On Wavelength Assignment in Optical Burst Switched Networks

Jing Teng, George N. Rouskas
Department of Computer Science, North Carolina State University, Raleigh, NC 27695-7534

***Abstract*— A fundamental assumption underlying most studies of optical burst switched (OBS) networks is that full wavelength conversion is available throughout the network. In practice, however, economic and technical considerations are likely to dictate a more limited and sparse deployment of wavelength converters in the optical network. Therefore, we expect wavelength assignment policies to be an important component of OBS networks. In this paper, we explain why wavelength selection schemes developed for wavelength routed networks are not appropriate for OBS. We then develop a suite of adaptive and non-adaptive policies for OBS switches. We also apply traffic engineering techniques to reduce wavelength contention through traffic isolation. Our performance study indicates that, in the absence of full conversion capabilities, intelligent choices in assigning wavelengths to bursts at the source can have a profound effect on the burst drop probability in an OBS network.**

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 control packet, 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 control packet. For a detailed description, evaluation, and comparison of the various OBS reservation protocols, the reader is referred to [3].

Over the last five years, research in OBS networks has rapidly progressed from purely theoretical investigations to prototypes and proof-of-concept demonstrations. Yet despite the multitude of directions that OBS research has taken, and the broad set of challenges it addresses, there is one fundamental assumption underlying most studies of OBS networks: namely, that full wavelength conversion is available throughout the network. The existence of wavelength conversion capability at optical switches has a profound effect on the performance of an OBS network, since it removes the wavelength continuity constraint. Without wavelength conversion, a switch can forward an incoming burst to an output port *if and only if* the wavelength carrying the burst is available (free) on the output port. Otherwise, wavelength contention arises and the incoming burst is dropped. By allowing a switch to forward an incoming burst to an output port as long as the port has at least one free wavelength, full wavelength conversion eliminates the wavelength continuity constraint altogether, and improves significantly the performance of the OBS network.

Currently, wavelength converters are expensive and complex devices, and this state of affairs is expected to continue in the foreseeable future. Therefore, it is widely expected that any wavelength conversion capabilities in the optical network will be limited and only sparsely deployed. This observation has two important consequences. First, any performance studies relying on the assumption of full wavelength conversion will underestimate the burst drop probability in the network, possibly by a substantial factor, and may also fail to correctly identify the real behavior and dynamics of the network. Second, the absence of (full) conversion necessitates the development of good and efficient wavelength assignment policies. Such policies are even more important in OBS networks than in wavelength routed (circuit-switched) optical networks, due to the fact that in the former, a burst is transmitted without first reserving resources along the path. Therefore, a burst may be dropped at any intermediate switch along its path, even as it enters its last hop before the destination, resulting in substantial waste of network resources.

Although there is a substantial amount of research addressing the wavelength assignment problem in circuit-switched optical networks (for instance, refer to [5] and references thereof), the same problem has received little attention in the context of OBS networks. Recently, a priority-based wavelength assignment (PWA) algorithm was presented in [4]; we discuss the PWA algorithm in detail in Section IV. Another recent work presented an algorithm to reduce wavelength contention in the OBS network by using some information regarding the routing paths [2]. Although these studies represent a step in the right direction, we feel that the issue of wavelength assignment in OBS network has not been adequately addressed, and that much remains to be done in order to develop a good understanding of the problem in all its aspects.

In this paper, we attempt to fill this gap by presenting a comprehensive study of wavelength assignment in OBS networks. Specifically, we develop a suite of wavelength selection

Proceedings of the First International Conference on Broadband Networks (BROADNETS'04)
0-7695-2221-1/04 $ 20.00 IEEE

policies and we evaluate their relative performance in terms of both burst drop probability and fairness with respect to burst path lengths. Our policies attempt to alleviate the effects of wavelength contention by using the wavelength dimension to isolate traffic from different sources that uses overlapping paths through the network. We present two methods to achieve traffic isolation: the first is based on traffic engineering approaches that take into account the network topology and the routing paths to reduce wavelength contention though traffic isolation, while the second uses adaptive selection strategies that respond to feedback from the network. We also show that by appropriately combining the two methods we can achieve substantial improvement in performance.

The next section discusses our main assumptions regarding the OBS network we study. In Section III, we explain why conventional wavelength assignment schemes are not appropriate for OBS networks, and we develop a traffic engineering approach to achieve traffic isolation. In Section IV, we present a number of adaptive, priority-based wavelength assignment schemes, and we show how to combine them with the traffic engineering approach. We present the results of an experimental study of the performance of the various wavelength assignment policies in Section V, and we conclude the paper in Section VI.

II. The OBS Network Under Study

We consider an OBS network with N switches, interconnected in a general topology. Each link in the network can carry burst traffic on any wavelength from a fixed set of W wavelengths, $\{\lambda_1, \lambda_2, \cdots, \lambda_W\}$. The network switches employ the JIT reservation scheme and the associated Jumpstart signaling protocol [1] for JIT OBS networks. We emphasize, however, that the wavelength assignment policies we develop 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 assume that there are no wavelength converters in the OBS network; however, our work can be extended to OBS networks with sparse conversion capabilities. A switch wishing to transmit a burst selects a free wavelength on the outgoing link for the transmission. The optical signal carrying the burst must then remain in the same wavelength on all the links along the path to the destination, unless an intermediate switch is capable of wavelength conversion. A wavelength contention arises when two bursts, which overlap in time, arrive at a switch on the same wavelength and need to use the same output port (outgoing link). We assume that switches have no buffers (electronic or optical) to store bursts; therefore, if the switch does not have any wavelength converters, one of the overlapping bursts is dropped. Consequently, wavelength selection at the source of the burst will critically affect the performance of the network in terms of burst drop probability.

The set of rules used by a switch in selecting the wavelength on which to transmit a burst define a wavelength assignment policy. Wavelength assignment is a hard problem that has been studied extensively in the context of wavelength routed networks [5]. Since wavelength assignment decisions must be made in real time, an efficient implementation approach is to have each switch order the W wavelengths in a wavelength list. When a switch has a new burst to send, it starts at the top of the list and transmits the burst on the first wavelength that is free on the desired outgoing link. Typically, all switches in the network will use the same policy (rules) to order the wavelengths. However, if the policy rules use information on the state of the network to rank wavelengths, the wavelength list at any given time may be different at various switches; furthermore, the wavelength list at a given node may change over time. This operation may result in different choices in wavelength assignment at various switches, and over time at the same switch. We also note that, a wavelength assignment policy is fully defined by describing the set of rules the network switches use to rank wavelengths.

