Delft University of Technology

Inter-domain routing in optical networks wavelength converters

Beshir, AA; Yannuzzi, M; Kuipers, FA

Citation (APA)
Beshir, AA., Yannuzzi, M., & Kuipers, FA. (2010). Inter-domain routing in optical networks wavelength converters. In C. Xiao, & JC. Olivier (Eds.), *Proceedings ICC 2010* (pp. 1-5). IEEE. https://doi.org/10.1109/ICC.2010.5502242

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Inter-domain Routing in Optical Networks with Wavelength Converters

Anteneh Beshir*, Marcelo Yannuzzi†, and Fernando Kuipers*

*Network Architectures and Services
Delft University of Technology
Mekelweg 4, 2628 CD Delft, the Netherlands
{A.A.Beshir, F.A.Kuipers}@tudelft.nl

†Advanced Network Architectures Lab (CRAAX)
Technical University of Catalonia
Vilanova i la Geltrú, Catalonia, Spain
yannuzzi@ac.upc.edu

Abstract—With the increasing deployment of wavelength-division multiplexing (WDM) optical networks, the need for advanced lightpath provisioning algorithms and protocols in a multi-domain setting is becoming evident. In order to increase efficiency by relaxing the wavelength continuity constraint in WDM optical networks, wavelength converters are often placed at certain nodes in the network. In this paper, we study the efficiency of using converters in a multi-domain setting. We have made simple but important modifications to existing optical inter-domain routing protocols in order to utilize the power of wavelength converters and have tested their performance. These modifications can be seamlessly integrated into these protocols (i.e., without changing their algorithmic aspects) to significantly reduce their blocking ratio. We also show that there is a clear performance difference among the considered protocols.

I. INTRODUCTION

Optical networks using wavelength-division multiplexing (WDM) technology are being widely deployed within domains. Future optical networks will require new protocols in order to route and support on-demand provisioning of lightpaths between different domains. Unlike traditional IP multi-domain networks, the study of optical multi-domain issues is at a very early stage. One important issue is what type of information should be exchanged among neighboring domains in order to increase efficiency. Previous works [9] [10], have proposed approaches where neighboring domains are able to exchange both *Network Reachability Information* (NRI), and highly aggregated *Path State Information* (PSI). However, the presence of wavelength converters is not analyzed in these works. Our main contribution in this paper is to seamlessly incorporate modifications to the protocols proposed in [9] and [10], so that wavelength converters are utilized.

In WDM optical networks without wavelength converters, a lightpath has to use the same wavelength all along its path. This implies that lightpath requests may be blocked, even though there are unused wavelengths. In order to decrease the blocking ratio, wavelength converters are employed. Moreover, the optical signal can be regenerated at converter nodes to extend its reach. There are different methods for sharing a pool of wavelength converters at a given node among the wavelengths of its different fiber links [4]. Due to its sharing efficiency, we assume a share-per-node approach, where there is a single bank of converters at a given switching node shared

by all its links and only wavelengths that need to be converted are directed to this bank.

Since wavelength converters are costly (yet usually more affordable than adding fibers in already existing networks), we assume that for inter-domain traffic in a given domain, the wavelength converters are placed at border optical cross-connects (OXCs). This assumption is a realistic representation of emergent multi-domain optical networks [6]. Due to the large amount of traffic that goes through border OXCs, putting wavelength converters at the border OXCs is expected to have a significant performance improvement.

II. RELATED WORK

In the literature, there are only few works dealing with optical multi-domain networks; there are even fewer works that study the effect of wavelength converters. The three relevant standardization bodies, namely, the International Telecommunications Union (ITU), the Internet Engineering Task Force (IETF), and the Optical Internetworking Forum (OIF) have analyzed some of the topics related to multi-domain optical networks. In 2002, the OIF proposed the Domain-to-Domain Routing Protocol (DDRP). The drawbacks of DDRP are that it represents a major change in the routing system and it is not suitable for path protection. The IETF has proposed the generalized multi-protocol label switching (GMPLS) framework, which extends the features of multi-protocol label switching (MPLS) for provisioning circuit-switched connections via label abstractions for wavelengths, timeslots, etc. The ITU-T has specified a broad-based automatic switched optical network (ASON) framework. However, most of the research surrounding GMPLS and ASON is limited to intra-domain routing.

