PROCEEDINGS
SPIE—The International Society for Optical Engineering

Advanced Fiber Communications Technologies

Leonid G. Kazovsky
Chair

3–4 September 1991
Boston, Massachusetts

Sponsored and Published by
SPIE—The International Society for Optical Engineering

SPIE
P
PROCEEDINGS SERIES

Volume 1579

SPIE (The Society of Photo-Optical Instrumentation Engineers) is a nonprofit society dedicated to the advancement of optical and optoelectronic applied science and technology.

The papers appearing in this book comprise the proceedings of the meeting mentioned on the cover and title page. They reflect the authors' opinions and are published as presented and without change, in the interests of timely dissemination. Their inclusion in this publication does not necessarily constitute endorsement by the editors or by SPIE.

Please use the following format to cite material from this book:
Author(s), "Title of Paper," *Advanced Fiber Communications Technologies,* Leonid G. Kazovsky, Editors, Proc. SPIE 1579, page numbers (1991).

Library of Congress Catalog Card No. 91-50766
ISBN 0-8194-0710-0

Published by
SPIE—The International Society for Optical Engineering
P.O. Box 10, Bellingham, Washington 98227-0010 USA
Telephone 206/676-3290 (Pacific Time) • Fax 206/647-1445

Printed in the United States of America.

CONFERENCE COMMITTEE

Conference Chair
Leonid G. Kazovsky, Stanford University

Program Committee
Robert W. Boyd, University of Rochester
P. Nalinaj Fernando, GEC Hirst Research Center (UK)
Bernard Glance, AT&T Bell Laboratories
Scott H. Hinton, AT&T Bell Laboratories
Michelle S. Malcuit, Lehigh University
Abdellatif Marrakchi, Bell Communications Research
Takao Matsumoto, NTT Transmission Systems Laboratories (Japan)
Yong K. Park, AT&T Bell Laboratories
P. T. Poggiolini, Stanford University
Harish R. Sunak, University of Rhode Island
Robert W. Tkach, AT&T Bell Laboratories

Session Chairs
Session 1—Optical Networks
Leonid G. Kazovsky, Stanford University
Karen Liu, IBM/Thomas J. Watson Research Center

Session 2—Novel Optical Transmission Systems
Bernard Glance, AT&T Bell Laboratories
Rabinder N. Madan, Office of Naval Research

Session 3—Performance Analysis of Optical Communication Systems
Brian M. Hendrickson, Rome Laboratories
P. T. Poggiolini, Stanford University

Session 4—Novel Components and New Effects in the Fiber
P. Nalinaj Fernando, GEC Hisrt Research Center (UK)
John M. Zavada, U.S. Army Research Office

INTRODUCTION

Optical fiber communications is now the technology of choice for many current applications. Innovative research and development continue unabated, and will undoubtedly lead to ever more exciting devices, systems, networks, and applications. This collection of papers on optical fiber communications technologies provides some insight into current research efforts, and should be of interest to both researchers and development engineers. This proceedings contains 25 papers from the United States, Italy, United Kingdom, Portugal, Taiwan, Germany, Brazil, China, and France, reflecting the international character of research in optical communications.

The papers are grouped into four sessions presented over two days. The four sessions deal, respectively, with optical networks, transmission systems, performance of optical links, and novel components and new effects in the fiber. The session on optical networks presents novel network ideas and experimental multi-Gb/s network demonstrations aimed at utilizing the huge bandwidth of single-mode fiber. The session on optical transmission focuses on novel coherent transmission methods and on a variety of new modulation, detection, and multiplexing techniques. The session on analysis of optical communication systems contains six papers analyzing the impact of various imperfections (laser phase noise, spontaneous emission noise of optical amplifiers, crosstalk, polarization-mode dispersion) on optical systems; two other papers deal with Reed-Solomon codes and high-speed (8-Gb/s) transmission over subcarrier-multiplexed systems. Finally, the last session deals with novel components and new effects in the fiber.

The credit for the broad coverage of the papers and strong international participation belongs to the authors, speakers, and members of the Program Committee. All of them worked hard to make the conference a success. We are especially grateful to Pierluigi Poggiolini of Stanford University and Karen Liu of IBM for their outstanding efforts.

