A New Analysis for Wavelength Translation in Regular WDM Networks[1]

Jonathan P. Lang and Emmanouel A. Varvarigos
Department of Electrical and Computer Engineering,
University of California, Santa Barbara, CA 93106.
jonathan@paros.ece.ucsb.edu
manos@ece.ucsb.edu

Vishal Sharma[2]
Tellabs Operations Inc.
Tellabs Research Center - Cambridge
1 Kendall Sq., Bldg. 100, Suite 202
Cambridge MA 02139
vsharma@trc.tellabs.com

Abstract

We present a new analysis of wavelength translation in regular, all-optical WDM networks, that is simple, computationally inexpensive, and accurate for both low and high network loads. In a network with k wavelengths per link, we model the output link by an auxiliary $M/M/k/k$ queueing system. We then obtain a closed-form expression for the probability P_{succ} that a session arriving at a node at a random time successfully establishes a connection from its source node to its destination node. Unlike previous analyses, which use the link independence blocking assumption, we account for the dependence between the acquisition of wavelengths on successive links of the session's path. Based on the success probability, we show that the throughput per wavelength increases superlinearly (as expected) as we increase the number of wavelengths per link; however, the extent of this superlinear increase in throughput saturates rather quickly. This suggests some interesting possibilities for network provisioning in an all-optical network. We verify the accuracy of our analysis via simulations for the torus and hypercube networks.

1. Introduction

Wavelength division multiplexing (WDM) exploits the terahertz bandwidth potential of optical fiber by dividing the total available bandwidth into a number of narrower channels. A critical functionality for the scalability and improved performance of multihop WDM networks is *wavelength translation* [2], which is the ability of network nodes to switch data from an incoming wavelength ϕ_i to an outgoing wavelength $\phi_j, j \neq i$. There are three natural classes of wavelength-routing nodes in this context: nodes with *full-wavelength translation* capability (see, for example [4], [10], and [7]), which can exchange an incoming wavelength with any outgoing wavelength, nodes with *limited-wavelength translation* capability (see, for example [23], [22], [16], and [18]), which can switch an incoming wavelength to a subset of the outgoing wavelengths, and nodes with *no-wavelength translation* capability (see, for example [5], [13], [8], [22], [6], and [20]), which can map each incoming wavelength only to the same outgoing wavelength; the so called wavelength-continuity constraint. The requirement of wavelength continuity restricts the routing flexibility and increases the probability of call blocking.

Recently, several researchers have examined wavelength translation in all-optical, circuit-switched networks, with full- or no-wavelength translation. A study of this literature reveals that although this initial work correctly identifies several parameters that affect the performance of wavelength translation (such as path length, number of wavelengths, switch size, network topology, and interference length), and provides useful qualitative insights into network behavior, several difficulties remain. One problem is of accurately accounting for the load correlation between the wavelengths on successive links of a session's path. Kovačević and Acampora [10] provided a model to compute the approximate blocking probability for Poisson input traffic in all-optical networks with and without wavelength translation. As pointed out by them, however, their model is inappropriate for sparse networks because it uses the link independence blocking assumption, which does not consider the dependence between the acquisition of wavelengths on successive links of a session's path. Furthermore, their model requires the iterative solution of a set of Erlang fixed-point equations, which is cumbersome and is known to be computationally expensive. Barry and Humblet [4] presented a new model that takes the link load dependence partially into account, but they assumed that a wavelength is used on a link independently of other wavelengths. Al-

[1] Research supported by DARPA under the MOST Project.

[2] This work was done while the author was a Post-Doctoral Researcher at the University of California, Santa Barbara.

though their simplified model makes good qualitative predictions of network behavior (predicting even some non-obvious behavior observed in simulations [19]), it is unable to predict the behavior of simulations with numerical accuracy. The analysis presented by Birman [7], on the other hand, uses a Markov chain model with state dependent arrival rates. Although more accurate than the previous models, its calculation of blocking probabilities involves modified reduced-load approximations and is computationally intensive. Furthermore, the analysis is tractable only for small networks with at most two or three hops per path and a modest number of wavelengths per link. Yet another difficulty with the previous analyses, which are based on the "blocked calls cleared" model, is that in addition to treating long connections unfairly, they also tend to overestimate the achievable throughput for a given blocking probability. The analysis presented in [18] for limited wavelength translation accounts partially for the link load dependence, and maintains fairness to all connections by retrying blocked sessions at a later time. Although this analysis can be applied to full-wavelength translation, the number of states grows exponentially with the degree of translation (in this case the number of wavelengths), and is impractical for large k.

The analysis that we present for studying wavelength translation in regular, all-optical WDM networks overcomes many of the difficulties highlighted above. We first present a general analysis applicable to any regular topology, and then apply it to study the performance of wavelength translation in the torus and hypercube networks. In our model, sessions or connection requests arrive independently at each node of the network according to a Poisson process with rate ν per unit time, and session destinations are uniformly distributed over the remaining nodes. A circuit between the source node and the destination node is established by sending a setup packet along a path determined at the source. We assume that at each attempt the setup packet randomly chooses a route to the destination from among the allowed shortest-path routes. If the setup packet is successful in establishing a connection, the wavelengths required by the session are reserved for the session duration. Otherwise, the session is randomly assigned a new time at which to try. This is done such that the combined arrival process of new sessions and retrials can be approximated by a Poisson process.

