Optical Layer Multicast: Rationale, Building Blocks, and Challenges

George N. Rouskas, North Carolina State University

Abstract

Optical layer multicast refers to the support of point-to-multipoint connections directly at the physical layer by employing passive devices capable of splitting the power of an input signal among several outputs. Optical multicast technology will enable a broad set of applications and will open new directions to network design. In this article we examine the underlying principles and essential components for a practical optical multicast service. We also present a set of key research challenges along with a survey of the literature.

Over the last decade, optical wavelength division multiplexing (WDM) transmission systems have been deployed in the Internet infrastructure to meet the ever-increasing demand for bandwidth. WDM technology was initially introduced as a means to increase the capacity of existing point-to-point fiber links. The result has been the creation of *opaque* optical networks in which the optical signal undergoes *optical-electronic-optical* (OEO) conversion or regeneration at every intermediate node in the network. As more advanced devices, such as *optical crossconnects* (OXCs), mature and become commercially available, it is possible to design and deploy *transparent* optical networks in which there is no signal regeneration or OEO conversion at intermediate nodes. An OXC can switch an optical signal arriving on a wavelength of an input fiber link to the same wavelength on an output fiber link. The main mechanism of transport in such a network is the *lightpath*, an optical channel on a particular wavelength that may span a number of fiber links (physical hops). Transparent networks are bit-rate-, protocol-, and format-independent, and can take full advantage of the available bandwidth by eliminating the need for per-hop packet forwarding.

In [1], the concept of a lightpath was generalized to that of a *light-tree*, which, like a lightpath, is a clear channel originating at a 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. As shown in Fig. 1, 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. Light-trees may be implemented by employing optical devices known as *power splitters* at the OXCs; we will refer to OXCs equipped with power splitters as *multicast-capable* OXCs (MC-OXCs). A power splitter has the ability to split an incoming signal, arriving at some wavelength λ, into up to n outgoing signals, $n \geq 2$; n is referred to as the *fanout* of the power splitter. Each of these n signals is then independently switched to a different output port of the MC-OXC. 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 MC-OXC. Also, to ensure the quality of each outgoing signal, the maximum fanout n of the power splitter may have to be limited to a small integer. If the OXC is also capable of wavelength conversion, each of the n outgoing signals may be shifted, independent of the others, to a wavelength different than the incoming wavelength λ. Otherwise, all n outgoing signals will be on the same wavelength λ. Note that incorporating power splitters within an OXC increases the network cost because of the large amount of power amplification and the difficulty of fabrication.

In this article we take a close look at the enabling technologies, rationale, and issues facing the provision of light-tree services in mesh WDM networks; the reader is referred to [2] for a survey of multicasting in shared-medium optical networks based on a broadcast star or bus topology. We present a sampling of the applications that can benefit from a light-tree service, and examine the building blocks for implementing such a service, in terms of both relevant hardware devices, and the algorithmic and protocol issues involved. We outline a broad set of challenges in the design and control of multicast optical networks, survey the recent literature, and offer a set of directions for future research. We then conclude the article.

Light-Tree Applications

A light-tree service has many applications and can offer significant advantages in several critical aspects of optical networks. The following list is not meant to be all-inclusive; rather, it is intended to illustrate the practical importance of light-trees in emerging transparent networks.

Optical multicast. An attractive feature of light-trees is the inherent ability to perform multicast in the optical domain. While performing multicast at a higher layer (e.g., the network layer) requires OEO conversion, native optical multicast improves performance by eliminating the need for the store-and-forward functions of packet-switched technology. Therefore, light-trees can be useful for transporting high-bandwidth real-time streams such as high-definition TV (HDTV), as well as for optical storage area networks (O-SANs). We note that TV signals are currently carried over distribution networks having a tree-like *physical* topology [3]; creating a *logical* tree topology (light-tree) over an arbitrary physical topology for

This material is based on work supported by the National Science Foundation under grant ANI-9805016.

the distribution of similar applications would be a natural next step.

Enhanced virtual connectivity. In opaque networks, the virtual degree of connectivity of each node is not tied to the number of its interfaces: electronic routing creates the illusion that a node can reach any other node in the network. In transparent networks, on the other hand, the degree of connectivity of each client node (e.g., IP router) connected to the optical core is limited by its physical degree (i.e., the number of its optical transceivers). Subject to wavelength availability, a light-tree service would enable a client node to reach a larger number of other client nodes than would be allowed by its physical degree, enhancing the virtual connectivity of the network.

