



A Comparison of Allocation Policies in Wavelength Routing Networks*

Yuhong Zhu, George N. Rouskas, Harry G. Perros

Department of Computer Science, North Carolina State University, Raleigh, NC 27695-7534

Received March 15, 2000; Revised April 28, 2000

Abstract. We consider wavelength routing networks with and without wavelength converters, and several wavelength allocation policies. Through numerical and simulation results we obtain upper and lower bounds on the blocking probabilities for two wavelength allocation policies that are most likely to be used in practice, namely, most-used and first-fit allocation. These bounds are the blocking probabilities obtained by the random wavelength allocation policy with either no converters or with converters at all nodes of the network. Furthermore, we demonstrate that using the most-used or first-fit policies gives an improvement on call blocking probabilities that is equivalent to employing converters at a number of nodes in a network with the random allocation policy. These results have been obtained for a wide range of loads for both single-path and general mesh topology networks. The main conclusion of our work is that the gains obtained by employing specialized and expensive hardware (namely, wavelength converters) can be realized cost-effectively by making more intelligent choices in software (namely, the wavelength allocation policy).

Keywords: wavelength division multiplexing, wavelength routing networks, call blocking probability, wavelength allocation

1 Introduction

Recent advances in wavelength division multiplexing (WDM) and optical switching make it possible to contemplate the deployment of wavelength routing networks that will provide backbone connectivity over wide-area distances and at very high data capacities [8,4]. A wavelength routing network consists of wavelength routers and the fiber links that interconnect them [6,11,7]. Wavelength routers are optical switches capable of routing a light signal at a given wavelength from any input port to any output port, making it possible to establish end-to-end lightpaths, i.e., direct optical connections without any intermediate electronics. The functionality of optical switches may be enhanced by employing wavelength converters, devices that are capable of shifting an incoming wavelength to a different outgoing wavelength [15]. Wavelength conversion is a desirable feature since it improves the performance of the network in terms of call blocking probability. However, this gain in performance must be weighted against the cost of wavelength converters.

While the operation of wavelength routing networks is expected to be similar to that of conventional

circuit-switched networks, several new issues arise which add significant complexity to the problems of design and performance evaluation of the former. Specifically, the existence of multiple distinct wavelengths makes it necessary to employ a wavelength allocation policy to assign an available wavelength to an incoming call. Similarly, the wavelength conversion feature gives rise to new problems associated with evaluating the benefits of conversion and optimally placing the converters at the various network nodes. Also, dynamic (or adaptive) routing is tightly coupled with wavelength allocation, since it involves a search over the available wavelengths in addition to a search over the possible paths for establishing a call.

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 [1,13,2,10,16,18]. The model presented in Barry and Humblet [1] is based on the assumption that the 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 Kovacevic and Acampora [13] it

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

was assumed that the statistics of link loads are mutually independent, an approximation that is not accurate for sparse network topologies. In [2] a Markov chain with state-dependent arrival rates was developed to model call blocking in arbitrary mesh topologies and fixed routing; this technique was extended to alternate routing in [10]. 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 [16] to compute recursively the 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 [18], 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 [3,13], and it was shown to perform better than random allocation, while an analytical overflow model for first-fit allocation was developed in [12]. A dynamic routing algorithm that selects the least loaded path-wavelength pair was also studied in [12], and in [14] an unconstrained dynamic routing scheme with a number of wavelength allocation policies was evaluated. Except in [16,17], all other studies assume that either all or none of the wavelength routers have wavelength conversion capabilities. The work in [16] 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 [17] 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 [13,2,10,18,12,14,17] make the assumption that link blocking events are independent and amount to the well-known link

decomposition approach [9]. Also, the development of some other 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, we believe that more accurate analytical tools are needed to efficiently evaluate the performance of these networks, as well as to overcome complex network design problems.

The authors have considered the problem of computing call blocking probabilities in meshed wavelength routing networks with fixed and alternate routing and random wavelength allocation in [20]. Unlike previous studies, we have developed an iterative path decomposition algorithm for analyzing arbitrary network topologies. Specifically, we analyze a given network by decomposing it into a number of single path sub-systems. These sub-systems are then analyzed in isolation using our algorithm for calculating the blocking probabilities in a single path in a wavelength routing network [19]. The individual solutions are appropriately combined to form a solution for the overall network. This 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 non-uniform traffic, in the sense that call request arrival rates can vary for each source-destination pair. Finally, our algorithms can compute call blocking probabilities in a mesh network where only a subset of arbitrarily selected nodes are capable of wavelength conversion.

In this paper, we study the blocking performance of several wavelength allocation policies for various network topologies and traffic patterns. Our main results are as follows. First, we show that the most-used and first-fit policies have very similar call blocking probabilities for all calls in a network, regardless of the number of hops used by the calls. Through numerical and simulation results we obtain upper and lower bounds on the blocking probabilities for two wavelength allocation policies that are most

likely to be used in practice, namely, most-used and first-fit allocation. These bounds are the blocking probabilities obtained by the random wavelength allocation policy with either no converters or with converters at all nodes of the network. This result is important since, for the random policy with or without converters, efficient analytical solutions have been developed for networks of large size. We also present results which indicate that the call blocking probabilities of the first fit and most-used policies is similar to that of the random policy when a number of converters is employed in the network. Overall, our results contradict previous studies, which have only concentrated on random wavelength allocation, and in which it was suggested that sparse wavelength conversion is beneficial to wavelength routing networks. Although we have identified regions of operation where converters do offer significant benefits, the regions are characterized by very high call blocking probabilities, and it is unlikely that networks will be designed to operate in these regions.

In Section 2 we study a single path in a wavelength routing network, and in Section 3 we consider both regular and irregular mesh network topologies. We conclude with a summary of our findings in Section 4.

2 A Single Path of a Wavelength Routing Network

We consider a single path of a wavelength routing network, such as the k -hop path shown in Fig. 1. A k -hop path consists of $k + 1$ nodes labeled $0, 1, \dots, k$, and hop $i, i = 1, \dots, k$, represents the link between nodes $i - 1$ and i . (Unless noted otherwise, the terms ‘hop’ and ‘link’ will be used interchangeably.) Each link in the path supports exactly W wavelengths, and each node is capable of transmitting and receiving

on any of the W wavelengths. We assume that calls arrive as a Poisson process. Let $\lambda_{ij}, j \geq i$, denote the arrival rate of calls that use hops i through j of the path, i.e., calls that originate at node $i - 1$ and terminate at node j . For instance, λ_{22} is the arrival rate of calls that only use hop 2 (that is, those arriving at node 1 and leaving at node 2), while λ_{12} is the arrival rate of calls using hops 1 and 2 (refer to Fig. 1). If the request can be satisfied, an optical circuit is established between the source and destination for the duration of the call. Call holding times are exponentially distributed with mean $1/\mu$. Also, let $\rho_{ij} = \lambda_{ij}/\mu$ denote the offered load of calls using hops i through j .

We define a ‘segment’ of a k -hop path as a sub-path consisting of one or more consecutive links of the original path. We let $n_{ij}, j \geq i$, be a random variable representing the number of calls using hops i through j that are currently active. We also let $f_{ij}, j \geq i$, be a random variable representing the number of wavelengths that are free on all hops i through j . We shall see shortly that random variables n_{ij} and f_{ij} are part of the state description of the Markov process corresponding to the k -hop path.

Some of the nodes in the path can be equipped with a wavelength converter. These nodes can switch an incoming wavelength to an arbitrary outgoing wavelength. If no wavelength converters are employed in the path, a call can only be established if the same wavelength is free on all the links used by the call. This is known as the wavelength continuity requirement, and it increases the probability of blocking for calls using multiple hops. If a call cannot be established due to lack of available wavelengths, the call is blocked. On the other hand, if a call can be accommodated, it is assigned one of the wavelengths that are available on the links used by the call. If there are multiple wavelengths available, a

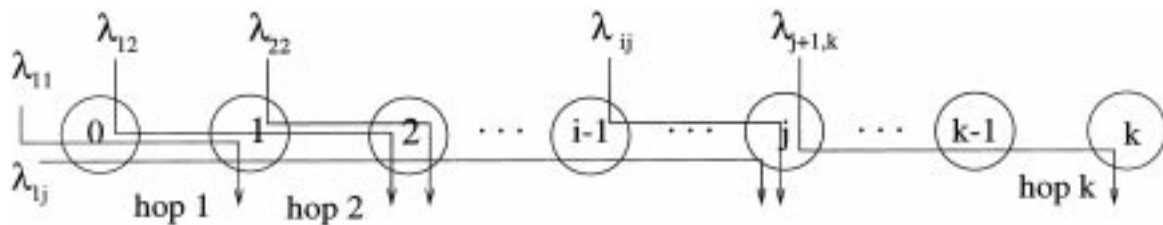


Fig. 1. A k -hop path.

wavelength allocation policy must be employed to select a wavelength for the call. Different selection policies lead to different call blocking probabilities. In this paper we investigate the following four wavelength allocation policies:

- *Random allocation*: a call is randomly assigned to one of the wavelengths that are available on all the links that will be used by the call. This policy has been extensively studied in the literature. As mentioned in the introduction, we have developed approximate analytical algorithms to evaluate the call blocking performance for a single path and mesh topologies [19,20].
- *Most-used allocation*: among the free wavelengths in the path, the one that is already in use on the largest number of links in the network is assigned to the call; ties are broken arbitrarily. The objective of the policy is to keep more wavelengths available for calls traveling over long paths.
- *Least-used allocation*: among the free wavelengths in the path, the call is assigned to the one that is currently used in the smallest number of links in the network, with ties broken arbitrarily. Intuitively, this policy results in “wavelength fragmentation,” and leads to higher blocking probability for calls traveling over long paths.
- *First-fit allocation*: the wavelengths on each link are given labels in a fixed order, and the call is assigned to the wavelength with the smallest label that is available on all the links it requires. The objective of this allocation scheme is to minimize wavelength fragmentation. As we shall show later, its performance is very close to that of the most-used policy, but it is easier to implement since there is no need to maintain information on the global use of wavelengths.

In a path with wavelength converters, each of the above allocation policies is used in order to assign a wavelength to the call within each segment of the path whose starting and ending nodes are equipped with converters. In addition to the above wavelength allocation policies, we will also consider the following case:

- *All-converter paths*: paths in which there are converters at all nodes. In this case, a call can be

established as long as at least one wavelength (not necessarily the same one) is free on each of the links required by the call, in a manner similar to conventional circuit-switching. Consequently, wavelength allocation is not an issue under an all-converter scenario, and all allocation policies, including the ones studied here, reduce to random allocation within each link.

In our study, we will use six different traffic load patterns to compare the four wavelength allocation policies against each other and against the all-converter scenario. The six patterns are representative of the wide range of loading situations that one expects to encounter in practice. Figs. 2 and 3 illustrate the six traffic patterns for a 10-hop path. Specifically, the figures plot the load ρ_l of each hop l , $l = 1, \dots, 10$, in the path, defined as the sum of the offered loads ρ_{ij} , $i \leq l \leq j$, for all calls that use hop l , for each load pattern. In the “uniform” pattern, all hops are equally loaded. The “bowl” (respectively, “inverted bowl”) pattern is such that the load decreases (respectively, increases) from hop 1 to hop 5, and then it increases (respectively, decreases) from hop 6 to hop 10. These patterns are shown in Fig. 2. The “ascending” and “descending” patterns are such that the load increases or decreases, respectively, from hop 1 to hop 10. Finally, in the oscillating pattern the load at each hop alternates between a low and a high value. The last three load patterns are shown in Fig. 3. To ensure that the results are comparable across the different patterns, the load values were chosen so that the total load (or, equivalently, the average load per hop) is the same for all patterns.

2.1 Policy Comparison for 2-Hop Paths

We will first study the blocking probabilities of the wavelength allocation policies for 2-hop paths, as shown in Fig. 4. The state space of these systems is small enough so that we can obtain exact numerical solutions for the call blocking probabilities.

2.1.1 Exact and Approximate Markov Processes

We have shown in [19] that the evolution of a 2-hop path with random wavelength allocation can be characterized by the four-dimensional Markov process $(n_{11}, n_{12}, n_{22}, f_{12})$. The first three random variables in the state description provide the number of active calls between the three source-destination pairs in the path, and the last random variable gives

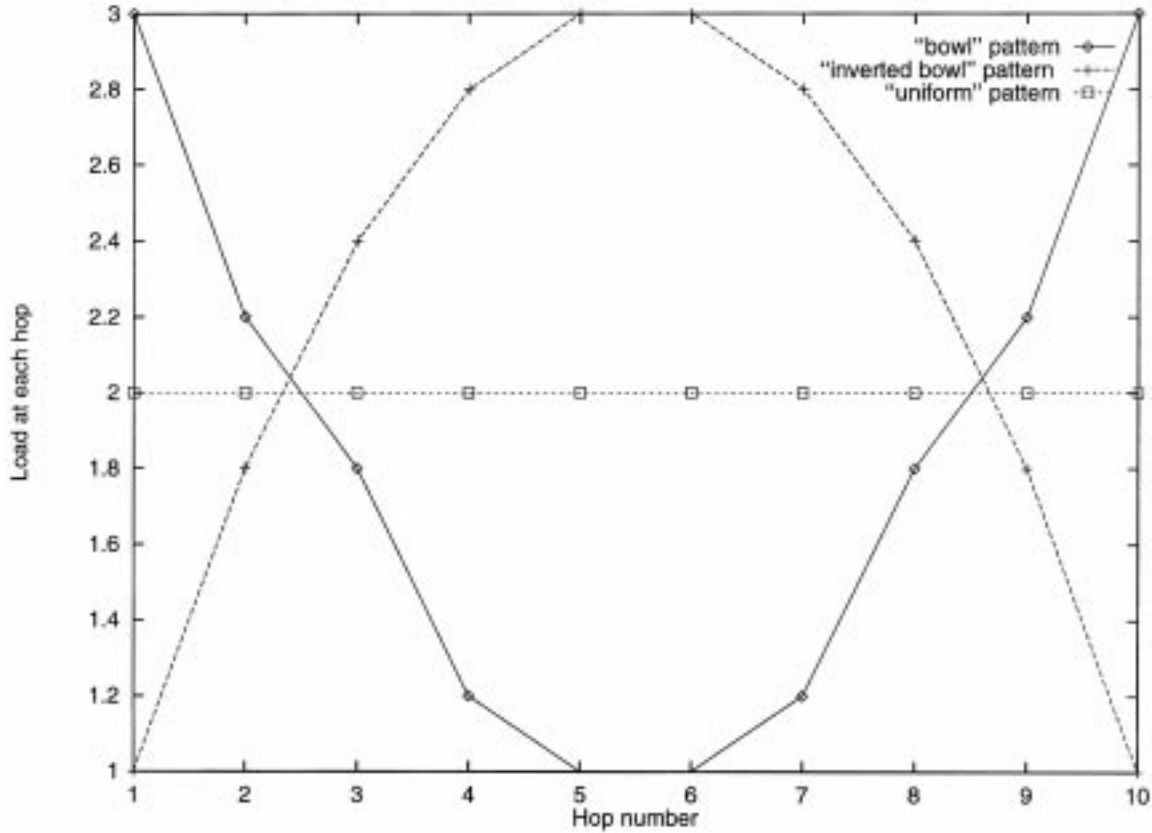


Fig. 2. The “bowl”, “inverted bowl”, and “uniform” load patterns.

the number of wavelengths that are free on both links of the path. The state transition diagram of this Markov process is shown in Fig. 5 for $W = 2$ wavelengths, and it is straightforward to see that the process is not time-reversible [19]. By modifying a few of the transition rates of this process, we were able to derive a time-reversible Markov process with the same state space, which has a product-form solution. If we let $G(W)$ denote the normalizing constant for a 2-hop path with W wavelengths per link, the solution of the approximate Markov process is given by [19]:

$$\pi_{\text{random}}(n_{11}, n_{12}, n_{22}, f_{12}) = \frac{1}{G(W)} \frac{\rho_{11}^{n_{11}} \rho_{12}^{n_{12}} \rho_{22}^{n_{22}}}{n_{11}! n_{12}! n_{22}!} \frac{\binom{f_{11}}{f_{12}} \binom{n_{11}}{f_{22} - f_{12}}}{\binom{n_{11} + f_{11}}{f_{22}}} \quad (1)$$

where $f_{11} = W - n_{11} - n_{12}$ and $f_{22} = W - n_{22} - n_{12}$. We have demonstrated in [19] that the blocking probabilities obtained through the product-form solution to the time-reversible Markov process are very close to the blocking probabilities obtained through the numerical solution to the original Markov process for a wide range of traffic loads.