The capacity of each link is divided into k bands, where each band corresponds to a distinct wavelength, and each node has full-wavelength translation capability. We model the k wavelengths on an outgoing link of a node by an auxiliary $M/M/k/k$ queueing system. Using the occupancy distribution of this system, we derive a closed-form expression for the probability P_{succ} of successfully establishing a circuit from a source node to a destination node. To evaluate this probability, we do not use the link independence blocking assumption, but instead account partially for the dependence between the acquisition of successive wavelengths on the path followed by a session. As will be seen in Section 2, our formulation is intuitive, analytically simple, and computationally inexpensive (it avoids, for example, Erlang fixed-point or reduced load approximations), and scales easily for larger network sizes and arbitrary k. We also note that our analysis applies equally well to multi-fiber networks, with no wavelength translation.

Using our analysis we show how, for the mesh and hypercube networks, the extent of improvement in achievable throughput, for a fixed P_{succ}, depends on the number of wavelengths k per link. We find that although the throughput per wavelength increases superlinearly with k, the incremental gain in throughput per wavelength (for a fixed P_{succ}) saturates rather quickly. This implies several interesting alternatives for network provisioning and network expansion in an all-optical network, some of which we discuss in Section 3. For instance, instead of building a network in which every node can translate between W wavelengths, it may be beneficial to build a network in which every node consists of W/k simpler switching elements operating in parallel, each of which translates only over k wavelengths. Similarly, a network designer may start with simple network nodes, with a few parallel channels, and grow them incrementally by adding more channels as network traffic increases.

The organization of the remainder of the paper is as follows. In Section 2, we first present a general analysis for wavelength translation applicable to any regular network, and then apply it to the torus and hypercube networks. In Section 3, we present our results for the success probability P_{succ} obtained from our analysis and compare them to those obtained via simulations, and we discuss our results. In Section 4, we present our conclusions.

2. Analysis For Regular Networks

In this section, we first present a general analysis for full wavelength translation in regular networks, and then apply it to analyze the performance of full wavelength translation in the torus and hypercube networks. The general formulation that we develop provides a method to analyze other regular topologies that have been proposed for building all-optical networks, such as the family of banyan networks, e.g., shufflenet ([1], [15]), wrapped butterfly networks ([18]), and deBruijn networks [19].

Our choice of the mesh and hypercube topologies reflects our interest in analyzing two significantly different types of networks. The torus is a sparse topology with a small (fixed) node degree and diameter that is linear in the mesh dimension, and is a natural choice for building wide-area networks (WANs). The hypercube, on the other hand, is a dense topology, with node degree and diameter that increase logarithmically with the number of nodes, and is an attractive topology for optical interconnects, especially given the recent advances in optical switch design (see for example, [3], [17], [14]).

In our model, external session requests are generated independently at each node of the network according to a Poisson process with rate ν, and their destinations are uniformly distributed over all nodes, except the source node. The holding time, or duration, of a session is exponentially distributed with mean $\overline{X}$, and connections are established by transmitting a setup packet from the source to the destination. Our routing scheme is oblivious, or non-adaptive; that is, the path followed by the session is chosen at the source, and fixed for the entire connection attempt. If the setup packet is successful in establishing the circuit, the wavelengths required by it are reserved for the duration of the session. Otherwise, the session is blocked and is reinserted into the input stream so that the combined process of exogenous arrivals and retrials can be approximated as a Poisson process, and all sessions are eventually served. By contrast, in the call dropping model used in previous analyses, sessions with longer path lengths are dropped with a higher probability (unfairly treating long connections), and the maximum throughput is overstated, especially at higher loads.

In full-wavelength translation, the switching of a new session arriving over an incoming link of a node on a wavelength $\phi_i, i = 1, \ldots, k$ or a new session originating at that node, depends on the availability of a wavelength $\phi_j, j = 1, \ldots, k$ on the desired outgoing link. A session is blocked and scheduled to retry only if all of the k wavelengths on the outgoing link are unavailable. We assume that a setup packet selects the desired outgoing wavelength from among the k wavelengths on the link with equal probability. In the subsequent sections, we define the auxiliary system that we use to model the outgoing link at a node, and obtain the probability P_{succ} of successfully establishing a circuit.

2.1. The Auxiliary System

We first describe a general framework for regular networks, and then present an analysis for the torus and hypercube topologies. We focus on setup packets emitted on outgoing link L, and define the *type* σ of a setup packet according to whether it belongs to a session originating at the node or according to the incoming link upon which it arrives. Consider a network G where a link u at a node v is denoted by (u, v). A $1-1$ function T defined over the set of nodes of G will be called an automorphism if link $(u, v) \in G$ is mapped to $(T(u), T(v)) \in G$. T will be called a fixed automorphism for L if it maps L to itself. We say that two incoming links l_1 and l_2 of node s belong to the same *spatial group* with respect to L if there exists a T such that $T(l_1) = l_2$. We use this mapping to partition the $d-1$ incoming links of a node (except for L) into l, $l \leq d-1$, different groups (or types), so that the links of each group l have the same spatial relationship with respect to outgoing link L. The total number of incoming links of type l is denoted by M_l (see Fig. 1(a)). Originating setup packets that are emitted on link L are defined as being of type $\sigma = 0$, while transit setup packets are defined as being of type $\sigma = 1, \ldots, l$ if the incoming links over which they arrive are of type σ. We also let $\gamma(\sigma)$ denote the rate at which setup packets of type σ are emitted on an outgoing link.

