Quantifying the benefit of configurability in circuit-switched WDM ring networks[1]

Brett Schein
Laboratory for Information
and Decision Systems, MIT
77 Massachusetts Ave, Rm 35-316
Cambridge, MA 02139
schein@mit.edu

Eytan Modiano
Laboratory for Information
and Decision Systems, MIT
77 Massachusetts Ave, Rm 33-219
Cambridge, MA 02139
modiano@mit.edu

Abstract—**We attempt to characterize the gain in traffic capacity that a reconfigurable network offers over a fixed topology network. We define the gain as the ratio of the maximum offered loads that the two systems can support for a given blocking probability. We develop a system model to analytically predict the blocking probability for both the fixed and reconfigurable systems. This model is different from previous models developed to analyze the blocking probability in WDM networks in that it accounts for a port limitation at the nodes.**

We study high bandwidth calls, where each call requires an entire wavelength. We find that reconfigurability offers a substantial performance improvement, particularly when the number of available wavelengths significantly exceeds the number of ports per node. In this case, we find that the gain approaches a factor of $\frac{N}{2}$ over a fixed topology unidirectional ring, and $\frac{N}{4}$ over a fixed topology bi-directional ring (where N is the number of nodes in the ring). We validate our model via simulation, and we find that it agrees strongly with the simulation results, particularly for a large number of ports per node. We also obtain upper and lower bounds on capacity for various ring topologies that give additional insight into the benefits of configurability.

Keywords—**Keywords: optical networks, WDM, reconfiguration, configurability.**

I. INTRODUCTION

We take a preliminary look at reconfigurability in circuit-switched wavelength division multiplexed (WDM) ring networks. In WDM networks, the physical topology consists of passive or configurable optical nodes interconnected with fiber links. In a fixed topology system, permanent lightpaths are set up between nodes to construct the logical topology of the network. Traffic is then routed on this fixed logical topology — the set of lightpaths are maintained regardless whether traffic is carried on them. In a reconfigurable topology, lightpaths can be dynamically reconfigured to reflect changes in traffic conditions. A reconfigurable network can thus adapt to changing traffic patterns.

To get an intuitive feeling for why reconfigurability can be advantageous, consider four nodes physically connected in a ring. Assume that each node has one port, that the fiber supports two wavelengths, λ_1 and λ_2, and that a call takes a full wavelength. A connected, fixed logical topology must take the form of a unidirectional ring, as pictured on the left in Figure 1. If a call is in progress from node 1 to node 3, and a call request arrives from node 2 to node 4, then that request must be blocked despite the fact that there is sufficient capacity on the fiber.

[1]This work was supported by DARPA under the Next Generation Internet program.

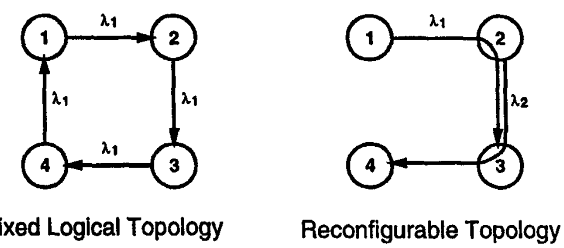

Fig. 1. Fixed Versus Reconfigurable Network Idea

In a reconfigurable system, both calls can be supported as shown on the right in Figure 1. The call between nodes 1 and 3 can be routed without requiring a port at node 2. The conceptual idea behind reconfiguration is that reconfiguring the logical topology of a network utilizes available wavelengths on a fiber without dedicated electronic ports, bypassing the electronic layer at most intermediate nodes. In the above example, node 2 was a bottleneck because it had to process the call between nodes 1 and 3 though that call was not intended for it. By reconfiguring the topology so that node 1 is directly connected to node 3 via a wavelength, this bottleneck was alleviated.

An important characterization of reconfigurability is the time scale in which lightpaths can be changed relative to changes in the offerred traffic pattern. A slowly reconfigurable network, tuned on the order of minutes or hours, can expect several new circuits to be placed and removed in the time required to reconfigure. Such slow reconfiguration is useful for adapting to predictable variations in the statistics of the offerred traffic. For example, if it is known that in a nationwide telephone network, traffic flows more heavily from east to west in the morning and from west to east in the evening, then the network can reconfigure lightpaths to reflect this. In a simplified sense, an optimized fixed logical topology can be designed for each variety of offerred traffic statistics, and then the topology can migrate to the appropriate optimized fixed logical topolgy. Similarly, in a packet-switched system, the logical topology of the network can be reconfigured for load balancing. Recently, topology design and reconfiguration algorithms have been developed for reducing the electronic processing load in WDM-based packet networks ([1], [2], [3], [4], [5]).

When we speak of reconfigurability in this paper, we assume that lightpaths can be changed within an acceptable delay at call setup. This requires the use of tunable lasers and configurable WDM switches that can be tuned on the order of tenths of a second. These physical components are more complex than their fixed counterparts. Therefore, although it is clear that reconfiguration offers a performance benefit, it is important to evaluate it carefully. In the case of a circuit-switched network, the benefit can be expressed in terms of increased traffic load that the network can support for a given blocking probability. Hence, in order to evaluate the benefit of reconfiguration, one must be able to compute the blocking probabilities for both the fixed topology and reconfigurable network.

Many researchers have studied blocking probabilities for circuit-switched WDM networks with or without wavelength changers ([6], [7], [8]). Earlier work assumed that wavelengths are the precious resource in the network. Therefore, in analyzing blocking probabilities, previous researchers assumed that a call request can be placed in the system if and only if a wavelength (or a series of wavelengths using wavelength changers) is available between the source and destination, thus ignoring the possibility of calls being blocked due to the lack of electronic resources. However, when considering a multi-hop circuit switched network, calls can be blocked when lightpaths are available. A call may be blocked because ports on the source or destination nodes are occupied or because an intermediate node has no ports available. In order to analyze the blocking probability in such a system, a model for blocking probability that takes both the wavelengths and electronic ports into account must be developed.

