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

Yuhong Zhu George N. Rouskas Harry G. Perros

TR-98-02

February 17, 1998

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 is fairly general and it allows non-uniform traffic, i.e., call request arrival rates and call holding times that can vary with the source-destination pair. It also 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 derive exact and approximate Markov process models, and we develop an iterative algorithm to analyze arbitrarily long paths by decomposing them into shorter segments connected in tandem which are studied in isolation. By considering a set of different traffic load patterns, we also demonstrate how our analytical techniques can be used to gain insight into the problem of converter placement in a wavelength routing network.

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

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

1 Introduction

A basic property of single mode optical fiber is its enormous low-loss bandwidth of several tens of terahertz (THz). Unfortunately, due to dispersive effects and limitations in optoelectronic component technology, single channel transmission is limited in speed to only a small fraction of the fiber capacity. To take full advantage of the potential of fiber, the use of wavelength division multiplexing (WDM) techniques has become the option of choice [4, 13], and WDM networks have been a subject of research both theoretically [24, 21, 3, 23] and experimentally [8, 9, 10, 11, 12]. Multiwavelength optical networks have the potential of delivering an aggregate throughput in the order of Terabits per second, and they appear as a viable approach to satisfying the ever-growing demand for more bandwidth per user [5, 22] on a sustained, long-term basis.

The wavelength routing mesh architecture appears promising for Wide Area Network (WAN) distances [21, 26, 7, 6]. 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 [24, 1], 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 [21, 7, 1]; in other words, once set up, a lightpath connection is expected to be long lived. This architecture is attractive for two reasons: first, the same wavelength can be used simultaneously at different parts of the network, and second, the signal power is channeled to the receiver and is not spread to the entire network as in the broadcast-and-select approach [17, 18, 20]. Hence, wavelength routing WANs can be highly scalable.

Given the installed base of optical fiber, and the maturing of optical component technology, it appears that current network technologies are transitory, and will eventually evolve to an all-optical, largely passive infrastructure. Although the long-term result of such an evolution is not yet clear, a feasible scenario for near-term large-scale all-optical networks has emerged in recent years [4, 13]. Wavelength routing WANs are an integral part of this scenario, since it is envisioned that they will act as the backbone that provides the interconnection for local area photonic sub-networks attached to them. The contribution of our work is the development of an analytical framework for evaluation of the performance of wavelength routing optical networks.

The call blocking performance of optical wavelength routing networks has been studied in

[24, 1, 16, 2, 25]. In [16] it was assumed that statistics of link loads are mutually independent, and thus, the model is not appropriate for sparse network topologies. 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. This approximation cannot capture the dynamic nature of traffic. The work in [2] used a Markov chain with state-dependent arrival rates to model call blocking in arbitrary mesh topologies. However, this approach is so computationally intensive that it can only be used for very small size networks in which paths have at most three links. Another model was presented in [24], where it was assumed that the load on link i of a path depends only on the load of link $i-1$. This work takes a probabilistic approach in modeling wavelength conversion by introducing the converter density, which represents the probability that a node is capable of wavelength conversion independently of other nodes in the network. While this approach works well when the objective is the estimation of the expected call blocking performance, it cannot be used to calculate the actual blocking probability on individual paths when the placement of converters is known, nor can it be used to compare various converter placement strategies. Finally, in [25], a study of call blocking under non-Poisson input traffic was presented, but it also assumes that link loads are statistically independent.

The above studies indicate that wavelength routing networks with arbitrary mesh topologies are rather hard to analyze. In this paper we present the first part of a new analytical framework which overcomes many of the problems mentioned above, and which permits us to accurately study the call blocking performance of wavelength routing networks. Specifically, we study a realistic model of a *single path* in a wavelength routing network, and we develop approximate analytical techniques, which are both accurate and efficient, for computing the call blocking probabilities for paths of arbitrary length.

Our work improves upon previously published work in several ways. First, we allow non-uniform traffic, in the sense that call request arrival rates and call holding times vary for each source-destination pair; in contrast, a uniform traffic distribution is assumed in most other studies. Second, our model accounts for the calls between any source-destination pair in the path (while recursive models such as the one in [24] only keep aggregate information about the effects of cross-traffic as the analysis proceeds from one step of the recursion to the next). As a result, our model can be used to compute the blocking probability not only for the end-to-end traffic, but also for the cross-traffic in a path. More importantly, our model captures the correlation among the loads on *all* links of a path (as opposed to the independence model in [16] and the Markovian correlation model in [24]). 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 exact Markov process models that capture the correlation of wavelength use among *all* links of a k-hop path with and without converters. We then show how to slightly modify these processes to obtain new Markov processes that are time-reversible and which have a closed-form solution that resembles the product form solution in queueing networks [15]. The solution to the new Markov processes provides an accurate approximation to the blocking probabilities obtained through the original ones. Because of their computational requirements, however, these Markov process models can only be directly applied to relatively short paths. For longer paths, we then develop iterative algorithms for computing the blocking probabilities by decomposing a path into a series of shorter segments connected in tandem. In particular, when wavelength converters are employed at some nodes of a path, we present a natural decomposition of a long path into segments, whereby each segment consists of the links between successive nodes with converters. Finally, we show how our analytical techniques can be used to gain insight into the problem of converter placement in a wavelength routing network.

Following this introduction, we present a model that captures the salient features of wavelength routing networks. In Section 3 we obtain exact and approximate Markov process models for computing the call blocking probabilities in a path with and without wavelength converters. In Section 4 we use the approximate Markov process models to develop iterative decomposition algorithms for long paths. In Section 5 we validate our decomposition algorithm through simulation, and we show how our techniques can be used to determine the best placement of converters in a path. We then conclude the paper in Section 6.

2 Wavelength Routing Network Model

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

In our model, we allow some of the nodes in the path to employ wavelength converters. These nodes can switch an incoming wavelength to an arbitrary outgoing wavelength. (When there are converters at all nodes, the situation is identical to that in classical circuit-switching networks, a special case of the more general scenario discussed here.) If no wavelength converters are employed

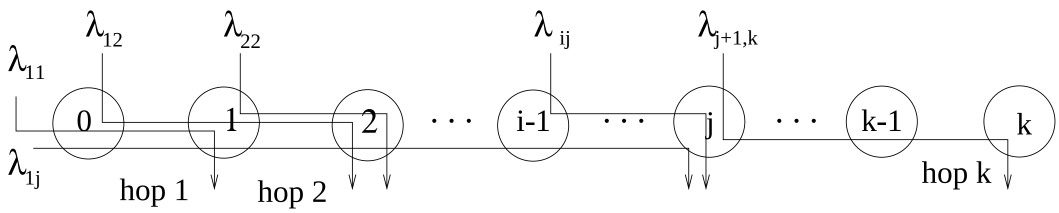

Figure 1: A k-hop path

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 available wavelengths, the call is blocked. On the other hand, if a call can be accommodated, it is randomly assigned one of the wavelengths that are available on the links used by the call [1]. Thus, we only consider the random wavelength assignment policy in this paper.

We define a path "segment" as a sub-path consisting of one or more consecutive links of the original path. (Unless noted otherwise, the terms "hop" and "link" will be used interchangeably.) The following notation will be used throughout the paper (refer to Figure 1).

- A k-hop path consists of $k+1$ nodes labeled $0, 1, \cdots, k$. Hence, hop $i, i = 1, \cdots, k$, represents 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, inclusive. For instance, λ_{22} is the arrival rate of calls that only use hop 2 (that is, those arriving at node 1 and leaving at node 2), while λ_{12} is the arrival rate of calls using hops 1 and 2 (i.e., those that arrive at node 0 and leave at node 2).

- $1/\mu_{ij}, j \geq i$, is the mean of the (exponentially distributed) holding time of calls using hops i through j, inclusive. Also, $\rho_{ij} = \lambda_{ij}/\mu_{ij}$ is the offered load of these calls.

- n_{ij}, $j \geq i$, is the number of calls using hops i through j, inclusive, 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, inclusive.

[1] In a path with wavelength converters, a wavelength is randomly assigned within each segment of the path consisting of links between successive nodes with converters.

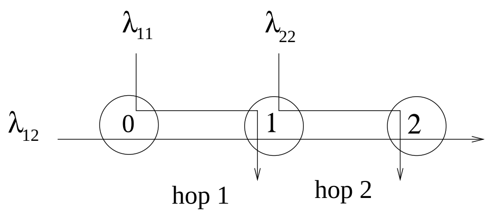

Figure 2: A two-hop path

3 Exact and Approximate Markov Process Models

In this section we present exact and approximate Markov process models that capture the dynamic behavior of a k-hop path in terms of wavelength use. The approximate Markov process has a closed-form solution for its steady-state probability, and it provides an accurate approximation to the blocking probabilities obtained through the exact Markov process. We first study paths without wavelength converters (in Section 3.1), and then we extend our results to paths with converters employed at some nodes (in Section 3.2).

