

Blocking in Wavelength Routing Networks, Part II: Mesh Topologies *

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

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

We study a class of circuit switched wavelength routing networks with fixed or alternate routing and with random wavelength allocation. We present an iterative path decomposition algorithm to evaluate accurately and efficiently the blocking performance of such networks with and without wavelength converters. The path decomposition approach naturally captures the correlation of both link loads and link blocking, giving accurate results for a wide range of loads and network topologies. Our model also allows non-uniform traffic, i.e., call request arrival rates that can vary with the source-destination pair, and it can be used when the location of converters is fixed but arbitrary.

Keywords: Wavelength division multiplexing, wavelength routing networks, call blocking probability, decomposition algorithms

1. Introduction

A wavelength routing network consists of wavelength routers and the fiber links that interconnect them [6, 5, 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, that is, 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 [13, 1]. Wavelength conversion is a desirable feature since it improves the performance of the network in terms of call blocking probability.

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, 11, 2, 9, 13, 15]. Other wavelength allocation schemes, as well as dynamic routing are harder to analyze. First-fit wavelength allocation was studied using simulation in [3, 11], and it was shown to perform better than random allocation, while an analytical overflow model for first-fit allocation was developed in [10]. A dynamic routing algorithm that selects the least loaded path-wavelength pair was also studied in [10], and in [12] an unconstrained dynamic routing scheme with a number of wavelength allocation policies was evaluated. Except in [13, 14], all other studies assume that either all or none of the wavelength routers have wavelength conversion capabilities.

Most of the approximate analytical techniques developed for computing blocking probabilities in wavelength routing networks [11, 2, 9, 15, 10, 12, 14] make the assumption that link blocking events are independent and amount to the well-known *link decomposition* approach [8], while the development of some techniques is based on the additional assumption that link loads are also independent. Link decomposition has been extensively used in conventional circuit switched networks where there is no requirement for the *same* wavelength to be used on successive links of the path taken by a call. The accuracy of these underlying approximations also depends on the traffic load, the network topology, and the routing and wavelength allocation schemes employed. While link decomposition techniques make it possible to study the qualitative behavior

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

of wavelength routing networks, more accurate analytical tools are needed to both evaluate efficiently the performance of these networks, as well as to tackle complex network design problems, such as selecting the optical switches where to employ wavelength converters.

In this paper we consider the problem of computing call blocking probabilities in mesh wavelength routing networks with fixed and alternate routing and random wavelength allocation. We develop an iterative *path decomposition* algorithm [8] for analyzing arbitrary network topologies. Specifically, we analyze a given network by decomposing it into a number of path sub-systems. Each sub-system is then analyzed in isolation using our algorithm for a single path in a wavelength routing network [17]. 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 arrival rates vary for each source-destination pair. Finally, our algorithm can compute call blocking probabilities in a mesh network where a fixed, but arbitrary, subset of nodes are equipped with wavelength converters.

The rest of the paper is organized as follows. In Section 2 we present a new path decomposition algorithm for analyzing a mesh wavelength routing network topology under both fixed and alternate routing. In Section 3 we validate our algorithm through simulation, and we conclude the paper in Section 4.

2. Path Decomposition Algorithm for Mesh Networks

2.1. Network Model

We consider a circuit-switched wavelength routing network with an arbitrary topology. There are N nodes and L (unidirectional) links in the network, with each link supporting W wavelengths. Call requests between a source node s and a destination node d arrive at the source node according to a Poisson process with a rate of λ_{sd} . The call holding times are exponentially distributed with mean $1/\mu$, and the quantity $\rho_{sd} = \lambda_{sd}/\mu$ is the offered load of the calls.

In wavelength routing networks, there are two parts in the routing problem. When a call request arrives, a path over which the connection will be established must first be determined. In this work we consider both fixed and alternate routing [8]. In *fixed* routing, each source-destination pair is assigned to a single path. If there are no wavelength converters in the path, a call is blocked if there is no wavelength which is free on all links of the path. If some nodes in the path employ wavelength converters, a call is blocked if no wavelength is free on all the links of any segment of the path consisting of the links between successive nodes with converters. In *alternate* routing, a set of paths (consisting of one primary path and one or more alternate paths) is assigned to each source-destination pair. This set is searched in a fixed order to find an available path for the call. Once a path is selected, one of the (possibly many) free wavelengths in the path must then be assigned to the call. As in [17], we only consider the random wavelength assignment policy in this work, whereby a call is allocated one of the available wavelengths in the selected path at random.