Unfortunately, calculating the blocking probability for a reconfigurable system is complex because, in general, calls are not electronically processed at every node. Unlike the fixed topology systems, there is no one-to-one correspondence between wavelengths on a fiber and node ports. A precise analysis therefore requires global information about the state of the network, resulting in a computationally intensive and uninsightful approach.

To avoid maintaining global state information, we develop a stochastic system model to analytically predict the blocking probability for both the fixed and reconfigurable systems. This model is based largely on that introduced by Barry and Humblet [6], which we extend to address a port limitation at the nodes. We develop an iterative computational method to calculate blocking probabilities in ring networks. We focus on ring networks because of their simplicity and ubiquity. We validate our model via simulation, and we find that the sustainable load predicted by our model at low blocking probability agrees strongly with the simulation results, particularly for a large number of ports. Finally, we develop upper and lower bounds for fixed topology, unidirectional rings and for reconfigurable rings with a large number of wavelengths. These bounds yield additional insight into the difference between configurable and fixed topology systems.

II. Physical Assumptions and Traffic Model

Our network consists of N nodes physically located in a ring and connected by fiber. Each fiber contains W wavelengths and each node has P electronic ports. Each electronic port consists of a transmitter and receiver that is tunable to any one of the W wavelengths. Furthermore, each node has a configurable WDM switch that can allow each wavelength to either bypass the node or be processed at a port at that node.

For a fixed topology system, all transmitters and receivers are fixed tuned to their chosen wavelengths when the system is constructed and are never changed. We consider two fixed topology systems — the unidirectional and the bidirectional ring topologies. The fixed topology systems use P wavelengths, all of which are processed at every node, and the P transceivers at each node are tuned to the same set of P wavelengths on the fiber. In the unidirectional ring, P lightpaths are set up between successive nodes in the ring, all in the same direction. In the bidirectional ring, $\frac{P}{2}$ lightpaths are set up between successive nodes in the ring in each direction. Thus the bidirectional ring can be considered two unidirectional rings of $\frac{P}{2}$ ports, routed in opposite directions.

For the reconfigurable topology, all unused ports at a node can be tuned to any unused wavelength in either direction. Furthermore, the WDM switch can be configured dynamically to have a wavelength either bypass or be processed at the node. We assume that calls require a full wavelength and that whenever a call request arrives, it is placed in the system if at all possible, provided that existing calls do not have to be rearranged. If a call request cannot be placed given the current calls in the system, it is blocked and departs from the system. In a reconfigurable network, calls can be placed in one of two ways. A call can be placed using a single wavelength from source to destination provided that one is available. If no single wavelength is free, a call can be placed using intermediate nodes and different wavelengths. In the latter case, the intermediate nodes essentially serve as wavelength changers, with wavelength changing accomplished via electronic multiplexing (for the fixed topology system, this is done at every intermediate node). Furthermore, when multiple routing options are available, calls are placed on the shortest path, use the fewest number of hops on this path, and randomly choose viable wavelengths.

Call requests arrive to the system as a Poisson process of rate Nr_0. Each call request has a uniformly distributed source and destination (different from the source). This is equivalent to an independent Poisson call request arrival process of rate $\frac{r_0}{N-1}$ between each source-destination pair. Call durations are exponentially distributed.

III. Approximate System Analysis

A. *Stochastic Model Structure*

With the given assumptions, the entire system can be represented as a single, finite, continuous-time Markov chain where each state represents a particular configuration of calls in progress around the ring. Though precise, this Markov chain is

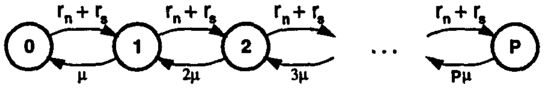

Fig. 2. Port Usage Model

too large to be computationally useful or insightful. We therefore develop a stochastic approach for estimating the blocking probability.

Our analytical approach relies on two main stochastic model components. The first component models the port usage at the nodes, and it is used for two purposes. First, it determines the probability that a source node has a free transmitter, a destination node has a free receiver, and that any intermediate nodes used for placing the call have a free transceiver. Second, it determines an important system parameter used by the second component, γ, the average number of calls sourced by a node. The second component models the wavelength usage along a single path through the ring (e.g., one of the two directions from source to destination). It uses updated information from the first component to revise estimates of blocking probabilities for calls in the rings. This second component is then used to update estimates of two rates required by the first component — the rate at which calls are accepted into the system, r_n, and the rate at which intermediate nodes are used in the placing of calls, r_s.

Our model can be viewed as a generalization of that in [6] with which we try to accomplish two things. First, we capture the effects of a limited number of ports available at each node to process calls (as a source, destination, or an intermediate hop). Second, we construct an iterative computational method to estimate the blocking probability without requiring important operating system parameters, such as the average utilization of wavelengths and the average number of hops that a call takes.

B. Stochastic Model Components

In order to obtain the call blocking probability without resorting to a global Markov chain, we use a number of approximations. We assume that the number of ports used at a node at any given time is independent of the number used at the other nodes. This approximation is used to analytically decouple the ports at different nodes. We let α be the probability that an arbitrary node has no free port available when a call request arrives to the system, and we assume that it is independent from node to node.

We estimate α directly from the first model component, the port usage model. This model component is represented as a continuous-time Markov chain where the state number equals the number of busy ports at a node, and there is one such independent chain for each node in the system. The model is drawn for a single node in Figure 2.

The port usage model is parametrized by two Poisson arrival rates, r_n and r_s, and an exponential call duration parameter μ. The Poisson processes parametrized by these arrival rates are assumed independent. The first arrival rate, r_n, represents the rate of accepting new calls into the system sourced at the given node. It is important to note that r_n is not equal to the offered traffic per node, r_o, since some calls are blocked. The second arrival rate, r_s, represents the rate of accepting new calls into the system using the given node as an intermediate hop in placing calls.

