All-Optical Networks with Sparse Wavelength Conversion

Suresh Subramaniam, *Student Member, IEEE*, Murat Azizoğlu, *Member, IEEE*, and Arun K. Somani, *Senior Member, IEEE*

***Abstract*—We study the effects of topological connectivity and wavelength conversion in circuit-switched all-optical wavelength-routing networks. A blocking analysis of such networks is given. We first propose an analytical framework for accurate analysis of networks with arbitrary topology. We then introduce a model for networks with a variable number of converters and analyze the effect of wavelength converter density on blocking probability. This framework is applied to three regular network topologies that have varying levels of connectivity: the ring, the mesh-torus, and the hypercube. The results show that either a relatively small number of converters is sufficient for a certain level of performance or that conversion does not offer a significant advantage. The benefits of conversion are largely dependent on the network load, the number of available wavelengths, and the connectivity of the network. Finally, the tradeoff between physical connectivity, wavelength conversion, and the number of available wavelengths is studied through networks with random topologies.**

I. Introduction

Wavelength-division multiplexing (WDM), used in conjunction with wavelength-routing, is a promising mechanism for information transport in future all-optical networks (AON's) [1]–[3]. Unlike broadcast-and-select networks, wavelength-routing networks offer the advantages of wavelength reuse and scalability and are thus suitable for wide-area networks (WAN's) [4]. These networks consist of wavelength-routing nodes interconnected by optical fibers. A wavelength-routing node is one that is capable of switching a signal dynamically, based on the input port and the wavelength on which the signal arrives [5]–[7]. Such a node can be implemented by a set of wavelength multiplexers and demultiplexers and a set of photonic switches (one per wavelength).

We consider circuit-switched networks. Call requests arrive at a node according to a random point process, and an optical circuit is established between the source and destination for the (random) duration of the call. Two different functionalities of the wavelength-routing nodes are important in this context: nodes with no wavelength conversion capability and nodes which can convert an incoming wavelength to an arbitrary outgoing wavelength. When all nodes have wavelength converters, the situation is analogous to trunk switching in digital telephony [8], as a call arriving on one trunk (wavelength) can be switched to any outgoing trunk, so long as one is available. On the other hand, in a network without any wavelength converting nodes, a call arriving at a certain wavelength on an input fiber has to be switched to an output fiber at the same wavelength. This requirement of wavelength continuity increases the probability of call blocking; to honor a call request, it is necessary that the *same* wavelength be free on all the links of the circuit.

There has been considerable interest in obtaining the call blocking performance of wavelength-routing networks [9]–[11]. The performance improvement with wavelength converters is of fundamental importance to quantify. This improvement depends on the topology of the network, the traffic demand, and the number of available wavelengths, among other factors. As the network becomes denser, one would expect the usefulness of converters to decrease, since the paths get shorter. In the limiting case with a link between every node pair, wavelength converters have no effect on the blocking performance, since all sessions are one-hop sessions.[1] On the other hand, a sparsely connected network tends not to mix calls well and thus causes a load correlation in successive links. This reduces the usefulness of wavelength converters [9]. The benefits of conversion are thus largely dependent on which of the above two effects dominates.

The analytical models proposed in the literature have considered networks in which there are no converters and those in which all nodes have converters. All-optical wavelength converters are being prototyped in research laboratories [12], and are likely to remain costly devices. Therefore, a more practical situation is one in which wavelength conversion capability is available in a relatively small fraction of nodes. We refer to such a network as one with *sparse wavelength conversion*. In [13], a node architecture that can perform a limited number of wavelength conversions is proposed, and a heuristic algorithm for dynamic routing is presented for

Manuscript received January 7, 1996; revised April 16, 1996; approved by IEEE/ACM Transactions on Networking Editor B. Mukherjee. This work was supported by the University of Washington Royalty Research Fund, and the National Science Foundation under Grants NCR-9309574, NCR-9502610, and NCR-9103485. An earlier version of this paper was presented at the IEEE INFOCOM'96 Conference, San Francisco, CA, March 26–28, 1996. The paper was selected by the Conference as one of its top papers and referred to the Transactions for possible publication after the Transactions' own independent review.

S. Subramaniam and M. Azizoğlu are with the Department of Electrical Engineering, University of Washington, Seattle, WA 98195 USA (email: suresh@ee.washington.edu).

A. K. Somani is with the Department of Electrical Engineering and the Department of Computer Science and Engineering, University of Washington, Seattle, WA 98195 USA.

Publisher Item Identifier S 1063-6692(96)06079-7.

[1] This assumes that the direct link is always used in this case. If alternate path routing is allowed, wavelength converters may still be of some benefit.

reducing the number of conversions. In this paper, we consider *converters* that can permute any set of incoming wavelengths to any set of outgoing wavelengths, and develop a model to analyze the effect of sparse wavelength conversion on the blocking probability. Specifically, we address the following question: For a given network topology, how many converters, if any, are needed to achieve a desired performance? Our analysis indicates that, in most cases, either a relatively small fraction of the nodes has to be equipped with wavelength conversion capability for good performance, or conversion does not offer great advantages.

The literature on the blocking analysis of networks with full or no wavelength conversion points out the difficulty in accurate accounting of load correlation. In [10], a model to compute the approximate blocking probability with Poisson traffic input is presented. As pointed out in [10], that model is inappropriate for networks with sparse topologies because it does not consider the correlation of wavelength use between successive links of a path. In [9], a new model that takes into account this dependence is presented. However, it is assumed that a wavelength is used on a link with a fixed probability independently of the other wavelengths, and therefore, the dynamic nature of the traffic is hidden. Another model presented in [11] considers Poisson input traffic and uses a Markov chain model with state-dependent arrival rates. It is more accurate, but the computation of blocking probabilities is very intensive, and the analysis is tractable only for networks with a few nodes.

In this paper, we present an analytical model of modest complexity for evaluating the blocking performance of networks with sparse wavelength conversion.[2] The model is shown to be accurate for a variety of network topologies by comparing the analytical results with those obtained using simulation. In Section II, we present a model for computing the blocking performance of networks without wavelength converters. The model takes into account the correlation of wavelength use on successive links. Section III extends this model to incorporate the effect of sparse wavelength conversion for an arbitrary topological connectivity. We then study the blocking performance of ring and mesh-torus networks as examples of sparse topologies in Sections IV and V, respectively, and the hypercube as an example of a dense topology in Section VI. We notice that there is an excellent match between the analytical and simulation results even when the topologies are sparse. The results of the study are counterintuitive and the effect of wavelength conversion on the blocking performance of a given network topology cannot be predicted *a priori* without an accurate analysis.

Wavelength-routing is most likely to be used in WAN's where a broadcast-and-select approach to switching ceases to be feasible [14]. The topology of such a network typically evolves into an irregular physical topology with arbitrary connection patterns. To model this practical situation, we consider random topologies in Section VII. We evaluate the ensemble average of the blocking probability and study the effects of connectivity and wavelength conversion. Our conclusions are presented in Section VIII.

[2] Our model includes networks without wavelength converters and networks with full wavelength conversion as special cases.

II. An Analytical Model for Path Blocking Performance

In this section, we develop an analytical model that has modest computational requirements and improves upon the previously proposed models of [9] and [10] by considering real-time input traffic and by incorporating the correlation between the wavelengths used on successive links of a multilink path. We first assume that there is no wavelength conversion. In Section III, we include the effect of wavelength conversion.

A. Model Assumptions

The following assumptions are used in our analytical model.

1) Call requests arrive at each node according to a Poisson process with rate λ. Each call is equally likely to be destined to any of the remaining nodes.
2) Call holding time is exponentially distributed with mean $1/\mu$; the offered load per station is $\rho = \lambda/\mu$.
3) The path used by a call is chosen according to a prespecified criterion (e.g., random selection of a shortest path), and does not depend on the state of the links that make up a path. The call is blocked if the chosen path cannot accommodate it. Alternate path routing is not allowed.
4) The number of wavelengths, F, is the same on all links. Each node is capable of transmitting and receiving on any of the F wavelengths. Each call requires a full wavelength on each link it traverses.
5) Wavelengths are assigned to a session randomly from the set of free wavelengths on the associated path.[3]

First, let us define a wavelength as "free" on a path if that wavelength is not used on any of the links constituting the path. A wavelength is "busy" on a path otherwise. The model in [10], henceforth called the *independence model*, assumes that the link loads are independent *and* that the wavelengths used on a link are uniformly distributed over the entire set of wavelengths, independently of all other links. These two assumptions result in an overestimation of the blocking probability, as shown in Section IV. The performance estimate is very crude (an inaccuracy of about two orders of magnitude) for sparsely connected networks such as the ring, and gets better with increasing connectivity. In this paper, we propose a model that can be used to estimate the performance accurately even for sparse networks.

B. Notation

We define the following probabilities[4] that will be used in obtaining the blocking probabilities.

- $Q(w_f) = \Pr\{w_f$ wavelengths are free on a link$\}$.
- $S(y_f \mid x_{pf}) = \Pr\{y_f$ wavelengths are free on a link of a path $| x_{pf}$ wavelengths are free on the previous link of the path$\}$.

[3] Other heuristic wavelength allocation strategies may provide better performance [15], but are considerably more difficult to analyze.

[4] We have attempted to make the notation easier to follow by using the subscript f for *free*, c for *continuing*, pf for *free on previous*, and ff for *free on first*.

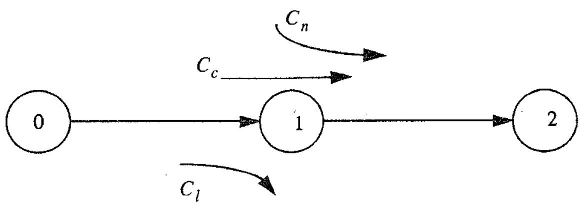

Fig. 1. Calls arriving and leaving on a two-hop path.

