the path that goes through $OXC13$ has a minimum $W_{13,12}(\Lambda_1) = 1$, whereas the one that goes through $OXC11$ has a minimum $W_{11,12}(\Lambda_1) = 3$. Then, the maximum between both of them is 3. The ENAW given in (2) is especially important between two border OXCs in a transit domain, since it captures the practical availability of the wavelength Λ_i within the domain. In addition, (2) offers highly aggregated network state information, so this is the intra-domain portion of the wavelength availability component of a PSI aggregate.

For the inter-domain portion, each OBGp+ node is aware of which wavelengths are actually being used on its inter-domain links, and it also knows which wavelengths are effectively available downstream through the PSI advertisements received

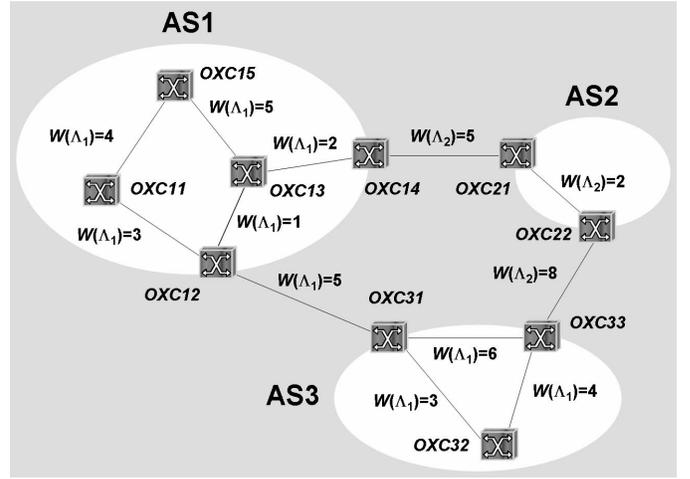


Fig. 1. Exploiting the PSI exchanged between routing domains.

from neighboring OBGp+ nodes. Let $W_{l_b, r_b}(\Lambda_i)$ denote the number of available wavelengths of type Λ_i in the inter-domain link between the local border node l_b , and a remote border node r_b . For instance, in Fig. 1 the OBGp+ nodes in AS1 are aware that $W_{12,31}(\Lambda_1) = 5$. Similarly, let $W_{r_b, d}^{adv}(\Lambda_i)$ denote the ENAW of type Λ_i between the remote border node r_b and the destination node d , advertised by r_b . Using these two inter-domain components and (2), an OBGp+ node advertises upstream that the ENAW between a local border node l_b and a distant destination node d is:

$$W_{l_b, d}^{adv}(\Lambda_i) = \min \left\{ W_{l_b, l'_b}(\Lambda_i), W_{l'_b, r_b}(\Lambda_i), W_{r_b, d}^{adv}(\Lambda_i) \right\} \quad (3)$$

For instance, in Fig. 1 the border node $OXC14$ advertises to its neighbor $OXC21$ in AS2 that the ENAW of type Λ_1 to reach $OXC32$ is:

$$W_{14,32}^{adv}(\Lambda_1) = \min_{\Lambda_1} \left\{ W_{14,12}, W_{12,31}, W_{31,32}^{adv} \right\} = \min\{2, 5, 4\} = 2. \quad (4)$$

In sum, the PSI received by an OBGp+ node for destination d is:

$$\Phi_{PSI}(d) = \left\{ W_{r_b, d}^{adv} \right\}_{\Lambda_i} \quad (5)$$

To advertise the PSI associated with the destinations contained in the NRI messages, we take advantage of the *Keepalive* messages exchanged between neighboring OBGp+ nodes. Similarly as in the case of BGP, OBGp+ nodes exchange *Keepalive* messages to confirm that neighboring nodes are still operative. In BGP, *Keepalive* messages are of fix length, consisting only of the 19-byte BGP header. In our

OBGP+ model, we extend the BGP Keepalive concept with the purpose of conveying PSI, when relevant PSI needs to be updated. In other words, the update of PSI is supported by the exchange of non-dummy Keepalive messages between routing domains. A major advantage of this strategy is that it does not increment the number of routing messages exchanged between domains.