We let \mathcal{R} denote the set of paths assigned to all the source-destination pairs. For fixed routing, $|\mathcal{R}| = N(N-1)$. If alternate routing with m paths (one primary and $m-1$ alternates) for each source-destination pair is used, then $|\mathcal{R}| = mN(N-1)$. We also let $P_{sd}^{(n)}, 1 \leq n \leq m$, denote the probability that a call originating at node s and terminating at node d will be blocked on the n -th path assigned to this source-destination pair.

2.2. Fixed Routing

We analyze a mesh network by decomposing it into a number of sub-systems where each sub-system is a single path. Each sub-system is analyzed in isolation using the analytical techniques developed in [17]. Specifically, a sub-system consisting of three links or less is analyzed by solving a time-reversible Markov process, which is obtained approximately from the underlying Markov process of the sub-system. Sub-systems




Figure 1. (a) Original network, (b) set \mathcal{R}' of paths into which the network is decomposed

longer than three hops are analyzed using an iterative decomposition algorithm to obtain the call blocking probabilities. The individual solutions are appropriately combined (as explained shortly) by modifying the call arrival rates to each sub-system to reflect the newly computed blocking probabilities. The process is repeated until all blocking probabilities converge within a prescribed tolerance.

Before we proceed we emphasize that the number of sub-systems into which the network is decomposed is significantly smaller than the total number $|\mathcal{R}|$ of paths. This is because many of the shorter paths are completely contained within other longer paths². Therefore, these shorter paths do not need to be considered as separate sub-systems. Instead, the blocking probability of these paths is obtained as a by-product of the computation of the blocking probabilities of a long sub-system. Since a k -hop path may contain up to $(k+2)(k-1)/2$ shorter paths as sub-paths, by selecting long paths as sub-systems we can drastically reduce the number of sub-systems into which the original network is decomposed.

The first step in analyzing a given network is to decompose it into a set $\mathcal{R}' \subseteq \mathcal{R}$ of paths such that: (1) no path $r \in \mathcal{R}'$ is contained within a path $q \in \mathcal{R}$, $q \neq r$, and (2) any path $q \in \mathcal{R}$ either belongs to \mathcal{R}' or is completely contained within a path $r \in \mathcal{R}'$. These two requirements ensure that a minimal set of sub-systems that includes all possible paths is used. We can construct such a set \mathcal{R}' by using the following steps. First, the paths in \mathcal{R} are sorted in a list in order of decreasing length. The first path r in the list is removed and inserted in \mathcal{R}' . Then, any sub-paths of r that are also in the list are removed from it. The process continues with the next path in the list and is repeated until the list becomes empty. It is straightforward to show that this algorithm will construct a set \mathcal{R}' which satisfies the above two properties. Figure 1(a) shows a 5-node network. Without loss of generality, we assume that shortest paths are used for fixed routing in this network. The set of sub-systems \mathcal{R}' obtained by applying the above algorithm to this network is shown in Figure 1(b). As we can see, while there are 20 source-destination pairs and corresponding paths in the network, only 10 path sub-systems are used. The blocking probability on the path from, say, node 1 to node 4, is obtained through the solution to the sub-system corresponding to the path from node 1 to node 3.

²We say that a path q is completely contained within another path r if q is a sub-path of r .

Once the set \mathcal{R}' of sub-systems has been selected, for each path $r \in \mathcal{R}'$ we determine the set of paths $\mathcal{S}(r) \subseteq \mathcal{R}$ that intersect, i.e., have at least one link in common, with path r . As an example, path (1,4,3) in Figure 1 intersects with path (4,3,5). The significance of set $\mathcal{S}(r)$ lies in the fact that the blocking probability experienced by calls using the links of path r may be affected by the calls using the links of a path $q \in \mathcal{S}(r)$, and vice versa. Thus, when we compute the solution for path r , we must appropriately modify the call arrival rates along this path to account for the effect of calls along paths that intersect with r . Note also that $q \in \mathcal{S}(r)$ implies that $r \in \mathcal{S}(q)$.

