Scalable Performance Evaluation of a Hybrid Optical Switch

Hai Le Vu, *Member, IEEE,* Andrew Zalesky, Eric W. M. Wong, *Senior Member, IEEE,* Zvi Rosberg, Murtaza S. Bilgrami, Moshe Zukerman, *Senior Member, IEEE,* and Rodney S. Tucker, *Fellow, IEEE*

***Abstract*— This paper provides new loss models for a hybrid optical switch combining optical circuit switching (OCS) and optical burst switching (OBS). Exact blocking probabilities are computed when 1) no priority is given to either circuits or bursts, and 2) circuits are given preemptive priority over bursts. Because it is difficult to exactly compute in realistic scenarios, we derive computationally scalable approximations for the blocking probability. The sensitivity of the analytical results to burst length and circuit holding time distributions is quantified by simulation. We demonstrate how the proposed approximations can be used for multiplexing gain evaluation of a hybrid switch. In addition, the extension of our single node model to a network model comprising OCS, OBS and hybrid switches is outlined.**

Index Terms– **Wavelength division multiplexing (WDM), hybrid optical switching, optical circuit switching (OCS), optical burst switching (OBS), blocking probability.**

I. Introduction

THE rapid advancement of wavelength-division multiplexing (WDM) technology has emerged as a promising means to open up the terahertz transmission bandwidth of optical fiber [36]. A core WDM network consists of many optical cross-connects (OXCs) interconnected through hundreds, possibly thousands, of fibers containing hundreds of wavelength channels. Three basic switching technologies have been proposed for WDM networks: optical circuit switching [10], [11], optical burst switching [1], [25], [30] and optical packet switching [3], [7].

In optical circuit switching (OCS) networks, traffic is delayed until it is confirmed that connections (or *lightpaths*) between source and destination pairs are established using two-way reservation signaling. Currently, OCS is mainly used in the backbone as point-to-point links (or transmission pipes) over long distances in a quasi-static configuration. As traffic volume grows with different requirements for data, video and voice traffic, OCS may not be sufficiently flexible in responding to dynamically varying and bursty traffic loads and service diversity [10], [31]. This motivates the idea of optical burst switching [25].

In optical burst switching (OBS) networks, traffic flows are gathered at edge routers located at the periphery of the WDM network [1]. Flows are then sorted according to destination and grouped into variable-sized elementary switching entities known as *bursts*. Before a burst is sent, a *control packet* is generated at the edge router and sent toward the destination to set up a lightpath. Upon its arrival at each OXC along the lightpath, the burst size and arrival time are read from the control packet and the burst is scheduled in advance to an appropriate outgoing wavelength. The burst itself is sent after a fixed delay, referred to as an *offset*, which is greater or equal to the total processing delay encountered by the control packet. The burst is blocked at any OXC along the lightpath if it cannot be scheduled to an appropriate outgoing wavelength [30].

The third switching technology, optical packet switching (OPS) [3], makes it possible to exploit single wavelength channels as shared resources to better utilize the huge bandwidth of WDM networks by allowing for statistical multiplexing of traffic flows. While OPS is rather ideal from a performance viewpoint, it is considered the most impractical of the above mentioned three switching technologies because it mandates the deployment of high-speed optical switches and bulky delay lines to enable optical buffering of packets.

In this paper, we consider a so-called *hybrid optical switching* network as an alternative network architecture

This work was supported by the Australian Research Council, and by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China [Project No. 9040928].
H. L. Vu, A. Zalesky, M. S. Bilgrami, M. Zukerman and R. S. Tucker are with the ARC Special Research Center for Ultra-Broadband Information Networks (CUBIN), Department of Electrical and Electronic Engineering, The University of Melbourne, Melbourne, VIC 3010, Australia. CUBIN is an affiliated program of National ICT Australia.
E. W. M. Wong is with the Optoelectronics Research Centre, Department of Electronic Engineering, City University of Hong Kong, Hong Kong SAR, China.
Z. Rosberg is with the Department of Communication Systems Engineering, Ben Gurion University, Beer-Sheva, 84105, Israel.

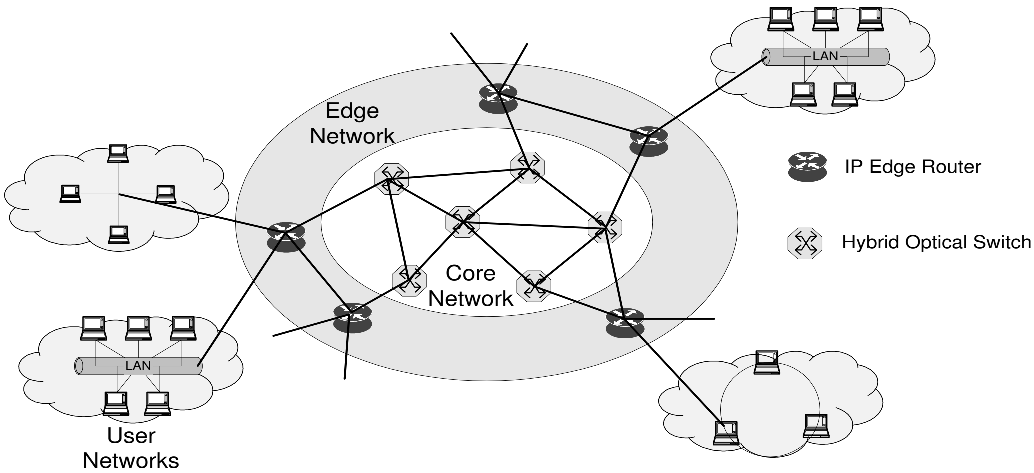

Fig. 1. Hybrid optical transport network architecture.

in which both OCS and OBS are used as the transmission mechanism [22], [35]. This architecture, shown in Fig. 1, comprises electronic routers (IP edge routers) located at the edge and so-called *hybrid optical switches* positioned inside the core WDM network. The hybrid switches are similar to OXCs but are also capable of accepting OBS traffic. There are several motivations for considering this hybrid optical switching network:

- OCS provides coarse access to bandwidth using the wavelength routing capability of the optical layer [31] and is justified only in the core network where there is a large volume of traffic between nodes. Therefore, adding OBS into the network will increase the network's flexibility and make it possible to establish point-to-point links beyond the current reach of the core network. Hybrid switches can provide an evolutionary path towards the introduction of an OCS-based network.
- Combining OCS and OBS allows a network to support the growing diversity of services. In particular, premium services such as delay-sensitive real-time traffic flows can establish a circuit on demand using OCS, while best-effort traffic can be delivered using OBS without any quality guarantee [35].
- A hybrid switching network that combines OBS and OCS is more efficient than having two separate networks. The efficiency will be achieved by reducing the maintenance and management overhead, as well as increasing traffic multiplexing.

The main contribution of this paper is to provide a computationally scalable analytical model for a single node in hybrid optical switching networks. The model can be used in performance evaluation, network dimensioning and traffic management. Although extending our single node model to a model of a complete network involving multiple hybrid optical switches is outside the scope of this paper, we discuss possible approaches to achieve this goal.

This paper is organized as follows. We develop a single node model in Section II. The single node model is then analyzed using a multidimensional Markov process in Section III. We consider the case where circuits are given preemptive priority over bursts in Section IV. Because computing the exact blocking probabilities is intractable, scalable approximations for both cases are thereafter derived. The complexity of the approximations and solving the exact blocking probabilities is shown in Section V. Furthermore, the sensitivity of the analytical results to burst length and circuit holding time distributions, as well as to the distribution of the inter-arrival times, is quantified by simulation in Section VI. We then show how our single node model can be used in dimensioning a switch in Section VII-A, and outline the extension of our model to the network case in Section VII-B. Finally, we conclude the paper in Section VIII.