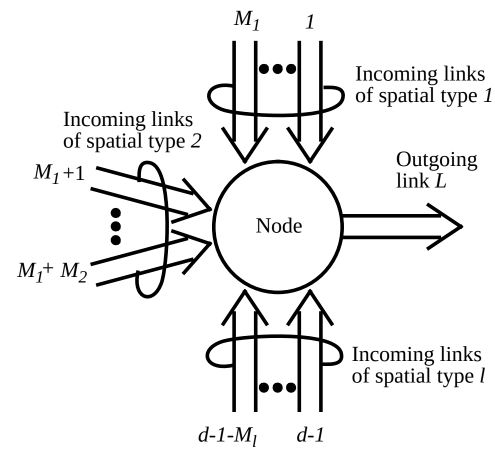

(a)

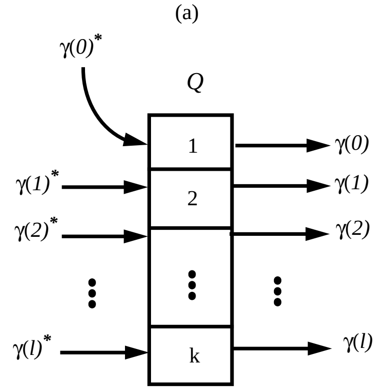

(b)

Figure 1. (a) We illustrate how the incoming links are related to the outgoing link at a node of the network, where each link has k wavelengths. (b) The auxiliary $M/M/k/k$ queueing system $\mathcal{Q}$.

We denote the state of an outgoing link by the vector $\overline{S} = (S_0, S_1, \ldots, S_l)$, where S_0 is the number of originating sessions on the link, and $S_\sigma, \sigma = 1, \ldots, l$, is the number of transit sessions of type σ using the link. Clearly, the set of feasible states of the outgoing link is given by $\mathcal{F} = \{\overline{S} : |\overline{S}| \leq k\}$, where $|\overline{S}| = S_0 + S_1 + \ldots + S_l$ is the cardinality of $\overline{S}$. We let $\pi(\overline{S})$ be the steady-state probability that an outgoing link is in state $\overline{S}$, and we approximate $\pi(\overline{S})$ as the stationary distribution of an auxiliary $M/M/k/k$ queueing system Q, defined as follows (see Fig. 1(b)). Customers of type σ, $\sigma = 0, 1, \ldots, l$, arrive to the system Q according to a Poisson process with rate $\gamma^*(\sigma)$, and ask for a server from among the k identical servers. If all k servers are busy, the customer is dropped, never to appear again. We require that the rate at which customers of type σ are accepted in the auxiliary system Q be the same as the rate $\gamma(\sigma)$ at which setup packets of type σ are emitted on an outgoing link in the actual system. For this to hold, we must have

$$\gamma^*(\sigma) = \frac{\gamma(\sigma)}{1 - \sum_{\overline{S}:|\overline{S}|=k} \pi(\overline{S})}. \tag{1}$$

To calculate the steady-state probabilities $\pi(\overline{S})$, for all feasible states, we write down the global balance equations of the Markov chain that corresponds to the auxiliary system Q :

$$\pi(\overline{S}) = \frac{\sum\limits_{\sigma} \gamma^*(\sigma)\pi(\overline{S} - e_\sigma) + \sum\limits_{\sigma:|\overline{S}|<k} \mu(S_\sigma + 1)\pi(\overline{S} + e_\sigma)}{\sum\limits_{\sigma} \mu S_\sigma + \sum\limits_{\sigma:|\overline{S}|<k} \gamma^*(\sigma)}, \tag{2}$$

where e_σ is a unit vector of dimension l whose σ^{th} component is one, "+" corresponds to componentwise addition, and $\overline{X} = 1/\mu$ is the average service time. Equations (1) and (2) together with the normalization condition $\sum_{\overline{S}\in\mathcal{F}} \pi(\overline{S}) = 1$, can be solved iteratively to obtain the steady-state probabilities $\pi(\overline{S})$ and the rates $\gamma^*(\sigma)$, $\sigma = 0, 1, \ldots, l$. These probabilities are used in the next section to obtain the probability of successfully establishing a circuit.

2.2. The Probability of Successfully Establishing a Circuit

To determine the probability P_{succ} that a session arriving at a random time successfully establishes a circuit from its source node to its destination node, we first find an approximate expression for the probability $P_{succ}(s,d)$ that a session with a given source-destination pair (s,d) is successful on a particular trial. We then average over all source-destination pairs to determine the probability P_{succ} that a session arriving to the network at a random time successfully establishes a circuit.

The path followed by a session with source destination pair (s,d), consists of an originating node followed by a sequence of transit nodes. Thus, the probability α_0 that a wavelength on the outgoing link of the originating node is available is simply the probability that at least one wavelength on that link is idle, and is given by

$$\alpha_0 = \sum_{\overline{S}:|\overline{S}|<k} \pi(\overline{S}). \tag{3}$$

At each transit node, the probability that a wavelength is available on an outgoing link L, given that a wavelength was available on an incoming link $L-1$, and that a transit setup packet proceeding from a wavelength on $L-1$ to a wavelength on L is of type σ can be found to be

$$\alpha_\sigma = \frac{\sum\limits_{S_\sigma=0}^{k-1} \sum\limits_{\overline{S}:|\overline{S}|<k} \pi(\overline{S})(1 - \frac{S_\sigma}{kM_\sigma})}{1 - \sum\limits_{S_\sigma=1}^{k} \sum\limits_{\overline{S}:|\overline{S}|=k} \pi(\overline{S})\frac{S_\sigma}{kM_\sigma}}, \tag{4}$$

where M_σ is the number of input links of type σ. The numerator in Eq. (4) is the sum of all of the state probabilities where at least one wavelength on outgoing link L is available, conditioned on the fact that the wavelength on which the transit setup packet arrives on link $L-1$ is available. The multiplicative factor $(1 - S_\sigma/kM_\sigma)$ is needed because the wavelengths in use on link L cannot be in use by sessions from the particular wavelength on link $L-1$ upon which the transit setup packet arrives. The denominator is one minus the sum of the probabilities of all states where link L is unavailable, conditioned on the fact that the wavelength on which the transit setup packet arrives on link $L-1$ is available.

