

Tracing an Optical Buffer's Performance: an Effective Approach

W. Rogiest, D. Fiems, K. Laevens, and H. Bruneel

SMACS Research Group, Ghent University (UGent); Sint-Pietersnieuwstraat 41;
B-9000 Gent; Belgium; Wouter.Rogiest@UGent.be

Abstract. Optical Burst Switching proposes a future-proof alternative to the current electronic switching in the backbone. The involved optical buffers are implemented with a set of Fiber Delay Lines, and suffer serious performance loss, when compared to RAM. Various existing models trace this loss, but either lack generality, accuracy, or effectiveness.

The optical buffer model we constructed is valid for general line lengths and burst sizes. An effective approach allowed to strongly reduce the solution's complexity, while remaining exact. This document presents the key formulas and performance graphs. The obtained model serves as a basic optimization tool, yielding results fast.

keywords: OBS, optical buffer, FDL, fiber delay lines, discrete-time optimization.

1 Introduction

In a decade where bandwidth-consuming web services (with YouTube.com as a fashionable example) and peer-to-peer data exchange are working their way into everyday life, backbone infrastructure needs to be prepared for this ever-increasing bandwidth demand. Although data packets travel through the backbone in the form of light, they are still converted into electricity at every hop, in order to extract header information, buffer them, convert them back to light and transmit them to the next hop. Since this conversion is expected to be the bottleneck in the near future, the research community proposes alternative switching approaches, such as Optical Burst Switching (OBS) [1] and Optical Packet Switching (OPS) [2].

Just like conventional switches, optical switches need to resolve contention. Even without internal blocking, this arises inevitably, whenever two or more bursts (or packets) head for the same output at the same time. In general, wavelength conversion and buffering offer the most viable solutions to date.

Since light cannot be frozen, optical buffering is implemented by sending the bursts through sufficiently long pieces of fiber, often referred to as Fiber Delay Lines (FDLs). An optical buffer is thus a set of FDLs, each with different lengths. This way of implementing, though implementable with off-the-shelf components, has two drawbacks if compared to electronic RAM memory. On the one hand, optical buffers cannot realize all possible delays, which results in performance loss. Incoming bursts can only undergo delays equal to the length of one of the lines. For this reason, some capacity will be lost on the outgoing channel because, even when some bursts (or packets) are

present in the buffer, they may not be available yet for transmission. The ensuing periods during which the outgoing channel remains unused, despite its availability, are referred to as voids. They account for additional waiting time for the burst, and thus lead to increased loss. On the other hand, optical buffers have a large physical size. For typical OBS specifications (10 Gbps link, 100 kbit burst sizes), it takes about 2 km of fiber to apply a delay for the duration of a burst, and the physical size of a buffer thus grows quickly with the number of fibers. As such, an optical buffer typically has a smaller storage capacity than a RAM buffer, which leads to further increased loss. Therefore, quantifying this loss by means of an analytic model, and tuning the design parameters for optimal performance, are main questions in the ongoing research.

A prime contribution in the study of optical buffer performance, is [3]. There, and in [4, 5], Callegati studied an optical buffer with fiber lengths that are a multiple of a basic unit D , that is, $0, D, 2D, 3D, \dots$ (that is, equidistant line lengths). The type of buffer he called a degenerate buffer, while for the basic unit D he coined the term granularity. The analysis explored an approximation based on the iteration of a classic $M/M/1/N$ buffer model, and yields approximate results for memoryless burst sizes. Extending the analysis to general burst sizes, and improving accuracy, [6] presents an accurate performance model for a degenerate buffer in discrete time, that is extended to continuous time in [7]. Murata and his co-authors [8] extend the approach of Callegati to multiwavelength systems, while in [9, 10], a method is provided to account for non-equidistant FDL lengths, that are not necessarily a multiple of the granularity. This approach is examined more thoroughly by Lambert and her co-authors in [11], and also, alternatively, in [12].

