Modeling the Performance of FDL Buffers with Wavelength Conversion

W. Rogiest[1,2], D. Fiems[1], K. Laevens[1] and H. Bruneel[1]
1: SMACS Research Group, TELIN Department, Ghent University (UGent)
Sint-Pietersnieuwstraat 41; B-9000 Ghent, Belgium; wouter.rogiest@UGent.be
2: Bell Labs, Alcatel-Lucent Bell NV
Copernicuslaan 50, B-2018 Antwerpen, Belgium

***Abstract*— In Optical Burst Switching and Optical Packet Switching, contention of bursts (or packets) can be dealt with most effectively through a combination of wavelength conversion and optical buffering. While this is generally accepted in the optical networking community, and validated through simulation, analytic performance results for optical buffers were limited to the single-wavelength case, and the performance gain from wavelength conversion was never traced analytically for general assumptions.**

Quantifying this gain analytically is the scope of the current contribution. Relying on generating functions, we developed a Fiber Delay Line (FDL) buffer model with wavelength conversion, which assumes the buffer located at the output of an optical switch, having access to multiple wavelengths. This document presents our model, validates its accuracy, and compares its output for different burst sizes (fixed or varying), scheduling policies and buffer sizes. Several numerical examples assess the applicability of our approximation, and show that our approach yields accurate results.

***Index Terms*— optical fiber delay lines, wavelength division multiplexing, queueing analysis, stochastic processes, buffers.**

I. Introduction

How are future networks to cope with ever-increasing traffic demands? With new, bandwidth-consuming applications (Video-On-Demand, interactive television over VDSL,...) heading towards us in the very near future, backbone infrastructure is bound to evolve fast, so as to keep up with demand.

State-of-the-art fiber links unleash huge capacities. However, whenever a burst (or packet) is switched to an alternate route, or buffered, it is first converted to the electronic domain, then handled electronically, and finally converted back to light. This becomes infeasible in the near future, for reasons of port count figures, required switching speeds and the related power consumption [1]. All-Optical Packet Switching (OPS) [2] and Optical Burst Switching (OBS) [3] promise to alleviate these problems, and have received considerable attention over the last years, as e.g. witnessed by research projects such as LASOR [1], [4], ATMOS [5], KEOPS [6], DAVID [7] and LASAGNE [8].

A key issue in optical switching is contention resolution. This arises inevitably, even without internal blocking, whenever two or more packets contend for the same outgoing channel at the same time. Possible remedies for this are deflection routing, wavelength conversion and buffering [9], of which the latter two are preferred. Wavelength conversion is an obvious choice, as it capitalizes on multiplexing gain. Buffering poses a challenging problem [4]: as light cannot be stored, it has to be delayed by means of Fiber Delay Lines (FDLs), which cannot realize all delay values exactly, but typically only a multiple of D, the so-called granularity of the system. Clearly, as not all delays are realizable, some capacity is lost on the outgoing channel, because, even when some packets or bursts are present in the buffer, they may not be available yet for transmission. The ensuing periods during which the channel remains unused, are called voids. This is opposed to the situation in a Random Access Memory (RAM) buffer, where bursts are available for access at any (random) instant, and no voids occur. This difference results in performance loss, that can be mitigated only if one chooses the involved design parameters with care.

The objective of this contribution is to evaluate the performance of a heterogeneous contention resolution scheme for optical switching that uses FDL buffering, in addition to wavelength conversion. Performance is measured in terms of the burst loss probability (BLP). The emphasis lies, as in previous work, on FDL buffering. However, as this works best in combination with wavelength conversion, we incorporate both into one model, by adopting the model of [10] without wavelength conversion, to accommodate for multiple wavelengths. In [11], [12], work has already been done on (limited) wavelength conversion, by means of simulations. In contrast, we study performance analytically, by a discrete-time queueing model, offering the benefit of being valid for a broad range of parameter values at once.

In literature, many authors already analyzed single-wavelength optical buffers: the fore-mentioned [10], but also [13]–[21]. Callegati [13] provides an approximate analysis for the optical buffer with exponentially-distributed burst sizes and inter-arrival times. Making use of general notions of queueing theory, an iterative procedure is constructed which yields performance measures for the infinite and finite optical buffer. An exact analysis of the discrete-time optical infinite-capacity buffer with geometric inter-arrival times and general burst-sizes is presented in [14]. In [18], the finite counterpart thereof is treated in an exact manner in both discrete-time and continuous-time. Also in [18], a closed-form solution for waiting times and loss probability is presented in case of upper-bounded burst-size distribution. The extension to generally distributed inter-arrival times and burst-size dis-

tributions is the subject of [10] and [15] for infinite and finite buffers respectively. Infinite optical buffers with batch Markovian arrivals and fixed burst-sizes were investigated in [17]. Finally, we mention that similar models are also analyzed in a completely different context. Lakatos [19], [20] studied the stability of the "optical buffer problem" in the context of airplane landings, where arriving airplanes must orbit the airport in case it is occupied by previously arrived airplanes. In this context, the granularity corresponds to the cycle time needed for one orbit, the airplanes to bursts, the single server to the airport, and the burst length itself to the landing times. A recent contribution in which also the authors have also been involved [21] presents stability results of a generalized setting with general line lengths, general inter-arrival time and burst size distribution.

In contrast to the single-wavelength case, multiple-wavelength buffers have attracted less attention. In [22], Murata and his co-authors are the first to deal with wavelength conversion through analysis, by extending Callegati's approach to multiple wavelengths. However, they only consider a situation where burst lengths are distributed according to a negative-exponential distribution, while our model allows for any burst-size distribution. Further, these authors assume an optical buffer with full wavelength conversion: buffer control can convert incoming bursts to every wavelength and every FDL, given that it is available. As discussed below, we have adopted full wavelength conversion in this paper as well. Optical buffers with limits on the number of wavelength converters [11] or on their wavelength conversion capabilities [23] have been studied by means of simulation, and, most recently, also analytically [24], [25]. In general, limitations on wavelength conversion capabilities can severely degrade performance. However, it is known that this degradation remains small if limitations are not too tight, that is, if there are a sufficient number of wavelength converters [11], or sufficient wavelength conversion capability [23].

