TERABIT LIGHTWAVE NETWORKS: THE MULTIHOP APPROACH

Anthony S. Acampora, Mark J. Karol, and Michael G. Hluchyj

Anthony S. Acampora *is director of the Transmission Technology Laboratory and* ***Mark J. Karol*** *is a member of technical staff in the Network Systems Research Department. Both are at AT&T Bell Laboratories in Holmdel, New Jersey, and are involved with basic and applied research on local communications networks.* ***Michael G. Hluchyj*** *is a former member of technical staff in the Network Systems Research Department. Mr. Acampora joined the company in 1968 and has a B.S.E.E., M.S.E.E., and Ph.D. in electrical engineering from the Polytechnic Institute of Brooklyn. Mr. Karol joined the company in 1985 and has a B.S.E.E. and a B.S. in mathematics from Case Western Reserve University, and an M.S.E.E., M.A., and Ph.D. from Princeton University.*

For multiuser communications networks, lightwave technology offers the potential to supply a pool of capacity—to be shared among users—far above that provided by any alternative technology. However, the bandwidth limitation of the electro-optic converters needed to attach each user to the optical medium prevents any one user from accessing more than a tiny fraction of the overall capacity. This paper discusses the problems with conventional approaches for tapping the capacity contained within the optical communications band. It then proposes a new network architecture that permits tapping lightwave's vast capacity potential without requiring a technological breakthrough. With this approach, it becomes possible to create networks that offer hundreds of thousands of gigabits-per-second total capacity, to be shared among users, each limited to a peak rate of 1 Gb/s.

Background

In long-distance point-to-point communications, lightwave has clearly emerged as the technology of choice.[1,2] For this application, the figure of merit for comparing alternative approaches and technologies traditionally has been the product of the achievable bit rate and the unrepeatered distance. It is in this regard that lightwave technology excels. What finally limits the figure of merit is the loss and dispersion in the optical transmission medium, with the following combining to increase the achievable figure-of-merit:

- Improvements in fiber fabrication processes
- The transitions from multimode to single-mode fiber and from multimode to single-mode lasers
- Operation near the zero-dispersion wavelength of 1.3 μm
- Availability of narrow-linewidth lasers
- The introduction of coherent demodulation in receivers with performance that approaches the quantum limit.

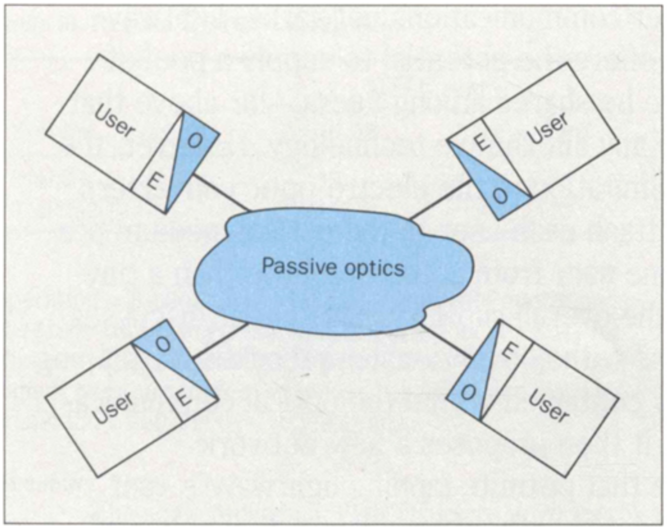

22

Figure 1. Ideal lightwave network. E/O is an electro-optic interface.

These impressive accomplishments have raised the obvious follow-up question: Are the advantages of lightwave technology limited in utility to long-distance point-to-point communications? In particular, what opportunities, advantages, and problems arise, if any, when lightwave technology is applied to the domain of multiuser communication networks? Perhaps the most dramatic appeal of lightwave technology stems from the opportunity to share an unprecedented amount of capacity among all network users.

If we limit the radius of the network to 20 km, then dispersion and transmission loss become much less of an issue, and the network could still interconnect users located within a good-sized city. Conservatively, if we estimate that the optical band lies in the 200-nm window centered at the zero-dispersion wavelength of 1.3 μm, we find that the medium possesses the enormous bandwidth of 35 terahertz (THz). In contrast, a "broadband" coaxial-cable-based system has a bandwidth of about 300 MHz (five orders of magnitude lower than fiber), and ordinary twisted-pair copper wiring has a bandwidth of only several megahertz (seven orders of magnitude lower). The appeal of lightwave is obvious.

But although the optical medium possesses enormous bandwidth, any one user may tap into—at most—a bandwidth of "only" several gigahertz because each user is limited by the bandwidth constraint that electro-optic access establishes. In a fundamental sense, the bandwidth associated with modulation of photons by electrons and electronic signals generated by photodetection will always be orders of magnitude lower than the medium's inherent bandwidth. Thus, perhaps the most basic issue to arise in multiuser lightwave networks is: How can we tap the optical medium's vast bandwidth potential and effectively share it among many users, each constrained by the far lower bandwidth of electro-optic conversion?

Figure 1, which shows a simplified version of an ideal lightwave network, depicts this situation. The network consists of three parts:

- The *passive optical media.* This is the "glass" portion of the network that contains all the bandwidth.
- The *generic "users"* themselves, electronic in nature. Each user may be a PBX (private branch exchange) trunk, LAN (local-area network) gateway, host computer, graphics workstation, high-resolution video camera or monitor, fast document scanner, high-speed printer, digital telephone, etc.
- The *electro-optic conversion* that changes electronic signals to optical signals, and vice versa.

The peak data rate that any user generates or receives can be no greater than that allowed by the electro-optic bandwidth constraint. (We envision a user interface in which traffic is converted to a packet format. User data—generated in bursts or streams—is broken into packets, and each packet's transmission rate is the maximum that the electro-optic bandwidth constraint permits.) We see that each user can access (transmit or receive) only a tiny portion of the optical bandwidth.

