Computing Blocking Probabilities in Multi-class Wavelength Routing Networks *

Sridhar Ramesh, George N. Rouskas, and Harry G. Perros

Department of Computer Science, North Carolina State University

Abstract. We present an approximate analytical method to compute efficiently 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 serve calls of different classes, with each call occupying one wavelength per link. Our approach involves two steps. The arrival process of calls on some routes is first modified slightly to obtain an approximate multi-class 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 networks may be extended to the multi-class case. Our work is a first step towards understanding the issues arising in wavelength routing networks that serve multiple classes of customers.

1 Introduction

A basic property of single mode optical fiber is its enormous low-loss bandwidth of several tens of Terahertz. However, due to dispersive effects and limitations in optical device technology, single channel transmission is limited 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, and WDM networks have been a subject of research both theoretically and experimentally [1]. Optical networks have the potential of delivering an aggregate throughput in the order of Terabits per second, and they appear as 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 Networks (WAN). The architecture consists of wavelength routers interconnected by fiber links. 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 [2], 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 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

* This work was supported by the NSF under grant ANI-9805016.

G. Pujolle et al. (Eds.): NETWORKING 2000, LNCS 1815, pp. 176–188, 2000.

to operate in circuit-switched mode. 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. Hence, wavelength routing WANs can be highly scalable.

Given the installed base of 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. A feasible scenario for near-term large-scale all-optical networks has emerged in [1], and it is envisioned that wavelength routing WANs will act as the backbone that provides interconnection for local area photonic sub-networks attached to them. The contribution of our work is the development of an approximate analytical framework for evaluating the performance of multi-class wavelength routing networks.

The problem of computing call blocking probabilities under static routing with random wavelength allocation and with or without converters has been studied in [3,4,5,6,7,8]. The model in [3] 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 does not capture the dynamic nature of traffic. In [4] it was assumed that statistics of link loads are mutually independent, an approximation that is not accurate for sparse network topologies. The work in [5] developed a Markov chain with state-dependent arrival rates to model call blocking in arbitrary mesh topologies and fixed routing; it was extended in [6] to alternate routing. A more tractable model was presented in [7] 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 [8], 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 [4], and it was shown to perform better than random allocation, while an analytical overflow model for first-fit allocation was developed in [9]. A dynamic routing algorithm that selects the least loaded path-wavelength pair was also studied in [9], and in [10] an unconstrained dynamic routing scheme with a number of wavelength allocation policies was evaluated. Except in [7,11], all other studies assume that either all or none of the wavelength routers have wavelength conversion capabilities. The work in [7] 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. A dynamic programming algorithm to determine the location of converters on a single path that minimizes average or maximum blocking probability was developed in [11] under the assumption of independent link loads.

Most of the approximate analytical techniques developed for computing blocking probabilities in wavelength routing networks [4,5,6,8,9,10,11] amount to the well-known *link decomposition* approach [12], 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 nodes where to employ converters.

We have developed an iterative *path decomposition* algorithm [13] for analyzing arbitrary network topologies. Specifically, we analyze a given network by decomposing it into a number of path sub-systems. These sub-systems are analyzed in isolation using our approximation algorithm for computing blocking probabilities in a single path of a network [14]. 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 where only a fixed but arbitrary subset of nodes are capable of wavelength conversion. Therefore, our algorithm can be an important tool in the development and evaluation of converter placement strategies. Also, in [15] 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 provide 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 have considered single-class wavelength routing networks. However, future networks will be utilized by a wide range of applications with varying characteristics in terms of their arrival rates and call holding times. In this paper, we present a method to extend the results in [14] and [13] to multi-class optical networks. To the best of our knowledge, this is the first time that multi-class wavelength routing networks are analyzed.

The development of our approximate techniques involves two steps. The arrival process of calls on some routes is first modified slightly to obtain a *modified multi-class network model*. Next, all classes of calls on a particular route are aggregated to give an *equivalent single-class model*. This equivalent model has the same call blocking probability on any given route as the modified multi-class network, and can be solved using our previous algorithms.

In Section 2, we describe the network under study. In Section 3, we explain how the modified multi-class model and the equivalent single-class model are obtained for a single path of a network. In Section 4, we describe a decomposition algorithm for mesh networks. Section 5 presents numerical results, and we conclude the paper in Section 6.

