Routing and Wavelength Assignment in Optical WDM Networks

George N. Rouskas
Department of Computer Science
North Carolina State University
Raleigh, NC 27695-7534
Phone: 919-515-3860
Fax: 919-515-7925
Email: rouskas@csc.ncsu.edu

Abstract

This article discusses the routing and wavelength assignment (RWA) problem in optical networks employing wavelenength division multiplexing (WDM) technology. Two variants of the problem are studied: static RWA, whereby the traffic requirements are known in advance, and dynamic RWA in which connection requests arrive in some random fashion. Both point-to-point and multicast traffic demands are considered.

Keywords: Wavelength division multiplexing (WDM), Optical networks, Routing and wavelength assignment (RWA), Virtual topology design, Optical multicast

1 Introduction to Optical WDM Networks

A basic property of single mode optical fiber is its enormous low-loss bandwidth of several tens of Terahertz. However, due to dispersive effects and limitations in optical device technology, single channel transmission is limited to only a small fraction of the fiber capacity. To take full advantage of the potential of fiber, the use of wavelength division multiplexing (WDM) technology has become the option of choice. With WDM, a number of distinct wavelengths are used to implement separate channels [1]. An optical fiber can carry several channels in parallel, each on a particular wavelength. The number of wavelengths that each fiber can carry simultaneously is limited by the physical characteristics of the fiber and the state of the optical technology used to combine these wavelengths onto the fiber and isolate them off the fiber. With currently available commercial technology, a few tens of wavelengths can be supported within the low-loss window at 1550 nm, but this number is expected to grow rapidly in the next few years. Therefore, optical fiber links employing WDM technology have the potential of delivering an aggregate throughput in the order of Terabits per second, enough to satisfy the ever-growing demand for more bandwidth per user on a sustained, long-term basis.

Unfortunately, due to the mismatch between aggregate fiber capacity and peak electronic processing speeds, simply upgrading existing point-to-point fiber links to WDM creates the well-known *electro-optic bottleneck* [2]: rather than achieving the multiterabit-per-second throughput of the fiber, one has to settle for the multigigabit-per-second throughput that can be expected of the electronic devices where the optical signals terminate. Overcoming the electro-optic bottleneck, therefore, involves the design of properly structured architectures to interconnect the fiber links. An optical WDM network is a network with optical fiber transmission links and with an architecture that is designed to exploit the unique features of fibers and WDM. Such networks offer the promise of an all-optical information highway capable of supporting a wide range of applications that involve the transport of massive amounts of data and/or require very fast response times. Such applications include video on demand and teleconferencing, telemedicine applications, multimedia document distribution, remote supercomputer visualization, and many more to come. Consequently, optical WDM networks have been a subject of extensive research both theoretically and experimentally [3, 4].

The architecture for wide-area WDM networks that is widely expected to form the basis for a future all-optical infrastructure is built on the concept of *wavelength routing*. A wavelength routing network, shown in Figure 1, consists of two types of nodes: *optical cross-connects (OXCs)*, which connect the fibers in the network, and *edge nodes* which provide the interface between non-optical

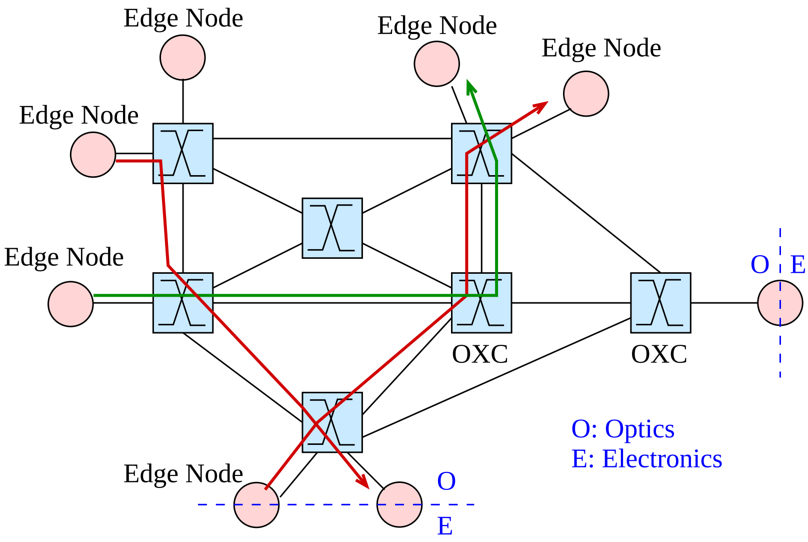

Figure 1: A wavelength routed WDM network

end systems (such as IP routers, ATM switches, or supercomputers) and the optical core. Access nodes provide the terminating points (sources and destinations) for the optical signal paths; the communication paths may continue outside the optical part of the network in electrical form.

The services that a wavelength routed network offers to end systems attached to edge nodes are in the form of *logical* connections implemented using *lightpaths*. Lightpaths (also referred to as λ-*channels*), are clear optical paths between two edge nodes, and are shown in Figure 1 as red and green directed lines. Information transmitted on a lightpath does not undergo any conversion to and from electrical form within the optical network, and thus, the architecture of the optical network nodes can be very simple because they do not need to do any signal processing. Furthermore, since a lightpath behaves as a literally transparent "clear channel" between the source and destination edge node, there is nothing in the signal path to limit the throughput of the fibers.

