Photonic Network Communications manuscript No.
(will be inserted by the editor)

The effect of Partial Conversion and Fiber Delay Lines in an OBS switch with a large number of wavelengths

Juan F. Pérez · Benny Van Houdt

Received: date / Accepted: date

Abstract In this paper, we analyze an OBS switch endowed with both wavelength converters (WCs) and fiber delay lines (FDLs) to resolve contention. We consider the case where the number of wavelengths is large by introducing a mean field model that provides exact results when the number of wavelengths tends to infinity. We have confirmed through simulations that the mean field model provides accurate approximations for switches with a large but finite number of wavelengths, which are of interest in view of wavelength division multiplexing (WDM). Furthermore, our model allows a very general behavior for the arrival process and the packet size distribution, as well as two different wavelength allocation policies: minimum horizon and minimum gap. Our results include a detailed analysis of the effect that these parameters have on the burst loss rate, and on the minimum number of WCs required to attain a zero loss rate as the number of wavelengths becomes large. We have found that at high loads there is little value in adding FDLs and, if included, shorter granularities result in fewer WCs required to achieve a zero loss rate. The inclusion of FDLs becomes more significant under mid loads and bursty traffic, where the addition of several FDLs may reduce the conversion requirements. Also, increasing the number of WCs under the minimum horizon policy may worsen the loss rate, while this is never the case for the minimum gap policy.

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).

Juan F. Pérez · Benny Van Houdt
Performance Analysis of Telecommunication Systems (PATS)
Department of Mathematics and Computer Science
University of Antwerp - IBBT
Middelheimlaan 1, B-2020 Antwerp, Belgium
E-mail: {juanfernando.perez, benny.vanhoudt}@ua.ac.be

Keywords Optical Switch · Partial Wavelength Conversion · Mean Field

1 Introduction

Nowadays optical fibers are able to carry a huge amount of information thanks to wavelength division multiplexing (WDM). With WDM several signals can be transmitted at the same time using different wavelengths, increasing the fiber capacity by tens or hundreds. To cope with this increasing capacity, optical burst switching (OBS) has been proposed as a solution to reduce the opto-electronic translations at the backbone network switches [16,18]. In OBS, only the header requires this translation, while the payload is processed in the optical domain. To resolve contention in the optical domain we consider two alternatives: wavelength conversion and optical buffering. Therefore, when an incoming packet requires transmission through a wavelength currently unavailable there are two options to prevent dropping the burst: it can be translated to a different wavelength using a wavelength converter; or it can be buffered using a Fiber Delay Line (FDL) until the wavelength becomes available. Furthermore, both solutions can be combined to allow a burst to be first converted to a different wavelength and then buffered, or vice versa. The main purpose of this paper is to analyze the effect of combining these solutions in an OBS switch with a large number of wavelengths.

Our Contribution: In this paper, we consider an OBS switch equipped with a pool of full-range wavelength converters and a set of FDLs per output port. Each of these ports has W wavelengths that can be used to simultaneously transmit the same number of bursts. The

number of wavelengths can be over a hundred due to the advances in WDM. To analyze the performance of this switch we introduce a mean field model that is exact when the number of wavelengths tends to infinity. Moreover, the performance of a switch with a large but finite number of wavelengths tends to that of the mean field as the number of wavelengths increases. Therefore, the mean field model can be used to approximate the performance of a switch with a large number of wavelengths, as has been confirmed by means of simulations. Furthermore, the time required to evaluate a particular scenario with the mean field model is typically a few seconds on a personal computer, compared to the long-lasting simulations that are needed when the number of wavelengths is large and the performance measures to evaluate are small (e.g., burst loss probability). As a result, the mean field model can be used to efficiently evaluate the effect of various parameters on the performance of the switch. These include, among others, the number of converters, the number and granularity of FDLs, the packet size distribution and the arrival process' burstiness. The main features of the mean field model can be summarized as follows: (i) contention is resolved by means of both wavelength conversion and optical buffering; (ii) the number of converters may vary between zero and the number of wavelengths, a property called *partial conversion*; (iii) two different policies are considered for wavelength allocation: minimum horizon and minimum gap; (iv) the arrival process is assumed to be a general Markovian arrival process (MAP) [8], which is able to represent general correlated inter-arrival times; (v) the burst size is assumed to follow a general distribution with finite support.

The main performance measure for the switch is the burst loss probability, which is the probability that an incoming burst has to be dropped due to the impossibility to transmit, convert or buffer it. Since the number of wavelengths is large, it is expected that the loss probability will decrease to a near-zero value if the number of converters is large enough, as in an Erlang loss system. Therefore, it becomes relevant to determine the minimum number of converters required to attain a near-zero loss probability. We can reformulate this by expressing the number of converters C in terms of the number of wavelengths W as $C = \sigma W$, where $\sigma \in (0, 1)$ (partial conversion). The goal is therefore to find the minimum value of σ such that the loss probability is almost zero, which will be denoted as σ^*. As will be shown later, with the mean field model we are able to compute the value of σ^* in a single run, which allows us to consider the effect that other parameters have on σ^*, specially the number and granularity of FDLs. Some of the insights we have gathered in this direction are:

1. If the system is under-dimensioned in terms of conversion resources, meaning $\sigma < \sigma^*$, 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.
2. The effect of the number of FDLs on the loss rate and σ^* is highly dependent on the load. While for high loads it may have little or no effect, for mid loads the addition of a few FDLs may reduce both the loss rate and σ^*.
3. Also, if the number of WCs is insufficient ($\sigma < \sigma^*$), increasing the number of FDLs may not improve the loss rate. However, under bursty traffic the effect of adding FDLs is more substantial, helping to reduce the conversion requirements.
4. The effect of the granularity on σ^* also depends on the load: while for low and mid loads the set of granularity values with the best performance depends on the packet size distribution, for high loads the performance is inversely proportional to the granularity. As a result, among the best possible values for the granularity under mid loads, the results favor the selection of a small granularity since this requires fewer converters to attain a near-zero loss probability at high loads.
5. The minimum horizon policy shows a consistently worse performance than its minimum gap counterpart. Moreover, if $\sigma < \sigma^*$ the addition of converters may worsen the loss rate under the minimum horizon policy.

Related work: The effect of wavelength conversion on a bufferless switch has been considered in [1,20] by means of analytical models. Separately, the performance of a switch equipped with FDLs but without converters has been treated in [9, 10, 19]. The analysis of a switch including both solutions turns out to be more complex since the multidimensional nature of a multi-wavelength switch has to be combined with the special queuing behavior of the optical buffer. The interaction of both wavelength conversion and FDLs has been analyzed by means of simulation models in [4,5,7]. 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 [2, 6, 14, 15]. It must be noted that a five-page version of this paper was presented in [13], where only an overview of the model and a few results were put forward.

The paper is organized as follows: Section 2 introduces the architecture and operation of the switch under analysis; Section 3 describes the mean field model

in detail; finally, Section 4 compares the results of the mean field model with results from the simulation of a finite system. This section also analyzes the effect of various parameters on the performance of the switch, with special emphasis on the effect of the allocation policies, the number and granularity of FDLs and the burstiness of the arrival process.

2 Switch Architecture and Operation

In this section, we describe the operation and main features of the optical switch, the wavelength allocation policies and some modeling issues relevant for the description of the switch. In this and the next sections we use the terms packet and burst interchangeably. 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 synchronous operation, as opposed to the asynchronous case, makes the switching matrix design simpler but requires packet synchronization and alignment [1, 5].

