

Optical Buffers, Batch Arrivals and Synchronization

W. Romiast, K. Laevens, D. Fiems and H. Bruneel
SMACS Research Group
Ghent University
Sint-Pietersnieuwstraat 41
B-9000 Ghent, Belgium
{womiast,kl,df,hb}@telin.ugent.be

Abstract— We present the analysis of a single-wavelength Fiber-Delay-Line buffer fed by a batch arrival process, and apply the results to study the influence of (internal) synchronization on performance. We rely on generating functions to develop a discrete-time queueing model that, unlike previously obtained results, is able to cope with concurrent arrivals. Without making any special assumptions on the burst or batch size distribution, we derive an expression for the scheduling horizon as seen by arriving batches in a system of infinite size. From that, we are able to derive e.g. a maximum tolerable arrival intensity. For finite systems, we use a heuristic to calculate the burst loss probability which is validated against simulations. Consequently, we employ the model assuming a Poisson batch size distribution, and investigate the influence of synchronization. Whether or not this turns out beneficial is shown to relate strongly to the burst size distribution.

Index Terms—fiber delay lines, loss probability, optical buffers, batch arrivals, synchronization

I. INTRODUCTION

Optical fibers have become the standard carrier for data transport over long distances. Nowadays, major cities are connected by dense wavelength division multiplexing (DWDM) links, enabling transmission capacities well beyond the Tbit/s. Packet switching over these optical links, however, requires that the transmission speeds over the links are matched by equivalent switching capacities in the nodes. As current packet switches perform data processing in the electronic domain, there is a growing discrepancy between channel capacity and switching capacity. Due to increased port count figures, required switching speeds and the associated power consumption, electronical switching might become unfeasible. Solutions like Optical Packet Switching (OPS) and Optical Burst Switching (OBS) aim to process data in the optical domain [1], [2], [3]. Both suffer from output port contention in the switches and therefore equally require contention resolution. One of several possible solutions [4], [5] is optical buffering, which uses pieces of Fiber Delay Lines (FDLs) to delay the light.

Over the last years, several authors have developed analytic models for optical buffers. One of the first was Callegati [6], [7], who studied an equidistant single-wavelength FDL buffer in a continuous-time setting assuming memoryless burst sizes and inter-arrival times. His results were later extended to the multi-wavelength case in [8]. In [9], [10], the loss probability is determined for an M/D/1/1 system with feedback. In [11], [12], [13], Almeida et al. consider more general inter-arrival times and burst sizes. Also, they take a look at more arbitrary sets of FDL lengths (see below). Such is also done in [14]. There, it is shown that the equidistant structure is optimal for lower loads, while the optimum for high loads depends on the specifics of the traffic.

In this paper, we present an analytic model for an optical buffer in a discrete-time setting, allowing for multiple arrivals within the same slot. In this respect, it is an extension of results in e.g. [15], or, more extensively, in [5], where the model was restricted to Bernoulli arrivals from slot to slot. This extension towards batch arrivals is new, and allows for a study of internal synchronization, a process at node level that can be implemented by retiming incoming traffic at the input. In optical network design, synchronization within the network nodes offers several advantages. Although the retiming itself is considered hard to implement (see, for example, [4]), it can reduce the complexity of the node's control logic, and make technologically demanding tasks such as header extraction more feasible. Further, a separate effect occurs within the buffer, and is addressed in this paper. The results presented below show that synchronization can improve the loss performance of an optical buffer spectacularly, especially in the case of fixed-sized bursts. In [16], the influence of synchronization on losses was first mentioned, as the comparison was made between performance in the synchronous versus the asynchronous time setting. However, there, accuracy dropped for long synchronization periods, as the restriction to single arrivals then becomes more and more unrealistic.

Our results now fully incorporate the effects of re-

timing, and are based on the exact synchronization of a Poisson arrival process. It is shown that, in the case of fixed-sized bursts, synchronization can mitigate losses drastically, while in other cases the loss might even increase.

