Path Switching in OBS Networks^{*}

Li Yang and George N. Rouskas

Department of Computer Science,
North Carolina State University,
Raleigh, NC, USA
{lyang, rouskas}@eos.ncsu.edu

Abstract. We investigate the concept of path switching in optical burst switched (OBS) networks and its potential to reducing the overall burst drop probability. With path switching, each source maintains a list of alternate paths to each destination, and uses information regarding the recent congestion status of the network links to rank the paths; it then transmits bursts along the least congested path. We present a suite of path switching strategies, each utilizing a different type of information regarding the link congestion status, and evaluate them using simulation. Our results demonstrate that, in general, path switching outperforms shortest path routing, and, depending on the path strategy involved, the network topology, and the traffic pattern, this improvement can be significant. We also present a new framework for the development of hybrid path switching strategies, which make routing decisions based on a weighted combination of the decisions taken by several independent path switching strategies. We present two instances of such hybrid strategies, one that assigns static weights and one that dynamically adjusts the weights based on feedback from the network.

1 Introduction

Optical burst switching (OBS) is a promising switching paradigm which aspires to provide a flexible infrastructure for carrying future Internet traffic in an effective yet practical manner. OBS transport is 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, whereas optical packet switching requires practical, cost-effective, and scalable implementations of optical buffering and optical header processing, which are several years away. OBS does not require buffering or packet-level parsing in the data path, and it is more efficient than circuit switching when the sustained traffic volume does not consume a full wavelength. The transmission of each burst is preceded by the transmission of a setup message 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

^{*} This work was supported in part by MCNC-RDI as part of the Jumpstart project.

has been set-up; instead it starts transmitting a data burst after a delay (referred to as “offset”), following the setup message. For a recent overview of the breadth and depth of current OBS research, the reader is referred to [1].

One of the most important issues in OBS networks is that of burst loss due to congestion, and appropriate mechanisms must be in place to manage the increased demand for resources during a period of congestion. Such mechanisms can be implemented either inside the network (at core switches) or at the edge. Contention-resolution schemes at the core of the network can be based on one of four orthogonal approaches, or a combination thereof: buffering, wavelength conversion, burst segmentation, or deflection [2, 3]. All these approaches require additional hardware or software components at each OBS switch, increasing their cost significantly; furthermore, practical implementations of some of these components require technology which may be several years from maturity.

At the network edge, one strategy that has the potential to improve burst contention, especially when wavelength conversion is unavailable or sparse, is wavelength assignment [4]. A traffic engineering approach to select paths for source routing so as to balance the traffic load across the network links was investigated in [5]. Finally, a dynamic scheme for selecting routes at the burst sources was proposed in [6]. We note that a similar technique, referred to as “end-to-end path switching” was proposed and evaluated recently for selecting one among a set of Internet paths [7]; the main finding of the study was that path switching can result in substantial improvement in packet loss.

In this paper, we undertake a comprehensive study of (end-to-end) path switching in OBS networks. The remainder of the paper is organized as follows. In Section 2 we discuss our assumptions regarding the OBS network we consider in our study. In Section 3 we describe path switching strategies each utilizing partial information about the network state to select one of a set of available paths to route bursts. In Section 4 we develop a framework for combining several path switching strategies into hybrid (or meta-) strategies which base their routing decisions on the decisions of multiple individual methods. In Section 5, we present simulation results to demonstrate the effectiveness and benefits of path switching, and we conclude the paper in Section 6.

2 The OBS Network Under Study

We will use $G = (V, E)$ to denote an OBS network. V is the set of switches, $N = |V|$, and $E = \{\ell_1, \ell_2, \dots, \ell_M\}$ is the set of unidirectional fiber links, $M = |E|$. Each link in the network can carry burst traffic on any wavelength from a fixed set of W wavelengths, $\{\lambda_1, \lambda_2, \dots, \lambda_W\}$. We assume that each OBS switch in the network has full wavelength conversion capabilities which are used in the case of wavelength contention. The network does not use any other contention resolution mechanism. Specifically, OBS switches do not employ any buffering, either electronic or optical, in the data path, and they do not utilize deflection routing or burst segmentation. Therefore, if a burst requires an output port at a

time when all wavelengths of that port are busy transmitting other bursts, then the burst is dropped.

