An Analysis of Limited Wavelength Translation in Regular All-Optical WDM Networks

Vishal Sharma, *Member, IEEE*, and Emmanouel A. Varvarigos, *Member, IEEE*

***Abstract*—We analyze limited-wavelength translation in regular all-optical, wavelength division multiplexed (WDM) networks, where up to W wavelengths, each of which can carry one circuit, are multiplexed onto a network link. All-optical wavelength translators with a limited translation range permit an incoming wavelength to be switched only to a small subset of the outgoing wavelengths. We focus on the wraparound mesh and hypercube WDM networks, and analyze the case where an incoming wavelength can be switched to one of k ($k = 2, 3$) outgoing wavelengths (called the *feasible wavelength set*). Our analysis captures the state of a feasible wavelength set at a network node, which allows us to obtain the probability that a session arriving at a node at a random time successfully establishes a connection from its source node to its destination node in each of these topologies. Based on this probability, we quantify the throughput and blocking performance of limited wavelength translation, and compare it to that of no wavelength translation and full wavelength translation. We demonstrate that in regular networks it can obtain most of the performance advantages of full translation at a fraction of the cost, and we present a simple, economical switch architecture to effect limited wavelength translation at a cost that is effectively independent of the number of wavelengths W in the system.**

***Index Terms*—All-optical network, hypercube, mesh, performance analysis, throughput, wavelength division multiplexing (WDM), wavelength translation.**

I. Introduction

RECENT interest in all-optical networks (e.g., [5], [3]) has focused attention on wavelength division multiplexing (WDM) as a promising technique for utilizing in a natural way the terahertz bandwidth of optical fiber. A critical functionality for the scalability and improved performance of multihop WDM networks is *wavelength translation*, that is, the ability of network nodes to switch data from an incoming wavelength ϕ_i to an outgoing wavelength ϕ_j, $j \neq i$. Two different classes of wavelength-routing nodes important in this context are: nodes with a *full-wavelength translation* capability, which translate an incoming wavelength to any outgoing wavelength and nodes with *no-wavelength translation* capability, which map each incoming wavelength to the same outgoing wavelength, the so called wavelength-continuity constraint (e.g., [4]). The requirement of wavelength continuity increases the probability of call blocking and can be avoided by the use of wavelength translation.

Although full-wavelength translation is desirable because it substantially decreases blocking probability ([2], [9]), it is difficult to implement in practice due to technological limitations. Since all-optical converters are still being prototyped in laboratories and are likely to remain expensive, researchers have turned their attention to searching for suitable alternatives. The analysis by Subramaniam *et al.* [14], found, for instance, that there is no significant degradation in network performance even when only a few (as opposed to all) of the network nodes have a full-wavelength translation (conversion) capability. A natural question that arises is, whether or not similar performance advantages can be obtained by using switches that have only a *limited wavelength translation* capability, where an incoming wavelength can be translated to only a small subset (as opposed to all) of the available outgoing wavelengths. This problem assumes practical significance when one considers that all-optical wavelength translators demonstrated in laboratories to date are, in general, capable only of limited translation (e.g., [22] and [20]).

Realizing this limitation, researchers have begun studying limited wavelength translation in a systematic way to quantify its advantages vis-a-vis no wavelength translation and full wavelength translation. Yates *et al.* [20] were the first to present a simple, approximate probabilistic analysis for single paths in isolation, while Gerstel *et al.* [23] were the first to examine limited wavelength translation for ring networks under a nonprobabilistic model. Recently, Wauters and Demeester [21] have provided new upper bounds on the wavelength requirements for a WDM network under a static model of the network load and have also looked at the problem under dynamic traffic conditions, while Ramaswami and Sasaki [13] have provided a nonprobabilistic analysis of the problem for ring networks and, under certain restrictions, for tree networks and networks of arbitrary topology. Although their models are valuable and make no assumptions about traffic behavior, they focus most often on the static (one-time) problem, where there is a set of source–destination pairs that have to be connected and the objective is to serve as many such connection requests as possible. Such a model does not capture the dynamic nature of a network, where requests arrive and leave at random instants of time continuously over an infinite time horizon.

In this paper, we analyze limited wavelength translation in regular all-optical WDM networks, such as a torus network or a hypercube network that uses either crossbar or *descending dimensions* switches. Descending dimensions switches are much

Manuscript received March 27, 2000; revised September 19, 2000. This work was supported by ARPA under the MOST and Thunder and Lightning Projects.

V. Sharma was with the Tellabs Research Center, Cambridge, MA 02139 USA. He is now with Jasmine Networks, Inc., San Jose, CA 95134 USA.

E. A. Varvarigos is with the Data Transmission and Networking Laboratory, Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA 93106 USA.

Publisher Item Identifier S 0733-8724(00)10601-2.

simpler to build and implement than the crossbar switches usually assumed for the hypercube. We consider *k-adjacent wavelength switching,* where an incoming wavelength can be translated only to a subset consisting of k of the W outgoing wavelengths (i.e., to $k-1$ wavelengths in addition to itself), and we call the set of output wavelengths the *feasible wavelength set.* We wish to answer the following question: For a given network topology, how much of the performance improvement provided by full-wavelength translation over no wavelength translation, can we achieve by using switches with only a limited translation capability? Namely, for a given probability of success, how much of the improvement in sustainable load provided by a network with full wavelength translation does a network with limited wavelength translation provide?

Our analysis for the torus and hypercube topologies demonstrates that k-adjacent wavelength switching with only $k = 2, 3$ suffices to give performance significantly superior to that obtained with no wavelength translation, and close to that obtained with full wavelength translation. The results that we obtain are very close to the corresponding simulation results, and predict network performance over a wider range of network loads than previous works. We find that the extent of translation k does not have to be a function of the total number of wavelengths W, as argued in previous works. This is because the models used by previous works are based either on trying multiple options simultaneously, or on using backtracking[1] when establishing the connection [20], [14], so that the improvements obtained in those cases increased both with k and with W. When a more realistic scenario is considered, where at each hop a setup packet examines only the outgoing wavelengths at that hop, instead of examining all possible paths from that hop onward, it turns out (see Sections II and III) that the relative improvement in performance when k is increased does not depend on W. Our results show that in this case relatively small values of k yield performance close to that of full wavelength translation.

We also look at an implementation of a limited wavelength translation switch, which is based on the concept of *bulk frequency conversions* [1], where several wavelength translations can be performed in parallel. We show that the switch is composed of simple components with a cost that depends mainly on the extent k of translation, and depends only weakly on W.

The remainder of the paper is organized as follows. We focus first on the torus as a representative network and analyze it in detail in Section II. In particular, in Section II-A, we explain the node model used, while in Sections II-B and II-C, we discuss the auxiliary queuing system that we use for the case of $k = 2$-adjacent wavelength switching. In Section II-D, we derive the expression for the probability that a session arriving at a random time successfully establishes a circuit. In Section III, we compare our analytic results for the success probability with those obtained by simulations for the torus, and discuss our results. In Section IV, we extend our previous analysis to hypercubic networks. Specifically, in Section IV-A, we consider the hypercube network with simple *descending dimensions* switches (or the *wrapped butterfly* network) and in Section IV-B, we consider the hypercube network with *cross-bar* switches. In Section V, we present a simple, cost effective implementation of a limited wavelength translation switch. Our conclusions appear in Section VI.

[1]Here, backtracking refers to the case where, at a given node, the setup packet of a session attempts to make a reservation, one by one, along each of the possible wavelength choices, until it finds a sequence of wavelengths along which it is able to establish a circuit all the way to its destination.

II. Limited Wavelength Translation in the Torus Network

In this section, we examine limited wavelength translation in a two-dimensional (2-D) wraparound mesh network, more popularly known as the *torus* network. We focus on the mesh because we expect the results obtained for it to be representative of the performance obtainable from practical networks, which also tend to have small degree and several paths of similar distance between each source–destination pair. (Other authors, e.g., [7] and [17], have also verified the accuracy of their analyses by applying it to the particular case of mesh and hypercube networks.)

A $p \times p$ wraparound mesh consists of p^2 nodes arranged along the points of a 2-D space with integer coordinates, with p nodes along each dimension. Two nodes (x_2, x_1) and (y_2, y_1) are connected by a (bidirectional) link if and only if for some $i = 1, 2$ we have $|x_i - y_i| = 1$ and $x_j = y_j$ for $j \neq i$. In addition to these links, wraparound links connecting node $(x_2, 1)$ with node (x_2, p), and node $(1, x_1)$ with node (p, x_1) are also present. The *routing tag* of a session with source node $x = (x_2, x_1)$ and destination node $y = (y_2, y_1)$, is defined as (t_2, t_1), where

$$t_j = \begin{cases} y_j - x_j & \text{if } |y_j - x_j| \le \left\lfloor \frac{p}{2} \right\rfloor \\ y_j - x_j - p \cdot \text{sgn}(y_j - x_j) & \text{if } |y_j - x_j| > \left\lfloor \frac{p}{2} \right\rfloor \end{cases} \quad (1)$$

for all $j \in \{1, 2\}$, and sgn(x) is the signum function, which is equal to $+1$ if $x \ge 0$, and equal to -1, otherwise.

In what follows, we analyze *$k = 2$-adjacent wavelength switching,* where each wavelength ϕ_j, $j = 0, 1, \cdots, W-1$ on an incoming link can be switched on the outgoing link, either to the same wavelength ϕ_j or to the next higher wavelength ϕ_{j+1}. (A similar analysis exits for the case of 3-adjacent wavelength switching [16].) We call the set $\{\phi_{j+1}, \phi_j\}$ of output wavelengths, the *feasible wavelength set* of input wavelength ϕ_j. For symmetry, we assume that the boundary wavelength ϕ_{W-1} can be switched to wavelengths ϕ_{W-1} and ϕ_0. Note that this is not merely an analytical convenience, but is also essential for performance so that the total load is distributed equally among all wavelengths on a link.

