IEEE
COMMUNICATIONS
SURVEYS

WAVELENGTH CONVERTERS IN DYNAMICALLY-RECONFIGURABLE WDM NETWORKS

JENNIFER M. YATES AND MICHAEL P. RUMSEWICZ,
THE ROYAL MELBOURNE INSTITUTE OF TECHNOLOGY
JONATHAN P. R. LACEY, HEWLETT-PACKARD LABORATORIES

ABSTRACT

In simple wavelength-division multiplexed (WDM) networks, a connection must be established along a route using a common wavelength on all of the links along the route. This constraint may be removed by the introduction of wavelength converters, which are devices which take the data modulated on an input wavelength and transfer it to a different output wavelength. Wavelength converters thus improve network blocking performance. However, the introduction of wavelength converters into WDM cross-connects increases the hardware cost and complexity. Thus, it is important to establish precisely what advantages wavelength converters offer WDM networks. There has been considerable interest in the literature in the performance improvements offered by the introduction of wavelength converters into dynamically-reconfigurable WDM networks. This article provides a review of the conclusions drawn from these investigations. The performance improvements offered by wavelength converters depend on a number of factors, including network topology and size, the number of wavelengths, and the routing and wavelength assignment algorithms used. We discuss these factors here. However, it has been shown that wavelength converters offer only modest performance improvements in many networks. We also consider networks with limited wavelength conversion, in which the set of allowable conversions at a network node is constrained by having limited numbers of wavelength converters, or by using non-ideal wavelength converters. Limited wavelength conversion has been shown to provide performance which is often close to that achieved with ideal wavelength conversion in networks with tunable transmitters and receivers.

Wavelength-Division Multiplexing (WDM) has emerged as a promising technique for opening up the Terahertz transmission bandwidth of single-mode optical fiber [1]. In WDM transmission, each data channel is modulated onto an optical carrier with a unique wavelength (or optical frequency). The optical carriers are then combined and transmitted on a single fiber. In this way, WDM not only enables the use of the enormous fiber bandwidth, but provides channels whose individual bandwidths are within the capacity of conventional electronic information-processing devices.

WDM technology is being extensively deployed on point-to-point links within transport networks in the United States, while WDM point-to-point links are soon to be deployed within Europe [2]. However, WDM promises advantages for switching and routing as well as for transmission. Optical cross-connects are currently being developed which can switch an entire wavelength from an input fiber to an output fiber so that large-bandwidth circuits can be routed through the network according to wavelength. High-speed, fixed-bandwidth, end-to-end connections called lightpaths can then be established between different nodes. Networks which use optical cross-connects to route lightpaths through the network are referred to as wavelength-routing networks. Wavelength-routing optical core networks are expected to evolve from the existing separate WDM transmission systems to form optical layers in future transport networks. These optical layers will provide switching, routing, and (potentially) restoration on a per-wavelength basis.

Future transport networks are expected to incorporate both electronic and optical switching, as depicted in Fig. 1. A variety of different user applications may be combined within the electronic switching layer and then transported over high-bandwidth optical “pipes” in the wavelength-routing optical network layer. Alternatively, high-bandwidth connections may be established via direct access to the optical network.

A simple wavelength-routing network with example connections is illustrated in Fig. 2. The nodes are interconnected by optical fibers, on which WDM signals are transmitted. The nodes are known as wavelength routers, and have the ability to route an incoming signal to an outgoing port according to the signal’s input port and wavelength.

Wavelength-routing networks employ “spatial reuse” of wavelengths, by allowing the same wavelength to be used by multiple lightpaths in the same network, providing that none of these lightpaths share a common link. This allows scalability of wavelength-routing networks, although this scalability may be limited in non-reconfigurable networks [3].

WAVELENGTH CONVERSION

In simple wavelength-routing networks, a lightpath between two nodes along a particular route must use a single wavelength on all hops, or links, within the route. This requirement is referred to as the wavelength continuity constraint. For instance, consider the two-link route shown in Fig. 3. Imagine that a connection is to be established between nodes 1 and 3 along a route which passes through a cross-connect at node 2. This connection can only be established if the same wavelength is available on both links. If only wavelength λ_1 is available on link 1, and only wavelength λ_2 is available on link 2, then the connection cannot be established.

The restriction imposed by the wavelength continuity constraint can be avoided by the use of wavelength conversion (also referred to as wavelength translation or wavelength changing). A wavelength converter is a device which takes as its input a data channel modulated onto an optical carrier with a wavelength λ_{in}, and produces at its output the same data channel modulated onto an optical carrier with a different wavelength λ_{out}. If wavelength converters are included in the cross-connects in WDM networks, connections can be established without the need to find an unoccupied wavelength which is the same on all the hops making up the route. For instance, if a wavelength converter was available at node 2 in Fig. 3, the connection could be established using wavelength λ_4 on link 1, and wavelength λ_3 on link 2. This means that networks with wavelength converters are equivalent to traditional circuit-switched networks. Wavelength converters thus result in improvements in network performance.

The original dream of many optical network pioneers was to build optically-transparent or all-optical networks [4], in which no optical to electronic conversions were performed between each source and destination.

Wavelength converters used in these networks must be all-optical wavelength converters. However, the extent to which all-optical, transparent transport networks will be used in the future still remains to be determined. Optical nonlinearities, chromatic dispersion, amplifier spontaneous emission, and other factors together limit the scalability of a transparent WDM network [5]. It appears that "3R" regeneration (reamplification, reshaping, and retiming) is required to build large, scalable WDM networks, and this is currently performed using optoelectronic (OE) regenerators, described below. Also, full 3R OE regeneration allows network operators to monitor signal quality by measuring bit-error rates. In situations where the signal quality is found to be degraded at the endpoint of a connection, regular bit-error-rate measurements allow the network operator to determine in which part of the network the degradation occurred [5].

Optoelectronic regenerators convert the modulated optical carrier into a baseband electrical signal using a photodetector, regenerate and amplify this electrical signal, and then use it to remodulate an output laser with the desired frequency (wavelength). If the frequency to which the output laser is tuned is different to the frequency of the input signal, then wavelength conversion is also performed. Thus, if optoelectronic regenerators are used in a WDM network, then wavelength conversion can be implemented without having to also introduce all-optical wavelength converters.

This argument suggests that wavelength conversion could be introduced into some networks without significant additional cost. However, the introduction of wavelength converters — whether they are all-optical or opto-electronic — is expected to significantly complicate the design of an optical cross-connect, because cross-connects without wavelength converters can be implemented by independently switching connections at different wavelengths. This is illustrated in Fig.

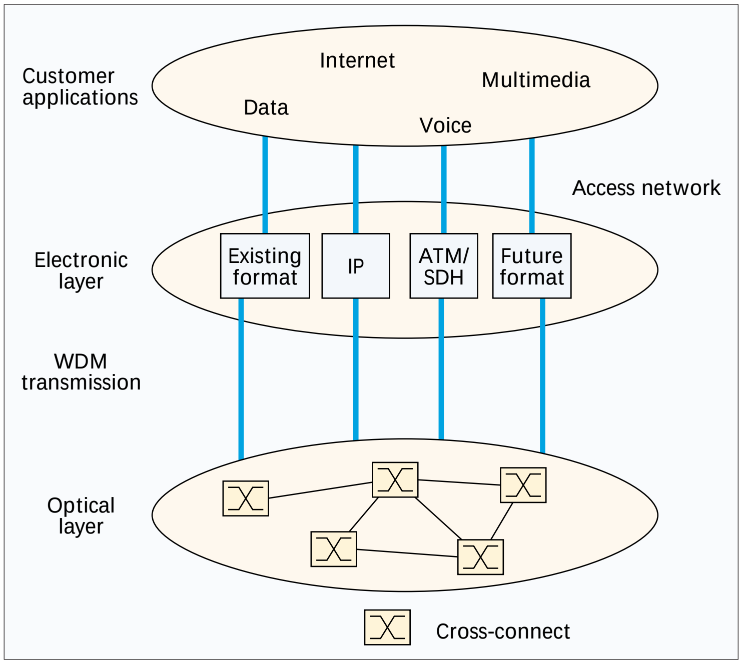

Figure 1. *Potential layered architecture for future telecommunications network.*

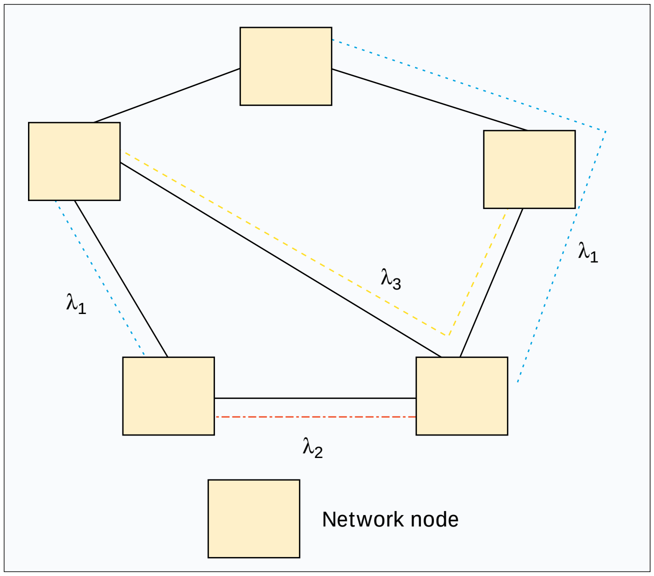

Figure 2. *An example wavelength-routing network.*

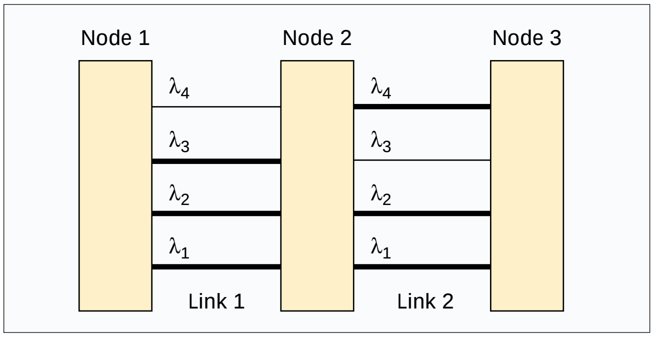

Figure 3. *Blocking in WDM networks. Bold lines represent wavelengths that are not available.*

4, where separate optical space switches are used for each wavelength. If there are M input and M output fibers, with W wavelengths on each fiber, then W separate $M \times M$ space switches are required to implement a cross-connect without wavelength converters. In contrast, a single $MW \times MW$ space switch is required to implement the cross-connect with wavelength converters shown in Fig. 5.