II. Node Model

Consider a single hybrid optical switch in the optical network shown in Fig. 1. Without loss of generality, let us consider only traffic flows directed from left to right. In the example shown, there are two incoming and three outgoing links connected to this switch. The architecture of the hybrid optical switch is detailed in Fig. 2. The switch controller receives incoming requests in the form of control packets on each incoming fiber [35]. We assume full wavelength conversion capabilities in the switching fabric and that the switching fabric is strictly non-blocking. The main difference between the architecture in Fig. 2 and the architecture of an OXC is that the controller can accept

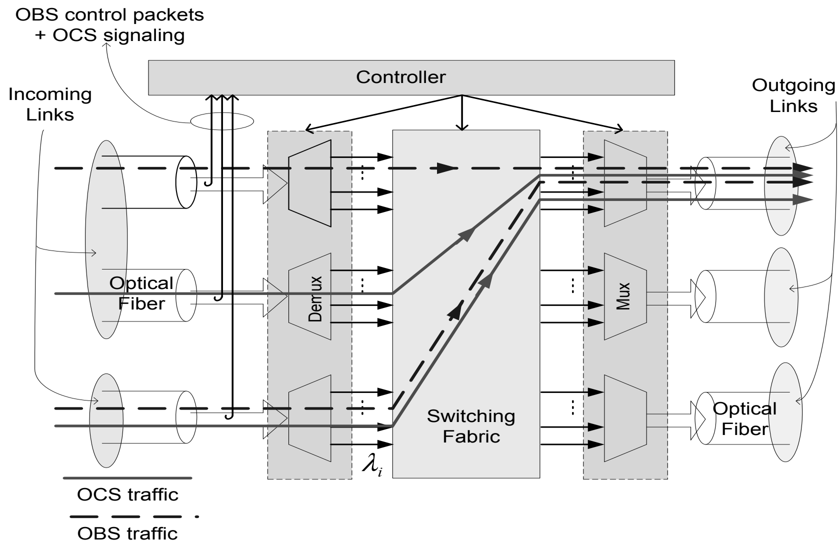

Fig. 2. Hybrid optical switch architecture.

and process both requests for the establishment of circuits as well as control packets for burst scheduling. Note that the incoming or outgoing links of the hybrid switch may contain a different number of fibers, each supporting many wavelengths. In the example shown in Fig. 2, the top incoming link from another node contains two fibers, while the bottom incoming link and all the outgoing links have one fiber each. Fig. 2 shows examples of transmission of OCS traffic and OBS traffic in progress.

Previous studies of hybrid optical switching in [22] have quantified the mean delay for bursts in a single link. Performance evaluation of hybrid packet/circuit switching in electronic SONET networks has been considered in [6], [13], [21], [23], [29], [34], [38], [39], and references therein. Here we develop an analytical model for a single switch/node within the hybrid optical network. For modeling purposes, effects related to the use of control packets in OBS, including offset, and reservation signaling in OCS are ignored.

Consider all the traffic flows coming from M input wavelengths from a number of incoming links that are directed to an outgoing link consisting of K wavelengths. Note that there is no loss when the number of input wavelengths, M, is lower than the number of wavelengths, K, in the outgoing link. We are therefore only interested in the case of $(0 < K \leq M)$, where loss can occur. In comparison, the assumption of Poisson arrivals used in previous studies of OBS performance [26] will incorrectly lead to some loss of traffic even in situations where $K \geq M$.

A request for burst transmission or circuit allocation (lightpath establishment) arrives randomly on one of the input wavelengths. If there is no available outgoing wavelength, then the request is denied and the corresponding burst is blocked (or lost) at the switch, or the requested lightpath is not established. Note that the data belonging to the dropped burst needs to be retransmitted by a higher layer protocol such as the Transmission Control Protocol (TCP).

On each input wavelength, the time period during which a burst is being received, or a circuit is allocated, is called an *on period*, and a continuous period of time between two successive on periods is called an *off period*. Here, an on period associated with a circuit allocation applies for the entire time period, including the setup time, for which a wavelength is exclusively dedicated to a circuit. We assume that the on and off periods are exponentially distributed, and the traffic streams on all input wavelengths are statistically identical. During the on period, the input wavelength is said to be *active*, and during the off period, it is said to be *inactive*.

An input wavelength may carry bursts some of the time and may be allocated to circuits at other times. We assume that a burst is transmitted on a wavelength for an exponentially distributed period of time with mean $1/\mu_b$, while a circuit is allocated for an exponentially distributed period of time with mean $1/\mu_c$. The off period is assumed to be exponentially distributed with mean $1/\lambda$. Upon termination of an off period, an on period associated with a circuit allocation will commence with probability p_c, and a burst transmission will commence with probability $p_b = 1-p_c$. Define $\lambda_c = \lambda p_c$ and $\lambda_b = \lambda p_b$.

Typically, $1/\mu_c \gg 1/\mu_b$ and $\lambda_b \gg \lambda_c$. The arrival of a request for circuit allocation may represent (but is not limited to) the following: 1) a request for setting up a private network, 2) dynamic capacity leasing on a wavelength-by-wavelength basis, and 3) online trading of bandwidth. The value $1/\mu_c$ will then represent the average actual holding/usage time associated with such wavelength bandwidth requests.

In principle, as discussed in [35], a lightpath established for a circuit between two edge routers may not be fully utilized. In such a case, an arriving request for a new circuit between the same edge routers may be accommodated by the old lightpath. Here we do not consider the second request as a new request, but rather assume that it simply increases the holding time of the existing lightpath. If the second request cannot be accommodated by an existing lightpath at the edge router, then it is considered as a new request in our model.

At first glance, it might seem that our single node model thus far defined is nothing more than a loss model covered by Engset [12] with K servers, M sources and two arrival classes. Hence, if priority is not given to any of the two classes, it seems that the standard Engset formula [12], [16] can be applied to compute blocking probabilities. However, the Engset formula will typically overestimate blocking probabilities because it allows a new arrival (either a burst or circuit) on the input wavelength while the burst is being blocked.

In practice, when a burst is blocked at a switch, the input wavelength carrying the blocked burst remains active until the end of the burst has arrived at the switch. Clearly new arrivals on that wavelength can only occur after the burst has been blocked. (See [37] for an example demonstrating the inaccuracy of the standard Engset formula in estimating burst blocking probabilities in this case.) During the period of time that a burst is being dumped at the switch, the input wavelength is said to be blocked. Herein, we refer to this input wavelength in the blocked state as a blocked wavelength. Thus, an input wavelength can either be active, inactive or blocked. Note that when a circuit request is blocked there is no dumping, and the circuit is assumed lost. In practice, depending on the application, a retry of a rejected request for circuit allocation may be conducted, possibly, by buffering and delaying packets waiting for circuit allocation at the edge router. In this paper, we do not consider such effects. See [15], [40] for information on retry models and their performance implications on circuit switching networks.

III. Analysis

Using the above node model, the blocking probability is now computed in the case where no priority is given to either circuits or bursts.

A. Exact Blocking Probability

Let the set of triples $\{(i,j,k) : i = 0,\dots,K; j = 0,\dots,K; k = 0,\dots,M-K; i+j \le K\}$ denote the states of the underlying Markov process, where i is the number of bursts in progress, j is the number of circuits in progress and k is the number of blocked input wavelengths on which bursts are being dumped. Furthermore, let $\pi_{i,j,k}$ denote the stationary distribution of the underlying Markov process. Under appropriate conditions, a unique stationary distribution exists and can be computed by solving the associated system of balance equations. Fig. 3 shows the rates of each ingoing and outgoing transition for an arbitrary state of the three dimensional Markov process for $i+j<K$.

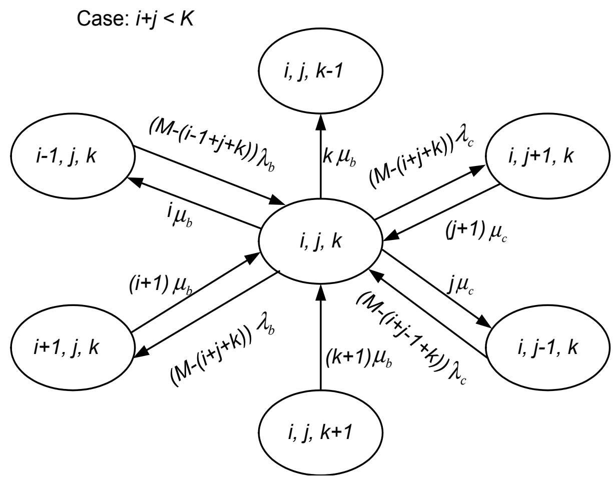

Fig. 3. State diagram for $i+j<K$.

The corresponding balance equations (for $i+j<K$) are:

$$\begin{aligned} & \pi_{i,j,k}\big((i+k)\mu_b + j\mu_c + (M-i-j-k)\lambda\big) \\ = \; & \pi_{i,j,k+1}(k+1)\mu_b \\ + \; & \pi_{i,j-1,k}\big(M-(i+j-1+k)\big)\lambda_c \\ + \; & \pi_{i,j+1,k}(j+1)\mu_c \\ + \; & \pi_{i-1,j,k}\big(M-(i-1+j+k)\big)\lambda_b \\ + \; & \pi_{i+1,j,k}(i+1)\mu_b, \end{aligned} \tag{1}$$

and for $i+j=K$,

$$\begin{aligned} & \pi_{i,j,k}\big((i+k)\mu_b + j\mu_c + (M-K-k)\lambda_b\big) \\ = \; & \pi_{i,j-1,k}(M-(K-1+k))\lambda_c \\ + \; & \pi_{i-1,j,k}(M-(K-1+k))\lambda_b \\ + \; & \pi_{i,j,k+1}(k+1)\mu_b \\ + \; & \pi_{i,j,k-1}(M-(K-1+k))\lambda_b. \end{aligned} \tag{2}$$

In (1) and (2), $\pi_{i,j,k} = 0$ for $(i,j,k) \notin \{(i,j,k) : i = 0,\dots,K; j = 0,\dots,K; k = 0,\dots,M-K; i+j \le K\}$.