The remainder of this paper is organized as follows. In Section II, we explicate the FDL buffer setting. We highlight three different possible scheduling policies: Round-Robin (RR), Join-the-Shortest-Queue (JSQ), and Random (RND). We also discuss the arrival process, and the modeling thereof in terms of probability generating functions (pgf's). In Section III, we analyze the system, in the case where we apply a RR or RND policy. We only foreground the crucial steps, and refer to the appendix for details. In Section IV, the analytic results are backed by simulations, and find that the applied heuristic is very accurate. Further, we present simulation results for JSQ. Considering different burst sizes (fixed or varying) and buffer sizes, we find that, for fixed burst sizes, RR approximates JSQ very well. We also focus on the impact of an increase in the number of wavelengths on loss performance. Finally, conclusions are drawn in Section V.

II. Stochastic Model

A. FDL Buffer

We consider an output buffer of a multiple-wavelength optical switch. Time is discrete, arrivals and departures are synchronized with respect to slot boundaries. Optical bursts arrive at the switch at the different inputs and on different wavelengths and are routed through a switching element to the different outputs and wavelengths. The switch is assumed to be an ideal non-blocking switch and the output buffer under consideration is dedicated to c wavelengths on an output link. If none of the destination output wavelengths are available, the bursts are delayed by sending them through one of the available FDLs which can carry at least c wavelengths as well. There is no FDL buffer recirculation; a burst can traverse at most one FDL and is dropped if no suitable FDL can be found. For a practical implementation of such a switch, the reader is referred to [26].

The FDL buffer can only realize delays that are multiples of the granularity D ($0 \times D$, $1 \times D$, $2 \times D, \ldots$). A buffer of this type is called degenerate, a term coined in [13]. We further define the size N of the buffer as the index of the largest delay line, and assume that also a delayless connection is present in the buffer; this brings the total number of connections in a buffer of size N to $N + 1$.

As we work in discrete-time, D is (as all involved variables) expressed in slots. To have an idea of typical values for the granularity, consider OBS, and a granularity equal to the mean burst size. Considering a 10 Gbps link per wavelength, and burst sizes of 100 kbit on average, we obtain a delay of about 10 μs. This corresponds to about 2 km of fiber for D, 4 km for $2 \times D, \ldots$ Also, notice that the discrete-time assumption is not restrictive. The slot length can be chosen, sufficiently small slot lengths yield accurate approximations for the corresponding continuous model; see [16] where a limit procedure is developed which translates discrete-time results into corresponding continuous-time results.

As mentioned in the previous section, an FDL buffer suffers from performance loss due to voids, that occur whenever the outgoing channel remains unused, while bursts are present in the buffer, but not yet available. Allowing for newly arriving bursts to be scheduled inside such voids can mitigate their impact, but calls for a far more complex scheduling (a void-filling policy) and allows for bursts to overtake one another, i.e., they might arrive out-of-order at the end nodes of the network. We, however, will always assign delays to newly arriving bursts, so as not to overlap with previous ones, and leave after all previously arrived ones.

B. Scheduling Policy

When buffer control detects that a number of bursts, say r ($0 \leq r$), want to be switched to the same output port at the same time, contention does not necessarily arise. Some of the wavelengths, say t ($0 \leq t \leq c$), are available immediately, while the other ones ($c - t$ in number) are already reserved by other bursts (some of them being transmitted, others waiting their turn in the buffer). As such, no bursts have to be buffered if $r \leq t$, and all can be transmitted directly, on separate wavelengths. When $r > t$, t bursts can be transmitted directly, and the others ($r - t$ in number) are queued. As c different wavelengths are available to queue for, a scheduling discipline has to be adopted. What queue to join—each wavelength represents a separate queue—is decided upon by the scheduling

discipline. We consider here three possibilities, that are well-known from classical (i.e., non-optical) buffers.

- Join-The-Shortest-Queue (JSQ): Here, in all circumstances, bursts join the shortest of the c queues. In other words, the burst is routed through the shortest FDL which is available on some wavelength. This policy is known to be optimal for a large class of service time distributions [27]. Of the three, this policy has the best performance in terms of loss and delay, but in return has highest implementation complexity since one needs to keep track of the number of bursts in each queue.
- Round-Robin (RR): Whenever a burst is sent to queue i, the next one is sent to queue $i+1$, and so on until queue c, that is followed by queue 1 again. Hence, the burst is routed to the shortest FDL possible, given the wavelength. Performance is less than for JSQ, but is sometimes comparable, as will be shown below. Further, hardware implementation complexity is low.
- Random (RND): Each burst is sent to one of the channels in random order. Hence, a wavelength is selected at random and the burst is routed to the shortest available FDL for this wavelength. Performance is worst in this case, because the load is not spread intelligently among different queues, and one queue can have an overflow, while another is almost empty. As such, performance is equal to that of c separate single-wavelength buffers, that receive a fraction $1/c$ of the total load. From the implementation point-of-view, this performance weakening occurs when bursts are passed on to the queues in a more or less random order.

Our analysis yields analytical results for RR and RND. Our method transforms the multi-wavelength optical buffer to c identical single-wavelength optical buffers, by studying the arrival process of the bursts at one of the single-wavelength buffers. This approach is not feasible for JSQ, which is known to be very hard to trace analytically, even in the case of classical buffers [28]. Therefore, JSQ performance is obtained by simulation.

C. Bernoulli Arrival Process

We assume that bursts arrive in the system according to a Bernoulli arrival process with parameter p, the probability that there is an arrival in a slot. This process is known to be memoryless — no process state is involved — which significantly simplifies the analysis. Bursts arrive in the system one by one, with inter-arrival times that are distributed geometrically, the geometric distribution being the discrete-time counterpart of the negative exponential distribution. In the following, we no longer consider the arrival process at the entire output port. Instead, we transform the system's arrival process to an arrival process at the level of a single-wavelength buffer. As such, we are able to model RND and RR with a queueing model for a single wavelength. For RR, this comes at the price that the arrival process for the single queue is no longer memoryless as shown below.