Let us now consider the same 2-hop path with the most-used wavelength allocation policy. This policy can be modeled as a Markov process with the same state description as the random policy case, i.e., $(n_{11}, n_{12}, n_{22}, f_{12})$. The key difference is that, under the most-used policy, if $n_{11} > n_{22}$, then we know that there is at least one wavelength that is used on hop 1 but not used on hop 2. Thus, an incoming call that uses the second hop only will be assigned a wavelength that is already in use on the first hop, and will cause a transition to state $(n_{11}, n_{12}, n_{22} + 1, f_{12})$; similarly for $n_{22} > n_{11}$ and incoming calls using only the first hop.

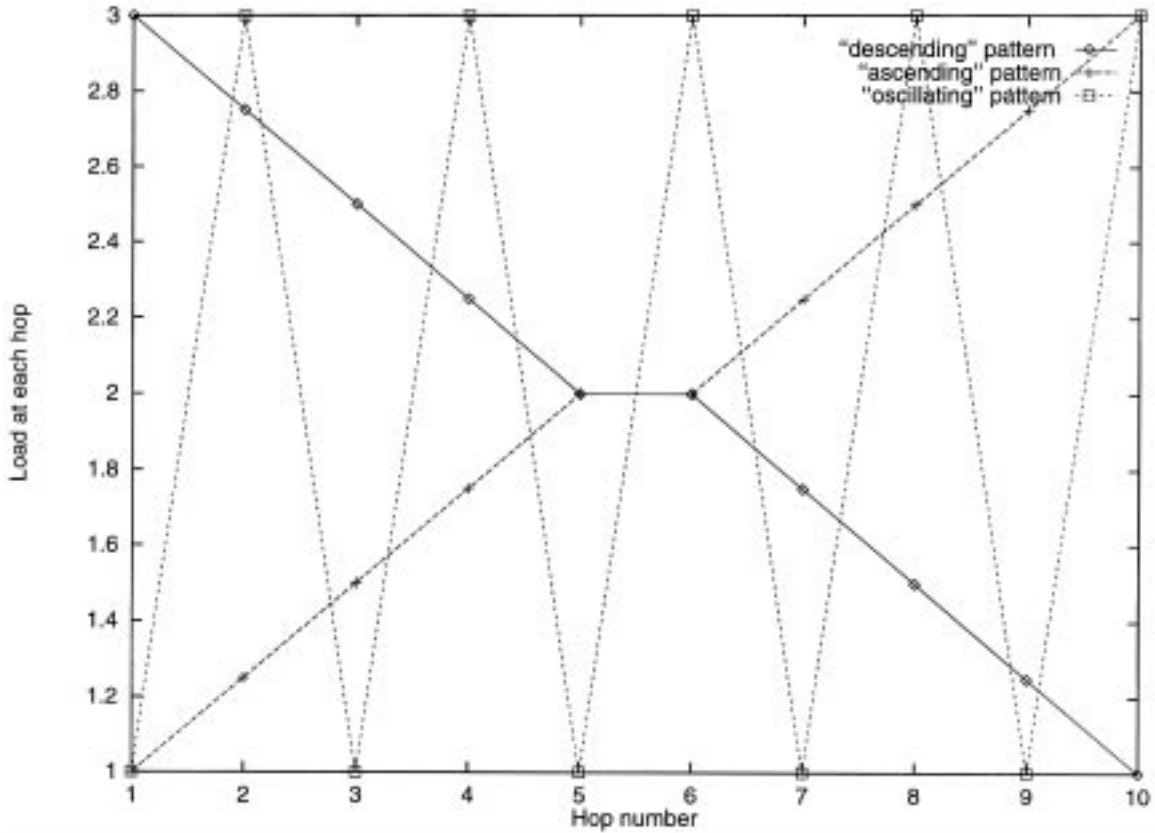


Fig. 3. The “ascending”, “descending”, and “oscillating” load patterns.

(Under the random wavelength allocation policy, the transition could be to either state $(n_{11}, n_{12}, n_{22} + 1, f_{12})$ or to state $(n_{11}, n_{12}, n_{22} + 1, f_{12} - 1)$ if the number of free wavelengths on both hops $f_{12} > 0$ and one of these wavelengths is assigned to the call.)

The state transition diagram of the Markov process for the most-used allocation policy is shown in Fig. 6

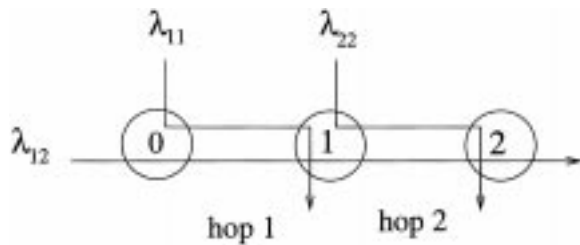


Fig. 4. A 2-hop path.

for a 2-hop path with $W = 2$ wavelengths. Again, it is straightforward to verify that this Markov process is not time-reversible. Comparing to Fig. 5, we note that despite having the same state space, the two processes differ in two ways. First, some of the transition rates are different; for instance the transition rate from state $(0,0,1,1)$ to state $(1,0,1,1)$ is equal to $\lambda_{11}/2$ for the random allocation, but λ_{11} for the most-used allocation. Second, some of the transitions are missing in the new Markov process. For example, there is a transition from state $(0,0,1,1)$ to state $(1,0,1,0)$ under random allocation in Fig. 5, but there is no such transition in Fig. 6. Furthermore, since there is a transition from state $(1,0,1,0)$ to state $(0,0,1,1)$ in Fig. 6, but no transition in the reverse direction, it is not possible to obtain an approximate time-reversible process by simply modifying some of the transition rates, as we did for the random wavelength allocation policy. Although we do not have an approximate

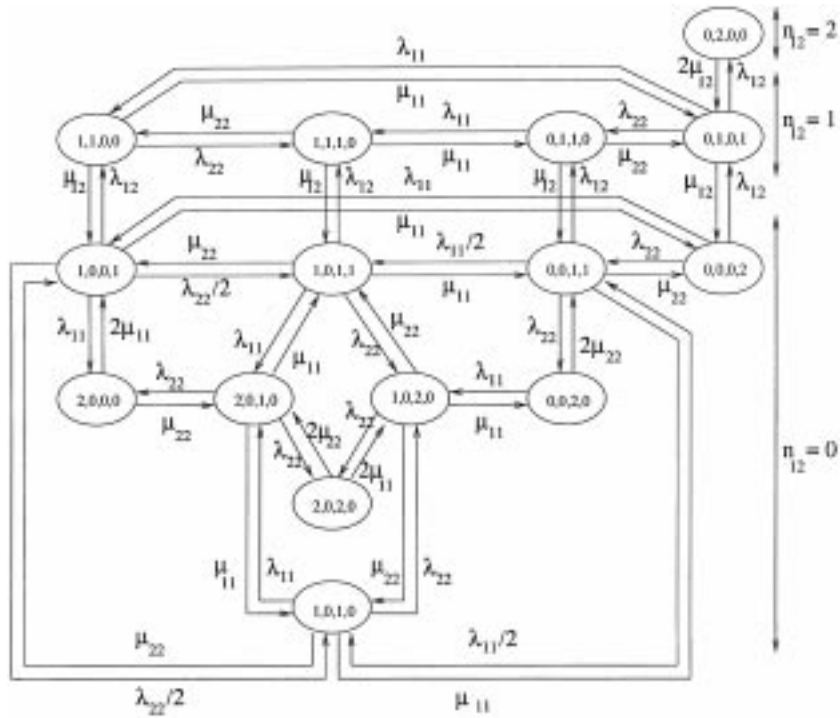


Fig. 5. State space $(n_{11}, n_{12}, n_{22}, f_{12})$ of a 2-hop path with $W = 2$ wavelengths (random allocation).

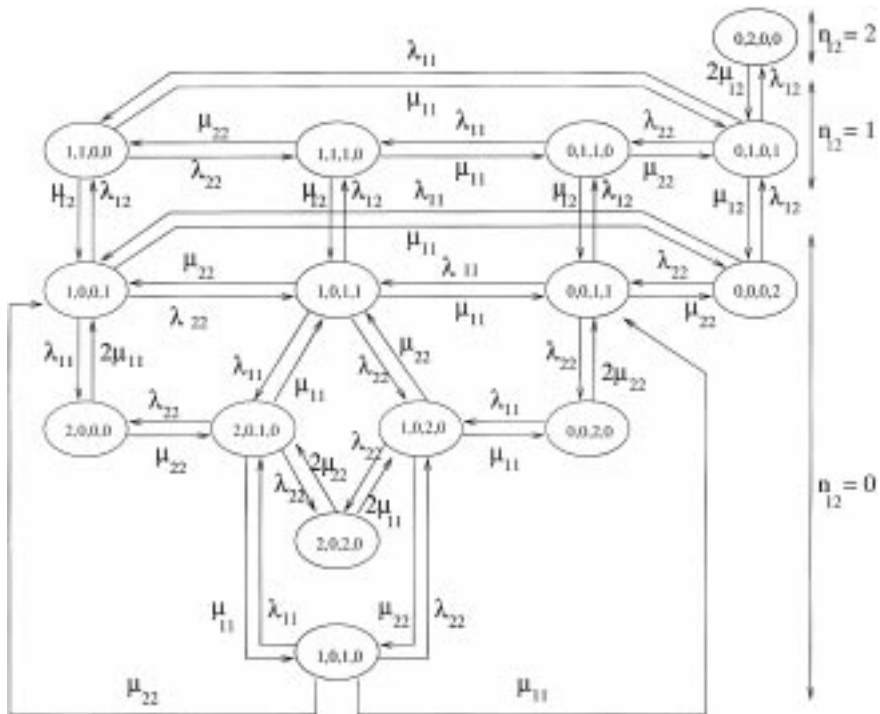


Fig. 6. State space $(n_{11}, n_{12}, n_{22}, f_{12})$ of a 2-hop path with $W = 2$ wavelengths (most-used allocation).

product-form solution for the most-used allocation policy, the state space for a 2-hop path is small enough so that the solution to the Markov process can be obtained numerically for up to $W = 20$ wavelengths.

Based on similar arguments, it can be determined that the least-used wavelength allocation policy can also be modeled by a Markov process with the state description $(n_{11}, n_{12}, n_{22}, f_{12})$. The state transition diagram for this process is shown in Fig. 7, and it can be easily verified that the process is not time-reversible.

If a converter is placed at node 1 of the 2-hop path shown in Fig. 4 (the only interesting possibility in this case), the system becomes a 2-hop all-converter path, and it can be described by the three-dimensional Markov process (n_{11}, n_{12}, n_{22}) . Random variable f_{12} becomes redundant because calls continuing on both hops can now use any of the $(W - n_{12} - n_{22})$ available wavelengths on the second hop. It is well-

known that this Markov process has the closed-form solution:

$$\pi_{cs}(n_{11}, n_{12}, n_{22}) = \frac{1}{G(W)} \frac{\rho_{11}^{n_{11}} \rho_{12}^{n_{12}} \rho_{22}^{n_{22}}}{n_{11}! n_{12}! n_{22}!}. \quad (2)$$

In Fig. 8 we show the state space of a 2-hop all-converter path with two wavelengths. Although this path is described by the above 3-dimensional Markov process, we include in the state description of Fig. 8 the random variable f_{12} to make it easier to compare to Figs. 5–7. For instance, the fact that there are no transitions into state $(1,0,1,0)$ in the figure can be explained by recalling that $f_{12} = 0$ (i.e., that no wavelength is free on both links of the path) implies that calls traversing both hops are blocked. However, since exactly one wavelength is free on each hop (even if it is not the same one), calls traversing both hops cannot be blocked in the all-converter path, and

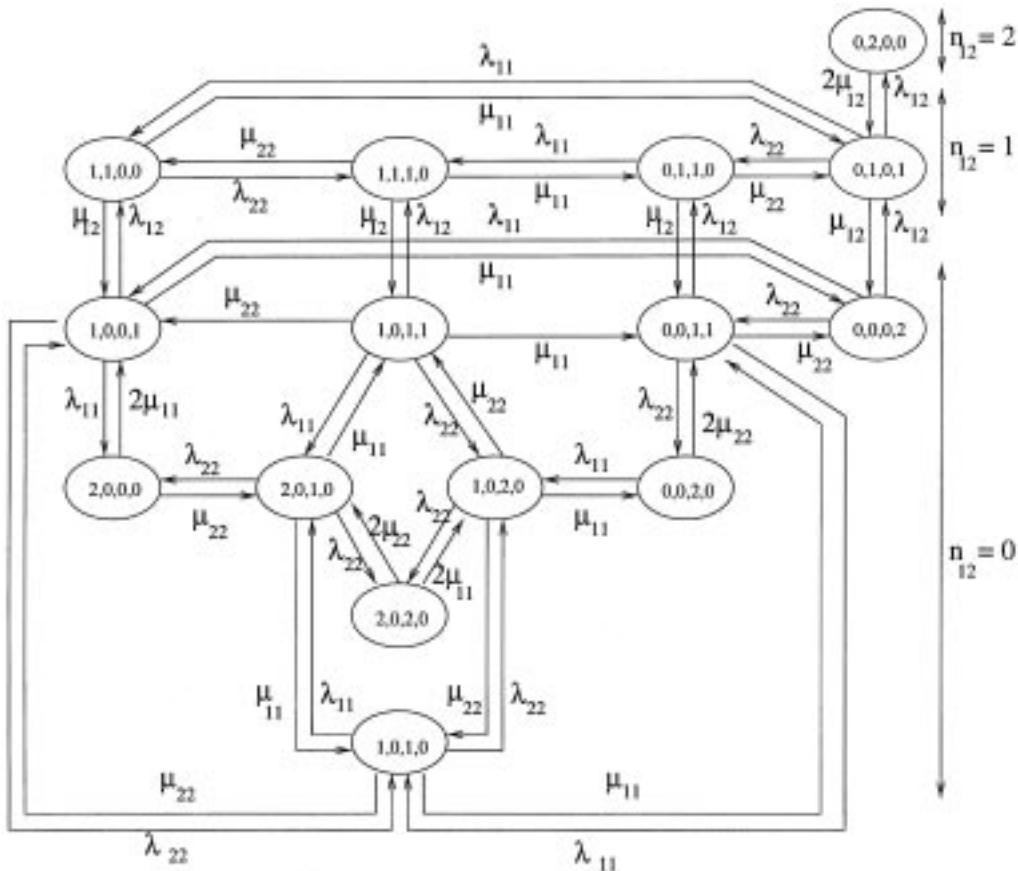


Fig. 7. State space $(n_{11}, n_{12}, n_{22}, f_{12})$ of a 2-hop path with $W = 2$ wavelengths (least-used allocation).