Introducing the normalization equation $\sum_{i,j,k} \pi_{i,j,k} = 1$ gives rise to a linearly independent system of equations, which can be solved with elementary methods to compute the stationary distribution.

The total load offered by both bursts and circuits is given by

$$T_o = \sum_{i,j,k} (M - i - j - k)(\lambda_b/\mu_b + \lambda_c/\mu_c)\pi_{i,j,k},$$

and the total load carried by both bursts and circuits is given by

$$T_c = \sum_{i,j,k} (i + j)\pi_{i,j,k}.$$

Note that T_o, T_c are expressed in units of wavelength capacity. Thus, the blocking probability for both circuits and bursts is the same and equal to $(T_o - T_c)/T_o$.

Solving the system of equations given by (1) and (2) is not scalable for large K and M. Two approximations for the blocking probability which are applicable for realistic values of K and M are now derived by reducing the dimensionality of the underlying Markov process.

B. First Approximation

The dimension of the underlying Markov process can be reduced from three to two by considering an approximation in which no distinction is made between a burst in progress and a circuit in progress. Because the blocking probability is equal for both circuits and bursts, we consider here a single entity which can be either a burst or a circuit. This way, we do not need to keep track of the number of bursts and circuits in the system, thus reducing the dimensionality of the problem. The holding time (transmission time) of this combined entity will be the weighted average of burst transmission and circuit holding time. This results in a slight loss of accuracy because the holding time of this entity is no longer exponential but assumed to be in our approximation. This is related to the discussion in Section VI on the sensitivity of the analytical results to burst length and circuit holding time distributions.

Let the set of doubles $\{(j,k) : j = 0,\dots,K; k = 0,\dots,M-K\}$ denote the states of the approximate Markov process, where j is the total number of bursts and circuits in progress and k is number of blocked input wavelengths. On each inactive input wavelength the off period is exponentially distributed with mean $1/\lambda$, where $\lambda = \lambda_b + \lambda_c$. Let $1/\mu^*$ be the modified mean on period, which is chosen as the weighted average given by $1/\mu^* = \lambda_b/(\lambda\mu_b) + \lambda_c/(\lambda\mu_c)$.

Let $\pi_{j,k}$ denote the stationary distribution of the approximate Markov process. Under appropriate conditions, a unique stationary distribution exists and can be computed by solving the following system of balance equations. For $j < K$,

$$\begin{aligned} & \pi_{j,k}\big(j\mu^* + k\mu_b + (M-j-k)\lambda\big) \\ = \; & \pi_{j-1,k}(M-j+1-k)\lambda \\ + \; & \pi_{j,k+1}(k+1)\mu_b \\ + \; & \pi_{j+1,k}(j+1)\mu^*, \end{aligned} \tag{3}$$

and for $j = K$,

$$\begin{aligned} & \pi_{K,k}\big(K\mu^* + k\mu_b + (M-K-k)\lambda_b\big) \\ = \; & \pi_{K-1,k}(M-K+1-k)\lambda \\ + \; & \pi_{K,k+1}(k+1)\mu_b \\ + \; & \pi_{K,k-1}(M-K+1-k)\lambda_b. \end{aligned} \tag{4}$$

In (3) and (4), $\pi_{j,k} = 0$ for $(j,k) \notin \{(j,k) : j = 0,\dots,K; k = 0,\dots,M-K\}$. Introducing the normalization equation $\sum_{j,k} \pi_{j,k} = 1$ gives rise to a linearly independent system of equations, which can be solved as before to compute the stationary distribution.

The total load offered is given by

$$\widetilde{T}_o = \sum_{j,k} (M - j - k)(\lambda/\mu^*)\pi_{j,k},$$

and the total load carried is given by

$$\widetilde{T}_c = \sum_{j,k} j\pi_{j,k}.$$

Thus, an approximation of the blocking probability for both circuits and bursts is equal to $(\widetilde{T}_o - \widetilde{T}_c)/\widetilde{T}_o$.

A cruder, yet more scalable, approximation is now derived, which bears much similarity to the standard Engset formula. In fact the approximation is based on the Engset formula with mean on period $1/\mu^*$ and a modified mean off period, which is determined by solving a fixed-point equation with repeated substitution.

C. Second Approximation

Observe that from the point of view of the switch, when the input wavelength is blocked and the burst is dumped, the input wavelength behaves as if it were inactive until the end of the burst has arrived at the switch. Therefore, the blocked input wavelength encounters a longer off period with mean equals to $(\lambda_b/\lambda)(1/\mu_b)+1/\lambda$. Let $P_{blocked}$ be the probability that all K wavelengths are busy at a time instant just before the arrival of a burst or circuit (and therefore its input wavelength is blocked) and let $1/\lambda^*$ be the modified mean off period given by

$$\frac{1}{\lambda^*} = (1 - P_{blocked})\frac{1}{\lambda} + P_{blocked}\left(\frac{\lambda_b}{\lambda\mu_b} + \frac{1}{\lambda}\right). \tag{5}$$

An input wavelength only dumps an arriving burst if there is a total of K bursts and circuits in progress. Therefore,

$$P_{blocked} = \text{Eng}(\lambda^*, \mu^*, M, K) \triangleq \frac{\binom{M-1}{K}(\lambda^*/\mu^*)^K}{\sum_{i=0}^{K}\binom{M-1}{i}(\lambda^*/\mu^*)^i},$$

which is the standard Engset formula. Since the blocking probabilities for the circuits and for the bursts are equal, $P_{blocked}$ is an approximation of the blocking probability in question.

The functional relation between $P_{blocked}$ and $1/\lambda^*$ expressed in (5) gives rise to a fixed-point equation. The fixed-point, i.e., consistent values for $P_{blocked}$ and $1/\lambda^*$, may be computed with the following repeated substitution algorithm.

Let $\lambda^*(0) = \lambda$. While $|\lambda^*(n) - \lambda^*(n-1)| > \epsilon$, $n \geq 1$, generate another iteration such that

$$1/\lambda^*(n+1) = 1/\lambda + \text{Eng}(\lambda^*(n), \mu^*, M, K)/\mu^*, \quad (6)$$

or,

$$\lambda^*(n+1) = \mu^*/\big(\mu^*/\lambda + \text{Eng}(\lambda^*(n), \mu^*, M, K)\big).$$

It is now proven that the repeated substitution algorithm must converge to the unique fixed-point of (5). Observe that the transformation from $\lambda(n)$ to $\lambda(n+1)$ is defined by the function $\Gamma(x)$, where

$$\Gamma(x) = \mu^*/\big(\mu^*/\lambda + \text{Eng}(x, \mu^*, M, K)\big), \quad x > 0.$$

Because $\text{Eng}(x, \mu^*, M, K)$ is increasing with x, $\Gamma(x)$ is a strictly decreasing function for $x \geq 0$ and approaches $\lambda\mu^*/(\lambda + \mu^*)$ as x approaches infinity. Therefore, $\Gamma(x) = x$ has a unique solution (fixed-point) from which it follows that (5) also has a unique solution.

By (6), $\lambda^*(0) = \lambda > \lambda^*(1)$ and $\lambda^*(0) = \lambda > \lambda^*(2)$. In fact, $\lambda > \lambda^*(n)$ for all $n > 0$. As $\Gamma(x)$ is a decreasing function, $\lambda^*(0) > \lambda^*(1)$ implies $\lambda^*(2) > \lambda^*(1)$. Hence, $\lambda^*(0) > \lambda^*(2) > \lambda^*(1)$, and for similar reasoning $\lambda^*(1) < \lambda^*(3) < \lambda^*(2)$. In general, $\lambda^*(n) > \lambda^*(n+2) > \lambda^*(n+1)$, for n even, and $\lambda^*(n) < \lambda^*(n+2) < \lambda^*(n+1)$, for n odd.