The OBS network employs source routing, in that the ingress switch determines the path of a burst entering the network. We assume the existence of a routing algorithm that is capable of computing a set of k , $k = 2 - 4$, alternate paths for each source-destination pair. Each source node maintains the list of paths for each possible destination, and is responsible for selecting the path over which a given burst will travel. Once the source has made a routing decision for a burst, the path is recorded in the setup message and it cannot be modified by downstream nodes (i.e., no deflection is allowed).

All source nodes use the same *path switching strategy* to make routing decisions on a per-burst basis. A path switching strategy is characterized by the metric used to rank the paths to a certain destination node. In general, the metric is designed to reflect the likelihood that a burst transmitted on a particular path will experience resource contention and be dropped before it reaches its destination. Whenever a new burst is ready for transmission, the source node selects the “best” path according to the metric used (with ties broken arbitrarily) and injects the burst into the network. We present a number of path switching strategies, and their associated metrics, in the next section.

The rank of each path maintained at a source node is updated dynamically based on information regarding the state of the network collected by the node. We assume that the control plane of the OBS network provides support for the collection and dissemination of information required by the path switching strategies. For instance, this information may be part of the feedback the source receives from the signaling protocol regarding the success or failure of each burst transmission [8]. Alternatively, the OBS switches may collect information and statistics regarding the (long-term) congestion status of their links, and use a link-state protocol to disseminate this information to the rest of the network. Since signaling and state dissemination protocols are required for a variety of network functions, the additional overhead due to the path switching strategies we propose in this paper is expected to be only moderate.

3 Pure Path Switching Strategies

A path switching strategy uses information about the current state of the OBS network to select one of a small number of routing paths for transmitting burst traffic between a source-destination pair. There are several different pieces of information that could be used to describe the congestion level in the network (for instance, link utilization, end-to-end path burst drop rate, etc.); and there are several ways in which this information can be combined into a metric to rank paths. It is unknown which types of information or what metrics perform best for path switching in terms of burst drop probability. In this section we present a suite of *pure* path switching strategies, i.e., strategies which use a single path selection method.

3.1 Weighted Bottleneck Link Utilization (WBLU) Strategy

The WBLU strategy ranks the candidate paths using information on link utilization. The motivation behind this strategy is to reduce or prevent contention by using paths with less utilized links. Consider a (directional) link ℓ of the OBS network, let $Succ(\ell, t)$ denote the set of bursts that have successfully traversed link ℓ until time t , and let T_i denote the length of burst i . The utilization $U(\ell, t)$ of link ℓ at time t is defined as:

$$U(\ell, t) = \frac{\sum_{i \in Succ(\ell, t)} T_i}{Wt} \quad (1)$$

where W is the number of available wavelengths; at time $t = 0$, we assume that the utilization $U(\ell, 0) = 0$ for all links ℓ .

Consider now a source-destination pair (s, d) , and let $\{\pi_z, z = 1, \dots, m\}$ be the set of m candidate paths for transmitting bursts from node s to node d . Let $\{\ell_k, k = 1, \dots, |\pi_z|\}$ be the set of links composing path π_z which has length (in number of hops) $|\pi_z|$. At time t , the WBLU strategy routes bursts from s to d along the path $\pi_{z^*}(t)$ whose index $z^*(t)$ is obtained using the following metric:

$$z^*(t) = \arg \max_{1 \leq z \leq m} \frac{1 - \max_{1 \leq k \leq |\pi_z|} U(\ell_k, t)}{|\pi_z|} \quad (2)$$

The numerator in the above expression is the available capacity of the bottleneck link in a given path π_z . Therefore, the WBLU strategy routes bursts along the path with the highest ratio of available bottleneck link capacity to path length. By taking the number of hops into account as in expression (2), we ensure that if the bottleneck link utilization is similar for two paths, then the shortest path is selected for routing; the longer path is preferred only if the utilization of its bottleneck link is significantly lower than that of the shorter one. We note that a similar metric for ranking paths was used in [9] as part of a routing and wavelength assignment algorithm for wavelength routed networks.