This paper is organized as follows. In Section II, we present the modelling of the optical buffer structure, and focus on an expression that captures the evolution in time of the scheduling horizon. From there, our approach will be to derive results for a system of infinite size ($M = \infty$) in Section III, such as the probability generating function (pgf) of the waiting time of batches, and the maximum tolerable arrival intensity. We then generalize and apply these results in Section IV to finite systems, through the use of heuristics. A separate section then investigates the influence of synchronization on performance. Conclusions are drawn in the last section.

II. MODEL

A. The FDL Buffer

Different types of optical buffers are presented in literature. The model we use here was first presented in [6], and consists of a feed-forward buffer, that has FDLs of equidistant lengths: corresponding with its index i , fiber i has a delay time of size $i \times D$. Here, D denotes the granularity. As our results will show, this is one of the critical design parameters of an FDL buffer. We further define the size M 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 M to $M + 1$. As is done in [14], we call a buffer of this type an equidistant FDL buffer.

This buffer is now considered as part of a network node, located at an output port. This output port is associated with a single outgoing channel. Whenever two or more bursts are switched to this channel, and contend for it at the same time, all but one have to be buffered. The FDL buffer cannot delay incoming bursts for an arbitrary period of time, but only for multiples of the granularity D . Each incoming burst is routed to the shortest of these FDLs such that the burst will not overlap on departure with bursts from the other FDLs. If such an FDL cannot be found, the burst is dropped. Each burst travels through its assigned FDL only once, and several bursts might be travelling through a single FDL at the same time (without overlapping, however). Typically, bursts are delayed for more time than strictly needed. This extra delay results in so-called voids, i.e., periods during which the output channel remains unused, despite the fact that the system is not empty. This leads

to an under-utilization of the outgoing channel, leading e.g. to increased loss.

B. Scheduling Horizon




Fig. 1. Evolution of the scheduling horizon.

We start out with the system of infinite buffer size. We consider arriving batches, which we number in the order of their arrival. We define the scheduling horizon as the time at which all bursts will have left the buffer, and the system will become empty again. Note that, due to the occurrence of voids, the scheduling horizon is in general larger than the unfinished work, and thus incorporates the effects of granularity. The variable H_k is derived thereof, and is defined as the scheduling horizon as seen by the k th arriving batch, just before its arrival. The arriving batch contains a number of bursts N_k ($N_k \geq 1$), each of which has an associated burst length $B_{k,i}$, $i = 1 \dots N_k$, which equals the time needed for its transmission. The time between the k th batch arrival and the next is captured by the batch inter-arrival time, denoted by T_k . The evolution of these variables is described by

$$H_{k+1} = \left[D \left\lceil \frac{H_k}{D} \right\rceil + \sum_{i=1}^{N_k-1} D \left\lceil \frac{B_{k,i}}{D} \right\rceil + B_{k,N} - T_k \right]^+ \quad (1)$$

as illustrated in Figure 1. The expression $\lceil x \rceil$ is equal to the smallest integer greater than or equal to x . The notation $[x]^+$ is shorthand for $\max(x, 0)$. When the k th batch sees a non-zero scheduling horizon H_k upon arrival, the first burst of the batch will have to be delayed for at least that amount to avoid contention. As the optical buffer can only realize delays that are a multiple of D , we obtain that the first burst has to be delayed for an amount $D \lceil H_k/D \rceil$. The second burst of the batch (in case there is one, i.e., $N_k > 1$) has to be delayed too, in order for it not to overlap with the first. Again, only multiples of D are realizable, so that the second burst needs an additional amount of delay $D \lceil B_{k,1}/D \rceil$. Applying the same reasoning to possible other bursts in the batch, we obtain that delaying and transmitting batch k pushes the scheduling horizon (just after arrival)

to $D \lceil H_k/D \rceil + \sum_{i=1}^{N_k-1} D \lceil B_{k,i}/D \rceil + B_{k,N}$. Taking into account the batch inter-arrival times T_k , and the possibility that the system becomes empty in between arrivals, one arrives at (1). This is illustrated in Figure 1, where the k th batch carries two bursts. For the given setting, with $D = 2$, $H_k = 3$, $B_{k,1} = 3$, $B_{k,2} = 2$ and $T_k = 3$, one can verify that indeed $H_{k+1} = 7$, namely $H_{k+1} = [2\lceil 3/2 \rceil + 2\lceil 3/2 \rceil + 2 - 3]^+ = [4 + 4 + 2 - 3]^+ = 7$.