OBGP (Optical BGP) is an extension of BGP that has been proposed to "glue" multi-domain optical networks [1], [3], [8]. The strength of this approach is that future optical networks will benefit from the advantages of the BGP-based routing model, such as scalability, clear administrative limits of routing domains, etc. However, besides inheriting the well-known disadvantages of BGP, a multi-domain routing model mainly based on the exchange of network reachability information, which is currently the case in BGP, may not be sufficient. This has initiated the proposal of different path state aggregation

schemes and updating policies at the inter-domain level for WDM optical networks [5] [9] [10].

In [9], the authors showed that by integrating only plain and highly aggregated PSI in OBGP (in the form of an extended protocol called OBGP+), it is possible to drastically improve its performance, without increasing the number or the frequency of routing updates exchanged between domains. In [10], a novel distributed route control model is proposed, which is based on the deployment of inter-domain routing agents (IDRAs). We refer to the routing protocol running among the IDRAs as an IDRAs-based routing protocol (IDRP). IDRP is able to significantly reduce the blocking ratio compared to that of OBGP. However, mechanisms to take advantage of the presence of wavelength converters in these protocols were not developed.

In this paper, we make simple but important modifications that will allow OBGP+ and IDRP to benefit from the use of wavelength converters. The modifications are simple in that the algorithmic details of these protocols are not affected, and they are important because a significant reduction in the blocking can be achieved due to these modifications. We also show the performance gain obtained by having wavelength converters at border OXCs, and compare the performances of OBGP, OBGP+, and IDRP in the presence of wavelength converters.

In Section III, we give a brief description of OBGP+ and IDRP. In Section IV, we show how these protocols can be modified to take into account the presence of wavelength converters. In Section V, we present simulation results comparing the performance of the three protocols and also the improvement associated with having wavelength converters at the border OXCs. Finally, we give conclusions in Section VI.

III. OBGP+ AND IDRP

The major advantage of our approach is that our modifications can be seamlessly integrated in OBGP+ and IDRP. In other words, the algorithmic details of these protocols can be reused since our modifications concern only the wavelength aggregation process. For completeness and in order to introduce the notation used in Section IV, we give a brief introduction to OBGP+ and IDRP. For a detailed description of these protocols, the reader is referred to [9] and [10].

OBGP+ is an improved version of OBGP in that PSI is advertised besides the usual NRI exchanged in OBGP; whereas IDRP is a novel optical routing protocol that allows the exchange of useful traffic engineering (TE) information.

A. Network Reachability Information (NRI)

NRI messages are triggered when a new destination becomes available, or an already known one becomes unreachable. The reachability information contained in the NRI messages conveyed by OBGP+ consists of:

1) The set of destination networks $\{d\}$ and their associated autonomous system (AS)-path.
2) 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.

3) A set of pairs $(\lambda_i, W(\lambda_i))$ available for each destination d, where λ_i denotes a particular wavelength, and $W(\lambda_i)$ denotes the maximum multiplicity of λ_i.

Unlike BGP/OBGP, the NRI exchanged among the IDRAs does not include the AS-path to reach a destination. In IDRP, rather than comparing candidate routes according to the length of the AS-path, the IDRAs use the TE information contained in the routing advertisements.

B. Aggregated Path State Information (PSI)

At a given OXC, PSI messages aggregate (i) intra-domain PSI; (ii) PSI related to the inter-domain links towards its downstream domains; and (iii) the already aggregated PSI contained in the inter-domain advertisements received from downstream domains. In OBGP+, the PSI is basically composed of aggregated wavelength availability information. In IDRP, the PSI is not only composed of aggregated wavelength availability information, but it also contains aggregated load information, which is represented by associating a cost with each candidate (path, wavelength) pair [10]. For notation purposes, we describe how the aggregated wavelength availability is computed.

