

A Comparison of Allocation Policies in Wavelength Routing Networks

Yuhong Zhu^a, George N. Rouskas^b, Harry G. Perros^b

^aLucent Technologies, Acton, MA

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

ABSTRACT

We consider wavelength routing networks with and without wavelength converters, and several wavelength allocation policies. We show through numerical and simulation results that the blocking probabilities for the random wavelength allocation and the circuit-switched case provide 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. Furthermore, we demonstrate that using the most-used or first-fit policies has an effect 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).

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 rates⁷. A wavelength routing network consists of wavelength routers and the fiber links that interconnect them. 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 wavelength¹¹. 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.

The problem of computing call blocking probabilities under static (fixed or alternate) routing with random wavelength allocation and with or without wavelength converters has been studied in^{1,9,2,6,12,14}. The model presented in¹ 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. In⁹ it was assumed that statistics of link loads are mutually independent, an approximation that is not accurate for sparse network topologies. In² 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 in⁶. 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 in¹² 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 in¹⁴, 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^{3,9}, and it was shown to perform better than random allocation, while an analytical overflow model for first-fit allocation was developed in⁸. A dynamic routing algorithm that selects the least loaded path-wavelength pair was also studied in⁸, and in¹⁰ an unconstrained dynamic routing scheme with a number of wavelength allocation policies was evaluated. Except in^{12,13}, all other studies assume that either all or none of the wavelength routers have wavelength conversion capabilities. The work in¹² takes a probabilistic approach

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

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. Finally, in¹³ 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^{9,2,6,14,8,10,13} make the assumption that link blocking events are independent and amount to the well-known *link decomposition* approach⁵. 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 tackle complex network design problems.

We have considered the problem of computing call blocking probabilities in mesh wavelength routing networks with fixed and alternate routing and random wavelength allocation in¹⁶. 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 network¹⁵. 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. Our algorithms can also 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. 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. We also demonstrate that the random policy and the circuit-switched case (i.e., a system with a converter in each node), for which analytical solutions exist for networks of large size, provide lower and upper bounds on the call blocking probability under the first-fit and most-used policies. 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.

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

2. A SINGLE PATH OF A NETWORK

We consider a single path of a wavelength routing network. 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 . 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 . 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 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.
- *Most-used allocation*: the wavelength that is already in use on the largest number of links in the path 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*: the call is assigned to the wavelength that is currently used in the smallest number of links in the path, with ties broken arbitrarily. This policy results in wavelength fragmentation, leading to higher blocking probability for calls using 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, the above allocation policies are used 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 these wavelength allocation policies, we will also consider the following case:

- *Circuit-switched paths*: paths in which there are converters at all nodes. In circuit-switching, 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. Consequently, wavelength allocation is not an issue under circuit-switching.

In our study, we have used a number of different traffic load patterns to compare the four wavelength allocation policies against each other and against circuit-switching. These patterns are representative of the wide range of loading situations that one expects to encounter in practice, and can be found in¹⁷. To ensure that the results are comparable across the different patterns, the load values were chosen so that the total load is the same for all patterns.

2.1. Policy Comparison for a Single Path of a Network

We have shown in¹⁵ that the evolution of a 2-hop path with random wavelength allocation can be characterized by the 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 the number of wavelengths that are free on both links of the path. The state transition diagram of this Markov process is shown in Figure 1 for $W = 2$ wavelengths, and it is straightforward to see that the process is not time-reversible¹⁵. 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. We have demonstrated in¹⁵ 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. (Under the random wavelength allocation policy, the

