Routing and Wavelength Assignment Methods in All-Optical Networks

A. Birman and A. Kershenbaum
IBM T.J. Watson Research Center, Yorktown Heights, NY 10589, USA

September 1, 1994

Abstract

For all-optical networks using a wavelength routing architecture and no wavelength converters we consider several methods for routing and wavelength assignment. These methods combine in various ways three main principles: logical links, wavelength reservation and alternate routes. The methods are evaluated and compared in two case studies. In this type of networks the traffic over multi-hop routes is susceptible to high blocking probabilities, which can interfere with the quality of service the user requires. It is shown that some of the routing and wavelength assignment methods studied have the potential to overcome this difficulty.

1 Introduction

We study all-optical networks with Wavelength Division Multiplexing and wavelength routing, targeted to wide-area applications. In such networks a node is able to route each wavelength on the incoming link to any outgoing link. The same wavelength, however, on two incoming links cannot be routed simultaneously onto a single outgoing link. There are no optical-to-electronic conversions, and thus there is no buffering at intermediate nodes. A number of prototyping projects involving such networks are under way [7].

Each connection is assigned a path through the network and a wavelength, which must be the same wavelength on all links in that path. Moreover, the assignment of wavelengths to paths must be such that no two paths which have a link in common are assigned the same wavelength. We thus consider all-optical networks without wavelength converters. We assume that requests arrive independently of one another and that the precise time of their arrivals is not known in advance, although we do assume that the average rate of arrivals is known. Thus, while they may in fact be requests for bandwidth for video transmission or file transfers, they are in this case simply requests for a single wavelength on each link in the path to which they are assigned. Connection requests can be thought of as *calls*, as in telephone networks, and we will refer to them as such in the remainder of the paper.

We assume that an arriving call is either assigned a wavelength if an idle wavelength is available on all links on the chosen route between its source and destination, or else it is rejected, or blocked, by the network. We also assume that a link topology is given and that the same set of wavelengths is in use on all links.

The requirement to allocate the same wavelength on all links of a route, also known as the wavelength continuity constraint, increases the blocking probability substantially for routes with many links [5, 3]. The use of wavelength converters eliminates this problem and reduces the network to a circuit-switched network. The difficulty with this approach is that wavelength converters are difficult to build and they add substantially to both the cost and complexity of the network. Another possible approach is to allow, at the time a call arrives, a re-assignment of wavelengths to calls which may sometimes result in a call being serviced instead of blocked [10]. For networks with unrestricted topology, however, this method may also add to the complexity and cost of operating the network while compensating only partially for the wavelength continuity constraint.

For the class of all-optical networks without wavelength converters we describe several methods for routing and wavelength assignment. These are: Fixed Routing, Fixed Alternate Routing, Fixed Routing over Logical Links, Fixed Routing with Wavelength Reservation, and Fixed Alternate Routing over Logical Links. These methods are based on three main principles: logical links, wavelength reservation and alternate routes. The methods are evaluated and compared in two case studies and it is shown that some of the routing and wavelength assignment methods studied have the potential to overcome the problem of blocking on multi-hop routes due to the wavelength continuity constraint.

While the routing problem problem in circuit switched networks has been extensively studied [8, 6], routing in the all-optical networks considered here is relatively new. The problem of optimally routing and assignment of wavelengths has been studied in [2, 5] where several heuristic algorithms were proposed and their performance evaluated through simulation. In [12] a lower bound on the blocking probability for any routing and wavelength assignment algorithm was given. An analytical method for calculating the blocking probability for some routing and wavelength assignment schemes is given in [3].

In Section 2 which follows we describe the routing and wavelength assignment methods under study. In Section 3 these methods are evaluated and compared using two case studies. Concluding remarks are found in Section 4.

2 Routing and Wavelength Assignment Methods

In this section we describe the methods we consider for assigning routes and wavelengths to calls on their arrival. Our principal goal is to compare these approaches with one another rather than to present an optimal implementation of any of them. In order to make a meaningful comparison, however, it is important to implement each approach in a reasonable way. We thus present a straightforward implementation of each approach, but justify its validity as we do so.