The arrival process at each wavelength is modeled as a Markovian Arrival Process (MAP) [8] characterized by the set of $m \times m$ matrices $\{B_0, B_1, \ldots, B_{L_{\max}}\}$, where $L_{\max}$ is the maximum packet length. 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 (i, j)-th entry 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 states of the underlying Markov chain are also referred to as the *phases* of the arrival process. The class of MAP processes has been used before to model the arrival process at a bufferless optical switch [1]. It includes many well-known processes as special cases, e.g., the discrete-time versions of the Poisson process, interrupted Poisson process (IPP), Markov modulated Poisson process (MMPP), etc. 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 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

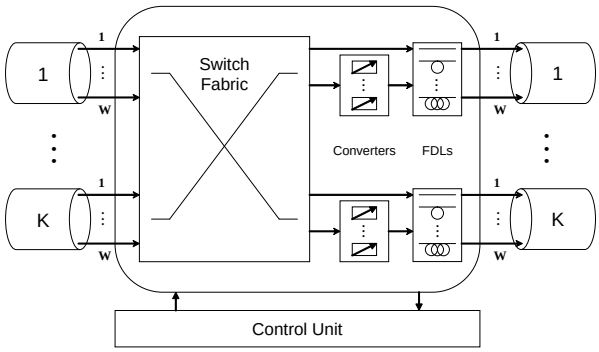

Figure 1 Optical switch with K input/output ports, W wavelengths, converters and FDLs

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 [7].

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 $D\left\lceil \frac{h}{D} \right\rceil$ slots, if $h \leq ND$. If the packet is of size L the new value of the horizon is $D\left\lceil \frac{h}{D} \right\rceil + L$. Notice, in this particular case the wavelength remains unused during $D\left\lceil \frac{h}{D} \right\rceil - h$ slots 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 in 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. 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*, which selects the wavelength with the minimum scheduling horizon; and *minimum gap*, which selects the wavelength with a horizon such that the allocation of a new packet generates a gap of minimum value among all available wavelengths. Recall, the gap is the difference between the horizon observed by an incoming packet and the actual delay that a packet assigned to the wavelength must face.

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 current state $i \in \{1, \ldots, m\}$ 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 the *loss probability* and is considered the main measure of performance.

3 Mean Field model

Our model is based on a general result for a system of interacting objects introduced in [11]. In our case, the system consists of two types of objects: wavelengths and converters. To describe the evolution of the system during a time slot we start with the state of the objects at the beginning of the time slot. Then we determine the transition matrices that describe the state transitions at each of the three steps: transmission, arrivals and reallocation. The matrices associated to these steps are S_k, A_k and Q_k, respectively, where the subscript k may be equal to w or c depending on whether the matrix describes the transition of a wavelength or a converter. These matrices are then used to build a complete description of the evolution of the switch at each time slot. At the beginning of slot t (before packet transmission) the state of a *single* wavelength can be described by the tuple $\{(H(t), J(t)), t \geq 0\}$, with $H(t)$ the scheduling horizon of the wavelength and $J(t)$ the phase of its arrival process. Its state space is the set $\{(i, j) | 0 \leq i \leq ND + L_{\max}, 1 \leq j \leq m\}$. Similarly, a converter can be described by its scheduling horizon $\{\bar{H}(t), t \geq 0\}$ with state space $\{i | 0 \leq i \leq L_{\max}\}$. We now define the evolution matrices for each of the three steps.

3.1 Step 1, packet transmission

In the first step (**S1**) the horizon of each busy wavelength and each busy converter is reduced by one, as they transmit (translate) part of the scheduled packets. Let T_n be the $(n+1) \times n$ matrix with entries

$$[T_n]_{ij} = \begin{cases} 1, i = j = 1, \\ 1, j = i - 1, i = 2, \ldots, n+1, \\ 0, \text{otherwise.} \end{cases}$$

Then the evolution of a single wavelength in **S1** is given by the transition matrix $S_w = T_{ND+L_{\max}} \otimes I_m$, where I_n is the identity matrix of size n and $\otimes$ denotes the Kronecker product. This product shows that packet transmissions affect the horizon value but not the phase of the arrival process. Accordingly, the matrix $S_c = T_{L_{\max}}$ contains the transition probabilities for a converter during **S1**.

3.2 Step 2, packet arrivals

The arrival of packets during the second step (**S2**) has no influence on the state of a converter; therefore, its transition matrix in this step is given by $A_c = I_{L_{\max}}$. Similarly, the matrix A_w describes the transition of a

single wavelength in **S2**, but its definition is more involved. If after **S1** the wavelength has a horizon less than or equal to ND, it can accept any incoming packet. On the other hand, if the scheduling horizon is greater than ND and a packet arrives, it cannot be buffered and becomes part of the extra-packets. To keep track of the size of the (possibly empty) set of extra-packets, the horizon and the phase of the arrival process, we separate the resulting state space after **S2** into two sets. The first set is $\{(i,j)|0 \leq i \leq ND + L_{\max}, 1 \leq j \leq m\}$, which captures two cases: the horizon was less than or equal to ND after **S1**, and the transition in **S2** results in a horizon equal to i and a phase of the arrival process equal to j; or the horizon was greater than ND but no packet is received. In this first set the wavelength holds zero extra-packets. The second set is $\{(ND + L_{\max} + i, k, j)|1 \leq i \leq L_{\max}–1, 1 \leq k \leq L_{\max}, 1 \leq j \leq m\}$, considering the case where the horizon was equal to $ND+i$ after **S1** and the arrival process (during **S2**) generates a packet of size k and makes a transition to phase j. The two sets of states are put together by imposing a lexicographic order, resulting in a transition matrix A_w of size $m(ND + L_{\max}) \times m(ND + L_{\max}^2 + 1)$ (since the horizon after step **S1** is at most $ND + L_{\max} - 1$).

To explicitly describe the matrix A_w we partition the state space before and after **S2** in levels. Before **S2**, level i is the set of states $\{(i,j)|1 \leq j \leq m\}$, for $0 \leq i \leq ND+L_{\max}-1$. After **S2**, there are two subsets of levels, corresponding to the two subsets of the state space described above: in the first subset, level i is the set of states $\{(i,j)|1 \leq j \leq m\}$, for $0 \leq i \leq ND + L_{\max}$; in the second subset, level (i,k) is the set of states $\{(ND+L_{\max}+i,k,j)|1 \leq j \leq m\}$, for $1 \leq i \leq L_{\max}-1$ and $1 \leq k \leq L_{\max}$. Let the matrix $A_w^{\{i,i'\}}$ contain the transition probabilities from level i to level i', for $0 \leq i \leq ND + L_{\max}-1$ and $0 \leq i' \leq ND+L_{\max}$. These matrices are given by

$$A_w^{\{i,i'\}} = \begin{cases} B_{i'},\ i = 0,\ 0 \leq i' \leq L_{\max}, \\ B_0,\ 1 \leq i = i' \leq ND + L_{\max} - 1, \\ B_k,\ i' = D\left\lceil \frac{i}{D} \right\rceil + k,\ 1 \leq i \leq ND,\ 1 \leq k \leq L_{\max}. \end{cases}$$

To define these matrices we consider each case separately: in the first case the wavelength is idle and can start the transmission of a newly arriving packet immediately; in the second case the wavelength is busy and receives no packet in this slot; in the last case the wavelength is busy and receives a new packet that can be buffered, increasing the scheduling horizon. Similarly, let the matrix $A_w^{\{i,(i',k')\}}$ contain the transition probabilities from level i to level (i',k'), for $0 \leq i \leq ND + L_{\max}-1$, $1 \leq i' \leq L_{\max}-1$ and $1 \leq k' \leq L_{\max}$. These matrices are given by

$$A_w^{\{ND+i,(i',k')\}} = B_{k'},$$

for $1 \leq i = i' \leq L_{\max} - 1$ and $1 \leq k' \leq L_{\max}$. These transitions correspond to the case where the wavelength is busy and receives a packet that cannot be buffered. Since all the possible transitions in **S2** have been covered, we now turn to the last step.