C. Traffic Model

To analyze (1), we need to impose restrictions on the involved variables T_k , N_k and $B_{k,i}$. We assume the T_k to form a sequence of iid (independent and identically distributed) random variables (rv's), having a common geometrical distribution. The batch sizes N_k and burst sizes $B_{k,i}$ also form a sequence of iid rv's, and can have a general distribution, thus being independent of the index k and, in the case of the burst size, also i . In our analysis, we will use the probability generating function (pgf) of the probability mass function (pmf) of the involved variables. The burst sizes $B_{k,i}$, e.g., have a common distribution of general form with a pgf

$$B(z) = \mathbb{E}[z^{B_{k,i}}] = \sum_{n=1}^{\infty} z^n \Pr[B_{k,i} = n]$$

The inter-arrival times were assumed to have a geometric distribution, implying that the associated pmf is

$$\Pr[T_k = n] = p\bar{p}^{n-1} \quad n \geq 1$$

and its resulting pgf

$$T(z) = \frac{pz}{1 - \bar{p}z}$$

where we adopted the standard notation $\bar{p} = 1 - p$. The mean inter-arrival time is $1/p$. For the expectation of T_k , N_k and $B_{k,i}$, we make the usual assumption that $\mathbb{E}[T_k], \mathbb{E}[N_k], \mathbb{E}[B_{k,i}] < \infty$.

III. ANALYSIS

A. Main Result

We begin by analyzing the infinite system. All derivations assume the system is stable. On this condition, to be discussed later on, the distributions of H_k converge, for $k \rightarrow \infty$, to a unique stochastic equilibrium distribution, independent of the initial system conditions. By H , we will denote a generic random variable following this equilibrium distribution. Likewise, we will drop the index k for other variables involved. This enables us to write (1) as

$$H \leftarrow \left[D \left\lceil \frac{H}{D} \right\rceil + \sum_{i=1}^{N-1} D \left\lceil \frac{B_i}{D} \right\rceil + B_N - T \right]^+ \quad (2)$$

where the arrow denotes convergence in distribution for $k \rightarrow \infty$. Furthermore, we will denote

$$\begin{aligned} H^* &= D \lceil H/D \rceil \\ B_i^* &= D \lceil B_i/D \rceil \end{aligned} \quad (3)$$

To solve (2), we will follow the same approach as in [15]. In the expression, we discern two types of non-linear effects. The first type we call the queueing effect, since it is related to the non-negativeness of the buffer content. The second type of non-linearity is called the FDL effect, since it is related to the finite granularity of the FDLs. It requires the analysis of the expressions in (3). Both effects will first be analyzed separately, their combination will then lead to the overall solution.

The queueing effect can be formulated as

$$H = [G - T]^+$$

where

$$G = H^* + \sum_{i=1}^{N-1} B_i^* + B_N \quad (4)$$

The variable T is the inter-arrival time of batches (and not of distinct bursts, as was the case in previous work). As T is geometrically distributed, and independent of G , the solution in terms of pgf's is a well known result in queueing theory (see, for example, [17]):

$$H(z) = \frac{p}{z - \bar{p}} G(z) + \frac{z - 1}{z - \bar{p}} K \quad (5)$$

with

$$G(z) = H^*(z) \frac{N(B^*(z))}{B^*(z)} B(z) \quad (6)$$

and K a constant that will be determined later on. Note that (6) follows directly from (4), because the random variables involved are independent.

