A Framework for Absolute QoS Guarantees in Optical Burst Switched Networks

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

***Abstract*— We consider the problem of supporting absolute QoS guarantees in terms of the end-to-end burst loss in OBS networks. We present a parameterized model for wavelength sharing which provides for isolation among different traffic classes while also making efficient use of wavelength capacity through statistical multiplexing. We develop a heuristic to optimize the policy parameters for a single link of an OBS network. We also develop a methodology for translating the end-to-end QoS requirements into appropriate per-link parameters so as to provide network-wide guarantees. Our approach is easy to implement, it can support a wide variety of traffic classes, and is effective in meeting the QoS requirements and keeping the loss rate of best-effort and overall traffic low.**

I. INTRODUCTION

Optical burst switching (OBS) [12] is a promising switching paradigm which aspires to provide a flexible infrastructure for carrying future Internet traffic in an effective yet practical manner. The transmission of each burst is preceded by the transmission of a setup message [1], whose purpose is to reserve switching resources along the path for the upcoming data burst. 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.

As OBS is becoming more widely accepted as a potential transport technology, supporting end-to-end quality of service (QoS) guarantees in OBS networks is arising as an important yet challenging issue. In general, there are two approaches to providing QoS guarantees [15]. In the *relative QoS* model, the service guarantees promised by the network provider to a given class of traffic are specified relative to the service guarantees of another class of traffic. Under the *absolute QoS model*, each priority class is guaranteed a worst-case service level that is independent of the service levels provided to other classes. Most of the 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 [14], intentionally dropping non-compliant bursts [2], 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 [15], 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.

In this paper we develop a general framework for absolute service guarantees to users of an OBS network in terms of the end-to-end burst loss. Inspired by earlier work on resource sharing [5], [6], we first present a parameterized model for wavelength sharing among traffic classes that can provide a desired degree of isolation while taking advantage of statistical multiplexing gains. Then, considering a single OBS link, we develop a heuristic for optimizing the policy parameters to support per-link absolute QoS guarantees for a given set of heterogeneous traffic classes. Finally, we develop a methodology for translating the end-to-end QoS requirements into appropriate per-link parameters so as to provide network-wide guarantees. Our approach is easy to implement and is effective in meeting the QoS requirements and keeping the loss rate of best-effort traffic low.

The paper is organized as follows. In Section II, we discuss the assumptions regarding the OBS network we consider in this study. In Section III we present a suite of parameterized wavelength sharing policies, and in Section IV we develop an algorithm for optimizing the policy parameters for a single OBS link. In Section V we extend our model to an OBS network and introduce an algorithm for determining near-optimal link policy parameters from the end-to-end QoS requirements, traffic statistics, and network properties. We present numerical results to validate our approach in Section VI, and we conclude the paper 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 wavelength from a fixed set of W wavelengths, $\{\lambda_1, \lambda_2, \cdots, \lambda_W\}$. We assume that each OBS node is capable of full wavelength conversion. The network does not use any other contention resolution mechanism, i.e., OBS nodes do not employ buffering, deflection routing, or burst segmentation.

The network supports P classes of traffic, where P is a small integer. 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 recorded in the setup message that precedes the burst transmission. Intermediate nodes make forwarding decisions by taking into account both the availability of resources and the information regarding the class of a burst. Specifically, an intermediate

1-4244-0425-8/06/$20.00 ©2006 IEEE

node may drop a burst of a lower priority class even when there are wavelengths available at its outgoing link.

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 long-run fraction of bursts from class i that are 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 set $B_P^{e2e} = 1.0$.

The objective of the network provider, and the one we consider in this work, 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}, while at the same time minimizing the loss rate of the best-effort class P.*

In order to achieve this objective, the nodes need to employ mechanisms to allocate wavelength resources to bursts of each class based on its load and worst-case loss requirement. Next, we develop a suite of wavelength sharing policies and evaluate their performance.

III. Wavelength Sharing Policies: The Single Link Case

In this section we consider a single link of an OBS network, and we present a set of policies to support different classes of traffic sharing the wavelength resources of the link. The techniques we propose allow for (limited) resource sharing among classes, but also offer each class varying degrees of protection from other classes. The ideas underlying our policies arise naturally in practice, and have been considered before: in the specific setting of memory allocation in network nodes [6], and in the more general context of resource sharing [5]. Our main contribution is to develop analytical methods to calculate the burst loss probability for the various traffic classes under each policy. The analytical methods are the first step towards the design of effective mechanisms to provide absolute end-to-end QoS guarantees in OBS networks, a task we undertake in the following two sections.

We assume that the (unidirectional) OBS link under study consists of W parallel wavelengths, and carries P classes of bursts. The policies we consider manage the wavelength space by associating with each traffic class a pair of values that impose bounds on the use of the link's transmission resources by the class:

- W_i^{max}, referred to as *wavelength upper bound for class* i, is the maximum number of wavelengths that may be occupied simultaneously by bursts of class i. Setting $W_i^{max} < W$ ensures that class i bursts will not consume all available wavelengths at any given time, thus providing a form of protection to other traffic classes from class i.
- W_i^{min}, referred to as *wavelength lower bound for class* i, is the minimum number of wavelengths set aside (reserved) by the link for class i bursts. Whenever $W_i^{min} > 0$, the lower bound guarantees that there is always space for a specified number of bursts from class i, in essence protecting this class in case other classes experience (transient or permanent) overload.

By specifying values for the pair of bounds (W_i^{min}, W_i^{max}) for each traffic class i, a policy may strike any desired balance between two conflicting objectives: *QoS protection*, through class separation, and *efficient utilization*, through sharing of wavelength resources.

We note that a *complete wavelength sharing* policy dictates that:

$$W_i^{min} \quad = \quad 0, \qquad W_i^{max} \quad = \quad W, \qquad i = 1, \cdots, P \tag{2}$$

Such a policy offers no protection, and cannot provide any differentiation among bursts with respect to loss guarantees. Therefore, we do not consider this policy here.

We now present four broad classes of policies as determined by the range of values that the lower and upper bounds, W_i^{min} and W_i^{max}, respectively, are allowed to take. We also present analytical models for computing the burst loss probability, assuming that the pair of values (W_i^{min}, W_i^{max}) for each class i are known in advance; how to determine these values so as to achieve the objective stated in Section II is the subject of Section IV. The models are derived based on the assumption that traffic class $i, i = 1, \cdots, P$, is characterized by a Poisson arrival rate λ_i, and mean holding time μ_i. We also let $\rho_i = \lambda_i / \mu_i$ denote the offered load of class i to the link.