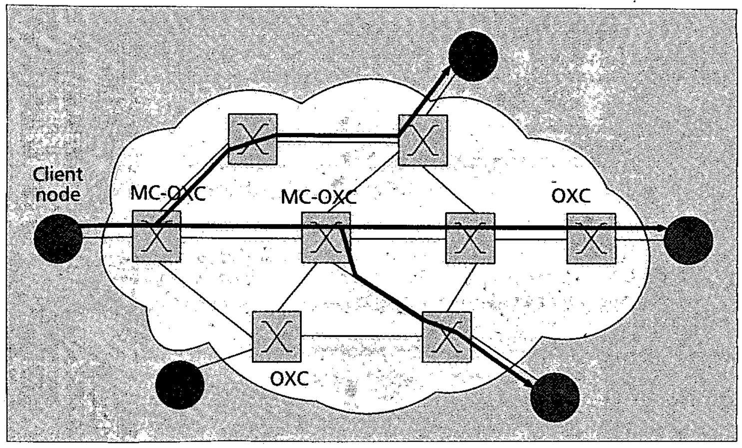

■ Figure 1. *A light-tree in a transparent optical network.*

Traffic grooming. Generalized multiprotocol label switching (GMPLS) makes it possible to tunnel a set of MPLS label-switched paths (LSPs) over a wavelength channel. Since switching at OXCs takes place at the granularity of a whole wavelength, a point-to-point lightpath allows the sharing of the wavelength bandwidth only among traffic streams between the same ingress and egress OXCs. The light-tree concept offers a way to overcome this constraint, since it allows for the grooming and tunneling of a number of lower-rate point-to-point LSPs to several destinations, regardless of the egress OXC to which these destinations attach.

1+1 optical layer protection. A protected lightpath consists of a *primary* (working) route and a dedicated, diversely routed *backup* (protection) route. Under 1+1 protection, the source signal is transmitted simultaneously on both the primary and backup paths. Upon failure of the primary path, the destination simply switches to the backup path to receive the signal. An efficient implementation of 1+1 protection would be to have the primary and backup paths be part of a light-tree rooted at the ingress MC-OXC.

Improved optical network performance. It was pointed out in [1] that since a light-tree is a more general construct than 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 using lightpaths only. Thus, for any given virtual topology problem, an optimal solution using light-trees is guaranteed to be at least as good as, and possibly an improvement over, the optimal solution obtained using only lightpaths. Consequently, incorporating the light-tree concept into network design can reduce optical hardware (e.g., transceiver) cost and lead to optical networks that make more efficient use of bandwidth [1, 4].

The Multicast Optical Network

Data Plane: Hardware Building Blocks

A transparent optical WDM network consists of some number N of nodes interconnected by fiber links. Each of the links is capable of carrying W wavelengths, and each of the nodes is equipped with a nonblocking OXC capable of optically switching signals at the granularity of a wavelength. In a multicast optical network, the OXC at (some of) the nodes is multicast-capable (MC-OXC). An MC-OXC consists of two key hardware components: the *power splitter* and the *splitter-and-delivery* (SaD) switch.

The power splitter provides the basic functionality for realizing multicast at the optical layer. A power splitter is a passive device capable of splitting an input optical signal without any knowledge of the optical characteristics (e.g., modulation, BER, etc.) of that signal. More specifically, an n-way splitter is a device that splits an input signal into n outputs, such that each of the n output signals is identical to the input signal in all respects except signal power. While this light splitting operation can be achieved at low cost by simply fusing fibers together, it contributes to power loss. For an ideal device, the power of each output is $(1/n)$th that of the original signal; in practice, the splitting operation introduces additional losses, and the power of each output is lower than that of the ideal case.

We distinguish between *fixed* and *configurable* power splitters. The former is a device with a fixed fanout n, such that an input signal is always split into n identical output signals. The main drawback of such a device is that it introduces unnecessary power losses when an input signal only needs to be split into a number of outputs less than n. A configurable power splitter, on the other hand, is such that an input signal can be split into any number n_{out} of output signals, $n_{out} = 1, \ldots, n$, where n is the maximum fanout of the device. Configurability is made possible by new devices such as the compact multimode interference couplers with tunable power splitting ratios recently reported in [5]. Configurable splitters eliminate unnecessary power loss by splitting an input signal into the exact number of outputs required by the application.