3.1 Paths With No Wavelength Conversion

3.1.1 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 $(n_{11}, n_{12}, n_{22}, f_{12})$. In addition, on each hop, the number of busy wavelengths plus the number of wavelengths that are free on both hops may not exceed the number W of available wavelengths. This condition leads to the following two constraints that the parameters of the Markov process must satisfy:

$$n_{11} + n_{12} + f_{12} \leq W \quad \text{and} \quad n_{12} + n_{22} + f_{12} \leq W \tag{1}$$

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

$$\underline{n} = (n_{11}, n_{12}, \cdots, n_{1k}, n_{22}, \cdots, n_{2k}, \cdots, n_{kk}, f_{12}, f_{13}, \cdots, f_{1k}, f_{23}, \cdots, f_{2k}, \cdots, f_{k-1,k}) \tag{2}$$

The first $\frac{k(k+1)}{2}$ parameters $n_{ij}, 1 \leq i \leq j \leq k$, in the state description (2) provide the number of active calls between all possible source-destination pairs in the path. The last $\frac{(k-1)k}{2}$ parameters $f_{ij}, 1 \leq i < j \leq k$, represent the number of wavelengths that are free on all segments of the path consisting of two or more links. The following constraints are imposed on the state space of Markov process $\mathcal{M}_k$:

$$f_{ij} \leq f_{i,j-1} \leq \cdots \leq f_{i,i+1}, \quad 1 \leq i < j \leq k \tag{3}$$

$$f_{ij} \leq f_{i+1,j} \leq \cdots \leq f_{j-1,j}, \quad 1 \leq i < j \leq k \tag{4}$$

$$\begin{cases} \sum_{j=1}^{k} n_{1j} + f_{12} \leq W, & l = 1 \\ \sum_{i=1}^{l} \sum_{j=l}^{k} n_{ij} + f_{l-1,l} + f_{l,l+1} - f_{l-1,l+1} \leq W, & l = 2, \cdots, k-1 \\ \sum_{i=1}^{k} n_{ik} + f_{k-1,k} \leq W, & l = k \end{cases} \tag{5}$$

The set of constraints (3) (respectively, (4)) account for the fact that, if a wavelength is free on all hops of a m-hop segment, then it is also free on the first (respectively, last) $m-1$ hops of the segment. The k constraints (5) ensure that the number of wavelengths (used or free) on each hop of the path does not exceed W. For instance, let us refer to the first constraint in (5); the others can be explained using similar arguments. The sum in the left hand side of the constraint corresponds to the total number of active calls that use hop 1, while f_{12} is the total number of wavelengths that are free on both hops 1 and 2. By definition (see Section 2 and constraint (3)), parameter f_{12} includes the free wavelengths on all path segments of two or mode hops starting with hop 1. Therefore, the terms in left hand side of this constraint must sum up to at most W (if they sum up to *less than* W, the difference represents the number of wavelengths that are free on hop 1 but are not free on hop 2, and hence, can only be assigned to calls using only hop 1).

We have the following lemma.

Lemma 3.1 *The stochastic process defined in (2) is a Markov process.*

Proof. See Appendix A. □

Markov process $\mathcal{M}_k$ captures the correlation of wavelength use on all links of a k-hop path, and it can be used to provide an exact solution for the probability that a call request will be blocked. Unfortunately, the large number of parameters in its state description makes it impossible to numerically solve the MC for anything but very short paths and small values of W. In addition, the transition rates of Markov process $\mathcal{M}_k$ are state-dependent. In Figure 3 we show the state space $(n_{11}, n_{12}, n_{22}, f_{12})$ and the transition rates of Markov process $\mathcal{M}_2$, for $W = 2$ wavelengths. From the figure we can see that there exists a sequence of states, $\underline{n_1}, \cdots, \underline{n_s}$, such that

$$r(\underline{n_1}, \underline{n_2}) r(\underline{n_2}, \underline{n_3}) \cdots r(\underline{n_{s-1}}, \underline{n_s}) r(\underline{n_s}, \underline{n_1}) \neq r(\underline{n_1}, \underline{n_s}) r(\underline{n_s}, \underline{n_{s-1}}) \cdots r(\underline{n_3}, \underline{n_2}) r(\underline{n_2}, \underline{n_1}) \tag{6}$$

where $r(\underline{n}, \underline{n'})$ is the transition rate from state $\underline{n}$ to state $\underline{n'}$. Two such sequences of states 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). Therefore, Kolmogorov's criterion for reversibility [14, Theorem 1.8] is not satisfied, and the Markov process is not

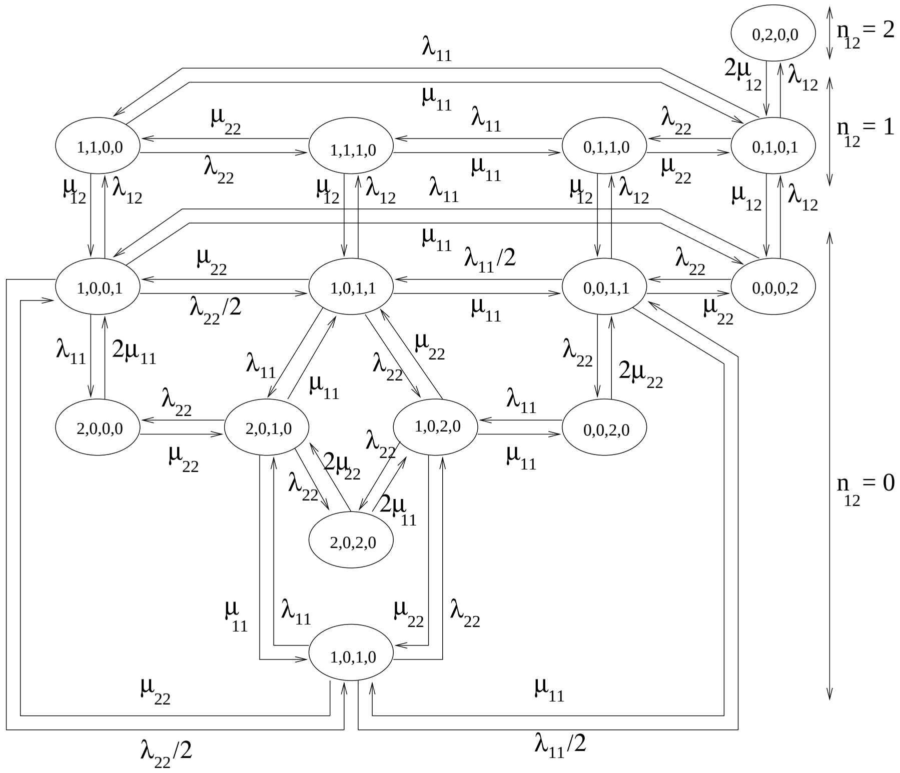

Figure 3: State space $(n_{11}, n_{12}, n_{22}, f_{12})$ of a 2-hop path with $W = 2$ wavelengths

time-reversible. It is straightforward to show that this result is true in general, and that Markov process $\mathcal{M}_k, k \geq 2$ is not time-reversible when $W > 1$.

In the following subsection we show how to modify $\mathcal{M}_k$ to obtain an approximate, time-reversible Markov process that has a closed-form solution.

3.1.2 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 of states for which (6) holds true; any other sequence that satisfies (6) contains one of these two sequences. We also note that both four-state sequences involve transitions that cause changes in the value of parameter n_{12}. Let us define $\mathcal{L}_{2,c}$ as the sub-chain of Markov process $\mathcal{M}_2$ that includes only the states for which the value of parameter $n_{12} = c$:

$$\mathcal{L}_{2,c} \quad = \quad \{(n_{11}, n_{12}, n_{22}, f_{12}) \ : \ n_{12} \ = \ c\} \ , \qquad c = 0, \cdots, W \tag{7}$$

Sub-chain $\mathcal{L}_{2,c}$ corresponds to a new system with $W - c$ wavelengths per hop, in which no calls using both hops ever arrive (that is, $\lambda_{12} = 0$ in this new system). Then, it can be easily verified

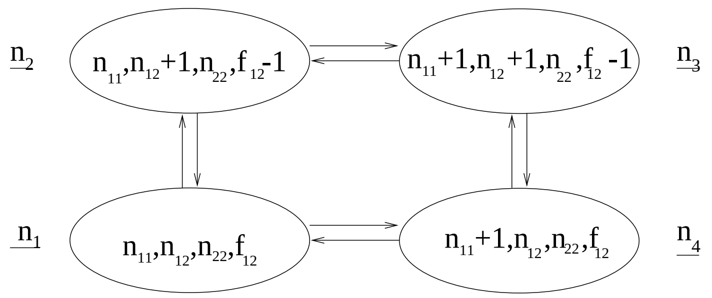

Figure 4: Four-state sequence with states from two different sub-chains $\mathcal{L}_{2,c}$

that Kolmogorov's criterion for reversibility is satisfied by any sequence of states that includes only states from a certain sub-chain $\mathcal{L}_{2,c}$.

On the other hand, let us consider the four-state sequence $\underline{n_1} = (n_{11}, n_{12}, n_{22}, f_{12})$, $\underline{n_2} = (n_{11}, n_{12}+1, n_{22}, f_{12}-1)$, $\underline{n_3} = (n_{11}+1, n_{12}+1, n_{22}, f_{12}-1)$, and $\underline{n_4} = (n_{11}+1, n_{12}, n_{22}, f_{12})$, shown in Figure 4, which includes states from two different sub-chains [2]. We have that:

$$\begin{array}{ll} r(\underline{n_1}, \underline{n_2}) = \lambda_{12}, & r(\underline{n_3}, \underline{n_4}) = (n_{12}+1)\mu_{12} \\ r(\underline{n_2}, \underline{n_1}) = (n_{12}+1)\mu_{12}, & r(\underline{n_4}, \underline{n_3}) = \lambda_{12} \end{array} \tag{8}$$

so these transition rates balance along the two directions. However, the rates of the other transitions do not balance, since [3]:

$$\begin{array}{ll} r(\underline{n_2}, \underline{n_3}) = \lambda_{11}\left(1 - \frac{f_{12}-1}{W-(n_{12}+1)-n_{11}}\right), & r(\underline{n_4}, \underline{n_1}) = (n_{11}+1)\mu_{11}\left(1 - \frac{W-n_{12}-n_{22}}{n_{11}+1}\right) \\ r(\underline{n_3}, \underline{n_2}) = (n_{11}+1)\mu_{11}\left(1 - \frac{W-(n_{12}+1)-n_{22}}{n_{11}+1}\right), & r(\underline{n_1}, \underline{n_4}) = \lambda_{11}\left(1 - \frac{f_{12}}{W-n_{11}-n_{12}}\right) \end{array} \tag{9}$$

It is due to transitions between states with different values of parameter n_{12} that Markov process $\mathcal{M}_2$ is not time-reversible.

Returning to Figure 3, we note that if the transition rate from state (1,0,1,1) to state (1,1,1,0) is changed to $2\lambda_{11}$ (from λ_{11}), then the Markov process becomes time-reversible. This is an important result because, as we will see in a moment, 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 Markov process. The rule for changing the transition rates is as follows.

[2] Similar arguments apply when $\underline{n_3}$ and $\underline{n_4}$ are obtained from $\underline{n_2}$ and $\underline{n_1}$, respectively, by incrementing the value of parameter n_{22} instead of that of parameter n_{11}.