Then α is the probability of being in state P, where all ports are busy, and is given by the well-known Erlang B formula.

$$\alpha = \frac{\frac{\left(\frac{r_n+r_s}{\mu}\right)^P}{P!}}{\sum_{j=0}^{P} \frac{\left(\frac{r_n+r_s}{\mu}\right)^j}{j!}}.$$

In addition to determining α, we use this port model to estimate γ, the average number of active calls sourced at a node in equilibrium. Denoting the steady state probabilities of the Markov chain by $\{p_k\}_{k=0}^{P}$,

$$\gamma = \sum_{k=0}^{P} \frac{r_n}{r_n + r_s} \cdot k p_k = \frac{r_n}{\mu} \cdot \left(1 - \frac{\frac{\left(\frac{r_n+r_s}{\mu}\right)^P}{P!}}{\sum_{j=0}^{P} \frac{\left(\frac{r_n+r_s}{\mu}\right)^j}{j!}} \right).$$

Note that γ equals the average state number of the Markov chain weighted by the fraction of ports used for sourcing a call.

Next, ignoring the port usage in the system, we want to model the wavelength usage when an arbitrary call request arrives. We use the local stochastic model defined in [6], albeit with a slight modification. This second stochastic component models the wavelength usage along a single possible path from source to destination. We will use this component, along with the estimates of γ and α, to estimate the probability that a new call can be routed from source to destination either directly on an unused wavelength or by using intermediate nodes to change wavelengths. Based on these blocking probabilities, we will update the estimates of r_n and r_s used in the first component.

We formally define the second component model, the wavelength usage model, along a single possible route from source node 0 to destination node L as follows, where λ represents an arbitrary wavelength:

1. All events on different wavelengths are statistically independent.
2. The marginal probability that λ is used on a link is ρ.
3. Given the state of link $i-1$, i.e., whether λ is used on link $i-1$, the state of link i is statistically independent of the states of links $1, 2, \ldots, i-2$.
4. Given that λ is used on hop $i-1$, we assume that the call on λ terminates at node i with probability P_t. Otherwise, it continues on the same wavelength.
5. If λ is not used on hop $i-1$, then a new call is sourced at node i on λ with probability P_n.
6. If λ is used on hop $i-1$ and the call using λ on hop $i-1$ terminates at node i, then a new call is sourced at node i on λ with probability P_n.

We have changed the wavelength model defined in [6] in assumption 2, where previously it was assumed that all wavelengths were used with probability P_n on the first link.

C. Calculating the Blocking Probability

Assuming a source and destination have a free port to place a call on a path through the ring, we use our second component model to estimate the blocking probability for a call request of path length L along a single path (possibly one of two in a bidirectional system). For the fixed topology systems, where all P usable wavelengths are processed at every node and thus every intermediate node will be used in placing a call, the probability that a call of path length 1 will be blocked due to a wavelength limitation is the probability that all wavelengths are used on the first hop,

$$P_{\text{fix},1}(1) = \rho^W.$$

Define $\text{P}_{\text{full}}(i|\overline{i-1})$ as the probability that all wavelengths on link i are being used given that at least one wavelength is free on link $i-1$. Similarly, $\text{P}_{\text{full}}(i)$ is the probability that all wavelengths on link i are being used, and $\text{P}_{\text{full}}(\bar{i}) = 1 - \text{P}_{\text{full}}(i)$. Finally, $\text{P}_{\text{full}}(i, i-1)$ is the probability that all wavelengths are used on both links i and $i-1$. Then from Bayes' rule and elementary probability,

$$\text{P}_{\text{full}}(i|\overline{i-1}) \cdot \text{P}_{\text{full}}(\overline{i-1}) = \text{P}_{\text{full}}(i) - \text{P}_{\text{full}}(i, i-1).$$
$$\text{P}_{\text{full}}(i, i-1) = \text{P}_{\text{full}}(i|i-1) \cdot \text{P}_{\text{full}}(i-1).$$

We determine these probabilities from our second stage model. In particular,

$$\begin{aligned}\text{P}_{\text{full}}(\overline{i-1}) &= 1-\rho^W,\\ \text{P}_{\text{full}}(i) &= \rho^W,\\ \text{P}_{\text{full}}(i-1) &= \rho^W,\\ \text{P}_{\text{full}}(i|i-1) &= (1-P_l+P_lP_n)^W.\end{aligned}$$

This yields

$$\text{P}_{\text{full}}(\bar{i}|\overline{i-1}) = 1 - \frac{\rho^W(1-(1-P_l+P_lP_n)^W)}{1-\rho^W}.$$

Finally, noting that $\text{P}_{\text{full}}(\bar{i}|\overline{i-1}, \overline{i-2}) = \text{P}_{\text{full}}(\bar{i}|\overline{i-1})$ we find for $L > 1$ that

$$\begin{aligned}P_{\text{fix},1}(L) &= 1 - \text{P}_{\text{full}}(\bar{1}) \cdot \prod_{i=2}^{L} \text{P}_{\text{full}}(\bar{i}|\overline{i-1})\\ &= 1-(1-\rho^W)\\ &\quad \cdot \left[1 - \frac{\rho^W(1-(1-P_l+P_lP_n)^W)}{1-\rho^W}\right]^{L-1}. \quad (1)\end{aligned}$$

The three parameters ρ, P_n, and P_l in (1) are defined above for the second component model. The parameter ρ is called the wavelength utilization, and it represents the fraction of time a wavelength is used. P_n represents the probability that a call is sourced at a node on a wavelength given that the wavelength would otherwise not be used on the following link (given by 5 and 6 in the model definition). P_l represents the probability that a call terminates at each successive node and is closely related to the average path length of calls in the system.