While much of the previous work is approximate [3, 6–8], here, exact performance measures are obtained. The model we constructed yields exact formulas, that offer the benefit of being valid for a broad range of parameter values at once. Like in [11], we worked in a discrete-time setting, with arbitrary sets of FDLs, and made no restrictive assumptions on the burst size distribution. Main difference with [11] is that we obtained a more effective modeling approach, by focusing on the waiting time only. In [11], the analysis is based on the evolution of the scheduling horizon, and computational complexity depends on the length of the longest delay line. By ruling out the scheduling horizon as a measure, our model's complexity is independent thereof, and depends solely on the number of delay lines, which is small in all practical cases. Our modeling thus allows to easily compute results for any burst size distribution, and any size, and this also in the case of very long delay lines. Therefore, it allows to approximate the continuous-time case arbitrary close without augmenting the calculation time.

In Sect. 2, we present the FDL buffer setting and modeling variables. The analysis of this model will be given in Sect. 3, proposing a simple Markov chain, and showing how to derive the waiting time and loss characteristics from this. Some numerical examples follow in Sect. 4. The latter show how the optical buffer's performance is strongly impacted by the buffer size, in a way that depends much on the burst size distribution.

2 Stochastic Model

Here, we first focus on the FDL buffer setting, to then move to the assumptions of the traffic model. Consequently, we look at the system equations, that are expressed in terms of the waiting time of an arbitrary burst.

2.1 Optical Buffer Setting

The optical buffer under consideration is a set of $N + 1$ Fiber Delay Lines, one of them with length $\omega_0 = 0$, and N lines with different lengths ω_i , $i = 1 \dots N$. As the set of lines are intended to resolve contention, it is necessary that contending bursts undergo different delays. Therefore, a useful FDL set never contains the same length twice, $\omega_i \neq \omega_j$ for $i \neq j$. For notational convenience, we denote the set of FDL lengths as $\Omega = \{\omega_0, \omega_1, \dots, \omega_N\}$, and we sort the line lengths ascendingly, $\omega_0 < \omega_1 < \dots < \omega_N$. Note that although a common choice is to choose equidistant lengths ($\omega_i = i \times D$, $i = 0 \dots N$), the analysis is done for arbitrary lengths.

This buffer is located at the output of a backbone switch, and is dedicated to a single outgoing wavelength. We consider bursts arriving at the buffer randomly, and possibly overlapping in time. Since there is only one wavelength to queue for, all overlap during transmission should be prevented. By means of a switching matrix that allows to send any burst to any of the $N + 1$ delay lines, buffer control exercises a FIFO (First-In-First-Out) scheduling discipline, and sends every burst to a sufficiently long delay line, so as not to overlap with the one-but-last burst. If such a (sufficiently long) delay line is present, the burst is accepted and enters; if not, the burst is dropped. The periods during which the system can accept any possibly arriving burst, are called available periods; the periods during which the system drops arriving bursts are called unavailable periods.

2.2 Arrivals and Acceptances

We work in a discrete-time setting, which implies that all random variables (rv's) and parameters (such as the ω_i) are expressed in multiples of a (generic) time slot length. We assume the bursts arrive in the system according to a Bernoulli arrival process, with parameter p . This implies that at the most one arrival occurs during a slot, and p gives the probability of such a burst arrival, for an arbitrary slot.

Arriving bursts are either accepted upon arrival (during available periods), or dropped (unavailable periods). We number the bursts in the order at which they arrive, but only assign an index to those bursts that are accepted. With each accepted burst k , we associate an inter-arrival time T_k , that captures the time between the k th arrival and the next burst arrival, being the arrival of (i) burst $k + 1$, if this next burst is accepted or (ii) a burst without number, if this next burst is dropped. For the assumed Bernoulli arrival process, these inter-arrival times T_k form a sequence of identical and independently distributed (iid) random variables (rv's) that have a common geometric distribution, with parameter p

$$\Pr[T_k = n] = t(n) = p \cdot \bar{p}^{n-1}, \quad n \geq 1. \quad (1)$$

Here, \bar{p} is the commonly adopted shorthand for $1 - p$. The inter-acceptance time A_k is defined as the time between the k th acceptance and the $(k + 1)$ th, and is in general larger (and never smaller) than T_k .

To track the system's performance in an easy way, we consider two mutually exclusive events, for an arbitrary accepted burst k .