Power splitters are the fundamental hardware devices used in implementing splitter-and-delivery (SaD) switches, which in turn are the main building blocks of MC-OXCs. The SaD switch design was first proposed in [6] and was later modified [7] in order to reduce cost and improve power efficiency. Figure 2 shows a SaD switch with P input ports and P output ports which utilizes configurable power splitters. The SaD switch consists of P power splitters, P^2 amplifiers, and P^2 2×1 photonic switches. All the input signals to the switch are on the *same* wavelength. The power splitter associated with input port $p, p = 1, \ldots, P$, can be configured to split the incoming signal into n_{out} output signals, $n_{out} = 1, \ldots, P$; note that $n_{out} = 1$ corresponds to no power splitting (i.e., no multicast), while $n_{out} = P$ corresponds to a broadcast operation. Each signal resulting from a splitting operation undergoes optical amplification, which partially compensates for the power loss due to light splitting. Finally, by configuring the corresponding n_{out} 2×1 photonic switches, each of the n_{out} signals resulting from the splitting operation is switched to the appropriate output port.

A $P \times P$ MC-OXC consists of a set of W $P \times P$ SaD switches, one for each wavelength. Figure 3 shows such an MC-OXC with W wavelengths. At each input port, a demultiplexer is used to extract individual wavelengths, which are then directed to the appropriate SaD switch. The space switching and/or splitting of signals is performed at the SaD switches, as shown in Fig. 2. Finally, P multiplexers, one at each output port, are employed to combine the W signals on individual wavelengths for transmission on an outgoing fiber. This MC-OXC design can realize any permutation of optical multicast

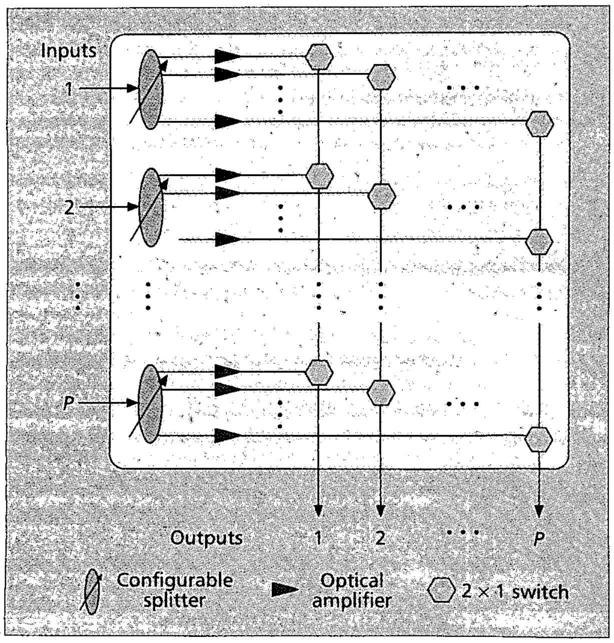

■ Figure 2. *A* P × P *SaD switch [6, 7].*

or unicast connections in a strictly nonblocking manner [7]. However, it requires a large number of power splitters and optical amplifiers within each of the *W* SaD switches, resulting in a complex and expensive fabrication process. We discuss alternative cost-effective MC-OXC designs in the next section.

Control Plane: Algorithmic Building Blocks

The fundamental control problem in optical networks is the *routing and wavelength assignment* (RWA) problem [8]. This problem arises due to the tight coupling between routing and wavelength selection in WDM networks: 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 RWA problem 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:

- *Wavelength continuity constraint*: In the absence of converters, a lightpath must use the same wavelength on all the links along its path from source to destination node.
- *Distinct wavelength constraint*: All lightpaths using the same link (fiber) must be allocated distinct wavelengths.

In an optical network that provides light-tree services, the corresponding control problem is referred to as *multicast* RWA (MC-RWA). MC-RWA bears many similarities to the RWA problem. Specifically, the 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 optical connection not just along the links of a single path but along the links of the whole light-tree. Depending on the nature of traffic demands, we also distinguish between static and dynamic MC-RWA problems. It is well-known [9] that optimal solutions to point-to-point RWA problems are not practically obtainable. With a more general construct (the light-tree) and hence a much larger search space, this observation is even more true for MC-RWA problems. The challenge in this case is to design MC-RWA algorithms that can cope with the increased complexity of the problem and yet produce good solutions; we discuss this issue in more detail in the next section.

In addition to MC-RWA algorithms, a set of signaling and control protocols must be implemented within the optical network to support the establishment and management of light-trees (and lightpaths). Overall, there are three components of the control plane that are crucial to setting up light-trees within the optical network:

Topology and resource discovery. The main purpose of discovery mechanisms is to disseminate network state information including optical resource usage (e.g., wavelength availability), network connectivity, link capacity and protection information, and special constraints (e.g., physical layer constraints such as power budget). Similar mechanisms are employed in protocols such as Open Shortest Path First (OSPF) and Intermediate System to Intermediate System (IS-IS). In order to support an effective light-tree service, additional information about the availability and location of MC-OXCs and relevant attributes (e.g., maximum splitter fanout) must be included in these protocols.