Leonid Kazovsky
Stanford University

ADVANCED FIBER COMMUNICATIONS TECHNOLOGIES

Volume 1579

CONTENTS

(continued)

ADVANCED FIBER COMMUNICATIONS TECHNOLOGIES

Volume 1579

ADVANCED FIBER COMMUNICATIONS TECHNOLOGIES

Volume 1579

SESSION 1

Optical Networks

Chairs

Leonid G. Kazovsky
Stanford University

Karen Liu
IBM/Thomas J. Watson Research Center

Invited Paper

Architectures and access protocols for multi-channel networks

Mark J. Karol

AT&T Bell Laboratories, 101 Crawfords Corner Rd.,
P.O. Box 3030, Room 4F-529, Holmdel, New Jersey 07733-3030

ABSTRACT

For multiuser, local and metropolitan area networks, lightwave technology provides an opportunity to share an enormous network capacity among all network users. However, the selection of an appropriate media access control algorithm may be more important than the fiber bandwidth and the transmitter power in determining the network throughput. We compare a variety of architectures and access protocols for multi-channel packet networks. Included is a recent architecture that uses coherent optics to address the media access issue.

1. INTRODUCTION

For multiuser, local and metropolitan area networks, lightwave technology provides an opportunity to share an enormous network capacity (potentially tens of terabits per second on a single fiber) among all network users. For example, Wavelength-Division Multiplexing (WDM) offers the potential to support many concurrent transmissions (on different wavelength channels) on a single optical fiber, and thereby move beyond the capacity limits imposed by electronic speeds. However, for multiuser packet communications with unpredictable user demands, the coordination and control of access to the WDM channels is difficult. The selection of an appropriate media access control algorithm is especially critical since it may be more important than the fiber bandwidth and the transmitter power in determining the network throughput.[1]

In this paper, we compare a variety of architectures and access protocols for multi-channel packet networks. We treat ring, bus, and star topologies, and examine the performance of various access protocols. Included is a recent architecture that uses coherent optics to address the media access issue. For motivation, though, we first describe the limitations of single-channel shared-medium networks.

2. SINGLE-CHANNEL SHARED-MEDIUM NETWORKS

The emerging FDDI[2,3] and IEEE 802.6 DQDB[4,5] standards are examples of "single-channel" shared-medium networks. Stations take turns accessing the shared-channel bandwidth as their transmission "requests" contend for a pool of dynamically-allocated bandwidth. As a result, the capacity available to each station falls off at least linearly as more stations are added to share the network capacity. The capacity per station actually falls off faster than linear because factors such as media access overhead and propagation delays result in less than 100% utilization of the bandwidth.[6,7]

The Fiber Distributed Data Interface (FDDI) - see Fig. 1 - consists of dual, counter-rotating 100-Mbit/s fiber-optic rings (to provide redundant data paths for reliability). Media access is controlled by a 100-Mbit/s token-passing protocol in which a token is passed from station to station, signifying the right to transmit information frames on the ring. Although a transmitting station must remove its own frames from the ring, pipelining is achieved because a station immediately releases the token after transmitting its frame(s). Consequently, multiple frames from several users may be simultaneously propagating around the ring. However, there is no concurrency: the throughput of an FDDI ring is at most 100 Mbit/s since only one station (the one with the token) can inject a new frame at any time. The achievable throughput is actually less than 100 Mbit/s because of the wasted bandwidth associated with passing the token from station to station. Using both rings (rather than saving one as a standby ring for fault recovery) allows an effective transmission rate of 200 Mbit/s.

 0-8194-0710-0/91/$4.00

The IEEE 802.6 Distributed Queueing Dual Bus (DQDB) consists of two 150-Mbit/s (or, initially, 45-Mbit/s) contra-flowing, unidirectional buses with slot generators at the head-ends that continuously send fixed-length time slots down the buses (see Fig. 2). Nodes access the time slots via a global distributed queueing algorithm. The operation of the algorithm is based on a Busy/Idle bit and a Request Field in the header of each time slot. In particular, a node requests a time slot on one of the buses ("downstream" toward an intended receiver) by setting a request bit on the opposite bus to inform "upstream" nodes that an additional segment of information is queued for access. Nodes continuously monitor the buses, counting the requests and the idle time slots that pass by, so that they always know how many time slots are reserved by "downstream" nodes. From this information, a node can determine its position in the distributed queue, and know when it has access to an idle time slot. As messages pass their destinations, the receivers copy the messages but do not "remove" the messages from the bus; the messages stay on the bus to its end.