In writing Eqs. (3) and (4) above, we do *not* assume that the probabilities of acquiring wavelengths on successive links of a session's path are independent. Instead, we account partially for the dependence between the acquisition of successive wavelengths on a session's path, by using the approximation that the probability of acquiring a wavelength on link L depends on the availability of a wavelength on link $L-1$ (in reality this probability depends, even though very weakly, on the availability of a wavelength on every link $1, 2, \ldots, L-1$ preceding link L). As the simulation results presented in Section 3 demonstrate, our approximation is a very good one, however.

The (conditional) probability of successfully establishing a circuit is then given by

$$P_{succ}(s,d) = \alpha_0 \prod_{\sigma=1}^{l} \alpha_\sigma^{h_\sigma(s,d)},$$

where $h_\sigma(s,d)$ is the number of hops on which the transit session is of type σ for a particular source-destination pair (s, d), and the product is taken over all types σ, $\sigma = 1, \ldots, l$, $l \leq d-1$, of transit sessions. For uniformly distributed destinations, the average probability of success for a new arrival P_{succ} can be written as

$$P_{succ} = \frac{1}{N(N-1)} \sum_{(s,d)} P_{succ}(s, d), \qquad (5)$$

where N is the total number of nodes in the network.

In our model, sessions that are not successful in establishing a circuit are blocked and reinserted into the input stream. Since sessions with longer path lengths are blocked and reattempted with higher probability than sessions with shorter paths, the destination distribution of the arrivals (both new and reattempting sessions) may no longer be the same as that of only new arrivals. The success probability of a random session arrival (averaged over all trials, both new and reattempting) $\overline{P}_{succ}$ can be written as

$$\overline{P}_{succ} = \sum_{(s,d)} P_{succ}(s, d) w(s, d),$$

where

$$w(s, d) = \frac{P_{succ}^{-1}(s, d)}{\sum_{(s,d)} P_{succ}^{-1}(s, d)}$$

is a weighting factor to account for the retrials, which reduces to

$$\overline{P}_{succ} = \frac{N(N-1)}{\sum_{(s,d)} P_{succ}^{-1}(s, d)}. \qquad (6)$$

Note that $\overline{P}_{succ}$ is the harmonic mean of the $P_{succ}(s, d)$ over all pairs (s, d), $s \neq d$.

2.3. P_{succ} for the Torus Network

In this section, we derive the expressions for P_{succ} in a torus network. The $p \times p$ torus consists of $N = p^2$ processors arranged along the points of a 2-dimensional space with integer coordinates with p processors along each dimension. In addition to the links connecting nodes that differ in exactly one coordinate, the torus also has wraparound links connecting the first and last node along each dimension. The *routing tag* of a session with source node $s = (s_1, s_2)$ and destination node $d = (d_1, d_2)$, is defined as $t = (t_1, t_2)$, where

$$t_j = \begin{cases} d_j - s_j, & \text{if } |d_j - s_j| \leq \lfloor \frac{p}{2} \rfloor; \\ d_j - s_j - p \cdot sgn(d_j - s_j), & \text{if } |d_j - s_j| > \lfloor \frac{p}{2} \rfloor, \end{cases}$$

for all $j \in 1, 2$, and where $sgn(x)$ is the signum function, which is equal to $+1$ if $x \geq 0$, and equal to -1, otherwise.

To define the feasible states of an outgoing link L, we note that in the torus network, we can partition the incoming links into two spatial types: *bend* types ($\sigma = 1$), where an incoming link is in one dimension and outgoing link L is in a different dimension; and *straight-through* types ($\sigma = 2$), where the incoming link is in the same dimension as outgoing link L. Originating sessions that are emitted on L are defined as being of type $\sigma = 0$. Transit sessions using L that arrive on an incoming link that is a bend type are defined as being of type $\sigma = 1$, and transit sessions using L that arrive on an incoming link that is a straight-through type are defined as being of type $\sigma = 2$. Therefore, for the torus network, the set of feasible states of the outgoing link is $\mathcal{F} = \{\overline{S} : \sum_{i=0}^{2} S_i \leq k\}$.

The rates $\gamma(\sigma)$ can be calculated using Little's Theorem and the symmetry of the torus network to be

$$\gamma(\sigma) = \begin{cases} \frac{\nu \overline{h}(p^2-1)}{4p\lceil \frac{p^2-1}{2} \rceil}, & \sigma = 0 \text{ (originating)}; \\ \frac{\nu \overline{h}(p-1)^2}{4p\lceil \frac{p^2-1}{2} \rceil}, & \sigma = 1 \text{ (bend)}; \\ \frac{\nu \overline{h}\lceil \frac{p-2}{2} \rceil \lfloor \frac{p-2}{2} \rfloor}{4p\lceil \frac{p^2-1}{2} \rceil}, & \sigma = 2 \text{ (straight)}, \end{cases} \qquad (7)$$

where $\overline{h}$ is the mean internodal distance of the torus and can be calculated to be

$$\overline{h} = \begin{cases} p/2, & \text{if } p \text{ is odd;} \\ (p/2)\frac{p^2}{p^2-1}, & \text{if } p \text{ is even.} \end{cases}$$

