

Blocking in Wavelength Routing Networks, Part I: The Single Path Case *

Yuhong Zhu George N. Rouskas Harry G. Perros

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

Abstract

We study a class of circuit-switched wavelength routing networks with and without wavelength converters, and we present the first part of a new analytical framework to accurately and efficiently evaluate the blocking performance of such networks. Our model allows non-uniform traffic, it accounts for the correlation among the loads at all links in a path, and it can be used when the location of converters is fixed but arbitrary. We first construct an exact Markov process that captures the behavior of a path in terms of wavelength use. We also obtain an approximate Markov process which has a closed-form solution that can be efficiently computed for short paths. We then develop an iterative algorithm to analyze approximately arbitrarily long paths. The algorithm decomposes a path into shorter segments which are then studied in isolation using the corresponding approximate Markov process. The individual solutions are appropriately combined to obtain a solution for the original path. Finally, we demonstrate how our analytical techniques can be used to gain insight into the problem of converter placement in wavelength routing networks.

1 Introduction

To take full advantage of the potential of optical fiber, the use of wavelength division multiplexing (WDM) techniques has become the option of choice, and WDM networks have been a subject of research both theoretically and experimentally. In particular, the wavelength routing mesh architecture appears promising for wide area distances. The network architecture consists of wavelength routers and fiber links that interconnect them. A wavelength router is capable of switching a light signal at a given wavelength from any input port to any output port. A router may also be capable of enforcing a shift in wavelength, in which case a light signal may emerge from the switch at a different wavelength than the one it arrived. By appropriately configuring the routers, all-optical paths (lightpaths) may be established between pairs of nodes in the network. Lightpaths represent direct optical connections without any intermediate electronics. Because of the long propagation delays, and the time required to configure the routers, wavelength routing WANs are expected to operate in circuit switched mode.

The call blocking performance of optical wavelength routing networks has been studied in [7, 1, 5, 2, 9]. In [5] it was assumed that statistics of link loads are mutually independent, while the model developed in [1] is based on the assumption that wavelength use on each link can be characterized by a fixed probability, independently of other wavelengths and links. The work in [2] used

a Markov chain with state-dependent arrival rates to model call blocking in arbitrary mesh topologies, a computationally intensive approach. Another model was presented in [7], where it was assumed that the load on link i of a path depends only on the load of link $i - 1$. A study of call blocking under non-Poisson traffic was presented in [9], but it assumes that link loads are statistically independent. The problem of optimal converter placement in a single path of a network was studied in [8], again assuming no correlation among link loads.

The above studies indicate that wavelength routing networks are rather hard to analyze. In this paper we present the first part of a new analytical framework which permits us to accurately study the call blocking performance of wavelength routing networks. Our work improves upon previously published work in several ways. First, we allow non-uniform traffic, i.e., call arrival rates may vary for each source-destination pair. Second, our model accounts for the calls between any source-destination pair in the path, and it can be used to compute the blocking probability not only for the end-to-end traffic, but also for the traffic along shorter segments of the path. More importantly, our model captures the correlation among the loads on *all* links of a path. Finally, the approximate analytical techniques we develop are very efficient and can be used for long paths one would encounter in realistic wide area networks.

The main results of our work are as follows. We first develop an exact Markov process model that captures the correlation of wavelength use among *all* links of a k -hop path with and without converters. We then show how to slightly modify this process to obtain an approximate Markov process model which is time-reversible and which has a closed-form solution that resembles the product form solution in queueing networks [4]. The solution to the time-reversible Markov process provides an accurate approximation to the blocking probabilities obtained through the original ones. Because of computational requirements, both the exact and the approximate Markov process models can only be applied to relatively short paths. For longer paths, we then develop an iterative algorithm for computing the blocking probabilities by decomposing a path into a series of shorter segments connected in tandem. Finally, we show how our analytical techniques can be used to gain insight into the problem of converter placement in a wavelength routing network.

In the following section we describe the wavelength routing network under study. In Section 3 we obtain the exact and approximate Markov process models for computing call blocking probabilities, and in Section 4 we develop an iterative decomposition algorithm for long paths. We present numerical results in Section

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




Figure 1: A k -hop path




Figure 2: A two-hop path

5, and we conclude in Section 6.

2 The Wavelength Routing Network

We consider a single path in a circuit-switched wavelength routing network, where each link supports W wavelengths. Call requests arrive according to a Poisson process with a rate that depends on the source-destination pair. If the request can be satisfied, an optical circuit is established between the source and destination for the duration of the call. Call holding times are exponentially distributed.

If a node in the path employs wavelength converters, it can switch an incoming wavelength to an arbitrary outgoing wavelength. (When there are converters at all nodes, the situation is identical to that in classical circuit-switching networks, a special case of the one studied here.) 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 call blocking. If a call cannot be established due to lack of wavelengths, the call is blocked. On the other hand, if a call can be accommodated, one of the wavelengths that are available on the links used by the call is randomly assigned to it. Thus, we only consider the random wavelength assignment policy.