- $U(z_c \mid y_f, x_{pf}) = \Pr\{z_c$ calls (wavelengths) continue to the current link from the previous link of a path $|x_{pf}$ wavelengths are free on the previous link, and y_f wavelengths are free on the current link$\}$.
- $R(n_f \mid x_{ff}, y_f, z_c) = \Pr\{n_f$ wavelengths are free on a two-hop path $|x_{ff}$ wavelengths are free on the first hop of the path, y_f wavelengths are free on the second hop, and z_c calls continue from the first to the second hop$\}$.
- $T^{(l)}(n_f, y_f) = \Pr\{n_f$ wavelengths are free on an l-hop path and y_f wavelengths are free on hop $l\}$.
- $p_l = \Pr\{$an l-hop path is chosen for routing$\}$.

C. The Conditional Free Wavelength Distribution

In addition to the assumptions stated in Section II-A, we assume the following. The load on link i of a path given the loads on links $1, 2, \cdots, i-1$, depends only on the load on link $i-1$. We, therefore, call our analytical model as the (Markovian) *correlation model*. The analysis in this paper differs from the one we presented in [16], which neglected link load correlation. We start by considering a two-hop path and deriving the conditional free wavelength distribution on the path. In Section II-D, we extend the analysis to determine the blocking probability on a path of arbitrary hop length. To model the correlation between the loads on the two links, we employ a three-dimensional Markov chain as follows. Referring to Fig. 1, let C_l be the number of calls that enter the path at node 0 and leave at node 1, let C_c be the number of calls that enter the path at node 0 and continue on to the second link, and let C_n be the number of calls that enter the path at node 1. Therefore, the number of calls that use the first link is $C_l + C_c$ and the number of calls that use the second link is $C_c + C_n$.

Since the number of calls on a link cannot exceed the total number of available wavelengths, F, we have $C_l + C_c \leq F$, and $C_c + C_n \leq F$. Suppose the arrival rate of calls that enter at node 0 and leave at node 1 is λ_e, and the arrival rate of calls that enter at node 0 and continue on to node 2 is λ_c. Let the corresponding Erlang loads be denoted by $\rho_e = \lambda_e/\mu$ and $\rho_c = \lambda_c/\mu$, where $1/\mu$ is the expected value of the exponentially distributed call holding time. By the assumption of uniform traffic distribution, the arrival rate of calls that enter the path at node 1 is the same as the arrival rate of calls that leave the path at node 1. C_l, C_c, and C_n can therefore be characterized by a three-dimensional (3-D) Markov chain, with the state space as shown in Fig. 2. Each state is represented by an integer triplet (c_l, c_c, c_n) (the circles in Fig. 2).

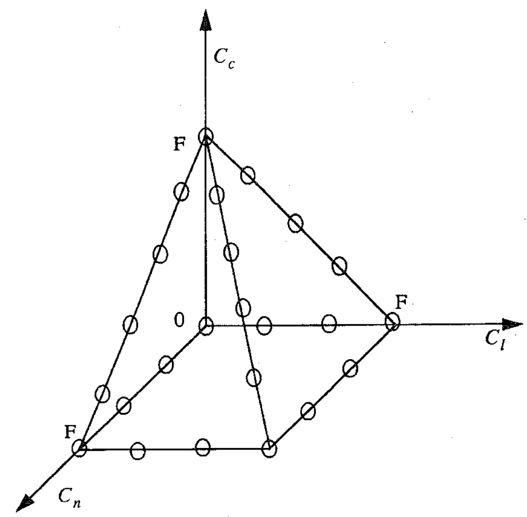

Fig. 2. The state space of the 3-D Markov chain.

We can now determine the steady-state probability of state (c_l, c_c, c_n) as [17]

$$\Pi(c_l, c_c, c_n) = \frac{\frac{\rho_e^{c_l}}{c_l!}\frac{\rho_c^{c_c}}{c_c!}\frac{\rho_e^{c_n}}{c_n!}}{\sum_{j=0}^{F}\sum_{i=0}^{F-j}\sum_{k=0}^{F-j}\frac{\rho_e^i}{i!}\frac{\rho_c^j}{j!}\frac{\rho_e^k}{k!}}, \quad \begin{array}{l} 0 \leq c_l + c_c \leq F \\ 0 \leq c_c + c_n \leq F \end{array}.$$

In the absence of wavelength converters, a call is blocked on a two-hop path if there is no free wavelength on the path. Now we compute the free wavelength distribution on a two-hop path. The conditional probability that n_f wavelengths are free on a two-hop path given x_{ff} wavelengths are free on the first link, y_f wavelengths are free on the second link, and z_c calls continue from the first link to the second link, is obtained by simple enumeration as

$$R(n_f \mid x_{ff}, y_f, z_c) = \frac{\binom{x_{ff}}{n_f}\binom{F - x_{ff} - z_c}{y_f - n_f}}{\binom{F - z_c}{y_f}}$$

for $\min(x_{ff}, y_f) \geq n_f \geq \max(0, x_{ff} + y_f + z_c - F)$, and is zero otherwise.

Recalling our notation in Section II-B, the probabilities $U(z_c \mid y_f, x_{pf})$, $S(y_f \mid x_{pf})$, and $Q(w_f)$ can be derived using Fig. 2 as follows:

$$\begin{aligned} &U(z_c \mid y_f, x_{pf}) \\ &= P(C_c = z_c \mid C_n + C_c = F - y_f, C_l + C_c = F - x_{pf}) \\ &= \frac{\Pi(F - x_{pf} - z_c, z_c, F - y_f - z_c)}{\sum_{x_c=0}^{\min(F - x_{pf}, F - y_f)} \Pi(F - x_{pf} - x_c, x_c, F - y_f - x_c)}, \\ &S(y_f \mid x_{pf}) \\ &= P(C_n + C_c = F - y_f \mid C_l + C_c = F - x_{pf}) \\ &= \frac{\sum_{x_c=0}^{\min(F - x_{pf}, F - y_f)} \Pi(F - x_{pf} - x_c, x_c, F - y_f - x_c)}{\sum_{x_c=0}^{F - x_{pf}} \sum_{x_n=0}^{F - x_c} \Pi(F - x_{pf} - x_c, x_c, x_n)} \end{aligned}$$

and

$$Q(w_f) = P(C_l + C_c = F - w_f) = \sum_{x_c=0}^{F-w_f} \sum_{x_n=0}^{F-x_c} \Pi(F - w_f - x_c, x_c, x_n). \quad (1)$$

Next, we use the conditional free wavelength distribution to obtain an expression for the blocking probability (in the absence of wavelength conversion) on a path of arbitrary hop length.

D. Blocking on a Multihop Path

Consider an l-hop path. A call is blocked if there is no free wavelength on the path at the time of call arrival. Suppose we know the joint probability, $T^{(l-1)}(x_{ff}, x_{pf})$, that there are x_{ff} free wavelengths on an $(l-1)$-hop path, and x_{pf} $(\geq x_{ff})$ wavelengths are free on the last hop of that path. Because of our assumption that the load on the lth hop is dependent only on the load on the $(l-1)$th hop, the probability of blocking on the l-hop path can be computed using the results for a two-hop path derived in Section II-C. This is done by viewing the first $l-1$ hops as the first hop and the lth hop as the second hop of a two-hop path. To complete the recursion, we need to determine the joint probability of n_f free wavelengths on the l-hop path and y_f free wavelengths on the lth hop.

Using the chain rule of probability, we have

$$T^{(l)}(n_f, y_f) = \sum_{x_{pf}=0}^{F} \sum_{x_{ff}=0}^{F} \sum_{z_c=0}^{\min(F-x_{pf}, F-y_f)} R(n_f \mid x_{ff}, z_c, y_f) \times U(z_c \mid y_f, x_{pf}) S(y_f \mid x_{pf}) T^{(l-1)}(x_{ff}, x_{pf}). \quad (2)$$

In writing the above, we have used the following facts.

i) The number of free wavelengths on the l-hop path is not dependent on the number of free wavelengths on hop $l-1$ when the number of free wavelengths on the path consisting of the first $l-1$ hops is given.
ii) The number of calls that continue from the first to the second link of a two-hop path depends only on the number of free wavelengths on each of the two hops.
iii) The number of free wavelengths on the second link of a two-hop path is dependent only on the number of free wavelengths on the first link.

The starting point of the recursion, $T^{(1)}(n_f, y_f)$ is zero when $n_f \neq y_f$, and is equal to the probability of having n_f free wavelengths on a link, $Q(n_f)$, given by (1), when $n_f = y_f$.

The probability of blocking on an l-hop path is simply the probability of finding no wavelength free on the l-hop path, and is, therefore, given by $\sum_{y_f=0}^{F} T^{(l)}(0, y_f)$. The network-wide blocking probability in the absence of converters is then computed as

$$P_b = \sum_{l=1}^{N-1} \sum_{y_f=0}^{F} T^{(l)}(0, y_f) p_l$$

for a network of N nodes.

E. Estimation of Parameters

The analysis in the last two subsections can be used to compute the call blocking probability in a network without wavelength converters. This analysis assumes that the hop-length distribution, p_l, and the arrival rates of calls at a link that continue on to the next link of a path and of those that do not, λ_c and λ_e respectively, are known. The hop-length distribution is a function of the topology and the routing algorithm and is easily determined for most regular topologies with shortest-path routing, as seen in the following sections.

Typically, the traffic in the network is specified in terms of the set of offered loads between station pairs. The call arrival rates at links have to be estimated from the arrival rates of calls to nodes. The complication in estimating the link arrival rates is that the entire offered load is not carried by the network because of call blocking. The probability of blocking is, in turn, dependent on the arrival rate to the links. This leads to a system of coupled nonlinear equations called the *Erlang map* [8]. While solving the Erlang map leads to a more accurate computation of blocking probabilities, the effect of blocking probabilities on the carried load can be neglected, especially when these probabilities are small. We take this approach in the rest of the paper to keep the analysis simple.

When the blocking probabilities are small, the link arrival rates, λ_e and λ_c, can be computed as follows. Consider a network with N nodes. Let the total arrival rate of calls at a link be γ $(=\lambda_e + \lambda_c)$. The sum of arrival rates of calls at all links in the network is γL, where L is the number of links in the network. Since a call uses a path of expected length $\bar{H} = \sum_{l=1}^{N-1} l p_l$, the sum of arrival rates at all links in the network is also $N\lambda\bar{H}$ where λ is the call arrival rate at a node. Thus

$$\gamma = \frac{N\lambda\bar{H}}{L}. \quad (3)$$

(For networks that are asymmetric, the arrival rates at all links may not be the same even if the node arrival rates are the same. In such cases, γ is the arrival rate averaged over all the links in the network.)