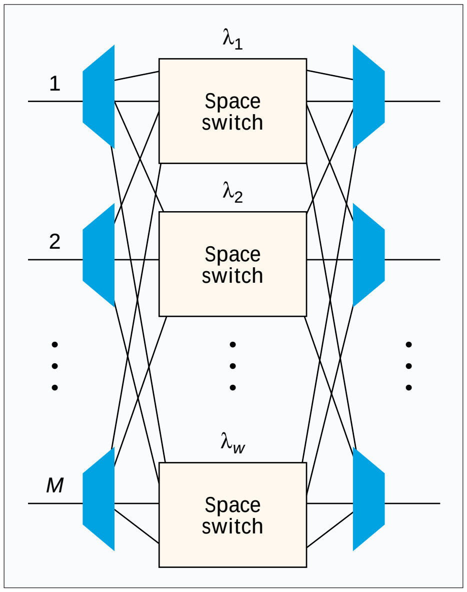

Figure 4. *Cross-connect without wavelength converters. There are* M *input and* M *output fibers, with* W *wavelength on each fiber, and the cross-connect uses* W *separate* M x M *space switches.*

We can conclude that wavelength converters are expected to increase the cost and complexity of WDM networks. It is thus important to establish precisely what advantages wavelength converters offer optical networks. Wavelength converters are expected to be of great importance in providing interoperability in multi-vendor environments [6], and may significantly simplify network management [7]. Wavelength converters may also improve network performance, allowing more efficient use of network resources. In this article we examine in detail the potential network performance improvements achieved through the introduction of wavelength converters.

STATIC AND DYNAMIC NETWORKS

It is difficult to predict the bandwidth requirements and statistical properties of the traffic that will be carried by future wavelength-routing WDM networks, but consideration of these factors is important for network design and analysis. For example, the bandwidth provided by a single lightpath is expected to far exceed the requirements of most individual calls. Thus, lightpaths may be used to carry a single high-bandwidth call established between two individual users, or a stream of traffic from many different users electronically multiplexed in the time domain. The users of a WDM transport network could thus be electronic switching equipment, such as SONET/SDH cross-connects, ATM switches or IP routers, or individual workstations or video servers [8].

The extent to which lightpaths are dynamically established due to time-varying fluctuations in the traffic demand is an important issue in network design and analysis, and is somewhat dependent on the bandwidth requirements just discussed. If the bandwidth of individual calls is significantly less than that of a single lightpath, then large numbers of calls may be time-division multiplexed onto each lightpath. Depending on the rate at which overall traffic demand varies, the lightpaths may remain relatively fixed over time, with only occasional changes in the lightpath allocations for restoration (fault recovery) or to follow slowly changing mean traffic requirements through the course of a day. Alternatively, in very dynamic environments, such as in the transportation of multiplexed Internet traffic, or if the bandwidths of individual calls are large in comparison with the lightpath bandwidth, lightpath requirements may vary considerably over time and lightpaths may thus be established on demand. Thus, the traffic offered to a wavelength-routing network may be either effectively static, with lightpath requirements fixed over time, or dynamic, with lightpaths established on demand. In a realistic network we may expect some combination of both of these cases — some lightpaths being established semi-permanently, while others are established and torn down as calls are offered and depart from the network.

Dynamic WDM networks may perform online or offline routing. In offline routing, all of the lightpath requests to be routed are known in advance, and the optimum routes and wavelength assignments are determined. If a new lightpath is to be included in an existing set of lightpaths, the wavelength assignments for all lightpaths may need to be recomputed and existing lightpaths rearranged. Optimized wavelength assignment can thus be achieved. However, it may not be practical to rearrange existing established lightpaths. In contrast, networks using online routing establish new lightpaths without changing the wavelength allocations of existing lightpaths.

In this article we consider the case in which the lightpaths are dynamically established using online routing, reflecting the arrivals and departures of large bandwidth connections. We refer to these networks offered dynamic traffic as dynamically-reconfigurable WDM networks with online routing. The performance benefits of wavelength converters in these networks have been a topic of intense interest within the literature. Similarly, the potential benefits of wavelength converters in reducing the number of required wavelengths in static networks, or equivalently in networks with offline routing, has been examined in a number of publications, including [9, 10, 13]. In the next section of this article we examine the factors governing the performance improvements offered by wavelength converters in dynamically-reconfigurable WDM networks with online routing. In this discussion we assume that the wavelength converters are ideal, in that any wavelength can be converted to any other wavelength, and that there is a full set of wavelength converters in every cross-connect in the network. We then continue in the section that follows to examine the performance improvements offered by non-ideal wavelength conversion, in which not every wavelength can be con-

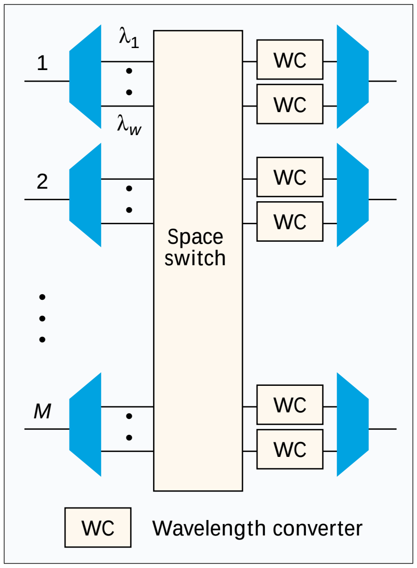

Figure 5. *Cross-connect with wavelength converters. There are* M *input and* M *output fibers, with* W *wavelengths on each fiber, and the cross-connect uses a single* MW x MW *space switch.*

verted to every other wavelength in every cross-connect. This may be the result of having limited numbers of wavelength converters in the WDM networks, or by using devices in which noise or other device limitations restrict the set of conversions which may be performed. We then examine other potential benefits of wavelength converters in WDM networks. Finally, we present our views on important future research topics.

WAVELENGTH CONVERTERS IN WDM NETWORKS

The potential benefits of wavelength converters in dynamically-reconfigurable WDM networks with online routing have been studied in [11–34]. The performance improvements offered by wavelength converters have been investigated both in terms of the improvements attained in network blocking probabilities for a fixed offered load, and in the increase in the offered load which can be supported for a fixed blocking probability. This second measure is of particular interest to telecommunications providers as it represents the possible increase in revenue for a given quality of service. Results presented in the literature have shown that the performance improvements obtained by introducing wavelength converters into WDM networks depend on a number of factors, including:

- Network topology and size [15–19]
- The number of wavelengths used on each link [13, 15, 16, 19, 20]
- The number of fibers on each link [13, 25, 26]
- The routing and wavelength assignment scheme used [16, 18, 22, 26, 28]
- The traffic arrival process [21, 32, 33]

We discuss these factors in this article.

It is important to note that in the work discussed below, when considering networks without wavelength converters (wavelength-continuous networks), tunable transmitters and receivers are imagined to be used, so that a connection may be established on any available wavelength. Unacceptable performance is generally obtained if fixed-frequency transmitters and receivers are used in wavelength-continuous networks. This contrasts with networks with ideal wavelength converters, in which the tunability of the transmitters and receivers has no effect on network performance.

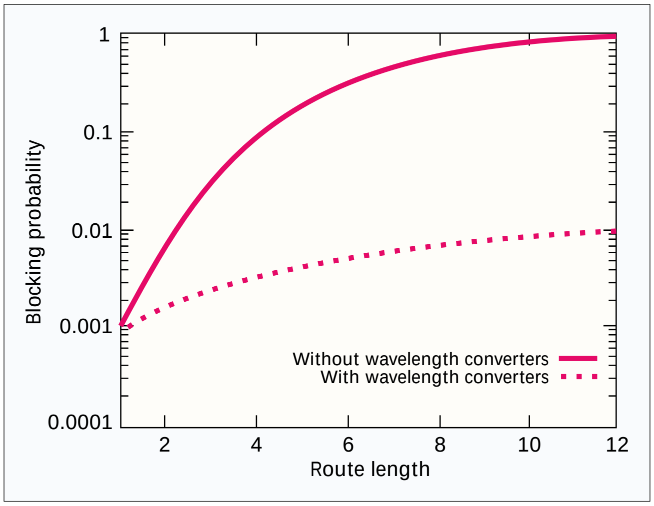

Figure 6. *Blocking probability vs. route length [15].*

TOPOLOGICAL DEPENDENCE

A number of authors have identified the network diameter as an important factor in improving the performance of WDM networks [15–19]. The network diameter is defined as the maximum over all pairs of nodes with non-zero offered loads, of the length in hops of their defined routes. For example, in fully-meshed networks with fixed routing, routes consist of only a single link and wavelength converters offer no performance improvements. However, as route lengths increase, the relative performance of networks with and without wavelength converters diverge. In general, blocking in a wavelength-continuous network increases with increasing route length, as it becomes increasingly difficult to locate a common wavelength on each hop of a route. The increase in blocking probability with increasing route length is considerably less dramatic in networks with wavelength converters, because a connection can access any wavelength on each link along a route. This is illustrated in Fig. 6 where blocking probability is calculated using Kovačević and Acampora's analytical model [15] and is plotted versus route length for a single route with 10 wavelengths on each link.