To tap the vast bandwidth potential, it is essential

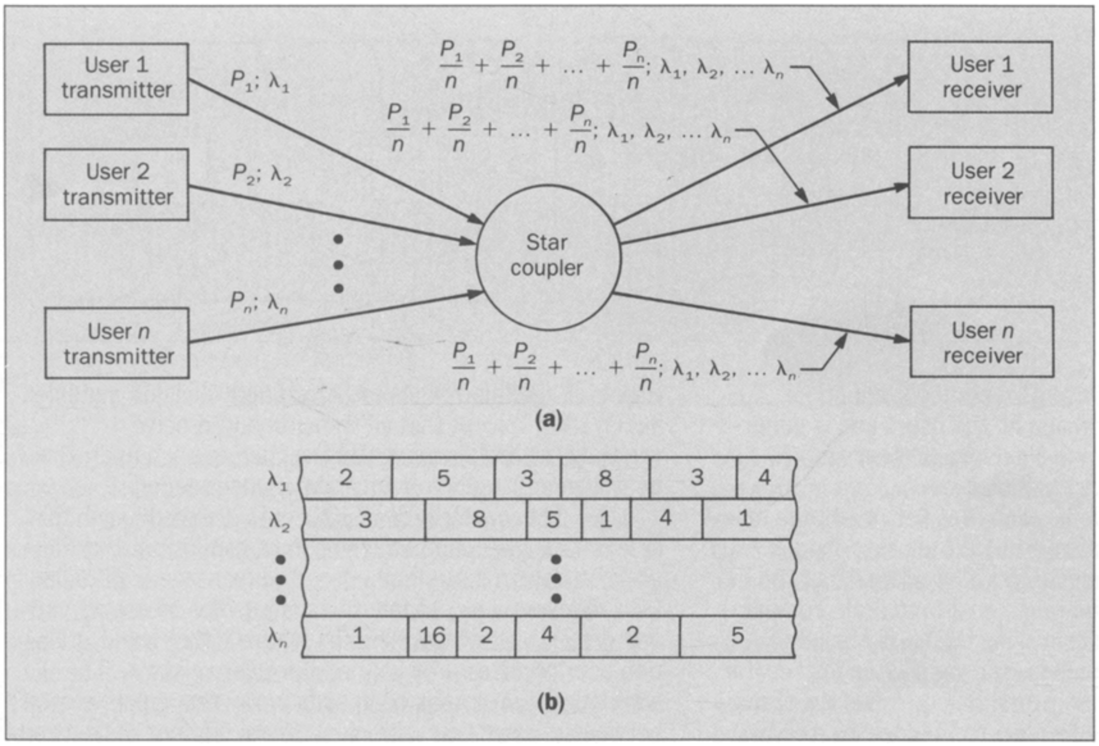

Figure 2. Conventional approach to achieving concurrency. (a) A wavelength-division multiplexed system; (b) interconnection pattern. *P* = power; λ = wavelength.

that a multitude of messages among the various users concurrently reside on the optical medium. Thus, the great challenge is to devise a networking arrangement that achieves a high level of concurrency. Simply put, we want to build a network that offers a total capacity, or throughput, measured in terabits per second from electronic components that operate in the gigabit-per-second range.

At first glance, it might appear that concurrency is easy to achieve with a straightforward, conventional approach. Simply assign to each user a unique wavelength on which each packet may be transmitted in the network. Figure 2a illustrates the approach.

Here, user j has been assigned to transmit on wavelength λ_j, $j = 1, 2, \dots, n$. All signals are linearly combined in a centrally located passive-star coupler, and the superimposed signals are made available to all receivers. Notice that splitting each signal n ways at a star coupler causes a power loss of $1/n$ on each signal toward each receiver. The star topology best minimizes the loss for this broadcast arrangement, while providing acceptable detectability.

Each receiver, then, simply tunes to the wavelength of interest, much like a listener tunes a radio receiver to pick up the desired station from the sea of stations being broadcasted. Figure 2b illustrates a typical interconnection pattern among the n users, showing the time sequence of the channels to which each receiver must tune. Two problems are apparent:

- Each receiver must be instructed about which channel must be tuned for each interval to receive the appropri-

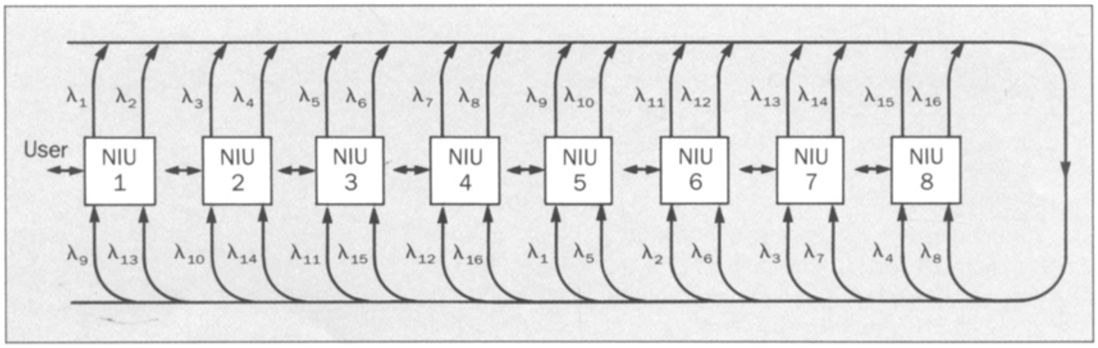

Figure 3. Multihop lightwave network. NIU = network interface unit; λ = wavelength.

ate messages. Such pretransmission coordination requires communication among the users and is generally done using a separate signaling or reservation channel that is common to all users.

- By far the larger problem, each receiver must tune rapidly over the entire optical band in the prescribed interconnection sequence. Optical receivers that can be tuned over such a broad range on a time scale comparable with the packet transmission rate (i.e., tune in several tens of nanoseconds, and then listen for several microseconds) are well beyond the current state of the art and may require a profound technological breakthrough.

Thus, we have a dilemma: Lightwave networks appear to offer unique possibilities by virtue of the vast bandwidth potential, but suffer unique constraints by virtue of limitations of the technology. The challenge, then, is to devise a new architectural approach to resolve this dilemma.

The multihop approach to lightwave networking, described in this paper, resolves this dilemma without requiring a technological breakthrough. In fact, after illustrating the approach with an embodiment that exploits today's technology, we conclude with an embodiment that can be realized with yesterday's technology.

The multihop approach is based on the observation that optical bandwidth is so plentiful that—for the first time—it is acceptable to "squander" some bandwidth if this permits tapping the rest. Like the approach of Figure 2, multihop relies on wavelength-division multiplexing (WDM), except that all transmit and receive wavelengths are permanently assigned to each user; that is, wavelength agility or tunability is not needed.

The wavelength assignments are made such that at least one user can receive each transmission. But that user, on determining that a received message is intended for a different user, simply retransmits the message onto one of its transmit wavelengths where either the destination user receives it or yet another user relays it. The wavelengths are assigned in such a way that a path exists between all user pairs and, for a given number of transmitters and receivers per user, the average number of hops a packet must take to reach its destination is close to the minimum. It is worth noting that multihop networks are common in packet communications systems—such as packet radio networks and local, metropolitan, and wide-area networks.[3-8]

Typical capacity or throughput results are as follows. Let the network contain about 900 users that each can transmit on two wavelengths and receive on two wavelengths. Let each transmitter and receiver operate at a peak rate of 1 Gb/s. Then, after accounting for the effects of multiple hops, the network will provide a total capacity of 200 Gb/s to be shared among all users. Moreover, each user enjoys a peak transmission rate of 1 Gb/s and can present a sustained, effective average-traffic load of about 220 Mb/s. The throughput of a multihop lightwave network scales linearly with the user transmission rate, so each

