Analytical Modeling of Offset-Induced Priority in Multiclass OBS Networks

Neil Barakat, *Student Member, IEEE,* and Edward H. Sargent, *Senior Member, IEEE*

***Abstract*—In this paper, we present for the first time an analytical model that quantifies the mechanism by which offset size affects priority in multiclass optical-burst switching (OBS) systems. Using the model, we derive an exact expression for the distribution of the number of bursts that contend with an arriving burst. The model is applicable to systems in which each class has an arbitrary burst-length distribution and an arbitrary offset size. We also derive accurate approximate expressions for the burst-blocking probability of premium-class traffic, as well as expressions for the sensitivity of premium-class performance to offset jitter and variations in the arrival rates of each class. In a case study, we find that scaling up a system in terms of the number of wavelengths and the traffic load significantly improves not only the burst-blocking performance of the premium class, but also its sensitivity to lower class traffic variations. We also use the model to dimension and provision the system to guarantee a minimum level of premium-class blocking and premium-class robustness to low-class load variations.**

***Index Terms*—Analytical model, blocking probability, optical-burst switching (OBS), optical networks.**

I. Introduction

OPTICAL-BURST switching (OBS) is a promising architecture for accommodating bursty traffic and efficiently using the huge amounts of available bandwidth in wavelength-division multiplexing (WDM) optical networks. OBS is seen as a compromise between optical-packet switching and optical-circuit switching, as it benefits from the flexibility and efficiency of the former, while requiring a level of optical technology complexity that is closer to the latter [1]–[4].

In OBS, packets are assembled electronically at the edge of the network into long bursts, which are transmitted through the OBS network core entirely within the optical domain. By separating the transmission of the burst header (or control packet) and the burst payload by a time offset, OBS eliminates the need for buffering of the burst during processing of the header and configuring of the switch. Further, because headers are transmitted on a separate control channel and are processed electronically in each node, there is no need for optical processing.

In this paper, we restrict our examination to just-enough-time (JET) OBS [5]. Each control packet contains information about the size of its offset and the duration of its burst. Upon arriving at a node, it attempts to reserve resources for the time between the start and the end of its burst. If sufficient resources are not available to accommodate the burst, it is dropped. A multiclass version of JET to support quality of service (QoS) has also been proposed. In multiclass JET OBS, high-priority bursts are given an extended *QoS offset*, which gives them a higher probability of successfully reserving resources at each hop in their path [6].

Simulation is often used to analyze OBS system performance. However, simulations can be very time consuming, particularly when examining systems with very low blocking rates. Analytical models allow researchers to examine systems under a very wide range of system parameters in a reasonable amount of time. Perhaps even more significantly, models also provide valuable insight that may not be available from simulation. Further, models can also be integrated into scheduling mechanisms to provision resources dynamically and improve system performance.

Because the offset of every burst in a single-class OBS system is the same, it can be neglected for the purpose of system-performance analysis, and classical Markovian queueing analysis can yield accurate performance predictions. In multiclass OBS systems, however, the different header-offset sizes of each class render the blocking processes nonmemoryless. This poses a problem, because classical Markovian queueing models cannot be directly applied. However, a number of researchers have been able to apply Markovian models to OBS systems with multiple classes by restricting their examination to particular types of networks, and by describing the network and traffic parameters such that the systems resemble a set of classical M/G/m/m (or M/G/m/D) queuing systems. The two most common simplifications are the assumption that the high-priority class is completely isolated from the low-priority class [6]–[8], and the assumption that the OBS system is work-conserving [7]–[10].

In a system with class *isolation*, the behavior of the low-class traffic will have no effect on the performance of the premium class. Isolation between classes can be achieved if the additional offset of the high class is selected to be larger than the maximum burst length of the low class. However, depending on the ingress traffic characteristics, the burst-aggregation scheme, and the latency requirements of the premium-class traffic, complete isolation may not be practical. For a network with N classes, the premium-class offset must be increased N-fold if isolation between each class is to be achieved. Thus, there is an active interest in modeling multiclass OBS systems in which only partial isolation between classes exists [9], [10].

A multiclass OBS system is said to be *work-conserving* if the overall rate of blocking is unaffected by the number of classes and the size of the QoS offsets assigned to them. However, a number of practical OBS systems are not work-conserving. For

Paper approved by K. Kitayama, the Editor for Optical Communication of the IEEE Communications Society. Manuscript received April 19, 2004; revised October 2, 2004. This paper was presented in part at the IEEE Canadian Conference on Communications, Niagara Falls, ON, Canada, May 2004.

The authors are with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON M5S 3G4, Canada (e-mail: neil.barakat@utoronto.ca; ted.sargent@utoronto.ca).

Digital Object Identifier 10.1109/TCOMM.2005.852845

example, in [9] and [11], it was demonstrated that systems in which low-class bursts are much longer than high-class bursts do not obey the conservation law. Thus, the assumption of work conservation imposes a restriction on the types of systems that can be analyzed.

We presented a general analytical model for multiclass OBS systems that assumed neither work conservation nor isolation between classes in [11]. However, the model was only applicable to single-channel OBS networks.

In this paper, we derive an exact model that quantifies the mechanisms by which extended offsets induce priority and affect performance in a wide range of OBS systems, including multichannel systems and systems without work conservation or isolation. We define the notation used in Section II, and describe qualitatively how extended offsets can be used to implement multiple classes in Section III. We introduce our new analytical framework in Section IV, and present an exact model that quantifies the manner in which extended offsets induce priority in OBS systems. In Section V, we derive accurate expressions for both the blocking probability of the highest priority class, and the sensitivity of this performance with respect to variations in low-class traffic load and offset jitter. Simulation results verifying the correctness and accuracy of the model are presented in Section VI, and the model is used to perform a case study examining the effect of wavelength number and offset size. Finally, future work is discussed and conclusions are drawn in Section VII.

II. Notation

The following notation will be used throughout this paper.

- λ_i is the average arrival rate of class-i bursts.
- L_i is a random variable representing the length of class-i bursts. F_{L_i} denotes its distribution function, f_{L_j} denotes its probability density function (pdf), $\overline{L}_i = E[L_i]$ denotes its mean, and $L_i^{\max}$ denotes its maximum value (if one exists).
- Ω_i is the size of the extended QoS offset of each class-i burst (it is assumed that the size of the header-processing offset is negligible, compared with Ω_i).
- $\delta_{ij} = \Omega_i - \Omega_j$ is the difference between size of the class-i and class-j offsets.
- $\Delta_{ij} = \delta_{ij}/\overline{L_j}$ is the normalized offset difference between class-i and class-j.
- β_{ij} denotes the number of class-j bursts in the contention window of a class-i burst. (The contention window is defined in Section IV.)
- f_{β_i} is the probability mass function (pmf) of the overall number of bursts in the contention window of an arbitrary class-i burst.
- $\eta_{ij|l_i}$ is the average number of class-j bursts in the contention window of a class-i burst of length l_i.
- $\eta_{i|l_i}$ is the overall average number of bursts in the contention window of a class-i burst of length l_i.
- Pb_i is the average burst-blocking probability of class-i.