Results presented in the literature have shown that wavelength converters generally provide more significant improvements in network performance in mesh topologies than in ring topologies [15, 19], despite the fact that large ring topologies having correspondingly large route lengths. This is illustrated

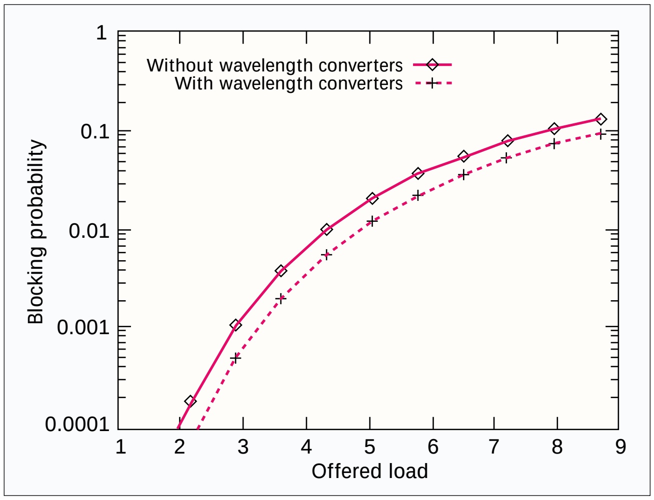

Figure 7. *Blocking probability vs. offered load for a unidirectional ring network with eight wavelengths on each link.*

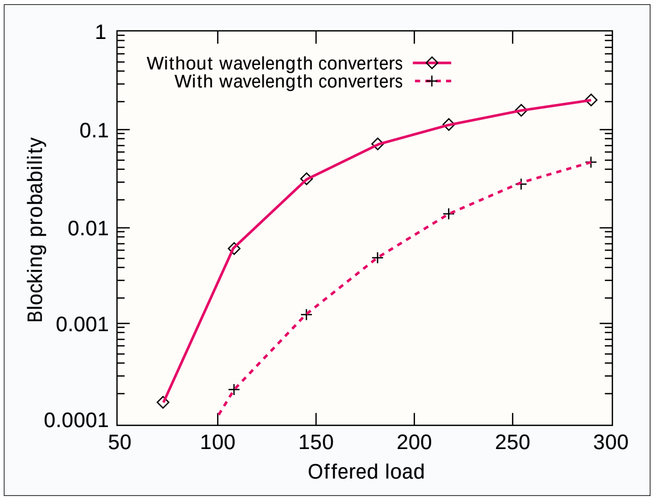

Figure 8. *Blocking probability vs. offered load for a mesh-torus topology [15] network with eight wavelengths on each link.*

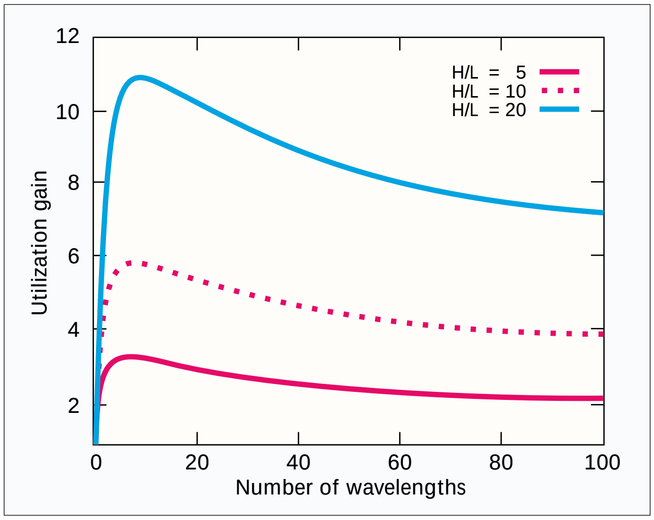

Figure 9. *Utilization gain vs. number of wavelengths for a blocking probability of 10^{-3} [16].*

in Figs. 7 and 8, respectively, for a 121-node unidirectional ring topology and an 11 x 11 mesh-torus topology [15]. Each of the networks has a total of 121 nodes, and eight wavelengths on each link. It can be seen that wavelength converters offer greater improvements in the blocking probability for a fixed offered load in the mesh-torus network than in the ring network. Similarly, if we consider a fixed blocking probability, then a larger proportion of extra traffic can be carried as a result of the introduction of wavelength converters in the mesh-torus network than in the ring network. For example, at a blocking probability of 10^{-3} , the introduction of wavelength converters allows an approximately 60 percent increase in the traffic which can be carried in the mesh-torus topology compared with a 14 percent increase in the traffic in the ring network. This is because in a ring topology, a large proportion of connections which use any given link also require the use of an adjacent link. This reduces the “mixing” of connections along a route [16], consequently reducing the need for wavelength conversion. We can illustrate this effect using an extreme example of a single traffic stream offered to a multiple-hop route. If we consider an H-hop route, then each traffic stream requires the use of every link along the H-hop route. If no wavelength converters are used along the route, each lightpath is established using a common wavelength on each link. At any point in time exactly the same wavelengths are allocated to lightpaths on each link. The blocking probability experienced thus reduces to the one-hop blocking probability and wavelength converters provide no improvement in the performance. Barry and Humblet [16] quantified this effect using the interference length, L, which they defined as the expected number of links shared by two lightpaths which share some link [16]. In the simple example outlined here of a single traffic stream offered to an H-hop route, the interference length is $L = H$.

Barry and Humblet analytically showed that for a route of fixed length, as the interference length increases, the benefit of wavelength converters decreases [16]. They identified the effective route length, defined as the ratio of the route length, H, to the interference length, L, as an important measure of the benefit of wavelength converters. For example, ring networks consist of large route lengths (large H), but also large interference lengths (large L) due to the high proportion of lightpaths which use adjacent links along a route. The effective route lengths are thus relatively small, and wavelength converters provide only marginal improvements in the average performance [15, 16, 19]. Highly connected networks such as the hypercube [19] and fully-meshed topologies have relatively short interference lengths but also have relatively short routes, again leading to small effective route lengths and marginal benefits for wavelength converters [19, 27]. In contrast, the mesh-torus topology has relatively long routes and short interference lengths, and wavelength converters thus provide significant performance improvements in large mesh-torus topologies. For example, Subramaniam, Azizoğlu, and Somani [19] showed that wavelength converters provide a reduction of up to eight orders of magnitude in the average network blocking probabilities in a 101 x 101 mesh-torus network with eight wavelengths.

In summary, wavelength converters generally offer only marginal performance benefits in networks with small diameters or with large interference lengths. However, wavelength converters may provide significant performance improvements in networks with large diameters and small interference lengths.

NUMBER OF WAVELENGTHS

Another factor in determining the benefits of wavelength converters is the number of wavelengths used on each link in the network. With the exception of [20, 21], discussions have focussed on networks with the same set of wavelengths provided on each link, so that each link has the same number of wavelengths.

The performance of single-wavelength networks is the same with or without wavelength converters. Kovačević and Acampora [15], Subramaniam, Azizoğlu, and Somani [19], and Wauters and Demeester [13] have shown that improvements in the blocking probability attained with wavelength converters become increasingly significant as the number of wavelengths increases, because the increased number of wavelengths allows increased mixing of connections.

However, as the number of wavelengths is increased, the offered load which can be supported in a network for a given blocking probability also increases. Barry and Humblet [16] examined the utilization gain for a single route, where they defined the utilization gain to be the ratio of the utilization with wavelength converters to the utilization without wavelength converters, for a fixed blocking probability. Intuitively, this measures the proportion of extra traffic which can be carried by introducing wavelength converters, and is of particular interest to telecommunications providers as it quantifies the possible increase in revenue for a given quality of service. This is illustrated in Fig. 9, where utilization gain is plotted versus number of wavelengths for different effective route lengths, H/L, and for a blocking probability of 10^{-3} [16]. Figure 9 shows that the utilization gain increases with increasing numbers of wavelengths when the number of wavelengths is small. However, as the number of wavelengths increases beyond a certain value (predicted using Barry and Humblet’s model [16] as being less than or equal to approximately 10 wavelengths for $H/L \leq 20$), the utilization gain slowly decreases, tending toward 1 as the number of wavelengths tends toward infinity. This makes sense intuitively as, due to trunking efficiency, the utilization achieved both with and without wavelength converters approaches 1 as the links approach infinite capacity [16]. Similar behavior is likely to be demonstrated in more complicated network topologies, as the carried load for a fixed blocking probability increases to infinity as the number of wavelengths tends toward infinity. Thus, the improvements that wavelength converters provide in increasing carried traffic become less significant as the number of wavelengths becomes large.

As discussed in the introduction, future optical WDM net-

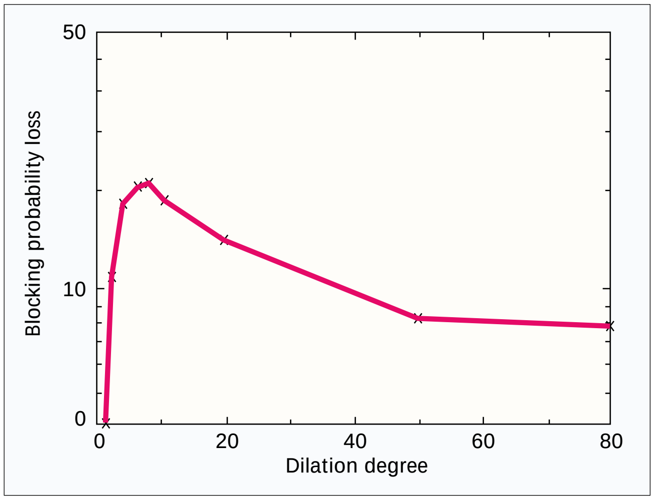

Figure 10. *Blocking probability gain vs. number of fibers per link for a 20-node ring and a blocking probability with wavelength converters of* 10^{-4} *[26].*

works are likely to support both static network traffic (in which the lightpath allocations do not vary significantly over time) and dynamic traffic (in which the lightpaths are established and released on demand). However, it is likely that the number of wavelengths allocated to support rapidly fluctuating lightpath demands will be relatively small, with most of the wavelengths in a WDM network supporting static lightpath allocations. Thus, most of the results published in the literature for dynamically-reconfigurable optical networks have considered relatively small numbers of wavelengths.

NUMBER OF FIBERS PER LINK

Multiple working fibers are often available on a single link, providing more capacity and better performance than single-fiber networks [13]. Wauters and Demeester [13], Jeong and Ayanoglu [25], and Subramaniam, Azizoglu, and Somani [26] have examined the benefits of wavelength converters in networks with multiple fibers on each link.

Wauters and Demeester [13] examined blocking in the 21-node, 26-link ARPA2 topology [14] when the product of the number of fibers and the number of wavelengths on a link is held constant. They showed that the benefits of wavelength converters rapidly disappear as the number of wavelengths decreases and the number of fibers per link correspondingly increase. For example, they showed that wavelength converters offer negligible performance improvements in the ARPA2 network with four fibers per link and four wavelengths per fiber. Subramaniam and Barry [26] similarly showed that the benefits of wavelength converters are reduced in a mesh-torus network with multiple fibers on each link as compared with networks with only a single fiber per link.

However, Subramaniam and Barry [26] have also examined multiple-fiber ring networks and have shown slightly contradicting results to those presented in [13, 25, 26] for more highly connected topologies. They measured the benefits offered by wavelength converters using the blocking probability loss, which they defined to be the ratio of the blocking probabilities without wavelength conversion to blocking probabilities with wavelength conversion. This is illustrated in Fig. 10 [26] for a 20-node ring with a fixed blocking probability with wavelength conversion of 10^{-4}. Figure 10 shows this ratio initially increasing with increasing numbers of fibers on each link, and then decreasing. Subramaniam and Barry [26] showed that wavelength conversion can potentially be more beneficial in multiple-fiber ring networks with moderate numbers of wavelengths than in single-fiber ring networks. As noted in [26], the reasons for this effect are not clear, and further investigation is required.