The OXCs provide the switching and routing functions for supporting the logical connections between edge nodes. An OXC takes in an optical signal at each of the wavelengths at an input port, and can switch it to a particular output port, independent of the other wavelengths. An OXC with N input and N output ports capable of handling W wavelengths per port can be thought of as W independent $N \times N$ switches. These switches have to be preceded by a wavelength demultiplexer and followed by a wavelength multiplexer to implement an OXC, as shown in Figure 2. Thus, an OXC can cross-connect the different wavelengths from the input to the output, where the connection pattern of each wavelength is independent of the others. By appropriately configuring the OXCs along the physical path, a logical connection (lightpath) may be established between any pair of

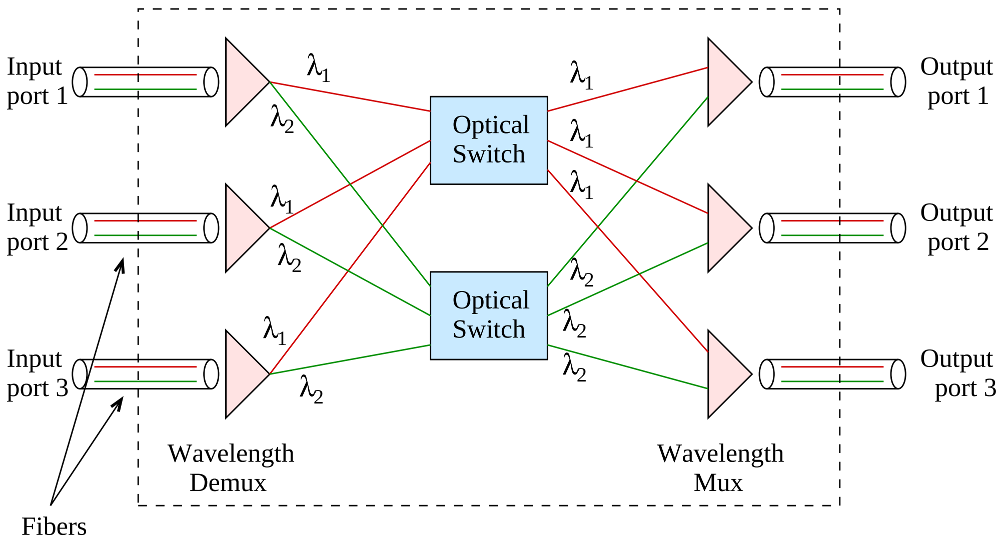

Figure 2: A 3×3 optical cross-connect (OXC) with two wavelengths per fiber

edge nodes.

A unique feature of optical WDM networks is the tight coupling between routing and wavelength selection. As can be seen in Figure 1, a lightpath is implemented by selecting a path of physical links between the source and destination edge nodes, and reserving a particular wavelength on each of these links for the lightpath. Thus, in establishing an optical connection we must deal with both routing (selecting a suitable path) and wavelength assignment (allocating an available wavelength for the connection). The resulting problem is referred to as the *routing and wavelength assignment (RWA)* problem [5], and is significantly more difficult than the routing problem in electronic networks. The additional complexity arises from the fact that routing and wavelength assignment are subject to the following two constraints:

1. *Wavelength continuity constraint:* a lightpath must use the same wavelength on all the links along its path from source to destination edge node. This constraint is illustrated in Figure 1 by representing each lightpath with a single color (wavelength) along all the links in its path.

2. *Distinct wavelength constraint:* all lightpaths using the same link (fiber) must be allocated distinct wavelengths. In Figure 1 this constraint is satisfied since the two lightpaths sharing a link are shown in different colors (wavelengths).

The RWA problem in optical networks is illustrated in Figure 3, where it is assumed that each fiber supports two wavelengths. The effect of the wavelength continuity constraint is represented

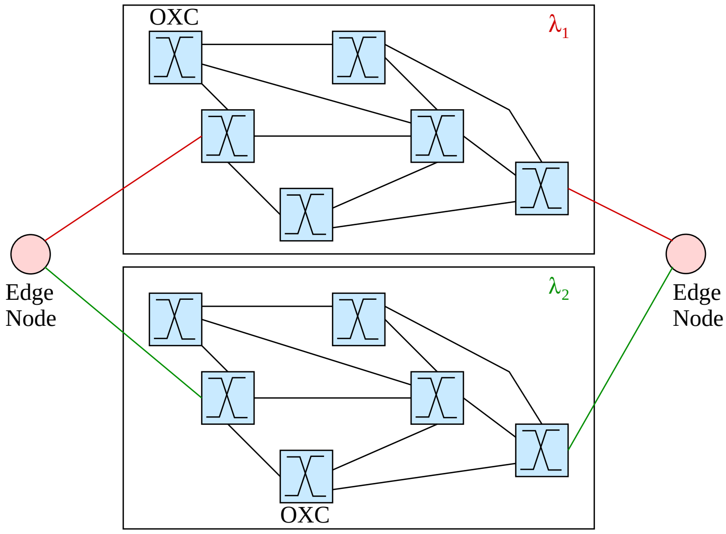

Figure 3: The RWA problem with two wavelengths per fiber

by replicating the network into as many copies as the number of wavelengths (in this case, two). If wavelength i is selected for a lightpath, the source and destination edge node communicate over the i-th copy of the network. Thus, finding a path for a connection may potentially involve solving W routing problems for a network with W wavelengths, one for each copy of the network.

