

Dynamic Lightpath Establishment in Wavelength-Routed WDM Networks

Hui Zang, *Sprint Advanced Technology Laboratories*

Jason P. Jue, *University of Texas at Dallas*

Laxman Sahasrabudde, *Summit Networks Inc.*

Ramu Ramamurthy, *Tellium Inc.*

Biswanath Mukherjee, *University of California, Davis*

ABSTRACT

In wavelength-routed WDM networks, a control mechanism is required to set up and take down all-optical connections. Upon the arrival of a connection request, this mechanism must be able to select a route, assign a wavelength to the connection, and configure the appropriate optical switches in the network. The mechanism must also be able to provide updates to reflect which wavelengths are currently being used on each link so that nodes may make informed routing decisions. In this work, we review control mechanisms proposed in the literature. We also investigate and compare two different distributed control mechanisms for establishing all-optical connections in a wavelength-routed WDM network: an approach based on link-state routing, and one based on distance-vector routing.

INTRODUCTION

In recent years, advances in wavelength-division multiplexing (WDM) technology have enabled the deployment of systems that are capable of providing large amounts of bandwidth. At the same time, systems utilizing optical wavelength routing switches (WRSs) or photonic cross-connects (PXC) have emerged that enable data to be switched entirely in the optical domain without requiring conversion to electronics at intermediate nodes. Configuring these optical devices across a network enables one to establish all-optical connections, or *lightpaths*, between source and destination nodes (Fig. 1).

In the absence of wavelength conversion devices, a lightpath must occupy the same wavelength on each link in its route, a restriction known as the *wavelength-continuity constraint*. Thus, given a set of lightpaths that need to be

established, and given a constraint on the number of wavelengths, it is necessary to determine both the routes over which these lightpaths should be established and the wavelengths that should be assigned to these lightpaths. This problem is known as the *routing and wavelength assignment (RWA)* problem.

In a wavelength-routed network, the traffic can be either static or dynamic:

- In a static traffic pattern, a set of lightpaths are set up all at once and remain in the network for a long period of time. The RWA problem for static traffic is known as the *static lightpath establishment (SLE)* problem. A review of approaches to the SLE problem may be found in [1].
- In a dynamic traffic pattern, a lightpath is set up for each connection request as it arrives, and the lightpath is released after some finite amount of time. The problem of lightpath establishment in a network with dynamic traffic demands is called the *dynamic lightpath establishment (DLE)* problem. This article focuses on the DLE problem.

With the rapid growth of the Internet, the bandwidth demand for data traffic is exploding. It is believed that dynamic lightpath establishment, or *on-demand lightpath establishment*, will enable service providers to respond quickly and economically to customer demands.

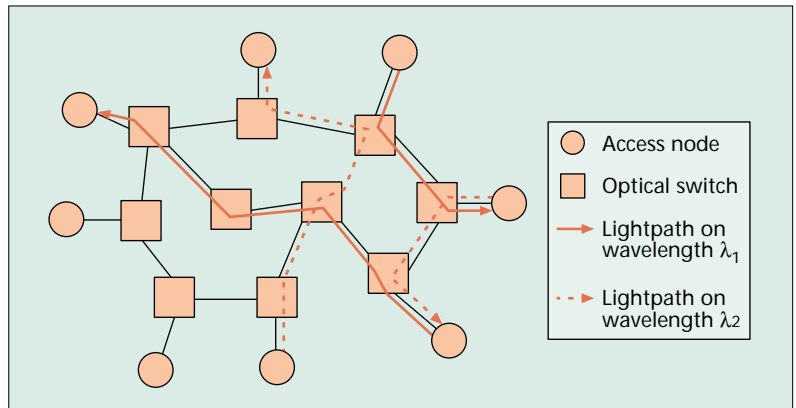
One of the challenges involved in designing wavelength-routed networks with dynamic traffic demands is to develop efficient algorithms and protocols for establishing lightpaths. The algorithms must be able to select routes and assign wavelengths to connections in a manner that efficiently utilizes network resources and maximizes the number of lightpaths established. Signaling protocols for setting up lightpaths must effectively manage the distribution

of control messages and network state information in order to establish a connection in a timely manner. Typically, a *network control and management protocol* is employed to perform the RWA and signaling tasks mentioned above.

Another issue in dynamic lightpath establishment is how to initiate requests for lightpath establishment and removal. There are a number of possible approaches for generating a connection request. For example, a customer can initiate a connection request by clicking on a Web page or calling up the service provider. A connection request can also be initiated by an IP router or other network control devices that can identify a demand between two PXC's. The problem of identifying a flow in a packet-based network, which is referred to as the *flow detection* problem, has been studied in the literature [2] and is beyond the scope of this article.

Recently, there has been much effort to develop the aforementioned protocols for providing on-demand establishment of lightpaths. More specifically, a number of organizations are currently working to develop standards for facilitating dynamic lightpath establishment in optical networks. The Optical Domain Service Interconnect (ODSI) group is working to standardize interfaces to allow client networks and devices to interact with the optical network. The ODSI framework does not specify how the lightpaths are actually established within the optical network, but simply defines how a client would request a lightpath or release a lightpath from the optical network. The current work of the ODSI is summarized in [3].

