Dynamic Wavelength Sharing Policies for Absolute QoS in OBS Networks

Li Yang, George N. Rouskas
Department of Computer Science, North Carolina State University, Raleigh, NC 27695-8206

Abstract— **We consider the problem of providing absolute QoS guarantees to multiple classes of users of an OBS network in terms of the end-to-end burst loss. We employ Markov decision process (MDP) theory to develop wavelength sharing policies that maximize throughput while meeting the QoS guarantees. The randomized threshold policies we obtain are simple to implement and operate, and make effective use of statistical multiplexing.**

I. Introduction

As optical burst switching [2] technology becomes more mature, 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, and a variety of schemes have been proposed, such as assigning an additional offset to higher priority bursts [13], intentionally dropping non-compliant bursts [5], and allowing in-profile bursts to preempt out-of-profile ones [9]. A study of absolute QoS guarantees in OBS networks can be found in [14], where two mechanisms were proposed to enforce a loss probability threshold for guaranteed traffic while reducing the loss rate of non-guaranteed traffic: an early dropping mechanism to selectively drop non-guaranteed traffic, and a wavelength grouping strategy to allocate wavelengths to priority traffic. Finally, the study in [8] differs from the above in that it considers delay, rather than burst drop probability, as the QoS parameter to be guaranteed.

We consider the problem of providing absolute QoS guarantees to multiple classes of users of an OBS network in terms of end-to-end loss. We employ Markov decision process (MDP) theory to develop wavelength sharing policies that maximize throughput 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 Section II, we describe our assumptions regarding the OBS network. In Sections III and IV we study a single link of an OBS network, and we apply MDP theory to obtain link wavelength sharing policies. In Section V, we show how to extend the results to an OBS network. We present numerical results in Section VI, and we conclude in Section VII.

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 of W wavelengths. We assume that OBS nodes are capable of full wavelength conversion, but the network does not use any other contention resolution mechanism. Specifically, nodes do not employ any buffering, either electronic or optical, in the data path, and they do not utilize deflection routing or burst segmentation.

The network supports P classes of traffic, where P is a small integer. Each 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 long-run fraction of bursts from class i that are dropped 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 $i < j$:

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

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

Once assembled, 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 recorded 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 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.

The objective we consider is to *ensure that the loss rate of class $i, i = 1, \cdots, P-1$, does not exceed its worst-case loss guarantee B_i^{e2e}, and at the same time maximize the weighted throughput in the network.* 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. Next, we use MDP theory to develop such mechanisms.

III. 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

This work was supported by the NSF under grant ANI-0322107.

1-4244-0357-X/06/$20.00 ©2006 IEEE
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2006 proceedings.

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 (resp., 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

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

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 [7]. 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 [7]. 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

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

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 shortly 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 type j customers rejected is no greater than $B_j^\ell,\ 1 \leq j \leq P-1$. Miller [10] 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 [6] extended Miller's study by adding the constraint that the blocking probability of type-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 [10], 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 [6] 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 [7], 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 [7].

Let us define the one-step reward and cost for the controls taken at each state. There is no reward for the departure state $\eta = 0, \cdots, W-1$. Define $r((n, \alpha_j), a), a \in A(n, \alpha_j)$ as the reward collected by the system in the arrival state (n, α_j). Specifically, we have that $r((n, \alpha_j), 0) = 0$ and $r((n, \alpha_j), 1) = r_j$. We define the one-step cost function c^j for class j as $c^j((n,\ \alpha_j), 0) = 1$, and $c^j(\eta) = c^j((n,\ \alpha_j), 1) = 0, 1 \leq j \leq P$. Thus, for each rejected class-j burst, the system accumulates one unit of cost. 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:

$$T(\pi) = \liminf_{t \to \infty} t^{-1} E^\pi \left[\sum_{i=0}^{N(t)-1} r(n[i], a[i]) \right] \tag{6}$$

where π is a stationary policy, E^π is the expectation operator for the policy π, and $N(t)$ is the number of events by time t. The fraction of rejected type-j bursts, $j = 1, \cdots, P-1$, is:

$$C^j(\pi) = \limsup_{t \to \infty} E^\pi \left[N_j^{-1}(t) \sum_{i=0}^{N(t)} c^j(n[i], a[i]) \right] \tag{7}$$

where $N_j(t)$ is the number of arrivals of type j bursts by time t. Recall that $B_j^\ell, j = 1, \cdots, P-1$, is the loss rate to