The aggregated wavelength availability information is obtained by computing the *Effective Number of Available Wavelengths* (ENAW) for each type of wavelength, both inside an AS and across ASs. Inside an AS, the aggregation process is as follows. Let u and v be a pair of OXCs inside an AS, $P(u,v)$ be a candidate path between u and v, and l be a link within the path $P(u,v)$. The ENAW of wavelength type λ_i between the OXCs u and v within a domain is computed as follows:

$$W_{u,v}(\lambda_i) = \max_{P(u,v)} \left\{ \min_{l \in P(u,v)} [W_l(\lambda_i)] \right\} \quad (1)$$

The rationale behind eq. (1) is that the ENAW of a wavelength λ_i along a path P, which is basically the number of lightpaths that can possibly be setup on P using λ_i, is determined by the value of λ_i at the bottleneck link, i.e., the link with the minimum number of λ_i along P. Among all the paths between u and v, the path with the largest ENAW is chosen.

The inter-domain part is composed of the unused wavelengths on the directly-connected inter-domain links of the OXC, and wavelengths that are available downstream, which are known through the PSI advertisements from neighboring OXCs. Let $W_{l_b,l'_b}(\lambda_i)$ be the ENAW of type λ_i between OXC l_b and a local border OXC l'_b, $W_{l'_b,r_b}(\lambda_i)$ be the number of free wavelengths of type λ_i in the inter-domain link between the local border OXC l'_b and a remote border OXC r_b; and $W_{r_b,d}^{adv}(\lambda_i)$ be the ENAW of type λ_i between the remote border OXC r_b and the destination OXC d, which is advertised by r_b or the IDRA of r_b. By combining these inter-domain components and eq. (1), the OXC advertises to upstream neighbors the ENAW between the local border OXC l_b and the destination OXC d as:

$$W_{l_b,d}^{adv}(\lambda_i) = \min\left\{ W_{l_b,l'_b}(\lambda_i), W_{l'_b,r_b}(\lambda_i), W_{r_b,d}^{adv}(\lambda_i) \right\} \quad (2)$$

IV. Wavelength Aggregation in the Presence of Wavelength Converters

In this section, we present one of the main contributions of the paper, which is the extension of OBGP+ and IDRP to deal with the presence of wavelength converters. Having wavelength converters relaxes the wavelength continuity constraint, thereby increasing the "availability" of wavelengths. We show that with simple but necessary modifications, this information can be incorporated in the wavelength aggregation process. Our approach does not entail too much overhead since the only additional information is the number of wavelength converters at the remote border router.

We identify two types of unoccupied wavelength channels at any given border OXC: *converter* and *non-converter channels*. A converter channel consists of different types of wavelengths on either side of the OXC, thus requiring wavelength conversion if it is to be used for lightpath establishment. A non-converter channel, on the other hand, is made up of the same wavelength on both sides of the OXC and does not require wavelength conversion. In this section, unless explicitly specified, wavelengths/channels refer to *unoccupied* wavelengths/channels.

Since wavelength converters are scarce, it is assumed that they are used only when absolutely necessary. Therefore, we first compute the number and type of non-converter channels the same way as in the case where there are no converters. Then, the remaining wavelengths on either side of the OXC are candidates of converter channels. However, since a single wavelength converter can translate only one input wavelength to another output wavelength, the number of unused wavelength converters also affects the possible number of converter channels. Usually, there are more candidate wavelengths than the possible number of converter channels. Hence, there should be a mechanism to pick a specific wavelength for each converter channel (e.g., first-fit, random-fit, etc.) before being advertised upstream. This approach provides a highly aggregated state information, while capturing the availability of wavelength channels.