We can classify wavelength assignment policies as adaptive or non-adaptive. In adaptive policies, the rules for ordering wavelengths take into account the network and traffic dynamics, hence, the order in which a given switch considers the wavelengths in search of a free one may change over time. In non-adaptive schemes, on the other hand, the order in which wavelengths are considered by each switch is neither dependent on, nor determined by, the prevailing network conditions. We emphasize that the rules of a non-adaptive policy may dictate a different wavelength list at different switches, or even a different ordering of wavelengths at a given switch over time; however, the rules must be independent of the network dynamics, although they may depend on certain properties of the network, such as topology or routing, that change at longer time scales.

Adaptive wavelength assignment policies depend on feedback from the network in order to adapt their rules to reflect the state of the network. This feedback can take many forms, depending on the specifics of the signaling protocol and the implementation details. The Jumpstart signaling protocol for JIT OBS networks [1] provides such feedback in the form of two messages. The CONNECT message is returned to the source of a burst by the destination switch, and indicates that the burst transmission was successful. The FAILURE message is sent to the source of a burst by an intermediate switch when the latter is forced to drop the burst; certain fields of the FAILURE message indicate the reason for dropping the burst, e.g., "output port unavailable." As a result, the source of a burst can determine whether the burst is successfully received or dropped, and in the latter case, where the drop occurred and whether the cause was wavelength contention. Some of the wavelength assignment policies we develop in this work rely on similar feedback from the network to adapt their rules.

In this paper, we only consider the wavelength assignment problem. For simplicity, we assume fixed-path routing, in that all bursts between a source-destination pair follow the same path. Our work does not preclude changes in the routing paths, however, we make the reasonable assumption that any such changes take place at time scales significantly longer than the diameter of the network.

III. Non-Adaptive Wavelength Assignment

A. First-Fit and Random

The First-Fit and Random wavelength assignment schemes are well-known and have been extensively studied in the context of wavelength routed networks [5]. In First-Fit, the W wavelengths are labeled arbitrarily and are listed in increasing order of label value, say, $\lambda_1, \lambda_2, \cdots, \lambda_W$. This order is identical at all network switches, and remains unchanged throughout the operation of the network. When a switch wishes to select a free wavelength for transmitting its burst, it searches the wavelength list in this order, until either a free wavelength is found and assigned to the burst, or the list is exhausted (in which case, we assume that the burst is dropped). The Random wavelength assignment policy works as follows. We assume that each switch maintains a list of the wavelengths that are busy on each of its outgoing links. Suppose that at a given time, a switch needs to select a wavelength for a burst whose outgoing link has W' free wavelengths, $W' \leq W$. If $W' = 0$, the switch drops the burst; otherwise, it randomly allocates one of the W' free wavelengths to the burst.

It is known that, in wavelength routed networks, where wavelength assignment decisions are based on complete knowledge of wavelength availability along the links of the path, First-Fit minimizes wavelength fragmentation and, hence, performs significantly better than Random in terms of blocking probability [5]. First-Fit is also simple to implement and does not require the exchange of any information among network switches regarding wavelength usage statistics.

However, in OBS networks, a switch must select a wavelength *without* any knowledge of the instantaneous wavelength occupancy of the links along the path. In this context, the First-Fit policy may in fact result in poor performance in terms of burst drop probability. In order to illustrate the problems associated with the First-Fit policy in OBS networks, consider the simple network shown in Figure 1: switches S_1 and S_2 transmit bursts which must travel over link e_3. The switches make wavelength assignment decisions using only local information, without any knowledge of the state of the link e_3 (due to the relatively short duration of bursts, any information that S_1 and S_2 may have regarding the state of link e_3 may already be out-of-date by the time they receive it). Since both switches search for a free wavelength in the same order, it is highly likely to pick the same wavelength, causing one of the bursts to be dropped at switch S_3. With the Random policy, on the other hand, the probability that both switches will select the same wavelength for the transmission is lower, leading to better performance. The performance results we present in Section V confirm this intuition; in fact, our study indicates that First-Fit is the worst policy by far, while Random performs significantly better in relative terms.

B. First-Fit-TE: Combining First-Fit and Traffic Engineering

We now present a modified version of the First-Fit wavelength assignment policy which is designed to overcome the shortcomings of the conventional First-Fit policy in OBS networks. In order to motivate our approach, let us return to the scenario depicted in Figure 1, and assume again that the W wavelengths on each link are labeled $\lambda_1, \cdots, \lambda_W$. It is not difficult to see that, among all wavelength assignment policies that use only local information at switches S_1 and S_2, the following policy would minimize the burst drop probability at switch S_3: one of the two switches (say, S_1) uses the First-Fit policy, and searches for a free wavelength in the order $\lambda_1, \lambda_2, \cdots, \lambda_W$, while the other switch (say, S_2) also uses the First-Fit policy, but searches for a free wavelength in the reverse order $\lambda_W, \lambda_{W-1}, \cdots, \lambda_1$. This policy minimizes the burst drop probability at switch S_3 because switches S_1 and S_2 will select the same wavelength (and thus, a burst will be dropped at switch S_3) if and only if all other wavelengths are busy transmitting bursts. In contrast, other policies using only local information at switches S_1 and S_2 (e.g., Random, conventional First-Fit, etc.) might select the same wavelength at both switches even while other wavelengths are free.

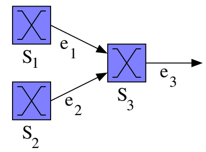

Fig. 1. First-Fit results in high burst drop probability at Switch S_3

While it is straightforward to identify the optimal wavelength assignment policy for the simple network of Figure 1, determining the optimal policy for a large network with a general topology is a difficult task. Therefore, we now present a new wavelength assignment policy that is similar to First-Fit, but uses information regarding the network topology and routing paths to improve upon conventional First-Fit in terms of the burst-drop probability; we will refer to this new policy as First-Fit-TE, where "TE" stands for "traffic engineering."

Consider an OBS network with general topology. The network consists of N switches, and each link can carry W wavelengths. The W wavelengths are labeled arbitrarily as $\lambda_1, \cdots, \lambda_W$, and this order is fixed and known at all N switches. Each switch $S_i, i = 1, \cdots, N$, is assigned a *start wavelength*, $start(i) \in \{\lambda_1, \cdots, \lambda_W\}$. The value of $start(i)$ is determined using a traffic engineering approach we describe shortly, and remains fixed throughout the operation of the network. Furthermore, it is possible that two different switches, S_i and $S_j, j \neq i$, be assigned the same start wavelength, $start(i) = start(j)$.