Ideally, the topology of the network and the routes should be chosen together, given an estimate of the traffic that will be offered to the network. If this is done, it is possible to create a topology and routes which are well suited to one another in the sense that the routes make good use of the available capacity or, conversely, that there is capacity in places where the routes can make use of it. A network created in this way will usually perform better than one where these two problems were solved separately, even if the solution to each subproblem is highly optimized.

We thus chose to get a good solution to both the topology and route determination problems simultaneously rather than trying to get an optimal solution to the routing problem in isolation. We used INTREPID [4], an automated tool which has been used for the design of a wide variety of networks and which simultaneously produces topologies and routes, to obtain solutions to this problem. INTREPID uses the MENTOR [9] algorithm to select links based on a goal of having a reasonable link utilization. The links it selects have utilizations which are high enough to be cost effective without being so high as to result in severely degraded performance. In order to identify such links, MENTOR also produces a route for each traffic stream so that it can be loaded onto the network, thereby allowing link utilizations to be computed and assessed. Here, we associated a cost with each potential link in the network and used these costs both to optimize the network cost and as a starting estimate for link lengths in a shortest path routing algorithm. Thus, the routes are shortest paths based on link lengths which, in turn, are based on (but not necessarily equal to) the link costs. Because the link costs include a significant fixed charge, the routes produced in this way are also min-hop routes.

2.1 Fixed Routing

The simplest approach is to assign each call to the single pre-determined route associated with the arrival stream, and to allow the call to use any wavelength which is free on all links of the route. To implement this approach, one must solve two problems. First, in setting up the network, one must associate a specific route with each call. Then, when a call arrives, one must assign it to a specific free wavelength on the route.

There are many approaches for generating routes [1, 6, 11] in order to obtain a good grade of service for all calls. All these approaches try to distribute the calls over the available network capacity to avoid creating bottlenecks; i.e, overloaded routes which block a lot of traffic. In doing so, however, one may be forced to send traffic over routes with more hops than the minimum necessary. Thus , there is a trade-off between these two factors.

There are several possible ways of selecting one wavelength from the set of idle wavelengths for an arriving call on its chosen route. The simplest would be to assign it at random. This would tend, however, to distribute the calls evenly over all wavelengths and hence to create the undesirable situation that there are different free wavelengths on different links in a route. A way of avoiding this would be to assign calls to the lowest available wavelength, making it more likely that a high wavelength would be available on all links simultaneously.

Another possibility along the same lines would be to assign single hop calls to the lowest available wavelength and to assign multi-hop calls to the highest one. Numerical results show that choosing the lowest wavelength among the set of idle wavelengths is best for obtaining low blocking probabilities accross all traffic streams.

2.2 Fixed Routing over Logical Links

In order to overcome the problem the multi-hop calls have of finding the same wavelength free on all links of their route, we can assign specific wavelengths to each route in advance.

Then, when a call arrives, we are guaranteed that if one of these wavelengths is available it will be available simultaneously on all links in the route.

More generally, we refer to a set of wavelengths on a sequence of links in tandem as a *logical link*. In conventional telecommunication networks, logical links are often created with multiplexors. Thus, for example, if we have two physical links in tandem and each has C channels, we could define a logical link between the end nodes by allocating C_1 channels exclusively to the end-to-end traffic. A call between the two end nodes will then be routed over the logical link instead of being switched through the middle node. This is attractive in telecommunication networks because multiplexors tend to be faster and less expensive than switches. In an all-optical network, we have the added advantage that creating logical links correlates the use of specific wavelengths on tandem links in a route. In general, a route may consist of any sequence of links, be they physical or logical. In the small examples considered in this paper, if a logical link is used then it constitutes the entire route. This allows us to study the effects of creating logical links more clearly without complicating the problem by involving route selection.

There are two problems with using logical links. The first is that the approach fragments the link capacity, creating smaller groups of wavelengths from the original larger group. This is a central trade-off considered in this paper. We can correlate the times a wavelength is free on tandem links in a route, but only at the expense of restricting the use of that wavelength to calls which use the logical links.

The second problem is that we must come up with a way of effectively assigning wavelengths to logical links, which is not trivial. When each route is defined as a single logical link and when all traffic streams have the same rate then it is best to partition the pool of wavelengths equally among all logical links on a physical links. We are able to use this approach in the case study of Section 3.2.

