PHOTONIC MULTIPLE-ACCESS NETWORKS: TOPOLOGIES

Ivan P. Kaminow

Ivan P. Kaminow *is head of the Photonic Networks and Components Research Department at AT&T Bell Laboratories (Crawford Hill) in Holmdel, New Jersey. He is responsible for research on high-speed optical devices and networks for computer communication and local distribution. Mr. Kaminow joined the company in 1954 and has a B.S.E.E. from Union College, Schenectady, New York; an M.S.E. from the University of California at Los Angeles; and both an A.M. and Ph.D. in applied physics from Harvard University.*

Lightwave technology is well-established for long-haul and point-to-point applications. Current research is aimed at extending these techniques to multiple-access computer networks and to the distribution plant. We are exploring the opportunities for interconnecting many users (more than 100) at high data rates (greater than 1 Gb/s) and the need for novel optoelectronic components to realize these photonic networks of the future. In this paper, we focus on how photonic components constrain topologies, while a companion paper[1] focuses on their effect on routing and multiplexing.

Introduction

Lightwave technology has revolutionized long-haul telecommunications. Today's pervasive use of optical fiber for digital transmission is based on efficient semiconductor lasers and low-loss single-mode fiber developed in the early 1970s. The next lightwave breakthrough may come in local applications: computer networks and loop distribution. As with long-haul lightwave, a full-scale revolution must build on the development of novel photonic components today.

Local multiple-access networks pose a much more difficult and diffuse target than point-to-point links, where the simple criterion is transport at the maximum bit rate over the greatest distance at the lowest cost. Photonic access networks can interconnect N equivalent users—each at bit rate B—in a computer network, or can connect N users at bit rate B to a central office, remote terminal, or computer center. (Panel 1 defines acronyms and terms used in this paper. For a comprehensive discussion of all aspects of lightwave technology, see Miller and Kaminow.[2] This paper was derived from Chapter 26 of their book.[3])

Optical fiber's chief advantage is its bandwidth capability. So, B will be large, say 100 Mb/s (megabits per second) or 1 Gb/s (gigabits per second) or more, and N will also be large, say 100 or more, yielding aggregate network throughputs NB that approach 1 Tb/s (terabit per second = 10^{12} bits per second).

But who are the potential customers for such services? Will they transfer occasional, but large files; or will they require real-time

Panel 1. Acronyms and Terms

α	fiber and connector loss between stations
B	bit rate
BER	bit-error rate
c	velocity of light
CATV	cable television
CD	collision detection
CSMA/CD	carrier-sense multiple access with collision detection
E/O	electrical to optical
FDDI	fiber distributed data interface
GaAs	gallium arsenide
IEEE	Institute of Electrical and Electronic Engineers
L	distance
LAN	local area network
LED	light-emitting diode
n	effective refractive index
N	number of equivalent stations
NB	aggregate network throughput
O/E	optical to electrical
P	optical power
P_R	minimum receiver power that provides a tolerable bit-error rate
P_T	available optical transmitter power
P_T/P_R	power margin
T	packet length (or packet duration)
t	time
TDM	time-division multiplexed
VCS	virtual-circuit switch

video interconnects? Is circuit or packet switching the best choice? Could analog modulation provide one-way cable television (CATV) distribution more economically than digital? And, who will devise the industry-wide protocols and standards to make these networks an acceptable product? As important as these questions are, it is too early to attempt intelligent answers until we can explore what is physically realizable and with what degree of complexity.

In this paper, we examine research efforts on the "physical layer" of photonic access networks. We show how photonic components differ from their electronic counterparts, and how these differences lead to constraints on topology and access protocols. The initial objective of photonic network research is to maximize B and N at reasonable cost for a circuit-switched network. The next step will be to incorporate the switching controls needed to adapt the most attractive network concept to wideband packet switching.

At the outset, note that we can realize many of optical fiber's advantages simply by replacing point-to-point links in wire networks with optical data links. Low loss and low dispersion, low initial and installation costs, and immunity from radio frequency interference all make optical links more attractive. However, no innovation is required for this substitution.

We want to learn how to use photonics to produce higher performance, lower cost networks. One problem with data links is the high cost of the electrical-to-optical (E/O) and optical-to-electrical (O/E) conversions at each end. Although the fibers themselves have large bandwidths, the electronic terminals often introduce bandwidth bottlenecks. Electronic buffering and processing further limit high-speed performance. Thus, the trick is to keep the bits in photonic form as long as possible.

For editorial convenience, we have divided this survey into two papers. This one discusses the constraints of photonic components on network topology. The other paper[1] discusses the constraints of photonic components on multiplexing and modulation formats. Currently, a passive star seems the most attractive topology, and optical frequency-division multiplexing with tunable filter or heterodyne demultiplexing and detection are the most interesting formats.