The First-Fit-TE wavelength assignment policy at switch $S_i, i = 1, \cdots, N$, operates as follows. When the switch has a new burst to transmit, it searches for a free wavelength in the order: $\lambda_{start(i)}, \lambda_{start(i)+1}, \cdots, \lambda_W, \lambda_1, \cdots, \lambda_{start(i)-1}$. The switch transmits the burst on the first free wavelength found, and drops it if all W are found busy. In other words, each switch follows a First-Fit policy, but, unlike the conventional First-Fit scheme that requires all nodes to use the same search sequence, under First-Fit-TE, the starting wavelength of the search sequence can be different for different switches.

Let $d(i, j)$ denote the distance between the start wavelengths of the two switches S_i and S_j in the sequence $\lambda_1, \cdots, \lambda_W$: $d(i,j) = start(j) \ominus start(i)$, where $\ominus$ denotes subtraction modulo-W. We note that, when the network is not heavily loaded, the wavelengths on which a switch S_i transmits its bursts will be close to its start wavelength $start(i)$. Therefore, the main idea behind the First-Fit-TE policy is to assign a start wavelength to each switch in the network in such a manner that, the higher the "interference" among bursts originating at two switches S_i and S_j, the higher the distance $d(i,j)$ between the start wavelengths of the two switches. In this context, we use the notion of "interference" as a measure of the likelihood that bursts generated by different switches will use the same link on the way to their respective destinations.

The level of "interference" among two switches depends on the network topology, the relative location of the switches in the network, the traffic characteristics and the routing algorithm. For instance, bursts from two switches located at diametrically opposite points in a large network are likely to use non-overlapping paths, while bursts originating at two neighboring switches may use paths with substantial overlap; we say that the former pair of switches has low "interference" while the latter pair has high "interference." We now formalize the concept of interference in a quantitative manner. In the following discussion, we assume that the network employs fixed routing so that bursts between a given source-destination pair always follow the same path; however, our main idea can be adapted to apply to other routing schemes.

Let Π_i denote the set of paths taken by bursts originating at switch S_i, $\Pi_i = \{\pi_{i1}, \pi_{i2}, \cdots \pi_{iN}\}$, where π_{ij} is the path from switch S_i to switch S_j. Let also γ_{ij} denote the traffic load from switch S_i to switch S_j. We define the *degree of interference* of a path π_{ij} and a switch S_k, denoted by $ID(\pi_{ij}, k)$, as the amount of traffic from switch S_i to S_j on the path π_{ij} that interferes with traffic originating from switch S_k:

$$ID(\pi_{ij}, k) = \begin{cases} \gamma_{ij}, & \pi_{ij} \text{ shares a link with a path in } \Pi_k \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

We also define the *interference level* between two switches S_i and S_j, which we will denote by $IL(i,j)$, as:

$$IL(i,j) = \begin{cases} \sum\limits_{\pi_{ik} \in \Pi_i} ID(\pi_{ik}, j) & i \neq j \\ 0, & i = j \end{cases} \tag{2}$$

That is, $IL(i,j)$ is the total amount of traffic originating at switch S_i which may interfere (through the use of common network links) with any traffic originating at switch S_j. Finally, we define the *combined interference level* $CIL(i,j)$ between two switches S_i and S_j as the total interference between the two switches:

$$CIL(i,j) = IL(i,j) + IL(j,i), \quad i \neq j \tag{3}$$

With the above definitions, the higher the combined interference level between two switches, the higher the likelihood that bursts from the two switches will share some network link. Therefore, to minimize the probability that bursts from the two switches will collide on a common link, we must ensure that they do not use the same wavelength. In other words, we must assign start wavelengths to the two switches that are far apart from each other. Conversely, if the interference level between two switches is low, their bursts are less likely to share links and collide; consequently, the start wavelengths of the two switches can be close to each other.

Given the interference levels $IL(i,j)$ for all pairs of switches (S_i, S_j) in the network, our objective is to determine the start wavelength $start(i)$ for each switch S_i so as to minimize the burst dropping probability in the network under the First-FIT-TE wavelength assignment policy defined earlier. It might be tempting to formulate this problem as an integer optimization problem and attempt to solve it using standard problem solvers. Unfortunately, it is not possible to express the objective function (i.e., the network-wide burst drop probability) analytically in terms of the problem variables. Even if we chose to formulate the problem in terms of a different objective function for which such an analytical expression is available, two issues would arise. First, there is the question of what would be an appropriate and relevant objective function; and second, even if we were to find an appropriate objective function, the complexity of the resulting problem would preclude the use of optimal solution methods for anything other than small, toy networks.

Instead, we use a simple heuristic to assign start wavelengths to the various switches, which we have found to work well in practice. The heuristic consists of three steps:

1) Partition the set of N switches in K groups (subsets), $g_1, g_2, \cdots, g_K$, such that there is little interference among switches in each group. All the switches in a given group $g_k, k = 1, \cdots, K$, will be assigned the same start wavelength.
2) Arbitrarily label the W wavelengths as $\lambda_1, \cdots, \lambda_W$, and let $x = W/K$ (note that x may not be integer). We evenly space the K start wavelengths across the W wavelengths, such that the k-th start wavelength is the wavelength labeled $\lambda_{1+\lfloor (k-1)x \rfloor}$.
3) We assign the K start wavelengths to each of the K groups so as to minimize the interference level among groups with adjacent start wavelengths.

We now explain the first and third steps of the heuristic.

Partitioning. Typically, partitioning problems with objective functions similar to the one we consider here (i.e., to minimize the interference among switches in each group) are hard optimization problems. Therefore, we use the following greedy heuristic to assign each switch to one of K groups. Let $N = LK + M, M < K$; in our heuristic, the first M groups will consist of $L+1$ switches, and the last $K - M$ groups of L switches. Consider group $g_k, k = 1, \cdots, K$. Initially, $g_k = \emptyset$. Select the switch S_i that has not been assigned to a group yet, such that S_i has the minimum total combined interference level, $\sum_{j=1}^{N} CIL(i,j)$, among unassigned switches. Let $g_k \leftarrow g_k \cup \{S_i\}$. Then, select the unassigned switch S_j that has the minimum combined interference level $CIL(i,j)$ with switch S_i, and let $g_k \leftarrow g_k \cup \{S_j\}$. Continue in this manner, selecting the next switch to add to g_k so as to minimize the overall combined interference level in the group, until the total number of switches in group g_k has been reached. If $k = K$, the algorithm stops; otherwise, it continues with group g_{k+1}.

Assignment of start wavelengths. Again, we use a greedy algorithm to assign start wavelengths to groups of switches in sequential order. First, note that the first start wavelength is always λ_1. We assign this wavelength as the start wavelength of the group g for which the total combined interference level among all switches in g and switches in any other group is minimum (over all K groups). Suppose now that the first $k, k < K$, start wavelengths have been assigned, and let g be the group that was the last to be assigned a start wavelength. Let g' denote the unassigned group such that the total interference among switches in g and switches in g', $\sum_{S_i \in g, S_g \in g'} IL(i,j)$, is minimum. Then, we assign the $(k+1)$-th start wavelength to group g'. The algorithm proceeds in this manner until all groups have been assigned start wavelengths.