Multiprotocol label switching (MPLS) is a control framework currently being developed as a standard to enable fast switching in IP networks. MPLS control mechanisms can be used to establish a label-switched path (LSP) between two nonneighboring IP routers, enabling packets to bypass IP routers at intermediate nodes. The primary signaling mechanism for establishing an LSP in MPLS is the label distribution protocol (LDP). The concept of MPLS can be extended to wavelength-routed optical networks as multiprotocol lambda switching (MP λ S). The Internet Engineering Task Force (IETF) is currently working on generalized multiprotocol label switching (GMPLS), a generalized control framework for establishing various types of connections, including lightpaths, in IP-based networks. Currently, there is much work being done by the IETF to modify existing routing and signaling protocols in order to support GMPLS. In particular, the IETF is focusing on enhancements to the Open Shortest Path First (OSPF) routing protocol, and the Constraint-Based Routing Label-Distribution Protocol (CR-LDP) and Resource Reservation Protocol (RSVP) signaling protocols. A summary of the proposed enhancements can be found in [4]. OSPF is a link state protocol in which the state of each link in the network is periodically broadcast to all nodes in the form of link state advertisements (LSAs). The nodes may then make routing decisions based on this information. RSVP is an IP protocol used for signaling the



■ Figure 1. A wavelength-routed WDM network.

resource requirements of an application to intermediate routers. CR-LDP is a protocol that enables the distribution of control messages for establishing LSPs. CR-LDP utilizes constraint-based routing and runs on top of TCP for reliability. While the current focus of the IETF is on a few specific protocols, GMPLS itself is not restricted to any single routing or signaling protocol. Furthermore, protocols such as OSPF, CR-LDP, and RSVP are flexible and lend themselves to the implementation of various routing and signaling schemes for lightpath establishment. In this article we discuss some of these underlying routing and signaling schemes that can be implemented within a GMPLS framework, and discuss these schemes from a performance perspective.

We also investigate and compare the performance of two distributed control and management protocols: a link state approach based on link state routing, and a distributed routing approach based on distance-vector routing. We choose the link state and distance-vector routing approaches for comparison because these are the primary approaches being considered by the industry today. There has been some work comparing the performance of these routing protocols for IP traffic [5]. In this article we investigate the performance of these approaches in a multi-wavelength network.

The rest of the article is organized as follows. In the next section we discuss dynamic routing and wavelength assignment algorithms for WDM networks. Then we review various signaling protocols for reserving resources and setting up lightpaths in a wavelength-routed network. Next, we describe the two control mechanisms and present numerical examples to compare the performance of these two approaches. Finally, we provide a summary and conclude the article.

DYNAMIC ROUTING AND WAVELENGTH ASSIGNMENT

When lightpaths are established and taken down dynamically, routing and wavelength assignment decisions must be made as connection requests arrive at the network. It is possible that, for a given connection request, there

Similar to the case of static lightpaths, the dynamic RWA problem can be decomposed into a routing subproblem and a corresponding wavelength assignment subproblem.

may be insufficient network resources to set up a lightpath, in which case the connection request will be blocked. The connection may also be blocked if there is no common wavelength available on all of the links along the chosen route. Thus, the objective in the dynamic situation is to choose a route and a wavelength that maximize the probability of setting up a given connection, while at the same time attempt to minimize blocking for future connections. Similar to the case of static lightpaths, the dynamic RWA problem can also be decomposed into a routing subproblem and a corresponding wavelength assignment subproblem.

Approaches to solving the routing subproblem can be categorized as either static or adaptive, and as utilizing either global or local network state information.

FIXED ROUTING AND FIXED-ALTERNATE PATH ROUTING

Two examples of algorithms which utilize static routes are fixed routing and fixed-alternate path routing. In fixed routing, a single fixed route is predetermined for each source-destination pair. In fixed-alternate path routing, multiple fixed routes are precomputed for each source-destination pair and stored in an ordered list at the source node's routing table. As a connection request arrives, one route is selected from the set of precomputed routes. Both of these approaches are much simpler to implement than adaptive routing schemes, but may suffer from higher connection blocking.

ADAPTIVE ROUTING BASED ON GLOBAL INFORMATION

Adaptive routing approaches increase the likelihood of establishing a connection by taking into account network state information. For the case in which global information is available, routing decisions may be made with full information as to which wavelengths are available on each link. In order to find an optimal route, a cost may be assigned to each link based on wavelength availability, and a least-cost routing algorithm may be executed.

Adaptive routing based on global information may be implemented in either a centralized or distributed manner. In a centralized algorithm a single entity, such as a network manager, maintains complete network state information, and is responsible for finding routes and setting up lightpaths for connection requests. Since a centralized entity manages the entire network, a high degree of coordination among nodes is not needed; however, a centralized entity becomes a possible single point of failure.

A distributed adaptive routing algorithm based on global information may be implemented in a number of ways. In a link state approach, each node in the network must maintain complete network state information [6]. Each node may then find a route for a connection request in a distributed manner. Whenever the state of the network changes, all of the nodes must be informed. Therefore, the establishment or removal of a lightpath in the net-

work may result in the broadcast of update messages to all nodes in the network. The need to broadcast update messages may result in significant control overhead, especially if lightpaths are being established and removed at a high rate. Furthermore, it is possible for a node to have outdated information, and for the node to make an incorrect routing decision based on this information.

A distance-vector or *distributed-routing* approach to routing with global information is also possible [7]. This approach does not require that each node maintain complete link-state information as in [6], but instead has each node maintain a routing table which indicates, for each destination and on each wavelength, the next hop to the destination and the distance to the destination. The approach relies on a distributed Bellman-Ford algorithm to maintain the routing tables. Similar to [6], the scheme also requires nodes to update their routing table information whenever a connection is established or taken down. This update is accomplished by having each node send routing updates to their neighbors periodically or whenever the status of the node's outgoing links changes. We evaluate the above two approaches from several performance metrics later.