We define a “segment” as a sub-path consisting of one or more consecutive links of the original path. We will use the following notation in this paper (refer to Figure 1):

- A k -hop path consists of $k + 1$ nodes labeled $0, 1, \dots, k$, and hop i , is the link between nodes $i - 1$ and i .
- $\lambda_{ij}, j \geq i$, is the Poisson arrival rate of calls that use hops i through j , i.e., they originate at node $i - 1$ and terminate at node j .
- $1/\mu$, is the mean of the (exponentially distributed) holding time of all calls. Also, $\rho_{ij} = \lambda_{ij}/\mu$ is the offered load of calls using hops i through j .
- $n_{ij}, j \geq i$, is the number of calls using hops i through j that are currently active in the network.
- $f_{ij}, j \geq i$, is the number of wavelengths that are free on all hops i through j .

3 Markov Process Models

3.1 Exact Markov Process Model

Let us first consider the 2-hop path (without converters) shown in Figure 2. It is straightforward to verify that the evolution of this system can be characterized by the four-dimensional Markov

process $(n_{11}, n_{12}, n_{22}, f_{12})$. Since, on each hop, the number of busy wavelengths plus the number of wavelengths that are free on both hops may not exceed W , the following two constraints must be satisfied:

$$n_{11} + n_{12} + f_{12} \leq W \quad \text{and} \quad n_{12} + n_{22} + f_{12} \leq W$$

This result can be generalized to k -hop paths, $k > 2$. Let \mathcal{M}_k denote the Markov process for a k -hop path. There are k^2 random variables in a state \underline{n} of \mathcal{M}_k , as follows:

$$\underline{n} = (n_{11}, \dots, n_{1k}, n_{22}, \dots, n_{2k}, \dots, n_{kk}, f_{12}, \dots, f_{1k}, f_{23}, \dots, f_{2k}, \dots, f_{k-1,k}) \quad (1)$$

The first $\frac{k(k+1)}{2}$ random variables $n_{ij}, 1 \leq i \leq j \leq k$, in (1) provide the number of active calls between all source-destination pairs in the path. The last $\frac{(k-1)k}{2}$ random variables $f_{ij}, 1 \leq i < j \leq k$, represent the number of free wavelengths on all segments consisting of two or more links. The constraints on the state space can be found in [10].

\mathcal{M}_k captures the correlation of wavelength use on all links of the path, and it can be used to provide an exact solution for the call blocking probability. However, the large number of random variables in its state description makes it impossible to numerically solve it for large k or W . In addition, the transition rates of \mathcal{M}_k are state-dependent. In Figure 3 we show the transition diagram of \mathcal{M}_2 for $W = 2$ wavelengths. From this figure we can see that there exist sequences of states for which Kolmogorov’s criterion for reversibility [3, Theorem 1.8] is not satisfied, and the Markov process is not time-reversible. Two such sequences are: $(1,1,1,0), (1,0,1,1), (1,0,0,1), (1,1,0,0)$, and $(1,1,1,0), (1,0,1,1), (0,0,1,1), (0,1,1,0)$. This result is true in general, and that $\mathcal{M}_k, k \geq 2$ is not time-reversible for $W > 1$.

Next, we show how to modify some of the transition rates of \mathcal{M}_k in order to obtain an approximate Markov process which is time-reversible and which has a closed-form solution.

3.2 Approximate Time-Reversible Markov Process Model

A closer examination of Figure 3 reveals that the two four-state sequences mentioned above are the shortest sequences for which Kolmogorov’s criterion is not satisfied. We also note that these sequences involve transitions that cause changes in the value of random variable n_{12} . Define $\mathcal{L}_{2,c}$ as the sub-chain of \mathcal{M}_2 that includes only states for which the value of random variable n_{12} is constant, i.e., $n_{12} = c$:

$$\mathcal{L}_{2,c} = \{(n_{11}, n_{12}, n_{22}, f_{12}) : n_{12} = c\}, \quad c = 0, \dots, W \quad (2)$$




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

Sub-chain $\mathcal{L}_{2,c}$ corresponds to a new with $W - c$ wavelengths per hop, in which no calls using both hops ever arrive ($\lambda_{12} = 0$). It can be easily verified that Kolmogorov's criterion for reversibility is satisfied by any sequence of states in sub-chain $\mathcal{L}_{2,c}$, and that \mathcal{M}_2 is not time-reversible due to transitions between states with different values of random variable n_{12} .

Returning to Figure 3, if the transition rate from state $(1,0,1,1)$ to state $(1,1,1,0)$ is changed to $2\lambda_{12}$, then the process becomes time-reversible. This is an important result because we can obtain a closed-form solution for the time-reversible process. However, when each hop supports more than 2 wavelengths, a larger number of transition rates must be modified to yield a time-reversible process. The rule for changing the transition rates can be found in [10].

