Performance Analysis of Optical Burst Switching Networks with and without Class Isolation

Neil Barakat and Edward H. Sargent
Edward S. Rogers School of Electrical and Computer Engineering
University of Toronto
10 King's College Road
Toronto, Ontario, Canada M5S 3L1
Email: neil.barakat@utoronto.ca, ted.sargent@utoronto.ca

***Abstract*— This paper presents an analytical model that evaluates the blocking probability of each service class in optical burst-switching networks. The model is applicable to systems with arbitrary burst length distributions and arbitrary-sized QoS header offsets. Thus, unlike previous models, it is applicable to the design and study of networks with a wide range of traffic characteristics, including systems in which higher classes are not necessarily isolated from lower classes and systems in which the conservation law does not necessarily hold. We derive explicit expressions for blocking probability both the cases of constant burst lengths and exponentially distributed burst lengths and verify the model's accuracy through simulation. We show the model to be accurate for a number of different traffic loads and class priorities. For an OBS system with two classes and a 1:10 ratio of high-priority to low-priority traffic, our model is able to predict accurately the blocking probability for each class, whereas the predictions from a model that assumes isolation deviates by as much as an order of magnitude from the simulation results for the higher priority class.**

I. INTRODUCTION

Emerging WDM technology allows huge amounts of data to be carried transparently in optical networks. Optical burst switching (OBS) is a means of accommodating bursty traffic and allowing sub-wavelength granularity in a transparently optical manner without the need for buffering or other complex functionality in the optical plane. This is achieved by decoupling of the control plane and the data plane. In OBS, IP packets are electronically aggregated at the edge of the network into bursts that are transmitted transparently through the network core. The core switches are configured in advance to accept the burst using control packets, which are transmitted before the bursts in an electronic control plane.

OBS is characterized by two main qualities. First, in order to accommodate the bounded, non-zero electronic processing time of control packets in core switches, OBS introduces a time offset between each burst's control packet and payload. This offset eliminates the need for buffering of bursts in the core of the network, which is typically required in optical packet-switched networks. Secondly, because of the relatively short durations of the burst compared to the network size, most OBS architectures uses one-way reservation schemes in order to maximize bandwidth efficiency and reduce latency. In such schemes, the burst is sent into the network without waiting for a reservation acknowledgement. In one such signaling protocol, called just-enough-time (JET) [1], each control packet carries information about the offset size and the length of the burst and attempts to reserve a corresponding bandwidth window. If such a window is not available due to previous reservations, the burst is dropped upon arriving at the switch.

An additional offset, which we term a *quality of service (QoS) offset*, can also be added to increase a burst's priority and implement priority classes [2]. If the total offset of the control packet for a class i burst is longer than the offset of a class j burst *plus* the *maximum* class j burst length, then class i will be isolated from class j [3]. In this paper we consider OBS networks that employ JET with multiple classes. We further assume that no buffering or wavelength conversion is available in the network core.

Analytical models have previously been proposed for OBS networks with multiple priority classes. In [4] and [3], the authors assume that the QoS offsets are sufficiently large to ensure class isolation and that the OBS system obeys the conservation law.[1] This leads to a recursive analytical formula that approximates the blocking probability of each class using the Erlang-B blocking formula. However, depending on the ingress traffic characteristics, the burst aggregation scheme, and the resulting maximum burst length, achieving class isolation for even two classes may impose a large delay on the higher priority class. For a network with N classes, this imposed delay must be increased N-fold if isolation between each class is to be achieved. Therefore, depending on the latency requirements of the traffic, complete isolation may not be practical. Further, it has been shown that there are a number of different traffic scenarios for which the conservation law does not hold [6], so it is desirable to have models that are accurate for these systems as well. This motivates a more general model that is accurate for any degree of class isolation and for systems that do not necessarily obey the conservation law.