1-4244-0357-X/06/$20.00 ©2006 IEEE
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2006 proceedings.

be guaranteed at this link ℓ. Then, the problem of maximizing the constrained weighted throughput can be formulated as:

$$\text{maximize} \qquad T(\pi) \tag{8}$$

$$\text{subject to} \qquad C^j(\pi) \quad \leq \quad B_j^\ell, \qquad 1 \leq j \leq P-1. \tag{9}$$

One might be tempted to apply the uniformization technique in [3, 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 [3] 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 [4]. 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. In the next section, we show how to make use of Linear Programming (LP) to solve the MDP problem in (8)-(9). The advantage of the LP approach is that we can add additional constraints easily [11].

B. Linear Programming Formulation

First, let us introduce the following notations: $p(n, \eta, a)$ is defined as the transition probability from state n to η if action a is taken; $z_{n,a}$ denotes the probability that action a is taken at state n per unit of time; and $\tau(n, a)$ denotes the sojourn time of state (n, a). 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.

A similar constrained optimization problem was considered in [7], in which a $(P+1)$-class system with finite state space $\mathcal{S}$ and finite action set A was studied. Rewards $r_j(s, a)$ are collected during sojourn times for each class $j = 0, \cdots, P$ at state $s \in \mathcal{S}$. The problem studied in [7] was to maximize the reward T_0 from class-0 customers, subject to the constraints that the reward T_j from class-j customers is no less than a given level $l_j, j = 1, \cdots, P$. If $s_0 \in \mathcal{S}$ is the initial state, then this problem can be expressed as:

$$\text{maximize} \quad T_0(s_0,\ \pi) \tag{10}$$

$$\text{subject to} \qquad T_j(s_0,\ \pi) \geq l_j, \quad 1 \leq j \leq P. \tag{11}$$

To solve (10)-(11), the following LP was formulated [7]:

$$\text{maximize} \quad \sum_{s \in \mathcal{S}} \sum_{a \in A(s)} r_0(s, a) z_{s,a} \tag{12}$$

subject to

$$\sum_{a \in A(s')} z_{s',a} - \sum_{s \in \mathcal{S}} \sum_{a \in A(s)} p(s, s', a) z_{s,a} = 0, \quad s' \in \mathcal{S} \tag{13}$$

$$\sum_{s \in \mathcal{S}} \sum_{a \in A(s)} r_k(s, a) z_{s,a} \geq l_k, \quad k = 1, \cdots, K \tag{14}$$

$$\sum_{s \in \mathcal{S}} \sum_{a \in A(s)} \tau(s, a) z_{s,a} = 1 \tag{15}$$

$$z_{s,a} \geq 0, \quad s \in \mathcal{S},\ a \in A(s) \tag{16}$$

It was shown that if z is the optimal solution for LP (12)-(16), then there exists a $P-$randomized stationary policy π in the form of $\pi(a'|s) \ = z_{s,a'} / \sum_{a \in A\{s\}} z_{s,a}$ which is optimal for problem (10)-(11). In addition, [6] studied the problem of maximizing the expected average reward of a $P-$class system subject to the blocking probability constraint on class-1 customers. That is,

$$\text{maximize} \qquad T(\pi) \tag{17}$$

$$\text{subject to} \qquad C^1(\pi) \leq B_1^\ell. \tag{18}$$

Since the departure process will not be blocked, as pointed out in [6], the problem defined in (17)-(18) satisfies the unichain condition. Replacing T_0 with T and setting $r_1(s,\ a) \ = -c(s,\ a)/\lambda_1$ and $l_1 \ = -B_1^\ell$, make the formulation in (17)-(18) the same as in (10)-(11), thus problem (17)-(18) is a special case of problem (10)-(11).

Returning to our problem, define $\Lambda = \sum_{j=1}^{P} \lambda_j$; then:

$$p(n, \eta, a) = \begin{cases} \frac{\lambda_j}{n\mu + \Lambda}, \eta = n, a = 0, 0 \leq n \leq W \\ \frac{\lambda_j}{n\mu + \Lambda}, \eta = n+1, a = 1, 0 \leq n \leq W-1 \\ \frac{n\mu}{n\mu + \Lambda}, \eta = n-1, a = 0, 1 \leq n \leq W \end{cases} \tag{19}$$