3.3 Step 3, packet conversion and reallocation

In this step (**S3**) the extra-packets that arrived in the previous step are reallocated using the available converters. To determine the evolution of a *single* wavelength or converter it is necessary to consider the state of the whole system (W wavelengths and C converters). It is important to stress however that we do not need to determine the joint evolution of multiple wavelengths or converters for the mean field result to apply. Let $w_i(t)$ be the $1 \times m$ vector whose j-th entry contains the number of wavelengths holding no extra-packets with horizon equal to i and phase of the arrival process equal to j after **S2**, for $0 \leq i \leq ND + L_{\max}$ and $1 \leq j \leq m$. Additionally, let the j-th entry of the $1 \times m$ vector $w_{(ND+L_{\max}+i,k)}(t)$ be the number of wavelengths at time t with horizon equal to $ND + i$ after **S1** that receive a packet of size k in **S2**, after which the phase of the arrival process is equal to j, for $1 \leq i \leq L_{\max}-1$, $1 \leq k \leq L_{\max}$ and $1 \leq j \leq m$. The vector

$$M^{W,(w)}(t) = \frac{1}{W}[w_0(t), \ldots, w_{ND+L_{\max}}(t), w_{(ND+L_{\max}+1,1)}(t), \ldots, w_{(ND+2L_{\max}-1,L_{\max})}(t)]$$

describes the state of all the wavelengths at time t before **S3** as fractions of the total number of wavelengths W. Analogously, let $c_i(t)$ be the number of converters with horizon equal to i at time t before **S3**, for $i = 0, \ldots, L_{max}$. The state of the converters at time t, as a fraction of the total number of converters, is therefore contained in the vector $M^{W,(c)}(t) = \frac{1}{C}[c_0(t), \ldots, c_{L_{\max}}(t)]$. The superscript W indicates that the system is composed of W wavelengths and $C = \sigma W$ converters.

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. The weights $\frac{1}{1+\sigma}$ and $\frac{\sigma}{1+\sigma}$ are

the proportion of wavelengths and converters, respectively, in relation to the total number of objects. Based on this vector, we can define the matrices $Q_w(M^W(t))$ and $Q_c(M^W(t))$, which contain the transition probabilities in **S3** under the *minimum horizon* policy for wavelengths and converters, respectively. The matrices $\bar{Q}_w(M^W(t))$ and $\bar{Q}_c(M^W(t))$ contain similar information for the *minimum gap* policy. However, to specify these matrices it is necessary to first determine the number and size of the extra-packets that can actually be converted, regardless the wavelength allocation policy.

Let $d_i(M^W(t))$ be the number of extra-packets of size i, for $1 \leq i \leq L_{\max}$, which is given by

$$d_i(M^W(t)) = \sum_{k=1}^{L_{\max}-1} w_{(ND+L_{\max}+k,i)}(t)\mathbf{1}_m,$$

where $\mathbf{1}_m$ is a column vector of size m with all its entries equal to one. Therefore, the total number of extra-packets is $d(M^W(t)) = \sum_{i=1}^{L_{\max}} d_i(M^W(t))$. Also, let $W_{ND}(M^W(t))$ be the number of wavelengths with horizon less than or equal to ND after **S2**, i.e.,

$$W_{ND}(M^W(t)) = \sum_{i=0}^{ND} w_i(t)\mathbf{1}_m.$$

The number of extra-packets that can actually be converted ($\hat{d}(M^W(t))$) is given by the minimum of three quantities: the number of packets to convert, the number of wavelengths with available buffering capacity, and the number of available converters, i.e.,

$$\hat{d}(M^W(t)) = \min\{d(M^W(t)), W_{ND}(M^W(t)), c_0(t)\}.$$

Since each wavelength with available buffering capacity receives at most one extra-packet, $\hat{d}(M^W(t))$ is also the number of wavelengths that receive an extra-packet in **S3**. The selection of these $\hat{d}(M^W(t))$ wavelengths is done using the *minimum horizon* or *minimum gap* policies. Once a wavelength is chosen to receive an extra-packet, the selection of the packet is done randomly among the $d(M^W(t))$ extra-packets. This means that the probability that a selected wavelength receives a packet of size k, for $1 \leq k \leq L_{\max}$, is $p_k(M^W(t)) = \frac{d_k(M^W(t))}{d(M^W(t))}$. Relying on these definitions, the purpose of the following subsections is to determine the transition matrices for both wavelength allocation policies.

3.3.1 Minimum Horizon

To determine the wavelengths that, under the *minimum horizon* (*minH*) policy, will receive the $\hat{d}(M^W(t))$ extra-packets, we need to define the quantities $\alpha_i(M^W(t))$ as the number of wavelengths with horizon less than or equal to i after **S2**, i.e., $\alpha_i(M^W(t)) = \sum_{k=0}^{i} w_k(t)\mathbf{1}_m$, for $0 \leq i \leq ND$. As the extra-packets are assigned to the wavelengths with the smallest horizons, we need to find an $h(M^W(t))$ such that

$$\alpha_{h(M^W(t))-1} < \hat{d}(M^W(t)) \leq \alpha_{h(M^W(t))}.$$

This means that the wavelengths with horizon strictly less than $h(M^W(t))$ receive one extra packet each, while those with a horizon strictly greater than $h(M^W(t))$ receive no extra-packets. The packets that cannot be accommodated in the wavelengths with horizons up to $h(M^W(t))-1$ are randomly assigned among the wavelengths with horizon equal to $h(M^W(t))$. Let $\theta(M^W(t))$ be the probability that a wavelength receives a packet in **S3** if its horizon is equal to $h(M^W(t))$. This is given by

$$\theta(M^W(t)) = \frac{\hat{d}(M^W(t)) - \alpha_{h(M^W(t))-1}}{w_{h(M^W(t))}(t)\mathbf{1}_m}.$$

Now we can define $r_i(M^W(t))$, the probability that a wavelength with horizon equal to i receives an extra-packet in **S3** under the *minH* policy, as

$$r_i(M^W(t)) = \begin{cases} 1, & 0 \leq i < h(M^W(t)), \\ \theta(M^W(t)), & i = h(M^W(t)), \\ 0, & h(M^W(t)) < i \leq ND. \end{cases}$$

Let $u_{ii'}(M^W(t))$ be the probability that a wavelength with horizon i and holding no extra-packets after **S2** ends up with a horizon equal to i' after **S3**, for $0 \leq i \leq ND + L_{\max}$ and $0 \leq i' \leq ND + L_{\max}$. These probabilities are given by

$$u_{ii'}(M^W(t)) = \begin{cases} 1 - r_i(M^W(t)), & 0 \leq i = i' \leq ND, \\ r_i(M^W(t))p_k(M^W(t)), & 0 \leq i \leq ND, \\ & i' = D\left\lceil \frac{i}{D} \right\rceil + k, \\ 1, & ND < i = i' \leq ND + L_{\max}. \end{cases}$$

Now let $u_{(i,k),i'}(M^W(t))$ be the probability that a wavelength with horizon $ND+i$ after **S1** and that received a packet of size k in **S2**, ends up with a horizon equal to i' after **S3**, for $1 \leq i \leq L_{\max}-1$, $1 \leq k \leq L_{\max}$ and $0 \leq i' \leq ND + L_{\max}$. Since such a wavelength keeps its horizon independently of whether the extra-packet is successfully reallocated or not, the transition probabilities are given by

$$u_{(i,k),ND+i'}(M^W(t)) = \begin{cases} 1, & 1 \leq i = i' \leq L_{\max}-1, \\ & 1 \leq k \leq L_{\max}, \\ 0, & \text{otherwise.} \end{cases}$$

Let $U(M^W(t))$ be the $(ND+L_{\max}^2+1)\times(ND+L_{\max}+1)$ matrix that describes the evolution of the horizon during **S3**. The first $ND+L_{\max}+1$ rows of this matrix have entries $u_{ii'}(M^W(t))$, for $0\le i\le ND+L_{\max}$. The remaining $(L_{\max}-1)L_{\max}$ rows are made by the entries $u_{(i,k),i'}(M^W(t))$ in lexicographic order, for $1\le i\le L_{\max}-1$ and $1\le k\le L_{\max}$. Therefore, the transition matrix for a single wavelength during **S3** is

$$Q_w(M^W(t)) = U(M^W(t))\otimes I_m,$$

making explicit that the allocation of extra-packets has no effect on the phase of the arrival process.