In OBS systems without complete isolation, the degree of isolation and the resulting performance of the OBS system

[1] In general, (work) conservation is a term used to describe scheduling mechanisms in which the priority discipline does not effect the overall average performance of the system [5]. If the conservation law holds for an OBS system, the overall rate of blocking is unaffected by the number of classes and the size of the QoS offsets assigned to them.

may depend quite strongly on the burst length distributions. This, coupled with the fact that burst aggregation mechanisms and their effect on burst-length distributions is still a very active area of research [7][8][9][10] motivates the development of models that are applicable and accurate for arbitrary burst length distributions.

In this paper, we present an analytical model for computing blocking probability of a multi-class OBS system. The model explicitly takes into account the control packet offset and the burst length distribution. As such, the results are accurate for an arbitrary number of classes with an arbitrary degree of class isolation. Further, since these probabilities are computed directly (i.e. not recursively as in previous models), the results apply equally well to work-conserving and non-work-conserving systems. Lastly, by allowing for closed-form, non-recursive expressions for blocking probabilities, our model may also be amenable for use in network design and operation.

In Section II and Section III, we derive the model. In Section IV, we use the model to obtain closed form expressions for blocking probabilities of each class for both the cases of fixed and exponential burst-length distributions. We demonstrate the model's accuracy and examine the performance of OBS networks with and without class isolation in Section V. Conclusions are drawn and future work is discussed in Section VI.

II. Modeling the Two-Burst Blocking Process

Since OBS uses advanced reservation, a burst can be blocked if a contending burst overlaps its head, its tail, both its head and tail, or some part of its middle, as shown in Fig. 1, where

- T_i is the arrival time of a given class i control packet
- L_i is a random variable representing the length of a class i burst
- ω_i is the offset time between the control packet and burst for class i (including the control packet length and assuming that processing-time offset is negligible compared to QoS offset).

We can represent all four types of blocking shown in Fig. 1 simultaneously, by recognizing that a given class i burst will be blocked by a class j burst, if all of the following events occur:

1) the control packet of the class j burst arrives before that of the class i burst
2) the start of the class j burst arrives before the end of the class i burst
3) the end of the class j burst arrives after the start of the class i burst.

We denote the intersection of these three events as B_{ij}. Thus we have

$$P[B_{ij}] = P[(T_j < T_i) \cap (T_j + \omega_j < T_i + \omega_i + L_i) \cap (T_j + \omega_j + L_j > T_i + \omega_i)]. \quad (1)$$

We define $\delta_{ij} \triangleq \omega_i - \omega_j$ as the offset difference between class i and class j. If bursts in class j arrive according to a Poisson process, then we can make use of the memoryless property [11] to write

$$P[B_{ij}] = P[(T_j - T_i) \in [\delta_{ij} - L_j, \min(0, L_i + \delta_{ij})]] \quad (2)$$

$$= P[\tau_j < \min(0, L_i + \delta_{ij}) - (\delta_{ij} - L_j)] \quad (3)$$

where τ_j is a random variable representing the interarrival time of class j bursts.

For the multi-class model in the next section, we will require the probability of $\overline{B_{ij}}$, the complement of B_{ij}. From (3), we have

$$P[\overline{B_{ij}}] = P[\tau_j > \min(0, L_i + \delta_{ij}) - (\delta_{ij} - L_j)] \quad (4)$$

$$= \begin{cases} P[\tau_j > L_j + L_i], & \delta_{ij} < 0, L_i < -\delta_{ij} \\ P[\tau_j > L_j - \delta_{ij}], & \delta_{ij} < 0, L_i \geqslant -\delta_{ij} \\ P[\tau_j > L_j - \delta_{ij}], & \delta_{ij} \geqslant 0 \end{cases} \quad (5)$$

