

Dimensioning an OBS switch with Partial Wavelength Conversion and Fiber Delay Lines via a Mean Field Model

Juan F. Pérez and Benny Van Houdt

Performance Analysis of Telecommunication Systems Research Group
Department of Mathematics and Computer Science
University of Antwerp - IBBT
Middelheimlaan 1, B-2020 Antwerp, Belgium
Email: {juanfernando.perez, benny.vanhoudt}@ua.ac.be

Abstract—In this paper we introduce a mean field model to analyze an optical switch equipped with both wavelength converters (WCs) and fiber delay lines (FDLs) to resolve contention in OBS networks. Under some very general conditions, that is, a general burst size distribution and any Markovian burst arrival process at each wavelength, this model determines the minimum number of WCs required to achieve a zero loss rate as the number of wavelengths becomes large. The mean field result is exact as the number of wavelengths goes to infinity and turns out to be very accurate for systems with (a few) hundred wavelengths, commonly occurring when using wavelength division multiplexing (WDM). Moreover, we show that if the number of WCs is underdimensioned, (i) periodic system behavior may occur (with the period being the greatest common divisor of the burst lengths) and (ii) increasing the number of WCs may even worsen the loss rate under the often studied minimum horizon allocation policy (as opposed to the minimum gap policy). Finally, we further demonstrate that in terms of the loss rate, including (more) FDLs may have little or no effect on the number of WCs required to achieve a near-zero loss, especially for higher loads.

I. INTRODUCTION

Optical burst switching (OBS) has been proposed as a solution to minimize the opto-electronic translations at the backbone network switches [1], [2]. As only the burst header requires this translation, the main part of the signal can be processed in the optical domain. In consequence, OBS enables the switches to catch up with the growing transmission capacity of the optical fibers, driven by wavelength division multiplexing (WDM). With WDM several signals can be sent at the same time using different wavelengths, increasing the fiber capacity by tens or hundreds. As the main part of the signal is processed in the optical domain, contention can be resolved using wavelength conversion or optical buffering. A wavelength converter allows an incoming burst to use a different wavelength for transmission if the one it used to enter the switch is unavailable. On the other hand, optical buffering is implemented using Fiber Delay Lines (FDLs) that allow an incoming packet to be delayed for a specific amount of time proportional to the length of each fiber.

a) Our Contribution: In this work we introduce a mean field model of an optical switch equipped with a pool of full-range wavelength converters and a set of FDLs per output port. The mean field model is exact when the number of wavelengths tends to infinity, while it is shown to be very accurate when compared to a finite system with a large number of wavelengths. This case is particularly relevant as WDM technology has increased the number of signals that a single fiber can carry to more than a hundred. Our model allows a general burst size distribution while the burst arrival process at each wavelength is modeled as a Markovian arrival process (MAP) [3]. This process is able to represent general correlated inter-arrival times. In order to select a wavelength for a specific incoming burst we consider two different allocation policies, the minimum horizon and minimum gap policies, as explained in Section II.

As the loss rate in an Erlang loss model decreases to zero as the number of servers becomes large, it is clear that a near-zero loss rate can be realized for WDM links with hundreds of wavelengths provided that there are plenty of wavelength converters (WCs) available. On the other hand, as switches with a high number of WCs are not very cost effective, limiting their numbers is important. Therefore, in switches with partial wavelength conversion, one typically has only $C = \sigma W$ converters, with $\sigma \in (0, 1)$ and W the number of wavelengths. Some important questions that arise are: (i) how to determine σ to achieve a near-zero loss and (ii) how is σ affected by the presence of FDLs. The main *new* insights gathered from the mean field model can be summarized as follows:

- 1) The mean field model allows us to determine σ in the general setting defined above (using only a single run).
- 2) If the number of WCs is underdimensioned, meaning σ is selected too small, periodic system behavior may occur, which is a very unwanted effect in any system. The period seems to equal the greatest common divisor of the burst lengths.
- 3) Moreover, if the number of WCs is too small, increasing the number of WCs may even worsen the loss rate under

the minimum horizon policy (which aims at minimizing the burst delays). This is not the case when the minimum gap policy is used.