With regard to the converters, only those with horizon equal to 0 may be affected during **S3** since these can be used to translate the $\hat{d}(M^W(t))$ extra-packets. Let $b_i(M^W(t))$ be the probability that an idle converter receives a packet of size i in **S3**, for $1\le i\le L_{\max}$. Also, let $b_0(M^W(t))$ be the probability that the converter remains idle. Clearly,

$$b_i(M^W(t)) = \begin{cases} \frac{c_0(t)-\hat{d}(M^W(t))}{c_0(t)}, & i=0,\\ \frac{\hat{d}(M^W(t))}{c_0(t)}p_i(M^W(t)), & 1\le i\le L_{\max}.\end{cases}$$

Therefore, the entries of the $L_{\max}\times(L_{\max}+1)$ transition matrix for a single converter in **S3** can be defined as

$$[Q_c(M^W(t))]_{ij} = \begin{cases} b_j(M^W(t)), & i=0,\ 0\le j\le L_{\max},\\ 1, & 1\le i=j\le L_{\max}-1,\\ 0, & \text{otherwise.}\end{cases}$$

3.3.2 Minimum Gap

In this section we determine the matrices $\bar{Q}_w(M^W(t))$ and $\bar{Q}_c(M^W(t))$ to describe the evolution of the system during **S3** under the *minimum gap* (*minG*) policy. Since the wavelength allocation policy has no effect on the state of the converters, $\bar{Q}_c(M^W(t)) = Q_c(M^W(t))$. To specify the transition matrix for a single wavelength we start by defining the gap function $v(h) = D\left\lceil\frac{h}{D}\right\rceil - h$, which is the size of the gap created when assigning a packet to a wavelength with horizon h, for $0\le h\le ND$. Now we can define $g_i(M^W(t))$ as the number of wavelengths with $v(\cdot)=i$, which is given by

$$g_i(M^W(t)) = \sum_{\{j|v(j)=i\}} w_j(t)\mathbf{1}_m, \quad 0\le i\le D-1.$$

In a similar way as in the previous section, we define $\gamma_i(M^W(t))$ as the number of wavelengths with $v(\cdot)\le i$, i.e., $\gamma_i(M^W(t)) = \sum_{j=0}^{i} g_j(M^W(t))$, for $0\le i\le D-1$. In this case we need to find an $x(M^W(t))$ such that

$$\gamma_{x(M^W(t))-1} < \hat{d}(M^W(t)) \le \gamma_{x(M^W(t))}.$$

Thus, $\gamma_{x(M^W(t))-1}$ extra-packets can be assigned to the wavelengths with $v(\cdot) < x(M^W(t))$. The packets that cannot be accommodated in these wavelengths are distributed among those with $v(\cdot) = x(M^W(t))$, while the rest of the wavelengths receive zero extra-packets. In this case, however, we use the *minH* policy to allocate the remaining $\hat{d}(M^W(t)) - \gamma_{x(M^W(t))-1}$ extra-packets among the wavelengths with $v(\cdot) = x(M^W(t))$ (as opposed to randomly). Since the horizon h can be expressed as $h = D\left\lceil\frac{h}{D}\right\rceil - v(h)$ and the wavelengths that may receive a packet have $v(\cdot) = x(M^W(t))$, we only need to focus on $l(h) = \left\lceil\frac{h}{D}\right\rceil$, which takes values between 0 and N. Let $f_i(M^W(t))$ be the number of wavelengths with horizon h such that $v(h)=x(M^W(t))$ and $l(h) = i$, for $0\le i\le N$. Also, let $\phi_i(M^W(t)) = \sum_{j=0}^{i} f_j(M^W(t))$ be the number of wavelengths with horizon h such that $v(h) = x(M^W(t))$ and $l(h)\le i$, for $0\le i\le N$. We then need to find a $y(M^W(t))$ such that

$$\phi_{y(M^W(t))-1} < \hat{d}(M^W(t)) - \gamma_{x(M^W(t))-1} \le \phi_{y(M^W(t))}.$$

Then, among the wavelengths with horizon h such that $v(h) = x(M^W(t))$, one extra-packet is assigned to the wavelengths with $l(h) < y(M^W(t))$, no extra-packet is assigned to those with $l(h) > y(M^W(t))$, and the remaining extra-packets are randomly assigned among the wavelengths with $l(h) = y(M^W(t))$. Therefore, the probability that a wavelength with horizon h such that $v(h) = x(M^W(t))$ and $l(h) = y(M^W(t))$ receives an extra-packet during **S3** is

$$\eta(M^W(t)) = \frac{\hat{d}(M^W(t)) - \gamma_{x(M^W(t))-1} - \phi_{y(M^W(t))-1}}{f_{y(M^W(t))}(M^W(t))}.$$

Now we can define $\bar{r}_i(M^W(t))$ as the probability that a wavelength with horizon equal to i receives an extra-packet in **S3** under the *minG* policy, given by

$$\bar{r}_i(M^W(t)) = \begin{cases} 1, & 0\le v(i) < x(M^W(t)),\\ 1, & v(i) = x(M^W(t)),\\ & l(i) < y(M^W(t)),\\ \eta(M^W(t)), & v(i) = x(M^W(t)),\\ & l(i) = y(M^W(t)),\\ 0, & \text{otherwise.}\end{cases}$$

Based on these probabilities we can build the matrix $\bar{U}(M^W(t))$ in the same manner as we did with the matrix $U(M^W(t))$ for the *minH* policy, but replacing the $r_i(M^W(t))$ by $\bar{r}_i(M^W(t))$, for $0\le i\le ND$. Thus, the transition matrix of a wavelength in **S3** under the *minG* policy is $\bar{Q}_w(M^W(t)) = \bar{U}(M^W(t))\otimes I_m$.

3.4 Computation of $M^W(t)$ for large W

In the previous sections we built the transition matrices related to each of the three main events (steps) in a slot, for wavelengths and converters separately. These matrices can be combined to describe the evolution of a single object as a discrete-time Markov chain (DTMC). We will observe the system just after **S2** and, therefore, the state of the wavelengths (resp. converters) at time t is described by the vector $w^W(t)$ (resp. $c^W(t)$). Since the order of the events is **S3**, **S1** and **S2**, the transition matrices of a single wavelength or converter under the *minH* policy are

$$K_k^W(\bar{m}) = Q_k(\bar{m})\, S_k\, A_k,\ k \in \{w, c\},$$

where $\bar{m}$ is an occupancy vector that represents any particular value of $M^W(t)$. The superscript W refers to the total number of wavelengths in the system. We now combine these two matrices into $K^W(\bar{m})$ to describe the evolution of a single object, which can be a wavelength or a converter, as a DTMC with two non-communicating classes

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

A similar construction can be made to determine the matrix $\bar{K}^W(\bar{m})$ for the *minG* policy.