Based on these interpretations, we set the required parameters as follows. $\overline{L}$ is the average path length of an accepted call and is initialized to $\frac{N}{2}$ for a unidirectional ring and $\frac{N}{4}$ for a bidirectional ring. $\overline{L}$ is later updated based on the estimated call blocking probabilities. Then,

$$\rho = \frac{N\gamma\overline{L}}{NW} = \frac{\gamma\overline{L}}{W}, \quad (2)$$

$$P_l = \frac{1}{\overline{L}}, \quad (3)$$

$$P_n = \frac{\gamma}{(1-\rho)W + \rho W P_l}. \quad (4)$$

The approximation in (2) is justified by noting that $(N\gamma\overline{L})$ is the average number of link wavelengths busy carrying calls in the system, while (NW) is the total number of link wavelengths available. This relation must be used with care during the numerical iteration before an accurate estimate of γ is available, but this is handled by restricting ρ to lie between 0 and 1. The parameter P_l in (3) is set so that the average path length of a call in progress is approximately $\overline{L}$. The parameter P_n in (4) is set by noting that, on average, $(1-\rho)W$ wavelengths are unused on the link prior to an arbitrary node, while $\rho W P_l$ calls terminate at the node. Therefore, we can set $\gamma = P_n((1-\rho)W + \rho W P_l)$.

The single path blocking probability for the reconfigurable system is more complex than that of the fixed topology system because calls in progress do not require many, if any, intermediate hops. Therefore, unlike the fixed topology systems, a busy wavelength on link i does not necessarily mean a port is used at node i for that call. Furthermore, an unused wavelength on link i does not necessarily mean that node i has a free port to process a new call.

One way to determine if it is possible to place a call in a reconfigurable system is to progress serially from source to destination, using the model to keep track of which wavelengths are used on successive links and using α to determine if switching wavelengths is possible at an intermediate node. To simplify the analysis and to simplify the estimation of how many intermediate hops are required to place a call, we assume

$$P_{\text{rec},1}(i|\overline{i-1}) = P_{\text{rec},1}(2|\bar{1}), \quad (5)$$

where $P_{\text{rec},1}(i|\overline{i-1})$ estimates the probability that a call of path length i will be blocked on this path due to a wavelength and intermediate node port limitation given that a call of path length $i-1$ *will not* be blocked. Herein we have assumed that the statistics of links $i-1$ and i, conditioned on being able to route a call a distance $i-1$, are identical to those of links 1 and 2, conditioned on being able to route a call over link 1.

Based on this assumption and the wavelength usage model, we derive the blocking probability as a function of call path length for the reconfigurable system. We define $P_{\text{rec},1}(L)$ as the probability that a call of path length L on a particular path will be blocked given that both the source and destination have a free port when the call request arrives. Then

$$P_{\text{rec},1}(1) = \rho^W. \tag{6}$$

Once we have calculated $P_{\text{rec},1}(2)$, then from Bayes' rule and elementary probability,

$$P_{\text{rec},1}(2|\overline{1}) = \frac{P_{\text{rec},1}(2) - P_{\text{rec},1}(1)}{1 - P_{\text{rec},1}(1)}.$$

Then using (5),

$$\begin{aligned}P_{\text{rec},1}(L) = P_{\text{rec},1}&(L-1)\\ &+(1 - P_{\text{rec},1}(L-1)) \cdot P_{\text{rec},1}(L|\overline{L-1})\\ &= P_{\text{rec},1}(1) \cdot (1 - P_{\text{rec},1}(2|\overline{1}))^{L-1}\\ &+ P_{\text{rec},1}(2|\overline{1}) \cdot \sum_{j=0}^{L-2} (1 - P_{\text{rec},1}(2|\overline{1}))^j.\end{aligned}$$

So we need only calculate $P_{\text{rec},1}(2)$ and use assumption (5) to determine $P_{\text{rec},1}(L)$ for all call path lengths $L \geq 2$. We determine $P_{\text{rec},1}(2)$ directly.

For the following discussion, in deriving $P_{\text{rec},1}(2)$, define the state of the network as the set of wavelengths used on links 1 and 2, as well as whether node 1 has a free port. Then partition the set of network states into the following three mutually exclusive and collectively exhaustive sets. Set $\mathcal{S}_1$ is the set of network states where all wavelengths are busy on link 1, all wavelengths are busy on link 2, or both. Set $\mathcal{S}_2$ is the set of network states where no single wavelength route is available, a route using two different wavelengths is available, but node 1 does not have a free port. The network state is in $\mathcal{S}_2$ if and only if every wavelength is used on either link 1 or link 2, there is at least one unused wavelength on link 1, there is at least one unused wavelength on link 2, and node 1 has no free port. Set $\mathcal{S}_3$ is the set of states where a call can be routed to node 2 either on a single wavelength or on two different wavelengths using node 1 as an intermediate hop.

The probability that the network state is in any of these sets can be found using our second component model and elementary probability. The call will be blocked if and only if the network state is in set $\mathcal{S}_1$ or set $\mathcal{S}_2$.

$$\begin{aligned}\text{Prob}(\mathcal{S}_1) &= \rho^W \cdot \left(2 - ((1-P_l) + P_l P_n)^W\right).\\ \text{Prob}(\mathcal{S}_2) &= \alpha\big((\rho + (1-\rho)P_n)^W\\ &\quad + (\rho(1 - P_l + P_l P_n))^W - 2\rho^W\big).\\ P_{\text{rec},1}(2) &= \text{Prob}(\mathcal{S}_1) + \text{Prob}(\mathcal{S}_2).\end{aligned}$$

Using the fact that $\rho = \rho(1 - P_l + P_l P_n) + (1-\rho)P_n$, we find

$$\begin{aligned}P_{\text{rec},1}(2) = \rho^W \cdot \Bigg(&\alpha\left(1 + (\frac{1}{\rho} - 1)P_n\right)^W\\ &+(1-\alpha)\left(2 - (1 - (\frac{1}{\rho} - 1)P_n)^W\right)\Bigg).\end{aligned} \tag{7}$$