Therefore, the sequence $\{\lambda^*(2n) : n \geq 0\}$ is decreasing and the sequence $\{\lambda^*(2n+1) : n \geq 0\}$ is increasing. Since $\text{Eng}(x, \mu^*, M, K)$ is strictly concave, $\Gamma(x) = x$ has a unique solution, and each sequence is bounded and monotonic, both sequences converge to the same unique fixed-point. ∎

IV. Preemptive Priority

We now extend the analysis for situations in which circuits are given preemptive priority over bursts. By preemptive priority, we mean that if no other wavelengths are available, a circuit is allowed to seize a wavelength being used by a burst in progress or reserved for a burst. This burst is then left without a wavelength and it must be blocked. With preemptive priority assigned to circuits, the availability of capacity for circuits is more predictable and easier to manage. Also, circuits may require priority because they may be assigned premium traffic which needs to meet certain quality of service (QoS) requirements. If a network is evolving from pure circuit switching to hybrid switching, it may be necessary to protect OCS traffic from OBS traffic in order to maintain existing OCS service levels. In such situations it would be desirable that a burst with a reasonably long burst offset time on a long-haul system, for example, does not interfere with setting up a circuit with relatively short round trip time.

A. Exact Blocking Probability

Here the states of the underlying Markov process are similar to that of the previous case when no priority is given. Thus, it is convenient to maintain the notation defined in the previous section. In fact it is only the set of balance equations defined by (2) that needs to be replaced to take into account that a circuit can preempt a burst if no other wavelengths are available. Note that (1) still holds. The set of balance equations defined by (2), which pertain to the case $i + j = K$, are replaced with the following. For $j = K$,

$$\begin{aligned}
& \pi_{i,j,k}\big((M-K-k)\lambda_b + (k+i)\mu_b + j\mu_c\big) \\
= \; & \pi_{i,j-1,k}(M-K+1-k)\lambda_c \\
+ \; & \pi_{i-1,j,k}(M-K+1-k)\lambda_b \\
+ \; & \pi_{i,j,k+1}(k+1)\mu_b \\
+ \; & \pi_{i,j,k-1}(M-K-k+1)\lambda_b \\
+ \; & \underbrace{\pi_{i+1,j-1,k-1}(M-K-k+1)\lambda_c}_{\text{circuit preempts burst}},
\end{aligned} \quad (7)$$

and for $j < K$,

$$\begin{aligned}
& \pi_{i,j,k}\big((M-K-k)\lambda + (k+i)\mu_b + j\mu_c\big) \\
= \; & \pi_{i,j-1,k}(M-K+1-k)\lambda_c \\
+ \; & \pi_{i-1,j,k}(M-K+1-k)\lambda_b \\
+ \; & \pi_{i,j,k+1}(k+1)\mu_b \\
+ \; & \pi_{i,j,k-1}(M-K-k+1)\lambda_b \\
+ \; & \underbrace{\pi_{i+1,j-1,k-1}(M-K-k+1)\lambda_c}_{\text{circuit preempts burst}}.
\end{aligned} \quad (8)$$

It is only the left-hand sides of (7) and (8) that differ. For $j < K$, there is less than K circuits in progress, hence an additional circuit can be admitted by preempting a burst in progress. This burst is then left without a wavelength and its remainder must be blocked. For $j = K$, there are K circuits in progress, hence an additional circuit cannot be admitted.

Under appropriate conditions a unique stationary distribution exists and can be computed by solving the system of balance equations defined by (1), (7) and (8). The normalization equation $\sum_{i,j,k} \pi_{i,j,k} = 1$ is required to ensure the system of balance equations is linearly independent.

Let T_o^b and T_o^c be the total load offered by bursts and circuits, respectively. Similarly, let T_c^b and T_c^c be the total load carried by bursts and circuits, respectively. Thus,

$$T_o^x = \sum_{i,j,k} (M - i - j - k)(\lambda_x/\mu_x)\pi_{i,j,k}, \quad x \in \{b, c\},$$

and,

$$T_c^b = \sum_{i,j,k} i\pi_{i,j,k}, \qquad T_c^c = \sum_{i,j,k} j\pi_{i,j,k},$$

where the index x represents an element of the index set $\{b, c\}$ referring to burst or circuit traffic, respectively. The stationary blocking probability is equal to

$$(T_o^x - T_c^x)/T_o^x, \; x \in \{b, c\},$$

and is strictly lower for circuits.

A scalable approximation for the exact blocking probability is now derived by decoupling the underlying Markov process according to bursts and circuits.

B. *Approximation for the Blocking Probability*

The approximation consists of two stages. Both stages are based on the standard Engset formula in which the mean off period is modified in much the same manner as in Subsection III-C. The first stage yields the exact blocking probability and state distribution for circuits. By state distribution, it is meant the set of probabilities $\{p_j : j = 0 \ldots, K\}$, where p_j is the probability that j circuits are in progress at steady-state. The second stage approximates the burst blocking probability by conditioning on the state distribution computed earlier.

Circuits cannot distinguish an active input wavelength belonging to a burst in progress from a blocked input wavelength in which a burst is being dumped. In both such cases the input wavelength appears *busy* and is assigned to a single state labelled the active state. The first stage of the approximation makes use of the fact that an input wavelength is either active or inactive from the viewpoint of a circuit.

On an inactive input wavelength a burst arrives with probability λ_b/λ, while a circuit arrives with probability λ_c/λ. The effect of burst arrivals can be exactly taken into account by modifying (increasing) the mean off period between two successive circuits. Let the modified mean off period between two circuits be $1/\lambda'$, which is given by

$$1/\lambda' = (\lambda_c/\lambda)(1/\lambda) + (\lambda_b/\lambda)(1/\lambda + 1/\mu_b + 1/\lambda'), \quad (9)$$

or,

$$1/\lambda' = 1/\lambda + (\lambda_b/\lambda_c)(1/\lambda + 1/\mu_b).$$

The term $1/\lambda + 1/\mu_b + 1/\lambda'$ in (9) is the mean off period given that the next arrival is a burst, which occurs with probability λ_b/λ, while the term $1/\lambda$ is the mean off period given that the next arrival is a circuit, which occurs with probability λ_c/λ. Thus, the exact circuit blocking probability is given by $\text{Eng}(\lambda', \mu_c, M, K)$, and the state distribution is given by

$$p_j = \frac{\binom{M}{j}(\lambda'/\mu_c)^j}{\sum_{i=0}^{K}\binom{M}{i}(\lambda'/\mu_c)^i}, \qquad j = 0, \ldots, K.$$

The second stage involves approximating the burst blocking probability by conditioning on the state distribution $\{p_j : j = 0 \ldots, K\}$. In particular, the burst blocking probability is computed given $j = 0, \ldots, K$, circuits are in progress using the approximation based on the Engset formula with modified off period derived in Subsection III-C. I.e., the approximation derived in Subsection III-C is applied $K + 1$ times to compute the burst blocking probability given $j = 0, \ldots, K$, circuits are in progress.

In particular, let $P_{blocked}(j)$, $j = 0, \ldots, K$, be the probability that an input wavelength is blocked given j circuits are in progress. Furthermore, let $1/\lambda^*(j)$, $j = 0, \ldots, K$, be the modified mean off period between two successive bursts given j circuits are in progress, which is given by

$$\begin{aligned} 1/\lambda^*(j) &= (1 - P_{blocked}(j))/\lambda_b \\ &+ P_{blocked}(j)(1/\mu_b + 1/\lambda_b). \end{aligned} \quad (10)$$

Given that j circuits are in progress, an arriving burst is only blocked if there is a total of $K - j$ bursts in progress. Therefore,

$$P_{blocked}(j) = \text{Eng}(\lambda^*(j), \mu_b, M - j, K - j).$$

The functional relation between $P_{blocked}(j)$ and $1/\lambda^*(j)$ expressed in (10) gives rise to a fixed-point equation. The fixed-point, i.e., consistent values for $P_{blocked}(j)$ and $1/\lambda^*(j)$, is computed with the same repeated substitution algorithm defined in Subsection III-C. Based on the earlier proof, the repeated substitution algorithm must converge to the unique fixed point of (10).

The repeated substitution algorithm is applied $K + 1$ times to compute $P_{blocked}(j)$, $j = 0, \ldots, K$.

The burst blocking probability is then approximated by un-conditioning the state distribution to give $\sum_{j=0}^{K} p_j P_{blocked}(j)$.

V. Computational Aspects

Because some of our approximations rely on an iterative procedure that terminates once a prescribed error criterion is satisfied, it is difficult to provide rigorous remarks on the computational complexity of our approximations. In particular, complexity largely depends on the approach used to solve the set of local balance equations. For this reason, instead of considering complexity, we provide the cardinality of the state-space for each of our approximations in Table I. This may give a rough indication of the potential computational savings relative to computing exact blocking probabilities.

TABLE I

Cardinality of State-Space

No Priority

Exact	$(1/2)(K^2+3K+2)(M-K+1)$
1st Approx.	$(K+1)(M-K+1)$
2nd Approx.*	K

Preemptive Priority

Exact	$(1/2)(K^2+3K+2)(M-K+1)$
Approx.*	$(1/2)K^2+(3/2)K$

Approximations marked with an asterisk in Table I make use of the repeated substitution algorithm defined in Subsection III-C. In practice, our numerical experimentation has revealed that the repeated substitution algorithm typically converges in less than 10 iterations for a stopping criterion of 10^{-8}. Table I gives an indication of the scalability of the approximations relative to computing the exact stationary probability.