We are now ready to present the decomposition algorithm used to analyze a wavelength routing network with an arbitrary topology. We illustrate the operation of the algorithm using the network shown in Figure 1. We show how to update the arrival rates along each path sub-system after each iteration of the algorithm by considering only paths (1,4,3) and (4,3,5). The other path sub-systems are handled in a similar way. Recall that $\lambda_{sd}, 1 \leq s, d \leq N$, are the arrival rates to the original network. We also let $\hat{\lambda}_{sd}$ denote the arrival rates used to solve the various path sub-systems; these rates are updated at the beginning of each iteration of the algorithm. As will be explained next, the rate $\hat{\lambda}_{sd}$ accounts for all calls in the original network that use the links between nodes s and d within a path r .

Using the algorithms in [17], we initially solve path (1,4,3) in Figure 1 in isolation using the following arrival rates:

$$\hat{\lambda}_{14} = \lambda_{14}, \quad \hat{\lambda}_{13} = \lambda_{13}, \quad \hat{\lambda}_{43} = \lambda_{43} + (1 - P_{45})\lambda_{45} \quad (1)$$

We note that only calls from node 1 to node 4 use link (1,4) of path (1,4,3), thus, the arrival rate of calls using this link as seen by the path sub-system (1,4,3) is given by the first expression in (1). Similarly, the second expression in (1) can be explained by the fact that only calls from node 1 to node 3 use both links of sub-system (1,4,3). On the other hand, the last expression in (1) for $\hat{\lambda}_{43}$ is slightly different because, in addition to calls from node 4 to node 3, calls from node 4 to node 5 also use the second link of path (1,4,3) since paths (1,4,3) and (4,3,5) intersect. Quantity P_{45} in (1) represents the current estimate of the probability that a call from node 4 to node 5 will be blocked on sub-system (4,3,5). For the first iteration, we use $P_{45} = 0$; how this value is updated in subsequent iterations will be discussed shortly. Therefore, the term $(1 - P_{45})\lambda_{45}$ represents the *effective* arrival rate of calls from node 4 to node 5 as seen by sub-system (1,4,3), since the fraction $P_{45}\lambda_{45}$ of these calls will be blocked in sub-system (4,3,5). Consequently, the right hand side of the third expression in (1) is the effective arrival rate of calls that use the link (4,3) of path (1,4,3) when the latter is considered in isolation.

We also solve path (4,3,5) in isolation by using the following arrival rates:

$$\hat{\lambda}_{45} = \lambda_{45}, \quad \hat{\lambda}_{35} = \lambda_{35}, \quad \hat{\lambda}_{43} = \lambda_{43} + (1 - P_{13})\lambda_{13} \quad (2)$$

The expressions in (2) can be explained using arguments similar to the ones used for the expressions in (1). In particular, the second term in the right hand side of the third expression in (2) represents the effective arrival rate of calls originating in sub-system (1,4,3) and using the link (4,3) of sub-system (4,3,5)

The solution to the path sub-systems (1,4,3) and (4,3,5) will yield an initial value for the probabilities P_{45} and P_{13} that a call using links (3,5) and (1,4), respectively, will be blocked. The new estimates for P_{45} and P_{13} are then used in expressions (1) and (2), respectively, to update the arrival rates for the two path sub-systems, the sub-systems are solved again and new estimates for the blocking probabilities are obtained, and so on. We repeat the process until the blocking probabilities for all calls in the original network converge within a certain tolerance.

A detailed description of our decomposition algorithm is provided in Figure 2. We note that this path decomposition algorithm not only accounts explicitly for the correlation of link loads among the various links of the network, but it also accounts for the fact that link blocking events are not independent. This is in sharp contrast to link decomposition algorithms for wavelength routing networks that have appeared in the

Decomposition Algorithm for Mesh Networks with Fixed Routing

Input: Network topology, set \mathcal{R} of paths for all source-destination pairs, and arrival rates λ_{sd}

Output: Call blocking probabilities P_{sd} for all source-destination pairs in the network