Finally, for the reconfigurable, bidirectional ring system, we have two possible routes from a given source node to a given destination node. If one possible route has path length L, then the other possible route has path length $N - L$. Since a pair of free ports at the source and destination can be used to place calls in either direction in a reconfigurable ring, the blocking probability for a call of path length L, conditioned on having a pair of free ports, is

$$P_{\text{rec},2}(L) = P_{\text{rec},1}(L) \cdot P_{\text{rec},1}(N-L).$$

From these equations we can estimate the net blocking probability averaged over call path lengths and both possible routes for the systems. Conditioned on having a free port at the source and the destination,

$$\begin{aligned}P_{\text{fix},1} &= \frac{1}{N-1}\sum_{l=1}^{N-1} P_{\text{fix},1}(l),\\ P_{\text{rec},2} &= \frac{1}{N-1}\sum_{l=1}^{N-1} P_{\text{rec},1}(l) \cdot P_{\text{rec},1}(N-l).\end{aligned}$$

The form of the net blocking probability for the fixed topology, bidirectional system, $P_{\text{fix},2}$, differs from that of the reconfigurable system, $P_{\text{rec},2}$, because a single port can be used in either direction for the reconfigurable system but not for the fixed topology system. $P_{\text{fix},2}$ can be expressed under various conditioning assumptions, e.g., by assuming ports are available at the source and destination on the shortest path, but we omit such expressions for brevity.

Finally we give the update equations for the parameters required by the first stage model. We estimate the rate of placing new calls into the system, r_n, as follows. The probability that both the source and destination each have a free port available when a call request arrives is $(1-\alpha)^2$. For the fixed topology unidirectional ring, the fraction of these calls that can then be placed is $(1 - P_{\text{fix},1})$, and thus

$$r_n = r_o(1-\alpha)^2(1 - P_{\text{fix},1}).$$

The rate r_n is updated similarly for the other two systems. For all three systems,

$$r_s = \sum_{l=2}^{N-1} r_n p_l \overline{s_l}. \tag{8}$$

In (8), p_l is the probability that a call has path length l. This distribution of accepted call lengths is derived directly from the blocking probabilities and is used to estimate $\overline{L}$. The parameter $\overline{s_l}$ is the average number of intermediate hops used in placing a call of path length l. For the fixed topology systems, $\overline{s_l} = (l-1)$. For the reconfigurable system, $\overline{s_l}$ is estimated using assumption (5) and then counting the average number of intermediate nodes *required* to place a call of path length l. Our estimate results in $\overline{s_l} \propto (l-1)$, but we omit the derivation for brevity.

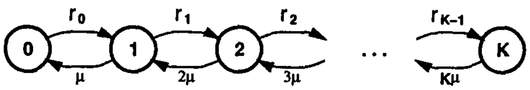

Fig. 3. A Class of Birth-Death Chains

IV. Bounds on Performance

In this section we develop upper and lower bounds on blocking probability. For most bounds we will use one core method — we will refer to this as the dominated rates technique. Consider a class of finite-state birth-death Markov chains with non-negative birth rates $\mathbf{r} = \{\mathbf{r_1}, \mathbf{r_2}, \dots, \mathbf{r_K}\}$ but fixed relative death rates $\mu_j = j\mu$, as in Figure 3. Both μ and the maximum state number K are fixed for the class, and we assume that $\mu > 0$. Note that we include cases where any (or even all) of the state birth rates equal 0. We denote by $\overline{X}$ the average state number for a particular member of this class, as calculated by the appropriate steady state probabilities. Note that the steady state probabilities always exist since $\mu > 0$. Consider two members of this class denoted by their birth rate vectors $\mathbf{r}_\mathrm{a}$ and $\mathbf{r}_\mathrm{b}$. Provided $\mathbf{r}_\mathrm{a} \geq \mathbf{r}_\mathrm{b}$ componentwise, it is easy to show that $\overline{X}_a \geq \overline{X}_b$.

Indeed, assume $\mathbf{r}_\mathrm{a} \geq \mathbf{r}_\mathrm{b}$ componentwise. Denote by $\mathrm{CDF}_X(x)$ the complementary distribution function of a random variable X, so $\mathrm{CDF}_X(x) = \mathrm{Prob}(X > x)$. Then the conclusion will follow when we prove that $\mathrm{CDF}_{X,a}(x) \geq \mathrm{CDF}_{X,b}(x)\ \forall\ x$ since $\overline{X} = \int_0^\infty \mathrm{CDF}(x)\, dx$ for any non-negative random variable. Assume that $r_\mathrm{a} \geq r_\mathrm{b} > 0$ componentwise. The latter strict inequality assumption is easily relaxed in the following proof.

Denote the steady state probabilities for the two member chains by $\{p_{x,\mathrm{a}}\}$ and $\{p_{x,\mathrm{b}}\}$. From the partial balance equations between states x and $x+1$, we know

$$\frac{p_{x,\mathrm{a}}}{p_{x+1,\mathrm{a}}} = \frac{r_{x,\mathrm{b}}}{r_{x,\mathrm{a}}} \cdot \frac{p_{x,\mathrm{b}}}{p_{x+1,\mathrm{b}}} \leq \frac{p_{x,\mathrm{b}}}{p_{x+1,\mathrm{b}}}\ \forall\, x \in [0, K-1], \tag{9}$$

$$\Rightarrow\ \frac{p_{x,\mathrm{a}}}{p_{x,\mathrm{b}}} \leq \frac{p_{x+1,\mathrm{a}}}{p_{x+1,\mathrm{b}}}\ \forall\, x \in [0, K-1]. \tag{10}$$

From (9) $\frac{p_{j,\mathrm{a}}}{p_{x,\mathrm{a}}} \leq \frac{p_{j,\mathrm{b}}}{p_{x,\mathrm{b}}}\ \forall\, j \leq x,\ x \in [0, K-1]$,

$$\Rightarrow\ \frac{\sum_{j=0}^{x} p_{j,\mathrm{a}}}{p_{x,\mathrm{a}}} \leq \frac{\sum_{j=0}^{i} p_{j,\mathrm{b}}}{p_{x,\mathrm{b}}}\ \forall\, x \in [0, K-1]. \tag{11}$$