For the FDL effect we find from previous work (see [5] for full detail) that

$$\begin{aligned} H^*(z) &= \sum_{k=0}^{D-1} \frac{1}{D} \frac{z^D - 1}{z\epsilon_k - 1} H(z\epsilon_k) \\ B^*(z) &= \sum_{k=0}^{D-1} \frac{1}{D} \frac{z^D - 1}{z\epsilon_k - 1} B(z\epsilon_k) \end{aligned} \quad (7)$$

where $\epsilon_k = \exp(j2\pi k/D)$, for $k = 0 \dots D-1$ (the D different complex D^{th} roots of unity).

Combining the solutions of the queueing effect (5) and FDL effect (7), we find (after some simplification in the numerator)

$$H^*(z) = \frac{K \cdot \left[\frac{\bar{p}^{D-1}(z^D - 1)}{z^D - \bar{p}^D} \right]}{1 - \left[\sum_{k=0}^{D-1} \frac{1}{D} \frac{z^D - 1}{z\epsilon_k - 1} \frac{pB(z\epsilon_k)}{z\epsilon_k - \bar{p}} \right] \cdot \frac{N(B^*(z))}{B^*(z)}}$$

This expression is the generating function of H^* , which is the waiting time of the first burst of an arriving batch. We further remark that this expression is similar to the formula in [15]. The difference is captured within the factor $N(B^*(z))/B^*(z)$, which is the pgf of $\sum_{i=1}^{N-1} B^*$. This sum is associated with the total amount of work brought about by all bursts that do not arrive last. In the case of single burst arrivals, $N(z) = z$, and this factor simplifies to 1, and we obtain again the expression in [15].

Next task is to determine the constant K occurring in the formula for $H^*(z)$. This can be done by demanding that $H^*(1)$ be one, as prescribed by the normalization condition. By applying the rules of de l'Hôpital, we find

$$K = \frac{1 - \bar{p}^D}{D\bar{p}^{D-1}} \cdot \left(\frac{1}{\bar{p}} + \sum_{k=1}^{D-1} \frac{B(\epsilon_k)}{\epsilon_k - \bar{p}} - N'(1) \cdot \left[B'(1) + \frac{D-1}{2} + \sum_{k=1}^{D-1} \frac{B(\epsilon_k)}{\epsilon_k - 1} \right] \right) \quad (8)$$

Now, the expressions for $H(z)$ and $H^*(z)$ can be obtained as a combination of the expressions above.

B. Further Derivations

The probability of finding the system empty upon arrival, relates to the constant K , as

$$H(0) = \lim_{k \rightarrow \infty} \Pr[H_k = 0] = \frac{K}{\bar{p}}$$

Demanding the infinite system to be stable, is equivalent to demanding this probability not to be zero. The condition $H(0) > 0$ defines a maximum tolerable arrival intensity p_{max} , which is the solution of

$$\frac{1}{p_{max}} = N'(1) \cdot \left[B'(1) + \frac{D-1}{2} + \sum_{k=1}^{D-1} \frac{B(\epsilon_k)}{\epsilon_k - 1} \right] - \sum_{k=1}^{D-1} \frac{B(\epsilon_k)}{\epsilon_k - \bar{p}_{max}} \quad (9)$$

We note that this only involves the mean of the batch size distribution $E[N] = N'(1)$, and any type of batch size distribution with the same mean thus produces the same p_{max} .

Presently we can also define an equivalent load,

$$\rho_{eq} = p \cdot \left\{ N'(1) \cdot \left[B'(1) + \frac{D-1}{2} + \sum_{k=1}^{D-1} \frac{B(\epsilon_k)}{\epsilon_k - 1} \right] - \sum_{k=1}^{D-1} \frac{B(\epsilon_k)}{\epsilon_k - \bar{p}} \right\}$$

that incorporates the effects of voids, and is 100% when $p = p_{max}$.

IV. HEURISTICS FOR THE BLP IN FINITE SYSTEMS

In this section, heuristics are derived for the finite system. They involve the tail probabilities of H in the infinite system, for which a dominant pole approximation is developed first.