Additionally, for a given burst-A, we say that:

- τ_A is the arrival instant of its control packet;
- Ω_A is the duration of its offset;

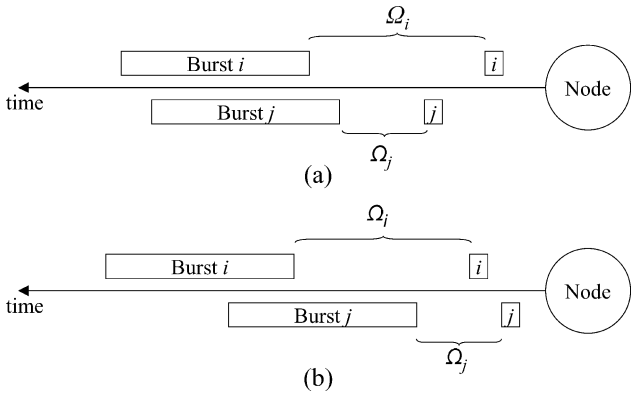

Fig. 1. Achieving priority with QoS offsets in JET OBS networks. In (a), burst A successfully reserves the channel due to its longer QoS offset, even though its start time is later than that of burst B. In a system without class isolation, high-priority bursts may be blocked by low-priority bursts, as shown in (b), where burst B successfully reserves the channel.

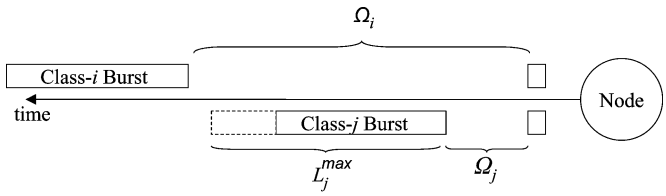

Fig. 2. Class-i offset is larger than the maximum burst length of class-j, so class-i is isolated from class-j.

- l_A is the duration of its burst payload;
- $H_A \equiv \tau_A + \Omega_A$ is the *head* or *start* of the burst;
- $\varepsilon_A \equiv \tau_A + \Omega_A + l_A$ is the *tail* or *end* of the burst.

III. QoS via Extended Offsets

Class differentiation in JET OBS networks can be implemented by assigning a *QoS offset* to higher class bursts [6]. In particular, bursts with longer offsets will tend to have higher priority than bursts with shorter offsets. To illustrate this, we describe a simple two-class system in Fig. 1. In Fig. 1(a), burst A and burst B arrive at the node at nearly the same time, and desire to be switched to the same outgoing channel. In a system with no offsets, bursts are serviced in the order of their arrival (first come, first serve) so burst B will successfully reserve the channel. However, in the system in Fig. 1(a), since burst A has a larger offset, its control packet arrives first and successfully reserves the channel before burst B. When the control packet of burst B arrives, it will see that the channel has already been reserved. In this way, the extended offset of burst A gives it implicit priority over burst B. This scheme can be extended to an arbitrary number of classes by assigning each class a longer offset than the offsets of all lower priority classes [12].

Despite having a shorter offset, burst B may still successfully reserve the bandwidth if its control packet arrives before that of burst A, as illustrated in Fig. 1(b). In certain systems, however, this event can be made impossible if the burst-A offset is made sufficiently large. More specifically, if the offset difference δ_{ij} between class i and class j is chosen to be longer than the maximum class-j burst length, class-i traffic will be completely isolated from class-j traffic [8]. This is illustrated in Fig. 2. Al-

though this result has been described intuitively in previously studies, it has never been formally proven. In Appendix II, we prove it in the context of our new analytical model framework.

While QoS offsets are an effective method for increasing the priority of premium-class traffic, the resulting improvement in blocking performance comes at the cost of increased premium-class delay. Although this delay is typically quite small (on the order of milliseconds), depending on the latency requirements of the premium-class traffic and the other sources of delay in the system, the maximum offset size could be limited. In [13], a framework is presented for selecting the premium-class offset size and the burst-assembly mechanism, such that both the premium-class blocking and delay requirements are simultaneously satisfied.

IV. Analytical Model

We now describe a general and exact analytical model that quantifies the effect of QoS offsets on the burst-contention process in multiclass OBS systems.

A. Model Framework: Contention Window and Maximum Perceived Load

In this section, we introduce the concept of the *contention window*, the *instantaneous perceived load*, and the *maximum perceived load* of a burst. We also prove that, for a burst from the highest priority class, the maximum perceived load and the number of bursts in the burst's contention window are equal.

Assume that there are two bursts arriving at nearly the same time at a node. If the first burst's control packet arrives before that of the second burst, and the two bursts overlap, the first burst is said to be in the *contention window* of the second burst. The contention window is the framework with which the rest of the results in this paper will be derived. Using the notation from Section II, we define the contention window of a class-A burst as follows.

Definition 1: Burst B is said to be in the *contention window* of burst A if all three of the following conditions are satisfied:

$$\tau_B < \tau_A \tag{1}$$

$$\tau_B + \Omega_B < \tau_A + \Omega_A + l_A \tag{2}$$

$$\tau_B + \Omega_B + l_B > \tau_A + \Omega_A. \tag{3}$$

We say that burst B *contends* with burst A. The conditions in (2) and (3) can be restated in terms of the start and end times of burst A and burst B as $H_B < \varepsilon_A$ and $\varepsilon_B > H_A$, respectively.

The contention window is a good reflection of how a burst "sees" the rest of the bursts in the system. By increasing the offset of class-i bursts, we reduce the number of bursts that contend with it, which is tantamount to increasing its priority. We quantify the relationship between offset size and contention-window occupancy with an exact model in the next section.

For the purpose of computing blocking probabilities, one would also like to know the maximum number of bursts that simultaneously overlap within the duration of a given burst A. To aid us in finding this value, we make the following definitions.

Definition 2: Given a time instance t that lies within the duration of a given burst A, the *instantaneous perceived load* of burst A at time t is denoted by $\ell_A(t)$, and is equal to the number of bursts that contend with burst A, and that overlap time t. Thus, $\ell_A(t)$ can be expressed as

$$\ell_A(t) = \sum_{i \in S} \phi_{A,i}(t), \quad t \in (H_A, \varepsilon_A) \tag{4}$$

where S is the set of all bursts in the system, and $\phi_{A,i}(t)$ is defined as

$$\phi_{A,i}(t) = \begin{cases} 1, & (\tau_i < \tau_A) \bigcap (H_i < t) \bigcap (\varepsilon_i > t) \\ 0, & \text{otherwise.} \end{cases} \tag{5}$$

Definition 3: The *maximum perceived load* of a given burst A is the maximum number of contending bursts that overlap any subinterval of the burst's duration. Thus, the maximum perceived load $\ell_A^{\max}$ can be expressed as

$$\ell_A^{\max} = \max_{t \in (H_A, \varepsilon_a)} \ell_A(t). \tag{6}$$

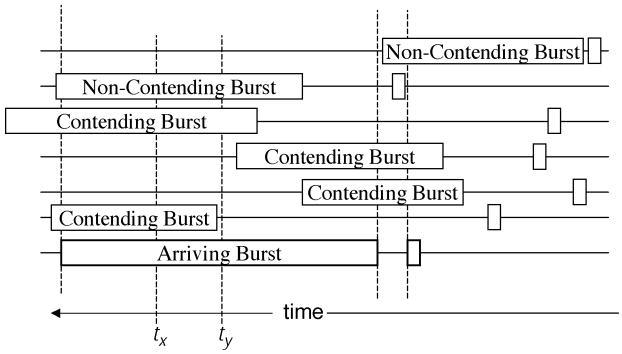

Fig. 3. For the system shown, there are four bursts in the arriving burst's contention window. The arriving burst's instantaneous perceived loads at t_x and t_y are two and one, respectively. The arriving burst's maximum perceived load is two.

The concepts of contention window, instantaneous perceived load, and maximum perceived load are illustrated in Fig. 3. In the figure, there are four bursts in the contention window of the arriving burst. The instantaneous perceived load at time t_x and time t_y are two and one, respectively. Since there are never more than two contending bursts simultaneously overlapping the arriving burst, the maximum perceived load is two.