2.3 Fixed Alternate Routing

If no free wavelength is available on the primary route defined for a given call, it is possible to offer it to one or more additional routes before blocking it. As in the case of fixed routing, we assume all routes have been defined beforehand. As in the case of fixed routing, we assign calls to the lowest free wavelength.

In implementing this approach, we use the routes defined in Section 2.1 as primary routes and also define one or more alternate routes for each multi-hop call. Whenever possible, we select an alternate route which is link-disjoint from the primary route (reducing the chance that of blocking on both routes because of congestion on a common link) and also has the minimum possible number of hops. When there is a conflict between these two objectives, we consider both possibilities. In the case study of Section 3.2 we offer one or more alternate routes only to multi-hop calls.

2.4 Fixed Routing with Wavelength Reservation

A logical link creates a partitioning of the set of wavelengths and thus prevents sharing. This leads to the undesirable situation that an arriving call may be blocked even though a free wavelength may exist along its physical path. This is particularly troublesome in cases where the call is a multi-hop call which is receiving a poor grade of service. On the other hand, we need not block a single hop call if there is an adequate number of free wavelengths currently available for the multi-hop calls.

Another approach is to use wavelength reservation, similar to trunk reservation in circuit-switched networks. Single-hop and multi-hop traffic have both access to all wavelengths. However, a single-hop call is assigned an idle wavelength only if the fraction of idle wavelengths from the entire set of wavelengths is greater than a specified threshold. This, in effect reserves capacity for the multi-hop routes when it is needed without unconditionally excluding the single-hop calls from access to any particular wavelength.

2.5 Fixed Alternate Routing over Logical Links

It is possible to combine the use of logical links, as described in Section 2.2 with the use of fixed alternate routing. We associate a pre-determined fixed route with each arrival stream, as in fixed routing. On some of these routes we may define some logical links, which need not exhaust the pool of available wavelengths on every link. For an arriving call we attempt to assign an idle wavelength in two steps. First, we attempt to find an idle wavelength on the logical link, if one is defined. Second, if unsuccessful at the first step, we attemp to find an idle wavelength on all the links in the route from the set of wavelengths which do not belong to a logical link. If such an idle wavelength is not found then the call is blocked.

While this approach can be generalized to more than one route, and to a more general interpretation of logical links, we chose to implement it in this 'minimal' way in order to keep things simple and thus make it easier to interpret the numerical results.

3 Numerical results

The wavelength assignment and routing schemes described in the previous section are now applied to two case studies. First we consider a simple network consisting of three links in tandem. Then we consider a mesh network with six nodes and nine links. All numerical results were obtained by simulation unless specified otherwise.

3.1 Three links in tandem

Consider a network consisting of three links in tandem. All links have the same capacity (number of wavelengths) $C = 12$. There are three 'local' traffic streams each involving a single link, and one end-to-end traffic stream involving all three links. The local traffic streams all have the same offered load while the offered load of the end-to-end traffic stream is 30% of the local offered load.

Numerical results [13] for the three links in tandem are shown in Table 1. There are four rows in the table, corresponding to several loadings of the network. The entry in column 'Offered load' consists of two numbers representing Erlangs- the first is the offered load of the local traffic, the second is the offered load of the end-to-end traffic. The column 'Link loading' represents (in percentages) the ratio between the offered load to each link and the capacity C. The remaining three columns show the probability of a call being blocked because there

Offered load (local/end-to-end)	Link loading (%)	Blocking Probability (%)		
		incr./incr.	incr./decr.	random/random
4.0/1.20	43.3	0.44/2.71	0.41/2.77	0.32/9.10
4.5/1.35	48.8	0.84/5.68	0.85/5.51	0.62/15.46
5.0/1.50	54.2	1.49/10.73	1.48/9.84	1.04/22.56
6.0/1.80	65.0	3.27/23.90	3.35/22.78	2.48/40.90

Table 1: Three links in tandem: blocking probabilities for three wavelength assignment schemes. Offered load is in Erlangs. C=12.

are no idle wavelengths on the links associated with its route. Blocking probabilities are given for three wavelength assignment schemes. All entries consist of two numbers- one for the local traffic, the other for the end-to-end traffic. The wavelength assignment schemes differ in the way one wavelength is picked from the pool of idle wavelengths in order to satisfy an arriving call. The wavelengths are assumed ordered in some way. In the first scheme, column 'incr./incr.', the pool of idle wavelengths is scanned in increasing order and the first idle wavelength is allocated whether the call is local or end-to-end. In the second scheme, column 'incr./decr.', the pool of idle wavelengths is scanned in increasing order for local calls but in decreasing order for end-to-end calls, the first idle wavelength is allocated. In the last scheme, column 'random/random', an idle wavelength is picked at random for both types of traffic.