A. Dominant Pole Approximation

For the dominant pole approximation, we start from the relation

$$\Pr[H > n] = \text{Res} \left[\frac{1}{z^{n+1}} \cdot \frac{H(z) - 1}{z - 1} \right]_{z=0} \quad (10)$$

This relation implies the calculation of the residu at $z = 0$. At this point, we need further assumptions on $H(z)$ to proceed. If we assume that $H(z)$ has no singularities other than isolated poles, we are able to move from (10) to

$$\Pr[H > n] = - \sum_l \text{Res} \left[\frac{1}{z^{n+1}} \cdot \frac{H(z) - 1}{z - 1} \right]_{z=z_l}$$

where the summation index l runs over all poles z_l of $H(z)$. This relation can now be approximated as

$$\Pr[H > n] \approx - \sum_k \text{Res} \left[\frac{1}{z^{n+1}} \cdot \frac{H(z) - 1}{z - 1} \right]_{z=z_k} \quad (11)$$

where the summation index k only runs over the poles z_k of $H(z)$ with smallest modulus.

Clearly, given (5), these so-called dominant poles are also the dominant poles of $G(z)$, and, given (6), the dominant poles of the product of $H^*(z)$, $B(z)$ and $N(B^*(z))/B^*(z)$. Therefore, the assumption made on $H(z)$ has implications for $B(z)$ and $N(z)$. A sufficient condition is to assume that $B(z)$ and $N(z)$ have no singularities (if any) other than dominant poles. This assumption includes the broad class of rational generating functions, and also frequently used others, such as the pgf of the Poisson distribution. Excluded is the class of heavy-tailed distributions. We note that the assumption allows us to apply the dominant pole approximation but that, even for heavy-tailed distributions, all additional derivations from the previous sections remain valid.

With these additional assumptions, we are in the position to use (11), retaining the dominant poles of $H^*(z)$. Their number is D , each has multiplicity one, and is of the form $z_k = z_0 \epsilon_k$ ($k = 0 \dots D-1$), with z_0 the positive real one. Applying this allows to write the

approximate relation as

$$\begin{aligned} \Pr[H > n] &\approx -\frac{\lim_{z \rightarrow z_0} [H^*(z)(z - z_0)]}{z_0^{n+1}} \\ &\cdot \sum_{k=0}^{D-1} \frac{1}{\epsilon_k^n} \frac{B(z_0 \epsilon_k)}{z_0 \epsilon_k - 1} \frac{p}{z_0 \epsilon_k - \bar{p}} \frac{N(B^*(z_0))}{B^*(z_0)} \\ &= \frac{C^*(n)}{z_0^n} \end{aligned} \quad (12)$$

where we introduced the notation $C^*(n)$ to lay emphasis on the quasi-geometrical tail decay, with decay rate z_0 . The function $C^*(n)$ is periodical, i.e. $C^*(n + m \cdot D) = C^*(n)$ ($m = 0, 1, 2, \dots$). The values that will be of interest here are

$$C^*(M \cdot D) = -\frac{1}{z_0} \lim_{z \rightarrow z_0} [H^*(z)(z - z_0)] \frac{D}{z_0^D - 1}$$

which can be found from (12), using the fact that z_0 is a pole of H^* , i.e.,

$$1 - \left[\sum_{k=0}^{D-1} \frac{1}{D} \frac{z_0^D - 1}{z_0 \epsilon_k - 1} \frac{p B(z_0 \epsilon_k)}{z_0 \epsilon_k - \bar{p}} \right] \cdot \frac{N(B^*(z_0))}{B^*(z_0)} = 0$$

Determining the limits involved, provides us with an approximative expression for $\Pr[H > n]$. This result we then use to estimate the BLP (burst loss probability), through

$$BLP \approx \frac{1 - \rho_{eq}}{\rho_{eq}} \cdot \frac{\Pr[H_\infty > M \cdot D]}{1 - \Pr[H_\infty > M \cdot D]}$$

which was also applied in [15]. This heuristic, denoted heuristic A, performs well for high losses, while a related one,

$$BLP \approx (1 - \rho_{eq}) \cdot \frac{\Pr[H_\infty > M \cdot D]}{1 - \Pr[H_\infty > M \cdot D]}$$

denoted heuristic B, performs better in case of low BLP.