We now consider the framework in [11] to compute $M^W(t)$ when W is large. The discussion is for the *minH* policy, but it applies *mutatis mutandis* for the *minG* policy. In [11] the authors show 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. The first condition for this result to hold is that the entries of the transition matrix of a single object $[K^W(\bar{m})]_{ij}$ converge uniformly to some $[K(\bar{m})]_{ij}$ on the set of all occupancy vectors when $W(1+\sigma) \to \infty$. In our model the transition matrix $K^W(\bar{m})$ is actually independent of the number of objects $W(1+\sigma)$. This can be seen by dividing all the quantities involved in the computation of the probabilities $u_{ii'}(M^W(t))$ and $b_i(M^W(t))$ by $W(1+\sigma)$. This means that $K(\bar{m}) = K^W(\bar{m})$. The second condition is that $[K(\bar{m})]_{ij}$ must be continuous in $\bar{m}$, which also holds for both allocation policies. Since both conditions are valid for the model described by the matrix $K(\bar{m})$, we can approximate the evolution of the system by means of the mean field, which is described by the vector $\mu(t)$, for $t \geq 0$. Let $\mu(t) = \left[\frac{1}{1+\sigma}\mu^{(w)}(t),\ \frac{\sigma}{1+\sigma}\mu^{(c)}(t)\right]$, for $t \geq 0$. The initial state of the wavelengths is defined as $\mu^{(w)}(0) = [\pi_B, 0, \ldots, 0]$, where the $1 \times m$ vector π_B is the stationary probability distribution of the Markovian arrival process. Similarly, the vector $\mu^{(c)}(0) = [1, 0, \ldots, 0]$ describes the initial state of the converters. The initial distribution is independent of the number of objects and establishes that all the wavelengths and converters are idle at time 0. Now, let the mean field model evolve as $\mu(t+1) = \mu(t)K(\mu(t))$, then, by [11, Theorem 4.1], for any fixed time t, almost surely,

$$\lim_{W \to \infty} M^W(t) = \mu(t).$$

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 $1 \times m(ND + L_{\max}^2 + L_{\max} + 1)$. Additionally, at each time slot the matrix $Q(\bar{m})$ (or $\bar{Q}(\bar{m})$) must be computed since it depends on the value of the occupancy vector. However, if the packet-size distribution is independent of the arrival process the description of the system after **S**2 can be simplified. In this case the probability distribution of the extra-packets' size is equal to the original packet-size distribution. Therefore it is not necessary to keep track of the size of the extra-packets but only their number, reducing the size of the occupancy vector to $1 \times m(ND + 3L_{\max})$. Also, the structure of the transition matrices in **S**1 can be exploited to further reduce the computation times.

In addition to approximate the state of a switch with a large number of wavelengths at any finite time t, we are particularly interested in its long-run behavior, but the mean field model is time-dependent and gives no information about the steady-state, if it exists. However, 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 converges to a unique steady state. When the conversion ratio is not enough to avoid packet losses the system shows a stationary periodic behavior. The length of the period has been 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. Actually, a formal proof appears to be hard since the evolution of the sequence of occupancy vectors $\{\mu(t)t \geq 1\}$ is determined by a non-homogeneous Markov chain, whose transition probabilities at time t are a nonlinear function of the occupancy vector. Moreover, it must be shown that $\{\mu(t)t \geq 1\}$ converges either toward a single point (if the number of converters is enough to prevent losses), or toward a set of points that the sequence will visit cyclically (if there are losses because of the lack of converters). Let

δ be the greatest common divisor of the possible packet sizes. As we do not know in advance if the conversion ratio is enough to prevent losses or not[1], we observe the system every δ time slots to check the difference in the entries of the occupancy vector, and we let it evolve until this difference is less than $\epsilon = 10^{-10}$. For each of the δ steady states we compute the performance measures, as shown in the next section, and their average is the value of the steady-state performance measures.

3.5 Computation of the performance measures

If time t corresponds to a steady state, then $d(M^W(t))$ is also the number of packets requiring conversion per slot in this steady state, which we call the *spill rate*. Similarly, $\hat{d}(M^W(t))$ is called the *conversion rate*, while $d(M^W(t)) - \hat{d}(M^W(t))$ is the *loss rate*. In a system with W wavelengths the total arrival rate is $W\lambda$, where λ is the arrival rate at each wavelength, given by

$$\lambda = \pi_B \sum_{k=1}^{L_{\max}} B_k \mathbf{1}_m = \pi_B (I_m - B_0)\mathbf{1}_m.$$

Therefore the spill probability p_{spill}, i.e., the probability that an incoming packet requires conversion, is given by $p_{\text{spill}} = \frac{d(M^W(t))}{W\lambda}$. Dividing the numerator and denominator by the number of objects $W(1+\sigma)$, we get

$$p_{\text{spill}} = \frac{\delta(M^W(t))}{\frac{\lambda}{1+\sigma}},$$

where $\delta(M^W(t)) = \frac{d(M^W(t))}{W(1+\sigma)}$ is independent of the number of objects. In a similar manner, we define $\hat{\delta}(M^W(t))$ as $\frac{\hat{d}(M^W(t))}{W(1+\sigma)}$, which allows us to define the conversion probability p_{conv} and the loss probability p_{loss} as

$$p_{\text{conv}} = \frac{\hat{\delta}(M^W(t))}{\frac{\lambda}{1+\sigma}},\ p_{\text{loss}} = \frac{\delta(M^W(t)) - \hat{\delta}(M^W(t))}{\frac{\lambda}{1+\sigma}}.$$

4 Results

The purpose of this section is two-fold: first, we illustrate the time-dependent behavior of the mean field model as well as its convergence toward a state that matches well with the results obtained by means of simulation. This will be considered in Section 4.1. Second, we make use of the mean field model to analyze the effect of the switch parameters on its performance. We consider the loss probability as the main measure of performance and put special emphasis on the minimum conversion ratio required to attain zero loss probability, referred to as σ^*. In a switch with a finite number of wavelengths, the goal is to determine a conversion ratio such that the loss probability stays below a certain predefined threshold. Since there are no analytical models available to determine the exact loss probability in a switch with a large number of wavelengths, FDLs and partial wavelength conversion, the only alternative is to rely on simulation. However, to estimate a very small loss probability using simulation requires long computation times since the event that must be observed (a packet loss) becomes very unlikely. One of the main advantages of the mean field model is the fast computation of the approximate loss probability and σ^* for any particular scenario with a large number of wavelengths. This allows the analysis of the effect that the various switch parameters have on these performance measures. Sections 4.2 and 4.3 deal with these issues, where the latter is concerned with the effect of the arrival process' burstiness on σ^*.

[1] Actually, by running the mean field model once with $\sigma = 1$, we can determine the required σ value at once.

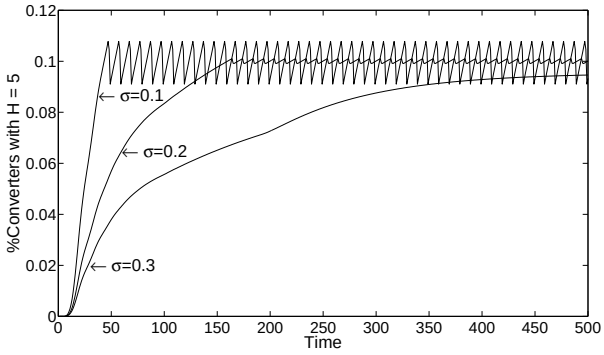

Figure 2 Time-dependent behavior of a switch with $N = 3$, $\rho = 0.8$, $D = 10$, geometric inter-arrival times and packet size equal to 10

4.1 Validation

Given the time-dependent character of the mean field model, there is a natural interest in the behavior of the state vector $\mu(t)$ as a function of time. In Figure 2 we illustrate this behavior using the fraction of converters with horizon equal to 5, i.e., $\mu_5^{(c)}(t)$. 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 we use the load $\rho = \lambda E[L]$, where $E[L]$ 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