We now explain how the ENAW is computed using Fig. 1, which shows an example network with two ASs, their border OXCs and the unoccupied wavelengths at each OXC. For AS1, l_b and l_b' represent its border nodes, whereas r_b is the node that is directly connected to AS1. The downstream AS (in this case AS2) advertises a set of available wavelengths to the upstream
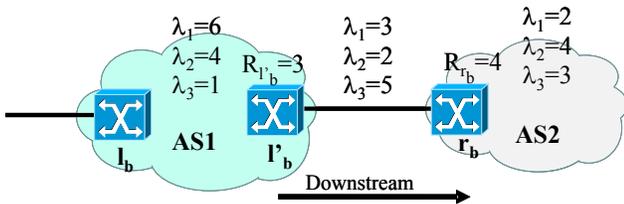

Fig. 1: An example depicting border OXCs of two domains connected by a single inter-domain link. This example shows the number of wavelengths and wavelength converters available at border OXCs of the two domains.

AS (in this case AS1). Let $W_{r_b,d}^{adv}(\lambda_i)$ be the advertised number of wavelengths of type λ_i from the downstream AS. Also, let $R^{adv} = R_{r_b}$ be the advertised number of available converters at r_b. $W_{l_b',r_b}(\lambda_i)$ is the number of wavelengths of type λ_i on the link between l_b' and r_b. This value is known to l_b' since the link is physically attached to it.

Thus, the number of non-converter channels of type λ_i at l_b' is:

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

In Fig. 1, $W_{l_b',d}(\lambda_1) = \min\{3,2\} = 2$, $W_{l_b',d}(\lambda_2) = \min\{2,4\} = 2$, and $W_{l_b',d}(\lambda_3) = \min\{5,3\} = 3$.

The remaining wavelengths can be part of converter channels at l_b'. The maximum number of possible converter channels is determined not only by the number of wavelengths that are *not* in the non-converter channels, but also by the number of available converters. Hence, it can be shown that the maximum number of converter channels is,

$$\min\left\{\left[\sum_i \left(W_{l_b',r_b}(\lambda_i) - W_{l_b',d}(\lambda_i)\right)\right], \quad (4)\right.$$
$$\left.\left[\sum_i \left(W_{r_b,d}^{adv}(\lambda_i) - W_{l_b',d}(\lambda_i)\right)\right], R^{adv}\right\}$$

In Fig. 1, the number of converter channels is: $\min\{\{(3-2)+(2-2)+(5-3)\}, \{(2-2)+(4-2)+(3-3)\}, 4\} = \min\{3,2,4\} = 2$.

For these converter channels, wavelengths are selected from the set $\{W_{l_b',r_b}(\lambda_i)\}\backslash\{W_{l_b',d}(\lambda_i)\}$, i.e., the set of wavelengths in $W_{l_b',r_b}(\lambda_i)$ that are not in the non-converter channels. Then, $W_{l_b',d}(\lambda_i)$ is updated so that it includes both the converter and non-converter channels before being advertised upstream. Let us assume that a random selection is used and the updated $W_{l_b',d}(\lambda_1) = 3$, $W_{l_b',d}(\lambda_2) = 2$, and $W_{l_b',d}(\lambda_3) = 4$.

Similarly, the number of non-converter channels of type λ_i at l_b is:

$$W_{l_b,d}(\lambda_i) = \min\left\{W_{l_b,l_b'}(\lambda_i), W_{l_b',d}(\lambda_i)\right\} \quad (5)$$

In Fig. 1, $W_{l_b,d}(\lambda_1) = \min\{6,3\} = 3$, $W_{l_b,d}(\lambda_2) = \min\{4,2\} = 2$, and $W_{l_b,d}(\lambda_3) = \min\{1,4\} = 1$.

The total number of converter channels at l_b is,

$$\min\left\{\left[\sum_i \left(W_{l_b,l_b'}(\lambda_i) - W_{l_b,d}(\lambda_i)\right)\right], \quad (6)\right.$$
$$\left.\left[\sum_i \left(W_{l_b',d}(\lambda_i) - W_{l_b,d}(\lambda_i)\right)\right], R_{l_b'}\right\}$$

where $R_{l_b'}$ is the number of converters at l_b'. In Fig. 1, this is equal to $\min\{\{(6-3)+(4-2)+(1-1)\}, \{(3-3)+(2-2)+(4-1)\}, 3\} = \min\{5,3,3\} = 3$. Let us assume that after randomly selecting from wavelengths that are not in the non-converter channels for the three converter channels, the updated $W_{l_b,d}^{adv}(\lambda_1) = 5$, $W_{l_b,d}^{adv}(\lambda_2) = 3$, and $W_{l_b,d}^{adv}(\lambda_3) = 1$.