Another form of adaptive routing is *least-congested-path* (LCP) routing [8]. The congestion on a link is measured by the number of wavelengths available on the link. Links that have fewer available wavelengths are considered more congested. The congestion on a path is indicated by the congestion on the most congested link in the path. Similar to alternate routing, for each source-destination pair a sequence of routes is preselected. Upon the arrival of a connection request, the least congested path among the predetermined routes is chosen. It has been shown in [8] that using shortest-path routing first and LCP second for tie breaking works better than using LCP alone.

Although routing schemes based on global knowledge must deal with the task of maintaining a potentially large amount of state information which changes constantly, these schemes often make the most optimal routing decisions if the state information is up to date. Thus, global-knowledge-based schemes may be well suited to networks in which lightpaths are fairly static and do not change much with time.

ADAPTIVE ROUTING BASED ON NEIGHBORHOOD INFORMATION

In the global-information-based LCP approach, all links on all candidate paths must be examined in choosing the least congested path. Either each node is required to maintain complete state information, or the information must be gathered in real time as the lightpath is established. A variant of LCP is proposed in [9], which only examines the first k links on each path (referred to as the source's neighborhood information), where k is a parameter of the algorithm. It has been shown that when $k = 2$, this algorithm can achieve similar performance to fixed-alternate routing.

ADAPTIVE ROUTING BASED ON LOCAL INFORMATION: DEFLECTION ROUTING

Another approach to adaptive routing with limited information is deflection routing, or alternate link routing [10]. This routing scheme chooses from alternate links on a hop-by-hop basis rather than from alternate routes on an end-to-end basis. The routing is implemented by having each node maintain a routing table that indicates, for each destination, one or more alternate outgoing links to reach that destination. These alternate outgoing links are pre-computed and may be ordered such that a connection request will preferentially choose certain links over other links as long as wavelength resources are available on those links. If resources are unavailable on the preferred link, an alternate link is chosen for the route. Other than a static routing table, each node will only maintain information regarding the status of wavelength usage on its own outgoing links. Hence, there are no update messages in the network, and control bandwidth demand is greatly decreased.

WAVELENGTH ASSIGNMENT

In general, if there are multiple feasible wavelengths between a source node and a destination node, a wavelength assignment algorithm is required to select a wavelength for a given lightpath. A more detailed review of wavelength assignment algorithms can be found in [1].

SIGNALING AND RESOURCE RESERVATION

In order to set up a lightpath, a signaling protocol is required to exchange control information among nodes and reserve resources along the path. In many cases, the signaling protocol is closely integrated with the routing and wavelength assignment protocols. Signaling and reservation protocols may be categorized based on whether the resources are reserved on each link in parallel, on a hop-by-hop basis along the forward path, or on a hop-by-hop basis along the reverse path. Protocols will also differ depending on whether or not global information is available.

PARALLEL RESERVATION

In [6], the control scheme reserves wavelengths on multiple links in parallel. The scheme, which is based on link-state routing, assumes that each node maintains global information on the network topology and the current state of the network, including information regarding the wavelengths being used on each link. Based on this global information, the node can calculate an optimal route to a destination on a given wavelength. The source node then attempts to reserve the desired wavelength on each link in the route by sending a separate control message to each node in the route. Each node that receives a reservation request message will attempt to reserve the specified wavelength, and will send either a positive or negative acknowledgment back to the source. If the source node receives positive acknowledgments from all of the nodes, it can establish the lightpath and begin communicating with the desti-

nation. The advantage of a parallel reservation scheme is that it shortens the lightpath establishment time by having nodes process reservation requests in parallel. The disadvantage is that it requires global knowledge, since both the path and the wavelength must be known ahead of time.

HOP-BY-HOP RESERVATION

An alternative to parallel reservation is hop-by-hop reservation in which a control message is sent along the selected route one hop at a time. At each intermediate node, the control message is processed before being forwarded to the next node. When the control message reaches the destination, it is processed and sent back toward the source node. The actual reservation of link resources may be performed while either the control message is traveling in the forward direction toward the destination, or the control message is traveling in the reverse direction back toward the source.

Forward Reservation — In forward reservation schemes, wavelength resources are reserved along the forward path to the destination on a hop-by-hop basis. The method of reserving wavelengths depends on whether or not global information is available to the source node. If the source node is maintaining complete state information, it will be aware of which wavelengths are available on each link. Assuming the state information is current, the source node may then send a connection setup message along the forward path, reserving the same available wavelength on each link in the path. The distributed routing approach [7] is an example of this scheme.