- 4) The number of FDLs in the system may have little or no effect on the required σ , meaning if the number of WCs is sufficiently large, there might be no use in incorporating FDLs (as far as the loss rate is concerned).
- 5) Even if the number of WCs is insufficient, increasing the number of FDLs may not improve the loss rate. Moreover, higher loads tend to decrease the usefulness of incorporating FDL buffers.

To the best of our knowledge, each of these conclusions is novel and of significant importance when designing optical switches with partial wavelength conversion and FDLs.

b) Related work: In previous studies analytical models have been used to evaluate the effect of wavelength conversion on a bufferless switch [4], [5], and to examine the performance of a switch equipped with FDLs but without converters [6]–[8]. The interaction of both wavelength conversion and FDLs has been analyzed by means of simulation models only in [9]–[11]. In these studies, as well as in the present paper, the converters are assumed to have full-range conversion, i.e., a burst can be converted to any wavelength. The case where the bursts can only be converted to a restricted set of wavelengths has been treated in [12]–[15].

This paper is organized as follows: in Section II we present the switch under analysis and the wavelength allocation policies; Section III gives a general description of the mean field model, while Section IV compares the results of the model with results from the simulation of a finite system. This section also analyzes the effect of the allocation policies and of various parameters on the performance of the switch.

II. THE OPTICAL SWITCH

The optical switch under analysis, shown in Figure 1, is made of a number of input/output ports, each one connected to a fiber with W wavelengths. The switch works in a synchronous manner, where the time is divided in equally-spaced slots and the state of the switch is observed at slot boundaries. The arrival process at each wavelength is modeled as a MAP [3] characterized by the set of $m \times m$ matrices $\{B_0, B_1, \dots, B_{L_{\max}}\}$, where L_{\max} is the maximum packet length (the terms packet and burst are used interchangeably). The MAP is driven by an underlying Markov chain with transition matrix $B = \sum_{k=0}^{L_{\max}} B_k$. For $k \geq 1$, the entry (i, j) of the matrix B_k is the probability that a packet of size k arrives and the underlying Markov chain makes a transition from i to j . Correspondingly, B_0 contains the transition probabilities of the underlying chain involving no arrivals. The class of MAP processes has been previously used to model the arrival process at a bufferless optical switch [4]. When a burst arrives it is switched to the corresponding output port using its own wavelength, called home wavelength. If the home wavelength is available for transmission in the output port, the burst starts transmission immediately. If the wavelength is already transmitting another burst or has scheduled the transmission

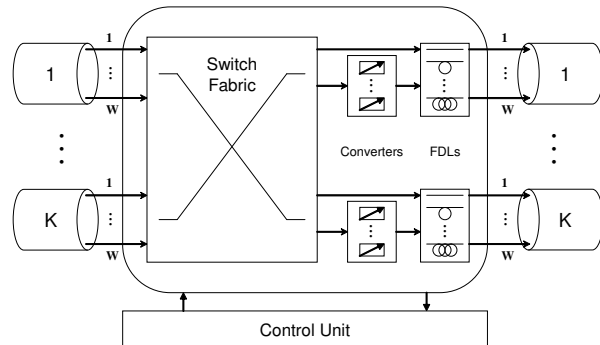


Figure 1. Optical switch with K input/output ports, W wavelengths, converters and FDLs

of a burst waiting in the FDL, the new packet is buffered using the FDL. In case the FDL has no available buffering capacity in that wavelength, the incoming burst is converted to a different wavelength using one of the available converters. If there are no idle converters or no wavelengths with available buffering capacity, the burst must be dropped. Thus, to resolve contention the switch first tries to buffer the signal and only if this is not possible it tries to convert it, aiming to minimize the converter usage, as the *minConv* strategy in [11].

