

Towards Efficient Optical Burst-Switched Networks without All-Optical Wavelength Converters

João Pedro^{1,2}, Paulo Monteiro^{1,3}, and João Pires²

¹ D1 Research, Siemens Networks S.A., R. Irmãos Siemens 1, 2720-093 Amadora, Portugal
{joao.pedro, paulo.monteiro}@siemens.com

² Instituto de Telecomunicações, Instituto Superior Técnico, Av. Rovisco Pais 1,
1049-001 Lisboa, Portugal
jpires@lx.it.pt

³ Instituto de Telecomunicações, Universidade de Aveiro, Campus Universitário de
Santiago, 3810-193 Aveiro, Portugal

Abstract. Optical Burst Switching (OBS) is a promising switching paradigm to efficiently support Internet Protocol (IP) packets over optical networks, under current and foreseeable limitations of optical technology. The prospects of OBS networks would greatly benefit, in terms of cost and ease of implementation, from limiting the wavelength conversion capabilities at the network nodes. This paper presents a framework for contention minimization/resolution combining traffic engineering in the wavelength domain to minimize contention in advance and optical buffering at the core nodes to resolve contention. Simulation results show that with the proposed contention minimization/resolution framework the large number of expensive all-optical wavelength converters used at the core nodes of an OBS network can be replaced by a moderate number of shared optical delay lines without compromising network performance.

Keywords: optical burst switching, traffic engineering, wavelength contention minimization, wavelength continuity constraint.

1 Introduction

The Optical Burst Switching paradigm [1] has attracted considerable interest as an optical networking architecture for efficiently supporting IP packet traffic, while exploiting the huge transmission capacity provided by optical fibres and Wavelength Division Multiplexing (WDM) technology [2]. OBS bandwidth utilization efficiency and technological requirements are in between those of coarse-grained Optical Circuit Switching (OCS) and fine-grained Optical Packet Switching (OPS). At the OBS edge nodes, multiple IP packets are assembled into bursts, which are the traffic units routed and switched inside the OBS core network. Thus, OBS provides sub-wavelength granularity, rendering higher bandwidth utilization efficiency than OCS in supporting IP traffic. Moreover, bandwidth for data burst transmission is reserved in advance and using out-of-band signalling, avoiding complex optical processing capabilities at the core nodes, which are mandatory in OPS networks.

OBS networks use one-way resource reservation mechanisms for setting up the resources for each burst transmission [1]. Therefore, the burst's ingress node starts transmitting soon after the burst has been assembled, instead of waiting for an acknowledgment of successful resource reservation in the entire burst path. Hence, two or more bursts may contend for the same resources at a core node. Given that unsolved contention leads to burst loss, degrading the network performance, it is clear that efficient contention resolution strategies are of paramount importance in these networks. Wavelength conversion was shown to be the most effective contention resolution strategy for OBS/OPS networks [3] and, as a result, most OBS proposals and studies assume the use of full-range wavelength converters at the network core nodes to resolve contention. However, because all-optical wavelength conversion devices are still undergoing research and development [4], they remain complex and expensive. Consequently, avoiding wavelength conversion would greatly reduce the complexity/cost of implementing OBS networks in the near future.

Recently, strategies for minimizing wavelength contention in advance have been investigated [5], [6], and even tested in an OBS network demonstrator [7], with the aim of reducing burst loss in the absence of wavelength conversion and optical buffering. Still, despite the reported performance improvements, an OBS network in these conditions only achieves low burst blocking probabilities at the expense of very small offered traffic loads [6], thus becoming a less attractive networking solution.

In this paper, we propose a novel framework for contention minimization and resolution in OBS networks without all-optical wavelength converters at the network nodes. Simulation results show that the proposed framework enables the design of efficient OBS networks employing a moderate number of shared Fiber Delay Line (FDL) buffers instead of a large number of expensive wavelength converters.

The remainder of the paper is organized as follows. Section 2 outlines the OBS contention resolution strategies and corresponding node architectures. A contention minimization/resolution framework tailored for OBS networks without wavelength converters is described in section 3, whereas section 4 assesses the performance of this framework and compares its network resource requirements with those of a network with wavelength conversion, for the same objective performance. Finally, section 5 presents some concluding remarks.