Ring Network

Figure 1 illustrates a ring network. The real topology may be a good deal more irregular than a circle, depending on the accessibility of stations. In its usual application, which uses a token-ring protocol for media access, a repeater that operates at the aggregate network rate is required at each station.

A *token*—i.e., a "1" or a "0" bit—is propagated in one direction from station to station. When a station has a packet to send to another station, it adds the address of the receiving station in a header and holds the combined packet in a buffer. The sending station reads the tokens as they go by until it receives an empty token, a 0. It then converts the 0 to a 1, a busy token, and appends the packet.

Intermediate stations repeat the bits in the packet and also "listen" for their own addresses. If a station recognizes its address in the packet header, it copies the packet. When the packet returns to the sender, it serves as an acknowledgment, and the sender removes it from the ring, after converting the token back to 0.

Commercial token rings use wire interconnections or optical data links to join stations at rates in the 10-Mb/s range. Actual network use is less than 10 Mb/s because of the time it takes an empty token to pass around the ring. This transit-time delay increases linearly with the number of stations. It includes propagation delay between stations and processing delay at each station, which must examine the token for every packet before repeating the bits to the next station.

A token-ring architecture is not especially attractive for a high-speed optical network (where B is about 1 Gb/s) because of the cost of high-speed repeater optoelectronics at each station and the packet-processing delay. In addition, at high bit rates, the packet time may be much shorter than the propagation time around the ring, unless a packet contains an unusually large number of bits. Efficient use of the ring with short packets may call for multiple tokens, which can lead to complex protocols. Increasing the number of bits per packet increases the packet time but places added burden on the high-speed buffer.

Reliability—if one station is disabled, or if the fiber breaks—is a problem in both fiber and wire rings.

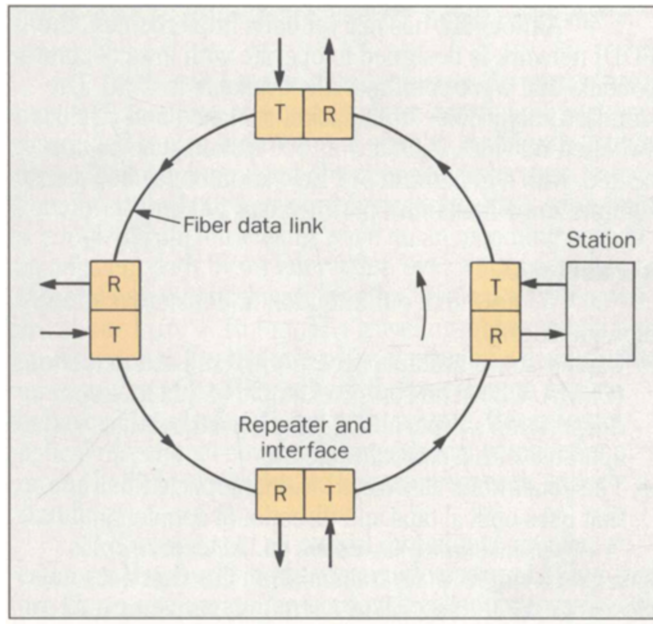

Figure 1. Unidirectional ring network. R and T represent the receiver and transmitter functions, respectively.

To address these reliability problems, a double-ring optical network can provide for bypass of defective stations and loop back around a fiber break. Each repeater has two inputs and two outputs connected to two rings that operate in opposite directions; this, of course, increases the cost.

The fiber distributed-data interface (FDDI)[4,5] is a standard proposed by the American National Standard Institute for a 100-Mb/s double-ring, time-division multiplexed (TDM), local-area network (LAN) that uses 1.3-μm (micrometer) multimode fiber and light-emitting diode (LED) data-links between stations. (A single-mode fiber, laser diode version will be standardized later.) This LAN is designed to provide both backbone services that interconnect lower speed LANs and back-end services that interconnect mainframe computers, mass storage systems, and other high-speed peripherals.

Although it has not yet been implemented, the FDDI network is designed to operate with low-cost components that were commercially available in 1986. The standard can provide both packet-switched and circuit-switched services. As many as 500 stations can be connected, with a maximum of 2 km (kilometers) between stations and a maximum perimeter of 100 km.

Bus Networks

We now turn our attention to bus network topology, specifically:

- Carrier-sense multiple access with collision detection (CSMA/CD). The classic example of this is Xerox Corporation's Ethernet® network protocol that operates with a coaxial cable bus.
- The constraints associated with an optical fiber bus that uses optical taps and directional couplers.

Ethernet Networks—CSMA/CD. A *bus network* is based on a single, main transmission line that does not close on itself. It may follow a straight, serpentine, or spiral path. Stations tap onto the bus to transmit or receive signals that are broadcast to stations on the bus.