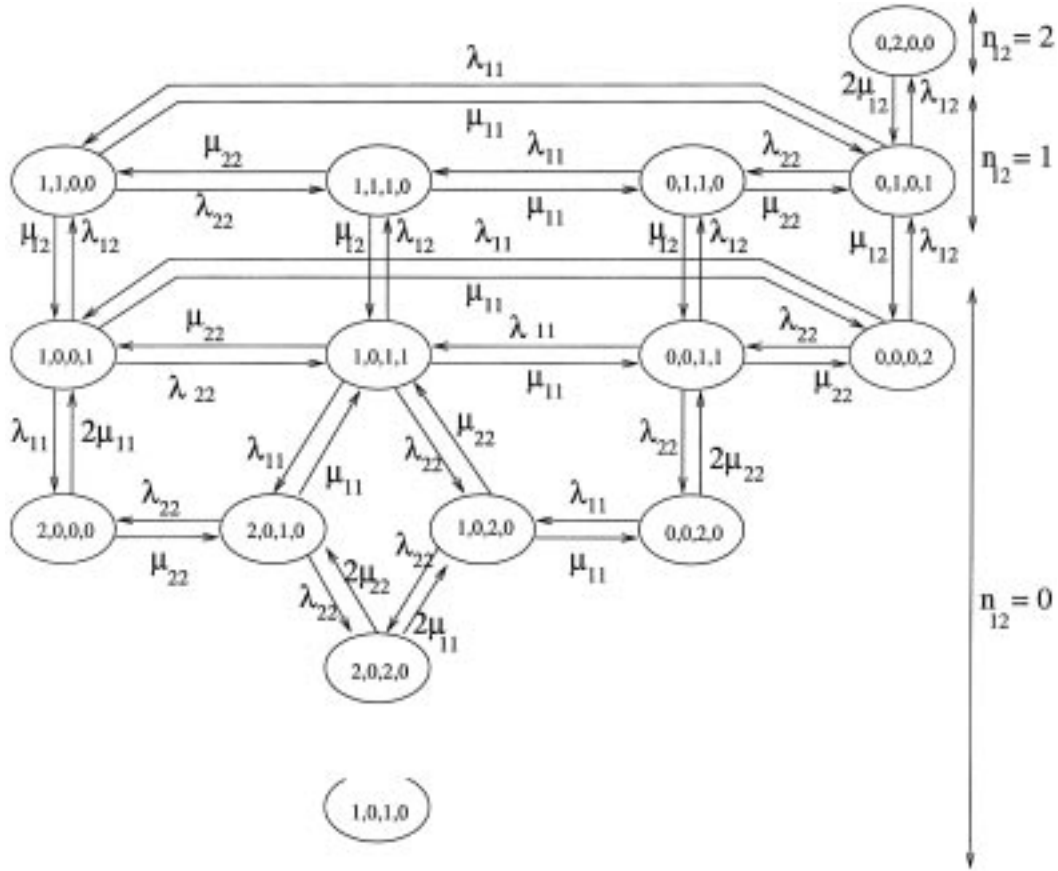


Fig. 8. State space (n_{11}, n_{12}, n_{22}) of a 2-hop path with $W = 2$ wavelengths (all-converter).

the system will never enter state $(1,0,1,0)$, but only state $(1,0,1,1)$.

The first-fit wavelength allocation policy can be modeled as a Markov process with W state variables (l_1, \dots, l_W) . Each random variable l_i corresponds to one of the W wavelengths, and can take one of five values representing the status of wavelength i on the two links of the path: 0, if the wavelength is free on both links, 1, if it is free on the first link and busy on the second, 2, if it is busy on the first link and free on the second, 3, if the wavelength is used by two different calls on each link, and 4, if it is used by a call traversing both links of the path. The state space of this Markov process is quite different than that in Fig. 5–8, and its states and transitions cannot be compared to those of the previous Markov processes. Furthermore, the size of the state space is in the order of W^5 , too large to obtain a numerical solution even for relatively small values of W . In view of this,

the blocking probabilities for this policy are obtained by simulation only.

2.1.2 Numerical Comparisons

Let us first consider the blocking probabilities of the random, most-used, least-used, and all-converter systems for calls traversing both links of the 2-hop path. In Figs. 5 to 8, the blocking states for these calls are those with $f_{12} = 0$, i.e., those states in which neither of the two wavelengths is free on both links. We also observe that, except for state $(1,0,1,0)$ at the bottom of each of the four figures, the transitions (and transition rates) in and out of all other blocking states are exactly the same for all four wavelength allocation policies. Consequently, we expect that the difference in the blocking probability experienced by calls traversing both links of the path under the different policies will

be mainly due to the steady-state probability of blocking state (1,0,1,0).

Referring to Fig. 8, we note that the corresponding Markov process never enters state (1,0,1,0). Thus, we expect that calls traversing both hops will experience the least blocking probability in an all-converter path. In Fig. 6 (most-used policy) we note that there are two transitions into state (1,0,1,0), and four transitions out of it. The blocking probability will be higher under this policy compared to the all-converter case. The Markov process in Fig. 5 (random policy) has two additional transitions into state (1,0,1,0) from states (0,0,1,1) and (1,0,0,1) with rates $\lambda_{11}/2$ and $\lambda_{22}/2$, respectively. Therefore, the blocking probability of these calls under the random policy will be higher than under the most-used policy. Finally, the Markov process in Fig. 7 (least-used policy) has the same transitions as the one in Fig. 5, but the transition rates into state (1,0,1,0) from states (0,0,1,1) and (1,0,0,1) are λ_{11} and λ_{22} , respectively. Therefore, we expect

that these calls will experience the highest blocking probability under the least-used policy.

We now note that the lower the blocking probability for calls traversing both hops, the larger the number of such calls accepted, and the larger the number of wavelengths they occupy, thus leaving fewer wavelengths available for calls using a single link (either the first or the second) of the path. Hence, we expect that the behavior of the four policies in terms of the blocking probability of calls using a single link of the path will be exactly the opposite of what was discussed above. Specifically, we expect the least-used policy to provide the lowest blocking probability for these calls, followed by the random, the most-used, and the all-converter policies, in that order.

The above conclusions, derived by direct comparison of the states of the Markov processes, are in agreement with intuition. We have confirmed these conclusions by numerically comparing the blocking probabilities of the various policies for 128 different

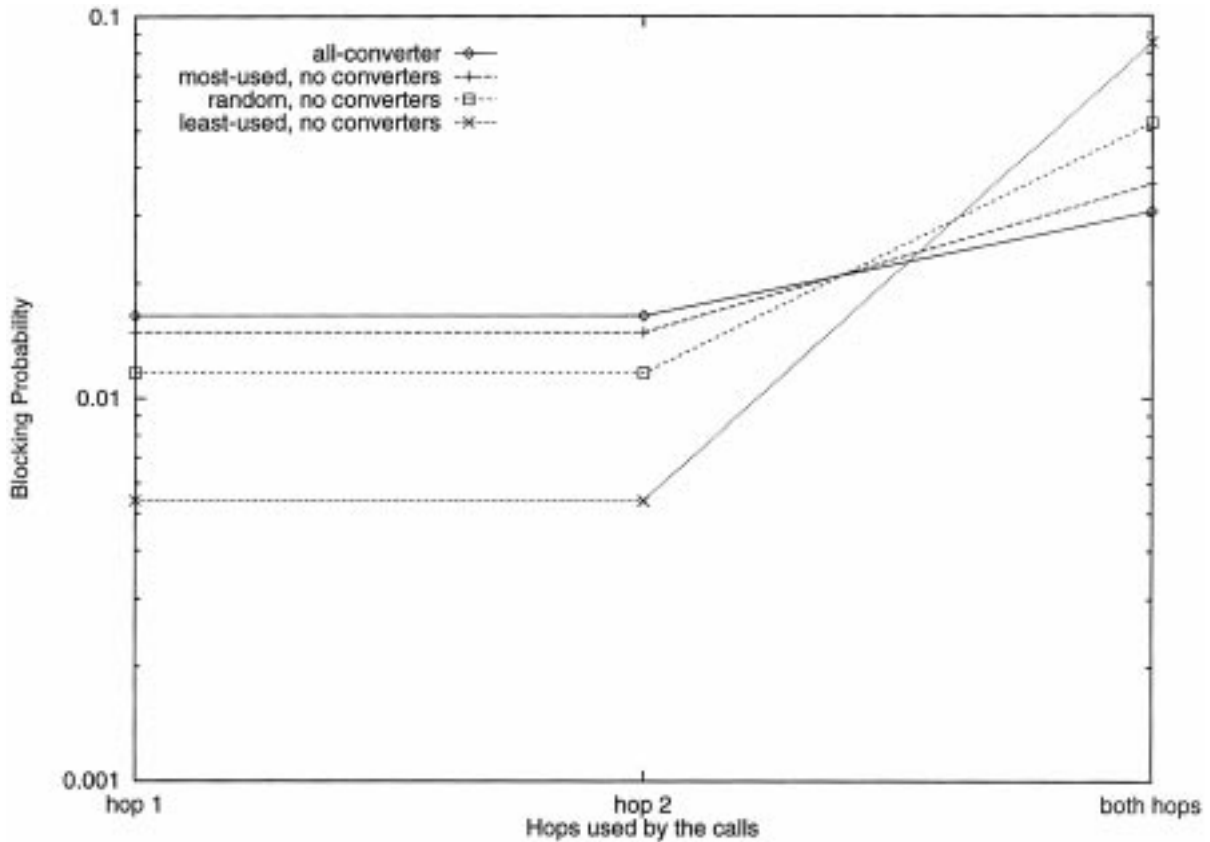


Fig. 9. Policy comparison, 2-hop path, uniform traffic pattern.

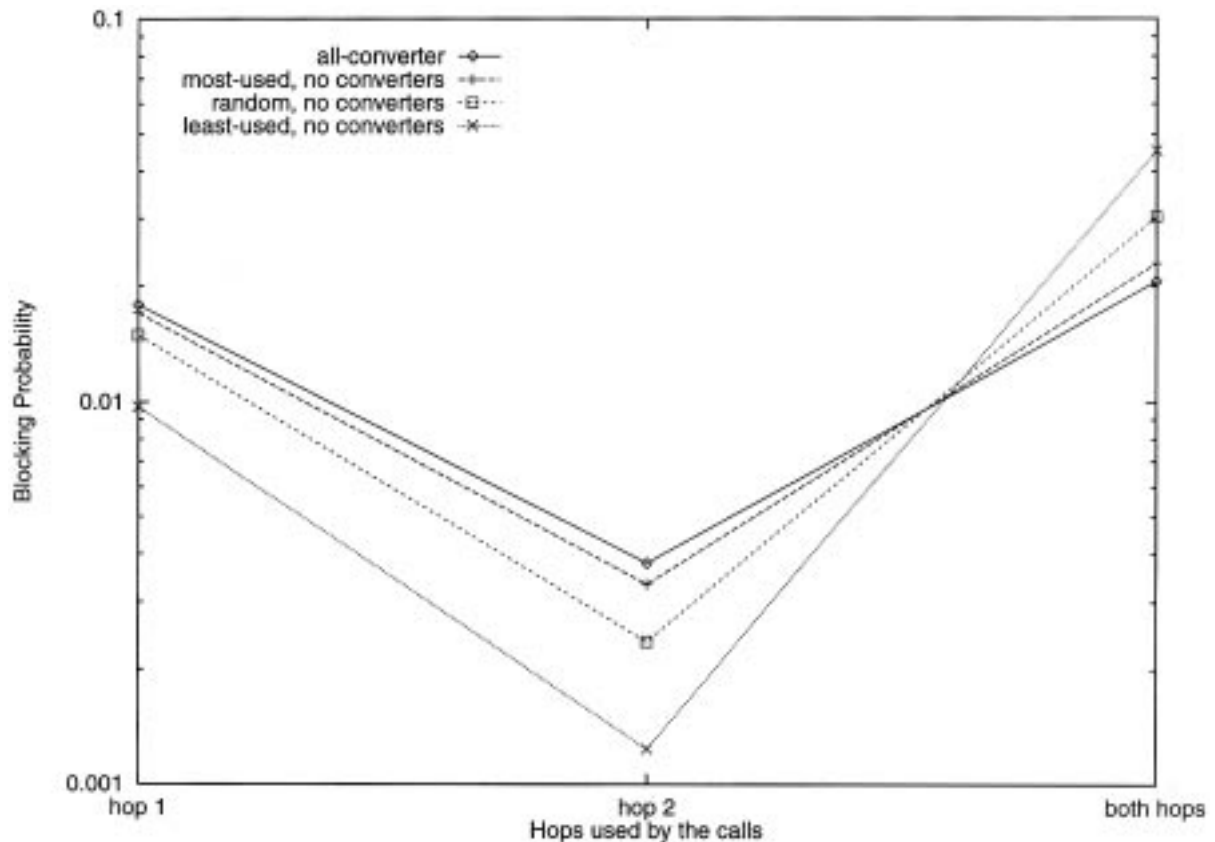


Fig. 10. Policy comparison, 2-hop path, descending traffic pattern.

load values and load scenarios similar to Figs. 2 and 3. Figs. 9 and 10 show results for two cases corresponding to a uniform and descending load pattern, respectively (see the patterns in Figs. 2 and 3) and for $W = 10$ wavelengths. More specifically, the arrival rates (refer also to Fig. 4) used to obtain the results in Fig. 9 were $\lambda_{11} = 0.2$, $\lambda_{12} = 0.1$, $\lambda_{22} = 0.2$, while for the results in Fig. 10 we used $\lambda_{11} = 3.0$, $\lambda_{12} = 2.0$, $\lambda_{22} = 2.0$. In both Figs. we plot the blocking probability for the three types of calls, namely, calls using the first hop only (label “hop 1” in the x -axis of the figures), calls using the second hop only (label “hop 2”), and calls using both hops (label “both hops”). We first note that the results are affected by the traffic pattern used. For instance, under uniform loading (Fig. 9), calls using the first hop only experience the same blocking probability as calls using the second hop only, while in the descending pattern (Fig. 10), due to the lower load offered to the second hop, the latter calls experience a much lower

blocking probability for all four policies. More importantly, the relative values of the blocking probabilities for the four policies are also consistent with our discussion above. Very similar results have been obtained for all 128 different load values that we have studied.