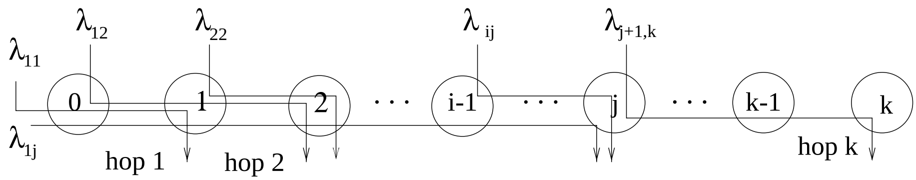

Fig. 1. A k-hop path

2 The Multi-class 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. 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. If the request can be satisfied, an optical circuit is established between the source and destination for the duration of the call. Further, calls between any two nodes may be of several classes. Without any loss of generality, we may assume that there are R classes of calls. Call holding times are assumed to be exponentially distributed, with a mean that depends on the class of the call.

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 classical circuit-switching networks, a special case of the more general scenario discussed here.) If no wavelength converters are employed in the path between the source and the destination, a call can only be established if the *same* wavelength is free on all the links used by the call. If there exists no such wavelength, the call is blocked. This is known as the *wavelength continuity* requirement (see [13,14]), 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 this paper.

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 (see Figure 1) consists of $k+1$ nodes. Nodes $(i-1)$ and i, $1 \le i \le k$, are said to be connected by link (hop) i. Calls originating at node $i-1$ and terminating at node j use hops i through $j, j \ge i \ge 1$, which we shall 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.

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

[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.

- $1/\mu_{ij}^{(r)}$ is the mean of the exponentially distributed service time of calls of class r that originate at node $(i-1)$ and terminate at node j. 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:

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

A closer examination of the process $\underline{X}_k(t)$ reveals that it is not time-reversible (see [14]). This result is true in general, when $k \geq 2$ and $W \geq 2$.

Since the number of states of process $\underline{X}_k(t)$ grows very fast with the number k of hops in the path and the number W of wavelengths, it is not possible to obtain the call blocking probabilities directly from the above process. Consequently, an approximate model was constructed in [14] 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. To illustrate our approach, let us consider the Markov process corresponding to a 2-hop path:

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

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:

$$\Lambda_{12}\left(n_{11}, n_{12}, n_{22}, f_{12}\right) = \lambda_{12} \frac{f_{12}(W - n_{12})}{f_{11} f_{12}} \quad (3)$$

The arrival process of other calls remain 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 [14] 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:

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

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 $N_{ij}(t) = n_{ij}$, $1 \leq i \leq j \leq 2$. It can be verified [14] 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, $k \geq 2$, with the modified state-dependent Poisson arrival process, the marginal distribution over the states for which $N_{ij}(t) = n_{ij}$, $1 \leq i \leq j \leq k$, is given by:

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

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 [16,12]. Further, 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 Multi-class Case

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

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

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 $\sum_r n_{ij}^{(r)} = s_{ij}$, $1 \leq i \leq j \leq r$, we obtain:

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

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 multi-class 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 multi-class 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 it is given by:

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

Then, the probability that the equivalent single-class network without converters is in state $S_{\underline{x}} = \left(s_{11}, s_{12}, \cdots, s_{1k}, s_{22}, \cdots, s_{kk}, f_{12}, f_{13}, \cdots, f_{(k-1)k}\right)$ is given by:

$$\pi\left(S_{\underline{x}}\right) = \left(\prod_{i,j} \frac{s_{ij}!}{\sigma_{ij}^{s_{ij}}}\right)\left(\prod_{l=2}^{k} \frac{\binom{f_{1(l-1)}}{f_{1l}}\left\{\prod_{m=2}^{l-1}\binom{f_{m(l-1)}-f_{(m-1)(l-1)}}{f_{ml}-f_{(m-1)l}}\right\}\binom{f_{ll}+n_{ll}-f_{(l-1)(l-1)}}{f_{ll}-f_{(l-1)l}}}{\binom{f_{ll}+n_{ll}}{f_{ll}}}\right) \quad (10)$$