To analyze the performance of the switch we can focus on a single output port as the incoming traffic is assumed to be uniformly distributed among the output ports. To describe the state of one of these ports we consider two types of objects: wavelengths and converters. The state of a single wavelength is described by the scheduling horizon, which is the time until all the packets already scheduled for transmission in that wavelength have left the switch. If the horizon is equal to 0 and a packet of size L arrives, it can start transmission immediately and the horizon increases to L . On the other hand, if the incoming burst finds a horizon equal to h , it will experience a delay of at least h units before actual transmission. As the buffering is carried out by a set of N FDLs, the possible delay a packet can experience depends on the length of these delay lines. Here we assume the N FDLs have linearly growing length with granularity D , i.e., the first line provides a delay of D time slots, the delay in the second is equal to $2D$, and the last line delays the packet for ND slots. With this setup an incoming packet that observes a scheduling horizon equal to h has to wait for $D \lceil \frac{h}{D} \rceil$ slots, if $h \leq ND$. If the packet is of size L the new value of the horizon is $D \lceil \frac{h}{D} \rceil + L$. Notice, in this particular case the wavelength remains unused for a length of $D \lceil \frac{h}{D} \rceil - h$ just prior to the packet transmission, we refer to this as a *gap*. If h is greater than ND the packet cannot be buffered in the FDL using the same wavelength and it must be reallocated to another wavelength with horizon less than or equal to ND .

A packet that cannot be buffered in its home wavelength, called an *extra-packet*, can be reallocated if there is both a wavelength with scheduling horizon no greater than ND and an available converter. Hence, it is necessary to check

the state of all the wavelengths and the converters. There are C converters per output port and the state of a single converter is also described by its scheduling horizon. In this case the converter has no buffering capacity, therefore its horizon reduces to the time required by the packet already in service to be completely converted to the other wavelength. Then, if an extra-packet of size L finds an available converter (and there is a wavelength with available buffering capacity) the horizon of the selected converter changes its value from 0 to L . Naturally, when this conversion occurs the horizon of the wavelength that receives the burst increases its value as described previously. An important assumption is that each wavelength with available buffering capacity can only receive one extra-packet during one slot, even if it has enough free FDLs to receive more than one additional packet. Removing this assumption would complicate both the possible set of wavelength allocation policies and its corresponding modeling aspects. The number of converters C per output port is determined as a fraction of the number of wavelengths W , i.e., $C = \sigma W$, where σ is the conversion ratio. If $\sigma = 0$ (resp. $\sigma = 1$) the switch is said to have null (resp. full) conversion. Here we assume that σ takes values between 0 and 1, which is called partial conversion. If an extra-packet finds an idle converter it has to choose a wavelength among those with horizon less than or equal to ND . This selection can be made using two different allocation policies: *minimum horizon* (*minH*), which selects the wavelength with the minimum scheduling horizon; and *minimum gap* (*minG*), which selects the wavelength with a horizon such that the allocation of a new packet generates a gap of minimum value.

III. THE MEAN FIELD MODEL

To model the evolution of the switch in a single slot we consider the following order of events: first, the busy wavelengths (resp. converters) transmit (resp. translate) part of the packet in service, reducing their horizons by one. Second, a new packet may arrive at each wavelength with a probability related to the phase of its arrival process; the packet is buffered if there is space available in its home wavelength, otherwise it becomes part of the set of extra-packets. Third, the extra-packets are converted to a different wavelength with available buffering capacity. Any extra-packet that does not find an available converter or a wavelength with buffering capacity must be dropped. The probability that a packet is dropped is called *loss probability* and is considered the main measure of performance.

Our model is based on a general result for a system of interacting objects introduced in [16]. In this case, the system consists of two types of objects: wavelengths and converters. Their evolution during a time slot is described by the matrices associated to each of the three steps: transmission (S_k), arrivals (A_k) and reallocation (Q_k). The subscript k may be equal to w or c depending on whether the matrix describes the transition of a wavelength or a converter. A thorough description of these matrices can be found in [17]. Even though S_k and A_k can be defined independently for

each wavelength or converter, the reallocation matrices Q_k clearly depend on the state of the whole system. To consider this we observe the system just before the reallocation step. The state of an individual wavelength includes its horizon, the phase of the arrival process, and the size of the extra-packet (if there is one). This information can be compactly held in a vector $M^{W,(w)}(t)$, whose components store the proportion of wavelengths in each state in slot t . Analogously, the vector $M^{W,(c)}(t)$ can be defined to hold the proportion of converters in each possible state (scheduling horizon) before reallocation in slot t . Therefore, the state of the complete system at time t can be described by the vector