For a k -hop path, $k \geq 2$, a time-reversible process \mathcal{M}'_k can be derived from \mathcal{M}_k as follows. \mathcal{M}'_k has the same state space and the same transitions as \mathcal{M}_k . The vast majority of its transition rates are the same as the respective transition rates of \mathcal{M}_k . However, to ensure that the new Markov process is time-reversible, the transition rates between some pairs of states must be appropriately modified as explained in [10].

\mathcal{M}'_k has a product-form solution for its steady-state probability [4]. Let $G_k(W)$ denote the normalizing constant for a k -hop path with W wavelengths per link. Then, the solution to \mathcal{M}'_3 , corresponding to a 3-hop path with state $\underline{n} = (n_{11}, n_{12}, n_{13}, n_{22}, n_{23}, n_{33}, f_{12}, f_{13}, f_{23})$ is given by:

$$\pi(\underline{n}) = \frac{1}{G_3(W)} \frac{\rho_{11}^{n_{11}} \rho_{12}^{n_{12}} \rho_{13}^{n_{13}} \rho_{22}^{n_{22}} \rho_{23}^{n_{23}} \rho_{33}^{n_{33}}}{n_{11}! n_{12}! n_{13}! n_{22}! n_{23}! n_{33}!} \frac{\binom{f_{11}}{f_{12}} \binom{n_{11}}{f_{22} - f_{12}}}{\binom{n_{11} + f_{11}}{f_{22}}}$$

$$\times \frac{\binom{f_{12}}{f_{13}} \binom{n_{22} + n_{12} + f_{22} - f_{12}}{f_{33} - f_{13}} \binom{f_{22} - f_{12}}{f_{23} - f_{13}} \binom{n_{22} + n_{12}}{f_{33} - f_{23}}}{\binom{n_{22} + n_{12} + f_{22}}{f_{33}} \binom{n_{22} + n_{12} + f_{22} - f_{12}}{f_{33} - f_{13}}}$$

where $f_{11} = W - n_{11} - n_{12} - n_{13}$, $f_{22} = W - n_{12} - n_{22} - n_{13} - n_{23}$, and $f_{33} = W - n_{13} - n_{23} - n_{33}$.

In general, the solution to \mathcal{M}'_k , $k > 3$, is the product of k^2 terms. The first $k(k+1)/2$ terms are of the form $\rho_{ij}^{n_{ij}}/n_{ij}$, corresponding to the random variables n_{ij} in (1). The last $k(k-1)/2$ combinatorial terms correspond to the dependent variables f_{ij} in the state description (1).

The significance of the new Markov process \mathcal{M}'_k will be illustrated in Section 5, where it will be shown that the blocking probabilities obtained through the product-form solution of \mathcal{M}'_k closely approximate the exact blocking probabilities obtained through the numerical solution of \mathcal{M}_k . In order for the closed-form solution to be useful, we need to have a computationally efficient procedure for calculating the normalizing constant. However, the computation of the normalizing constant in time that is polynomial in W and k has turned out to be a very difficult task. In Section 4, we describe an iterative decomposition algorithm that can be used to efficiently obtain the blocking probabilities for long paths.

When the path employs wavelength converters, it is straightforward to derive exact and approximate Markov processes similar to (but simpler than) the ones described above. Specifically, in [10] we present a method which can be used to obtain exact and approximate call blocking probabilities in a k -hop path when the placement of converters is known.

4 Decomposition Algorithm for Long Paths

Let K denote the largest integer such that the closed-form solution to Markov process \mathcal{M}'_K can be computed within a reasonably short amount of time. Consider a k -hop path. If $k \leq K$, the path can be analyzed approximately by solving the corresponding Markov process \mathcal{M}'_k . If, however, $k > K$, the approximate closed-form solution cannot be used directly. In this section, we develop an iterative decomposition algorithm to analyze paths of length greater than K . We first consider paths without converters, and we then extend the algorithm to handle wavelength conversion.

4.1 Paths With No Wavelength Conversion

We analyze a k -hop path, $k = lK + m$, $l \geq 1, m < K$, by decomposing it into one m -hop segment and l K -hop segments in tandem. Each segment is first analyzed in isolation using the corresponding Markov process \mathcal{M}'_m or \mathcal{M}'_K . The arrival rates of calls originating in a segment but terminating in another segment are accounted for by increasing the arrival rate of calls in the individual segments. The individual solutions are appropriately combined to obtain an initial value for the blocking probability of calls that traverse more than one segment. Using these initial estimates, the arrival rates to each segment are modified and each segment is again solved in isolation in order to obtain a new solution. These new individual solutions are again combined to update the blocking probability of calls traversing multiple segments. This is repeated until the blocking probabilities converge.