One of the best known protocols for gaining access to the bus is the Ethernet network protocol, or CSMA/CD. In the original Ethernet network (which operated at 3 Mb/s and was later upgraded to 10 Mb/s), a coaxial cable served as the bus, and a tap consisted of a small antenna probe inserted through a hole in the outer conductor of the bus coaxial cable. This antenna is part of the transceiver; a cable physically attaches the transceiver to the coaxial cable on one side and to the station on the other side.

The antenna, which is much smaller than the 30m (meter) wavelength of a 10-MHz (megahertz) signal, radiates bidirectionally on the bus. Because the probe penetration is small, reflections on the bus are small and the shunt impedance is large. Today's taps are a bit more sophisticated but behave similarly. Even with this small coupling, station amplifiers can restore the signal strength to the sensitivity of electronic receivers [about –60 dBm (decibels above 1 milliwatt)]. Problems with reflections, attenuation, and dispersion would begin to appear at rates above 100 Mb/s.

A station with information to send collects and forms a packet that contains both data and a header with the receiver address. Then, the station senses the bus to see if another station is transmitting; if not, the station sends its packet. All stations listen as the packets go by. If a station recognizes its address, it copies the packet, which finally gets absorbed at either end of the bus.

If a station—call it A—starts to send a packet at time $t = 0$, then another station—call it B—a distance L away, will not sense A until $t = t_1 = L/(c/n)$ because of the electromagnetic propagation delay. (Here, c is the velocity of light, and n is the effective refractive index.) And, if station B starts to send its packet just before $t = t_1$, then station A will not sense B until $t = 2t_1$.

If the packet length T is greater than $2t_1$, then station A will detect the collision (CD) before completing the packet transmission. It then stops transmitting. Station B will have detected a collision at $t = t_1$ and also stopped transmitting. After a random delay, each station will again try to send its packet. After successful transmission, the packet is removed from the storage buffer.

During periods of heavy demand, many collisions and retransmissions may occur, which results in poor bus utilization. A critical constraint on the CSMA/CD network is: If all collisions are to be detected before packet transmission is completed, then the packet duration T must be greater than the maximum round-trip propagation delay, i.e.,

$$T \geq 2L/(c/n) \qquad (1)$$

where L is the bus length.

Several difficulties arise in trying to extend CSMA/CD to fiber-optic networks, including:

- Limitations imposed by optical taps restrict the number of stations (as discussed next).
- A collision with a packet from a remote station, whose

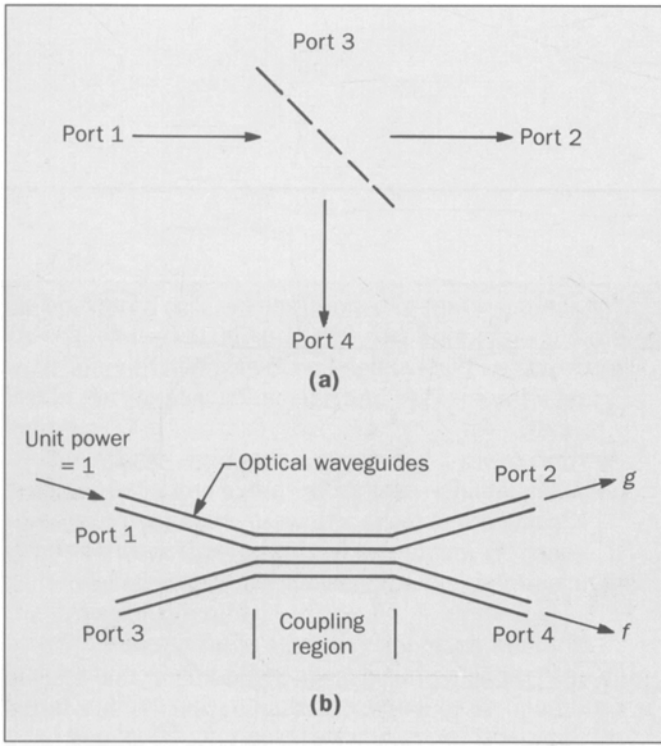

Figure 2. Optical taps: (a) beam splitter and (b) directional coupler. These are functionally equivalent four-port devices. *f* and *g* are fractions of power at the output ports for the coupled and incident guides, respectively.

signal will be weak, is difficult to sense in the presence of the strong optical signal from the local station. Later, we describe an experimental, 10-Mb/s fiber Ethernet network that uses an "active star" to sidestep the tap and collision detection problems imposed by the bus architecture.

At high speeds (B about 1 Gb/s), the requirement for long packets that contain many bits is a severe constraint, according to equation (1). As an example, consider $L = 1$ km, $n = 1.5$, and $B = 1$ Gb/s. Then, we require $T \geq 10$ μs (microseconds), and $BT = 10^4$ is the minimum number of bits per packet. This many bits per packet requires a large high-speed buffer and imposes restrictions on applications.