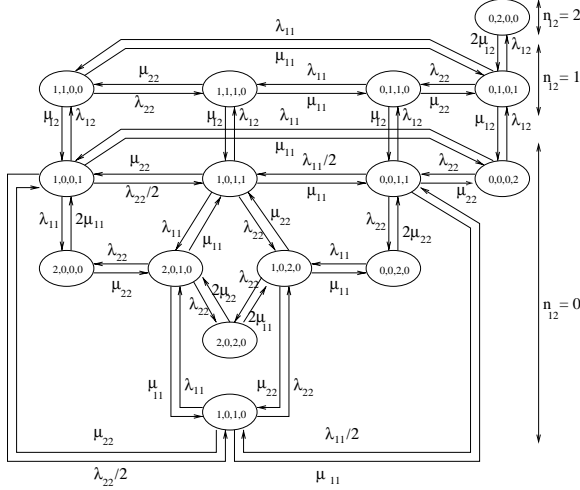


Figure 1. State space $(n_{11}, n_{12}, n_{22}, f_{12})$ of a 2-hop path with $W = 2$ wavelengths (random allocation)

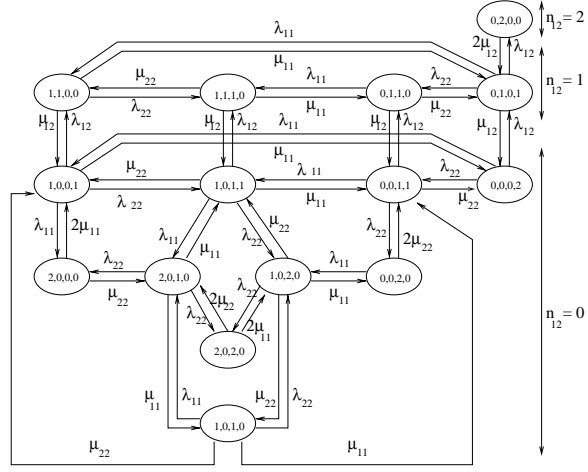


Figure 2. State space $(n_{11}, n_{12}, n_{22}, f_{12})$ of a 2-hop path with $W = 2$ wavelengths (most-used allocation)

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 Figure 2 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 Figure 1, 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 Figure 1, but there is no such transition in Figure 2. Furthermore, since there is a transition from state $(1,0,1,0)$ to state $(0,0,1,1)$ in Figure 2, 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 policy. Although we do not have an approximate 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.

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 Figure 3, and it can be easily verified that the process is not time-reversible.

If a converter is placed at node 1 of a 2-hop path (the only interesting possibility in this case), the system becomes equivalent to a 2-hop circuit-switched 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 using 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 process has a closed-form solution. In Figure 4 we show the state space of a 2-hop circuit switched path with two wavelengths. Although this path is described by the above 3-dimensional process, we include in the state description of Figure 4 the variable f_{12} to make it easier to compare to Figures 1–3. 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 using both hops cannot be blocked in the circuit-switched path, and the system will never enter state $(1,0,1,0)$, only state $(1,0,1,1)$.

The first-fit wavelength allocation policy can also be modeled as a Markov process, but the size of its 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.

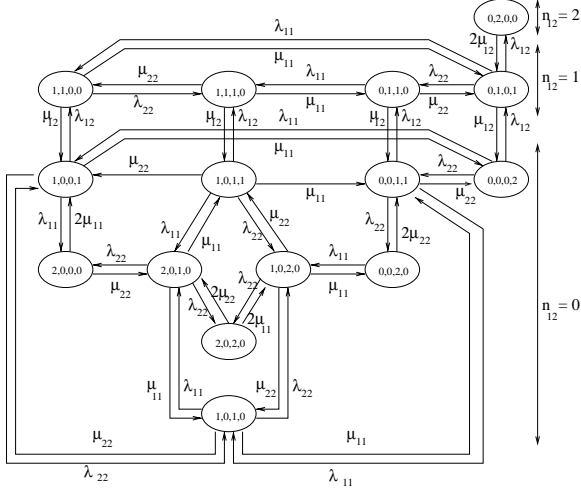


Figure 3. State space $(n_{11}, n_{12}, n_{22}, f_{12})$ of a 2-hop path with $W = 2$ wavelengths (least-used allocation)

