Optimal Wavelength Sharing Policies in OBS Networks Subject to QoS Constraints

Li Yang and George N. Rouskas, *Senior Member, IEEE*

***Abstract*— We consider the general problem of optimizing the performance of OBS networks with multiple traffic classes subject to strict (absolute) QoS constraints in terms of the end-to-end burst loss rate of each guaranteed class of traffic. We employ Markov decision process (MDP) theory to obtain optimal wavelength sharing policies for two performance objectives, namely, maximization of weighted network throughput and minimization of the loss rate of best-effort traffic, while meeting the QoS guarantees. The randomized threshold policies we obtain are simple to implement and operate, and make effective use of statistical multiplexing. In particular, the threshold randomization feature enables the policies to allocate bandwidth at arbitrarily fine sub-wavelength granularity, hence making effective use of the available network capacity.**

***Index Terms*— Optical burst switching (OBS) networks, wavelength reservations, quality of service, Markov decision process, randomized threshold policies.**

I. Introduction

OPTICAL burst switching (OBS) is a technology positioned between wavelength routing (i.e., circuit switching) and optical packet switching. All-optical circuits tend to be inefficient for traffic that has not been groomed or statistically multiplexed, and optical packet switching requires practical, cost-effective, and scalable implementations of optical buffering and optical header processing, which are several years away. OBS is a technical compromise that does not require optical buffering or packet-level parsing, and it is more efficient than circuit switching when the sustained traffic volume does not consume a full wavelength. The transmission of each burst is preceded by the transmission of a *setup* (also referred to as burst header control) message, whose purpose is to inform each intermediate node of the upcoming data burst so that it can configure its switch fabric in order to switch the burst to the appropriate output port. An OBS source node does not wait for confirmation that an end-to-end connection has been set-up; instead it starts transmitting a data burst after a delay (referred to as *offset*), following the transmission of the setup message. For an excellent survey of OBS literature, the reader is referred to [7].

As OBS technology becomes more mature [3], supporting end-to-end quality of service (QoS) guarantees in OBS networks is arising as an important yet challenging issue. Most recent research in this area has focused on relative service differentiation, in which the QoS experienced by a class of users is specified relative to the QoS of other classes. Several complementary schemes for relative service differentiation have been proposed, such as assigning an additional offset to higher priority bursts [16], intentionally dropping non-compliant bursts [6], and allowing in-profile bursts to preempt out-of-profile ones [10], as well as combinations thereof.

An alternative approach is to guarantee each priority class a worst-case level of service (e.g., in terms of burst loss rate) that is independent of the service levels provided to other classes. A comprehensive study of absolute QoS guarantees in OBS networks can be found in [17], where several mechanisms were proposed to enforce a loss probability threshold for guaranteed traffic while reducing the loss rate of non-guaranteed traffic, including: an early dropping mechanism to drop non-guaranteed traffic selectively, a wavelength grouping strategy to allocate wavelengths to priority traffic, and a path clustering algorithm. In earlier work [14], we also considered the problem of providing QoS guarantees to multiple classes of users of an OBS network in terms of end-to-end loss. We developed a parameterized model for wavelength sharing among traffic classes that can provide a desired degree of isolation while achieving substantial statistical multiplexing gains. For a single link, we developed a heuristic for optimizing the policy parameters to support absolute QoS guarantees for a given set of heterogeneous traffic classes. We also developed a methodology for translating the end-to-end QoS requirements into appropriate per-link parameters so as to provide network-wide guarantees. The generalized wavelength sharing (GWS) policies in [14] are easy to implement, can be applied to a wide variety of traffic classes, and are effective in meeting the QoS of priority traffic.

In this paper, we consider the general problem of optimizing the performance of multi-class OBS networks subject to strict (absolute) QoS constraints in terms of the burst loss rate of each guaranteed class of traffic. We employ Markov decision process (MDP) theory to develop optimal wavelength sharing policies for two distinct performance objectives, while meeting specified levels of QoS. Specifically, we obtain optimal randomized threshold policies that are easy to configure and operate, and which make effective use of statistical multiplexing. A unique feature of our policies is their ability, due to threshold randomization, to allocate bandwidth at arbitrarily fine sub-wavelength granularity; to the best of our knowledge, no previously published resource allocation scheme for OBS networks has this ability. Consequently, these policies are quite effective in their use of the available network capacity, and

Manuscript received January 5, 2007; revised August 1, 2007. This material is based upon work supported in part by MCNC-RDI as part of the Jumpstart project and by the NSF under grant ANI-0322107.

L. Yang is with Scalable Network Technologies, Inc., Los Angeles, CA 90045 (email: lyang@scalable-networks.com).

G. N. Rouskas is with the Department of Computer Science, North Carolina State University, Raleigh, NC 27695-7534, USA (email: rouskas@ncsu.edu).

Digital Object Identifier 10.1109/JSAC-OCN.2007.029707.

outperform our earlier generalized wavelength sharing (GWS) policies, which in turn were shown to outperform previously proposed schemes [14].

The remainder of this paper is organized as follows. In Section II, we state our assumptions regarding the OBS network. In Sections III and IV we study a single link of a network under the objective of maximizing the weighted throughput subject to QoS constraints. We then apply MDP theory to obtain optimal wavelength sharing policies for a given link. In Section V we demonstrate how to apply this approach to a different objective, namely, to minimize the loss rate of best-effort traffic subject to the QoS constraints of priority classes. In Section VI, we extend our results to an OBS network. We present numerical results in Section VII, and we conclude the paper in Section VIII.

II. The OBS Network Under Study

We consider an OBS network with N nodes. Each link in the network can carry burst traffic on any wavelength from a fixed set of W wavelengths. We assume that each OBS node is capable of full wavelength conversion, hence an incoming burst can be forwarded on any wavelength available at its output port regardless of the wavelength on which it arrived. The network does not use any other contention resolution mechanism. Specifically, OBS nodes do not employ any buffering, either electronic or optical, in the data path, and they do not utilize deflection routing or burst segmentation. Therefore, if a burst requires an output port at a time when all wavelengths of that port are busy transmitting other bursts, then the burst is dropped.

The network supports P classes of traffic, where P is a small integer. Each traffic class $i, i = 1, \cdots, P-1$, is characterized by a worst-case *end-to-end* loss guarantee B_i^{e2e}. Parameter B_i^{e2e} represents the upper bound on the long-run fraction of bursts from class i that may be dropped by the network before reaching their destination. Without loss of generality, we assume that bursts of class i have more stringent loss requirements than bursts of class j, when $i < j$; in other words:

$$B_i^{e2e} \quad < \quad B_j^{e2e}, \qquad 1 \leq i < j \leq P \tag{1}$$

Bursts of class P are not associated with any worst-case loss guarantee; consequently, we will refer to class P as the *best-effort* class, and, for convenience, we let $B_P^{e2e} = 1.0$. In addition, each class is associated with a weight r_j, which is a measure of the importance of this class to the network. In general, we have that $r_j > r_{j+1}, 1 \leq j \leq P-1$, since higher priority users are likely to pay more for service.

Once assembled at the edge of the network, a burst is assigned to one of the P classes; the mechanism for assigning bursts to traffic classes is outside the scope of our work. The class to which a burst belongs is encoded in the setup (control) message that precedes the burst transmission. We assume that intermediate nodes make forwarding decisions by taking into account both the availability of resources (e.g., the number of free wavelengths at an output port) and the information regarding the class of a burst. Specifically, an intermediate node may drop a burst of a lower priority class even when there are wavelengths available at its outgoing link.