Finally, in Fig. 11 we compare the most-used and first-fit policies for the same arrival rates as those used for Fig. 10. As before, the blocking probabilities of the most-used policy were obtained through a numerical solution to the corresponding Markov process, while the values for the first-fit policy were obtained through simulation. We observe that the blocking probabilities of the first-fit policy are almost identical to those of the most-used policy for all three types of calls. This result can be explained by noting that both policies attempt to maximize the number of wavelengths that are available for calls that use both hops of the 2-hop path by reducing the “fragmentation” of the set of wavelengths. The most-used policy assigns to an

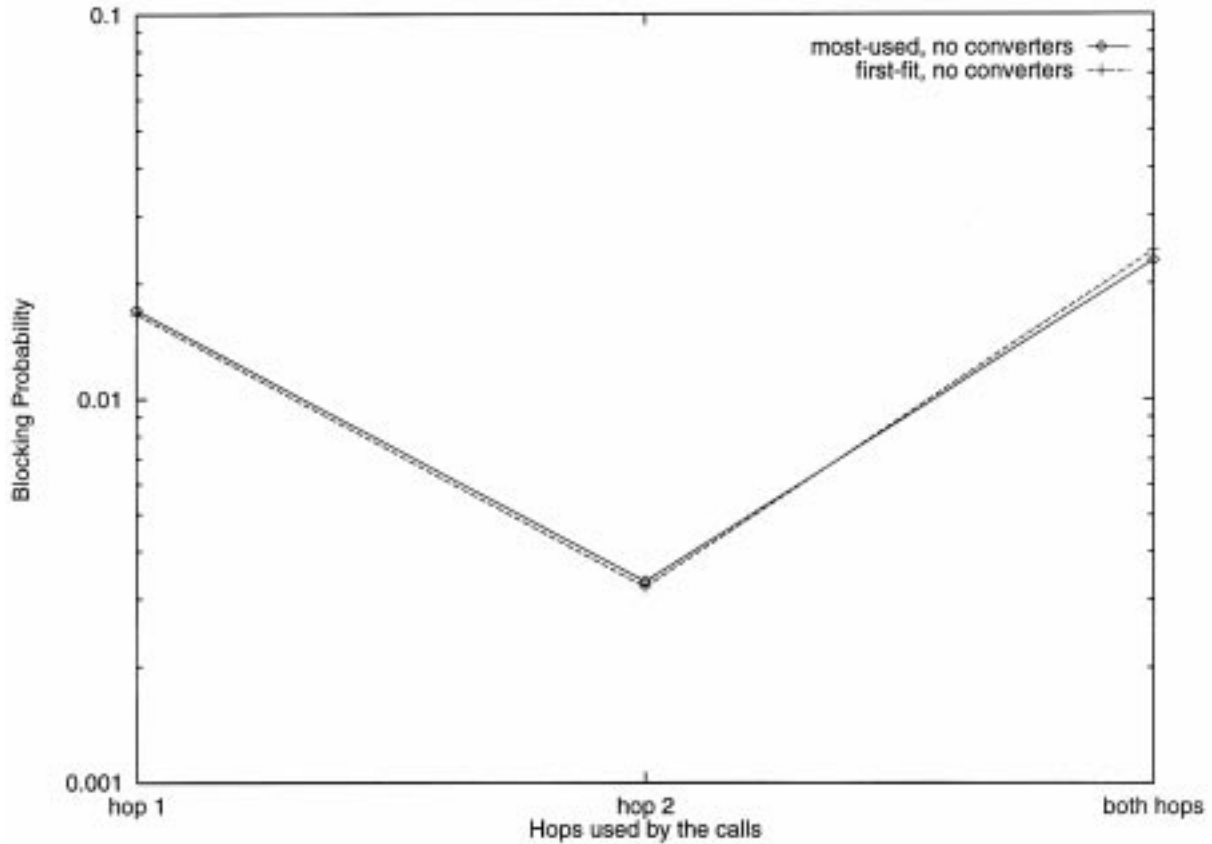


Fig. 11. Most-used vs. first-fit allocation, 2-hop path, descending traffic pattern.

incoming call that requires a single hop of the path a wavelength that is already used on the other hop, if such a wavelength exists. On the other hand, the first-fit policy attempts to achieve the same goal by searching the set of wavelengths in a fixed order, thus increasing the chances that a wavelength used on a single hop will be assigned to an incoming call using the other hop. As can be seen from Fig. 11, the most-used policy is slightly better, but overall the blocking probability values of the two policies are very close. Similar results have been obtained for all 128 traffic loads that we have studied.

2.2 Policy Comparisons for Longer Paths

Consider a k -hop path, $k > 2$, with the random wavelength allocation policy. Paths consisting of four links or less can be analyzed approximately by solving the corresponding time-reversible Markov process. For instance, a 3-hop path can be modeled by

the 9-dimensional Markov process $(n_{11}, n_{12}, n_{13}, n_{22}, n_{23}, n_{33}, f_{12}, f_{13}, f_{23})$ whose solution can be written down as a straightforward generalization of expression (1). Paths longer than four hops are analyzed using the iterative decomposition algorithm in [19] to obtain the call blocking probabilities. The analytical techniques developed in [19] are both accurate and efficient, and can be used when the path employs converters at nodes arbitrarily chosen. When all nodes in a k -hop path employ converters (the all-converter case), the call blocking probabilities can be also obtained by using a straightforward generalization of expression (2). For very large k , when the computation of the normalizing constant becomes computationally expensive, a decomposition algorithm similar to the one in [19] can be used.

Let us now consider the most-used wavelength allocation policy. It is relatively easy to derive an exact Markov process to model a k -hop path, $k > 2$.

However, the number of random variables in the state description of the process for a k -hop path, $k > 2$, grows very large, and therefore it is difficult to obtain the call blocking probabilities numerically by directly solving the exact Markov process. Furthermore, developing an iterative algorithm for analyzing long paths by decomposing them into 2-hop path sub-systems which can be solved in isolation, similar to the algorithm developed for the random policy in [19], has turned out to be a difficult task. For such an algorithm, it is crucial to have an accurate estimate of the blocking probability due to the wavelength continuity requirement for calls traversing more than one sub-system. Since the different wavelengths are not equally utilized, as under random allocation, it is difficult to derive an approximate expression for blocking due to the wavelength continuity requirement that is accurate for a wide range of loads. In view of all these, the results presented in this section for the

most-used policy have been obtained by simulation.

For similar reasons, we have used simulation to obtain the call blocking probabilities for paths with the least-used and first-fit wavelength allocation policies. (While approximate analytical techniques based on overflow traffic have been developed for the first-fit policy in [12,14], these techniques are inaccurate since the link blocking events are taken to be independent, an assumption that is not true over a wide range of traffic loads.) Since the 95% confidence intervals are very tight, to increase the readability, we have decided not to plot them along with the simulation results.

2.2.1 Numerical Comparisons

In this section we present results for 6-hop and 10-hop paths, since the length of these paths (in hops) is representative for future backbone wavelength

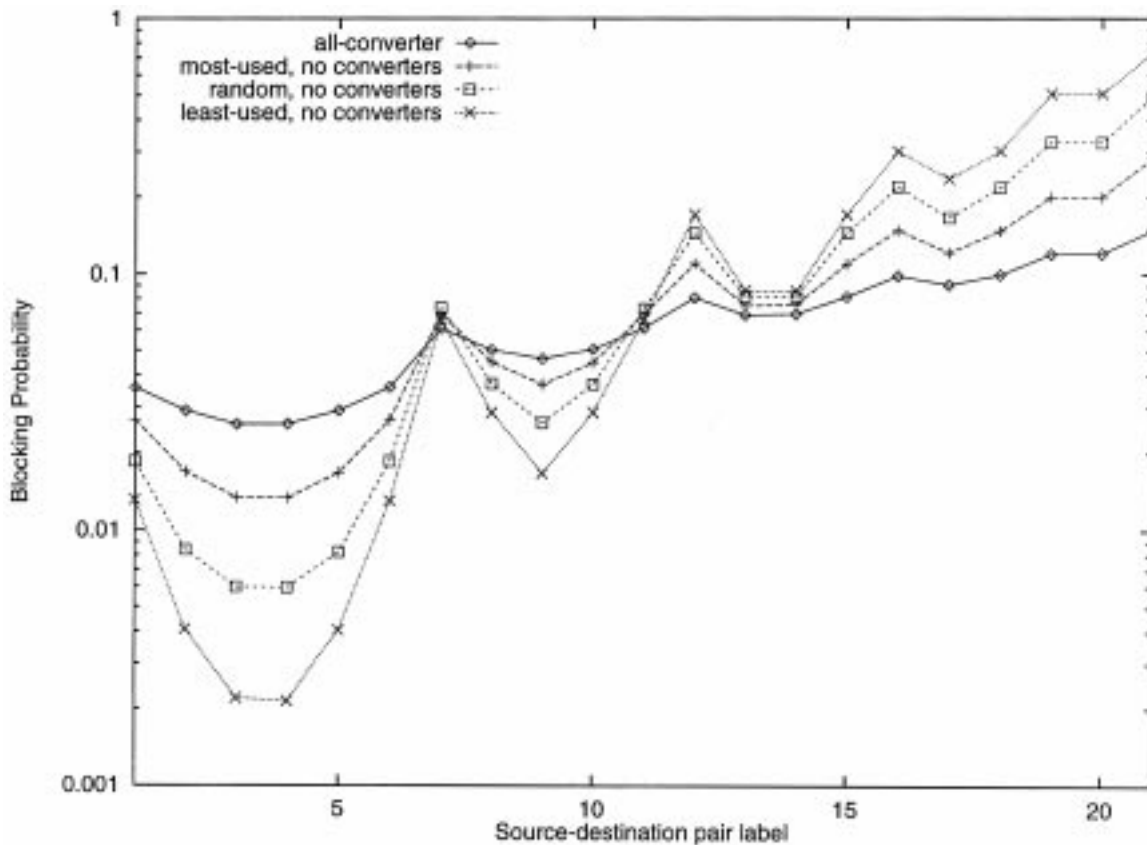


Fig. 12. Policy comparison, 6-hop path, uniform traffic pattern.

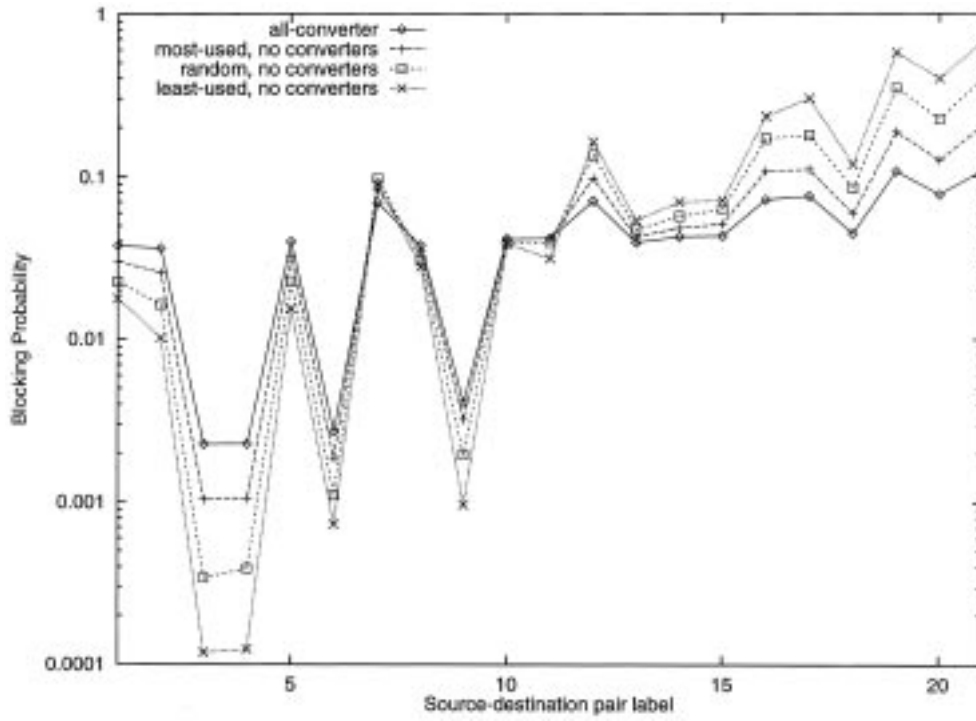


Fig. 13. Policy comparison, 6-hop path, bowl traffic pattern.

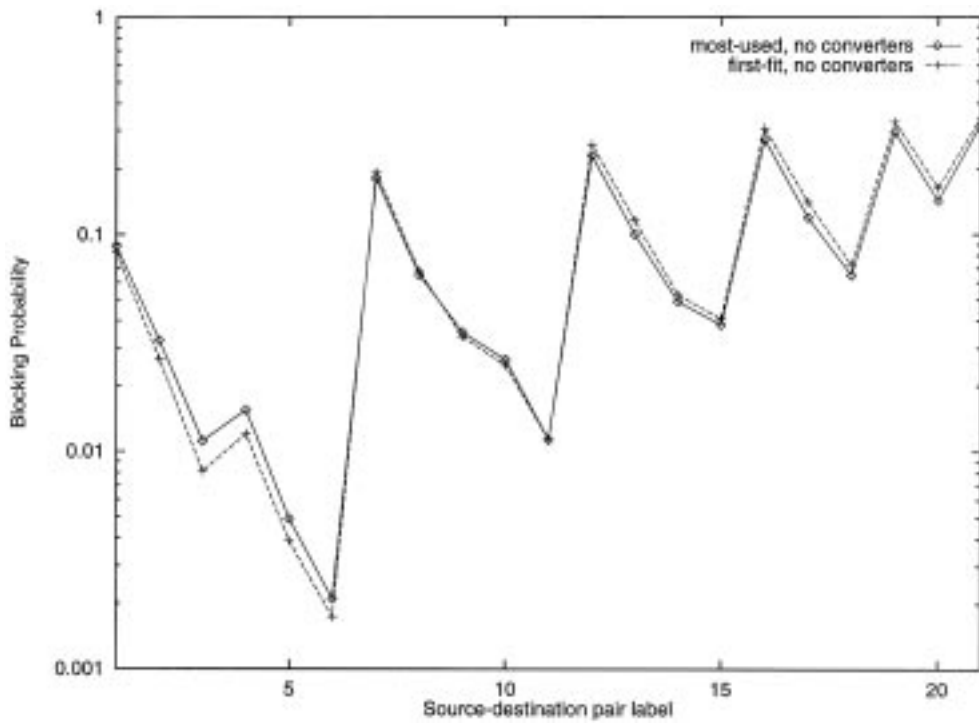


Fig. 14. Most-used vs. first-fit allocation, 6-hop path, descending traffic pattern.

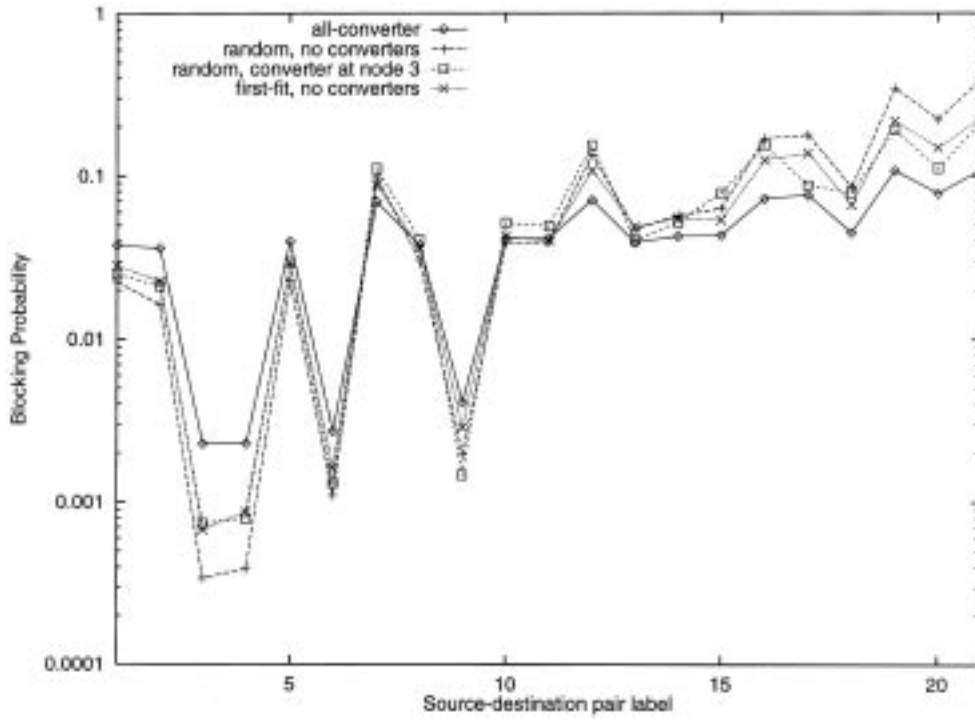


Fig. 15. First-fit policy vs. random policy with converters, 6-hop path, bowl traffic pattern.