- RND: Here, the arrivals in the buffer are selected from a Bernoulli arrival process in a random manner. Therefore, the same type of arrival process occurs at the level of the single-wavelength buffer, and T is also distributed geometrically with parameter p/c. The probability mass function (pmf) and probability generating function (pgf) of the inter-arrival times are given by,

$$\Pr[T_n = k] = \frac{p}{c}\left(1-\frac{p}{c}\right)^{k-1} \quad (k \geq 1)\,,$$

$$T(z) = \frac{pz}{c-(c-p)z}\,.$$

- RR: Here, for every c arrivals in the system, an arrival is selected to join the given buffer. As such, the random variable T is the sum of c different geometrically-distributed random variables. This is the so-called negative-binomial distribution (sometimes referred to as the Pascal distribution), which is the discrete-time counterpart of the Erlang distribution. The pmf and pgf are given by,

$$\Pr[T_n = k] = \binom{k-1}{c-1} p^c\,(1-p)^{k-c} \quad (k \geq c)\,,$$

$$T(z) = \left(\frac{pz}{1-\bar{p}z}\right)^c\,.$$

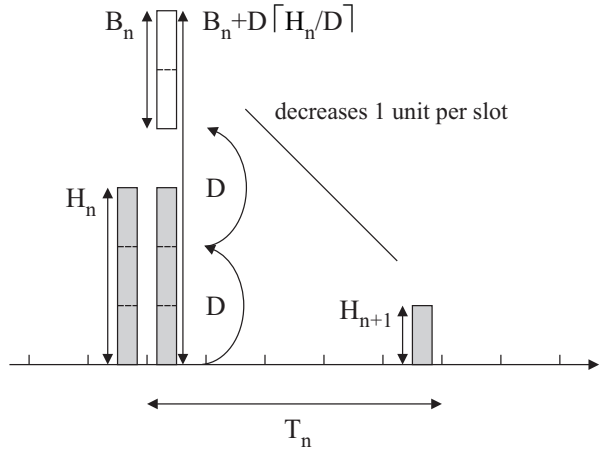

Fig. 1. Evolution of H_n in time. A burst B_n arrives, and has to wait for $D\lceil H_n/D\rceil$, so as to avoid contention with previously arrived bursts.

D. Scheduling Horizon

Focusing on a single wavelength, we now introduce the scheduling horizon which captures the "state" of the single-wavelength buffer. The scheduling horizon at the beginning of a slot is defined as the number of slots till all bursts that are present in the optical buffer at that slot, are transmitted. Now, assume that arriving bursts are numbered and let H_n denote the scheduling horizon upon arrival of the nth burst. Since no void filling is allowed, an arriving burst n has to be buffered for at least H_n slots, the time needed for all previous bursts in that queue to be transmitted. Since there is no FDL for each possible H_n, bursts are scheduled to wait for a time period W_n, that is a multiple of D, and is sufficiently long, i.e., $W_n \geq H_n$. Depending on the size N, the burst is either queued ($W_n \leq ND$) or dropped ($W_n > ND$). W_n is the waiting time

in the given queue, of the nth burst. Mathematically, we have that

$$W_n = D \cdot \left\lceil \frac{H_n}{D} \right\rceil \tag{1}$$

The expression $\lceil x \rceil$ is shorthand for the smallest integer greater than or equal to x. Further, the nth burst has a burst size B_n, that amounts to the time needed for its transmission. The time between its arrival, and the next, is captured by the inter-arrival time T_n. As a result, the evolution of the involved variables can be captured by

$$H_{n+1} = [W_n + B_n - T_n]^+ \,, \tag{2}$$

as illustrated in Fig. 1. The expression $[x]^+$ is shorthand for $\max(0, x)$. Equations (1) and (2) together fully capture the behavior of the single-wavelength optical buffer.

E. Burst sizes

To analyze (1) and (2), we already imposed that the bursts arrive in accordance with a Bernoulli arrival process. Additionally it is imposed that the burst sizes B_n constitute a series of independent and identically distributed random variables. For further use, let $B(z)$ denote their common probability generating function. Although the analysis holds for any burst size distribution, we consider the following distributions for the numerical examples.

- Fixed burst sizes. Bursts thus are distributed according to a deterministic distribution, with parameter B, equal to the mean burst size $\mathrm{E}[B_n]$. The pmf and associated pgf are

$$\Pr[B_n = k] = \delta_{k-B} \quad (k \geq 1)\,, \quad B(z) = z^B \,.$$

  where δ_i denotes the Kronecker delta, that is one if $i = 0$, and zero everywhere else.
- Varying burst sizes. Bursts are distributed according to a geometric distribution, with parameter q, and mean $\mathrm{E}[B_n] = 1/q$. The pmf and corresponding pgf are given by,

$$\Pr[B_n = k] = q\,(1-q)^{k-1} \quad (k \geq 1)\,,$$

$$B(z) = \frac{qz}{1-(1-q)z} \,.$$

III. Analysis

In this section, we analyze the queue associated with a single wavelength, making use of the pgf approach. First, we assume that the FDL buffer size N is infinite. Then, we use a heuristic to move to finite systems, and derive an expression for the burst loss probability in the finite FDL buffer. As the c different queues all display similar behavior, the performance measures are valid for the entire system.

A. Exact Result for Infinite System

Assume that the infinite size optical buffer is stable, this means that the load is sufficiently low such that the scheduling horizon does not grow unbounded. One can show that this is the case if

$$\mathrm{E}\left[\left\lceil \frac{B_n - T_n}{D} \right\rceil\right] < 0\,. \tag{3}$$

Under this condition, the distributions of H_n (and likewise, for other variables involved) converge, for $n \to \infty$, to a unique stochastic equilibrium distribution, that no longer relates to the initial condition of the system. Associated with this distribution is a common pgf $H(z)$. In the Appendix, an exact expression is derived, under the assumption that the inter-arrival times have a common rational pgf. Clearly, this assumption is satisfied for the inter-arrival time distributions for RND and RR.