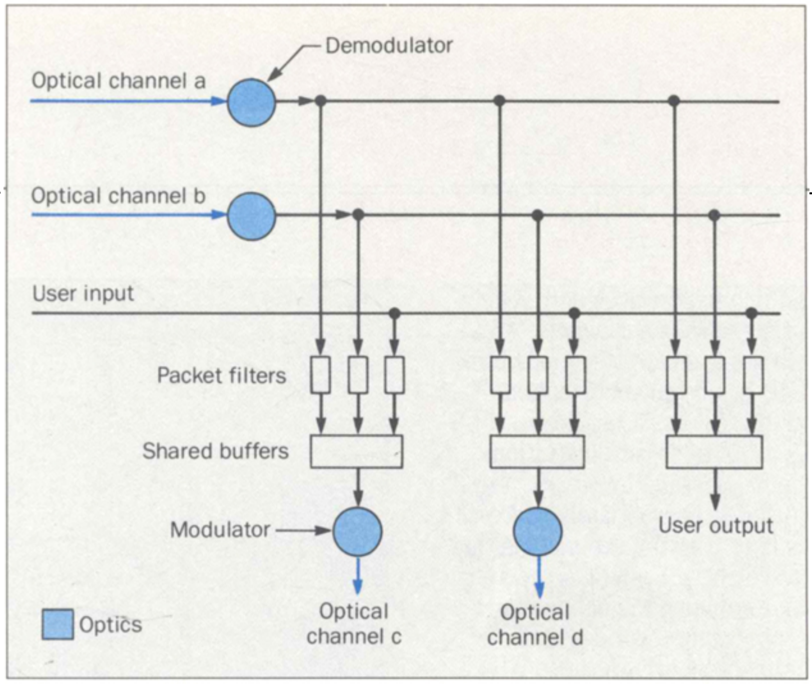

Figure 4. Block diagram of two-receiver, two-transmitter network interface unit.

user can present an average load of about 22 Mb/s if the transmission rate is 100 Mb/s.

Each user needs equipment similar to that for a simple network in which all users share a single 1-Gb/s channel. However, for such a single-channel network, a total capacity of only 1 Gb/s is to be shared among all users, and each of 900 users connected to the network could present an average traffic load of only 1.1 Mb/s. Thus, by doing little more than permanently assigning different users to different wavelengths, both network capacity and sustainable throughput per user are tremendously increased. Moreover, as more users attach to the network, each brings its own channels, or resources, increasing the total capacity.

We will illustrate the multihop concept first for the case where each network port is equipped with two fixed-wavelength transmitters and two fixed-wavelength receivers. This then will be generalized to an arbitrary number p of transmitters and receivers. The benefits of such a generalization include:

1. The ability to concentrate the traffic of several users through each port, which increases the sustainable throughput per user
2. The ability to grow the network modularly without rearranging existing channel assignments
3. Improved reliability.

Finally, we shall show how to implement the multihop concept without wavelength-division multiplexing, thereby permitting the use of yesterday's optical technology. The same approach, coincidentally, also provides a way to realize a large, centrally located space-division packet switch as well as a distributed-lightwave network.

The Multihop Approach

In the multihop approach to lightwave networks, wavelength-division multiplexing is the key to tapping the bandwidth potential of the optical medium, allowing simultaneous transmissions among multiple, electronic speed-limited users. But to avoid the problems of pretransmission coordination and wavelength agility common in conventional approaches, a multihop network uses only fixed-wavelength transmitters and receivers. Transmitting a packet from one user to another may require routing the packet through intermediate users, each repeating the packet on a new wavelength, until the packet is finally transmitted on a wavelength that the destination user receives. In other words, a packet may have to take multiple hops to reach its destination.

Although the physical topology for a multihop network can take a variety of forms (e.g., star, bus, or tree), for simplicity, Figure 3 illustrates an eight-user multihop network in a bus topology. The users interface to the unidirectional optical bus through network interface units (NIUs) distributed along the bus. Each NIU has two fixed-wavelength transmitters and two fixed-wavelength receivers. The example requires sixteen WDM channels. All information transmitted onto the bus passively loops back past every NIU, but each NIU hears only information transmitted on channels for which it has a receiver.

The scheme works as follows. For illustration, suppose user 1 wants to send a message to user 5. NIU 1 receives the message from user 1, appropriately addressed with a header that identifies user 5 as the destination, and transmits the message on wavelength λ_1. NIU 5 receives the message, recognizes (by examining the destination address) that it is intended for "local reception," and passes it to user 5 through the user port. Because NIU 5 is (permanently) assigned to receive on one of the two channels on which NIU 1 is assigned to transmit, namely λ_1, messages pass from NIU 1 to NIU 5 with just a single hop through the network.

But if user 1 wants to send a message to user 2, we recognize that NIU 2 does not receive any of the transmit channels assigned to NIU 1. However, if we use NIU 5 as an intermediate-relay node, we can establish a two-hop path between NIU 1 and NIU 2. Specifically, when NIU 1 receives a message from user 1 that is destined for user 2, it transmits the message on wavelength λ_1. When NIU 5 receives the message, it recognizes that the message is not intended for local reception and, with user 2 as the destination address, retransmits the message over wavelength λ_{10}. NIU 2 is assigned to receive wavelength λ_{10}, so it receives the message, and then passes it on to user 2 through the user port.

As a final example, consider messages that user 1 sends to user 7. These messages will require three passes through the network, using first NIU 5 and then NIU 2 as intermediate-relay nodes. NIU 1 transmits to NIU 5 using wavelength λ_1, which demodulates the message and remodulates it onto λ_{10}. NIU 2 receives the message over λ_{10} and relays it to the destination NIU 7 using wavelength λ_3.

Channels are permanently assigned to users such that after one or more passes through the network, a message is assured reception by the proper user's receiver. In the example just presented, three hops is the maximum

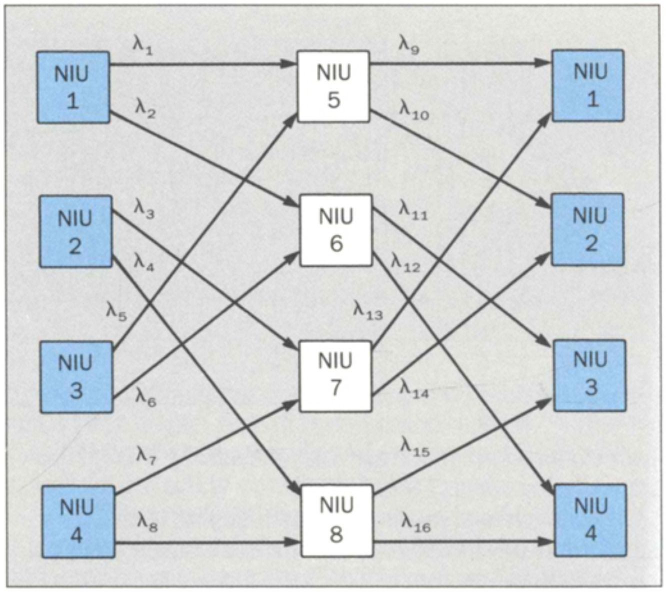

Figure 5. Connectivity of network interface units (NIUs), for the example in Figure 3. λ **= wavelength.**

required of any source-destination pair.

