Computing Blocking Probabilities in Multiclass Wavelength Routing Networks

SRIDHAR RAMESH, GEORGE N. ROUSKAS, and HARRY G. PERROS
North Carolina State University

We present an approximate analytical method to evaluate efficiently and accurately the call blocking probabilities in wavelength routing networks with multiple classes of calls. The model is fairly general and allows each source-destination pair to service calls of different classes, with each call occupying one wavelength per link. Our approximate analytical approach involves two steps. The arrival process of calls on some routes is first modified slightly to obtain an approximate multiclass network model. Next, all classes of calls on a particular route are aggregated to give an equivalent single-class model. Thus, path decomposition algorithms for single-class wavelength routing networks may be readily extended to the multiclass case. This article is a first step towards understanding the issues arising in wavelength routing networks that serve multiple classes of customers.

Categories and Subject Descriptors: G.3 [**Mathematics of Computing**]: Probability and Statistics—*Markov processes*; I.6.4 [**Simulation and Modeling**]: Model Validation and Analysis

General Terms: Performance

Additional Key Words and Phrases: Blocking Probabilities, Multiclass networks, Wavelength division multiplexing (WDM), Wavelength routing

1. INTRODUCTION

A basic property of single mode optical fiber is its enormous low-loss bandwidth of several tens of terahertz (THz). Unfortunately, due to dispersive effects and limitations in optoelectronic component technology, single channel transmission is limited in speed to only a small fraction of the fiber capacity. To take full advantage of the potential of fiber, the use of wavelength division multiplexing (WDM) techniques has become the option of choice [Brackett 1996; Green 1996], and WDM networks have been a subject of research both theoretically [Subramaniam et al. 1996; Ramaswami and Sivarajan 1996] and experimentally [Hall et al. 1996; Wagner et al. 1996]. Multiwavelength optical networks have the potential of delivering an aggregate throughput on the order of terabits per second, and they appear to be a viable approach to satisfying the ever-growing demand for more bandwidth per user on a sustained long-term basis.

The wavelength routing mesh architecture appears promising for wide area network (WAN) distances [Mukherjee et al. 1996; Wauters and Demeester 1996; Chlamtac et al. 1992]. The network architecture consists of wavelength

This work was supported by the NSF under grant ANI-9805016.
Authors address: Department of Computer Science, North Carolina State University, Raleigh, NC 27695-7534; e-mail: rouskas@csc.ncsu.edu.

routers and fiber links that interconnect them. A wavelength router is capable of switching a light signal at a given wavelength from any input port to any output port. A router may also be capable of enforcing a shift in wavelength [Ramamurty and Mukherjee 1998], in which case a light signal may emerge from the switch at a different wavelength than the one it arrived. By appropriately configuring the routers, all-optical paths (lightpaths) may be established between pairs of nodes in the network. Lightpaths represent direct optical connections without any intermediate electronics. Because of the long propagation delays, and the time required to configure the routers, wavelength routing WANs are expected to operate in circuit-switched mode [Ramaswami and Sivarajan 1996; Mukherjee et al. 1996]. This architecture is attractive for two reasons: the same wavelength can be used simultaneously at different parts of the network; and the signal power is channeled to the receiver and is not spread to the entire network as in the broadcast-and-select approach [Ramaswami 1993]. Hence, wavelength routing WANs can be highly scalable.

Given the installed base of optical fiber, and the maturing of optical component technology, it appears that current network technologies are transitory, and will eventually evolve to an all-optical, largely passive infrastructure. Although the long-term result of such an evolution is not yet clear, a feasible scenario for near-term large-scale all-optical networks has emerged in recent years [Brackett 1996; Green 1996]. Wavelength routing WANs are an integral part of this scenario, since it is envisioned that they will act as the backbone that provides the interconnection for local area photonic subnetworks attached to them. The contribution of our work is the development of an approximate, but accurate and efficient, analytical framework for evaluation of the performance of multiclass wavelength routing optical networks.

Pankaj and Gallager [1996] study lower bounds on the number of wavelength channels required in all-optical networks with and without wavelength converters, in order to solve an arbitrary permutation routing problem in a nonblocking manner. Our approach lies in solving for the blocking probabilities in a network with a fixed number of wavelengths and a given routing scheme. The problem of computing call blocking probabilities under static (fixed or alternate) routing with random wavelength allocation and with or without wavelength converters has been studied in Barry and Humblet [1996]; Kovacevic and Acampora [1996]; Birman [1996]; Harai et al. [1997]; and Subramaniam et al. [1996; 1997]. The model presented in Barry and Humblet [1996] is based on the assumption that wavelength use on each link is characterized by a fixed probability, independently of other wavelengths and links, and thus, it cannot capture the dynamic nature of traffic. In Kovacevic and Acampora [1996] it was assumed that statistics of link loads are mutually independent, an approximation that is not accurate for sparse network topologies. The work in Birman [1996] developed a Markov chain with state-dependent arrival rates to model call blocking in arbitrary mesh topologies and fixed routing; it was extended in Harai et al. [1997] to alternate routing. While more accurate, this approach is computationally intensive and can only be applied to networks of small size in which paths have at most three links. A more tractable model was presented in Subramaniam et al. [1996] to recursively compute blocking probabilities assuming that the load on link i of a path depends only on the load of link $i - 1$. Finally, a study of call blocking under non-Poisson input traffic was presented in Subramaniam et al. [1997], under the assumption that link loads are statistically independent.

Other wavelength allocation schemes, as well as dynamic routing are harder to analyze. First-fit wavelength allocation was studied using simulation in Chlamtac et al. [1992] and Kovacevic and Acampora [1996], and it was shown to perform better than random allocation, while an analytical overflow model for first-fit allocation was developed in Karasan and Ayanoglu [1998]. A dynamic routing algorithm that selects the least loaded path–wavelength pair was also studied in Karasan and Ayanoglu [1998]; and in Mokhtar and Azizoglu [1998], an unconstrained dynamic routing scheme with a number of wavelength allocation policies was evaluated. Except in Subramaniam et al. [1996; 1998], all other studies assume that either all or none of the wavelength routers have wavelength conversion capabilities. The work in Subramaniam et al. [1996] takes a probabilistic approach in modeling wavelength conversion by introducing the converter density, which represents the probability that a node is capable of conversion independently of other nodes in the network. While this approach works well when the objective is the estimation of the expected call blocking performance, it cannot be used to calculate the actual blocking probability on individual paths when the placement of converters is known, nor can it be used to compare various converter placement strategies. Finally, in Subramaniam et al. [1998] a dynamic programming algorithm to determine the location of converters on a single path that minimizes average or maximum blocking probability was developed under the assumption of independent link loads.