In the torus network, the success probability for a session arriving to the network at a random time instant depends on its routing tag (t_1, t_2). Using shortest-path routing, the path followed by a session with routing tag (t_1, t_2) will make $I(t_1, t_2) - 1$ bends along the way, and will go straight-through for a total of $|t_1| + |t_2| - I(t_1, t_2)$ steps, where $I(t_1, t_2) \in \{1, 2\}$ is the number of non-zero entries in (t_1, t_2). The probability of successfully establishing a connection is given by

$$P_{succ}(t_1, t_2) = \alpha_0 \cdot \alpha_1^{I(t_1,t_2)-1} \cdot \alpha_2^{|t_1|+|t_2|-I(t_1,t_2)}, \qquad (8)$$

where α_0 is given by Eq. (3), and α_1 and α_2 are given by Eq. (4) with $M_1 = 2$ and $M_2 = 1$, respectively. For uniformly distributed destinations, the average probability of success for a new session can be written as

$$P_{succ} = \begin{cases} \frac{\alpha_0\left[\left(1 + 2\alpha_1\left(\frac{1-\alpha_2^{\lfloor p/2 \rfloor}}{1-\alpha_2}\right)\right)^2 - 1\right]}{\alpha_1(p^2-1)}, & \text{p odd} \\ \frac{\alpha_0\left[\left(1 + \alpha_1\left(\frac{1-\alpha_2^{\lfloor p/2 \rfloor}}{1-\alpha_2}\right)\right]}{\alpha_1(p^2-1))} \\ + \frac{\alpha_0\left[\alpha_1\left(\frac{1-\alpha_2^{\lfloor p/2 \rfloor - 1}}{1-\alpha_2}\right)\right)^2 - 1\right]}{\alpha_1(p^2-1)}. & \text{p even} \end{cases} \qquad (9)$$

The average probability of success for a random arrival (both new and reattempting) can be written as

$$\overline{P}_{succ} = \begin{cases} \beta\big[\big(1 + \frac{2}{\alpha_1}\big(\frac{1-\alpha_2^{-\lfloor p/2 \rfloor}}{1-\alpha_2^{-1}}\big)\big)^2 - 1\big], & \text{p odd} \\ \beta\big[\big(1 + \frac{1}{\alpha_1}\big(\frac{1-\alpha_2^{-\lfloor p/2 \rfloor}}{1-\alpha_2^{-1}}\big) \\ \quad + \frac{1}{\alpha_1}\big(\frac{1-\alpha_2^{-\lfloor p/2 \rfloor + 1}}{1-\alpha_2^{-1}}\big)\big)^2 - 1\big], & \text{p even} \end{cases} \quad (10)$$

where $\beta = \frac{\alpha_0(p^2-1)}{\alpha_1}$.

2.4. P_{succ} for the Hypercube Network

In this section, we give the expressions for P_{succ} in a hypercube network with crossbar switches at the nodes. Each node of a 2^d-node hypercube network can be represented by a binary string $(x_1, x_2, \ldots, x_d)$, where two nodes are connected via a bidirectional link if their binary representations differ in one bit. The routing tag of a session with source s and destination d is just the bitwise XOR of the binary representations of the source and destination. To define the feasible states of an outgoing link L in the 2^d-node hypercube network with crossbar switches, we observe that due to symmetry, all of the incoming links are of the same spatial type. Originating sessions that are emitted on L are defined as being of type $\sigma = 0$, while transit sessions using L are defined as being of type $\sigma = 1$. Therefore, for the hypercube, the set of feasible states of the outgoing link is $\mathcal{F} = \{\overline{S} : S_0 + S_1 \leq k\}$.

The rates $\gamma(\sigma)$ at which setup packets of type σ are emitted on a link can now be found to be

$$\gamma(\sigma) = \begin{cases} \frac{\nu}{d}, & \sigma = 0 \text{ (originating);} \\ \frac{\nu[(d-2)2^{d-1}+1]}{d(2^d-1)}, & \sigma = 1 \text{ (transit).} \end{cases} \quad (11)$$

Assuming that a session at its source is at a distance of i hops from its destination, the probability of successfully establishing a connection is given by

$$P_{succ}(n) = \alpha_0 \cdot \alpha_1^{i-1}, \quad (12)$$

where α_0 is given by Eq. (3) and α_1 is given by Eq. (4) with $M_1 = d - 1$.

For uniformly distributed destinations, the average probability of success for a new arrival can be written as

$$\begin{aligned} P_{succ} &= \frac{1}{2^d-1}\sum_{i=1}^{d}\binom{d}{i}P_{succ}(i) \\ &= \frac{\alpha_0}{\alpha_1(2^d-1)}\Big[(1+\alpha_1)^d - 1\Big]. \end{aligned} \quad (13)$$

The average probability of success for a random arrival (both new and reattempting) can be written as

$$\begin{aligned} \overline{P}_{succ} &= (2^d-1)\Big(\sum_{i=1}^{d}\binom{d}{i}\frac{1}{P_{succ}(i)}\Big)^{-1} \\ &= \frac{\alpha_0(2^d-1)}{\alpha_1}\Big[\Big(1+\frac{1}{\alpha_1}\Big)^d - 1\Big]^{-1}. \end{aligned} \quad (14)$$

3. Analytical and Simulation Results

In this section, we present our analytical and simulation results for full wavelength translation in the torus and hypercube networks.

We first compare the probability of success for a new arrival P_{succ} and for a random arrival $\overline{P}_{succ}$ predicted by our analysis with that obtained via simulations for the torus and the hypercube networks, for the number of wavelengths k ranging from 1 to 4 (see Fig. 2 and Fig. 3, respectively). Each simulation point was obtained by averaging over 2×10^6 successes, after discarding initial transients. We observe that there is close agreement between the simulations and the analytically predicted values over the entire range of applicable input rates, which is a significant improvement over previous analyses. Despite its accuracy, our analysis is intuitive and computationally inexpensive, and considerably simpler than the analyses available in the literature, which allows it to scale easily for large k.