To illustrate our approach, let us consider the two network topologies shown in Figures 3 and 4. To simplify the presentation, we assume that the traffic load $\gamma_{ij} = \gamma = 1$ for all switch pairs (S_i, S_j). The 4×4 torus network of Figure 3 has a regular topology, while the 16-node network of Figure 4, has an irregular topology which is obtained by augmenting the 14-node NSFNet topology through the addition of two fictitious switches, S_1 and S_{16}. For each topology, we first run Dijkstra's algorithm to compute the shortest path for each pair of switches. We then computed the interference level $IL(i,j)$ for each pair of switches (S_i, S_j) using expressions (2) and (1), after letting $\gamma_{ij} = 1$ for all i, j. Tables III-B and III-B list the interference levels for each pair of switches in the torus and NSFNet topologies respectively. Assuming that $K = 8$, i.e., that we partition the 16 switches into 8 groups of size 2, the groups for the torus network are: $\{1,11\}$,$\{2,12\}$, $\{3,9\}$,$\{4,10\}$, $\{5,15\}$,$\{6,16\}$, $\{7,13\}$, and $\{8,14\}$. For the NSFNet, on the other hand, the eight groups are: $\{1,14\}$, $\{2,16\}$, $\{3,13\}$, $\{4,8\}$, $\{5,12\}$, $\{6,9\}$, $\{7,11\}$, and $\{10,15\}$. Also, the start wavelength assigned to each group of switches in the NSFNet is shown in Figure 2.

IV. Adaptive Wavelength Assignment Schemes

In adaptive wavelength assignment schemes, the order in which each switch uses to search for an available wavelength changes over time in response to the state of the network and prevailing traffic conditions. A common mechanism to implement adaptive wavelength assignment, which we adopt in this work, is by assigning a *priority* to each wavelength. At any given instant, the priority of a wavelength reflects the likelihood that a burst transmission on this wavelength will be successful. The priorities are updated periodically based on feedback from the network. Specifically, when a switch determines that a burst transmitted on a particular wavelength has been successfully received, it increases the priority of the wavelength; conversely, if the burst is dropped inside the network, the priority of the wavelength carrying the burst is

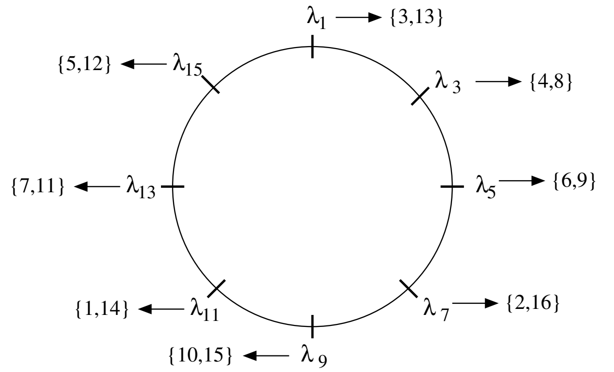

Fig. 2. Start wavelength for each group of switches, NSFNet, $W = 16$ wavelengths, $K = 8$ groups

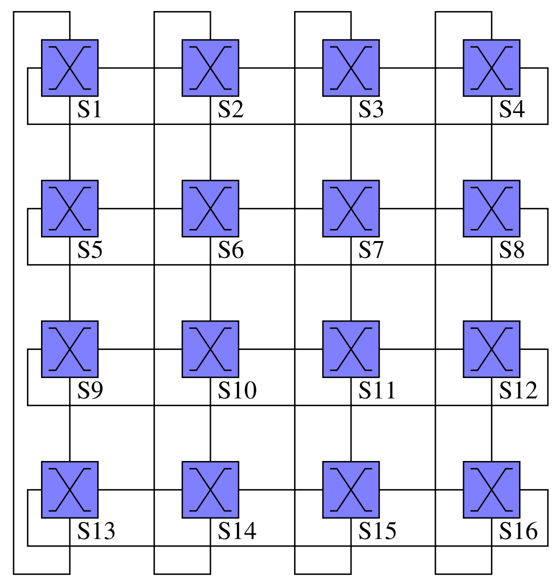

Fig. 3. The 4×4 torus network

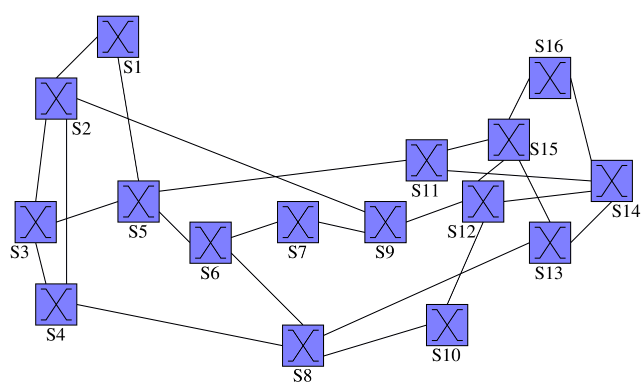

Fig. 4. 16-node topology from the 14-node NSFNet

TABLE I

INTERFERENCE LEVELS $IL(i, j)$ FOR THE 4×4 TORUS NETWORK

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	0	12	6	11	11	9	4	6	4	2	0	3	14	5	4	4
2	12	0	10	8	9	11	8	6	4	10	5	0	6	14	12	1
3	5	8	0	14	3	6	10	9	0	7	10	6	3	8	10	11
4	14	11	12	0	8	7	9	12	3	0	5	10	7	6	9	12
5	9	8	4	6	0	12	6	9	11	3	1	9	8	2	0	1
6	5	14	6	6	14	0	8	8	7	10	3	4	3	7	6	0
7	2	7	14	12	7	14	0	14	2	4	10	10	0	3	5	8
8	8	6	12	14	14	12	14	0	11	2	10	14	6	0	4	12
9	8	2	0	2	9	4	1	5	0	8	3	14	10	4	3	4
10	7	12	5	0	9	7	5	3	14	0	9	7	11	14	10	1
11	0	2	10	5	3	7	14	12	3	14	0	11	0	12	12	9
12	3	0	7	8	5	3	10	11	11	7	12	0	4	2	7	14
13	12	6	4	6	3	1	0	1	8	3	0	7	0	5	5	12
14	11	13	10	7	6	3	3	0	9	13	6	2	11	0	13	1
15	4	6	12	8	0	2	6	3	2	12	10	6	4	13	0	5
16	7	2	8	11	1	0	8	8	3	1	9	14	6	2	7	0