2 OBS Contention Resolution and Node Architecture

The performance of OBS networks is mainly hampered by contention at the network nodes, which arises whenever two or more data bursts, overlapping in time, are directed to the same wavelength on the same output fibre. Contention can be resolved using strategies acting in one or several of three domains: wavelength, time, and space. Thus, contending bursts can be converted to other wavelengths available at the output fibre, delayed using optical buffers, or deflected to other output fibres/links of the node. Wavelength conversion requires all-optical wavelength converters, whereas optical buffering demands FDLs, which unlike Random-Access Memory (RAM) only provide fixed delays. Deflection routing does not require additional hardware, but its effectiveness is heavily dependent on network topology and traffic pattern.

The complexity of the OBS network nodes depends on the strategies employed for contention resolution in the wavelength and time domains. Fig. 1 (a) shows the key blocks of the simplest node architecture, comprising a fast space switch matrix with a number of transit ports given by two times MW , where M is the number of input and output fibres and W is the number of wavelengths per fibre. Fig. 1 (b) depicts the most common OBS node architecture, which includes a set of MW dedicated wavelength converters, one per output transit port. It has been shown in [8] that significant savings in the number of converters can be achieved, without noticeable performance degradation, by sharing the converters among all output transit ports, as illustrated in Fig. 1 (c). However, this is realized at the expense of increasing the size of the space switch matrix by adding $2C$ transit ports, where C is the number of shared converters. Fig. 1 (d) depicts a node using F shared FDL buffers for resolving contention in the time domain, which also requires a larger space switch matrix. Node architectures combining wavelength converters with FDL buffers can also be devised [3], [8].