1. begin
 2. From \mathcal{R} construct the set of path sub-systems \mathcal{R}' into which the network will be decomposed, as described in Section 2.2
 3. For each $r \in \mathcal{R}'$ construct the set $\mathcal{S}(r) = \{q \in \mathcal{R}' \mid q \text{ intersects with } r\}$
 4. $h \leftarrow 0$ // Initialization step
 $P_{sd}(h) \leftarrow 0 \quad \forall s, d$ // All blocking probabilities initialized to zero
 5. $h \leftarrow h + 1$ // h -th iteration
 For each path $r = (r_1, r_2, \dots, r_k) \in \mathcal{R}'$ do // compute the arrival rates for this iteration
 For each path $q = (q_1, \dots, r_i, \dots, r_j, \dots, q_m) \in \mathcal{S}(r)$ that intersects with r from node r_i to r_j do
 // Calls using path q affect the blocking probability of calls using path r ; the call arrival rate
 // seen by path r must be increased appropriately to account for the effect of these calls
 $\hat{\lambda}_{r_i, r_j}(h) \leftarrow \hat{\lambda}_{r_i, r_j}(h) + (1 - P_{q_1, q_m}(h - 1))\lambda_{q_1, q_m}$
 Solve each path sub-system $r \in \mathcal{R}'$ using the algorithms in [17] to obtain new values for the
 blocking probabilities $P_{sd}(h)$
 7. Repeat from Step 5 until the blocking probabilities converge
 8. end of the algorithm
-

Figure 2. A summary of the path decomposition algorithm

literature (e.g., see [11,2,9,15,10,12,14]) which compute the blocking probability along a path by assuming that blocking events on each link of the path are independent. In all the cases we have studied, we have found that the algorithm converges in only a few (less than ten) iterations, and that the blocking probabilities obtained closely match simulation results (the performance of the decomposition algorithm will be discussed in detail in Section 3).

2.3. Alternate Routing

In order to improve the call blocking performance, a source-destination pair (s, d) may be assigned m paths, one primary and $m - 1$ alternates, which are searched in a fixed order. Typically, in implementations, the m shortest paths from s to d in the physical topology are used. If a call is blocked on the primary path, the first alternate path is examined. If available wavelengths exist on this path, the call is established. Otherwise, the next alternate path is examined, and so on. In other words, the traffic offered to alternate path $i, i = 2, \dots, m$, is the *overflow* traffic from path $i - 1$. The call is blocked if no free wavelength can be found on any of the m paths, i.e., if it overflows from the last alternate path.

Although the traffic offered to the primary path for source-destination pair (s, d) is Poisson with rate λ_{sd} , it is clear that the overflow traffic offered to the alternate paths is not Poisson. The overflow model is a well-known model that has been studied extensively in the literature, and moment matching techniques have been used to analyze blocking probabilities in circuit-switched networks with alternate routing [8]. Overflow models have also been used in the study of blocking probabilities in wavelength routing networks in [10,12]. Below we describe our approach to computing call blocking probabilities in networks with alternate routing by assuming that there is one primary and one alternate path per source-destination pair. In order to make use of the path decomposition algorithm developed in the previous subsection, we will assume that overflow traffic is also Poisson with an appropriate rate.

Let \mathcal{R} denote the set of primary and alternate paths for all node pairs, with $|\mathcal{R}| = 2N(N-1)$. From \mathcal{R} we construct the set of path sub-systems \mathcal{R}' as described in the previous subsection. We analyze the path sub-systems in \mathcal{R}' using the algorithm of Figure 2 to obtain an initial estimate of the call blocking probabilities $P_{sd}^{(1)}$ and $P_{sd}^{(2)}$ for the primary and alternate paths, respectively. The arrival rate for the overflow traffic offered to alternate paths is simply taken to be the product of the arrival rate of the traffic to the primary path times the blocking probability of this path. Also, if a primary path r intersects with an alternate path q , the arrival rate on the alternate path q (primary path r) is taken into account when solving path r (path q). This approach captures the effect that calls established over alternate (primary) paths have on calls established over primary (alternate) paths.

Once estimates for the blocking probabilities $P_{sd}^{(1)}$ and $P_{sd}^{(2)}$ have been obtained, an estimate of the blocking probability of calls for the source-destination pair (s, d) can be computed as

$$P_{sd} = P_{sd}^{(1)} \times P_{sd}^{(2)} \quad (3)$$

These estimates are used to update the arrival rates of calls to the network, and the decomposition is solved again. This process is repeated until the blocking probabilities P_{sd} in (3) converge for all s, d .

Extensive validation has shown that, despite the assumption that overflow traffic is Poisson, this iterative path decomposition approach is quite accurate for both regular and irregular topologies, and for low to moderate loads. Although we do expect our decomposition algorithm to yield less accurate results under high or very high offered loads (due to the peakedness of the overflow traffic), it is unlikely that wavelength routing networks will be operated at loads which will result in high call blocking probabilities. Our approximate approach works well for blocking probability values as high as 0.5, which we feel is significantly higher than the blocking probabilities that can be tolerated in such an environment.