A. The Node Model

In our model, connection requests or unidirectional sessions arrive independently at each node of the torus over an infinite time horizon according to a Poisson process of rate ν, and their destinations are distributed uniformly over all nodes, except the source node. Each session wishes to establish a circuit (or connection) to its destination node for a duration equal

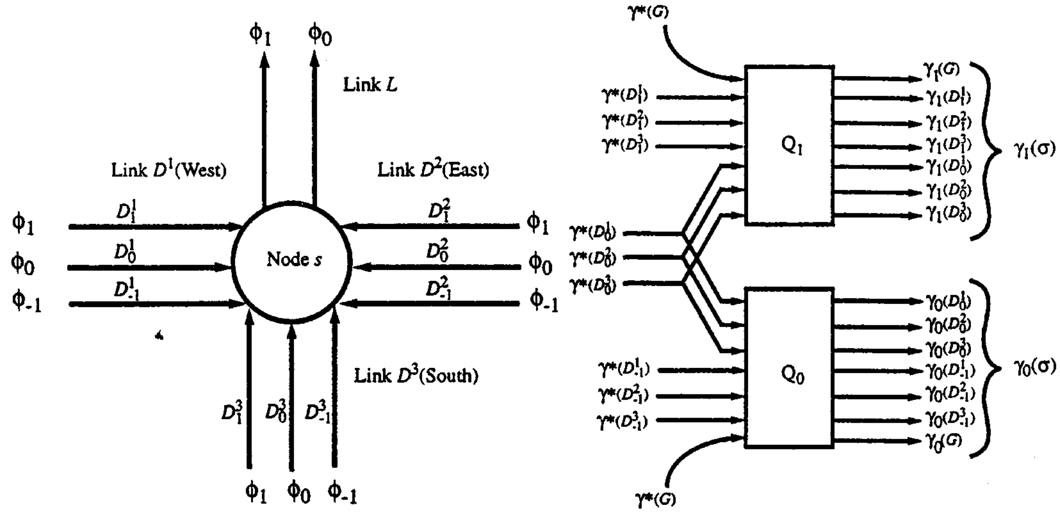

(a) (b)

Fig. 1. (a) Relationship of the set of incoming wavelengths to the set of outgoing wavelengths at a node s of the network when 2-adjacent wavelength switching is used. (b) The auxiliary system $\hat{\mathcal{Q}}$ (see Section II-B), which consists of two subsystems $\mathcal{Q}_1$ and $\mathcal{Q}_0$, each consisting of one server and no waiting space. The rates $\gamma^*(\sigma)$ are chosen so that rate at which customers of type σ depart from the subsystem l is $\gamma_l(\sigma)$, $l = 0, 1$.

to the holding time X of the session, which is exponentially distributed with mean $\overline{X}$, and does so by transmitting a setup packet along a path to its destination. Setup packets use *dimension-order* routing to establish the circuit, that is, they either traverse all links along dimension 1 (horizontal) followed by all links along dimension 2 (vertical) or visa versa. In our model, links are bidirectional, so all incident links of a node (and their wavelengths) can be used simultaneously for transmission and reception. Our routing scheme is oblivious, or nonadaptive; that is, the path followed by the session is chosen at the source, and the setup packet insists on that path as it progresses from its source to its destination. If the setup packet of a session is successful in establishing the circuit, the wavelengths required by it are reserved for the duration of the session. Otherwise, the session is reinserted randomly into an input queue for the node (which is ordered as per the times at which the sessions in it must retry), and tries again when its turn arrives. Thus, a session may try several times before it is finally successful. The interarrival time between sessions inserted in the input queue is several times the interarrival time between external arrivals, so that the combined arrival process of the new sessions and the reinserted sessions can still be approximated as a Poisson process. Observe that this is an analytical convenience that simplifies modeling and our subsequent analysis (see Sections II-A to II-D), without detracting from our goal of obtaining the probability that a new session arriving at a *random* time is successful in establishing a lightpath. In a practical system, however, it is expected that blocked sessions will be allowed to try more frequently. As we argue later, such a system would not affect the blocking probability for new sessions arriving at a random time, even though it would of course have smaller delay and would be preferable.

By contrast, in the call dropping model used in other analyses, sessions with longer path lengths are dropped with a higher probability. The call dropping model, in addition to treating long connections unfairly, also tends to overestimate the maximum throughput achievable, by favoring connections that require fewer hops.

In 2-adjacent wavelength switching, the switching of a new session arriving on a wavelength ϕ_j on an incoming link of a node (or a new session originating at the node that wishes to use wavelength ϕ_j on its first outgoing link) depends only on the availability of wavelengths ϕ_j and ϕ_{j+1} on the outgoing link, since those are the only wavelengths that ϕ_j can be switched to. Because of symmetry, we focus, without loss of generality, on an incoming session that arrives over a particular wavelength, say ϕ_0, and wants to use outgoing link L of a node. Such a session can only be switched to wavelengths ϕ_0 or ϕ_1 on link L. Since our eventual goal is to evaluate the probability that a session arriving on wavelength ϕ_0 on an incoming link finds wavelengths ϕ_0 and ϕ_1 on the outgoing link busy, we proceed by enumerating all the possible ways in which these wavelengths can be busy. Therefore, as depicted in Fig. 1(a), we concentrate only on all those wavelengths on the remaining incoming links that can be switched either to wavelength ϕ_0 or ϕ_1 on link L. This allows us to characterize the state of the wavelengths ϕ_0 and ϕ_1 on outgoing link L, by the incoming (link, wavelength) pairs that are at any time using wavelengths ϕ_0 and ϕ_1 on link L. The following development, in essence, formalizes the approach outlined above.

As shown in Fig. 1, D_1^1, D_1^2, and D_1^3 denote wavelength ϕ_1 on the incoming links D^1, D^2, and D^3, respectively, which can be switched to wavelength ϕ_1 on link L. Similarly, D_{-1}^1, D_{-1}^2, and D_{-1}^3 denote wavelength ϕ_{-1} on links D^1, D^2, and D^3, respectively, which can be switched to wavelength ϕ_0 on link L. (In our notation, wavelength ϕ_{-1} corresponds to wavelength ϕ_{W-1}.) Finally, D_0^1, D_0^2, and D_0^3 denote wavelength ϕ_0 on the incoming links D^1, D^2, and D^3, respectively, which can be switched either to wavelength ϕ_0 or to wavelength ϕ_1 on link L.

We define the joint state of a pair of outgoing wavelengths according to whether they are being used or not by ongoing connections and, if they are being used, according to the incoming link and wavelength over which these connections arrive. The state of the outgoing wavelengths ϕ_1 and ϕ_0 is therefore represented by the pair $\overline{S} = (S_1, S_0)$, where S_1 and S_0 denote those wavelengths on the incoming links that are being switched to wavelengths ϕ_1 and ϕ_0 respectively on link L. A wavelength on link L may also be idle or be used by a new session originating at node s. We denote by G the state in which a wavelength is being used by a newly generated session at a node, and by I the state in which a wavelength is idle. Based on this, the possible states of wavelengths ϕ_0 and ϕ_1 of link L are

$$S_0 \in \{D^1_{-1}, D^1_0, D^2_{-1}, D^2_0, D^3_{-1}, D^3_0, I, G\} \triangleq \mathcal{S}_0 \quad \text{(2a)}$$

and

$$S_1 \in \{D^1_0, D^1_1, D^2_0, D^2_1, D^3_0, D^3_1, I, G\} \triangleq \mathcal{S}_1. \quad \text{(2b)}$$

We denote by Ω the set of states of a feasible wavelength set, where

$$\Omega = \{(S_1, S_0) | \Lambda_1 \in \mathcal{S}_1, S_0 \in \mathcal{S}_0, S_0 \neq S_1 \text{ or } S_0 = S_1 = G, \text{ or } S_0 = S_1 = I\}$$

and we denote by $\pi(\overline{S})$ the steady-state probability that a feasible wavelength set on link L is in state $\overline{S}$. The states (D^1_0, D^1_0), (D^2_0, D^2_0), and (D^3_0, D^3_0), where two outgoing wavelengths are being simultaneously used by a session arriving on a wavelength on an incoming link, are infeasible, since we only consider unicast communication, where an incoming session is switched to a single outgoing wavelength. The states (I, I) and (G, G) however are feasible, and correspond to both outgoing wavelengths being idle or to both being used by originating sessions, respectively. We use the convention that the probability $\pi(\overline{S})$ of an infeasible state $\overline{S}$ is zero. The number of states in Ω for a d-dimensional mesh can be found to be $16d^2 - 2d + 1$ when $k = 2$-adjacent wavelength switching is used, for a total of 61 states for the torus network ($d = 2$). The case $k = 3$, where an incoming wavelength can be switched to two wavelengths in addition to itself is also examined in [16]; in that case the number of feasible states of a d-dimensional mesh is $(6d-1)^3 - (30d-7)(2d-1)$, for a total of 1172 states for a torus network.

B. The Auxiliary System

To develop our analysis further, we focus on setup packets emitted on wavelength ϕ_0 or ϕ_1 of link L, and we define the *type* σ of a setup packet (and of the corresponding session) according to whether it belongs to a session originating at the node or according to the incoming link and wavelength D^i_j, $i = 1, 2, 3$, $j = 1, 0, -1$, upon which it arrives. Therefore, the set of possible types σ is $\mathcal{S}_0 \cup \mathcal{S}_1 - \{I\}$. We also let $\gamma_l(\sigma)$ denote the rate at which setup packets of type σ are emitted on an outgoing wavelength l, $l = \phi_0, \phi_1$, at a node of the torus. Our problem now is to calculate $\pi(\overline{S})$, $\overline{S} \in \Omega$. To do so we approximate $\pi(\overline{S})$ as the stationary distribution of an auxiliary system $\hat{\mathcal{Q}}$ defined as follows [see Fig. 1(b)].

The system $\hat{\mathcal{Q}}$ consists of two subsystems $\mathcal{Q}_1$ and $\mathcal{Q}_0$, each of which has a single server and no waiting space. Customers of type σ arrive at the system $\hat{\mathcal{Q}}$ according to a Poisson process with rate $\gamma^*(\sigma)$. Customers of type $\sigma = D^i_1$, $i = 1, 2, 3$ (that is, those arriving on wavelength ϕ_1 on either one of the three incoming links) ask for the server of $\mathcal{Q}_1$, while those of type $\sigma = D^i_{-1}$, $i = 1, 2, 3$ (that is, those arriving on wavelength ϕ_{-1} on either one of the three incoming links) ask for the server of $\mathcal{Q}_0$, and these customers are lost if the corresponding server is busy. Customers of type $\sigma = D^i_0$, $i = 1, 2, 3$ (that is, those arriving on wavelength ϕ_0) ask for server $\mathcal{Q}_0$ or $\mathcal{Q}_1$ with equal probability. If the server they ask for is busy, they proceed to the other server, and are dropped only if both servers are busy.