A. Wavelength Partitioning (WP)

The WP policy partitions the wavelength space such that each of the P traffic classes has dedicated access to a subset of the W wavelengths. More specifically, the wavelength bounds for the traffic classes are defined as:

$$0 \ < \ W_i^{min} \ = \ W_i^{max} \ = \ W_i \ < \ W, \quad i = 1, \cdots, P \tag{3}$$

with the additional constraint that the sum of the number of wavelength dedicated to each class must equal the number of available wavelengths: $\sum_{i=1}^{P} W_i = W$. More specifically, bursts arriving at a link following the WP policy are handled as follows:

> *when a class-i burst arrives, if the number n_i of wavelengths busy with class-i bursts is less than W_i, the burst is transmitted on any free wavelength; otherwise, it is simply dropped.*

Clearly, the WP policy and the complete sharing policy defined by expression (2) are at the opposite ends of the spectrum of possible wavelength sharing policies.

The WP policy was considered earlier in the context of OBS networks in [15], where it was referred to as dynamic wavelength grouping (DWG). We adopt it here as a baseline policy against which to compare the policies we present next. A link using the WP policy operates as P independent $M/M/m/m$ queueing systems, one per traffic class. The drop probability B_i for class-i bursts can be computed using the well-known Erlang-B formula:

$$B_i = \frac{{\rho_i}^{W_i}/W_i!}{\sum_{j=0}^{W_i} {\rho_i}^j/j!} \tag{4}$$

WP is easy to implement, as at any time t, one only needs to keep track of the number of wavelengths occupied by bursts of each class. Its main drawback is the lack of statistical multiplexing of bursts from different classes, which can lead to a substantial increase in the number of wavelengths required to guarantee a given level of QoS for each class. As suggested in [6], the performance of complete partitioning can be improved if some sharing of resources is introduced. In the next three subsections we describe policies which provide different levels of wavelength sharing among the various traffic classes.

B. Wavelength Sharing with Maximum Occupancy (WS-Max)

In this scheme, all classes share the whole wavelength space, but we restrict the level of sharing by imposing an upper bound, $W_i^{max} < W$, on the number of wavelengths that class i can use at any given time. On the other hand, the wavelength lower bound for each class is set to zero:

$$0 = W_i^{min} < W_i^{max} < W, \quad i = 1, \cdots, P \tag{5}$$

To allow for wavelength sharing, the sum of the wavelength upper bounds over all traffic classes must exceed the number of available wavelengths, i.e., $\sum_{i=1}^{P} W_i^{max} > W$. More formally, the WS-Max policy operates as follows:

> *when a class-i burst arrives, if the number n_i of wavelengths busy with class-i bursts is less than W_i^{max}, the burst is transmitted on any free wavelength if one exists; otherwise the burst is dropped.*

C. Wavelength Sharing with Minimum Provisioning (WS-Min)

The wavelength sharing with minimum provisioning (WS-Min) permanently allocates a number W_i^{min} of wavelengths to class i, and allows the remaining wavelengths to be shared by all classes. In other words, the wavelength lower and upper bounds are defined as:

$$0 < W_i^{min} < W_i^{max} = W, \quad i = 1, \cdots, P \tag{6}$$

with the additional constraint that the sum of wavelength lower bounds be less than the number of wavelengths W, in order to allow for sharing among classes: $\sum_{i=1}^{P} W_i^{min} < W$. The operation of an OBS link with the WS-Min policy is specified as follows:

> *when a class-i burst arrives to find the link at state $P(\underline{n}) = (n_1, \cdots, n_P)$, it is transmitted on any free wavelength if the number n_i of wavelengths busy with class-i bursts is less than the maximum number of wavelengths that class i may use at that time:*

$$n_i < \max \left\{ W_i^{min}, W - \sum_{k \neq i} \max \left\{ n_k, W_k^{min} \right\} \right\} \tag{7}$$

> *Otherwise, the burst is dropped.*

D. Wavelength Sharing with Minimum Provisioning and Maximum Occupancy (WS-MinMax)

The WS-Max policy prevents any single traffic class from occupying all available wavelengths of the OBS link by imposing the wavelength upper bounds W_i^{max}. However, it cannot provide guarantees to a given class since it is possible for a few aggressive classes to consume all of the link's transmission resources. The WS-Min policy, on the other hand, can guarantee a minimum level of performance to each traffic class through the wavelength lower bounds W_i^{min}. However, it does not impose any constraints on the shared wavelengths, which may lead to unfair utilization of these resources. The wavelength sharing with minimum provisioning and maximum occupancy (WS-MinMax) combines the features of the WS-Max and WS-Min policies to provide per-class QoS guarantees and high link utilization.

The WS-MinMax policy reserves a number W_i^{min} wavelengths to be used exclusively by class i, but it also restricts the number of wavelengths that can be occupied simultaneously by class-i bursts to W_i^{max}:

$$0 < W_i^{min} < W_i^{max} < W, \quad i = 1, \cdots, P \tag{8}$$

In addition, the following constraints are imposed on the wavelength lower and upper bounds to ensure a certain level of wavelength sharing among the traffic classes:

$$\sum_{i=1}^{P} W_i^{min} < W, \quad \sum_{i=1}^{P} W_i^{max} > W \tag{9}$$

Since WS-MinMax is a generalization of both WS-Min and WS-Max, its operation can be described as follows:

> *when a class-i burst arrives to find the link at state $P(\underline{n}) = (n_1, \cdots, n_P)$, it is transmitted on any free wavelength if the number n_i of wavelengths busy with class-i bursts is less than the maximum number of wavelengths that class i may use at that time:*

$$n_i < \min \left\{ W_i^{max}, W - \sum_{k \neq i} \max \left\{ n_k, W_k^{min} \right\} \right\} \tag{10}$$

> *Otherwise, the burst is dropped.*

To obtain the burst drop probability under the WS-MinMax policy, we observe that the state of the OBS link can be described by the vector $\underline{n} = (n_1, \cdots, n_P)$, where n_i is a nonnegative random variable denoting the number of class-i bursts. The evolution of the link is described by a Markovian process whose feasible state space is defined by the following expression:

$$S = \left\{ \begin{array}{c} \underline{n} \mid 0 \leq \sum_{i=1}^{P} \max\left\{W_i^{min}, n_i\right\} \leq W, \\ 0 \leq n_i \leq W_i^{max}, i = 1, \cdots, P \end{array} \right\} \qquad (11)$$

The steady state probability of the Markovian process has a product form solution [7]. Let $C^{-1}(W, \underline{W}^{min}, \underline{W}^{max})$ denote the inverse of the normalizing constant for an OBS link with W wavelengths and vectors of lower and upper wavelength bounds $\underline{W}^{min}$ and $\underline{W}^{max}$, respectively. An effective algorithm for calculating the normalizing constant (and, consequently, the steady-state blocking probabilities) for a class of resource-sharing models was proposed in [3]. This algorithm is based on the numerical inversion approach introduced in [4]. In this work, we adopt the direct method in [3], which is appropriate for the system sizes we consider, and we calculate the normalizing constant via the appropriate P-fold nested sum.