Having obtained the total arrival rate per link, we now estimate the arrival rates λ_e and λ_c. A plausible estimate for the probability of a call leaving the network at a given node is $1/\bar{H}$. Now, consider an intermediate node of a path and define an *exit link* of this node as an outgoing link that does not return to the previous node of the path. Suppose there are k exit links per node. We assume that if a call does not leave the network at the node, it chooses one of the k exit links arbitrarily.[5] Therefore, given a path and a link of the path, the arrival rate of calls that continue on to the next link of the path can be estimated as

$$\lambda_c = \gamma \frac{1 - 1/\bar{H}}{k}. \quad (4)$$

Calls that do not continue on to the next link of the path either leave the network or continue on a different link from

[5] For irregular topologies where the number of exit links may be different for different nodes, k could be taken as the average number of exit links per node.

the given node. The arrival rate of such calls at the given link is simply given as

$$\lambda_e = \gamma - \lambda_c. \tag{5}$$

It is worthwhile observing at this point that the correlation model we have just presented subsumes the independence model that assumes that link loads are independent. When λ_c is set to zero and λ_e is set to γ, our model greatly simplifies and reduces to the model presented in [10]. The independence model is shown to be inaccurate for sparse topologies in Section IV, and, more importantly, the correlation model is shown to be quite accurate.

III. Sparse Wavelength Conversion

In the previous section, we assumed that wavelength converters are absent in the network. Wavelength converters improve the blocking performance by allowing a circuit to be established as long as *some* wavelength is available on each link of the desired path. To enhance the blocking performance of the network, wavelength converters are placed at some of the nodes. All the previous analyses in the literature have considered networks without any converters or networks with converters at every node. An interesting design alternative that has not been considered previously is one in which wavelength conversion is available in a subset of network nodes to achieve a balance between cost and performance. An architecture for network nodes which convert wavelengths is presented in [18], and can be used in networks with sparse wavelength conversion.

We model a network with sparse wavelength conversion by assuming that a node is capable of wavelength conversion with probability q independently of the other nodes. q is called the *conversion density* of the network. The number of converter nodes in an N-node network is thus binomially distributed with an average of Nq converters, and the blocking performance we obtain is the ensemble average of the blocking probability over this distribution. This probabilistic approach enables a single parameter characterization of the wavelength conversion density and eliminates the unpleasant task of evaluating the performance for each number and placement of converters.

Note that we assume that a node can either convert any set of wavelengths to any other, or cannot convert any wavelength. However, our analysis also applies to the case of limited wavelength conversion per node discussed in [13]. In particular, we predict the ensemble average blocking performance of the share-per-node architecture of [13], where the number of conversions at each node is binomially distributed with a mean of FDq, D being the number of links per node.

Consider an l-hop path in a network with conversion density q and let the nodes on the path be numbered $s, 1, 2, \cdots, l-1, d$, as shown in Fig. 3. Let the call blocking probability on the path be denoted by $P_b^{(l)}(q)$. We recursively compute the call blocking probabilities on paths of different hop lengths. The idea behind the recursion is as follows. Suppose node i is the last converter on the l-hop path. A call is not blocked on the path if a) the call is not blocked on the first i hops of the path

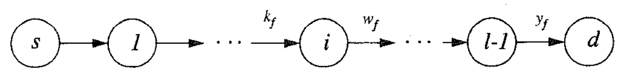

Fig. 3. An l-hop path with nodes numbered $s, 1, 2, \cdots, l-1, d$.

and b) there is a wavelength that is free on the last $l-i$ hops of the path. These two events, however, are not independent because the probability of blocking on the last $l-i$ hops depends on the number of free wavelengths on hop $(i+1)$. This depends on the number of free wavelengths on hop i, which, in turn, is dependent on whether the call is blocked on the first i hops or not. To analyze this situation exactly, we introduce the following probabilities. Let

- $V^{(l)}(n_f, y_f \mid w_f) = \Pr\{n_f$ wavelengths are free on an l-hop path and y_f wavelengths are free on hop $l \mid w_f$ wavelengths free on the first hop of the path and no converters along the path$\}$;
- $W^{(l)}(y_f \mid w_f) = \Pr\{y_f$ wavelengths are free on hop l of an l-hop path $\mid w_f$ wavelengths free on the first hop of the path and no converters along the path$\}$;
- $P_b^{(l)}(q, y_f) = \Pr\{$a call is blocked on an l-hop path *and* y_f wavelengths free on hop l, when the conversion density is $q\}$.

Equation (2) can then be modified to give

$$\begin{aligned} &V^{(l)}(n_f, y_f \mid w_f) \\ &= \sum_{x_{pf}=0}^{F} \sum_{z_c=0}^{Z} \sum_{x_{ff}=0}^{x_{pf}} R(n_f \mid x_{ff}, z_c, y_f) \\ &\quad \times U(z_c \mid y_f, x_{pf}) S(y_f \mid x_{pf}) V^{(l-1)}(x_{ff}, x_{pf} | w_f) \end{aligned} \tag{6}$$

for $l = 2, 3, \cdots, N-1$, where $Z = \min(F - x_{pf}, F - y_f)$. For a one-hop path

$$V^{(1)}(n_f, y_f \mid w_f) = \begin{cases} 1, & \text{if } n_f = y_f = w_f \\ 0, & \text{otherwise.} \end{cases}$$

Furthermore, by summing over all possible values of n_f in (6), we obtain

$$W^{(l)}(y_f \mid w_f) = \sum_{n_f=0}^{y_f} V^{(l)}(n_f, y_f \mid w_f).$$

The range of the variable n_f in the above equation is only up to y_f because the number of free wavelengths on an l-hop path without converters cannot be higher than the number of free wavelengths on any hop, in particular, hop l.

On a one-hop path, a call is blocked if and only if there is no free wavelength on the hop. Therefore

$$P_b^{(1)}(q, y_f) = \begin{cases} Q(0), & \text{if } y_f = 0 \\ 0, & \text{otherwise} \end{cases}$$

where $Q(w_f)$ is the probability that w_f wavelengths are free on a link, as given by (1).

For $l \geq 2$, the joint probability, $P_b^{(l)}(q, y_f)$, can be computed recursively by conditioning on the disjoint events that node i is the last converter on the given path of l hops,

$i = 1, 2, \cdots, l-1$. Thus

$$\begin{aligned} P_b^{(l)}(q, y_f) &= P(\text{Blocking } \textit{and } y_f \text{ free on last hop} \mid \text{no converters} \\ &\quad \text{in path}) \cdot P(\text{no converters in the path}) \\ &+ \sum_{i=1}^{l-1} P(\text{Blocking } \textit{and } y_f \text{ free on last hop} \mid \text{node } i \text{ last} \\ &\quad \text{converter}) \cdot P(\text{node } i \text{ last converter}). \end{aligned} \tag{7}$$

When there are no converters, the joint probability of a call being blocked and y_f wavelengths being free on the last hop of an l-hop path is given by $T^{(l)}(0, y_f)$ [see (2)]. The first term in (7) is thus $T^{(l)}(0, y_f)(1-q)^{l-1}$. The joint probability, $Y^{(l)}(i, q, y_f)$, of a call being blocked and y_f wavelengths being free on the last hop of an l-hop path when node i is the last converter is obtained in the Appendix. The resulting $P_b^{(l)}(q, y_f)$ is given by

$$\begin{aligned} P_b^{(l)}(q, y_f) &= T^{(l)}(0, y_f)(1-q)^{l-1} \\ &+ \sum_{i=1}^{l-1} Y^{(l)}(i, q, y_f) q (1-q)^{(l-i-1)} \end{aligned} \tag{8}$$

where $Y^{(l)}(i, q, y_f)$ is given in terms of $P_b^{(i)}(q, y_f)$ in (A.1) of the Appendix. The network call blocking probability when the conversion density is q is then given by

$$P_b(q) = \sum_{l=1}^{N-1} \sum_{y_f=0}^{F} P_b^{(l)}(q, y_f) p_l.$$

Given a network topology, we first determine the hop-length distribution, p_l, and the number of exit links, k. Then, (3)–(5) are used to obtain the call arrival rates. The blocking performance can then be studied by using the analysis presented in Section II and this section. In the next four sections, we evaluate the call blocking performance for the ring, the mesh-torus, the hypercube, and random topologies, respectively.

IV. Blocking in a Ring Network

In this section, we apply the above analysis to a unidirectional ring network. The ring is the most sparsely connected network with a given number of nodes[6], and we are interested in finding the effect of conversion density on such a network.

To verify the accuracy of the proposed analytical framework, we performed a simulation study of the call blocking performance in ring networks of different sizes. Each data point in the simulation was obtained using 10^6 call arrivals. We simulated the no-converter ($q = 0$) and full-converter ($q = 1$) cases and obtained the blocking probability to compare with the analytical results.

For an N-node unidirectional ring, the probability that an l-hop path is used for routing a session is, $p_l = \frac{1}{N-1}$ for $1 \le l \le N-1$, and k is 1. First, we plot the call blocking probability against the load per station for a 100-node network when the number of wavelengths per fiber are five and 20, in Fig. 4(a) and (b), respectively. Analytical and simulation results are plotted for the no converter case ($q = 0$), and the full-converter case ($q = 1$). In both cases, the converter case curves lie below the corresponding no-converter case curves. The close match between the analytical and simulation results indicates that the model is adequate in analytically predicting the performance of even very sparse networks. For comparison, we also plot the analytical results when $\lambda_e = \lambda \bar{H}$ and $\lambda_c = 0$ (independence model). It is seen that the independence model severely overestimates the blocking probability.[7] We observe that wavelength converters are more useful when the number of wavelengths per fiber is larger and the load is lower. This can be explained as follows. When the number of wavelengths is larger, blocking occurs primarily not due to a lack of resources (wavelengths) but due to the inability to use those resources efficiently in the absence of conversion. Thus, converters are more useful when the number of wavelengths is larger. Under heavy loads, blocking occurs primarily due to a lack of sufficient number of wavelengths and the presence of converters does not have as much effect as at lighter loads.