Next-Accept. The burst that arrives just after the k th burst is accepted, and counted as burst $k + 1$. The inter-acceptance time A_k is identical to T_k , and thus, it follows from (1) that

$$\Pr[A_k = n | \text{Next-Accept}] = t(n) = \bar{p} \cdot p^{n-1}, \quad n \geq 1. \quad (2)$$

Next-Drop. The burst that arrives just after burst k is dropped. Now, the burst following burst k is not assigned an index, and possibly, even more bursts are dropped before burst $k + 1$ is accepted. The inter-acceptance time A_k clearly differs from the inter-arrival time, and has a more complicated probability distribution. Luckily, the latter need not be tracked, and we rely on an additional measure: the reactivation time \tilde{A}_k , defined as the time between the end of the unavailable period following the k th burst, and the arrival of burst $k + 1$. Note that the reactivation time is only relevant if the associated burst k effectively causes burst loss, by driving the system into an unavailable state. Invoking the memoryless nature of the arrival process, the reactivation time is easy to trace, and intimately linked to the inter-arrival times (1),

$$\Pr[\tilde{A}_k = n | \text{Next-Drop}] = \tilde{t}(n) = \bar{p} \cdot p^n, \quad n \geq 0. \quad (3)$$

The complementary use of the series of random variables A_k and \tilde{A}_k suffices to capture the timing aspects of arrival and acceptance, relevant for our analysis.

2.3 General Burst Sizes

For the characterization of the burst sizes, we adopt the mentioned numbering of bursts, and so the k th burst has a burst size B_k . The burst sizes, just like the inter-arrival times, form a sequence of iid rv's with a common distribution, but have no further restriction on their distributions. Therefore, we consider general probabilities

$$\Pr[B_k = n] = b(n), \quad n \geq 1, \quad (4)$$

that are arbitrary, except for the conditions that any useful probability mass function has to comply with: $0 \leq b(n) \leq 1$, $\sum_n b(n) = 1$.

2.4 System Equations

As mentioned, the system's evolution can be captured by means of the waiting time of a burst only. Still using the same numbering, we associate the waiting time W_k with the k th burst, and define it as the time between the acceptance of burst k , and the start of its transmission. Again considering that either the next burst is accepted or dropped, we have

Next-Accept. The condition for this to happen, in terms of waiting times and FDL lengths, is that the burst that arrives just after the k th burst can be provided a sufficiently long delay, that is,

$$W_k + B_k - T_k \leq \omega_N.$$

Then, A_k equals T_k , and

$$W_{k+1} = \lceil W_k + B_k - A_k \rceil_{\Omega}, \quad (5)$$

where we adopted the notation $\lceil x \rceil_{\Omega} = \inf\{y \in \Omega, y \geq x\}$, $x \leq \omega_N$.

Next-Drop. Now, the burst that arrives just after the k th burst can not be provided a sufficiently long delay, and

$$W_k + B_k - T_k > \omega_N.$$

As a result of this (and of the memoryless nature of the arrival process), the waiting time of burst $k + 1$ no longer relates to W_k , and W_{k+1} is characterized by the reactivation time, through

$$W_{k+1} = \lceil \omega_N - \tilde{A}_k \rceil_{\Omega}. \quad (6)$$

These two system equations (5) and (6), together with their respective probability mass functions (2) and (3), provide the input for the analysis.

3 Analysis

In this section, the limited set of waiting times serves a state variable for a Markov chain, of which we trace the transition probabilities. Inversion yields the waiting time probabilities of accepted bursts, and this in turn allows to extract the loss ratio.

3.1 Markov Chain for Waiting Time

Before delving into the analysis, we note that the waiting time can only take on $N + 1$ different $\omega_i \in \Omega$. Therefore, it is an attractive state variable for a Markov chain approach to the system.