The wavelength continuity constraint may be relaxed if the OXCs are equipped with *wavelength converters* [6]. A wavelength converter is a single input/output device that converts the wavelength of an optical signal arriving at its input port to a different wavelength as the signal departs from its output port, but otherwise leaves the optical signal unchanged. In OXCs without a wavelength conversion capability, an incoming signal at port p_i on wavelength λ can be optically switched to any port p_j, but must leave the OXC on the same wavelength λ. With wavelength converters, this signal could be optically switched to any port p_j on some other wavelength λ'. That is, wavelength conversion allows a lightpath to use different wavelengths along different physical links.

Different levels of wavelength conversion capability are possible. Figure 4 illustrates the differences for a single input and single output port situation; the case for multiple ports is more complicated but similar. *Full wavelength conversion* capability implies that any input wavelength may be converted to any other wavelength. *Limited wavelength conversion* [7] denotes that each input wavelength may be converted to any of a specific set of wavelengths, which is not the set of all wavelengths for at least one input wavelength. A special case of this is *fixed wavelength conversion*, where each input wavelength can be converted to exactly one other wavelength. If each wavelength is "converted" only to itself, then we have no conversion.

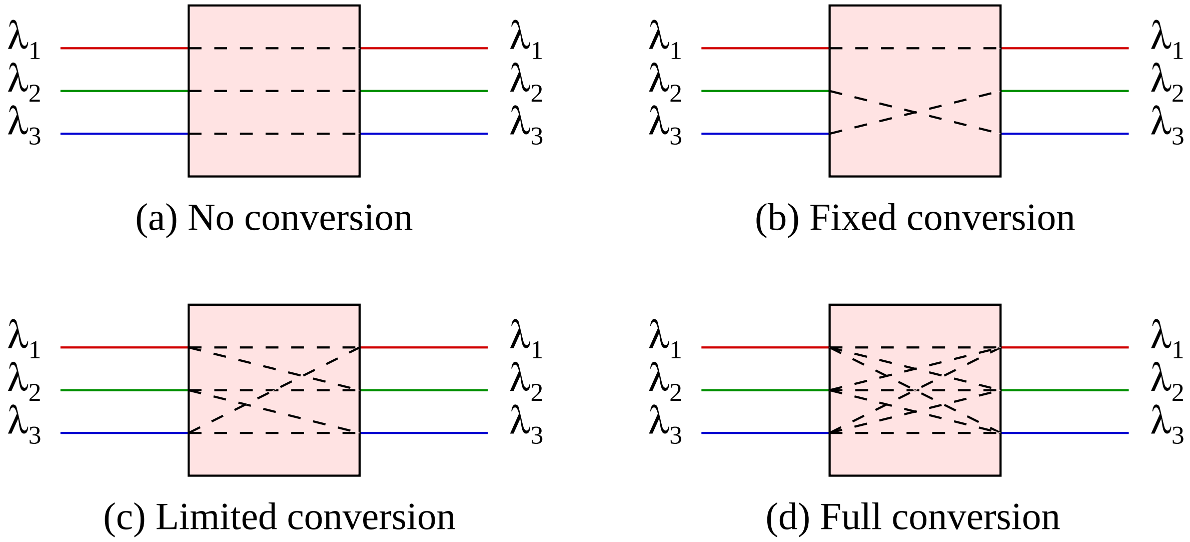

Figure 4: Wavelength conversion

The advantage of full wavelength conversion is that it removes the wavelength continuity constraint, making it possible to establish a lightpath as long as each link along the path from source to destination has a free wavelength (which could be different for different links). As a result, the RWA problem reduces to the classical routing problem, that is, finding a suitable path for each connection in the network. Referring to Figure 3, full wavelength conversion collapses the W copies of the network into a single copy on which the routing problem is solved. On the other hand, with limited conversion, the RWA problem becomes more complex than with no conversion. To see why, note that employing limited conversion at the OXCs introduces links between *some* of the network copies of Figure 3. For example, if wavelength λ_1 can be converted to wavelength λ_2 but not to wavelength λ_3, then links must be introduced from each OXC in copy 1 of the network to the corresponding OXC in copy 2, but not to the corresponding OXC in copy 3. When selecting a path for the connection, at each OXC there is the option of remaining at the same network copy or moving to another one, depending on the conversion capability of the OXC. Since the number of alternatives increases exponentially with the number of OXCs that need to be traversed, the complexity of the RWA problem increases accordingly.

Wavelength conversion (full or limited) increases the routing choices for a given lightpath (i.e., makes more efficient use of wavelengths), resulting in better performance. Since converter devices increase network cost, a possible middle ground is to use *sparse conversion*, that is, to employ converters in some, but not all, OXCs in the network. In this case, a lightpath must use the same wavelength along each link in a segment of its path between OXCs equipped with converters, but it may use a different wavelength along the links of another such segment. It has been shown that by implementing full conversion at a relatively small fraction of the OXCs in the network is sufficient to achieve almost all the benefits of conversion [8, 9].

Routing and wavelength assignment is the fundamental control problem in optical WDM networks. Since the performance of a network depends not only on its physical resources (e.g., OXCs, converters, fibers links, number of wavelengths per fiber, etc.) but also on how it is controlled, the objective of an RWA algorithm is to achieve the best possible performance within the limits of physical constraints. The RWA problem can be cast in numerous forms. The different variants of the problem, however, can be classified under one of two broad versions: a static RWA, whereby the traffic requirements are known in advance, and a dynamic RWA, in which a sequence of lightpath requests arrive in some random fashion. Sections 2 and 3 discuss the static and dynamic versions, respectively, of the RWA problem, and present some algorithms to solve them. Finally, Section 4 presents the multicast RWA problem and algorithms to build *light-trees* that connect a source edge node to multiple destinations.