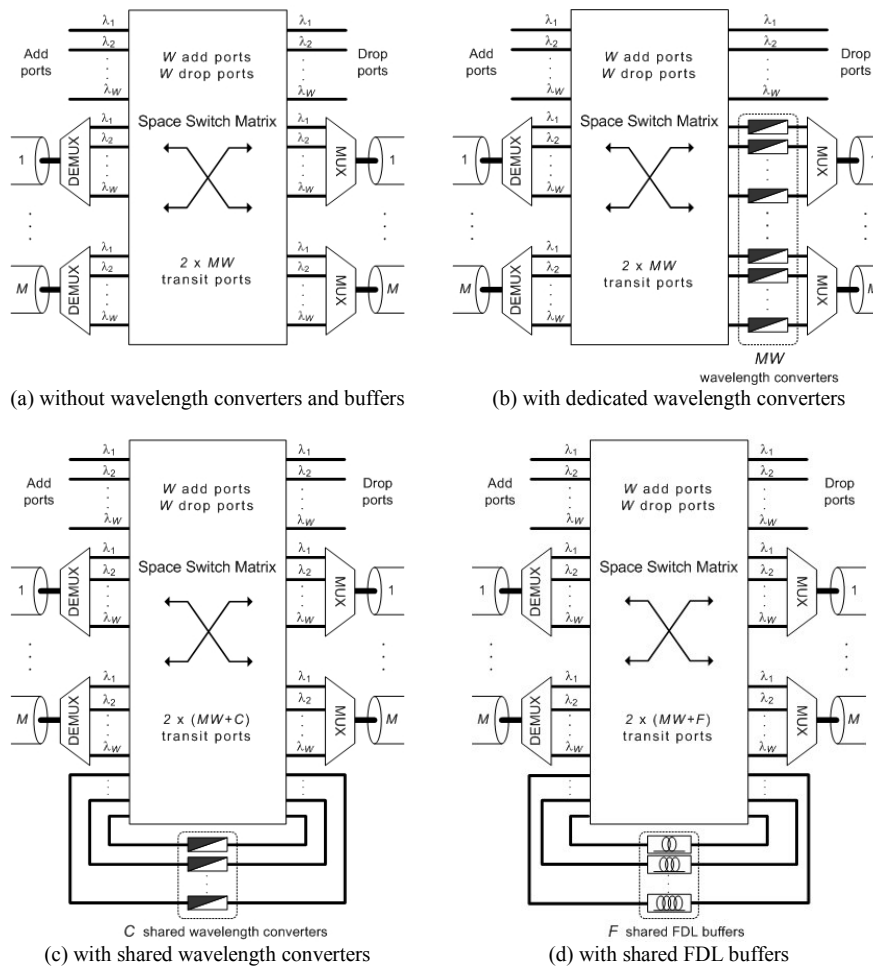


Fig. 1. OBS network node architectures.

In view of the above simplified node architectures, the complexity and cost of the OBS network depends of: (i) the switch matrices size, (ii) the number of wavelength converters, and (iii) the number of FDL buffers. Moreover, since a FDL buffer is basically an optical fibre with a length designed to provide a specific delay, it is simpler than an all-optical wavelength converter. This fact motivated us to investigate the feasibility of an OBS network using only shared FDL buffers to resolve contention. In the following, we describe a framework for contention minimization and resolution tailored for networks based on this architecture and show evidence that, using a moderate number of FDLs, it enables the same bandwidth utilization efficiency of an OBS network using dedicated full-range wavelength converters.

3 Contention Minimization and Resolution Framework for OBS Networks under the Wavelength Continuity Constraint

In all-optical networks, the absence of wavelength conversion at the network nodes imposes the wavelength continuity constraint on the data path, that is, the same wavelength must be used in all links of the path, from the ingress node to the egress node. The resulting wavelength assignment problem has been extensively studied in the context of OCS networks [9], where wavelength availability in the entire path is known and exploited by the wavelength assignment strategy. However, due to the one-way nature of resource reservation in OBS networks, the ingress node knows the wavelengths availability on its output links, but it is not aware of their availability on the downstream links of the burst path. Thus, bursts going through overlapping paths can be assigned the same wavelength by their ingress nodes, resulting in contention for that wavelength at some common link. The probability of wavelength contention can be minimized in advance by exploiting the following principle [5]: if two or more burst paths share one or more network links, contention on those links will be reduced if each burst path preferably uses wavelengths different from those preferred by the other overlapping burst paths. In practice, this is achieved by maintaining at the burst path's ingress node an optimized priority-based ordering of the wavelengths and using it to search for an available wavelength (on its output fibre link) for transmitting data bursts towards the egress node.

The work in [5] introduced a strategy for optimizing the wavelength orderings using information of network and traffic conditions that usually remain unchanged over relatively long time scales, such as network topology, routing paths, and average offered traffic load between nodes. Recently, we proposed a new strategy [6] that uses the same input information and was shown to be significantly more effective in minimizing wavelength contention than that of [5]. In the following, we describe a strategy simpler than that of [6], but that was found to attain similar performance.

Consider an OBS network with N nodes, L unidirectional links, and W wavelengths per link. Let Π denote the set of paths used to transmit bursts, and let E_i denote the set of links traversed by path $\pi_i \in \Pi$. Let also γ_i denote the average traffic load offered to burst path π_i . The extent of wavelength contention from bursts of a path π_i on bursts of a path π_j is expected to increase with both the average traffic load offered to the paths and the number of common links. Define the interference level of π_i on π_j as

$$I(\pi_i, \pi_j) = \gamma_i |E_i \cap E_j|, \quad i \neq j, \quad (1)$$

where $|E_i \cap E_j|$ denotes the number of links shared by both paths.