Decomposition Algorithm for Paths Without Converters

A k -hop path, $k = K + m$, $m \leq K$, is decomposed into a K -hop segment (Segment 1) and an m -hop segment (Segment 2). Segment 1 consists of nodes 0 to K , and Segment 2 consists of nodes K to $K = m$ of the original path. λ_{ij} refer to the call arrival rates in the original path, whereas $\lambda_{ij}^{(n)}$ refer to call arrival rates in Segment n , $n = 1, 2$.

1. begin
2. $h \leftarrow 0$ //Initialization step
 - // $p_{ij}^{(1)}(h)$ is the blocking prob. of calls using hops
 - // i to j of Segment 1; $F_{ij}^{(1)}$ is the average number
 - // of free wavelengths on hops i to j of Segment 1
 - $p_{ij}^{(1)}(h) \leftarrow 0$, $F_{ij}^{(1)} \leftarrow W$, $1 \leq i \leq j \leq K$
 - // $p_{ij}^{(2)}(h)$ is the blocking prob. of calls using hops
 - // i to j of Segment 2; $F_{ij}^{(2)}$ is the average number
 - // of free wavelengths on hops i to j of Segment 2
 - $p_{ij}^{(2)}(h) \leftarrow 0$, $F_{ij}^{(2)} \leftarrow W$, $1 \leq i \leq j \leq m$
 - // $q_{ij}(h)$ is the conditional probability that an inter-
 - // segment call will be blocked in Segment 2, given
 - // that it has found a free wavelength in Segment 1
 - $q_{ij}(h) \leftarrow 0$, $1 \leq i \leq K < j \leq K + m$
3. $h \leftarrow h + 1$ // h -th iteration
4. // Segment 1
 - $\lambda_{ij}^{(1)}(h) \leftarrow \lambda_{ij}$, $1 \leq i \leq j < K$
 - // Include the effective arrival rate of calls
 - // continuing to Segment 2
 - $\lambda_{iK}^{(1)}(h) \leftarrow \lambda_{iK}$
 - $+ \sum_{j=K+1}^{K+m} \lambda_{ij} (1 - q_{ij}(h-1))$, $1 \leq i \leq K$
 - Solve Segment 1 to obtain $p_{ij}^{(1)}(h)$ and $F_{ij}^{(1)}(h)$
5. // Segment 2
 - $\lambda_{ij}^{(2)}(h) \leftarrow \lambda_{K+i, K+j}$, $1 < i \leq j \leq m$
 - // Include the effective arrival rate of calls
 - // continuing from Segment 1
 - $\lambda_{1j}^{(2)}(h) \leftarrow \lambda_{K+1, K+j}$
 - $+ \sum_{i=1}^K \lambda_{i, K+j} (1 - p_{iK}^{(1)}(h-1))$, $1 \leq j \leq m$
 - Solve Segment 2 to obtain $p_{ij}^{(2)}(h)$ and $F_{ij}^{(2)}(h)$
6. // Conditional blocking prob. of inter-segment calls
 - $q_{ij}(h) \leftarrow p_{1, j-K}^{(2)}(h) + (1 - p_{1, j-K}^{(2)}(h))$
 - $\times \left(1 - \frac{F_{iK}^{(1)}(h)}{W}\right)^{F_{1, j-K}^{(2)}(h)}$, $1 \leq i \leq K < j \leq K + m$
7. Repeat from Step 3 until convergence
8. end of the algorithm

Figure 4: Decomposition algorithm for long paths




Figure 5: (a) A 4-hop path and (b) its decomposition into two 2-hop segments in tandem

A summary of our iterative algorithm is provided in Figure 4. We now describe the algorithm using the 4-hop path in Figure 5(a), decomposed into the two 2-hop in Figure 5(b). Let λ_{ij} , $1 \leq i \leq j \leq 4$, be the arrival rates of calls to the original 4-hop path, and let $\lambda_{11}^{(1)}$, $\lambda_{12}^{(1)}$, $\lambda_{22}^{(1)}$ and $\lambda_{11}^{(2)}$, $\lambda_{12}^{(2)}$, $\lambda_{22}^{(2)}$ denote the arrival rates of calls in the first and second segments, respectively. The interpretation of the arrival rates in the segments is as follows. $\lambda_{12}^{(1)}$ accounts for all the calls in the original 4-hop path that originate at node 0 and terminate at nodes 2 or higher; similarly for $\lambda_{22}^{(1)}$. Also, $\lambda_{12}^{(2)}$ accounts for all calls in the original path that originate at nodes 2 and lower and terminate at node 4; similarly for $\lambda_{11}^{(2)}$.