In this work we consider a broad class of constrained optimization problems, that address the requirements of both the network providers and the users, and can be expressed as:

> *optimize a performance objective related to the traffic-carrying capacity of the network while ensuring that the loss rate of class $i, i = 1, \cdots, P-1$, does not exceed its worst-case loss guarantee B_i^{e2e}.*

In order to achieve this objective, the network nodes need to employ appropriately designed mechanisms to allocate wavelength resources to bursts of each class based on its load and worst-case loss requirement. In the rest of this paper, we use Markovian decision process (MDP) theory to develop optimal wavelength sharing policies for two performance objectives:

1) maximize the weighted throughput of the network, and
2) minimize the loss rate of the best-effort class P.

More generally, our work can be extended in a straightforward manner to obtain optimal policies for any objective that can be expressed as the long-run average of rewards collected at state transitions, similar to expression (10) in the next section.

III. MDP Model of a Single OBS Link

Let us first consider a single link of an OBS network with W wavelengths. Class j bursts arrive to the link according to a Poisson process with rate λ_j. The service time of bursts is assumed to have an exponential distribution with mean $1/\mu$ that is independent of the class of the burst. Let n_j denote the number of class j bursts in progress (i.e., receiving service) on the link. Since the service rate does not depend on the traffic class, we can use the total number of bursts $n = \sum_{j=1}^{P} n_j$ to describe the system state at any point in time [1]. Intuitively, since there is no difference in the service rates, once a burst is admitted to service, the future system evolution is not affected by the class of this burst. Therefore, the evolution of the system is described by the Markov model $\{n(t),\ t \geq 0\}$; for the sake of simplicity, we will omit the index t whenever there is no ambiguity. Transitions in the state are either due to an arrival or a service completion event. We will use α_j (respectively, δ_j) to denote the arrival (respectively, departure) of a class-j burst.

A *control policy* determines the action to be taken at arrival events. We let $A(n, \alpha_j) \in \{0, 1\}$ denote the set of actions when a class-j burst arrives to find the system in state n. Action $a = 0$ means that the arrival is rejected, and $a = 1$ that the arrival is accepted. If the system is full (all wavelengths are occupied), then the only action available is $a = 0$, thus

$$A(W, \alpha_j) = 0, \qquad j = 1, \cdots, P. \tag{2}$$

If the system is not full, an arriving burst may be dropped if the free wavelengths are reserved for other classes of traffic:

$$A(n, \alpha_j) \in \{0, 1\},\ \ n = 0, \cdots, W-1,\ \ j = 1, \cdots, P. \tag{3}$$

There is no control at departure epochs, hence we let

$$A(n,\ \delta_j) = 0,\ \ n = 1, \cdots, W,\ \ j = 1, \ldots, P. \tag{4}$$

We consider the set of *stationary control policies* in this work. The definition of a stationary policy can be found in [9]. In essence, the controls of the stationary policy at each state are history-independent and do not change with time t. There are two commonly used types of stationary policies [9]. A *randomized stationary policy* π, defined on the state space $\mathcal{S}$, is such that the policy probabilistically selects one of a set of actions at each state. We let $\pi(a|s)$ denote the probability that an action $a \in A(s)$ is chosen at state s; clearly, $\pi(A(s)|s) = 1,\ s \in \mathcal{S}$. A randomized stationary policy π is called k-randomized stationary, $k = 0, 1, \cdots$, if

$$\sum_{s \in \mathcal{S}} \sum_{a \in A\{s\}} \mathbf{1}\{\pi(a|s) > 0\} \leq |\mathcal{S}| + k. \qquad (5)$$

In other words, there exist at most k states at which the number of control actions chosen by π is greater than 1.

A *deterministic stationary policy* is equivalent to a 0-randomized stationary policy: $A(s)$ reduces to a singleton, and we use the action $\pi(s)$ at each state s to describe the policy.

IV. Throughput Maximization Subject to QoS Constraints: The Single Link Case

Consider a single OBS link ℓ, and let B_j^ℓ denote the loss guarantee for class-j traffic *on this link.* The quantities B_j^ℓ should not be confused with the end-to-end guarantees B_j^{e2e} in (1); we will discuss in Section VI how to obtain B_j^ℓ from B_j^{e2e}.

Our objective is to determine an optimal stationary control policy that maximizes the expected sum of the class-based rewards earned by the system, subject to the constraints that the fraction of class-j customers rejected is no greater than $B_j^\ell,\ 1 \leq j \leq P-1$. Miller [11] studied the problem of maximizing the expected sum of class-based awards in a $M/M/c/N$ system (similar to our OBS link), without imposing any constraints on the blocking probabilities. He showed that, for each class, the optimal policy is of threshold form, i.e., for each class j there is a critical level M_j such that no customers of class j are admitted if the total occupancy $n \geq M_j$; he also showed that $M_j \geq M_i, j < i$, i.e., higher priority classes have higher thresholds. Feinberg and Reiman [8] extended Miller's study by adding the constraint that the blocking probability of class-1 customers not exceed a given value. They showed that for this single-constraint problem, the optimal policy has a threshold structure similar to that in [11], but one of the thresholds may be randomized: for a particular state s, the optimal policy chooses the threshold M with probability p and the threshold $M+1$ with probability $1-p$. We discuss these works in more detail later.

A. Constrained MDP (CMDP) Formulation

The P-class problem we study is more general than that in [8] when $P > 2$, as there are $P-1$ constraints, one for each of the $P-1$ guaranteed classes. In this section we formulate the problem as a constrained Markov decision process.

Since our system does not block departures, the state $n = 0$ (corresponding to an empty system) can be reached from any other state with probability 1. Therefore, the system satisfies the *unichain* condition [9], which requires for every stationary policy π, the transition matrix defined by π to form a Markov chain on the state space with one ergodic class and a (possibly empty) set of transient classes. Consequently, the optimal policy is independent of the initial distribution [9].

We now define the *one-step reward* and *one-step cost* functions for the controls taken at each state. We assume that a reward r_j is collected provider upon accepting a burst of class $j, j = 1, \cdots, P$. There is no reward whenever an arriving burst is rejected, or for the departure state $\eta = 0, \cdots, W-1$. Define $r(n, \alpha_j)$ as the reward collected by the system in the arrival state (n, α_j); similarly, define $r(n, \delta_j)$ as the reward collected at departure state (n, δ_j). We have that:

$$r(n, \alpha_j) = \begin{cases} r_j, & A(n, \alpha_j) = 1 \\ 0, & A(n, \alpha_j) = 0 \end{cases} \quad j = 1, \cdots, P, \qquad (6)$$

$$r(n, \delta_j) = 0, \quad j = 1, \cdots, P. \qquad (7)$$

We define the one-step cost function c^j for class j as:

$$c^j(n, \alpha_j) = \begin{cases} 0, & A(n, \alpha_j) = 1 \\ 1, & A(n, \alpha_j) = 0 \end{cases} \quad j = 1, \cdots, P, \qquad (8)$$

Thus, for each rejected class-j burst, the system accumulates one unit of cost. There is no cost associated with departure instants, hence

$$c^j(n, \delta_j) = 0, \quad j = 1, \cdots, P. \qquad (9)$$

We also define the cost function C^j as the fraction of class j bursts being rejected. Since the MDP satisfies the unichain condition, the reward and cost functions are independent of the initial state.