With a single user attached to each NIU, Figure 4 shows a functional block diagram of a two-transmitter, two-receiver NIU. The NIU functions as a three-by-three nonblocking packet switch: two inputs and two outputs for the optical channels, and an electronic input and output for the user. An additional input and output can be included for control information.

The incoming optical signals for the two assigned receive wavelengths are demodulated to baseband electrical signals that, together with the generated user data, enter the three-by-three nonblocking packet switch on separate inputs. The destination address in the header of each arriving packet determines to which of the three outputs the packet is to be routed. Buffers are provided at the outputs, because multiple packets may arrive on different inputs destined for the same output. Two outputs drive optical modulators for the two

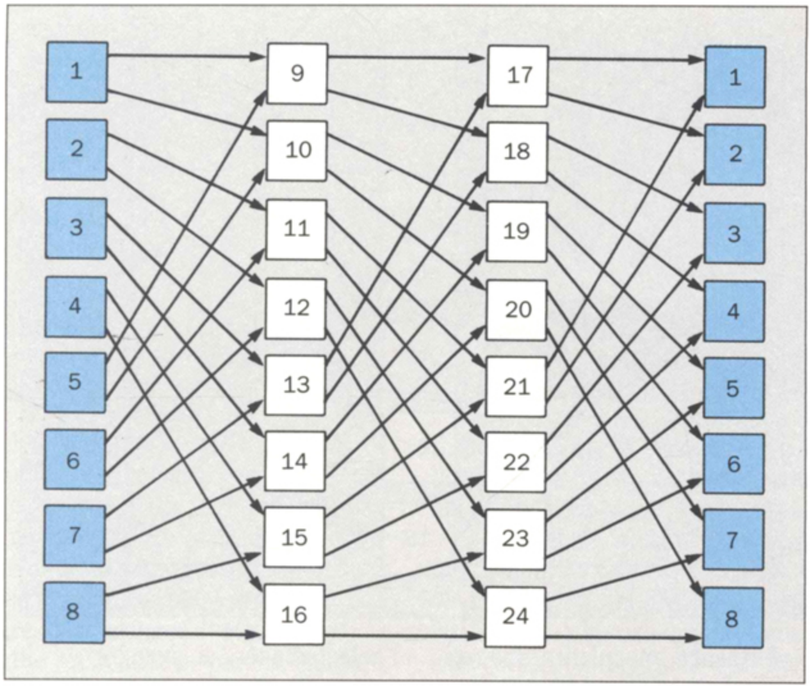

Figure 6. Multihop connectivity of 24 network interface units.

assigned transmit wavelengths at the NIU. The third output goes directly to the user.

Figure 5 shows the connectivity graph associated with the example of Figure 3. The connectivity graph indicates all (ordered) pairs of NIUs that are separated by a single hop. Specifically, for each NIU in the network, we draw a directed arc from an NIU that transmits on a particular wavelength to all NIUs that receive the wavelength. In this connectivity graph (Figure 5), the eight NIUs are placed in two columns of four NIUs each. The assignment of transmit and receive wavelengths is such that the interconnection pattern between columns is a perfect shuffle,[9] with the right side of the second column connected to the left side of the first column as if the entire graph were wrapped around a cylinder.

Figure 6 shows the connectivity graph for a 24-NIU multihop network. Here, using 48 WDM channels, 24 NIUs are arranged in three columns of eight NIUs each. In general, for two transmitters and two receivers per NIU, $n = k2^k$ NIUs ($k = 2, 3, \ldots$) are arranged in k columns of 2^k NIUs each, making use of $W = k2^{k+1}$ WDM channels.[10]

The recirculating perfect-shuffle connectivity graph can be further generalized for the case of p transmitters and p receivers ($p = 1, 2, \ldots$) per NIU.[11] Figure 7 shows an 18-NIU connectivity graph with $p = 3$ and $k = 2$. The resulting pattern of arcs between adjacent columns is called a p-shuffle,[12] being a generalization of the ($p = 2$) perfect shuffle.[9] A (p,k) recirculating perfect-shuffle connectivity graph has $n = kp^k$ NIUs ($k = 2, 3, \ldots; p = 1, 2, \ldots$) in k columns of p^k NIUs each, using $W = kp^{k+1}$ WDM channels.

Although the transmit and receive wavelengths for the NIUs in a multihop network could be assigned using other connectivity graphs, the class of recirculating perfect-shuffle graphs has several desirable features

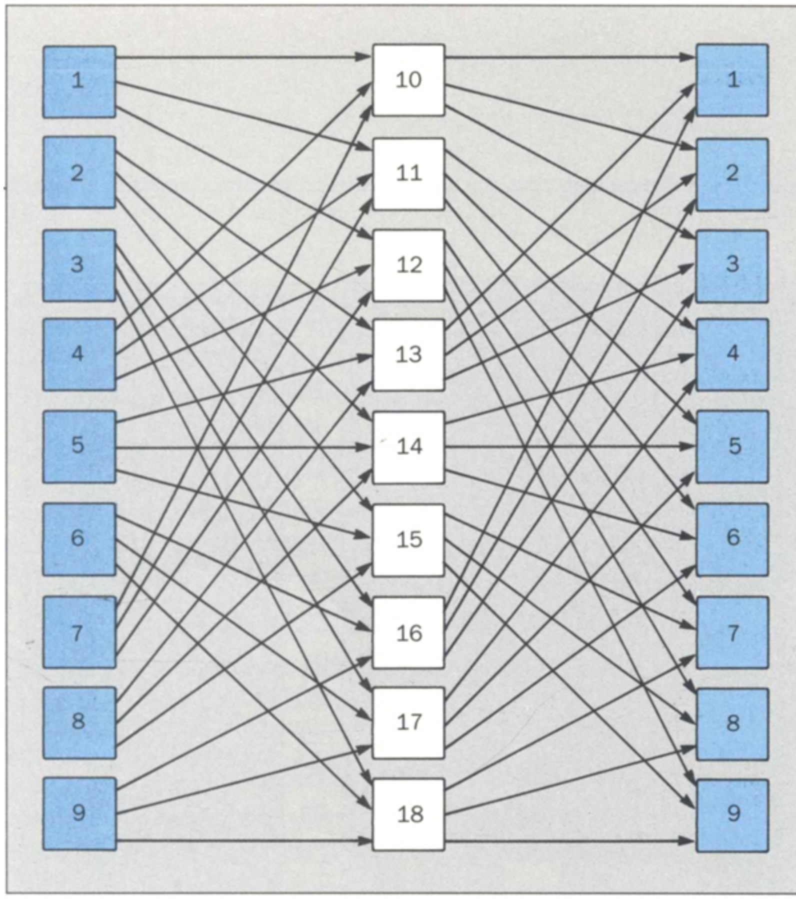

Figure 7. Multihop connectivity of 18 network interface units (p = 3, k = 2).

related to the network's performance, modularity, and reliability. We will now focus on performance and then cover modularity and reliability in the next section.