The expectation in the stability condition (3) is not easy to evaluate. However, the analysis of the optical buffer provides a practical way to obtain stability conditions. We can characterize the case of unbounded growth by

$$\lim_{n\to\infty} \Pr[H_n = 0] = H(0) = 0\,,$$

the expression for $H(0)$ being obtained during the analysis. From this, one defines a maximum tolerable arrival intensity λ_{max}, that puts an upper limit to the arrival intensity λ, defined as $1/\mathrm{E}[T]$. Using the results of the Appendix, it follows that λ_{max} is the solution to the implicit expression (implicit, as p, or, λ, also occurs in the expression for $T(z)$),

$$\frac{1}{\lambda_{max}} = \mathrm{E}[B] + \frac{D-1}{2} + \sum_{k=1}^{D-1} \frac{B(\epsilon_k)T(1/\epsilon_k)}{\epsilon_k - 1} \tag{4}$$

The symbols ϵ_k represent the D different complex D^{th} roots of one, see the Appendix. The solution is thus function of the FDL granularity D, and the (complete) pgf's of both inter-arrival and burst-size distribution, and can be found from (4) with a simple bisection algorithm. Related, we define an equivalent load, that is given by

$$\rho_{eq} = \lambda \cdot \left(\mathrm{E}[B] + \frac{D-1}{2} + \sum_{k=1}^{D-1} \frac{B(\epsilon_k)T(1/\epsilon_k)}{\epsilon_k - 1} \right)$$

Hence, the effect of voids is incorporated into an altered definition of the load. Remark that ρ_{eq} equals one when the arrival intensity reaches λ_{max}, which in general happens for a (classic) load $\rho = \lambda_{max}\,\mathrm{E}[B]$ smaller than one. As mentioned, infinite size optical buffers are thus unstable for lower loads, compared to classical (RAM) buffers.

Given the expression for $H(z)$, moments of the steady-state scheduling horizon are easily obtained by means of the moment-generating property of probability generating functions. For example, the mean scheduling horizon upon arrival of a burst in steady state is given by $H'(1)$. We however mainly focus on the burst loss probabilities which are the subject of the next subsection.

B. Approximate Result for the Finite System

Results up to now are valid for an optical buffer of infinite size. Now, we derive the loss probability for a buffer of finite size, with $N+1$ delay lines. Like we did before in [14], [16] (there referred to as "heuristic B"), we will rely on a heuristic to make this transition, that is defined as

$$BLP \approx (1-\rho_{eq}) \cdot \frac{\Pr[H > N \cdot D]}{1 - \Pr[H > N \cdot D]}$$

Here, H denotes the steady-state scheduling horizon in an infinite optical buffer and ρ_{eq} is the above-mentioned equivalent load. This heuristic is motivated by the similarity in shape of tail probability curves for infinite size buffers and burst loss probability curves for finite capacity buffers. The accuracy of this heuristic is demonstrated in section IV.

The tail probabilities $\Pr[H > N \cdot D]$ can be obtained in many different ways; here, we choose to apply a dominant pole approximation. This approximation exploits the fact that the asymptotic behavior of curves closely relates to the poles of the probability generating functions with the smallest radius, the so-called dominant poles. We first impose some (rather general) restrictions on the distribution of $B(z)$: assume that B has a rational pgf. We adopt this convention from here on, and note that it poses no problem for our application. Remark, however, that the model up to now is valid for any $B(z)$ with $\mathrm{E}[B] < \infty$. Under this additional assumption, tail probabilities have a quasi-geometrical tail decay, with decay rate z_0, as

$$\Pr[H > N \cdot D] \approx \frac{cst}{z_0^{N \cdot D+1}}$$

with z_0 the (single) dominant pole along the positive real axis. Its value can easily be determined by means of a simple bisection algorithm. The constant follows from the application of residue theory and is, in its final form, given by

$$cst = -\frac{1}{z_0} \lim_{z \to z_0} \left(W(z) \cdot (z - z_0)\right) \frac{D}{z_0^D - 1}$$

The limit in the above can easily be calculated explicitly. The function $W(z)$ relates to $H(z)$, see Appendix, and both have the same dominant poles. There are D of them, of the form $|z_0| \cdot \epsilon_k$, $k = 0 \ldots D-1$. How to obtain them is explained in more detail in [10].

Applying the heuristic to the result of the appendix, we are able to obtain numerical results for both RR and RND.

IV. Numerical Comparison

At this point, we evaluate the performance of different multi-wavelength buffer settings. Our scope here is twofold. On one hand, we want to assess the accuracy of our analytic results for RR and RND. On the other hand, we study the relation between RR and JSQ, that will show to be intimate for fixed-sized bursts. To do this, we consider all three scheduling disciplines, varying and fixed burst sizes, and two buffer sizes. In the remainder, the slot length is fixed to 100 ns.

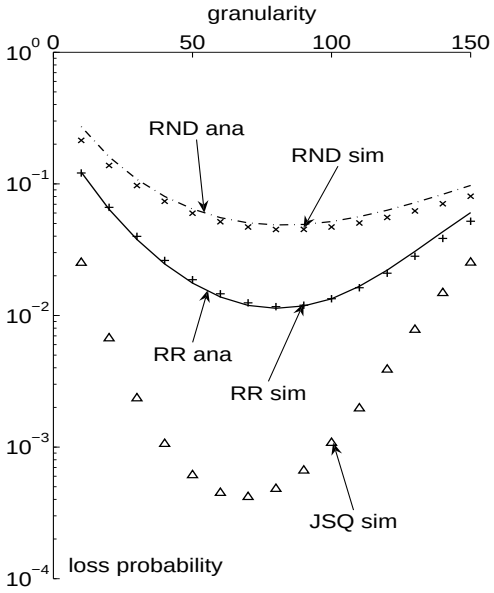

Fig. 2. For geometrically-distributed burst sizes, the three wavelength assignment algorithms lead to much difference in terms of performance: JSQ outperforms RR, that in its turn outperforms RND. This figure was obtained for $\mathrm{E}[B] = 100$ slots, $\rho = 60\%$, $c = 4$, and $N = 10$.