The first case corresponds to an arriving burst being dropped, the second to an arriving burst being admitted, and the third to a burst departure. Regarding sojourn times, we have:

$$\tau(n, a) = \begin{cases} (n\mu + \Lambda)^{-1}, & a = 0, 0 \leq n \leq W \\ ((n+1)\mu + \Lambda)^{-1}, & a = 1, 0 \leq n \leq W-1 \end{cases} \tag{20}$$

Then, the LP corresponding to (8)-(9) can be formulated as:

$$\text{maximize} \qquad \sum_{n \in \mathcal{N}} \sum_{a \in A(n)} r(n, a) z_{n,a} \tag{21}$$

subject to

$$\sum_{a \in A(\eta)} z_{\eta,a} - \sum_{n \in \mathcal{N}} \sum_{a \in A(n)} p(n, \eta, a) z_{n,a} = 0, \quad \eta \in \mathcal{N} \tag{22}$$

$$\sum_{n \in \mathcal{N}} \sum_{a \in A(n)} c^j(n, a) z_{n,a} \leq B_j^\ell \times \lambda_j, \quad j = 1, 2 \tag{23}$$

$$\sum_{n \in \mathcal{N}} \sum_{a \in A(n)} \tau(n, a) z_{n,a} = 1 \tag{24}$$

$$z_{n,a} \geq 0, \quad a \in A(n),\ n \in \mathcal{N} \tag{25}$$

Equation (21) represents the objective to maximize, i.e., the weighted expected throughput. Equation (22) maintains the flow balance for each state, while expression (23) represents the constraints in terms of the time-average burst drop rate for each class. Note that the left hand side of (23) is the cost for each class over the time t, while the term in the right hand side represents the imposed threshold on the time-average cost. Expression (24) simply states that the summation of the selection probability over all states and controls is equal to 1.

We now see that our problem defined in (8)-(9) is also a special case of (10)-(11), and satisfies the unichain condition. Thus, there exists an optimal policy $\pi^\star$ in the form of (26), where z is optimal for the LP (21)-(24), stating the probability $\pi(a|n)$ for each action $a \in A\{n\}$ chosen at state n. The optimal policy is P-randomized, thus there are at most P states such that $0 < \pi(a|n) < 1$.

$$\pi(a=1|n) = \begin{cases} z_{n,a=1}/\sum_{a=0}^{1} z_{n,a}, & \sum_{a=0}^{1} z_{n,a} \neq 0 \\ \mathbf{1}\{a'=a\}, \text{any } a' \in A(i), & \text{otherwise} \end{cases} \quad (26)$$

$$\text{and} \qquad \pi(a=0|n) = 1 - \pi(a=1|n). \qquad (27)$$

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 we draw a random number θ uniformly in $[0,1]$; the burst will be admitted if $\theta \leq \pi[(n,\ \alpha_j),\ a=1]$, otherwise it will be rejected.

C. Structure of the Optimal Policy

In [6], the authors analyzed the structure of the optimal policy which maximizes the expected average reward subject to the constraint that the blocking probability of type 1 customers is no greater than a given threshold. They proved that the probabilities π dictated by the optimal policy conform to expressions (28)-(30). For our problem, we have also noticed that the optimal policy has the same properties described in (28)-(30); however, we have not been able to prove this result yet.

$$\begin{aligned} \pi[(n,\alpha_1), a=1] &= 1, n=0,\cdots W-1, && (28) \\ \pi[(n,\alpha_j), a=1] &\geq \pi[((n+1),\alpha_j), a=1], && \\ n=0,\cdots,W-2 \quad &\text{and} \quad j=1,\cdots,P && (29) \\ \pi[(n,\alpha_j), a=1] &\geq \pi[(n,\alpha_{j+1}), a=1], && \\ n=0,\cdots,W-1 \quad &\text{and} \quad j=1,\cdots,P-1. && (30) \end{aligned}$$

Expression (28) 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 (29), 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 (30) 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).

Expression (29) 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 (30) 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.

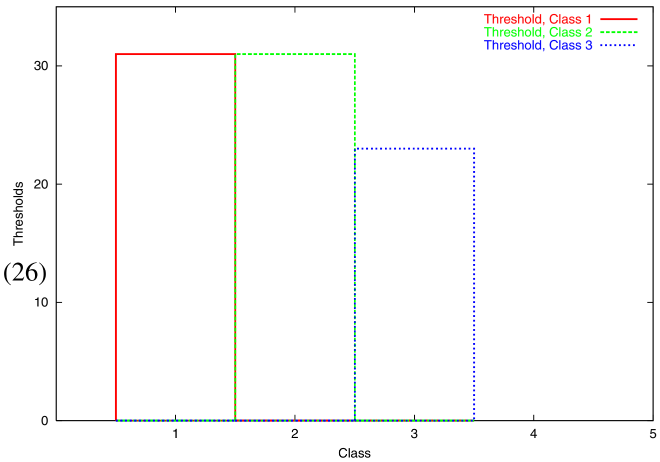

Fig. 1. Class thresholds, link load = 32 Erlang

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.