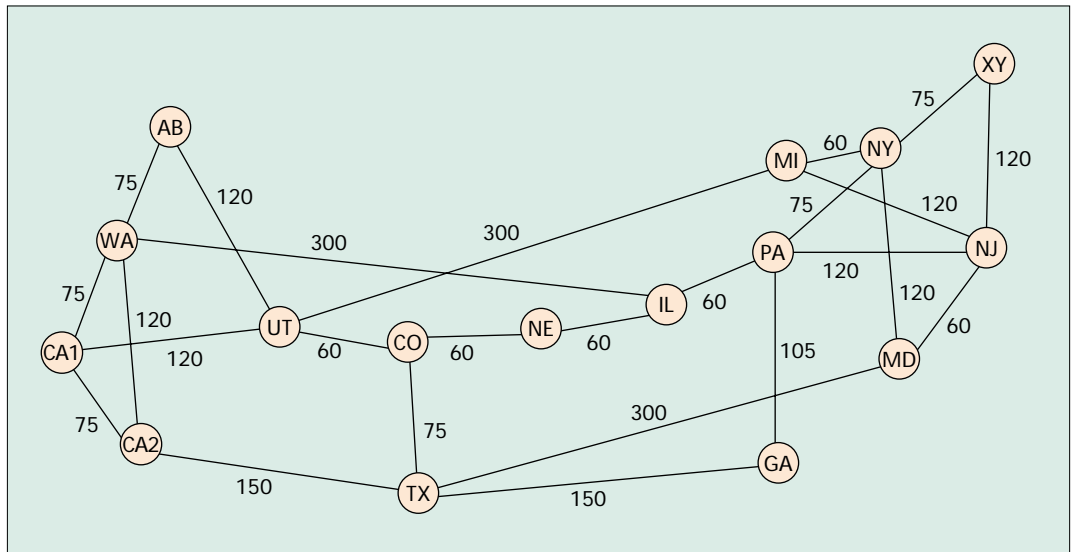
For the case in which a node only knows the status of its immediate links, the source node may utilize a conservative reservation approach, choosing a single wavelength and sending out a control message to the next node attempting to reserve this wavelength along the entire path; however, there is no guarantee that the selected wavelength will be available along every link in the path. If the wavelength is blocked, the source node may select a different wavelength and reattempt the connection. The limitation of this approach is that it may result in long setup times, since it may take several attempts before a node can establish a lightpath.

An alternate approach to maximizing the likelihood of establishing a lightpath in a forward reservation scheme is to use an aggressive reservation scheme which overreserves resources [11]. When a reservation message arrives at a node, the node reserves all wavelengths available on all of the links traversed so far. When the reservation message reaches the destination node, the destination then chooses one wavelength from the wavelengths reserved along the entire path and releases the reservations on the remaining wavelengths.

The drawback of this forward reservation scheme is that network resources are overreserved for a short period of time, which may lead to the blocking of subsequent connection requests and lower network utilization.

The advantage of a parallel reservation scheme is that it shortens the lightpath establishment time by having nodes process reservation requests in parallel. The disadvantage is that it requires global knowledge, since both the path and the wavelength must be known ahead of time.

Load is measured in Erlangs, which can be calculated by multiplying the connection arrival rate with the average connection holding time. Hence the load refers to the average number of connections measured at any instance of time in the network if there is no blocking.



■ Figure 2. NSFNET: a nationwide backbone network.

Backward Reservation — To prevent over-reservation of resources altogether, reservations may be made after the control message has reached the destination and is headed back to the source. Such reservation schemes are referred to as *backward reservation* schemes [11]. In backward reservation, the source node sends control packets to the destination without reserving any resources. These control packets will collect information about wavelength usage along one or more paths, and the destination will then utilize this information to decide on a route and a wavelength. The destination then sends a reservation message to the source nodes along the chosen route, and this reservation message will reserve the appropriate network resources along the way [9, 11]. One possible drawback of a backward reservation scheme is that if multiple connections are set up simultaneously, it is possible that a wavelength available on a link in the forward direction will be taken by another connection request and no longer be available when the reservation message traverses the link in the reverse direction.

TWO CONNECTION MANAGEMENT APPROACHES

PROTOCOLS

We compare two connection management approaches over various metrics. The first, the link state approach, is proposed in [6]. The second, the distributed-routing approach, is presented in [7]. Both approaches have been described in previous sections.

We assume the signaling messages are delivered in a packet-switched control network. This control network is implemented on an out-of-band supervisory channel that operates on its own wavelength. The control layer has the same topology as the physical network,¹ and all packets are routed by shortest paths.

Since the control network is packet-switched, the signaling method described in [6] consumes much control bandwidth and involves longer

connection setup delay than a hop-by-hop signaling method [7]. In order to look into the routing issues more closely, we modify the signaling method in [6] to be hop-by-hop signaling. Moreover, in the modified link state approach, each node computes only its next hop based on the topology information. This is different from [7] in which source routing is applied (i.e., the full route is determined at the source node). Therefore, the only difference between the compared approaches is how the routing information is updated and RWA is performed.

The following summarizes the two approaches under consideration:

- **Routing:** In both approaches, routing is done with global information. However, in the link state approach, each node maintains a database of network topology and the wavelength state on each link; LSAs are used to update the topology and wavelength usage information. In the distributed routing approach, a distance-vector protocol is executed to keep the routing tables up to date.
- **Wavelength assignment:** A first-fit approach is used in both cases.
- **Signaling procedures:** A similar signaling procedure is used in both approaches. After the source node determines the route or next hop, it sends out a RESERVE message to its next hop. Each intermediate node will examine the requested resources. If the resources are available, the node will reserve the resource and send a RESERVE message to the next hop; otherwise, the node sends a RESERVE-NACK back to the source. After the destination receives the RESERVE message, it checks whether or not it has a spare receiver on the requested wavelength. If the node has a spare receiver, it sends a RESERVE-ACK back to its previous hop; otherwise, it sends a RESERVE-NACK. The switches are configured when a node receives a RESERVE-ACK. The node is also responsible for delivering the RESERVE-ACK to its previ-

¹ In the control layer, we can set up the lightpaths in a way that every lightpath occupies a single fiber link. Hence, the virtual topology in the control layer is exactly same as the physical topology.

ous hop on the route. If a node receives a RESERVE-NACK, it releases the reserved resources. When the source receives a RESERVE-NACK, it performs RWA again and attempts to set up the connection on another route and wavelength. If such a route and wavelength cannot be found, the connection is blocked. When the source receives a RESERVE-ACK, the connection setup is successful, and the source may begin sending data over the connection. In order to prevent a connection request from being reattempted too many times, a parameter M is used to control the maximum number of times a connection may be attempted. A connection is blocked after the M th attempt fails.