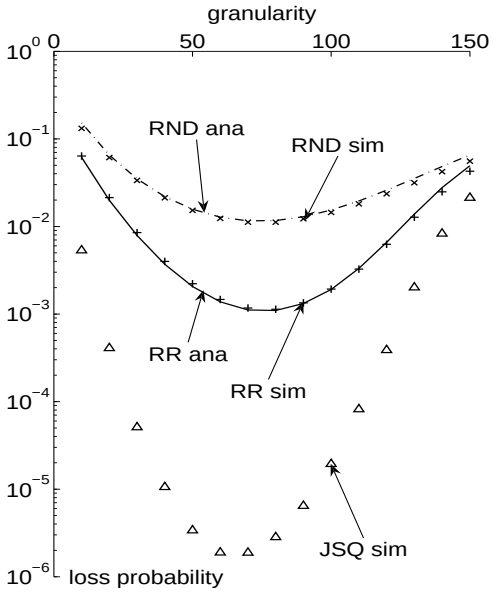

Fig. 3. Identical setting as in Fig. 2, except that here, $N = 20$.

A. Varying Burst Sizes

We first consider varying burst sizes, with a mean length of 100 slots, in Fig. 2 (for buffer size $N = 10$) and Fig. 3 ($N = 20$). Making the same assumptions as in Section II-A (100 kbit burst, 10 Gbps link per wavelength), this corresponds to a mean burst size $\mathrm{E}[B] = 10$ μs, and granularity values ranging from 10 to 150 slots, about 3 km of fiber. The load is fixed to 60%. Simulation results (sim) are calculated for D equal to multiples of 10 slots, and analytic results (ana) for all intermediate integer values. The number of wavelengths is four ($c = 4$).

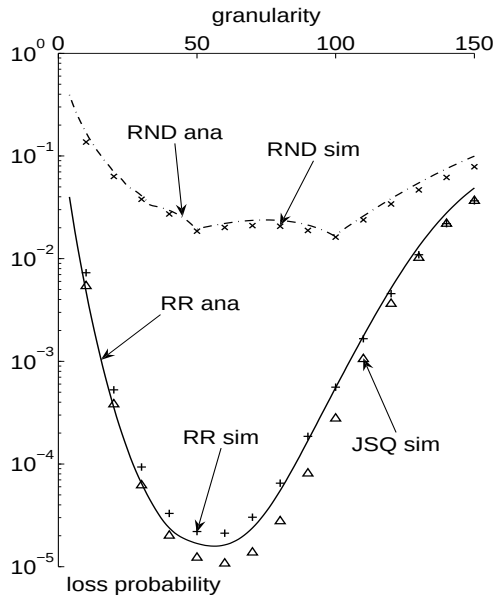

Fig. 4. For deterministic burst sizes, JSQ and RR yield much better performance than RND. However, the difference between JSQ and RR is much less pronounced than in the case of geometrically-distributed burst sizes (Fig. 2 and 3). This figure was obtained for $\mathrm{E}[B] = 100$ slots, $\rho = 60\%$, $c = 4$, and $N = 10$.

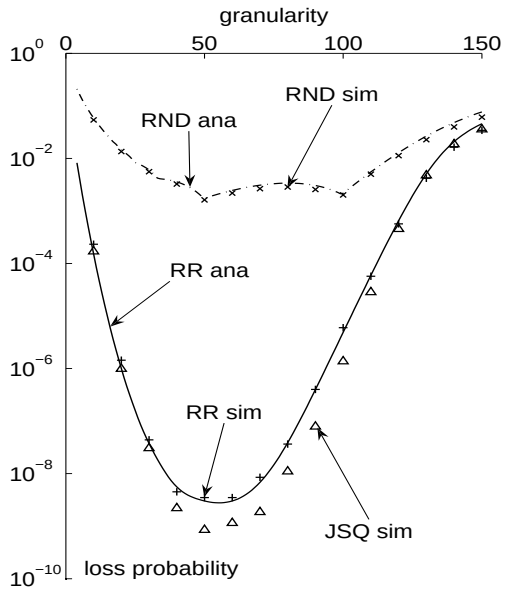

Fig. 5. Identical setting as in Fig. 4, except that here, $N = 20$.

Obviously, the main difference between the figures is in the range of the BLP, and they further display a similar behavior. In both figures, the analytic results for RND and RR both match simulation results very well. This asserts the functioning of our heuristic for varying burst sizes.

From the figure, it is obvious that RND is a policy that is to be avoided, as both RR and JSQ perform significantly better. Comparing JSQ and RR, the simulation shows that JSQ outperforms RR by far, if burst sizes vary. Therefore, it is understandable that JSQ often is the scheduling discipline of practical interest. This can be understood intuitively, if one realizes that the next queue (as selected in RR) seldom is the shortest queue (as selected in JSQ) if burst sizes vary. This is different for fixed burst sizes.

B. Fixed Burst Sizes

For burst sizes fixed to 10 μs (or, 100 slots), we obtain Fig. 4 ($N = 10$) and Fig. 5 ($N = 20$). Again, granularity ranges from 1 to 150 slots, the load is 60%, the number of wavelengths is four, simulation results (sim) are calculated for multiples of 10 slots, and analytic results (ana) for the whole range of D.

Similar to the case of varying burst sizes, the impact of buffer size is only in the range of the BLP, and both figures display a similar behavior. As for the optimal granularity, well-pronounced minima are observed for RND around the burst size and around half of the burst size. Further, simulation results for RND and RR assert the functioning of our heuristic again, now for fixed-sized bursts.

The RND scheduling discipline is again the one to avoid. The main difference with varying bursts, is that the gap in performance between JSQ and RR is really small, and this is true for the whole range of D. Because this is so for classical buffers, it comes not entirely as a surprise. However, since optical buffers differ from classical buffers in a non-trivial way (one can think of the voids, non-existent for classical buffers), these results show that this is also valid for optical buffers. Thus, the intuition applies that to select the next queue (RR) often comes down to selecting the shortest one (JSQ), if burst sizes are fixed.