Performance. For a recirculating perfect-shuffle connectivity graph, a path that has at most $2k - 1$ hops exists from each NIU to all others. Figure 8 illustrates this for the ($p = 2$, $k = 3$) 24-NIU connectivity graph, where we show a spanning tree for routing packets from NIU 1 to all other NIUs. Because of the symmetric nature of the graph, we can draw a topologically equivalent spanning tree for all other NIUs. Notice from Figure 8 that the depth of the spanning tree grows as $2k - 1$, which is about proportional to the base p logarithm of the number of NIUs. More specifically, we have

$$\text{Number of NIUs } h \text{ hops from source} = \begin{cases} p^h & h = 1, 2, \ldots, k-1 \\ p^k - p^{h-k} & h = k, k+1, \ldots, 2k-1 \end{cases}$$

so that the expected number of hops between two randomly selected users is given by

$$E[\textit{number of hops}] = \frac{kp^k(p-1)(3k-1) - 2k(p^k - 1)}{2(p-1)(kp^k - 1)} \qquad (1)$$

If the network is uniformly loaded (i.e., each user generates an equal amount of traffic destined for all other users) and all WDM channels are loaded equally, the channel efficiency η for the network is given by

$$\eta = \frac{1}{E[\textit{number of hops}]} = \frac{2(p-1)(kp^k - 1)}{kp^k(p-1)(3k-1) - 2k(p^k - 1)} \qquad (2)$$

Figure 9 shows the channel efficiency η as a function of the connectivity graph degree p for various numbers

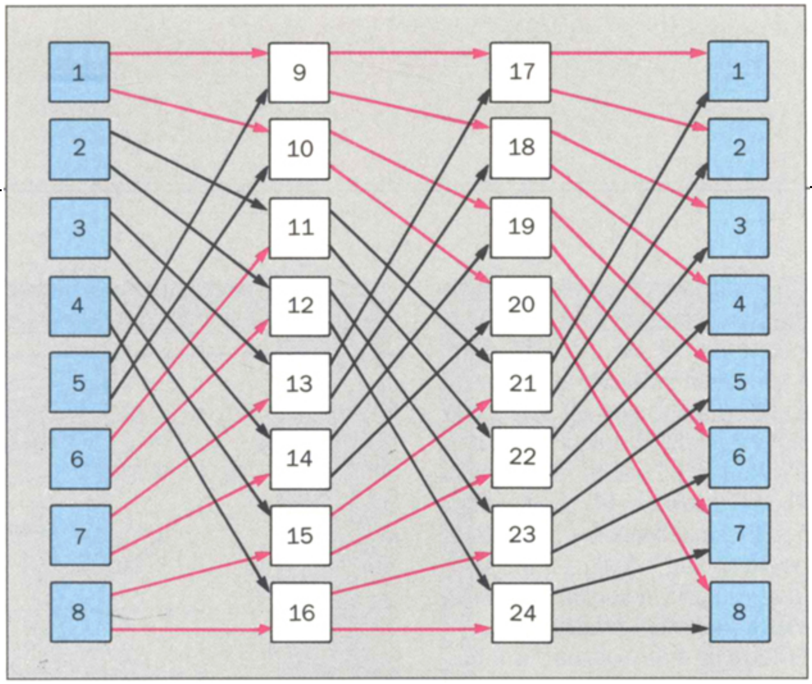

Figure 8. Spanning tree (paths in color) for multihop connectivity graph of 24 network interface units ($p = 2, k = 3$).

of columns k. Also plotted in Figure 9, for each k, is an upper bound on the channel efficiency for any regular directed graph of degree p. To obtain the upper bound, we assume that an ideal (but not necessarily realizable) connectivity graph exists with an associated routing such that each NIU has a perfect p-ary spanning tree for routing packets and all channels are load balanced.[11] Notice how close the performance of the recirculating, perfect-shuffle connectivity graph is to that of the ideal graph. For large p, Equation (2) becomes

$$\eta = \frac{2}{3k - 1} \quad \text{for } p = \infty \qquad (3)$$

which shows clearly how the channel efficiency decreases as the inverse of the number of columns k.

Our main concern, however, is not with the channel efficiency, but with achieving concurrency in the network as measured by the network throughput. The total network throughput and the throughput per NIU are simply $\eta W = \eta k p^{k+1}$ and ηp, respectively, normalized to the channel transmission rate. Table I shows, for various p and k, some of the possible network configurations in terms of the number of NIUs and total network throughput with all kp^{k+1} WDM channels operating at 1 Gb/s. (All throughput values given in this paper scale linearly with the user transmission rate.) Observe from Table I that total throughputs in the hundreds and even thousands of gigabits per second are achievable even though each user is limited to 1 Gb/s.

Figure 10 shows the throughput per NIU as a function of the number of NIUs in the network for various values of p, the number of transmitters and receivers per NIU. Each curve is parameterized in k, with the top left side corresponding to $k = 2$. For comparison, we also show the throughput per user of a single-channel network that operates at 100-percent media-access efficiency. The throughput per user for the single-channel network decreases as $1/n$ versus the nearly $1/\log_p n$ decrease for the multihop network. In particular, with 896 NIUs, the throughput per user of a single-channel network is 1.1 Mb/s versus 222 Mb/s (i.e., a factor of 200 greater) for the multihop network.

Propagation delay is related to the geographical extent of the network, but multiuser lightwave networks are generally characterized by a large ratio of propagation delay to packet transmission time. For example, with a transmission rate of 1 Gb/s over a 10-km length of fiber, the propagation delay is 50 times the transmission time of a 1000-bit packet (assuming the speed of light in fiber is $2 \cdot 10^8$ m/s). Reference 10 presents an approximate expression for the mean queueing delay and buffer overflow probability.

Although the focus of this paper is on the dedicated-

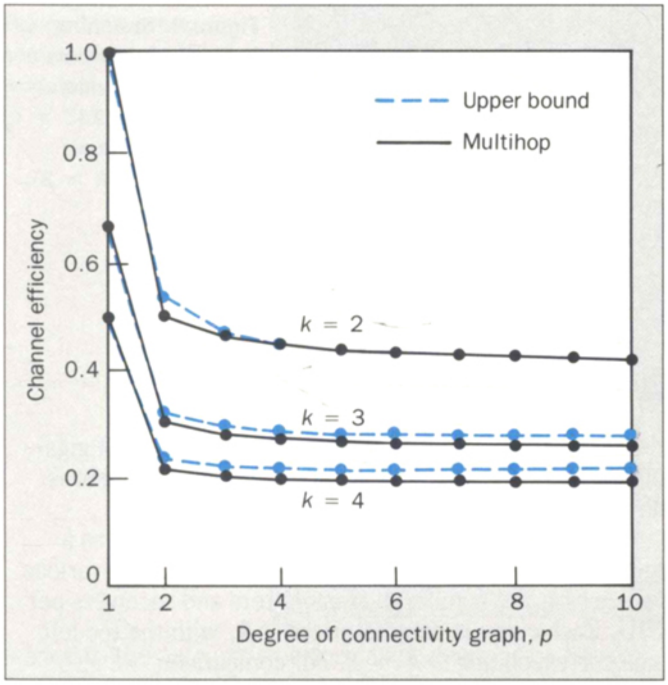

Figure 9. Multihop channel efficiency. *k* = number of columns in connectivity graph.