B. Numerical Examples

For ease of notation, we introduce an additional random variable A , defined as the number of arrivals in a random slot, which is perhaps a more natural way of describing the arrival process in discrete-time. Of course, A , T and N are related. One can easily show that $\Pr[A = 0] = \bar{p}$, and $\Pr[A = n | A > 0] = \Pr[N = n]$, $n \geq 1$, and $A(z) = \bar{p} + p \cdot N(z)$, so that $A(z)$ determines $T(z)$ and $N(z)$.

We now apply the obtained results to special cases of the burst and batch size distribution. In Figure 2, the BLP as a function of granularity is considered, for a deterministic burst size distribution, with value $E[B] =$




Fig. 2. Burst loss probability for Poisson batch size distribution.

$B = 50$. The batch size distribution is chosen Poisson. In that case, A is distributed as

$$\Pr[A = k] = e^{-\lambda} \frac{\lambda^k}{k!}, \quad k = 0, 1, 2, \dots \quad (13)$$

where the parameter λ denotes the average number of arrivals per slot. Further, the load was fixed to $\rho = 60\%$. For buffer sizes $M = 5$ and $M = 20$, results from simulation, heuristic A and heuristic B are compared. As was remarked before in [15], heuristic B performs better than A around the optima, and attains a very high accuracy for large buffer sizes. Heuristic A performs better than B only for high values of the granularity. The curves differ only little from the case of Bernoulli arrivals, as the probability of observing more than a single arrival, $\Pr[A > 1]$, is rather small. Results for other values of the load, not included here, reveal similar plots. The accuracy of heuristic B remains high, especially when the BLP drops below 10^{-2} .

To further assess the accuracy of our heuristics, we also take a look at a more exotic case of a deterministic batch size distribution, with $E[N] = N = 2$. Figure 3 shows curves of the BLP as a function of granularity, again for a fixed load of $\rho = 60\%$, with buffer sizes $M = 5$ and $M = 20$. The match between heuristic B and simulations is again very good. As in the case of the single arrivals, local optima occur, but, here, performance is even more sensitive to the value of D .

Having sufficient confidence in the accuracy of the heuristics, we now move on to the study of synchronization.




Fig. 3. Burst loss probability for deterministic batch size distribution ($N = 2$).

V. SYNCHRONIZATION

A. Imposing a Slotted Structure




Fig. 4. Synchronization effects on traffic.

By synchronization we intend the transformation from an asynchronous, i.e., unslotted setting, to a setting with a slot length Δ , at node level. This means that, within a network, this synchronization uses a local clock, and only has impact on traffic within the node.

The process is illustrated in Figure 4. In the initial setting, depicted on the upper time axis, arrivals can occur at the input at any instant, and are assumed to happen one at a time. In the synchronized environment, depicted on the lower time axis, bursts are forced to arrive at slot boundaries by imposing a small delay, which is uniformly distributed between 0 and Δ . Due to this retiming effect, several bursts can now arrive at the same slot boundary, i.e., they can arrive in batches.

For the burst sizes, a different effect comes about. Once synchronized, traffic patterns are captured using discrete variables, that express an integer number of time slots involved. Burst sizes are thus virtually increased to

an integer multiple of slots. This effect, that adds to the burst size an amount ranging from 0 to Δ , is denoted padding, as indicated in gray in Figure 4.