Most of the approximate analytical techniques developed for computing blocking probabilities in wavelength routing networks [Kovacevic and Acampora 1996; Birman 1996; Harai et al. 1997; Subramaniam et al. 1997; Karasan and Ayanoglu 1998; Mokhtar and Azizoglu 1998; Subramaniam et al. 1998] make the assumption that link blocking events are independent and amount to the well-known *link decomposition* approach [Girard 1990], while the development of some techniques is based on the additional assumption that link loads are also independent. Link decomposition has been extensively used in conventional circuit-switched networks where there is no requirement for the *same* wavelength to be used on successive links of the path taken by a call. The accuracy of these underlying approximations also depends on the traffic load, the network topology, and the routing and wavelength allocation schemes employed. While link decomposition techniques make it possible to study the qualitative behavior of wavelength routing networks, more accurate analytical tools are needed to evaluate the performance of these networks efficiently, as well as to tackle complex network design problems, such as selecting the optical switches where wavelength converters are to be employed.

We have considered the problem of computing call blocking probabilities in mesh wavelength routing networks with fixed and alternate routing and random wavelength allocation. Unlike previous studies, we have developed an iterative *path decomposition* algorithm [Zhu et al. 1999c] for analyzing arbitrary network topologies. Specifically, we analyze a given network by decomposing it into a number of path subsystems. These subsystems are analyzed in isolation using our approximation algorithm for computing blocking probabilities in a single path in a wavelength routing network [Zhu et al. 1999b]. The individual solutions are appropriately combined to form a solution for the overall network, and the process repeats until the blocking probabilities converge. Our approach accounts for the correlation of both link loads and link blocking events, giving accurate results for a wide range of loads and network topologies. It also allows nonuniform traffic, in the sense that call request arrival rates can vary for each source-destination pair. Our algorithm can compute call blocking probabilities in a mesh network where only a fixed but arbitrary subset of nodes is capable of wavelength conversion. Therefore, our algorithm can be an important tool in the development and evaluation of converter placement strategies. Finally, in Zhu et al. [1999a] we studied the first-fit and most-used wavelength allocation policies, and we showed that they have almost identical performance in terms of blocking probability for all calls in the network. We also demonstrated that the blocking probabilities under the random wavelength allocation policy with no converters and with converters at all nodes provides upper and lower bounds for the values of the blocking probabilities under the first-fit and most-used policies.

All previous studies of call blocking probabilities discussed here have only considered single-class wavelength routing networks. It is expected, however, that a wide range of future applications with varying characteristics in terms of their arrival rates and call holding times will be utilized. In this article, we present a method to extend the results in Zhu et al. [1995b; 1999c] to multiclass optical networks. While multiclass circuit-switched networks have been studied in the literature [Baynat and Dallery 1993], to the best of our knowledge, this is the first time that multiclass wavelength routing networks are analyzed.

The development of our approximate analytical techniques involves two steps. The arrival process of calls on some routes is first modified slightly to obtain a *modified multiclass network model*. Next, all classes of calls on a particular route are aggregated to give an *equivalent single-class model*. This equivalent single-class model has the same call blocking probability on any given route as the modified multiclass network, and can be easily solved due to the existence of a product-form solution. Simulation results further show that these results approximate the blocking probability of the original multiclass network with good accuracy. Thus, the solutions for single-class networks developed in Zhu et al. [1995b; 1999c] may be extended to the multiclass case.

In Section 2, we describe the multiclass wavelength routing network under study. In Section 3, we explain how the modified multiclass model and the equivalent single-class model are obtained for a single path of a wavelength routing network. In Section 4, we describe a decomposition algorithm for long paths and mesh network topologies. Section 5 gives numerical comparisons of performance measures obtained from the approximate model and from the exact model through simulations. We conclude the paper in Section 6.

2. THE MULTICLASS WAVELENGTH ROUTING NETWORK

We consider a wavelength routing network with an arbitrary topology. Each link in the network supports exactly W wavelengths, and each node is capable of transmitting and receiving on any of these W wavelengths. In our model, we allow some of the nodes in the network to employ wavelength converters. These nodes can switch an incoming wavelength to an arbitrary outgoing wavelength. (When there are converters at all nodes, the situation is identical to that in

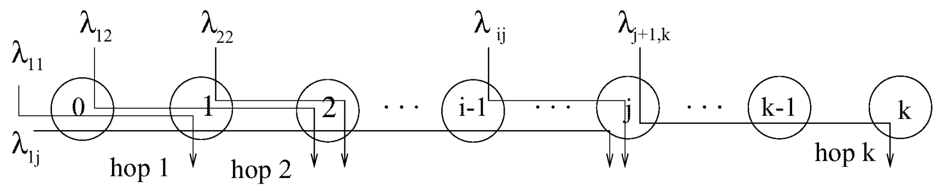

Fig. 1. A *k*-hop path.

classical circuit-switching networks, a special case of the more general scenario discussed here.) Call requests between a source and a destination node arrive at the source according to a Poisson process with a rate that depends on the source-destination pair. When a call arrives, a path over which the connection will be established first must be determined. In this article, we consider both fixed and alternate routing.

In fixed routing, each source-destination pair is assigned to a single path. If there are no wavelength converters in the path, a call is blocked if there is no wavelength that is free on all links of the path. If some nodes along the path employ converters, the call is blocked if no wavelength is free on all links of any segment of the path consisting of the links between successive nodes on the path with converters. This is known as the *wavelength continuity* requirement, and it increases the probability of call blocking. On the other hand, if a call can be accommodated, it is randomly assigned one of the wavelengths that are available on the links used by the call.[1] Thus, we only consider the random wavelength assignment policy.

In alternate routing, a set of paths (consisting of one primary path and one or more alternate paths) is assigned to each source-destination pair. Upon arrival of a call, this set is searched in a fixed order to find an available path for the call. If the request can be satisfied, an optical circuit is established between the source and destination for the duration of the call, which is referred to as the call holding time, or the call service time. Call holding times are assumed to be exponentially distributed, with a mean that depends on the class of the call.