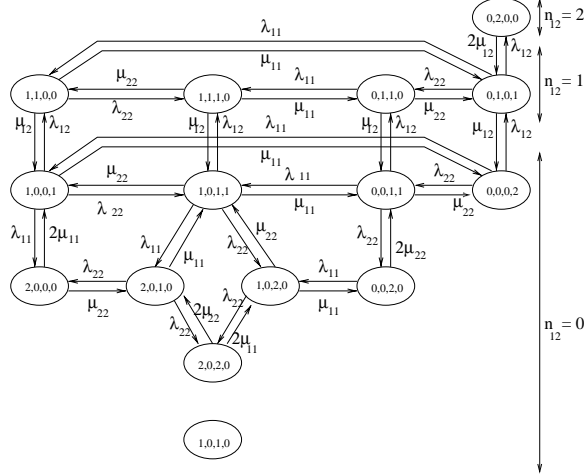


Figure 4. State space (n_{11}, n_{12}, n_{22}) of a 2-hop path with $W = 2$ wavelengths (circuit-switched)

2.1.1. Numerical Comparisons

Let us first consider the blocking probabilities of the random, most-used, least-used, and circuit-switched systems for calls traversing both links of the 2-hop path. In Figures 1 to 4, 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 cases. 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 Figure 4, 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 a circuit-switched path. In Figure 2 (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 circuit-switched case. The Markov process in Figure 1 (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 Figure 3 (least-used policy) has the same transitions as the one in Figure 1, 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 circuit-switched 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 load values. Figures 5 and 6 show results for two cases corresponding to a uniform and descending load pattern, respectively, and for $W = 10$ wavelengths. More specifically, the arrival rates used to obtain the results in Figure 5 were $\lambda_{11} = 0.2, \lambda_{12} = 0.1, \lambda_{22} = 0.2$, while for the results in Figure 6 we used $\lambda_{11} = 3.0, \lambda_{12} = 2.0, \lambda_{22} = 2.0$. In both figures 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

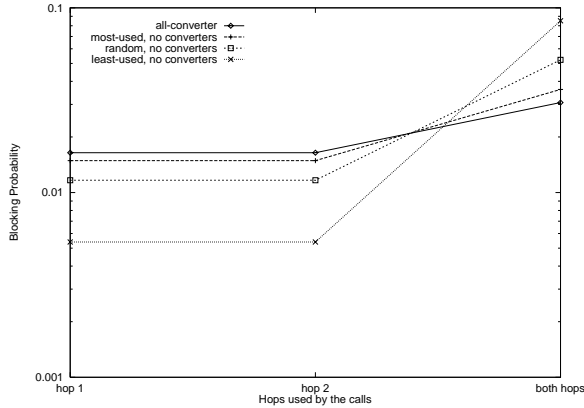


Figure 5. Policy comparison, 2-hop path, uniform traffic pattern

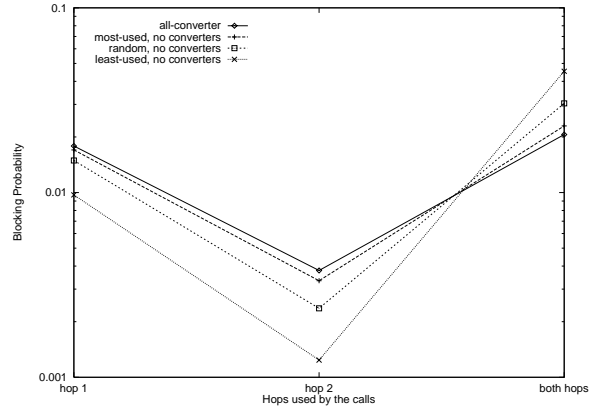


Figure 6. Policy comparison, 2-hop path, descending traffic pattern

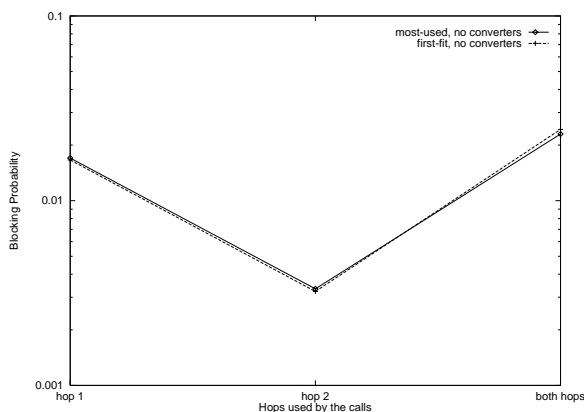


Figure 7. Most-used vs. first-fit allocation, 2-hop path, descending traffic pattern