ROUTING AND WAVELENGTH ASSIGNMENT ALGORITHMS

The blocking probabilities experienced in a wavelength-continuous WDM network depend on the routing and wavelength assignment schemes used. Barry and Humblet suggested that a good routing algorithm in wavelength-continuous networks should consider route length in hops (H), interference lengths (L), and link congestion [16]. They suggested that routes chosen to minimize H/L may provide better performance than other choices. A good routing algorithm for wavelength-continuous networks will reduce the benefits of wavelength converters.

Due to ease of analysis, many of the analyses examining the benefits of wavelength converters in WDM networks have assumed random wavelength assignment [13, 15–17, 19, 25, 26, 29, 30]. However, numerous other wavelength assignment algorithms have been proposed in the literature. Careful wavelength assignment in a wavelength-continuous network can lead to improved performance, again reducing the benefits of wavelength converters.

Some of the proposed wavelength-assignment algorithms are:

- Random wavelength assignment allocates a new connection to a wavelength which is randomly chosen from among the set of available wavelengths.
- First-fit wavelength assignment [35] is implemented by predefining an order on the wavelengths. Wavelengths are searched in this order and a new connection is established on the first available wavelength.
- Most-used wavelength assignment [36] (also referred to as the pack scheme [22]) allocates a new connection to the wavelength that is used on the greatest number of fibers in the network. If several available wavelengths share the same maximum usage, the wavelength with, say, the lowest index is chosen. If instead of using the most-used wavelength, we allocate a connection to the least-used wavelength, we implement least-used wavelength assignment (also referred to as the spread scheme in [22]).
- The MaxSum wavelength assignment algorithm was recently proposed by Subramaniam and Barry in [26]. This algorithm attempts to minimize network blocking by minimizing the effect of establishing a new connection. Using the MaxSum algorithm, the effect of establishing a new connection is measured in terms of the number of routes whose capacities decrease by one [26].

Mokhtar and Azizoğlu [22] and Karasan and Ayanoglu [28] used network simulations to investigate the performance of the least-used, random, most-used, and first-fit wavelength assignment algorithms in networks with a single fiber on each link. Adaptive Unconstrained Routing was used in [22], with each wavelength searched for the shortest route available at connection establishment. In contrast, Karasan and Ayanoglu [28] used shortest path routing. The results showed that the least-used heuristic provides the worst performance of those considered, as it reduces the probability of finding an available route for connection establishment by distributing the load over all of the wavelengths. Random wavelength assignment effectively equalizes the load on each wavelength, providing performance which is only marginally better than that achieved using the least-used heuristic.

Mokhtar and Azizoğlu [22] and Karasan and Ayanoglu [28] showed that improved network performance is attained using the most-used wavelength allocation algorithm, particularly in networks with large numbers of wavelengths where increased "mixing" occurs and a good wavelength assignment scheme can be effective. However, this algorithm requires global knowledge of the network state. In the ARPA2 and random topologies examined, first-fit wavelength assignment gives performance which is close to that achieved using the most-used wavelength allocation algorithm. Additionally, the first-fit assignment scheme requires knowledge of the state of only the links along the route, thus providing a compromise between the knowledge required about the current network state and the performance attained.

Subramaniam and Barry [26] examined the performance of other wavelength assignment schemes in single-fiber and multiple-fiber networks. They showed that their proposed MaxSum algorithm outperforms any of the previously proposed algorithms in a single-fiber mesh-torus topology, and provides similar performance in a single-fiber ring topology to the most-used heuristic, which achieves the best performance from among the other proposed algorithms [26]. The drawback of the algorithm, however, is the computational complexity involved in determining a new connection's wavelength allocation.

The differences between wavelength assignment schemes are accentuated in a ring topology with multiple fibers on each link [26], with MaxSum again shown to provide better performance than other proposed algorithms [26]. In contrast, as the number of fibers used on each link increases in the mesh-torus networks, the performance of the various heuristics converges to the performance of the network with wavelength converters, suggesting that the choice of heuristic is not very important and that wavelength converters offer few advantages in a multiple-fiber mesh-torus network [26] (as discussed previously).

Birman and Kershenbaum [18] have investigated wavelength reservation and protection threshold schemes in single-fiber networks. Wavelength reservation refers to the dedication of a specific wavelength on each link along a route to a traffic stream, while the technique of protection threshold allows the assignment of single-hop traffic to an idle wavelength only if the number of idle wavelengths on the link is at or above a predefined threshold (similar to trunk reservation in circuit-switched networks). These techniques were designed to improve the blocking of multiple-hop connections in wavelength-continuous networks. Birman and Kershenbaum showed that the wavelength reservation and protection threshold schemes can improve the blocking of multiple-hop connections, but that this is at the expense of the blocking experienced by the single-hop connections [18]. These techniques were also shown to degrade the average blocking in the networks examined, and are thus not considered to be effective wavelength assignment schemes.

TRAFFIC DISTRIBUTION

The majority of the analyses of dynamically-reconfigurable WDM networks have assumed a Poisson connection (lightpath) arrival process and negative exponential holding times. These are common assumptions for modeling telephone (voice) traffic in traditional circuit-switched networks [37]. However, they may only be applicable in dynamically-reconfigurable WDM networks for individual calls with very large bandwidths. It is expected that these assumptions are not good approximations for lightpaths carrying time-division multiplexed traffic streams.

Given such uncertainties, Subramaniam *et al.* [32], Yates [21], and Späth and Bodamer [33] have investigated the benefits offered by wavelength converters in networks with example non-Poisson input traffic models. Specifically, Subramaniam *et al.* [32] used a Bernoulli–Poisson–Pascal model to model different traffic in mesh-torus and binary hypercube topologies. In contrast, Yates [21] used an approximate analytical model to examine blocking along a single route with traffic described using various traffic models similar to those in [32]. In both works, the traffic was examined in terms of its peakedness, i.e., the ratio of the traffic variance to the mean. Subramaniam *et al.* [32] and Yates [21] showed that wavelength converters provide smaller improvements in blocking probabilities in networks as the traffic becomes increasingly peaked (so that the ratio of the traffic variance to the mean increases). To illustrate why this occurs, consider a single hop (link) being offered increasingly peaked, or bursty, traffic. Each burst results in a rapid increase in the number of used wavelengths. When blocking does occur, many calls are likely to be blocked, even though the overall fraction of time spent in the blocking state may be relatively small. Recall that blocking probability is a measure of call, rather than time, congestion of the system. Thus, to maintain a constant level of blocking with increasing peakedness, it is necessary to decrease the overall arrival rate. As a result, the probability that a significant number of wavelengths on the hop are used decreases as peakedness increases. Looking then at a network without wavelength converters, as peakedness is increased, the probability that more than one hop has many wavelengths in use decreases. In such situations, little benefit is to be obtained via the use of wavelength converters. However, Subramaniam *et al.* [32] also predicted that in the mesh-torus and hypercube topologies with 10 wavelengths on each link, the network utilization gain is actually higher with peaked traffic than with Poisson traffic. The approximate analytical model used by Yates [21] agrees with this conclusion when the number of wavelengths is small. However, if the number of wavelengths is increased, this analytical model predicts the utilization gain decreasing with increasing traffic peakedness. Thus, for large numbers of wavelengths it is predicted that wavelength converters provide less significant performance improvements for more peaked traffic measured both in terms of blocking probabilities and utilization gain. However, the relative benefits of wavelength converters are dependent on the actual traffic model used. Further, if the traffic is less peaked than Poisson traffic, Yates [21] predicted that the wavelength converters might offer increasingly significant improvements in the utilization gain.

Späth and Bodamer [33] took an alternative approach to that in [32] and [21] by considering networks with various lightpath arrival request distributions and generally distributed holding times. Their results indicated that network performance is nearly independent of the lightpath holding time distribution, but is a strong function of the lightpath request arrival process.

MULTIWAVELENGTH TDM NETWORKS

One type of network which could lead to non-Poisson lightpath requests is one in which calls are time-division multiplexed onto lightpaths. How time-division multiplexing (TDM) and WDM are used in a network depends on the cross-connects used in the optical layer and on the electronics used in the electronic switching layer of the future transport network depicted in Fig. 1. For example, if the state of the space switches in the optical cross-connects can be rapidly reconfigured, then space switching can be performed on a per timeslot basis. We refer to this as fast space switching. In this section we discuss the impact of fast space switching on network operation and performance.

An initial investigation of networks which utilize both TDM and WDM was performed by Yates in [21]. Two different scenarios were considered, depending on whether fast space switching can be performed. If fast space switching is not available in an optical network, then high-bandwidth lightpaths are established between individual optical nodes, as described in the first section. If simple electronic multiplexers are used in the electronic layer of Fig. 1, then calls between common source-destination optical nodes can be electronically multiplexed in the time domain onto common lightpaths. Calls established between different source-destination optical nodes must use different lightpaths. In this way, the use of electronic cross-connects and add-drop multiplexers can be minimized, potentially leading to simpler, cheaper networks. Similarly, if lightpaths are allocated on demand to individual corporate customers, for example, these customers may themselves multiplex individual calls onto a single lightpath.

Yates [21] showed that in these networks, wavelength converters may offer more significant performance improvements as the number of calls which can be multiplexed onto a single lightpath increases. For example, in a 5 x 5 mesh-torus network with 15 wavelengths per link, the introduction of wavelength converters increased the traffic which can be supported with a blocking probability of 10^{-3} by approximately 19 percent when a single call is carried on each lightpath. This can be compared with an 81 percent increase in traffic which can be supported with the introduction of wavelength converters in the same network topology when a maximum of four calls can be multiplexed onto each lightpath. However, [21] shows that if the capacity of each lightpath is sufficiently large compared with call bandwidths, then a static allocation of lightpaths often provides better blocking performance than when lightpaths are allocated on demand. This is because if multiple wavelengths are allocated on demand to one source-destination pair, but are not well utilized, then this may prevent other lightpaths between different source-destination pairs being established, increasing the overall blocking probability.

In contrast, if fast space switching is available, either optically or electronically, then different time slots on a common wavelength can be switched from a common input fiber to different output fibers [21, 34]. Networks using fast space switching generally offer better blocking than networks without fast space switching, at the cost of more complicated cross-connects [21]. In such networks, timeslot interchange (TSI) and wavelength conversion may both be used to improve network performance. A timeslot interchanger is a device which can rearrange the order of the time slots in a (single wavelength) traffic stream passing through it. TSI reduces blocking in TDM networks in the same way as wavelength converters reduce blocking in WDM networks.