Route computation and wavelength selection. This component employs RWA/MC-RWA algorithms and traffic engineering functions to select an appropriate route and wavelength for a requested lightpath/light-tree. The operation of the RWA/MC-RWA algorithms is guided by the optical network *traffic engineering database*, which is built and updated using the information collected by the topology and resource discovery module.

Connection management. Connection management is concerned with setup and teardown of lightpaths/light-trees, as well as coordination of protection switching in case of failures.

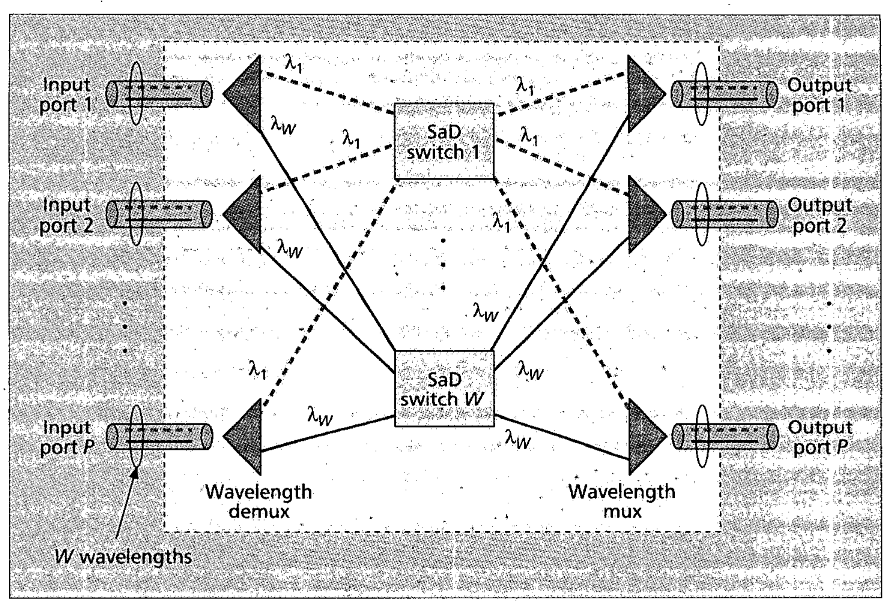

■ Figure 3. *A* P × P *MC-OXC based on the SaD switch architecture [6],* W *wavelengths.*

Example protocols under the umbrella of GMPLS include the constrained-based routing label distribution protocol (CR-LDP), and enhancements to the Resource Reservation Protocol for Traffic Engineering (RSVP-TE).

Currently, most of the work on defining control plane protocols for optical networks, or augmenting existing protocols to handle optical resource information, takes place under the auspices of the Internet Engineering Task Force (IETF). This activity reflects a convergence of the optical networking and Internet communities to develop technology built around a single common framework, namely, GMPLS, for controlling both IP and optical network elements. For an overview of GMPLS and related resource discovery and connection management protocols, the reviewer is referred to [10].

Challenges in Multicast Optical Networks

While we have reached the point where there is proper appreciation of optical multicast and its applications, many open questions exist, and much serious work remains to be performed. In this section we discuss a number of research areas and summarize recent relevant work. The reader should keep in mind that this is an area of ongoing research, and we expect that many significant results are forthcoming.

Power-Efficient Optical Network Design

It has been widely recognized [11] that physical layer impairments must be taken into account when routing optical connections in transparent networks. One of the most important measures of optical quality of service (QoS) is signal power: if the power level falls below a threshold determined by the receiver sensitivity, the information carried by the signal cannot be recovered. For a signal traveling over a light-tree, the light splitting operations at the MC-OXCs are the main contributors to power loss. While the effects of light splitting may be partially mitigated by amplification, incorporating large numbers of optical amplifiers within the hardware components of an MC-OXC increases the cost and difficulty of fabrication. Two approaches have been taken in the literature to deal with the issues that power splitters introduce: one approach [7, 12], focuses on including power efficiency as one of the primary objectives of network design, and is the subject of this subsection; while another emphasizes light-tree routing algorithms that are power-budget-aware, and is discussed later. We note that the two approaches are complementary, and can both be employed in the same optical network.

The primary goal of power-efficient design in the context of multicast optical networks is to minimize (and even eliminate) any unnecessary power splitting operations. A secondary but equally important objective is to minimize network cost. Furthermore, this design philosophy can be applied at either the MC-OXC level or on a networkwide scale, as explained next.