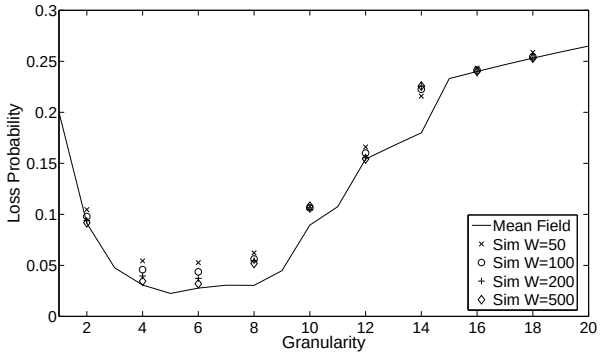

(a) *minimum horizon* policy

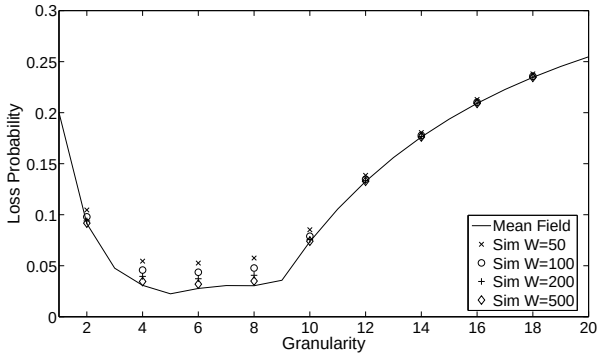

(b) *minimum gap* policy

Figure 3 Mean field model vs. simulation for a switch with $N = 5$, $\rho = 0.8$, $\sigma = 0.1$, packet size equal to 10 and geometric inter-arrival times

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

Since the mean field model tends toward a (possibly periodic) steady state, we can use this state to compute the performance measures of the system as indicated in Section 3.5. A first question to address is how the mean field model approximates the behavior of a finite system. In Figure 3 we compare the results of the mean field model with results from simulation of a switch with 50, 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. As can be expected, the simulations require long execution times to obtain a small confidence interval for the loss probability. 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). For both policies, the accuracy of the mean field approximation depends on the granularity, being almost exact for granularities well above the packet size. For granularities between 2 and 10, the performance of the finite system smoothly converges to that of the mean field for the *minG* policy. On the other hand, the convergence for the *minH* policy shows different patterns for different granularity values. Notably, when $D = 14$ the performance of the finite system does not appear to converge to the mean field. We have observed that, in this case the loss probability of a sequence of finite systems with increasing number of wavelengths first increases and then decreases toward the loss probability of the mean field, but this only occurs when the finite system has a few thousand wavelengths. On the other hand, the *minG* policy aims at minimizing the gaps in the FDLs, and an increase in the number of wavelengths directly implies more options to allocate an extra-packet while creating the smallest possible gap. Therefore, increasing the number of wavelengths will reduce the loss probability under this policy and the convergence will be smoother than under the *minH* policy. As a result, when approximating the performance of a finite system with a given number of wavelengths, the mean field model is expected to be less accurate under the *minH* than under the *minG* policy.

4.2 Combined effect of FDLs and WCs

One of the main characteristics of the mean field model is its ability to include both partial wavelength conversion and buffering as solutions for contention resolution. We exploit this feature in this section by analyzing the effect of three main parameters: the conversion ratio σ, the number of FDLs N and their granularity D, as well as the wavelength allocation policy. The arrival process is assumed to be geometric as the effect of burstiness

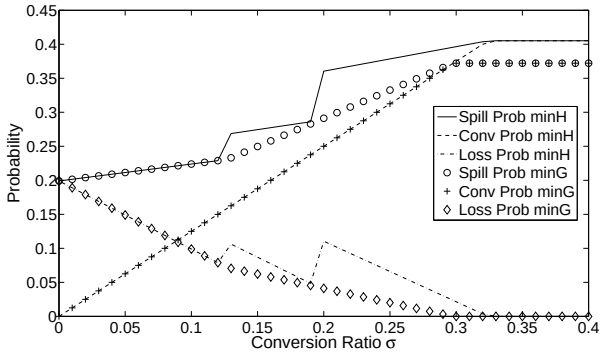

Figure 4 Comparison of policies for a switch with $N = 3$, $\rho = 0.8$, $D = 10$, geometric arrivals and packet size equal to $\{8, 12\}$

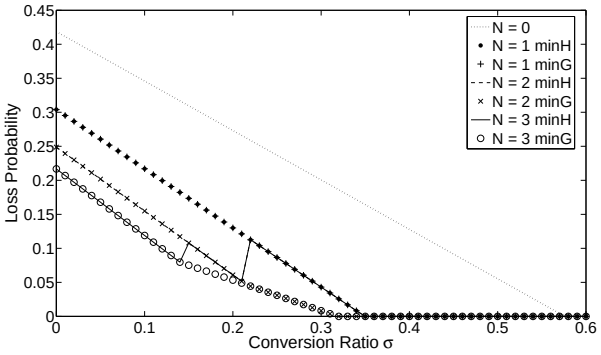

Figure 5 Comparison of policies for a switch with $\rho = 0.8$, $D = 10$, geometric arrivals and packet size equal to $\{5, 15\}$

in the arrival process will be the topic of the next section. We start by comparing 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 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, i.e., when spill and conversion probabilities are equal, the switch no longer experiences losses due to the lack of converters. Notice, we can determine the σ value where the loss rate becomes zero (σ^*) 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. From the figure we observe that the *minG* policy requires a smaller conversion ratio than the *minH* policy to reach the point where spill and conversion probabilities are the same. 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 Figure 5 as well as 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, the reallocated packets go first to the wavelengths with horizon 0 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.

A relevant issue in the design of an optical switch is the influence of the number of FDLs on the loss probability. Figure 5 shows the loss probability as a function of the conversion ratio, for a variable number of FDLs and both allocation policies. The packet size can be 5 or 15 with equal probability, the load is 0.8 and the granularity is 10. The effect of adding FDLs on the loss probability depends on the conversion ratio. If the conversion ratio is large enough, then adding more FDLs has *no effect*. However, the conversion ratio σ where the loss rate drops to zero does depend on N. For instance, in Figure 5, having $N = 1$ FDLs allows us to use significantly fewer WCs compared to having zero FDLs, while increasing N to 2 has a smaller effect, and an additional FDL has no effect (as a buffer capacity of $N = 2$ suffices with $C = 0.3W$ WCs). If σ is such that the switch has losses due to the lack of converters, then the addition of buffering capacity might reduce the losses substantially. However, adding an extra FDL might also have no effect at all, even if the switch presents losses. This is clear in Figure 5 for $\sigma = 0.25$, where the loss with two FDLs is lower than with one, but the addition of a third makes no difference.