Yates [34] showed that in a network with large effective route lengths, a small number of wavelengths, and a relatively large number of time slots per wavelength, TSI and wavelength conversion together can offer significant performance improvements and wavelength conversion alone can offer moderate performance improvements. However, when the effective route length is small, TSI and wavelength conversion together offer only small performance improvements, and wavelength conversion alone generally offers insignificant performance improvements. The most significant conclusion, though, is that independent of the effective route length, TSI alone provides almost all of the performance improvements achieved using both TSI and wavelength conversion [34]. By contrast, in a network with a large number of wavelengths and a relatively small number of time slots per wavelength, the above result can be reinterpreted, replacing TSI everywhere by wavelength conversion, and vice versa.

ALTERNATE ROUTING

The majority of analyses of wavelength-continuous WDM networks have focussed on networks with fixed routing, in which only a single route is defined between each source and destination. If this route is not available when a lightpath request arrives, the lightpath is blocked.

Significant performance improvements are often obtained in circuit-switched networks, and equivalently in networks with wavelength converters, if alternate routing is introduced. Birman [17], Mokhtar and Azizoğlu [22], Karasan and Ayanoglu [28], Harai, Murata and Miyahara [29], Chan and Yum [30], and Ramamurthy and Mukherjee [31] have investigated the performance of wavelength-continuous single-fiber networks with alternate routing. They have considered various routing and wavelength assignment schemes, and have shown that the performance of wavelength-continuous networks is strongly dependent on the schemes chosen.

Karasan and Ayanoglu [28] proposed a Maximum H/L Routing (MHLR) scheme which chooses a route from among the k shortest routes which maximizes the ratio of the route length to the interference length (H/L). First-fit wavelength assignment is used in the wavelength-continuous networks. As the number of possible alternate routes, k, increases, H/L increases and the wavelength-continuous blocking probability was shown to increase with increasing k in the network topology considered in [28]. In contrast, the blocking probability with wavelength converters decreases with increasing k, due to the increased set of possible routes through the network for establishing a lightpath. Thus, the performance improvements offered by wavelength converters increases.

The MHLR algorithm is an extreme example designed to increase the benefits of wavelength converters. Karasan and Ayanoglu [28] also considered Least-Loaded Routing (LLR). In networks with wavelength converters, the LLR algorithm chooses a route from the k possible alternate routes which maximizes the minimum number of available wavelengths on any given link along a route. In a wavelength-continuous network, the route and wavelength are chosen together to maximize the minimum number of fibers on which a wavelength is available on a link. The most-used wavelength and the shortest path are used as tie-breakers.

The LLR algorithm was shown to decrease blocking in the example single-fiber network with and without wavelength conversion. However, the performance improvements obtained with the introduction of wavelength conversion were shown to be comparable to those obtained using the MHLR algorithm, and significantly greater than with fixed routing. This is because having k alternate routes, as compared with a single choice, results in increased average route lengths and also in reduced interference lengths [28]. This leads to an increased effect of the wavelength continuity constraint on network blocking. Similarly, Ramamurthy and Mukherjee [31] showed the benefits of wavelength conversion increasing with number of alternate routes in networks employing fixed-alternate routing, in which the different alternate routes have a predefined order and lightpaths are established on the first available alternate route.

Chan and Yum [30] examined blocking in a fully-meshed network with Least-Congested Path (LCP) routing, random wavelength assignment and trunk reservation. When a connection request is made, the number of available wavelengths on each link of each possible route is determined. The connection is established on the route with the maximum number of free wavelengths, where the number of free wavelengths along a route is defined as the minimum of the number of available wavelengths of all links constituting the route. Ties are broken by examining the next most highly used links, and so on.

A direct link exists between each source and destination in a fully-meshed network, with alternate routes consisting of two links. Wavelength converters provide no performance improvements in a fully-meshed network with fixed routing as all routes consist of only a single hop. However, Chan and Yum [30] showed that if alternate routing is introduced, wavelength converters can reduce blocking probabilities at low offered loads. However, the performance improvements are relatively small due to the short route lengths.

Analyses of traditional circuit-switched networks (equivalent to networks with wavelength converters) have shown that alternate routing without trunk reservation can produce higher blocking than fixed routing in a fully-meshed network operating with high offered loads [38]. This results from having two regimes of operation, in which there are times when much of the traffic is alternately routed, requiring two links per connection, and times when much of the traffic is directly-routed, requiring only a single link. However, Chan and Yum [30] and Yates *et al.* [27] observed that wavelength-continuous networks which use LCP routing and fixed alternate routing respectively with connections preferentially established on direct routes experience lower blocking probabilities than networks with wavelength converters. This results from the inherent alternate route blocking of the wavelength-continuity constraint. As the number of wavelengths used on a link increases, the probability of blocking for a two-hop route also increases due to the wavelength continuity constraint, reducing the number of alternately routed connections. This allows more directly routed one-hop connections to be established, reducing the overall blocking probability [30].

NETWORKS WITH DIFFERENT NUMBERS OF WAVELENGTHS ON DIFFERENT LINKS

Yates *et al.* [20, 21] considered the performance improvements offered by wavelength converters in WDM networks with different numbers of wavelengths on different links. In networks without wavelength converters and with different sets of wavelengths on different links, a connection can only be established on a wavelength which exists on every link along the desired route. In contrast, the introduction of wavelength converters allows connections access to every available wavelength on every link of the route.

It was shown in [20] that wavelength converters can often provide more significant performance improvements in networks with different sets of wavelengths available on different links than in networks in which the same set of wavelengths exists on each link. This is because in the wavelength-continuous case, connections can only be established on a wavelength which exists on all links along a route. In a network with different sets of wavelengths on different links, this reduces the number of wavelengths on which connections can be established, potentially resulting in significantly higher blocking. However, the performance improvements achieved with wavelength converters is a strong function of the loads offered to each link and the wavelength assignment scheme used in the networks without wavelength converters. Careful choice of wavelength assignment can again significantly reduce the need for wavelength conversion [20].

FAIRNESS

As previously discussed, wavelength converters may reduce average network blocking. However, the wavelength continuity constraint may also introduce significant unfairness in the blocking probabilities. In networks without wavelength converters, many-hop connections often experience significantly higher blocking than connections with shorter routes [14] as it is more difficult to locate a common wavelength on each hop of the route. Bouillet and Bala [23 , 24] showed that even though wavelength converters provide only marginal improvements in the average blocking in a ring topology, they can dramatically increase the fairness as blocking probability becomes a significantly smaller function of route length. They illustrated this in an 11-node ring with 32 wavelengths and using first-fit wavelength assignment. The variation in blocking between the longest and shortest routes in the network without wavelength converters is significantly greater (up to 100 times) than the variation in blocking with wavelength converters. More dramatically, simulations of 15 interconnected rings with 13 nodes on each ring and 32 wavelengths on each link showed that the variation in blocking between the longest and shortest routes can be up to 10000 times greater without wavelength converters than with wavelength converters.

SUMMARY

In this section we have discussed the important factors which determine the performance improvements offered by wavelength converters in WDM networks. In most network topologies, wavelength converters generally provide only modest improvements (approximately 10–40 percent [39]) in the traffic which can be supported for a given quality of service (blocking probability). However, more significant performance improvements may be available in networks with large effective route lengths, such as in meshed networks.

LIMITED WAVELENGTH CONVERSION

In the previous section we discussed the performance improvements offered by wavelength converters in WDM networks. However, it was assumed that a full set of ideal wavelength converters was available at every cross-connect in the network. That is, it was assumed that in every cross-connect, any input wavelength could be converted to any output wavelength.

In this section we discuss networks with limited wavelength conversion, in which we no longer assume that any input wavelength can be converted to any output wavelength in every cross-connect in the network. This may be the result of placing wavelength converters at a limited selection of cross-connects in the network, using limited numbers of wavelength converters in each cross-connect, or using wavelength converters whose performance limits the set of allowable conversions.

LIMITED NUMBER OF WAVELENGTH CONVERTERS IN EACH NODE

Architectures for optical cross-connects with limited numbers of wavelength converters have been proposed by Lee and Li in [14] and by Parys *et al.* in [42]. In all of these architectures, wavelength conversion is performed after space switching, as illustrated in Fig. 5. The architectures vary, however, according to how the wavelength converters are shared within the cross-connect. This can be done on a per node basis [14], a per link basis [14], or the wavelength converters can be dedicated to selected wavelengths on selected links [42].

The simplest limited cross-connect architecture provides wavelength converters on only selected wavelengths on selected links [42]. The resulting architecture is similar to that shown in Fig. 4, but with some of the wavelength converters removed. The blocking performance of dynamically-reconfigurable networks which incorporate these cross-connects has not been investigated. However, it has been shown by Parys *et al.* [42] that only marginal improvements can be achieved using this cross-connect architecture in a network with static lightpath requirements and offline routing due to its very limited conversion capability. Similar results could be expected in dynamically-reconfigurable networks with online routing, as the cross-connect design is fairly inflexible in terms of the conversions which can be performed. However, this requires further investigation to verify this conclusion.

Lee and Li [14] considered architectures in which banks of wavelength converters are provided in each cross-connect, with either a single bank of converters shared by all lightpaths passing through a node, or converter banks provided for each output link. Connections are either switched through the wavelength converter banks, where they undergo wavelength conversion, or they bypass the wavelength converter bank and remain on the same wavelength. Lee and Li [14] used the ARPA2 network topology with 16 wavelengths per link to investigate the performance of networks which use these cross-connects. They showed that for small numbers of wavelength converters in each cross-connect, increasing the number of converters improves network performance. However, after a certain threshold, the performance improvements offered by further increasing the number of wavelength converters are marginal. For example, Lee and Li considered wavelength converters shared on a per-node basis, with the same number of wavelength converters in each cross-connect, independent of the cross-connect size. They showed that, for the loads considered, between 4 and 12 wavelength converters are required in each node to provide almost all of the performance improvements offered by having a wavelength converter for every output wavelength on every link. This compares with an average of 40 wavelength converters required per node to provide full wavelength conversion. Lee and Li thus concluded that only a small number of wavelength converters are required in WDM cross-connects to obtain almost the full performance benefits of wavelength converters.

Lee and Li [14] used a fixed number of wavelength converters in each node, and made no attempt to optimize the placement of wavelength converters in the network. Iness [43] proposed a simple heuristic for placing wavelength converters in cross-connects within a network. The effectiveness of this heuristic requires investigation, and new heuristics should be proposed if future optical networks are to have limited numbers of wavelength converters in a cross-connect.