Calls between any two nodes may be of several classes. In the general case, the number of classes of calls on a given route may not be the same as the number of classes on another route. However, to simplify the model, we can add some *dummy* classes of calls on any given route such that the arrival rate of calls of each dummy class is zero and the mean holding time is some finite quantity. Thus, without any loss of generality, we may assume that there are R classes of calls on each and every route. Each class of calls has its own Poisson arrival rate and exponential holding time.

Since many of our results are developed in terms of a single path in a wavelength routing network, we introduce some relevant notation. A k-hop path, such as the one shown in Figure 1, consists of $k+1$ nodes. Call these node 0, node 1, $\cdots$, node k. nodes $(i-1)$ and i, $1 \leq i \leq k$, are said to be connected by link i. A segment is a subpath that includes links i through j, $j \geq i \geq 1$. Calls

[1]In a path with wavelength converters, a wavelength is randomly assigned within each segment of the path whose starting and ending nodes are equipped with converters.

originating at node $i-1$ and terminating at node j use the above segment, which we denote by the pair (i, j). Calls between these two nodes may belong to one of R classes, and these calls are said to use *route* (i, j). We also define the following parameters.

—$\lambda_{ij}^{(r)}$, $j \geq i$, $1 \leq r \leq R$, is the Poisson arrival rate of calls of class r that originate at node $(i-1)$ and terminate at node j.

—$1/\mu_{ij}^{(r)}$ is the mean of the exponentially distributed service time of calls of class r that originate at node $(i-1)$ and terminate at node i. We also let $\rho_{ij}^{(r)} = \lambda_{ij}^{(r)}/\mu_{ij}^{(r)}$.

—$N_{ij}^{(r)}(t)$ is the number of active calls at time t on segment (i, j) belonging to class r.

—$F_{ij}(t)$ is the number of wavelengths that are free on all hops of segment (i, j) at time t. A call that arrives at time t and uses route (i, j) is blocked if $F_{ij}(t) = 0$.

3. BLOCKING PROBABILITIES IN A SINGLE PATH OF A NETWORK

3.1 The Single-Class Case

In this section we briefly review some of our previous results for a path of a single-class wavelength routing network. Consider the k-hop path shown in Figure 1. Let the state of this system at time t be described by the k^2-dimensional process:

$$\underline{X}_k(t) = (N_{11}(t), N_{12}(t), \cdots, N_{kk}(t), F_{12}(t), F_{13}(t), \cdots, F_{1k}(t), F_{23}(t), \cdots, F_{(k-1)k}(t)). \tag{1}$$

A closer examination of the process $\underline{X}_k(t)$ reveals that it is not time-reversible (see Zhu et al. [1999b]). This result is true in general, when $k \geq 2$ and $W \geq 2$.

The number of states of the Markov chain for the 2-hop case grows as W^4. In the k-hop case, the state-space grows as W^{k^2}. The computational complexity of a brute-force approach grows nearly as the cube of the size of the state-space. An exact solution would thus become computationally infeasible even for *modest* values of W and k.

Consequently, an approximate model was constructed in Zhu et al. [1999b] to analyze a single-class, k-hop path of a wavelength routing network. The approximation consists of modifying the call arrival process to obtain a time-reversible Markov process that has a closed-form solution. In a time-reversible Markov process, the steady-state probabilities of being in a given state can be expressed as the product of the marginal probabilities over the state variables. In other words, $P(n_{11}, n_{12}, n_{22}, f_{12})$ may be written as $H \times \Pi_{i,j} P_{ij}(n_{ij}) \times Q_{12}(f_{12})$, where $P_{ij}(.)$ represents the marginal distribution of $N_{ij}(t)$, $Q_{12}(.)$ represents the marginal distribution of $F_{12}(t)$, and H is some normalizing constant. Computing these marginal probabilities is fairly straightforward, as shown later, and requires constant time. Evaluating the normalizing constant for the k-hop case involves $O(W^{k(k+1)/2})$ steps. Thus, there is a considerable reduction in computational complexity in constructing a time-reversible Markov chain.

To illustrate our approach, let us consider the Markov process corresponding to a 2-hop path:

$$\underline{X}_2(t) = (N_{11}(t), N_{12}(t), N_{22}, F_{12}(t)). \tag{2}$$

We now modify the arrival process of calls that use both hops (a Poisson process with rate λ_{12} in the exact model) to a state-dependent Poisson process with rate Λ_{12} given by

$$\Lambda_{12}\,(n_{11}, n_{12}, n_{22},\, f_{12}) = \lambda_{12}\frac{f_{12}(W - n_{12})}{f_{11}\, f_{12}}. \tag{3}$$

The arrival process of other calls remains as in the original model. As a result, we obtain a new Markov process $\underline{X}'_2(t)$ with the same state space and the same state transitions as process $\underline{X}_2(t)$, but which differs from the latter in some of the state transition rates.

We made the observation in Zhu et al. [1999b] that under the new arrival process (3) for calls using both hops, the Markov process $\underline{X}'_2(t)$ is time-reversible and the stationary vector π is given by

$$\pi(n_{11}, n_{12}, n_{22},\, f_{12}) = \frac{1}{G_2(W)}\frac{\rho_{11}^{n_{11}}}{n_{11}!}\frac{\rho_{12}^{n_{12}}}{n_{12}!}\frac{\rho_{22}^{n_{22}}}{n_{22}!} \times \frac{\binom{f_{11}}{f_{12}}\binom{n_{11}}{W - n_{12} - n_{22} - f_{12}}}{\binom{W - n_{12}}{W - n_{12} - n_{22}}}, \tag{4}$$

where $G_k(W)$ is the normalizing constant for a k-hop path with W wavelengths.

Let $P(n_{11}, n_{12}, n_{22})$ be the marginal distribution over the states for which $N_{ij}(t) = n_{ij}$, $1 \le i \le j \le 2$. It can be verified [Zhu et al. 1999b] that

$$P(n_{11}, n_{12}, n_{22}) = \frac{1}{G_2(W)}\frac{\rho_{11}^{n_{11}}}{n_{11}!}\frac{\rho_{12}^{n_{12}}}{n_{12}!}\frac{\rho_{22}^{n_{22}}}{n_{22}!} \tag{5}$$

Likewise, for a k-hop path, $k \ge 2$, with the modified state-dependent Poisson arrival process, the marginal distribution over the states for which $N_{ij}(t) = n_{ij}$, $1 \le i \le j \le k$, is given by

$$P(n_{11}, n_{12}, \cdots, n_{kk}) = \frac{1}{G_k(W)} \prod_{\{(i,j)|1\le i\le j\le k\}} \frac{\rho_{ij}^{n_{ij}}}{n_{ij}!} \tag{6}$$

It is easily seen that this distribution is the same as in the case of a network with wavelength converters at each node. An interesting feature of having wavelength converters at every node is that the network has a product-form solution even when there are multiple classes of calls on each route, as long as call arrivals are Poisson, and holding times are exponential [Kelly 1979; Girard 1990]. Furthermore, when calls of all classes occupy the same number of wavelengths, we can aggregate classes to get an equivalent single class model with the same steady-state probability distribution over the aggregated states, as we show next.