The Markov chain we consider consists of $N + 1$ states, that correspond to $N + 1$ possible waiting times ω_i , $i = 0 \dots N$. It is characterized by a transition matrix with probabilities β_{ij} ,

$$\beta_{ij} = \Pr[W_{k+1} = \omega_j | W_k = \omega_i], \quad 0 \leq i, j \leq N.$$

For ease of notation, we introduce $\omega_{-1} = -\infty$. We split β_{ij} in two separate contributions, that correspond to the events discussed in Sect. 2.4.

$$\begin{aligned} \beta_{ij} &= \Pr[\omega_i + B_k - T_k \leq \omega_N, \omega_j = \lceil \omega_i + B_k - A_k \rceil_{\Omega}] \\ &\quad + \Pr[\omega_i + B_k - T_k > \omega_N, \omega_j = \lceil \omega_N - \tilde{A}_k \rceil_{\Omega}] \\ &= \Pr[\omega_{j-1} - \omega_i < B_k - T_k \leq \omega_j - \omega_i] \\ &\quad + \Pr[B_k - T_k > \omega_N - \omega_i] \Pr[\omega_N - \omega_{j-1} > \tilde{A}_k \geq \omega_N - \omega_j]. \end{aligned} \quad (7)$$

Since the burst sizes B_k and inter-arrival times T_k only occur as $B_k - T_k$, we introduce the series of random variables $U_k = B_k - T_k$, and its cumulative distribution function (CDF) $U(n)$. Taking into account the distribution of B_k and A_k , calculations show that $U(n)$ simplifies to

$$U(n) = \Pr[B_k - A_k \leq n] = \sum_{i=1}^n b(i)(1 - \bar{p}^{i-n-1}) + \bar{p}^{-n-1}B(\bar{p}), \quad (8)$$

where the sum over n disappears if $n \leq 0$, and $B(z)$ is the probability generating function of B_k , defined as $B(z) = \mathbb{E}[z^{B_k}] = \sum_{n=1}^{\infty} b(n)z^n$. Similarly, we consider (3) to obtain the CDF of $\tilde{A}_k, \tilde{A}(n)$, as

$$\tilde{A}(n) = \Pr[\tilde{A}_k \leq n] = 1 - \bar{p}^{n+1}, \quad n \geq 0, \quad (9)$$

and zero when $n < 0$. Adopting these notations, (7) can be stated as

$$\begin{aligned} \beta_{ij} = & U(\omega_j - \omega_i) - U(\omega_{j-1} - \omega_i) \\ & + [1 - U(\omega_N - \omega_i)][\tilde{A}(\omega_N - \omega_{j-1} - 1) - \tilde{A}(\omega_N - \omega_j - 1)] \end{aligned} \quad (10)$$

Combining the last three expressions, we obtain an explicit expression for the coefficients β_{ij} , in terms of the (given) $\omega_i, b(n)$ and p .

From here, we obtain the vector of the waiting times as the normalized left eigenvector $[w(n)]$ of the matrix $[\beta_{ij}]$, associated with the eigenvalue 1, that is to satisfy $w(n) = \sum_{j=0}^N w(j)\beta_{jn}, 0 \leq n \leq N$. This eigenvector can easily be obtained numerically, posing no problem for the small N we are interested in. It contains the $N + 1$ different steady-state waiting time probabilities

$$\lim_{k \rightarrow \infty} \Pr[W_k = \omega_n] = \Pr[W = \omega_n] = w(n), \quad 0 \leq n \leq N. \quad (11)$$

From this, we can also define a mean waiting time $\mathbb{E}[W] = \sum_{i=1}^N w(i)\omega_i$.

3.2 Loss Ratio

Up to now, we considered only bursts that were accepted, and even chose to leave the dropped bursts unnumbered. We now focus on the burst loss ratio (LR), defined as the fraction of arriving bursts that is dropped, and study the unavailable period, associated with an accepted burst k , in two cases. If on the one hand, the arrival of burst k does not push the system into unavailability, then the unavailable period following burst k equals zero. This implies that, the time slot after the arrival of burst k , a new arrival can already be accepted. In terms of the involved rv's, this means that $W_k + B_k - 1 \leq \omega_N$.