Initially, we solve the first segment in isolation using

$$\lambda_{12}^{(1)} = (1 - q_{14})\lambda_{14} + (1 - q_{13})\lambda_{13} + \lambda_{12} \quad (3)$$

$$\lambda_{22}^{(1)} = (1 - q_{24})\lambda_{24} + (1 - q_{23})\lambda_{23} + \lambda_{22} \quad (4)$$

$$\lambda_{11}^{(1)} = \lambda_{11} \quad (5)$$

Quantity q_{ij} , $1 \leq i \leq 2 < j \leq 4$, represents the current estimate of the conditional probability that a call using hops i through j (where i lies in the first segment and j lies in the second segment) will be blocked in the second segment, given that a free wavelength for the call exists within the first segment. Initially, we use $q_{ij} = 0$ for all i and j ; these values are updated in subsequent iterations as described shortly. Thus, the term $(1 - q_{14})\lambda_{14}$ in (3) represents the *effective* arrival rate of calls using all four hops, as seen by the first segment; similarly for the term $(1 - q_{13})\lambda_{13}$. Expression (4) for $\lambda_{22}^{(1)}$ includes similar terms that account for the effective arrival rate of calls which originate at node 1 and terminate at nodes 2 or higher. But expression (5) for $\lambda_{11}^{(1)}$ does not include any such terms, since this type of calls do not involve calls in the original path that continue to the second segment.

The solution to the first segment yields an initial value for the probability $p_{ij}^{(1)}$, $1 \leq i \leq j \leq 2$, that a call using hops i through j of the first segment will be blocked within the segment. Therefore, the effective arrival rate of calls originating at node 0 and

terminating at node 4 that is offered to the second segment can be initially estimated as $\lambda_{14}(1 - p_{12}^{(1)})$, while the effective rate of calls originating at node 1 and terminating at node 4 can be estimated as $\lambda_{24}(1 - p_{22}^{(1)})$. We can now solve the second segment using

$$\lambda_{12}^{(2)} = \lambda_{14}(1 - p_{12}^{(1)}) + \lambda_{24}(1 - p_{22}^{(1)}) + \lambda_{34} \quad (6)$$

$$\lambda_{11}^{(2)} = \lambda_{13}(1 - p_{12}^{(1)}) + \lambda_{23}(1 - p_{22}^{(1)}) + \lambda_{33} \quad (7)$$

$$\lambda_{22}^{(2)} = \lambda_{44} \quad (8)$$

$\lambda_{12}^{(2)}$ in (6) represents the effective arrival rate of calls using the last two hops of the 4-hop path, as seen by the second segment. Expression (7) for $\lambda_{11}^{(2)}$ can be explained using similar arguments. Expression (8) for $\lambda_{22}^{(2)}$ contains only one term since it does not involve calls that originate in segment 1. The solution to the second segment provides an estimate of the blocking probabilities $p_{ij}^{(2)}$, $1 \leq j \leq 2$, of calls traversing hops 1 and 2 of the second segment.

We can now obtain new values for the conditional blocking probabilities q_{ij} used in (3), (4). Consider a call using hops i through j , where i lies in the first segment and j in the second segment. Given that free wavelengths exist on hops i through 2 (i.e., the call makes it through the first segment), the call will be blocked if (a) there is no free wavelength in the links it uses in the second segment, or (b) there exist free wavelengths in the second segment, but they are not the same as the free wavelengths in the first two hops.

The probability of the first event occurring is equal to $p_{1,j-2}^{(2)}$, obtained through the solution of the second segment. The probability of the second event is equal to $(1 - p_{1,j-2}^{(2)})Q$, where Q represents blocking due to the wavelength continuity requirement. Probability Q cannot be computed exactly since each segment is solved independently of the other, and thus, it is not possible to determine whether a wavelength which is free in one segment will also be free in the other. An approximate value for Q can be obtained from the average number of free wavelengths in the two segments. Let $F_{ij}^{(1)}$ (resp., $F_{ij}^{(2)}$) be the current estimate of the average number of free wavelengths on hops i through j of segment 1 (resp., segment 2); this estimate is obtained from the solution to the segment. Consider a call using hops i through j , $1 \leq i \leq 2 < j \leq 4$, of the original path, which captures one of the free wavelengths in hops i through 2 of segment one. Because of the random wavelength assignment policy, the probability that any given wavelength is assigned to this call can be approximated by $F_{i2}^{(1)}/W$. The probability that this wavelength is not free in the next segment (given that there are free wavelengths in that segment) can be approximated by $(1 - F_{i2}^{(1)}/W)F_{1,j-2}^{(2)}$. Finally, the conditional blocking probability q_{ij} for this call, given that free wavelengths exist for the call in the first segment, is approximated by:

$$q_{ij} = p_{1,j-2}^{(2)} + \left(1 - p_{1,j-2}^{(2)}\right) \left(1 - F_{i2}^{(1)}/W\right) F_{1,j-2}^{(2)} \quad (9)$$

The new estimates for q_{ij} are used in expressions (3) to (5) to update the arrival rates for the first segment. The first segment is

solved again, the estimates for $p_{ij}^{(1)}$ are updated and used in expressions (6) to (8), and so on. We iterate in this fashion until the blocking probabilities for all calls in the original path converge within a certain tolerance. In our study, we have found that the algorithm converges in only a few iterations even for long paths, and that the blocking probabilities obtained closely match simulation results.