We will now prove some useful results that describe the manner in which low-priority bursts can contend with high-priority bursts.

Theorem 1: Suppose that a class-j burst is in the contention window of a class-i burst, and that class j is of equal or lower priority than class i. Then the class-j burst overlaps the head of the class-i burst.

Proof: If a class-j burst is in the contention window of a class-i burst, from (1), we have that $\tau_j < \tau_i$. Additionally, if class j is of lower or equal priority than class i, we have $\Omega_j \leqslant \Omega_i$. Thus, we have that $\tau_i + \Omega_i > \tau_j + \Omega_j$. From (3), we also have that $\tau_i + \Omega_i < \tau_j + \Omega_j + l_j$. Thus, $\tau_i + \Omega_i \in (\tau_j + \Omega_j, \tau_j + \Omega_j + l_j)$, so the head of the class-i burst lies between the head and the tail of the class-j burst. ■

Corollary 1: Assume that class i has higher priority than class j, and that there are n class-j bursts in the contention window of a class-i burst. Then all n of them overlap the head of the class-i burst. The proof follows directly from *Theorem 1*.

Theorem 2: Assume that class i has the highest priority of all classes in the system. Then the maximum perceived load of a class-i burst is equal to the total number of bursts that lie in its contention window.

Proof: Assume that exactly n contending bursts overlap the head of the arriving class-i burst. Also assume that its maximum perceived load is m. Thus, we must have $n \leqslant m$. If $n < m$, then there exists at least one burst in the contention window that does not overlap the head of the class-i burst. By *Theorem 1*, since class i is the highest priority class, this is not possible. Thus, we must have that $n = m$. From *Corollary 1*, it follows that the number of bursts in the contention window of the class-i burst is also equal to m. ∎

Thus, for the highest priority class in the network, the maximum perceived load is exactly equal to the total number of bursts in the contention window. This result enables us to derive an expression for the blocking probability of the highest priority class of a multiclass OBS network in Section V.

B. Exact Analytical Model of Burst Overlap Process

In this section, we derive an exact analytical model for the distribution of the number of bursts in the contention window of a given class-i burst. We start by deriving an expression for the probability that there are m class-j bursts in the contention window of a class-i burst of length l_i. In general, an expression for this distribution is elusive, because of the dependency between the arrival instants of the m class-j bursts. However, assuming that bursts arrive according to a Poisson process, we can render the arrival instants independent and identically distributed by conditioning on the event that the class-j bursts arrive in a given time interval [14]. More specifically, if we condition on the event that k class-j bursts arrive in the interval $[\tau_i - \Delta/2, \tau_i + \Delta/2]$, the resulting conditioned arrival instants are independent and identically, uniformly distributed over that interval. We can then proceed to find the desired distribution based on the probability that any one of these k bursts lie in our class-i contention window.

We define β_{ij} as the number of class-j bursts that lie in the contention window of an arbitrary class-i burst, and y_{ij}^{Δ} as the number of class-j arrivals in $[\tau_i - \Delta/2, \tau_i + \Delta/2]$, so we have

$$\mathrm{P}[\beta_{ij}=m|L_i=l_i]=\sum_{k=0}^{\infty}\mathrm{P}\left[\beta_{ij}=m|L_i=l_i, y_{ij}^{\Delta}=k\right]\mathrm{P}\left[y_{ij}^{\Delta}=k\right]. \tag{7}$$

Without loss of generality, we select Δ to be sufficiently large $(\Delta \to \infty)$ such that there is negligible probability that a class-j burst that arrives outside of $[\tau_i - \Delta/2, \tau_i + \Delta/2]$ lies in the class-i burst's contention window. Then the only contributing terms in (7) are those for which $k \geq m$, so we write

$$\mathrm{P}[\beta_{ij}=m|L_i=l_i]=\sum_{k=m}^{\infty}\mathrm{P}\left[\beta_{ij}=m|L_i=l_i, y_{ij}^{\Delta}=k\right]\mathrm{P}\left[y_{ij}^{\Delta}=k\right]. \tag{8}$$

The conditioned arrival instants and the lengths of each class-j burst are independent. The k arriving bursts can, therefore, be considered independently, and the number of bursts in the contention window follows a binomial distribution. Further, given that class-j bursts arrive according to a Poisson process with mean rate λ_j, we have

$$\mathrm{P}[\beta_{ij}=m|L_i=l_i]=\sum_{k=m}^{\infty}\binom{k}{m}\left(A_{ij|\Delta,l_i}\right)^m \times\left(1-A_{ij|\Delta,l_i}\right)^{k-m}\frac{(\lambda_j\Delta)^k}{k!}e^{-\lambda_j\Delta} \tag{9}$$

where we define $A_{ij|\Delta,l_i} \equiv P[\beta_{ij}=1|L_i=l_i, y_{ij}^{\Delta}=1]$. After grouping terms and shifting indexes, we are left with

$$\mathrm{P}[\beta_{ij}=m|L_i=l_i]=\frac{\eta_{ij|l_i}^m e^{-\eta_{ij|l_i}}}{m!} \tag{10}$$

where we define $\eta_{ij|l_i} \equiv \lambda_j \Delta A_{ij|\Delta,l_i}$ as the mean number of class-j bursts in the contention window of a class-i burst of length l_i.

Equation (10) implies that the number of class-j bursts in the contention window of a given class-i burst with length l_i is distributed according to a Poisson random variable, with mean equal to $\eta_{ij|l_i}$. Since the burst-arrival processes of each class are independent of each other, and since the sum of Poisson random variables is also Poisson, we have the following conclusion.

> If bursts from each class in a multiclass OBS system arrive according to a Poisson process, the distribution of the number of class-j bursts and the distribution of the overall number of bursts in the contention window of a class-i burst are Poisson random variables, with means of $\eta_{ij|l_i}$ and $\eta_{i|l_i} = \sum_j \eta_{ij|l_i}$, respectively.

This result is true for an arbitrary number of classes with arbitrary QoS offset values and arbitrary burst-length distributions. This includes not only multiclass systems, but also systems in which bursts have different offset values corresponding to the amount of residual header-processing time.

Expressions for the mean occupancy $\eta_{i|l_i}$ can be derived as (see Appendix I)

$$\eta_{ij|l_i}=\begin{cases}\lambda_j\cdot\overline{L_j}\cdot\left[1-R_{L_j}(\delta_{ij})\right], & \delta_{ij}\geq 0\\ \lambda_j\cdot(\overline{L_j}+l_i), & \delta_{ij}<0,\ l_i<-\delta_{ij}\\ \lambda_j\cdot(\overline{L_j}-\delta_{ij}), & \delta_{ij}<0,\ l_i\geq-\delta_{ij}\end{cases} \tag{11}$$

where $\overline{L_j}$ is the mean burst length of class j, δ_{ij} is the offset difference between class i and class j, and $R_{L_j}(\delta_{ij})$ is the distribution function for the residual life of the class-i burst-length distribution.

Equation (11) implies that the contention window occupancy of a given low-class burst depends on its length. In multiclass OBS systems, short low-class bursts will experience better performance than long low-class bursts because they can more easily fit into the spaces (or voids) left by premium-class reservations. Although this low-class burst-selection effect has been observed in previous OBS simulation studies [12], to our knowledge, this paper is the first to provide a framework in which this effect can be analytically quantified. Graphs that clearly illustrate this effect are included in our numerical study in Section VI.