3.2 Weighted Link Congestion (WLC) Strategy

The objective of the WLC strategy is to route bursts along the path that is most likely to lead to a successful transmission. To this end, the source uses information on link congestion along each path to infer the burst drop rate of the path. This strategy assumes the existence of a link-state protocol that disseminates information on link congestion.

Let $N_{succ}(\ell, t)$ (respectively, $N_{drop}(\ell, t)$) denote the number of bursts that have been successfully transmitted along (respectively, dropped at) link ℓ up to time t . We define the congestion level $c(\ell)$ of link ℓ at time t as the fraction of bursts that have been dropped at the link:

$$c(\ell, t) = \frac{N_{drop}(\ell, t)}{N_{drop}(\ell, t) + N_{succ}(\ell, t)} \quad (3)$$

We assume that at time $t = 0$, the congestion $c(\ell, 0) = 0$ for all links ℓ .

Let π_z be a candidate path for routing bursts between a source-destination pair (s, d) , consisting of links $\ell_1, \dots, \ell_{|\pi_z|}$. Assuming that link drop probabilities are independent, at time t the probability that a burst will be dropped along this path can be calculated as:

$$b(\pi_z, t) = 1 - \prod_{1 \leq i \leq |\pi_z|} (1 - c(\ell_i, t)) \quad (4)$$

The weighted link congestion (WLC) strategy routes bursts from s to d along the path $\pi_{z^*(t)}$ whose index $z^*(t)$ is obtained using the following metric:

$$z^*(t) = \arg \max_{1 \leq z \leq m} \frac{1 - b(\pi_z, t)}{|\pi_z|} \quad (5)$$

As in expression (2), this metric takes the number of hops of each path into account, in order to ensure that longer paths are preferred over shorter ones only when they offer a substantial improvement in drop probability.

3.3 End-to-End Path Priority-Based (EPP) Strategy

The EPP strategy is similar in spirit to the WLC strategy in that it also attempts to route bursts along paths with low drop probability. However, rather than relying on information on individual link congestion levels to infer the burst drop probability, this strategy requires the source to directly measure this probability from feedback messages it receives from the network regarding the status of each burst transmission.

Consider the source-destination pair (s, d) , and let π_z be one of the m candidate paths for this pair as before. Let $N_z(t)$ denote the total number of bursts that have been transmitted (successfully or unsuccessfully) from s to d on path π_z up to time t . The EPP strategy assigns a priority $prio(\pi_z, t)$ to path π_z at time t which is updated each time a new burst is transmitted on this path, and is recursively defined as:

$$prio(\pi_z, t) = \begin{cases} 1.0, & t = 0 \\ \frac{prio(\pi_z, t-1) \times N_z(t-1) + 1}{N_z(t-1) + 1}, & \text{burst transm. success at time } t \\ \frac{prio(\pi_z, t-1) \times N_z(t-1)}{N_z(t-1) + 1}, & \text{burst transm. failed at time } t \end{cases} \quad (6)$$

$N_z(t)$ is also updated as: $N_z(t) = N_z(t-1) + 1$ each time a new burst is transmitted on path π_z , with $N(0) = 0$. In the above expressions, the time index t refers to the time the source receives feedback from the network regarding the outcome (success or failure) of the most recent burst transmission along path π_z ; similarly, index $t-1$ refers to the time feedback was received regarding the immediately previous burst transmission over the same path. The priority of a path remains unchanged in the interval $[t-1, t)$. The priority of a path in expression (6) is simply the fraction of bursts that have been successfully transmitted along this path up to time t ; hence, the range of path priorities is in $(0, 1)$.

At time t , the EPP strategy routes bursts from s to d along the path $\pi_{z^*}(t)$ whose index $z^*(t)$ is obtained using the following metric:

$$z^* = \begin{cases} z, & \text{prio}(\pi_z, t) - \text{prio}(\pi_x, t) > \Delta \quad \forall x \neq z \\ \arg \max_{1 \leq z \leq m} \frac{\text{prio}(\pi_z, t)}{|\pi_z|}, & \text{otherwise} \end{cases} \quad (7)$$

The threshold Δ reflects the degree of confidence in the selection of a given path for routing paths. If we are sufficiently confident that a path is better than others in terms of burst drop probability, then the selection is based solely on path priorities. Otherwise, we discount the priority of each path by its length, and we select a path based on the discounted priorities.