A major drawback of using limited numbers of wavelength converters in each cross-connect in a WDM network is that the resulting cross-connect, such as proposed by Lee and Li in [14], is often more complicated than cross-connects with wavelength converters for every wavelength on every link, as extra switching capability is required to allow connections to either pass through a wavelength converter, or to bypass it. The control of networks using these cross-connects may also be more complicated.

WAVELENGTH CONVERTERS IN SELECTED NODES

An alternative to using a limited number of wavelength converters in each cross-connect is to use a combination of cross-connects with a full set of wavelength converters and cross-connects with no wavelength converters [40]. Subramaniam, Azizoğlu, and Somani [19], Wauters *et al.* [7] and Iness [43] have analytically modeled WDM networks with wavelength converters placed in only some nodes within the network.

Subramaniam, Azizoğlu, and Somani [19] examined the effect of having limited, or sparse, wavelength conversion by providing comparisons of blocking probabilities in networks with varying conversion densities, where they defined the conversion density to be the probability that any given cross-connect contains wavelength converters [19]. They showed that in relatively large ring topologies with reasonable numbers of wavelengths (illustrated for a 100-node ring with 10 wavelengths per link), the blocking probability initially rapidly decreases with increasing conversion density and then levels off after a certain point. However, in ring networks with only a few nodes or with few wavelengths on each link, the blocking probability decreases more gradually with increasing conversion density. They also showed that at light loads the number of nodes which can be supported with a constant blocking probability and a constant offered load per node initially increases with increasing conversion density. However, as the conversion density increases further, the rate of increase in the number of nodes which can be supported reduces.

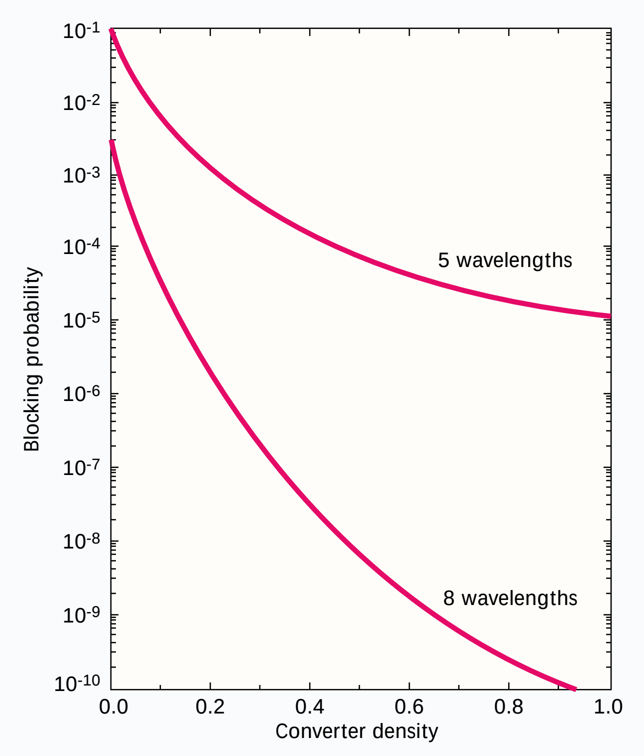

Figure 11. *Blocking probability vs. converter density for a 101 x 101 mesh-torus topology with 5 and 8 wavelengths.*

Similar observations were made in a mesh-torus topology. In large mesh-torus topologies, the blocking probability initially decreases dramatically with increasing conversion density. As the conversion density increases further, the improvement in performance with increases in the conversion density become less dramatic. This is illustrated in Fig. 11 for a 101 x 101 mesh-torus topology with 5 and 8 wavelengths [19]. However, in smaller mesh-torus topologies, and in the hypercube topologies considered, steady improvements in blocking probability are achieved with increasing conversion density.

Bouillet and Bala [23] have examined the performance improvements offered by using wavelength converters at only selected nodes in four-fiber bidirectional WDM ring networks [23]. They assumed that wavelength converters are placed in nodes at regularly spaced intervals. They measured the improvements in the fairness of the blocking probabilities obtained by having only limited numbers of wavelength-converting nodes. As discussed earlier, they measured the fairness of a network as the ratio of the average blocking probability

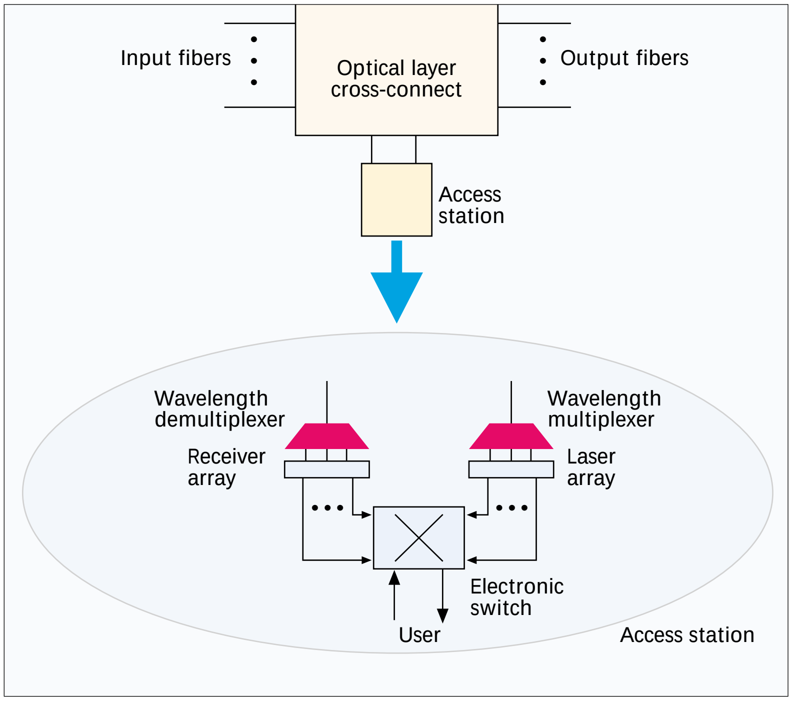

Figure 12. *An optical layer cross-connect with associated access station [44]. Optoelectronic wavelength conversion can be performed by the access station by receiving the data modulated onto a wavelength and switching the data back to the access node's transmitter for retransmission on a different wavelength.*

on the longest routes to the average blocking on the shortest routes. They showed that in a network consisting of 15 WDM interconnected rings with 13 nodes per ring and 32 wavelengths per link, adding wavelength converters to 10 percent to 20 percent of nodes in the network results in considerable improvements in the fairness, with the addition of more wavelength converters providing only small further improvements. They added that having wavelength converters at only the boundaries of the interconnected rings may be sufficient to obtain almost the full benefits of wavelength converters. However, they also observed almost linear improvement in fairness in smaller rings as the number of wavelength converters was increased.

Wauters *et al.* [7] examined the performance of the COST 239 European Optical Network with wavelength converters used in 6 of the 19 nodes to create non-overlapping network partitions [7]. Connections within each partition must be established without wavelength conversion, while wavelength conversion is performed in passing between partitions. Wauters *et al.* showed that the use of wavelength converters in 6 of the 19 nodes within the network reduces the blocking probability to almost that achieved with wavelength converters in every node in the network [7].

The results discussed above did not use optimal placement of wavelength converters in the WDM networks considered. Subramaniam, Azizoglu and Somani [41] and Iness [43] considered the problem of wavelength converter placement in WDM networks and in many networks showed significant differences in the performance obtained using optimal converter placement as compared with random [41] or average [43] wavelength converter placement. Subramaniam, Azizoglu and Somani [41] showed that uniform spacing of wavelength converters is optimal for the end-to-end performance when link loads are uniform and statistically independent. They also provided solutions based on dynamic programming for the optimal placement of wavelength converters along routes with non-uniform traffic and in which the loads on different links are not statistically independent. The dynamic programming techniques were also applied to ring and bus networks.

Iness [43] proposed a simple heuristic based on output-link congestion for wavelength converter placement. This heuristic was shown to work well at high loads in the example networks considered. However, further investigation is required to verify the applicability of the heuristic to other network topologies and at various network loads.

Subramaniam, Azizoğlu, and Somani [19], Bouillet and Bala [23] and Wauters *et al.* [7] have thus concluded that it may rarely be necessary to place wavelength converters in every node of a WDM network. However, it remains an important open problem to optimally place wavelength converters within arbitrary network topologies.

OPTOELECTRONIC WAVELENGTH CONVERSION AT NETWORK ACCESS STATIONS

Kovačević and Acampora [44] proposed using optoelectronic wavelength conversion performed by receiving and retransmitting an optical signal using an access station connected to a cross-connect. Their architecture has access stations connected to each cross-connect via a fiber carrying WDM. Each station has a limited number of transmitters and receivers which are used either to establish or terminate a connection, or, in this proposed architecture, to perform wavelength conversion. This is illustrated in Fig. 12.

The use of limited numbers of transmitters and receivers at the access station limits the number of conversions which can be performed and the number of connections which can be established or terminated at an access station. On the other hand, if wavelength conversion were performed in the optical cross-connect, the limited number of transmitters and receivers could be used purely for connection establishment and termination. However, network simulation results for mesh-torus topologies showed that significant performance improvements can be obtained using this form of optoelectronic wavelength conversion as compared with no conversion, particularly for large networks with relatively few wavelengths. Kovacevic and Acampora showed that in a 101 x 101 mesh-torus network with five wavelengths on each link, wavelength conversion at network access stations provides almost all of the performance benefits obtained using full wavelength conversion in every cross-connect in the network.

NETWORKS WITH NON-IDEAL WAVELENGTH CONVERTERS

Even if the number of wavelength converters in each cross-connect is sufficient to provide full conversion in each cross-connect, wavelength conversion may be limited by non-ideal performance of wavelength converters.

A number of techniques for wavelength conversion have been proposed in the literature. Some of the most promising of these techniques are cross-phase modulation in semiconductor optical amplifiers (SOAs), cross-gain modulation in SOAs and opto-electronic wavelength conversion [49]. However, the transparency of these techniques is limited [11]. An alternative technique which provides modulation-independent (transparent) wavelength conversion is four-wave mixing

(FWM) in semiconductor optical amplifiers (SOAs) [50, 51]. However, the performance of SOA FWM wavelength converters is a strong function of the difference between the wavelength converter's input and output frequencies [50, 51]. That is, for a particular input frequency (wavelength), conversion to some output frequencies results in an output signal which is significantly degraded. This leads to limitations in the conversions which can be performed.