We define $\lambda(P_{succ}, k)$ to be the throughput per node per wavelength of a full-wavelength translation system with k wavelengths, when the probability of success is equal to P_{succ}. Similarly, we define $P_{succ}(\lambda, k)$ to be the probability of success of a full-wavelength translation system with k wavelengths, when the arrival rate per node per wavelength is equal to $\lambda = \nu/k$. To quantify the performance of full-wavelength translation for varying k, we also define the *incremental per-wavelength throughput gain* $\Delta\lambda(k_1, k_2)$ of a full-wavelength translation system with k_2 wavelengths, over a system with k_1 wavelengths, for a given P_{succ}, to be

$$\Delta\lambda(k_1, k_2) = \frac{\lambda(P_{succ}, k_2) - \lambda(P_{succ}, k_1)}{\lambda(P_{succ}, k_1)} \times 100\,\%. \quad (15)$$

We also define the *incremental probability of success gain* $\Delta P_{succ}(k_1, k_2)$ of a full-wavelength translation system with k_2 wavelengths, over a system with k_1 wavelengths, for a given λ, to be

$$\Delta P_{succ}(k_1, k_2) = \frac{P_{succ}(\lambda, k_2) - P_{succ}(\lambda, k_1)}{P_{succ}(\lambda, k_1)} \times 100\,\%. \quad (16)$$

The throughput and probability of success gains measure the degree of improvement that a full-wavelength translation system with k_2 wavelengths provides over a similar system with k_1 wavelengths.

four wavelengths per link achieves a 114% gain in throughput per wavelength over a system with two wavelengths per link. (Similar results also hold for the incremental success gains in both topologies.)

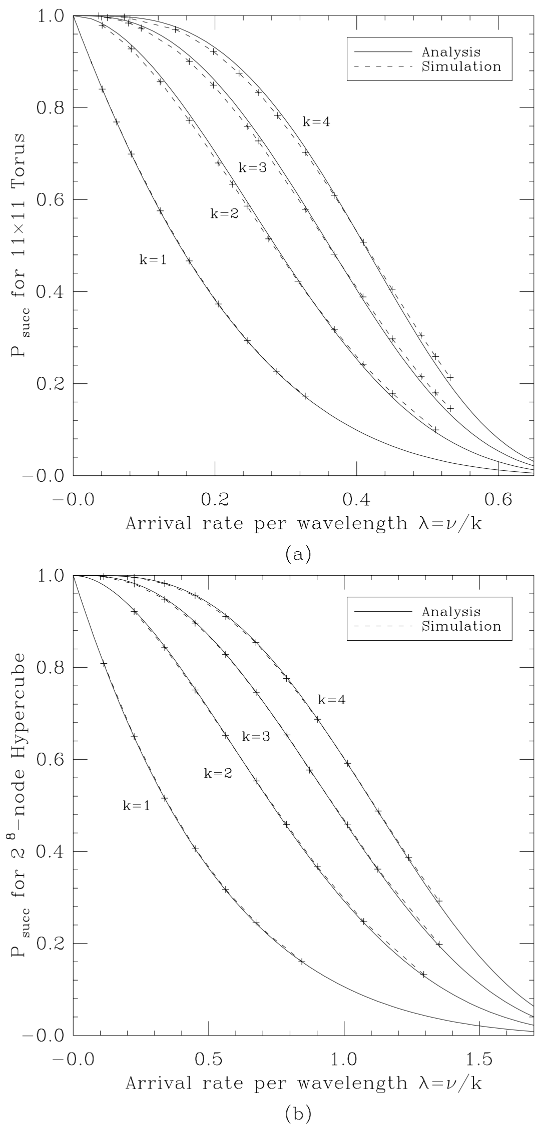

Figure 2. Analytical and simulation results for P_{succ} versus the arrival rate per wavelength ν/k, for (a) an 11×11 torus, and (b) 2^8-node hypercube network.

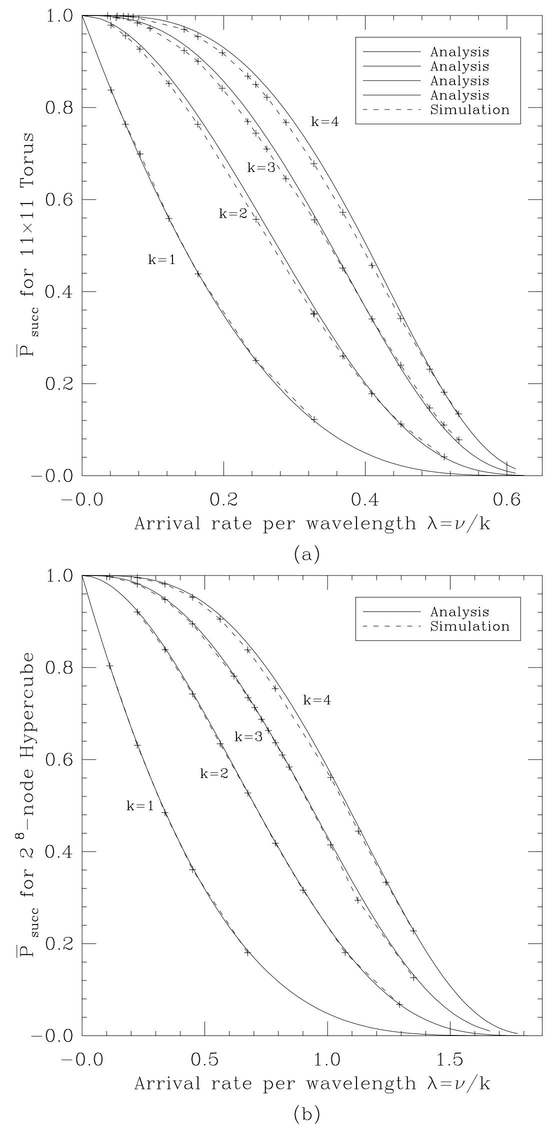

Figure 3. Analytical and simulation results for $\overline{P}_{succ}$ versus the arrival rate per wavelength ν/k, for (a) an 11×11 torus, and (b) 2^8-node hypercube network.