Finally, AS1 advertises $W_{l_b,d}^{adv}(\lambda_i)$ and $R^{adv} = R_{l_b}$ to upstream domains. However, without the modified wavelength

aggregation process (see eq. (2)), AS1 would have instead advertised $W_{l_b,d}^{adv}(\lambda_1) = 2$, $W_{l_b,d}^{adv}(\lambda_2) = 2$, and $W_{l_b,d}^{adv}(\lambda_3) = 1$.

In [9] [10], it is proposed to piggyback *Keepalive* messages that are exchanged between neighboring OXCs with PSI messages. In this approach, keepalive messages are, just like in BGP, exchanged to notify if the neighboring node is still operative. However, unlike in BGP, the keepalive messages are extended to convey PSI messages. A major advantage of this strategy is that it does not increase the number of routing messages exchanged between domains. In this paper, we employ the same approach.

V. RESULTS AND DISCUSSION

In this section, we present simulation results that compare the performance of OBGP, OBGP+ and IDRP. Our performance metrics are the *Blocking Ratio* (BR) of inter-domain lightpath requests, and the number of routing messages exchanged to achieve this blocking ratio. To this end, we have conducted extensive simulations using OPNET. In our simulations, we have used a PAN-European topology, which was introduced in [2] as a reference topology suitable for a PAN-European fiber-optic network. The network consists of 28 domains and 41 inter-domain links, and the nodes were chosen in such a way that some of the main European Internet Exchange Points are included.

Inside each domain of the PAN European network, we placed a random number of OXCs, which is equal to or higher than the number of inter-domain links of that domain. There are 18 source and 10 destination OXCs randomly located covering the entire PAN European network in such a way that each domain has one source or destination OXC. In other words, we simulate inter-domain traffic which is transferred between domains. Each link in the network consists of 5 fibers and each fiber has 14 wavelengths.

In our simulation, traffic was modeled according to a Poisson distribution with exponentially distributed inter-arrivals. The blocking ratio and routing messages are collected under different traffic loads, varying from 100 up to 300 Erlangs. In order to evaluate the impact of the frequency of updates in the PSI messages, we have tested three scaled and normalized *Keepalive Update Interval* (K_T) of the Keepalive messages: $K_T = 1$, $K_T = 3$, and $K_T = 5$ units. In terms of the availability of converters, we have considered three scenarios: no converters, 5 converters and 10 converters at each border OXC of the domains in the network. For each case, the results are the averages of over 30 randomly generated PAN European network configurations. These network configurations are different from each other in the network topology inside each domain, and the location of source and destination OXCs over the entire network.

Due to space constraints, we are able to show only some of the results. Figs. 2 and 3 show the efficiency of using wavelength converters in OBGP+ and IDRP for $K_T = 1$. Similar results have been obtained for $K_T = 3$ and $K_T = 5$. Table I shows the improvement factor (**IF**) in the blocking ratios of OBGP+ and IDRP over OBGP and the number of

messages generated under traffic values 200, 250 and 300 Erlangs for 5 converters. Similar results have been obtained for 10 converters. The following observations can be made from our results.