This decomposition algorithm is similar in spirit to the decomposition algorithms developed for tandem queueing networks with finite capacity queues (see [6]). The algorithm can be easily extended to handle paths decomposed into more than two segments. It is well known in decomposition algorithms that the larger the individual sub-systems that have to be analyzed in isolation, the better the accuracy of the decomposition algorithm. Thus, as we mentioned earlier, we have decided to decompose a path in segments of the largest size K for which we can efficiently analyze the Markov process \mathcal{M}'_K , plus, possibly, a segment of smaller size, if the path length is not a multiple of K .

4.2 Paths With Wavelength Conversion

The addition of $l < k$ converters leads to a natural decomposition of a k -hop path into $l + 1$ segments, each consisting of the links between successive nodes where converters are employed. Given such a decomposition, the blocking probability of calls spanning several segments depends only on the number of calls within each segment, *not* on the actual wavelengths used by those calls, and there is no wavelength continuity requirement between segments. Each segment is analyzed in isolation as described above. Specifically, if it is feasible, we analyze the segment's underlying approximate time-reversible Markov process \mathcal{M}'_k . Otherwise, we analyze it using the decomposition algorithm in Section 4.1.

As an example, consider a k -hop path with a converter located at node $K < k$. This path can be analyzed using the algorithm in Figure 4 after making a single modification: in Step 6, the expression for the conditional blocking probabilities is changed to $q_{ij} = p_{1,j-K}^{(2)}$. The second term in this expression represents blocking due to the wavelength continuity requirement, and since $Q = 0$ in this case, it drops out. The algorithm can be extended in a straightforward way to handle paths with more than one converters.

5 Numerical Results

5.1 Validation of the Time-Reversible Process

In Figure 6 we plot the blocking probability of calls for each source-destination pair in a 2-hop path without converters, against the number W of wavelengths per hop. For each type of call we show two curves. The first curve is obtained through a numerical solution of the exact Markov process \mathcal{M}_2 , and is referred to in the figure as "exact solution." The second curve is obtained from the closed-form solution of the approximate Markov process \mathcal{M}'_2 , and is referred to as "approximate solution". The overall behavior of the two curves is similar for all types of calls, and the approximate blocking probability is always very close to the exact value.

Figure 7 presents results for a 3-hop path. We only plot the blocking probability of calls for three of the six source-destination pairs,




Figure 6: Exact and approximate blocking probabilities of the various calls in a 2-hop path




Figure 7: Exact and approximate blocking probabilities of various calls in a 3-hop path

namely, calls that traverse all three hops, calls that use only the last two hops, and calls that use only the first hop. The blocking probability curves for the other three types of calls are very similar to the ones shown in Figure 7. Again, we observe that the values of the blocking probabilities obtained through the closed-form solution of the time-reversible Markov process \mathcal{M}'_3 are very close to the exact numerical values obtained from the Markov process \mathcal{M}_3 . However, the figure does not include values for the exact blocking probability when $W > 4$ because of the state space explosion of the exact Markov process.

5.2 Validation of the Decomposition Algorithm

We now present results for 6-hop and 10-hop paths with $W = 10$ wavelengths, with and without converters. We have used the following values for the traffic parameters:

$$\mu = 1, \forall i, j, \quad \lambda_{ij} = \begin{cases} 0.1 \text{ or } 0.3, & i < j \\ \lambda, & i = j \end{cases} \quad (10)$$

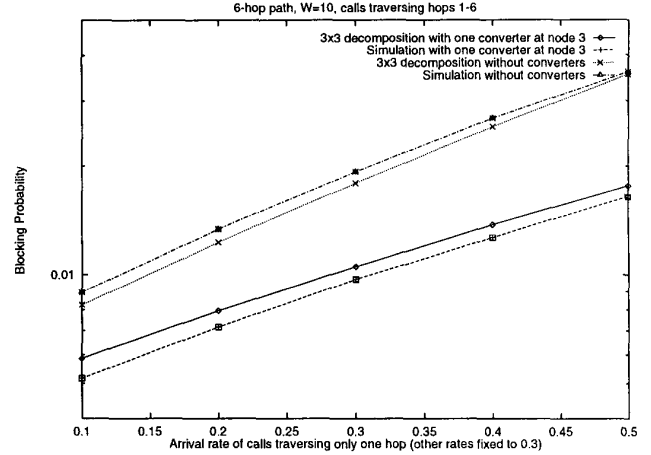


Figure 8: Blocking probability of calls traversing all links of a 6-hop path with $W = 10$

Figures 8 to 11 plot the call blocking probability as a function of λ , the arrival rate of calls traversing exactly one hop.