[3] To obtain $r(\underline{n_2}, \underline{n_3})$ we multiply the arrival rate λ_{11} by the probability that the wavelength assigned to the new call is one of the wavelengths used on the second hop, or equivalently, the probability that the new state will have the same value for parameter f_{12}. To obtain $r(\underline{n_4}, \underline{n_1})$, we multiply the rate $(n_{11}+1)\mu$ by the probability that the terminating call will free a wavelength that is used on the second hop, and thus the new state has the same value for f_{12}. Similarly for the other rates.

Consider all states $\underline{n} = (n_{11}, c, n_{22}, f_{12})$ of sub-chain $\mathcal{L}_{2,c}$ with $n_{11} > 0$ and $n_{22} > 0$, for which there is a transition with rate $r(\underline{n}, \underline{n'}) = \lambda_{12}$ to a state $\underline{n'} = (n_{11}, c+1, n_{22}, f_{12}-1)$ of sub-chain $\mathcal{L}_{2,c+1}$. If these transition rates are changed to:

$$r'(\underline{n}, \underline{n'}) \;=\; \lambda_{12}\,\frac{f_{12}(W - n_{12})}{f_{11} f_{22}} \;=\; \lambda_{12}\,\frac{f_{12}(W - n_{12})}{(W - n_{11} - n_{12})(W - n_{22} - n_{12})} \qquad (10)$$

then we obtain a new Markov process which is time-reversible.

In fact, it is straightforward to show that if the transition rate $r(\underline{n_1}, \underline{n_2})$ in (8) is modified according to this rule, then the four-state sequence in Figure 4 will satisfy Kolmogorov's criterion for reversibility.

The above observations can be generalized to a k-hop path, $k \geq 2$. Consider a sub-chain $\mathcal{L}_{k,\underline{c}}$ of $\mathcal{M}_k$ which includes all states of the Markov process for which the number of active calls using two or more hops is constant:

$$\mathcal{L}_{k,\underline{c}} \;=\; \{\underline{n} \in \mathcal{M}_k \;:\; n_{ij} = c_{ij} = \text{constant},\; i < j\} \;,\;\; \underline{c} = (c_{12}, \cdots, c_{1k}, c_{23}, \cdots, c_{2k}, \cdots, c_{k-1,k}) \quad (11)$$

Sub-chain $\mathcal{L}_{k,\underline{c}}$ models a k-hop path in which there are no arrivals of calls using two or more hops (i.e., $\lambda_{ij} = 0$ for $i < j$), and in which each hop supports a fixed number of wavelengths (that can be different from the number of wavelengths supported by other hops). Then, it is straightforward to verify that any sequence of states that includes only states from sub-chain $\mathcal{L}_{k,\underline{c}}$ satisfies Kolmogorov's criterion for reversibility, but there exist sequences that include states from different sub-chains that violate this criterion.

We now introduce a new time-reversible Markov process $\mathcal{M}'_k$ which is derived from $\mathcal{M}_k$ as follows. The new Markov process $\mathcal{M}'_k$ has the same state space and the same transitions as $\mathcal{M}_k$. The vast majority of the transition rates of the new Markov process are the same as the respective 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. Consider the states $\underline{n}$ of sub-chain $\mathcal{L}_{k,\underline{c}}$ for which there exists a transition with rate $\lambda_{lm}, l < m$, to a state $\underline{n'}$ of another sub-chain $\mathcal{L}_{k,\underline{c'}}, c \neq c'$:

$$\mathcal{N}_{k,\underline{c}} \;=\; \{\underline{n} \in \mathcal{L}_{k,\underline{c}} \;:\; \exists\, i,j,l,m,\;\; l \leq i < j \leq m,\;\; n_{ii} > 0, n_{jj} > 0, n_{lm} < W, f_{lm} > 0\} \qquad (12)$$

The transition rate $r(\underline{n}, \underline{n'}) = \lambda_{lm}$ in Markov process $\mathcal{M}_k$. In the new Markov process $\mathcal{M}'_k$ the transition rate is changed to:

$$r'(\underline{n}, \underline{n'}) \;=\; \lambda_{lm}\frac{f_{lm}\left(\sum_{i=1}^{l} n_{il} + f_{ll}\right)\left(\sum_{i=1}^{l+1} n_{i,l+1} + f_{l+1,l+1}\right)\cdots\left(\sum_{i=1}^{m-1} n_{i,m-1} + f_{m-1,m-1}\right)}{f_{ll}\; f_{l+1,l+1}\;\cdots\; f_{m,m}} \quad (13)$$

Markov process $\mathcal{M}'_k$ has a closed-form solution for its steady-state probability that resembles the product form solution in queueing networks [15]. Let $G_k(W)$ denote the normalizing constant for a k-hop path with W wavelengths per link. Then, the solution of Markov process $\mathcal{M}'_2$ corresponding to the 2-hop path in Figure 2 is [4]:

$$\pi(n_{11}, n_{12}, n_{22}, f_{12}) = \frac{1}{G_2(W)} \frac{\rho_{11}^{n_{11}} \rho_{12}^{n_{12}} \rho_{22}^{n_{22}}}{n_{11}! n_{12}! n_{22}!} \frac{\binom{W - n_{11} - n_{12}}{f_{12}} \binom{n_{11}}{W - n_{12} - n_{22} - f_{12}}}{\binom{W - n_{12}}{W - n_{12} - n_{22}}} \tag{14}$$

while the solution to Markov process $\mathcal{M}'_3$ corresponding to a 3-hop path with state description $\underline{n} = (n_{11}, n_{12}, n_{13}, n_{22}, n_{23}, n_{33} f_{12}, f_{13}, f_{23})$ is given by:

$$\begin{aligned} \pi(\underline{n}) &= \frac{1}{G_3(W)} \frac{\rho_{11}^{n_{11}} \rho_{12}^{n_{12}} \rho_{13}^{n_{13}} \rho_{22}^{n_{22}} \rho_{23}^{n_{23}} \rho_{33}^{n_{33}}}{n_{11}! n_{12}! n_{13}! n_{22}! n_{23}! n_{33}!} \frac{\binom{W - n_{11} - n_{12} - n_{13}}{f_{12}} \binom{n_{11}}{W - n_{12} - n_{22} - n_{13} - f_{12}}}{\binom{W - n_{12}}{W - n_{12} - n_{22}}} \\ &\times \frac{\binom{W - n_{11} - n_{12}}{f_{12}} \binom{n_{11}}{W - n_{12} - n_{22} - f_{12}}}{\binom{W - n_{12}}{W - n_{12} - n_{22}}} \frac{\binom{W - n_{11} - n_{12}}{f_{12}} \binom{n_{11}}{W - n_{12} - n_{22} - f_{12}}}{\binom{W - n_{12}}{W - n_{12} - n_{22}}} \end{aligned} \tag{15}$$

We can write down the solution to any Markov process $\mathcal{M}'_k$, $k > 3$, by a straightforward generalization of expressions (14) and (15). Specifically, the solution to $\mathcal{M}'_k$, for any k, is the product of k^2 terms as follows. The first $k(k+1)/2$ terms are of the form $\rho_{ij}^{n_{ij}}/n_{ij}$, and each corresponds to one of parameters n_{ij} in the state description (2). The last $k(k-1)/2$ terms are combinatorial terms, each corresponding to one of the dependent variables f_{ij} in the state description (2).

The significance of the new Markov process $\mathcal{M}'_k$ will be illustrated in Section 5 and Figures 7 and 8, where it will be shown that the blocking probabilities obtained through the closed-form solution of $\mathcal{M}'_k$ closely approximate the exact blocking probabilities obtained through the numerical solution of $\mathcal{M}_k$. Hence, our main contribution in this area is the development of a time reversible Markov process, with a closed-form solution for its steady-state probability that can be easily written down, which accurately captures the correlation of wavelength use among the links of a path.

[4] The closed-form solution (14) is similar to the one presented in [24]. However, there are several important differences in the two approaches. First, the solution in [24] was derived by considering a three-dimensional Markov process (n_{11}, n_{12}, n_{22}), while as we have shown, the fourth parameter f_{12} is necessary to completely characterize a 2-hop path. Second, we have shown that the closed-form solution (14) is the exact solution to an approximate Markov process; in contrast, in [24, Section II.C] this result was derived as a solution to the original 2-hop problem, without explicitly stating that it is an approximation. Finally, whereas only a 2-hop path was studied in [24], our approach is far more general, and it leads to an approximate closed-form solution for any k-hop path, $k \geq 2$ (see also (15)).

Of course, the main concern in any product-form solution is the computation of the normalizing constant. Using brute force enumeration, we can obtain the product form solution of $\mathcal{M}'_k$ for up to 25 wavelengths, when the number of hops is $k = 3$, and for up to 8 wavelengths when $k = 4$. (As a comparison, for $k = 3$ we can obtain the numerical solution of $\mathcal{M}_k$ only up to $W = 4$.) Despite our efforts, we have not been able to devise a recursive procedure that would allow us to compute the normalizing constant in time that is polynomial in W and k. Therefore, a different approximate method is needed for paths longer than four hops. In Section 4 we describe an iterative decomposition algorithm that can be used to efficiently obtain the blocking probabilities for paths of arbitrary length.

3.2 Paths With Wavelength Conversion

Let us now turn our attention to paths in which wavelength converters are employed at some nodes. We show how the results of Section 3.1 can be extended to this case, leading to a method which can be used to obtain the call blocking probabilities in a k-hop path when the placement of converters is known.

Let us again refer to Figure 2, and let us assume that wavelength converters are located at node 1 (the only interesting possibility in this case). We immediately see that the 2-hop system can now be described by the three-dimensional Markov process (n_{11}, n_{12}, n_{22}). Parameter f_{12} is no longer necessary because wavelength continuity is not required, and calls continuing on both hops can now use *any* of the $(W - n_{12} - n_{22})$ available wavelengths on the second hop. In other words, the 2-hop system with a converter at node 1 becomes equivalent to a 2-hop circuit-switched path.

In the general case, consider a k-hop path, $k \geq 2$, with converters employed at one or more nodes. This path can be modeled by a new Markov process which is simpler than $\mathcal{M}_k$, as follows. The new Markov process has the same $\frac{k(k+1)}{2}$ parameters n_{ij} as $\mathcal{M}_k$, but some of the parameters f_{ij} are no longer present in the state description. More specifically, let us consider the case when a converter is employed at node $l, 0 < l < k$, of the path. Then, parameters $f_{ij}, i \leq l < j$, which are required for $\mathcal{M}_k$, are not part of the state description of the new Markov process. Because of the converter at node l, a call using hops i through j, inclusive, can now be completed as long as there is at least one free wavelength on hops i through l, and at least one free wavelength on hops $l + 1$ through j. Therefore, parameters f_{il} and $f_{l+1,j}$ (which remain part of the state description) provide all the information needed to determine whether the call can be completed, making f_{ij} redundant.