The simple media access controls of FDDI and IEEE 802.6 DQDB exploit the sequential ordering of stations on a ring topology and bus topology, respectively. However, their capacity is limited in comparison to the optical bandwidth potential. To achieve a total throughput greater than an individual station's transmission rate, multiple stations must, on average, be able to successfully inject new packets into the network at any time. If at most one station is permitted to transmit new packets into the network at any time (as in token protocols such as FDDI), or if a time slot on a bus is used at most once (as in DQDB), then the network throughput - shared by all stations - is no more than the transmission rate. Concurrency is needed to increase the network capacity beyond the electronic transmission rate and achieve perhaps hundreds-of-Mbit/s throughput *per station*.

Focusing now on the network architectures, note in Fig. 1 that FDDI signals are regenerated at each station through electro-optic conversions. Also, although DQDB uses logical buses for the media access, the physical implementation in Fig. 3 shows that signals are actually regenerated at each node. That is, DQDB is (most likely) implemented as an active optical bus, not passive tapping of a broadcast bus. In both FDDI and IEEE 802.6 DQDB, lightwave technology is used for its superior transmission properties to interconnect nodes via point-to-point links. The purpose of this paper is to describe some architectures and access protocols for multi-channel networks that incorporate the unique systems opportunities offered by lightwave technology.

3. INCREASING THE CAPACITY OF RING AND BUS NETWORKS

One way to provide some concurrency and increase the capacity of ring and bus networks is to remove packets from the network at their destinations, rather than let them propagate around the entire ring (as in FDDI) or to the end of the bus (as in DQDB). Destination removal of packets allows the bandwidth to be spatially reused. We will illustrate this concept first for a ring topology and then for a bus topology.

Figure 4 shows a bidirectional buffer insertion ring.[8-10] There are no tokens; a station can insert its own packets at any time, provided it has room in its shift register (buffer) to store packets that may arrive while it's transmitting. Because of its point-to-point connections, there are no inefficiencies resulting from media access of shared channels. Buffering "resolves" the contention and prevents collisions, without the need for a token. Also, for low latency, a station does not have to store an entire packet before forwarding it to the next station. Stations do, however, delay relaying incoming packets by the length of the address field so that packets can be removed from the ring at their destinations, thereby immediately making slots available for other packet transmissions. Removing packets at their destinations permits spatial reuse of the ring bandwidth. With a uniform source-destination traffic distribution, the maximum achievable throughput is approximately $8R$, where R is the transmission rate. $2R$ is due to the presence of two rings, each operating at rate R. The additional factor of 4 is achieved because there are on average four new packets simultaneously on each ring; each packet on average travels only one-fourth around the ring before it is removed by its destination.

Similarly, the throughput of an IEEE 802.6 DQDB system is increased if a time slot is used multiple times as it flows down the bus (i.e., once a time slot reaches its destination node, the time slot becomes available for reuse by other downstream source-destination node pairs). This enhancement to DQDB is referred to as Destination Release of slots,[11,12] and is similar in spirit to the destination removal of packets in register (buffer) insertion networks. Figure 5 illustrates an "erasure node" implementation of Destination Release. Erasure nodes "erase" all the slots that have been marked as "read" so that they can be reused downstream for other transmissions.

Erasure nodes also have the attractive feature that they may be used in public networks to erase traffic that has left a particular user's community of interest.

Besides "destination removal," another way to increase the throughput of a single-channel network is to simply increase the transmission rate. This, however, first assumes that it is technologically feasible to operate at the higher rate. Second, even if it is possible to run the system at a higher rate, it may require the use of much more expensive components. Besides the higher cost associated with faster transmission rates, the efficiency of the access protocol also might be lower. Lower-speed protocols may not be appropriate for higher speeds, as some protocols become less efficient as the transmission rate increases.[6]

Alternatively, suppose the network uses r fibers in parallel, and every station has a transceiver for each of r fibers (as in Fig. 6). If multiple packets can be simultaneously transmitted on the fibers, then this parallelism is in some ways equivalent to increasing the transmission rate. The "extra" fibers may already be installed (for reliability and network upgrades), or cost only incrementally more than a single fiber. However, (i) each station must be equipped with a large number of (expensive) transmitters and receivers; (ii) each station must handle or "process" all network traffic that passes, regenerating the signals and keeping copies of the messages destined for itself; and (iii) many "hops" and electro/optic conversions are necessary. Furthermore, it may be difficult to coordinate a large population of distributed stations attempting to share the r high-speed channels.