2 Static Routing and Wavelength Assignment

If the traffic patterns in the network are reasonably well-known in advance and any traffic variations take place over long time scales, the most effective technique for establishing optical connections (lightpaths) between edge nodes is by formulating and solving a static RWA problem. For example, static RWA is appropriate for provisioning a set of semipermanent connections. Since these connections are assumed to remain in place for relatively long periods of time, it is worthwhile to attempt to optimize the way in which network resources (e.g., physical links and wavelengths) are assigned to each connection, even though optimization may require a considerable computational effort.

A solution to the static RWA problem consists of a set of long-lived lightpaths which create a *logical* (or *virtual*) topology among the edge nodes. This virtual topology is embedded onto the physical topology of optical fiber links and OXCs. Accordingly, the static RWA problem is often referred to as the *virtual topology design* problem [10]. In the virtual topology, there is a directed link from edge node s to edge node d if a lightpath originating at s and terminating at d is set up (refer also to Figure 1), and edge node s is said to be “one hop away” from edge node d in the virtual topology, although the two nodes may be separated by a number of physical links. The type of virtual topology that can be created is usually constrained by the underlying physical topology. In particular, it is generally not possible to implement fully connected virtual topologies: for N

edge nodes this would require each edge node to maintain $N-1$ lightpaths and the optical network to support a total of $N(N-1)$ lightpaths. Even for modest values of N, this degree of connectivity is beyond the reach of current optical technology, both in terms of the number of wavelengths that can be supported and in terms of the optical hardware (transmitters and receivers) required at each edge node.

In its most general form, the RWA problem is specified by providing the physical topology of the network and the traffic requirements. The physical topology corresponds to the deployment of cables in some existing fiber infrastructure, and is given as a graph $G_p(V, E_p)$, where V is the set of OXCs and E_p is the set of fibers that interconnect them. The traffic requirements are specified in a traffic matrix $\mathbf{T} = [\rho p_{sd}]$, where ρp_{sd} is a measure of the long-term traffic flowing from source edge node s to destination edge node d [11]. Quantity ρ represents the (deterministic) total offered load to the network, while the p_{sd} parameters define the distribution of the offered traffic.

Routing and wavelength assignment are considered together as an optimization problem using mixed integer programming (MIP) formulations. Usually, the objective of the formulation is to minimize the maximum congestion level in the network subject to network resource constraints [10, 12]. While other objective functions are possible, such as minimizing the average weighted number of hops or minimizing the average packet delay, minimizing network congestion is preferable since it can lead to linear programming (MILP) formulations. While we do not present the RWA problem formulation here, the interested reader may refer to [11, 10, 12]. These formulations turn out to have extremely large numbers of variables, and are intractable for large networks. This fact has motivated the development of heuristic approaches for finding good solutions efficiently.

Before we describe the various heuristic approaches, we note that the static RWA problem can be logically decomposed into four subproblems. The decomposition is approximate or inexact, in the sense that solving the subproblems in sequence and combining the solutions may not result in the optimal solution for the fully integrated problem, or some later subproblem may have no solution given the solution obtained for an earlier subproblem, so no solution to the original problem may be obtained. However, the decomposition provides insight into the structure of the RWA problem and is a first step towards the design of effective heuristics. Assuming no wavelength conversion, the subproblems are as follows.

1. **Topology Subproblem:** Determine the logical topology to be imposed on the physical topology, that is, determine the lightpaths in terms of their source and destination edge nodes.

2. **Lightpath Routing Subproblem:** Determine the physical links which each lightpath con-

sists of, that is, route the lightpaths over the physical topology.

3. **Wavelength Assignment Subproblem:** Determine the wavelength each lightpath uses, that is, assign a wavelength to each lightpath in the logical topology so that wavelength restrictions are obeyed for each physical link.

4. **Traffic Routing Subproblem:** Route packet traffic between source and destination edge nodes over the logical topology obtained.

A large number of heuristic algorithms have been developed in the literature to solve the general static RWA problem discussed here or its many variants. Overall, however, the different heuristics can be classified into three broad categories: (1) algorithms which solve the overall MILP problem sub-optimally, (2) algorithms which tackle only a subset of the four subproblems, and (3) algorithms which address the problem of embedding regular logical topologies onto the physical topology.

Suboptimal solutions can be obtained by applying classical tools developed for complex optimization problems directly to the MILP problem. One technique is to use LP-relaxation followed by rounding [13]. In this case, the integer constraints are relaxed creating a non-integral problem which can be solved by some linear programming method, and then a rounding algorithm is applied to obtain a new solution which obeys the integer constraints. Alternatively, genetic algorithms or simulated annealing [14] can be applied to obtain locally optimal solutions. The main drawback of these approaches is that it is difficult to control the quality of the final solution for large networks: simulated annealing is computationally expensive and thus, it may not be possible to adequately explore the state space, while LP-relaxation may lead to solutions from which it is difficult to apply rounding algorithms.

Another class of algorithms tackles the RWA problem by initially solving the first three subproblems listed above; traffic routing is then performed by employing well-known routing algorithms on the logical topology. One approach for solving the three subproblems is to maximize the amount of traffic that is carried on one-hop lightpaths, i.e., traffic that is routed from source to destination edge node directly on a lightpath. A greedy approach taken in [15] is to create lightpaths between edge nodes in order of decreasing traffic demands as long as the wavelength continuity and distinct wavelength constraints are satisfied. This algorithm starts with a logical topology with no links (lightpaths) and sequentially adds lightpaths as long as doing so does not violate any of the problem constraints. The reverse approach is also possible [16]: starting with a fully connected logical topology, an algorithm sequentially removes the lightpath carrying the smallest traffic flows until no constraint is violated. At each step (i.e., after removing a lightpath), the traffic routing subproblem is solved in order to find the lightpath with smallest the flow.