- Update procedures: Both approaches use incremental updates. In the link state approach, each LSA contains the information about one channel on one link. In the distributed routing approach, each node keeps a copy of every neighbor's routing table, and each update message only contains the recently changed entries in the sending node's routing table.

COMPARISON

We compare the performance of the two approaches via simulation on the nationwide network NSFNET (Fig. 2). NSFNET has 16 nodes and 25 links, and the link lengths range from 750 to 3000 km. Each link is bidirectional fiber, and the number on the links in Fig. 2 represent the length of the links in units of 10 km.

We assume the following:

- The number of wavelengths on each link, W , is 8.
- The traffic is uniformly distributed among all node pairs.
- Connection holding time is exponentially distributed with mean 100 ms.
- Message processing time at a node, P , is 10 μ s.
- The time to configure, test and set up a cross-connect, C , is 500 μ s.
- The time to transmit or switch a packet in the control network, R , is 0.
- The shortest path obtained in adaptive routing is defined as the path with the minimum number of hops. Under uniform traffic and low load, the average propagation delay between two nodes is $D = 14.7$ ms and the average hop distance is $H = 2.28$. Signaling messages are routed via the path with the shortest propagation delay in the control network.
- The number of transceivers on each wavelength at each node, TR , is a parameter of the simulation; $TR = 1, 2, 3$.
- No reattempt is performed when a connection is blocked.

We obtained simulation results over the NSFNET topology by simulating a total of 10^4 connection requests that arrive and depart from the network over time. To study the network's behavior under different loads, the arrival rate of connection requests is varied as a parameter in the simulation. Load is measured in Erlangs, which can be calculated by multiplying the con-

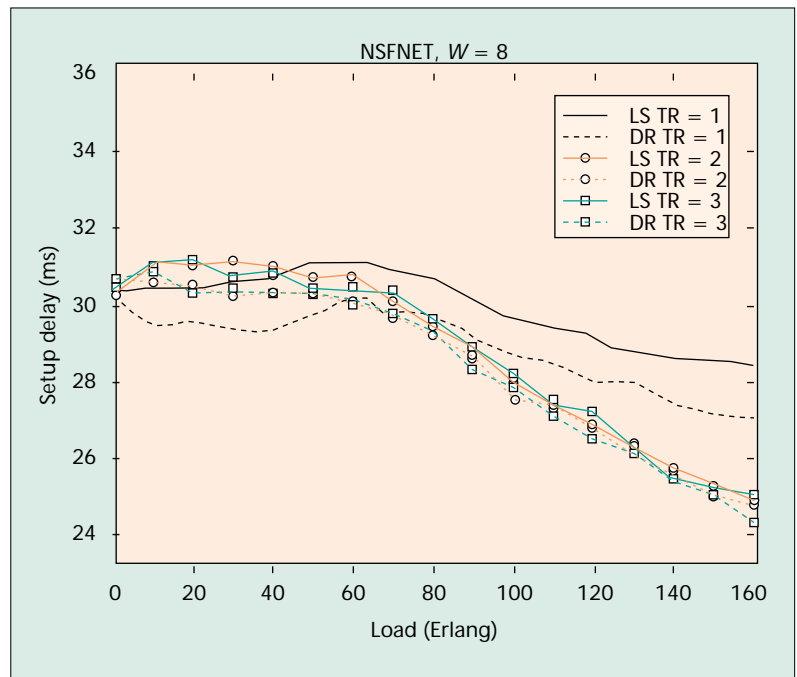


Figure 3. Connection setup delay vs. load for NSFNET and 8 wavelengths.

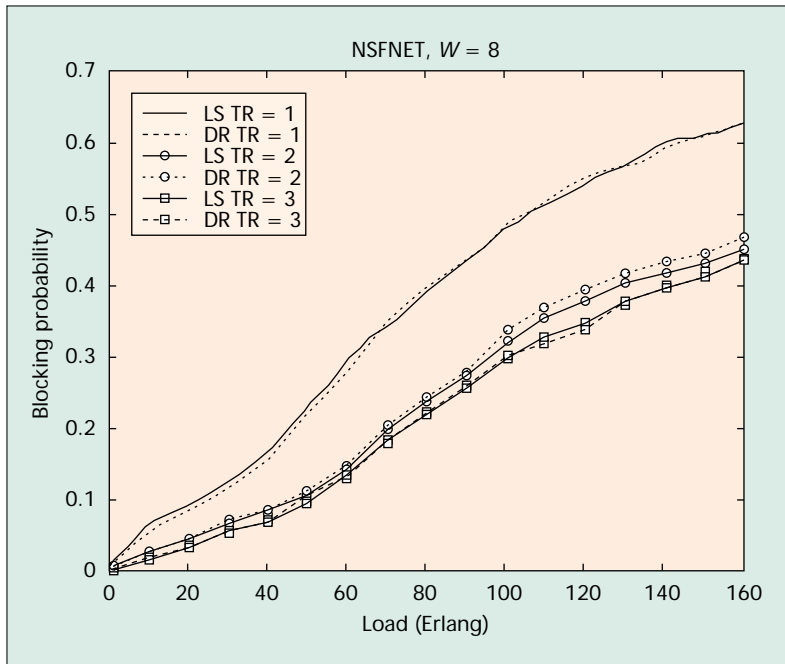
nection arrival rate with the average connection holding time. Hence the load refers to the average number of connections measured at any instance of time in the network if there is no blocking.