Fig. 5(a) shows how the blocking probability changes with wavelength conversion density, q, for several values of F for a 20-node (solid lines) and a 100-node (dashed lines) ring network when the network load is 2 Erlangs (a load of 0.1 per station for the 20-node network and a load of 0.02 per station for the 100-node network). Two important observations can be made from these curves. Firstly, when the network has more nodes and/or the number of wavelengths is higher, the blocking probability drops fast initially with conversion density and then rapidly levels off at a certain point. Secondly, the density at which the performance begins to level off increases marginally as the number of wavelengths increases. In contrast, when the network is smaller, the blocking probability decreases more gradually. In either case, the decrease in blocking due to conversion density is almost insignificant. This leads to the observation that adding converters is not the solution to increasing the performance in very sparse networks.

We have also simulated the performance of a ring network with a fixed number Nq of converters which are randomly placed over the ring. We do not show the results of this experiment in order not to impair the clarity of Fig. 5(a). These results indicate that the ensemble average is a very good approximation for this scenario.

Finally, we plot in Fig. 5(b) the number of stations that can be supported for a blocking probability of 10^{-3} against the conversion density, for different loads. When the load per station is high (0.05 Erlangs), conversion density has very little effect on the number of stations that can be supported. This is again due to the limitation of the resources and not due to lack of efficient utilization of those resources. When the load becomes lighter (0.01 Erlangs), wavelength conversion helps initially in increasing the utilization but the performance begins to level off when the limit on resources is approached.

[6] A star network is equally sparse in terms of the number of links, but has a lower average hop-length.

[7] The blocking probabilities are somewhat overestimated in all the curves because of the slight decrease in carried load due to blocking, which we have ignored.

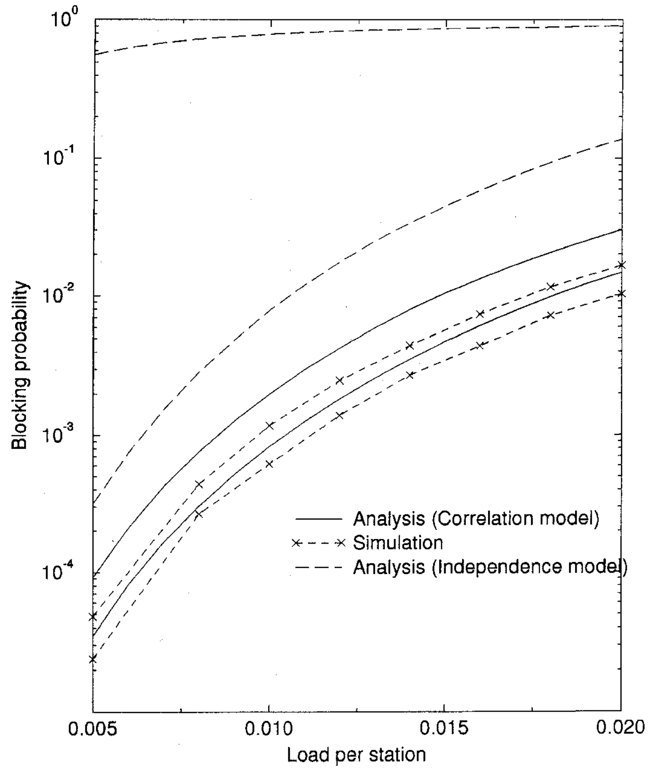

(a)

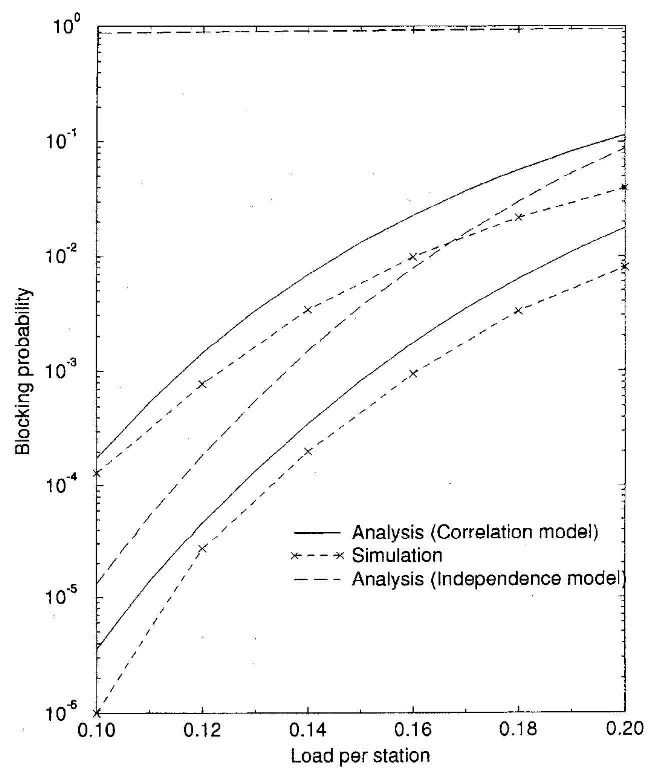

(b)

Fig. 4. The blocking probability versus the load per station for a 100-node ring network: (a) $F = 5$ and (b) $F = 20$.

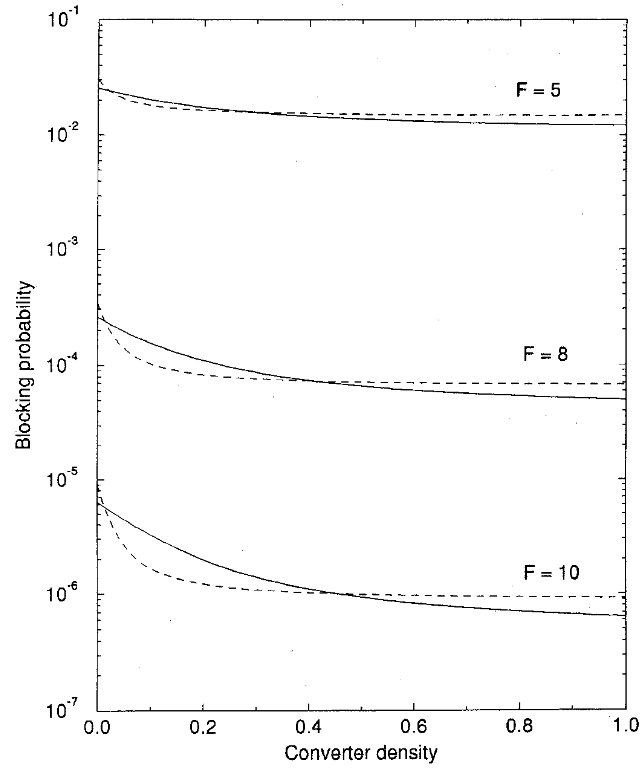

(a)

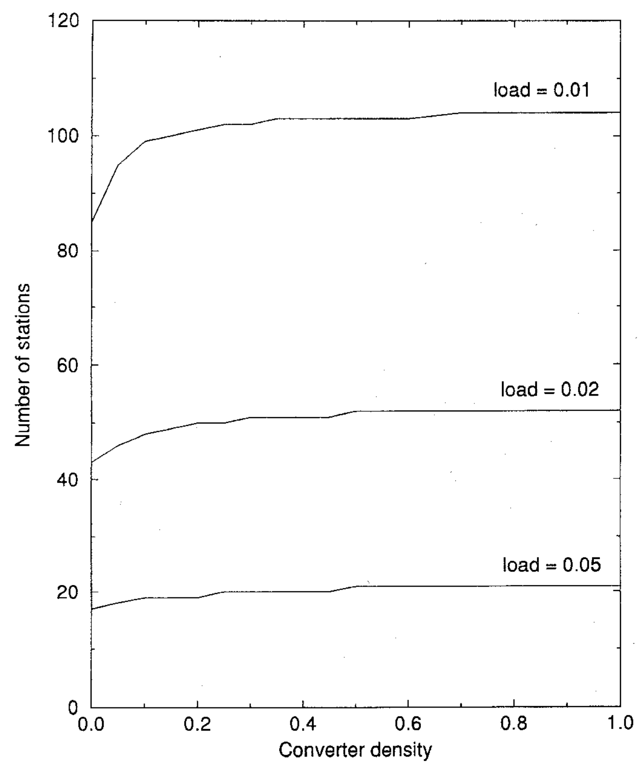

(b)

Fig. 5. (a) Blocking probability versus the converter density for a ring network with a total network load of 2 Erlangs. Solid lines: $N = 20$ and $\rho = 0.1$; dashed lines: $N = 100$ and $\rho = 0.02$. (b) The number of stations supported versus converter density for a ring network. $F = 5, P_b = 10^{-3}$.

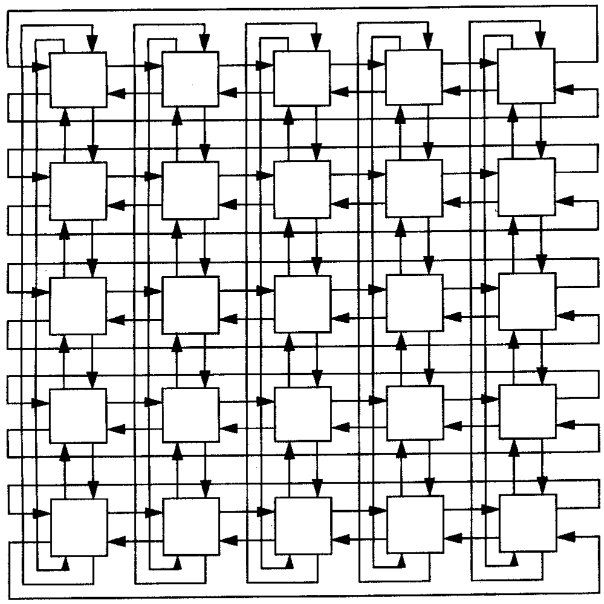

Fig. 6. A 5 × 5 bidirectional mesh-torus network.