In Fig. 4, we illustrate the analytically predicted probability of success P_{succ} versus the arrival rate per wavelength ν/k, for k ranging from 1 to 16, for both the torus and hypercube networks. In Table 1, we show the per-wavelength incremental throughput gains for an 11×11 torus network and for a 2^8-node hypercube network, for two values of P_{succ}. For example, in the torus network, for $P_{succ} = 0.9$, using a full-wavelength translation system with two wavelengths per link achieves a 320% gain in throughput per wavelength over a system with one wavelength per link (i.e, with no wavelength translation). Likewise, a system with

As is evident from Fig. 4 and Table 1, for a given P_{succ}, the throughput per wavelength increases with increasing k (therefore, the network throughput increases superlinearly with k). The linear increase in throughput is because of the increase in capacity, while the superlinear increase is due to the greater flexibility in establishing a circuit when a larger number of wavelengths is available. The incremental gain in achievable throughput per wavelength $\Delta\lambda(k_1, k_2)$, however, decreases rapidly with increasing k. This result is in agreement with the results presented in [9] and [18], where the authors had found that the incremental gain in throughput per unit of capacity obtainable by using links with larger capacity (alternately, links with a greater number of wavelengths) diminishes as the capacity k per link (alternately, the number of wavelengths k per link) increases.

Torus

P_{succ}	$\Delta\lambda(1,2)$	$\Delta\lambda(2,4)$	$\Delta\lambda(4,8)$	$\Delta\lambda(8,16)$
0.9	320%	114.3%	55.6%	31.43%
0.5	86.7%	46.4%	25.6%	14.6%

Hypercube

P_{succ}	$\Delta\lambda(1,2)$	$\Delta\lambda(2,4)$	$\Delta\lambda(4,8)$	$\Delta\lambda(8,16)$
0.9	432%	125.6%	60%	33%
0.5	107%	51.7%	27.3 %	17%

Table 1. Incremental per-wavelength throughput gains for a 11×11 torus and a 2^8-node hypercube.

The above observation leads to some interesting design options when building an all-optical network. It suggests, for example, that when W wavelengths per link are available, simply building a network node with full wavelength translation capability over W wavelengths may not be the most efficient option. This is because it may be possible to build a network node with W/k simpler switching elements operating in parallel, each switching between a non-intersecting subset of k wavelengths, that achieves performance comparable to that of the W-wavelength system at a much lower cost.

To illustrate this, we consider the general wavelength convertible switch architecture presented in Fig. 2 of [11], which consists of an optical switch followed by a system of wavelength converters. The optical switch may be implemented as a non-blocking crossbar switch or as a multistage switch. Alternately, the entire switch architecture may be implemented as a single system that switches in both the space and wavelength domains, using either a rearrangeably non-blocking Twisted Bene's network, as proposed by Yoo and Bala [12], or the more general architectures proposed by Thompson and Hunter [21]. Assuming the implementation in [12] (to be specific) with W/k parallel groups of wavelengths, where each wavelength is able to translate only to the k wavelengths within the same group, the component cost (in terms of elementary 2×2 switches) is $\frac{W}{k} kd \log(kd) - \frac{kd}{2} = Wd \log(kd) - \frac{kd}{2}$, and the increase in blocking probability becomes negligeable as k increases beyond some point. This suggests that a network designer may initially choose to build the network with nodes that have a small number of parallel channels (groups), with k wavelengths per channel, and may gradually expand the nodes as network traffic grows, by adding more parallel channels.

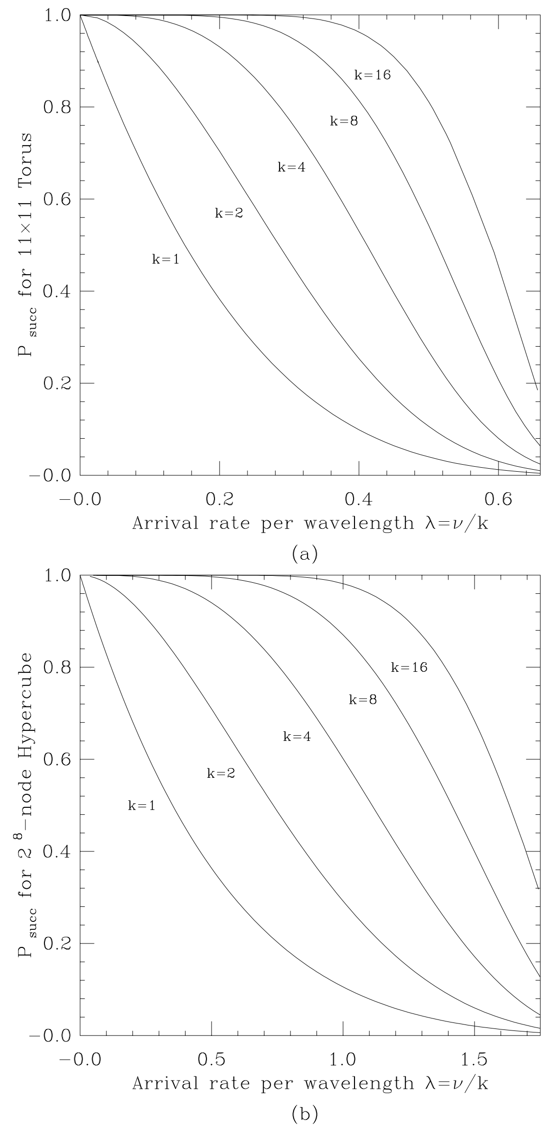

Figure 4. The probability of success P_{succ} obtained from our analysis for (a) an 11×11 torus and (b) 2^8-node hypercube network, for k varying from 1 to 16.