In Section V we shall show that when the RWA algorithm supporting OBGP+ exploits the highly aggregated PSI in (5), it is possible to achieve drastic reductions in the number of blocked inter-domain lightpath requests compared to that obtained with OBGP.

IV. ROUTING AND WAVELENGTH ASSIGNMENT STRATEGY

Similarly as BGP, OBGP is essentially a shortest AS-path routing algorithm that exchanges NRI, but it does not handle PSI. Our OBGP+, however, handles the highly aggregated PSI supplied by the ENAW introduced in Section III-B. Accordingly, each OBGP+ node computes and advertises the ENAW along the candidate paths, as described in (2) and (3). Algorithm 1 shows a simplified version of the OBGP+ decision process. This algorithm is the result of a set of enhancements that we introduced to OBGP [1].

From Algorithm 1, it is clear that OBGP+ is essentially a “shortest AS-path highest ENAW” RWA algorithm, given that it usually prefers the shortest AS-path (step 2 of the algorithm),

Algorithm 1 OBGP+({ $P(s, d)$, Λ_i , M_{Λ_i} , z_i })

Input: $\{P(s, d)\}$ - set of paths between nodes s and d
 Λ_i - a particular wavelength on path $P(s, d)$
 M_{Λ_i} - Multiplicity of wavelength Λ_i on path $P(s, d)$
 z_i - ENAW of type Λ_i along the path $P(s, d)$
Output: (P^{best} , Λ^{best}) - The best lightpath between s and d

- 1: Choose the (path, wavelength) pair with the highest local preference (LOCAL_PREF) /* As in BGP */
 - 2: If the LOCAL_PREFs are equal, choose the shortest AS-path and assign the wavelength with the highest ENAW among the ones available on that path. If more than one wavelength has the same (highest) ENAW along the shortest AS-path, choose the wavelength with the lowest identifier i
 - 3: If the AS-path lengths are equal choose the (path, wavelength) pair associated with the highest ENAW
 - 4: If the ENAWs are equal prefer external paths over internal paths
 - 5: If the paths are still equal prefer the one with the highest ENAW to the next-hop OXC (i.e., to the OXC r_b in the neighboring domain)
 - 6: If more than one path is still available run OBGP tie-breaking rules /* As in BGP */
-

but if more than one candidate lightpath exists, then it chooses the one with the highest ENAW (step 3).

V. PERFORMANCE EVALUATION

The aim of this section is to contrast the performance of OBGP+ against OBGP. Our interest here is to compare two different performance metrics, namely, the *Blocking Ratio* (BR) of inter-domain lightpath requests, and the number of routing messages exchanged to achieve this blocking.

To this end, we have conducted extensive simulations using OPNET Modeler [15]. The simulation results presented here can be reproduced using the modules that we developed in OPNET, which are available online from [16].

The inter-domain scenario chosen for the trials was the complete PAN European network topology illustrated in Fig. 2 [17]. This multi-domain network is composed by 28 domains and 41 inter-domain links. 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 inside each domain. This approach guarantees that each inter-domain link of a domain in the PAN is supported by a different border OXC. In this setting, we have randomly placed 18 sources and 10 destinations, covering in this way 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 100 times, and the results shown here are the averages over those 100 settings for both the BR and the number of routing messages exchanged during the simulation run-time. It is worth emphasizing that these results are the outcome of more than 900 hours of event-driven simulations, so our criterion to choose a maximum of 6 OXCs per-domain reflects the trade-off between the size of the network, and the time needed to run the tests.