- Increasing the update interval K_T causes more blocking because a higher value of K_T means that the PSI is not accurate enough since messages are exchanged less frequently. In fact, a major advantage of embedding PSI messages in Keepalive messages is that when K_T is decreased so as to improve the responsiveness of OXC neighbors, PSI messages will be updated more frequently.
- IDRP always significantly outperforms both OBGP+ and OBGP (whereas OBGP+ outperforms OBGP). This is due to the fact that IDRP additionally utilizes aggregated load information. In fact, for 10 converters and $K_T = 1$, IDRP achieves a blocking ratio of less than 0.1% for all simulated traffic values. The 0.1% blocking ratio is a threshold recommended by the IST FP6 NOBEL project [7] for optical networks in order to support real-time and streaming applications.
- The total number of messages generated decreases as more wavelength converters are used in the network. The reason for this is that in the presence of wavelength converters, the wavelength continuity constraint is relaxed and there will be more wavelengths available along a path. Therefore, it is less likely for the wavelengths of a path to be exhausted fast, thereby triggering reachability messages and path exploration.
- The blocking ratios for IDRP and OBGP+ decrease as more wavelength converters are placed in the network. But this is not the case in OBGP (results not shown here) if it always chooses the wavelength with the lowest identifier (First-Fit) along the shortest path. Such a first-fit approach increases conflicts as different OXCs tend to simultaneously choose lower identifier wavelengths, while higher identifier wavelengths are available. The situation is worsened as the number of converters in the network is increased, since the "availability" of these lowest indexed wavelengths is also increased, thereby exacerbating the possibility of conflicts. This situation can be avoided by choosing wavelengths randomly (Random-Fit) instead of always choosing lower indexed wavelengths.

VI. CONCLUSIONS

In this paper, we have made simple but important modifications to two inter-domain optical protocols, namely, OBGP+ and IDRP, to handle the presence of wavelength converters. We have also performed extensive simulations comparing the performance of OBGP (Optical BGP) and these protocols. The results obtained in a PAN European network show that IDRP significantly outperforms OBGP+ and OBGP, and OBGP+ outperforms OBGP. The performance metrics in the simulation were blocking ratio and the number of messages generated (for a duration of one week).

From these results, it can be inferred that the exchange of aggregated path state information (PSI), and the presence of

	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
IF (OBGP+)	3673.38	68.50	10.97	572.01	50.79	10.30	315.59	38.55	9.69
IF (IDRP)	13395.80	172.95	25.14	827.13	103.67	18.42	518.14	69.97	15.69
Traffic (Erlangs)	Routing Messages OBGP	Routing Messages OBGP+	Routing Messages IDRP	Routing Messages OBGP	Routing Messages OBGP+	Routing Messages IDRP	Routing Messages OBGP	Routing Messages OBGP+	Routing Messages IDRP
100	7,393,754	4,106,295	3,316,670	6,538,899	4,045,308	3,262,196	5,810,185	3,996,540	3,204,790
150	8,433,843	3,999,754	3,362,538	7,066,636	3,920,216	3,278,214	6,113,651	3,830,622	3,195,828
200	9,139,267	4,002,240	3,377,721	7,593,317	3,881,279	3,260,521	6,381,400	3,775,545	3,149,897
250	9,149,884	4,025,679	3,378,519	7,410,155	3,916,994	3,239,177	6,323,175	3,823,118	3,110,352
300	9,420,468	4,771,478	3,455,367	7,433,076	4,469,614	3,279,149	6,254,482	4,252,166	3,121,277

TABLE I: Improvement Factors (**IF**) in the blocking ratios of OBGP+ and IDRP over OBGP for 200, 250, and 300 Erlangs, and overall number of routing messages exchanged for 5 converters.
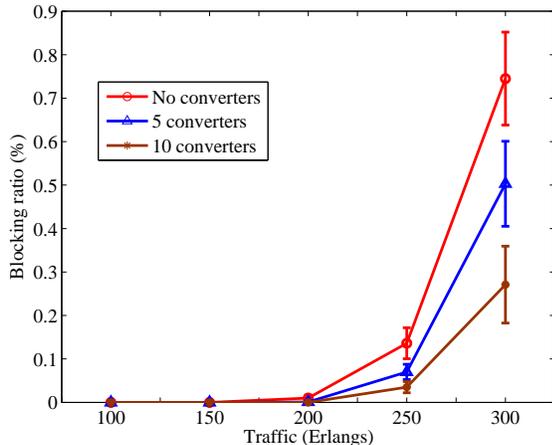