At the MC-OXC level, the objective is to ensure that optical signals traversing a given MC-OXC do not suffer unnecessary losses. Recall that an MC-OXC (refer to Fig. 3) consists of a number of SaD switches, such as the one in Fig. 2. One of the first SaD switch architectures [6] required all connections (even point-to-point connections) passing through the MC-OXC to face P-way splitting, where P is the number of output ports of the MC-OXC. Clearly, this design introduces excessive signal losses even to connections that need no splitting, or require n-way splitting, $n < P$. A better approach would be to incorporate configurable power splitters in the SaD architecture, as in the design of Fig. 2, which was considered in [13]. Such a design provides each multicast connection with the exact fanout

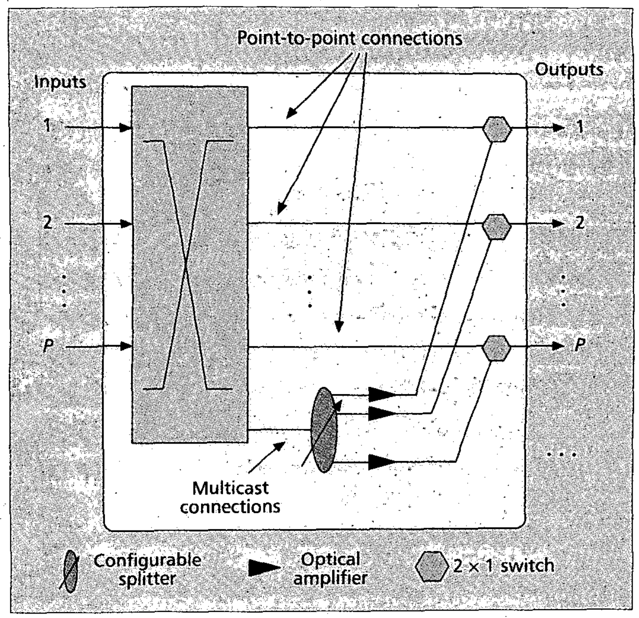

■ Figure 4. *A* $P \times P$ *SaD switch with splitter sharing.*

required, altogether eliminating unnecessary power losses. However, it requires more splitters and amplifiers than the one in [7], which is discussed next, and additional research is required to evaluate the relative merits of the two approaches.

A different SaD switch architecture was introduced in [7], which addressed the power loss issue in two ways. First, point-to-point connections are switched separately from multicast connections, and thus do not undergo any splitting. Second, the architecture is based on the concept of *splitter sharing*. Specifically, rather than including a power splitter and P amplifiers at each input port of the SaD switch, as in Fig. 2, there is a single power splitter (of fanout P) and P amplifiers *shared* by all input ports. This architecture is illustrated in Fig. 4. It is shown in [7] that this architecture significantly reduces the cost and fabrication complexity of the SaD switch. These savings in cost, however, come at the expense of higher blocking for multicast connections. In particular, the switch of Fig. 4 can only accommodate a single multicast connection at a time. The switch design in Fig. 2, on the other hand, can accommodate any number of simultaneous multicast connections, as long as their destination output port sets do not intersect. Therefore, splitter sharing is expected to perform well when multicast connections represent a relatively small fraction of the total number of connections in the network. We also note that, between the two extremes of one splitter per input port (Fig. 2) and a single splitter shared among all input ports (Fig. 4), there is a wide range of alternative designs with varying trade-offs between cost and multicast connection blocking.

At the network level, it has been observed [12–14] that not all OXCs in the network need to be multicast-capable. A network in which only a subset of the nodes are MC-OXCs is referred to as one with *sparse light splitting*. In this case, an important design objective is to minimize the number of MC-OXCs with minimal effects on the performance experienced by multicast connections. Such an objective is directly related to network cost minimization since, due to the splitters and associated amplifiers, the cost of an MC-OXC can be significantly higher than that of a non-multicast-capable OXC. The study in [12] demonstrated that, for a wide range of network topologies and traffic patterns, only about 50 percent of the OXCs need to be multicast-capable; including additional MC-OXCs only slightly enhances the performance. Sparse light splitting introduces new research problems that have only recently received attention. For instance, good strategies for placing the MC-OXCs in the network are needed; a first step in this direction was taken in [12], but more studies are needed. Also, new routing algorithms that take the constraint on the availability of MC-OXCs into account must be designed; we discuss some work in this direction later. Another issue that has not been sufficiently explored is the relationship between sparse light splitting and sparse wavelength conversion, and the benefits of collocating converters and splitters.