We observe that the probability that a class-i burst would be dropped at an arbitrary time is equal to one minus the probability that class i can be allocated one wavelength at that time. Let $\underline{1}_i$ denote a P-element vector with all elements equal to zero, except the element at position i which is equal to one. Then, the probability that a class-i burst will be dropped at an arbitrary time can be represented as:

$$B_i \quad = \quad 1 - \frac{C^{-1}(W-1, \underline{W}^{min} - \underline{1}_i, \underline{W}^{max} - \underline{1}_i)}{C^{-1}(W, \underline{W}^{min}, \underline{W}^{max})} \qquad (12)$$

Due to Poisson arrivals, (12) also represents the probability that an arriving class-i burst will be dropped.

The burst loss probability expression can be easily adapted to either WS-Min (by fixing the wavelength upper bound of each class to W) or WS-Max (by fixing the wavelength lower bound of each class to 0).

IV. Policy Optimization

We now present a method for selecting the wavelength lower and upper bounds so as to keep the burst drop probabilities below a desired threshold. Our goal is to control the level of resource sharing at the link level in a near-optimal manner in order to achieve absolute QoS differentiation among the traffic classes.

We again consider an OBS link with W wavelengths and P classes of traffic. Each class i is characterized by a worst-case *link* (or *one-hop*) loss guarantee B_i^{ℓ}, which corresponds to the fraction of bursts from class i that are dropped by the link in the long run. We defer to the next section the issue of translating the end-to-end loss guarantees B_i^{e2e} to appropriate link loss guarantees B_i^{ℓ}. As before, we assume that class i has stricter QoS requirements than class $j > i$:

$$B_i^{\ell} \quad < \quad B_j^{\ell}, \qquad 1 \leq i < j \leq P \qquad (13)$$

Traffic class P, the best-effort class, has no associated worst-case loss guarantee, and we let $B_P^{\ell} = 1.0$.

Under the WP policy, the OBS link reserves $W_i (= W_i^{min} = W_i^{max})$ wavelengths for the exclusive use of class-i bursts. Let $Erl^{-1}(\rho, B)$ denote the inverse Erlang-B formula, which returns the number of wavelengths required for the drop probability not to exceed B, when the load is equal to ρ. As it was pointed out in [15], each guaranteed class i must be allocated W_i wavelengths such that:

$$W_i \quad = \quad Erl^{-1}(\rho_i, B_i^{\ell}), \qquad i = 1, \cdots, P-1 \qquad (14)$$

As long as the number of reserved wavelengths, $W_{res} = \sum_{i=1}^{P-1} W_i$, is less than the number W of wavelengths, the best-effort class, class P, will use the remaining unreserved wavelengths. If, however, $W_{res} \geq W$, then it is not feasible to carry the offered traffic mix with the given link capacity using the WP policy. In this case, it may still be possible to meet the QoS requirements of the guaranteed classes and also carry the best-effort class without additional capacity, by exploiting the statistical multiplexing gains achievable by the wavelength sharing policies. For the remainder of this section, we will focus only on the WS-MinMax policy.

Let ρ_i and B_i^{ℓ} be the offered load and link loss guarantee, respectively, of traffic class $i, i = 1, \cdots, P$ (with $B_P^{\ell} = 1.0$). Our objective is to determine the optimal pair of wavelength bounds (W_i^{min}, W_i^{max}) for each class so as to minimize the burst loss probability B_P of the best-effort traffic while keeping the burst loss probability B_i of each guaranteed class $i, i = 1, \cdots, P-1$ below B_i^{ℓ}. More formally, this optimization problem can be stated as:

minimize: B_p
subject to:

$$B_i \leq B_i^{\ell}, \qquad i = 1, \cdots, P-1 \qquad (15)$$

$$0 \leq W_i^{min} \leq W_i^{max} \leq W, \qquad i = 1, \cdots, P \qquad (16)$$

$$W_i^{min}, W_i^{max} : \text{integer}, \qquad i = 1, \cdots, P \qquad (17)$$

where $B_i, i = 1, \cdots, P$, are obtained from expression (12).

Clearly, the above is an integer optimization problem with a nonlinear objective function and nonlinear constraints (15). Furthermore, important mathematical properties such as monotonicity and convexity have not been established for this type of objective function [5]. Since existing optimization tools (e.g., CPLEX) are not appropriate for this problem and an exhaustive search of the entire space of candidate solutions is computationally prohibitive, next we develop a greedy local search heuristic to obtain a near-optimal solution.

A. The Local Search Heuristic

The main idea of the heuristic is to attempt to decrease the value of the objective function (i.e., the drop probability of the best-effort class P), by slightly increasing at each iteration the drop probability of one of the guaranteed classes, say, class $i, i = 1, \cdots, P-1$. However, the algorithm ensures that at the end of the iteration, the loss guarantee of class i will not be violated. The algorithm manipulates the values of the drop probabilities by adjusting the wavelength lower and upper bounds of classes i and P at each iteration. For the selected guaranteed class i, in particular, the algorithm attempts to increase its drop probability by searching in directions which

$$\mathcal{L}(k+1) = \begin{cases} \{(W_m^{min}(k), W_m^{max}(k)-1), (W_m^{min}(k)+1, W_m^{max}(k)-1)\}, & W_m^{min}(k) = 0 \\ \{(W_m^{min}(k)-1, W_m^{max}(k)-1), (W_m^{min}(k)-1, W_m^{max}(k)), & \\ (W_m^{min}(k), W_m^{max}(k)-1), (W_m^{min}(k)+1, W_m^{max}(k)-1), & \\ (W_m^{min}(k)-1, W_m^{max}(k)+1)\}, & W_m^{min}(k) > 0 \end{cases} \quad (18)$$

(1) reduce its maximum usage of wavelengths, (2) reduce its minimum allocation of wavelengths, or both.