VI. Numerical Evaluation

In this section, we quantify via simulation the accuracy of our analytical results and approximations, as well as their sensitivity to non-exponentially distributed on and off periods. Note that the traditional Engset model is not sensitive to the on and off period distributions [16]. In our analysis, however, a different model is used and therefore the sensitivity needs to be examined. We considered the following cases:

- Exponentially distributed on and off periods (to verify the correctness of our analytical results and approximations).
- Gamma distributed on periods and exponentially distributed off periods (to test the sensitivity of the analytical results to the distribution of the on period).
- Gamma distributed on and off periods (to examine the accuracy of the analytical results when both on and off periods are non-exponential).

Let λ_c/μ_c and λ_b/μ_b be the circuit and burst *traffic intensity per input wavelength*, respectively. When comparing between models involving Gamma versus exponential distributions, we fit the respective means of the on and off periods, as well as of the burst and circuit traffic intensity per input wavelength. All data points generated by simulation are plotted with their respective 95% confidence intervals, which are based on the Student's t-distribution [4]. Plots are presented for the blocking probability versus the *normalized traffic intensity*, defined as $(M/K)(\lambda_b/\mu_b + \lambda_c/\mu_c)$.

Define the parameter $S \geq 1$ as μ_b/μ_c, which represents the factor by which the mean circuit holding period is greater than that of the period required to send a burst, and the parameter $1 \geq R \geq 0$ by λ_c/λ_b. Notice that the mean inter-arrival time between two consecutive circuits may include many burst transmissions. Furthermore, for the case in which priority is not given to either bursts or circuits, regardless of the choice of parameters S and R, the blocking probabilities of a burst and circuit are the same and are only a function of the total traffic load. Herein, we set $R = 0.01$, $S = 100$ and $\mu_c = 1$. By setting $\mu_c = 1$, we normalize all time units with respect to the mean circuit holding time. Thus, for a given normalized traffic intensity, the corresponding parameters $\lambda_b, \mu_b, \lambda_c, \mu_c$ are determined using the above $\{R, S, \mu_c\}$ values. Knowing $\lambda_b, \mu_b, \lambda_c and \mu_c$, we are then able to calculate the blocking probabilities based on results developed in Section III and Section IV.

We consider the following scenarios: 1) $M = 5, K = 3$ with an equal ratio of burst-to-circuit traffic intensity per input wavelength, and 2) $M = 30, K = 10$ with various ratios of burst-to-circuit traffic intensity per input wavelength. For the first scenario, the exact and approximate blocking probabilities versus the normalized traffic intensity are shown in Figs. 4(a) and 4(c). The corresponding probabilities for the second scenario are shown in Figs. 4(b) and 4(d), where the ratio of burst-to-circuit traffic intensity per input wavelength is set to be one-to-two. Other ratios of burst-to-circuit traffic intensity per input wavelength, such as one-to-one or two-to-one, result in similar plots which are omitted here for brevity.

In all scenarios studied, we observe that the approximate blocking probabilities are in agreement with the results for the exact blocking regardless of the values of M

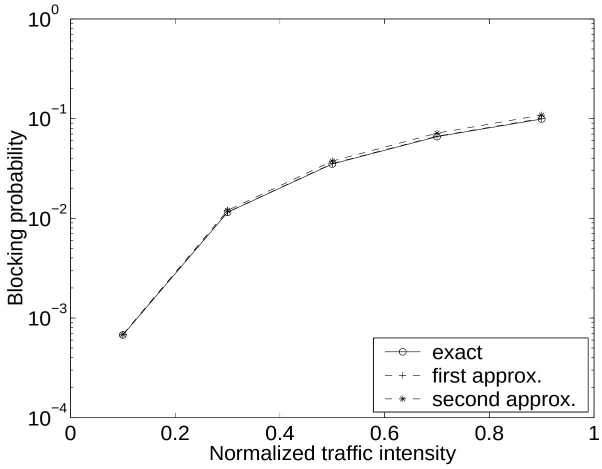

(a) No priority; $M = 5, K = 3$.

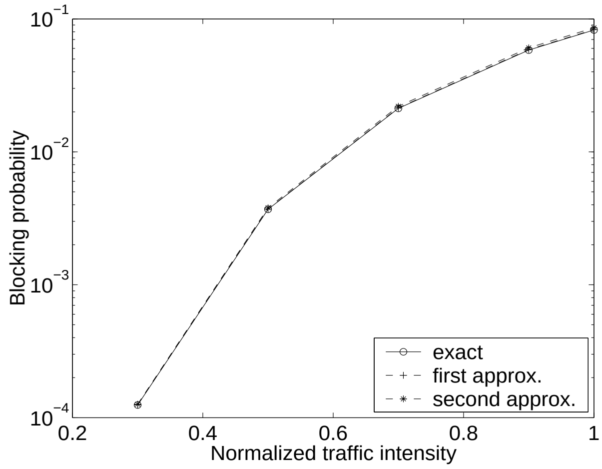

(b) No priority; $M = 30, K = 10$.

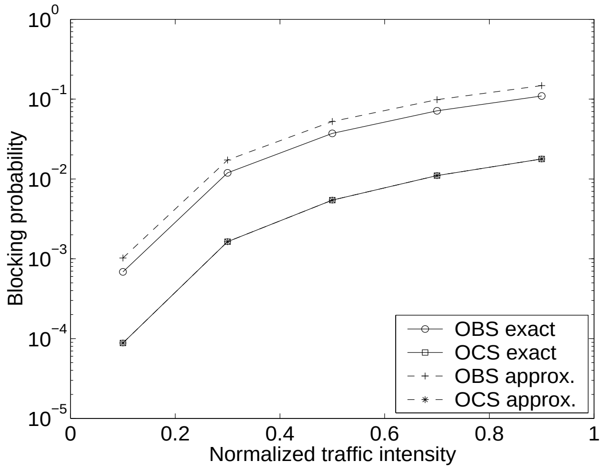

(c) Preemptive priority; $M = 5, K = 3$.

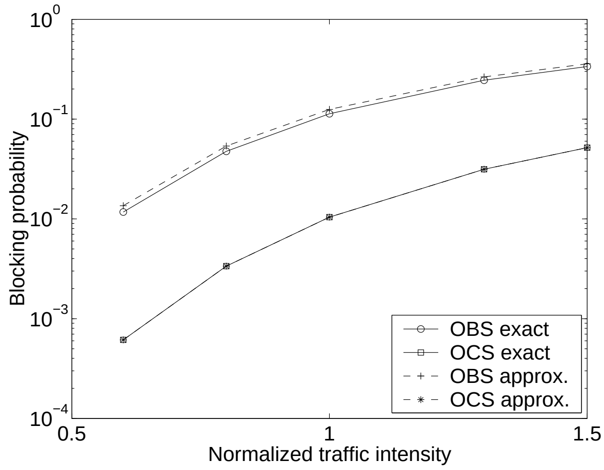

(d) Preemptive priority; $M = 30, K = 10$.

Fig. 4. Blocking probability versus normalized traffic intensity: exact and approximate results.

and K, and the ratio of burst-to-circuit traffic intensity per input wavelength. In particular, when no priority is given to circuits, the difference between the approximate and exact values divided by the exact value (referred to as *relative error*) is around 1% and 5% using the first and second approximation, respectively. This error can also be observed for various values of M up to 35 while keeping the blocking probability between 10^{-3} and 10^{-2}.

Our numerical results show that when circuits are given preemptive priority over bursts, the approximation provides the exact circuit blocking probability (as expected), and it provides a tight upper-bound for bursts (Figs. 4(c), 4(d)).

Figs. 5(a), 5(c), 5(e) and Figs. 5(b), 5(d), 5(f) show plots generated by simulation in which the on and off periods are exponentially and Gamma distributed. Exact analytical results are also plotted for comparison. For all cases studied, we observe that the analytical results are within the 95% confidence intervals of their simulation counterparts, which indicates that the analytical model is not too sensitive to non-exponentially distributed on and off periods.

In the next section, we demonstrate the scalability of our approximations; in particular, we consider cases involving hundreds of wavelengths. Computing exact blocking probabilities for such cases is intractable, however, our approximations yield estimates within a few seconds.