On the other hand, if the unavailable period following burst k is larger than zero, then $W_k + B_k - 1 > \omega_N$. Now, it takes the system a number of slots equal to $W_k + B_k - \omega_N - 1$, to become available again. The last measure is the unavailable period following burst k , under the condition $W_k + B_k - 1 > \omega_N$. Combination of both cases leads to the conclusion that the unavailable period, following burst k , is given

by $(W_k + B_k - \omega_N - 1)^+$, where $(x)^+$ is shorthand for $\max\{0, x\}$. Invoking the memoryless nature of the arrival process, we can write down an expression for $E[X_k]$, the average number of lost bursts during the unavailable period following burst k ,

$$E[X_k] = p \cdot E[(W_k + B_k - \omega_N - 1)^+]. \quad (12)$$

With (4) and (11), this becomes

$$E[X_k] = p \cdot \left(E[B] + E[W] - \omega_N - 1 - \sum_{i=0}^N w(i) \sum_{j=1}^{\omega_N - \omega_i} b(j)(j + \omega_i - \omega_N - 1) \right).$$

Now, it suffices to note that, with every accepted burst, a number of $E[X_k]$ bursts on average is dropped, resulting in a burst loss ratio (LR)

$$\text{LR} = E[X_k]/(1 + E[X_k]).$$

4 Numerical Results and Discussion

With the above results at hand, one can easily study the impact of the various design parameters on loss performance. More specifically, one wants to determine optima for the granularity, which are values that yield a minimal burst loss ratio. While similar curves already occur in [6] for $N = 20$, the approximation applied there lost accuracy for small N . As such, the examples given here yield additional information, for the case of smaller (more realistic) buffer sizes. We look at four examples.

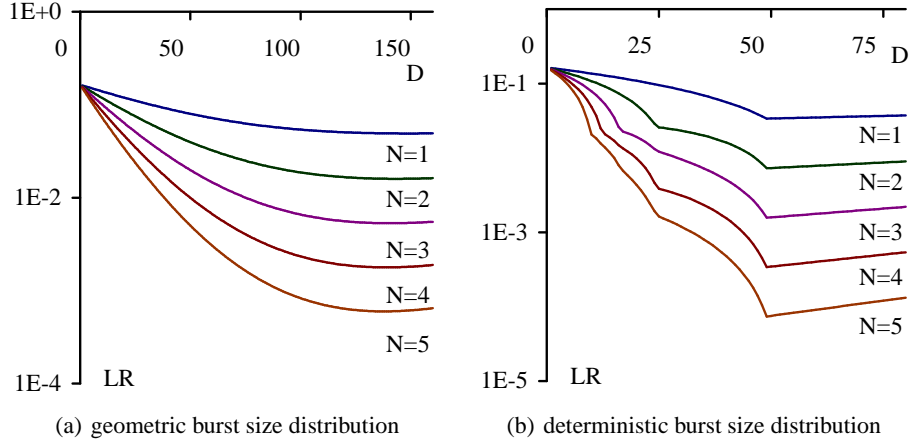


Fig. 1. Loss ratio as function of the FDL granularity (slots), for various buffer sizes N , burst size distributions (a) and (b), both with $E[B] = 50$ slots, for a load $\rho = 0.2$.

In both panes of Fig. 1, we assess the impact of a small buffer size on loss performance. They show the burst loss ratio (LR) as a function of the granularity D , for an

equidistant delay line setting, $E[B] = 50$, and a load of $\rho = 0.2$. The left pane displays the situation for geometric burst size distribution, with the expected rise of the loss ratio, when the buffer size N decreases. Also, it can be seen that the LR lowers for increasing granularity, and only starts to rise again for D larger than about 150 slots. Considering the five curves together, the figure suggests that the optimal granularity, for a given load, only slightly increases when the buffer size approaches its minimum of $N = 1$. Results for higher loads, not included here, confirm that the influence of diminishing buffer size on optimal granularity is but weak, especially if compared to the impact of variations in the load. The latter is illustrated in [6], and is much stronger than the influence observed here.