The third approach to RWA is to start with a given logical topology, thus avoiding to directly solve the first of the four subproblems listed above. Regular topologies are good candidates as logical topologies since they are well understood and results regarding bounds and averages (e.g., for hop lengths) are easier to derive. Algorithms for routing traffic on a regular topology are usually simple, so the traffic routing subproblem can be trivially solved. Also, regular topologies possess inherent load balancing characteristics which are important when the objective is to minimize the maximum congestion.

Once a regular topology is decided on as the one to implement the logical topology, it remains to decide which physical node will realize each given node in the regular topology (this is usually referred to as the *node mapping* subproblem), and which sequence of physical links will be used to realize each given edge (lightpath) in the regular topology (this *path mapping* subproblem is equivalent to the lightpath routing and wavelength assignment subproblems discussed earlier). This procedure is usually referred to embedding a regular topology in the physical topology. Both the node and path mapping subproblems are intractable, and heuristics have been proposed in the literature [16, 17]. For instance, a heuristic for mapping the nodes of shuffle topologies based on the gradient algorithm was developed in [17].

Given that all the algorithms for the RWA problem are based on heuristics, it is important to be able to characterize the quality of the solutions obtained. To this end, one must resort to comparing the solutions to known bounds on the optimal solution. A comprehensive discussion of bounds for the RWA problem and the theoretical considerations involved in deriving them can be found in [10]. A simulation-based comparison of the relative performance of the three classes of heuristic for the RWA problem is presented in [12]. The results indicate that the second class of algorithms discussed earlier achieve the best performance.

3 Dynamic Routing and Wavelength Assignment

Under a dynamic traffic scenario, edge nodes submit to the network requests for lightpaths to be set up as needed. Thus, connection requests are initiated in some random fashion. Depending on the state of the network at the time of a request, the available resources may or may not be sufficient to establish a lightpath between the corresponding source-destination edge node pair. The network state consists of the physical path (route) and wavelength assignment for all active lightpaths. The state evolves randomly in time as new lightpaths are admitted and existing lightpaths are released. Thus, each time a request is made, an algorithm must be executed in real time to determine whether it is feasible to accommodate the request, and, if so, to perform routing and wavelength assignment.

If a request for a lightpath cannot be accepted because of lack of resources, it is blocked.

Because of the real-time nature of the problem, RWA algorithms in a dynamic traffic environment must be very simple. Since combined routing and wavelength assignment is a hard problem, a typical approach to designing efficient algorithms is to decouple the problem into two separate subproblems: the routing problem and the wavelength assignment problem. Consequently, most dynamic RWA algorithms for wavelength routed networks consist of the following general steps:

1. Compute a number of candidate physical paths for each source-destination edge node pair and arrange them in a path list.
2. Order all wavelengths in a wavelength list.
3. Starting with the path and wavelength at the top of the corresponding list, search for a feasible path and wavelength for the requested lightpath.

The specific nature of a dynamic RWA algorithm is determined by the number of candidate paths and how they are computed, the order in which paths and wavelengths are listed, and the order in which the path and wavelength lists are accessed.

Let us first discuss the routing subproblem. If a *static* algorithm is used, the paths are computed and ordered independently of the network state. With an *adaptive* algorithm, on the other hand, the paths computed and their order may vary according to the current state of the network. A static algorithm is executed off-line and the computed paths are stored for later use, resulting in low latency during lightpath establishment. Adaptive algorithms are executed at the time a lightpath request arrives and require network nodes to exchange information regarding the network state. Lightpath set up delay may also increase, but in general adaptive algorithms improve network performance.

The number of path choices for establishing an optical connection is another important parameter. A *fixed* routing algorithm is a static algorithm in which every source-destination edge node pair is assigned a single path. With this scheme, a connection is blocked if there is no wavelength available on the designated path at the time of the request. In *fixed-alternate* routing, a number $k, k > 1$, of paths are computed and ordered off-line for each source-destination edge node pair. When a request arrives, these paths are examined in the specified order and the first one with a free wavelength is used to establish the lightpath. The request is blocked if no wavelength is available in any of the k paths. Similarly, an adaptive routing algorithm may compute a single path, or a number of alternate paths at the time of the request. A hybrid approach is to compute

k paths off-line, however, the order in which the paths are considered is determined according to the network state at the time the connection request is made (e.g., least to most congested).

In most practical cases, the candidate paths for a request are considered in increasing order of *path length.* Path length is typically defined as the sum of the weights assigned to each physical link along the path, and the weights are chosen according to some desirable routing criterion. Since weights can be assigned arbitrarily, they offer a wide range of possibilities for selecting path priorities. For example, in a static (fixed-alternate) routing algorithm, the weight of each link could be set to 1, or to the physical distance of the link. In the former case, the path list consists of the k minimum-hop paths, while in the latter the candidate paths are the k minimum-distance paths (where distance is defined as the geographic length). In an adaptive routing algorithm, link weights may reflect the load or "interference" on a link (i.e., the number of active lightpaths sharing the link). By assigning small weights to least loaded links, paths with larger number of free channels on their links rise to the head of the path list, resulting in a *least loaded* routing algorithm. Paths that are congested become "longer" and are moved further down the list; this tends to avoid heavily loaded bottleneck links. Many other weighting functions are possible.