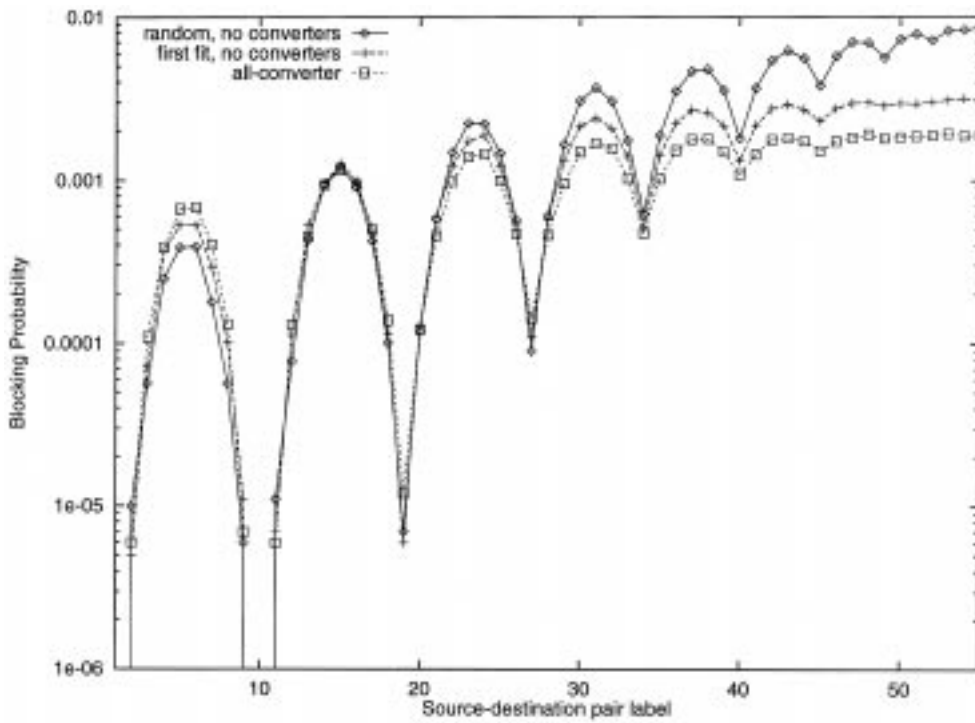


Fig. 16. Policy comparison, 10-hop path, inverted bowl traffic pattern.

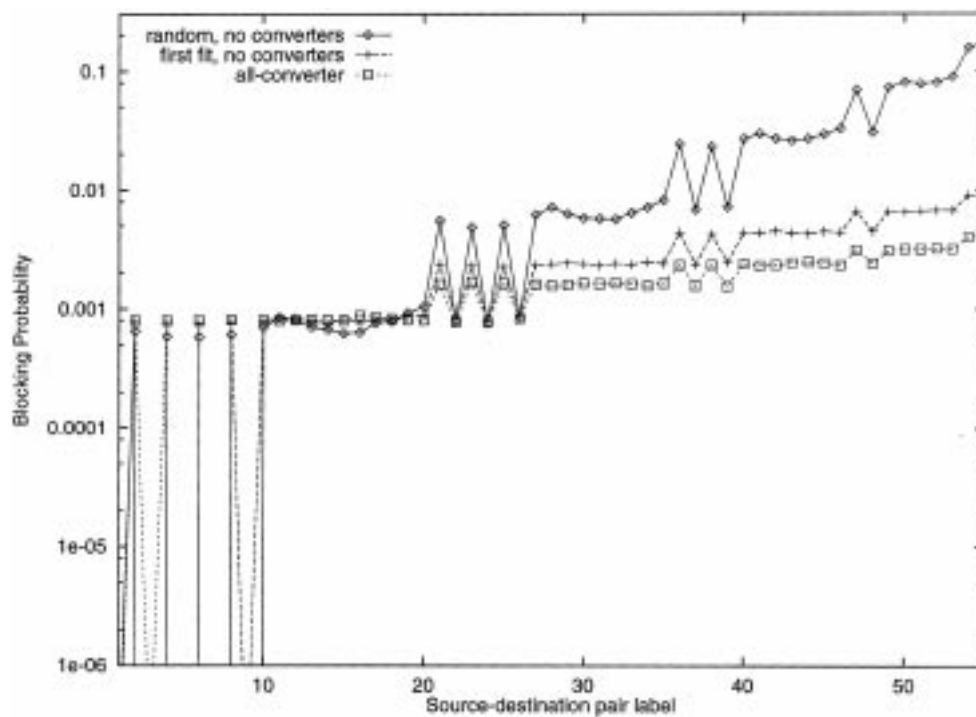


Fig. 17. Policy comparison, 10-hop path, oscillating traffic pattern.

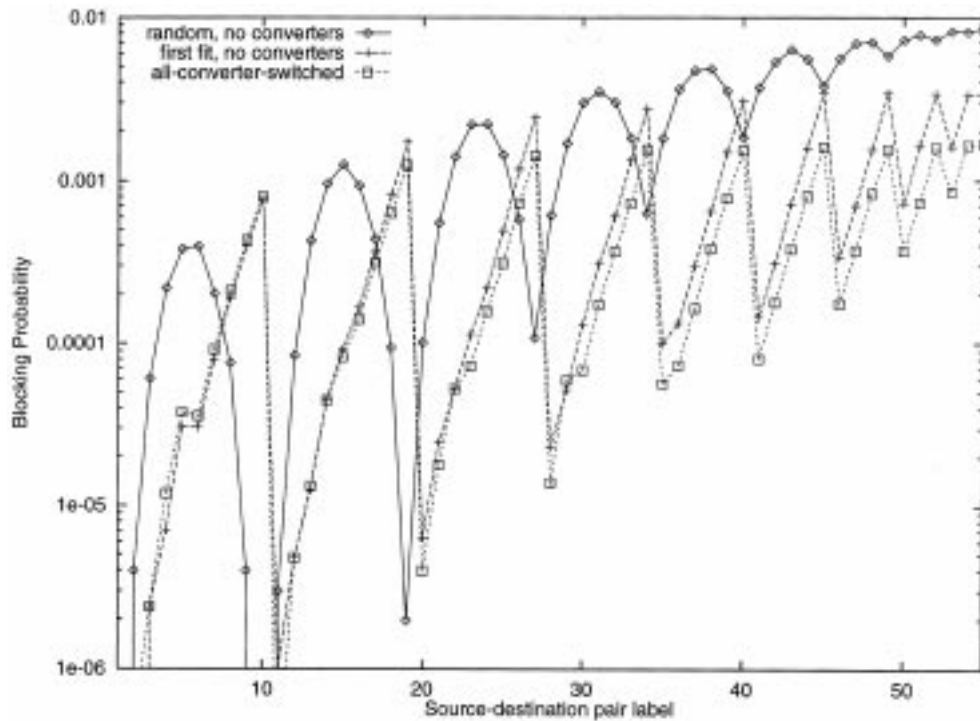


Fig. 18. Policy comparison, 10-hop path, ascending traffic pattern.

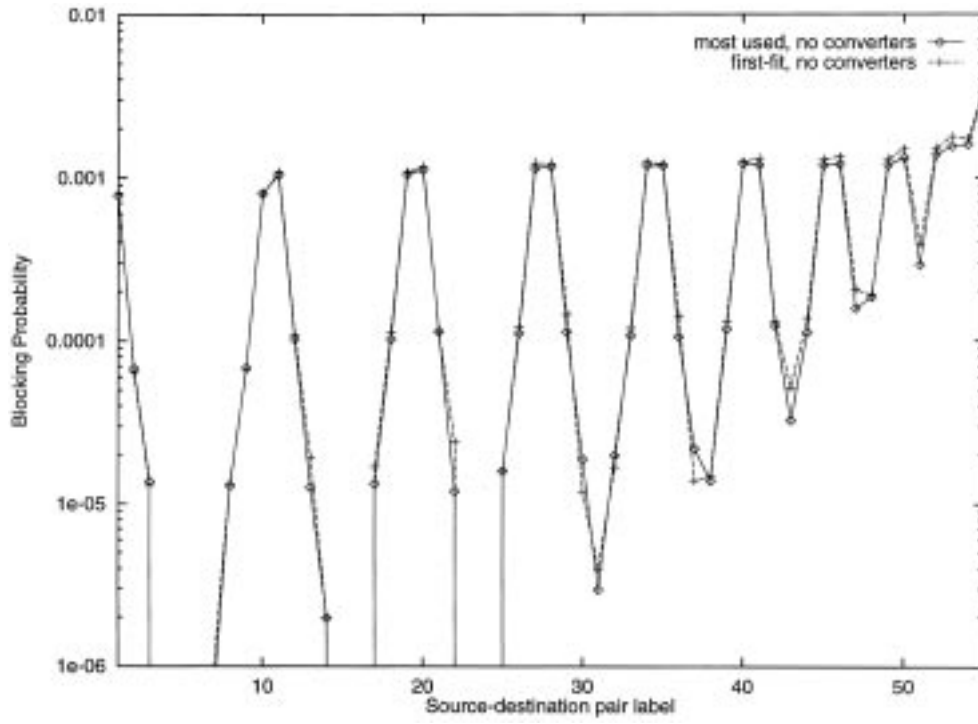


Fig. 19. Most-used vs. first-fit allocation, 10-hop path, bowl traffic pattern.

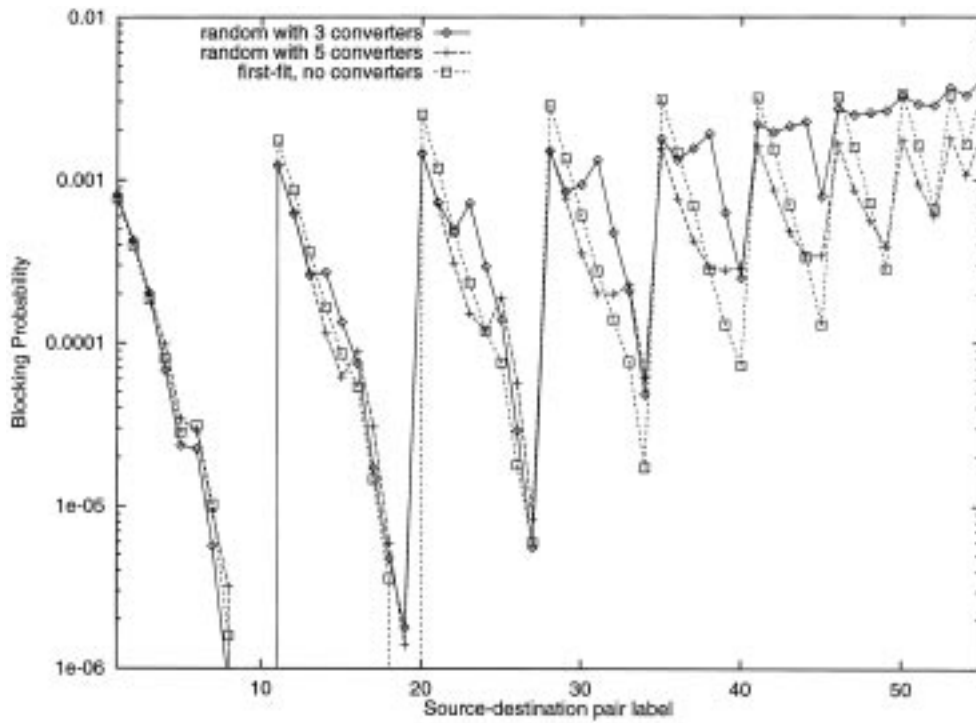


Fig. 20. First-fit policy vs. random policy with converters, 10-hop path, descending traffic pattern.

routing networks. Figs. 12 to 15 correspond to a 6-hop path, and Figs. 16 to 20 are for a 10-hop path.

In Figs. 12 and 13 we compare the blocking probabilities for the four policies (random, most-used, least-used, and all-converter) under the uniform and bowl traffic patterns, respectively. In both figures we plot the blocking probability for the twenty one different types of calls in a 6-hop path, numbered 1 through 21 on the x -axis. The calls have been numbered so that numbers 1 through 6 correspond to calls that traverse only a single hop in the path, that is $x = 1$ is for calls using hop 1, $x = 2$ for calls using hop 2, etc. Numbers 7 to 11 on the x -axis correspond to calls that traverse exactly two hops in the path, that is $x = 7$ is for calls traversing hops 1 and 2, etc. Numbers 12 to 15 on the x -axis correspond to calls traversing exactly three hops, and so on.

From Figs. 12 and 13 we observe that the relative behavior of the four policies is similar to that shown in Figs. 9 and 10 despite the fact that the traffic patterns in these figures are very different. Specifically, for calls using one or two hops (calls one through eleven in the figures), the least-used policy provides the lowest blocking probability, followed by the random policy, the most-used policy, and the all-converter case. However, for calls traversing three or more hops, the situation is reversed. The same behavior was observed for other traffic patterns and paths of different length. We also note that, under the least-used policy, the blocking probability of calls using multiple hops increases significantly, and that the average blocking probability over all calls is higher than other policies. Therefore, we will not consider the least-used policy any further.

In Fig. 14 we compare the most-used and first-fit policies for a 6-hop path assuming a descending traffic pattern. Again, as in Fig. 11, we find that the two policies give almost identical blocking probabilities, not only for the end-to-end call, but for all calls, regardless of the number of hops used by the calls. Very similar results have been obtained for all traffic patterns studied. Therefore, in the rest of the paper we will concentrate on the first-fit policy, since its implementation does not require that the network nodes maintain information about the global use of wavelengths. Using simulation results we will show that the values of the blocking probabilities obtained with this policy are bounded by the blocking probability values obtained by the random policy without converters and the all-converter case (i.e., the

random policy with converters at all nodes of a path). Also, we will demonstrate that, for calls traversing multiple hops, the gain (in terms of reduction in the blocking probabilities) obtained by employing the first-fit instead of the random policy is roughly equivalent to using the random policy and deploying converters in the network.

In Fig. 15 we compare the blocking probabilities of a 6-hop path obtained using the first-fit policy to the random policy with no converters, the random policy and one converter at node 3, and the all-converter case (i.e., the random policy and a converter at each node). As we can see, the first-fit policy has an effect similar to that of using the random policy and employing a converter in the path. This is a general result that has been observed for a wide range of traffic loads, and will be discussed below in more detail.

Figs. 16 to 20 present results for a 10-hop path and various traffic patterns. In Figs. 16 to 18 we compare the first-fit policy to the random (no converters) and all-converter cases for the inverter bowl, oscillating, and ascending traffic patterns, respectively. Two interesting observations can be made from these three figures. First, the blocking probability values of the first-fit policy are always between the corresponding values of the random and all-converter cases. In other words, the blocking probability values under the random and all-converter cases provide lower and upper bounds for the blocking performance of the first-fit policy. Note that the random policy provides a lower bound for calls using one or two hops, and an upper bound for calls traversing three or more hops; the reverse is true for the all-converter path. Second, it appears that the first-fit policy is quite effective in reducing the blocking probability of calls traveling over multiple hops (which are the ones that experience the highest blocking probability under the random policy) close to the level of the all-converter case.

In Fig. 19 we compare the first-fit to the most-used policy for the descending traffic pattern. As before, the blocking probability values of the two policies match for all types of calls. Finally, in Fig. 20 we attempt to quantify the effect of the first-fit policy in terms of ‘‘number of converters.’’ Specifically, we plot the blocking probability values for the first-fit policy as well as those of a random policy with either three or five converters. The converters are placed at nodes in a way that minimizes the blocking probability of calls traveling over all 10-hops, using

the techniques developed in [19]. Note that, by employing converters at some of the nodes, the blocking probability of calls traversing multiple hops improves, since converters reduce the requirement that the same wavelength be used on all hops of the path taken by the call. However, this improvement is at the expense of calls using a single hop, which now experience higher blocking probability. As we can see, the effect of using the first-fit policy in place of the random policy has an effect similar to employing converters.

3 Mesh Wavelength Routing Networks

In this section we compare further the effects of the first-fit policy to the random policy with converters, by studying two network topologies: a regular 5×5 torus network, and the NSFNET irregular topology.

3.1 The 5×5 Torus Network

We consider the 5×5 torus network shown in Fig. 21, with $W = 10$ wavelengths per link. Since there are 600 source-destination pairs in this network, it is impossible to present numerical results for all of them.

We present, therefore, the call blocking probabilities for only 24 different source-destination pairs, namely, those with node 1 as the source. Because of the regular topology, the selected pairs are a representative sample of the various source-destination pairs. Similar to previous figures, the source-destination pairs have been labeled such that numbers 1 through 4 correspond to pairs for which a 1-hop path is used, numbers 5 to 12 correspond to pairs for which a 2-hop path is used, and so on, as shown in Table 1.