4. Conclusions

We presented a new general analysis for wavelength translation in all-optical regular networks that is intuitive, simple, and computationally inexpensive. Our analysis does not use the link independence blocking assumption, and is more accurate than previous analyses over a wider range of network loads. We verified our analysis for the hypercube and torus topologies, and found that although the throughput per wavelength increases with an increase in the number of wavelengths, this increase saturates quickly. This observation leads to some interesting possibilities for provisioning an all-optical network from a performance-cost perspective. One direction for future work could be to obtain functions that can appropriately measure the cost and the control complexity of the switch, so that different implementations of all-optical networks may be compared in a fair manner.

5. References

[1] Acampora, A., "The multichannel multihop lightwave network," *Proc. IEEE GLOBECOM*, Tokyo, Japan, vol. 3, pp. 1459–67, Nov. 15–18, 1987.

[2] Acampora, A., "The scalable lightwave network," *IEEE Comm. Mag.*, vol. 32, no. 12, pp. 36–42, Dec. 1994.

[3] Ali, M. F., Guizani, M., "A new design methodology for optical hypercube interconnection network," in *Proc. First Int'l Conf. on Algorithms and Architectures for Parallel Processing*, Australia, vol. 2, pp. 920–922, Apr. 1995.

[4] Barry, R. A., and Humblet, P., "Models of blocking probability in all-optical networks with and without wavelength changers," *IEEE J. Select. Areas Comm.*, vol. 14, no. 5, pp. 858–867, June 1996.

[5] Barry, R. A., *Wavelength routing for all-optical networks*, Ph.D. dissertation, Dept. of Electrical Engineering and Computer Science, MIT, Cambridge, Mass., 1993.

[6] Baroni, S., Bayvel, P., Midwinter, J. E., "Wavelength requirements in dense wavelength-routed optical transport networks with variable physical connectivity," *Electron. Lett.*, vol. 32, no. 6, pp. 575–576, Mar. 1996.

[7] Birman, A., "Computing approximate blocking probabilities for a class of all-optical networks," *IEEE J. Select. Areas Commun.*, vol. 14, no. 5, pp. 852–857, June 1996.

[8] Kaminow, I. P., Doerr, C. R., Dragone, C., et al, "A wideband all-optical WDM network," *IEEE J. Select. Areas Commun.*, vol. 14, no. 4, pp. 780–89, June 1996.

[9] Koch, R. "Increasing the size of a network by a constant factor can increase performance by more than a constant factor," *Twenty-ninth Annual Symp. Foundations of Computer Science*, pp. 221–230, Oct. 1988.

[10] Kovačević, M., and Acampora, A., "Benefits of wavelength translation in all-optical clear-channel networks," *IEEE J. Select. Areas Commun.*, vol. 14, no. 4, pp. 868–880, June 1996.

[11] Lee, K. C., and Li, V. O. K., "A wavelength convertible optical network," *IEEE /OSA J. Lightwave Tech.*, vol. 11, no. 5/6, pp. 962–970, May/June 1993.

[12] Yoo, S. J. B., and Bala, K., "Parametric wavelength conversion and cross-connect architecture for transparent all-optical networks," in *Proc. SPIE Conf. on All-optical Commun.*, Boston, Mass., Nov. 1996.

[13] Nagatsu, N., Hamazumi, Y., Sato, K., "Number of wavelengths required for constructing large-scale optical path networks," *Electronics and Communications in Japan*, Part 1, vol. 78, no. 9, pp. 1–10, Sept. 1995.

[14] Nishimura, S., Inoue, H., Hanatani, S., Matsuoka, H., et al, "Optical interconnections for the massively parallel computer," *IEEE Photon. Technol. Lett.*, vol. 9, no. 7, pp. 1,029–1,031, July 1997.

[15] Pankaj, R. K., and Gallager, R. G., "Wavelength requirements of all-optical networks," *IEEE ACM Trans. Networking*, vol. 3, no. 3, pp. 269–279, June 1995.

[16] Ramaswami, R., and Sasaki, G. H., "Multiwavelength optical networks with limited wavelength conversion," *Proc. INFOCOM'97*, Kobe, Japan, April 1997.

[17] Shionoya, S., "A study on multiwave interconnection networks," *Record of Elec. and Commun. Eng. Conversazione*, Tohoku Univ., vol. 65, pp. 265–266, Dec. 1996.

[18] Sharma, V., and Varvarigos, E. A., "An analysis of limited wavelength translation in regular all-optical WDM networks," submitted *IEEE Trans. on Comm.* (Short version appeared in Proc. IEEE INFOCOM'98, vol. 2, pp. 893–901.)

[19] Sivarajan, K. N., and Ramaswami, R., "Lightwave networks based on de Bruijn graphs," *IEEE/ACM Trans. Networking*, vol. 2, no. 1, pp. 70–79, Feb. 1994.

[20] Subramanium, S, and Barry, R. A., "Wavelength assignment in a fixed routing network," *Proc. ICC'97*, Montreal, Que., Canada, vol. 1, pp. 406-10, June 8–12, 1997.

[21] Thompson, R. A., and Hunter, D. K, "Elementary photonic switching modules in three divisions," *IEEE J. Select. Areas Comm.*, vol. 14, no. 2, pp. 362–373, Feb. 1996.

[22] Wauters, N., Demeester, P., "Wavelength translation in optical mult-wavelength multi-fibre transport networks," *Int'l Jour. of Optoelectronics*, vol. 11, no. 1, pp. 53–70, Jan.–Feb. 1997.

[23] Yates, J., Lacey, J., Everitt, D., and Summerfield, M., "Limited-range wavelength translation in all-optical networks," Proc. IEEE INFOCOM'96, San Francisco, CA, March 26-28, 1996, vol. 3, pp. 954-961.