In Figures 8 and 9 we present results for a 6-hop path with and without converters, and for $W = 10$. We show results for only two source-destination pairs in the 6-hop path. The blocking probability of calls traversing all six hops in the path is plotted in Figure 8, and the blocking probability of calls traversing hops 1 through 4 of the path is shown in Figure 9. In both figures, the value of λ is varied from 0.1 to 0.5, while the arrival rate of all other calls is fixed to 0.3. Each figure contains two sets of plots, one for the 6-hop path without converters, and one for the same path with a single converter. Each set consists of two plots, one for the results from our decomposition algorithm, and the other for the simulation results. We analyzed the 6-hop path without converters by decomposing it into two 3-hop segments. We refer to this as a “ 3×3 decomposition”. When there is a single converter in the path, it is placed at node 3, also resulting in a 3×3 decomposition.

As the value of λ increases, the blocking probability of both types of calls increases. We also note that, when there is a converter at node 3, the blocking probability for both types of calls is significantly lower than when there is no converter. Both these results are expected. Furthermore, the values of the blocking probability obtained through our iterative decomposition algorithm are very close to the values obtained through simulation. Results similar to the ones shown in these figures were obtained for calls for all other source-destination pairs, and for a wide range of traffic loads.

Figures 10 and 11 show results for a 10-hop path with $W = 10$. Again, we only plot the blocking probability of calls for two source-destination pairs against the arrival rate λ , as the latter increases from 0.05 to 0.21, while all other arrival rates are fixed to 0.1. Figure 10 shows the blocking probability of calls traversing all ten hops of the path, while Figure 11 presents the blocking probability of calls using hops 2 through 6. As before, there are two sets of plots in each figure, one for the 10-hop path without converters, and one for the same path employing three con-




Figure 9: Blocking probability of calls traversing links 1 through 4 of a 6-hop path with $W = 10$

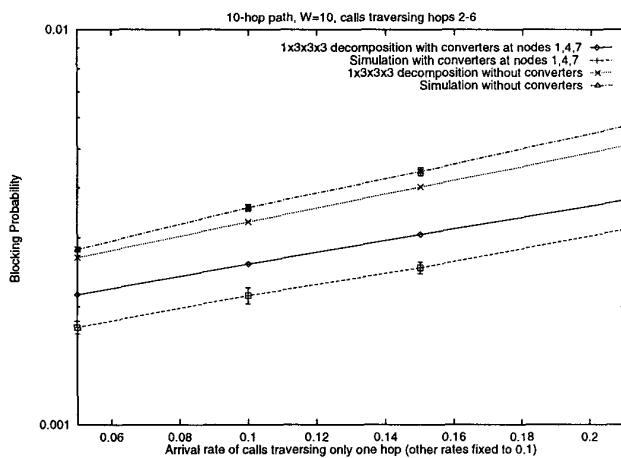


Figure 11: Blocking probability of calls traversing links 2 through 6 of a 10-hop path with $W = 10$




Figure 10: Blocking probability of calls traversing all links of a 10-hop path with $W = 10$

verters. For the no-converter case, in addition to simulation results, we present the blocking probability values obtained through a $1 \times 3 \times 3 \times 3$ decomposition into a 1-hop segment followed by three 3-hop segments. For the converter case, the three converters are assumed to be at nodes 1, 4, and 7, a configuration that also results in a $1 \times 3 \times 3 \times 3$ decomposition.

The general behavior of the curves in Figures 10 and 11 as λ increases is very similar to that of the curves in Figures 8 and 9. Regarding the accuracy of our decomposition algorithm, we note that the approximation results closely match the simulation results for both paths (with and without converters), and for both types of calls. Overall, we have found that the decomposition algorithm gives accurate results for long paths and for a wide range of traffic loads. We have also observed that the algorithm always converges in a few iterations, taking one minute for a 10-hop path, while the simulation takes several hours.

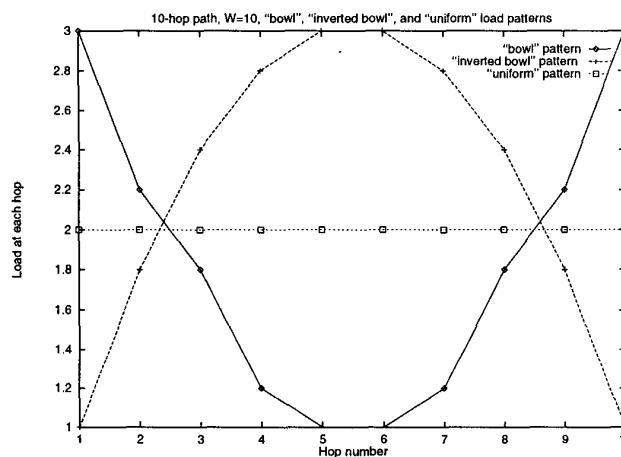


Figure 12: The “bowl”, “inverted bowl”, and “uniform” load patterns

5.3 Converter Placement

We now consider the problem of determining the best placement of l converters on a k -hop path, $k > l$, that minimizes the blocking probability of calls that travel over all k hops. To find the best converter placement we first enumerate all possible ways of placing l converters on a k -hop path; then, we calculate the blocking probability of interest for each alternative using the decomposition algorithm. The best placement is the one with the minimum such probability.