$$M^W(t) = \left[\frac{1}{1+\sigma} M^{W,(w)}(t), \frac{\sigma}{1+\sigma} M^{W,(c)}(t) \right],$$

which is called the occupancy vector and contains the fraction of objects in each state, including both wavelengths and converters. Based on this vector, we can define the transition matrices $Q_w(M^W(t))$ and $Q_c(M^W(t))$ for the reallocation step, whose characterization depends on the allocation policy [17]. The evolution of a wavelength or a converter can be described as a discrete-time Markov chain (DTMC) with transition matrix

$$K_k^W(M^W(t)) = Q_k(M^W(t)) S_k A_k, \quad k \in \{w, c\}.$$

We now combine these two matrices into $K^W(M^W(t))$ to describe the evolution of a single object, which can be a wavelength or a converter, as a DTMC with two non-communicating classes

$$K^W(M^W(t)) = \begin{bmatrix} K_w^W(M^W(t)) & 0 \\ 0 & K_c^W(M^W(t)) \end{bmatrix}.$$

To compute $M^W(t)$ when W is large we rely on [16], where it is shown that, under some mild conditions, a system of interacting objects converges to its mean field when the number of objects is large. The mean field is a time-dependent deterministic system that can be used to approximate the behavior of a system with a large number of objects. Those conditions are valid for our model [17] and therefore we can approximate the evolution of the system via the mean field, which is described by the vector $\mu(t)$, for $t \geq 0$, and the kernel $K(\cdot) = K^W(\cdot)$. The vector $\mu(0)$ is initialized to describe an empty system and it evolves as $\mu(t+1) = \mu(t)K(\mu(t))$. Then, by [16, Theorem 4.1], for any fixed time t , almost surely, $\lim_{W \rightarrow \infty} M^W(t) = \mu(t)$.

Using the mean field model we can compute the state of the system at time t by performing t vector-matrix multiplications, where the vector is of size $1 \times m(ND + L_{\max}^2 + L_{\max} + 1)$. We are particularly interested in the long-run behavior of the switch but the mean field model is time-dependent and gives no additional information about the steady-state behavior, if it exists. We have numerically observed that when the conversion ratio is large enough to prevent losses caused by the lack of available converters, the state of the system seems to converge to a unique steady state. When the conversion ratio is not enough to avoid packet losses the system shows a stationary

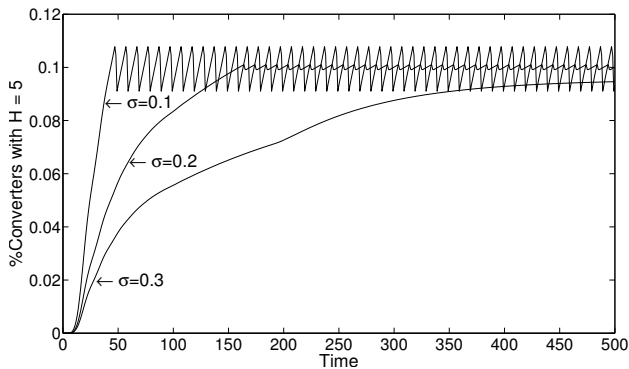
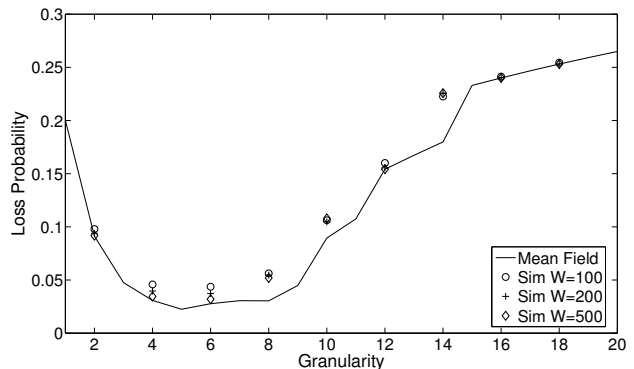