It is now straightforward to show that the Markov process for a k-hop path that employs wavelength converters is not time-reversible, except for $k = 2$. We can then use an approach similar to the one we followed in Section 3.1 to modify some of the transition rates of this process in order to obtain an approximate, time-reversible Markov process which has a closed-form solution. Thus, we have developed exact and approximate Markov process models for computing the call blocking probabilities in paths with wavelength conversion capabilities.

4 Decomposition Algorithm for Long Paths

Let K denote the largest integer such that the normalizing constant of the closed-form solution to Markov process $\mathcal{M}'_k$ can be computed within a reasonably short amount of time (currently, $K = 3$ or 4, depending on the number W of wavelengths per link, but this value may increase if we can improve our method for computing the normalizing constant). 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, on the other hand, $k > K$, 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. As before, 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 l K-hop segments and one m-hop segment in tandem. Each segment is first analyzed in isolation using the corresponding Markov process $\mathcal{M}'_k$. The arrival rates of calls originating (terminating) in a segment but terminating (originating) in another segment are accounted for by increasing the arrival rate of calls of the individual segments. The individual solutions can then be 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 to obtain a new solution. These new individual solutions are again combined to update the blocking probability of calls traversing multiple segments. The process is repeated until the blocking probabilities converge.

A detailed description of our iterative algorithm is provided in Figure 5. We now illustrate the operation of the algorithm by referring to Figure 6 which shows a 4-hop path and its decomposition into two 2-hop segments. We let $\lambda_{11}, \lambda_{12}, \cdots, \lambda_{44}$, be the arrival rates of calls to the original 4-hop path, and we use $\lambda^{(1)}_{11}, \lambda^{(1)}_{12}, \lambda^{(1)}_{22}$ and $\lambda^{(2)}_{11}, \lambda^{(2)}_{12}, \lambda^{(2)}_{22}$ to denote the arrival rates of calls to the first

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) connected in tandem, such that node K of the original path is the last node of segment 1 and the first node of segment 2. λ_{ij} refer to the call arrival rates in the original path, while $\lambda_{ij}^{(1)}$ (respectively, $\lambda_{ij}^{(2)}$) refer to call arrival rates in the first (resp., second) segment.

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

Figure 5: Decomposition algorithm for long paths

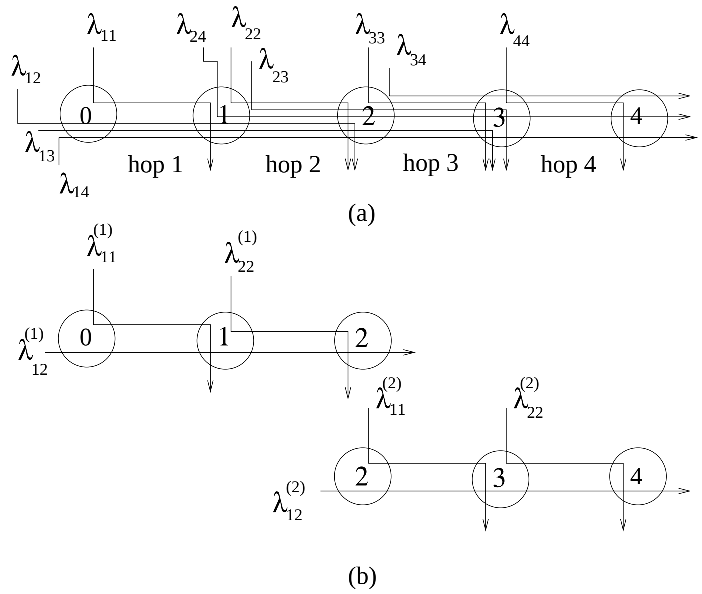

Figure 6: (a) A 4-hop path and (b) its decomposition into two 2-hop segments in tandem

and second segments, respectively. However, the interpretation of the arrival rates for the 2-hop segments is somewhat different under our decomposition approach. Specifically, $\lambda_{12}^{(1)}$ accounts for all the calls of the original 4-hop path that originate at node 0 and terminate at node 2 or higher; similarly for $\lambda_{22}^{(1)}$. On the other hand, $\lambda_{12}^{(2)}$ accounts for all calls of the original path that originate at nodes 2 and lower and terminate at node 4; similarly for $\lambda_{11}^{(2)}$. Below, we describe briefly the main steps of our algorithm.

Initially, we solve the first segment in isolation by using

$$\lambda_{12}^{(1)} = (1-q_{14})\lambda_{14} + (1-q_{13})\lambda_{13} + \lambda_{12} \tag{16}$$

$$\lambda_{22}^{(1)} = (1-q_{24})\lambda_{24} + (1-q_{23})\lambda_{23} + \lambda_{22} \tag{17}$$

$$\lambda_{11}^{(1)} = \lambda_{11} \tag{18}$$

Quantity q_{ij}, $1 \le i \le 2 < j \le 4$, represents the current estimate of the conditional probability that a call using hops i through j (where i lies within the first segment and j lies within the second segment) will be blocked given that a free wavelength for the call exists within the first segment. For the first iteration, we use $q_{ij} = 0$ for all i and j; how these values are updated in subsequent iterations will be described shortly. Thus, the term $(1-q_{14})\lambda_{14}$ in (16) represents the *effective* rate of calls using all four hops, as seen by the first segment; similarly for the term $(1-q_{13})\lambda_{13}$. Expression (17) for $\lambda_{22}^{(1)}$ includes similar terms that account for the effective rate of calls which originate at node 1 and terminate at nodes 2 or higher. On the other hand, expression (18) for $\lambda_{11}^{(1)}$

does not include any such terms, since this type of calls in the first segment do not involve calls of the original path continuing to the second segment.

The solution to the first segment will yield an initial value for the probability $p_{ij}^{(1)}$, $1 \le i \le j \le 2$, that a call using hops i through j of the first segment will be blocked within the segment. Therefore, the effective 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 by using

$$\lambda_{12}^{(2)} = \lambda_{14}(1-p_{12}^{(1)}) + \lambda_{24}(1-p_{22}^{1)}) + \lambda_{34} \tag{19}$$

$$\lambda_{11}^{(2)} = \lambda_{13}(1-p_{12}^{(1)}) + \lambda_{23}(1-p_{22}^{1)}) + \lambda_{33} \tag{20}$$

$$\lambda_{22}^{(2)} = \lambda_{44} \tag{21}$$

Based on the above discussion, $\lambda_{12}^{(2)}$ in (19) represents the effective rate of calls using the last two hops of the 4-hop path, as seen by the second segment. Expression (20) for $\lambda_{11}^{(2)}$ can be explained using similar arguments, while expression (21) for $\lambda_{22}^{(2)}$ contains only one term since, as seen in Figure 6, it does not involve calls continuing from segment 1. The solution to the second segment will provide an estimate of the blocking probabilities $p_{ij}^{(2)}$, $1 \le j \le 2$, of calls traversing hops 1 and 2 of the second segment (i.e., hops 3 and 4 of the original path).

We can now obtain new values for the conditional blocking probabilities q_{ij}, $1 \le i \le 2 < j \le 4$, used in (16) and (17), as follows. Consider a call using hops i through j, where i lies in the first segment and j lies in the second segment. Given that at least one free wavelength exists on hops i through 2 (i.e., the call successfully makes it through the first segment), the call will be blocked if

1. there is no free wavelength on the links it uses in the second segment, or

2. there do exist free wavelengths on the second segment, but they are not the same as the free wavelengths on the first two hops.

The probability of the first event occurring is equal to $p_{1,j-2}^{(2)}$, which is 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 parameter Q represents blocking due to the wavelength continuity requirement. Unfortunately, the 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. However, we can obtain an approximate value for this probability from the average number of free wavelengths in the two segments, as follows. Let $F_{ij}^{(1)}$ (respectively, $F_{ij}^{(2)}$)

be the current estimate of the average number of free wavelengths on hops i through j of segment 1 (respectively, segment 2); this estimate is obtained as part of the solution to the segment. Consider a call using hops i through j, $1 \le i \le 2 < j \le 4$, 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 $\frac{F^{(1)}_{i2}}{W}$. The probability that this wavelength is not free in the next segment (given that there are free wavelengths in that segment) can be approximated as $\left(1 - \frac{F^{(1)}_{i2}}{W}\right)^{F^{(2)}_{1,j-2}}$. Finally, the conditional blocking probability q_{ij} for this call, given that free wavelengths exist for the call within the first segment, can be approximated as:

$$q_{ij} = p^{(2)}_{1,j-2} + \left(1 - p^{(2)}_{1,j-2}\right) Q = p^{(2)}_{1,j-2} + \left(1 - p^{(2)}_{1,j-2}\right) \left(1 - \frac{F^{(1)}_{i2}}{W}\right)^{F^{(2)}_{1,j-2}} \tag{22}$$

The new estimates for q_{ij} are then used in expressions (16) – (18) to update the arrival rates for the first segment, the first segment is solved again, the estimates for $p^{(1)}_{ij}$ are updated and used in expressions (19) – (21), and so on. We repeat the process until the blocking probabilities for all calls of the original path converge within a certain tolerance. In all cases studied, we have found that the algorithm converges in only a few (under ten) iterations even for long paths, and that the blocking probabilities obtained closely match simulation results (more on this in Section 5).

The decomposition algorithm just described is similar in spirit to the decomposition algorithms developed for tandem queueing networks with finite capacity queues (see [19]). This algorithm can be easily extended to handle paths decomposed into more than two segments. We note that when employing the decomposition algorithm, the eventual selection of the segment size will depend on the following factors: length of the original path, how efficiently we can calculate the normalizing constant for the Markov process $\mathcal{M}'_k$ associated with each segment, and how accurate the approximate solution of $\mathcal{M}'_k$ is. 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 in the beginning of this subsection, we have decided to decompose a path in segments of the largest size K for which we can efficiently compute the normalizing constant, plus, possibly, a segment of smaller size, if the path length is not a multiple of K.