The numerical results in Table 1 show that there is only a slight difference between the first two schemes. Scanning the idle wavelength in decreasing order for the multi-hop traffic seems to improve the blocking probability somewhat, but the effect is minor. A significant degradation occurs, however, if the idle wavelength is chosen at random. The end-to-end traffic in this case is blocked more often, which makes this scheme less desirable then the first two.

As the offered load increases, all blocking probabilities increase, as one would expect. However, the rate of increase for the end-to-end traffic relative to the rate of increase for the local traffic is not the same for different schemes. For the 'incr./decr.' scheme these rates are about the same. For the 'incr./incr.' scheme the rate of increase for the end-to-end traffic is larger than for the local traffic, while for the 'random' scheme the opposite is true. For practical purposes only the 'ordered' schemes, that is 'incr./incr.' and 'incr./decr.', are of interest. For Fixed Routing using these schemes the blocking probability increases much faster with the offered load than desirable in a real network, and thus the need for the routing and wavelength assignment methods of Section 2.

3.2 A mesh network

We consider the network in Figure 1 with six nodes and nine links. The links are numbered from 0 to 8. Note that nodes DAL and HOU are joined by two links, 4 and 8. The number of wavelengths on each link is $C = 24$. There are fifteen traffic streams, one for each pair of nodes. The offered load of each traffic stream is 5 Erlang.

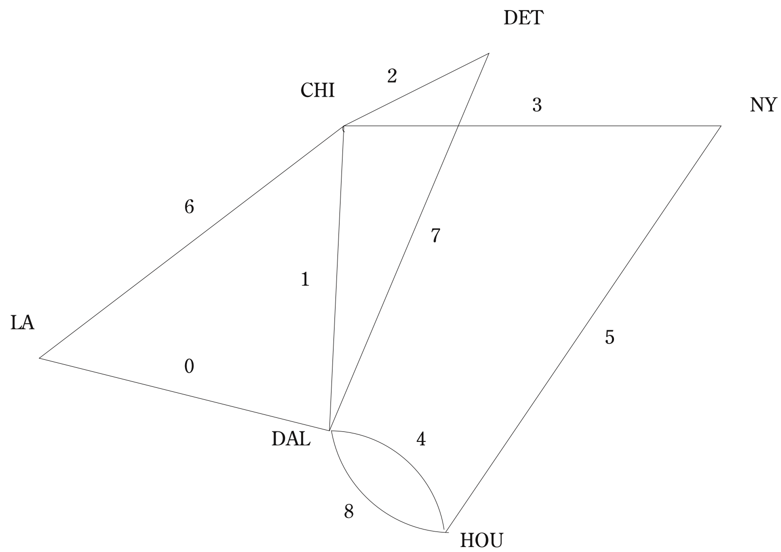

Figure 1: A mesh network

In the routing and wavelength assignment schemes we will consider below, one or more routes are associated with each traffic stream, and these routes are described in Table 2. The first column shows all the traffic streams. Column 'dir' shows the direct route for the traffic stream. If one or more alternate routes are defined for a traffic stream these appear in columns 'alt1' and 'alt2'.

Numerical results [13] for this network are shown in Table 3. Each row in the table corresponds to one of the traffic streams, as described in the entry of the first column. The next four columns show the call blocking probabilities under four routing and wavelength assignment schemes. Column FR shows the blocking under Fixed Routing, FRLL stands for Fixed Routing over Logical Links, FRWR stands for Fixed Routing with Wavelength Reservation, and FAR stands for Fixed Alternate Routing. Finally, for comparison purposes, the last two columns show the call blocking probabilities for a similar network with circuit-switching and two routing schemes: Fixed Routing (FR) and Fixed Alternate Routing (FAR).