When path lengths are sums of of link weights, the k-shortest path algorithm [18] can be used to compute candidate paths. Each path is checked in order of increasing length, and the first that is feasible is assigned the first free wavelength in the wavelength list. However, the k shortest paths constructed by this algorithm usually share links. Therefore, if one path in the list is not feasible, it is likely that other paths in the list with which it shares a link will also be infeasible. To reduce the risk of blocking, the k shortest paths can be computed so as to be pairwise link-disjoint. This can be accomplished as follows: when computing the i-th shortest path, $i = 1, \cdots, k$, the links used by the first $i-1$ paths are removed from the original network topology and Dijkstra's shortest path algorithm [19] is applied to the resulting topology. This approach increases the chances of finding a feasible path for a connection request.

Let us now discuss the wavelength assignment subproblem which is concerned with the manner in which the wavelength list is ordered. For a given candidate path, wavelengths are considered in the order in which they appear in the list to find a free wavelength for the connection request. Again, we distinguish between the static and adaptive cases. In the static case, the wavelength ordering is fixed (e.g., the list is ordered by wavelength number). The idea behind this scheme, also referred to as *first-fit*, is to pack all the in-use wavelengths towards the top of the list so that wavelengths towards the end of the list will have higher probability of being available over long continuous paths. In the adaptive case, the ordering of wavelengths is typically based on usage. Usage can be defined either as the number of links in the network in which a wavelength is currently used, or as the number of active connections using a wavelength. Under the *max-reuse* method, the most used wavelengths are considered first (i.e., wavelength are considered in order of decreasing usage). The rationale behind this method is to reuse active wavelengths as much as possible before trying others, packing connections into fewer wavelengths and conserving the spare capacity of less-used wavelengths. This in turn makes it more likely to find wavelengths that satisfy the continuity requirement over long paths. Under the *min-reuse* method, wavelengths are tried in the order of increasing usage. This scheme attempts to balance the load as equally as possible among all the available wavelengths. However, min-reuse assignment tends to "fragment" the availability of wavelengths, making it less likely that the same wavelength is available throughout the network for connections that traverse longer paths.

The max-reuse and min-reuse schemes introduce communication overhead because they require global network information in order to compute the usage of each wavelength. The first-fit scheme, on the other hand, requires no global information, and since it does not need to order wavelengths in real-time, it has significantly lower computational requirements than either max-reuse or min-reuse. Another adaptive scheme that avoids the communication and computational overhead of max-reuse and min-reuse is *random* wavelength assignment. With this scheme, the set of wavelengths that are free on a particular path is first determined. Among the available wavelengths, one is chosen randomly (usually with uniform probability) and assigned to the requested lightpath.

We note that in networks in which all OXCs are capable of wavelength conversion, the wavelength assignment problem is trivial: since a lightpath can be established as long as at least one wavelength is free at each link and different wavelengths can be used in different links, the order in which wavelengths are assigned is not important. On the other hand, when only a fraction of the OXCs employ converters (i.e., a sparse conversion scenario), a wavelength assignment scheme is again required to select a wavelength for each segment of a connection's path that originates and terminates at an OXC with converters. In this case, the same assignment policies discussed above for selecting a wavelength for the end-to-end path can also be used to select a wavelength for each path segment between OXCs with converters.

The performance of a dynamic RWA algorithm is generally measured in terms of the call blocking probability, that is, the probability that a lightpath cannot be established in the network due to lack of resources (e.g., link capacity or free wavelengths). Even in the case of simple network topologies (such as rings) or simple routing rules (such as fixed routing), the calculation of blocking probabilities in WDM networks is extremely difficult. In networks with arbitrary mesh topologies, and/or when using alternate or adaptive routing algorithms, the problem is even more complex. These complications arise from both the link load dependencies (due to interfering lightpaths) and the dependencies among the sets of active wavelengths in adjacent links (due to the wavelength continuity constraint). Nevertheless, the problem of computing blocking probabilities in wavelength routed networks has been extensively studied in the literature, and approximate analytical techniques which capture the effects of link load and wavelength dependencies have been developed in [8, 9, 20]. A detailed comparison of the performance of various wavelength assignment schemes in terms of call blocking probability can be found in [21].

Though important, average blocking probability (computed over all connection requests) does not always capture the full effect of a particular dynamic RWA algorithm on other aspects of network behavior, in particular, *fairness*. In this context, fairness refers to the variability in blocking probability experienced by lightpath requests between the various edge node pairs, such that lower variability is associated with a higher degree of fairness. In general, any network has the property that longer paths are likely to experience higher blocking than shorter ones. Consequently, the degree of fairness can be quantified by defining the *unfairness factor* as the ratio of the blocking probability on the longest path to that on the shortest path for a given RWA algorithm. Depending on the network topology and the RWA algorithm, this property may have a cascading effect which can result in an unfair treatment of the connections between more distant edge node pairs: blocking of long lightpaths leaves more resources available for short lightpaths, so that the connections established in the network tend to be short ones. These shorter connections "fragment" the availability of wavelengths, and thus, the problem of unfairness is more pronounced in networks without converters since finding long paths that satisfy the wavelength continuity constraint is more difficult than without this constraint.