In our study we have used two traffic patterns. For the first pattern, the call arrival rates were selected such that

$$\lambda_{sd} = \begin{cases} 0.4, & \text{if the length of the path from } s \text{ to } d \text{ is } 1 \\ 0.3, & \text{if the length of the path from } s \text{ to } d \text{ is } 2 \\ 0.2, & \text{if the length of the path from } s \text{ to } d \text{ is } 3 \\ 0.1, & \text{if the length of the path from } s \text{ to } d \text{ is } 4 \end{cases} \quad (3)$$

This selection of arrival rates was intended to capture the locality of traffic that has been observed in many networks. The utilization of each link in the network for these arrival rates is in the range $[3.140, 3.144]$. The tight range of link utilizations can be explained by the fact that both the topology and the traffic load are symmetric. For the second pattern (which we will refer to as the random pattern), each arrival rate λ_{sd}

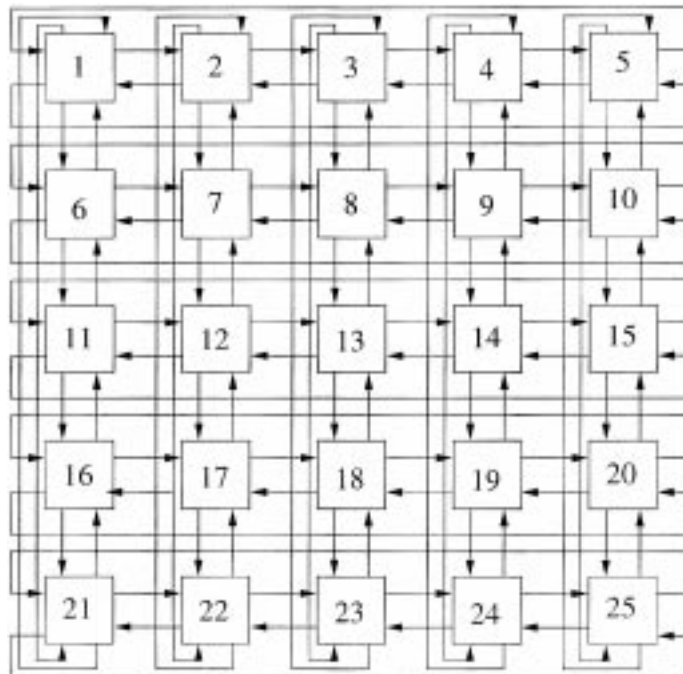


Fig. 21. The 5×5 bidirectional mesh torus network.

Table 1. Selected source-destination pairs for the torus network.

Pair	(1,2)	(1,5)	(1,6)	(1,21)	(1,3)	(1,4)	(1,7)	(1,10)
Label	1	2	3	4	5	6	7	8
Shortest Path Length	1	1	1	1	2	2	2	2
Pair	(1,11)	(1,16)	(1,22)	(1,25)	(1,8)	(1,9)	(1,12)	(1,15)
Label	9	10	11	12	13	14	15	16
Shortest Path Length	2	2	2	2	3	3	3	3
Pair	(1,17)	(1,20)	(1,23)	(1,24)	(1,13)	(1,14)	(1,18)	(1,19)
Label	17	18	19	20	21	22	23	24
Shortest Path Length	3	3	3	3	4	4	4	4

was selected from a uniform distribution in the range (0.1, 0.4).

In Fig. 22 we compare the blocking probabilities obtained through the most-used and first-fit policies for the pattern based on locality of traffic. From the figure, we observe that calls using a single hop (labels 1 to 4 in the figure) experience the lowest blocking

probability, calls traveling over two hops have the next lowest blocking probability, and so on. (The fact that the blocking probability values are almost the same for all calls using the same number of hops is due to the symmetry of both the topology and the traffic pattern.) We also see that the two policies result in almost identical blocking probability values for all

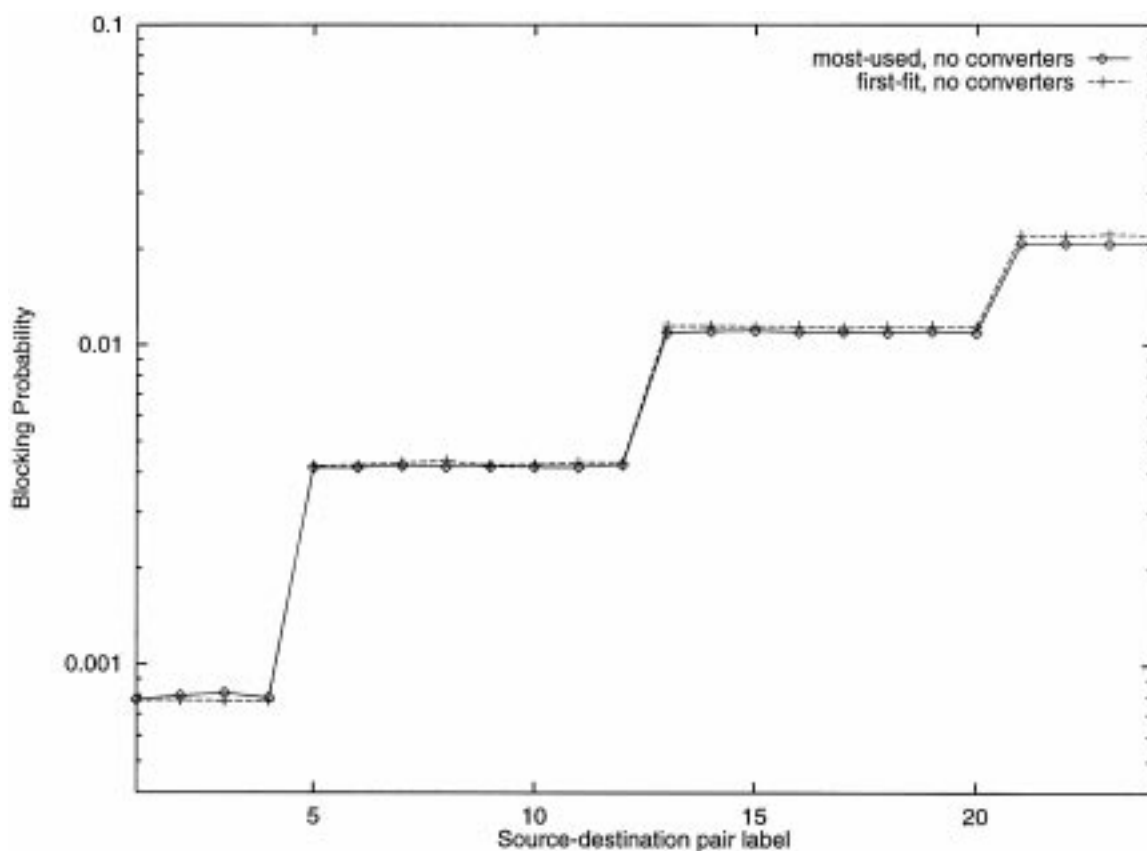


Fig. 22. Most-used vs. first-fit allocation, 5×5 torus network, traffic pattern based on locality.

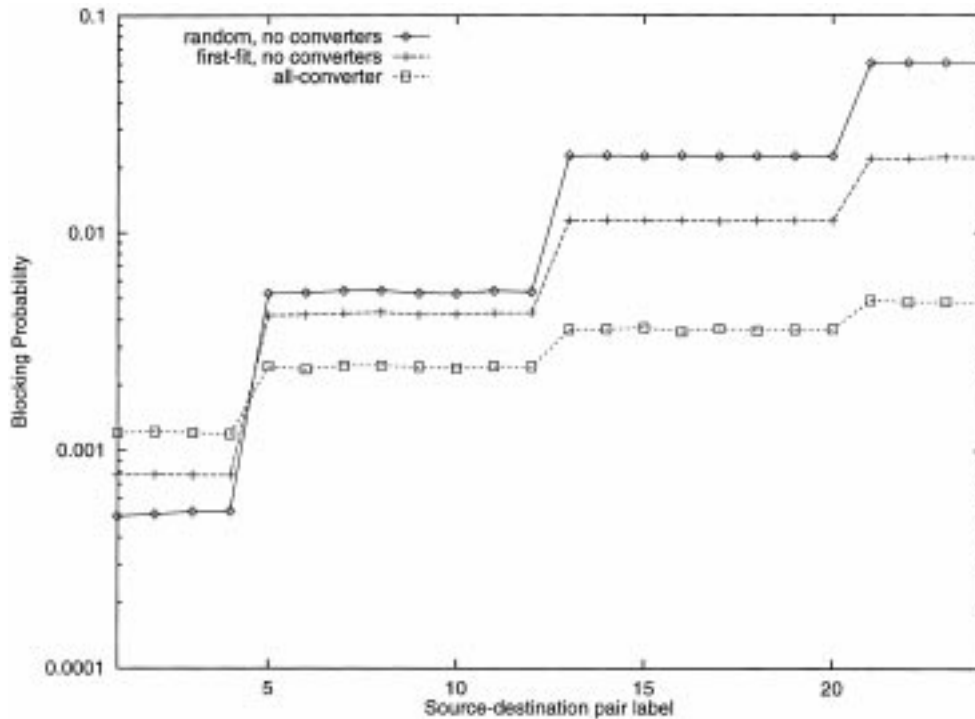


Fig. 23. Policy comparison, 5×5 torus network, traffic pattern based on locality.

calls, further confirming our claim that the (simpler) first-fit policy can be used as a quite accurate approximation of the most-used policy. Similar results, not shown here, have been obtained for the random traffic pattern.

In Figs. 23 and 24, we compare the first-fit policy to the random policy with no converters and the all-converter case, under the two traffic patterns. It is clear from both figures that the blocking probability values for the first-fit policy fall between those for the other two cases, a behavior which is consistent with our earlier results on single paths. However, there is also an important difference in the two figures. The first-fit policy appears to have a significant effect for the traffic pattern based on locality (Fig. 23), in that the blocking probability values for calls using multiple hops drops significantly from the corresponding values under the random policy with no converters. This effect, however, is less pronounced in Fig. 24 for the random traffic pattern. This difference can be explained by noting that the values of the blocking probability for calls 13 and higher in Fig. 24 are more than 0.1, about an order of magnitude higher

than the values for the corresponding calls in Fig. 23. At such high values, not many wavelengths are available for these calls, and as a result, the actual wavelength allocation policy used will have little effect on the blocking probability. It is at these high blocking probability values that using converters at all nodes (the all-converter case) will help. However, it is unlikely that realistic networks will be designed to operate in this region.

In Figs. 25 and 26 we compare the first-fit policy to the random policy with 4 and 12 converters employed in the torus network (note that these values correspond to 16% and 48%, respectively, of the network nodes having converters). As can be seen using the first-fit policy is roughly equivalent to employing a significant number of converters in the network. Another important observation is that the converters introduce an uneven effect on the blocking probabilities of the various calls. Specifically, calls whose path includes a converter experience a rather dramatic drop in the blocking probability. However, the effect of converters on the blocking probability of other calls is considerably smaller. The first-fit policy, on the other

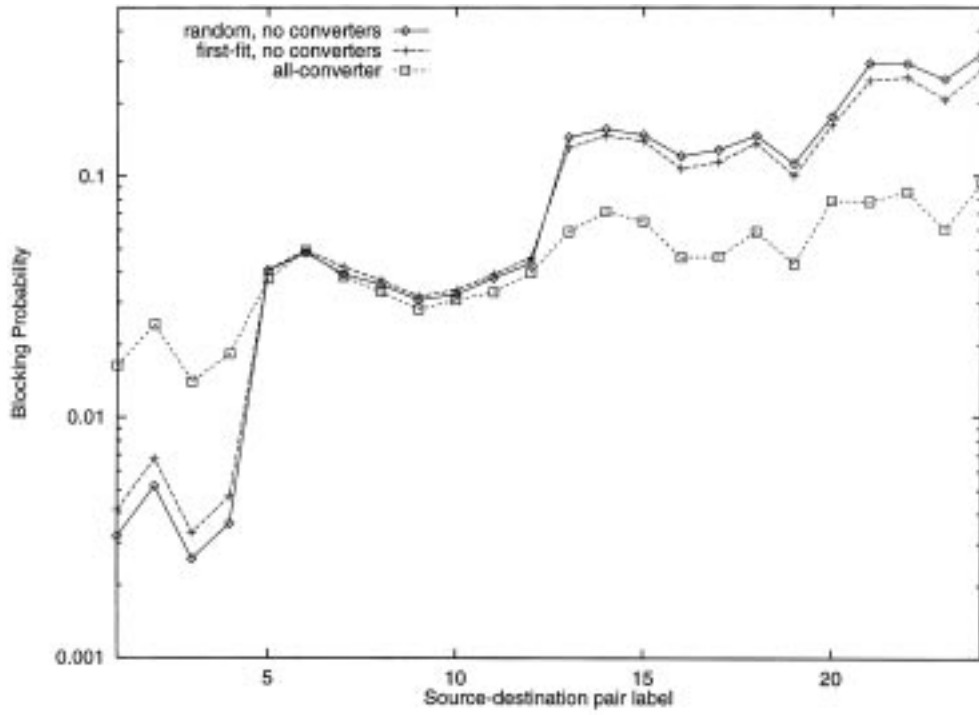


Fig. 24. Policy comparison, 5×5 torus network, random traffic pattern.

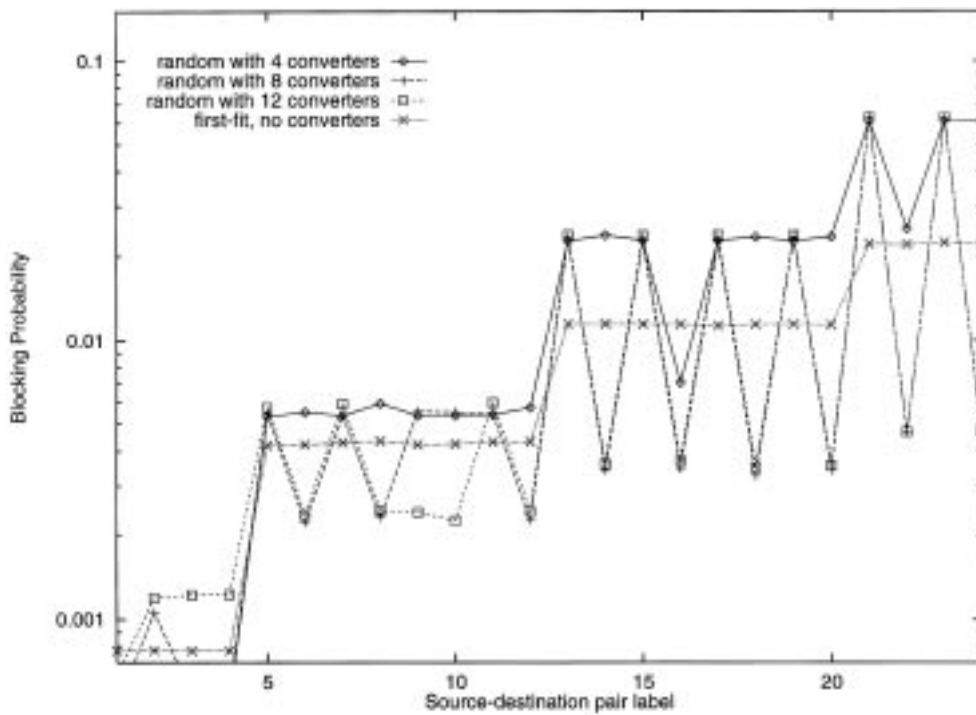


Fig. 25. First-fit policy vs. random policy with converters, 5×5 torus network, pattern based on locality.