Define the long-run average reward earned by the system:

$$T(\pi) = \liminf_{t \to \infty} t^{-1} E^\pi \left[\sum_{i=0}^{N(t)-1} r(n[i], a[i]) \right] \qquad (10)$$

where π is a stationary policy, E^π is the expectation operator for the policy π, $N(t)$ is the number of events by time t, $n[i]$ is the state of the system just prior to occurrence of event i, and $a[i] = 0, 1$, is the action of the policy π at event i. The fraction of rejected class-j bursts, $j = 1, \cdots, P-1$, is:

$$C^j(\pi) = \limsup_{t \to \infty} E^\pi \left[N_j^{-1}(t) \sum_{i=0}^{N(t)} c^j(n[i], a[i]) \right] \qquad (11)$$

where $N_j(t)$ is the number of arrivals of class-j bursts by time t. Recall that $B_j^\ell, j = 1, \cdots, P-1$, is the loss rate to be guaranteed at this link ℓ. Then, the problem of maximizing the constrained weighted throughput can be formulated as:

$$\text{maximize} \quad T(\pi) \qquad (12)$$

$$\text{subject to} \quad C^j(\pi) \leq B_j^\ell, \quad 1 \leq j \leq P-1. \qquad (13)$$

One might be tempted to apply the uniformization technique in [4, Chapter 6] to the continuous-time MDP we defined earlier in this section in order to obtain a discrete-time MDP; and then apply the Policy-Iteration algorithm [4] to obtain the optimal policy. Unfortunately, we cannot apply the uniformization approach here, since our constraints may lead to randomized policies, under which the uniformization technique does not apply, as explained in [5]. The uniformization introduces fictitious transitions from a state to itself in the new Markov chain $\hat{X}$, which do not exist in the original process X.

The randomization allows for the possibility of changes in the action at fictitious transitions in $\hat{X}$ which are not available in X. Thus, there is the possibility that the usual uniformization technique fail to yield the same reward for $\hat{X}$ as for X.

A similar constrained optimization problem was considered in [9], in which a $(P+1)$-class system with finite state space $\mathcal{S}$ and finite action set A was studied. The optimal solution was obtained from a linear programming (LP) formulation; the formulation is omitted, but can be found in [9]. Also, the problem of maximizing the expected average reward of a $P-$class system subject to the blocking probability constraint on class-1 customers was studied in [8], and was shown to be a special case to the one in [9].

Our objective is to find the probability $\pi(a|n)$ that an action $a \in A\{n\}$ is chosen at state $n \in \mathcal{N}$, as dictated by the optimal stationary policy. As we have shown in [13], [15], our problem defined in (12)-(13) is also a special case of the one in [9], and satisfies the unichain condition. Thus, according to the results of [9], there exists an optimal policy $\pi^\star$ in the form of:

$$\pi(a=1|n) = \begin{cases} z_{n,a=1} / \sum_{a=0}^{1} z_{n,a}, & \sum_{a=0}^{1} z_{n,a} \neq 0 \\ \mathbf{1}\{a' = a\}, \text{any } a' \in A(i), & \text{otherwise} \end{cases} \quad (14)$$

$$\text{and} \qquad \pi(a=0|n) = 1 - \pi(a=1|n). \qquad (15)$$

The above expressions state the probability $\pi(a|n)$ for each action $a \in A\{n\}$ chosen at state n. In particular, $z_{n,a}$ in (14) denotes the probability that action a is taken at state n per unit of time, and its value is obtained from the optimal solution to an appropriate LP [13], [15] along the lines of the one in [9]. The optimal policy is P-randomized, thus there are at most P states such that $0 < \pi(a|n) < 1$.

The optimal policy $\pi^\star$ works as follows. If the system state is n and a class-j burst arrives, the burst will be admitted if $\pi[(n, \alpha_j), a = 1] = 1$; it will be rejected if $\pi[(n, \alpha_j), a = 0] = 1$. If $0 \leq \pi[(n,\ \alpha_j),\ a = 1] \leq 1$, then the burst will be admitted with probability $\pi[(n,\ \alpha_j),\ a = 1]$.

B. Structure of the Optimal Policy

In [8], the authors analyzed the structure of the optimal policy which maximizes the expected average reward subject to the constraint that the blocking probability of class-1 customers is no greater than a given threshold. They proved that the probabilities π dictated by the optimal policy conform to the following expressions:

$$\pi[(n, \alpha_1), a = 1] \quad = \quad 1, n = 0, \cdots W-1, \qquad (16)$$

$$\pi[(n, \alpha_j), a = 1] \quad \geq \quad \pi[((n+1), \alpha_j), a = 1],$$
$$n = 0, \cdots, W-2 \quad \text{and} \quad j = 1, \cdots, P \qquad (17)$$

$$\pi[(n, \alpha_j), a = 1] \quad \geq \quad \pi[(n, \alpha_{j+1}), a = 1],$$
$$n = 0, \cdots, W-1 \quad \text{and} \quad j = 1, \cdots, P-1. \qquad (18)$$

Expression (16) states that bursts of class 1 (the highest priority class) are always admitted as long as there are available resources in the system. According to expression (17), the optimal policy is such that the probability that a class-j burst will be admitted (i.e., action $a = 1$ is taken) is a non-increasing function of the system occupancy n. Finally, expression (18) states that the probability that an arriving burst is admitted at a given state n is a non-increasing function of the burst class (i.e., bursts of lower priority have lower probability to be admitted than bursts of higher priority at a given state). For our problem, we have also noticed that the optimal policy has the same properties described in (16)-(18); however, we have not been able to prove this result yet.

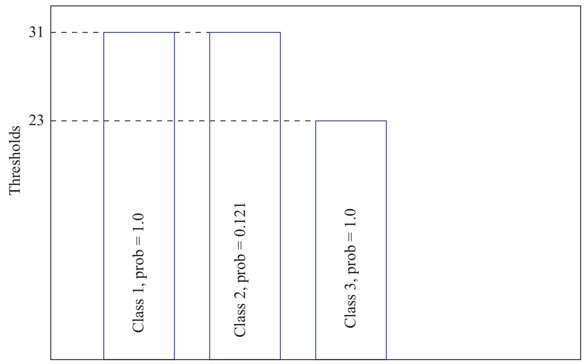

Fig. 1. Class thresholds, link load = 32 Erlang.

Expression (17) implies that for each class j, there is at most one state $M_j < W$ where $0 < \pi[(M_j, \alpha_j), a = 1] < 1$; we refer to this as the *threshold state* for class j. If a class-j burst arrives to find fewer than M_j bursts in the system, the burst is always accepted, and if it arrives to find more than M_j bursts, it is always rejected. If, on the other hand, the burst arrives to find exactly M_j bursts being served, then it is accepted with probability $\pi[(M_j, \alpha_j), a = 1]$, and it is rejected otherwise. Similarly, expression (18) implies that the threshold states are such that $M_j \geq M_{j+1}, j = 1, \cdots, P-1$, i.e., higher priority bursts are accepted in a larger number of states than lower priority ones.

To illustrate the structure of the optimal randomized threshold policy, we consider a single OBS link with $W = 32$ wavelengths and $P = 3$ classes of traffic. Classes 1 and 2 require a link loss guarantee of $B_1^\ell = 10^{-3}$ and $B_2^\ell = 10^{-2}$, respectively. We assume that class-1 (respectively, class-2) bursts represent 20% (respectively, 30%) of the traffic, and the remaining traffic is best-effort. We let the rewards r_j for admitting a class-j burst take the values: $r_1 = r_2 = 2, r_3 = 1$.