V. Throughput Maximization Subject to QoS Constraints: The OBS Network

Typically, users (applications) are interested in the end-to-end loss, rather than loss at individual links. Assuming that the end-to-end loss guarantees B_j^{e2e} are given for all guaranteed classes j, we have developed an algorithm for obtaining the link rate guarantees B_j^ℓ for all links ℓ in the network; due to space constraints, we omit a detailed description of the algorithm, which can be found in [12, Chapter 4]. The

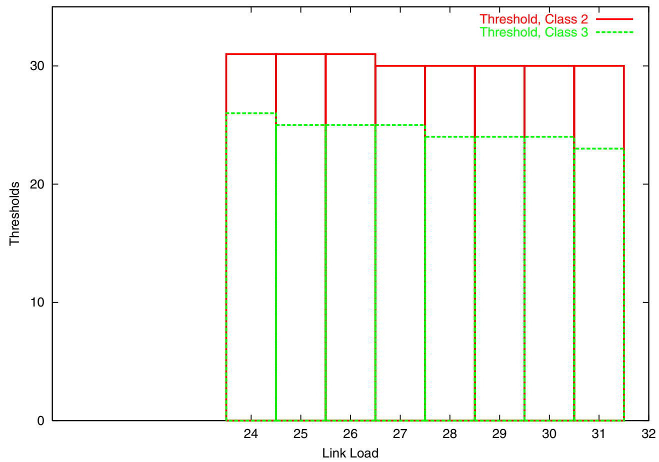

Fig. 2. Class 1 and 2 thresholds vs. link load

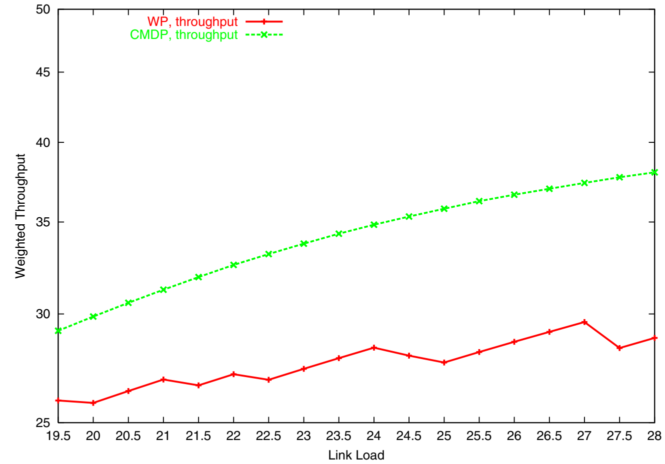

Fig. 3. Single link with $W = 32$ wavelengths and $P = 3$ traffic classes

algorithm takes into account the routing paths to determine the link rate guarantees such that the end-to-end guarantees are satisfied regardless of the specific path taken by the bursts. Once the values of B_j^ℓ are obtained for all links ℓ, we use the method we described in the previous section to obtain the optimal randomized policy for each link in the network. This approach for tackling the problem for the network as a whole, while sub-optimal, is necessitated by the fact that the state space of the Markov process describing the whole network increases exponentially with the number of links, making the problem of determining an optimal policy intractable.

VI. Numerical Results

A. A Single OBS Link

Consider a link with $W = 32$ wavelengths, $P = 3$ classes of traffic, and the same loss rate guarantees, traffic mix, and reward values as in the example we used in Section IV-C for the results in Figures 1 and 2. We compare the following two policies in terms of the overall weighted throughput that they achieve, subject to the QoS (loss rate) constraints:

1) **CMDP policy.** This is the optimal randomized threshold policy obtained through the constrained MDP (CMDP) formulation we developed in Section IV.
2) **Wavelength partitioning (WP) policy.** The WP policy partitions the set of wavelengths into P sets (in this case, $P = 3$), and each class is given exclusive use of all wavelengths in its own subset. We consider the WP policy which maximizes $\sum_{j=1}^{3} r_j \lambda_j (1 - b_j)/\mu$, subject to $b_1 \leq B_1^\ell$ and $b_2 \leq B_2^\ell$, 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 15-40% higher than that of the WP policy. This result is due to statistical multiplexing: the CMDP policy makes effective use of multiplexing, but the WP 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 WP throughput curve is non-monotonic. The latter is due to the fact that the WP 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 a much finer granularity, as it can appropriately adjust the probabilities of the threshold states for each class, giving it much more flexibility and increased efficiency in utilizing the available resources.