With the observation that JSQ resembles RR for fixed-size bursts, two implications go. On one hand, considering also that RR has lower implementation complexity than JSQ, a hardware designer can choose implementing RR instead of JSQ. On the other hand, disregarding the choice between RR and JSQ, the analytical model obtained for RR, offers a good approximation for both RR and JSQ.

C. Multiplexing Gain

In Fig. 6 (varying burst sizes) and Fig. 7 (fixed burst sizes), we focus on the impact of wavelength conversion on loss performance. Therefore, we consider FDL buffers with one, two and four wavelengths (c =1, 2 or 4). For a single wavelength, no scheduling discipline is specified (as there is only one channel to queue for), while for two and four wavelengths, we consider both RR and JSQ. For RND, the results are also shown (albeit implicitly), as they are identical to the results for $c = 1$, independent of the number of available wavelengths. As such, RND does not benefit from multiplexing gain in any way.

We assume a buffer size $N = 10$, a mean burst size $\mathrm{E}[B] = 100$ slots (or, 10 μs), and a load of 60%. As we have assessed the accuracy of our heuristic in the above, we only show the analytic results for RND and RR (ana), for integer values of D. (Simulation results for $c = 2$ not included here, display a good match.) Simulation results for JSQ (sim) are calculated for multiples of 10 slots.

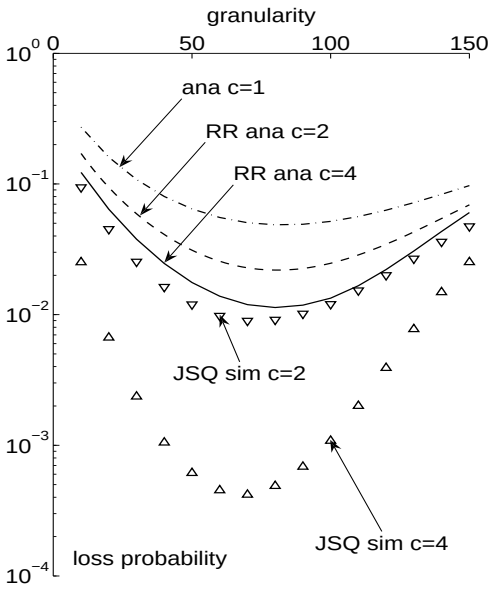

Fig. 6. For geometrically-distributed burst sizes, JSQ performs much better than RR, which in its turn outperforms RND. The difference between JSQ and RR is much larger than in case of deterministic burst sizes, see Fig. 7. This figure was obtained for $\mathrm{E}[B] = 100$ slots, $\rho = 60\%$, and $N = 10$.

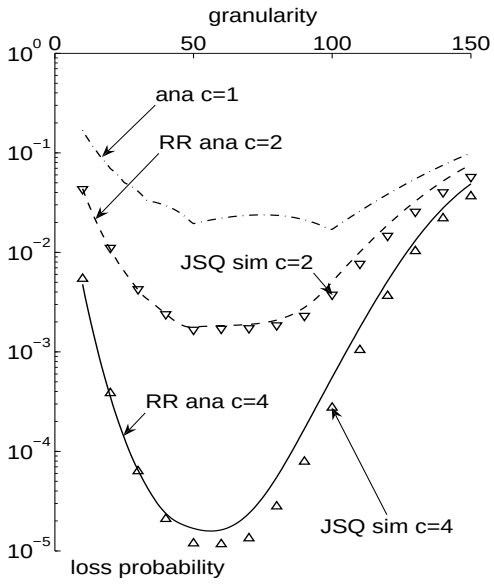

Fig. 7. Identical setting as in Fig. 6, except that here, burst sizes have deterministic distribution. Again, both JSQ and RR perform much better than RND; however, the performance difference between JSQ and RR is much less than in case of geometrically-distributed burst sizes, see Fig. 6. This figure was obtained for $\mathrm{E}[B] = 100$ slots, $\rho = 60\%$, and $N = 10$.

As a reference, we also mention the loss in case wavelength conversion is applied without FDL buffering. For both varying and fixed burst sizes, the BLP is then approximately 0.38 for one server, 0.25 for two, and 0.14 for four servers. These results are obtained by the well-known Erlang-B formula, with 1, 2 and 4 servers respectively. Although this formula by definition only applies in a continuous-time setting, it serves a good approximation in this specific case.

Figure 6 shows loss performance for varying burst sizes. For both RR and JSQ, performance does benefit from multiplexing gain, and the curves lower for increasing c. Results not included here, demonstrate a further decrease in loss probability, when one adds even more wavelengths. In this respect, an FDL buffer behaves similarly to a classical multiplexer.

Clearly, the gain is a lot more pronounced for JSQ than for RR: loss for JSQ lowers significantly for two wavelengths, while RR needs four wavelengths to obtain a comparable lowering. For RR and JSQ, we note how the optimum for the granularity (for which BLP is minimal) lowers a little, but not much, as the number of wavelengths increases. Results not shown here, illustrate that this optimum is more sensitive to the load, than to the number of wavelengths.

Figure 7 displays the burst loss probability for fixed burst sizes. Again, RND does not benefit at all from multiplexing gain. Loss performance for JSQ and RR is further very similar, and both clearly capitalize on multiplexing gain. Here, the optimal granularity for JSQ and RR drops markedly as the number of wavelengths increases.

V. Conclusions

In this contribution, we presented an analytic model for a multi-wavelength optical buffer, and its application for performance assessment. We have considered three different scheduling disciplines: Random (RND), Round-Robin (RR), and Join-the-Shortest-Queue (JSQ). These handle either varying- or fixed-sized bursts. For RND and RR, we developed an analytic model, while for JSQ, we relied on simulation.

In the case of varying burst sizes, we found that performance is matched accurately in the case of RND and RR. Further, we found that the performance of JSQ can in no case be approximated by that of RR, as they differ too much. We also found that the optimal granularity is not much influenced by the number of wavelengths.