4.2 Paths With Wavelength Conversion

The iterative algorithm described above can also be used for paths with converters. We note, however, that the addition of $l < k$ converters leads to a natural decomposition of a k-hop path into $l+1$ segments, with each segment consisting of the links between successive nodes where converters are employed. Given such a decomposition, the blocking probability of calls spanning several segments now depends only on the number of calls within each segment (similar to the circuit switching case), *not* the actual wavelengths used by those calls. Hence, the probability that a call spanning multiple segments will be successfully completed becomes equal to the product of the probabilities of finding a free wavelength (not necessarily the same one) within each segment. Therefore, we choose to decompose a path into segments between successive nodes with converters. These segments do not contain converters, and they can be analyzed in isolation as described previously. If one of these segments is too long, it can be analyzed using the iterative decomposition algorithm in Figure 5.

Based on these observations, an overall solution can be obtained by using a modified version of the iterative algorithm developed for the no-converter case. In particular, a k-hop path with a single converter located at node $K < k$ can be analyzed using the algorithm in Figure 5 after making a single modification as follows: in Step 6, the expression for the conditional blocking probabilities must be changed to $q_{ij} = p^{(2)}_{1,j-K}$. The second term in the expression of Step 6 in Figure 5 (as well as in expression (22)) 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

We now demonstrate the accuracy of the blocking probabilities obtained with our analytical techniques by comparing to either numerical or simulation results. We then apply our decomposition algorithm in order to determine the optimal placement of converters in a path.

5.1 Accuracy of the Approximate Markov Process

In Figure 7 we plot the blocking probability of calls for each source-destination node 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 as "exact solution" in the figure. 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". As we can see, the overall behavior of the two curves is similar for all types of calls, and, more importantly, the approximate blocking probability is always very close to the exact value.

Figure 8 is similar to Figure 7, but presents results for a 3-hop path. We only plot the blocking probability of calls for three of the six source-destination pairs, namely, calls that traverse all three hops, calls that use only the last two hops, and calls that use only the first hop. The behavior of the blocking probability of the other three types of calls is very similar to that shown in Figure 8. Again, we observe that the values of blocking probability obtained through the closed-form solution to Markov process $\mathcal{M}'_3$ are an accurate approximation of the exact values obtained from process $\mathcal{M}_3$. However, the figure does not include any values for the exact blocking probability in a 3-hop path when $W > 4$ because the size of the exact Markov process is so large in this case that we could not solve it numerically. (For the same reason we do not show any comparison results for 4-hop paths.)

Overall, the results shown in Figures 7 and 8 indicate that it is appropriate to use the approximate Markov process when analyzing short paths. We emphasize that exact Markov process cannot be solved numerically for a large number of wavelengths. Also, for the same values of k and W, the closed-form solution of the approximate Markov process can be obtained significantly faster (up to an order of magnitude) than the numerical solution to the exact process.

5.2 Validation of the Decomposition Algorithm

The figures in this subsection show the results of applying our iterative decomposition algorithm to 6-hop and 10-hop paths with $W = 10$ wavelengths, and compare them to results obtained through simulation. We include results for paths with and without converters. Because of the very large number of parameters (i.e, the arrival rates λ_{ij} and mean call holding times $1/\mu_{ij}$) that can potentially be varied, for the results presented here we have decided to use the following values:

$$\mu_{ij} \quad = \quad 1, \ \forall\, i,j, \qquad \text{and} \qquad \lambda_{ij} \quad = \quad \begin{cases} 0.3, & i < j \\ \lambda, & i = j \end{cases} \tag{23}$$

In other words, we let the mean holding time to be equal to one for all calls, we fix the arrival rate of all calls traversing two or more hops to the value 0.3, and we also set the arrival rate of calls traversing exactly one hop to λ. Figures 9 – 12 plot the call blocking probability as we vary the value of λ.

In Figures 9 and 10 we present results for a 6-hop path with and without converters, and for $W = 10$. We note that there are 21 different source-destination pairs in a 6-hop path, making it impossible to present results for all of them here. Thus, we show results for only two source-destination pairs. The blocking probability of calls traversing all six hops in the path is plotted in Figure 9, while the blocking probability of calls traversing hops 1 through 4 of the path is shown

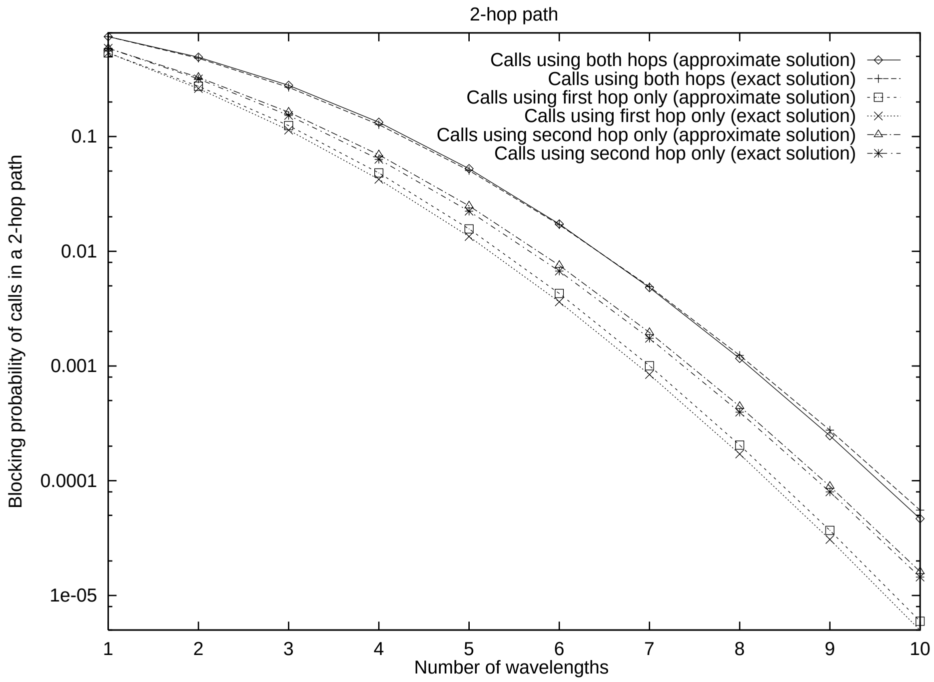

Figure 7: Exact and approximate blocking probability of the various calls in a 2-hop path

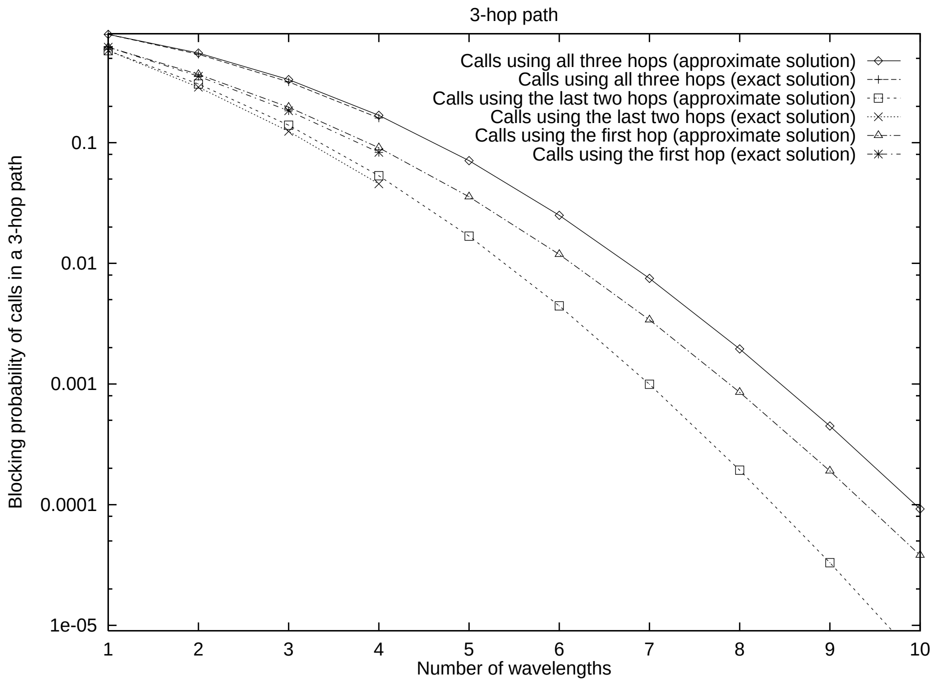

Figure 8: Exact and approximate blocking probability of various calls in a 3-hop path

in Figure 10. 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, as described above. Each figure contains two sets of plots, one for a path without converters, and one for a path with a single converter. Each set consists of two plots, one corresponding to results from our decomposition algorithm, and one corresponding to simulation results. Recall that, for $W = 10$ wavelengths, we can efficiently compute the closed-form solution of the approximate Markov process for a 3-hop path, but not for a 4-hop path. Hence, we let $K = 3$ in the decomposition algorithm of Figure 5, and we analyze a 6-hop path without converters by decomposing it into two 3-hop segments. We refer to this decomposition as a "3×3 decomposition". For the case when there is a single converter in the path, we decided to place the converter at node 3, since this arrangement will also result in a 3×3 decomposition.

From Figures 9 and 10 we observe that as the value of the load 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. The most important observation from these figures, however, is the fact that 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.

The results in Figures 11 and 12 are very similar to those in Figures 9 and 10, but they correspond to a 10-hop path with $W = 10$ wavelengths. 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 0.25. Figure 11 shows the blocking probability of calls traversing all ten hops of the path, while Figure 12 presents the blocking probability of calls using hops 2 through 6. As before, there are two sets of plots, one for a path without converters, and one for a path employing three converters. For the no-converter case, in addition to simulation results, we present the blocking probability values obtained through two decomposition approaches for a 10-hop path: a $3 \times 3 \times 3 \times 1$ decomposition into three 3-hop segments followed by a 1-hop segment, and a "dual" $1 \times 3 \times 3 \times 3$ decomposition. For the converter case, the three converters are assumed to be at nodes 3, 6, and 9, a configuration that also results in a $3 \times 3 \times 3 \times 1$ decomposition.