We have used 5 fibers per-link, and 12 wavelengths per-fiber throughout the entire PAN European network. In order to assess the impact of the frequency of update in the PSI, we have used different *Keepalive Update Intervals* K_T during the trials. K_T

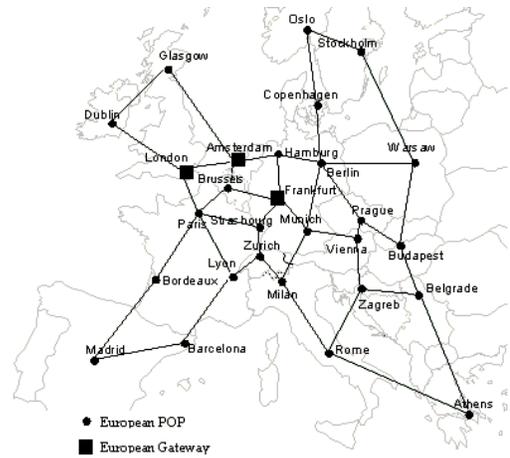


Fig. 2. PAN European network topology.

corresponds to the time interval between the delivery of non-dummy Keepalive messages conveying PSI. At present, most implementations of BGP use a default Keepalive value of 60 seconds, and three consecutive Keepalive messages need to be lost so that a BGP router proceeds to shutdown a BGP session. In our simulations we have tested three different scaled and normalized values: $K_T = 1$, $K_T = 3$, and $K_T = 5$ units through the simulation run-time. Clearly, the higher the values of K_T , the more time is needed by OBGP+ nodes to detect an react when a neighbor becomes inoperative. Therefore, a major advantage of conveying PSI piggy-backed in Keepalive messages is that low values of K_T are desired both to increase the responsiveness between OBGP+ neighbors as well as to support updating PSI more frequently.

As shown in Fig. 3, 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 3 shows the BR and the standard deviation of inter-domain lightpath requests obtained with OPNET, for the different traffic loads, and the different Keepalive update intervals K_T . Clearly, OBGP+ outperforms OBGP, and it becomes evident that even minor PSI, like the one proposed in this paper, is enough to drastically reduce the blocking obtained $\forall K_T$. Whereas OBGP experiences blocking for all traffic loads tested, OBGP+ starts to show some negligible blocking only after reaching 200 Erlangs.

Figure 3 also shows that the performance of both OBGP and OBGP+ degrade when K_T increases. The reason for this in the case of OBGP is the following. In OBGP [1], a source node is not aware of the subset of wavelengths $\mathcal{W}(P)$ that are no longer available along the different segments of an AS-path P . A source OBGP node will receive a reachability message indicating the withdrawal of path P , only after all the candidate wavelengths in that path have been consumed. A simple way to considerably reduce the BR experienced by OBGP is to update the subset $\mathcal{W}(P)$ through the Keepalive messages exchanged between OBGP neighbors. This approach not only offers more granular and updated NRI at the source OBGP node, but also, allows us to contrast the performance of OBGP and OBGP+ under more fair conditions. Our implementation of OBGP in OPNET follows this approach. As expected, Fig. 3 confirms that the blocking experienced by OBGP increases with K_T . The corollary is that in a regular implementation of OBGP (i.e., one where the subset $\mathcal{W}(P)$ is not conveyed and updated through the Keepalive messages) the BR yield by OBGP becomes rather independent of K_T , but it is much higher than that shown for $K_T = 5$ in Fig. 3.

In order to quantify the reductions supplied by OBGP+ in terms of blocking, we define the following *Improvement Factor* (IF):

$$IF \triangleq \left(\frac{BR^{(OBGP)}}{BR^{(OBGP+)}} \right)_{\text{Traffic (Erlangs)}} \quad (6)$$

Table I summarizes the improvement factor IF for 200, 250, and 300 Erlangs, as well as the number of routing messages exchanged for the different traffic loads and update intervals tested. The results show that:

- Up to 200 Erlangs, OBGP+ is able to reduce the BR by more than one order of magnitude.
- For 250 Erlangs, OBGP+ is able to reduce the BR by approximately one order of magnitude.
- Even for the highest traffic load assessed, i.e., 300 Erlangs, OBGP+ achieves an improvement factor that varies approximately between 3 and 5.