Ideally, for any given routing and wavelength assignment scheme, the blocking probabilities should be roughly the same accross the traffic streams, since usually we want to provide for the multi-hop traffic the same grade of service provided for single-hop traffic. As the numerical results show Fixed Routing (column FR) does not display this desired property. The LA-HOU traffic, for example, has a blocking probability of 2.84% while the CHI-DAL traffic has practically no blocking.

The FRLL scheme attempts to improve the situation by allocating blocks of wavelengths to specific traffic streams. While the results are good for some traffic streams which are allocated a block of 16 wavelengths, they are worse for others for which only 8 wavelengths are available. The results in this column were calculated using Erlang's formula: E(5,8)=0.0700, E(5,16)=0.0000.

Traffic req. (source-dest.)	Route		
	dir	alt1	alt2
LA-CHI	6	-	-
LA-DET	0,7	6,2	-
LA-NY	3,6	0,8,5	0,1,3
LA-DAL	0	-	-
LA-HOU	0,4	0,8	6,1,4
CHI-DET	2	-	-
CHI-NY	3	-	-
CHI-DAL	1	-	-
CHI-HOU	1,4	3,5	-
DET-NY	2,3	7,8,5	7,4,5
DET-DAL	7	-	-
DET-HOU	7,8	7,4	2,3,5
NY-DAL	5,8	3,1	-
NY-HOU	5	-	-
DAL-HOU	4	-	-

Table 2: Routes defined for the mesh network: dir= direct route, alt1= first alternate route, alt2= second alternate route.

The Fixed Routing with Wavelength Reservation (FRWR) scheme attempts to decrease the blocking on the multi-hop routes by increasing the blocking on single-hop routes. A threshold is associated with each link; a wavelength is allocated to a single-hop traffic stream only if the fraction of idle wavelengths to C is above the threshold. For Table 3 two threshold values were used: 10% and 20%. Links 0,3,4,7,8 are given 10%, while the rest are given 20%. The numerical results show that the largest blocking probability (3.02% for the LA-DET traffic) has decreased slightly. We are still short, however, of the goal which is to bring the blocking probabilities for multi-hop traffic in line with the single-hop traffic.

Column FAR of Table 3 shows blocking probabilities for Fixed Alternate Routing. The alternate routes are described in Table 2. As shown, multi-hop traffic is given one or two alternate routes in order to improve its grade of service. The results show a significant improvement in blocking probabilities for multi-hop traffic. Moreover, the largest blocking probability is now well under 1%.

3.3 Hybrid methods

As can be seen from the results in the preceding section, it is possible to obtain reasonable performance using fixed alternate routing even without wavelength conversion or wavelength reservation. These results were obtained with offered load of 5 erlangs, which correspond to an average channel utilization of roughly 50%.

As the offered load rises, however, we would expect that performance would deteriorate to some extent. In Table 4 below, we see results when the offered load is raised to 6 erlangs

Traffic req. (source-dest.)	B_{WA} (%)				B_{CS} (%)	
	FR	FRLL	FRWR	FAR	FR	FAR
LA-CHI	0.00	0.00	0.21	0.01	0.01	0.01
LA-DET	3.02	7.00	2.23	0.02	1.75	0.00
LA-NY	1.06	7.00	0.56	0.08	0.92	0.02
LA-DAL	0.39	7.00	2.15	0.55	0.60	0.65
LA-HOU	2.84	7.00	2.12	0.18	1.40	0.05
CHI-DET	0.00	0.00	0.21	0.01	0.02	0.05
CHI-NY	0.60	7.00	2.83	0.56	0.70	0.71
CHI-DAL	0.00	0.00	0.14	0.00	0.02	0.02
CHI-HOU	0.95	7.00	0.53	0.02	0.79	0.01
DET-NY	1.14	7.00	0.69	0.03	0.57	0.01
DET-DAL	0.54	7.00	2.74	0.56	0.71	0.72
DET-HOU	1.01	7.00	0.61	0.03	0.66	0.01
NY-DAL	0.05	7.00	0.02	0.00	0.02	0.00
NY-HOU	0.00	0.00	0.21	0.00	0.00	0.00
DAL-HOU	0.55	7.00	2.68	0.56	0.55	0.61

Table 3: Blocking probabilities for wavelength routing (B_{WA}) vs. circuit-switching (B_{CS}). Wavelength assignment schemes: Fixed Routing (FR), Fixed Routing over Logical Links (FRLL), Fixed Routing with Wavelength Reservation (FRWR) and Fixed Alternate Routing (FAR). Circuit-switching schemes: Fixed Routing (FR) and Fixed Alternate Routing (FAR). Network with 6 nodes, 9 links and 15 traffic requirements. Offered load is 5 Erlangs for all requirements. C=24.

for each traffic requirement.