The heuristic works as follows. Let $(W_i^{min}(k), W_i^{max}(k))$ denote the pair of wavelength lower and upper bounds for class $i, i = 1, \cdots, P$, at the end of iteration k. Let also $B_i(k)$ denote the burst drop probability of class i at the end of the k-th iteration, as computed by expression (12). At the start of the $(k+1)$-th iteration, the algorithm computes the ratio $\frac{B_i(k)}{B_i^\ell}$ for each guaranteed class $i, i = 1, \cdots, P-1$. This ratio is a measure of how close the long-term burst drop probability of a class is to its link loss guarantee. Let m be the class for which $\frac{B_m(k)}{B_m^\ell}$ is minimum among all guaranteed classes. Note that the constraint in (15) corresponding to class m has the largest relative slack among all such constraints. In the current (i.e., $(k+1)$-th) iteration, the algorithm will modify the wavelength lower and upper bounds of classes m and P in an attempt to lower the burst drop probability $B_P(k+1)$ of the best-effort class at the expense of class-m bursts which may experience a higher drop probability $B_m(k+1)$ (the latter, however, is not allowed to exceed B_m^ℓ). The algorithm does not modify the wavelength lower and upper bounds of any other class during this iteration.

Let us now describe how the algorithm attempts to increase the burst drop probability of guaranteed class m that was selected at the beginning of the $(k+1)$-th iteration. Let $(W_m^{min}(k), W_m^{max}(k))$ be the pair of wavelength lower and upper bounds for this class at the end of the k-th iteration. At the end of the $(k+1)$-th iteration, the algorithm will determine new bounds $(W_m^{min}(k+1), W_m^{max}(k+1))$ for this class. In order to bound the computational requirements of each iteration, the heuristic limits the set of candidate values for $(W_m^{min}(k+1), W_m^{max}(k+1))$ that it considers to a small neighborhood around $(W_m^{min}(k), W_m^{max}(k))$; this is the "local search" feature of the algorithm. Specifically, the local neighborhood examined during the $(k+1)$-th iteration is defined in expression (18), shown at the top of the page. Hence, the wavelength lower and upper bounds of class m will not be adjusted by more than one unit (up or down) at any iteration, preventing large changes in the drop probabilities from one iteration to the next.

For each pair (w_m^{min}, w_m^{max}) in the local neighborhood set $\mathcal{L}(k+1)$, and using the same wavelength lower and upper bounds $(W_i^{min}(k), W_i^{max}(k))$ as at the end of the the previous iteration for all guaranteed classes $i \neq m$, we determine through expression (12) a pair of wavelength lower and upper bounds (w_P^{min}, w_P^{max}) for the best-effort class that minimizes its burst drop probability B_P and does not violate any of the loss guarantees. Among these, we select the pairs (w_m^{min}, w_m^{max}) and (w_P^{min}, w_P^{max}) corresponding to the minimum B_P as the values for $(W_i^{min}(k+1), W_i^{max}(k+1))$ and $(W_P^{min}(k+1), W_P^{max}(k+1))$, respectively. For all other classes we let $W_i^{min}(k+1) = W_i^{min}(k)$ and $W_i^{max}(k+1) = W_i^{max}(k)$, at the end of iteration $k+1$. The algorithm proceeds similarly with the next iteration, and terminates when no improvement in the value of the objective function B_P is possible.

To fully specify the algorithm, we need to determine initial values for the wavelength lower and upper bounds of each class. We use the information regarding the loss guarantees $B_i^\ell, i = 1, \cdots, P-1$, to start the algorithm from an appropriate initial solution. Let W_i denote the number of wavelengths returned by the inverse Erlang-B formula for guaranteed class i. At the beginning of the algorithm, for the guaranteed classes we let:

$$W_i^{min}(0) = W_i, W_i^{max}(0) = \min\{2W_i, W\}, i = 1, \cdots, P-1 \quad (19)$$

while for the best-effort class we set $W_P^{min}(0)$ and $W_P^{max}(0)$ to the pair of values that minimizes $B_P(0)$ while not violating constraints (15).

A step-by-step description of the local search algorithm is provided in Figure 1. Our experimental results indicate that the algorithm converges to a local optimum after only a few iterations.

V. Wavelength Sharing Policies in an OBS Network

We now consider an OBS network with P traffic classes, where each link operates under the WS-MinMax policy. Typically, applications specify their QoS requirements in terms of an end-to-end loss guarantee, and we assume that each class i is associated with an end-to-end loss rate threshold B_i^{e2e}; without loss of generality, we let:

$$B_1^{e2e} < B_2^{e2e} < \cdots < B_{P-1}^{e2e} < B_P^{e2e} = 1.0 \quad (20)$$

The main issue we address in this section is how to optimize the parameters of the WS-MinMax policy at each link, so that the network will meet the end-to-end loss requirements of the guaranteed classes while minimizing the loss probability of the best-effort class P.

Consider any link of the network, and recall that in order to apply the policy optimization algorithm in Figure 1 we need to determine the link offered load ρ_i and link loss rate guarantee B_i^ℓ for each class i. The offered load ρ_i can be determined in several different ways. If the network uses fixed routing, and making the reasonable assumption that link drop probabilities are relatively small, we can approximate ρ_i by summing the amount of class-i traffic offered by source-destination pairs

WS-MinMax Policy Optimization for an OBS Link

Input: An OBS link with W wavelengths, P traffic classes, offered load ρ_i and burst loss guarantee $B_i^\ell, i = 1, \cdots, P$ ($B_P^\ell = 1.0$)

Output: Pair of wavelength lower and upper bounds $(W_i^{min}, W_i^{max}), i = 1, \cdots, P$, such that $B_i \le B_i^\ell, i = 1, \cdots, P-1$, and B_P is minimized