Figure 1 plots the thresholds for each class when the overall link load $\rho = 32$ Erlang. As we can see, the threshold for class 1 is $M_1 = 31$ and $\pi[(M_1, \alpha_1), a = 1] = 1.0$; therefore, as long as there is a free wavelength in the system, class 1 bursts are always admitted. The threshold for class 2 is $M_2 = 31$, and $\pi[(M_2, \alpha_2), a = 1] = 0.121$. Hence, class-2 burst will be always admitted if the number of bursts being served is less than 31; if there are exactly 31 bursts in service at the time a class-2 burst arrives, it is admitted with probability 0.121, and it is rejected with probability 0.879. The threshold for class 3 is the lowest, $M_3 = 23$, and $\pi[(M_3, \alpha_3), a = 1] = 1$; thus class 3 bursts are admitted if $n \leq M_3$.

Figure 2 plots the class thresholds as a function of link load. Since the threshold of class 1 is always $M_1 = 31$, we only plot the thresholds of class 2 and 3, respectively. As expected, the thresholds of both classes decrease with the increase in traffic load, in order to ensure that the loss rate for class 1 does not exceed the given threshold B_1^ℓ.

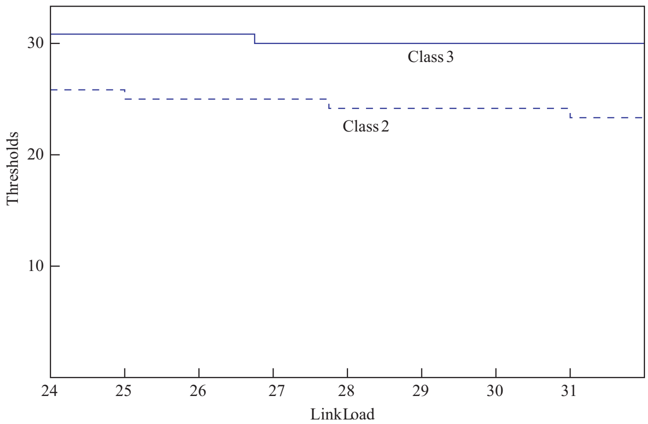

Fig. 2. Class 1 and 2 thresholds vs. link load.

V. Best-Effort Loss Minimization Subject to QoS Constraints: The Single Link Case

Let us now consider the minimization of the loss rate of the best-effort traffic class P, rather the maximization of weighted throughput, as the optimization objective; as in the previous section, the optimization is subject to the constraints that the loss rate of each priority class $i, i = 1, \cdots, P-1$, not exceed its guarantee B_i^ℓ. This objective aims to maximize the total (unweighted) amount of traffic carried by the network subject to the QoS constraints. We now show how to obtain optimal randomized policies for this objective.

The problem of minimizing the loss rate of best-effort traffic subject to the QoS constraints of guaranteed classes can be formulated as a CMDP similar to the one described in Section IV-A. Specifically, define the weights (rewards) r_j associated with the various traffic classes as:

$$r_j = \begin{cases} 0, & j = 1, \cdots, P-1 \\ 1, & j = P \end{cases} \tag{19}$$

In this case, the long-run average reward (10) earned by the system is equivalent to the number of class-P (best-effort) bursts accepted per unit of time; the cost functions C^j representing the fraction of class-j bursts, $j = 1, \cdots, P-1$, remain the same as in expression (11). Consequently, the optimization objective (12) of the CMDP formulation in expressions (12)-(13) is to maximize the fraction of accepted class-P bursts, or, equivalently, to minimize the long-term average loss rate of best-effort traffic. Since the new weights in (19) do not affect the nature of the CMDP or the properties of the problem, the techniques we presented in the previous section can be applied directly to obtain optimal randomized policies for this objective.

We emphasize that while we only consider two optimization objectives in this work, namely, maximization of weighted throughput and minimization of best-effort loss rate, the CMDP formulation and solution approach of the previous section are quite general. As such, they may be used to obtain optimal wavelength sharing policies for a wide range of objectives, subject to QoS constraints, as long as the objectives and QoS constraints can be expressed as functions of the policy actions at burst arrival and departure events similar to expressions (10) and (11), respectively.

VI. Wavelength Sharing Policies Subject to QoS Constraints: The OBS Network Case

Typically, users (applications) are interested in the end-to-end loss, rather than loss at individual links. Let us assume that the end-to-end loss guarantees B_j^{e2e} are given for all guaranteed classes j. One approach to obtaining optimal wavelength sharing policies for the OBS network would be to formulate the problem as a CMDP for the network as a whole, in a manner similar to the one we presented in Section IV. However, there are several challenges with such an approach. For instance, the action associated with a policy (i.e., accept or reject a burst) must be modified since accepting a new burst at the edge of the network does not mean that the burst will not be later rejected at a downstream link before reaching its destination. Therefore, determining the reward or cost of each policy action may require keeping track of the history of the burst as it traverses the network, hence destroying the Markovian property. Even if this challenge were possible to overcome, the state space of the process describing the whole network increases exponentially with the number of links, making the problem of determining an optimal policy intractable.

Instead, we follow a sub-optimal approach that consists of two phases and decomposes the network-wide problem into a set of independent single-link problems. In the first phase, we translate the end-to-end loss guarantees B_j^{e2e} into link loss guarantees $B_j^\ell, j = 1, \cdots, P-1$. In order to keep link management and configuration functions simple, we insist that link loss guarantees B_i^ℓ depend only on the traffic class j, *not* on the link ℓ; in other words, each traffic class is provided the *same* loss guarantee on all network links. The translation of the end-to-end guarantees into link guarantees is performed using the algorithm we developed in [14]. The algorithm determines the link loss guarantees such that the end-to-end guarantees are satisfied regardless of the specific path taken by the bursts, by assuming that the length of all paths are equal to the diameter of the network. Consequently, the network is over-provisioned, in the sense that the per-link guarantees are quite stringent; as a result, the actual end-to-end loss rate of the guaranteed classes may be somewhat lower than the end-to-end guarantee, as seen in Section VII-B.2. For a detailed description of the algorithm, the reader is referred to [14].

Once the values of B_j^ℓ are obtained for all classes j, in the second phase we use the techniques we described in Section IV to obtain the optimal randomized policy for each network link, *independently* of the other links. The simulation study we present in the following section demonstrates that this approach for tackling the problem for the network as a whole, while sub-optimal, provides good results and outperforms previously proposed solutions.

VII. Numerical Results

In this section, we compare the following two policies:

1) **CMDP policy.** This is the optimal randomized threshold policy from the constrained MDP (CMDP) formulation we developed in Section IV. Specifically, we used the simplex method to solve efficiently the corresponding linear program [13], [15] and obtain the optimal policy.

2) **Generalized wavelength sharing (GWS) policy.** We introduced the parameterized GWS policy in [14], and showed how to obtain near-optimal values for its parameters under this constrained objective.

We emphasize that the GWS family of policies is of particular interest due to the following observations: (1) they belong to a different class than the policies obtained through the CMDP approach, (2) they include many previously proposed schemes (including complete sharing and wavelength partitioning) as special cases; and (3) they lead to efficient product-form solutions for the steady-state probabilities (thus, obviating the need for time-consuming simulation).