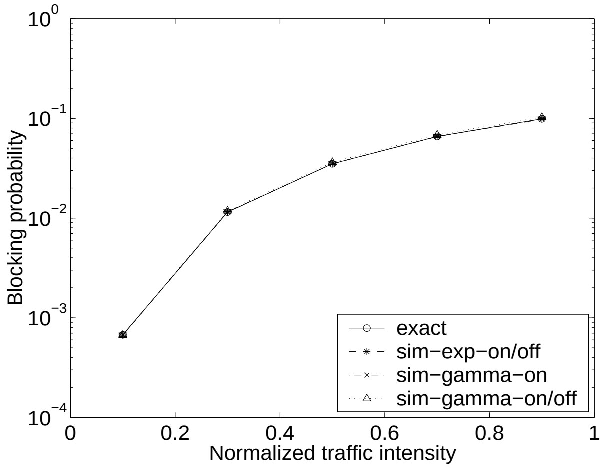

(a) No priority; $M = 5, K = 3$.
Exponential ON/OFF; Gamma ON/OFF; Gamma ON.

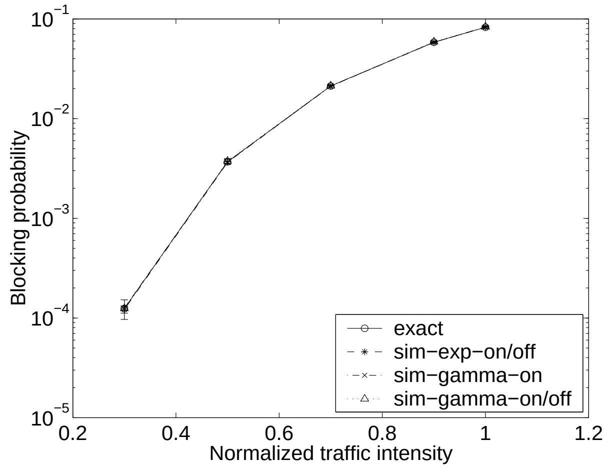

(b) No priority; $M = 30, K = 10$.
Exponential ON/OFF; Gamma ON/OFF; Gamma ON.

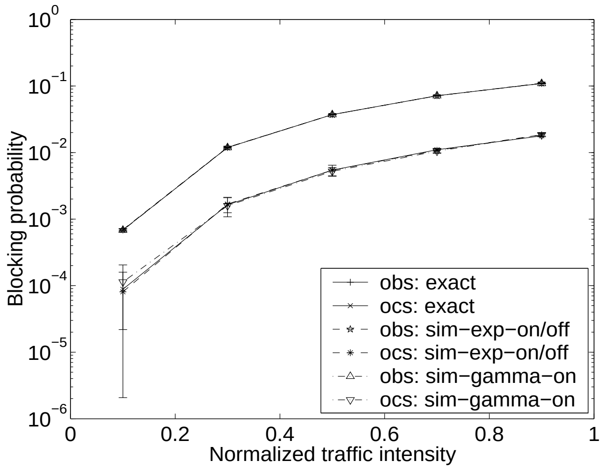

(c) Preemptive priority; $M = 5, K = 3$.
Exponential ON/OFF; Gamma ON.

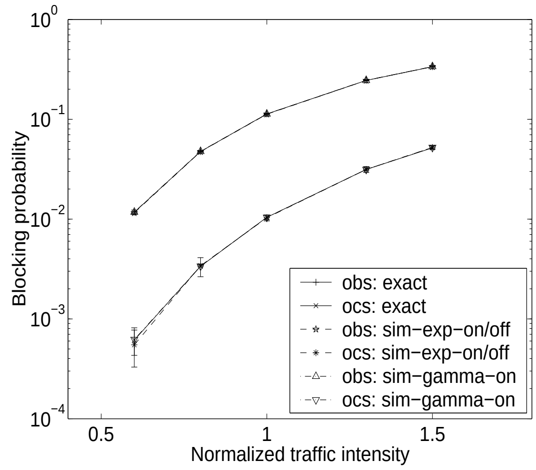

(d) Preemptive priority; $M = 30, K = 10$.
Exponential ON/OFF; Gamma ON.

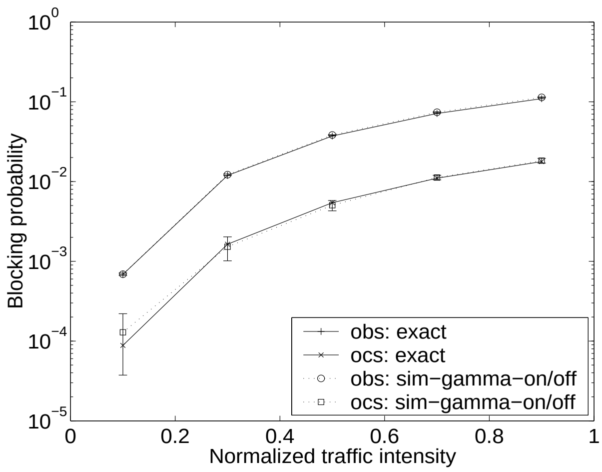

(e) Preemptive priority; $M = 5, K = 3$.
Gamma ON/OFF.

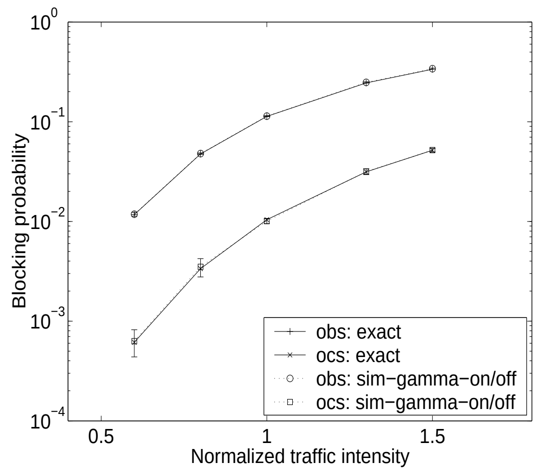

(f) Preemptive priority; $M = 30, K = 10$.
Gamma ON/OFF.

Fig. 5. Blocking probability versus normalized traffic intensity: exact and simulation results.

VII. Use and Extensions

In this section, we first demonstrate how our model can be used to evaluate the number of wavelengths on an outgoing link required to meet a pre-specified maximum blocking probability. We then use our model to evaluate the multiplexing gain achieved by hybrid optical switching. Finally, we outline how our single node model can be extended to evaluate performance of a network comprising OCS, OBS and hybrid switches.

A. Link Dimensioning and Multiplexing Gain

In order to determine the minimum number of wavelengths on an outgoing link of the switch to meet specified blocking probability requirements, we formulate an optimization problem based on two separate scenarios.

TABLE II
Link Dimensioning

	Given				Minimize
	# input	on period	off period	block. thresh.	# output
First Scenario					
OCS	M_C	$1/\mu_C$	$1/\lambda_C$	β_C	K_C
OBS	M_B	$1/\mu_B$	$1/\lambda_B$	β_B	K_B
Second Scenario					
HOS	M_H	$1/\mu_H$	$1/\lambda_H$	β_H	K_H

In the first scenario, we consider separately an OCS and OBS switch, while in the second scenario, we consider a single hybrid optical switch (HOS). The parameters associated with each switch (OBS, OCS and HOS) are summarized in Table II. For example, there are M_C input wavelengths directed to an outgoing link of the OCS switch, and the on and off periods on each input wavelength are exponentially distributed with mean $1/\mu_C$ and $1/\lambda_C$, respectively. For this OCS switch, our objective is to find the minimum number of outgoing wavelengths K_C such that the blocking probability is under a certain threshold β_C. Given M_C, λ_C and μ_C, we define the following optimization problem:

$$\begin{aligned} \textbf{minimize} \quad & K_C \\ \textbf{subject to} \quad & B_C \leq \beta_C, \end{aligned}$$

where B_C is the resulting blocking probability at the OCS switch. A similar optimization problem is defined to determine the minimum number of outgoing wavelengths K_B for an OBS switch in the first scenario.

Let $\rho_C = \lambda_C/\mu_C$ and $\rho_B = \lambda_B/\mu_B$ be the traffic intensity per input wavelength of the OCS and OBS switches, respectively, and set $\rho_C = \rho_B$. To ensure an equitable comparison with the hybrid switching case, in the second scenario, the same blocking probability threshold is specified for OCS and OBS; thus, we set $\beta_C = \beta_B = 10^{-4}$. Here, we use a similar setting $\lambda_C/\lambda_B = 0.01$, $\mu_B/\mu_C = 100$, and $\mu_C = 1$ as in Section VI. To solve the optimization problems, the blocking probability of the separate OCS switch is calculated using the standard Engset formula, and that of the separate OBS switch is calculated using our approximation described in Section III assuming no OCS traffic.