By conditioning over all possible values of L_i, we can combine the first two cases in (5) and write

$$P[\overline{B_{ij}}] = \begin{cases} \int_0^{-\delta_{ij}} P[\tau_j > L_j + l_i] f_{L_i}(l_i) dl_i \\ \quad + \int_{-\delta_{ij}}^{\infty} P[\tau_j > L_j - \delta_{ij}] f_{L_i}(l_i) dl_i, & \delta_{ij} < 0 \\ P[\tau_j > L_j - \delta_{ij}], & \delta_{ij} \geqslant 0. \end{cases} \quad (6)$$

$$= \begin{cases} \int_0^{-\delta_{ij}} P[\tau_j > L_j + l_i] f_{L_i}(l_i) dl_i \\ \quad + P[\tau_j > L_j - \delta_{ij}](1 - F_{L_i}(-\delta_{ij})), & \delta_{ij} < 0 \\ P[\tau_j > L_j - \delta_{ij}], & \delta_{ij} \geqslant 0 \end{cases} \quad (7)$$

where $f_{L_i}(l)$ and $F_{L_i}(l)$ are the probability density and distribution functions for the length of class i bursts respectively, and $P[\tau_j > L_j + l_i]$ and $P[\tau_j > L_j - \delta_{ij}]$ can be evaluated using (24) (see appendix).

III. Multi-Class Blocking Model

In this section, without loss of generality, we assume that the network has N classes of traffic labelled $1, \ldots, N$ such that $\delta_{ij} > 0$ for all $i < j$. Thus, class 1 has the highest priority, class 2 has the second highest, etc. We are interested in finding P_{bi} the average blocking probability of a class i burst. When a control packet arrives, it attempts to reserve a time-slot of bandwidth to accommodate its burst. We call this time-slot the *reservation window* (*RW*) of the burst. When a control packet arrives, if one or more burst lie in its RW, the bursts is blocked. Let us denote the event that k bursts have arrived in the reservation window of a class i burst as a_i^k. Using this definition we have

$$P_{bi} = 1 - P[\text{no bursts lie in RW of class } i \text{ burst}] \quad (8)$$

$$= 1 - P[a_i^0] - \sum_{k=1}^{\infty} P[a_i^k \cap \text{all } k \text{ are blocked}] \quad (9)$$

$$= 1 - P[a_i^0] - \sum_{k=1}^{\infty} \Big(P[k \text{ bursts blocked} \mid a_i^k] \cdot P[a_i^k] \Big). \quad (10)$$

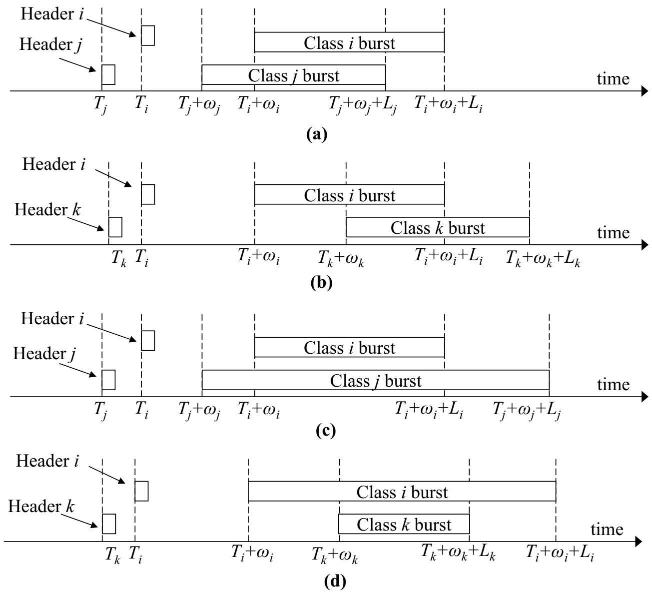

Fig. 1. Types of Blocking in OBS network. (a) Contending burst overlaps head of blocked burst. (b) Contending burst overlaps tail of blocked burst. (c) Contending burst overlaps both head and tail of blocked burst. (d) Contending burst overlaps middle of blocked burst. Note that (b) and (d) are only possible if the priority class of the contending burst is higher than that of the blocked burst (i.e. if $\omega_k > \omega_i$)