TABLE II

INTERFERENCE LEVELS $IL(i, j)$ FOR THE 16-NODE NSF NETWORK

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	0	13	13	9	3	3	11	14	13	5	3	4	3	1	12	3
2	13	0	13	14	7	7	6	3	2	2	9	3	2	3	11	2
3	13	13	0	6	11	14	5	3	3	13	1	2	1	13	11	1
4	11	12	8	0	14	10	10	4	2	3	14	11	10	6	14	11
5	5	9	8	14	0	14	14	13	1	11	5	5	4	13	13	4
6	5	9	14	7	14	0	8	4	1	14	4	5	4	14	7	4
7	9	7	3	5	14	13	0	14	13	7	3	6	5	4	9	5
8	14	8	6	3	3	2	14	0	14	9	3	11	10	3	9	11
9	13	2	4	2	1	1	13	14	0	14	12	12	11	10	4	12
10	6	2	13	2	13	14	7	7	14	0	4	8	6	10	2	7
11	4	11	1	14	6	4	4	4	12	4	0	12	11	8	13	12
12	5	3	3	10	5	4	5	10	11	11	13	0	10	14	5	14
13	3	1	1	6	3	2	3	6	11	5	12	10	0	8	3	14
14	1	3	13	5	13	14	3	4	3	10	12	13	10	0	2	12
15	14	11	8	14	8	6	9	6	4	2	13	6	5	2	0	5
16	3	1	1	8	3	2	3	9	11	8	12	14	14	11	3	0

decreased. Typically, every switch in the OBS network uses the same algorithm to set the priority of wavelengths, and maintains locally a list of the W wavelengths in decreasing order of priority.

A wavelength assignment scheme based on priorities was presented in [4], and was referred to as "priority wavelength assignment" (PWA). This work assumes a single, fixed path for each source-destination pair (S_i, S_j) which all bursts from switch S_i to S_j follow. Under PWA, each switch S_i in the OBS network maintains locally a priority value for each wavelength-destination pair; in other words, switch S_i assigns a priority to each tuple $(\lambda_w, S_j), w = 1, \cdots, W, i \neq j = 1, \cdots, N$. The priority of tuple (λ_w, S_j) is set to the ratio of the number of bursts which have been successfully transmitted from S_i to S_j on wavelength λ_w (along the fixed path associated with this pair of switches) over the total number of bursts transmitted from S_i to S_j on the same wavelength. When switch S_i needs to transmit a burst to S_j, it considers the wavelengths in decreasing order of priority of the corresponding tuples (λ_w, S_j) and uses the first free one. Depending on the outcome of the transmission, the switch then updates the priority of the tuple. It is shown in [4] that, under low load, PWA performs better than the Random wavelength assignment policy in terms of burst drop probability; under high load, on the other hand, it performs only marginally better than Random.

We now introduce two additional PWA schemes which differ from the the one presented in [4] in two ways. First, a priority is associated with each wavelength in a different way than in [4], resulting in a trade-off between complexity (in both space and time) and performance. Second, our notion of priority, and the manner in which it is incremented and decremented, are different than the one in [4]. Next, we describe the operation of the new PWA schemes, and then we define the priority values and the way they are updated. In our discussion, we will use $p(\bullet)$ to denote the priority function.

The first scheme, which we call "PWA-Link," works as follows. Each switch S_i maintains a priority value for each wavelength-link pair, i.e., for each tuple $(\lambda_w, e), w = 1, \cdots, W, e \in E$, where E is the set of links in the network. Whenever the switch wishes to transmit a burst to some switch

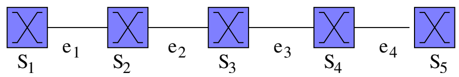

Fig. 5. A linear network to explain the difference between PWA, PWA-Link

S_j over path $\pi = \{e_1, e_2, \cdots, e_k\}$, it computes the wavelength-path priorities $p(\lambda_w, \pi)$ by adding up the corresponding wavelength-link priorities along the path links: $p(\lambda_w, \pi) = \sum_{e \in \pi} p(\lambda_w, e), \; w = 1, \cdots, W$. The switch considers the wavelengths in decreasing order of $p(\lambda_w, \pi)$, and transmits the burst on the first free wavelength, with ties broken arbitrarily. Upon learning the outcome of the transmission, the switch: *increments* the priority of the links, if any, on which the burst was successfully transmitted; *decrements* the priority of the link, if any, at which it was dropped due to contention; and *maintains* the priority of any other links (e.g., links following the one where the burst was dropped). We will explain shortly how the priorities are incremented or decremented.

PWA-Link operates at finer granularity and uses more information than PWA in making wavelength assignments, therefore one might expect that it would lead to better performance; indeed, numerical results to be presented in the next section confirm this observation. To explain the difference in performance, let us consider the simple linear network shown in Figure 5, and suppose that switch S_1 transmits a burst to switch S_5 on some wavelength λ_w. Suppose further that the burst is dropped at switch S_4. Under PWA, the priority of the tuple (λ_w, S_5) is decremented, without taking into account the fact that the burst transmission was successful on the first three links of the path from S_1 to S_5; indeed, $p(\lambda_w, S_5)$ is decremented by the same amount regardless of which switch in the path dropped the burst. In PWA-Link, on the other hand, this additional information is used in updating the priorities of the wavelength-link tuples. Since the burst was successful on links e_1, e_2, and e_3, the priorities $p(\lambda_w, e_1)$, $p(\lambda_w, e_2)$, and $p(\lambda_w, e_3)$ are incremented, while $p(\lambda_w, e_4)$ is decremented. Note that by increasing the priorities of λ_w on the first three links, this wavelength will move up the list with respect to burst transmissions to switches S_2, S_3, and S_4, as it should, since the burst reached all three switches successfully.

The second scheme we propose is simpler than both PWA and PWA-Link, and we will refer to it as "PWA-λ." With this scheme, each switch S_i assigns a priority value $p(\lambda_w)$ to every wavelength $\lambda_w, w = 1, \cdots, W$. When switch S_i successfully transmits a burst on wavelength λ_w, the priority $p(\lambda_w)$ is incremented regardless of the destination of the burst or the path traveled. Otherwise, the priority of the wavelength is decremented. Intuitively, PWA-λ will perform worse than either PWA or PWA-Link in terms of burst drop probability, but it is simpler and easier to implement.