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

$$s_{ij} = \sum_{r=1}^{R} n_{ij}^{(r)} \quad 1 \le i < j \le k, \qquad \sigma_{ij} = \sum_{r=1}^{R} \rho_{ij}^{(r)} \quad 1 \le i < j \le k \quad (11)$$

3.3 Blocking Probabilities in the Multi-class 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 are no wavelengths that are free on all hops along route (i,j) (see the *PASTA* theorem in [17]). Thus, we have:

$$Q_{ij}^{(r)} = \lim_{\tau\to\infty} \frac{\int_{t=0}^{\tau} I_{\{F_{ij}(t)=0\}}dt}{\tau}, \quad \text{where} \quad I_{\{F_{ij}(\tau)=0\}} = \begin{cases} 1, \text{ if } F_{ij}(\tau)=0 \\ 0, \text{ otherwise} \end{cases} \quad (12)$$

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:

$$A_{ij}^{(r)}(\tau)\Lambda_{ij}^{(r)}\left(\underline{X}(\tau)\right) = \lambda_{ij} \frac{F_{ij}(\tau)\prod_{k=i}^{j-1}\left(\sum_{l=1}^{k} N_{lk}(\tau)+F_{kk}(\tau)\right)}{F_{ii}F_{i+1,i+1}\cdots F_{jj}} \quad (13)$$

Note that the modified arrival process satisfies the criterion:

$$\frac{\Lambda_{ij}^{(r_1)}(\underline{x})}{\Lambda_{ij}^{(r_2)}(\underline{x})} = \frac{\lambda_{ij}^{(r_1)}}{\lambda_{ij}^{(r_2)}} \quad 1 \leq r_1, r_2 \leq r \tag{14}$$

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 are no wavelengths that are free on all hops of route (i,j). Therefore:

$$\mathcal{P}_{ij}^{(r)}(\underline{x}) = \lim_{\tau \to \infty} \frac{\int_{t=0}^{\tau} I_{\{F_{ij}(t)=0, \underline{X}(t)=\underline{x}\}} dt}{\int_{t=0}^{\tau} I_{\{\underline{X}(t)=\underline{x}\}} dt} = \begin{cases} 1, & \text{if} \quad f_{ij} = 0 \\ 0, & \text{otherwise} \end{cases} \tag{15}$$

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

$$P_{ij}^{(r)} = \frac{\sum_{\underline{x}} \Lambda_{ij}^{(r)}(\underline{x}) \pi(\underline{x}) \mathcal{P}_{ij}^{(r)}(\underline{x})}{\sum_{\underline{x}} \Lambda_{ij}^{(r)}(\underline{x}) \pi(\underline{x})} = \frac{\sum_{\{\underline{x} | f_{ij}=0\}} \Lambda_{ij}^{(r)}(\underline{x}) \pi(\underline{x})}{\sum_{\underline{x}} \Lambda_{ij}^{(r)}(\underline{x}) \pi(\underline{x})} \tag{16}$$

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 multi-class path.

4 Blocking Probabilities in Mesh Topologies

The solution to single-class networks with wavelength converters at an arbitrary subset of nodes has been presented in [14,13]. 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 in (11), the same algorithms may be used to calculate blocking probabilities for multi-class networks. Specifically, we use these steps for a network with R classes of calls:

1. **Path decomposition:** Decompose the multi-class mesh network topology into L single-path sub-systems using the algorithm in [13].
2. **Time-reversible process approximation:** For each single-path sub-system, 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 sub-system, apply the transformations in (11) to obtain an equivalent single-class path sub-system.
4. **Calculation of blocking probabilities:** For each path sub-system, obtain the blocking probabilities as follows. If the path is at most three hops long, use expressions (16) and (10) directly. If the path sub-system is longer than three hops, analyze it by decomposing it into 2- or 3-hop paths which are solved in isolation, and combine the individual solutions to obtain the blocking probabilities along the original longer path (see [14]).
5. **Convergence:** Repeat Steps 2 to 4, after appropriately modifying the original arrival rates to each single-path sub-system to account for the new values of the blocking probabilities obtained in Step 4 (see [13]), 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.