For practical networks, blocking events are relatively rare. Thus, $P[k \text{ bursts blocked} \mid a_i^{(k)}] \ll P[a_i^{(0)}]$, so we can ignore all higher order terms in (10), yielding

$$P_{bi} \approxeq 1 - P[a_i^0] = 1 - P\left[\bigcap_{j=1}^{N} \overline{B_{ij}}\right] = 1 - P\left[\bigcap_{j=1}^{i-1} \overline{B_{ij}} \bigcap_{j=i}^{N} \overline{B_{ij}}\right] \quad (11)$$

where we have divided the above intersection into two subsets, corresponding to the effects of higher-priority traffic, and the effects of equal-priority and lower-priority traffic. From (7), we can see that for $i \leqslant j$, or equivalently when $\delta_{ij} \geqslant 0$, all of the $\overline{B_{ij}}$ are independent since each depends only on τ_j and L_j. So we have

$$P_{bi} = 1 - \bigcap_{j=1}^{i-1} P[\overline{B_{ij}}] \cdot P\left[\prod_{j=i}^{N} \overline{B_{ij}}\right]. \quad (12)$$

When $\delta_{ij} < 0$, all of the $\overline{B_{ij}}$ term depend on L_i, and thus independence is not assured. However, in this paper, we make the approximation[2] that they are independent, and write

$$P_{bi} = 1 - \prod_{j=1}^{i-1} P[\overline{B_{ij}}] \prod_{j=i}^{N} P[\overline{B_{ij}}]. \quad (13)$$

[2]The accuracy of the resulting model, as described in V, attests to the validity of this approximation.

We then substitute (7) into (13), to yield

$$P_{bi} = 1 - \prod_{j=1}^{i-1} \left(\int_0^{-\delta_{ij}} P[\tau_j > L_j + l_i] f_{L_i}(l_i) dl_i + [1 - F_{L_i}(-\delta_{ij})] P[\tau_j > L_j - \delta_{ij}] \right) \cdot \prod_{j=i}^{N} P[\tau_j > L_j - \delta_{ij}]. \quad (14)$$

Equation (14) gives the expression for the average blocking probability of a class i burst when there are N classes of traffic, each with Poisson arrivals, and each with an arbitrary burst-length distribution and offset time.

IV. Exponential and Constant Burst Lengths Distributions

We now use the model derived in Section II and III to obtain closed-form expressions for the blocking probability of each class in networks with constant and exponential burst-length distributions.

A. Exponential Burst Lengths

In this section, we assume that the burst lengths for each class follows an exponential distribution. Thus, the probability density function f_{L_k} and corresponding distribution function F_{L_k} for class k bursts are

$$f_{L_k}(l) = \frac{1}{\overline{l_k}} e^{-l/\overline{l_k}} \quad \text{and} \quad F_{L_k}(l) = 1 - e^{-l/\overline{l_k}} \quad (15)$$

where $\overline{l_k}$ is the mean burst length of class k. P_{bi} can be found directly by substituting (25) and (15) into (14) and solving the resulting integral to yield

$$P_{bi} = 1 - \frac{\prod_{j=1}^{i-1} \left[\frac{1 + \overline{l_i} \lambda_j e^{\delta_{ij}(\lambda_j + 1/\overline{l_j})}}{1 + \overline{l_i} \lambda_j} \right] \prod_{j=i}^{N} \left[1 + \rho_j (1 - e^{-\delta ij/\overline{l_j}}) \right]}{\prod_{j=1}^{N} [1 + \rho_j]} \quad (16)$$

where λ_k is the mean arrival rate of class k bursts, and where $\rho_k = \overline{l_k} \lambda_k$ is the offered load of class k.