3.2 The Multiclass Case

Let us now consider a k-hop path with wavelength converters at all nodes, and with R classes of calls. If $\lambda_{ij}^{(r)}$, $1 \le i \le j \le k$, $1 \le r \le R$, is the arrival rate of calls of class r on route (i, j), and $1/\mu_{ij}^{(r)}$, $1 \le i \le j \le k$, $1 \le r \le R$, is the mean of the exponential holding time of calls of class r, the probability of being in state $\underline{n} = (n_{11}^{(1)}, n_{11}^{(2)}, \cdots, n_{11}^{(R)}, n_{12}^{(1)}, \cdots, \cdots, n_{kk}^{(R)})$ is given by:

$$P(\underline{n}) = \frac{1}{G_k(W)} \left(\prod_{\{(i,j)|1 \le i \le j \le k\}} \prod_{r=1}^{R} \frac{\left(\rho_{ij}^{(r)}\right)^{n_{ij}^{(r)}}}{n_{ij}^{(r)}!} \right). \tag{7}$$

Let $\sigma_{ij} = \sum_r \rho_{ij}^{(r)}$ and $s_{ij} = \sum_r n_{ij}^{(r)}$. As defined, s_{ij} is the total number of calls of all classes that use segment (i, j) of the path, and σ_{ij} is the total offered load of these calls. Taking the summation of (7) over all states such that $\sum_r n_{ij}^{(r)} = s_{ij}$, $1 \le i \le j \le r$, we obtain:

$$P'(s_{11}, s_{12}, \cdots, s_{kk}) = \sum_{\{\underline{n}|\sum_r n_{ij}^{(r)} = s_{ij}\}} P(\underline{n}) = \frac{1}{G_k(W)} \prod_{\{(i,j)|1 \le i \le j \le k\}} \frac{\sigma_{ij}^{s_{ij}}}{s_{ij}!} \tag{8}$$

Observe that this is identical to the solution (6) for the single-class case obtained by substituting σ_{ij} by ρ_{ij} and s_{ij} by n_{ij} in (8).

Based on the above results, we conclude that by employing class aggregation on a multiclass path with converters at all nodes, we obtain a system equivalent to a single-class path with converters. In Section 3.1, we showed that the modified single-class wavelength routing network without converters has a steady-state marginal distribution similar to the exact single-class network with converters. We now show that a modified multiclass network without wavelength converters can also be subjected to class aggregation that results in an equivalent single-class model. The modification applied to the arrival process of calls is similar to the single-class case, and is given by expression (9):

$$\Lambda_{ij}^{(r)}(\underline{x}) = \lambda_{ij}^{(r)} \frac{f_{ij}\left(\sum_{l=1}^{i} s_{li} + f_{ii}\right)\left(\sum_{l=1}^{i+1} s_{l(i+1)} + f_{(i+1)(i+1)}\right) \cdots \left(\sum_{l=1}^{j-1} s_{l(j-1)} + f_{(m-1)(m-1)}\right)}{f_{ii}\, f_{(i+1)(i+1)} \cdots f_{jj}} \tag{9}$$

Then, the probability that the equivalent single class network without converters is in state

$$S_{\underline{x}} = \left(s_{11}, s_{12}, \cdots, s_{1k}, s_{22}, \cdots, s_{kk}, f_{12}, f_{13}, \cdots, f_{(k-1)k}\right) \tag{10}$$

is given by expression (11):

$$\pi(S_{\underline{x}}) =$$

$$\left(\prod_{i,j} \frac{s_{ij}!}{\sigma_{ij}^{s_{ij}}}\right) \prod_{l=2}^{k} \frac{\binom{f_{1(l-1)}}{f_{1l}} \left\{\prod_{m=2}^{l-1} \binom{f_{m(l-1)} - f_{(m-1)(l-1)}}{f_{ml} - f_{(m-1)l}}\right\} \binom{f_{ll} + n_{ll} - f_{(l-1)(l-1)}}{f_{ll} - f_{(l-1)l}}}{\binom{f_{ll} + n_{ll}}{f_{ll}}}. \tag{11}$$

Once again, the parameters of the single-class model are given by:

$$s_{ij} = \sum_{r=1}^{R} n_{ij}^{(r)} \quad 1 \le i < j \le k \tag{12}$$

$$\sigma_{ij} = \sum_{r=1}^{R} \rho_{ij}^{(r)} \quad 1 \le i < j \le k \tag{13}$$

3.3 Blocking Probabilities in the Multiclass Case

Since the arrival rate of calls of each class on each route is Poisson, the blocking probability, $Q_{ij}^{(r)}$, of a call of class r using route (i, j) is just the fraction of time that there is no wavelength that is free on all hops along route (i, j) (see the *PASTA* theorem in Wolff [1982]). Thus, we have:

$$Q_{ij}^{(r)} = \lim_{\tau\to\infty} \frac{\int_{t=0}^{\tau} I_{\{F_{ij}(t)=0\}} dt}{\tau}, \tag{14}$$

where

$$I_{\{F_{ij}(\tau)=0\}} = \begin{cases} 1 & \text{if } F_{ij}(\tau) = 0 \\ 0 & \text{otherwise} \end{cases}. \tag{15}$$

As can be seen, the blocking probability is class-independent.