For example, if many bits were continuously available (as in a large file transfer or real-time video), then BT would not be restrictive. But if only occasional bursts, each containing fewer than BT bits of data, were available in time T (as in some interactive applications), then a packet defined by BT might carry many empty bits.

Optical Taps and Directional Couplers. An *optical tap* may take the form of a beam splitter (partial mirror) or, for guided waves, a directional coupler. As Figure 2 illustrates, beam splitters and directional couplers are functionally equivalent, four-port devices. In either case, light is coupled from a traveling wave on an incoming path to an adjacent path by an interaction over a coupling distance of many optical wavelengths. Typical wavelengths are about 1 μm = 10^{-6} m. The traveling-wave coupling is unidirectional. On the other hand, a tap on a coaxial transmission line at 10 MHz, where the wavelength is 30m, behaves like a bidirectional point source. Because the optical directional coupler is an essential component of most photonic networks, its behavior is worth understanding.

To construct an optical directional coupler, we can form the single-mode waveguide pattern of Figure 2b on a substrate (an approach called *integrated waveguide optics*), or fuse two single-mode optical fibers together with their cores slightly separated—which, in either case, allows the wave functions in neighboring guides to overlap. If unit power is incident on port 1, then the fraction f appears at port 4, g at port 2, and nil at port 3. Because the coupler has fourfold symmetry, power that is incident on any other port is split in the same way.

In practical couplers, a small fraction h of the power is lost inside the coupler or reflected at port 3. Then,

$$f + g + h = 1 \qquad (2)$$

We can adjust the coupling ratio f by varying the coupling length. The best excess loss values today are $(1-h) = 0.98$ [$10 \log (1-h) = -0.1$ dB]. Typical coupling ratios are $f = 0.5$ (−3 dB) and $f = 0.1$ (−10 dB), with $f \approx h = 0.02$ (−17 dB) as a practical lower limit.

Optical bus. The directionality of optical taps and the linear topology of the bus require separate couplers

Figure 3. Folded unidirectional bus. The dangling ends of couplers are terminated in matched impedances. R and T represent the receiver and transmitter functions, respectively; and α represents loss from fiber and connectors. A repeater at the head end may be provided in an active bus.

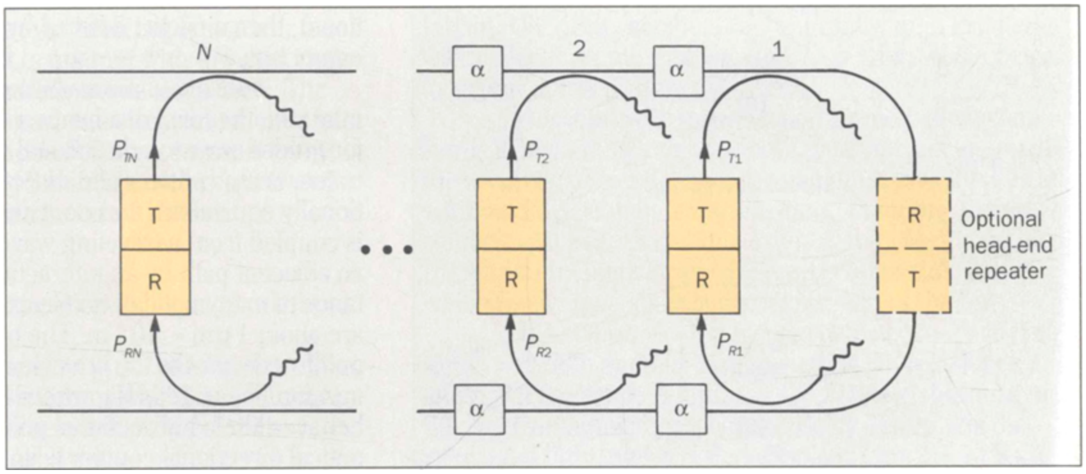

for the transmitter and for the receiver at each station, as illustrated in Figure 3 for a folded bus with N stations. Each box labeled α represents fiber and connector loss between adjacent stations. The number of stations on the bus is limited because the losses α and h add in series between distant stations. Several options are possible:

- A repeater may or may not be provided at the head end. (One is provided in Figure 3.)
- The couplers may be identical. Or, their coupling ratios may be tailored to increase with the distance from the head end, so that the power at each receiver is the same.

A repeatered bus with tailored coupling ratios yields the maximum number of stations for a given power margin P_T/P_R, where P_T is the available optical transmitter power and P_R is the minimum receiver power that provides a tolerable bit-error rate (BER). But a tailored-coupling bus does not appear to be a practical solution for several reasons:

- For a typical "allowed number of stations," the weakest coupler at the head end would require a coupling ratio comparable with the excess coupling loss.
- The precise coupling ratios needed may be difficult to manufacture, and expensive to stock in the wide variety required.
- If new stations are added, all the coupling ratios on the bus must be rearranged. Because we cannot adjust currently available fused-fiber couplers in the field, such a network would be difficult to alter.