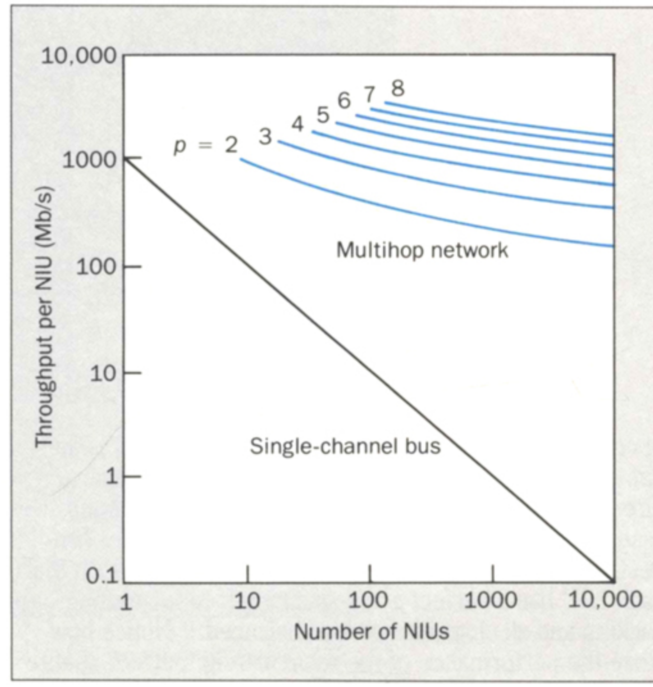

Figure 10. Multihop throughput per network interface unit (NIU) for values of *p* (number of transmitters and receivers per NIU).

channel multihop network (i.e., each arc in the connectivity graph represents a separate WDM channel), it is worth noting that the same connectivity graphs apply to the use of shared channels among the NIUs. Figure 11 shows the same connectivity graph as Figure 5, but here only four WDM channels are shared among the eight NIUs. Each NIU has only one fixed-wavelength transmitter and one fixed-wavelength receiver, with NIU pairs (1,3), (2,4), (5,7), and (6,8) each sharing a common channel. In general, with a connectivity graph of degree p, each NIU still requires only one transmitter and one receiver, but now p NIUs share each of the kp^{k-1} channels. Because multiple NIUs are transmitting on a common channel, we have to contend with a multiple access problem and the possible inefficiencies that result. In the next section, we show how—with concentration—one can have the equivalent of one transmitter and one receiver per user with dedicated channels.

Concentration, Modularity, and Reliability

A characteristic of fully distributed networks, such as found today in many local-area networks, is network interface circuitry that is located close to each user. Although this may minimize network interface equipment, one must also consider issues such as media-access efficiency, network capacity, security, maintenance, and administration. Networks with desirable characteristics along these lines typically have more of a centralized or, at least, hierarchical structure (see, for example, reference 4).

Figure 12 illustrates how a multihop network

Table I. Multihop Total Throughput, 1-Gb/s User Transmission Rate

(k,p)	Number of NIUs	Total throughput (Gb/s)
(2,2)	8	8.000
(2,3)	18	24.81
(3,2)	24	14.72
(2,4)	32	56.69
(2,5)	50	108.4
(4,2)	64	27.62
(2,6)	72	184.8
(2,7)	98	290.6
(3,3)	81	68.21
(5,2)	160	52.73
(3,4)	192	208.1
(4,3)	324	193.6
(3,5)	375	498.4
(6,2)	384	101.9
(3,6)	648	1,021
(7,2)	896	198.7
(4,4)	1,024	791.8
(3,7)	1,029	1,877
(5,3)	1,215	560.1
(4,5)	2,500	2,380
(6,3)	4,374	1,640
(5,4)	5,120	3,071
(4,6)	5,184	5,867
(4,7)	9,604	12,604
(7,3)	15,309	4,834
(5,5)	15,625	11,573
(6,4)	24,576	12,037
(5,6)	38,880	34,305

NOTE: k = number of columns in connectivity graph; p = number of transmitters (receivers) per NIU; NIU = network interface unit.

might be physically configured to support a large number of users with a smaller number of network interface units. Each NIU attaches to users on one side via separate access lines (possibly fiber) and to the other NIUs on the other side via dedicated WDM channels that are combined at a star coupler. Each user receives only that information destined for it; the NIUs forward packets that require multiple hops in the network. The NIUs can be geographically distributed—housed in satellite closets in a building or "remote terminals" in a city—or centrally located in a common equipment frame, such as in a central office.

Besides improved network security, monitoring, and maintenance, this approach can achieve the equivalent of a single transmitter and single receiver per user, while increasing the network capacity and eliminating the media access associated with shared channels. Specifically, suppose each NIU serves p users and contains p transmitters and p receivers (i.e., one pair per user) to form a p-ary multihop network with the other $kp^k - 1$ NIUs in the network. (The degree of concentration can be selected to yield a desired "equivalent number" of transmitters and receivers per user.) The throughput per user is then simply the channel efficiency η, as given in Equation (2), times the channel transmission rate.

For a 1-Gb/s channel transmission rate, Figure 13 shows the throughput per user as a function of the number of users for various NIU sizes. For comparison, we see that the throughput per user of a single-channel network—if we assume 100-percent media-access efficiency—decreases linearly with the number of users. For multihop, the throughput per user decreases (nearly) proportional to the logarithm of the number of users. Figure 13 also illustrates the performance advantages associated with increasing p (i.e., the NIU size) to grow the network while keeping k small.