Next, we focus on the call blocking probabilities in the modified model. The arrival process of calls of class r on route (i, j) is a state-dependent Poisson process whose rate at time τ, $A_{ij}^{(r)}(\tau)$ is a function of the state $\underline{X}(\tau)$ of the process, and is given by

$$A_{ij}^{(r)}(\tau) = \Lambda_{ij}^{(r)}(\underline{X}(\tau)) = \lambda_{ij} \frac{F_{ij}(\tau) \prod_{k=i}^{j-1} \left(\sum_{l=1}^{k} N_{lk}(\tau) + F_{kk}(\tau)\right)}{F_{ii} F_{i+1,i+1} \cdots F_{jj}}. \tag{16}$$

Note that the modified arrival process satisfies the criterion:

$$\frac{\Lambda_{ij}^{(r_1)}(\underline{x})}{\Lambda_{ij}^{(r_2)}(\underline{x})} = \frac{\lambda_{ij}^{(r_1)}}{\lambda_{ij}^{(r_2)}} \quad 1 \le r_1, r_2 \le r \tag{17}$$

By applying the *PASTA* theorem conditioned on being in state $\underline{x}$, the conditional call blocking probability, $\mathcal{P}_{ij}^{(r)}(\underline{x})$, of calls of class r on route (i, j) is given by the fraction of time spent in state $\underline{x}$ in which there is no wavelength that is free on all hops of route (i, j). Therefore:

$$\mathcal{P}_{ij}^{(r)}(\underline{x}) = \lim_{\tau\to\infty} \frac{\int_{t=0}^{\tau} I_{\{F_{ij}(t)=0, \underline{X}(t)=\underline{x}\}} dt}{\int_{t=0}^{\tau} I_{\{\underline{X}(t)=\underline{x}\}} dt} = \begin{cases} 1 & \text{if } f_{ij} = 0 \\ 0 & \text{otherwise} \end{cases}. \tag{18}$$

Let $P_{ij}^{(r)}$ be the unconditional probability that a call of class r, on route (i, j) gets blocked in the modified multiclass model. This is given by:

$$P_{ij}^{(r)} = \frac{\sum_{\underline{x}} \Lambda_{ij}^{(r)}(\underline{x})\pi(\underline{x})\mathcal{P}_{ij}^{(r)}(\underline{x})}{\sum_{\underline{x}} \Lambda_{ij}^{(r)}(\underline{x})\pi(\underline{x})} = \frac{\sum_{\{\underline{x}\mid f_{ij}=0\}} \Lambda_{ij}^{(r)}(\underline{x})\pi(\underline{x})}{\sum_{\underline{x}} \Lambda_{ij}^{(r)}(\underline{x})\pi(\underline{x})} \tag{19}$$

and can also be seen to be independent of the class r. Thus, by computing the blocking probability on the equivalent single-class path, we can obtain the solution to the multiclass path.

4. BLOCKING PROBABILITIES IN LONG PATHS AND MESH TOPOLOGIES

The solution to single-class wavelength routing networks involving paths with a large number of hops, and for networks with mesh topologies, with wavelength converters at an arbitrary subset of nodes, has been presented in Zhu et al. [1999b; 1999c]. Recall the evaluation of the normalizing constant of the product-form solution has a computational complexity given by $O(W^{(k(k+1)/2})$. When k is large, a brute-force computation of the normalizing constant becomes computationally infeasible. So, we apply path decomposition.

This solution involves decomposition of the network into short path segments with two or three hops, and analyzing these approximately using expression (4). The solutions to individual segments are appropriately combined to obtain a value for the blocking probability of calls that traverse more than one segment. The effect of the wavelength continuity requirement is captured by an approximate *continuity factor* that is used to increase the blocking probability of calls continuing to the next segment to account for the possible lack of common free wavelengths in the two segments. The process repeats until the blocking probabilities converge. By applying the transformations given in (12) and (13), the same algorithms may be used to calculate blocking probabilities for multiclass networks.

Specifically, we use the following steps to compute the blocking probabilities of a wavelength routing network with R classes of calls.

(1) *Path decomposition.* Decompose the multiclass mesh network topology into L single-path subsystems using the algorithm in Zhu et al. [1999c].

(2) *Time-reversible process approximation.* For each single-path subsystem, modify the arrival process as given by expression (9) to obtain an approximate time-reversible Markov process for the path.

(3) *Class aggregation.* For each subsystem, apply the transformations in (12) and (13) to obtain an equivalent single-class path subsystem.

(4) *Calculation of blocking probabilities.* For each path subsystem, obtain the blocking probabilities as follows. If the path is at most three hops long, use expressions (19) and (11) directly. If the path subsystem is longer than three hops, analyze it by decomposing it into 2- or 3-hop paths which are solved

Table I. Arrival and Service Parameters for a 2-Hop Path

Route	Class 1		Class 2		Class 3	
(i, j)	λ	μ	λ	μ	λ	μ
(1, 1)	3.0	8.0	4.0	8.0	3.0	6.0
(1, 2)	3.0	4.0	4.0	5.0	3.0	6.0
(2, 2)	5.0	8.0	5.0	5.0	3.0	6.0

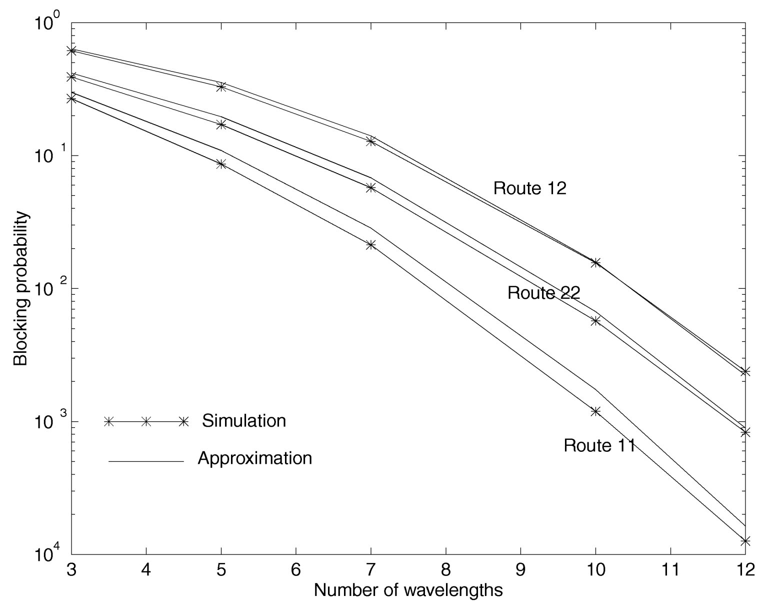

Fig. 2. Blocking probabilities for a 2-hop path and the parameters shown in Table I.

in isolation, and combine the individual solutions to obtain the blocking probabilities along the original longer path (see Zhu et al. [1999b]).

(5) *Convergence.* Repeat Steps 2 to 4, after appropriately modifying the original arrival rates to each single-path subsystem to account for the new values of the blocking probabilities obtained in Step 4 (see Zhu et al. [1999c]), until the blocking probabilities converge within a certain tolerance.