A more practical approach is to use identical couplers or a few different coupler types. But the latter approach supports fewer than half as many stations, and the receivers must also have a wider dynamic range.

As a simple example of how the optical coupler restricts the number of stations, consider a receiver bus with identical couplers. This corresponds to the lower half of the bus in Figure 3 with a head-end repeater. The repeater output, P_T, is applied directly on the bus, while $P_R = P_{RN}$ is the output of the coupler at the Nth (last) station. Then, the optimum f is about $1/N$, and the power margin limits N. (See Reference 2.)

Let us take $P_T = 0$ dBm and $P_R = -40$ dBm (which is reasonable for a good receiver at 100 Mb/s), $h = 0.023$ (– 0.1 dB) excess coupler loss, and $\alpha = 0.045$ (– 0.2 dB) connector and fiber loss. Then, for the repeatered case, we find that $N = 64$ stations and $f = 0.016$ (–18 dB). In the unrepeatered case, we are limited to fewer stations ($N = 16$), and $f = 0.067$ (−12 dB). In either case, about the same amount of power is coupled to each receiver as is lost in the coupler, connector, and fiber; i.e., efficiency is marginal.

These estimates suggest that a simple unrepeatered bus is limited to tens of stations, which may still be satisfactory for some networks or for segments of a larger network that has mixed topologies. Of course, if additional repeaters are allowed, the number of stations can grow considerably.

An even more attractive approach might be to include optical amplifiers at each station or after the number of stations that reduces the signal by an optimum fraction.[6] Semiconductor or rare-earth-doped fiber optical amplifiers[7-10] offer gains of more than 20 dB with a reasonable noise figure at lower cost and complexity than a repeater. The network performance of optical amplifiers is being studied widely to determine bandwidth, crosstalk, and saturation constraints, with promising results for high-speed, wide optical band operation. Because these devices amplify in both directions (they are the gain elements of lasers with the feedback eliminated), the network must contain nonreciprocal isolators to avoid the ringing attributed to spurious reflections.

Mesh Networks

A *mesh network* consists of a distribution of active nodes connected by point-to-point links. Unlike ring or bus networks, more than two links generally connect each node. The network of telephone central offices is an example. Mesh networks have tremendous aggregate bandwidth, even though each node and link has a small bandwidth. Furthermore, many alternate paths exist between nodes to assure reliable communications links despite failures of several nodes and links.

A regular mesh network, known as the Manhattan street network,[11] has been proposed as a topology for a packet LAN. Each active node has two inputs and two outputs, as in the bidirectional ring. In the conceptual topology, loops of fiber that run north and south interconnect columns of nodes, and loops that run east and west interconnect rows of nodes. Nodes are placed at the intersections of avenues and streets, as in Manhattan.

The topology has considerable flexibility in joining communities of interest in closed loops. Because two packets may arrive at a node simultaneously, provision must be made for storage and control at each node. Alternatively, the need for storage can be avoided by the following protocol: If two packets enter the node simultaneously, each is assigned an output, even if the output is not the one that leads to the shortest route; nevertheless, the packet is eventually sent to its destination after effectively being stored enroute. A double-ring network has the reliability rerouting feature of this mesh. But in the ring, each repeater must operate at the network bandwidth, while in the mesh, each repeater operates at the terminal bandwidth.

The multihop network[12] is another mesh configuration that incorporates arrays of multiport nodes interconnected by data links. Several jumps, with storage at intermediate nodes, are usually required to interconnect two stations.

While a mesh network has large throughput and good reliability, it may be very expensive to implement where high-speed terminals, large buffers, and O/E or E/O conversions are required at each node. Self-routing optical nodes, in which an optical packet finds its way through a node without conversion, are now being explored[13] and may make mesh networks more attractive.

Active Star Networks

The star topology lends itself naturally to physical applications—such as the wiring of a building or a neighborhood—where many diverse paths must be followed to all offices or homes. Further, it is often convenient to have a central node at which to place a large mainframe computer, file server, or network controller.