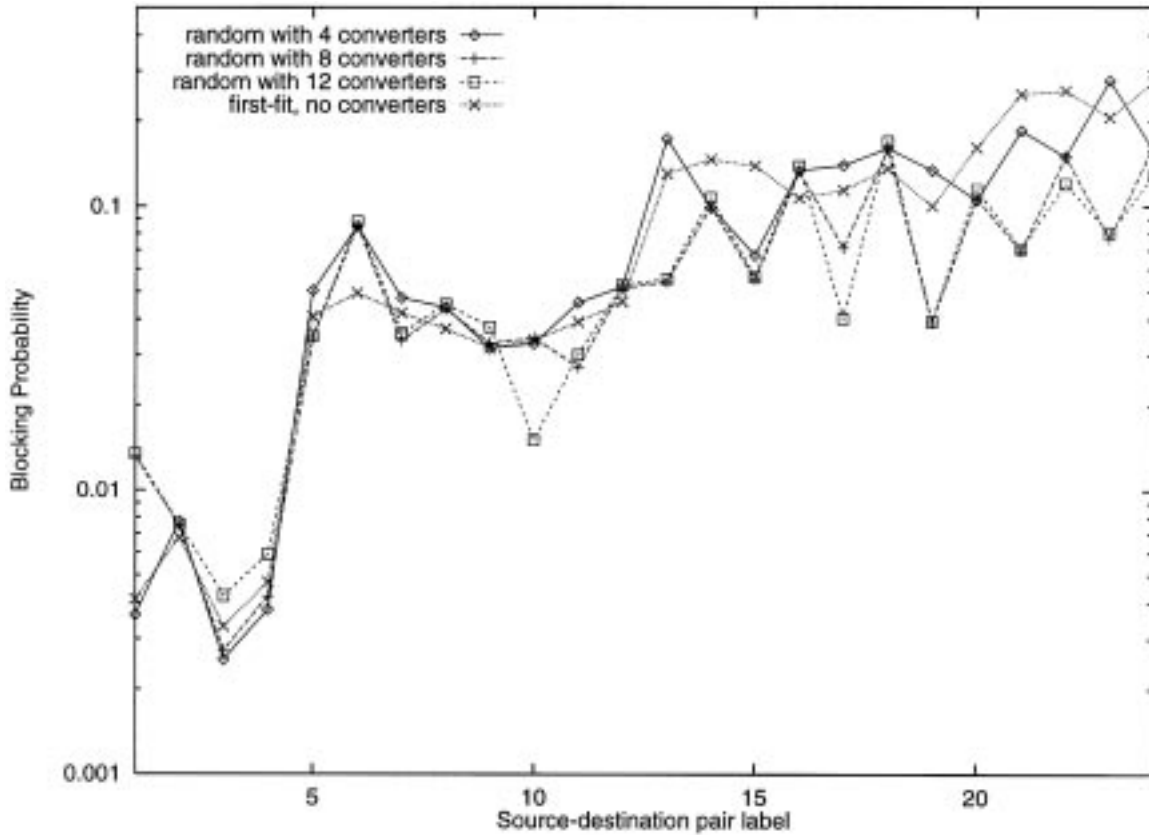


Fig. 26. First-fit policy vs. random policy with converters, 5 × 5 torus network, random traffic pattern.

hand, decreases evenly the blocking probability for calls traveling over long paths. This result is evident not only in Figs. 25 and 26 but in the results for the 6-hop and 10-hop paths, as well as for the NSFNET discussed next.

3.2 The NSFNET Topology

We have also considered a realistic example of a backbone network with an irregular topology, namely, the NSFNET shown in Fig. 27. Since we used the traffic data reported in [5], following that study, we have also augmented the 14-node NSFNET topology with two fictitious nodes, nodes 1 and 16 in Fig. 27, to capture the effect of NSFNET’s connections to Canada’s communication network, CA*net. The resulting topology consists of 16 nodes and a total of 240 source-destination pairs. As in the previous subsection, we only present detailed results for the blocking probabilities of only a small number of calls, those involving nodes along the path (3,5,6,7,9,12,15,16). (We note, however, that the

shortest path used by some of these calls is not a sub-path of (3,5,6,7,9,12,15,16); for instance, the shortest path for calls between nodes 3 and 15 is (3,5,11,15).) There are 28 source-destination pairs in this path, and in Figs. 28 to 32 they have been labeled so that numbers 1 to 7 refer to pairs one-hop paths, numbers 8 to 15 correspond to pairs with two-hop paths, etc. (refer to Table 2).

We have used two different traffic patterns with the NSFNET topology. The first traffic pattern is similar to one of the two patterns used with the torus network. Specifically, the arrival rate λ_{sd} for a source-destination pair (s,d) is given

$$\lambda_{sd} = \begin{cases} 0.5, & \text{if the length of the path from } s \text{ to } d \text{ is 1} \\ 0.4, & \text{if the length of the path from } s \text{ to } d \text{ is 2} \\ 0.3, & \text{if the length of the path from } s \text{ to } d \text{ is 3} \\ 0.2, & \text{if the length of the path from } s \text{ to } d \text{ is 4} \end{cases} \quad (4)$$

The second traffic pattern was designed to reflect actual traffic statistics collected on the NSFNET backbone network, as reported in the traffic matrix in

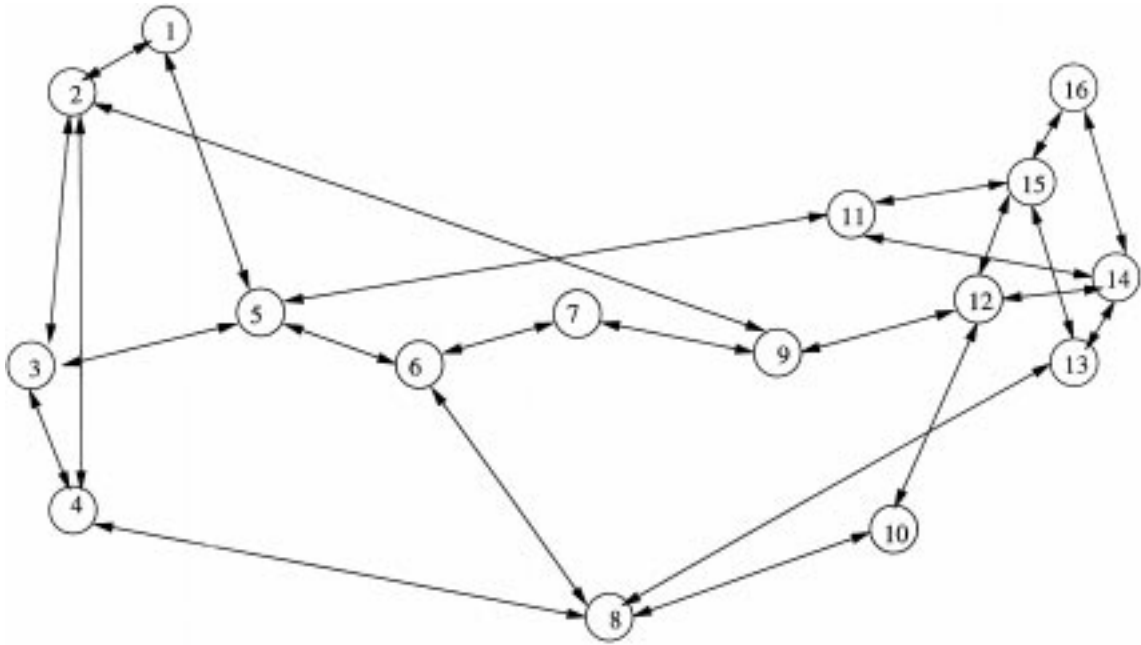


Fig. 27. The NSFNET topology.

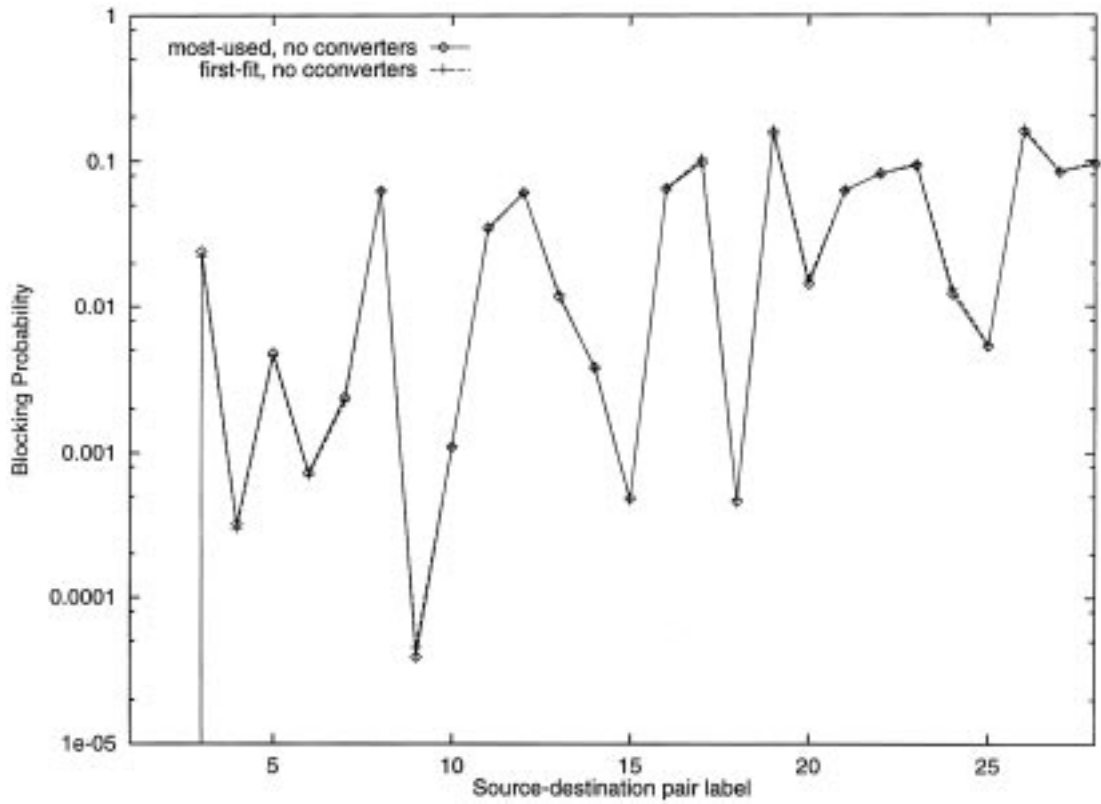


Fig. 28. Most-used vs. first-fit allocation, NSFNET, pattern based on actual traffic.

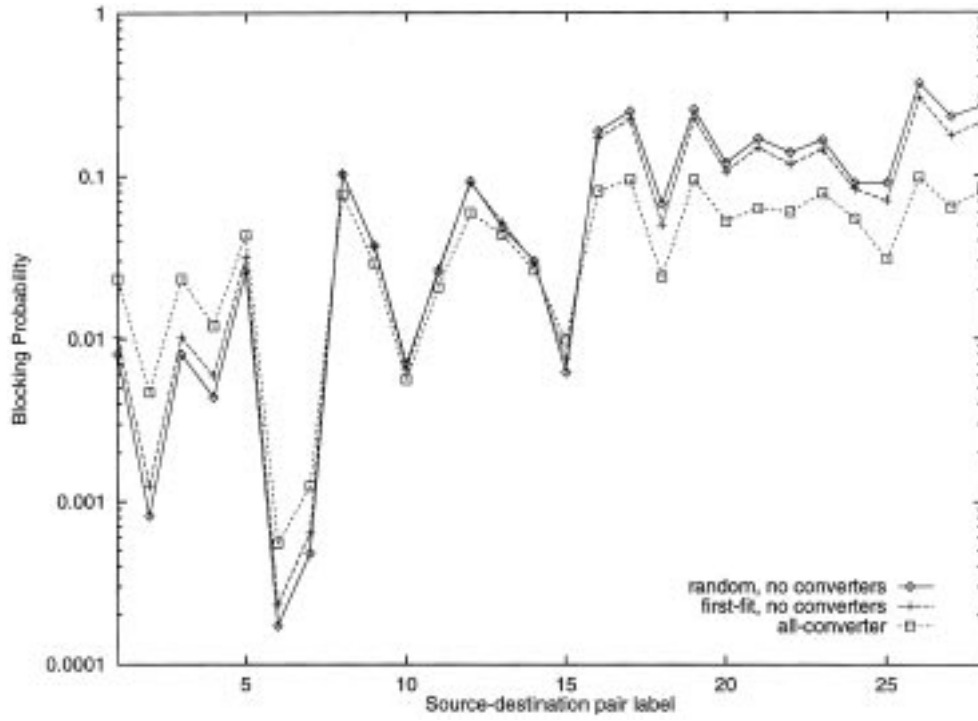


Fig. 29. Policy comparison, NSFNET, traffic pattern based on locality.

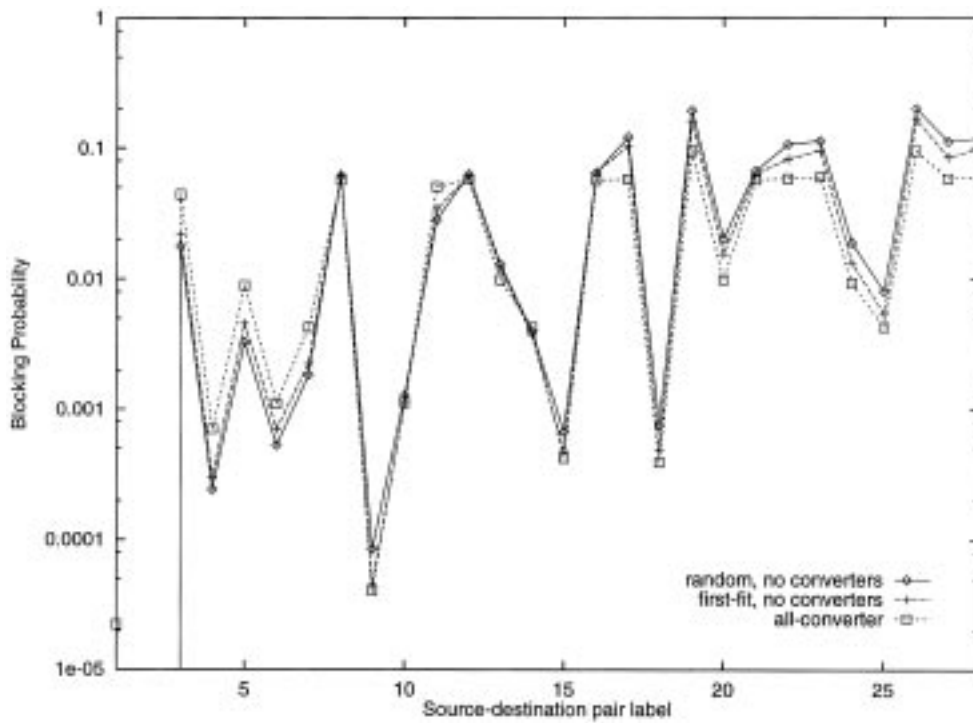


Fig. 30. Policy comparison, NSFNET, pattern based on actual traffic.