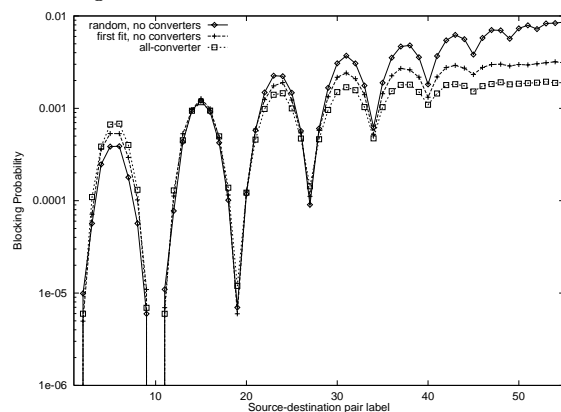


Figure 8. Policy comparison, 10-hop path, inverted bowl traffic pattern

pattern used. For instance, under uniform loading (Figure 5), calls using the first hop only experience the same blocking probability as hops using the second hop only, while in the descending pattern (Figure 6), 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.

In Figure 7 we compare the most-used and first-fit policies for the same arrival rates as those used for Figure 6. 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 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 Figure 7, 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.

The relative behavior of the four policies for longer paths is very similar to shown in Figures 5 and 7 for a wide range of traffic patterns. Specifically, for calls using one or two hops of a path only, the least-used policy provides the lowest blocking probability, followed by the random policy, the most-used policy, and the circuit-switched case. However, for calls traversing three or more hops of the path, the situation is reversed. Due to space limitations, results for paths longer than two hops are omitted, but can be found in¹⁷. Since, under the least-used policy, the

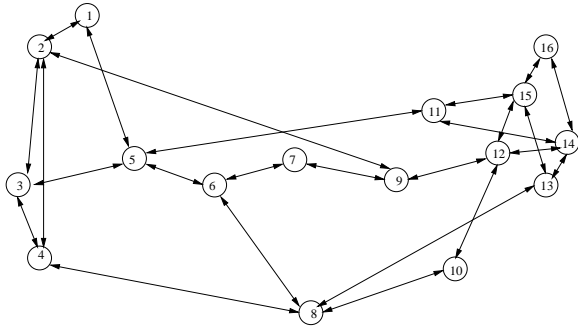


Figure 9. The NSFNET topology

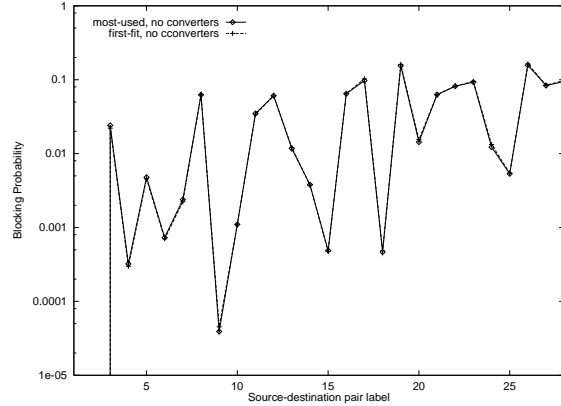


Figure 10. Most-used vs. first-fit allocation, NSFNET, pattern based on actual traffic

blocking probability of calls using multiple hops increases significantly, we will not consider the least-used policy any further.

Finally, in Figure 8 we compare the first-fit policy to the random (no converters) and circuit-switched cases for a 10-hop path (more results for 10-hop paths can be found in¹⁷). Two interesting observations can be made. First, the blocking probability values of the first-fit policy are always between the corresponding values of the random and circuit-switched cases. In other words, the blocking probability values under the random and circuit-switched cases provide lower and upper bounds for the blocking performance of the first-fit policy. Second, 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 circuit-switched case. Very similar results have been obtained for the other traffic patterns¹⁷.