V. Blocking in a Mesh-Torus Network

We next consider a bidirectional $M \times M$ mesh-torus network with $N = M^2$ nodes such as shown in Fig. 6. A mesh-torus network is more connected than a ring but the average-hop length is still large. For simplicity in computing the hop-length distribution, we consider only odd values of M. We assume static routing where the route for a connection is randomly chosen from one of the many shortest path routes available. With this routing scheme, there are $4i$ nodes at distance i and $4(M-j)$ nodes at distance j from any node, where $1 \le i \le \frac{M-1}{2}$ and $\frac{M-1}{2} < j \le M-1$. Therefore, we have

$$p_l = \begin{cases} \frac{4l}{M^2-1}, & 1 \le l \le \frac{M-1}{2}, \\ \frac{4(M-l)}{M^2-1}, & \frac{M-1}{2} < l \le M-1 \end{cases}$$

and the average hop-length is $O(M)$. The number of exit links per node, k, is three. The blocking performance is then analyzed by using the results of Sections II and III.

In Fig. 7, we plot the simulation and analytical (independence and correlation models) results for a 101 × 101 mesh-torus network with five wavelengths per fiber for the no-converter ($q = 0$) and full-converter ($q = 1$) cases. The results of our analytical model and the simulation results match very closely. The independence model is less accurate than our model but not significantly so, indicating that the load correlation between successive links is very high in sparse networks and decreases as the network becomes more densely connected. We observe from the figure the tremendous improvement in performance with wavelength conversion, unlike in the case of the ring.

In Fig. 8(a) and (b), we show the effect of conversion density on blocking probability for the 11 × 11 and a 101 × 101 bidirectional mesh-torus networks, respectively. As in the ring, we observe that conversion helps more when there are more wavelengths per fiber (the blocking probability drops more steeply with q as the number of wavelengths increases).

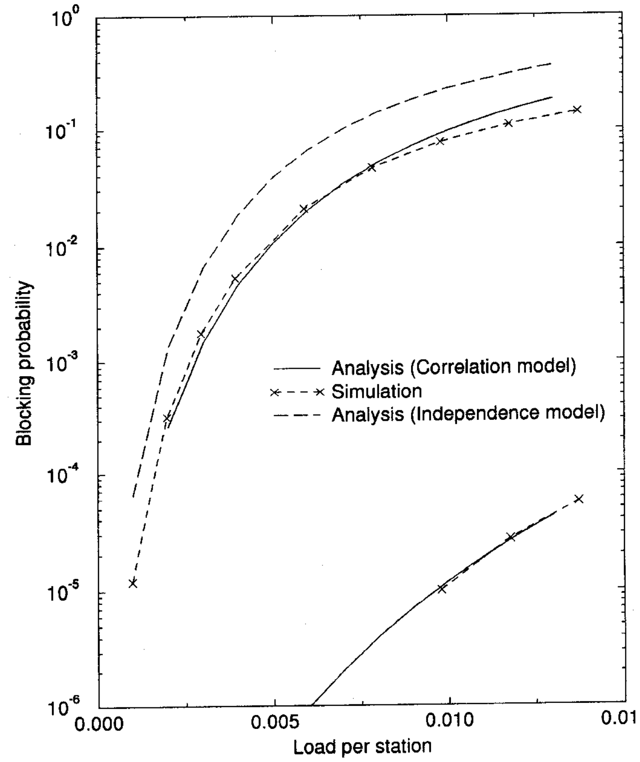

Fig. 7. The blocking probability versus the load per station for a 101 × 101 mesh-torus network with five wavelengths per fiber. Simulation results are reproduced from [10] with permission from the authors.

Furthermore, the advantages of wavelength conversion are much higher in a larger network and as the network size increases, performance increases dramatically initially with conversion density. For example, we observe from Fig. 8(b) that, when $F = 5$, the blocking probability drops from 10^{-1} to 10^{-3} as the conversion density increases from zero to about 0.2, and then decreases more gradually. When $F = 8$, a decrease in blocking probability of two orders of magnitude occurs as q increases from zero to about 0.1. This not only suggests that having a converter at every node may be unnecessary to achieve a certain performance but also that the proportion of converter nodes required is a function of the size of the network and the number of wavelengths per fiber. However, unlike in the ring, when there are a large number of wavelengths, the performance does not level off sharply with increasing conversion density. This suggests that the utilization of converters is affected more by the large hop-lengths in the mesh than by the load correlation that dominates in the ring.

Fig. 9 shows how the number of stations that can be supported increases with the conversion density. At lighter loads, the advantages of conversion are much more significant than in the case of the ring. (Note that the square root of the number of stations is plotted on the Y-axis.) This is due to the better mixing of traffic in a mesh-torus although the paths are shorter than those in the ring.

VI. Blocking in a Hypercube Network

We next analyze a well-connected network, the binary hypercube (called hypercube in the rest of the paper). An

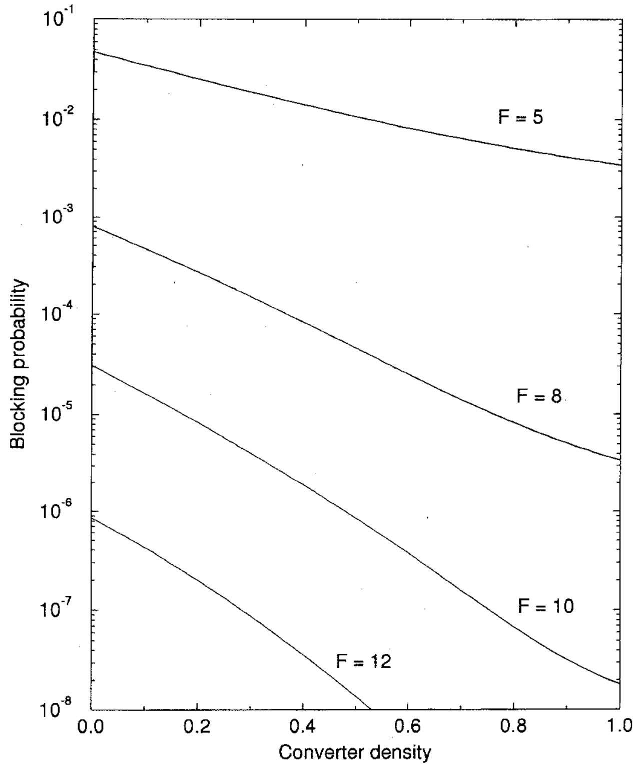

(a)

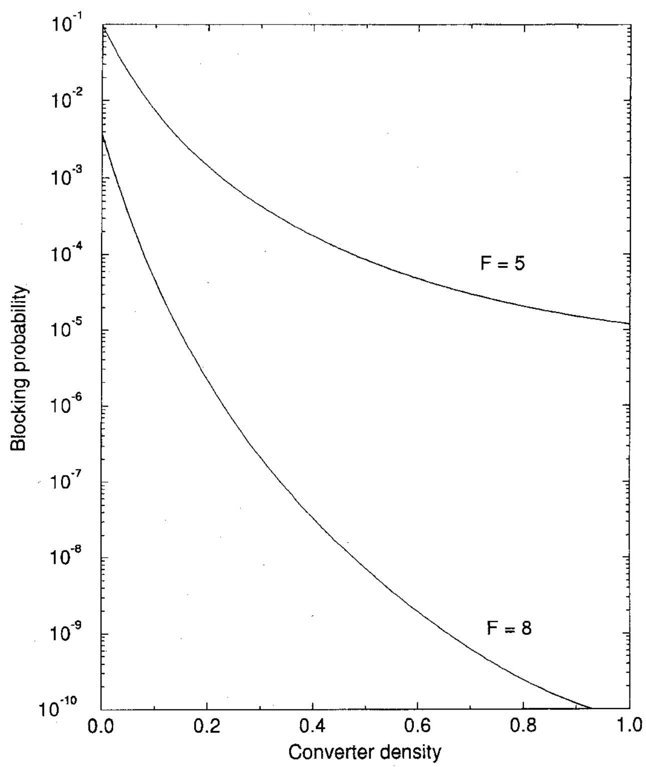

(b)

Fig. 8. The blocking probability versus conversion density for (a) a 11 × 11 mesh-torus with $\rho = 0.5$ and (b) a 101 × 101 mesh-torus with $\rho = 0.01$.

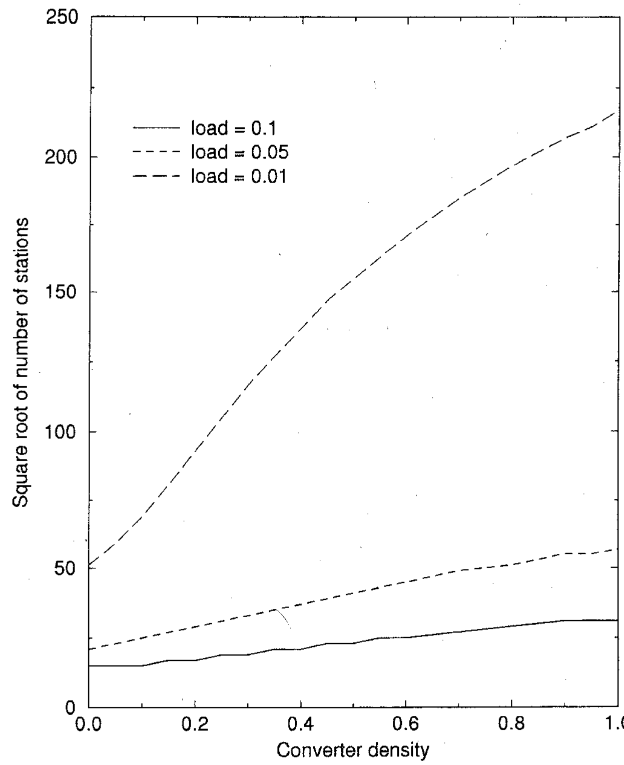

Fig. 9. The square root of the number of stations supported vs. converter density for a bidirectional mesh-torus network. $F = 5, P_b = 10^{-3}$.