The general behavior of the curves in Figures 11 and 12 as λ increases are very similar to that of the curves in Figures 9 and 10. Regarding the accuracy of our analytical techniques, we note that the results of our decomposition algorithm closely match the simulation results for the path with converters. For the path without converters, the blocking probability values for calls using hops 2 to 6 (in Figure 12 obtained through the decomposition algorithm, although the $3 \times 3 \times 3 \times 1$

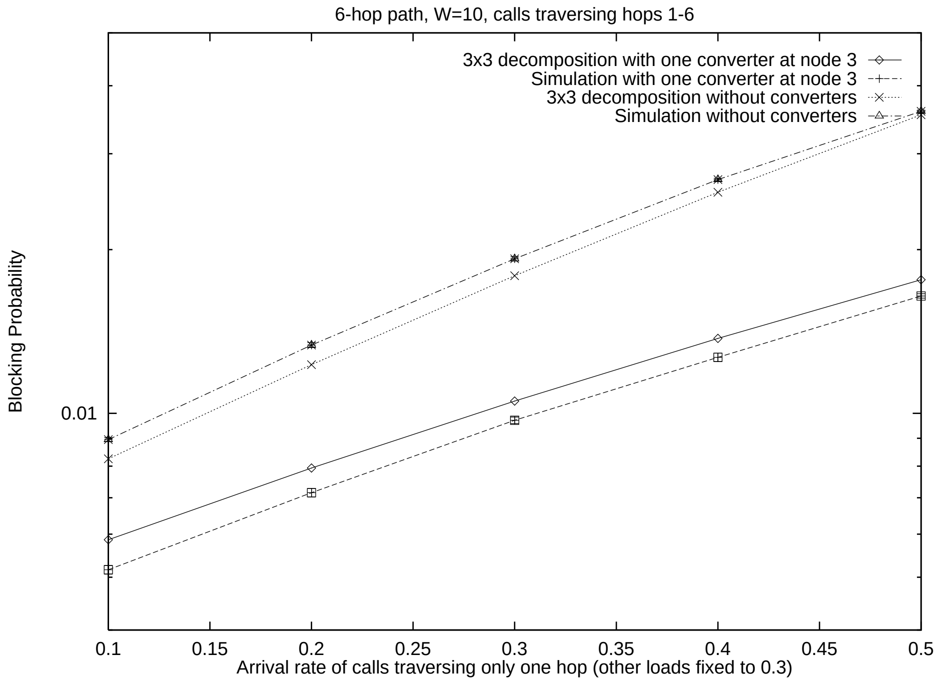

Figure 9: Blocking probability of calls traversing all links of a 6-hop path with $W = 10$

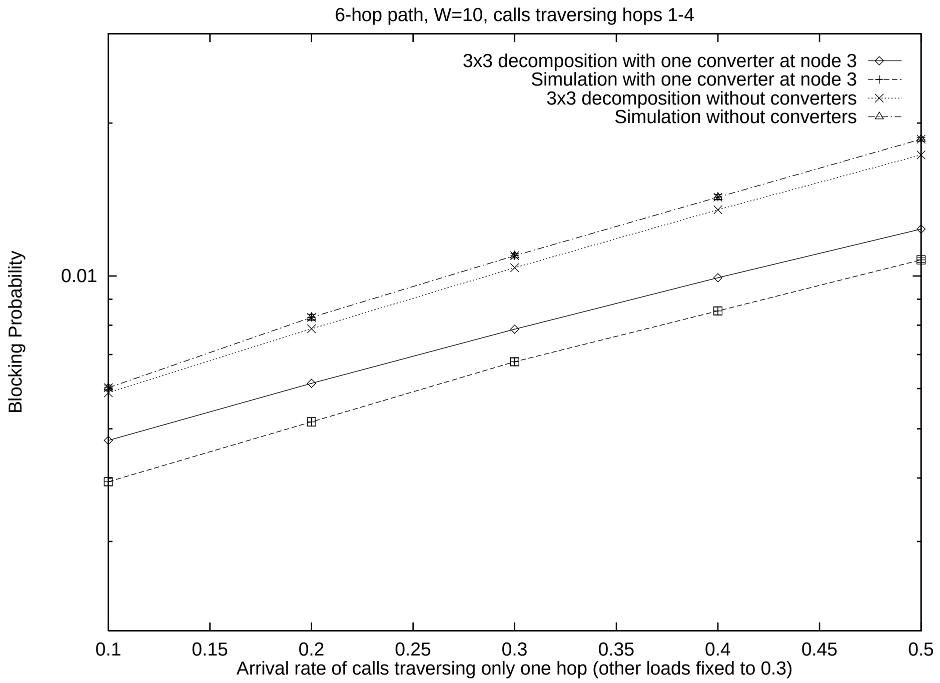

Figure 10: Blocking probability of calls traversing links 1 through 4 of a 6-hop path with $W = 10$

decomposition is more accurate than the $1 \times 3 \times 3 \times 3$ decomposition in this case. The situation is reversed when one considers calls traversing all ten hops in Figure 11, since the $1 \times 3 \times 3 \times 3$ decomposition gives more accurate results.

Compare decomposition method to simulation. In particular, mention that the decomposition algorithm converges in only a handful of iterations and takes a few minutes, while the simulation takes several hours.

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 optimizes a given objective function. It is important to emphasize that our goal is simply to demonstrate how the analytical techniques developed earlier can be used to gain insight into the problem of converter placement in a wavelength routing network. A comprehensive study of this potentially difficult problem is outside the scope of this paper and will be undertaken in future work. All the results presented in this section are obtained by using the following approach:

- The performance parameter of interest in a k-hop path is the blocking probability of the calls that travel over all k hops (i.e., the ones with arrival rate λ_{1k}, refer also to Figure 1). Thus, our objective is to determine the location of l converters that minimizes the blocking probability of these calls.
- To find the best converter placement we apply a straightforward method as follows. We enumerate all possible ways of placing l converters on a k-hop path, and we evaluate the blocking probability of interest for each using the decomposition algorithm presented in previous sections. The optimum placement is the one with the minimum such probability.

We note that calls traveling over all k hops of a k-hop path will experience the highest blocking probability. Hence, minimizing this parameter effectively minimizes the maximum blocking probability of any call. On the other hand, there are $\binom{k-1}{l}$ ways of placing l converters at the $k-1$ internal nodes of a k-hop path. Therefore, we expect our straightforward enumeration approach to work well for reasonably long paths (e.g., $k \leq 20$). For longer paths, when this method becomes prohibitively time consuming, heuristic approaches aimed at determining a near-optimal placement of converters are needed. The specification and evaluation of such heuristics will be the subject of future research.

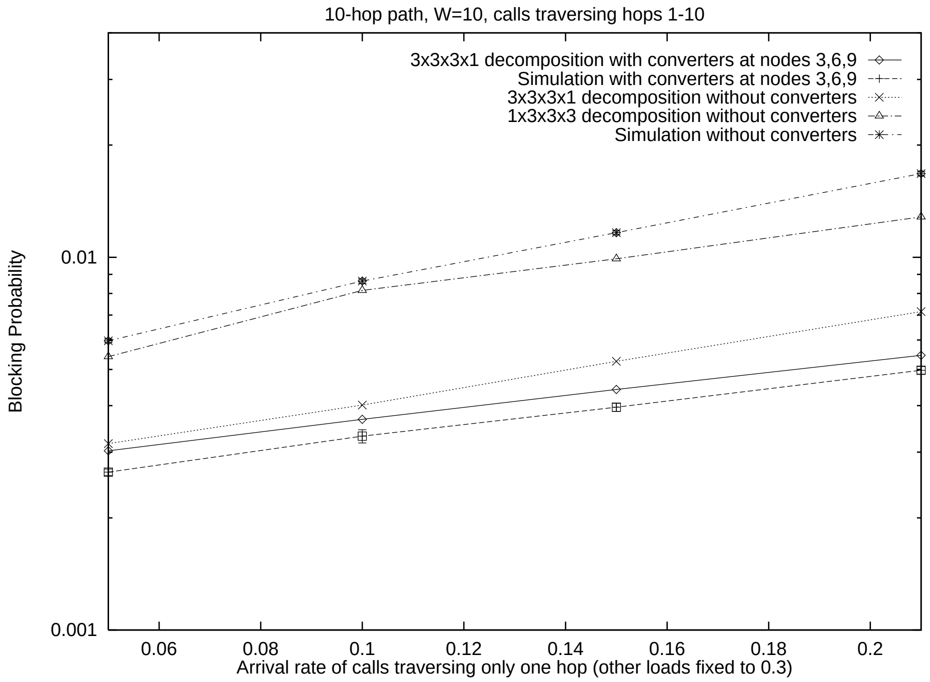

Figure 11: Blocking probability of calls traversing all links of a 10-hop path with $W = 10$

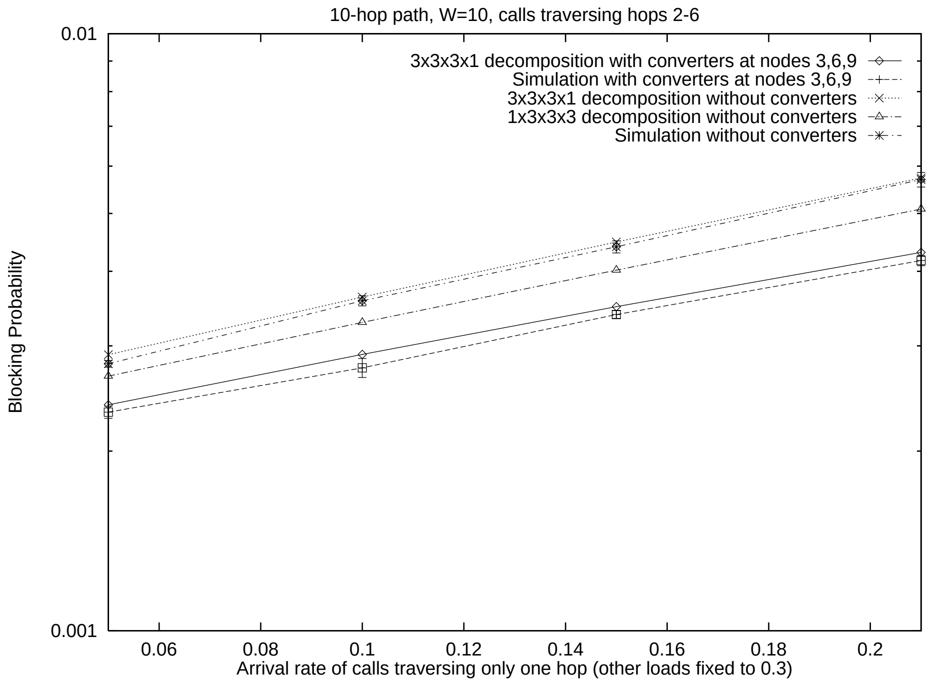

Figure 12: Blocking probability of calls traversing links 2 through 6 of a 10-hop path with $W = 10$

For our study, we consider a 10-hop path with $W = 10$ wavelengths, and six different traffic load patterns. Figures 13 and 14 plot the load ρ of each hop in the path for each load pattern [5]. The six patterns are representative of the wide range of loading situations one expects to encounter in practice. In the "uniform" pattern, all hops are equally loaded. The "bowl" (respectively, "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 fop 10. These patterns are shown in Figure 13. The "ascending" and "descending" patterns are such that the load increases or decreases, respectively, from hop 1 to hop 10. Finally, in the oscillating pattern the load at each hop alternates between a low and a high value. The last three load patterns are shown in Figure 14. To ensure that the results are comparable across the different patterns, the load values were chosen so that the total network load (or, equivalently, the average load per hop) is the same for all patterns.