Yates *et al.* [21, 45, 46] examined the performance of networks which use non-ideal wavelength converters. In particular, their work focussed on examining the performance of networks which used SOA FWM wavelength converters. This work recognized that connections are effectively blocked in WDM networks if the received signal quality is so poor that many bits are received in error. Thus, analyses of WDM networks with SOA FWM wavelength converters must account for blocking due to either insufficient capacity or unacceptable signal quality.

Although the performance of networks which use SOA FWM wavelength converters is a strong function of the device characteristics and required signal quality, Yates *et al.* [21,45,46] showed that in networks with tunable transmitters and receivers, even very limited wavelength converters can offer performance which is close to that achieved with ideal wavelength converters. They also showed that the introduction of SOA FWM wavelength converters can allow networks with fixed-frequency transmitters and receivers to perform adequately, but that better performance can often be achieved, with simpler hardware, by introducing tunable transmitters and receivers rather than SOA FWM wavelength converters [21, 46]. Networks with SOA FWM wavelength converters and tunable transmitters and fixed receivers (or fixed transmitters and tunable receivers) are more complex than equivalent networks with fixed transmitters and receivers. However, their performance is often better than that achieved in networks with no wavelength converters and tunable transmitters and receivers. Thus, non-ideal wavelength converters can provide performance which is close to that achieved with ideal wavelength converters, but this is a strong function of the device characteristics, the acceptable signal quality requirements and the tunability of the transmitters and receivers.

Gerstel, Sasaki and Ramaswami [47, 48] also examined the performance of dynamically-reconfigurable ring networks with limited wavelength conversion. Their analysis contrasted with many presented in the literature in that they investigated the maximum number of connections which could be supported in a dynamically-reconfigurable ring network without blocking. In examining this worst-case scenario, Gerstel, Sasaki, and Ramaswami [47] showed that as the number of nodes in a network becomes large, the worst-case behavior of networks with limited wavelength conversion can be significantly better than that with no wavelength conversion. However, in smaller networks, Gerstel, Ramaswami and Sasaki [48] showed that using their wavelength assignment algorithms, the guaranteed loads that can be supported with only limited wavelength conversion are significantly smaller than those achieved with ideal wavelength conversion. Improved wavelength assignment algorithms are thus required to guarantee high throughputs with limited wavelength conversion when dimensioning a dynamically-reconfigurable ring network with zero blocking.

OTHER POTENTIAL BENEFITS OF WAVELENGTH CONVERSION

Although the performance improvements offered by wavelength converters are often marginal, such as in a ring topology, wavelength converters could potentially be used in WDM networks to simplify network management. For example, wavelengths can be assigned on a link-by-link basis in networks with wavelength converters, eliminating the need for the wavelength assignment algorithms used in wavelength-continuous networks and thus simplifying network management [7]. Alternatively, subnetworks may be operated without wavelength converters, with wavelength converters used only at the interfaces between subnetworks to isolate the wavelength management, eliminating the need for subnetworks to share information regarding current wavelength allocations [6]. If the different subnetworks have different operators, wavelength conversion at the subnetwork interfaces removes the undesirable requirement for information regarding the current wavelength allocations to be shared between competitors. Wavelength converters can also be placed at the borders of networks to convert signals to the wavelengths used within the WDM network and to provide tunability for connections accessing the network using fixed-frequency transmitters and receivers. This may be particularly important in providing connections for existing SONET and ATM networks which do not operate at ITU standard wavelengths [52].

In networks without wavelength converters, absolute precision of the wavelength (or frequency) is required throughout the network [53] so that transmitters and receivers in different nodes operate at common frequencies and can be used to establish connections. In contrast, wavelength precision is only required between the link terminating nodes in networks with wavelength converters. This leads to relaxed requirements in terms of optical devices and wavelength precision for networks with wavelength converters [53].

One advantage of not using wavelength converters in a WDM network is that optical cross-connects without wavelength converters, as depicted in Fig. 4, are more modular and thus more easily upgraded than networks which incorporate the wavelength converting cross-connects depicted in Fig. 5 [54].

FUTURE RESEARCH

As discussed in the section "Multiwavelength TDM Networks," several calls will often be multiplexed onto a single lightpath, rather than allocating entire lightpaths to each call. This multiplexing may be performed by electronic multiplexers, add-drop multiplexers, digital cross-connects or routers (eg. IP routers). These different technology choices for the electronic layer of future networks (for example, in Fig. 1) will impact on how traffic is multiplexed onto lightpaths. This will affect the traffic distribution describing the lightpath requirements and consequently, the performance improvements attainable using wavelength converters. An initial investigation into some aspects of this problem was provided in [21] and discussed in "Multiwavelength TDM Networks." However, further work is required to investigate the performance improvements attainable using wavelength converters in networks incorporating both electronic and optical layers.

The rapid growth in IP traffic means that the efficient transport of IP traffic over WDM networks has recently become a topic of intense research interest [55]. It is important to understand where lightpaths should be established in an IP-over-WDM network, and how and when they should be reconfigured. If lightpaths are to be dynamically reconfigured to support high-bandwidth IP connections, then the results presented in this article are applicable to IP-over-WDM networks. However, the lightpath request and holding time distributions are again important parameters, and require investigation. The effect which lightpath blocking will have on IP congestion and packet loss will also require investigation.

Another question which will impact on whether wavelength converters are introduced into a network is transparency. As discussed in the introduction, the original dream of many optical network pioneers was to build transparent, all-optical networks. However, the extent to which transparency can or should be achieved in a WDM network is uncertain [5], and requires further investigation. If optoelectronic regenerators are included in WDM cross-connects, then the network performance improvements and simplified network management introduced by wavelength converters may be significant enough to justify the additional cost of increasing the space switch size to use the regenerators as wavelength converters. However, if optically transparent cross-connects are to be used, then the marginal performance improvements and simplified network management obtained using wavelength converters may not be significant enough to justify the extra cost of introducing all-optical wavelength converters throughout the network. Thus, the transparency of a WDM network is an important issue to be resolved before decisions can be made as to whether to introduce wavelength converters into the network.

If a network operator decides to deploy an all-optical network with all-optical wavelength converters, then the limitations of these devices should be considered. An initial investigation into the performance of networks which use wavelength converters based on four-wave mixing in semiconductor optical amplifiers has been presented in [45, 46]. However, other types of all-optical wavelength converters are expected to have different limitations imposed by the different physical laws governing their performance. The performance of WDM networks using such wavelength converters should be investigated, and the different wavelength converters compared. Finally, it has been shown that most of the performance improvements obtained using a full set of wavelength converters in every cross-connect of a WDM network can be achieved using limited numbers of wavelength converters [7, 14, 19, 43]. However, the optimal placement of these wavelength converters in an arbitrary network is an important issue which needs to be thoroughly investigated before limited wavelength conversion is implemented.

CONCLUSIONS

The performance improvements offered by wavelength converters in dynamically-reconfigurable WDM networks have been examined in [11-34]. These investigations have shown that in most network topologies, wavelength converters offer only modest performance improvements. However, in networks with large route lengths and small interference lengths (measured as the expected number of links shared by two lightpaths which share some link), wavelength converters can provide significant performance improvements.

A number of authors have considered networks with limited wavelength conversion, in which the set of allowable conversions which can be performed at a network node is constrained by having limited numbers of wavelength converters [7, 19, 23, 43, 44] or by using non-ideal wavelength converters [21, 45–48]. Limited wavelength conversion has been shown to provide performance which is often close to that achieved with ideal wavelength conversion in networks with tunable transmitters and receivers [7, 19, 21, 23, 43–46].

To summarize, the extensive use of wavelength converters in the majority of future WDM networks is unlikely to be justified solely on performance-related grounds.

ACKNOWLEDGMENTS

The authors thank Suresh Subramaniam for providing Figs. 10 and 11. Jennifer Yates and Michael Rumsewicz also thank Ericsson Australia for their financial support.

REFERENCES

[1] R. Ramaswami and K. Sivarajan, Optical Networks: A Practical Perspective, Morgan Kaufmann Publishers, 1998.

[2] *IEEE Commun. Mag.,* vol. 36, no. 2, Feb. 1998.

[3] R. Barry, "Tutorial on the Performance of WDM Optical Networks," *Optical Fiber Commun. Conf.*, Feb. 1997.

[4] P. Cochrane, R. Heckingbottom and D. Heatley, "The Hidden Benefits of Optical Transparency," *IEEE Commun. Mag.,* vol. 32, no. 9, pp. 90–97, Sept. 1994.

[5] R. Tkach, E. Goldstein, J. Nagel and J. Strand, "Fundamental Limits of Optical Transparency," *Optical Fiber Communication Conf.,* pp. 161–162, Feb. 1998.

[6] L. Clark, "Evolution of Optical Transport in the AT&T Network," *Optical Fiber Communication Conf.,* p. 70, Feb. 1998.

[7] N. Wauters, W. Van Parys, B. Can Caenegem and P. Demeester, "Reduction of Wavelength Blocking through Partitioning with Wavelength Converters," *Optical Fiber Commun. Conf.,* pp. 122–123, Feb. 1997.

[8] S. Finn and R. Barry, "Optical Services in Future Broadband Networks," *IEEE Network,* vol. 10, no. 6, pp. 7–13, Nov./Dec. 1996.

[9] A. Aggarwal, A. Bar-Noy, D. Coppersmith, R. Ramaswami, B. Schneider and M. Sudan, "Efficient Routing and Scheduling Algorithms for Optical Networks," *5th Annual ACM-SIAM Symp. on Discrete Algorithms*, pp. 412–423, Jan. 1994.

[10] R. Barry and P. Humblet, "Bounds on the Number of Wavelengths Needed in WDM Networks," *LEOS Summer Topical Mtg. Digest,* pp. 21–22, 1992.

[11] B. Ramamurthy and B. Mukherjee, "Wavelength Conversion in WDM Networking," *IEEE JSAC,* vol. 16, no. 7, pp. 1061–1073, Sept. 1998.

[12] E. Karasan and E. Ayanoglu, "Performance of WDM Transport Networks," *IEEE JSAC,* vol. 16, no. 17, pp. 1081–1096, Sept. 1998.

[13] N. Wauters and P. Demeester, "Wavelength Translation in Optical Multi-wavelength Multi-fiber Transport Networks," *Int'l J. of Opto-electronics,* vol. 11, no. 1, pp. 53–70, Jan./Feb. 1997.