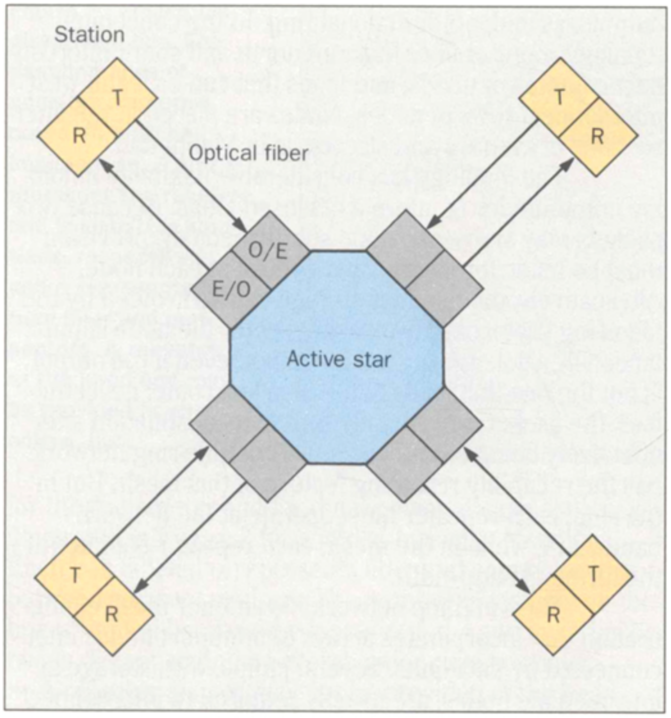

Figure 4. Active star network. Optical-to-electrical (O/E) and electrical-to-optical (E/O) converters must be provided at the star. R and T represent the receiver and transmitter functions, respectively.

The active star shown in Figure 4, in which all incoming optical signals are converted to electrical and then converted back to optical for outgoing signals, permits central control—as opposed to distributed control as in the ring and bus. Then, after any contention protocol has been applied, packets with appropriate address headers may be broadcast to all stations from the central node.

Fiber Ethernet Network. The problems associated with a fiber-optic bus topology for implementing the CSMA/CD protocol, as outlined earlier, can be avoided with an active star—as in Xerox Corporation's Fibernet II network.[14,15]

The main problems with the bus topology are the limited number of users allowed and the difficulty in detecting collisions between weak and strong packets. In the Fibernet II network, transceivers that are compatible with the Ethernet protocol are linked to the active star repeater by dual data links that operate with gallium arsenide (GaAs) LEDs at 10 Mb/s. Collision detection and broadcast retransmission are provided electrically by the star. A 1-MHz square-wave optical signal that cannot be confused with 10-Mb/s packets notifies stations of collisions.

Datakit® VCS. The protocol[16,17] for the AT&T Datakit virtual circuit switch (VCS) provides virtual circuit switching in that a reliable data path is set up for each session, and packet retransmission because of collisions is not required. Remote stations that may consist of mainframe computers, concentrators that bring together many terminals, or gateways to other networks are connected by 8-Mb/s fiber-optic data links to individual electronic modules at the central node. These modules plug into two electronic buses that are short (about 1m) compared to a packet length (16 bytes).

In the module, packets are formed and stored with a header that contains the source address. When the packet is complete (it has the full number of bytes, or a fixed waiting period for added bytes has passed), the module transmits its binary address on the contention bus while listening for bits transmitted by others. If the module transmits a 1 and hears a 1, it transmits the next address bit; if it transmits a 0 and hears a 0, it transmits the next bit. But if it transmits a 0 and hears a 1, it stops transmission, having lost the contention. This process is equivalent to a logical OR operation and assigns the contention to the highest address. The winner transmits the packet on the contention bus in the next time frame.

The switch at the end of the bus replaces the source address with the destination address and transmits the packet on the broadcast bus, where the appropriate module records the destination address and sends the packet to the remote station over the fiber link. Because the switch establishes a correspondence between source and destination at the beginning of a session (as in a circuit switch), source modules need not know the bus position of destination modules. The switch has a directory of positions and terminal names.

If we were to go to very high bit rates, the physical bus length might no longer be short compared with a packet, and collisions caused by delays might upset the "perfect scheduling" of packets. Although Acampora and Hluchyj have proposed methods[17] for overcoming this limitation, the electronic-circuit costs and electrical reflections on the bus may limit the effectiveness of a centralized bus at very high data rates.

Passive Stars and Splitters

Figure 5 illustrates a passive $N \times N$ optical fiber star. In the ideal case, all N input and N output ports are equivalent. If station i transmits power P_T, then P_T/N will appear at each output port (and none at the input ports) and will be broadcast equally to all stations, including i. In practice, of course, the power division may not be equal, output distribution may vary with the input port, and there will be some excess loss.

To construct a simple multimode-fiber star, we can butt N input and N output fibers to a block of glass, in analogy with Figure 5. Single-mode stars are more complicated. They can be fabricated as arrays of −3 dB directional couplers interconnected by fibers that must cross (making an integrated optic version difficult). As Figure 6a illustrates, the −3 dB coupler is itself a 2×2 star. Four of these couplers can be combined (as in Figure 6b) to form a 4×4 star. In a similar way, a large $N \times N$ coupler can be built up by cross connecting $rs \times s$ couplers with $sr \times r$ couplers, where $N = rs$.[18] A novel integrated-optics design promises[19] to provide inexpensive, high-performance stars for N of about 100.