Connection Setup Delay — Connection setup delay is the time required to establish a connection once a connection request arrives. When the network is very lightly loaded (i.e., the shortest path is available for all connections) and there are no reattempts, the following elements contribute to the average setup delays for both connection management approaches:

- 2 propagation delays from source to destination node, $2D$
- $2H + 1$ message processing delays, $(2H + 1)P$
- Switch configuration time, C

Hence, the connection setup time under light load is $T_D = 2D + (2H + 1)P + C = 30.02$ ms. Figure 3 plots the setup delays vs. load in NSFNET, with different numbers of transceivers (TR). We observe that when the load is very low, the setup delays are fairly close to this lower bound. As the load increases, the shortest paths may become unavailable, and longer paths must be selected. Therefore, the connection setup delay may increase as load increases. However, the call blocking rate also increases when the load increases. A connection that spans more hops is more likely to be blocked than a connection that spans fewer hops; thus, as the load continues to increase, the connection setup delays will decrease.

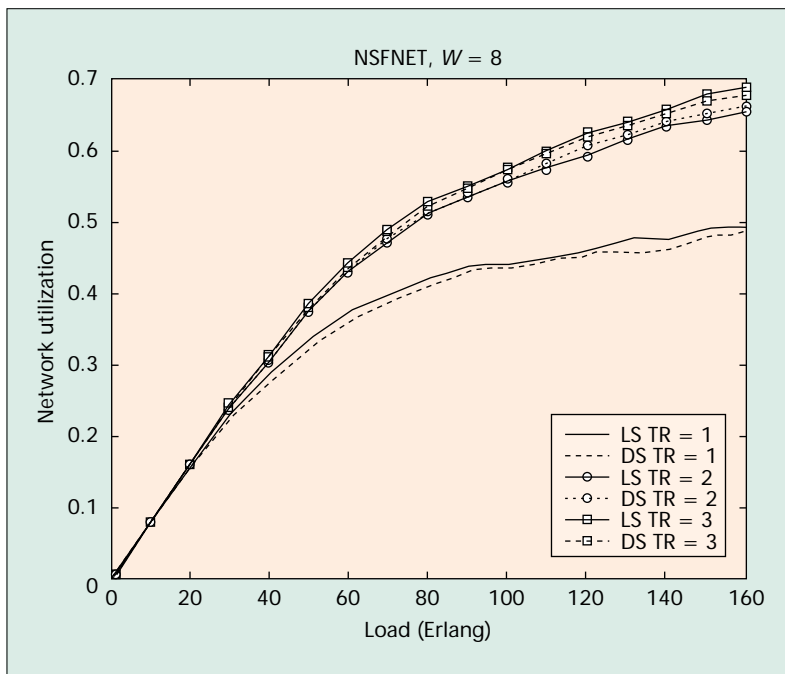
It is interesting to observe that the distributed routing approach gives lower connection setup delays than the link state approach. Both approaches attempt to find the path with the minimal hop count. In the link state approach, if there exist multiple paths to the destination, one



■ Figure 4. Blocking probability vs. load for NSFNET and 8 wavelengths.

is chosen randomly. However, in the distributed routing approach, the routing tables are exchanged between neighboring nodes; thus, the routing table from a path with a shorter propagation delay will reach a node first. Therefore, among paths of the same hop count, the path with the shortest propagation delay will be saved in the routing table.

Blocking Probability — Blocking probability refers to the probability that a connection cannot be established due to resource contention along the desired route. Figure 4 plots the block-



■ Figure 5. Network utilization vs. load for NSFNET and 8 wavelengths.

ing probabilities vs. load for the two approaches. It has been shown that blocking in the link state approach is slightly lower than in the distributed-routing approach under low load, but slightly higher under certain high load. These differences are due to the facts that under low load, the link-state approach has more accurate routing information which comes from shorter stabilizing delays (Fig. 6); under high load, both approaches may not have up-to-date information on routing, but setup delays in the link-state approach are higher. Hence, resources are reserved for a longer period of time for each connection in the link state approach under high load.

Figure 5 plots the network utilization vs. load obtained through simulation. When each node has only one transceiver on each wavelength ($TR = 1$), we observe that the network saturates at around 50 percent for a load of 160 Erlangs (where around 60 percent of the connections are blocked). When more transceivers are available ($TR > 1$), the network utilization is close to 70 percent for a load of 160 Erlangs (where around 45 percent of the connections are blocked). This performance is not a limitation of the routing approaches, but rather of the number of transceivers at each node (when $TR = 1$), as well as the wavelength-continuity constraint.

Stabilizing Time — Stabilizing time is the time required for nodes to update topology information after a connection has been established or taken down. In the link state approach, the stabilizing time is equal to the time it takes for a node's update message (LSA) to be delivered to the farthest node, which we denote as T_i for node i . T_i can be computed as follows. For each node i in the network, and $j \neq i$, find the shortest route by the minimal propagation delay from i to j . Denote the number of hops on this route H_{ij} and the propagation delay d_{ij} . If the time to transmit/switch an LSA is R , then the time it takes for node i 's LSA to reach j is $H_{ij}R + d_{ij} = d_{ij}$, since we assume $R = 0$.