Figure 2. Time-dependent behavior of a switch with $N = 3$, $\rho = 0.8$, $D = 10$, geometric IATs and packet size equal to 10

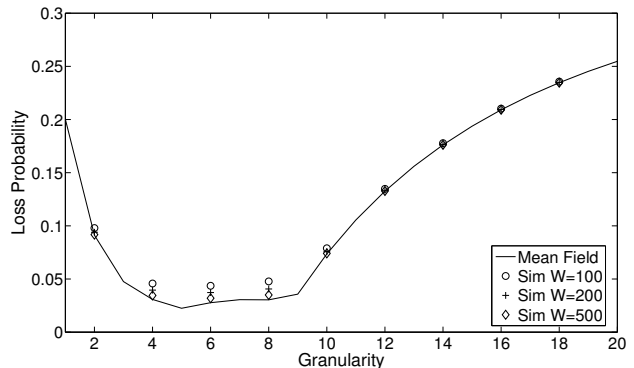
periodic behavior. The length of the period was observed to be the greatest common divisor of the possible packet sizes. Even though we do not provide a formal proof of this fact, the results presented in the next section, as well as many others not included here, support this observation [17].

IV. RESULTS

In this section we first concentrate in the long-run behavior of the model, showing the periodic and non-periodic cases. Next, we compare the results of the mean field model with estimates from the simulation of a switch with a finite number of wavelengths. Finally, we analyze the effect of various parameters on the switch performance. In Figure 2 we illustrate the time-dependent behavior of the mean field model using the fraction of converters with horizon equal to 5. The selection of this value is arbitrary as all the other entries in the state vector behave in a similar manner. To fix the arrival rate λ we use the load $\rho = \lambda \bar{L}$, where \bar{L} is the expected value of the packet size. In this scenario the switch has $N = 3$ FDLs per output port, the load ρ is 0.8, the granularity is $D = 10$, the burst length equals 10, the inter-arrival times (IATs) follow a geometric distribution (meaning $B_0 = 1 - 0.8/10 = 0.92$ and $B_{10} = 0.8/10 = 0.08$), the policy is *minG* and the conversion ratio is between 0.1 and 0.3. As can be seen, when the conversion ratio is equal to 0.1 the state of the converters is highly variable and after a short warm-up period it adopts a periodic behavior. When the conversion ratio rises to 0.2 the warm-up period becomes longer and the state of the converters is clearly less variable, but the period is exactly the same and equal to the packet size, in this case 10 slots. Finally, if the conversion ratio is equal to 0.3 no losses are caused by lack of converters. In this case the warm-up period is even longer but the system reaches a unique steady state. A similar behavior has been observed in all the experiments performed (including the simulations), with a periodic steady state and period equal to the greatest common divisor of the possible packet sizes. This periodic behavior arises when the conversion ratio is not enough to prevent packet losses. This is an important observation as it indicates that an underdimensioned number of WCs leads to a periodic system behavior. If there are plenty



(a) *minimum horizon* policy



(b) *minimum gap* policy

Figure 3. Mean field model vs. simulation for a switch with $N = 5$, $\rho = 0.8$, $\sigma = 10$, packet size equal to 10 and geometric IATs

of converters to translate any extra-packet, the system seems to converge to a unique steady state, as in Figure 2 for $\sigma = 0.3$.

A relevant issue for the mean field model is how it approximates the behavior of a finite system. Here we compare the results of the mean field model with results from simulation of a switch with 100, 200 and 500 wavelengths. The estimates from simulations have confidence intervals with half width less than 1% of the mean, obtained with the batch-means method. Figure 3 shows how the performance of the finite system tends to that of the mean field model, getting closer as the number of wavelengths increases. In this scenario, as in many others, the convergence for the *minG* policy, shown in Figure 3(b), is smoother than for the *minH* policy, shown in Figure 3(a). This is useful since the *minG* policy tends to use the buffer capacity in a more efficient manner as shown further on.