B. Fixed Burst Lengths

Here we assume that all bursts from class i have a fixed length equal to $\overline{l_i}$ and compute P_{bi}, the blocking probability for class i. Proceeding directly from (5), and using the fact that the class i burst interarrival time τ_i is simply an exponential random variable with mean $1/\lambda_i$, we have

$$P[\overline{B_{ij}}] = \begin{cases} e^{-\lambda_j(\overline{l_i} + \overline{l_j})}, & \delta_{ij} < 0, l_i < -\delta_{ij} \\ e^{-\lambda_j(\overline{l_j} - \delta_{ij})}, & \delta_{ij} < 0, l_i \geqslant -\delta_{ij} \\ e^{-\lambda_j(\overline{l_j} - \delta_{ij})}, & 0 \leqslant \delta_{ij} \leqslant l_j \\ 1, & \delta_{ij} > l_j. \end{cases} \quad (17)$$

If we again define $\rho_i = \lambda_i \overline{l_i}$ as the offered load of class i then we can write

$$P[\overline{B_{ij}}] = e^{-\rho_j(1 - d_{ij})} \quad (18)$$

where

$$d_{ij} = \begin{cases} 1, & \delta_{ij} \geqslant \overline{l_j} \\ \delta_{ij}/\overline{l_j}, & -\overline{l_i} < \delta_{ij} < \overline{l_j} \\ \overline{l_i}/\overline{l_j}, & \delta_{ij} \leqq -\overline{l_i} \end{cases} \quad (19)$$

and d_{ij} can be naturally thought of as the *degree of isolation* between class i and class j. Finally, substituting (18) into (13), we have

$$P_{bi} = 1 - \frac{\exp\left[\sum_{j=1}^{N} \rho_j d_{ij}\right]}{\exp\left[\sum_{j=1}^{N} \rho_j\right]}. \quad (20)$$

V. Accuracy of Model and Simulation Results

We simulated a two-class OBS node to demonstrate the accuracy of the analytical model and examine the effects of offered load and offsets on network performance. For each simulation, depending on the offered load, between two million and one hundred million bursts were simulated. Bursts arrived according to a Poisson process with an average burst length of 10 time units. Two different burst-length distributions were simulated: exponential and fixed-length. We found that our model was accurate for any degree of class isolation and over a large range of offered loads.

A. Modeling the Effects of Different Degrees of Class Isolation

For the results in this section, the overall traffic load (sum of class 1 and class 2 traffic) was 5×10^{-3} Erlangs, and the ratio of class 1 (higher class) traffic to class 2 traffic was 0.1.

An expression for the blocking probability for two classes with exponential burst lengths was found by setting $N = 2$ in (16), yielding

$$P_{b1} = 1 - \frac{1 + \rho_2(1 - e^{-\delta_{12}/\overline{l_2}})}{(1+\rho_1)(1+\rho_2)}$$
$$P_{b2} = 1 - \frac{1 + \overline{l_2}\lambda_1 e^{\delta_{21}(\lambda_1 + 1/\overline{l_1})}}{(1+\rho_1)(1+\rho_2)(1+\overline{l_2}\lambda_1)}. \quad (21)$$

Similarly, for the case of fixed-length bursts, we set $N = 2$ in (20) to yield

$$P_{b1} = 1 - \frac{e^{\rho_2 d_{12}}}{e^{\rho_1+\rho_2}}, \quad P_{b2} = 1 - \frac{e^{\rho_1 d_{21}}}{e^{\rho_1+\rho_2}}. \quad (22)$$

Fig. 2 compares the simulation results to our model and a model that assumes conservation between classes [4]. CL1 and CL2 refer to class 1 and class 2 respectively. We plot the blocking probability for each class as we vary the QoS offset of class 1 traffic, while keeping the class 2 QoS offset equal to zero. The results for fixed and exponential length burst are shown in Fig. 2(a) and Fig. 2(b) respectively. On the abscissa, we show the normalized QoS offset difference, $\Delta_{12} = \delta_{12}/\overline{l_2}$, which we have defined as the ratio of the difference of the class 1 and class 2 offsets to the mean burst length of class 2.