An $N \times N$ star can also be employed as a $1 \times N$ splitter or an $N \times 1$ combiner (as required in tree-network designs), if we use only one input or output port. Alternatively, a corporate array of −3 dB couplers (which requires fewer couplers) can serve as a $1 \times N$ splitter. In

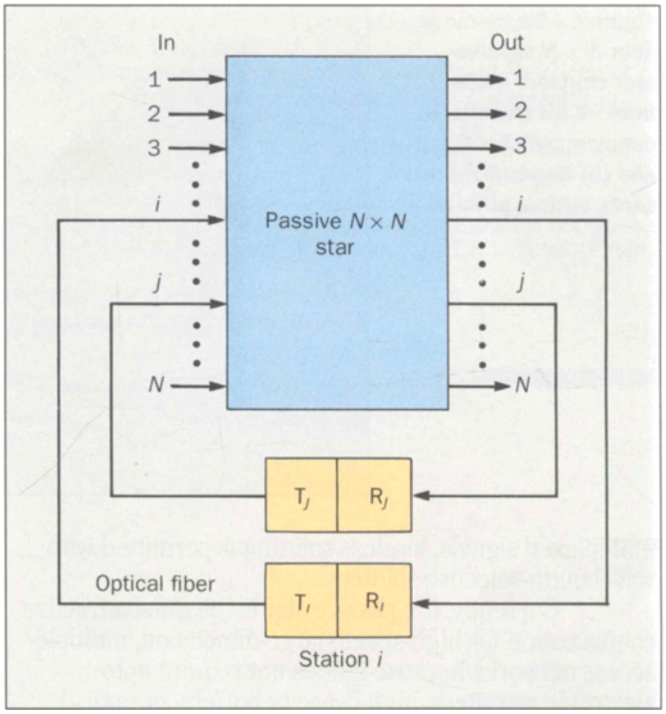

Figure 5. Passive star network. The $N \times N$ star has N single-mode optical-fiber input ports and N optical-fiber output ports. As in the active star (Figure 4), two fibers connect a remote station with the star. R_x and T_x represent the receiver and transmitter functions, respectively, for station x.

either case, recall that these are reciprocal devices, so:

- For unit power incident on a $1 \times N$ splitter, then $1/N$ will appear at each output.
- For unit power incident on any port of an $N \times 1$ combiner, then $1/N$ will appear at the single fiber output.

Thus, if we form an $N \times N$ splitter by connecting an $N \times 1$ with a $1 \times N$ through a single length of fiber, then the splitting loss will be $1/N^2$. These constraints hold for passive single-mode devices at a fixed wavelength; for wavelength-

Figure 6. Single-mode fiber $N \times N$ passive-star couplers made from −3 dB directional couplers: (a) 2 × 2 and (b) 4 × 4. P represents optical power.

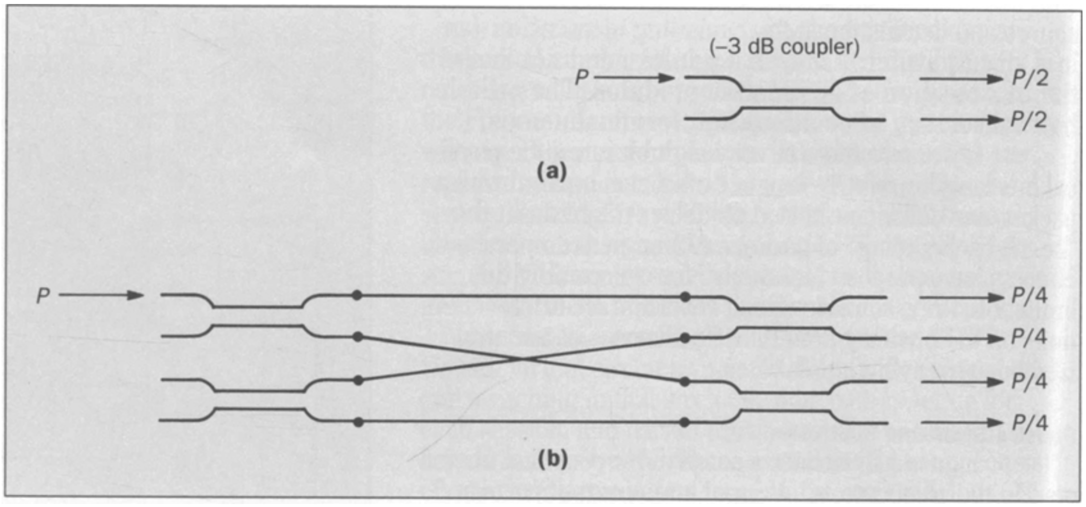

multiplexed signals, lossless splitting is permitted with wavelength-selective splitters.