procedure **PolicyOpt**
begin
$k \leftarrow 0$ // iteration index
for $i = 1$ to $P-1$ **do** // initialization
  $W_i^{min}(k) \leftarrow Erl^{-1}(\rho_i, B_i^\ell)$; $W_i^{max}(k) \leftarrow \min\{2W_i^{min}(k), W\}$
$(W_P^{min}(k), W_P^{max}(k)) \leftarrow$ pair of values that minimizes $B_P(0)$ without violating constraints (15)
repeat // main iteration
  $k \leftarrow k+1$
  Let m be the class with the minimum value of $\frac{B_i(k-1)}{B_i^\ell}, i = 1, \cdots, P-1$
  $\mathcal{L}(k) \leftarrow$ the local neighborhood from expression (18)
  $B \leftarrow 1.0$ // temporary variable
  for each $(w_m^{min}, w_m^{max}) \in \mathcal{L}(k)$ **do** // update the wavelength bounds of classes m and P
    $(w_P^{min}, w_P^{max}) \leftarrow$ pair of values that minimizes B_P without violating constraints (15)
    if $B_P < B$ **then**
      $W_m^{mim}(k+1) \leftarrow w_m^{min}$; $W_m^{max}(k+1) \leftarrow w_m^{max}$; $W_P^{min}(k+1) \leftarrow w_P^{min}$; $W_P^{max}(k+1) \leftarrow w_P^{max}$
  for $i = 1$ to $P-1, i \neq m$ **do** // wavelength bounds of other classes remain the same
    $W_i^{min}(k+1) \leftarrow W_i^{min}(k)$; $W_i^{max}(k+1) \leftarrow W_i^{max}(k)$
until B_P cannot be decreased any further
if $\sum_{i=1}^{P-1} W_i^{min} \ge W$ **then return** error // cannot meet QoS guarantees
else return $(W_i^{min}(k), W_i^{max}(k)), i = 1, \cdots, P$
end

Fig. 1. Local search heuristic for policy optimization

whose path uses this link. Alternatively, the OBS node at the head of the link may periodically measure the amount of class-i traffic passing through.

Let us now turn our attention to the problem of determining the per-link loss rate guarantees B_i^ℓ from the end-to-end guarantees $B_i^{e2e}, i = 1, \cdots, P-1$. Consider the burst traffic between a certain source-destination pair and let h denote the number of links (hops) in the path. Let us further make the common assumption that link drop probabilities are independent. In this case, we can guarantee that the end-to-end loss requirement of traffic class B_i^{e2e} for this source-destination pair will be met by letting the loss thresholds at each of the h links equal to:

$$B_i^\ell(h) \;=\; 1 - \exp\left(\frac{\ln(1 - B_i^{e2e})}{h}\right), \; i = 1, \cdots, P-1 \quad (21)$$

Note, however, that a link may carry class-i traffic from several source-destination pairs using paths of different lengths. Let D denote the diameter of the network. One possible way of dealing with this issue would be to subdivide class-i traffic into D subclasses, where each subclass h corresponds to class-i traffic traveling over an h-link path. While theoretically possible, the computational requirements of such an approach would be prohibitive in practice, due to the explosion in the number of traffic classes involved in evaluating expression (12) and the corresponding increase in the running time of the policy optimization algorithm.

A simple solution to this problem was suggested in [15], where it was proposed to set the loss guarantee at each link to the value $B_i^\ell(D)$ obtained by using the diameter D of the network in place of h in expression (21). This simple approach has the additional advantage that the values of B_i^ℓ are identical for all links of the network. A limitation of this method is that by using the diameter of the network in the above expression will result in over-provisioning link resources to guaranteed classes. Consequently, the network resources may not be sufficient to meet the QoS requirements of all classes, and/or the best-effort class may suffer losses that are unnecessarily high [15]. To alleviate the over-provisioning effect, it would be possible to partition the network into clusters whose diameter does not exceed a predefined threshold, and apply the above method to paths within each cluster. Maintaining multiple clusters, on the other hand, requires the use of intelligent partitioning techniques, increases complexity, and results in different per-link loss thresholds for each class.

We now propose another approach which is relatively simple to implement and specifies the same loss rate requirement B_i^ℓ at all links of the network. Let $\bar{H}$ denote the average number of hops, over all source-destination pairs, of a path in the network, and let $B_i^\ell(\bar{H})$ be the corresponding value of expression (21). Note that since $\bar{H} < D$, then $B_i^\ell(\bar{H}) > B_i^\ell(D)$. The first step in our approach is to check whether letting $B_i^\ell(\bar{H})$ as the per-link loss rate guarantee B_i^ℓ for class $i, i = 1, \cdots, P$ is sufficient to meet the end-to-end QoS. To this end, we compute the network-wide end-to-end burst loss probability of class-i traffic as [10]:

$$\mathcal{B}_i = \frac{\sum_{l \in E} B_i^\ell \times \rho_i^{(l)}}{\sum_{s,d} \rho_i^{(s,d)}}, \quad i = 1, \cdots, P-1 \tag{22}$$

where E is the set of links in the OBS network, $\rho_i^{(l)}$ is the total load of class-i traffic offered to link l, and $\rho_i^{(s,d)}$ is the class-i traffic load generated by source-destination pair (s, d). If $\mathcal{B}_i < B_i^{e2e}$ for all guaranteed classes i, we let $B_i^\ell = B_i^\ell(\bar{H})$ for all links in the network, and we stop: this value of per-link loss guarantee is sufficient to meet the end-to-end QoS requirements of all classes, as well as to ensure a low value for the end-to-end loss rate of the best-effort class P.

If, on the other hand, there is some class i for which $\mathcal{B}_i > B_i^{e2e}$, then we need to impose more stringent per-link guarantees in order to meet the end-to-end QoS requirements. We now observe that the feasible values of the per-link guarantee for class i are in the range $[B_i^\ell(D), B_i^\ell(\bar{H})]$. A natural approach for searching this range of values is to perform a binary search, where at each step with let $B_i^\ell, i = 1, \cdots, P$, be the midpoint $B_i^{mid} = (B_i^{min} + B_i^{max})/2$ of the current interval $[B_i^{min}, B_i^{max}]$, where initially we let $[B_i^{min}, B_i^{max}] = [B_i^\ell(D), B_i^\ell(\bar{H})]$. If, using expression (22), this value B_i^{mid} is sufficient to meet the end-to-end QoS requirements, the search continues in the interval $[B_i^{mid}, B_i^{max}]$; otherwise, it continues in the interval $[B_i^{min}, B_i^{mid}]$. This binary search algorithm repeats in this manner until the length of the search range becomes sufficiently small, i.e., until $B_i^{max} \le B_i^{min} \times \epsilon$, where $\epsilon > 1$ is a small constant. At that point, we let the per-link loss guarantee $B_i^\ell = B_i^{min}, i = 1, \cdots, P-1$.

The details of the binary search algorithm are in Figure 2. For comparisons involving vectors, if any one element of the vector violates the comparison conditions, then the vector itself is assumed to also violate them.

VI. Numerical Results

A. Policy Optimization at a Single OBS Link

Let us first 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. While there are no guarantees associated with best-effort class 3, it is desirable to keep its burst drop probability as low as possible provided that doing so does not lead to a violation of the QoS requirements of the two priority classes.