We first provide results for 3-hop paths which are the basic blocks of our decomposition algorithm (results for 2-hop paths can be found in [18]). In Table 1 we show the arrival and service rates for calls on each route $(i, j), 1 \leq j \leq 3$, of a 3-hop path. There are $R = 3$ classes of calls for each route. In Figure 2, we plot the blocking probability 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 $(2, 3)$, that is, those using the last two hops, and calls on Route $(1, 3)$ using all three hops of the path. As we can see, the blocking probability decreases as W increases, as expected. We also observe that calls on Route $(1, 3)$ (i.e., calls using all three hops of the path) experience the highest blocking probability, again as expected. Most importantly, 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 number of classes, indicating that our approximate method is accurate over a wide range of network characteristics.

Next, we consider a network with topology similar to the NSF network, shown in Figure 3. There are 16 nodes, and 240 uni-directional routes. There are three 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 2. The blocking probabilities are plotted in Figure 4 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 4 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

Table 1. 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

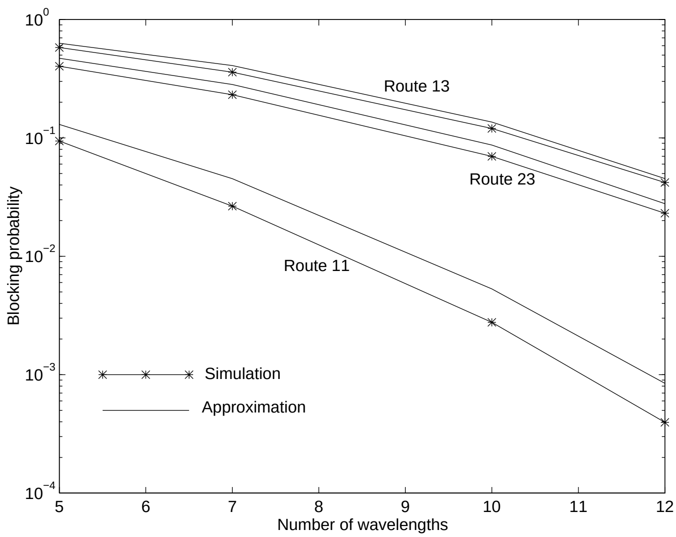

Fig. 2. Blocking probabilities for a 3-hop path and the parameters shown in Table 1

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, 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 2,4 have been obtained for a wide range of traffic loads and different classes of calls, and for other network topologies (see [18]). 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. The approximate technique affords a significant reduction in the time for computation of blocking probabilities. For instance, the simulation program took approximately 100 minutes while running on a Sun-sparc Ultra 10 workstation, to compute call blocking