Case 1 of 2: If $p_{x,\mathrm{a}} \leq p_{x,\mathrm{b}}$ then (11) implies

$$\sum_{j=0}^{x} p_{j,\mathrm{a}} \leq \sum_{j=0}^{x} p_{j,\mathrm{b}},$$
$$\Rightarrow\ \mathrm{CDF}_{X,a}(x) \geq \mathrm{CDF}_{X,b}(x).$$

Case 2 of 2: If $p_{x,\mathrm{a}} > p_{x,\mathrm{b}}$, then (10) implies

$$p_{j,\mathrm{a}} > p_{j,\mathrm{b}}\ \forall\, j > x,$$
$$\Rightarrow\ \sum_{j=x+1}^{K} p_{j,\mathrm{a}} > \sum_{j=x+1}^{K} p_{j,\mathrm{b}},,$$

or, in other symbols, $\mathrm{CDF}_{X,a}(x) > \mathrm{CDF}_{X,b}(x)$.
When we relax the condition that $r_{x,\mathrm{b}} > 0$, we need only redefine the top indices in the above summations. We have thus shown that $\overline{X}_a \geq \overline{X}_b$ whenever $r_\mathrm{a} \geq r_\mathrm{b}$.

We can now use Little's Law to relate the time-average accepted arrival rate, r_acc, and the average state number, $\overline{X}$, by $\overline{X} = \frac{r_\mathrm{acc}}{\mu}$. Using this technique, we can derive bounds on the blocking probability since $\mathrm{P_{block}} = 1 - \frac{r_\mathrm{acc}}{r_\mathrm{offer}}$.

Finally, we introduce notation for the Erlang B formula. Define for all $x > 0,\ P = 1, 2, \dots$,

$$\mathrm{P_{block,M/M/P/P}}(x) = \frac{\frac{(x)^P}{P!}}{\sum_{j=0}^{P} \frac{(x)^j}{j!}}.$$

Note that $\mathrm{P_{block,M/M/P/P}}$ is continuous and increasing in x for all fixed P.

A. Lower Bound for Fixed Topology, Unidirectional Ring

$\mathrm{P_{block}}$ is the blocking probability for all calls in the ring. We use symmetry arguments to equate $\mathrm{P_{block}}$ with the blocking probability for calls that cross an arbitrary but fixed link, say link i. Using the dominated rates technique, we then prove the intuitive statement that the blocking probability for calls crossing link i is lower bounded by the blocking probability for these calls when no other calls are in the system. Without these other calls, a call request that crosses link i will be blocked if and only if it is blocked on link i. The offered load to link i in the unidirectional ring is $r_\mathrm{o} \cdot \frac{N}{2}$, and in this case link i behaves exactly as an $M/M/P/P$ queue. To demonstrate the dominated rates technique explicitly, consider a Markov chain corresponding to link i, which has $P+1$ states, within the class of Markov chains with $K = P$. The state number corresponds to the number of calls currently supported on the link. The arrival rate for each state equals the rate of placing new calls using the link, and this rate is bounded by the offered load to the link, $r_\mathrm{o} \cdot \frac{N}{2}$. Therefore,

$$\mathrm{P_{block}} \geq \mathrm{P_{block,M/M/P/P}}\left(\frac{N r_\mathrm{o}}{2\mu}\right).$$

B. Upper Bound for Fixed Topology, Unidirectional Ring

We again use the dominated rates technique, but this time we bound the accepted call arrival rate as a function of the number of calls active in the system. We thus take a global view and consider a global Markov chain rather than one corresponding to a particular link. The Markov chain class is defined by μ and $K = NP$, since we could conceivably simultaneously hold P calls of length one for each node in the system.

We note here that this birth-death Markov chain precisely represents the system model, and this can be rigorously justified as follows. Consider a Markov chain with one unique state corresponding to every unique configuration of active calls in the system. Group these states into sets that correspond to the same number of active calls x in the system and write the global balance equations corresponding to each of these sets. There will be an associated accepted call rate r_x for each of these macrostates. Furthermore, $r_{x,\min} \leq r_x \leq r_{x,\max}$, where $r_{x,\min}$ and $r_{x,\max}$ are the minimum and maximum accepted call rate, respectively, amongst the microstates with exactly x calls in progress [9, Ch.3, App.A].

Now let x be the number of calls in progress in the ring. If $x < P$, then any arriving call request can be accepted since every node and link necessarily has room for at least one call. Thus $r_x = Nr_o$ whenever $x < P$. Now note that we can simply set $r_x = 0$ whenever $x \geq P$ to obtain a bound on the performance. Indeed in this case we can see that the bounded Markov chain precisely represents an $M/M/P/P$ queue with arrival rate Nr_o. We have thus found that

$$P_{\text{block}} \leq P_{\text{block,M/M/P/P}}(\frac{Nr_o}{\mu}).$$

C. Bounds for the Reconfigurable Topology, Bidirectional Ring

Here we develop bounds by assuming that calls are never blocked due to a wavelength limitation. With N nodes and P ports per node, this requires $W \geq N \cdot P$ wavelengths be available.

We begin by developing a lower bound similar to that derived for the fixed topology, unidirectional ring. P_{block} is the blocking probability for calls in the ring. By symmetry, P_{block} is the blocking probability for all calls destined to a particular destination node D. P_{block} is lower bounded by the blocking probability of calls destined for D when no other calls are in the system. This simplified system behaves exactly as an $M/M/P/P$ queue with offered load r_o. Specifically,

$$P_{\text{block}} \geq P_{\text{block,M/M/P/P}}(\frac{r_o}{\mu}).$$

This argument is rigorously justified using the same technique of dominated accepted call rates, where $\{r_x\} \leq r_o$ componentwise.