Finally, we require that the rate at which customers of type σ are accepted in the auxiliary system $\hat{\mathcal{Q}}$ be the same as the rate $\gamma_l(\sigma)$ at which setup packets are emitted from an incoming (link, wavelength) pair σ to the outgoing wavelength ϕ_l in the actual system. For this to hold, we must have (3), shown at the bottom of the page. Equation (3) basically says that the rate $\gamma_l(\sigma)$ at which setup packets arriving on an incoming (link, wavelength) pair σ are *successful* (that is, they obtain their required wavelength ϕ_l on the outgoing link) is equal to the rate $\gamma^*(\sigma)$ at which they arrive on the incoming (link, wavelength) pair times the probability that their required outgoing wavelength is available.

$$\gamma^*(\sigma) = \begin{cases} \dfrac{2\gamma_l(\sigma)}{\displaystyle\sum_{i \in \mathcal{S}_1, i \neq \sigma} \pi(i, I) + \sum_{j \in \mathcal{S}_0, j \neq \sigma} \pi(I, j) - \pi(I, I)}, & \text{for } \sigma \in \{D^1_0, D^2_0, D^3_0\}; \\ \dfrac{\gamma_0(\sigma)}{\displaystyle\sum_{i \in \mathcal{S}_1} \pi(i, I)}, & \text{for } \sigma \in \{D^1_{-1}, D^2_{-1}, D^3_{-1}\}; \\ \dfrac{\gamma_1(\sigma)}{\displaystyle\sum_{i \in \mathcal{S}_0} \pi(I, j)}, & \text{for } \sigma \in \{D^1_1, D^2_1, D^3_1\}; \\ \dfrac{\gamma_0(\sigma)}{\displaystyle\sum_{i \in \mathcal{S}_1} \pi(i, I)} = \dfrac{\gamma_1(\sigma)}{\displaystyle\sum_{i \in \mathcal{S}_0} \pi(I, j)} & \text{for } \sigma = G \end{cases} \quad (3)$$

C. Calculation of Arrival Rates for the Auxiliary System

To complete our development, what remains is that we obtain the rates $\gamma_l(\sigma)$ at which setup packets that arrive over the incoming (link, wavelength) pair σ are emitted over outgoing wavelength ϕ_l. We do so in this section.

To calculate the rates $\gamma_l(\sigma)$, we observe that each wavelength of a network link can be in one of four states [16]: State 0, which corresponds to the case where the wavelength is idle; State 1, which corresponds to the case where a wavelength on the outgoing link L at a node s is used by a session originating at node s, that is, a session of type $\sigma = G$ (or a *class 1* session); State 2, which corresponds to the case where a wavelength on L is used by a session that is going straight at node s, that is, a session of type $\sigma = D_j^3$, $j = 0, 1, -1$ (or a *class 2* session); and State 3, which corresponds to the case where a wavelength on L is used by a session that is making a turn at node s, that is, a session of type $\sigma = D_j^i$, $i = 1, 2$ and $j = 0, 1, -1$ (or a *class 3* session). Analogously, setup packets that place a link in state 1, 2, or 3, will be referred to as being of class 1, 2, or 3, respectively.

In our routing scheme, destinations are distributed uniformly over all nodes (excluding the source node) and no sessions are dropped (recall that blocked sessions are randomly reinserted into the input stream). Thus, wavelength utilization is uniform across all wavelengths of the network, and the probabilities q_0, q_1, q_2, and q_3 that a wavelength is in state 0, 1, 2, or 3 can be obtained simply by counting all possible ways in which a wavelength can be in each of these states and normalizing appropriately. (Their calculation is given in the Appendix.) We denote by ρ the total rate at which setup packets that are finally successful are emitted on an outgoing wavelength of a link, and by ρ_i, $i = 1, 2, 3$ the total rate at which class i setup packets that are finally successful are emitted on a given outgoing wavelength of a link. Clearly, $\rho = \rho_1 + \rho_2 + \rho_3$. Applying Little's Therem to the system that consists of class i sessions, and again to the entire network, we obtain

$$\rho = \rho_1 + \rho_2 + \rho_3 = \frac{q_1}{\overline{X}} + \frac{q_2}{\overline{X}} + \frac{q_3}{\overline{X}} = \frac{\nu \overline{h}}{4W} \tag{4}$$

where $\overline{h}$ is the mean internodal distance of the torus, which can be calculated to be

$$\overline{h} = \begin{cases} p/2, & \text{if } p \text{ is odd} \\ (p/2)\dfrac{p^2}{p^2-1}, & \text{if } p \text{ is even.} \end{cases} \tag{5}$$

Referring once again to the actual system depicted in Fig. 1(a), we see that the rates $\gamma_l(\sigma)$, at which setup packets of type σ are emitted on wavelength ϕ_l, $l = 0$ or 1, of an outgoing link L, can be related in the following way to the rates ρ_i, $i = 1, 2, 3$ at which class i setup packets are emitted on outgoing wavelength l (we write the equations for outgoing wavelength ϕ_1; the case of wavelength ϕ_0 is similar)

$$\rho_3 = \gamma_1(D_0^1) + \gamma_1(D_1^1) + \gamma_1(D_0^2) + \gamma_1(D_1^2) \tag{6a}$$

$$\rho_2 = \gamma_1(D_0^3) + \gamma_1(D_1^3) \tag{6b}$$

and

$$\rho_1 = \gamma_1(G). \tag{6c}$$

Using the symmetry of the rates on the wavelengths of incoming links, we obtain

$$\gamma_l(\sigma) = \begin{cases} \dfrac{\rho_3}{4}, & \text{for } \sigma \in \{D_1^1, D_0^1, D_{-1}^1, D_1^2, D_0^2, D_{-1}^2\} \\ \dfrac{\rho_2}{2}, & \text{for } \sigma \in \{D_1^3, D_0^3, D_{-1}^3\} \\ \rho_1, & \text{for } \sigma = G. \end{cases} \tag{7}$$

Equation (7) when substituted into (2) and (3) gives us the arrival rates for the auxiliary system of Fig. 2(b).

The steady-state probabilities $\pi(\overline{S})$, $\overline{S} \in \Omega$, of the auxiliary system (which can be obtained numerically without approximations) are used to approximate the probability that the wavelengths ϕ_0 and ϕ_1 (or, more generally, ϕ_j and ϕ_{j+1}) of the outgoing link L are in state $\overline{S}$. We are now very close to our final goal of obtaining the probability that a setup packet arriving on a certain wavelength on an incoming link finds the wavelengths of its feasible wavelength set (on the outgoing link) busy. In the following section, we use the state probabilities $\pi(\overline{S})$, $\overline{S} \in \Omega$ just obtained to derive an expression for the probability that a session successfully establishes a circuit to its destination.

D. Probability of Successfully Establishing a Circuit

The probability that a new session arriving at a random time successfully establishes a circuit from its source node to its destination node depends on its routing tag (t_2, t_1), which is essentially the number of hops that a session has to traverse along the vertical and horizontal axes to get from its source node to its destination node, and was formally defined in Section II. We denote by $P_{\rm succ}(t_2, t_1)$ the probability that a new session will succeed in establishing a connection on a particular trial, given that it has a routing tag (t_2, t_1). Using the approach in [16], in this section we will first find an approximate expression for $P_{\rm succ}(t_2, t_1)$, and from it an expression for $P_{\rm succ}$ the probability that a new session trying to establish a circuit at a random time is successful.

The path followed by a session with routing tag (t_2, t_1) will make $I(t_2, t_1) - 1$ turns along the way, and will go straight for a total of $|t_2| + |t_1| - I(t_2, t_1)$ hops, where $I(t_2, t_1) \in \{0, 1, 2\}$ is the number of nonzero entries in (t_2, t_1). Thus, the probability α_0 that a wavelength on the first link of the path is available is the probability that the chosen wavelength on the outgoing link is idle, and is given simply by

$$\alpha_0 = \sum_{i \in \mathcal{S}_1} \pi(i, I). \tag{8}$$

At each hop at which a session does not make a turn, the probability that a wavelength ϕ_j or ϕ_{j+1} on the next link L is available given that a wavelength ϕ_j on the previous link $L-1$ used by the session was available, is

$$\begin{aligned}&\Pr(\phi_j \text{ or } \phi_{j+1} \text{ on } L \text{ available } | \phi_j \text{ on vertical link } \\ &\quad L-1 \text{ available}) \\ &= \Pr(\phi_j \text{ or } \phi_{j+1} \text{ on } L \text{ available } | \\ &\quad \text{neither } \phi_j \text{ nor } \phi_{j+1} \text{ on } L \text{ used by } \phi_j \\ &\quad \text{on vertical link } L-1) \\ &= \frac{\sum_{\substack{i\in\mathcal{S}_1 \\ i\neq D_0^3}} \pi(i, I) + \sum_{\substack{l\in\mathcal{S}_0 \\ l\neq D_0^3}} \pi(I, l) - \pi(I, I)}{1 - \sum_{i\in\mathcal{S}_1} \pi(i, D_0^3) - \sum_{l\in\mathcal{S}_0} \pi(D_0^3, l)} \triangleq \alpha. \end{aligned} \tag{9}$$

The numerator in the above equation is the sum over all states where either wavelength ϕ_j or ϕ_{j+1} is idle, excluding states where either wavelength is in use by a session arriving on wavelength ϕ_j on a vertical link. The denominator is simply one minus the sum over those states where either wavelength ϕ_j or ϕ_{j+1} is in use by a session arriving on wavelength ϕ_j on a vertical link. The latter case must be excluded since the session under consideration is arriving on wavelength ϕ_j of a vertical link, implying that wavelengths on the outgoing link could not already be in use by a session on wavelength ϕ_j.