Within the context of the concentrator, we also have a convenient way of modularly growing a multihop network. Figure 14 shows that the connectivity graph for a ($p = 2$, $k = 2$) 8-NIU multihop network is a subgraph of the ($p = 3$, $k = 2$) 18-NIU connectivity graph. This holds true for any given k, as p is increased. Hence, one

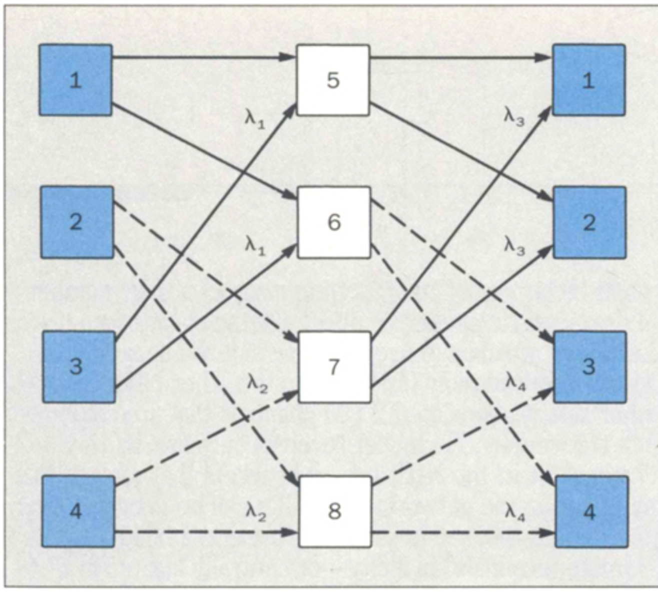

Figure 11. Multihop connectivity graph for eight users with four shared channels (p = 2, k = 2). λ = wavelength.

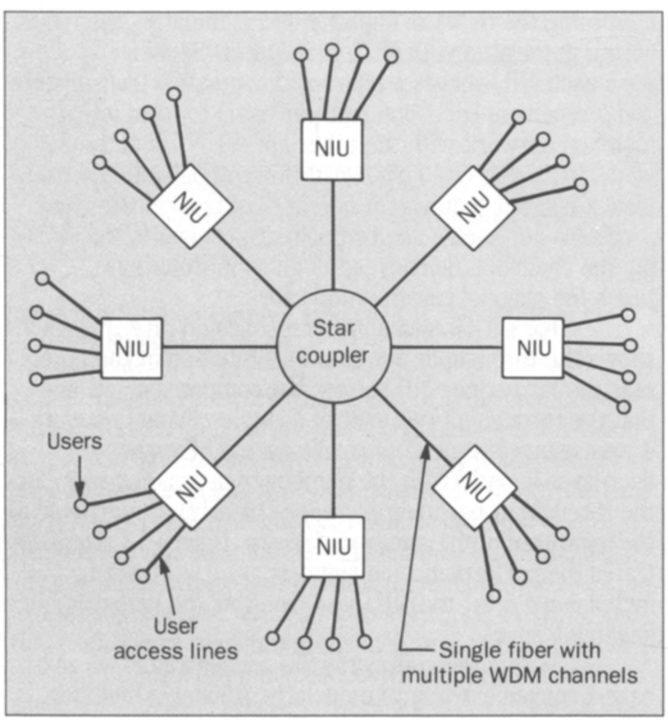

Figure 12. A multihop physical topology with network interface units (NIUs) as concentrators. WDM = wavelength-division multiplexing.

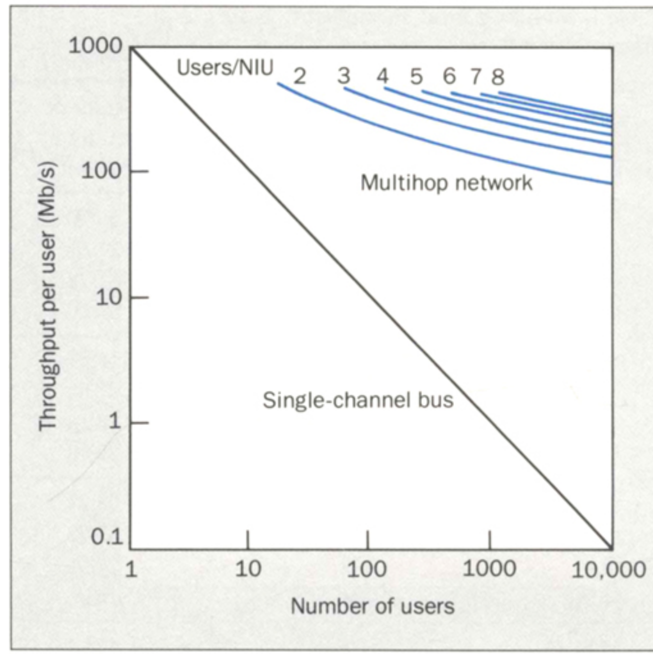

Figure 13. User throughput with concentrators; one transmitter and one receiver per user. NIU = network interface unit.

can grow from small to large multihop networks by adding new NIUs and adding (or enabling) additional transmitters and receivers to existing NIUs. As long as the number of columns k in the recirculating perfect-shuffle connectivity graph remains fixed, growing the network by increasing p requires no changes in the assignment of transmit and receive wavelengths to existing NIUs. To ensure that a path exists between all NIU pairs, while growing a multihop network from one value of p to the next, requires adding at most k new NIUs to the network (one per column in the connectivity graph) at any one time. Thus, the growth in the network can take place in small increments.

Finally, the cylindrical nature of the connectivity graph permits us a degree of fault tolerance. If, for example, NIU 5 in Figure 5 were to fail, the normal path from

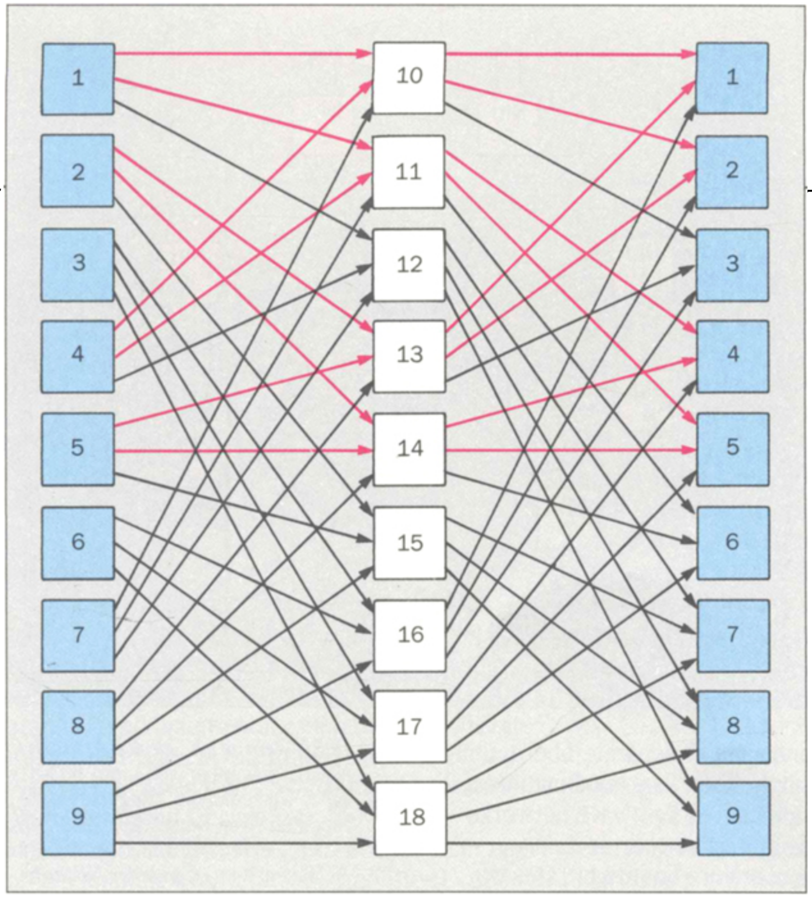

Figure 14. Modular growth of multihop networks. The paths in color represent a connectivity graph for a multihop network with eight network interface units.