An alternative architecture for realizing multicast at the optical layer is based on the *tap-and-continue* (TaC) concept introduced in [15]. This architecture does not require power splitters at the OXCs, only tapping devices. In a TaC architecture, an OXC may tap an input signal and direct a small fraction of its power to a local client; it then switches the remaining power of the signal to an output port. In order to perform optical multicast, one must determine a path that starts at the source and visits all destination nodes in some order. It is shown in [15] that the problem of finding such a path of minimum length is intractable, and an algorithm for this multidestination routing problem is developed.

MC-RWA Algorithms

If the traffic patterns in the network are reasonably well known in advance and any traffic variations take place over long timescales, the most effective technique for establishing optical connections (lightpaths or light-trees) is by solving a static MC-RWA problem. Since these connections are assumed to remain in place for relatively long periods of time, it is worthwhile to optimize the way in which network resources (e.g., physical links and wavelengths) are assigned to each connection, even though optimization may require considerable computational effort. Typically, routing and wavelength assignment are considered together as an optimization problem using integer programming formulations. An instance of the static MC-RWA problem with the objective of minimizing the number of wavelengths was considered in [16], and a set of heuristics was presented. On the other hand, the objective of the MC-RWA problem formulation in [17] was to maximize the total number of established multicast connections. Rather than providing heuristic algorithms for solving the integer programming problem, bounds on the objective function were presented by relaxing the integer constraints.

The dynamic MC-RWA problem is encountered during the real-time network operation phase, and involves the dynamic provisioning of light-trees. Specifically, users submit to the network, in some random fashion, requests for light-trees to be set up as needed. Depending on the state of the network at the time of the request, the available resources may or may not be sufficient to establish a light-tree that spans the requested source node and destination set. The network state consists of the physical route and wavelength assignment for all active light-trees, and it evolves randomly in time as new light-trees are admitted and existing ones released. Thus, each time a request is made, the network must determine whether it is feasible to accommodate the request, and if so, to perform routing and wavelength assignment. If a request for a light-tree cannot be accepted due to lack of resources, it is blocked.

Because of the real-time nature of the problem, MC-RWA algorithms in a dynamic traffic environment must be simple and fast. 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 light-tree routing problem and the wavelength assignment problem. This approach has been taken in [17–19] to study the blocking performance of optical networks with dynamic multicast traffic. All three studies decouple the routing and wavelength assignment problems, and consider alternate routing strategies, whereby a number of trees may be used for each multicast connection. When a request for a connection arrives, the associated trees are considered in some order, and the connection is blocked if no free wavelength is found in any of the trees. The study in [17] considered the first-fit wavelength allocation policy, in which for each tree, wavelengths are also considered in a fixed order. In [19], on the other hand, the random wavelength assignment was studied; in addition, multiple classes of requests were considered. A different approach was taken in [18], where an iterative approximation algorithm was developed for completely connected networks under random wavelength assignment. Since, in general, network topologies are not completely connected, the results of [18] can be used as lower bounds for more general topologies.

Most of the studies presented so far use existing routing algorithms to build light-trees. Next, we make the case for new algorithms that take into account optical layer constraints, and we discuss recent research in this direction.

Light-Tree Routing Under Optical Layer Constraints

While the problem of establishing a light-tree that spans a given source and a set of destination nodes bears some similarities to the Steiner tree problem [20], the nature of optical multicast introduces several new issues and complexities. Splitting an optical signal introduces losses, a problem not encountered in electronic packet- or circuit-switched networks, and thus not addressed by existing routing tree algorithms. Even in the presence of optical amplifiers, this signal loss imposes a hard upper bound on the number of times a signal can be split. Furthermore, since optical amplifiers also amplify noise levels, there is a limit on the number of amplifiers a signal may traverse. Also, in optical networks with sparse light splitting, a feasible light-tree may not exist. These new constraints require novel approaches for routing light-trees.

The problem of routing multicast optical connections in networks with sparse light splitting was studied in [14]. To deal with the fact that a feasible multicast tree may not exist for a given source and destination set, the concept of a *light-forest* was introduced, and four new algorithms were developed. 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 path out of this OXC are removed. The destinations on the path not removed continue to receive the signal via this OXC. However, new paths are created from some MC-OXC on the tree to destination nodes in the removed paths. In general, the routing algorithms assume unlimited fanout capacity at MC-OXCs; also, each tree of a given light-forest is assigned a different wavelength since the source transmits data on all trees simultaneously.