Let us now consider the space and time complexity of implementing the three PWA schemes at each switch. PWA-λ requires $O(W)$ memory to record the priority information, where W is the number of wavelengths. It also needs $O(1)$ time to update the priority value of a wavelength once the relevant feedback from the network has been received, and $O(\log W)$ time to maintain a sorted priority list. PWA requires $O(WN)$ memory for recording priority values, where N is the number of switches in the network; and it takes constant time to update the priority of a wavelength-destination pair. It also takes $O(\log W)$ time to maintain a sorted priority list, since it only needs to have one such list of W elements (wavelengths) for each of the N destinations. Finally, PWA-Link needs $O(W \mid E \mid)$ memory for the priority values, where E is the set of links in the network. When the feedback regarding a burst transmission is received, the switch must update the priority of all wavelength-link pairs along the path, and this operation takes time $O(\Delta)$, where Δ denotes the diameter of the network. The computational overhead for maintaining a sorted list per destination is $O(k \log W)$, where k is the number of paths overlapping with the path to this destination. As we can see, the three PWA schemes represent a tradeoff between implementation complexity and performance, with PWA-Link being the best performing but most complex, PWA-λ the worst performing but easiest to implement, and pure PWA occupying the middle ground in both metrics.

We now turn our attention to the priority function and the increment and decrement operations used to update the wavelength priorities. The priority of a wavelength-destination (λ_w, S_j) pair in [4] was defined as the fraction of transmissions to destination S_j on wavelength λ_w that have been successful. However, our experimental investigations indicate that this measure may not be appropriate because of disparities in the rates of change in priority over time and across wavelengths. Specifically, while initially the rate of change is relatively large, the rate of change diminishes over time: once the number of bursts transmitted on a wavelength becomes relatively large, each additional transmission has a negligible effect on the priority, regardless of the outcome. As a result, once the network has been in operation for a while and the priorities have settled, it will take a long time for priorities to adapt to any changes in the traffic or network dynamics, during which bursts will use suboptimal wavelengths and the burst drop probability will be high. Furthermore, the rate of change in priority can be different for different wavelengths, possibly substantially so. For instance, consider two wavelengths that have the same priority but one has been used substantially more often that the other for transmitting bursts. In this case, at each step (i.e., burst transmission), the priority of the wavelength that has been used more frequently will change by a small amount in either direction, while the the priority of the less frequently used wavelength will change by a larger amount. Given that both these properties are undesirable, it would be preferable to use a priority scheme in which the rate of change at each update is not affected by length of time or frequency of use of a wavelength.

In our work, the priorities $p(\bullet)$ are taken to be real numbers in the range $[1, W]$, and are initialized to $W/2$. We use an "Additive Increase, Additive Decrease" (AIAD) scheme to update the priorities, with increment Inc and decrement Dec. Specifically, after the result of a burst transmission has become known, a switch takes the following steps.

- If the burst transmission was successful, the appropriate priority (or priorities, in the case of PWA-Link) are incremented as follows: $p(\bullet) \leftarrow \max\{p(\bullet) + Inc,\ W\}$.
- Otherwise, the appropriate priorities are decremented as: $p(\bullet) \leftarrow \min\{p(\bullet) - Dec,\ 1\}$.

We have conducted a large number of experiments to determine the best values of Inc and Dec to use with the AIAD scheme. Our results indicate that the performance of the PWA policies is best when $Inc < Dec$, and Inc takes values from $0.2 \sim 0.4$, while the value of Dec is in the range $0.8 \sim 1.2$.

A. Combining PWA and Traffic Engineering

We now present a small modification to the PWA schemes to incorporate the traffic engineering approach of Section III-B. The modification is applied at initialization time only, while the operation of the PWA schemes remains identical to the one described above. Recall that, in the original PWA schemes, all priorities are initialized to $W/2$. Therefore, initially all wavelengths are indistinguishable from each other with respect to transmission preference. The modification is to use different initial priorities at each switch, so that different switches will be forced to use different wavelengths for interfering traffic from the very beginning. If the initial values are determined appropriately, this approach has the following benefits over pure PWA: (1) the initial burst losses will be avoided; (2) the switches will settle to preferred wavelengths faster; and (3) the network will achieve better overall traffic isolation.

Similar to the First-Fit-TE wavelength assignment policy, we arbitrarily order the W wavelengths as $\lambda_1, \cdots, \lambda_W$, and we assign start wavelengths to the switches as we described in Section III-B. Consider some switch S_i, and let $start(i)$ be its start wavelength. Let also $next$ denote the next wavelength (modulo-W) that is assigned as the start wavelength of another switch; in other words, the wavelengths $\lambda_{start(i)\oplus 1}, \cdots, \lambda_{next\ominus 1}$ are not assigned as start wavelengths for any switch ($\oplus$ and $\ominus$ denote addition and subtraction, respectively, modulo-W). Then, at switch S_i all priorities involving wavelengths $\lambda_{start(i)}, \cdots, \lambda_{next\ominus 1}$ are initialized to $W/2 + Inc$, while the priorities of all other wavelengths are initialized to $W/2$, as before. As a result, the switch will initially give preference to wavelengths $\lambda_{start(i)}, \cdots, \lambda_{next\ominus 1}$ when transmitting bursts. The operation of the PWA schemes is not affected in any other way.

We will use the terms "PWA-TE," "PWA-Path-TE," and "PWA-λ-TE to refer to the versions of PWA, PWA-Path, and PWA-λ, respectively, in which wavelength priorities are initialized in the manner described above.

V. Numerical Results

We consider two 16-node network topologies, the 4×4 torus network shown in Figure 3, and the NSF network in Figure 4. Our goal is to compare via simulation the various wavelength assignment schemes in terms of (1) overall (network-wide) burst drop probability, and (2) burst drop probability as a function of path length. Since, in an OBS network without wavelength converters, the drop probability may increase with the number of hops a burst has to traverse, it is important to achieve some degree of fairness among bursts that travel over paths of different length. In our simulations, the burst arrival process of each switch is Poisson and the burst length is exponentially distributed with mean $1/\mu$. For simplicity, we also assume that bursts originating at a given switch are equally likely to be destined to any of the other switches. We used the method of batch means to estimate the burst drop probability; each of the simulation runs lasts until 400,000 bursts have been transmitted by the whole 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.