In Figures 15 and 16 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 blocking probabilities shown in Figures 15 and 16 correspond to the load patterns of Figures 13 and 14, respectively. The optimal location of converters for the six different load patterns is shown in Table 1.

From Figures 15 and 16 it is immediately obvious that the blocking probability drops with the number of converters, as expected. However, after the initial steep drop, the curves in general flatten as the number of converters increases. This behavior is consistent with the results of an earlier work [24] which studied the effect of converter density on the blocking probability. But the results presented here shed new light into the problems of determining the importance of converters and their location, which have previously been studied only under uniform traffic loading.

A striking observation from Figures 15 and 16 is that the effect of converters on the blocking probability is strongly dependent on the actual traffic pattern. Let us compare, for example the curves for the "inverted bowl" (in Figure 15) and "ascending" (in Figure 16) load pattern. The former curve is essentially flat: the blocking probability decreases from a value of 0.0041 for one converter to a value of 0.0028 for five converters. In contrast, increasing the number of converters has a dramatic effect for the "ascending" load pattern, with the blocking probability dropping from a value of 0.0083 for one converter to 0.0018 for five converters, a 78% decrease.

Let us now refer to Table 1 which shows the optimal node location of converters for the different traffic patterns. We first note that the results are in agreement with intuition. The table indicates,

[5]The load ρ for hop l is the sum of the offered loads $\rho_{ij}, i \leq l \leq j$, for all calls that use hop l.

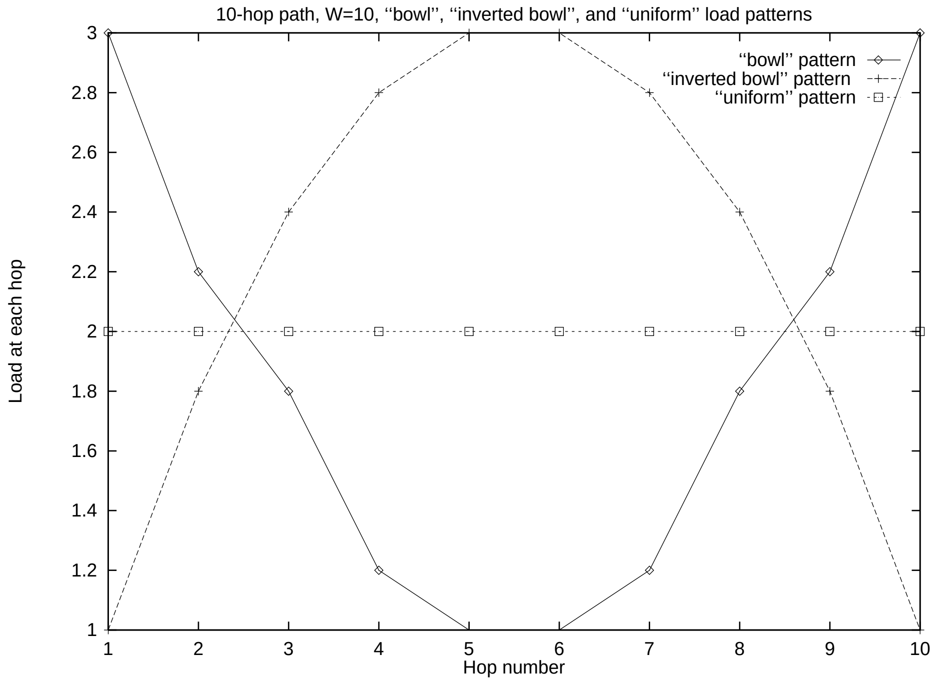

Figure 13: The "bowl", "inverted bowl", and "uniform" load patterns

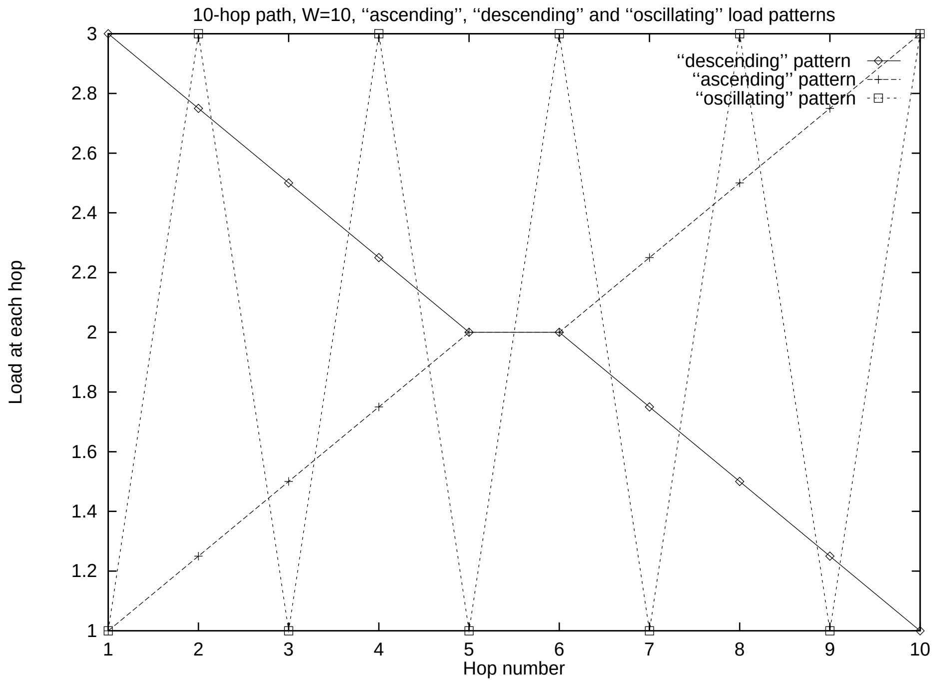

Figure 14: The "ascending", "descending" and "oscillating" load patterns

Table 1: Optimal node location of converters for the various load patterns

Load Pattern	Number of Converters				
	1	2	3	4	5
bowl	8	18	789	1589	12589
inverted bowl	4	45	457	3456	34567
uniform	7	28	479	1279	12379
ascending	8	79	689	5789	56789
descending	2	12	123	1235	12345
oscillating	7	58	479	4789	24789

for example, that converters be placed at the middle of the path for the "inverted bowl" pattern, and towards the end and beginning of the path for the "ascending" and "descending" patterns, respectively. However, we also observe that the optimal placement also depends strongly on the load pattern: compare, for instance, the converter locations for the "ascending" and "descending" patterns. The fact that optimal location depends on the load suggests that in a dynamic network environment where the traffic pattern varies 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 presented a new analytical framework to accurately and efficiently evaluate the call blocking performance in a single path of a wavelength routing network. We derived exact and approximate Markov process models, and we developed an iterative algorithm to analyze arbitrarily long paths by decomposing them into shorter segments connected in tandem which are studied in isolation. We also demonstrate how our analytical techniques can be used to gain insight into the problem of converter placement in such a network.

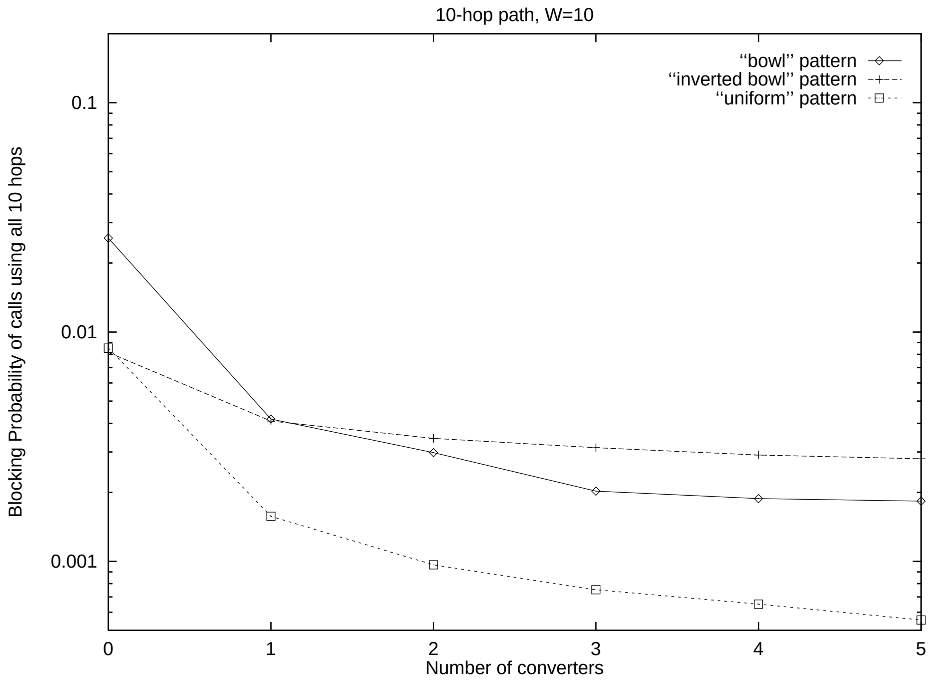

Figure 15: Blocking probability under optimal converter placement for the "bowl", "inverted bowl", and "uniform" load patterns