$N(=2^n)$-node hypercube network has n outgoing links per node. We assume that one of the shortest paths between the source and destination nodes is chosen for routing a session. Using this routing scheme, we have

$$p_l = \frac{1}{N-1}\binom{n}{l}, \quad 1 \leq l \leq n$$

and the average hop-length is $O(n)$. The number of exit links per node, $k = n - 1$. The blocking performance is then analyzed using the results of Sections II and III.

Fig. 10(a) and (b) show how the performance varies with the load per station for a 32-node and a 1024-node hypercube, respectively. Again, the converter case curves lie below the corresponding no-converter case curves. The analysis and simulation results match very closely. Because of very low load correlation between successive links, the independence model also predicts the performance accurately, though not as well as the correlation model. Since hop-lengths are small in a hypercube network, we do not expect converters to be very useful. The figures corroborate this intuitive observation.

In Fig. 11, we plot the blocking probability against the converter density for a 1024-node hypercube with a load of 0.1 Erlangs per station. The previous observation of converters helping more when there are more wavelengths per fiber holds for the hypercube as well. We also see that the performance improves dramatically with an increase in the number of wavelengths. This is a direct consequence of the fact that the hypercube nodes have very high degrees, and therefore, the

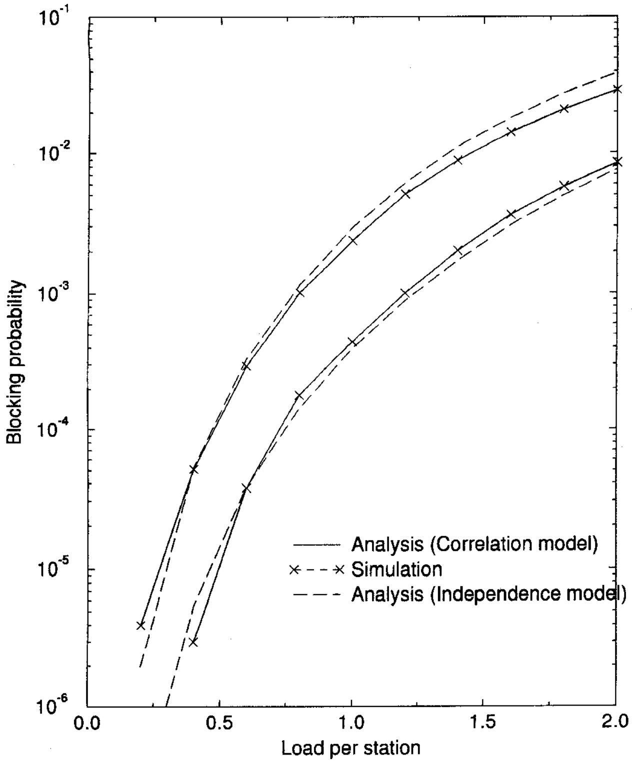

(a)

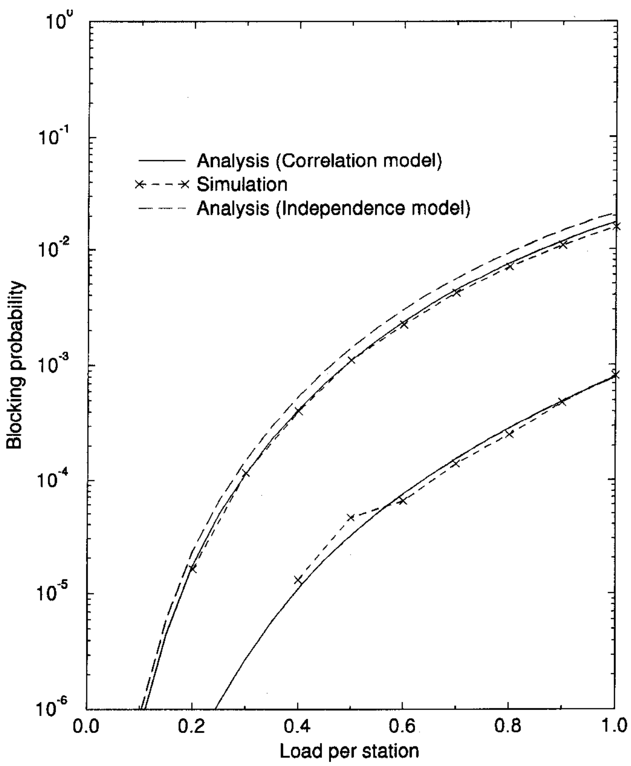

(b)

Fig. 10. The blocking probability versus the load per station for a hypercube network when the number of wavelengths per fiber is five. (a) $N = 32$ and (b) $N = 1024$.

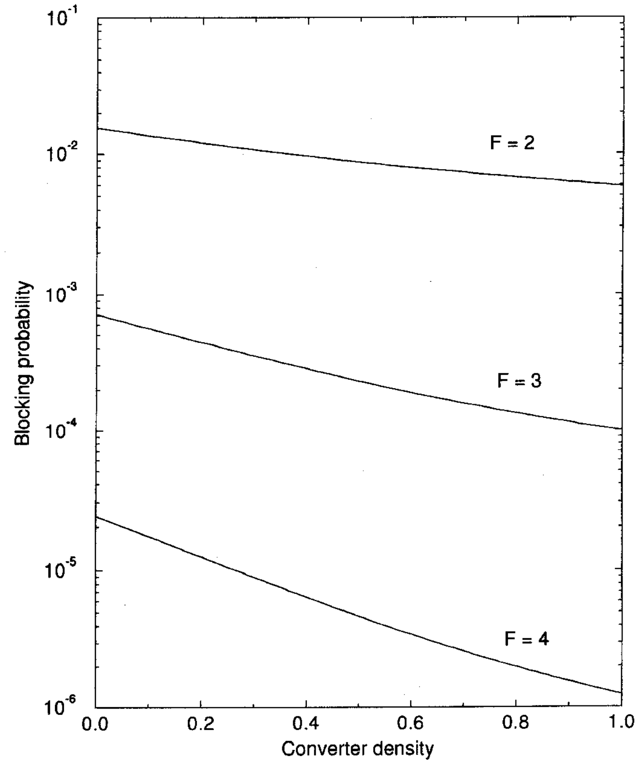

Fig. 11. The blocking probability versus the converter density for a 1024-node hypercube network with a load of 0.1 Erlangs per station.

total number wavelengths in the network grows very rapidly with increasing wavelengths per link.

The results of the last three sections lead to the remarkable observation that the usefulness of converters is a complex function of connectivity. On one hand, when connectivity is low as in the ring, hop-lengths are large and converters are expected to be more useful. On the other hand, the load correlation between successive links is high and this tends to reduce the usefulness of converters [9]. This latter effect dominates the hop-length effect in the ring. In the hypercube, the hop-lengths are small. Therefore, even though the load correlation is negligible, converters do not offer significant advantages. In the mesh network, the load correlation is fairly low while hop-lengths are large enough so that converters improve performance dramatically. The benefits of conversion are thus dependent on which of the two effects dominates in a given topology and are difficult to predict *a priori* without a detailed analysis such as the one we have presented.

VII. Blocking in Random Topologies

We studied the above networks to gain an understanding of how conversion density affects the blocking performance under varying degrees of network connectivity. In this section, we study the performance of a more likely topology for a WAN, namely an irregular topology. It is very difficult to analytically predict the performance of a given irregular topology with a particular placement of wavelength converters. However, it is possible to analyze the (ensemble) average

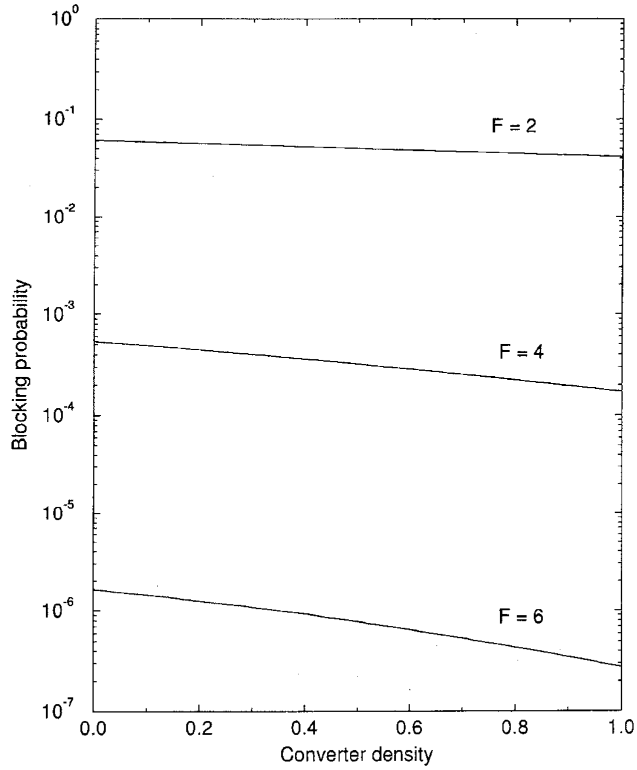

Fig. 12. The blocking probability versus the converter density for a random network. $N = 100$, $b = 10$, $\rho = 1$.