For fixed-sized bursts, we found that our model can estimate performance very well for RND and RR. As we found that JSQ performance is close to that of RR, our model is also a good approximation for JSQ. Finally, we also observed that the optimal granularity drops markedly, for RR and JSQ, as the number of wavelengths increases.

The same methodology can also be applied to the more general setting of $GI/G/c$ optical buffers whereby the probability generating function of the inter-arrival times is a rational function. However, it is not clear if the JSQ approximation is sufficiently accurate in this more elaborate case. For this, further research is required. Other future work further includes capturing the performance gap between RR and JSQ for varying burst sizes in a more accurate way.

Appendix

In this appendix, we derive an expression for $H(z)$, the equilibrium pgf of the scheduling horizon H as seen by arrivals. We assume that the system is stable (and can thus reach equilibrium), that $\mathrm{E}[B]$ is finite, and that the pgf of T is rational. The evolution of this random variable is described

by (1) and (2), and illustrated in Fig. 1. There, index n denotes the nth arrival; here, we will no longer write down this index explicitly. The overall solution for H involves tackling two equations separately and combining the results.

A. Two Effects

Equation (2) captures the queueing effect. First of all, we emphasize that $W + B$ and T are statistically independent, which is essential to our current analysis. To solve (2), we rewrite the rational pgf of T as

$$T(z) = \frac{N(z)}{P(z)} = \frac{\sum_{i=0}^{N} n_i z^i}{\sum_{i=0}^{P} p_i z^i} .$$

In other words, we label the numerator as a polynomial $N(z)$ of degree N, and the denominator as a polynomial $P(z)$ of degree P. Explicitly writing all involved pgf's in terms of power series, reordering several finite sums, and rebundling these sums in pgf's, yields

$$H(z) = W(z)B(z)T(z^{-1}) + T^*(1) - T^*(z^{-1}) . \quad (5)$$

Here, $T^*(z)$ is a rational function (not a pgf, however, as $T^*(1) \neq 1$) that has the same denominator as $T(z)$, $P(z)$, but a different numerator $N^*(z)$, $N^*(z) \neq N(z)$, of degree N or less. For details on the analysis, the interested reader is referred to [10]. As we do not need the exact expression of $N^*(z)$ in the following, we move to the second effect.

Equation (1) captures the FDL effect. By some standard z-transform manipulations and the identity $\frac{1}{D}\sum_{k=0}^{D-1} \epsilon_k^m = \delta_{\lceil m/D \rceil - m/D}$, we find [14],

$$W(z) = \sum_{k=0}^{D-1} \frac{1}{D} \frac{z^D - 1}{z\epsilon_k - 1} H(z\epsilon_k) \quad (6)$$

where the symbols ϵ_k denote the D different complex Dth roots of 1, i.e., $\epsilon_k = e^{j2\pi k/D}$, $k = 0 \dots D-1$.

B. Combining the results

Plugging (6) in (5), and solving for $W(z)$ yields,

$$W(z) = \frac{\sum_{k=0}^{D-1} \frac{1}{D} \frac{z^D-1}{z\epsilon_k - 1} \left\{ T^*(1) - T^*((z\epsilon_k)^{-1}) \right\}}{1 - \sum_{k=0}^{D-1} \frac{1}{D} \frac{z^D-1}{z\epsilon_k - 1} B(z\epsilon_k) T((z\epsilon_k)^{-1})} .$$

Here we used the equality $W(z\epsilon_k) = W(z)$, which follows from the observation that waiting times are always integer multiples of D. This is no solution for $W(z)$ yet, as we do not know the exact expression for $T^*(z)$. However, making use of Rouché's theorem, the formula simplifies, and $W(z)$ can be expressed in terms of only $B(z)$ and $T(z)$; more details are given in [10]. As already mentioned, we do not need the explicit form of $T^*(z)$. For $N \leq P$ the final result reads,

$$W(z) = \frac{K^*(z^D - 1)}{\hat{R}(z^D)} \prod_{i=1}^{P-1} \frac{z^D - \beta_i}{1 - \beta_i} \quad (7)$$

with

$$\hat{R}(z^D) = \prod_{j=1}^{J} \left(1 - z^D \gamma_j^D\right)^{m_j} \cdot \left\{ 1 - \sum_{k=0}^{D-1} \frac{1}{D} \frac{z^D - 1}{z\epsilon_k - 1} B(z\epsilon_k) T((z\epsilon_k)^{-1}) \right\} .$$

Here, β_i, $i = 0 \dots P-1$, are the P different zeroes within $|z| \leq 1$ of the numerator in (7), one of which (β_0) equals one. The values γ_j, $j = 1 \dots J$, are the J different poles of $T(z)$, each with associated multiplicity m_j, $m_j \in \mathbb{N}_0$. A similar expression can be obtained for $N < P$.

Finally, the constant K^* is determined by requiring normalization of $W(z)$, i.e., $W(1) = 1$, and results in

$$K^* = \prod_{j=1}^{J} \left(1 - \gamma_j^D\right)^{m_j} \cdot \left\{ \frac{\mathrm{E}[T] - B_{eq}}{D} \right\} ,$$

with B_{eq} the equivalent burst length, defined as

$$B_{eq} = \mathrm{E}[B] + \frac{D-1}{2} + \sum_{k=1}^{D-1} \frac{B(\epsilon_k) T(\epsilon_k^{-1})}{\epsilon_k - 1} . \quad (8)$$

Note that the two last terms in (8) reflect the effect of the voids. Now, it suffices to substitute $W(z)$ in (5), using (7), to obtain an explicit formula for $H(z)$, which was the aim of this Appendix.

ACKNOWLEDGMENT

D. Fiems is a post-doctoral fellow with the Research Foundation, Flanders (F.W.O.-Vlaanderen).

REFERENCES

[1] D. Blumenthal, "Photonic Packet Switching and Optical Label Swapping," *Optical Networks Magazine*, vol. 2, no. 6, pp. 54–65, 2001.