Similarly, at each hop at which a session makes a turn, the probability that a wavelength ϕ_j or ϕ_{j+1} on the next link L is available given that a wavelength ϕ_j on the previous link $L-1$ used by the session was available, is

$$\begin{aligned}&\Pr(\phi_j \text{ or } \phi_{j+1} \text{ on } L \text{ available} | \phi_j \text{ on horizontal link } \\ &\quad L-1 \text{ available}) \\ &= \Pr(\phi_j \text{ or } \phi_{j+1} \text{ on } L \text{ available} | \\ &\quad \text{neither } \phi_j \text{ nor } \phi_{j+1} \text{ on } L \text{ used by } \phi_j \\ &\quad \text{on horizontal link } L-1) \\ &= \frac{\sum_{\substack{i\in\mathcal{S}_1 \\ i\neq D_0^1}} \pi(i, I) + \sum_{\substack{l\in\mathcal{S}_0 \\ l\neq D_0^1}} \pi(I, l) - \pi(I, I)}{1 - \sum_{i\in\mathcal{S}_1} \pi(i, D_0^1) - \sum_{l\in\mathcal{S}_0} \pi(D_0^1, l)} \triangleq \beta. \end{aligned} \tag{10}$$

The logic for the above equation is exactly the same as that discussed for (9). In writing (9) and (10) above, we do *not* assume that the probabilities of acquiring wavelengths on successive links of a session's path are independent. Instead, we account partially for the dependence between the acquisition of successive wavelengths on a session's path by using the approximation that the probability of acquiring a wavelength on link L depends on the availability of the corresponding wavelength on link $L-1$ (in reality this probability depends, even though very weakly, on the availability of corresponding wavelengths on all links $1, 2, \cdots, L-1$ preceding link L). As the simulation results presented in the next section demonstrate, our approximation is a very good one, however.

The (conditional) probability of successfully establishing a connection is then given by

$$P_{\text{succ}}(t_2, t_1) = \alpha_0 \cdot \alpha^{|t_2|+|t_1|-I(t_2,t_1)} \cdot \beta^{I(t_2,t_1)-1}.$$

For uniformly distributed destinations, we have $P_{\text{succ}} = \sum_{t_2, t_1} P_{\text{succ}}(t_2, t_1)/(p^2-1)$, which after some algebra, yields

$$P_{\text{succ}} = \frac{\alpha_0}{\beta(p^2-1)} \left[\left(1 + 2\beta \left(\frac{1-\alpha^D}{1-\alpha} \right) \right)^2 - 1 \right] \tag{11}$$

for p odd, and

$$\begin{aligned} P_{\text{succ}} = \frac{\alpha_0}{\beta(p^2-1)} \Bigg[\Bigg(1 + \beta \left(\frac{1-\alpha^D}{1-\alpha} \right) \\ + \beta \left(\frac{1-\alpha^{D-1}}{1-\alpha} \right) \Bigg)^2 - 1 \Bigg] \end{aligned} \tag{12}$$

for p even, where $D = \lfloor p/2 \rfloor$, and α_0, α and β are given by (8), (9), and (10), respectively. The probability that a session is blocked is therefore given by

$$P_{blk} = 1 - P_{\text{succ}}. \tag{13}$$

As mentioned in Section II-A, two possible models for retrials are: our model or the *delayed retrial* model, where blocked sessions try again after a random time, and a *persistent retrial* model, where blocked sessions retry continuously. Since in the current formulation retrials have zero (or negligible) cost, the network or link utilization will be similar in both cases implying that the blocking probabilty P_{blk} for new arrivals will also be the same in the two cases. The delay for the persistent retrial scheme would be better, so this scheme could be expected to be used in practice.

III. Results and Discussion

In this section, we present simulation and analytical results for the torus network for three different cases: the case of no wavelength translation (or 1-adjacent wavelength switching); the case of limited wavelength translation using k-adjacent wavelength switching, where $k = 2, 3$; and the case of full wavelength translation (or W-adjacent wavelength switching) in a WDM network with W wavelengths per link (fiber). We note that full-wavelength translation provides the best achievable performance (in terms of the realizable probability of success for a given arrival rate per wavelength, or in terms of the realizable throughput per wavelength for a given probability of success) for a given number of wavelengths W per fiber. When no wavelength translation is used, the different wavelengths on a link do not interact with one another. Thus, an all-optical network with W wavelengths per fiber is essentially equivalent to W disjoint single-wavelength networks operating in parallel. To obtain the probability of success in this case, it is therefore enough to focus attention on any one of the W independent parallel networks, for which the analysis given in [16] applies.

We define the degree of translation δ of a k-adjacent wavelength switching system with W wavelengths per fiber to be

$$\delta = \frac{k-1}{W-1} \times 100\%.$$

Thus, $\delta = 100\%$ corresponds to the case of full wavelength translation (or W-adjacent wavelength switching), while $\delta =$

0% corresponds to the case of no wavelength translation (or 1-adjacent wavelength switching).

We define $P_{\rm succ}(\lambda, k)$ to be the probability of success in a k-adjacent wavelength switching system when the arrival rate λ per node per wavelength is equal to ν/W, and we define $\lambda(P_{\rm succ}, k)$ to be the throughput per node per wavelength of a k-adjacent wavelength switching system when the probability of success is equal to $P_{\rm succ}$. To quantify the performance of limited wavelength translation vis-a-vis full- or no-wavelength translation we also define the *throughput efficiency* $\Delta\lambda(P_{\rm succ}, k)$ of a k-adjacent wavelength switching scheme with W wavelengths per fiber, for a given probability of success $P_{\rm succ}$, to be

$$\Delta\lambda(P_{\rm succ}, k) = \frac{\lambda(P_{\rm succ}, k) - \lambda(P_{\rm succ}, 1)}{\lambda(P_{\rm succ}, W) - \lambda(P_{\rm succ}, 1)} \times 100\% \quad (14)$$

and the *success efficiency* $\Delta P_{\rm succ}(\lambda, k)$ of a k-adjacent wavelength switching system, for a given arrival rate per node per wavelength λ, to be

$$\Delta P_{\rm succ}(\lambda, k) = \frac{P_{\rm succ}(\lambda, k) - P_{\rm succ}(\lambda, 1)}{P_{\rm succ}(\lambda, W) - P_{\rm succ}(\lambda, 1)} \times 100\%. \quad (15)$$

The throughput and success efficiencies represent the degree of improvement (over no wavelength translation) in the throughput and in the probability of success respectively, that is obtained when limited wavelength translation with k-adjacent wavelength switching is used, as a percentage of the improvement obtained when full wavelength translation is used. For $k = W$ (full wavelength translation), we get $\Delta\lambda(P_{\rm succ}, k) = 100\%$ and $\Delta P_{\rm succ}(\lambda, k) = 100\%$, while for $k = 1$ (no wavelength translation), we get $\Delta\lambda(P_{\rm succ}, k) = 0\%$ and $\Delta P_{\rm succ}(\lambda, k) = 0\%$ or no improvement.

In Fig. 2 we present the results comparing our analysis and simulations for the probability of success $P_{\rm succ}$ plotted versus the arrival rate per node per wavelength $\lambda = \nu/W$ when limited wavelength translation to only one or two additional wavelengths (i.e, $k = 2, 3$) is permitted. The analytical results for $k = 3$-adjacent wavelength switching shown here, were obtained using the analysis presented in [16]. Each point in our simulation was obtained by averaging over 2×10^6 successes. Observe that limited translation to only one or two adjacent wavelengths provides a considerable fraction of the improvement that full wavelength translation provides over no wavelength translation. These benefits are summarized in Table I, where we illustrate the throughput and success efficiencies for a $p \times p$ torus ($p = 11$) for a few selected points.

Also the benefits of wavelength translation diminish as the extent of translation k increases, and eventually appear to saturate. We see therefore that our analysis and simulations predict that limited translation of small range, i.e., $k = 2$ or 3, gives most of the benefits obtained by full wavelength translation, where $k = W$. See for instance Fig. 2(b), which also illustrates the network performance for $k = W = 20$ wavelengths. As is evident from the plots, increasing the extent of translation k beyond some value leads to diminishing returns. (In fact, based on our observations, we feel that increasing k beyond the diameter of

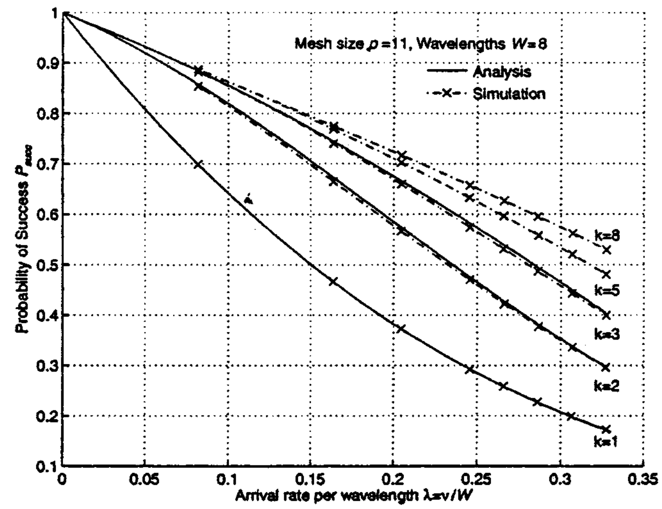

(a)

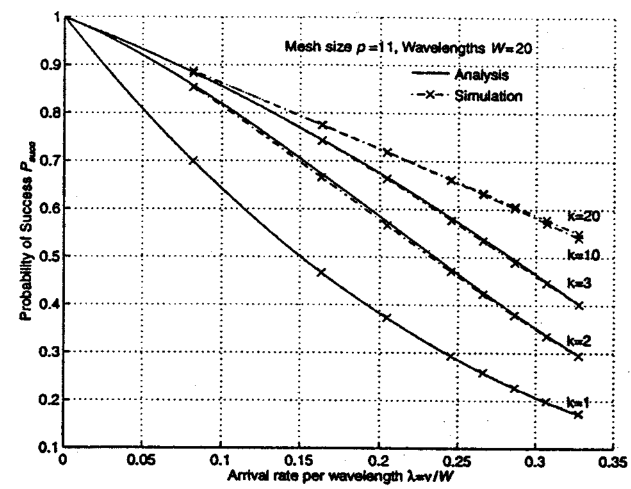

(b)

Fig. 2. Success probability $P_{\rm succ}$ versus the arrival rate per wavelength λ, for a mesh of size $p = 11$, for $W = 8$ and $W = 20$ wavelengths per link. All calculations and simulations have been performed with session holding times exponentially distributed with mean $\overline{X} = 1.0$.

the network will result in only a negligible increase in the probability of success.) Our results have recently also been corroborated by Tripathi and Sivarajan [17], who found that for arbitrary topology networks also, translation to $k = 2, 3$ adjacent wavelengths gave almost all of the benefits of full wavelength translation. The difference between their work and ours, however, is that the analysis in [17] still uses the "blocked calls cleared" model, and involved modified reduced-load approximations and so is tractable only for small networks in the regime of very low blocking probability.