In the second scenario, for the HOS switch, we set $M_H = M_C + M_B$, $1/\lambda_H = 2/(\lambda_C + \lambda_B)$ and $1/\mu_H = \lambda_C/(2\lambda_H\mu_C) + \lambda_B/(2\lambda_H\mu_B)$. This way $\rho_H \triangleq \lambda_H/\mu_H = (\rho_C + \rho_B)/2$, indicating that the total traffic intensity of the hybrid switch $M_H\rho_H$ is equal to the sum of the traffic intensity to the separate OCS and OBS switches. We then aim to find the minimal number of wavelengths on the output link required to meet the overall blocking probability requirement of 10^{-4}. This leads to the following optimization problem:

$$\begin{aligned} \textbf{minimize} \quad & K_H, \\ \textbf{subject to} \quad & B_H \leq \beta_H, \end{aligned}$$

where $\beta_H = 10^{-4}$. Here the blocking probability is calculated based on the approximations developed in Section III.

We define the gain of hybrid switching as the percentage reduction in the number of wavelengths achieved by multiplexing the OCS and OBS traffic into a hybrid switch relative to using separate OCS and OBS switches to meet pre-specified OCS and OBS blocking probability requirements. The gain is given by $100 \times \Delta K/K_H$ [%], where $\Delta K = (K_C + K_B) - K_H$, and K_C, K_B, K_H are optimal solutions obtained by solving the above problems. The hybrid switching gain is shown in Fig. 6(a) as a function of total number of input wavelengths ($M_H = M_C + M_B$), each of which has traffic intensity of ($\rho_H = (\rho_C + \rho_B)/2$). It can be seen that the gain decreases when the number of input wavelengths or the traffic intensity per wavelength increases. Observe that the multiplexing gain is between 10 and 20% at a moderate traffic intensity per input wavelength ($\rho_H = 0.5, 0.7$) and less than 200 input wavelengths. At high traffic intensity per input wavelength ($\rho_H = 1$), the hybrid switch can still achieve around 5% multiplexing gain with a few hundred input wavelengths. Similar observations can be made in Fig. 6(b) where we plot the hybrid switching gain versus the traffic intensity per input wavelength for various values of M_H.

Note that given the small blocking probability requirement, all three systems (OCS, OBS and hybrid switches)

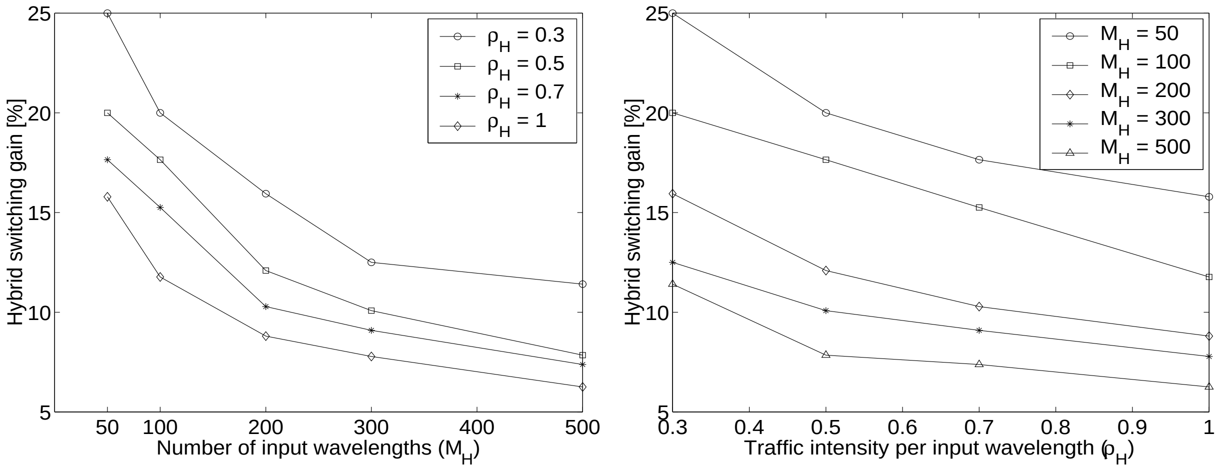

(a) Hybrid switching gain as a function of total number of input wavelengths, M_H.

(b) Hybrid switching gain as a function of traffic intensity per input wavelength, ρ_H.

Fig. 6. Hybrid switch dimensioning and multiplexing gain.

are behaving as Engset systems. When the number of input wavelengths (M_H) is large relative to the number of outgoing wavelengths (K_H) on the output link, all these systems behave like an M/M/k/k system [20], and it is known that in such a system, the multiplexing gain is insignificant if the offered load and the number of outgoing wavelengths increases while the blocking probability is kept fixed [2]. This is consistent with the plot shown in Fig. 6(a) where the multiplexing gain decreases with M_H, and with ρ_H.

B. Network Performance Modeling

In the following we describe in abstract terms an extension to approximately compute the blocking probability perceived by an ingress and egress router pair in a network consisting of OCS, OBS and hybrid optical switches. (Note that a pure OCS node, or a pure OBS node, is a special case of a hybrid optical switch.) A depiction of the architecture underlying a network consisting of hybrid optical switches is shown in Fig. 1. For simplicity, we will assume no priority is given to circuits and each ingress-egress pair is assigned a single fixed lightpath.

The main task in realizing this extension involves amending our single node approximations to dispose of the homogeneity assumption that is inherent to them. In particular, we can no longer assume that the load offered by each input wavelength that is incident to a single node is equal. This is because there may be multiple links, each of which is traversed by a different set of lightpaths, incident to a given node. Although relaxing the homogeneity assumption is mathematically tractable, for the sake of clarity, we have upheld it in the previous sections.

The closely allied problem of computing the blocking probability perceived by each ingress-egress pair in large circuit-switched networks has featured prominently in the literature [8], [14], [19], [32]. Although circuit-switched networks admit a simple product-form solution, computing the normalization constant of the product-form is often intractable. As a result, the reduced-load approximation was popularized in 1964 [9] and has remained a cornerstone of network performance evaluation. It is the reduced-load approximation that we propose to use for our purposes. Other approximations such as Monte-Carlo summation [27] and numerical inversion of the generating function of the normalization constant are not appropriate because hybrid networks do not admit a product-form solution.

Although the reduced-load approximation is usually used in the context of Poisson arrivals, it has been studied for state-dependent arrivals and in particular finite source Engset-type arrivals [17], [18], as is pertinent to our situation. For finite source Engset-type arrivals, the reduced-load approximation relies on two key assumptions:

1) The distribution of the number of busy wavelengths in a link is mutually independent of any other link.
2) The total traffic offered to a link comprises several independent on/off processes that may have been thinned owing to blocking at preceding links.

The first assumption is commonly referred to as the independence assumption which allows for decoupling of a network into its constituent links.

The first step is to compute the blocking probability perceived by a burst/circuit at each link since the independence assumption permits each link to be treated as an independent entity. For Poisson arrivals, this is usually accomplished with the Erlang B formula [26], however, in our case, we have arrivals that follow a rather complex birth-and-death process originated from finite non-homogenous input wavelengths.

To dispose of this impeding homogeneity assumption, we propose to amend the single node approximation described in Subsection III-C to allow for multiple classes of traffic, where each class corresponds to an ingress-egress pair. This amendment follows straightforwardly by using the generalized Engset formula [17] instead of its one-dimensional counterpart. The generalized Engset formula yields the blocking probability perceived by a burst/circuit of each class of traffic offered to a link. It can be computed efficiently via a generalization of the Kaufman-Roberts recursion, convolutional algorithms, fast Fourier transform algorithms [27], [28] or the unified asymptotic approximation [24]. The same iterative procedure described in Subsection III-C would still be used, however, there are now as many free variables as there are classes of traffic, where each free variable represents the blocking probability perceived by a burst/circuit of a given class.

The second step of the reduced-load approximation is to compute the reduced-load offered by bursts as well as circuits to each link. Consider an arbitrary ingress-egress pair associated with a lightpath that traverses N links. The burst load offered to the n^{th} link on its lightpath is reduced according to the blocking probability perceived by bursts at the $(1,\dots,n-1)^{th}$ links. And the circuit load offered to the n^{th} link on its lightpath is reduced according to the blocking probability perceived by circuits at the $(1,\dots,n-1,n+1,\dots,N)^{th}$ links. See [26], [33] for further details regarding the reduction of burst load and [14], [19] for further details regarding the reduction of circuit load.

It can be seen that step one is dependent on the outcome of step two and vice-versa. This dependency is synonymous with the reduced-load approximation and is usually resolved via an iterative procedure. At each iteration, link blocking probabilities and offered loads are updated according to step one and step two, respectively. The iterative procedure terminates as soon as a prescribed error criterion is satisfied. It is then a simple matter to compute the blocking probability perceived by each ingress-egress pair as a function of link blocking probabilities (see [17], [18] for further details).