Figures 6 and 7 plot the burst drop probability of the wavelength assignment schemes we described in Sections III and IV for the 4 torus network and for $W = 64$ wavelengths. Figure 6 shows the drop probability for low traffic load, while Figure 7 shows the performance of the network under moderate and high loads. From the figures, we see that First-Fit is always worse than Random, confirming our previous analysis. We also make two important observations. First, the adaptive, priority-based schemes (i.e., PWA, PWA-Link, and PWA-λ) perform better than Random (with the exception of PWA-λ at low loads less than 0.1). In particular, PWA-Link, which uses more detailed information than the original PWA, is the best of the three adaptive schemes, PWA is the second best, while PWA-λ, which uses the least amount of information, is the worst of the three. The second observation is that the traffic engineering approach we described earlier to achieve traffic isolation and reduce traffic interference, when combined with any wavelength assignment scheme, static or adaptive, leads to a significant decrease in burst drop probability. The most dramatic impact is with the First-Fit scheme, in which case First-Fit-TE has a burst drop probability that is up to two orders of magnitude lower than the plain First-Fit policy. Similar decreases (although of smaller magnitude) can be observed for the PWA-TE, PWA-Link-TE, and PWA-λ-TE schemes over the respective non-TE versions. Overall, we find that the best approach to wavelength assignment in OBS networks is to combine adaptive, priority-based schemes with our traffic engineering approach. Interestingly, we find that PWA-TE is the best performing scheme, having lower burst drop probability than even the PWA-Link-TE scheme over a wide range of load values (note that, in contrast, PWA-Link performs much better than PWA). We believe that this result is due to the fact that, in PWA-TE, the wavelength priorities are adjusted by considering the whole path of a burst, not individual links as in PWA-Link-TE, and this operation is more compatible with the traffic engineering approach we take.

Figures 8 and 9 are similar to the previous two figures, but compare the burst drop probability of the nine wavelength assignment schemes for the NSFNet. The burst drop probability is higher than in the torus network for a given load, since the NSFNet topology is (1) more sparsely connected than the torus network, and (2) irregular, and thus, without the inherent load balancing properties of the torus topology. As a result, certain