A. A Single OBS Link

Consider a single OBS link with $W = 32$ wavelengths and $P = 3$ classes of traffic. Classes 1 and 2 require a link loss guarantee $B_1^\ell = 10^{-3}$ and $B_2^\ell = 10^{-2}$, respectively. We assume that class-1 (respectively, class-2) bursts represent 20% (respectively, 30%) of the traffic, and the remaining traffic is best-effort. There is no guarantee associated with the best-effort traffic.

1) Constrained Maximization of the Weighted Throughput: We let the reward values for the three classes be $r_1 = r_2 = 2$, $r_3 = 1$, and we compare the CMDP and GWS policies in terms of the overall weighted throughput that they achieve, subject to the QoS (loss rate) constraints. In particular, we consider the GWS policy which maximizes the weighted throughput

$$T = \sum_{j=1}^{3} r_j \lambda_j (1 - b_j)/\mu, \quad (20)$$

subject to:

$$b_1 \le B_1^\ell, \quad b_2 \le B_2^\ell, \quad (21)$$

where b_j is the blocking probability for guaranteed class j under this policy.

Figure 3 plots the weighted throughput against the link load. The CMDP policy throughput is 5-15% higher than that of the GWS policy. This result is due to statistical multiplexing: the CMDP policy makes effective use of multiplexing, but the GWS policy does not allow any sharing of wavelengths among classes. Also, the CMDP throughput increases smoothly and almost linearly with the link load, whereas the GWS throughput curve is non-monotonic. The latter is due to the fact that the GWS policy has a granularity of one wavelength; as the load increases, it may have to shift one or more wavelengths to higher priority classes, resulting in a decrease in throughput as these wavelengths may not be utilized efficiently. The CMDP policy, on the other hand, has the ability to allocate bandwidth at an arbitrarily fine (i.e., sub-wavelength) granularity by appropriately adjusting the probabilities of the threshold states for each class. This unique feature affords the CMDP policy a substantial flexibility in allocating bandwidth, and hence a much higher degree of efficiency in utilizing the available resources.

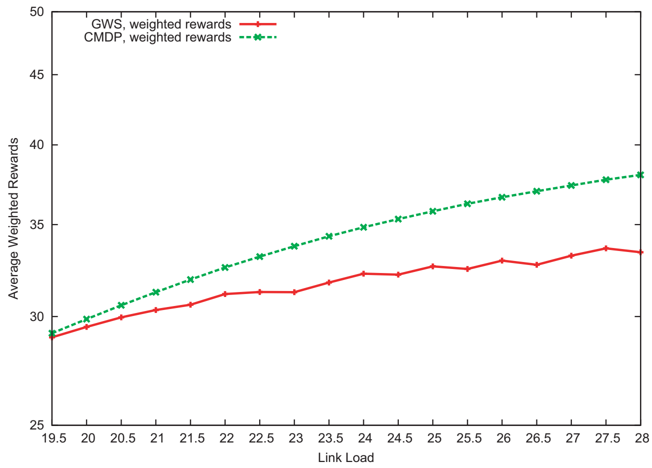

Fig. 3. Weighted throughput comparison, single link, $W = 32$, $P = 3$.

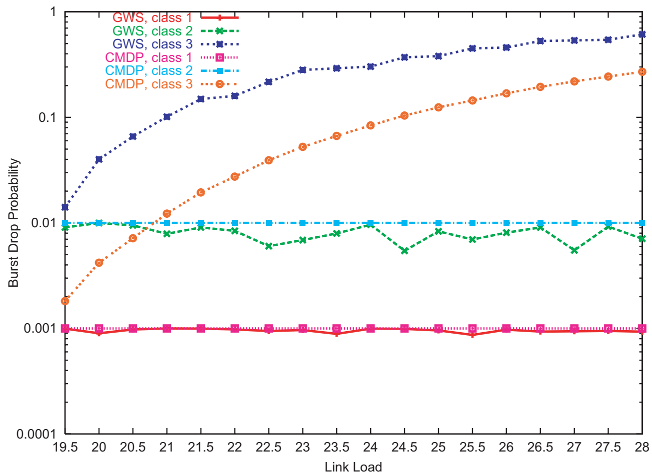

Fig. 4. Burst drop rate comparison, single link, $W = 32$, $P = 3$.

2) Constrained Minimization of the Best-Effort Traffic Loss: We also consider the objective of reducing the loss of the best-effort class. As we discussed in Section V, the class rewards for this objective are $r_1 = r_2 = 0$, $r_3 = 1$. We assume the same traffic and QoS parameters above.

Figure 4 displays the burst drop probability of each class under the GWS and CMDP policies. As we can see, the requirements of classes 1 and 2 can be guaranteed under both policies. On the other hand, the burst loss for class 3 increases with the link load, as expected. But whereas the class-3 burst loss under the GWS policy is the highest across all load values shown in the figure, under the CMDP policy, class-3 burst loss is 70-80% lower for low to moderate traffic loads; while at high loads, the burst loss rate of best-effort traffic under the CMDP policy is roughly half that under GWS policy. These results reflect the significant statistical multiplexing gains that are possible with the CMDP policy. Note also that the CMDP curves are much smoother than the GWS curves that exhibit a seesaw behavior; this result is again due to the finer granularity of the CMDP policy.

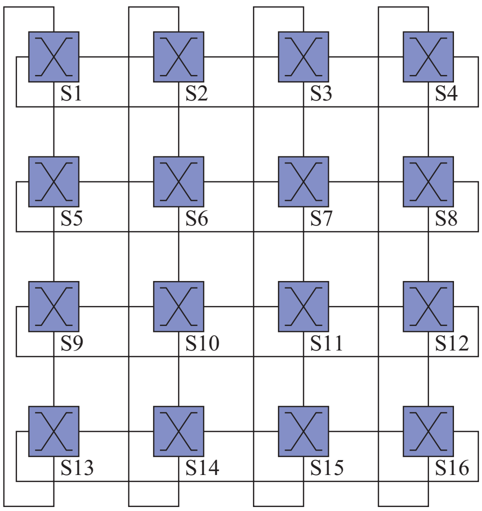

Fig. 5. The 4×4 torus network.

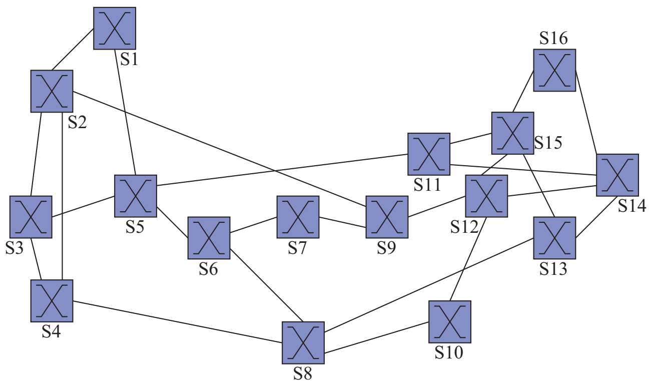

Fig. 6. The 16-node topology based on the 14-node NSFNet.

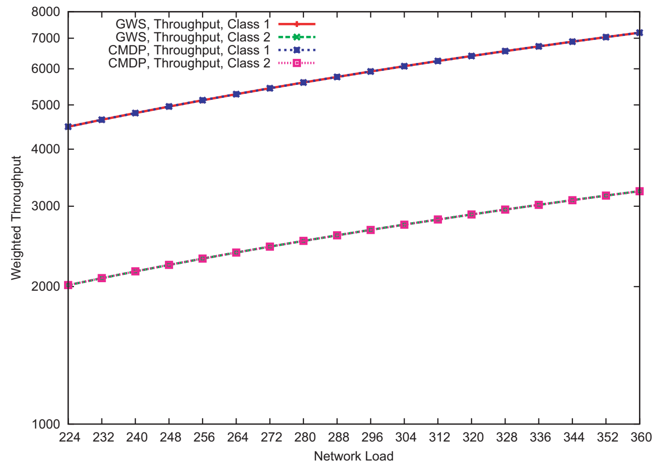

Fig. 7. NSF network, uniform traffic pattern, classes 1 and 2.