B. Poisson Batch Sizes and Padded Burst Sizes

With this definition of synchronization, we can now look at the impact on the number of arrivals per slot A , and the batch size B . The number of arrivals A per slot in a synchronized setting is derived from the number of arrivals A that can be expected to arrive in the asynchronous setting, in a time period of length Δ . In general, this derivation is a non-trivial one, and it is only in some special cases, such as a Poisson arrival process, that synchronization results in a number of arrivals that is independent and identically distributed from slot to slot. In that case, A has a Poisson distribution, with parameter λ , and pmf as in (13). We find that A remains Poisson distributed regardless of the slot length, with parameter λ scaling linearly with this slot length. This allows us to study the impact of synchronization within the single framework of Poisson batch arrivals. We note that this is not possible if A would have a Bernoulli distribution, as a Bernoulli arrival process is not the result of a synchronization process. For that reason, the comparison made in [16] became less and less accurate with increasing Δ .

As for the padding of burst sizes, we remark that no padding takes place, whenever the burst size is an integer multiple of the slot length. In other situations, padding can account for serious loss, as it increases the mean value of the burst size, and thus the load. Instances of both situations will be given in the next paragraph.

C. Numerical Examples

In this section, we take a look at two synchronization settings. Both assume a Poisson arrival process in the asynchronous setting, that is then synchronized up to a varying slot length Δ . As mentioned above, the modelling thereof implies using Poisson batch sizes and padded burst sizes. The load was fixed to $\rho = 60\%$, with a buffer size $M = 20$. Recall that this load applies before any synchronization is done. Based on the results in paragraph IV-B, we will only show heuristic B. (Simulation results not included here, again showed that this heuristic was accurate.)

In Figure 5, we present the BLP for the case of geometric burst sizes, as a function of the granularity. The mean burst size before synchronization is $E[B] = 50\mu s$, which corresponds to approximately 60 KiB. (A rather high value like this one might occur in OBS, where




Fig. 5. Influence of synchronization on burst loss probability, for geometric burst size distribution.

bursts are typically large, as they consist of an aggregation of packets gathered at the edge nodes.) Different slot lengths, ranging from $1\mu s$ to $50\mu s$, are applied, and the mean burst size increases due to padding. Clearly, synchronization does not have a large impact on burst losses, and the curves overlap almost completely. Only for large slot lengths ($\Delta = 25\mu s$ and $\Delta = 50\mu s$), the BLP alters a little bit. When considering the effects of retiming and padding separately, we found that, while retiming actually better performance, padding annihilates this benefit.

In Figure 6, a similar setting is assumed, now with a deterministic burst size of $B = E[B] = 50\mu s$. Clearly, synchronization now has a huge impact on performance, especially for large slot lengths. If we consider the loss for $D = B = 50\mu s$, we find that losses are reduced more than 10^6 times. Looking for the cause of this drastic reduction, we find that, due to the good choice of slot lengths, burst sizes could always be expressed as a multiple of slots. Therefore, no padding comes about, and B is $50\mu s$ in both the asynchronous and synchronized setting. On the other hand, the benefits of retiming are bigger than in the case of a geometric burst size distribution, and, for a slot length $\Delta = D = B = 50\mu s$, voids are even cancelled out completely. Then, the optical buffer's loss performance is equal to that of a RAM (Random Access Memory) buffer of the same size.

Taking into account both examples, we conclude that




Fig. 6. Influence of synchronization on burst loss probability, for deterministic burst size distribution.

synchronization can lower losses significantly if the burst size is fixed. In that case, we consider retiming a viable and effective way to mitigate losses drastically, and cost and effectiveness should be weighed to decide whether or not to implement this type of retiming. Further, synchronization is found to be of little merit (with respect to buffer performance) if burst sizes are distributed exponentially. Although not investigated upon explicitly, the intermediary case where only a limited number of burst sizes plays a role is expected to also allow for a bettering in loss performance. This benefit is expected to be not as drastic as in the case of fixed-sized bursts, yet still significant, if the synchronization slot length is chosen well.

VI. CONCLUSIONS

Our analysis allowed to model an optical buffer fed by batch arrivals. By constructing the pgf of the scheduling horizon, we derived several measures of interest, such as the maximum tolerable arrival intensity in an infinite system, valid for general batch and burst sizes. Using heuristics, we obtained an estimate of the burst loss probability in a finite system. Its accuracy was illustrated in two cases, where it showed to match simulation results well.