In this subsection, we compare two policies in terms of their effectiveness in meeting the above objective:

1) The WP policy, described in Section III-A and also considered in [15], reserves W_i wavelengths for the exclusive use of class-i bursts. For each guaranteed class $i, i = 1, 2$, the number W_i of wavelengths is determined by the inverse Erlang-B formula (14).
2) The WS-MinMax policy, described in Section III-D, which associates a pair of wavelength lower and upper bounds (W_i^{min}, W_i^{max}) with each traffic class. The values of these bounds are obtained by running the policy optimization algorithm in Figure 1.

Figure 3 plots the burst drop probability against the link load ρ, in Erlang, for the three classes of traffic under the two policies, WP and WS-MinMax; it also plots the average burst drop probability over all three classes of traffic. For this figure, we assume that class-1 (respectively, class-2) bursts represent 20% (respectively, 30%) of the traffic, and the remaining traffic is best-effort; in other words, $\rho_1 = 0.2\rho$, $\rho_2 = 0.3\rho$, and $\rho_3 = 0.5\rho$. As we can see, both policies ensure that the burst loss rate for classes 1 and 2 is kept below the loss requirement of 10^{-3} and 10^{-2}, respectively. On the other hand, the burst loss for class 3 increases with the link load ρ, as expected. But whereas class 3 burst loss under the WP policy is quite high across all load values shown in the figure, under the WS-MinMax policy, class 3 burst loss is one to two orders of magnitude lower for low to moderate traffic loads; even at high loads, the burst loss rate of best-effort traffic under the WS-MinMax policy is one-half that under the WP policy. More importantly, the WS-MinMax policy reduces the overall burst drop rate significantly, with a corresponding substantial increase in throughput (not shown here due to space constraints).

The above result can be explained by noting the two main shortcomings of the WP policy. First, the policy does not allow any statistical multiplexing: it partitions the available link capacity into three sets of wavelengths, each dedicated to carrying bursts in one of the three traffic classes. The WS-MinMax policy, on the other hand, is much more flexible in allocating the link capacity to the three traffic classes. Although it does dedicate a number of wavelengths (equal to the wavelength lower bound) to each of the two guaranteed classes, it does allow for a certain degree (as determined by the policy optimization algorithm in Figure 1) of wavelength sharing among the three classes. The corresponding statistical multiplexing gains contribute to a decrease in the burst loss rate of best-effort, as well as overall, traffic. Hence, the WS-MinMax policy is significantly more efficient and effective in utilizing the available network resources than WP.

A second problem is that the WP policy allocates bandwidth at the granularity of a whole wavelength; as a result, it often overprovisions the guaranteed classes. This is evident from the behavior of the burst loss curves for the guaranteed classes under the WP policy in Figure 3. Consider, for instance, class 1. As we can see, the burst loss initially increases with the link load, but when the load goes from 21 to 21.5 Erlang, the burst loss drops. This behavior is due to the fact that up to 21 Erlang, the WP policy allocates a certain number w wavelengths to class 1 traffic, but at 21.5 Erlang it allocates $w + 1$ wavelengths. In this case, the same number $w + 1$ wavelengths are allocated for loads greater than 21.5

Per-Link Loss Guarantee Optimization for an OBS Network
Input: An OBS network with diameter D and average path length $\bar{H}$, P classes of traffic, and end-to-end loss guarantee vector $\underline{B}^{e2e} = (B_1^{e2e}, \cdots, B_{P-1}^{e2e})$
Output: Per-link loss guarantee vector $\underline{B}^{\ell} = (B_1^{\ell}, \cdots, B_{P-1}^{\ell})$ such that the end-to-end loss guarantees are met and the end-to-end burst loss probability of the best-effort class is minimized

procedure **LinkGuaranteeOpt**
begin
// initialize the search range using expression (21)
$\underline{B}^{min} \leftarrow (B_1^{\ell}(D), \cdots, B_{P-1}^{\ell}(D))$
$\underline{B}^{max} \leftarrow (B_1^{\ell}(\bar{H}), \cdots, B_{P-1}^{\ell}(\bar{H}))$
while $\underline{B}^{max} > \underline{B}^{min} \times \epsilon$ **do** // binary search
    $\underline{B}^{mid} \leftarrow \left(\underline{B}^{min} + \underline{B}^{max}\right)/2$
    $\underline{\mathcal{B}} \leftarrow (\mathcal{B}_1, \cdots, \mathcal{B}_{P-1})$ from expression (22) with $\underline{B}^{\ell} = \underline{B}^{min}$
    if $\underline{\mathcal{B}} < \underline{B}^{e2e}$ **then** // attempt to increase the link guarantees to decrease B_P
        $\underline{B}^{min} \leftarrow \underline{B}^{mid}$
    else // must decrease the link guarantees
        $\underline{B}^{max} \leftarrow \underline{B}^{mid}$
end while
return $\underline{B}^{min}$
end

Fig. 2. Binary search algorithm for selecting the per-link loss guarantees

Erlang, hence the burst loss for class 1 continues to increase after the drop. Similar observations can be made for the burst loss curve of class 2. The WS-MinMax policy, on the other hand, by virtue of the wavelength sharing it allows, is able to allocate the link capacity at a finer granularity than a whole wavelength. Consequently, it "allocates" just enough capacity to each of the guaranteed classes to meet their loss requirements. Observe also that the burst loss for the guaranteed classes is generally higher under the WS-MinMax policy than under WP. In essence, the WS-MinMax policy reduces the loss rate of best-effort traffic by increasing the loss rate of the guaranteed classes just enough, so as not to violate the corresponding requirement.

For Figure 4, we fix the class 1 and class 2 load to $\rho_1 = 4$ Erlang and $\rho_2 = 6$ Erlang, respectively. The figure plots the burst loss rate of all classes under the WP and WS-MinMax policies against the load ρ_3 of the best-effort class, as the latter varies from 10 to 16.5 Erlang. Since the load of the guaranteed classes is constant, the WP policy allocates them the same number of wavelength regardless of the load of best-effort traffic; as a result, the burst loss of the two guaranteed classes is the same under the WP policy across the range of ρ_3 values. The WS-MinMax policy, on the other hand, adjusts the wavelength lower and upper bounds of the two guaranteed classes depending on the value of ρ_3, hence the behavior of the corresponding burst loss curves is non-monotonic. As a result, the WS-MinMax policy is able to reduce significantly the overall loss rate, and that of the best-effort traffic, without violating the loss requirements of the guaranteed classes.