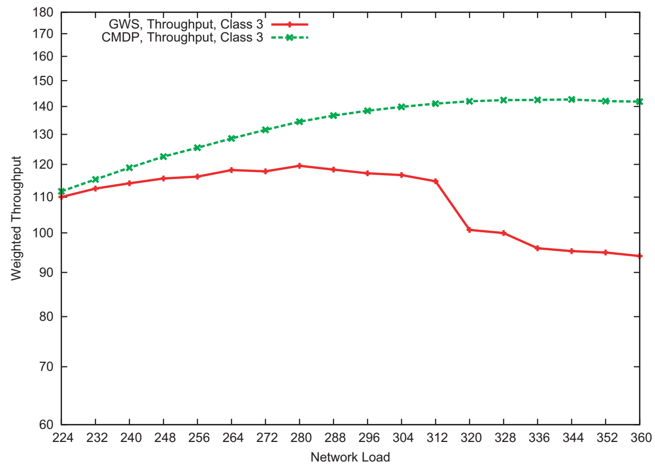

Fig. 8. NSF network, uniform pattern, class 3 (best-effort).

B. The NSF and Torus Networks

We now use simulation to demonstrate the effectiveness of the CMDP policies in providing end-to-end guarantees. For all results shown, we estimated 95% confidence intervals using the method of batch means. We used the simulator we developed as part of the Jumpstart project [12], [15]. The simulator accounts for all the details of the Jumpstart OBS signaling protocol [3], which employs the just-in-time (JIT) reservation scheme, including all messages required to set up the path of a burst and feedback messages from the network; the Jumpstart signaling protocol has been implemented in a proof-of-concept testbed on the Advanced Technology Demonstration Network (ATDNet) [2]. The burst size is assumed exponentially distributed with mean equal to one, and is taken as the unit of time in the simulation. The offset of each burst is set to the product of the number of hops on the path of the burst times the processing time at each intermediate switch, plus the switch configuration time at the last switch. We do not use larger offsets for bursts of higher priority; instead, priority is taken into account when resolving contention according to the CMDP and GWS policies, respectively.

In our experiments, we used two 16-node networks. The 4×4 torus network shown in Figure 5 is based on a regular topology, while the network in Figure 6 is based on an irregular topology derived from the 14-node NSF network. Each link in either network carries $W = 32$ wavelengths, and there are $P = 3$ classes of traffic. Classes 1 and 2 require an end-to-end loss guarantee of $B_1^{e2e} = 10^{-3}$ and $B_2^{e2e} = 10^{-2}$, respectively; class 3 is the best-effort class and does not require any guarantees. We assume shortest path routing, and we consider two traffic patterns:

1) **Uniform pattern:** each switch generates the same traffic load, and the traffic from a given switch is uniformly distributed to other switches.
2) **Distance-dependent pattern:** the amount of traffic between a pair of switches is inversely proportional to the minimum number of hops between these two switches.

1) Constrained Maximization of the Weighted Throughput: We first let the reward values be $r_1 = 100, r_2 = 30, r_3 = 1$. These values are selected to reflect the fact that higher priority traffic is worth more, in terms of revenue, to the network provider. Figures 7 and 8 plot the weighted throughput of the NSF network for the GWS and CMDP policies and the uniform traffic pattern, against the network load; similar results were obtained for the distance-dependent pattern. As Figure 7 shows, the throughput for the guaranteed classes is almost identical under the two policies. The main difference between the policies is in the throughput of the best-effort class, which is up to 50% higher under the CMDP policy, as

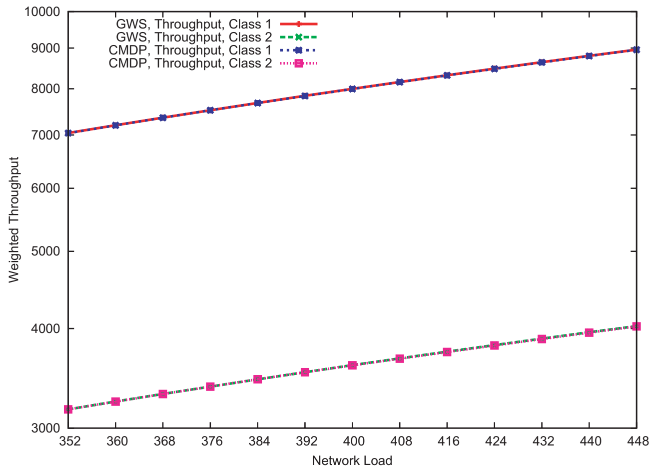

Fig. 9. Torus network, distance-dependent pattern, classes 1 and 2.

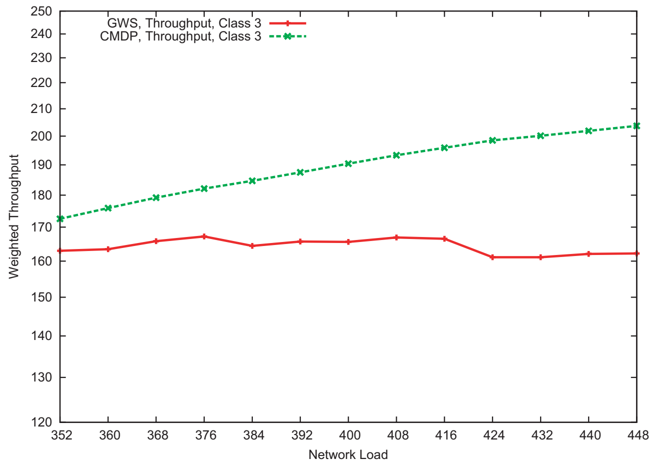

Fig. 10. Torus network, distance-dependent traffic, class 3 (best-effort).

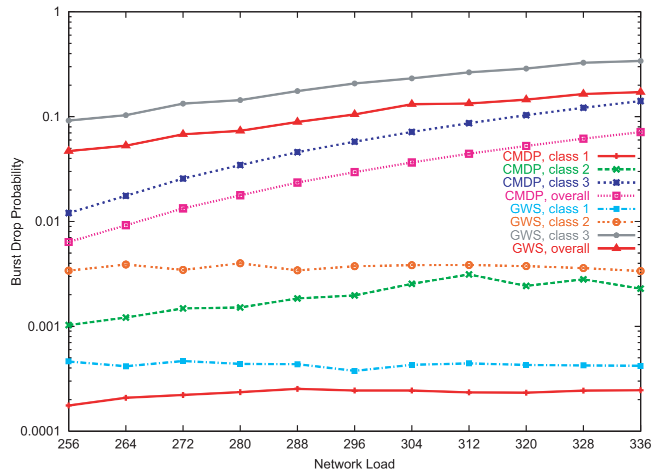

Fig. 11. NSF network, uniform traffic.