Complete isolation is achieved for fixed-length bursts when $\Delta_{12} \geqslant 1$. Beyond this point, the blocking probabilities are constant. For the exponential case, there is no maximum burst length, so complete isolation cannot be achieved. However, for the parameters simulated, the two classes are nearly isolated

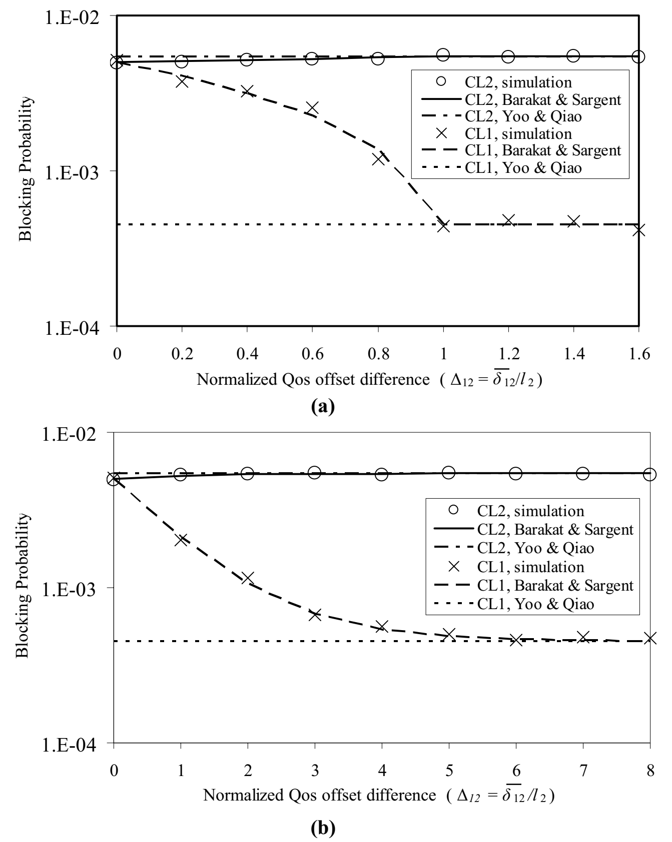

Fig. 2. Blocking probability versus QoS offset of high class for two-class OBS node. Bursts' lengths are fixed in (a) and exponentially distributed in (b).

for $\Delta_{12} > 8$. Subsequent simulations showed that value is dependent on not only the burst-length distribution, but also on the relative loads of the higher and lower class traffic. Thus, the nature of the lower class traffic must be carefully considered when selecting the size of the QoS offsets to achieve isolation.

For high values of Δ_{12}, when class 1 is well isolated from class 2, both models are quite accurate. However, as Δ_{12} decreases and class 1 becomes less isolated from class 2, the blocking experienced by class 1 and class 2 increases and decreases respectively. Because the model in [4] assumes complete isolation between classes, its computed blocking probabilities remain constant and are inaccurate for small values of Δ_{12}, deviating from the simulation results by as much as an order of magnitude. By contrast, the model presented in this paper makes no such assumption, and its output curve agrees closely with the simulation for all values of Δ_{12}. This ability to model accurately the blocking behavior in the regime of low Δ_{12} is particularly useful for systems with long bursts or low latency requirements, where complete isolation may not be achievable.

B. Effect of Offered Load on Accuracy of Model

In this section, we examine the accuracy of our analytical model over a large range of offered loads. We considered a network with three classes of traffic and exponential burst-length distributions. We fixed the ratio of Class 2 traffic to Class 1 traffic and the ratio of class 3 traffic to class 2 traffic at 10, while varying the overall offered load from 10^{-4} to 10^{3}.