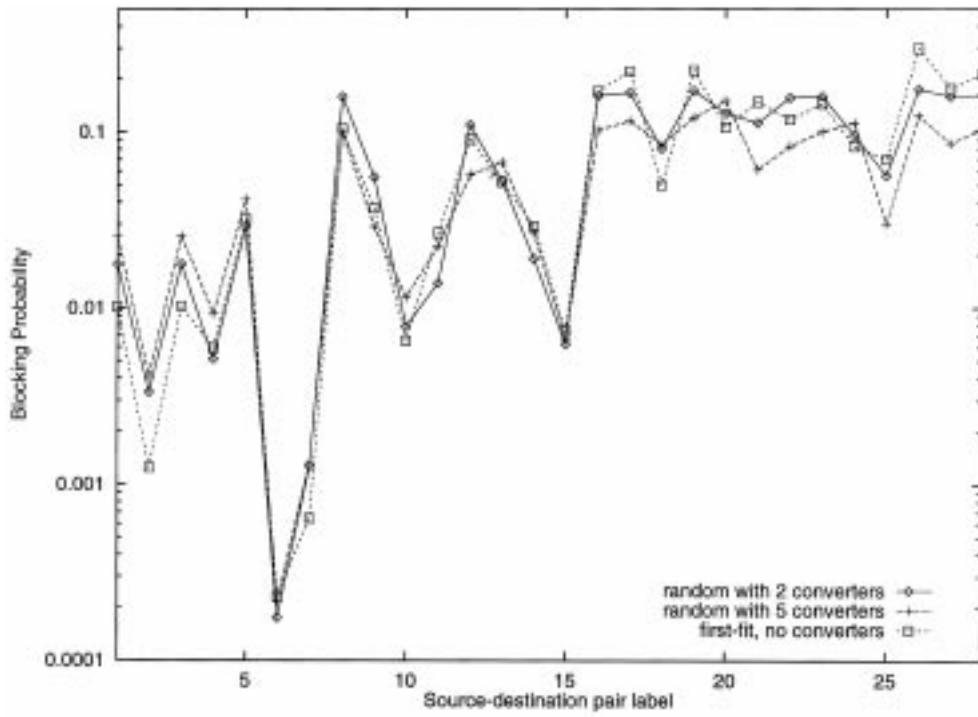


Fig. 31. First-fit policy vs. random policy with converters, NSFNET, traffic pattern based on locality.

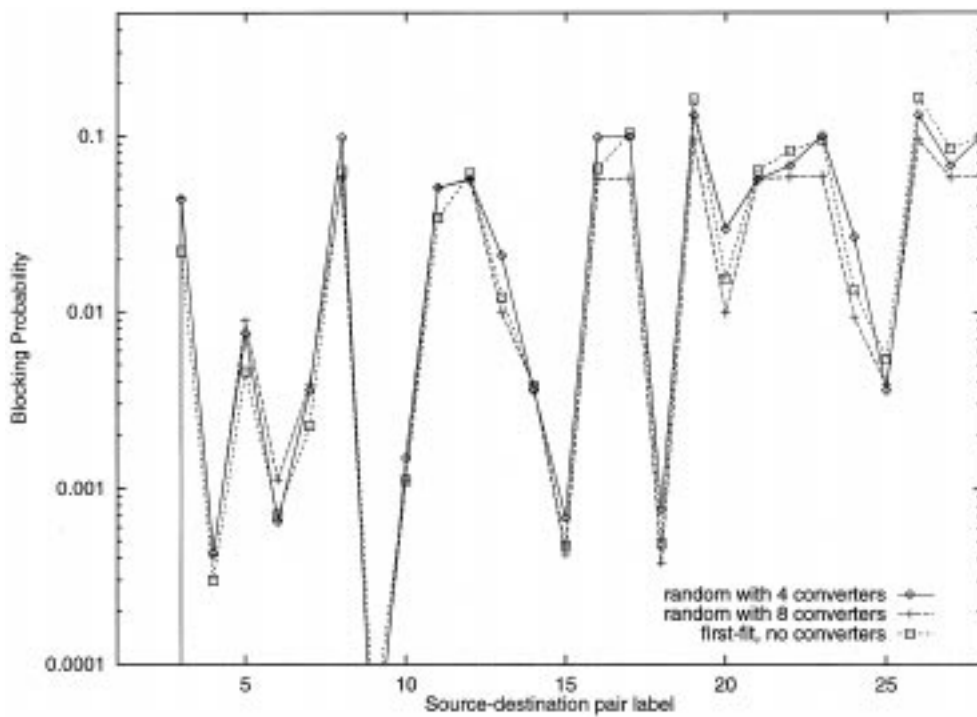


Fig. 32. First-fit policy vs. random policy with converters, NSFNET, pattern based on actual traffic.

Table 2. Selected source-destination pairs for the NSFNET topology.

Pair	(5,6)	(15,16)	(6,7)	(12,15)	(9,12)	(7,9)	(3,5)	(5,15)	(5,7)	(6,9)
Label	1	2	3	4	5	6	7	8	9	10
Shortest										
Path Length	1	1	1	1	1	1	1	2	2	2
Pair	(12,16)	(9,15)	(7,12)	(3,6)	(3,9)	(5,16)	(5,12)	(5,9)	(6,15)	(6,12)
Label	11	12	13	14	15	16	17	18	19	20
Shortest										
Path Length	2	2	2	2	2	3	3	3	3	3
Pair	(9,16)	(7,15)	(3,15)	(3,12)	(3,7)	(6,16)	(7,16)	(3,16)		
Label	21	22	23	24	25	26	27	28		
Shortest										
Path Length	3	3	3	3	3	4	4	4		

[5, Fig. 6]. The data in this traffic matrix represent the measured number of bytes transferred between two nodes in the NSFNET backbone within a certain 15-minute interval. Clearly, this data, collected over a packet-switched network, cannot be directly applied to a circuit-switched wavelength routing network, such as the one considered in this work. However, our intention is simply to capture the relative traffic demands among the different source-destination pairs. To this end, we first divided the entries of the matrix in [5, Fig. 6] by the link capacity (T3 links) to obtain the ‘‘offered load’’ ρ_{sd} per source-destination pair. Since the resulting values were too small, we multiplied them by a constant to obtain reasonable values for the offered load. Then, assuming that all calls have a mean holding time $1/\mu = 1$, the offered load values become the arrival rates λ_{sd} used in the experiments. As a result, the relative values of these arrival rates reflect the relative traffic requirements among the different source-destination pairs according to the specific traffic pattern reported in [5].

Our results are presented in Figs. 28 to 32. Fig. 28 compares the first-fit to the most-used policies. Figs. 29 and 30 demonstrate that the random and all-converter cases provide upper and lower bounds on the performance of the first-fit policy, and Figs. 31 and 32 attempt to quantify the effect of the first-fit policy in terms of the number of converters. The converters were placed in the network using the optimization techniques in [20]. The overall behavior of the graphs shown in these figures is very similar to that discussed earlier for the torus network and the single path case, indicating that our observations and conclusions are valid for a wide range of network topologies and traffic patterns.

4 Concluding Remarks

We have studied the blocking performance of several wavelength allocation policies for various single path and network topologies and under various traffic patterns. Our conclusions can be summarized as follows:

- We have shown that the most-used and first-fit policies have very similar call blocking probabilities for all calls in a network, regardless of the number of hops used by the calls. The two policies tend to favor calls using long paths at the expense of calls using short paths. This is a desirable feature, since calls traversing multiple hops experience higher blocking probability. However, the most-used policy requires that the network nodes exchange information about the network-wide usage of wavelengths, while the first-fit policy only relies on a fixed ordering of wavelengths, thus making it significantly easier to implement.
- We have also demonstrated that the random policy without converters and with converters at all nodes provide lower and upper bounds on the call blocking probability under the first-fit (or most-used) policy. Specifically, for calls using one or two hops, the random policy without converters provides a lower bound and the all-converter case provides an upper bound, while for calls using longer paths the bounds are reversed.
- We have presented results which indicate that the call blocking probabilities under the first-fit policy are similar to that under the random

policy but employing a number of converters (between 15% to 50% of the number of nodes) in the network. More importantly, in most cases, introducing the first-fit policy results in a decrease in the blocking probability of calls traveling over multiple hops to a level very close to the blocking probability experienced under the all-converter case. Note that, in terms of implementation, there is no significant difference between the first-fit and random policies. Consequently, the gains obtained by employing specialized (and expensive) hardware can be realized by making more intelligent choices in software.

- It also appears that the benefits of the first-fit policy diminish at high loads (blocking probability values of 0.1 or more). It is in these situations that employing converters would benefit calls traversing a large number of hops. However, the number of converters to be employed in this case must be very large, close to the number of nodes in the network, and even if all nodes contain converters the blocking probability will remain at (reduced but) high levels. Since it is unlikely that future wavelength routing networks will be designed to operate with such high call blocking probabilities, attempting to reduce the call blocking probabilities in this case may not be of practical importance.

Overall, our results appear to contradict previous studies which have indicated that “sparse” wavelength conversion capabilities (i.e., selective placement of converters in a subset of network nodes) will be beneficial to wavelength routing networks. Those studies measured the improvement obtained by employing converters in conjunction only with the random wavelength allocation policy. We have shown that an equivalent improvement can be achieved merely by using appropriate allocation policies such as first-fit or most-used.

References

- [1] R. A. Barry, P. A. Humblet, Models of blocking probability in all-optical networks with and without wavelength changers, *IEEE Journal Selected Areas in Communications*, vol. 14, no. 5, (June 1996), pp. 858–867.
- [2] A. Birman, Computing approximate blocking probabilities for a class of all-optical networks, *IEEE Journal Selected Areas in Communications*, vol. 14, no. 5, (June 1996), pp. 852–857.
- [3] I. Chlamtac, A. Ganz, G. Karmi, Lightpath communications: An approach to high bandwidth optical WANS, *IEEE Transactions on Communications*, vol. 40, no. 7, (July 1992), pp. 1171–1182.
- [4] R. Cruz, G. Hill, A. Kellner, R. Ramaswami, G. Sasaki, Y. Yamabayashi (eds.), Special issue on optical networks, *IEEE Journal Selected Areas in Communications*, vol. 14, no. 5, (June 1996).
- [5] B. Mukherjee, et al., Some principles for designing a wide-area WDM optical network, *IEEE/ACM Transactions on Networking*, vol. 4, no. 5, (October 1996), pp. 684–696.
- [6] R. E. Wagner, et al., MONET: Multiwavelength optical networking, *Journal of Lightwave Technology*, vol. 14, no. 6, (June 1996), pp. 1349–1355.
- [7] S. P. Monacos, et al., All-optical WDM packet networks, *Journal of Lightwave Technology*, vol. 14, no. 6, (June 1996), pp. 1356–1370.
- [8] M. Fujiwara, M. Goodman, M. O’Mahony, O. Tonguz, A. Willner, (eds.), Special issue on multiple wavelength technologies and networks, *Journal of Lightwave Technology*, vol. 14, no. 6, (June 1996).
- [9] A. Girard, *Routing and Dimensioning in Circuit-Switched Networks*. Reading, MA, (Addison Wesley, 1990).
- [10] H. Harai, M. Murata, H. Miyahara, Performance of alternate routing methods in all-optical switching networks. In *Proceedings of INFOCOM ’97*, IEEE, (April 1997), pp. 517–525.
- [11] I. P. Kaminow, C. R. Doerr, C. Dragone, T. Koch, U. Koren, A. A. M. Saleh, A. J. Kirby, C. M. Ozveren, B. Schoffield, R. E. Thomas, R. A. Barry, D. M. Castagnozzi, V. W. S. Chan, B. R. Hemenway, D. Marquis, S. A. Parikh, M. L. Stevens, E. A. Swanson, S. G. Finn, R. G. Gallager, A wideband all-optical WDM network, *IEEE Journal Selected Areas in Communications*, vol. 14, no. 5, (June 1996), pp. 780–799.
- [12] E. Karasan, E. Ayanoglu, Effects of wavelength routing and selection algorithms on wavelength conversion gain in WDM optical networks, *IEEE/ACM Transactions on Networking*, vol. 6, no. 2, (April 1998), pp. 186–196.
- [13] M. Kovacevic, A. Acampora, Benefits of wavelength translation in all-optical clear-channel networks, *IEEE Journal Selected Areas in Communications*, vol. 14, no. 5, (June 1996), pp. 868–880.
- [14] A. Mokhtar, M. Azizoglu, Adaptive wavelength routing in all-optical networks, *IEEE/ACM Transactions on Networking*, vol. 6, no. 2, (April 1998), pp. 197–206.
- [15] B. Ramamurty, B. Mukherjee, Wavelength conversion in WDM networking, *IEEE Journal Selected Areas in Communications*, vol. 16, no. 7, (September 1998), pp. 1061–1073.
- [16] S. Subramaniam, M. Azizoglu, A. Somani, All-optical networks with sparse wavelength conversion, *IEEE/ACM Transactions on Networking*, vol. 4, no. 4, (August 1996), pp. 544–557.
- [17] S. Subramaniam, M. Azizoglu, A. K. Somani, On the optimal placement of wavelength converters in wavelength-routed networks. In *Proceedings of INFOCOM ’98*, IEEE (April 1998), pp. 902–909.

- [18] S. Subramanian, A. K. Somani, M. Azizoglu, R. A. Barry, A performance model for wavelength conversion with non-poisson traffic. In Proceedings of INFOCOM '97, IEEE, (April 1997), pp. 500–507.
- [19] Y. Zhu, G. N. Rouskas, H. G. Perros, Blocking in wavelength routing networks, Part I: The single path case. In Proceedings of INFOCOM '99, IEEE, (March 1999), pp. 321–328.
- [20] Y. Zhu, G. N. Rouskas, H. G. Perros, Blocking in wavelength routing networks, Part II: Mesh topologies. In Proceedings of the Sixteenth International Teletraffic Congress (ITC 16), Elsevier Science, (June 1999), pp. 1321–1330.

Yuhong Zhu received the B.S. degree in Electrical Engineering from Beijing University in 1990, and the Ph.D. degree in Computer Science from the North Carolina State University, Raleigh, in 1999. He is currently a Member of Technical Staff in the Optical Network Group at Lucent Technologies, Acton, MA.

His interests include WDM networks, IP Multicast, MPLS and network performance evaluation.



George N. Rouskas received the Diploma in Electrical Engineering from the National Technical University of Athens (NTUA), Athens, Greece, in 1989, and the M.S. and Ph.D. degrees in Computer Science from the College of Computing, Georgia Institute of Technology, Atlanta, GA, in 1991 and 1994, respectively.

He joined the Department of Computer Science, North Carolina State University in August 1994, and has been an Associate Professor since July 1999. His research interests include network architectures and protocols, optical networks, multicast communication, and performance evaluation.

He is a recipient of a 1997 NSF Faculty Early Career



Development (CAREER) Award, and a co-author of a paper that received the Best Paper Award at the 1998 SPIE conference on All-Optical Networking. He also received the 1995 Outstanding New Teacher Award from the Department of Computer Science, North Carolina State University, and the 1994 Graduate Research Assistant Award from the College of Computing, Georgia Tech. He is a co-guest editor for the IEEE Journal on Selected Areas in Communications, Special Issue on Protocols and Architectures for Next Generation Optical WDM Networks, and is on the editorial boards of the IEEE/ACM Transactions on Networking and the Optical Networks Magazine. He is a member of the IEEE, the ACM and of the Technical Chamber of Greece.

Harry G. Perros received the B.Sc. degree in Mathematics in 1970 from Athens University, Greece, the M.Sc. degree in Operational Research with Computing from Leeds University, UK, in 1971, and the Ph.D. degree in Operations Research from Trinity College, Dublin, Ireland, in 1975. From 1976 to 1982 he was an Assistant Professor in the Department of Quantitative Methods, University of Illinois at Chicago. In 1979 he spent a sabbatical term in INRIA, Rocquencourt, France. In 1982 he joined the Department of Computer Science, North Carolina State University, as an Associate Professor, and since 1988 he is a professor. During the academic year 1988–89 he was on a sabbatical leave of absence first at BNR, Research Triangle Park, North Carolina, and subsequently at the University of Paris 6, France. Also, during the academic year 1995–96 he was on a sabbatical leave of absence at Nortel, Research Triangle Park, North Carolina.

He has published extensively in the area of performance modeling of computer and communication systems, and has organized several national and international conferences. He also published a monograph entitled "Queueing networks with blocking: exact and approximate solutions", Oxford Press. He is the chairman of the IFIP W.G. 6.3 on Performance of Communication Systems. His current research interests are in the areas of optical networks and their performance, and software performance evaluation.