We then find j for each node i such that d_{ij} is maximum. Thus, we have

$$T_i = \max_j d_{ij}.$$

The average stabilizing time in the link state approach will then be $(1/N)\sum_i T_i = 23.2$ ms for NSFNET, where N is the number of nodes in the network. From simulation, the average stabilizing delay is observed to be in the range [22.8 ms, 23 ms] for NSFNET, as plotted in Fig. 6. The stabilizing delay for the distributed routing approach is studied in simulation only, since the delay for this case is difficult to model. We notice from Fig. 6 that the distributed routing approach has larger stabilizing delays than the link state approach in most cases. This is because in the distributed routing approach, it usually takes several rounds of information exchange for the network to stabilize, especially in response to "bad" news. The stabilizing time for the distributed routing scheme decreases as the load increases, and may be lower than that of the link state approach under very high load. This is because

network resource utilization is fairly high at these loads, and most routes are unavailable in the network. Hence, a change in wavelength usage does not affect the rest of the network as much as for lower loads.

We also notice that in the distributed routing approach, the stabilizing delays increase as the number of transceivers decreases. This is because when each node has fewer transceivers, a change of state (connection being set up or taken down) will have a larger impact on the rest of the nodes. For example, when $TR = 1$, any connection being set up from or to a node means other nodes cannot connect to this node on this wavelength; hence, this information has to be propagated to every node in the network. When $TR > 1$, this information may not be necessary for some nodes.

CONCLUSION

The establishment of lightpaths in wavelength-routed WDM networks requires the implementation of control and management protocols to perform routing and wavelength assignment functions, as well as to exchange signaling information and reserve resources. In this article we present some of the routing, wavelength assignment, and signaling protocols for establishing lightpaths in a wavelength routed network.

In routing and wavelength assignment algorithms for dynamic lightpaths, the goal is to minimize the number of blocked connections. The performance of these algorithms depends on the amount of state information available to each node. If global information is known, the routing and wavelength assignment decisions can be nearly optimal; however, it may be difficult to maintain complete up-to-date information in a very dynamic environment.

The performance of signaling protocols for reserving wavelengths along a lightpath will depend on whether or not global information is available, and whether or not multiple connection requests may be attempted simultaneously. For the case in which global information is available, reservations may be made either in parallel or on a hop-by-hop basis, with parallel reservations leading to lower connection establishment times. When only local information is available, wavelength selection may be combined with the reservation scheme. Reservations may be made in the forward or backward direction. In forward reservation, overreservation of wavelength resources may lead to higher blocking for other connections, while in backward reservation, it is possible that a previously available link in the route will be taken by another connection request.

We compare two control and management approaches: link state and distributed routing. We observe that the link state approach has advantages in shorter stabilizing delays. Under low load, it also has lower blocking probability than the distributed routing approach. Distributed routing has advantages in shorter connection set-up delays and, under high load, lower blocking probability.

The link state approach also has an advantage in traffic engineering. Since each node

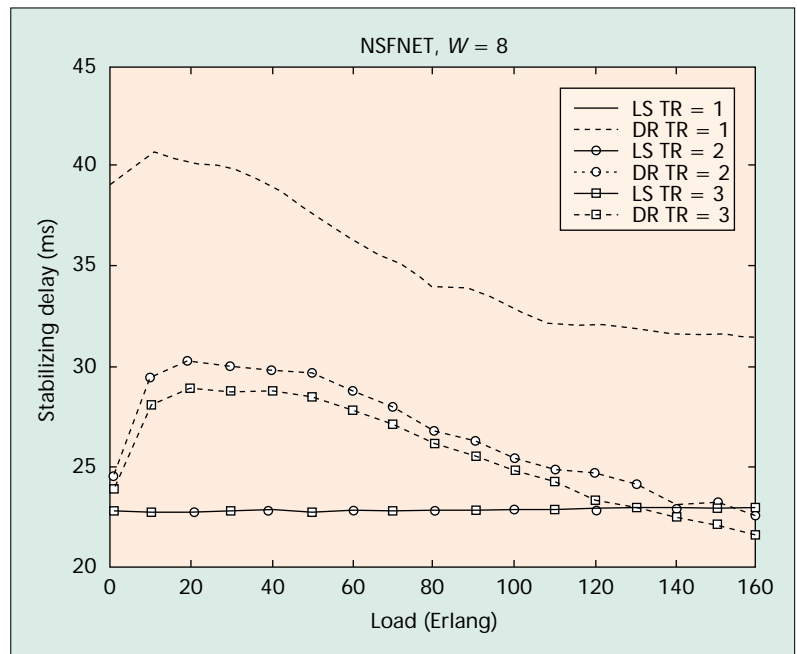


Figure 6. Stabilizing delay v.s load for NSFNET and 8 wavelengths.

maintains global information about the network, explicit routing can be implemented. This attribute can add more fault tolerance to the network. For example, it is simple and fast to compute two link-disjoint routes at the source node. It also makes shared protection possible with the knowledge of full network topology.