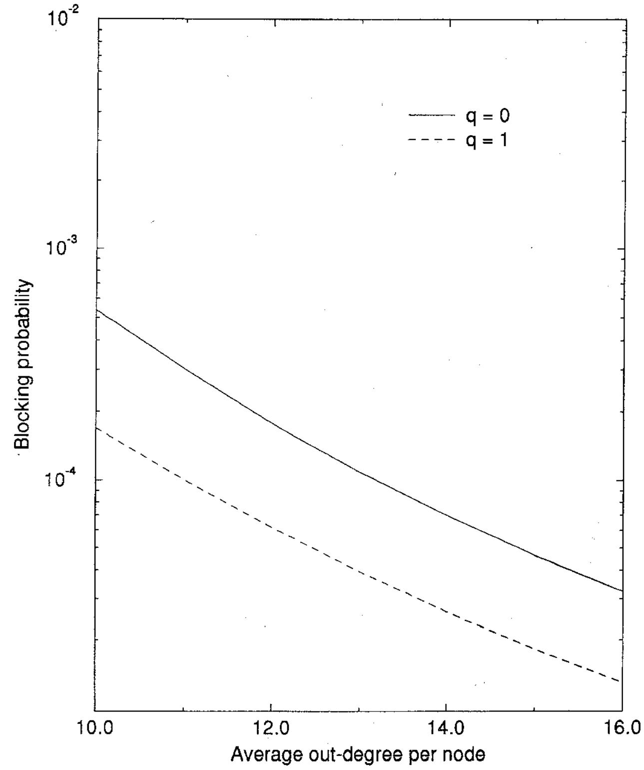

Fig. 13. The blocking probability versus the average out-degree per node for a random network. $N = 100$, $F = 4$, $\rho = 1$.

performance of all irregular network topologies characterized by a given connectivity parameter with a converter distribution characterized by another parameter.

We employ the following model for a random topology of N nodes. Each of the possible $N(N-1)$ directional links exists in the network with probability $\frac{b}{N-1}$, independently of all other links. Thus, each node has an average of b outgoing links. This model does not ensure that the network is connected. However, the model lends itself to easy analysis of the hop-length distribution and is used in this paper.

We again assume that one of the shortest paths between the two end-nodes is chosen arbitrarily for a connection. The hop-length distribution is easily obtained for the regular topologies discussed but is extremely difficult to compute exactly for a random network. Exact computation of the probabilities is trivial for one and two hop paths but for longer paths, the complexity is exponential in the network size [19]. Asymptotic expressions for the shortest-path distribution are given in [20]. In this paper, a slight variation of an approximation given in [21] is used. The approximation has been observed to be accurate for large values of N in simulations.

Given a node, let h_l be the average number of nodes that are reached on the lth hop from the given node, and let Γ_l be the average fraction of nodes (excluding the given node) that can be reached in l or fewer hops. The hop-length distribution is computed using the following set of recurrence relations, where $p_l = h_l/(N-1)$ is the probability of an l-hop path

$$h_l = (N-1)(1-\Gamma_{l-1})\left[1 - \left(1 - \frac{b}{N-1}\right)^{h_{l-1}}\right]$$

$$\Gamma_l = \Gamma_{l-1} + \frac{h_l}{N-1}$$

with $h_1 = b$, and $\Gamma_1 = b/(N-1)$.

The average number of exit links per node can be shown to be

$$k = \frac{\left(1 - \frac{1}{N-1}\right)b}{1 - \left(1 - \frac{b}{N-1}\right)^{N-2}}.$$

We assume a value of b that is sufficiently high to keep the probability of an unconnected network small. For $N \leq 500$, a value of $b \geq 10$ almost ensures connectivity.

We obtain the blocking performance using the analysis of Sections II and III. Fig. 12 shows that wavelength converters do not affect the performance much in a densely connected random network. A small increase in the number of wavelengths per link causes a significant improvement in the blocking performance, as shown in Fig. 12. This is because of the large number of links in the network. Because of the abundance of the number of available wavelengths in the network, the blocking probability never levels off with increasing conversion density. However, the decrease in blocking with increasing conversion density is only marginal because of the short hop-lengths.

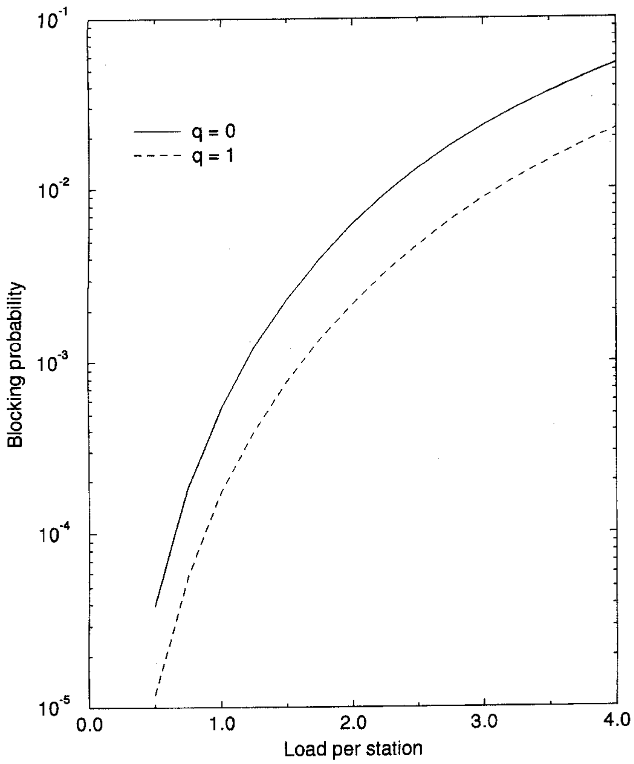

Fig. 14. The blocking probability versus the load per station for a random network. $N = 100$, $F = 4$, $b = 10$.

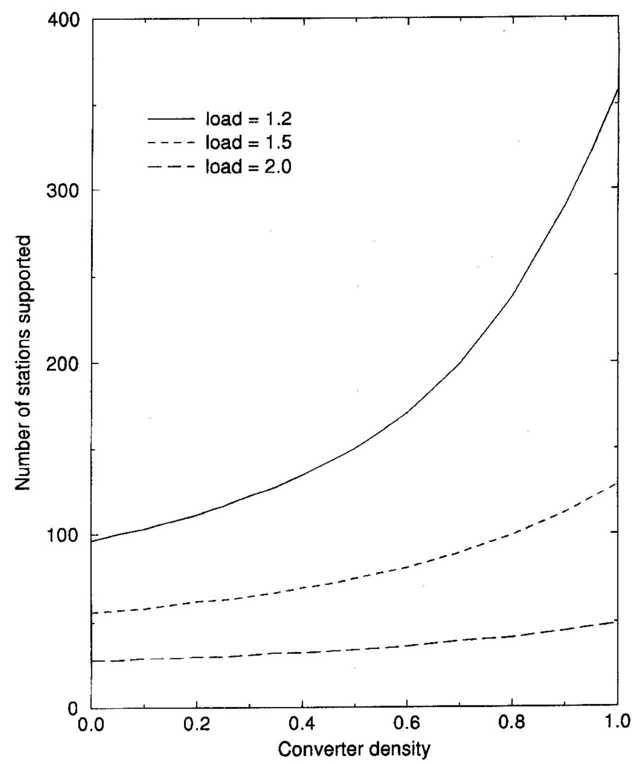

Fig. 15. The number of stations supported versus the converter density for a random network. $F = 4$, $b = 10$.

Fig. 13 shows the dramatic improvement in performance with network connectivity. In Fig. 14, we have plotted the blocking probability against the load per station for the no-converter and full-converter cases when the average out-degree per node is ten.

Fig. 15 shows how the maximum number of stations that can be supported so that the blocking probability is below 10^{-3} varies with conversion density for different loads. As before, when the loads are heavy, conversion does not help much. But when the load is light, conversion is beneficial. However, the benefits of conversion are not as great as in a less dense network such as the mesh.

The average amount of resources available per node is the average number of wavelengths available per node. For the random network considered here, this is equal to the product of the average number of links per node and the number of available wavelengths per link. Fig. 16 shows how the blocking performance varies with the number of wavelengths per link while the network becomes less dense so that the average amount of resources in the network is held constant. We see that, when the network is fairly dense ($b \geq 10$ for a 500-node network), it is always better to have fewer links with more wavelengths per link.

VIII. Conclusion

We have introduced the concept of sparse wavelength conversion in wavelength-routed, all-optical, circuit-switched networks. To evaluate the blocking performance of such networks, we have developed an analytical model of modest complexity. We have shown the model to be accurate for a variety of network topologies by comparing the analytical results with results from simulations. We applied this model to analyze three regular network topologies with varying degrees of connectivity—ring, mesh-torus, and hypercube.

An important conclusion of our study is that the usefulness of wavelength converters depends on the connectivity of the network in a manner that cannot be predicted by intuition. When the connectivity is low, such as in the ring, converters are not very useful because of the high load correlation. This high correlation implies that the expected number of links shared by any two sessions is large; conversion merely permutes the wavelengths used by the sessions without greatly increasing the ability to accommodate a new session. When the connectivity is high, such as in the hypercube or a densely connected random topology, converters are not very beneficial because of small hop-lengths, despite low load correlation and significant traffic mixing. This indicates that in a WAN with an irregular topology and a large number of links, the number of wavelengths is much more important than wavelength conversion capability. A mesh-torus network with a degree of connectivity between those of the ring and the hypercube, remarkably, offers great advantages when wavelength converters are present. This is because the link load correlation is almost negligible while the hop-lengths are large compared to the more densely connected hypercube.

The performance improves rapidly as the conversion density increases from zero, but the rate of improvement typically *decreases* with increasing conversion density. The conversion

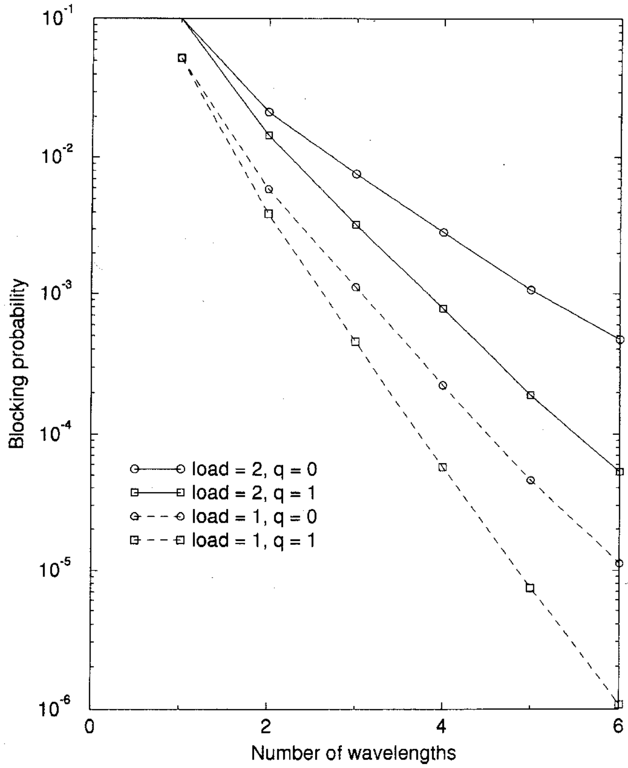

Fig. 16. The blocking probability versus the number of wavelengths per link for a random network. $N = 500$, $F \times b = 64$.

density beyond which the performance improves only marginally depends on the number of wavelengths and the network load. This suggests that placing a converter at every node is rarely necessary, if at all. In general, wavelength converters are more effective when the number of wavelengths is larger and when the load is lower.