[2] R. Van Caenegem, D. Colle, M. Pickavet, P. Demeester, K. Christodoulopoulos, K. Vlachos et al, "The Design of an All-Optical Packet Switching Network," *IEEE Communications Magazine*, vol. 45(11), pp. 52–61, 2007.

[3] C. Qiao and M. Yoo, "Optical Burst Switching–a New Paradigm for an Optical Internet," *Journal on High-Speed Networks*, vol. 8, pp. 69–84, 1999.

[4] E. Burmeister, D. Blumenthal, and J. Bowers, "A Comparison of Optical Buffering Technologies," *Optical Switching and Networking (OSN)*, vol. 5(1), pp. 10–18, 2008.

[5] F. Masetti, J. Benoit, F. Brillouet, J. Gabriagues, and et al., "High Speed, High Capacity ATM Optical Switches for Future Telecommunication Transport Networks," *IEEE Journal on Selected Areas in Communications*, vol. 14, pp. 979–998, June 1996.

[6] C. Guillemot, M. Renaud, P. Gambini, C. Janz, and et al., "Transparent Optical Packet Switching: the European ACTS KEOPS Project Approach," *IEEE/OSA Journal of Lightwave Technology*, vol. 12, pp. 2117–2134, December 1998.

[7] L. Dittmann, C. Develder, D. Chiaroni, F. Neri, and et al., "The European IST Project DAVID: A Viable Approach Toward Optical Packet Switching," *IEEE Journal on Selected Areas in Communications*, vol. 7, pp. 1026–1040, September 2003.

[8] F. Ramos, E. Kehayas, J. M. Martinez, R. Clavero, J. Marti, L. Stampoulidis et al, "IST-LASAGNE: Towards All-Optical Label Swapping Employing Optical Logic Gates and Optical Flip-Flops," *IEEE Journal of Lightwave Technology*, vol. 23(10), pp. 2993–3011, 2005.

[9] G. Rouskas and H. Perros, "A Tutorial on Optical Networks," *NETWORKING Tutorials 2002*, pp. 155–194, 2002.

[10] W. Rogiest, K. Laevens, J. Walraevens, and H. Bruneel, "Analyzing a Degenerate Buffer with General Inter-Arrival and Service Times in Discrete Time," *Queueing Systems*, vol. 56 (3-4), pp. 203–212, 2007.

[11] C. M. Gauger, "Optimized Combination of Converter Pools and FDL Buffers for Contention Resolution in Optical Burst Switching," *Photonic Network Communications*, vol. 8, no. 2, pp. 139–148, 2004.
[12] T. Zhang, K. Lu, and J. P. Jue, "Architectures and Performance of Fiber Delay Line Buffers in Packet-Based Multifiber Optical Networks," in *Proceedings of Optical Fiber Communication Conference, Technical Digest. OFC/NFOEC*, 2005.
[13] F. Callegati, "Optical Buffers for Variable Length Packets," *IEEE Communications Letters*, vol. 4, no. 9, pp. 292–294, 2000.
[14] K. Laevens and H. Bruneel, "Analysis of a Single-Wavelength Optical Buffer," in *Proceedings of Infocom 2003 (San Francisco)*, 2003.
[15] R. C. Almeida, J. U. Pelegrini, and H. Waldman, "A Generic-Traffic Optical Buffer Modeling for Asynchronous Optical Switching Networks," *IEEE Communications Letters*, vol. 9, no. 2, pp. 175–177, 2005.
[16] W. Rogiest, K. Laevens, D. Fiems, and H. Bruneel, "A Performance Model for an Asynchronous Optical Buffer," *Performance Evaluation*, vol. 62, no. 1-4, pp. 313–330, 2005.
[17] 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.
[18] W. Rogiest, J. Lambert, D. Fiems, B. Van Houdt, H. Bruneel, and C. Blondia, "A unified model for Synchronous and Asynchronous FDL Buffers allowing Closed-Form Solution," *Performance Evaluation*, vol. 66(7), pp. 343–355, 2009.
[19] L. Lakatos, "On a Simple Continuous Cyclic-Waiting Problem," *Annales Univ. Sci. Budapest., Sect. Comp. 14*, pp. 105–113, 1994.
[20] L. Lakatos, "On a Simple Discrete Cyclic-Waiting Queueing Problem," *Journal of Mathematical Sciences (New York)*, vol. 4, no. 92, pp. 4031–4034, 1998.
[21] W. Rogiest, E. Morozov, D. Fiems, K. Laevens, and H. Bruneel, "Stability of Single-Wavelength Optical Buffers," Accepted for publication in European Transactions on Telecommunications (ETT), 2009.
[22] M. Murata and K. Kitayama, "Ultrafast Photonic Label Switch for Asynchronous Packets of Variable Length," in *Proceedings of INFOCOM 2002 (New York)*, 2002.
[23] V. Puttasubbappa and H. Perros, "Performance Analysis of Limited-Range Wavelength Conversion in an OBS Switch," *Telecommunication Systems*, vol. 31, pp. 227–246, March 2006.
[24] J. Pérez and B. Van Houdt, "Wavelength Allocation in an Optical Switch with a Fiber Delay Line Buffer and Limited-Range Wavelength Conversion," *Telecommunication Systems*, vol. 41, no. 1, pp. 37–49, 2009.
[25] J. Pérez and B. Van Houdt, "Dimensioning an OBS switch with Partial Wavelength Conversion and Fiber Delay Lines via a Mean Field Model," in *Proceedings of INFOCOM 2009 (Rio De Janeiro)*, 2009.
[26] Y. Xiong, M. Vandenhoute, and H. Cankaya, "Control Architecture in Optical Burst-Switched WDM Networks," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 10, pp. 1828–1851, 2000.
[27] R. Weber, "On the Optimal Assignment of Customers to Parallel Servers," *Journal of Applied Probability*, vol. 15, no. 2, pp. 406–413, 1978.
[28] H. Lin and C. Raghavendra, "Approximating the Mean Response Time of Parallel Queues with JSQ Policy," *Computers & Operations Research*, vol. 23, no. 8, pp. 733–740, 1996.