Table 2. Arrival and service parameters for the NSF Network

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

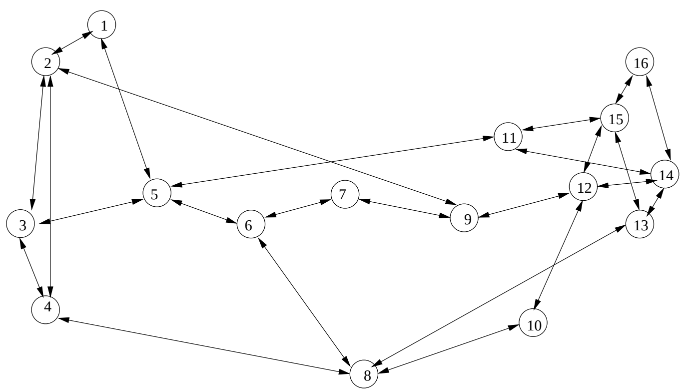

Fig. 3. The NSF network

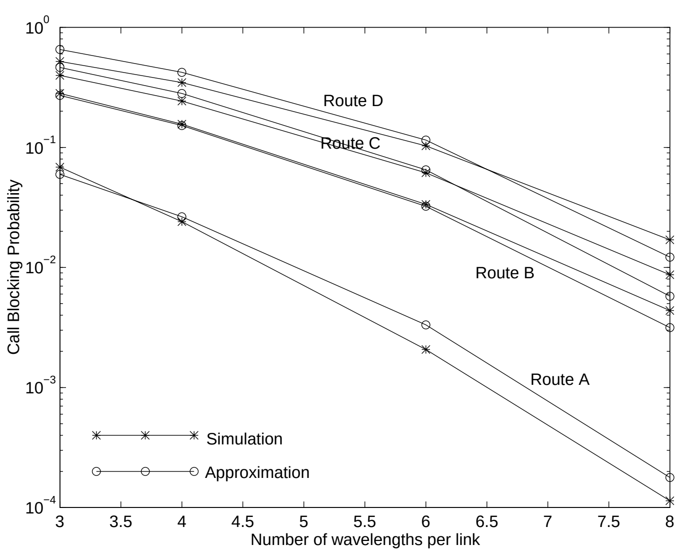

Fig. 4. Blocking probabilities for the NSF network

probabilities for the network with topology similar to the NSF network. The approximation method took less than a minute for the same network. In addition, the simulation program runs until there are at least 100000 call arrivals of every class. For call blocking probabilities less than 10^{-4}, the simulation program must consider greater than 10^6 call arrivals of each class, resulting in an even longer computation time.

6 Concluding Remarks

We have considered the problem of computing call blocking probabilities in multi-class wavelength routing networks which employ the random wavelength allocation policy. Our approach consists of modifying the call arrival process to obtain an approximate multi-class 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.

References

1. A. Hill, A. Salek, and K. Sato (Eds.). Special issue on high-capacity optical transport networks. *IEEE JSAC*, 16(7), Sep. 1998.
2. B. Ramamurty and B. Mukherjee. Wavelength conversion in WDM networking. *IEEE JSAC*, 16(7):1061–1073, Sep. 1998.
3. R. Barry and P. Humblet. Models of blocking probability in all-optical networks with and without wavelength changers. *IEEE JSAC*, 14(5):858–867, June 1996.
4. M. Kovacevic and A. Acampora. Benefits of wavelength translation in all-optical clear-channel networks. *IEEE JSAC*, 14(5):868–880, June 1996.
5. A. Birman. Computing approximate blocking probabilities for a class of all-optical networks. *IEEE JSAC*, 14(5):852–857, June 1996.
6. H. Harai, M. Murata, and H. Miyahara. Performance of alternate routing methods in all-optical switching networks. *Proc. INFOCOM '97*, pp. 517–525, April 1997.
7. S. Subramaniam, M. Azizoglu, and A. Somani. All-optical networks with sparse wavelength conversion. *IEEE/ACM Trans. Network.*, 4(4):544–557, Aug. 1996.
8. S. Subramanian *et al.* A performance model for wavelength conversion with non-poisson traffic. *Proc. INFOCOM '97*, pp. 500–507, April 1997.
9. E. Karasan and E. Ayanoglu. Effects of wavelength routing and selection algorithms on wavelength conversion gain in WDM optical networks. *IEEE/ACM Trans. Network.*, 6(2):186–196, April 1998.
10. A. Mokhtar and M. Azizoglu. Adaptive wavelength routing in all-optical networks. *IEEE/ACM Trans. Network.*, 6(2):197–206, April 1998.
11. S. Subramaniam, M. Azizoglu, and A. K. Somani. On the optimal placement of wavelength converters in wavelength-routed networks. *Proc. INFOCOM '98*, pp. 902–909, April 1998.
12. A. Girard. *Routing and Dimensioning in Circuit-Switched Networks.* Addison Wesley, Reading, MA, 1990.
13. Y. Zhu, G. N. Rouskas, and H. G. Perros. Blocking in wavelength routing networks, Part II: Mesh topologies. *Proc. ITC 16*, pp. 1321–1330, June 1999.
14. Y. Zhu, G. N. Rouskas, and H. G. Perros. Blocking in wavelength routing networks, Part I: The single path case. *Proc. INFOCOM '99*, pp. 321–328, March 1999.

15. Y. Zhu, G. N. Rouskas, and H. G. Perros. Bounds on the blocking performance of allocation policies in wavelength routing networks and a study of the effects of converters. TR-99-01, NCSU, Jan. 1999.
16. F. P. Kelly. *Reversibility and Stochastic Networks.* John Wiley & Sons, NY, 1979.
17. R. W. Wolff. Poisson arrivals see time averages. *Oper. Res.*, 30(2):223–231, 1982.
18. S. Ramesh, G. N. Rouskas, and H. G. Perros. Computing call blocking probabilities in multi-class wavelength routing networks. TR-99-08, NCSU, May 1999.