We can derive a similar upper bound. The blocking probability for an arbitrary source-destination pair (S, D) is the probability that an arriving call request (S, D) is blocked at the source, blocked at the destination, or both. Then

$$\begin{aligned} P_{\text{block}} &= \text{Prob(blocked at } S \cup \text{blocked at } D) \\ &\leq \text{Prob(blocked at } S) + \text{Prob(blocked at } D) \\ &\leq 2P_{\text{block,M/M/P/P}}(\frac{r_o}{\mu}). \end{aligned}$$

Rigorous justification of the latter inequality is straightforward but omitted for brevity.

At low blocking probability (e.g., $P_{\text{block}} = 0.01$) these two bounds are very close to one another for $P > 2$ and reasonably close when $P = 2$. This allows us to approximate the large W reconfigurable system as a single M/M/P/P queue, and this behavior is well-understood. We are able to obtain much tighter bounds for $P = 1$, but these are omitted for brevity.

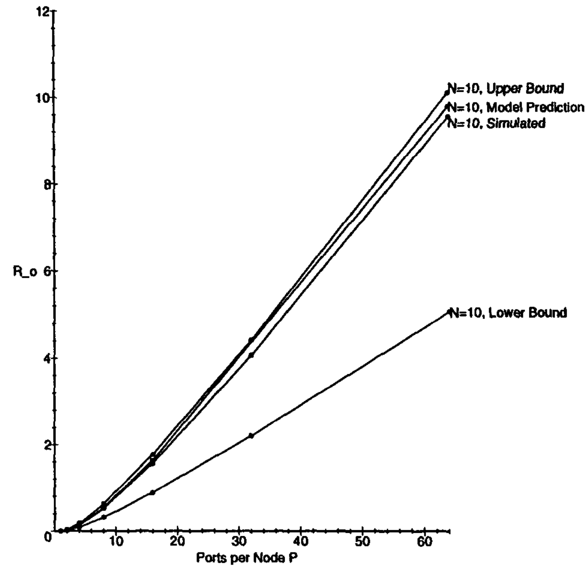

Fig. 4. Fixed Topology, Unidirectional Ring System, N=10, Pb = 0.01

D. Comparing the Bounds

We have been able to develop bounds strictly in terms of an $M/M/P/P$ queue with different offered rates but the same call termination rates. The blocking probability of such an $M/M/P/P$ queue is an increasing function of the offered rate, so this is a particularly nice way to compare the performance of the various topologies. As throughout the paper, performance is measured in terms of offered load at a fixed blocking probability. A potential source of confusion is that a lower bound on blocking probability is equivalent to an upper bound on r_o at a fixed P_{block}.

Recall that we define the gain as the ratio of offerred loads supported by two systems at the same blocking probability. Using the bounds we developed, the gain of a bidirectional reconfigurable system relative to a unidirectional fixed topology system is upper bounded by N. This is independent of the number of ports P and of the blocking probability P_{block}. Because the upper bound for the reconfigurable topology was developed for large W, a gain of N is more optimistic when $W < N \cdot P$. The large wavelength gain is lower bounded by about $\frac{N}{2}$ since the two bounds for the reconfigurable system are close at low blocking probability. As mentioned above, much

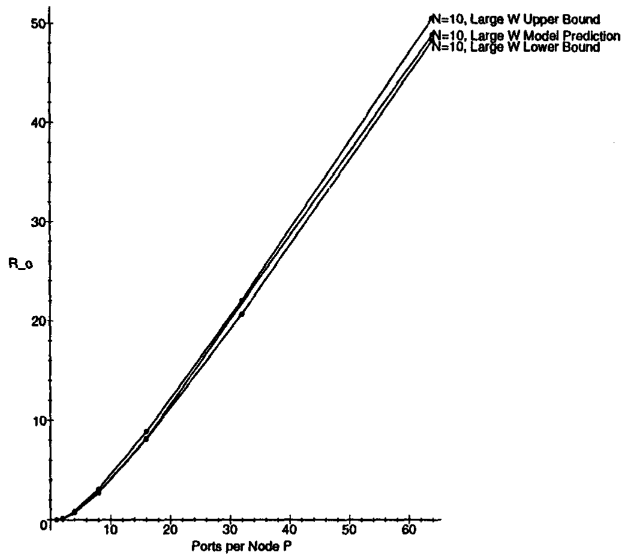

Fig. 5. Unlimited Wavelength Reconfigurable System, N=10, Pb = 0.01

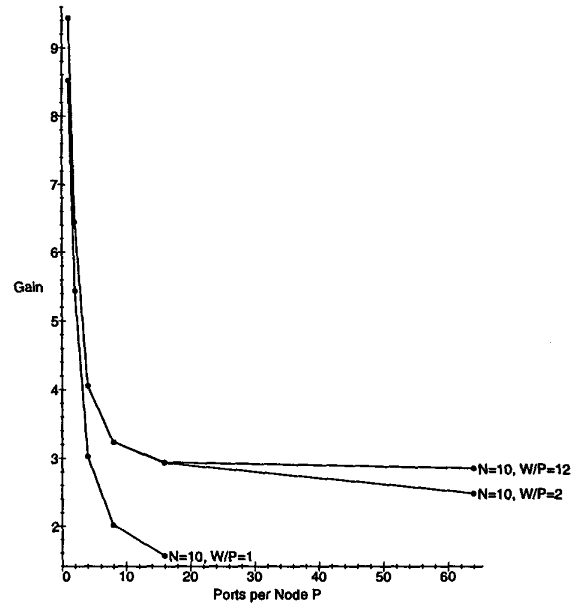

Fig. 6. Gain for N=10, Pb = 0.01

tighter bounds can be developed for the reconfigurable system when $P = 1$, yielding a much better theoretical estimate of the gain.

V. Results and Discussion

We simulated both fixed topology systems and found that the model calculations agree very closely with simulation when $P > 1$ and are slightly optimistic when $P = 1$. In Figure 4 we present four curves for the fixed topology, unidirectional system with $N = 10$ at a fixed blocking probability of 0.01. The curves specify the offerred load supported per node at $P_{block} = 0.01$ as a function of the number of ports per node, P. The four curves correspond to the analytical model prediction from Section III-C, simulation, and the lower and upper bounds from Sections IV-A and IV-B.