4 Hybrid Path Switching Strategies

Each pure path switching strategy uses only one piece of information in reaching a decision, and this information provides only a limited “view” of the network state. In this section, we focus on hybrid strategies which, at each burst transmission instant, combine the decisions of several pure strategies into an overall decision in the hope of improving the accuracy of the path selection process and improve the overall burst drop probability. In general, a hybrid strategy emulates a set of pure strategies which run independently of each other “on the side.” The motivation for this approach is to combine the different partial “views” of the network state in a way that improves the performance.

The principles underlying the hybrid path switching strategies are based on ideas from the domain of machine learning [10, 11]. Specifically, it has been shown [10] that the *ensemble* decision reached by a set of voters is more accurate than the decision of any individual voter, provided that each voter reaches a decision in a manner that is largely independent of other voters. In the context of path switching in an OBS network, a pure path switching strategy corresponds to a voter, and the selection of a path corresponds to a (routing) decision. A strategy is “correct” if transmitting the burst over the path selected by the strategy is successful, and it is “wrong” if the burst is dropped along the path before it reaches its destination. We can think of the overall burst drop probability of a strategy as its “error rate,” i.e., the fraction of time the method is incorrect in successfully selecting a path for a burst. Obviously, the drop probability *overestimates* the real error rate of the strategy, since the fact that a burst is dropped along a given path does not necessarily imply that the burst would have been successful had another path been chosen. We expect that making routing decisions by considering several different views simultaneously will lead to better performance in terms of burst drop probability.

In the remainder of this section, we consider a single pair (s, d) ; our observations apply similarly to all other source-destination pairs. The source s maintains $m > 1$ candidate paths, $\pi_1, \dots, \pi_m$, for routing bursts to d . For ease of presentation we will drop any references to the pair (s, d) .

Let us assume that there are n pure path switching strategies available, $S_1, S_2, \dots, S_n$. A strategy S_i takes as input some information regarding the net-

work state, and produces a probability distribution $p_i^{(z)}$ over the indices of the candidate paths. The probability $p_i^{(z)}$, $z = 1, \dots, m$, represents the degree of confidence that strategy S_i has in selecting candidate path π_z for routing the burst traffic. Obviously, we have: $p_i^{(1)} + p_i^{(2)} + \dots + p_i^{(m)} = 1$.

A hybrid strategy H assigns a probability distribution q_i over the n pure path switching strategies $S_1, \dots, S_n$. The probability q_i represents the degree of confidence of the hybrid strategy H that strategy S_i is correct in its selection of a path. Again, we have that: $q_1 + q_2 + \dots + q_n = 1$. Then, the expected confidence of the hybrid strategy in selecting candidate path π_z is:

$$E_z = \sum_{i=1}^n q_i p_i^{(z)} \quad z = 1, \dots, m \quad (8)$$

Therefore, the decision of the hybrid strategy H is to route bursts along the path π_{z^*} with the maximum expected confidence, i.e., the one with index z^* :

$$z^* = \arg \max_{1 \leq z \leq m} E_z \quad (9)$$

4.1 Majority Binary Voting (MBV) Strategy

Majority binary voting (MBV) is the simplest hybrid strategy. Let us assume that there are n pure strategies available, $S_1, \dots, S_n$, where n is odd. Each strategy S_i makes a binary decision for each of the m candidate paths: whether to select it for routing bursts or not. Formally, the probability distribution $p_i^{(z)}$ returned by each strategy S_i is as follows:

$$p_i^{(z)} = \begin{cases} 1, & S_i \text{ selects path } \pi_z \\ 0, & \text{otherwise} \end{cases} \quad i = 1, \dots, n, \quad z = 1, \dots, m \quad (10)$$

The path selected by the hybrid MBV strategy is the one with the most number of votes. This strategy assumes a uniform distribution q_i over the set $\{S_i\}$.