5. NUMERICAL RESULTS

In this section, we validate the approximate method described in Section 4 by comparing the blocking probabilities for each route as obtained from the approximate method with those obtained through simulation of the exact model.

5.1 A Single Path of a Network

We first provide results for 2- and 3-hop paths, since these systems represent the basic blocks of our decomposition algorithm. In Table I we show the arrival and service rates for calls on each route (i, j), $1 \leq j \leq 2$, of a 2-hop path. There are $R = 3$ classes of calls for each route. In Figure 2 we plot the blocking

Table II. Arrival and Service Parameters for a 3-Hop Path

Route	Class 1		Class 2		Class 3	
(i, j)	λ	μ	λ	μ	λ	μ
(1, 1)	3.0	3.0	1.0	3.0	1.0	3.0
(1, 2)	1.0	6.0	2.0	6.0	1.0	6.0
(1, 3)	1.0	2.0	1.0	2.0	2.0	2.0
(2, 2)	1.0	3.0	1.0	3.0	2.0	3.0
(2, 3)	1.0	2.0	1.0	2.0	3.0	2.0
(3, 3)	3.0	6.0	1.0	6.0	1.0	6.0

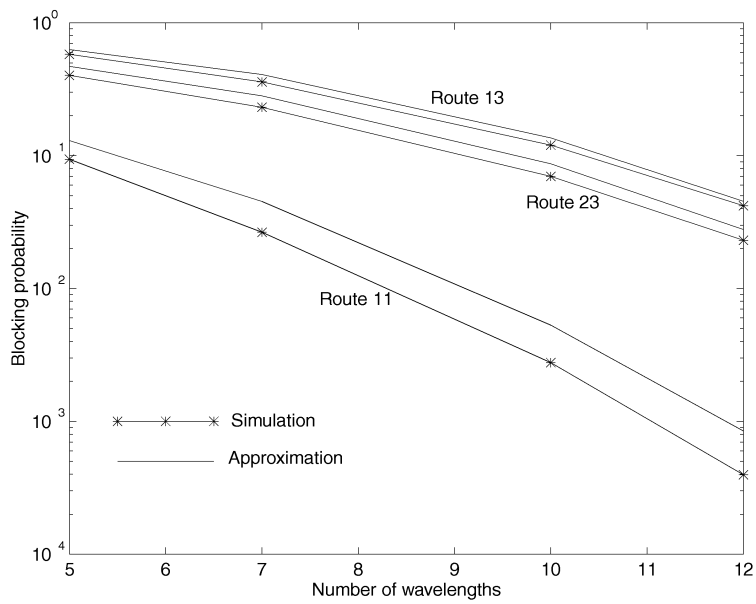

Fig. 3. Blocking probabilities for a 3-hop path and the parameters shown in Table II.

probability of calls along each of the three possible routes (recall that blocking probabilities are class-independent) against the number W of wavelengths supported by the links of the path. As we can see, the blocking probability decreases as W increases, as expected. We also observe that calls on Route (1, 2) (i.e., calls using both hops of the path) experience the highest blocking probability, since they have to compete against calls using either the first or the second hop of the path. Also, calls on Route (2, 2) (i.e., those using the second hop only) experience higher blocking probability than those using Route (1, 1), a direct consequence of the fact that the offered load of calls on Route (2, 2) is higher than that of calls on Route (1, 1) (refer to Table I). Most important, however, we can see that there is good agreement between the values of the blocking probabilities obtained through our analytical technique and those obtained through simulation. Similar results have been obtained for different values for the arrival and service parameters and for different numbers of classes, indicating that our approximate method is accurate over a wide range of network characteristics.

We now consider a 3-hop path with $R = 3$ classes of calls, and the arrival and service rates shown in Table II. In Figure 3, we plot the blocking probability

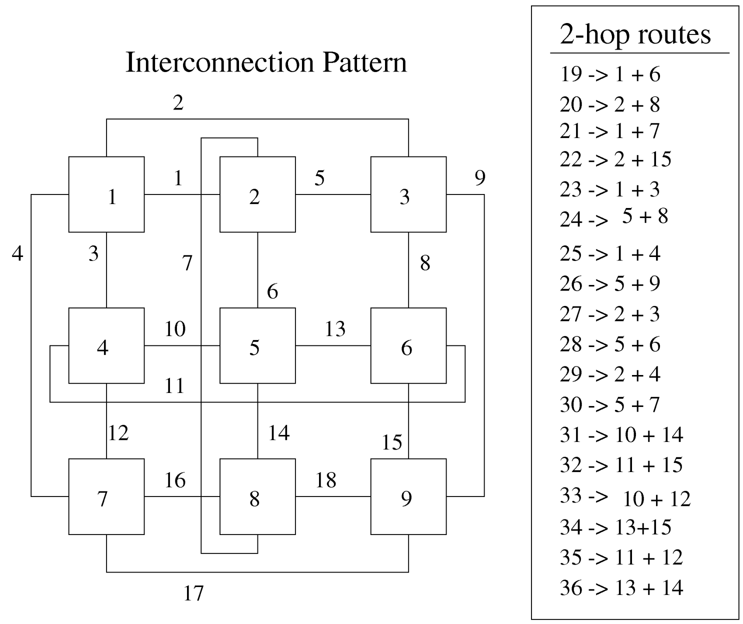

Fig. 4. The 3 × 3 torus network.

against the number W of wavelengths for three (out of the six) types of calls in this path: calls using Route (1, 1), (i.e., the first hop of the path), calls on Route (2, 3) (i.e., those using the last two hops), and calls on Route (1, 3) (using all three hops of the path). The behavior of the blocking probability curves in Figure 3 is similar to that in Figure 2. Specifically, the blocking probability decreases as W increases, while it increases with the number of hops a call must traverse. We also note that the curves derived analytically closely track those obtained by simulation.

From Figures 2 and 3 we conclude that for 2- and 3-hop multiclass paths (the building blocks of our decomposition algorithm), the class aggregation and the time-reversible Markov process (derived through the modification (9) of the arrival process) can be used to compute call blocking probabilities accurately. In the next subsection we demonstrate that the same is true for mesh topologies.

5.2 Mesh Topologies

In this section we provide results for two different mesh networks: a 3 × 3 regular torus topology, and the NSFNET irregular topology, shown in Figures 4 and 5, respectively.