In Figure 5 we present three corresponding curves for the large W reconfigurable system at the same fixed blocking probability. The derived bounds are very close for all P at low blocking probability. Since our model prediction lies between the relatively tight upper and lower bounds, we did not simulate the reconfigurable topology.

We find remarkable agreement between our model predictions, simulation, and the analytical bounds. In all three systems, the size of the ring N has little impact on the accuracy of the model or the bounds.

Our main results are plotted in Figures 6 and 7. We plot the gain of the reconfigurable system over the bidirectional system for $N = 10$ and for $N = 100$ nodes at a blocking probability of $P_{block} = 0.01$. The gain is plotted as a function of the number of ports per node P, and the curves are parametrized by the fraction of excess wavelengths available relative to the number of ports, i.e., $\frac{W}{P}$ is constant as P is varied. As expected, the gain increases (to a limit) as $\frac{W}{P}$ increases.

Though not presented graphically, we find that without the wavelength limitation for the reconfigurable system (that is, when $W \geq N \cdot P$), the gain of the reconfigurable system over the fixed topology, unidirectional ring is approximately $\frac{N}{2}$ when $P > 1$. The gain is slightly smaller when $P = 1$. Similarly, the gain over a bidirectional, fixed topology ring is approximately $\frac{N}{4}$. This linear relationship between the gains and N is quite insensitive to N.

There is an intuitive explanation for this gain relationship. In a reconfigurable system with large W, the scarce system resources are the $N \cdot P$ ports. A call requires one full port, an output port at the source and an input port at the destination. In a fixed topology system, since there is a one-to-one relationship between usable wavelengths and node ports, the scarce system resources are the $N \cdot P$ usable wavelengths on the links. Since the average call length in a unidirectional system is $\frac{N}{2}$, a call requires $\frac{N}{2}$ times the resources in this fixed topology. Similarly, since the average call length in a bidirectional system is $\frac{N}{4}$, a call requires $\frac{N}{4}$ times the resources in this fixed topology.

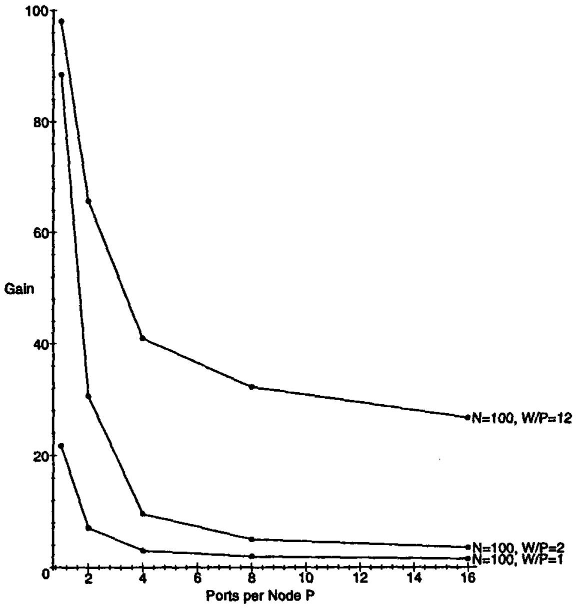

Fig. 7. Gain for N=100, Pb = 0.01

Focusing next on systems where $W = P$, we see that the gain of a reconfigurable system becomes relatively small as the number of ports per node, P, increases. This is a result of increased link utilization due to decreased statistical variation. At $P_{block} = 0.01$, the gain is about 1.5 (a 50% increase in capacity) for large P, for both $N = 10$ and $N = 100$ nodes.

VI. Conclusion

This paper presents a new analytical model for blocking probability in circuit-switched WDM ring networks in order to quantify the benefits of configurability. We verify the accuracy of our model via simulation as well as a number of upper and lower bounds. Our model predicts a significant increase in network capacity, particularly when there are more wavelengths than electronic ports per node. The latter case is likely to become prevalent as WDM systems with more and more wavelengths are being deployed. Hence, we conclude that a configurable WDM network can offer a significant performance benefit.

The work in this paper is preliminary. We believe that our stochastic model can be extended to more general topologies and traffic patterns, though the analysis will be complicated without a regular topological structure.

References

[1] I. Baldine and G. Rouskas, "Dynamic reconfiguration policies for WDM networks," in *Proceedings — IEEE INFOCOM 1999*, New York, N.Y., March 1999.

[2] J.-F. Labourdette and A. Acampora, "Logically rearrangeable lightwave networks," *IEEE Trans. on Comm.*, vol. 39, no. 8, pp. 1223–30, 1996.

[3] A. Narula-Tam and E. Modiano, "Dynamic load balancing for WDM-based packet networks," in *Proceedings — IEEE INFOCOM 2000*, Tel Aviv, Israel, March 2000.

[4] G. Rouskas and M. Ammar, "Dynamic reconfiguration in multi-hop lightwave networks," *Journal of High Speed Networks*, vol. 4, no. 3, pp. 221–38, 1995.

[5] R. Ramaswami and K. Sivarajan, "Design of logical topologies for wavelength-routed optical networks," *IEEE JSAC*, vol. 14, no. 5, pp. 840–51, June 1996.

[6] R. Barry and P. Humblet, "Models of blocking probability in all-optical networks with and without wavelength changers," *IEEE JSAC*, vol. 14, pp. 858–867, 1996.

[7] A. Birman, "Computing approximate blocking probabilities in a class of all-optical networks," *IEEE JSAC*, vol. 14, no. 5, pp. 852–7, 1996.

[8] Y. Zhu, G. Rouskas, and H. Perros, "Blocking in wavelength routing networks, part I: The single path case," in *Proceedings — IEEE INFOCOM 1999*, New York, N.Y., March 1999.

[9] D. Bertsekas and R. Gallager, *Data Networks*, Prentice-Hall, Englewood Cliffs, NJ, second edition, 1992.