Several studies [8, 9, 20] have examined the influence of various parameters on blocking probability and fairness, and some of the general conclusions include the following:

- Wavelength conversion significantly affects fairness. Networks employing converters at all OXCs sometimes exhibit orders of magnitude improvement in fairness (as reflected by the unfairness factor) compared to networks with no conversion capability, despite the fact that the improvement in overall blocking probability is significantly less pronounced. It has also been shown that equipping a relatively small fraction (typically, 20-30%) of all OXCs with converters is sufficient to achieve most of the fairness benefits due to wavelength conversion.
- Alternate routing can significantly improve the network performance in terms of both overall blocking probability and fairness. In fact, having as few as three alternate paths for each connection may in some cases (depending on the network topology) achieve almost all the benefits (in terms of blocking and fairness) of having full wavelength conversion at each OXC

with fixed routing.

- Wavelength assignment policies also play an important role, especially in terms of fairness. The random and min-reuse schemes tend to "fragment" the wavelength availability, resulting in large unfairness factors (with min-reuse having the worst performance). On the other hand, the max-reuse assignment policy achieves the best performance in terms of fairness. The first-fit scheme exhibits a behavior very similar to max-reuse in terms of both fairness and overall blocking probability, and has the additional advantage of being easier and less expensive to implement.

4 Multicast Routing and Wavelength Assignment

In Sections 2 and 3, we considered static and dynamic RWA algorithms, respectively, for establishing lightpaths in optical networks. In [22], the concept of a lightpath was generalized into that of a *light-tree*, which, like a lightpath, is a clear channel originating at given source node and implemented with a single wavelength. But unlike a lightpath, a light-tree has multiple destination nodes, hence it is a point-to-multipoint channel. The physical links implementing a light-tree form a tree, rooted at the source node, rather than a path in the physical topology, hence the name. The study in [22] focused on virtual topology design (i.e., static RWA) for point-to-point traffic and observed that, since a light-tree is a more general representation of a lightpath, the set of virtual topologies that can be implemented using light-trees is a superset of the virtual topologies that can be implemented only using lightpaths. Thus, for any given virtual topology problem, an optimal solution using light-trees is guaranteed to be at least as good and possibly an improvement over the optimal solution obtained using only lightpaths. Furthermore, it was demonstrated that by extending the lightpath concept to a light-tree, the network performance (in terms of average packet hops) can be improved while the network cost (in terms of the number of optical transmitters/receivers required) decreases.

Light-trees are implemented by employing optical devices known as *power splitters* [2] at the OXCs. A power splitter has the ability to split an incoming signal, arriving at some wavelength λ, into up to m outgoing signals, $m \geq 2$; m is referred to as the *fanout* of the power splitter. Each of these m signals is then independently switched to a different output port of the OXC. Note that due to the splitting operation and associated losses, the optical signals resulting from the splitting of the original incoming signal must be amplified before leaving the OXC. Also, to ensure the quality of each outgoing signal, the fanout m of the power splitter may have to be limited to a small number. If the OXC is also capable of wavelength conversion, each of the m outgoing signal may be shifted, independently of the others, to a wavelength different than the incoming wavelength λ. Otherwise, all m outgoing signals must be on the same wavelength λ.

While [22] considered mainly point-to-point traffic, another attractive feature of light-trees is the inherent capability for performing multicasting in the optical domain (as opposed to performing multicasting at a higher layer, e.g., the network layer, which requires electro-optic conversion). Such wavelength routed light-trees are useful for transporting high-bandwidth, real-time applications such as high-definition TV (HDTV). Therefore, OXCs equipped with power splitters will be referred to as *multicast capable* OXCs (MC-OXCs). Note that, just like with converter devices, incorporating power splitters within an OXC is expected to increase the network cost because of the large amount of power amplification and the difficulty of fabrication.

With the availability of MC-OXCs and the existence of multicast traffic demands, the problem of establishing light-trees to satisfy these demands arises. We will call this problem the *multicast routing and wavelength assignment (MC-RWA)* problem. MC-RWA bears many similarities to the RWA problem discussed in Sections 2 and 3. Specifically, the tight coupling between routing and wavelength assignment remains, and even becomes stronger: in the absence of wavelength conversion the same wavelength must be used by the multicast connection not just along the links of a single path but along the links of the whole light-tree. Since the construction of optimal trees for routing multicast connections is by itself a hard problem [23], the combined MC-RWA problem becomes even harder. Depending on the nature of traffic demands, we also distinguish between static and dynamic MC-RWA problems. As we already know, optimal solutions for the point-to-point RWA problems are not practically obtainable, and with a more general construct (the light-tree) and hence a much larger search space, this is going to be even more true for the MC-RWA problems. In general, the approaches to tackling the static and dynamic MC-RWA problems are similar to the ones we described for the static and dynamic RWA problems, respectively. The challenge in this case is to design heuristics that can cope with the increased complexity of the problem and yet produce good solutions. In the following we summarize the recent work on multicasting in optical networks, but the reader should keep in mind that this is an area of current research.

The benefits of multicasting in wavelength routed optical networks were first demonstrated in [24]. Specifically, it was shown that using light-trees (spanning the source and destination nodes) rather than individual parallel lightpaths (each connecting the source to an individual destination) requires fewer wavelengths and consumes a significantly lower amount of bandwidth. In [25] both the static and the dynamic MC-RWA problems were studied. A MILP formulation that maximizes the total number of multicast connections was presented for the static MC-RWA problem. Rather than providing heuristic algorithms for solving the MILP, bounds on the objective function were presented by relaxing the integer constraints. The dynamic MC-RWA problem, on the other hand, was solved by decoupling the routing and wavelength assignment problems. A number of *alternate* trees was constructed for each multicast connection using existing routing algorithms. When a request for a connection arrives, the associated trees are considered in a fixed order. For each tree, wavelengths are also considered in a fixed order (i.e., the first-fit strategy). The connection is blocked if no free wavelength is found in any of the trees associated with the multicast connection.