The problem of constructing light-trees under optical layer power budget constraints has been studied in [13, 21]. The work in [21] extended one of the algorithms of [14] to account for power losses along the links of the trees in a light-forest. On the other hand, the focus of [13] is on new routing algorithms that can guarantee a certain level of quality for the signals received by the destination nodes. A new constrained light-tree routing problem was defined, and a set of constraints were introduced on the source-destination paths to account for the power losses (due to both splitting and attenuation) at the optical layer. A number of variants of the problem were investigated, their complexity was characterized, and a suite of corresponding routing algorithms was developed; one of the algorithms is appropriate for networks with sparse light splitting and/or limited splitting fanout. One of the most important findings of [13] is that in order to guarantee an adequate signal quality and scale to large destination sets, light-trees must be as balanced as possible (i.e., the difference between the maximum and minimum depth among all leaf nodes is minimized). It was demonstrated that existing algorithms tend to construct highly unbalanced trees, and are thus expected to perform poorly in an optical network setting. The new algorithms, on the other hand, were designed to construct balanced trees that, in addition to having good performance in terms of signal quality, also ensure a certain degree of fairness among destination nodes.

Despite this early work on light-tree routing, various aspects of the problem remain uninvestigated. Specifically, there is a need for additional research in the area of routing algorithms that take into account constraints unique to optical layer multicast, including those due to physical impairments (e.g., loss or dispersion) and optical device characteristics (e.g., noise issues caused by the amplifiers in the signal paths). Furthermore, multicast routing is tightly coupled with certain network design issues; for instance, amplifier placement under multicast traffic is a difficult and open issue. We also note that most of the work in this area has been experimental in nature, with little formal analysis of the heuristics developed. Thus, there is a need for a more formal and systematic approach to light-tree routing under optical layer constraints.

Multicast Connection Management

An important component of an operational network is control plane functionality for supporting the network design objectives and automating the process of light-tree (and lightpath) establishment. Currently, a number of standardization activities addressing the control plane aspects of optical networks are underway within the IETF and the Optical Internetworking Forum (OIF). Most of these activities take place within the control framework provided by GMPLS. In particular, a number of existing protocols are being extended to support GMPLS and/or to carry network state information relevant to optical networks, including label distribution protocols (RSVP-TE, CR-LDP) and routing protocols (OSPF, IS-IS).

While ongoing work represents an important first step toward seamless integration of Internet and optical network technologies, additional work is required to develop practical control and signaling protocols for the management of multicast connections. For instance, label distribution protocols must be extended to set up and tear down light-trees in addition to point-to-point lightpaths. Of major concern in defining such protocols is the creation of loops when constructing a light-tree. Routing loops can cause serious problems in optical networks, and extreme care must be taken so that such loops do not occur. (In comparison, consider that routing loops frequently occur with existing Internet multicast protocols such as DVMRP.) Also, resource discovery protocols must be extended to carry information relevant to optical multicast, while routing protocols must be modified to incorporate path selection mechanisms that take optical layer constraints into account. Another important area of research involves the investigation of protection and restoration issues for survivable light-tree communication. Therefore, we expect the control plane and standardization activities to further develop as the underlying technology matures and our collective understanding of optical multicast advances.

Concluding Remarks

Optical multicast is an important and exciting research area that is likely to become more essential in the future. The new constraints and realities imposed by the optical layer and WDM technology will certainly affect our long-held assumptions regarding issues such as routing, control, and signaling, which have been developed for mostly opaque electronic networks. As we carefully rethink these issues in the context of transparent (or almost transparent) optical networks, protocols, algorithms, and network design methodologies will certainly evolve to better accommodate the new technology.