3. MESH WAVELENGTH ROUTING NETWORKS

In this section we consider the NSFNET irregular topology in Figure 9. Results for other topologies can be found in¹⁷. Since we use the traffic data reported in⁴, following that study, we have augmented the 14-node NSFNET topology with nodes 1 and 16 in Figure 9. We present detailed results for the blocking probabilities of calls involving nodes along the path (3,5,6,7,9,12,15,16). There are 28 source-destination pairs in this path, and in Figures 10 to 14 they have been labeled so that numbers 1 to 7 refer to pairs with one-hop paths, numbers 8 to 15 correspond to pairs with two-hop paths, etc.

We have used two traffic patterns. For the first pattern, the call arrival rates are $\lambda_{sd} = 0.6 - l$, where l is the length of the shortest path from s to d . The second traffic pattern was designed to reflect actual traffic statistics collected on the NSFNET backbone network⁴. Clearly, this data, collected over a packet-switched network, cannot be directly applied to a circuit-switched wavelength routing network. 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 in⁴ by the link capacity 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.

Figure 10 compares the first-fit to the most-used policies, and we again see that that the two policies result in almost identical blocking probability values for all calls. Figures 11 and 12 demonstrate that the random and circuit-switched cases provide upper and lower bounds on the performance of the first-fit policy, similar to the single-path cases studied above. Finally, in Figures 13 and 14 we compare the first-fit policy to the random policy with converters. The converters were placed in the network using the optimization techniques in¹⁶. As can be seen, using the first-fit policy is roughly equivalent to employing a significant number of converters in the network. The overall behavior of the graphs in these figures is very similar to the single path case as well as other topologies (see¹⁷), indicating that our observations and conclusions are valid for a wide range of topologies and traffic patterns.

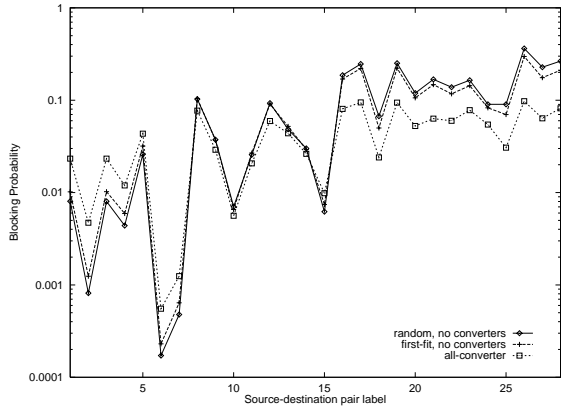


Figure 11. Policy comparison, NSFNET, traffic pattern based on locality

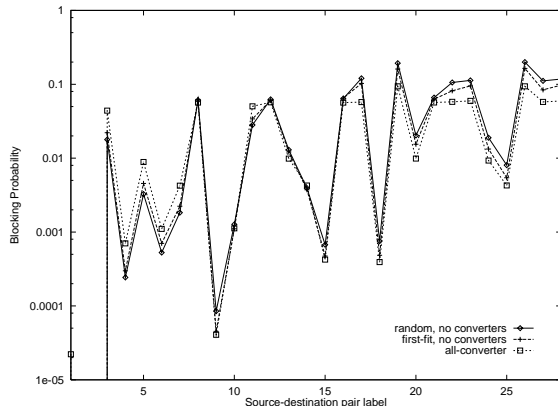


Figure 12. Policy comparison, NSFNET, pattern based on actual traffic

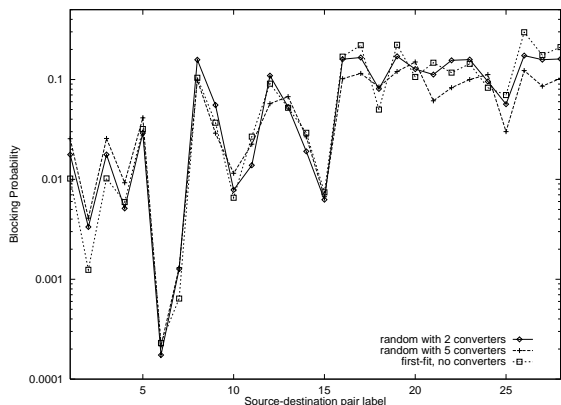


Figure 13. First-fit policy vs. random policy with converters, NSFNET, traffic pattern based on locality