Finally, the problem of constructing trees for routing multicast connections was studied in [26] independently of wavelength assignment, under the assumption that not all OXCs are multicast capable, i.e., that there is a limited number of MC-OXCs in the network. Four new algorithms were developed for routing multicast connections under this *sparse light splitting* scenario. While the algorithms differ slightly from each other, the main idea to accommodate sparse splitting is to start with the assumption that all OXCs in the network are multicast capable and use an existing algorithm to build an initial tree. Such a tree is infeasible if a non-multicast-capable OXC is a branching point. In this case, all but one branches out of this OXC are removed, and destination nodes in the removed branches have to join the tree at a MC-OXC.

References

[1] T. E. Stern and K. Bala. *Multiwavelength Optical Networks*. Prentice Hall, Upper Saddle River, New Jersey, 2000.

[2] B. Mukherjee. *Optical Communication Networking*. McGraw-Hill, 1997.

[3] O. Gerstel, B. Li, A. McGuire, G. N. Rouskas, K. Sivalingam, and Z. Zhang (Eds.). Special issue on protocols and architectures for next generation optical WDM networks. *IEEE Journal Selected Areas in Communications*, 18(10), October 2000.

[4] G.-K. Chung, K.-I. Sato, and D. K. Hunter (Eds.). Special issue on optical networks. *Journal of Lightwave Technology*, 18(12), December 2000.

[5] H. Zang, J. P. Jue, and B. Mukherjee. A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks. *Optical Networks*, 1(1):47–60, January 2000.

[6] B. Ramamurty and B. Mukherjee. Wavelength conversion in WDM networking. *IEEE Journal Selected Areas in Communications*, 16(7):1061–1073, September 1998.

[7] V. Sharma and E. A. Varvarigos. Limited wavelength translation in all-optical WDM mesh networks. In *Proceedings of INFOCOM '98*, pages 893–901. IEEE, March 1999.

[8] Y. Zhu, G. N. Rouskas, and H. G. Perros. A path decomposition algorithm for computing blocking probabilities in wavelength routing networks. *IEEE/ACM Transactions on Networking*, 8(6):747–762, December 2000.

[9] S. Subramaniam, M. Azizoglu, and A. Somani. All-optical networks with sparse wavelength conversion. *IEEE/ACM Transactions on Networking*, 4(4):544–557, August 1996.

[10] R. Dutta and G. N. Rouskas. A survey of virtual topology design algorithms for wavelength routed optical networks. *Optical Networks Magazine*, 1(1):73–89, January 2000.

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

[12] E. Leonardi, M. Mellia, and M. A. Marsan. Algorithms for the logical topology design in WDM all-optical networks. *Optical Networks*, 1(1):35–46, January 2000.

[13] D. Banerjee and B. Mukherjee. A practical approach for routing and wavelength assignment in large wavelength-routed optical networks. *IEEE Journal Selected Areas in Communications*, 14(5):903–908, June 1996.

[14] B. Mukherjee *et al.* Some principles for designing a wide-area WDM optical network. *IEEE/ACM Transactions on Networking*, 4(5):684–696, October 1996.

[15] Z. Zhang and A. Acampora. A heuristic wavelength assignment algorithm for multihop WDM networks with wavelength routing and wavelength reuse. *IEEE/ACM Transactions on Networking*, 3(3):281–288, June 1995.

[16] I. Chlamtac, A. Ganz, and G. Karmi. Lightnets: Topologies for high-speed optical networks. *Journal of Lightwave Technology*, 11:951–961, May/June 1993.

[17] S. Banerjee and B. Mukherjee. Algorithms for optimized node placement in shufflenet-based multihop lightwave networks. In *Proceedings of INFOCOM '93*. IEEE, March 1993.

[18] E. Lawler. *Combinatorial Optimization: Networks and Matroids*. Holt, Rinehart and Winston, 1976.

[19] D. Bertsekas and R. Gallager. *Data Networks*. Prentice Hall, Inc., Englewood Cliffs, NJ, 1992.

[20] E. Karasan and E. Ayanoglu. Effects of wavelength routing and selection algorithms on wavelength conversion gain in WDM optical networks. *IEEE/ACM Transactions on Networking*, 6(2):186–196, April 1998.

[21] Y. Zhu, G. N. Rouskas, and H. G. Perros. A comparison of allocation policies in wavelength routing networks. *Photonic Network Communications*, 2(3):265–293, August 2000.

[22] L. H. Sahasrabuddhe and B. Mukherje. Light-trees: Optical multicasting for improved performance in wavelength-routed networks. *IEEE Communications*, 37(2):67–73, February 1999.

[23] S. L. Hakimi. Steiner's problem in graphs and its implications. *Networks*, 1:113–133, 1971.

[24] R. Malli, X. Zhang, and C. Qiao. Benefit of multicastsing in all-optical networks. In *Proceedings of SPIE*, volume 3531, pages 209–220, November 1998.

[25] G. Sahin and M. Azizoglu. Multicast routing and wavelength assignment in wide-area networks. In *Proceedings of SPIE*, volume 3531, pages 196–208, November 1998.

[26] X. Zhang, J. Y. Wei, and C. Qiao. Constrained multicast routing in WDM networks with sparse light splitting. *Journal of Lightwave Technology*, 18(12):1917–1927, December 2000.