Currently, the passive star is the most attractive configuration for high-speed, large-dimension, multiple-access networks because it does not require optoelectronic repeaters, high-capacity buffers, or optical taps. The parallel paths through the star reduce the losses due to α and h encountered in the bus and allow more stations. The principal constraint is the intrinsic $1/N$ splitting loss. However, with a reasonable margin of 35 dB (where we allow a generous 15 dB for fiber, connector, and other losses), more than 100 users could be supported at bit rates that exceed 1 Gb/s. Optical amplifiers within the ranks of the star or external to it can further extend N without limiting B or requiring optoelectronic repeaters.

In the companion paper,[1] we review research demonstrations of optical star networks that are based on various channel multiplexing approaches.

References

1. I. P. Kaminow, "Photonic Multiple Access Networks: Routing and Multiplexing," *AT&T Technical Journal*, Vol. 68, No. 2, March/April 1989, pp. 72-86.
2. *Optical Fiber Telecommunications II*, S. E. Miller and I. P. Kaminow (eds.), Academic Press, New York, 1988.
3. I. P. Kaminow, "Photonic Local Networks," *Optical Fiber Telecommunications II*, S. E. Miller and I. P. Kaminow (eds.), Academic Press, New York, 1988, Chapter 26.
4. F. E. Ross, "FDDI—A Tutorial," *IEEE Communications Magazine*, Vol. 24, No. 5, May 1986, pp. 10-17.
5. W. E. Burr, "The FDDI Optical Data Link," *IEEE Communications Magazine*, Vol. 24, No. 5, May 1986, pp. 18-23.
6. P. S. Henry, R. A. Linke, and A. H. Gnauck, "Transmission Systems Principles," *Optical Fiber Telecommunications II*, S. E. Miller and I. P. Kaminow (eds.), Academic Press, New York, 1988, Chapter 21.
7. T. Saitoh and T. Mukai, "Recent Progress in Semiconductor Laser Amplifiers," *Journal of Lightwave Technology*, Vol. 6, 1988, pp. 1656-1664.
8. I. P. Kaminow and R. S. Tucker, "Mode-controlled Semiconductor Lasers," *Guided-Wave Optoelectronics*, T. Tamir (ed.), Springer-Verlag, New York, 1988, Chapter 5.
9. P. Urquhart, "Review of Rare Earth Doped Fibre Lasers and Amplifiers," *IEE Proceedings, Part J*, Vol. 135, 1988, pp. 385-407.
10. C. R. Giles, E. Desurvire, J. R. Talman, J. R. Simpson, and P. C. Becker, "2-Gb/s Signal Amplification at $\lambda = 1.53$ μm in an Erbium-Doped Single-Mode Fiber Amplifier," *Journal of Lightwave Technology*, Vol. 7, 1989, pp. 651-656.
11. N. F. Maxemchuk, "The Manhattan Street Network," *Globecom '85*, IEEE Global Telecommunications Conference, New Orleans,

Louisiana, December 4-5, 1985, pp. 252-261.
12. A. S. Acampora, M. J. Karol, and M. G. Hluchyj, "Terabit Lightwave Networks: The Multihop Approach," *AT&T Technical Journal*, Vol. 66, No. 6, November/December 1987, pp. 21-34.
13. P. A. Perrier and P. R. Prucnal, "Self-clocked Optical Control of a Self-routed Photonic Switch," *Journal of Lightwave Technology*, Vol. 7, to be published 1989.
14. R. V. Schmidt, E. G. Rawson, R. E. Norton, S. B. Jackson, and M. D. Bailey, "Fibernet II: A Fiber Optic Ethernet," *IEEE Journal on Selected Areas in Communications*, Vol. SAC-1, 1983, pp. 702-711.
15. E. G. Rawson, "The Fibernet II Ethernet-compatible Fiber-optic LAN," *Journal of Lightwave Technology*, Vol. LT-3, 1985, pp. 496-501.
16. A. G. Fraser, "Towards a Universal Data Transport System," *IEEE Journal on Selected Areas in Communications*, Vol. SAC-1, No. 5, November 1983, pp. 803-816.
17. A. S. Acampora and M. G. Hluchyj, "A New Local Area Network Architecture Using a Centralized Bus," *IEEE Communications Magazine*, Vol. 22, No. 8, August 1984, pp. 12-21.
18. T. Hermes, B. Hoen, J. Saniter, and F. Schmidt, "LOCNET—A Local Area Network Using Optical Switching," *Journal of Lightwave Technology*, Vol. LT-3, 1985, pp. 467-471.
19. C. Dragone, "Efficient $N \times N$ Star Couplers Using Fourier Optics," *Journal of Lightwave Technology*, Vol. LT-7, 1989, pp. 479-489.

(Manuscript received January 26, 1989)