Let $1 \leq P(\pi_i, \lambda_j) \leq W$ denote the priority of wavelength λ_j on path π_i , that is, the ingress node of π_i only assigns λ_j to a burst on this path if the first $W - P(\pi_i, \lambda_j)$ wavelengths, ordered by decreasing priority, are not available. Based on the problem inputs, the strategy determines the interference level of each path on every other path and reorders the paths of Π such that if $i < j$ the following condition holds

$$\sum_{\pi_k \in \Pi} I(\pi_i, \pi_k) > \sum_{\pi_k \in \Pi} I(\pi_j, \pi_k) \text{ or } \sum_{\pi_k \in \Pi} I(\pi_i, \pi_k) = \sum_{\pi_k \in \Pi} I(\pi_j, \pi_k) \text{ and } |E_i| > |E_j|. \quad (2)$$

Initially, no priorities are assigned to the wavelengths, that is, $P(\pi_i, \lambda_j) = 0$ for all π_i and λ_j . The following steps are executed for every priority $1 \leq p \leq W$ in decreasing order, and for every ordered path $\pi_i \in \Pi$:

(S1) Let $\Lambda = \{\lambda_j : P(\pi_i, \lambda_j) = 0, 1 \leq j \leq W\}$ denote the initial set of candidate wavelengths containing all wavelengths that were not assigned a priority on π_i . If $|\Lambda| = 1$ go to (S7).

(S2) Let $P_\Lambda = \{p_k : \exists \pi_l, l \neq i, P(\pi_l, \lambda_j) = p_k, |E_l \cap E_i| > 0, \lambda_j \in \Lambda\}$ denote the set of priorities already assigned to candidate wavelengths on paths that overlap with π_i .

(S3) Let $\rho = \min_{\lambda_j \in \Lambda} \{\max\{P(\pi_l, \lambda_j) : l \neq i, |E_l \cap E_i| > 0\}\}$ be the lowest priority from the set of the highest priorities assigned to candidate wavelengths on the paths that use links of π_i . Update the set of candidate wavelengths as follows

$$\Lambda \leftarrow \Lambda \setminus \{\lambda_j : \exists \pi_l, l \neq i, P(\pi_l, \lambda_j) > \rho, |E_l \cap E_i| > 0, \lambda_j \in \Lambda\}. \quad (3)$$

If $|\Lambda| = 1$ go to (S7).

(S4) Let $C(e_m, \lambda_j) = \sum \{\gamma_l : E_l \supset e_m, |E_l \cap E_i| > 0, P(\pi_l, \lambda_j) = \rho\}$ be the cost associated with wavelength λ_j on link $e_m \in E_i$. Thus, the minimum cost among the highest costs associated with the candidate wavelengths on the links of path π_i is $\alpha_e = \min_{\lambda_j \in \Lambda} \{\max\{C(e_m, \lambda_j) : e_m \in E_i\}\}$. Update the set of candidate wavelengths as follows

$$\Lambda \leftarrow \Lambda \setminus \{\lambda_j : \exists e_m, C(e_m, \lambda_j) > \alpha_e, e_m \in E_i, \lambda_j \in \Lambda\}. \quad (4)$$

If $|\Lambda| = 1$ go to (S7).

(S5) Let $C(\pi_i, \lambda_j) = \sum_{e_m \in E_i} C(e_m, \lambda_j)$ be the cost associated with wavelength λ_j on π_i . Thus, $\alpha_\pi = \min_{\lambda_j \in \Lambda} C(\pi_i, \lambda_j)$ is the minimum cost among the costs associated with the candidate wavelengths on π_i . Update the set of candidate wavelengths as follows

$$\Lambda \leftarrow \Lambda \setminus \{\lambda_j : C(\pi_i, \lambda_j) > \alpha_\pi, \lambda_j \in \Lambda\}. \quad (5)$$

If $|\Lambda| = 1$ go to (S7).

(S6) Update the set of priorities assigned to the candidate wavelengths as follows

$$P_\Lambda \leftarrow P_\Lambda \setminus \{p_k : p_k \geq p, p_k \in P_\Lambda\}. \quad (6)$$

If $|P_\Lambda| > 0$ go to (S3). Else, randomly select a candidate wavelength $\lambda_j \in \Lambda$.