Fig. 2: Average blocking ratio and standard deviation. Comparison of different number of wavelength converters for OBGP+ ($K_T = 1$).
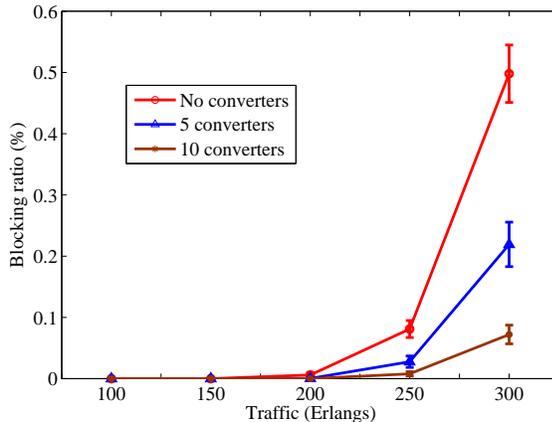

Fig. 3: Average blocking ratio and standard deviation. Comparison of different number of wavelength converters for IDRP ($K_T = 1$).

wavelength converters at border OXCs improve the blocking ratio and the number of messages generated significantly. In fact, using IDRP with enough wavelength converters, it is possible to achieve the 0.1% blocking ratio threshold that is recommended by the IST FP6 NOBEL project [7] for optical networks to support real-time and streaming applications. The decrease in the blocking ratio is obtained without an increase in the total number of messages exchanged, because we have employed a strategy of piggybacking PSI updates in the Keepalive messages exchanged between neighboring IDRAs/OBGP+ nodes.

A venue for future research can be to investigate the performance of the inter-domain routing protocols in larger inter-domain networks, and the optimization of the placement of wavelength converters in inter-domain networks.

REFERENCES

[1] M. Blanchet, F. Parent, and B. St-Arnaud, "Optical BGP (OBGP): IDRA lightpath provisioning," IETF draft, ietf-draft-parent-obgp-01, Mar. 2001.
[2] S. De Maesschalc, D. Colle, I. Lievens, M. Pickavet, P. Demeester, C. Mauz, M. Jaeger, R. Inkret, B. Mikac, and J. Derkacz, "Pan-European optical transport networks: an availability-based comparison," *Photonic Network Communications,* vol. 5, no. 3, pp. 203-225, May 2003.
[3] M. J. Francisco, L. Pezoulas, C. Huang, and I. Larnbadaris, "End-to-end signaling and routing for optical IP networks," *Proc. of IEEE ICC*, Apr. 2002.
[4] K.-C. Lee and V.O.K. Li, "A wavelength-convertible optical network," *Journal of Lightwave Technology*, vol. 11, no. 5/6, May/June 1993.
[5] G. Liu, C. Ji, and V. Chan, "On the scalability of network management information for inter-domain light-path assessment," *IEEE/ACM ToN*, vol. 13, no. 1, pp. 160-172, Feb. 2005.
[6] Q. Liu, M.A. Kök a, N. Ghani, and A. Gumaste, "Hierarchical routing in multi-domain optical networks," *Computer Communications*, vol. 30, no. 1, pp. 122-131, Dec. 2006.
[7] NOBEL: A Next Generation Optical Networks for Broadband European Leadership, "IST Integrated Project", http://www.ist-nobel.org.
[8] L. Wang, H. Zhang, and L. Zheng, "Reducing the OBGP protection switching time in WDM mesh networks," *Proc. of OFC*, Mar. 2006.
[9] M. Yannuzzi, X. Masip-Bruin, S. Sanchez-Lopez, and E. Marin-Tordera, "OBGP+: An improved path-vector protocol for multi-domain optical networks, " *Optical Switching and Networking, Elsevier*, vol. 6, issue 2, pp. 111-119, Apr. 2009.
[10] 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," *IEEE Communications Magazine*, vol. 46, no. 6, pp. 104-111, June 2008.