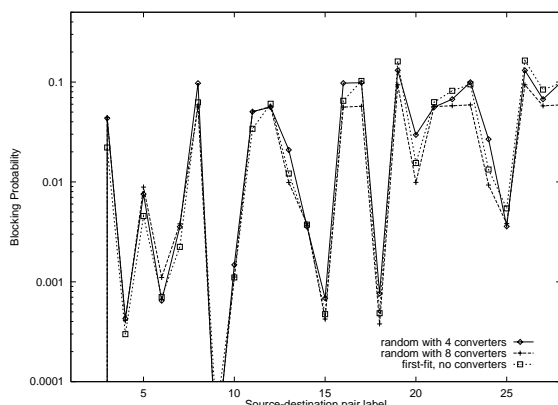


Figure 14. First-fit policy vs. random policy with converters, NSFNET, pattern based on actual traffic

4. CONCLUDING REMARKS

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 multiple paths at the expense of calls using a single path. This is a desirable feature, since calls traversing multiple paths experience the highest 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, and is significantly easier to implement.

We have also shown that the random policy and the circuit-switched case provide 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 provides a lower bound and the circuit-switched 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 those under the random policy but employing a number of converters in the network. 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 circuit-switched case. Since, in terms of implementation, there is no significant difference between the first-fit and random policies, the gains obtained by employing 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. 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 at such high call blocking probabilities, reducing the call blocking probabilities in this case may not be of practical importance.

While previous studies of “sparse” wavelength conversion have 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 and P. A. Humblet. Models of blocking probability in all-optical networks with and without wavelength changers. *IEEE J-SAC*, 14(5):858–867, June 1996.
2. A. Birman. Computing approximate blocking probabilities for a class of all-optical networks. *IEEE J-SAC*, 14(5):852–857, June 1996.
3. I. Chlamtac, A. Ganz, and G. Karmi. Lightpath communications: An approach to high bandwidth optical WANS. *IEEE Trans. Commun.*, 40(7):1171–1182, July 1992.
4. B. Mukherjee *et al.* Some principles for designing a wide-area WDM optical network. *IEEE/ACM Trans. Networking*, 4(5):684–696, October 1996.
5. A. Girard. *Routing and Dimensioning in Circuit-Switched Networks*. Addison Wesley, Reading, MA, 1990.
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. A. Hill, A. Salek, and K. Sato (Eds.). Special issue on high-capacity optical transport networks. *IEEE J-SAC*, 16(7), Sep. 1998.
8. E. Karasan and E. Ayanoglu. Effects of wavelength routing and selection algorithms on wavelength conversion gain in wdm optical networks. *IEEE/ACM Trans. Networking*, 6(2):186–196, April 1998.
9. M. Kovacevic and A. Acampora. Benefits of wavelength translation in all-optical clear-channel networks. *IEEE J-SAC*, 14(5):868–880, June 1996.
10. A. Mokhtar and M. Azizoglu. Adaptive wavelength routing in all-optical networks. *IEEE/ACM Trans. Networking*, 6(2):197–206, April 1998.
11. B. Ramamurty and B. Mukherjee. Wavelength conversion in WDM networking. *IEEE J-SAC*, 16(7):1061–1073, Sep. 1998.
12. S. Subramaniam, M. Azizoglu, and A. Somani. All-optical networks with sparse wavelength conversion. *IEEE/ACM Trans. Networking*, 4(4):544–557, Aug. 1996.
13. 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, Apr. 1998.
14. S. Subramaniam, A. K. Somani, M. Azizoglu, and R. A. Barry. A performance model for wavelength conversion with non-poisson traffic. *Proc. INFOCOM '97*, pp. 500–507, Apr. 1997.
15. 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.
16. 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.
17. 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. Tech. Rep. TR-99-01, NCSU, Raleigh, NC, Jan. 1999.