(S7) Assign priority p to the candidate wavelength $\lambda_j \in \Lambda$ on path π_i , that is

$$P(\pi_i, \lambda_j) = p. \quad (7)$$

Each execution of the above steps, selects the wavelength that will be used with priority p on path π_i . Therefore, these steps are executed $W \cdot |\Pi|$ times to determine all wavelength orderings. The highest priority wavelengths are the first to be selected and preference in wavelength selection is given to the burst paths with larger interference level over the other paths. The wavelength selection is made by first defining a set of candidate wavelengths, and then successively reducing their number until there is only one. The criterions used to reduce this set size are such that the candidate wavelength selected minimizes the *priority interference* on the links of the path. In other words, the wavelength selected is one that has preferably been assigned low priorities on paths that share links with the burst path π_i . In the following, this greedy strategy is named Heuristic Minimum Priority Interference (HMPI).

Upon assembling a data burst, directed to a given egress node, the ingress node will search for an available wavelength using the wavelength ordering optimized for the corresponding path. Since a burst transmitted over a high priority wavelength will have greater chances of avoiding wavelength contention, we use a proper scheduling algorithm at the ingress node to increase the amount of traffic carried on these wavelengths. The ingress burst scheduling algorithm is allowed to delay assembled bursts at the electronic buffer of the ingress node by any amount of time bounded to D_{ingress} . Hence, the algorithm will first try to assign the highest priority wavelength to the burst if this wavelength is available during the required time for burst transmission starting at some time in the future not exceeding the additional D_{ingress} delay. If this wavelength is still not available, the algorithm will proceed in the same manner for the second highest priority wavelength and so on, until either an available wavelength is found or all wavelengths have been searched.

In view of the wavelength continuity constraint, contention at a transit node can only be resolved by delaying contending bursts at the shared FDL buffers. Several FDL buffer configurations have been presented [3]. Here, we adopt a pragmatic approach to optical buffering. Thus, only single-wavelength FDL buffers are used, which cannot accommodate bursts overlapping in time using different wavelengths, but avoid a larger switch fabric and extra multiplexing and demultiplexing equipment required by multi-wavelength buffers. Moreover, since a larger delay demands either a longer fibre or buffer recirculation, increasing optical signal degradation and the amount of extra hardware for compensating it, we impose a lower bound D_{min} and an upper bound D_{max} on the FDL delay, irrespective of the number of buffers per node, and inhibit buffer recirculation. Hence, in a node with F shared FDL buffers, the FDLs delay is uniformly distributed between D_{min} and D_{max} . Contention is resolved by delaying the contending burst using the smallest available FDL buffer that enables the burst to be transmitted through the next link of its path using the same wavelength.

4 Results and Discussion

The performance of OBS networks using dedicated wavelength converters, shared wavelength converters, and shared FDL buffers, is evaluated in this section using network simulation [10]. A 10-node ring topology with a uniform traffic pattern is used in the performance studies. Moreover, the network employs Just Enough Time (JET) resource reservation [1] and bursts are routed through the shortest paths. In all cases, each of the $L=20$ unidirectional links has $W=32$ wavelengths, a wavelength capacity of $B=10$ Gb/s, a switch fabric configuration time of $10 \mu\text{s}$, and an average burst size of 100 kB, which gives an average burst duration of $T=80 \mu\text{s}$. A negative exponential distribution is used for both burst size and burst interarrival time. The average offered traffic load normalized to the network capacity is given by

$$\Gamma = \frac{\sum_{\pi_i \in \Pi} \gamma_i \cdot |E_i|}{L \cdot W \cdot B}. \quad (8)$$

The OBS ring network using shared wavelength converters (SWC) or shared FDL buffers (SFB) is designed to match the bandwidth utilization efficiency of using dedicated wavelength converters (DWC). This is achieved by determining the number of shared converters or FDL buffers required to attain the same average burst blocking probability at the expense of the same average offered traffic load than that of using dedicated converters. In this case, we set the objective average burst blocking probability to 10^{-3} and 10^{-4} , which is obtained with dedicated wavelength converters for an average offered traffic load of $\Gamma=0.40$ and $\Gamma=0.47$, respectively.

Fig. 2 and Fig. 3 plot the average burst blocking probability of the 10-node ring network as a function of the number of shared wavelength converters and the number of shared FDL buffers, respectively. In the latter case, the minimum and maximum FDL delays are set to $D_{\min}=T$ and $D_{\max}=10T$. The average burst blocking probabilities were obtained by simulating 20 independent burst traces for $\Gamma=0.40$ and $\Gamma=0.47$.