A closed-form expression for the blocking probability of each class when burst lengths are exponentially distributed was found by setting $N = 3$ in (16) to yield

$$P_{b1} = 1 - \frac{[1+\rho_2(1-e^{-\delta_{12}/\overline{l_2}})][1+\rho_3(1-e^{-\delta_{13}/\overline{l_3}})]}{(1+\rho_1)(1+\rho_2)(1+\rho_3)}$$
$$P_{b2} = 1 - \frac{[1+\rho_3(1-e^{-\delta_{23}/\overline{l_3}})][1+\overline{l_2}\lambda_1 e^{\delta_{21}(\lambda_1+1/\overline{l_2})}]}{(1+\rho_1)(1+\rho_2)(1+\rho_3)(1+\overline{l_2}\lambda_1)}$$
$$P_{b3} = 1 - \frac{[1+\overline{l_3}\lambda_1 e^{\delta_{31}(\lambda_1+1/\overline{l_3})}][1+\overline{l_3}\lambda_2 e^{\delta_{32}(\lambda_2+1/\overline{l_3})}]}{(1+\rho_1)(1+\rho_2)(1+\rho_3)(1+\overline{l_3}\lambda_1)(1+\overline{l_3}\lambda_2)}. \quad (23)$$

Fig. 3(a) shows results when the offset of class 1, class 2, and class 3 were set to 0, 30, and 60 respectively. For the results in Fig. 3(b), the offsets for class 1, class 2, and class 3 were 0, 20, and 40 respectively. Thus, Fig. 3(a) corresponds to a network in which each class is well isolated from the classes below it, and Fig. 3(b) corresponds to a network for which classes are not isolated from one another.

In both sets of simulation results in Fig. 3, we see that the relative priority and relative blocking probability of each class remains unchanged as the load varies up to 1, above which, the blocking probability of all classes begins to asymptotically approach 1. Therefore, we conclude that the QoS offset mechanism for class differentiation scales very well to arbitrary traffic loads.

In both Fig. 3(a) and Fig. 3(b), for offered loads < 1, our model predicts the blocking probability almost exactly for all three classes. For offered loads $\geqslant 1$, the predicted blocking probability deviates slightly from the simulation. This is due to the first order approximation in (11), which assumes that the blocking rate in then network $\ll 1$. When the overall offered load is large enough such that the resulting blocking rate is greater than 0.1, the above assumption does not hold, and the predictions from our model begin to deviate slightly from the simulations. However, this deviation is quite small even as the blocking rates approach 1. Thus, we conclude that, although our model is very accurate for all levels of blocking, it is most well suited for networks in which the blocking rate is less than 10%, which includes the vast majority of practical communication networks.

VI. Conclusions

In this paper, we derived an analytical model that evaluates the blocking probability for each class in a multi-class OBS system with arbitrary burst length distributions. Using simulation, we showed that the model accurately computes the average blocking probability experienced by multiple classes of traffic for both the cases of isolated and non-isolated traffic classes. The model was also shown to be very accurate for all values of offered load simulated.