By contrast, using the model that they had considered, Yates *et al.* [20] had predicted based on their simulations that limited wavelength translation of degree $\pm 50\%$ [which corresponds approximately to $k = (W-1)/2$ in our notation, since Yates *et al.* define the degree of translation by the number of translations allowed on *either side* of the input wavelength] would be

TABLE I
QUANTIFYING THE BENEFITS OBTAINED WITH LIMITED WAVELENGTH TRANSLATION, WITH $k = 2, 3$

p	W	$\Delta P_{succ}(0.25, 2)$	$\Delta P_{succ}(0.25, 3)$	$\Delta\lambda(0.7, 2)$	$\Delta\lambda(0.7, 3)$
11	5	51%	86%	61%	87%
11	20	47%	79%	49%	76%

needed to give performance nearly equal to that of full wavelength translation. Yates *et al.*, however, obtained improved results by using a model where a new session attempts to establish a connection by trying different starting wavelengths, and, in fact, by examining all alternate routes emanating from each of those wavelengths. Such improvements are based essentially on trying multiple options simultaneously, or alternatively on using backtracking while establishing the connection, which would not be practical in many cases in view of roundtrip delays and the blocking of other sessions that will occur during the setup phase, if it is done through multiple setup packets. (This is because, when multiple wavelengths are tried in parallel, capacity must be held on multiple paths, until the node initiating the multiple setups learns which of the alternate paths was successful, and removes partial or possibly complete reservations along other paths. The longer the roundtrip delays, the longer this takes and the longer the capacity reserved on the multiple paths blocks future sessions. This could be partially mitigated by intelligent, distributed reservation and connection control schemes [11], [19].) The improvements obtained by trying different options increase with increasing W and k. They increase with k because at each node a session examines all k options (which would be impractical when the roundtrip delay is comparable to session holding time, but could be practical otherwise), and they increase with W because at the first hop a session examines all W wavelengths on its outgoing link. Indeed if the link utilization information available at a source is accurate and connection setup takes zero time (that is, the roundtrip delay $\approx$0), then the results of Yates' *et al.* correspond to finding a path when running the routing algorithm, and it would be possible in this case to gain even further by trying not only alternative wavelengths, but also alternative topological paths from the source to the destination.

Our model, on the other hand, considers the more practical (in many cases) situation, where, at its first hop, a new session selects a wavelength upon which to attempt and at each subsequent hop looks only at the availability of the wavelengths at that hop and not of those beyond, and therefore does not try to "see the future," as it were. Our analysis illustrates that a favorable trade-off may result between the extent of translation k and the fraction of the performance of full wavelength translation that is achievable by using limited wavelength translation. Furthermore, as is evident from the results presented in Fig. 2, for a large range of network loads, limited translation of relatively small degree suffices to give performance comparable to that obtained by full wavelength translation. In Section V, we will present a simple switch architecture capable of performing

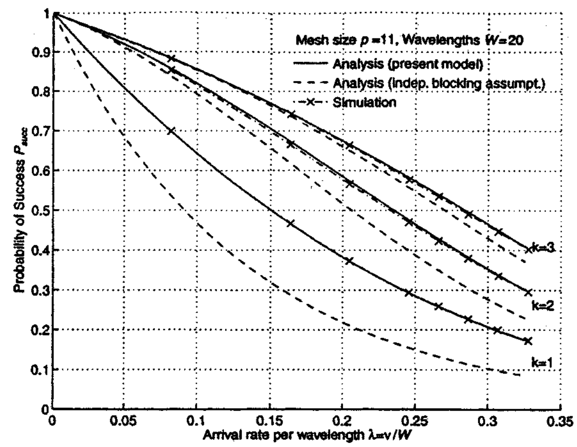

Fig. 3. Comparing the probability of success $P_{\rm succ}$ for the torus network, obtained by using the link independence blocking assumption and the present analysis. For small degree networks like the mesh the independence assumption is seen to break down at larger loads, where as the present analysis accurately captures network performance over the entire range of loads.

limited wavelength translation at a cost that is effectively independent of W.

In Fig. 3 we compare the probability of success obtained via our analytical and simulation results with that obtained by using the link independence blocking assumption, for the torus network. As is evident from the plots, the present analysis is very accurate for a large range of networks loads, including heavy loads; the independence blocking assumption, however, becomes inaccurate at heavy loads. The significant difference for the mesh between the present analysis and an analysis using the independence blocking assumption can be attributed to the small node degree and large diameter (and mean internodal distance) of the mesh. A small degree leads to less mixing of traffic and therefore a nonnegligible dependence between the probability of acquisition of successive wavelengths on the path followed by the setup packet of a session. A large diameter tends to amplify the inaccuracy in the probability of acquisition of each link, since some paths are long. By contrast, as we will see in Section IV-B, the large node degree of the hypercube and its smaller diameter (and average hop length) enable the independence blocking assumption also to predict the performance accurately in that case, though not as well as the present model. Similar observations about the accuracy of the independence blocking model for hypercube networks vis-a-vis torus networks were also made in [14].

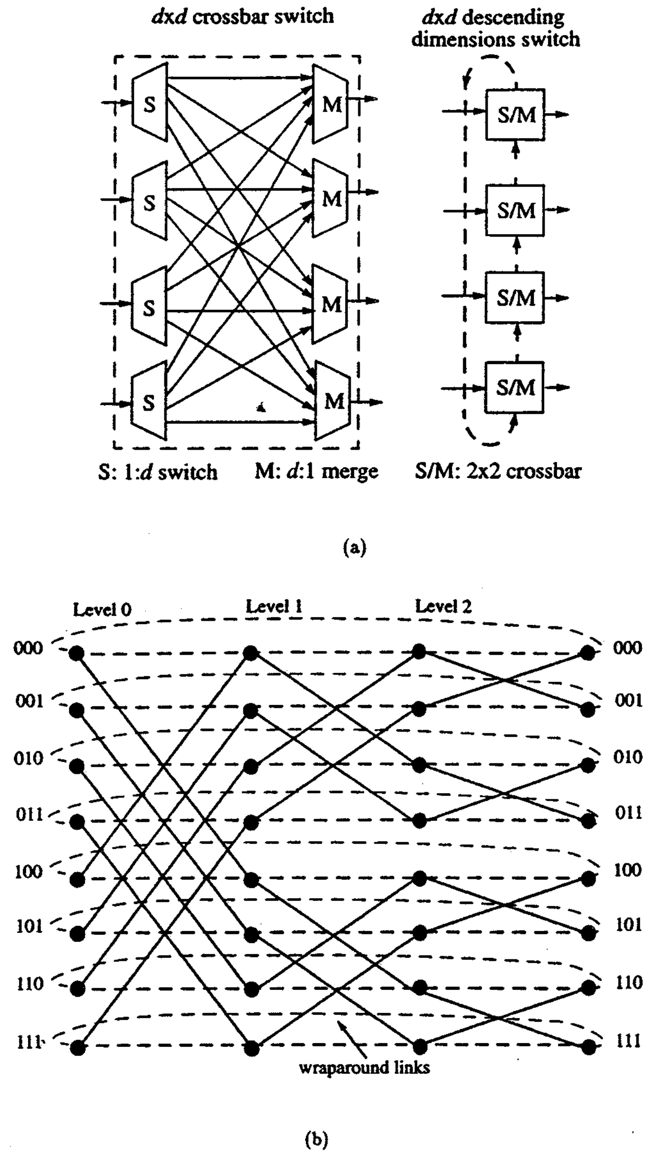

Fig. 4. (a) A $d \times d$ crossbar switch, a $d \times d$ descending-dimensions switch, and the modules out of which they are constructed. (b) A wrapped butterfly. The straight links of the wrapped butterfly correspond to the internal links in a descending dimensions switch, which we have shown with dashed lines above. The cross links of the wrapped butterfly correspond to the external links of the descending dimensions switch.

IV. Limited Wavelength Translation in the Hypercube

In this section, we focus on limited wavelength translation, using $k = 2$-adjacent wavelength switching, in a 2^d-node hypercube network. In Section IV-A we consider a hypercube network that uses simple *descending dimensions* switches, while in Section IV-B we consider a hypercube network with crossbar switches.

A. Hypercube with Descending Dimensions Switch

For the 2^d-node hypercube network, the descending dimensions switch (see Fig. 4) uses $\Theta(d)$ wires as opposed to the $\Theta(d^2)$ wires of a crossbar switch, and is simpler, faster, and cheaper than a crossbar switch ([18]). Furthermore, it uses simple 2 : 2 switch/merge modules, which also have the advantage of being the basic building blocks for several wavelength routing networks (such as wavelength-interchanging cross connect networks (WIXC); [1]).

The nodes and the links (external and internal; see Fig. 4) of a 2^d-node hypercube network with descending dimensions switches map in a 1-1 fashion to the nodes and links of a $d2^d$-node *wrapped butterfly* network [10]. We will henceforth refer to them simply as the wrapped butterfly. The $d2^d$-node wrapped butterfly network has d levels and 2^d nodes per level and is obtained by merging the first and the last levels of an ordinary butterfly [Fig. 4(b)]. All nodes of the wrapped butterfly are seen as potential sources and destinations; however, sessions have a destination at the same level with the source, so a circuit always traverses exactly d hops. The analysis presented here for the wrapped butterfly applies (with simple modifications) to other regular networks with two input links and two output links per node. It is therefore applicable to other multistage banyan networks like the perfect shuffle and Debruijn networks that are important choices for wavelength routed networks [12], [15].