Last, by integrating the expression in (10) over all possible burst lengths, we can yield the following pmf for the total

number of bursts in the contention window of class-i bursts averaged over all burst lengths:

$$f_{\beta_i}(m) = \int_{l_i} \frac{\eta_{i|l_i}^m e^{-\eta_{i|l_i}}}{m!} f_{L_i}(l_i) dl_i. \tag{12}$$

The model presented above is based purely on the traffic characteristics of the OBS system and makes no assumption about the OBS node architecture or the contention resolution capabilities of each node. The model's sole assumption is that bursts from each class arrive according to a Poisson process. In addition to allowing for tractable results, this assumption has also been found to be a reasonably valid for a large number of practical OBS systems. In particular, in systems that use timer-bases burst aggregation [15]–[17], the aggregate burst-arrival process has been found to closely resembles a Poisson process [18], [19].

V. Performance Modeling of Premium-Class Traffic

In this section, we apply the model in Section IV to derive an analytical expression for the burst-dropping probability of the highest priority class of traffic. We then derive expressions for the sensitivity of the highest-priority class's performance with respect to offset jitter and variations in the arrival rate of each class.

We assume that traffic arrives from C independent traffic streams, each corresponding to a different class and with a different QoS offset. Without loss of generality, we assume that $\Omega_i > \Omega_j$ for all $i < j$, such that class 0 has the highest priority, class 1 has the second highest, and class $C-1$ has the lowest. The total number of wavelength channels in the system is W. In this section, we examine the performance of OBS systems with full wavelength conversion and no fiber delay line (FDL) buffers.[1]

A. Estimating Premium-Class Burst-Loss Probability

In the system under consideration, a burst from the highest class is blocked if its control packet arrives to find that no single channel is available for the entire duration of its burst. This will occur if there is at least one subinterval in the burst duration over which all W channels are occupied. In networks in which $Pb_i \ll 1$ for all i, this can be closely approximated by the event that the maximum perceived load of the burst is greater than or equal to W.

By applying *Theorem 2*, we can approximate the burst-dropping probability of class 0 by the probability that W or more bursts lie in the contention window of an arbitrary class-0 burst. Using (12), we have

$$\begin{aligned} Pb_0 &\cong \sum_{k=W}^{\infty} f_{\beta_0}(k) \\ &= 1 - \sum_{k=0}^{W-1} \int_{l_0=0}^{\infty} \frac{\eta_{0|l_0}^m e^{-\eta_{0|l_0}}}{m!} f_{L_0}(l_0) dl_0. \end{aligned} \tag{13}$$

[1]Although we only consider full wavelength conversion systems here, the model in Section IV is directly applicable to any OBS system, regardless of the contention resolution strategy employed. The study of more general systems, such as those with partial wavelength conversion and/or FDL buffering, is left for future study.

From (11), it is evident that $\eta_{0|l_0}$ does not actually depend on l_0, so we can easily perform the integration in (13) to yield

$$Pb_0 = 1 - \sum_{k=0}^{W-1} \frac{\eta_0^m e^{-\eta_0}}{m!} \tag{14}$$

where we have replaced $\eta_{0|l_0}$ by η_0 for added clarity.

Thus, the probability of blocking for the highest class in a multiclass OBS system in which bursts arrive according to a Poisson process can be closely approximated by the complementary distribution function of an appropriately parameterized Poisson random variable.

B. Sensitivity Analysis

In this section, we derive direct expressions for the *sensitivity* of premium-class traffic, with respect to changes in the arrival rates in the system and to offset jitter.

1) Sensitivity to Lower Class Arrival-Rate Variation: When provisioning an OBS system, it may not be possible to exactly predict the arrival rate of each class *a priori*. Also, the arrival rate of each class may vary over short time scales. Thus, the sensitivity of the premium-class performance with respect to fluctuations in the lower class arrival rates provides an additional metric to measure the robustness of the QoS provided and the level of premium-class isolation.

We can derive an expression for the sensitivity of the class-0 blocking rate with respect to the class-j arrival rate as follows:

$$\frac{\partial Pb_0}{\partial \lambda_j} = \frac{\partial Pb_0}{\partial \eta_0} \cdot \frac{\partial \eta_0}{\partial \lambda_j}. \tag{15}$$

Substituting (14) and (11) into (15) yields

$$\frac{\partial Pb_0}{\partial \lambda_j} = \frac{\eta_0^{W-1} e^{-\eta_0}}{(W-1)!} \cdot \overline{L_j} \cdot \left[1 - R_{L_j}(\delta_{0j})\right]. \tag{16}$$

2) Sensitivity to Offset Jitter and Offset Variation: In real OBS systems, the size of the offsets of a burst will shrink after each hop by an amount equal to the electronic processing time of the control packet. Thus, at each node, burst from a given class do not all have the same offset, but follow some discrete offset distribution that depends of the residual path length of all arriving bursts of that class. Furthermore, queueing of control packets may contribute a further source of offset jitter in systems with a large number of control packets. In certain systems, this offset variability could lead to a degradation of premium-class performance. We can write an expression for the sensitivity of the premium-class blocking probability with respect to variations in offset size, as follows:

$$\begin{aligned} \frac{\partial Pb_0}{\partial \delta_{0j}} &= \frac{\partial Pb_0}{\partial \eta_0} \cdot \frac{\partial \eta_0}{\partial \delta_{0j}} \\ &= \frac{\eta_0^{W-1} e^{-\eta_0}}{(W-1)!} \cdot \lambda_j \left[1 - F_{L_j}(\delta_{0j})\right]. \end{aligned} \tag{17}$$

When class i is isolated from class j, $\delta_{ij} > L_j^{\max}$. Thus, $R_{L_j}(\delta_{0j}) = 1$ and $F_{L_j}(\delta_{0j}) = 1$, so the sensitivity with respect to arrival-rate variation and to offset variation are both zero, as expected.

VI. Numerical Results and Model Validation

In this section, we use simulation to verify the correctness of the exact analytical model presented in Section IV and the

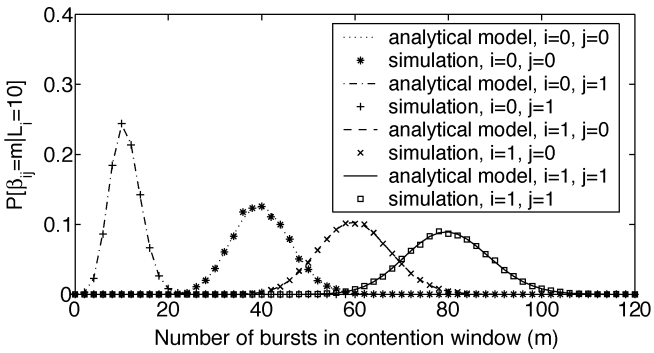

Fig. 4. PMF for the number of class-j bursts in the contention window of a class-i burst of length $L_i = 10$ for $ij \in \{00, 01, 10, 11\}$. Bursts are exponentially distributed with mean lengths for class 0 and class 1 of 20 and 40, respectively, arrival rates of 2 and 4, respectively, and an offset difference of $\delta_{01} = 40$.

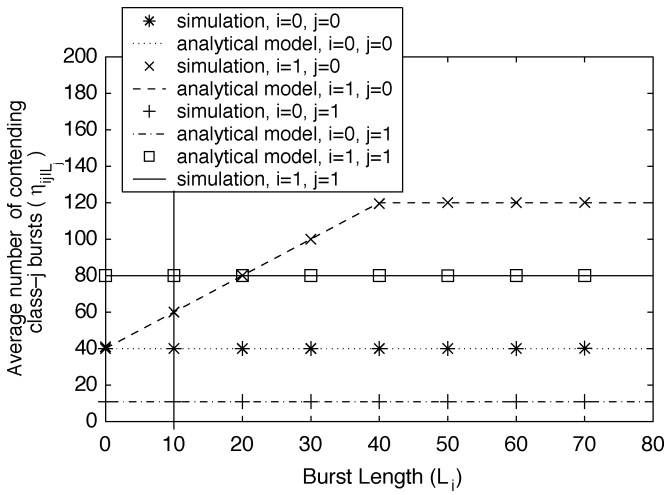

Fig. 5. Average number of class-j bursts in the contention window of a class-i burst for different values of class-i burst length L_i for $ij \in \{00, 01, 10, 11\}$. Traffic parameters are the same as those in Fig. 4.

accuracy of the premium-class blocking formulas presented in Section V. We then present a design case study, in which we vary the size of the premium-class offset and the number of wavelengths in the system.