The results of this paper point out the importance of a detailed analysis of a given network topology in determining the number and placement of wavelength converters. While we have shown that having converters at a small fraction of nodes is typically sufficient for a desired performance, the problem of converter placement remains to be studied.

Our model has assumed the use of fixed shortest path routing for establishing a session. The effect of alternate or dynamic routing on the benefits of conversion is an issue that merits attention. Finally, blocking probability is but one performance measure and other performance measures (such as throughput and delay in a packet-switched network) could be considered to study the usefulness of wavelength conversion.

Appendix

In this Appendix, we obtain the joint probability, $Y^{(l)}(i, q, y_f)$, of a call being blocked on an l-hop path and y_f wavelengths being free on the last hop of the path when the conversion density is q and node i is the last converter. This probability can be written as $Y^{(l)}(i, q, y_f) = \Pr\{y_f$ wavelengths free on hop $l \mid$ node i last converter $\} - \Pr\{$No blocking on first i and on last $l - i$ hops and y_f free wavelengths on hop $l \mid$ node i last converter$\}$.

The probability that y_f wavelengths are free on hop l is $Q(y_f)$ and is statistically independent of conversion density. Referring to Fig. 3, $\Pr\{$No blocking on first i and on last $l - i$ hops and y_f free wavelengths on hop $l \mid$ node i last converter$\} = \sum_{k_f=0}^{F} \sum_{w_f=0}^{F} P(A)P(B)P(C)$, where the events A, B, and C are defined as follows.

A: Call is not blocked on the first i hops and k_f wavelengths are free on hop i.

B: w_f wavelengths are free on hop $i+1$ given k_f wavelengths are free on hop i.

C: Call is not blocked on the last $l - i$ hops and y_f wavelengths are free on hop l given that w_f wavelengths are free on hop $i+1$ and there are no converters along the $(l - i)$-hop path.

Now

$$P(A) = Q(k_f) - P_b^{(i)}(q, k_f)$$
$$P(B) = S(w_f \mid k_f),$$

and

$$P(C) = W^{(l-i)}(y_f \mid w_f) - V^{(l-i)}(0, y_f \mid w_f).$$

Substituting for $P(A)$, $P(B)$, and $P(C)$, we can now write

$$\begin{aligned} &Y^{(l)}(i, q, y_f) \\ &= Q(y_f) - \left\{ \sum_{k_f=0}^{F} \left(Q(k_f) - P_b^{(i)}(q, k_f)\right) \cdot \right. \\ &\qquad \sum_{w_f=0}^{F} S(w_f \mid k_f) \cdot \\ &\qquad \left. \left(W^{(l-i)}(y_f \mid w_f) - V^{(l-i)}(0, y_f \mid w_f)\right) \right\}. \end{aligned} \tag{A.1}$$

This expression is used in (8) to evaluate the blocking probability of networks with sparse wavelength conversion.

References

[1] R. Ramaswami, "Multiwavelength lightwave networks for computer communication," *IEEE Commun. Mag.*, vol. 31, no. 2, pp. 78–88, Feb. 1993.

[2] I. Chlamtac, A. Ganz, and G. Karmi, "Lightpath communications: An approach to high bandwidth optical WAN's," *IEEE Trans. Commun.*, vol. 40, no. 7, pp. 1171–1182, July 1992.

[3] B. Mukherjee, "WDM-based local lightwave networks—Part I: Single-hop systems," *IEEE Network*, vol. 6, no. 3, pp. 12–27, May 1992.

[4] P. E. Green Jr., *Fiber Optic Networks*, Prentice Hall, 1992.

[5] M. S. Goodman, "Multiwavelength networks and new approaches to packet switching," *IEEE Commun. Mag.*, vol. 27, no. 10, pp. 27–35, Oct. 1989.

[6] C. A. Brackett, "Dense wavelength division multiplexing networks: principles and application," *IEEE J. Select. Areas Commun.*, vol. 8, no. 6, pp. 948–964, Aug. 1990.

[7] R. A. Barry and P. A. Humblet, "On the number of wavelengths and switches in all-optical networks," *IEEE Trans. Commun.*, vol. 42, no. 2/3/4, pp. 583–591, Feb./Mar./Apr. 1994.

[8] A. Girard, *Routing and Dimensioning in Circuit-Switched Networks*, Addison-Wesley, 1990.

[9] R. A. Barry and P. A. Humblet, "Models of blocking probability in all-optical networks with and without wavelength changers," in *Proc. IEEE INFOCOM'95*, Apr. 1995, pp. 402–412.

[10] M. Kovačević, and A. S. Acampora, "On wavelength translation in all-optical networks," in *Proc. IEEE INFOCOM'95*, Apr. 1995, pp. 413–422.
[11] A. Birman, "Computing approximate blocking probabilities for a class of all-optical networks," in *Proc. of IEEE INFOCOM'95*, Apr. 1995, pp. 651–658.
[12] B. S. Glance *et al.*, "New advances in optical components needed for FDM optical networks," *IEEE/OSA J. Lightwave Technol.*, vol. 11, no. 5/6, pp. 882–890, May/June 1993.
[13] K. C. Lee and V. O. K. Li, "A wavelength-convertible optical network," *IEEE/OSA J. Lightwave Technol.*, vol. 11, no. 5/6, pp. 962–970, May/June 1993.
[14] S. B. Alexander *et al.*, "A precompetitive consortium on wide-band all-optical networks," *IEEE/OSA J. Lightwave Technol.*, vol. 11, no. 5/6, pp. 714–735, May/June 1993.
[15] A. Mokhtar and M. Azizoğlu, "Dynamic route selection and wavelength assignment in all-optical networks," in *Proc. 8th Annual LEOS Meeting*, Oct. 1995, pp. 220–221.
[16] S. Subramaniam, M. Azizoğlu, and A. K. Somani, "Connectivity and sparse wavelength conversion in wavelength-routing networks," in *Proc. IEEE INFOCOM'96*, Mar. 1996, pp. 148–155.
[17] D. Bertsekas and R. Gallager, *Data Networks*, Prentice Hall, 1992.
[18] R. Ramaswami and K. N. Sivarajan, "Routing and wavelength assignment in all-optical networks," *IEEE/ACM Trans. Networking*, vol. 3, no. 5, pp. 489–500, Oct. 1995.
[19] G. A. Corea and V. G. Kulkarni, "Shortest paths in stochastic networks with arc lengths having discrete distribution," *Networks*, vol. 23, no. 3, pp. 175–183, May 1993.
[20] S. K. Walley, H. H. Tan, and A. M. Viterbi, "Limit distributions for the diameter and the shortest path hop count in random graphs with positive integer edge costs," in *Proc. IEEE INFOCOM'94*, June 1994, pp. 618–627.
[21] C. Rose, "Mean internodal distance in regular and random multihop networks," *IEEE Trans. Commun.*, vol. 40, no. 8, pp. 1310–1318, Aug. 1992.

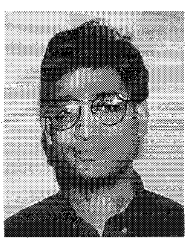

Suresh Subramaniam (S'95) received the B.E. degree in electronics and communication engineering from Anna University, Madras, India, in 1988, and the M.S.E.E. degree from Tulane University, New Orleans, LA, in 1993.

He was a Hardware Design Engineer at HCL Ltd., Madras, India, from 1988 to 1990, and a graduate student in electrical communication engineering at the Indian Institute of Science, Bangalore, India, in 1990–1991. He is currently pursuing the Ph.D. degree in electrical engineering at the University of Washington, Seattle. His research interests are in the areas of optical and ATM networks.

Murat Azizoğlu (S'86–M'91) received the B.S. degree from Middle East Technical University, Ankara, Turkey, in 1985, the M.S. degree from the Ohio State University, Columbus, in 1987, and the Ph.D. degree from Massachusetts Institute of Technology, Cambridge, in 1991, all in electrical engineering.

He was a Research Assistant at the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, in 1988–1991. He was a Member of Technical Staff at Bell Communications Research, Morristown, NJ, in 1991, and an Assistant Professor in the Department of Electrical Engineering and Computer Science at the George Washington University, Washington, DC, in 1991–1994. In September 1994, he joined the Department of Electrical Engineering at the University of Washington, Seattle, as an Assistant Professor. His teaching and research interests are in high-speed networks, communication theory and optical communications.

Dr. Azizoğlu received a National Science Foundation Career Award in 1995.

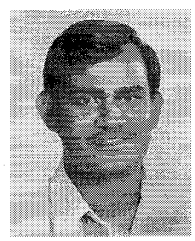

Arun K. Somani (S'83–M'83–S'84–M'85–SM'88) received the M.S.E.E. and Ph.D. degrees in electrical engineering from the McGill University, Montreal, Canada, in 1983 and 1985, respectively.

He is a Professor of Electrical Engineering and Computer Science and Engineering at the University of Washington, Seattle. Prior to that, he worked as a Scientific Officer for the Government of India, New Delhi, from 1974 to 1982. During this period, he designed and developed an anti-submarine warfare system for the Indian Navy. His research interests are in the areas of fault tolerant computing, computer interconnection networks, parallel computer systems, and design and analysis of algorithms. He designed and implemented the "Proteus" architecture, a multiprocessor system for automated classification of objects based on a generalized enhanced hypercube reconfigurable interconnection network exploring coarse grain parallelism.