Further, it was shown that synchronization benefits loss performance drastically when burst sizes are fixed. Together with other advantages (e.g., a simplification of the control logic), this might justify the additional cost of implementing retiming within an optical node.

In future work, we intend to study the impact of correlation in the arrival process on system performance. This research could then reveal the impact of synchronization within this context. The case where multiple channels serve a shared FDL buffer is also of great practical interest. As the analysis will run into the same mathematical difficulties as that of conventional multi-server queues, this promises to be a very challenging problem.

REFERENCES

- [1] M. Listanti, V. Eramo and R. Sabella, Architectural and technological issues for future optical internet networks, *IEEE Communications Magazine* 38(9), pp. 82-92, September 2000.
- [2] A. Jourdan, D. Chiaroni, E. Dotaro, G.J. Eilenberger, F. Masetti and M. Renaud, The perspective of optical packet switching in IP-dominant backbone and metropolitan networks, *IEEE Communications Magazine* 39(3), pp. 136-141, March 2001.
- [3] T. El-Bawab and J.-D. Shin, Optical packet switching in core networks: between vision and reality, *IEEE Communications Magazine* 40(9), pp. 60-65, September 2002.
- [4] G.N. Rouskas and H.G. Perros, A Tutorial on Optical Networks, *NETWORKING Tutorials* 2002, pp. 155-194.
- [5] K. Laevens and H. Bruneel, Analysis of Fiber Delay Line optical buffers, *Proceedings of the Workshop on Next Generation Networks: Architecture, Protocols, Performance* (Belize, 2005).
- [6] F. Callegati, Optical buffers for variable length packets, *IEEE Communications Letters* 4(9), pp. 292-294, 2000.
- [7] F. Callegati, Approximate modeling of optical buffers for variable length packets, *Photonic Network Communications* 3(4), pp. 383-390, 2001.
- [8] M. Murata and K. Kitayama, Ultrafast photonic label switch for asynchronous packets of variable length, *Proceedings of Infocom 2002* (New York, June 2002).
- [9] A. Rostami and S. S. Chakraborty, On Performance of Optical Buffers with Fixed Length Packets, *Proceedings of WOCN 2005* (Dubai, March 2005).
- [10] A. Rostami and S. S. Chakraborty, On Performance of Optical Buffers With Specific Number of Circulations, *IEEE Photonics Technology Letters*, 17(7), pp. 1570-1572, July 2005.
- [11] H. Waldman, R. Almeida, and J. Pelegri, An infinite granularity bound on the performance of delay-line buffering, *Proceedings of ONDM 2003* (Budapest, February 2003).
- [12] R. Almeida, J. Pelegri, and H. Waldman, Optical buffer modelling for performance evaluation considering any packet inter-arrival time distribution, *Proceedings of ICC 2004* (Paris, June 2004).
- [13] R. Almeida, J. Pelegri, and H. Waldman, A generic-traffic optical buffer modeling for asynchronous optical switching networks. *IEEE Communications Letters*, 9(2), pp. 175-177, February 2005.
- [14] J. Lambert, B. Van Houdt and C. Blondia, Single-wavelength optical buffers: non-equidistant structures and preventive drop mechanisms, *Proceedings of NAEC 2005* (Riva Del Garda, 2005).
- [15] K. Laevens and H. Bruneel, Analysis of a single-wavelength optical buffer, *Proceedings of Infocom 2003* (San Francisco, April 2003).
- [16] W. Rogiest, K. Laevens, D. Fiems and H. Bruneel, Analysis of an Asynchronous Single-Wavelength FDL Buffer, *Proceedings of ITC 19*, pp. 1917-1926 (Beijing, August 2005).
- [17] H. Bruneel and B. Kim, *Discrete-time models for communication systems including ATM*, Kluwer Academic Publishers, 1993.