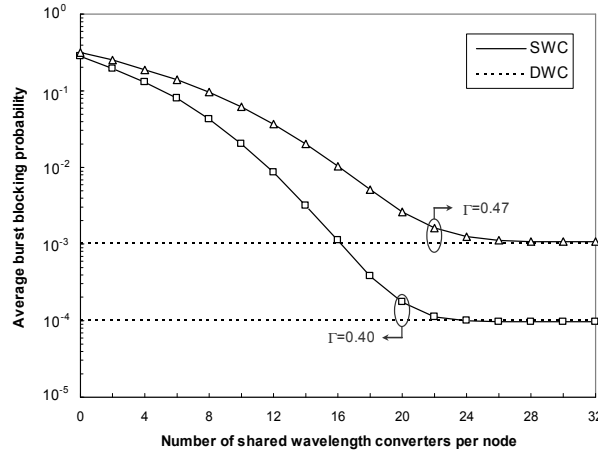


Fig. 2. Network performance using shared wavelength converters.

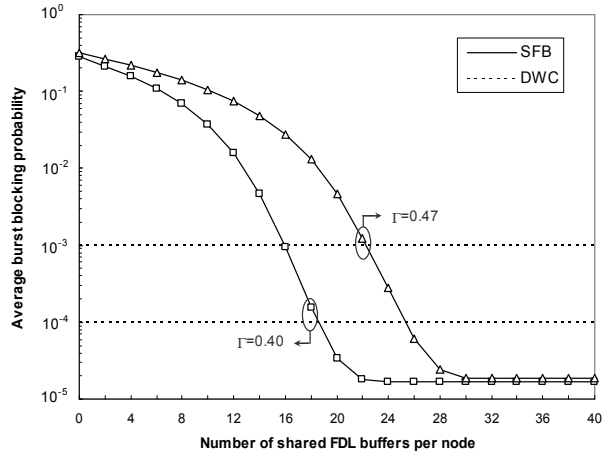


Fig. 3. Network performance using shared FDL buffers.

The curves in these plots show that as the number of shared wavelength converters or shared FDL buffers is increased, the average burst blocking probability is reduced and eventually the bandwidth utilization efficiency of using dedicated wavelength converters is achieved. The results also suggest that larger numbers of either shared converters or FDL buffers are needed when the objective burst blocking probability is higher. Moreover, in both cases the curves tend to a lower boundary of burst loss performance. In the SWC node architecture, this boundary corresponds exactly to the performance of DWC, whereas in the SFB node architecture the lower boundary is a result of imposing a maximum FDL delay irrespective of the number of FDL buffers.

Fig. 4 plots the performance of the 10-node ring network using shared FDL buffers and the HMPI strategy to optimize the wavelength orderings of each burst path. The minimum and maximum FDL delays are also set to $D_{\min}=T$ and $D_{\max}=10T$, whereas the maximum ingress delay is given by $D_{\text{ingress}}=T$.




Fig. 4. Network performance using shared FDL buffers and the HMPI strategy.

These curves exhibit a decrease of several orders of magnitude in terms of average burst blocking probability, as compared to the ones of Fig. 3. More to the point, they show a remarkable decrease on the number of shared FDL buffers required to match the bandwidth utilization efficiency of using dedicated wavelength converters. For instance, for an objective average burst blocking probability of 10^{-4} the SFB node architecture needs around 19 shared FDL buffers, whereas with the use of wavelength orderings optimized with the HMPI strategy, this number is reduced to only 5. Thus, buffer requirements are reduced by almost a factor of 4. The reported FDL buffer savings are due to the effectiveness of the priority-based wavelength assignment and burst scheduling strategy in minimizing contention in the wavelength domain.

In order to gain further insight on the complexity of each of the node architectures under study, Fig. 5 presents the number of transit switching ports, the number of wavelength converters, and the number of FDL buffers per node that is needed to achieve an average burst blocking probability of 10^{-4} when the average offered traffic load is 0.40.