References

[1] L. H. Sahasrabuddhe and B. Mukherjee, "Light-trees: Optical Multicasting for Improved Performance in Wavelength-routed Networks," *IEEE Commun. Mag.*, vol. 37, no. 2, Feb. 1999, pp. 67–73.
[2] D. Thaker and G. N. Rouskas, "Multi-Destination Communication in Broadcast WDM networks: A survey," *Opt. Net.*, vol. 3, no. 1, Jan./Feb. 2002, pp. 34–44.
[3] A. Dutta-Roy, "An Overview of Cable Modem Technology and Market Perspectives," *IEEE Commun. Mag.*, vol. 39, no. 6, June 2001, pp. 81–88.
[4] R. Malli, X. Zhang, and C. Qiao, "Benefit of Multicasting in All-Optical Networks," *Proc. SPIE*, vol. 3531, Nov. 1998, pp. 209–20.
[5] J. Leuthold and C. H. Joyner, "Multimode Interference Couplers with Tunable Power Splitting Ratios," *IEEE/OSA J. Lightwave Tech.*, vol. 19, no. 5, May 2001, pp. 700–6.
[6] W. S. Hu and Q. J. Zeng, "Multicasting Optical Cross Connects Employing Splitter-and-Delivery Switch," *IEEE Photonics Tech. Lett.*, vol. 10, July 1998, pp. 970–72.
[7] M. Ali and J. Deogun, "Power-efficient Design of Multicast Wavelength-Routed Networks," *IEEE JSAC*, vol. 18, no. 10, 2000, pp. 1852–62.
[8] H. Zang, J. P. Jue, and B. Mukherjee, "A Review of Routing and Wavelength Assignment Approaches for Wavelength-Routed Optical WDM Networks," *Opt. Net.*, vol. 1, no. 1, Jan. 2000, pp. 47–60.
[9] R. Dutta and G. N. Rouskas, "A Survey of Virtual Topology Design Algorithms for Wavelength Routed Optical Networks," *Opt. Net.*, vol. 1, no. 1, Jan. 2000, pp. 73–89.
[10] A. Banerjee *et al.*, "Generalized Multiprotocol Label Switching: An Overview of Routing and Management Functions," *IEEE Commun. Mag.*, vol. 39, no. 1, Jan. 2001, pp. 144–50.
[11] J. Strand, A. L. Chiu, and R. Tkach, "Issues for Routing in the Optical Layer," *IEEE Commun. Mag.*, Feb. 2001, pp. 81–96.
[12] M. Ali and J. Deogun, "Allocation of Splitting Nodes in Wavelength-routed Networks," *Photonic Net. Commun.*, vol. 2, no. 3, Aug. 2000, pp. 245–63.
[13] G. N. Rouskas, "Light-Tree Routing Under Optical Layer Power Budget Constraints," *Proc. 17th IEEE Comp. Commun. Wksp.*, Oct. 14–16, 2002.
[14] X. Zhang, J. Y. Wei, and C. Qiao, "Constrained Multicast Routing in WDM Networks with Sparse Light Splitting," *J. Lightwave Tech.*, vol. 18, no. 12, Dec. 2000, pp. 1917–27.
[15] M. Ali and J. Deogun, "Cost-efficient Implementation of Multicasting in Wavelength-routed Networks," *IEEE/OSA J. Lightwave Tech.*, vol. 18, no. 12, Dec. 2000, pp. 1628–38.
[16] X.-H. Jia *et al.*, "Optimization of Wavelength Assignment for QoS Multicast in WDM Networks," *IEEE Trans. Commun.*, vol. 49, no. 2, Feb. 2001, pp. 341–50.
[17] G. Sahin and M. Azizoglu, "Multicast Routing and Wavelength Assignment in Wide-Area Networks," *Proc. SPIE*, vol. 3531, Nov. 1998, pp. 196–208.
[18] A. E. Kamal and A. K. Al-Yatama, "Blocking Probabilities in Circuit-switched Wavelength Division Multiplexing Networks Under Multicast Service," *Perf. Eval.*, vol. 47, no. 1, 2002, pp. 43–71.
[19] S. Ramesh, G. N. Rouskas, and H. G. Perros, "Computing Call Blocking Probabilities in Multi-class Wavelength Routing Networks with Multicast Traffic," *IEEE JSAC*, vol. 20, no. 1, Jan. 2002, pp. 89–96.
[20] S. L. Hakimi, "Steiner's Problem in Graphs and Its Implications," *Networks*, vol. 1, 1971, pp. 113–33.
[21] K.-D. Wu, J.-C. Wu, and C.-S. Yang, "Multicast Routing with Power Consideration in Sparse Splitting WDM Networks," *Proc. IEEE ICC*, 2001, pp. 513–17.

Biography

GEORGE N. ROUSKAS [SM] (rouskas@csc.ncsu.edu) is a professor of computer science at North Carolina State University. He received a diploma in computer engineering from the National Technical University of Athens (NTUA), Greece, in 1989, and M.S. and Ph.D. degrees in computer science from the College of Computing, Georgia Institute of Technology, Atlanta, in 1991 and 1994, respectively. He is a recipient of a 1997 NSF Faculty Early Career Development (CAREER) Award, and a co-author of a paper that received the Best Paper Award at the 1998 SPIE Conference on All-Optical Networking. He also received the 1995 Outstanding New Teacher Award from the Department of Computer Science, North Carolina State University, and the 1994 Graduate Research Assistant Award from the College of Computing, Georgia Tech. During the 2000–2001 academic year he spent a sabbatical term at Vitesse Semiconductor, Morrisville, North Carolina, and in May and June 2000 he was an invited professor at the University of Evry, France.