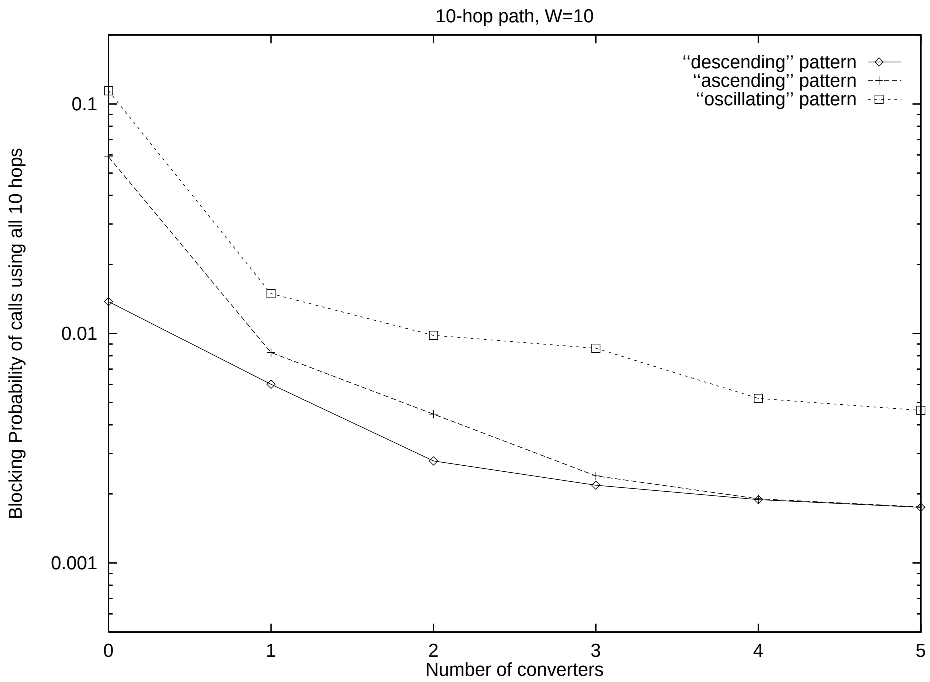

Figure 16: Blocking probability under optimal converter placement for the "ascending", "descending" and "oscillating" load patterns

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] M. S. Borella and B. Mukherjee. Efficient scheduling of nonuniform packet traffic in a WDM/TDM local lightwave network with arbitrary transceiver tuning latencies. *IEEE Journal on Selected Areas in Communications*, 14(5):923–934, June 1996.

[4] C. A. Brackett. Is there an emerging consensus on WDM networking? *Journal of Lightwave Technology*, 14(6):936–941, June 1996.

[5] C. E. Catlett. In search of gigabit applications. *IEEE Communications Magazine*, 30(4):42–51, April 1992.

[6] 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.

[7] B. Mukherjee *et al.* Some principles for designing a wide-area optical network. In *Proceedings of INFOCOM '94*, pages 110–119. IEEE, June 1994.

[8] E. Hall *et al.* The Rainbow-II gigabit optical network. *IEEE Journal Selected Areas in Communications*, 14(5):814–823, June 1996.

[9] L. Kleinrock *et al.* The supercomputer supernet testbed: A WDM-based supercomputer interconnect. *Journal of Lightwave Technology*, 14(6):1388–1399, June 1996.

[10] R. E. Wagner *et al.* MONET: Multiwavelength optical networking. *Journal of Lightwave Technology*, 14(6):1349–1355, June 1996.

[11] S. P. Monacos *et al.* All-optical WDM packet networks. *Journal of Lightwave Technology*, 14(6):1356–1370, June 1996.

[12] T. K. Chiang *et al.* Implementation of STARNET: A WDM computer communications network. *IEEE Journal Selected Areas in Communications*, 14(5):824–839, June 1996.

[13] P. E. Green. Optical networking update. *IEEE Journal Selected Areas in Communications*, 14(5):764–779, June 1996.

[14] F. P. Kelly. *Reversibility and Stochastic Networks.* John Wiley & Sons, New York, 1979.

[15] L. Kleinrock. *Queueing Systems, Volume 1: Theory.* John Wiley & Sons, New York, 1975.

[16] 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.

[17] B. Mukherjee. WDM-Based local lightwave networks Part I: Single-hop systems. *IEEE Network Magazine*, pages 12–27, May 1992.

[18] B. Mukherjee. WDM-Based local lightwave networks Part II: Multihop systems. *IEEE Network Magazine*, pages 20–32, July 1992.

[19] H. Perros. *Queueing Networks with Blocking: Exact and Approximate Solutions.* Oxford University Press, 1994.

[20] R. Ramaswami. Multiwavelength lightwave networks for computer communication. *IEEE Communications Magazine*, pages 78–88, February 1993.

[21] R. Ramaswami and K. N. Sivarajan. Design of logical topologies for wavelength-routed optical networks. *IEEE Journal Selected Areas in Communications*, 14(5):840–851, June 1996.

[22] M. N. Ransom and D. R. Spears. Applications of public gigabit networks. *IEEE Network Magazine*, March 1992.

[23] K. Sivalingam and J. Wang. Media access protocols for WDM networks with on-line scheduling. *IEEE/OSA Journal of Lightwave Technology*, 14(6):1278–1286, June 1996.

[24] 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.

[25] S. Subramanian, 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.

[26] N. Wauters and P. Demeester. Design of the optical path layer in multiwavelength cross-connected network. *IEEE Journal Selected Areas in Communications*, 14(5):881–892, June 1996.

A Proof of Lemma 3.1

Proof. We will show that the state descriptor in (2) provides all the information necessary to predict the future evolution of a k-hop system. Let us first consider the arrival of a call that uses hops i through $j \geq i$, inclusive, and which finds the system in state $\underline{n}$. If $i < j$, the number of wavelengths that are free on hops i through j is given by parameter f_{ij}. If $i = j$, the number f_{ii} of wavelengths that are free on hop i, is not included in the state description of MC $\mathcal{M}_k$, but it can be easily derived as

$$f_{ii} \quad = \quad W \; - \sum_{1 \leq l \leq i \leq m \leq k} n_{lm} \tag{24}$$

If $f_{ij} = 0$, the call is blocked, and the system remains at state $\underline{n}$. If $f_{ij} > 0$, however, the call will be randomly assigned one of the f_{ij} free wavelengths, and the system will make a transition to state $\underline{n'}$ such that:

$$n'_{lm} \quad = \quad \begin{cases} n_{lm} + 1, & l = 1 \text{ and } j = m \\ n_{lm}, & \text{otherwise} \end{cases} \quad , \quad f'_{ij} \quad = \quad f_{ij} - 1 \tag{25}$$

In addition, the values of some other parameters f_{lm} may be affected, depending on the particular wavelength assigned to the call. For instance, let us assume that $i = j$, and consider a parameter $f_{lm}, l \leq i \leq m$, in state $\underline{n}$ that has a non-zero value. Obviously, $f_{lm} \leq f_{ii}$, since the set of free wavelengths on hops l through m is a subset of the set of free wavelengths on hop i alone. Because of the random wavelength assignment policy, with probability $\frac{f_{lm}}{f_{ii}}$ the new state $\underline{n'}$ will have $f'_{lm} = f_{lm} - 1$, while with probability $1 - \frac{f_{lm}}{f_{ii}}$ it will have $f'_{lm} = f_{lm}$. Similarly when $i < j$. Thus, upon the arrival of a call request, the information in the descriptor of the current state $\underline{n}$ is sufficient to determine all possible new states $\underline{n'}$, as well as the corresponding transition rates.

Let us now consider the completion of a call that uses hops i through $j \geq i$, inclusive, when the system is in state $\underline{n}$ with $n_{ij} > 0$. The system will make a transition to a new state $\underline{n'}$ such that:

$$n'_{lm} \quad = \quad \begin{cases} n_{lm} - 1, & l = 1 \text{ and } j = m \\ n_{lm}, & \text{otherwise} \end{cases} \quad , \quad f'_{ij} \quad = \quad f_{ij} + 1 \tag{26}$$

As with a call arrival, the values of some other parameters f_{lm} may be affected. For simplicity, let us again assume that $i = j$; similar arguments apply when $i < j$. Because of the memoryless property, the completed call is equally likely to be any of the n_{ii} active calls on hop i. Consider parameter $f_{lm}, l \leq i \leq m$, and parameters $f_{l,i-1}$ and $f_{i+1,m}$. Obviously, $f_{lm} \leq f_{l,i-1}$ and $f_{lm} \leq f_{i+1,m}$. If $f_{lm} = f_{l,i-1}$ or $f_{lm} = f_{i+1,m}$, the transition due to the completion of the call will not affect the value of f_{lm}. But if $f_{lm} < f_{l,i-1}$ and $f_{lm} < f_{i+1,m}$, if the newly freed wavelength is the same as one of the free wavelengths on hops l through $i - 1$ *and* on hops $i + 1$ through m, then the new state will have $f'_{lm} = f_{lm} + 1$. The probability of the last event can be computed as follows. Of the $f_{l,i-1}$ wavelengths, $f_{l,i}$ are also free on hop i, so $f_{l,i-1} - f_{l,i}$ are not free on hop i. Similarly, of the $f_{i+1,m}$ wavelengths, $f_{i+1,m} - fi, m$ are not free on hop i. The probability that the newly freed wavelength is one of $f_{l,i-1} - f_{l,i}$ is simply $\frac{f_{l,i-1}-f_{l,i}}{n_{ii}}$, and the probability that it is one of $f_{i+1,m} - f_{i,m}$ is $\frac{f_{i+1,m}-f_{i,m}}{n_{ii}}$. Therefore, the probability that the newly freed wavelength will increase the number f_{lm} of wavelengths that are free on hops l is $\frac{f_{l,i-1}-f_{l,i}}{n_{ii}} \frac{f_{i+1,m}-f_{i,m}}{n_{ii}}$. Again, we have sufficient information to determine the set of possible new states $\underline{n}'$ and the transition rates upon the termination of an active call. $\square$