Let us first consider the 3 × 3 network in Figure 4. Each of the nine nodes is directly linked with four others to form a toroid interconnection pattern. There are 72 source-destination pairs in the network, and 36 bidirectional routes. Each route requires at most two hops. Routes 1 through 18 comprise single-hop paths, while routes 19 through 36 consist of 2-hop paths, as in Figure 4. Once again we study this network with three classes of calls. The arrival and service rates of calls of a particular class on each route were assumed identical, and are given in Table III. We study the blocking probability as a function of the

Table III. Arrival and Service Parameters for the 3 × 3 Mesh

	Class 1	Class 2	Class 3
λ	0.4	0.6	0.5
μ	3.0	5.0	6.0

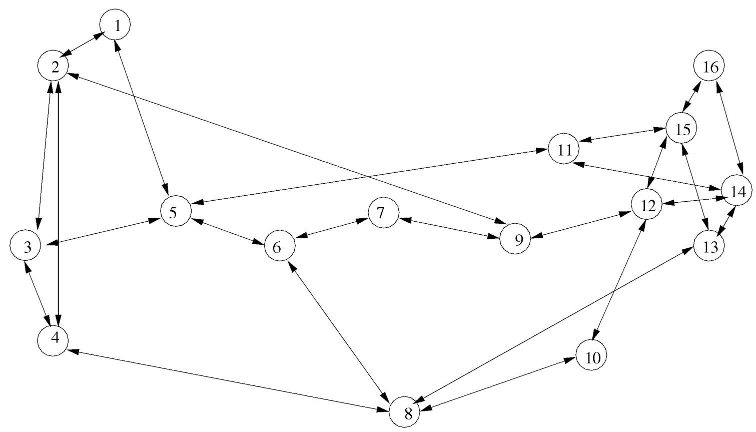

Fig. 5. The NSF network.

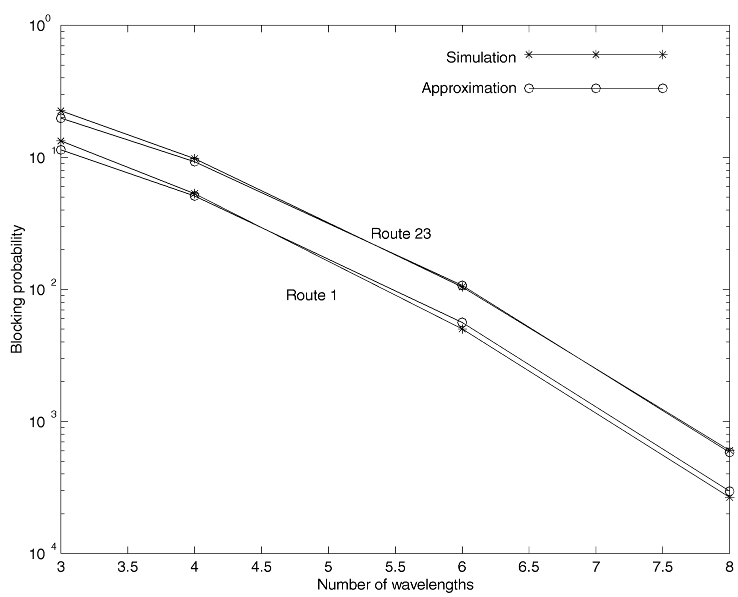

Fig. 6. Blocking probabilities for the 3 × 3 torus network.

number of wavelengths on each link for two routes: Route 1 (the single-hop path from node 1 to node 2) and Route 23 (the 2-hop path from node 2 to node 4). The approximation and simulation results are plotted in Figure 6, and we observe that they are in good agreement.

Next, we consider a network with topology similar to the NSF network, shown in Figure 5. There are 16 nodes, and 240 unidirectional routes. There are three

Table IV. Arrival and ServiceParameters for the NSF Network

	Class 1	Class 2	Class 3
λ	0.1	0.2	0.15
μ	3.0	6.0	4.0

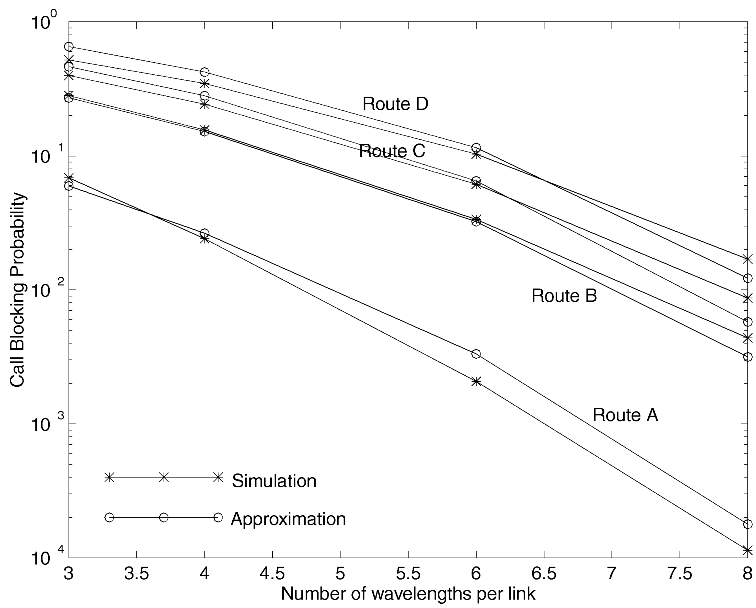

Fig. 7. Blocking probabilities for the NSF network.

classes of calls on each route. The arrival and service rates of calls of a particular class are the same on each route, and are shown in Table IV. The blocking probabilities are plotted in Figure 7 for four routes, as a function of the number of wavelengths on each link. Route A is a single-hop route from nodes 1 to 5. Route B has two hops, connecting node 1 to node 3 via node 2. Route C has three hops, connecting node 1 to node 4 via nodes 2 and 3. Route D has four hops, connecting node 4 to node 5 via nodes 3, 2, and 1.

From Figure 7 we can see that the length of the path used by a call considerably affects the blocking probability experienced by the call, an observation that is consistent with all our previous results in this section. Specifically, for a given number W of wavelengths, the blocking probability increases with the number of hops in a route, such that calls on Route A (a single-hop path) have the lowest blocking probability while calls on Route D (a four-hop path) the highest. Further, as in all previous topologies, the results indicate that our approximation method can be used to estimate accurately the blocking probabilities for all calls in the network.