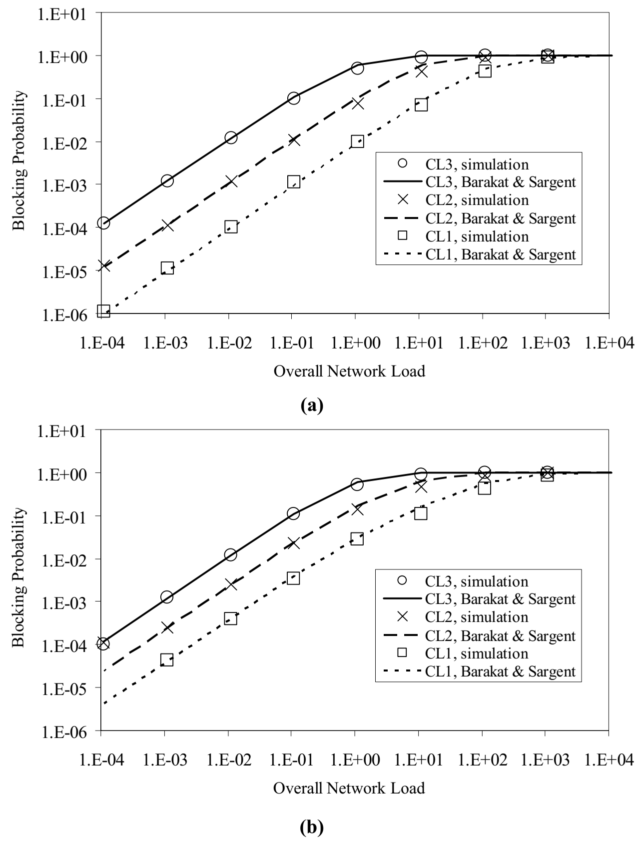

Fig. 3. Blocking probability versus overall network load with three classes. Offsets are chosen such that each class is isolated from lower classes in (a) and not isolated from lower classes in (b).

Future work includes extending the model to include the effects of wavelength conversion. Further, the model's ability to derive closed-form expressions for blocking implies that it can be used in future studies to aid in network design and operation. For example, by inverting the expressions in Section IV, offsets could be chosen based on desired blocking rates of each class, and traffic policing could be performed at the edge of the network to ensure that these blocking rates are achieved.

Appendix

If X and Y are non-negative random variables with joint distribution $f_{XY}(xy)$ and c is a constant,

$$P[X > Y + c] = \begin{cases} \int_0^\infty \int_0^{x-c} f_{XY}(x,y)dy\,dx & c < 0 \\ \int_0^\infty \int_{y+c}^\infty f_{XY}(x,y)dx\,dy & c \geqslant 0. \end{cases} \quad (24)$$

If X and Y are independent exponential random variables with mean $1/a_x$ and $1/a_y$ respectively,

$$P[X > Y + c] = \begin{cases} \frac{a_y + a_x(1-e^{a_y c})}{a^y + a^x} & c < 0 \\ \frac{a_y e^{-a_x c}}{a^y + a^x} & c \geqslant 0. \end{cases} \quad (25)$$

0-7803-7975-6/03/$17.00 (C) 2003

References

[1] M. Yoo and C. Qiao, "Just-enough-time (JET): a high speed protocol for bursty traffic in optical networks," in *Proc. IEEE/LEOS Summer Topical Meetings*, pp. 26–27, 1997.

[2] M. Yoo, C. Qiao, and S. Dixit, "Optical burst switching for service differentiation in the next-generation optical Internet," *IEEE Commun. Mag.*, vol. 39, pp. 98–104, Feb. 2001.

[3] H. L. Vu and M. Zukerman, "Blocking probability for priority classes in optical burst switching networks," *IEEE Commun. Lett.*, vol. 6, pp. 214–216, May. 2002.

[4] M. Yoo and C. Qiao, "Supporting multiple classes of services in IP over WDM networks," in *Proc. GLOBECOM '99*, pp. 1023–1027, 1999.

[5] N. K. Jaiswal, *Priority Queues*. Academic Press, New York, 1968.

[6] K. Dolzer and C. Gauger, "On burst assembly in optical burst switching networks - a performance evaluation of Just-Enough-Time," in *Proc. 17th International Teletraffic Congress*, pp. 149–160, Sept. 2001.

[7] A. Ge, F. Callegati, and L. Tamil, "On optical burst switching and self-similar traffic," *IEEE Commun. Lett.*, vol. 4, pp. 98–100, Mar. 2000.

[8] V. Vokkarane, K. Haridoss, and J. Jue, "Threshold-based burst assembly policies for QoS support in optical burst-switched networks," in *Proc. Opticomm '02*, pp. 125–136, 2002.

[9] J. White, M. Zukerman, and H. L. Vu, "A framework for optical burst switching network design," *IEEE Commun. Lett.*, vol. 6, pp. 268–270, Jun. 2002.

[10] M. Izal and J. Aracil, "On the influence of self similarity on optical burst switching traffic," in *Proc. Globecom '02*, pp. 2308–2312, 2002.

[11] L. Kleinrock, *Queueing Systems*. New York: John Wiley & Sons, Inc., 1975.

0-7803-7975-6/03/$17.00 (C) 2003