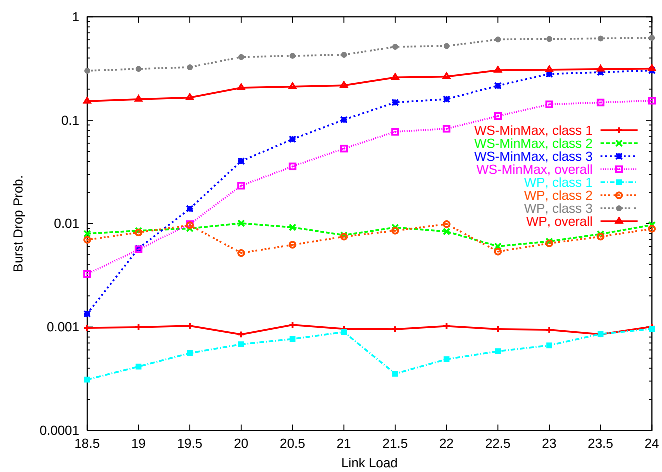

Fig. 3. Single link with $W = 32$ wavelengths and $P = 3$ traffic classes, $\rho_1 = 0.2\rho, \rho_2 = 0.3\rho, \rho_3 = 0.5\rho$

B. End-to-End QoS Guarantees in an OBS Network

We now use simulation to demonstrate the effectiveness of our wavelength sharing policies to provide end-to-end guarantees. We use the simulator that was developed as part of the Jumpstart project [11]. The simulator accounts for all the details of the Jumpstart OBS signaling protocol [1] which employs the Just-In-Time (JIT) reservation scheme. (We emphasize, however, that the wavelength sharing policies we present and evaluate in this work are independent of the specifics of the reservation protocol, and can be deployed alongside either the JET or the Horizon reservation schemes.) We use the method of batch means to estimate the burst drop

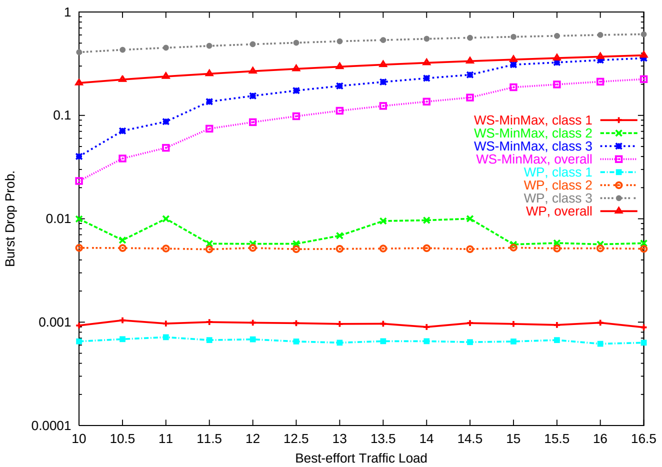

Fig. 4. Single link with $W = 32$ wavelengths and $P = 3$ traffic classes, $\rho_1 = 4$ Erlang, $\rho_2 = 6$ Erlang

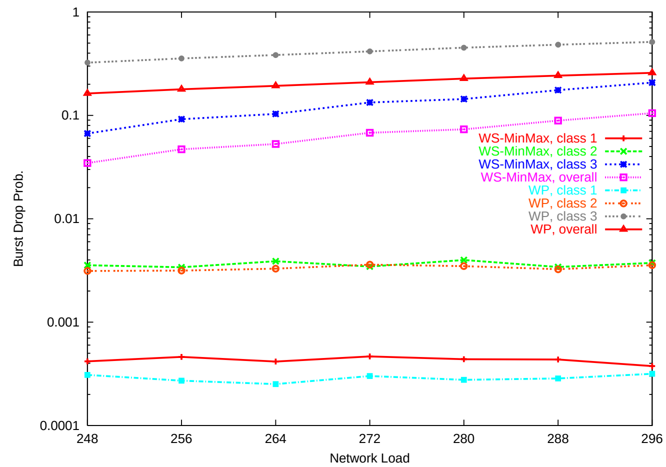

Fig. 5. NSFNet, $W = 32$ wavelengths, $P = 3$ traffic classes, uniform pattern, B_i^ℓ obtained from (21) with $h = D$

probability, with each simulation run lasting until 6×10^5 bursts have been transmitted in the entire network. We have also obtained 95% confidence intervals for all our results; however, they are so narrow that we omit them from the figures we present in this section in order to improve readability.

In our study, we consider a 4×4 regular topology torus network, and a 16-node network based on an irregular topology derived from the 14-node NSF network; the topologies can be found in [13]. We assume shortest path routing, and we consider two different traffic patterns:

- **Uniform pattern:** each switch generates the same traffic load, and the traffic from a given switch is uniformly distributed to other switches.
- **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.

We again assume that 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 $B_1^{e2e} = 10^{-3}$ and $B_2^{e2e} = 10^{-2}$, respectively; class 3 is the best-effort class and does not require any loss guarantees. We also note that the diameter of both the NSFNet and the torus networks is equal to 4, while the average hop distance of the two networks, used in the optimization algorithm in Figure 1, is $\bar{H}_{NSF} = 2.283$ and $\bar{H}_{torus} = 2.133$.

In Figure 5, we plot the overall burst drop probability, as well as that of the three classes of traffic, under the two policies, WP and WS-MinMax, for the NSFNet with the uniform traffic pattern. The results shown were obtained by setting the loss guarantee at each link of the network to the value obtained by using the diameter $D = 4$ of the network in place of parameter h in expression (21); this is the approach suggested in [15]. Figure 6 shows similar results for the torus network. Our observations regarding the relative behavior of the two policies, WP and WS-MinMax, from the two figures are similar to the ones we discussed in the previous section. Specifically, both policies guarantee that the burst loss of classes 1 and 2 is kept below the corresponding requirements,

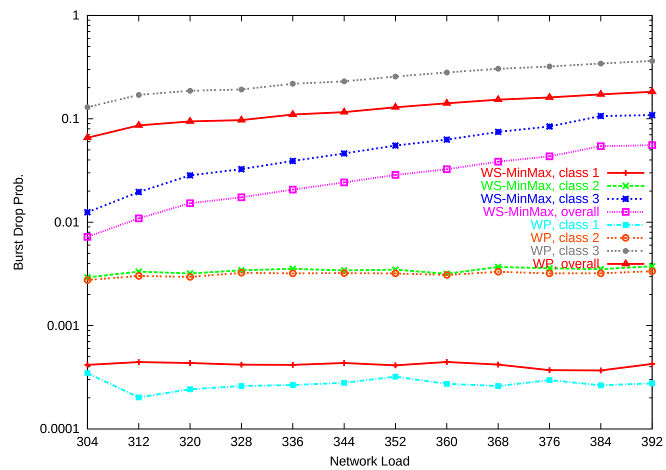

Fig. 6. Torus, $W = 32$ wavelengths, $P = 3$ traffic classes, uniform pattern, B_i^ℓ obtained from (21) with $h = D$

but the WS-MinMax policy achieves a burst loss for the overall and best-effort traffic that is significantly less than that under the WP policy. However, we also observe that using the diameter $D = 4$ to obtain the link-loss guarantees results in overprovisioning of the network for the guaranteed classes. Indeed, the network-wide burst loss of class 1 (respectively, class 2) is significantly less than the required guarantee of 10^{-3} (respectively, 10^{-2}).