Table I also confirms that OBGP+ always needs less overall number of routing messages than OBGP. The reason for this is twofold. First, because PSI updates are never triggered between OBGP+ neighbors. Instead, they are piggy-backed in the Keepalive messages used in both OBGP and OBGP+.

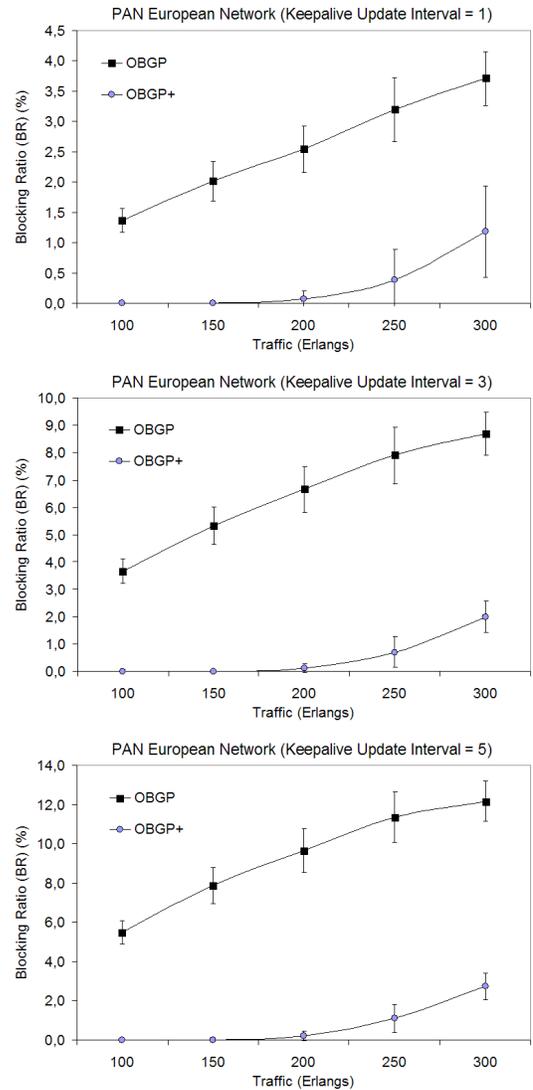


Fig. 3. Blocking ratio and standard deviation: Comparison between OBGP and OBGP+ for the different traffic loads and Keepalive Update Intervals (K_T).

	Keepalive Update Interval ($K_T = 1$)			Keepalive Update Interval ($K_T = 3$)			Keepalive Update Interval ($K_T = 5$)						
	200 Erlangs	250 Erlangs	300 Erlangs	200 Erlangs	250 Erlangs	300 Erlangs	200 Erlangs	250 Erlangs	300 Erlangs				
Improv. Factor (IF)	40.42	8.27	3.13	57.74	11.18	4.37	46.50	10.28	4.46				
Traffic (Erlangs)	Routing Messages OBGP	Routing Messages OBGP+	Routing Messages OBGP	Routing Messages OBGP+	Routing Messages OBGP	Routing Messages OBGP+	Routing Messages OBGP	Routing Messages OBGP+	Routing Messages OBGP	Routing Messages OBGP+			
	100	6,599,824	4,462,457	5,584,039	4,410,330	4,878,077	4,337,744	150	7,853,874	4,386,462	6,487,209	4,311,110	5,534,610
200	8,606,449	4,492,165	6,894,827	4,395,868	5,826,772	4,299,337	250	9,058,395	5,081,878	7,067,793	4,890,581	5,897,537	4,715,436
250	9,255,697	6,999,270	6,982,511	6,049,242	5,738,895	5,436,567	300						

TABLE I

IMPROVEMENT FACTOR IN THE BLOCKING REQUESTS FOR 200, 250, AND 300 ERLANGS, AND OVERALL NUMBER OF ROUTING MESSAGES EXCHANGED.