For each simulation in this section, two classes and a total of one hundred million bursts were simulated. Bursts arrived according to a Poisson process. and both Gaussian and exponential burst-length distributions were simulated.

A. Accuracy of Contention Model

In Section IV, we showed that the number of bursts in a given burst's contention window follows a Poisson distribution, whose mean is a function of the offset sizes, arrival rates, and burst-length distributions of each traffic class. In order to verify this conclusion and the correctness of our exact model, we simulated an OBS system with two classes. Bursts from class 0 and class 1 were exponentially distributed with mean lengths of 20 and 40, respectively; bursts from class 0 and class 1 arrived according to a Poisson process with means of 2 and 4, respectively; and the offset difference between class 0 and class 1 was 40. The results are shown in Fig. 4, where we graph $P[\beta_{ij} = m | L_i = 10]$, the pmf of the number of class-j bursts in the contention window of arriving class-i bursts of length 10. Both simulation results[2] and analytical results based on (10) and (11) are shown for all four combinations of i and j. The effect of the QoS offset difference between the two classes is clear. Even though the low-class load is four times that of the premium class, far fewer low-class bursts contend with high-class bursts due to the premium-class QoS offset, as the curve for β_{01} is significantly lower than the curve for β_{00}. The simulation results match the Poisson distribution curves from the analytical model extremely well, thus verifying the correctness of the model.

B. Low-Class Burst-Length Dependence

The results in Fig. 4 give the pmf for the number of bursts in the contention window for bursts of length $L_i = 10$. One may also be interested in examining similar curves for different burst lengths. In order to examine how a burst's length affects its performance, in Fig. 5, we plot $\eta_{ij|l_i}$, the average number of class-j bursts in the contention window of a class-i burst, for different values of burst length l_i. The intersections of the curves and the vertical line at $l_i = 10$ correspond to the means of the distributions illustrated in Fig. 4.

For all curves in Fig. 5, the simulation matches the analysis almost exactly. The curves corresponding to $\eta_{00|l_0}$, $\eta_{01|l_0}$, and $\eta_{11|l_1}$ are constant-valued, and the curve corresponding to $\eta_{10|l_1}$ increases linearly until $l_1 = \delta_{01} = 40$, where it becomes constant. These results are expected from the expressions in (11), and graphically illustrate the mechanism by which short low-class bursts are favored over long low-class bursts in multiclass OBS networks.

C. Accuracy of Premium-Class Blocking-Probability Formulas

In this section, we examine the accuracy of the expression for the class-0 burst-dropping probability presented in Section V-A. We simulated an OBS system with eight wavelengths, full wavelength conversion, no FDL buffers, and two traffic classes. Bursts from class 0 and class 1 arrived according to a Poisson process. The burst-length distribution of class 0 followed a Gaussian distribution, with a mean of 1 and a variance of 0.1. The distribution of class-1 burst lengths was Gaussian, with mean and variance of 2 and 0.2, respectively.[3]

For Fig. 6, the offset difference between class 0 and class 1 was 1 (i.e., 50% of the class-1 mean burst length). The figure plots the simulated burst-blocking probability of each class and the analytical blocking probability for class 0, as the arrival rate of the system is increased, while keeping the ratio of the class-1 and class-2 arrival rates fixed at 2. Using these traffic parameters, class 1 is not isolated from class 0, and the system is non-work-conserving.

[2]Because the burst-length distribution is continuous, each simulation curve in Fig. 4 corresponds to $L_i = 10 \pm 0.1$.

[3]The use of Gaussian burst-length distributions implies that negative burst lengths were possible. To avoid this nonsensical scenario, the simulator discarded any negative bursts that were generated. However, for the mean and variance used, the probability of a negative burst was extremely small (approximately 10^{-23}), so this had negligible impact on the system's performance or the models' accuracy.

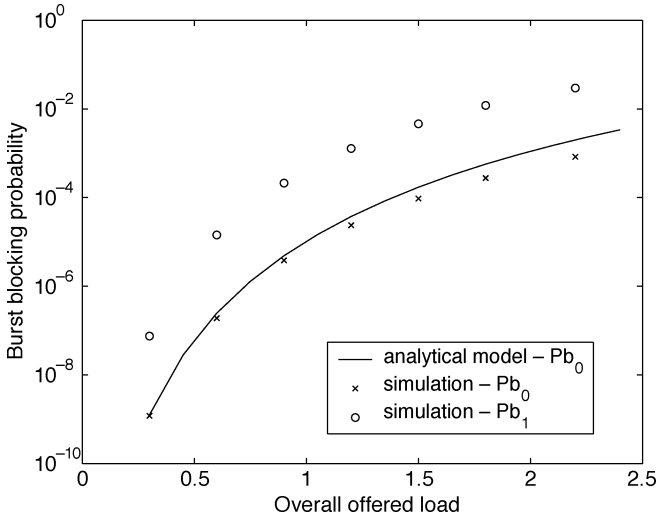

Fig. 6. Burst-blocking probability for each class versus overall traffic load for an eight-wavelength system with two classes of traffic and Gaussian burst-length distributions. The model is extremely accurate when $Pb_0 \ll 1$ and $Pb_1 \ll 1$, and reasonably accurate for higher burst-blocking probability values.

When the system is operating at low-to-moderate burst-dropping probabilities, the analytical formula predicts the premium-class burst-dropping probability very accurately. The predictions become less accurate as the traffic load increases and the low-class burst-blocking probability approaches one. This is expected, as the single assumption made in deriving (14) was that $Pb_i \ll 1$ for all i. However, even in this regime, the predictions from the analytical formula are still reasonably accurate, differing from the simulation by less than half of an order of magnitude.

Fig. 7(a) plots results from a simulation with similar parameters as those in Fig. 6, with the class-0 and class-1 arrival rates fixed at 0.5 and 1, respectively. We varied the offset difference between class 0 and class 1, and graphed the resulting premium-class burst-blocking probability Pb_0 versus the normalized offset difference $\Delta_{01} \equiv \delta_{01}/\overline{L_1}$. When the offset difference between the classes is zero, Pb_0 and Pb_1 are equal. As the offset of the premium class is increased, it becomes more isolated from the lower class, and its blocking probability decreases accordingly. When the normalized offset difference is larger than 1.6, increasing the class-0 offset has very little effect on its blocking probability, as class 0 is almost completely isolated from class 1. For all points simulated, the analytical model predicts the premium-class blocking probability very accurately.

Fig. 7(a) also plots $\partial Pb_0/\partial\lambda_1$, the sensitivity of class 0 with respect to variations in the class-1 arrival rate. As mentioned in Section V-B, $\partial Pb_0/\partial\lambda_1$ can be used as a good measure of the robustness of the offset-based class-differentiation scheme, and a quantitative measure of the isolation of class 0. From Fig. 7(a), one can conclude that class 0 is well isolated from class 1 when $\partial Pb_0/\partial\lambda_1$ is between one and two orders of magnitude lower than Pb_0.