NIU 1 to NIU 2 would be disrupted. However, the packet could still be routed to NIU 2 via the path 1-6-4-7-2. When NIU 5 fails, users attached to NIU 5 are cut off from the network, but users attached to all other NIUs feel the effects of the failure only through (perhaps) longer source-destination routing paths. At least p NIUs must fail, and in certain combinations, before users attached to nonfailed NIUs can be cut off from the network. For example, in Figure 5, if NIU 5 and NIU 6 both fail, NIU 1 and NIU 3 are cut off.

Alternative routing paths may also be useful even when NIUs have not failed. The existence of multiple paths (often of the same length) from a source to a destination provides the opportunity to route packets around NIUs that are temporarily congested with a large backlog of packets.

A Non-WDM Approach to Multihop Networks

Thus far, we have considered multihop networks for which the channels required by the connectivity graph were created by multiplexing different carriers onto a common fiber. However, as Figure 15 illustrates, the "WDM channels" referred to could instead be separate fibers, with a fiber-patch panel replacing the star coupler at the center (see, for example, Figure 12). The patch panel is used to make the appropriate interconnections among the incoming and outgoing NIU fibers according to the recirculating, perfect-shuffle connectivity graph. Because all connections in the network are strictly point-to-point, we can make use of yesterday's technology in building a high-capacity lightwave network.

Conclusions

We have proposed a multihop approach for achieving concurrency in distributed-lightwave networks. Multihop networks avoid two serious drawbacks of standard multichannel approaches:

- The requirement of wavelength-agile transmitters or receivers
- Pretransmission coordination between two users that want to communicate.

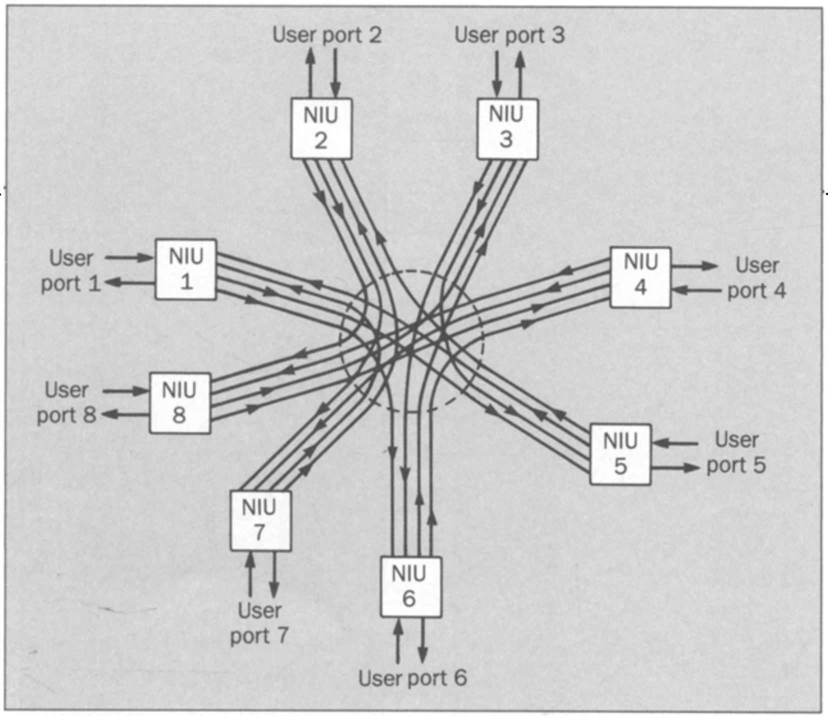

Figure 15. An approach to multihop networks without using wavelength-division multiplexing. The circled area (center) is a fiber-patch panel. NIU = network interface unit.

By creating dedicated channels out of separate fibers, multiple WDM channels on a single fiber, or a combination of both, one can build highly concurrent lightwave networks with hundreds or thousands of gigabits-per-second throughput, even while the users are limited to rates of 1 Gb/s or lower. Although transmitting a packet from one user to another may require routing the packet through intermediate NIUs, the network connectivity is specifically designed to:

- Achieve efficient use of the channel bandwidth.
- Allow modular growth of the network from small to large configurations.
- Provide a degree of network reliability.

References

1. P. S. Henry, "Introduction to Lightwave Transmission," *IEEE Communications Magazine*, Vol. 23, No. 5, May 1985, pp. 12-16.
2. T. Li, "Advances in Optical Fiber Communications: An Historical Perspective," *IEEE Journal on Selected Areas in Communications*, Vol. SAC-1, April 1983, pp. 356-372.
3. R. E. Kahn et al., "Advances in Packet Radio Networks," *Proceedings of the IEEE*, Vol. 66, No. 11, November 1978, pp. 1468-1496.
4. A. S. Acampora and M. G. Hluchyj, "A New Local Area Network Architecture Using a Centralizd Bus," *IEEE Communications Magazine*, Vol. 22, No. 8, August 1984, pp. 12-21.
5. A. G. Fraser, "Toward a Universal Data Transport System," *IEEE Journal on Selected Areas in Communications*, Vol. SAC-1, No. 5, November 1983, pp. 803-816.
6. N. F. Maxemchuk, "Regular Mesh Topologies in Local and Metropolitan Area Networks," *AT&T Technical Journal*, Vol. 64, No. 7, September 1985, pp. 1659-1685.
7. A. Tanenbaum, *Computer Networks*, Prentice-Hall, Englewood Cliffs, N. J., 1981.
8. I. Chlamtac and A. Ganz, "Toward Alternative High Speed Networks: The Swift Architecture," *INFOCOM '87*, San Francisco, California, March 1987.
9. H. S. Stone, "Parallel Processing with the Perfect Shuffle," *IEEE Transactions on Computers*, Vol. C-20, No. 2, February 1971, pp. 153-161.
10. A. S. Acampora, "A Multichannel Multihop Local Lightwave Network," *GLOBECOM '87*, Tokyo, Japan, November 1987.
11. M. G. Hluchyj and M. J. Karol, "ShuffleNet: An Application of Generalized Perfect Shuffles to Multihop Lightwave Networks," to be presented at *INFOCOM '88*, New Orleans, Louisiana, March 1988.
12. J. H. Patel, "Performance of Processor-Memory Interconnections for Multiprocessors," *IEEE Transactions on Computers*, Vol. C-30, No. 10, October 1981, pp. 771-780.

(Manuscript received June 4, 1987)