REFERENCES

- [1] H. Zang, J. P. Jue, and B. Mukherjee, "A Review of Routing and Wavelength Assignment Approaches for Wavelength-Routed Optical WDM Networks," *Opt. Net. Mag.*, vol. 1, no. 1, Jan. 2000, pp. 47–60.
- [2] B. S. Davie, P. Doolan, and Y. Rekhter, *Switching in IP Networks: IP Switching, Tag Switching and Related Technologies*, Morgan Kaufmann, 1998, pp. 88–89, 194–95.
- [3] A. Copley, "Optical Domain Service Interconnect (ODSI): Defining Mechanisms for Enabling On-Demand High-Speed Capacity from the Optical Domain," *IEEE Commun. Mag.*, vol. 38, no. 10, Oct. 2000, pp. 168–74.
- [4] A. Banerjee *et al.*, "Generalized Multiprotocol Label Switching: An Overview of Routing and Management Enhancements," *IEEE Commun. Mag.*, vol. 39, no. 1, Jan. 2001, pp. 144–50.
- [5] R. Perlman, *Interconnections, Second Edition: Bridges, Routers, Switches, and Internetworking Protocols*, Addison Wesley Prof. Comp. Series, Oct. 1999, pp. 320–25.
- [6] R. Ramaswami and A. Segall, "Distributed Network Control for Optical Networks," *IEEE/ACM Trans. Net.*, vol. 5, no. 6, Dec. 1997, pp. 936–43.
- [7] H. Zang *et al.*, "Connection Management for Wavelength-Routed WDM Networks," *Proc., IEEE GLOBECOM '99*, Rio de Janeiro, Brazil, Dec. 1999, pp. 1428–32.
- [8] K. Chan and T. P. Yum, "Analysis of Least Congested Path Routing in WDM Lightwave Networks," *Proc., IEEE INFOCOM '94*, Toronto, Canada, vol. 2, Apr. 1994, pp. 962–69.
- [9] L. Li and A. K. Somani, "Dynamic Wavelength Routing Using Congestion and Neighborhood Information," *IEEE/ACM Trans. Net.*, vol. 7, no. 5, Oct. 1999, pp. 779–86.
- [10] J. P. Jue and G. Xiao, "An Adaptive Routing Algorithm with a Distributed Control Scheme for Wavelength-Routed Optical Networks," *Proc., 9th Int'l. Conf. Comp. Commun. Net.*, Las Vegas, NV, Oct. 2000, pp. 192–97.
- [11] X. Yuan *et al.*, "Distributed Control Protocols for Wavelength Reservation and Their Performance Evaluation," *Phot. Net. Commun.*, vol. 1, no. 3, 1999, pp. 207–18.

The establishment of lightpaths in wavelength-routed WDM networks requires the implementation of control and management protocols to perform routing and wavelength assignment functions.

BIOGRAPHIES

HUI ZANG (hzang@sprintlabs.com) received her B.S. degree in computer science from Tsinghua University, China, in 1997, her M.S. degree in computer science from the University of California, Davis in 1998, and her Ph.D. in computer science from the University of California, Davis in June 2001. From June 1999 to September 1999 she worked at IBM Almaden Research Center, San Jose, California, as a summer intern. From October 1999 to September 2000 she worked at Sprint Advanced Technology Labs, Burlingame, California, as a part-time intern, where she became a principal R&D engineer in October 2000. Her research interests include photonic packet switching, WDM network control and management, and fault tolerance in WDM networks.

JASON P. JUE (jjue@utdallas.edu) received his B.S. degree in electrical engineering (with honors) from the University of California, Berkeley in 1990, his M.S. degree in electrical engineering from the University of California, Los Angeles in 1991, and his Ph.D. in electrical and computer engineering from the University of California, Davis in 1999. In 1999 he joined the faculty at the University of Texas, Dallas, where he is currently an assistant professor of computer science and a member of the Center for Advanced Telecommunications Systems and Services. His research interests are in the area of high-speed communication networks, with an emphasis on the design and performance evaluation of WDM optical networks.

LAXMAN SAHASRABUDDHE (laxman@summitnetworks.net) received his B.Tech. degree from the Indian Institute of Technology, Kanpur, in 1992, his M.Tech. degree from the

Indian Institute of Technology, Madras, in 1994, and his Ph.D. from the University of California, Davis in 1999. He is the recipient of the Best Doctoral Dissertation Award from the College of Engineering, University of California, Davis, for his research on WDM optical networks. From 1999 to 2000 he was with Amber Networks as an embedded software engineer. Currently, he is the principal optical routing architect at Summit Networks Inc.

RAMU RAMAMURTHY (RRamamurthy@tellium.com) is currently a senior network architect at Tellium, Inc., where he works on the design of algorithms and protocols for dynamic provisioning and restoration in optical networks. Prior to this, he was a research scientist at Telcordia Technologies, where he worked on network control and management of IP/WDM optical networks. He holds a B.Tech. degree from the Indian Institute of Technology, Madras, and M.S. and Ph.D. degrees from the University of California, Davis.

BISWANATH MUKHERJEE received his B.Tech. (Hons.) degree from the Indian Institute of Technology, Kharagpur, in 1980 and his Ph.D. from the University of Washington, Seattle, in June 1987. While at Washington, he held a GTE Teaching Fellowship and a General Electric Foundation Fellowship. In July 1987 he joined the University of California, Davis, where he has been a professor of computer science since July 1995 and served as chairman of computer science from September 1997 to June 2000. He is author of the textbook *Optical Communication Networks* (McGraw-Hill, 1997), which received the Association of American Publishers, Inc.'s 1997 Honorable Mention in Computer Science. His research interests include lightwave networks, network security, and wireless networks.