On the right pane of Fig. 1, we have the same setting, for fixed burst sizes, $E[B] = B = 50$. Again, a smaller buffer suffers more loss, but now the optima alter in a more surprising way. More precisely, the “notches” at $D = (B - 1)/n$, $n = 1, 2, \dots$, known from previous work, are not uncountable (as was the case in [6] for an infinite-sized buffer), but are limited in number to N . Exactly N notches occur for each curve, at $D = (B - 1)/n$, $n = 1, 2, \dots, N$. Since it is known from [6] that these values correspond to optima (with $D = B - 1$ being the optimum for low load, and $D = (B - 1)/2$, then $(B - 1)/3, \dots$ for increasing load), we verified and found that the set of potential optima, for a load $0 \leq \rho \leq 1$, is indeed limited to (at the maximum) the number of fibers. Also, it was found that the same optimum shift as known from [6] takes place, but now only over the available values: first $D = B - 1$ for low load, then $D = (B - 1)/2$ for increasing load if $N \geq 2$, and so on if $N \geq 3 \dots$

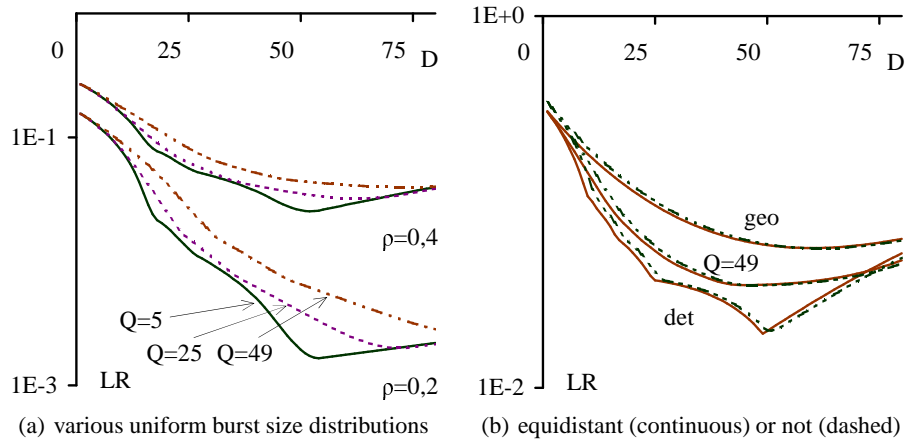


Fig. 2. Loss ratio as function of the FDL granularity (slots), for $E[B] = 50$ slots, with (a) uniform burst size distributions (various radii Q) and (b) different burst size distributions for an equidistant and non-equidistant setting, load $\rho = 0.5$.

To verify if the optima for deterministic burst sizes also apply to varying burst sizes, we consider a uniform burst size distribution with radius Q , that has a mean burst size

$E[B] = 50$, and is uniform within the range $[50 - Q, 50 + Q]$. For small Q , this distribution resembles the deterministic distribution. As such, this setting allows to verify what influence variances on the burst size have on loss performance. In Fig. 2, the left pane compares the performance of an optical buffer of size $N = 3$ for three uniform distributions, having a narrow range ($Q = 5$, range $[45, 55]$), an intermediate range ($Q = 25$, range $[25, 75]$) and a broad range ($Q = 49$, range $[1, 99]$), and this for load $\rho = 0.2$, $\rho = 0.4$, respectively. For the narrow-ranged one, the curves look very similar to those of the deterministic distribution, and the same optimum around $D = B - 1$ shows. The curves for the intermediate-ranged and broad-ranged case show that increasing Q makes the granularity optimum shift toward higher values, at least for load $\rho = 0.2$ and $\rho = 0.4$. Curves not included here, for higher load, show that the optima for a narrow range (small Q) concentrate around the limited set $D = (B - 1)/n$, $n = 1, 2, \dots, N$, known from the deterministic distribution, while for larger Q , the optimum only gradually decreases.