VIII. Conclusion

In this paper, we have developed new models to analyze the performance of a hybrid optical switch combining OBS and OCS. Exact blocking probabilities have been derived for the cases in which no priority is given to either bursts or circuits, and circuits are given preemptive priority over bursts. The main contribution of this paper has been the derivation of computationally scalable and accurate approximations for estimating blocking probabilities for these two cases. We have demonstrated by simulation that our analysis can still provide accurate approximations for cases where the on and off periods are Gamma distributed. The utility of the approximations is that they provide a means to provision capacity in optical hybrid switching networks. Furthermore, using our approximations, we have demonstrated that significant multiplexing gain can be achieved by hybrid switching, and we have outlined an extension of our single node model to a network comprising OBS, OCS and hybrid switches.

References

[1] T. Battestilli, and H. Perros, "An Introduction to Optical Burst Switching," *IEEE Communs. Mag.*, vol. 41, Aug. 2003, pp. S10–S15.

[2] S. A. Berezner, A. E. Krzesinski and P. G. Taylor, "On the Inverse of Erlang's Function," *J. Appl. Prob.*, vol. 35, 1998, pp. 246–252.

[3] D. J. Blumenthal, P. Pruncal, and J. Sauer, "Photonic Packet Switches – Architectures and Experimental Implementations," *IEEE Proceedings*, November 1994, pp. 1650–1667.

[4] S. K. Bose, *An Introduction to queueing systems*, Kluwer Academic/Plenum Publisher, New York 2002.

[5] D. Cavendish, "Evolution of Optical Transport Technologies: From SONET/SDH to WDM," *IEEE Communs. Magazine*, June 2000, pp. 164–172.

[6] H. Che, S. Li, and A. Lin, "Adaptive Resource Management for Flow-Based IP/ATM Hybrid Switching Systems," *IEEE/ACM Trans. on Networking*, vol. 6, no. 5, October 1998, pp. 544–557.

[7] R. Cruz, and J. Tsai, "COD: Alternative Architectures for High Speed Packet Switching," *IEEE Trans. Networking*, vol. 4, February 1996, pp. 11–21.

[8] S. P. Chung, A. Kashper and K. W. Ross, "Computing approximate blocking probabilities for large loss networks with state-dependent routing," *IEEE/ACM Trans. on Networking*, vol. 1, no. 1, Feb. 1993, pp. 105–115.

[9] R. B. Cooper and S. Katz, "Analysis of alternate routing networks with account taken of nonrandomness of overflow traffic," Technical Report, Bell Telephone Lab. Memo., 1964.

[10] M. Duser, and P. Bayvel, "Analysis of a Dynamically Wavelength Routed Optical Burst Switched Network Architecture," *IEEE J. Lightwave Tech.*, vol. 20, no. 4, April 2002, pp. 574–585.

[11] M. Duser, and P. Bayvel, "Performance of a Dynamically Wavelength Routed Optical Burst Switched Network," *IEEE Photon. Tech. Letters*, vol. 14, no. 2, February 2002, pp. 239–241.

[12] T. Engset, "Die Wahrscheinlichkeitsechnung zur Bestimmung der Wahleranzahl in automatischen Fernsprechamtern," *Elektrotechnische Zeitschrift*, vol. 39, no. 31, Aug. 1918, pp. 304–306.

[13] M. J. Fischer, and T. C. Harris, “A Model for Evaluating the Performance of an Integrated Circuit and Packet Switched Multiplex Structure,” *IEEE Trans. on Communs.*, vol. COM-24, 1976.

[14] A. Girard, *Routing and Dimensioning in Circiuit-Switched Networks*, Addison-Wesley, 1990.

[15] P. Tran-Gia and M. Mandjes, “Modeling of Customer Retrial Phenomenon in Cellular Mobile Networks,” *IEEE J. of Selec. Areas in Communs.*, vol. 15, no. 8, Oct. 1997, pp. 1406–1414.

[16] J. Hui, *Switching and Traffic Theory for Integrated Broadband Networks*, Kluwer Academic Press, 1990.

[17] J. Karvo, “Generalised Engset system - properties and an application,” Tech. Report, Helsinki University of Technology.

[18] J. Karvo, J. Virtamo, S. Aalto and O. Martikainen, “Blocking of dynamic multicast connections,” *Proceedings of the 4th INFORMS Telecomm. Conference*, May 1998.

[19] F. P. Kelly, “Blocking probabilities in large circuit-switched networks,” *Advances in Applied Probability*, vol. 18, 1986, pp. 473–505.

[20] L. Kleinrock, *Queueing Systems; Volume 1: Theory*, John Wiley & Sons, 1975.

[21] R. H. Kwong, and A. Leon-Garcia, “Performance Analysis of an Integrated Hybrid-Switched Multiplex Structure,” *Performance Evaluation*, vol. 4, 1984, pp. 81–91.

[22] G. M. Lee, B. Wydrowski, M. Zukerman, J. K. Choi, and C. H. Foh, “Performance Evaluation of an Optical Hybrid Switching System,” *Proceedings of IEEE GLOBECOM 2003*, vol. 5, Dec. 1-5, 2003, pp. 2508–2512.

[23] A. Leon-Garcia, R. H. Kwong, and G. F. Williams, “Performance Evaluation Methods for an Integrated Voice/Data Link,” *IEEE Trans. on Communs.*, vol. COM-30, no. 8, August 1982, pp. 1848–1858.

[24] D. Mitra and J. M. Morrison, “Erlang capacity and uniform approximations for shared unbuffered resources,” *IEEE/ACM Trans. on Networking*, vol. 2, no. 6, Dec. 1994, pp. 558–570.

[25] C. Qiao, and M. Yoo, “Optical Burst Switching (OBS): A New Paradigm for an Optical Internet,” *J. High Speed Nets.*, vol. 8, Jan. 1999, pp. 69–84.

[26] Z. Rosberg, H. L. Vu, M. Zukerman and J. White, “Performance analyses of optical burst-switching networks,” *IEEE J. of Selec. Areas in Communs.*, vol. 21, no. 7, Sept. 2003, pp. 1187–1197.

[27] K. W. Ross, *Mutliservice loss models for broadband telecommunications networks*, Springer, 1995.

[28] K. W. Ross and D. Tsang, “The stochastic knapsack problem,” *IEEE Trans. Communications*, vol. 37, July 1989, pp. 740–747.

[29] K. Sriram, P. Varshney, and J. Shanthikumar, “Discrete-Time Analysis of Integrated Voice/Data Multiplexers With and Without Speech Activity Detectors,” *IEEE J. of Selec. Areas in Communs.*, vol. SAC-1, December, 1983, pp. 1124–1132.

[30] J. Turner, “Terabit Burst Switching,” *J. High Speed Nets.*, vol. 8, Mar. 1999, pp. 3–16.

[31] S. Verma, H. Chaskar, and R. Ravikanth “Optical Burst Switching: A Viable Solution for Terabit IP Backbone,” *IEEE Network*, November/December, 2000, pp. 48–53.

[32] W. Whitt, “Blocking when service is required from several facilities simulataneously,” *AT&T Technical Journal*, vol. 64, 1985, pp. 1807–1856.

[33] I. Widjaja, “Performance analysis of burst admission-control protocols,” *IEE Proc.-Communications*, vol. 142 no. 1, Feb. 1995, pp. 7–14.

[34] G. F. Williams, and A. Leon-Garcia, “Performance Analysis of Integrated Voice and Data Hybrid-Switched Links,” *IEEE Trans. on Communs.*, vol. COM-32, June 1984, pp. 695–706.

[35] C. Xin, C. Qiao, Y. Ye and S. Dixit, “A hybrid optical switching approach,” *Proceedings of IEEE GLOBECOM 2003*, vol. 7, Dec. 2003, pp. 3808–3812.

[36] H. Yoshimura, K. Sato, and N. Takachio, “Future Photonic Transport Networks Based on WDM Technologies,” *IEEE Communs. Magazine*, Feb. 1999, pp. 74–81.

[37] M. Zukerman, E. W. M. Wong, Z. Rosberg, G. M. Lee, and H. L. Vu, “On Teletraffic Application to OBS,” *IEEE Communs. Letters*, vol. 8, no. 2, Feb. 2004, pp. 116–118.

[38] M. Zukerman, “Bandwidth Allocation for Bursty Isochronous Traffic in a Hybrid Switching System,” *IEEE Trans. on Communs.*, vol. 37, no. 12, Dec. 1989, pp. 1367–1371.

[39] M. Zukerman, “Circuit allocation and overload control in a hybrid switching system,” *Computer Networks and ISDN Systems*, vol. 16, no. 4, 1989, pp. 281–298.

[40] M. Zukerman and C. M. Lee, “Performance Bounds for Cellular Mobile Communications Networks with Repeated Attempts,” *Proceedings of IEEE VTC 2001*, May 2001, pp. 996–1000.