Fig. 7(b) plots results from a simulation with the same parameters as those in Fig. 7(a), except that burst lengths followed an exponential distribution. Again, the model's predictions match the simulation results very closely. Although the curves are similar to those in Fig. 7(a) when Δ_{01} is large enough that class 0

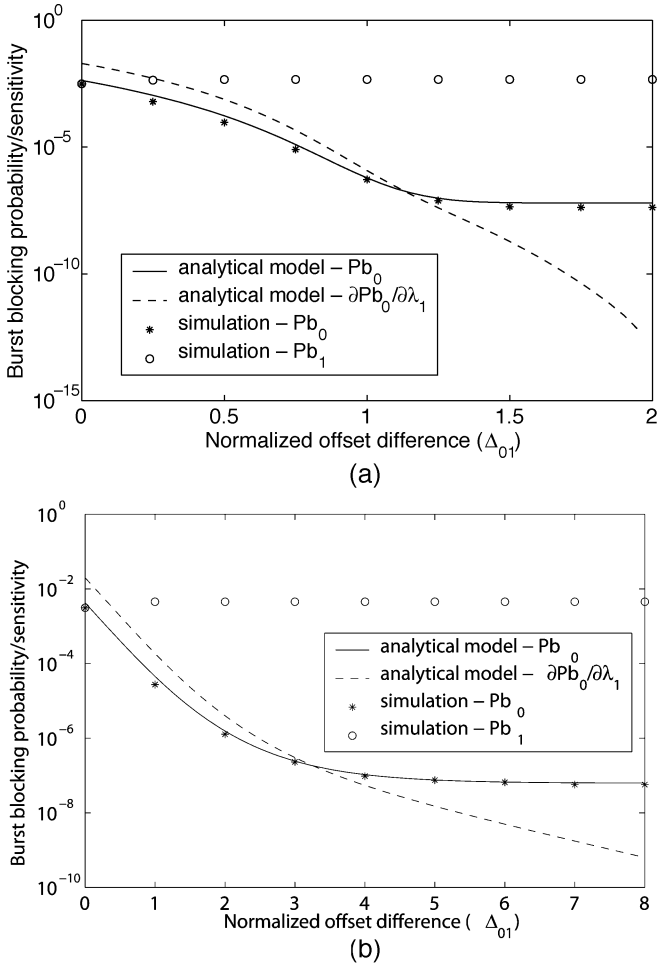

Fig. 7. Burst-blocking probability versus normalized offset difference in a two-class OBS system. Bursts have Gaussian-distributed burst lengths in (a), and exponentially distributed burst lengths in (b).

is isolated from class 1, it is interesting to note the difference in the partial-isolation regime. Because the tail of the exponential distribution is heavier than that of the Gaussian distribution, the probability of long class-1 bursts is higher. As such, to achieve near-complete class isolation, a value of $\Delta_{01} > 6$ is required in Fig. 7(b), compared with $\Delta_{01} > 1.6$ in Fig. 7(a). Again, isolation occurs when the sensitivity with respect to λ_1 is between one and two orders of magnitude lower than Pb_0.

The model presented in Section V-A is approximate, because it ignores the effect of second-order blocking in the system. As such, it gives an upper bound on the blocking probability of the premium class. This is verified by the curves in Figs. 6 and 7. The fact that the formula gives a tight upper bound makes it an attractive tool for providing worst-case, premium-class burst-blocking probability guarantees through QoS offset provisioning.

D. Design Case Study: The Effect of the Network Scaling on Premium-Class Performance

In this section, we use our model to examine the effect of network scaling on the performance and the level of isolation of premium-class traffic. We simulated an OBS system with two classes of traffic and Gaussian burst-length distributions. The mean burst

lengths of class 0 and class 1 were 1 and 2, respectively, and the variances of the class-0 and class-1 burst lengths were 0.1 and 0.2, respectively. Bursts arrived according to a Poisson process, and the arrival rate of class 1 was twice that of class 0. While varying the number of wavelengths in the system from 8 to 24, the load of each class was scaled to keep the load per wavelength constant at 0.625. Because of the large number of data points and the extremely low blocking probabilities, simulation would have been prohibitively time consuming for this analysis, so having an accurate analytical model was essential.

In Fig. 8(a), we plot the premium-class burst-dropping probability as a function of the normalized offset difference Δ_{01} and the number of wavelengths in the system. The burst-blocking probability decreases as we increase the offset, as was similarly observed in Section VI-C. It also decreases as the number of wavelengths is increased due to the well-known network-scaling effect.

Fig. 8(b) plots $\partial Pb_0/\partial\lambda_1$, the sensitivity of the class-0 burst-dropping probability with respect to variations of the low-class arrival rate versus the normalized offset difference and the number of wavelengths in the system. As Δ_{01} increases, class 0 becomes more isolated from class 1, and the sensitivity decreases. Also, as the number of wavelengths increases, the sensitivity $\partial Pb_0/\partial\lambda_1$ decreases, especially for large values of normalized offset difference.

Thus, we can conclude that scaling up an OBS network improves not only the burst-dropping probability of the premium class, but also the level of isolation between the high- and low-class bursts. This conclusion is especially important for systems in which the premium-class traffic has a low delay tolerance. In such systems, where the size of the premium-class offset is restricted, designers could use aggregation to scale up the network, and thereby improve the level of isolation between high- and low-class traffic without increasing the high-class offset.

To illustrate how the results in Fig. 8 could be used to dimension and provision an OBS system, we present the following network-design problem. Assume that we require that the premium-class burst-blocking probability is less than 10^{-10}. In order to ensure that the system was robust with respect to variations in the low-class traffic load, we also require that the sensitivity with respect to the low-class arrival rate is less than 10^{-12}. These design constraints are represented by the horizontal planes in Fig. 8(a) and (b). Thus, the allowable region of operation can be found by taking the intersection of the regions where the curves are below these planes. The resulting set of allowable offset-wavelength combinations is shown in Fig. 8(c). We observe that at least 18 wavelengths are needed, and that increasing the number of wavelengths decreases the size of the minimum required premium-class offset. The 18-wavelength system requires a normalized offset difference of at least 1.5 to meet the premium-class blocking and sensitivity requirements, while the 24-wavelength system requires a normalized offset value of only 1.1.