We now compare the spill, conversion and loss probabilities for both allocation policies. In Figure 4 these three quantities are shown for a switch with $N = 3$ FDLs, granularity $D = 10$, load equal to 0.8, geometric arrivals and packet size with equally probable values 8 and 12. For both policies the conversion probability increases linearly with the number of converters up to a point from which it no longer increases. During the interval where this probability increases the converters are the bottleneck of the system, and therefore they are busy all the time. When the switch has enough converters to translate any extra-packet, the switch no longer experiences losses due

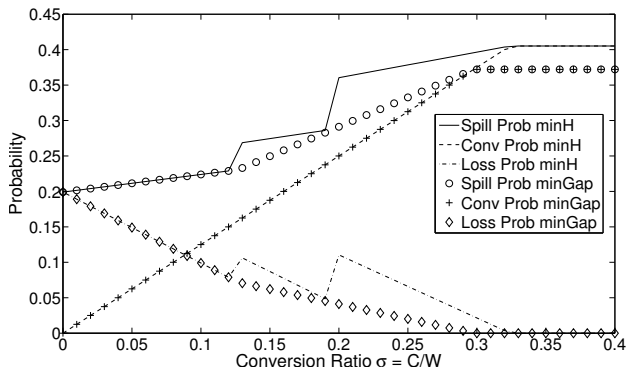


Figure 4. Comparison of policies for a switch with $N = 3$, $\rho = 0.8$, $D = 10$, geometric arrivals and packet size equal to $\{8, 12\}$

to the lack of converters. Notice, we can even determine the σ value where the loss rate becomes zero by running the mean field model once with $\sigma = 1$ and noting the percentage of busy converters, solving the dimensioning problem of WCs in a single run. The *minG* policy requires a smaller conversion ratio to reach the point where spill and conversion probabilities are the same than the *minH* policy. Furthermore, from this point on the spill probability under *minH* is larger than under *minG*, confirming the well-known result that *minH* is less efficient in managing the buffering resources (FDLs). An observation that can be made from Figure 4, also found in many other experiments, is the existence of jumps in the spill and loss probabilities as a function of the conversion ratio, for the *minH* policy. These jumps are closely related to the discrete nature of the FDLs and the way the *minH* policy reallocates the extra-packets. As this policy selects the wavelengths with minimum horizon and, if the number of converted packets is larger than the number of wavelengths with horizon 0, the packets are sent to the wavelengths with horizon equal to 1. However, this allocation creates large gaps (of size $D - 1$) in the wavelengths that receive the converted packets. This implies that the gap size distribution is affected in a bad manner, reducing the capacity of the wavelengths and causing the spill probability to increase. Hence, the jump in the spill probability, and therefore in the loss probability, is caused by an increase in the conversion ratio that makes the system able to convert more packets than the wavelengths with horizon equal to 0 are able to admit. This jump can be seen in Figure 4 when σ goes from 0.12 to 0.13. The other jumps occur similarly when the conversion ratio goes from a value in which the reallocated packets can be handled by the wavelengths with horizon less than or equal to iD to a value in which they cannot, for $1 \leq i \leq N$. Notice that the number of jumps is at most equal to N but might be less than this value.

Many other results can be obtained from the mean field model, as shown in [17]. Of particular interest to dimension the switch is the combined effect of the number of FDLs and the conversion ratio on the loss probability. We have found that the conversion ratio where the loss rate drops to zero σ^* may

decrease when the number of FDLs increases. Furthermore, if σ is such that the switch has losses due to the lack of converters, then using one or two FDLs might reduce the losses substantially. However, adding an extra FDL can also have no effect at all, even if the switch causes losses. The actual effect of an additional FDL strongly depends on the load. With the mean field model we are able to find $N - 1$ thresholds for the load value such that above threshold i having more than $N - i$ FDLs has no effect on σ^* , for $1 \leq i \leq N - 1$.

ACKNOWLEDGMENT

This work has been supported by the FWO-Flanders through project ‘‘Stochastic modeling of optical buffers and switching systems based on Fiber Delay Lines’’ (G.0538.07).