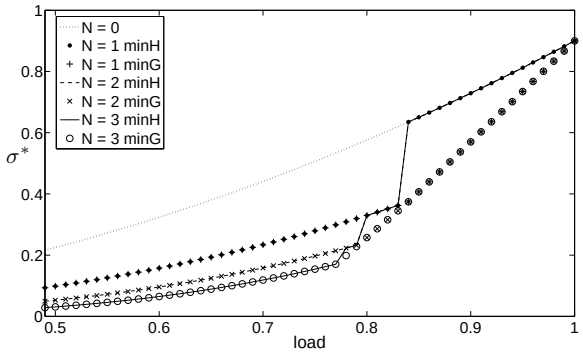

Figure 6 Comparison of policies for a switch with $D = 9$, geometric arrivals and packet size uniformly distributed between 5 and 15

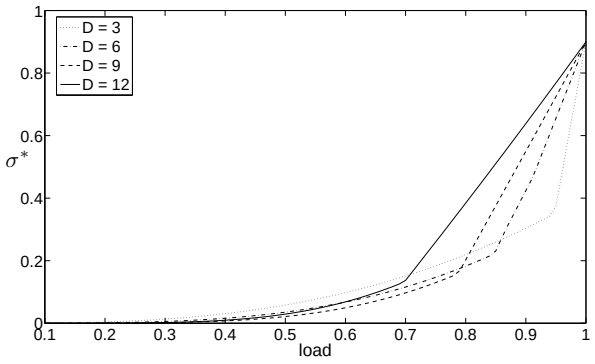

(a) B = 10

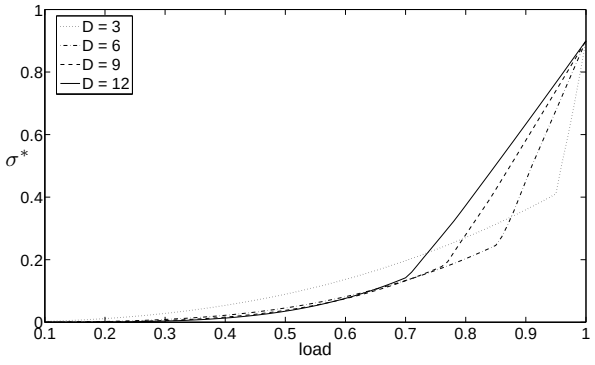

(b) $B = \{5, 15\}$

Figure 7 Effect of the granularity on σ^* for a switch under *minG* policy, geometric arrivals and three FDLs

As stated before, we can determine the value of σ at which the loss probability drops to zero (σ^*) in a single run of the mean field model. In Figure 6 we illustrate how the load affects the value of σ^* for both policies. In this case the IATs follow a geometric distribution, the packet size is uniformly distributed between 5 and 15, and the granularity is 9. As expected, a higher load implies a larger σ^*. Also, for high loads the *minG* policy requires a smaller conversion ratio to achieve zero losses than the *minH* policy. In relation to the number of FDLs, it is clear that the addition of one FDL reduces the value of σ^* for the *minG* policy, but the effect of additional FDLs depends on the load. For high loads, there is no difference in having one or more FDLs, while for middle and low loads the addition of FDLs may reduce the value of σ^*. This behavior can be explained as follows. If the switch has enough converters to prevent losses and the load is one, the probability that a wavelength has a horizon less than ND after **S1** is almost zero in steady state. When the load diminishes, the probability that the horizon is between $(N-1)D$ and $ND-1$ smoothly increases, but for values less than $(N-1)D$ it remains close to zero. To obtain a positive probability of having a wavelength with horizon less than $(N-1)D$ it is necessary for the load to go below a certain threshold, which in Figure 6 corresponds to 0.83. This behavior is independent of the value of N, explaining why the addition of more than one FDL has no effect on the conversion ratio required to achieve zero losses for loads over 0.83 in this scenario. Similar thresholds can be found for the values of the load required to have a positive probability that a wavelength has a horizon between $(i-1)D$ and $iD-1$, for $1 \leq i \leq N$. Hence, for loads above these thresholds having more than $N-i+1$ FDLs has no effect on σ^*. These thresholds coincide with the location of the jumps for the *minH* policy, but under this policy the probability of having a horizon less than ND is zero if $\sigma \geq \sigma^*$ and the load is greater than 0.83. If the load goes below this value, the probability of a horizon between $(N-1)D$ and $ND-1$ suddenly becomes positive and takes similar values to those of the *minG* policy. Therefore, both policies reach a similar σ^* at $\rho = 0.83$, but the *minG* policy does it in a smooth manner while the *minH* policy shows a big reduction in σ^* when the load goes from from 0.84 to 0.83. We may conclude that incorporating one or two FDLs may result in a significant cost reduction, as fewer WCs are needed. However, the results suggest that additional FDLs have little use as they affect the required number of FDLs in a less profound manner, especially for higher loads. The effect of the number of FDLs on σ^* will be discussed again in the next section when looking at the effect of burstiness in the arrival process.

The granularity of the FDL is a parameter with a significant influence on the loss probability for the single-wavelength buffer, as shown in [3, 9, 19]. In Figure 7 we illustrate the effect of the granularity on σ^* for two different packet-size distributions. For clarity reasons the results are only shown for the *minG* policy. The corresponding results for *minH* have a similar behavior as a function of D but include the jumps already shown in previous figures, redounding in a worse per-

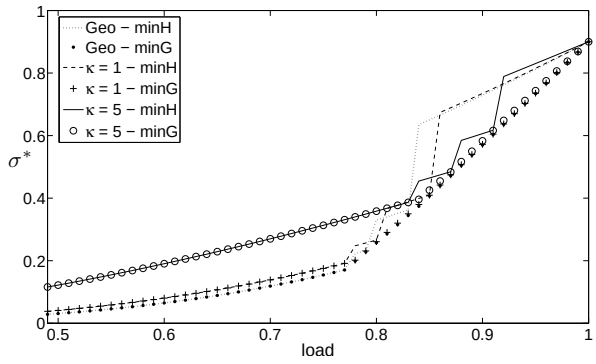

Figure 8 Comparison of policies for a switch with an ON-OFF arrival process, three FDLs, $D = 8$ and packet size equal to $\{5, 15\}$

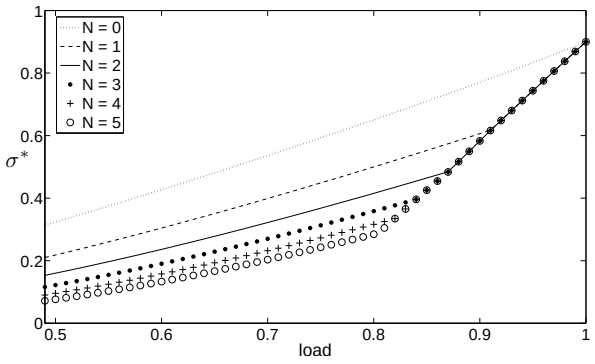

Figure 9 σ^* for a switch under *minG* policy with an ON-OFF arrival process, $D = 9$ and packet size uniformly distributed between 5 and 15

formance. In Figure 7(a) we consider the case where the packet size is fixed and equal to 10 slots. Here we can identify two main regions depending on the load. In the first region, low and mid loads, the optimal granularity is $D = 9$, as has been suggested before for fixed packet size. However, in the second region (high loads) this is no longer the case. At a load around 0.7 there is a pronounced change in the slope of the curve corresponding to $D = 12$. This change rapidly puts this curve above the others, making it the one with the highest requirements in terms of converters. A similar change in slope is suffered by the other curves, in an order that is inversely proportional to the granularity. Therefore, for high loads (in this case above 0.85) a lower granularity means a smaller σ^*. A similar behavior is observed in Figure 7(b), where the packet size can be 5 or 15 with equal probability. In this case however there is almost no difference among the granularities between 6 and 12 along the first region of the load range. This result agrees with previous observations related to the larger set of optimal granularities when the packet size is not fixed. Therefore, among the best possible values for the granularity in the first region, the results just described favor the selection of a small granularity since this requires fewer converters to attain a near-zero loss probability at high loads.