3. Numerical Results

We have applied our iterative decomposition algorithm to a variety of regular and irregular topologies and a wide range of traffic loads (see [16]). Due to lack of space, we only present results for the NSFNET topology shown in Figure 3. Since we use traffic data reported in [4], we have augmented the 14-node NSFNET topology by two nodes, nodes 1 and 16 in Figure 3, 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. We present detailed results for the blocking probabilities of calls involving nodes along the path (3,5,6,7,9,12,15,16). The 28 source-destination pairs in this path, along with the corresponding shortest path lengths and the labels used in Figures 4 through 7 are shown in Table 1. A summary of the results for the whole network can be found in [16].

We have used two different traffic patterns with the NSFNET topology. The first traffic pattern is such that the arrival rate λ_{sd} for a source-destination pair (s, d) is given by:

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

This selection of arrival rates is intended to capture the locality of traffic that has been observed in many networks. The second traffic pattern was designed to reflect actual traffic statistics collected on the NSFNET backbone network, as reported in the traffic matrix in [4, Figure 6]. The data in this traffic matrix represent the measured number of bytes transferred from a node s to a node d in the NSFNET backbone within a 15-minute interval. This data 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




Figure 3. The NSFNET topology

demands among the different source-destination pairs. To this end, we first divide the entries of the matrix in [4, Figure 6] by the link capacity (T3 links) to obtain the “offered load” ρ_{sd} per source-destination pair. Since the resulting values are too small, we multiply 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 [4].

Figures 4 and 5 present the call blocking probabilities for the selected pairs of Table 1 and the first traffic pattern. Figure 4 shows results for fixed routing, while Figure 5 shows results for alternate routing with two alternate paths per source-destination pair. Because of the irregular topology, the alternate paths for some of the calls are 6 hops long. The link utilizations are in the range [1.846, 5.668] with an average of 3.494 under fixed routing, while for alternate routing they are slightly higher, in the range [1.964, 5.722] with an average of 3.646.

We also observe that calls established over longer paths tend to experience higher blocking probability than calls using short paths. However, because of the irregular topology, the blocking probability can be significantly affected by the actual load along the path taken by a call. For instance, we observe in Figure 4 that the blocking probabilities of calls established over, say, 1-hop paths vary widely depending on the number of other calls using the same path. Regarding the accuracy of the decomposition algorithm, we note that the curve obtained analytically closely follows the simulation curve for the 28 source-destination pairs shown in Figures 4 and 5. We also note that the assumption that overflow traffic is Poisson does not appear to affect the performance of our algorithm when alternate routing is used.

Figures 6 and 7 are similar to Figures 4 and 5, respectively, except that they present results for the second traffic pattern derived from the traffic statistics presented in [4]. The utilization under this traffic pattern is in the range [0.015, 8.059] with an average of 3.976 for fixed routing, and in the range of [0.014, 9.231] with an average of 4.168 for alternate routing with one alternate path per call. As we can see, the relative behavior of the two curves (obtained through the analytical techniques and simulation, respectively) in Figures 6 and 7

Table 1
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		

is very similar to that in Figures 4 and 5, and all our previous conclusions regarding the accuracy of our decomposition algorithm are still valid, despite the fact that some of the blocking probability values are as high as 0.5. For high blocking probability values the analytical results underestimate the simulation results, while for lower blocking probability values the analytical results overestimate the simulation results. Despite this behavior, the analytical and simulation results are always very close even at high loads.

Another interesting observation from Figures 6 and 7 is that, with the exception of a few source-destination pairs, using an alternate path does not have a significant impact on the call blocking performance. Consider, for instance, the source-destination pairs (6,7) and (12,16) (labels 3 and 11 in the figures). The primary path for pair (6,7) is short (one hop) but the alternate path is quite long (five hops), thus, using an alternate path does not improve the blocking probability experienced by this pair. On the other hand, both the primary and the alternate paths for pair (12,16) are two hops long, thus, using an alternate path reduces the overall blocking probability for this pair by about four orders of magnitude! In general, source-destination pairs for which the alternate paths are longer than the primary paths experience only a slight drop (if any) in blocking probability under alternate routing. Since the traffic load offered to the network is high and the blocking probabilities also high, adding more traffic to the network through alternate paths does not improve the performance except for pairs for which the alternate path is short and not highly utilized, as is the case with pair (12,16).