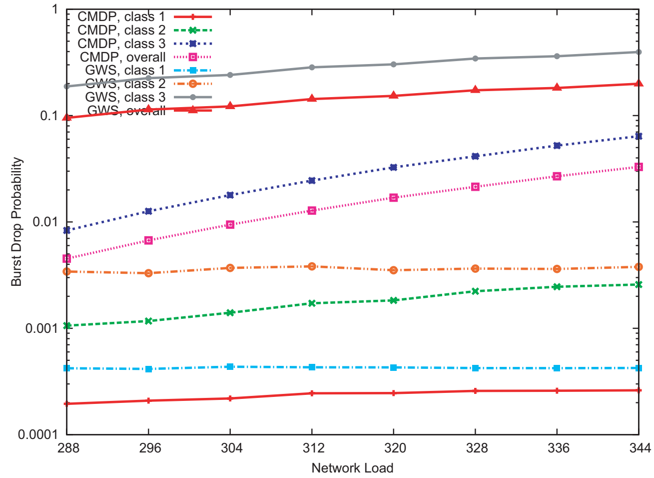

Fig. 12. NSF network, distance-dependent traffic.

shown in Figure 8. This result can be explained by noting that when the reward of guaranteed traffic is high, both the GWS and CMDP policies admit as much of this traffic as possible subject to the loss constraints; hence the weighted throughput for guaranteed classes will be the same under both policies. The CMDP policy, however, achieves higher throughput for the best-effort class due to statistical multiplexing gains as well as the finer granularity at which it can allocate wavelengths among the traffic classes. Also note that with the GWS policy, class 3 throughput decreases as the load increases from 280-360 Erlang. This behavior is due to the saturation of the bottleneck links: as the load increases, an increasing number of links have no wavelengths available for class-3 bursts, as resources are reserved to satisfy the QoS of guaranteed classes. On the other hand, due to statistical multiplexing, the CMDP policy can provide service to the best-effort traffic even at high loads; however, class-3 throughput saturates at very high loads, as resources are needed for the guaranteed classes.

Figures 9-10 are similar to Figures 7-8, respectively, but present results for the Torus network under the distance-dependent traffic pattern and the same QoS parameters. Specifically, The relative performance of the CMDP and GWS curves are very similar to the one observed for the NSF network. Indeed, the only significant difference between the two networks is due to the fact that the regular Torus network can accommodate a higher offered load, and hence achieve a higher weighted throughput, than the irregular NSF network, the traffic-carrying capacity of which is limited by the existence of bottleneck links.

2) Constrained Minimization of the Best-Effort Traffic Loss: Finally, we compare the CMDP and GWS policies under the objective of minimizing the loss rate of best-effort traffic subject to the usual QoS constraints for classes 1 and 2. Figures 11 and 12 plot the results under the uniform and distance-dependent traffic pattern, respectively, for the NSF network; Figures 13 and 14 show similar results for the Torus network. Each figure plots the burst drop probability against the network load and contains a set of four curves for each policy (CMDP or GWS): one burst drop probability curve for each traffic class, and one curve for the overall burst drop probability (i.e., the average burst drop probability over the three classes) in the network.

As we can see, the relative behavior of the two sets of curves for each policy is similar across all four figures. Specifically, both the CMDP and GWS policies are able to meet the QoS requirements for the two priority classes across the range of loads shown, implying that both policies are

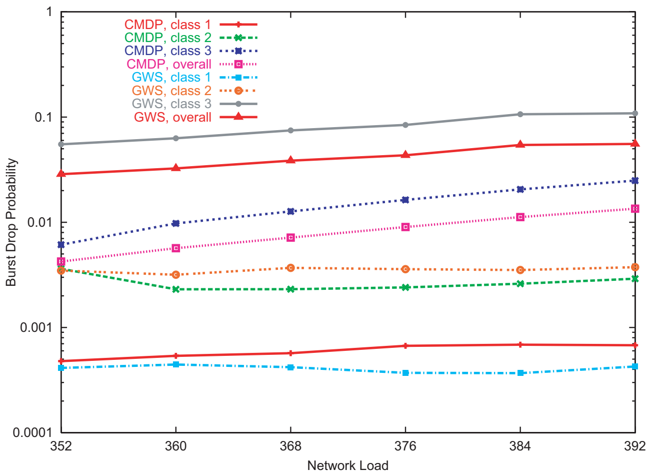

Fig. 13. Torus network, uniform traffic.

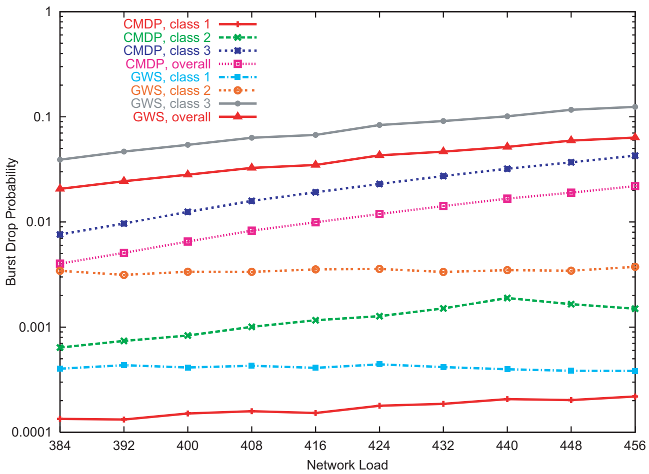

Fig. 14. Torus network, distance-dependent traffic.

effective in providing QoS guarantees. However, the results also demonstrate that the CMDP policy is more efficient in utilizing the available capacity and sharing it among the three classes. In particular, the loss rate of best-effort traffic, as well as the overall loss rate, is significantly lower under the CMDP policy: at low load, the reduction in the best-effort and overall loss rate over the GWS policy is close to one order of magnitude lower, while at moderate to low load the reduction can be up to 40-50%. More importantly, in most cases the CMDP policy also achieves lower loss rates for the two guaranteed classes, hence it improves the performance of the network across all traffic classes.

The relative performance of the two policies is due to two factors. First, the CMDP policy is optimal among the class of randomized threshold policies considered here. On the other hand, the parameters of the GWS policy are computed using a heuristic, as the associated optimization problem is intractable [14]; it is possible, therefore, that an optimal GWS policy might perform closer to the CMDP policy. Second, as we mentioned earlier, the CMDP randomized threshold policy allocates bandwidth at sub-wavelength granularity by adjusting the probabilities of the threshold states; while the GWS policy operates at the granularity of a whole wavelength. The advantage of the GWS policy is that it yields product-form solutions for the link burst loss probabilities [14], whereas such solutions are not possible for the CMDP policy.

C. Discussion of Results

Overall, the results we presented in this section demonstrate that appropriate CMDP policies can be effective in optimizing the overall performance of the network with respect to various objectives, while meeting specified levels of QoS for individual traffic classes. The CMDP policies were also shown to outperform previously proposed schemes across the network topologies and traffic patterns considered here. In particular, by selecting the rewards for the various classes appropriately, the CMDP policy can be adapted to optimize a wide range of objectives. In general, when the reward obtained for guaranteed classes is much higher than that for the best-effort class (e.g., as in Figures 7-10), the increase in throughput over the GWS policy is relatively small (as seen in Figures 8 and 10), since both policies increase the amount of guaranteed traffic as much as possible, subject to the loss constraints. However, as the weight of the best-effort class increases (e.g., as in Figures 3 and 11-14), the increase in weighted throughput compared to GWS becomes significant.