4.2 Weighted Non-binary Voting (WNV) Strategy

MBV restricts the pure path switching strategies to vote for a single path. Non-binary voting allows each pure strategy S_i to assign a degree of confidence to each candidate path π_z through a probability distribution $p_i^{(z)}$. One approach to obtaining the probability distribution is to normalize the values $v_i^{(z)}$ (e.g., priority, congestion level, etc.) assigned to the various paths by strategy S_i :

$$p_i^{(z)} = \frac{v_i^{(z)}}{\sum_{l=1, \dots, m} v_i^{(l)}} \quad i = 1, \dots, n, \quad z = 1, \dots, m \quad (11)$$

The weighted non-binary voting (WNV) strategy further assigns a probability distribution q_i over the set of pure strategies $\{S_i\}$, and reaches a decision

using expressions (8) and (9). The main motivation for using a non-uniform distribution q_i is the fact that each pure strategy results in a different burst drop probability; furthermore, the relative performance of the various pure strategies depends on system parameters such as the network topology, the traffic load and pattern, etc. In general, the performance of the hybrid strategy depends strongly on the choice of weights, with the best performance achieved when the weights reflect the relative error rate of the pure strategies.

4.3 Dynamic Weighted Non-binary Voting (DWNV) Strategy

Under WNV, the probability distribution q_i over the set of pure strategies $\{S_i\}$ remains fixed at all times. One problem with such an approach is the difficulty in appropriately selecting the weights q_i . Instead, it would be desirable to have a method for dynamically adjusting the probability distribution q_i in real time in a way that minimizes the overall burst drop probability; in this case, the probability distribution q_i would also converge to the optimal one. The dynamic weighted non-binary voting (DWNV) strategy is designed to achieve this objective.

Let $q(t) = (q_1(t), \dots, q_n(t))$ be the probability distribution at time t , and let $B(t, q(t))$ be the burst drop probability of the hybrid strategy at time t when the current distribution is $q(t)$. Our objective is to obtain the distribution $q(t+1)$ at time $t+1$ such that the burst drop probability is minimized (the time indices refer to the times a burst is ready to be transmitted). In other words, we need to select the distribution $q^*(t+1)$ such that: $q^*(t+1) = \arg \min_{q(t+1)} B(t+1, q(t+1))$. However, it is not possible to solve the above optimization problem directly. We therefore employ a heuristic to dynamically update the q -distribution. We assume that the confidence $c_i(t)$ in the decision of a strategy S_i is reversely proportional to its burst drop probability $b_i(t)$ at time t :

$$c_i(t) = \frac{1}{b_i(t) + \epsilon} \quad i = 1, \dots, n \quad (12)$$

where ϵ is a smoothing value to avoid division by zero when $b_i = 0$. Based on the confidence c_i of choosing strategy S_i , we compute the new weight q_i as:

$$q_i(t+1) = \frac{c_i(t)}{\sum_{l=1, \dots, n} c_l(t)} \quad i = 1, \dots, n \quad (13)$$

The computation of each of the expressions (13) warrants further discussion. The overall burst drop probability $B(t, q(t))$ of the hybrid policy is calculated at the source node using feedback from the network. However, it is not possible for the source node to calculate directly the burst drop probability $b_i(t)$ of each pure strategy S_i as required by (13). To see why a direct calculation of $b_i(t)$ is not possible, consider what happens if the hybrid strategy adopts a decision that is different than the decision of some pure strategy S_i . In this case, the feedback received by the source provides information regarding the decision made by the hybrid policy but no information regarding the decision made by S_i ; hence, the source has no way of knowing with certainty whether the burst transmission would have been successful had it used the path selected by S_i instead.

To overcome this difficulty, we use the following approach to compute the burst drop probability $b_i(t)$ for a pure strategy S_i whose decision at time t does not coincide with the decision of the hybrid strategy. Let π be the path chosen by S_i , and let $prio(\pi, t)$ be the priority of (burst drop probability along) this path; this priority is computed in the course of the operation of the hybrid policy as in expression (6). Then, we use $prio(\pi_t)$ to update the drop probability of strategy S_i , by making the approximation that the outcome of routing a burst over path π at time t will be failure with probability $1 - prio(\pi, t)$ and success with probability $prio(\pi, t)$. Therefore, the burst drop probability for any pure strategy S_i whose decision at time t is to use path π , is updated as follows:

$$b_i(t+1) = \begin{cases} 0, & t = 0 \\ \frac{b_i(t) \times N}{N+1}, & \text{hybrid strategy chose } \pi \wedge \text{burst success} \\ \frac{b_i(t) \times N+1}{N+1}, & \text{hybrid strategy chose } \pi \wedge \text{burst dropped} \\ \frac{b_i(t) \times N + (1 - prio(\pi, t))}{N+1}, & \text{hybrid strategy did not choose path } \pi \end{cases}$$