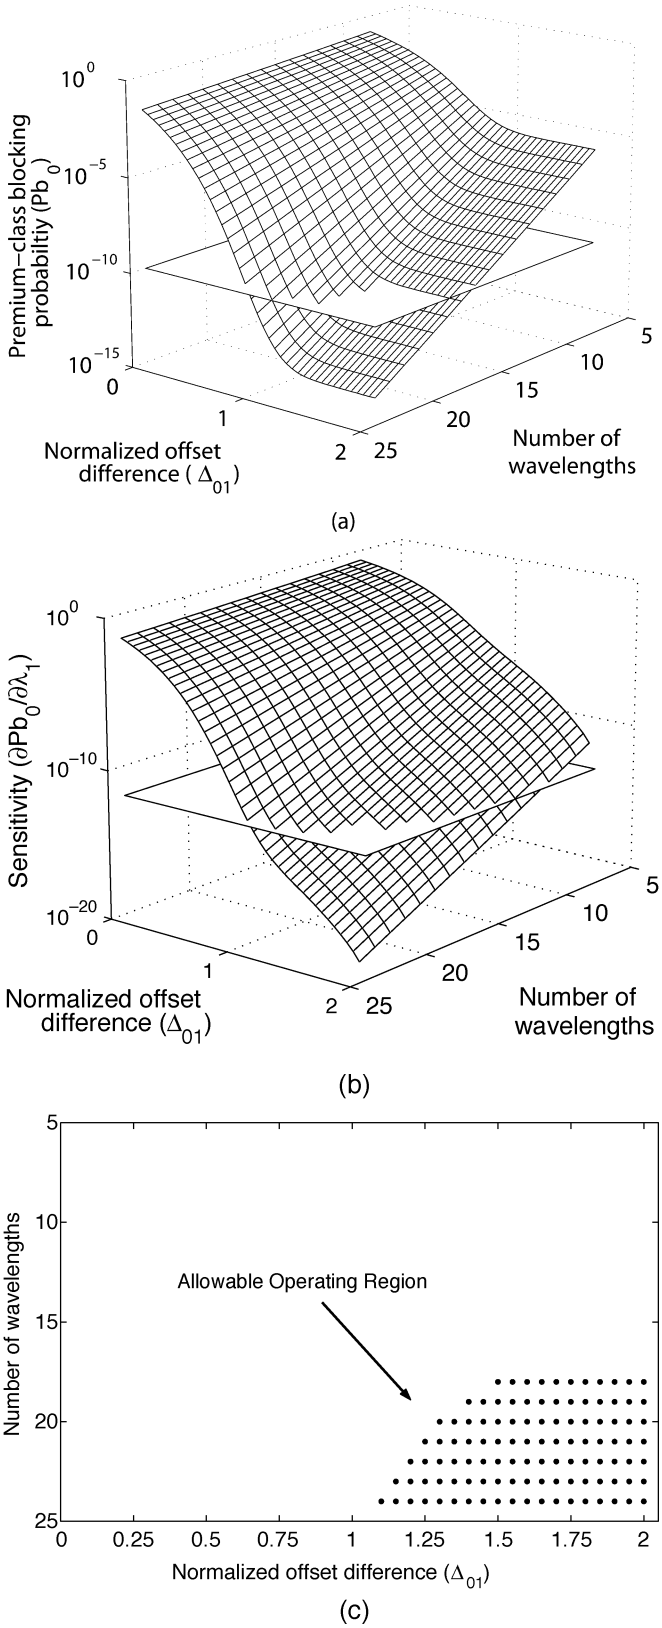

Fig. 8. Performance of high-class traffic as a function of normalized offset difference Δ_{01} and number of wavelengths in a two-class OBS system. Burst lengths have a Gaussian distribution. The blocking probability Pb_0 is shown in (a), and the sensitivity with respect to low-class arrival rate $\partial Pb_0/\partial\lambda_1$ is shown in (b). The allowable region of operation for a blocking probability less than 10^{-10}, and a sensitivity less than 10^{-12} is shown in (c).

VII. Conclusions and Future Work

In this paper, we proposed for the first time an analytical model that quantifies the mechanism by which QoS offsets induce priority in multiclass OBS systems. We introduced the concepts of the burst contention window and the maximum per-

ceived load of a burst, and we showed that the distribution of the number of bursts that contend with a given burst follows a Poisson distribution. We then used our model to derive accurate formulas for the blocking probability of premium-class traffic and the sensitivity of the premium-class blocking probability to variations in the arrival rates and offsets of each class.

We verified the correctness of our model and the accuracy of our burst-blocking formula using simulation, and performed a case study examining the effect of network scaling on premium-class performance. We found that scaling-up a system tends to improve the performance of the premium-class traffic, both in terms of the burst-blocking probability achieved and of the level of isolation achieved for a given set of QoS offsets.

Future work in involves applying the model to systems with partial wavelength conversion, systems in which limited FDL buffering is available, and multinode OBS networks.

APPENDIX I
AVERAGE CONTENTION-WINDOW OCCUPANCY

Here we derive an expression for the mean number of class-j bursts that lie in the contention window of a class-i burst, with length l_i for arbitrary classes i and j. Let T_j be a uniform random variable on the interval $[\tau_i - \Delta/2, \tau_i + \Delta/2]$ that represents the conditioned arrival time of a class-j burst. From our definition of $\eta_{ij|l_i}$ and using the definition of the contention window in Section IV-A, we have

$$\begin{aligned}\eta_{ij|l_i} &= \lambda_j \cdot \Delta \cdot \mathrm{P}\Big[T_j < \tau_i \bigcap \tau_j + \Omega_j < \tau_i + \Omega_i \\ &\qquad + l_i \bigcap T_j + \Omega_j + L_j > \tau_i + \Omega_i\Big] \\ &= \lambda_j \cdot \Delta \cdot \mathrm{P}\left[T_j \in (\tau_i + \delta_{ij} - L_j, \right. \\ &\qquad \left. \min(\tau_i, \tau_i + \delta_{ij} + l_i))\right]. \end{aligned} \quad (18)$$

To evaluate the right-hand side of (18), we divide the solution into two regions, corresponding to class-i having equal or higher priority than class-j ($\delta_{ij} \geqslant 0$), and class-i having lower priority than class-j ($\delta_{ij} < 0$).

1) $\delta_{ij} \geqslant 0$: Since $l_i > 0$, we have

$$\begin{aligned}\eta_{ij|l_i} &= \lambda_j \cdot \Delta \cdot \mathrm{P}\left[T_j \in (\tau_i + \delta_{ij} - L_j, \tau_i)\right] \\ &= \lambda_j \cdot \Delta \cdot \int_{l_j=0}^{\infty} \mathrm{P}[T_j \in (\tau_i + \delta_{ij} - l_j, \tau_i)] f_{L_j}(l_j) dl_j. \end{aligned} \quad (19)$$

The integrand in (19) will be zero when $l_j < \delta_{ij}$. Making use of the fact that the distribution of T_j is uniform over an interval of length Δ, we can write

$$\begin{aligned}\eta_{ij|l_i} &= \lambda_j \cdot \Delta \cdot \int_{l_j=\delta_{ij}}^{\infty} \frac{\tau_i - (\tau_i + \delta_{ij} - l_j)}{\Delta} f_{L_j}(l_j) dl_j \\ &= \lambda_j \left\{ \int_{l_j=\delta_{ij}}^{\infty} l_j f_{L_j}(l_j) dl_j - \delta_{ij}\left[1 - F_{L_j}(\delta_{ij})\right] \right\} \\ &= \lambda_j \left\{ \overline{L}_j - \int_{l_j=0}^{\delta_{ij}} l_j f_{L_j}(l_j) dl_j - \delta_{ij}\left[1 - F_{L_j}(\delta_{ij})\right] \right\}. \end{aligned} \quad (20)$$

We can make use of the convenient relation for nonnegative random variables [20]

$$\int_0^a x f_X(x) dx = a\left[F_X(a) - 1\right] + \int_0^a 1 - F_X(x) dx \quad (21)$$

to simplify (20) to yield the following expression for $\eta_{ij|l_i}$:

$$\eta_{ij|l_i} = \lambda_j \cdot \overline{L}_j \cdot \left[1 - R_{l_j}(\delta_{ij})\right] \quad (22)$$

where $R_{L_j}(l_a)$ is the distribution function for the residual life of L_j [21]

$$R_{L_j}(l_a) \equiv \frac{1}{\overline{L}_j} \int_0^{l_a} 1 - F_{L_j}(l_j) dl_j. \quad (23)$$

From *Theorem 3*, class i is isolated from class j when $\delta_{ij} > L_j^{\max}$. In this case, we have $R_{L_j}(\delta_{ij}) = 0$ and $\eta_{ij|l_i} = 0$, as expected.