Results similar to the ones presented in Figures 2 through 7 have been obtained for a wide range of traffic loads and different classes of calls. Our main conclusion is that our approximate analytical technique can be applied to compute the call blocking probabilities in wavelength routing networks of realistic size and topology.

6. CONCLUDING REMARKS

We have considered the problem of computing call blocking probabilities in multiclass wavelength routing networks which employ the random wavelength allocation policy. Trying to solve this problem by brute-force techniques is computationally infeasible. A straight path network with k hops, W wavelengths ($W^{R\times k^2}$) and R classes of calls could result in a state-space whose size grows. Our approach consists of modifying the call arrival process to obtain an approximate multiclass network model, using class aggregation to map this to an equivalent single-class network, and employing path decomposition algorithms on the latter to determine the call blocking probabilities (which are class-independent). The equivalent single-class model can be solved in $O(W^{k(k+1)/2})$ steps for a straight-path network with k hops, W wavelengths and R classes of calls. By applying path decomposition, we further break the network up into paths with only two or three hops. Thus, the computational complexity of the solution reduces to $O(k \times W^6)$. Our work is a first step towards realizing wavelength routing networks that can serve multiple classes of customers.

REFERENCES

Barry, R. A. and Humblet, P. A. 1996. Models of blocking probability in all-optical networks with and without wavelength changers. *IEEE J. Sel. Areas Commun. 14*, 5, 858–867.

Baynat, B. and Dallery, Y. 1996. A product-form approximation method for general closed queueing networks with several classes of customers. *Perform. Eval. 24*, 3, 165–188.

Birman, A. 1996. Computing approximate blocking probabilities for a class of all-optical networks. *IEEE J. Sel. Areas Commun. 14*, 5, 852–857.

Brackett, C. A.. 1996. Is there an emerging consensus on WDM networking? *J. of Lightw. Tech. 14*, 6(June), 936–941.

Chlamtac, I., Ganz, A., and Karmi, G. 1992. Lightpath communications: An approach to high bandwidth optical WANS. *IEEE Trans. Commun. 40*, 7, 1171–1182.

Mukherjee, B., Banerjee, D., Ramamurthy, S., and Mukherjee, A., 1996. Some principles for designing a wide-area WDM optical network. *IEEE/ACM Trans. Netw. 4*, 5, 684–696.

Hall, E. 1996. The Rainbow-II gigabit optical network. *IEEE J. Sel. Areas Commun. 14*, 5, 814–823.

Wagner, R. E. 1996. MONET: Multiwavelength optical networking. *J. Lightw. Tech. 14*, 6(June), 1349–1355.

Girard, A. 1990. *Routing and Dimensioning in Circuit-Switched Networks*. Addison-Wesley, Reading, MA.

Green, P. E. 1996. Optical networking update. *IEEE J. Sel. Areas Commun. 14*, 5, 764–779.

Harai, H., Murata, M., and Miyahara, H. 1997. Performance of alternate routing methods in all-optical switching networks. In *Proceedings of on IEEE INFOCOM 1997*(Kobe City, Japan, Apr.). IEEE Computer Society Press, Los Alamitos, CA, 517–525.

Karasan, E. and Ayanoglu, E. 1998. Effects of wavelength routing and selection algorithms on wavelength conversion gain in WDM optical networks. *IEEE/ACM Trans. Netw. 6*, 2, 186–196.

Kelly, F. P. 1979. *Reversibility and Stochastic Networks*. John Wiley and Sons, New York, NY.

Kovacevic, M. and Acampora, A. 1996. Benefits of wavelength translation in all-optical clear-channel networks. *IEEE J. Sel. Areas Commun. 14*, 5, 868–880.

Mokhtar, A. and Azizolu, M. 1998. Adaptive wavelength routing in all-optical networks. *IEEE/ACM Trans. Netw. 6*, 2, 197–206.

Pankaj, R. K. and Gallager, R. G. 1995. Wavelength requirements of all-optical networks. *IEEE/ACM Trans. Netw. 3*, 3(June), 269–280.

Ramamurty, B. and Mukherjee, B. 1998. Wavelength conversion in WDM networking. *IEEE J. Sel. Areas Commun. 16*, 7, 1061–1073.

RAMASWAMI, R. 1993. Multiwavelength lightwave networks for computer communication. *IEEE Commun. Mag.* (Feb.), 78–88.

RAMASWAMI, R. AND SIVARAJAN, K. N. 1996. Design of logical topologies for wavelength-routed optical networks. *IEEE J. Sel. Areas Commun. 14*, 5, 840–851.

SUBRAMANIAM, S., AZIZOGLU, M., AND SOMANI, A. K. 1996. All-optical networks with sparse wavelength conversion. *IEEE/ACM Trans. Netw. 4*, 4, 544–557.

SUBRAMANIAM, S., AZIZOGLU, M., AND SOMANI, A. K. 1998. On the optimal placement of wavelength converters in wavelength-routed networks. In *Proceedings on IEEE INFOCOM*. 902–909.

SUBRAMANIAN, S., SOMANI, M., AZIZOGLU, M., AND BARRY, R. A. 1997. A performance model for wavelength conversion with non-poisson traffic. In *Proceedings on IEEE INFOCOM 1997*(Kobe City, Japan, Apr.). IEEE Computer Society Press, Los Alamitos, CA, 500–507.

WAUTERS, N. AND DEMEESTER, P. 1996. Design of the optical path layer in multiwavelength cross-connected network. *IEEE J. Sel. Areas Commun. 14*, 5, 881–892.

WOLFF, R. W. 1982. Poisson arrivals see time averages. *Oper. Res. 30*, 2, 223–231.

ZHU, Y., ROUSKAS, G. N., AND PERROS, H. G. 1999. Blocking in wavelength routing networks, Part I: The single path case. In *Proceedings on INFOCOM.* IEEE Press, Piscataway, NJ, 321–328.

ZHU, Y., ROUSKAS, G. N., AND PERROS, H. G. 1999. Blocking in wavelength routing networks, Part II: Mesh topologies. In *Proceedings of the Sixteenth International on Teletraffic Congress* (ITC 16). Elsevier, Amsterdam, The Nethlands, 1321–1330.

ZHU, Y., ROUSKAS, G. N., AND PERROS, H. G. 1999. Bounds on the blocking performance of allocation policies in wavelength routing networks and a study of the effects of converters. TR-99-01.

Received September 1999; Revised April 2000; Accepted May 2000