Second, OBGP tends to exhaust the available wavelengths along the shortest AS-path before switching to an alternative path. This triggers network reachability messages and path exploration after paths become blocked. Conversely, OBGP+ explicitly considers the ENAW in the RWA algorithm when two or more paths exhibit the same AS-path length, so it is able to provide a much better traffic distribution than OBGP, with drastic reductions in the BR, and hence, less network reachability messages need to be exchanged.

VI. CONCLUSIONS AND FUTURE WORK

In this paper we have shown that by endowing OBGP with the capability to compute, aggregate, and convey only minor path state information, it is possible to drastically reduce its blocking ratio. We have also shown that these significant improvements can be achieved without needing to exchange more routing messages than in a regular implementation of OBGP. In fact, OBGP+ reduces the number of routing messages exchanged between routing domains, given that by decrementing the blocking, it is possible to reduce the exchange of network reachability messages and path exploration when blocking starts to occur. This is possible due to the strategy of piggy-backing PSI updates in the Keepalive messages exchanged between OBGP+ neighbors.

Despite these promising findings, more research is needed in this direction. Our main results and conclusions apply to a rather small multi-domain optical network (the PAN European topology shown in Fig. 2), so further studies are needed to analyze the performance of the proposals made here in a large-scale environment composed by thousands of ASs. We plan to explore this in the near future.

Overall, multi-domain RWA approaches like the one proposed here, offer a promising line of work to address the trade-off between obtaining a low blocking ratio, and keeping the path state information as limited as possible.

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] M. J. Francisco, L. Pezoulas, C. Huang, and I. Lambadaris. "End-To-End Signaling and Routing for Optical IP Networks". In *Proceedings of IEEE ICC*, New York, USA, April 2002.
- [3] L. Wang et al. "A Novel OBGP-based mechanism for Lightpath Establishment in WDM Mesh Networks". In *Proceedings of ECOC 2003*, Rimini, Italy, September 2003.
- [4] L. Wang, H. Zhang, and L. Zheng. "Reducing the OBGP protection switching time in WDM mesh networks". In *Proceedings of OFC*, Anaheim, CA, USA, March 2006.
- [5] M. Yannuzzi, X. Masip-Bruin, and O. Bonaventure. "Open Issues in Interdomain Routing: A Survey". *IEEE Network*, 19(6):49-56, November-December 2005.
- [6] Future Internet, "The Future Networked Society: A white paper from the EIFFEL Think-Tank," available from: <http://future-internet.eu/>.
- [7] Future Internet Design: <http://www.nets-find.net/>.
- [8] International Telecommunications Union (ITU): <http://www.itu.int/>.
- [9] Internet Engineering Task Force (IETF): <http://www.ietf.org/>.
- [10] Optical Internetworking Forum (OIF): <http://www.oiforum.com/>.
- [11] G. Bernstein et al. "Domain to Domain Routing using GMPLS, OSPF Extension V1.1(Draft)". OIF2002.23.06, July 2002.
- [12] G. Liu, C. Ji, and V. Chan. "On the Scalability of Network Management Information for Inter-Domain Light-Path Assessment". *IEEE/ACM ToN*, 13(1), February 2005.
- [13] M. Yannuzzi, S. Sánchez-López, X. Masip-Bruin, J. Solé-Pareta, and J. Domingo-Pascual. "A Combined Intra-Domain and Inter-Domain QoS Routing Model for Optical Networks". In *Proceedings of IFIP/IEEE ONDM*, Milan, Italy, February 2005.
- [14] A. Farrel, J. P. Vasseur, and A. Ayyangar. "A Framework for Inter-Domain Multiprotocol Label Switching Traffic Engineering". IETF RFC 4726, November 2006.
- [15] Opnet Technologies, Inc. <http://www.opnet.com/>.
- [16] Opnet Modules developed: <http://www.cba.upc.edu/opnet>.
- [17] R. Hulsermann et al. "A Set of Typical Transport Network Scenarios for Network Modelling". *ITG Fachbericht*, issue 182, 2004.