2) $\delta_{ij} < 0$: The value of $\min(\tau_i, \tau_i + \delta_{ij} + l_i)$ will depend on the magnitude of $\delta_{ij} + l_i$. Proceeding directly from (18) and again making use of the fact that the distribution of T_j is uniform over an interval of length Δ, we have

$$\begin{aligned}\eta_{ij|l_i} &= \begin{cases} \lambda_j \cdot \Delta \cdot \int_{l_j=0}^{\infty} \frac{(l_i + l_j)}{\Delta} f_{L_j}(l_j) dl_j, & \delta_{ij} + l_i < 0 \\ \lambda_j \cdot \Delta \cdot \int_{l_j=0}^{\infty} \frac{(l_j - \delta_{ij})}{\Delta} f_{L_j}(l_j) dl_j, & \delta_{ij} + l_i \geqslant 0 \end{cases} \\ &= \begin{cases} \lambda_j \cdot (\overline{L_j} + l_i), & l_i < -\delta_{ij} \\ \lambda_j \cdot (\overline{L_j} - \delta_{ij}), & l_i \geqslant -\delta_{ij}. \end{cases} \end{aligned} \quad (24)$$

APPENDIX II
OBS ISOLATION THEOREM

If the offset size difference between a high and low class of traffic is larger than the maximum burst duration of the low class, then the high class will be completely isolated from the low class. While this result is stated in a large number of studies on multiclass OBS, and while a number of these justify its correctness using intuitive arguments, to to our knowledge, a formal proof has never been presented. Thus, we now present a simple proof using our contention-window framework.

Theorem 3: Given two classes i and j in a multiclass OBS system, if $\delta_{ij} > L_j^{\max}$, then no class-j burst can contend with a class-i burst. Class i is therefore *isolated* from class j.

Proof: Let $\delta_{ij} > L_j^{\max}$. If $\tau_j < \tau_i$, then we have that $\tau_j + \Omega_j + L_j < \tau_i + \Omega_j + L_j < \tau_i + \Omega_j + L_j^{\max} < \tau_i + \Omega_j + \delta_{ij} = \tau_i + \Omega_i$. Therefore, the conditions in (1) and (3) cannot be simultaneously satisfied when $\delta_{ij} > L_j^{\max}$. ■

REFERENCES

[1] T. Battestilli and H. Perros, "An introduction to optical burst switching," *IEEE Commun. Mag.*, vol. 41, pp. S10–S15, Aug. 2003.

[2] C. Qiao, "Labeled optical burst switching for IP-over-WDM integration," *IEEE Commun. Mag.*, vol. 38, pp. 104–114, Sep. 2000.

[3] T. Chen, C. Qiao, and X. Yu, "Optical burst switching (OBS): A new area in optical networking research," *IEEE Network*, vol. 18, pp. 16–23, May 2004.

[4] C. Qiao and M. Yoo, "Optical burst switching (OBS)—A new paradigm for an optical internet," *J. High Speed Network*, vol. 8, pp. 69–84, Mar. 1999.

[5] 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*, 1997, pp. 26–27.

[6] 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.
[7] M. Yoo and C. Qiao, "Supporting multiple classes of services in IP over WDM networks," in *Proc. Globecom*, Dec. 1999, pp. 1023–1027.
[8] H. L. Vu and M. Zukerman, "Blocking probability for priority classes in optical burst switching networks," *IEEE Commun. Lett.*, vol. 6, no. 5, pp. 214–216, May 2002.
[9] K. Dolzer, C. Gauger, J. Späth, and S. Bodamer, "Evaluation of reservation mechanisms for optical burst switching," *AEÜ Int. J. Electron. Commun.*, vol. 55, pp. 18–26, Jan. 2001.
[10] W.-H. So, Y.-H. Cha, S.-S. Roh, and Y.-C. Kim, "Offset time decision (OTD) algorithm for guaranteeing the requested QoS of high priority traffic in OBS networks," in *Proc. APOC*, Nov. 2001, pp. 286–295.
[11] N. Barakat and E. H. Sargent, "An accurate model for evaluating blocking probabilities in multiclass optical burst switching systems," *IEEE Commun. Lett.*, vol. 8, no. 2, pp. 119–121, Feb. 2004.
[12] K. Dolzer and C. Gauger, "On burst assembly in optical burst switching networks—A performance evaluation of just-enough-time," in *Proc. 17th Int. Teletraffic Congr.*, Sep. 2001, pp. 149–160.
[13] N. Barakat and E. H. Sargent, "On optimal ingress treatment of delay-sensitive traffic in multiclass OBS systems," in *Proc. 3rd Int. Workshop Opt. Burst Switching*, Oct. 2004, [CD-ROM].
[14] L. Kleinrock, *Queueing Systems—Vol. 1: Theory*. New York: Wiley, 1975.
[15] J. White, M. Zukerman, and H. L. Vu, "A framework for optical burst switching network design," *IEEE Commun. Lett.*, vol. 6, no. 6, pp. 268–270, Jun. 2002.
[16] Y. Xiong, M. Vandenhoute, and H. Cankaya, "Control architecture in optical burst-switched WDM networks," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 10, pp. 1838–1851, Oct. 2000.
[17] A. Ge, F. Callegati, and L. Tamil, "On optical burst switching and self-similar traffic," *IEEE Commun. Lett.*, vol. 4, no. 3, pp. 98–100, Mar. 2000.
[18] M. Izal and J. Aracil, "On the influence of self-similarity on optical burst switching traffic," in *Proc. Globecom*, Nov. 2002, pp. 2308–2312.
[19] T. Tachibana, T. Ajima, and S. Kasahara, "Round-robin burst assembly and constant transmission scheduling for optical burst switching networks," in *Proc. Globecom*, Dec. 2003, pp. 2772–2776.
[20] A. Leon-Garcia, *Probability and Random Processes for Electrical Engineering*, 2nd ed. Reading, MA: Addison-Wesley, 1994.
[21] G. R. Grimmett, *Probability and Random Processes*, 3rd ed. Oxford, U.K.: Oxford Univ. Press, 2001.

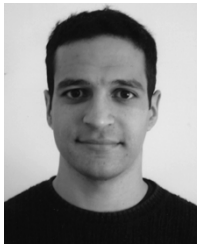

Neil Barakat (S'98) received the B.A.Sc. (Honors) degree from McGill University, Montreal, QC, Canada, in 1999, and the M.Sc. degree from the University of Toronto, Toronto, ON, Canada, in 2002, both in electrical engineering. He is currently working toward the Ph.D. degree in electrical engineering at the University of Toronto.

His current area of research is optical networks with a focus on architectures and protocols for burst-switched WDM networks.

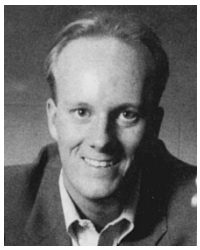

Edward H. Sargent (S'97–M'98–SM'03) received the B.Sc.Eng. degree in engineering physics from Queen's University, Kingston, ON, Canada, in 1995, and the Ph.D. degree in electrical and computer engineering (photonics) from the University of Toronto, Toronto, ON, Canada, in 1998.

He currently holds a Canada Research Chair at the University of Toronto, a position he was awarded in 2000 for creating "...a new type of laser that unites many sophisticated optical devices onto a single, integrated photonic chip. His research links the emerging concept of the photonic circuit with the exploding field of fiber optic networks."

Dr. Sargent was named "one of the world's top young innovators" by *MIT's Technology Review* in 2003. In 2002, he was honored by the Canadian Institute for Advanced Research as one of Canada's top 20 researchers under age 40. Also in 2002, he was the recipient of the Outstanding Engineer Award of the IEEE of of Canada "...for groundbreaking research in applying new phenomena and materials from nanotechnology toward transforming fiber-optic communications systems into agile optical networks." He received the 1999 NSERC Silver Medal for his doctoral research on the lateral current injection laser.