We first consider a circuit-switched environment. The column labelled C1 shows the loss probability for each traffic stream when fixed alternate routing and circuit switching are used. We see that four of the traffic streams now have losses in excess of 1%, our nominal objective. These are single-hop routes with no alternate routes. When we give each of these traffic streams an alternate route, we obtain the results shown in the column labelled C2. As can be seen, all traffic streams again have losses below 1%.

The conclusions we draw from this are that the original network is reasonably stable even if the traffic streams increase by 20%, but that the loss probability rises noticeably for some traffic streams without alternate routes. By providing these traffic streams with an alternate route, however, good performance is regained.

Next, we assume a wavelength routing architecture without wavelength converters and with the same alternate routes as in column C2. These results are shown in column C3. As can be seen, the network remains stable, but several traffic streams now have losses above 1%. Thus, we conclude that the use of wavelength converters, which would take us to the circuit-switched environment and the results in column C2, do help somewhat.

The degradation in performance between columns C2 and C3 is due to the constraint for multi-hop traffic streams to find the same wavelength free on both links in their paths. We attempt to compensate for this by creating logical links for the multi-hop traffic streams. These results are shown in the column labelled C4.

In this case, we created two sets of logical links. The first set was created from wavelengths 0 and 1 of each of the links. Specifically we created logical links (0,7), (1,8), (2,3), (4,5) and (6). Thus, channels 0 and 1 on links 0 and 7, for example, are used as a logical link to carry traffic between LA and Detroit. This link becomes the first route for that traffic stream. Similarly, the (1,8) logical link becomes the primary route for traffic streams between Chicago and Houston.

In a similar way, we create another set of logical links from channels 2 and 3. Specifically, we create links (0,8), (3,6), (4,7), (1), (2) and (5) and use these links as primary paths for traffic streams between their endpoints.

As can be seen, creating logical links helped noticeably, although not as much as wavelength converters did. We did not try to more finely optimize the logical links or to alter the alternate routes to try to improve performance further, but we believe it would be possible to do so.

The conclusion we draw from the experiments we have done thus far is that it is indeed possible to operate an optical network without wavelength converters and still get good performance (in terms of loss probability) even when the loads on the network are moderate. The more one wants to load up the network, however, the more carefully one will have to choose the routes and logical channels. Thus, we can think of wavelength converters not as essential to the operation of the network, but merely as an alternative to slightly lower link utilizations (and hence, somewhat higher link costs) or somewhat more complex routing algorithms. Obviously, these issues need to be explored further via additional analysis and experimentation and it is our intention to continue our investigation along these lines.

4 Concluding Remarks

In this paper we studied a class of all-optical networks using a wavelength routing architecture and no wavelength converters. We considered several methods for routing and wavelength assignment: Fixed Routing, Fixed Routing over Logical Links, Fixed Routing with Wavelength Reservation, Fixed Alternate Routing, and Fixed Alternate Routing over Logical Links. We evaluated and compared these methods in two case studies.

In the case study involving three links in tandem it became clear that with Fixed Routing the blocking probabilities for multi-hop routes are significantly higher than for single-hop routes. For practical considerations we need to control the discrepancy between the quality of service on various routes. In the case study involving a mesh network we evaluated the blocking probabilities for four routing and wavelength assignment methods. While the FR scheme displays again higher blocking probabilities over multi-hop routes, the other routing methods try, in different ways, to overcome this problem.