REFERENCES

- [1] J. Qiao and M. Yoo, ‘‘Optical burst switching: A new paradigm for an optical Internet,’’ *Journal of High-Speed Networks*, vol. 8, pp. 69–84, 1999.
- [2] J. Turner, ‘‘Terabit burst switching,’’ *Journal of High-Speed Networks*, vol. 8, pp. 3–16, 1999.
- [3] G. Latouche and V. Ramaswami, *Introduction to Matrix Analytic Methods in Stochastic Modeling*, ser. ASA-SIAM Series on Statistics and Applied Probability. Philadelphia, PA: SIAM, 1999.
- [4] N. Akar, E. Karasan, and K. Dogan, ‘‘Wavelength converter sharing in asynchronous optical packet/burst switching: an exact blocking analysis for markovian arrivals,’’ *IEEE J. Sel. Areas Commun.*, vol. 24, pp. 69–80, 2006.
- [5] L. Xu, H. Perros, and G. N. Rouskas, ‘‘A queueing network model of an edge optical burst switching node,’’ in *Proceedings of the IEEE Infocom 2003*, 2003.
- [6] B. Van Houdt, K. Laevens, J. Lambert, C. Blondia, and H. Bruneel, ‘‘Channel utilization and loss rate in a single-wavelength Fibre Delay Line (FDL) buffer,’’ in *Proceedings of IEEE Globecom 2004*, 2004.
- [7] K. Laevens, M. Moeneclaey, and H. Bruneel, ‘‘Queueing analysis of a single-wavelength fiber-delay-line buffer,’’ *Telecommunication Systems*, vol. 31, pp. 259–287, 2006.
- [8] J. Lambert, B. Van Houdt, and C. Blondia, ‘‘Queues with correlated inter-arrival and service times and its application to optical buffers,’’ *Stochastic Models*, vol. 22, no. 2, pp. 233–251, 2006.
- [9] F. Callegati, W. Cerroni, G. Corazza, C. Develder, M. Pickavet, and P. Demeester, ‘‘Scheduling algorithms for a slotted packet switch with either fixed or variable lengths packets,’’ *Photonic Network Communications*, vol. 8, pp. 163–176, 2004.
- [10] F. Callegati, W. Cerroni, C. Raffaelli, and P. Zaffoni, ‘‘Wavelength and time domain exploitation for QoS management in optical packet switches,’’ *Computer Networks*, vol. 44, pp. 569–582, 2004.
- [11] C. M. Gauger, ‘‘Optimized combination of converter pools and FDL buffers for contention resolution in optical burst switching,’’ *Photonic Network Communications*, vol. 8, pp. 139–148, 2004.
- [12] N. Akar, E. Karasan, G. Muretto, and C. Raffaelli, ‘‘Performance analysis of an optical packet switch employing full/limited range share per node wavelength conversion,’’ in *Proceedings of IEEE Globecom 2007*, 2007.
- [13] K. Dogan, Y. Gunulay, and N. Akar, ‘‘A comparative study of limited range wavelength conversion policies for asynchronous optical packet switching,’’ *Journal of Optical Networking*, vol. 6, pp. 134–145, 2007.
- [14] V. Puttasubba and H. Perros, ‘‘Performance analysis of limited-range wavelength conversion in an OBS switch,’’ *Telecommunication Systems*, vol. 31, pp. 227–246, 2006.
- [15] J. F. Pérez and B. Van Houdt, ‘‘Wavelength allocation in an optical switch with a fiber delay line buffer and limited-range wavelength conversion,’’ accepted for publication in *Telecommunication Systems*.
- [16] J. Le Boudec, D. McDonald, and J. Munding, ‘‘A generic mean field convergence result for systems of interacting objects,’’ in *Proc. 4th Int. Conf. on the Quantitative Evaluation of SysTems (QEST 2007)*, Edinburgh, UK, 16–19 Sep. 2007, pp. 3–15.
- [17] J. F. Pérez and B. Van Houdt, ‘‘A mean field model for dimensioning an OBS switch with partial wavelength conversion and fiber delay lines,’’ University of Antwerp, Tech. Rep. TR09/01, 2009.