B. The NSF Network

We now compare the CMDP and WP policies in terms of weighted throughput on the NSF network. Each link 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 let the reward values be $r_1 = 100, r_2 = 30, r_3 = 1$. We assume shortest path routing, and we consider two traffic patterns: under the *uniform pattern*, the traffic from each node is uniformly distributed to other nodes, while under the *distance-dependent pattern*, the amount of traffic between a pair of nodes is inversely proportional to the distance between them.

Figures 4-5 plot the weighted throughput for the WP and CMDP policies with the uniform traffic pattern, against the network load. As Figure 4 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 30-100% higher under the CMDP policy, as shown in Figure 5. This result is due to the statistical multiplexing gains of the CMDP policy, as well as the finer granularity with which it can allocate wavelengths among the traffic classes. Also note that with the WP policy, class 3 throughput decreases as the load increases from 240-312 Erlang, and starts to increase after that. 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. Beyond 312 Erlang,

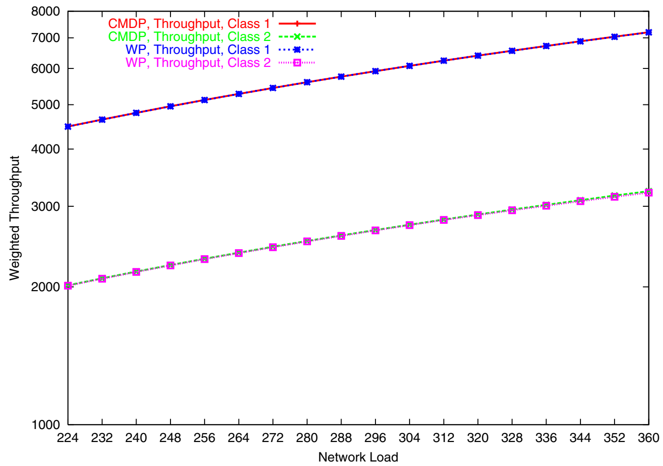

Fig. 4. NSF network, uniform traffic pattern, classes 1 and 2

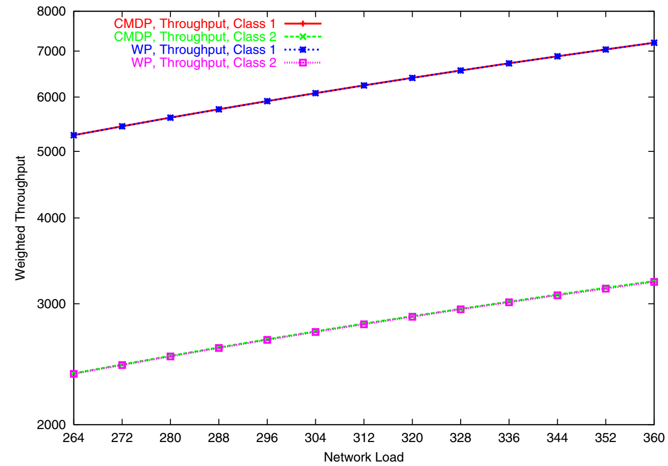

Fig. 6. NSF network, distance-dependent pattern, classes 1 and 2

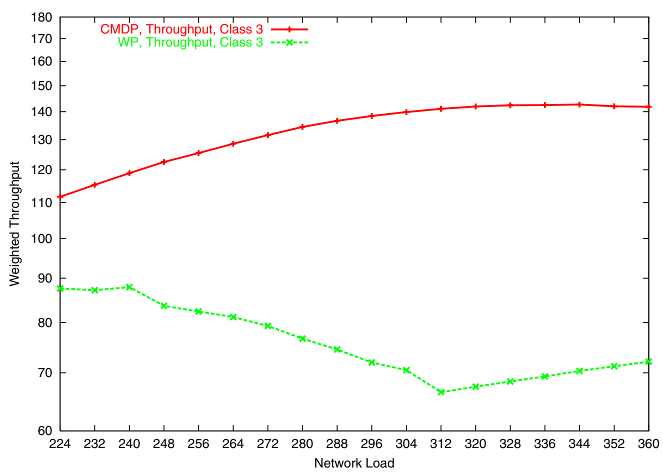

Fig. 5. NSF network, uniform pattern, class 3 (best-effort)