5 Numerical Results

In this section, we use simulation to investigate the performance benefits of path switching in OBS networks. We use the method of batch means to estimate the burst drop 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; however, they are so narrow that we omit them from the figures. We used two 16-node networks: a 4×4 torus network, and a 16-node network based on the 14-node NSF network. All the figures plot the burst drop probability against the “normalized network load” ρ_W , which is obtained by dividing the total load offered to the network by the number W of wavelengths: $\rho_W = \sum_{ij} \rho_{ij} / W$.

Let us first investigate the performance improvement that is possible with pure path switching over shortest path routing. We assume that each source has to select among $m = 2$ candidate paths to each destination; these are the two shortest link-disjoint paths in the network for the given source-destination pair. We compare four routing schemes: (1) Shortest-path (SP) routing, (2) WBLU path switching, (3) WLC path switching, and (4) EPP path switching.

Figures 1(a)-(b) plot the burst drop probability of the above four routing schemes for the NSF network with uniform traffic. Figure 1(a) (respectively, Figure 1(b)) plots the burst drop probability for low (respectively, high) loads. As we can see, all three path switching strategies perform consistently better than SP routing throughout the load range considered in the figures; the only exception is at very high loads, where the high burst drop probability is due to a saturated network. This result demonstrates the benefits of path switching.

Another important observation is that none of the three path switching strategies is a clear winner over the entire range of loads shown. In general, WBLU performs the best at low loads, EPP is the best strategy at high loads, while the burst drop probability of WLC is between the values of the other two

Fig. 1. Burst drop prob., NSF network, uniform traffic (a) Low load (b) High load

Fig. 2. Burst drop prob., Torus network, uniform traffic (a) Low load (b) High load

strategies. At low network loads, most links have low utilization, and avoiding the few highly utilized (bottleneck) links can significantly improve the burst drop probability. Since the WBLU strategy takes account the bottleneck link utilization in determining the burst path, it performs well at low loads. At high loads, the EPP strategy outperforms the WBLU and WLC strategies. This behavior can be explained by the manner in which the three strategies update their path decisions. Under EPP, path priorities are updated immediately upon the receipt of feedback messages from the network, whereas the WBLU and WLC strategies update their routing decisions periodically. The period of update for WBLU and WLC is independent of the network load. With the EPP strategy, however, as the load increases, the rate of feedback from the network increases accordingly, providing a more accurate view of the network state and resulting in better routing decisions.

The performance of the four routing methods for the Torus network and uniform traffic is shown in Figures 2(a)-(b). The WLC and EPP strategies perform consistently better than SP routing, and in fact the burst drop probability of EPP is significantly lower than that of both WLC and SP across the whole range

Fig. 3. Burst drop prob., NSF network, uniform traffic (a) Low load (b) High load

from low to high loads. The WBLU strategy, on the other hand, is only slightly better than SP at low loads, and slightly worse than SP at high loads. This result is due to the fact that WBLU makes a routing decision based only on the utilization of the bottleneck link. In a symmetric topology such as the Torus, WBLU leads to routing oscillations which tend to hurt the overall performance. In our experiments, we have observed that the oscillations persist throughout the simulation, and that they become worse as the offered load increases. In the asymmetric NSF network, on the other hand, we have observed that the routing decision of WBLU oscillates at first, but it later settles down to a fixed path. The only exception is at very high loads when the bottleneck links are saturated, in which case WBLU keeps oscillating among the candidate paths; this is reflected in Figure 1(b) for a load of 16, when WBLU performs worse than SP routing.