The FRLL method attempts to improve the situation by allocating blocks of wavelengths to specific traffic streams. While this can improve the performance over some traffic streams other traffic streams did poorly. The FRWR scheme attempts to decrease blocking on multi-hop routes by limiting the resources allocated to traffic on single-hop routes. A threshold mechanism is used similar to trunk reservation in circuit-switched networks. While the

Traffic req. (source-dest.)	C1 (FR,CS)	C2 (FAR,CS)	C3 (FAR)	C4 (FAR,LL)
LA-CHI	0.22	0.36	0.12	0.20
LA-DET	0.06	0.16	1.02	0.82
LA-NY	0.34	0.48	2.68	2.22
LA-DAL	3.74	0.04	0.04	0.04
LA-HOU	0.40	0.46	1.86	1.54
CHI-DET	0.22	0.38	0.18	0.10
CHI-NY	4.08	0.24	0.20	0.46
CHI-DAL	0.06	0.18	0.04	0.08
CHI-HOU	0.18	0.20	0.70	0.74
DET-NY	0.24	0.28	1.80	1.88
DET-DAL	3.84	0.00	0.00	0.10
DET-HOU	0.10	0.12	0.76	0.80
NY-DAL	0.08	0.10	0.38	0.38
NY-HOU	0.16	0.28	0.12	0.18
DAL-HOU	2.90	0.02	0.00	0.04

Table 4: Hybrid methods. Network with 6 nodes, 9 links and 15 traffic requirements. Offered load is 6 Erlangs for all requirements. C=24.

highest blocking observed in the FR case is somewhat alleviated the results are not much better. The FAR scheme attempts to improve the situation by providing alternate routes to the multi-hop traffic which experiences higher blocking. This scheme seems to work and we were able to control the quality of service over multi-hop routes.

The principles involved in the FRLL, FRWR and FAR scheme are different and could be combined in various ways to produce other, and possibly better, routing and wavelength assignment schemes. We explored the possibility of combining the FAR and FRLL schemes and showed that the two approaches combined well, yielding results which were an improvement over the schemes applied individually. At any rate, the results obtained here already indicate that the type of networks under investigation, using a wavelength routing architecture and no wavelength converters, can be made to work.

Acknowledgment. Thanks to Rajiv Ramaswami for his support and guidance throughout the research.

References

[1] J. Ash, R. Cardwell, and R. Murray. Design and optimization of networks with dynamic routing. *BSTJ*, 60:1787–1820, 1981.

[2] K. Bala, T.E. Stern, and K. Bala. Algorithms for routing in a linear lightwave network. *Proc. INFOCOM'91*, pages 1–9, 1991.

[3] A. Birman. Computing approximate blocking probabilities for a class of all-optical networks (preprint). 1994.

[4] R. Cahn, P.C. Chang, P. Kermani, and A. Kershenbaum. INTREPID: An integrated network tool for routing, evaluation of performance, and interactive design. *IEEE Communications Magazine*, pages 40–47, July 1991.

[5] I. Chlamtac, A. Ganz, and G. Karmi. Lightpath communications: An approach to high bandwidth optical wan's. *IEEE Transactions on Communications*, 40:1171–1182, 1992.

[6] A. Girard. *Routing and Dimensioning in Circuit Switched Networks.* Addison-Wesley, Reading, Mass., 1990.

[7] M.J. Karol, C. Lin, G. Hill, and K. Nosu, eds. *J. Lightwave Tech.: Special Issue on Broadband Optical Networks.* May 1993.

[8] F.P. Kelly. Routing and capacity allocation in networks with trunk reservation. *Math. of Operat. Res.*, 15:771–792, 1990.

[9] A. Kershenbaum, P. Kermani, and G. Grover. MENTOR: An algorithm for mesh network topological optimization and routing. *IEEE Trans. on Comm.*, 39:503–513, 1991.

[10] K.-C. Lee and V.O.K. Li. A circuit rerouting algorithm for all-optical networks. In *IEEE INFOCOM'94*, pages 954–961, June 1994.

[11] D. Mitra and J.B. Seery. Comparative evaluations of randomized and dynamic routing strategies for circuit-switched networks. *IEEE Transactions on Communications*, 39:102–116, 1991.

[12] R. Ramaswami and K.N. Sivarajan. Optimal routing and wavelength assignment in all-optical networks. In *IEEE INFOCOM'94*, pages 970–979, June 1994.

[13] H.D. Schwetman. CSIM Reference Manual. Technical report, Microelectronics and Computer Technology Corp., May 1992.