4. Concluding Remarks

We have presented a new path decomposition algorithm to evaluate accurately and efficiently the call blocking performance of wavelength routing network with an arbitrary topology. Our algorithm is applicable to networks with either fixed or alternate routing and random wavelength allocation. Our iterative algorithm analyzes the original network by decomposing it into single path sub-systems. These sub-systems are analyzed in isolation by using our previous algorithms for a single path of wavelength routing networks, and the individual results are appropriately combined to obtain a solution for the overall network. Our algorithm can also be applied to the problem of converter placement in wavelength routing networks.




Figure 4. Blocking probability for selected source-destination pairs in the NSFNET with $W = 10$ and fixed routing (first traffic pattern)




Figure 5. Blocking probability for selected source-destination pairs in the NSFNET with $W = 10$ and two alternate paths per pair (first traffic pattern)




Figure 6. Blocking probability for selected source-destination pairs in the NSFNET with $W = 10$ and fixed routing (second traffic pattern)




Figure 7. Blocking probability for selected source-destination pairs in the NSFNET with $W = 10$ and two alternate paths per pair (second traffic pattern)

REFERENCES

1. R. A. Barry and P. A. Humblet. Models of blocking probability in all-optical networks with and without wavelength changers. *IEEE Journal Selected Areas in Communications*, 14(5):858–867, June 1996.
2. A. Birman. Computing approximate blocking probabilities for a class of all-optical networks. *IEEE Journal Selected Areas in Communications*, 14(5):852–857, June 1996.
3. I. Chlamtac, A. Ganz, and G. Karmi. Lightpath communications: An approach to high bandwidth optical WANS. *IEEE Transactions on Communications*, 40(7):1171–1182, July 1992.
4. B. Mukherjee *et al.* Some principles for designing a wide-area WDM optical network. *IEEE/ACM Transactions on Networking*, 4(5):684–696, October 1996.
5. I. P. Kaminow *et al.* A wideband all-optical WDM network. *IEEE Journal Selected Areas in Communications*, 14(5):780–799, June 1996.
6. R. E. Wagner *et al.* MONET: Multiwavelength optical networking. *Journal of Lightwave Technology*, 14(6):1349–1355, June 1996.
7. S. P. Monacos *et al.* All-optical WDM packet networks. *Journal of Lightwave Technology*, 14(6):1356–1370, June 1996.
8. A. Girard. *Routing and Dimensioning in Circuit-Switched Networks*. Addison Wesley, Reading, MA, 1990.
9. H. Harai, M. Murata, and H. Miyahara. Performance of alternate routing methods in all-optical switching networks. In *Proceedings of INFOCOM '97*, pages 517–525. IEEE, April 1997.
10. E. Karasan and E. Ayanoglu. Effects of wavelength routing and selection algorithms on wavelength conversion gain in wdm optical networks. *IEEE/ACM Transactions on Networking*, 6(2):186–196, April 1998.
11. M. Kovacevic and A. Acampora. Benefits of wavelength translation in all-optical clear-channel networks. *IEEE Journal Selected Areas in Communications*, 14(5):868–880, June 1996.
12. A. Mokhtar and M. Azizoglu. Adaptive wavelength routing in all-optical networks. *IEEE/ACM Transactions on Networking*, 6(2):197–206, April 1998.
13. S. Subramaniam, M. Azizoglu, and A. Somani. All-optical networks with sparse wavelength conversion. *IEEE/ACM Transactions on Networking*, 4(4):544–557, August 1996.
14. S. Subramaniam, M. Azizoglu, and A. K. Somani. On the optimal placement of wavelength converters in wavelength-routed networks. In *Proceedings of INFOCOM '98*, pages 902–909. IEEE, April 1998.
15. S. Subramaniam, A. K. Somani, M. Azizoglu, and R. A. Barry. A performance model for wavelength conversion with non-poisson traffic. In *Proceedings of INFOCOM '97*, pages 500–507. IEEE, April 1997.
16. Y. Zhu, G. N. Rouskas, and H. G. Perros. Blocking in wavelength routing networks, Part II: Mesh topologies. Technical Report TR-98-10, North Carolina State University, Raleigh, NC, October 1998.
17. Y. Zhu, G. N. Rouskas, and H. G. Perros. Blocking in wavelength routing networks, Part I: The single path case. In *Proceedings of INFOCOM '99*. IEEE, March 1999.