We now consider the WNV and DWNV hybrid strategies. Each hybrid strategy utilizes four routing strategies in making its decision: SP routing and the WBLU, WLC, and EPP pure path switching strategies. In order to characterize the performance of hybrid path switching, we compare the following three routing schemes: (1) WNV path switching, which assigns static weights (q -distributions) to each of the four pure strategies. We have found that different weights perform differently for each of the two topologies. Therefore, after some experimentation, we have used the following weights: for the NSF network, all weights are equal to $1/4$ (a uniform distribution), while for the Torus network, the weights are: $1/8$ (SP and WBLU), $1/4$ (WLC), and $1/2$ (EPP). (2) DWNV path switching, in which the weights of all pure strategies are initially equal, but they are adjusted dynamically during the operation of the network as in Section 4. (2) Best pure strategy, in which bursts are sent along the path determined by the pure strategy with the best performance among the four strategies SP, WBLU, WLC, and EPP. If one pure strategy is best across some range of loads while another is best across a different range, we present both strategies.

Figures 3(a)-(b) show results for the NSF network with uniform traffic; WBLU has the best performance among the pure path switching strategies at low loads, while EPP is the best pure strategy at high loads. The hybrid WNV path

switching scheme improves the burst drop probability over both pure strategies; in effect, the WNV curve tracks the best of the WBLU or EPP curves. This result confirms our intuition that taking into account several different views of the network state increases the performance. When the weight of each pure strategy is adjusted dynamically to reflect the real-time network performance, as accomplished by DWNV, the burst drop probability is further improved. Our experiments indicate that through dynamic adjustments, the weights assigned to each pure strategy by a source-destination pair are fine-tuned depending on the source-destination pair and the traffic pattern. This tuning procedure can be viewed as a “dynamic optimization” process that allows the hybrid DWNV strategy to achieve a final set of weights that is near-optimal in the sense of minimizing the burst drop probability.

6 Concluding Remarks

We considered the problem of multipath routing in OBS networks and we developed a suite of path switching strategies, each utilizing one type of information regarding the network state to select one of a set of paths to route a given burst. We presented a probabilistic framework for hybrid path switching strategies which make routing decisions by taking into account the decisions of multiple pure strategies. We can summarize our results as follows. **(1)** Pure path switching strategies can reduce the burst drop probability compared to shortest path routing. **(2)** The performance improvement depends on the congestion information utilized by the strategy, the network topology, and the load; in many cases, the improvement over shortest path routing can be dramatic. **(3)** Hybrid strategies can be used to further improve the performance. However, if one pure strategy clearly outperforms all others, then a hybrid strategy may not provide any improvement. In this case, it is best to simply use the most successful pure strategy instead. **(4)** The performance is optimized when the weights assigned to the pure strategies by a hybrid strategy can be appropriately selected. Otherwise, a hybrid strategy that dynamically adjusts the weights performs best.

References

1. Chen, Y., Qiao, C., Yu, X.: Optical burst switching: A new area in optical networking research. *IEEE Network* **18** (2004) 16–23
2. Yao, *et al.*: A unified study of contention-resolution schemes in optical packet-switched networks. *Journal of Lightwave Technology* **21** (2003) 672–683
3. Vokkarane, V., *et al.*: Burst segmentation: An approach for reducing packet loss in optical burst switched networks. In: *Proc. IEEE ICC*. (2002) 2673–2677
4. Teng, J., Rouskas, G.N.: On wavelength assignment in optical burst switched networks. In: *Proceeding of BROADNETS 2004*. (2004) 24–33
5. Teng, J., Rouskas, G.N.: Routing path optimization in optical burst switched networks. In: *Proc. ONDM 2005*. (2005, to appear)

6. Thodime, G., *et al.*: Dynamic congestion-based load balanced routing in optical burst-switched networks. In: Proc. Globecom. Volume 5. (2003) 2628–2632
7. Tao, S., *et al.*: Exploring the performance benefits of end-to-end path switching. In: Proc. ACM SIGMETRICS 2004, New York, NY, USA (2004) 418–419
8. Baldine, I., Rouskas, *et al.*: **JumpStart**: A just-in-time signaling architecture for WDM burst-switched networks. IEEE Communications Magazine **40** (2002) 82–89
9. Chu, X.W., *et al.*: A dynamic RWA algorithm in a wavelength-routed all-optical network with wavelength converters. In: Proc. IEEE INFOCOM. (2003) 1795–1802
10. Hansen, L., Salamon, P.: Neural network ensembles. IEEE Transactions on Pattern Analysis and Machine Intelligence **12** (1990) 993–1001
11. Dietterich, T.G.: Machine learning research: Four current directions. AI Magazine **18** (1997) 97–136