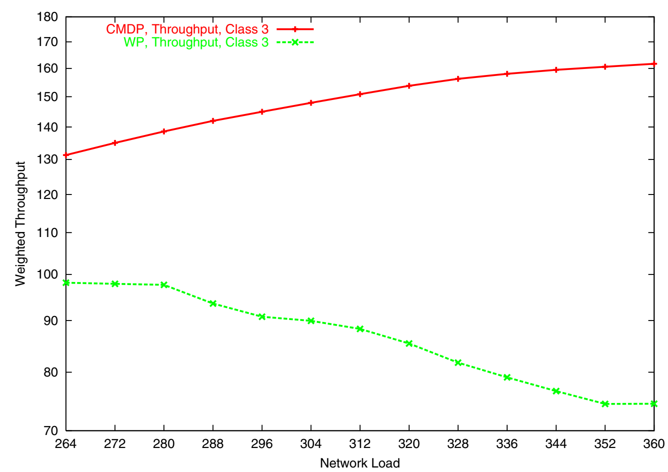

Fig. 7. NSF network, distance-dependent pattern, class 3 (best-effort)

the throughput increase along paths with no bottleneck links begins to dominate, and throughput starts to increase again. 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. Similar observations regarding the relative performance of the CMDP and WP policies can be made from Figures 6-7 which plot the weighted throughput with the non-uniform traffic pattern.

VII. Concluding Remarks

We have developed wavelength sharing policies for OBS networks to maximize throughput while providing absolute QoS guarantees in terms of the end-to-end burst loss. Our randomized threshold policies are practical to implement and operate in a distributed manner.

References

[1] E. Altman, T. Jimenez, and G. Koole. On optimal call admission control in a resource-sharing system. *IEEE Trans. Comm.*, 49:1659–1668, 2001.

[2] I. Baldine, G. N. Rouskas, H. G. Perros, and D. Stevenson. `Jump-Start`: A just-in-time signaling architecture for WDM burst-switched networks. *IEEE Communications Mag.*, 40(2):82–89, February 2002.

[3] Dimitri P. Bertsekas. *Dynamic Programming : deterministic and stochastic models*. Prentice-Hall, 1987.

[4] F. Beutler and K. Ross. Uniformization for semi-markov decision processes under stationary policies. *J. Applied Probab.*, 24:644–656, 1987.

[5] Y. Chen, M. Hamdi, and D.H.K. Tsang. Proportional QoS over OBS network. In *GLOBECOM '01*, pages 1510–1514, Nov. 2001.

[6] E. Feinberg and M. I. Reiman. Optimality of randomized trunk reservation. *Probability in the Eng. Info. Sci.*, 8:463–489, 1994.

[7] E. A. Feinberg. Constrained semi-markov decision processes with average rewards. *ZOR-Math. Methods of Oper. Res.*, 39:257–288, 1994.

[8] J. Liu and N. Ansari. Forward resource reservation for QoS provisioning in OBS systems. In *GLOBECOM '02*, 2002.

[9] C-H. Loi, W. Liao, and D-N. Yang. Service differentiation in optical burst switched networks. In *GLOBECOM '02*, pages 2313–2317, 2002.

[10] B. Miller. A queueing reward system with several customer classes. *Management Science*, 16:234–245, 1969.

[11] Sheldon Ross. *Introduction to Stochastic Dynamic Programming*. Academic Press, Inc., 1982.

[12] Li Yang. *Congestion Control and Quality-of-Service (QoS) Provisions for Optical Burst Switched Networks*. PhD thesis, NCSU, May 2006.

[13] M. Yoo, C. Qiao, and S. Dixit. QoS performance of optical burst switching in IP-over-WDM networks. *Journal on Selected Areas in Communications*, 18(10):2062–2071, October 2000.

[14] Q. Zhang, V. Vokkarane, J. Jue, and B. Chen. Absolute QoS differentiation in optical burst-switched networks. *IEEE Journal on Selected Areas in Communications*, 22(9):1781–1795, November 2004.

1-4244-0357-X/06/$20.00 ©2006 IEEE
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2006 proceedings.