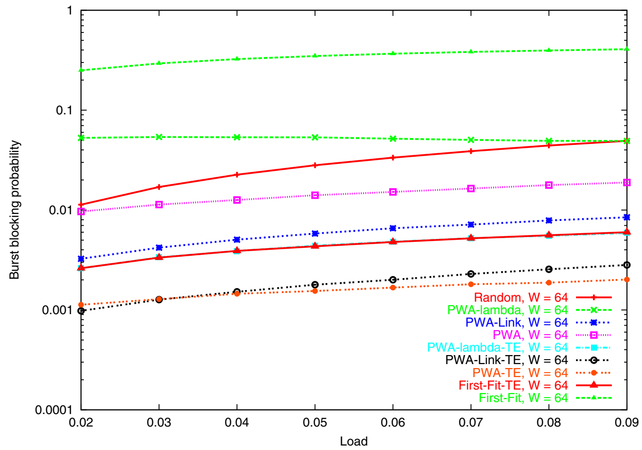

Fig. 6. Burst drop probability, 4×4 torus network, low load

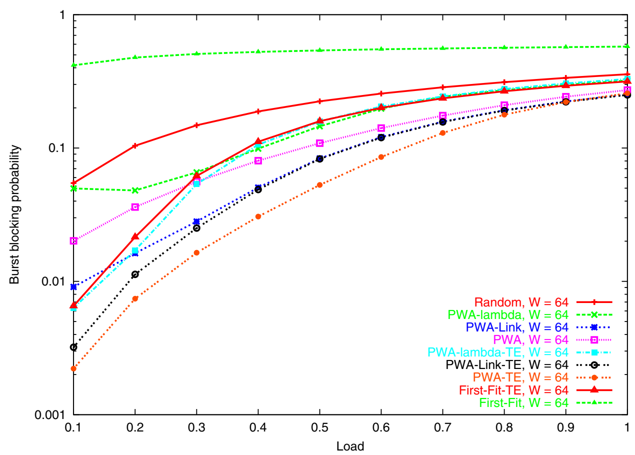

Fig. 7. Burst drop probability 4×4 torus network, moderate and high load

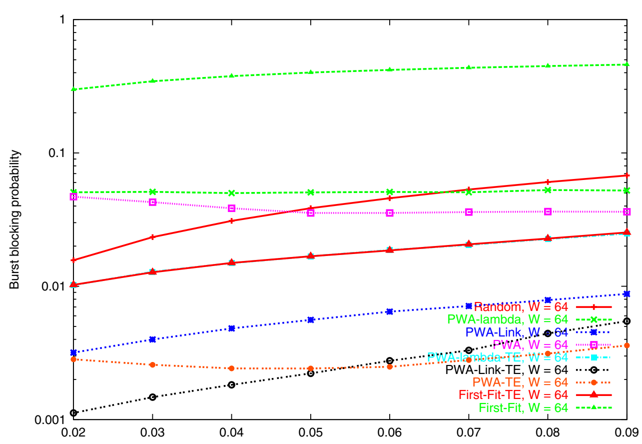

Fig. 8. Burst drop probability, NSFNet, low load

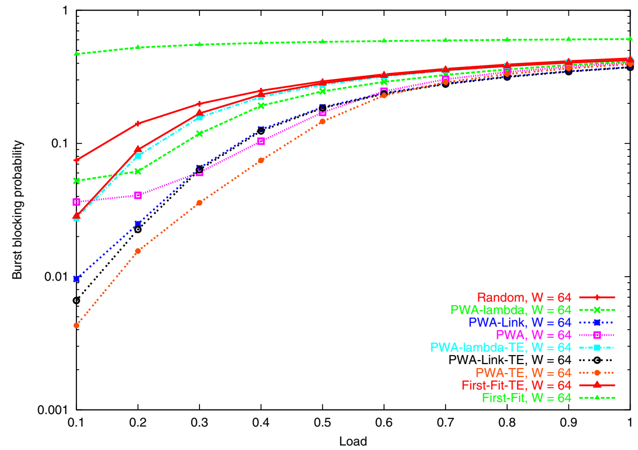

Fig. 9. Burst drop probability, NSFNet, moderate and high load

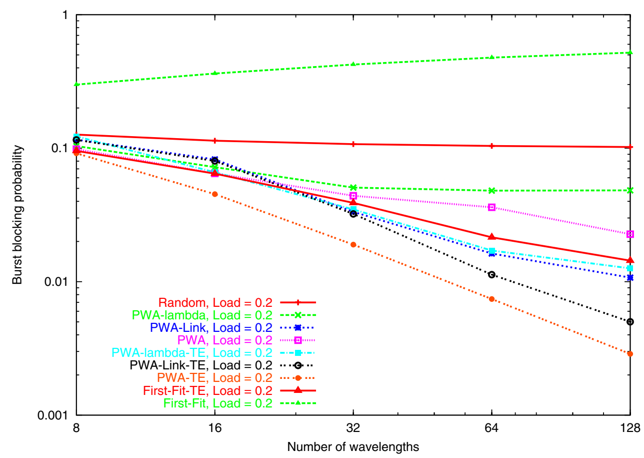

Fig. 10. Burst drop probability, 4×4 torus network, load = 0.2

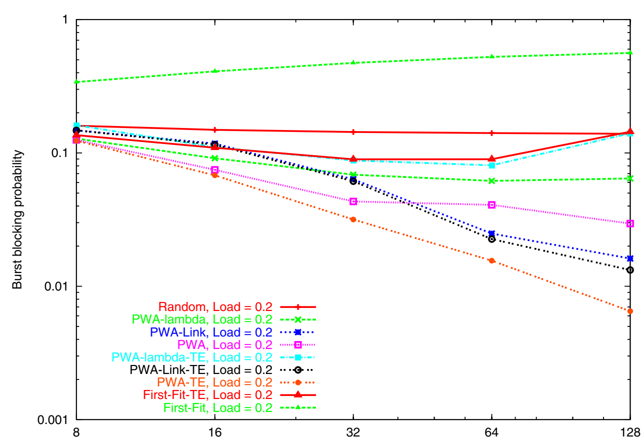

Fig. 11. Burst drop probability, NSFNet, load = 0.2

links may become heavily congested when using shortest path routing, leading to higher burst drop probability. The relative performance of the wavelength assignment schemes is very similar to the one we observed for the torus network: adaptive, priority-based schemes are better than static ones, and incorporating traffic isolation through traffic engineering leads to a decrease in drop probability. As before, PWA-TE is the best policy overall, except at very low loads. We also note that, at very high loads, the performance of all policies is similar; this is due to the fact that, at such high loads, burst dropping is mostly due to the lack of wavelengths.

Now let us consider the gain as we increase the number W of available wavelengths. Figures 10 and 11 plot the burst drop probability of the nine schemes as the number of wavelengths increases from 8 to 128, for the torus and NSFNet, respectively. The load per wavelength in the network is kept constant at 0.2 for these experiments. As we can see, the burst drop probability of First-Fit increases, and it remains mostly unchanged in the case of Random. These results are expected: Random distributes the bursts randomly to the various wavelengths, but since the load per wavelength is constant, there is little change in overall drop probability; while First-Fit attempts to use the same few first wavelengths, thus an increase in overall load, as W increases, results in higher drop probability. For the other schemes, in general, the drop probability decreases with the number of wavelengths, up to a point. Finally, we note that PWA-TE and PWA-Link-TE are the two schemes that show a consistent drop in burst drop probability for the range of wavelengths considered here. Since these are the best performing schemes overall, this result indicates that a combination of adaptive policies with traffic engineering is the best approach to take advantage of wavelength resources in the OBS network.

Figures 12 and 13 plot the burst drop probability as a function of the number of hops in a burst's path, for the torus and NSFNet topologies, respectively. As expected, the burst drop probability increases with the length of the path. However, while for some schemes (e.g., Random and First-Fit) there can be a difference of two orders of magnitude between the drop probability of bursts traversing one hop versus bursts that travel four hops, the difference is less acute when schemes employing adaptive policies with traffic engineering are used. Therefore, our approach not only improves the overall burst drop probability, it also increases fairness. Again, PWA-TE is the best performing scheme in terms of fairness.

Overall, we find that adaptive policies perform better than non-adaptive ones; and that applying traffic engineering to achieve traffic isolation can further improve the performance. PWA-TE has performed the best over all the experiments we have conducted, with PWA-Link-TE a close second. Since PWA-TE is relatively easier to implement and involves fewer computations and memory requirements, it is the best choice for OBS networks with no wavelength conversion capabilities.

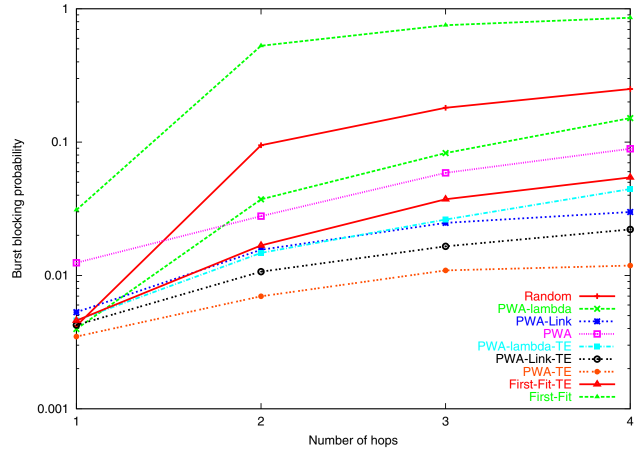

Fig. 12. Burst drop probability, 4×4 torus network, load = 0.2, W = 64

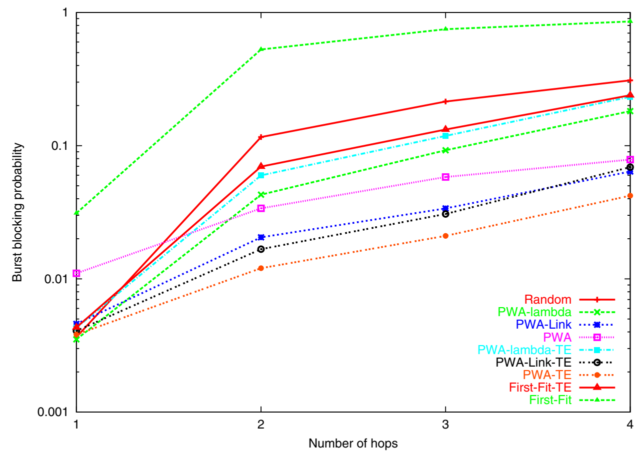

Fig. 13. Burst drop probability, NSF network, load = 0.2, W = 64

VI. Concluding Remarks

We studied the wavelength assignment problem in OBS networks, and proposed a suite of policies to achieve low burst drop probability and to attain fairness among bursts with different path lengths. We are currently investigating similar traffic engineering concepts to further reduce traffic interference in an OBS network, in particular by designing appropriate routing algorithms.

References

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

[2] J. Li and C. Qiao. Schedule burst proactively for optical burst switching networks. *Computer Networks*, 44(5):617-629, April 2004.

[3] J. Teng and G. N. Rouskas. A comparison of the JIT, JET, and Horizon wavelength reservation schemes on a single OBS node. In *Proc. of the First International Workshop on Optical Burst Switching*, October 2003.

[4] X. Wang, H. Morikawa, and T. Aoyama. Priority-based wavelength assignment algorithm for burst switched WDM optical networks. *IEICE Transactions on Communications*, E86-B(5):1508–1514, May 2003.

[5] Y. Zhu, G. N. Rouskas, and H. G. Perros. A comparison of allocation policies in wavelength routing networks. *Photonic Network Communications*, 2(3):265–293, August 2000.

Proceedings of the First International Conference on Broadband Networks (BROADNETS'04)
0-7695-2221-1/04 $ 20.00 IEEE