In order to alleviate the overpovisioning problem, we used the optimization procedure in Figure 1 to determine an appropriate value for the link-loss guarantee $B_i^\ell, i = 1, 2$, given the corresponding end-to-end loss guarantee B_i^{e2e}. The simulation results are shown in Figures 7 and 8, for the NSFNet and torus networks, respectively. Comparing to Figures 5 and 6, we can see that using a higher value for B_i^ℓ results in a higher end-to-end burst loss probability for class 1 and class 2 bursts, as expected. However, the burst loss of the guaranteed classes is kept well below their requirements. Furthermore, the burst loss of best-effort traffic is reduced, as its bursts can use additional wavelength resources that were previously dedicated to the

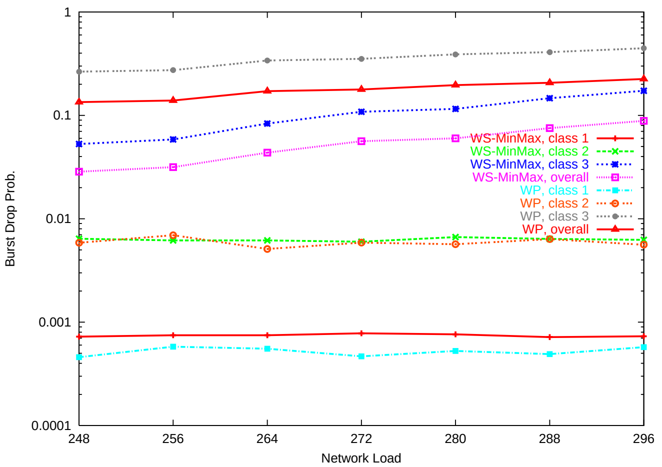

Fig. 7. NSFNet, $W = 32$ wavelengths, $P = 3$ traffic classes, uniform pattern, B_i^ℓ obtained by the optimization procedure in Figure 1

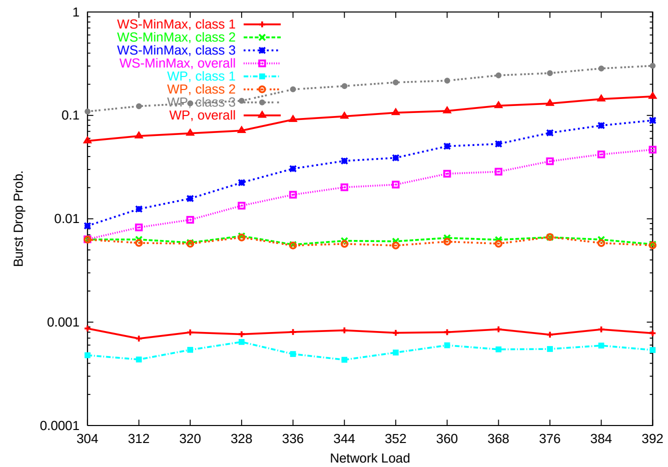

Fig. 8. Torus, $W = 32$ wavelengths, $P = 3$ traffic classes, uniform pattern, B_i^ℓ obtained by the optimization procedure in Figure 1

guaranteed traffic; as a result, the overall burst loss is also reduced.

Similar results from the distance-dependent pattern are omitted due to space constraints; they can be found in [13].

VII. Concluding Remarks

We have presented a framework for supporting absolute QoS guarantees in OBS networks, consisting of a link wavelength sharing model, and a method to translate end-to-end loss guarantees into per-link guarantees. Our approach is effective and efficient in managing the wavelength resources, is simple to implement, and outperforms previously proposed methods.

References

[1] 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 Magazine*, 40(2):82–89, February 2002.

[2] Y. Chen, M. Hamdi, and D.H.K. Tsang. Proportional QoS over OBS network. In *Proceedings of IEEE GLOBECOM*, volume 3, pages 1510–1514, November 2001.

[3] G. L. Choudhury, K. K. Leung, and W. W. Att. An algorithm to compute blocking probabilities in multi-rate multi-class multi-resource loss models. *Advances in Applied Probability*, 27:1104–1143, 1995.

[4] G. L. Choudhury, K. K. Leung, and W. Whitt. Calculating normalization constants of closed queueing networks by numerically inverting their generating functions. *Journal of the ACM*, 1993.

[5] G. L. Choudhury, K. K. Leung, and W. Whitt. Efficiently providing multiple grades of service with protection against overloads in shared resources. *AT&T Technical Journal*, 74:50–63, 1995.

[6] F. Kamoun and L. Kleinrock. Analysis of shared finite storage in a computer network node environment under general traffic conditions. *IEEE Transactions on Communications*, 28(7), July 1980.

[7] F. P. Kelly. *Reversibility and Stochastic Networks*. John Wiley & Sons, New York, 1979.

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

[9] C-H. Loi, W. Liao, and D-N. Yang. Service differentiation in optical burst switched networks. In *Proceedings of IEEE GLOBECOM*, volume 3, pages 2313–2317, 2002.

[10] J. Teng and G. N. Rouskas. Routing path optimization in optical burst switched networks. In *Proceedings of the 9th Conference on Optical Network Design and Modeling (ONDM 2005)*, pages 1–10, February 2005.

[11] Jing Teng. *A Study of Optical Burst Switched Networks with the Jumpstart Just-In-Time Signaling Protocol*. PhD thesis, North Carolina State University, Raleigh, NC, August 2004.

[12] L. Xu, H. G. Perros, and G. N. Rouskas. Techniques for optical packet switching and optical burst switching. *IEEE Communications*, 39(1):136–142, January 2001.

[13] Li Yang. *Congestion Control and Quality-of-Service (QoS) Provisions for Optical Burst Switched Networks*. PhD thesis, North Carolina State University, Raleigh, NC, May 2006.

[14] 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.

[15] 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.