We consider a 10-hop path with $W = 10$, and three different traffic load patterns. Figure 12 plots the load of each hop in the path for each pattern. In the “uniform” pattern, all hops are equally loaded, while the “bowl” (resp., “inverted bowl”) pattern is such that the load decreases (resp., increases) from hop 1 to hop 5, and then it increases (resp., decreases) from hop 6 to hop 10. The load values were chosen so that the total network load is the same for all patterns.

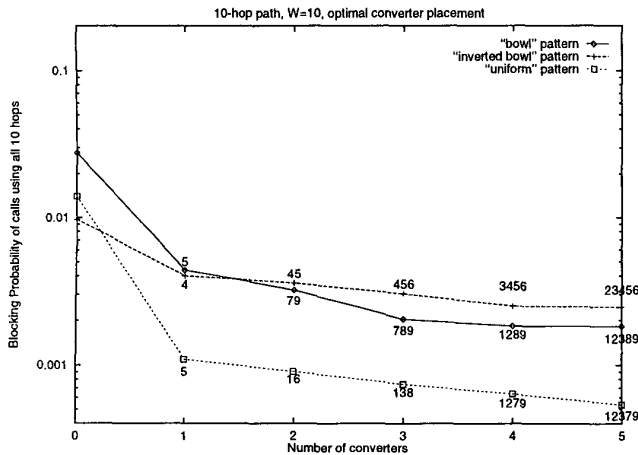


Figure 13: Blocking probability and optimal converter placement for the load patterns of Figure 12

In Figure 13 we plot the blocking probability of calls using all 10 hops of the path for the optimal placement of l converters, $1 \leq l \leq 5$. For comparison purposes, we also plot the blocking probability of these calls on a path without converters (the values for zero converters in these figures). The optimal location of the converters for each load pattern is also given next to each point of the curves.

As expected, the blocking probability drops as the number of converters increases. However, after an initial steep drop, the curves in general flatten as the number of converters increases. This behavior is consistent with the results of earlier work [7, 8]. We also observe that the effect of converters on the blocking probability is strongly dependent on the actual traffic pattern. Regarding the optimal node location of converters for the different traffic patterns, we first note that the results are intuitively obvious. The figure indicates, for example, that converters be placed at the middle of the path for the "inverted bowl" pattern. However, we observe that the optimal placement also depends strongly on the load pattern. This result suggests that in a dynamic environment where traffic patterns vary over time, there is no single assignment of converters to nodes that will work well for all possible loads. Consequently, simple optimization approaches, such as the one considered here, that seek to minimize the blocking probability under a specific traffic pattern may lead to poor performance if the pattern changes. Instead, more comprehensive approaches to the converter placement problem are needed, such as providing bounds for the blocking probability over a wide range of load patterns.

6 Concluding Remarks

We have presented a new framework to accurately and efficiently evaluate the call blocking performance in a single path of a wavelength routing network. We have derived exact and approximate Markov process models, and we have developed an iterative approximation algorithm to analyze long paths by decomposing them into shorter segments which are studied in isolation. We have also applied our techniques to the problem of converter placement in such a network.

References

- [1] R. A. Barry and P. A. Humblet. Models of blocking probability in all-optical networks with and without wavelength changers. *IEEE J. Selected Areas in Commun.*, pp. 858–867, June 1996.
- [2] A. Birman. Computing approximate blocking probabilities for a class of all-optical networks. *IEEE J. Selected Areas in Commun.*, pp. 852–857, June 1996.
- [3] F. P. Kelly. *Reversibility and Stochastic Networks*. John Wiley & Sons, New York, 1979.
- [4] L. Kleinrock. *Queueing Systems, Volume 1: Theory*. John Wiley & Sons, New York, 1975.
- [5] M. Kovacevic and A. Acampora. Benefits of wavelength translation in all-optical clear-channel networks. *IEEE J. Selected Areas in Commun.*, pp. 868–880, June 1996.
- [6] H. Perros. *Queueing Networks with Blocking: Exact and Approximate Solutions*. Oxford University Press, 1994.
- [7] S. Subramaniam, *et al.* All-optical networks with sparse wavelength conversion. *IEEE/ACM Transactions on Networking*, pp. 544–557, August 1996.
- [8] S. Subramaniam, *et al.* On the optimal placement of wavelength converters in wavelength-routed networks. *INFOCOM '98*, pp. 902–909, April 1998.
- [9] S. Subramaniam, *et al.* A performance model for wavelength conversion with non-Poisson traffic. *INFOCOM '97*, pp. 500–507, April 1997.
- [10] Y. Zhu, G. N. Rouskas, and H. G. Perros. Blocking in wavelength routing networks, Part I: The single path case. TR-98-02, NCSU, Raleigh, NC, February 1998.