The analysis for the wrapped butterfly parallels the development for the wraparound mesh network, which we have presented in detail in Section II. In the wrapped butterfly network, each session always uses exactly d hops to reach its destination, so the probability of success is independent of the routing tag, and may be written directly as

$$P_{\text{succ}} = q_0 \cdot \alpha^{d-1} \tag{16}$$

where α and q_0 can be calculated in a way similar to that in Section II. The detailed calculations appear in [16].

Fig. 5(a) illustrates analytical and simulation results for the probability of success P_{succ} for a 64-node butterfly network, with $W = 8$ wavelengths per link, for various k. Fig. 5(b) illustrates how the probability P_{succ} obtained by using the link independence blocking assumption differs from that obtained via the present analysis, which is very accurate for a large range of network loads. The accuracy of the independence blocking assumption for the butterfly network is somewhat better than that for the mesh for the case $k = 1$, and 2, and is quite good for the case $k = 3$. This is because for the butterfly network there is a tradeoff between the opposing effects of its node degree (which being small, accentuates the difference between the actual performance and that predicted by the independence assumption) and its diameter (which being small, mitigates this difference).

B. Hypercube with Cross-Bar Switches

In this section, we examine 2-adjacent wavelength switching in a 2^d-node hypercube network with cross-bar switches at the nodes. While the wrapped butterfly network, analyzed in the previous section, is an example of a sparse topology, the hypercube network, considered here, is an example of a dense topology, with an average hop length about half that in a butterfly network of the same size.

The relationship between the incoming and outgoing wavelengths for the hypercube network, is similar to that illustrated

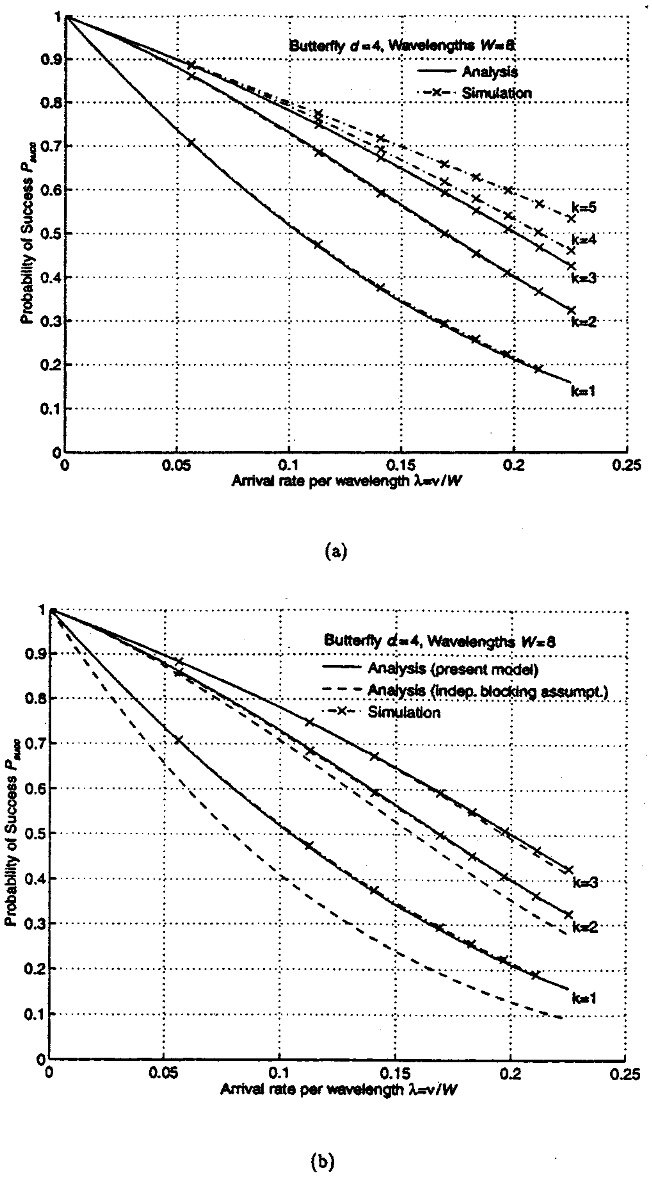

Fig. 5. (a) Probability of successfully establishing a connection P_{succ} versus the arrival rate per wavelength λ, for a butterfly of 64 nodes, for $W = 8$ wavelengths per link and for various k. (b) The probability P_{succ} as obtained by the independence blocking assumption and the present analysis.

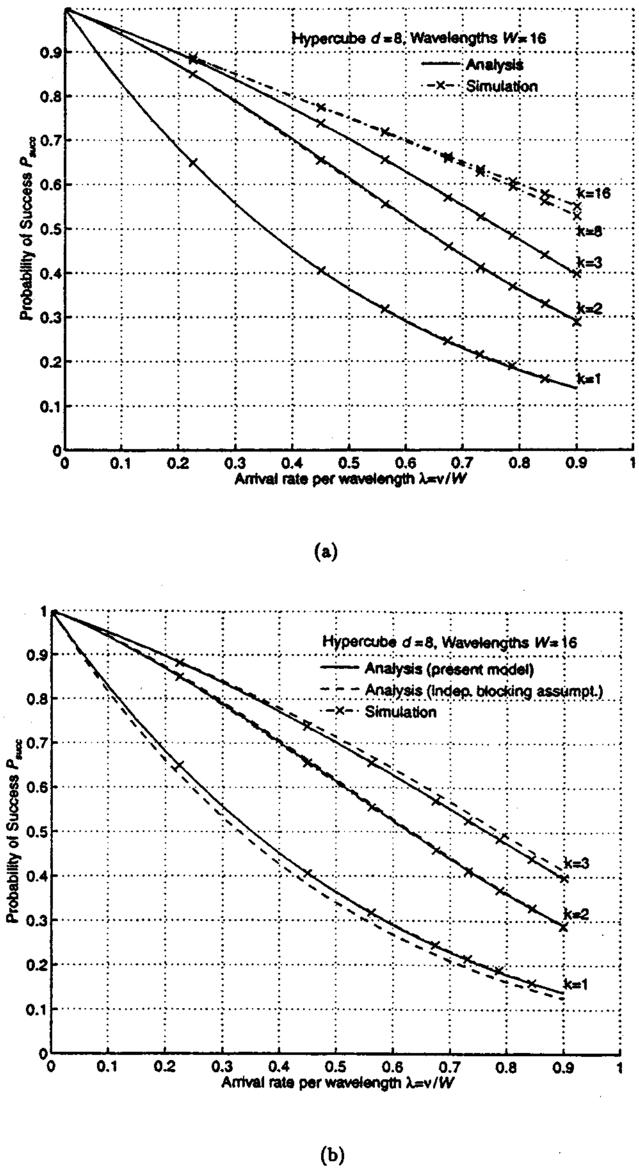

Fig. 6. Probability of successfully establishing a connection P_{succ} versus the arrival rate per wavelength λ, for an eight-dimensional (8-D) hypercube, with crossbar switches, hypercube of dimension $d = 8$, for $W = 8$ wavelengths per link and various values of k.

in Fig. 1, except that each outgoing link is fed by $d-1$ incoming links. For the hypercube network too, the average probability of success can be calculated to be [16]

$$\begin{aligned} P_{\text{success}} &= \frac{1}{2^d - 1} \sum_{t_{d-1}t_{d-2}\cdots t_0} P_{\text{success}}(t_{d-1}t_{d-2}\cdots t_0) \\ &= \frac{q_0}{\alpha(2^d-1)}\left[(1+\alpha)^d - 1\right] \end{aligned} \tag{17}$$

where α and q_0 can be calculated in a way similar to that in Section II, and appear in [16].

Fig. 6(a) illustrates the probability of success P_{succ} for an 8-D hypercube, having $W = 8$ wavelengths per link, for various k, while Fig. 6(b) illustrates how the probability of success P_{succ} obtained by using the link independence blocking assumption differs from that obtained via the present analysis.

V. Implementation of a Limited Wavelength Translation Switch

In this section, we provide a simple, economical implementation of a switch capable of performing up to $k = 3$-adjacent wavelength switching. We will see shortly that the cost of our switch depends mainly on the extent k of translation and is largely independent of the total number W of wavelengths in the system; a significant fact, considering that we have already

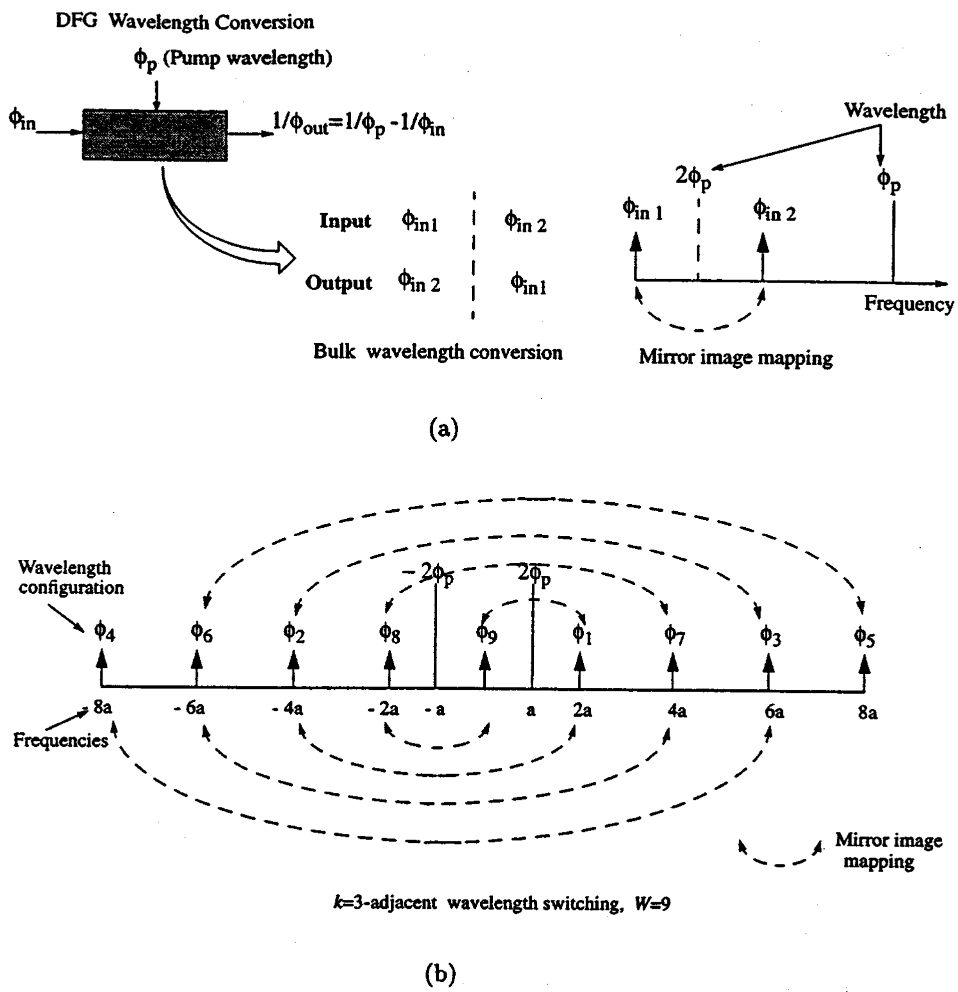

Fig. 7. (a) Parametric *difference frequency generation* (DFG) wavelength convertor. The DFG convertor can perform bulk wavelength conversions. (b) The arrangements of the wavelengths and the corresponding frequency values for realizing $k = 3$-adjacent wavelength switching with $W = 9$ wavelengths.