Fig. 5. Node resource requirements for an average burst blocking probability of 10^{-4} and an average offered traffic load of 0.40.

The node resource requirements plotted in Fig. 5 show that the use of dedicated wavelength converters requires a very large number of converters, although with the benefit of avoiding an increase on the number of transit switching ports. On the other hand, sharing the wavelength converters allows reducing their number to less than half, but demands a significant increase on the number of transit switching ports. The node architecture with shared FDL buffers needs slightly less transit switching ports and successfully replaces the expensive wavelength converters by simpler FDL buffers. However, the best compromise between the expansion of the space switch matrix and the number of either wavelength converters or FDL buffers is clearly given by using the priority-based wavelength assignment and burst scheduling strategy for minimizing wavelength contention combined with shared FDL buffers to resolve contention in the time domain.

5 Conclusions

In view of the current limitations of all-optical wavelength conversion technology, OBS networks would greatly benefit, in terms of cost and ease of implementation, from avoiding the use of wavelength converters at the network nodes. This paper has presented a framework from contention minimization/resolution combining traffic engineering in the wavelength domain to minimize contention in advance and optical buffering at the core nodes to resolve contention. Moreover, a performance study with node architectures that use either dedicated wavelength converters, shared wavelength converters, or shared FDL buffers to resolve contention was presented. Simulation results show that the proposed contention minimization/resolution framework only requires a moderate number of shared FDL buffers and a small increase in the space switch matrix to achieve the same bandwidth utilization efficiency of a network using dedicated full-range wavelength converters. Therefore, it can contribute to lower the complexity/cost of deploying OBS networks in the near future.

Acknowledgments. The authors acknowledge the financial support from Siemens Networks S.A. and Fundação para a Ciência e a Tecnologia (FCT), Portugal, through research grant SFRH/BDE/15584/2006.

References

1. C. Qiao and M. Yoo, "Optical Burst Switching (OBS) – A new paradigm for an optical Internet", *Journal of High Speed Networks*, vol. 8, no. 1, pp. 69-84, January 1999.
2. R. Ramaswami and K. Sivarajan, *Optical Networks: A Practical Perspective*, 2nd edition. San Francisco, CA: Morgan Kaufmann, 2002.
3. S. Yao, B. Mukherjee, S. Yoo and S. Dixit, "A unified study of contention-resolution schemes in optical packet-switched networks", *IEEE Journal of Lightwave Technology*, vol. 21, no. 3, pp. 672-683, March 2003.
4. A. Poustie, "Semiconductor devices for all-optical signal processing", in *Proc. of ECOC 2005*, paper We3.5.1, Glasgow, Scotland.
5. J. Teng and G. Rouskas, "Wavelength selection in OBS networks using traffic engineering and priority-based concepts", *IEEE Journal of Selected Areas in Communications*, vol. 23, no. 8, pp. 1658-1669, August 2005.
6. J. Pedro, P. Monteiro, and J. Pires, "Wavelength contention minimization strategies for optical burst-switched networks", in *Proc. of IEEE GLOBECOM 2006*, paper OPNp1-5, San Francisco, USA.
7. Y. Sun, T. Hashiguchi, V. Minh, X. Wang, H. Morikawa, and T. Aoyama, "Design and implementation of an optical burst-switched network testbed", *IEEE Communications Magazine*, vol. 43, no. 11, pp. s48-s55, November 2005.
8. C. Gauger, "Performance of converter pools for contention resolution in optical burst switching", in *Proc. of SPIE OptiComm 2002*, pp. 109-117, Boston, USA.
9. H. Zang, J. Jue, and B. Mukherjee, "A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks", *Optical Networks Magazine*, vol 1, no. 1, pp. 47-60, January 2000.
10. J. Pedro, J. Castro, P. Monteiro, and J. Pires, "On the modelling and performance evaluation of optical burst-switched networks", in *Proc. of IEEE CAMAD 2006*, pp. 30-37, Trento, Italy.