[14] K.C. Lee and V. Li, "A Wavelength-Convertible Optical Network," *IEEE/OSA J. of Lightwave Technol.*, vol. 11, no. 5/6, pp. 962–970, May/June 1993.

[15] M. Kovačević and A. Acampora, "Benefits of Wavelength Translation in All-Optical Clear-Channel Networks," *IEEE JSAC,* vol. 14, no. 5, pp. 868–880, June 1996.

[16] R. Barry and P. Humblet, "Models of Blocking Probability in All-Optical Networks with and without Wavelength Changers," *IEEE JSAC,* vol. 14, no. 5, pp. 858–867, June 1996.

[17] A. Birman, "Computing Approximate Blocking Probabilities for a Class of All-Optical Networks," *IEEE JSAC,* vol. 14, no. 5, pp. 852–857, June 1996.

[18] A. Birman and A. Kershenbaum, "Routing and Wavelength Assignment Methods in Single-Hop All-Optical Networks with Blocking," *IEEE INFOCOM,* pp. 431–438, April 1995.

[19] S. Subramaniam, M. Azizoğlu and A. Somani, "All-Optical Networks with Sparse Wavelength Conversion," *IEEE Trans. on Networking,* vol. 4, no. 4, pp. 544–557, August 1996.

[20] J. Yates, M. Rumsewicz and J. Lacey, "Wavelength Conversion in Networks with Differing Link Capacities," *IEEE GLOBECOM,* pp. 2315–2320, Nov. 1998.

[21] J. Yates, "Performance Analysis of Dynamically-Reconfigurable Wavelength-Division Multiplexed Networks," Ph.D. Thesis, The University of Melbourne, Australia, 1997. The thesis may be obtained at http://www.serc.rmit.edu.au/~jyates/papers

[22] A. Mokhtar and M. Azizoğlu, "Adaptive Wavelength Routing in All-Optical Networks," *IEEE/ACM Trans. on Networking,* vol. 6, no. 2, pp. 197–206, April 1998.

[23] E. Bouillet and K. Bala, "The Benefits of Wavelength Interchange in WDM Rings," *IEEE Int'l Conf. on Communications,* pp. 411–415, June 1997.

[24] K. Bala, E. Bouillet, and G. Ellinas, "Benefits of "Minimal" Wavelength Interchange in WDM Rings," *OSA/IEEE Optical Fiber Commun. Conf.,* pp. 120–121, Feb. 1997.

[25] G. Jeong and E. Ayanoglu, "Comparison of Wavelength-Interchanging and Wavelength-Selective Cross-Connects in Multiwavelength All-Optical Networks," *IEEE INFOCOM*, vol. 1, pp. 156–163, March 1996.

[26] S. Subramaniam and R. Barry, "Wavelength Assignment in Fixed Routing WDM Networks," *IEEE Int'l Conf. Commun.*, pp. 406–410, June 1997.

[27] J. Yates, M. Rumsewicz, J. Lacey and D. Everitt, "Modeling Blocking Probabilities in WDM Networks with Fixed Alternate Routing," *Int'l Conf. on Telecommun.*, pp. 1367–1372, April 1997.

[28] E. Karasan and E. Ayanoglu, "Effects of Wavelength Routing and Selec-

tion Algorithms on Wavelength Conversion Gain in WDM Optical Networks," *IEEE/ACM Trans. on Networking*, vol. 6, no. 2, pp. 186–196, April 1998.

[29] H. Harai, M. Murata and H. Miyahara, "Performance of Alternate Routing Methods in All-Optical Switching Networks," *IEEE INFOCOM*, pp. 517–525, April 1997.

[30] K. Chan and T. Yum, "Analysis of Least Congested Path Routing in WDM Lightwave Networks," IEEE INFOCOM, pp. 962–969, 1994.

[31] S. Ramamurthy and B. Mukherjee, "Fixed-Alternate Routing and Wavelength Conversion in Wavelength-Routed Optical Networks," *IEEE GLOBECOM*, pp. 2295–2302, Nov. 1998.

[32] S. Subramaniam, A. Somani, M. Azizoglu and R. Barry, "A Performance Model for Wavelength Conversion with Non-Poisson Traffic," *IEEE INFOCOM*, vol. 1, pp. 500–507, April 1997.

[33] J. Späth and S. Bodamer, "Routing of Dynamic Poisson and Non-Poisson Traffic in WDM Networks with Limited Wavelength Conversion," *European Conf. on Optical Commun.*, pp. 359–360, Sept. 1998.

[34] J. Yates, J. Lacey and D. Everitt, "Blocking in Multiwavelength TDM Networks," 4th Int'l Conf. on Telecommunication Systems, Modeling and Analysis, pp. 535–541, March 1996.

[35] I. Chlamtac, A. Ganz and G. Karmi, "Lightpath Communications : An Approach to High Bandwidth Optical WAN's," *IEEE Trans. on Commun.,* vol. 40, pp. 1171–1182, July 1992.

[36] I. Chlamtac, A. Ganz and G. Karmi, "Purely Optical Networks for Terabit Communication," *IEEE INFOCOM,* 1989.

[37] R. Cooper, *Introduction to Queueing Theory,* North Holland, 1981.

[38] K. Ross, *Multiservice Loss Models for Broadband Telecommunication Networks,* Springer, 1995.

[39] R. Ramaswami, "Optical Networking," Tutorial from IEEE INFOCOM, March 1996.

[40] C. Brackett, "Dense Wavelength Division Multiplexing Networks: Principles and Applications," *IEEE JSAC,* vol. 8, no. 6, pp. 948–964, August 1990.

[41] S. Subramaniam, M. Azizoglu and A. Somani, "On the Optimal Placement of Wavelength Converters in Wavelength-Routed Networks," *IEEE INFOCOM,* April 1998.

[42] W. Parys, B. Can Caenegem, B. Vandenberghe and P. Demeester, "Meshed Wavelength-Division Multiplexed Networks Partially Equipped with Wavelength Converters," *Optical Fiber Commun. Conf.,* pp. 359–360, Feb. 1998.

[43] J. Iness, "Efficient Use of Optical Components in WDM Networks," Ph.D. Thesis, University of California, Davis, 1997. The thesis may be obtained at http://networks.cs.ucdavis.edu/Pubs.html.

[44] M. Kovačević and A. Acampora, "Electronic Wavelength Translation in Optical Networks," *IEEE/OSA J. Lightwave Technol.,* vol. 14, no. 6, pp. 1161–1169, June 1996.

[45] J. Yates, J. Lacey, D. Everitt and M. Summerfield, "Limited-Range Wavelength Translation in All-Optical Networks," *IEEE INFOCOM,* vol. 3, pp. 954–961, March 1996.

[46] J. Yates, J. Lacey, M. Summerfield and D. Everitt, "Blocking in WDM Networks with Noisy Four-Wave Mixing Wavelength Converters," Australian Conf. on Optical Fiber Technology, pp. 9–12, Dec. 1996.

[47] O. Gerstel, G. Sasaki and R. Ramaswami, "Dynamic Channel Assignment for WDM Optical Networks with Little or No Wavelength Conversion," *34th Annual Allerton Conf. on Commun.,* Control and Computing, pp. 32–43, Oct. 1996.

[48] O. Gerstel, R. Ramaswami and G. Sasaki, "Benefits of Limited Wavelength Conversion in WDM Ring Networks," *Optical Fiber Commun. Conf.,* pp. 119–120, Feb. 1997.

[49] T. Durhuus, B. Mikkelsen, C. Joergensen, S. Danielsen and K. Stubkjaer, "All-Optical Wavelength Conversion by Semiconductor Optical Amplifiers," *IEEE/OSA J. of Lightwave Technol.,* vol. 14, no. 6, pp. 942–954, June 1996.

[50] J. Zhou, N. Park, K. Vahala, M. Newkirk and B. Miller, "Four-Wave Mixing Wavelength Conversion Efficiency in Semiconductor TravelingWave Amplifiers Measured to 65nm of Wavelength Shift," *IEEE Photon. Technol. Lett.,* vol. 6, no. 8, pp. 984–987, Aug. 1994.

[51] J. Zhou, N. Park, J. Dawson, K. Vahala, M. Newkirk and B. Miller, "Efficiency of Broadband Four-Wave Mixing Wavelength Conversion Using Semiconductor Traveling-Wave Amplifiers," *IEEE Photon. Technol. Lett.,* vol. 6, no. 1, pp. 50–52, Jan. 1994.

[52] J. Ryan, R. Kent, "WDM: North American Deployment Trends," *IEEE Commun. Mag.,* vol. 36, no. 2, pp. 40–44, Feb. 1998.

[53] K. Sato, S. Okamoto and H. Hadama, "Network Performance and Integrity Enhancement with Optical Path Layer Technologies," *IEEE JSAC,* vol. 12, no. 1, pp. 159–170, Jan. 1994.

[54] L. Lin, E. Karasan and R. Tkach, "Layered Switch Architecture for High Capacity Optical Transport Networks," *IEEE JSAC,* vol. 16, no. 17, pp. 1074–1080, Sept. 1998.

[55] G. Saussy, "Untangling the Web of Fiber-Optic Internet Backbone Approaches," *Lightwave,* pp. 61–66, Sept. 1998.

BIOGRAPHIES

JENNIFER YATES (jyates@serc.rmit.edu.au) received B.E. and B.Sc. degrees from the University of Western Australia in 1994, and a Ph.D. from the University of Melbourne, Australia in 1998. She is currently working with the Software Engineering Research Centre at the Royal Melbourne Institute of Technology University, Melbourne, Australia. Her research interests are in the performance engineering of optical networks and the internet.

JONATHAN LACEY received a Bachelor of Engineering from Melbourne University in 1987, and a Ph.D. from Cambridge University in 1993 for work on photosensitive gratings in optical fibers. Between 1993 and 1998, he was with the Australian Photonics Co-operative Research Centre's Photonics Research Laboratory at Melbourne University, where his research examined wavelength-division multiplexed (WDM) telecommunications networks, systems and devices. In September 1998 he joined Hewlett-Packard Laboratories in Palo Alto, California.

MICHAEL RUMSEWICZ received his Ph.D. in Applied Mathematics from the University of Adelaide in 1989 for his thesis in Applied Probability and Queueing Theory. From 1988 until 1994 he worked at Bellcore in the Network Performance and Control Group. In 1994 he joined the Software Engineering Research Centre of The Royal Melbourne Institute of Technology, where he leads the network performance research programme. His main research interests are in overload control for distributed systems, performance of Wavelength Division Multiplexing networks and scheduling algorithms for real-time applications.