shown that the performance of limited wavelength translation compares favorably to that of full-wavelength translation.

The architecture that we present makes use of the *difference frequency generation* (DFG) wavelength convertors of Yoo *et al.* [1], which utilize the parametric wavelength conversion process to realize *bulk conversions,* that is, to perform several wavelength translations at once. In particular, if we consider the DFG convertor as a black-box [Fig. 7(a)] that is pumped at a pump wavelength ϕ_p, the incoming signal on wavelength $\phi_{\rm in1}$ is converted to a wavelegth $\phi_{\rm out1}$, where $1/\phi_{\rm out1} = 1/\phi_p - 1/\phi_{\rm in1}$. The bulk conversion property implies that a second input signal at a wavelength $\phi_{\rm in2}$ is simultaneously converted to another wavelength $\phi_{\rm out2}$, where $1/\phi_{\rm out2} = 1/\phi_p - 1/\phi_{\rm in2}$. As shown in the figure, when $\phi_{\rm out1} = \phi_{\rm in2}$ and $\phi_{\rm out2} = \phi_{\rm in1}$, the conversion process is equivalent to interchanging $\phi_{\rm in1}$ and $\phi_{\rm in2}$ about a virtual wavelength $2\phi_p$, whose frequency is set to be in the middle of the corresponding frequencies of $\phi_{\rm in1}$ and $\phi_{\rm in2}$; this is the *mirror image mapping* property of the DFG wavelength convertor.

The bulk conversion and mirror-image mapping properties can be exploited to realize a $k = 3$-adjacent wavelength switching system with W wavelengths, which is shown in Fig. 8 [frequencies in the horizontal axis of Fig. 8(b) are relative to a given freqency]. If the wavelengths are arranged as shown in Fig. 7(b) (for $W = 9$) and their corresponding frequencies are assigned the values shown (where a is some constant), it can be seen that the mirror-image mapping about the two pump wavelengths ϕ_p and $-\phi_p$ enables each wavelength ϕ_i, $i = 1, 2, \cdots, 9$ (with the exception of ϕ_4 and ϕ_5) on an incoming link to be mapped on the outgoing link either to itself, or to the wavelengths ϕ_{i+1} or ϕ_{i-1} that lie above or below it, respectively. To interchange wavelengths ϕ_4 and ϕ_5, we require a separate single-pump wavelength converter, or two variable-input fixed-output wavelength converters (see Fig. 8). Although we have depicted the case $W = 9$, it is easy to see that the pattern of Fig. 7(b) can be extended for any (odd) W (a parallel implementation exists for even W), due to the bulk conversion property, *without* any corresponding increase in the complexity of the wavelength convertors of the switch. (The complexity of the multiplexers and demultiplexers depends on W, but that is essentially the same for either a full-or-limited-wavelength translation system.) Thus, the cost of our switch (which requires only three single-pump DFG wavelength convertors; alternately, one double pump [1] and

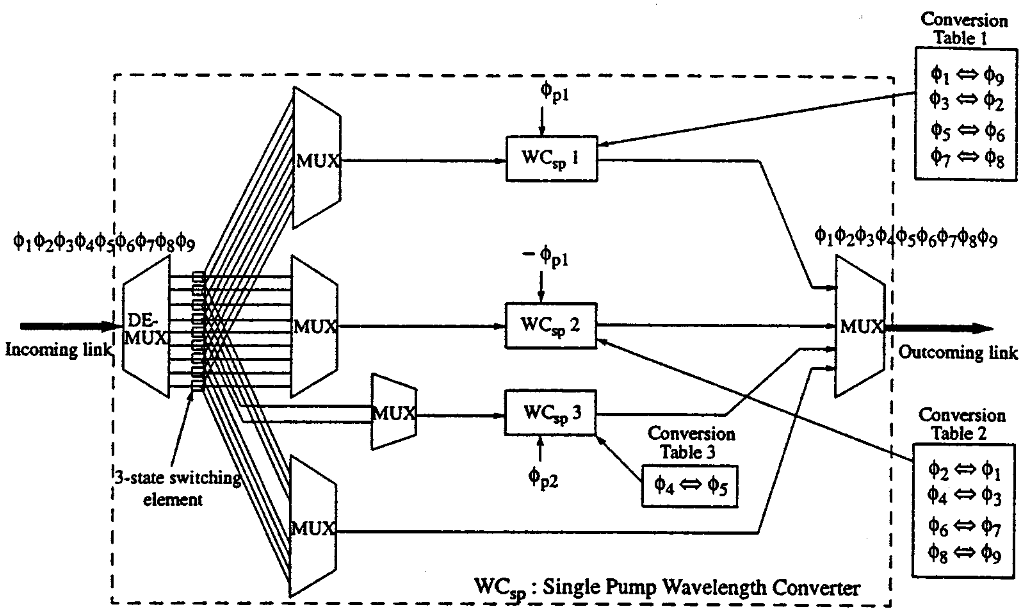

Fig. 8. Switch architecture for performing $k = 3$-adjacent wavelength switching. The cost of the switch is essentially a constant (three single-pump bulk wavelength converters), which is independent of the number of wavelengths W in the system.

one single pump wavelength convertor), is largely independent of the total number of wavelengths W per link or fiber. The 3-state switching element for each wavelength is set at connection setup time, depending on the particular translation desired. We observe finally that the same architecture can also effect $k = 2$-adjacent wavelength switching. (We do not present a separate architecture for the case $k = 2$, since its complexity turns out to be close to that of the architecture just given for the case $k = 3$.)

VI. Conclusion

We examined the case of limited wavelength translation in several wavelength routed, all-optical, WDM regular networks, and demonstrated that limited wavelength translation of fairly small degree is sufficient to obtain benefits comparable to those obtained by full wavelength translation. An important conclusion of our analysis was that the *efficiency* with which capacity is used depends on the extent of translation k and not on the total number of wavelengths W per fiber, as long as $W \geq k$ (equivalently, the throughput of the network for a given blocking probability increases with k, and increases only linearly, and not superlinearly, with the number of wavelengths W).

Increasing W with k fixed does not increase the probability of success P_{succ} for the same value $\lambda = \nu/W$ of admissible throughput per node per wavelength. This seems at first view to contrast with the observation made in [14], where the authors had found that for sparse wavelength conversion, converters help more (in improving efficiency) the more wavelengths there are per fiber. The model used in [14], however, essentially relies on backtracking or trying multiple options simultaneously, which may not be a feasible option in WAN's with large propagation delays, as discussed in Section III.

Our results also indicate that the network performance improves as the extent of translation k increases, but that the rate of improvement decreases with increasing k, and becomes negligible when k is equal to the diameter of the network.

We also presented the design of a simple switch architecture for performing $k = 2$, 3-adjacent wavelength switching. Our switch uses DFG bulk wavelength converters, which unlike four-wave mixing converters do not create additional cross terms that interfere with other WDM channels [1]. The cost of our switch was seen to be a function only of the extent k of translation.

Our analysis also answers partially a question raised in [16], namely, would the throughput of a circuit-switched regular network with k $(k > 1)$, circuits per link be expected to improve by more than just a factor of k? Since in our case the success probability depends on the extent of translation k and is independent of the number of wavelengths W per link ($W \geq k$), we can also view the P_{succ} curves in Figs. 4–6 for $k = 2, 3$ as representing the probability of successfully establishing a connection in a circuit switched network with a capacity of k circuits per link. We observe that with $k = 2, 3$ circuits per link, the improvement in throughput over a network with a capacity of only one circuit per link is considerably greater than by just a factor of k. (If the improvement obtained was only a factor of k, the curves for $k = 2$ and $k = 3$ would overlap the curves for $k = 1$, which is not the case.) However, increasing k further, say, from 2 to 3, gives diminishing returns. This observation agrees with the result obtained by Koch [8], who proved that the maximum total throughput of a $d2^d$-node butterfly with a capacity of k circuits per link is $\Theta(2^d/d^{1/k})$, and, therefore, that the increase in throughput (or efficiency) obtainable by using links with larger capacity diminishes as the capacity k per link increases.

Finally, our analysis was based on the use of fixed shortest path routing to establish a connection. The effect of alternate or dynamic routing (or deflection routing) when combined with limited wavelength translation is a possible area for future work. Furthermore, our focus was on a circuit-switched optical path layer. Castanon *et al.* [6] have recently looked at the performance of full wavelength translation in packet-switched network; a corresponding analysis for limited wavelength translation in packet-switched networks remains to be done. Also, the success probability is only one measure of network performance. A study that looks at other measures of network performance, such as the input queuing and connection delays, to evaluate the benefits of limited wavelength translation, in both circuit-switched and packet-switched networks, would also be useful.