The performance advantage of the CMDP policy is due to its fine (i.e., sub-wavelength) granularity, which allows it to allocate wavelength resources efficiently among the various classes of traffic and exploit the resulting statistical multiplexing gains. This fine granularity also permits the policy to gradually re-allocate resources among traffic classes as network conditions (e.g., load) change, resulting in the smooth dynamics evident in Figures 3, 8, and 10. Also, the CMDP policy is optimal within the class of randomized threshold policies, hence it may serve as a benchmark for other heuristic schemes. Finally, from a practical point of view, the randomized threshold policy is easy to implement in core network nodes.

VIII. Concluding Remarks

We have considered the general problem of optimizing the performance of multi-class OBS networks subject to strict QoS constraints in terms of the burst loss rate of each priority class. We have formulated the problem as a constrained Markov decision process (CMDP), and we have developed techniques to obtain optimal wavelength sharing policies under two performance objectives: maximization of the weighted throughput and minimization of the best-effort loss rate. The CMDP policies are in the class of randomized threshold policies, hence they make it possible to allocate bandwidth at fine (sub-wavelength) granularity; they are also practical to implement and operate in a distributed manner. Consequently, these policies can be an effective optimization tool in contexts where it is important to provide isolation among multiple classes of users while simultaneously making efficient use of available network resources.

References

[1] E. Altman, T. Jimenez, and G. Koole, "On optimal call admission control in a resource-sharing system," *IEEE Trans. Commun.*, vol. 49, no. 9, pp. 1659–1668, Sept. 2001.

[2] I. Baldine, M. Cassada, A. Bragg, G. Karmous-Edwards, and D. Stevenson, “Just-in-time optical burst switching implementation in the ATDnet all-optical networking testbed,” in *Proc. Globecom 2003*, vol. 5, pp. 2777–2781, Dec. 2003.
[3] I. Baldine, G. N. Rouskas, H. G. Perros, and D. Stevenson, “JumpStart: a just-in-time signaling architecture for WDM burst-switched networks,” *IEEE Commun. Mag.*, vol. 40, no. 2, pp. 82–89, Feb. 2002.
[4] D. P. Bertsekas, *Dynamic Programming: Deterministic and Stochastic Models*. Prentice-Hall, 1987.
[5] F. J. Beutler and K. W. Ross, “Uniformization for semi-Markov decision processes under stationary policies,” *J. Applied Probabilities*, vol. 24, pp. 644–656, 1987.
[6] Y. Chen, M. Hamdi, and D. H. K. Tsang, “Proportional QoS over OBS network,” in *Proc. IEEE GLOBECOM*, vol. 3, pp. 1510–1514, Nov. 2001.
[7] Y. Chen, C. Qiao, and X. Yu, “Optical burst switching: a new area in optical networking research,” *IEEE Network*, vol. 18, no. 3, pp. 16–23, May/June 2004.
[8] E. Feinberg and M. I. Reiman, “Optimality of randomized trunk reservation,” *Probability in the Engineerng and Informational Sciences*, vol. pp. 463–489, 1994.
[9] E. A. Feinberg, “Constrained semi-markov decision processes with average rewards,” *ZOR-Mathematical Methods of Operations Research*, vol. 39, pp. 257–288, 1994.
[10] C.-H. Loi, W. Liao, and D.-N. Yang, “Service differentiation in optical burst switched networks,” in *Proc. IEEE GLOBECOM 2002*, vol. 3, pp. 2313–2317.
[11] B. Miller, “A queueing reward system with several customer classes,” *Management Science*, vol. 16, pp. 234–245, 1969.
[12] J. Teng, “A study of optical burst switched networks with the Jumpstart Just-In-Time Signaling Protocol,” Ph.D. thesis, North Carolina State University, Raleigh, NC, Aug. 2004.
[13] L. Yang and G. N. Rouskas, “Dynamic wavelength sharing policies for absolute QoS guarantees in OBS networks,” in *Proc. IEEE Globecom 2006*.
[14] L. Yang and G. N. Rouskas, “Generalized wavelength sharing policies for absolute QoS guarantees in OBS networks,” *IEEE J. Sel. Areas Commun.*, Supplement on Optical Communications and Networking, vol. 25, no. 4, pp. 93–104, April 2007.
[15] Li Yang, “Congestion control and quality-of-service (QoS) provisions for optical burst switched networks,” Ph.D. thesis, North Carolina State University, Raleigh, NC, May 2006.
[16] M. Yoo, C. Qiao, and S. Dixit, “QoS performance of optical burst switching in IP-over-WDM networks,” *IEEE J. Sel. Areas Commun.*, vol. 18, no. 10, pp. 2062–2071, Oct. 2000.
[17] Q. Zhang, V. Vokkarane, J. Jue, and B. Chen, “Absolute QoS differentiation in optical burst-switched networks,” *IEEE J. Sel. Areas Commun.*, vol. 22, no. 9, pp. 1781–1795, Nov. 2004.

Li Yang earned a Ph.D. in Computer Science from North Carolina State University in 2006, an MS in Computer Science from Beijing University of Posts and Telecoms (BUPT) in 2000, and a BS in Computer Engineering from Xi’an Institute of Posts and Telecoms in 1997. She worked for a year as a software engineer in Alcatel Shanghai Bell-BUPT Research and Development Center in Beijing, China. She also worked as a software intern at Intel Raleigh Design Center. Dr. Yang joined Scalable Networks as a Research Scientist in 2006.

George N. Rouskas (S ’92, M ’95, SM ’01) is a Professor of Computer Science at North Carolina State University. He received the Diploma in Computer Engineering from the National Technical University of Athens (NTUA), Athens, Greece, in 1989, and the M.S. and Ph.D. degrees in Computer Science from the College of Computing, Georgia Institute of Technology, Atlanta, GA, in 1991 and 1994, respectively. During the 2000-2001 academic year he spent a sabbatical term at Vitesse Semiconductor, Morrisville, NC, and in May 2000, December 2002, and July 2006 he was an Invited Professor at the University of Evry, France. His research interests include network architectures and protocols, optical networks, multicast communication, and performance evaluation. He is also co-editor of the book *Traffic Grooming for Optical Networks: Foundations and Techniques*, to be published by Springer in 2008.

Dr. Rouskas received the 2004 ALCOA Foundation Engineering Research Achievement Award, and the 2003 NCSU Alumni Outstanding Research Award. He is a recipient of a 1997 NSF Faculty Early Career Development (CAREER) Award, and the recipient of the 1994 Graduate Research Assistant Award from the College of Computing, Georgia Tech. He is a co-author of two papers that received Best Paper Awards at the 1998 SPIE conference on All-Optical Networking and at the 2006 CSNDSP conference. Dr. Rouskas is especially proud of his teaching awards, including his induction in the NCSU Academy of Outstanding Teachers in 2004, and the Outstanding New Teacher Award he received from the Department of Computer Science in 1995.

Dr. Rouskas is founding co-editor-in-chief of Optical Switching and Networking (OSN), an Elsevier journal, he has served on the editorial boards of the *IEEE/ACM Transactions on Networking*, *Computer Networks*, and *Optical Networks*, and he was a co-guest editor for the *IEEE Journal on Selected Areas in Communications*, Special Issue on Protocols and Architectures for Next Generation Optical WDM Networks, published in October, 2000. He was technical program co-chair of the Networking 2004 conference, general co-chair of the IEEE LANMAN 2005 workshop, program chair of the IEEE LANMAN 2004 workshop, and program co-chair of the Traffic Grooming workshop 2004. He is serving as general co-chair for BROADNETS 2007. He is a member of the ACM and of the Technical Chamber of Greece.