4.3 Effect of the arrival process' burstiness

In this last section we analyze the effect of the burstiness in the arrival process on the minimum conversion ratio to attain zero losses σ^*. This is possible due to the versatility of the arrival process assumed by the model (MAP). In particular we consider an ON-OFF process with two states, where in one state the process generates arrivals with geometric IATs while in the other state no arrivals are generated. This kind of arrival process has been previously used to model the arrival process in an optical switch [15,19,20]. The duration of the ON and OFF periods (the sojourn time of the chain in each state) is geometrically distributed with the mean duration of the OFF periods being κ times that of the ON periods. A simple measure of the burstiness of an arrival process is the ratio between its peak rate and its mean rate [12,17]. For geometric IATs these two rates are equal and the ratio is one. For the ON-OFF process the peak rate is q (the rate of the geometric IATs during the ON periods), the mean rate is $\frac{q}{\kappa+1}$ and the ratio is $\kappa+1$. Therefore, increasing the value of κ while keeping the load fixed increases the burstiness of the process, which is expected since the same number of arrivals will occur in shorter time intervals (ON periods), followed by longer silent (OFF) periods. Figure 8 shows the effect of the burstiness on σ^*. As expected, the increase in the burstiness implies a higher conversion ratio to attain zero losses, with a large difference for the range of mid loads. Here again the *minH* policy shows large jumps when increasing the load, compared to the smooth behavior of the *minG* policy. In this case however we can compare the effect of the allocation policy versus that of the burstiness. From Figure 8 we see that when the load goes above 0.85 the conversion requirements for the *minH* policy under geometric arrivals becomes up to 50% higher than those for the *minG* policy under an ON-OFF arrival process with $\kappa = 5$. This is an important difference in conversion requirements and reveals a significantly worse performance of the *minH* policy even under non-bursty traffic.

We have seen that the *minG* policy requires significantly less conversion resources than its *minH* counterpart. Therefore, we now focus on the *minG* policy and analyze the effect of the number of FDLs under bursty traffic, as illustrated in Figures 9 and 10. Figure 9 shows this effect for a switch where the packet size is uniformly distributed between 5 and 15, and the gran-

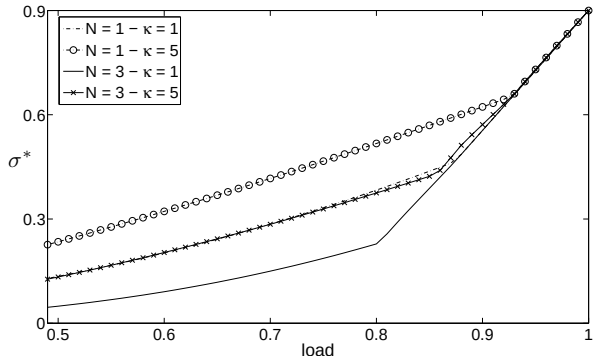

Figure 10 σ^* for a switch under $minG$ policy with an ON-OFF arrival process, $D = 8$ and packet size equal to $\{5, 15\}$

ularity is 9. This is the same scenario as in Figure 6, the only difference being that now the arrival process is ON-OFF with $\kappa = 5$. A first comparison between the two figures yields the expected result of higher σ^* when the arrival process is ON-OFF instead of geometric. A more relevant observation is that the addition of FDLs under bursty traffic has a more significant effect than it had under geometric arrivals. For instance, adding a second or a third FDL produces a larger reduction on σ^* under bursty traffic. As well as in the geometric case, this difference vanishes when the load becomes sufficiently high, giving no advantage in placing additional FDLs. However, the range of loads for which placing additional FDLs makes a difference is larger in the bursty traffic case. An example of this is the placement of the second FDL. While placing a second FDL makes no difference for loads below 0.83 under geometric IATs, it is worth doing so for loads up to 0.92 under bursty traffic. Therefore, the addition of FDLs under bursty traffic not only reduces the conversion requirements in a more significant manner than under non-bursty traffic, but this reduction is valid for a larger range of loads, increasing the value of additional FDLs.

We conclude by taking a final look at the effect of the FDLs on the minimum conversion ratio to attain zero losses. From Figure 10 it is evident that increasing the number of FDLs, in this case from one to three, has a significant effect on reducing σ^*. Furthermore, this reduction is as strong as to make the conversion requirements for the case with $N = 1$ and $\kappa = 1$ similar to those of the case with $N = 3$ and $\kappa = 5$. Consequently, it is possible, at least in part, to compensate the effect of the burstiness on σ^* by including additional FDLs, supporting a switching solution that combines both conversion and buffering resources to resolve contention, especially under bursty traffic.

References

1. Akar, N., Karasan, E., Dogan, K.: Wavelength converter sharing in asynchronous optical packet/burst switching: an exact blocking analysis for markovian arrivals. IEEE J. Sel. Areas Commun. **24**, 69–80 (2006)
2. Akar, N., Karasan, E., Rafaelli, C.: Fixed point analysis of limited range share per node wavelength conversion in asynchronous optical packet switching systems. Photonic Network Communications **18**, 255–263 (2009)
3. Callegati, F.: Approximate modeling of optical buffers for variable length packets. Photonic Network Communications **3**, 383–390 (2001)
4. Callegati, F., Cerroni, W., Corazza, G., Develder, C., Pickavet, M., Demeester, P.: Scheduling algorithms for a slotted packet switch with either fixed or variable lengths packets. Photonic Network Communications **8**, 163–176 (2004)
5. Callegati, F., Cerroni, W., Rafaelli, C., Zaffoni, P.: Wavelength and time domain exploitation for QoS management in optical packet switches. Computer Networks **44**, 569–582 (2004)
6. Dogan, K., Gunulay, Y., Akar, N.: A comparative study of limited range wavelength conversion policies for asynchronous optical packet switching. Journal of Optical Networking **6**, 134–145 (2007)
7. Gauger, C.M.: Optimized combination of converter pools and FDL buffers for contention resolution in optical burst switching. Photonic Network Communications **8**, 139–148 (2004)
8. Latouche, G., Ramaswami, V.: Introduction to Matrix Analytic Methods in Stochastic Modeling. ASA-SIAM Series on Statistics and Applied Probability. SIAM, Philadelphia, PA (1999)
9. Laevens, K., Moeneclaey, M., Bruneel, H.: Queueing analysis of a single-wavelength fiber-delay-line buffer. Telecommunication Systems **31**, 259–287 (2006)
10. Lambert, J., Van Houdt, B., Blondia, C.: Queues with correlated inter-arrival and service times and its application to optical buffers. Stochastic Models **22**(2), 233–251 (2006)
11. Le Boudec, J., McDonald, D., Mundinger, J.: A generic mean field convergence result for systems of interacting objects. In: Proc. 4th Int. Conf. on the Quantitative Evaluation of SysTems (QEST 2007), pp. 3–15. Edinburgh, UK (2007)
12. Michiel, H., Laevens, K.: Teletraffic engineering in a broadband era. In: Proceedings of the IEEE, vol. 85, pp. 2007–2033 (1997)
13. Pérez, J.F., Van Houdt, B.: Dimensioning an OBS switch with partial wavelength conversion and fiber delay lines via a mean field model. In: Proceedings of the IEEE Infocom 2009, Rio de Janeiro, Brazil, pp. 2651 – 2655 (2009)
14. Pérez, J.F., Van Houdt, B.: Wavelength allocation in an optical switch with a fiber delay line buffer and limited-range wavelength conversion. Telecommunication Systems **41**(1), 37–49 (2009)
15. Puttasubbapa, V., Perros, H.: Performance analysis of limited-range wavelength conversion in an OBS switch. Telecommunication Systems **31**, 227–246 (2006)
16. Qiao, J., Yoo, M.: Optical burst switching: A new paradigm for an optical Internet. Journal of High-Speed Networks **8**, 69–84 (1999)
17. Robertazzi, T.: Computer Networks and systems. Springer (2000)
18. Turner, J.: Terabit burst switching. Journal of High-Speed Networks **8**, 3–16 (1999)
19. Van Houdt, B., Laevens, K., Lambert, J., Blondia, C., Bruneel, H.: Channel utilization and loss rate in a single-wavelength Fibre Delay Line (FDL) buffer. In: Proceedings of IEEE Globecom 2004, vol. 3, pp. 1900 – 1906 (2004)
20. Xu, L., Perros, H., Rouskas, G.N.: A queueing network model of an edge optical burst switching node. In: Proceedings of the IEEE Infocom 2003, pp. 2019 – 2029 (2003)