APPENDIX
DERIVATION OF WAVELENGTH STATE PROBABILITIES

A. Probabilities q_i, $i = 1, 2, 3$ for the Torus

The probability q_3 that a wavelength on an outgoing link L is used by a turning session is given by

$$q_3 = K \sum_{\substack{t_2=-\lfloor p/2 \rfloor \\ t_2 \neq 0}}^{\lfloor p/2 \rfloor} \sum_{\substack{t_1=-\lfloor p/2 \rfloor \\ t_1 \neq 0}}^{\lfloor p/2 \rfloor} 1 = K(p-1)^2. \tag{A.1}$$

Similarly, the probability q_2 that a wavelength on an outgoing link L is used by a straight-through session is given by

$$\begin{aligned} q_2 &= K \left[\sum_{\substack{t_2=-\lfloor p/2 \rfloor \\ t_2 \neq 0}}^{\lfloor p/2 \rfloor} \sum_{\substack{t_1=-\lfloor p/2 \rfloor \\ t_1 \neq 0}}^{\lfloor p/2 \rfloor} (|t_2| + |t_1| - 2) \right. \\ &\qquad \left. + \binom{2}{1} \sum_{\substack{t_2=-\lfloor p/2 \rfloor \\ t_2 \neq 0}}^{\lfloor p/2 \rfloor} (|t_2| - 1) \right] \\ &= K \cdot \left\lceil \frac{p-2}{2} \right\rceil \left\lfloor \frac{p-2}{2} \right\rfloor. \end{aligned} \tag{A.2}$$

The remaining probability q_1 that a wavelength on link L is used by a session originating at a node is given by

$$q_1 = K(p^2 - 1). \tag{A.3}$$

Finally, a simple application of Little's Theorem gives

$$1 - q_0 = \frac{\nu \overline{X} h}{4W}. \tag{A.4}$$

Equations (A.1)–(A.4) together with the condition $\sum_{i=0}^{3} q_i$ can be used to calculate K, so that we finally get

$$q_1 = \frac{\nu \overline{X} h}{4W} \cdot \frac{p^2 - 1}{p \cdot \left\lceil \frac{p^2-1}{2} \right\rceil} \tag{A.5a}$$

$$q_2 = \frac{\nu \overline{X} h}{4W} \cdot \frac{\left\lceil \frac{p-2}{2} \right\rceil \left\lfloor \frac{p-2}{2} \right\rfloor}{p \cdot \left\lceil \frac{p^2-1}{2} \right\rceil} \tag{A.5b}$$

and

$$q_3 = \frac{\nu \overline{X} h}{4W} \cdot \frac{(p-1)^2}{p \cdot \left\lceil \frac{p^2-1}{2} \right\rceil} \tag{A.5c}$$

where $\lfloor x \rfloor$ is the largest integer less than or equal to x, and $\lceil x \rceil$ is the smallest integer greater than or equal to x. Equations (A.5a)–(A.5c) hold for both p odd and p even.

REFERENCES

[1] N. Antoniades, K. Bala, and S. J. B. Yoo *et al.*, "A parametric wavelength interchanging cross-connect architecture," *IEEE Photon. Technol. Lett.*, vol. 8, pp. 1382–1384, Oct. 1996.

[2] R. A. Barry and P. Humblet, "Models of blocking probability in all-optical networks with and without wavelength changers," *IEEE J. Select. Areas Commun.*, vol. 14, pp. 858–867, June 1996.

[3] J. M. H. Elmirghani and H. T. Mouftah, "All-optical wavelength conversion: Technologies and applications in DWDM networks," *IEEE Commun. Mag.*, vol. 38, pp. 86–92, Mar. 2000.

[4] S. Baroni, P. Bayvel, and J. E. Midwinter, "Wavelength requirements in dense wavelength-routed optical transport networks with variable physical connectivity," *Electron. Lett.*, vol. 32, pp. 575–576, Mar. 1996.

[5] R. A. Barry, V. W. S. Chan, and K. L. Hall *et al.*, "All-Optical Network Consortium-ultrafast TDM networks," *IEEE J. Select. Areas Commun.*, vol. 14, pp. 999–1101, June 1996.

[6] G. A. Castanon, O. K. Tonguz, and A. Bononi, "Benefits of wavelength translation in datagram networks," *Electron. Lett.*, vol. 33, pp. 1567–1568, Aug. 1997.

[7] H. Harai, M. Murata, and H. Miyahara, "Performance analysis of wavelength assignment policies in all-optical networks with limited-range wavelength conversion," *IEEE J. Select. Areas Commun.*, vol. 16, pp. 1051–1060, Sept. 1998.

[8] R. Koch, "Increasing the size of a network by a constant factor can increase performance by more than a constant factor," in *29th Annu. Symp. Founds. Comp. Sci.*, Oct. 1988, pp. 221–230.

[9] M. Kovačević and A. Acampora, "Benefits of wavelength translation in all-optical clear-channel networks," *IEEE J. Select. Areas Commun.*, vol. 14, pp. 868–880, June 1996.

[10] F. T. Leighton, *Introduction to Parallel Algorithms and Architectures—Arrays, Trees, Hypercubes*. San Mateo, CA: Morgan Kaufman.

[11] Y. Mei and C. Qiao, "Efficient distributed control protocols for WDM all-optical networks," in *Proc. 6th Int. Conf. Computer Commun. and Networks*, Las Vegas, NV, Sept. 1997, pp. 150–153.

[12] R. K. Pankaj and R. G. Gallager, "Wavelength requirements of all-optical networks," *IEEE ACM Trans. Networking*, vol. 3, no. 3, pp. 269–279, June 1995.

[13] R. Ramaswami and G. H. Sasaki, "Multiwavelength optical networks with limited wavelength conversion," in *Proc. INFOCOM'97*, Kobe, Japan, Apr. 1997.

[14] S. Subramaniam, M. Azizoglu, and A. K. Somani, "All-optical networks with sparse wavelength conversion," *IEEE/ACM Trans. Networking*, vol. 4, no. 4, pp. 544–557, Aug. 1996.

[15] K. N. Sivarajan and R. Ramaswami, "Lightwave networks based on de bruijn graphs," *IEEE/ACM Trans. Networking*, vol. 2, no. 1, pp. 70–79, Feb. 1994.

[16] V. Sharma, "Efficient communication protocols and performance analysis for gigabit networks," Ph.D. dissertation, Department ECE, Univ. of California, Santa Barbara, Aug. 1997.

[17] T. Tripathi and K. Sivarajan, "Computing approximate blocking probabilities in wavelength routed all-optical networks with limited range wavelength conversion," in *Proc. Infocom'99*, vol. 1, Mar. 1999, pp. 329–336.

[18] E. A. Varvarigos and D. Bertsekas, "Performance of hypercube routing schemes with and without buffering," *IEEE/ACM Trans. Networking*, vol. 2, no. 3, pp. 299–311, June 1994.

[19] E. A. Varvarigos and V. Sharma, "An efficient reservation virtual circuit protocol for gigabit networks," *Comput. Netw. and ISDN Syst.*, vol. 30, no. 12, pp. 1135–1156, July 1998.

[20] J. Yates, J. Lacey, D. Everitt, and M. Summerfield, "Limited-range wavelength translation in all-optical networks," in *Proc. IEEE INFOCOM'96*, vol. 3, Mar. 1996, pp. 954–961.

[21] N. Wauters and P. Demeester, "Wavelength translation in optical multi-wavelength multi-fiber transport networks," *Int. J. Optoelectronics*, vol. 11, no. 1, pp. 53–70, Jan.–Feb. 1997.
[22] J. Zhou, N. Park, and K. Vahala *et al.*, "Four-wave mixing wavelength conversion efficiency in semiconductor traveling-wave amplifiers measured to 65 nm of wavelength shift," *IEEE Photon. Technol. Lett.*, vol. 6, pp. 984–987, Aug. 1994.
[23] O. Gerstel, R. Ramaswami, and G. H. Sasaki, "Dynamic channel assignment for WDM optical networks with little or no wavelength conversion," in *Proc. 34th Annual Allerton Conf. on Commun., Control, and Signal Processing*, Monticelo, IL, Oct. 1996, pp. 32–43.

Vishal Sharma (M'98) received the B.Tech degree in electrical engineering from the Indian Institute of Technology, Kanpur, in 1991. He received the M.S. degrees in signals and systems and computer engineering both in 1993 and the Ph.D. degree in electrical and computer engineering from the University of California Santa Barbara in 1997.

From 1998 to 2000, he was with the Tellabs Research Center , Cambridge, MA, where he worked on the architecture and design of high-speed backbone routers, models for IP quality of service, and extensions to multi-protocol label switching (MPLS) for survivability, and optical and TDM networks. He is now Principal Architect with Jasmine Networks, Inc., San Jose, CA. His current research interests include the design and analysis of high-speed communication networks, models for service differentiation in IP networks, and dynamic optical networks.

Dr. Sharma is a member of the IEEE Computer and Professional Communication Societies and the ACM Sigcomm. He is active in the MPLS and Traffic Engineering Working Groups of the IETF.

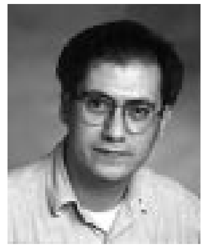

Emmanouel A. Varvarigos (M'00) received the Dipl. Eng. from the National Technical University of Athens, Athens, Greece, in 1988, the M.S. degree in electrical engineering in 1991, and the Ph.D. degree in electrical engineering and computer science in 1992, both from the Massachusetts Institute of Technology, Cambridge, MA.

In 1990, he was with Bell Communications Research, Morristown, NJ. In 1992, he joined the Department of Electrical and Computer Engineering, University of California, Santa Barbara, where he is currently an Associate Professor. In 1998–1999, during his sabbatical, he was a Visiting Associate Professor at the Technical University of Delft, The Netherlands, and the University of Patras, Greece. His research interests are in the areas of algorithms and protocols for data networks, performance evaluation, all-optical networks, parallel and distributed computation, and mobile communications.

Dr. Varvarigos received the First Panhellenic Prize in the Greek Mathamatical Olympiad in 1982, the Technical Chamber of Greece Award four times (1984-1988), and the National Science Foundation Research Initiation Award in 1993. He was Co-Organizer of the 1996 Workshop on Advanced Communication Systems and Networking.