Choosing non-equidistant lengths for the delay lines can in some cases provide better performance. As is shown in [11] for deterministic burst sizes, a non-equidistant set of FDLs can outperform an equidistant set of the same size. However, it turns out that this happens only when the load rises above a certain value (for example 60.17%, for $N = 10$, $B = 20$, $D = 19$). Further, even when the non-equidistant one outperforms the equidistant one, the performance gain is rather small. This said, non-equidistant settings remain interesting, since for more general assumptions (correlated arrivals, multi-wavelength output), the performance gain might be larger. For the right pane of Fig. 2, we chose non-equidistant FDL lengths in a way similar to [11], with shortened lengths for the largest lines. An equidistant set (continuous curves) and non-equidistant set (dashed curves) are considered, for a buffer size $N = 5$, load $\rho = 0.5$, and the burst size distributions geometric, deterministic and uniform ($Q = 49$). The non-equidistant set has FDL lengths $D, 2D - 2, 3D - 3, 4D - 4, 5D - 8$. The curves show how the non-equidistant set just outperforms the equidistant one, for geometric and uniform burst size distribution, while the opposite is true for a deterministic burst size distribution. Although not included here, figures for the same setting, for a load of $\rho = 0.3$, $\rho = 0.6$ and $\rho = 0.8$ resp., show the same qualitative result, while the performance difference itself always remains small.

5 Conclusions

Unlike previous work, we have studied an optical buffer by considering only the waiting times, that correspond to the lengths of the Fiber Delay Lines (FDLs). This straightforward approach allowed us to obtain exact results for a small computational load, especially for small buffers. Without posing restrictions on the lengths of these lines, we constructed a model valid for a Bernoulli arrival process, and a general burst size distribution, based on the analysis of the involved Markov chain. The results of the latter we used to obtain (i) the steady-state waiting time probabilities and, by considering the unavailable periods, (ii) the loss ratio.

The presented performance graphs show the impact of reduced buffer size on performance, and the associated optimal value for the granularity. While the optimum for

geometric burst size distributions hardly changes when we lower the buffer size, the optimum for a deterministic burst size distribution does change. Also, a study of a uniform distribution showed how smaller and larger variation around an average burst size impacts performance. Finally, performance results for a non-equidistant set were compared to those of an equidistant set.

Concluding, the model presented here is rather elementary in its approach, and therefore provides a basic tool, ready to use for optimization studies. The extension of the model toward (i) multi-wavelength systems, (ii) correlated arrivals, we consider challenging, and would surely lead to a more profound insight in the operation of optical buffers.

References

1. C. Qiao and M. Yoo. Optical burst switching—a new paradigm for an optical internet. *Journal on High-Speed Networks*, 8:69–84, 1999.
2. T. El-Bawab and J.-D. Shin. Optical packet switching in core networks: between vision and reality. *IEEE Communications Magazine*, 40:60–65, September 2002.
3. F. Callegati. Optical buffers for variable length packets. *Communications Letters, IEEE*, 4(9):292–294, 2000.
4. F. Callegati. On the design of optical buffers for variable length packet traffic. In *Proceedings of the Ninth International Conference on Computer Communications and Networks (Las Vegas)*, pages 448–452, 2000.
5. F. Callegati. Approximate modeling of optical buffers for variable length packets. *Photonic Network Communications*, 3(4):383–390, October 2001.
6. K. Laevens and H. Bruneel. Analysis of a single-wavelength optical buffer. In *Proceedings of Infocom 2003 (San Francisco)*, 2003.
7. W. Rogiest, K. Laevens, D. Fiems, and H. Bruneel. A performance model for an asynchronous optical buffer. *Perform. Eval.*, 62(1-4):313–330, 2005.
8. M. Murata and K. Kitayama. Ultrafast photonic label switch for asynchronous packets of variable length. In *INFOCOM 2002 (New York)*, 2002.
9. R. C. Almeida, J. U. Pelegri, and H. Waldman. Optical buffer modelling for performance evaluation considering any packet inter-arrival time distribution. In *IEEE International Conference on Communications, 2004*, volume 3, pages 1771–1775, 2004.
10. R. C. Almeida, J. U. Pelegri, and H. Waldman. A generic-traffic optical buffer modeling for asynchronous optical switching networks. *Communications Letters, IEEE*, 9(2):175–177, 2005.
11. J. Lambert, B. Van Houdt, and C. Blondia. Single-wavelength optical buffers: non-equidistant structures and preventive drop mechanisms. In *Proceedings of NAEC 2005 (Riva Del Garda)*, pages 545–555, 2005.
12. J. Lambert, B. Van Houdt, and C. Blondia. Queues with correlated inter-arrival and service times and its application to optical buffers. *Stochastic Models*, 22(2):233–251, 2006.