We will now focus on multi-connected rings (and buses) that use multiple fibers (or wavelengths) in parallel to achieve large throughputs, but which only require a small number (e.g., two) of (fixed-wavelength) transmitters and receivers per station. The objective is to keep the user interface as simple and cheap as possible.

First, however, note that rather than add more fibers to create many communication channels, an alternative is to exploit the vast optical fiber bandwidth and use WDM. WDM requires more expensive transmit/receive components at the nodes, and so the decision to use either WDM or multiple fibers will depend on factors such as the geographical extent of the network. We discuss WDM approaches for ring and star topologies in Sections 5 and 6, respectively.

4. MULTI-CONNECTED RING NETWORKS

Figure 7 shows a block diagram of a multi-connected ring network.[13-19] There are many fibers in the ring to support a large number of concurrent transmissions and provide a large network capacity, but each Network Interface Unit (NIU) has only a small number (e.g., two) of transmitters and receivers to keep costs low. Each NIU directly accesses only a small subset of the channels, with most fibers bypassing a NIU. An NIU is a small (e.g., 3x3) electronic packet router that handles only a fraction of the total network traffic. Each NIU transmits on, and receives from, specific point-to-point optical channels, and packets are relayed through intermediate NIUs to reach their destinations. Only a small number of hops are needed to reach any destination. Typically, the number of hops is proportional to $\log N$, where N is the number of nodes. There are many possible network connection patterns; selection is based on performance, growability, reliability, and simplicity.

Multi-connected ring (and bus) networks offer substantially better performance than single-channel systems, without increasing the number of transmitters and receivers per node. More fibers are needed, but the larger bundle of fibers (more than two) is likely to be already installed in an FDDI or IEEE 802.6 cable, for either reliability or to provide spare capacity. If not, the channels can be formed by multiplexing more than one wavelength on the fibers, or by adding more cables. The network capacity is typically proportional to $4FR$ with F fiber rings and a transmission rate R.

Besides the additional fibers (or wavelengths), there also is the expense of small 3x3 electronic packet routers at each node. Routing decisions are made, in hardware, at each NIU to direct packets to their appropriate output ports, and buffering is provided to resolve the contention that occurs when multiple packets arrive destined for the same output port. Hot potato (or deflection) routing[20,21] is another possible way to resolve this output port contention, and eliminate or reduce the amount of buffering.

5. RING-SHAPED PASSIVE-BUS OPTICAL NETWORKS

The architectures presented so far do not exploit the large bandwidth potential of optical fiber. WDM, though, provides a way to tap the bandwidth. In this section, we describe a ring-shaped passive-bus WDM system. In Section 6, we cover WDM star topologies.

Figure 8 illustrates a block diagram of a dual ring-shaped passive-bus optical network.[22] There are two closed circular loops of fiber: one for clockwise transmissions and one for counterclockwise transmissions. For easy network growth and fault tolerance, each node physically accesses the fibers via passive taps, with only one passive tap (per fiber) needed for both transmitter and receiver. All taps are identical and not wavelength-selective. Dedicated point-to-point links are formed by assigning to each node one fixed transmit wavelength and one fixed receive wavelength (per fiber). The transmit and receive wavelengths are assigned to form a connectivity pattern in which no point-to-point link passes through more passive taps than the link power budget allows. Also, on both fibers, each node receives a different wavelength than it transmits on, so that a node's transmissions don't interfere with the node's receiver.

As in the previous section, packets may need to be routed through intermediate nodes (i.e., regenerated and wavelength translated) in a multihop fashion; this regeneration is accomplished using the same optical transmitters and receivers needed anyway for accessing the fibers. A 3×3 electronic packet switch performs this routing function, using address information in the packet header. Reference [22] proposes an appropriate connectivity pattern for a dual ring-shaped passive-bus optical network. Using this wavelength assignment pattern, the network capacity is 96 Gbit/s for a uniform traffic model, assuming a 30-dB link power budget, ideal 11-dB couplers, 24 wavelengths, and a 1-Gbit/s transmission rate.

A ring-shaped passive bus exploits the power-splitting losses of fiber-optic passive taps so that there is negligible "self interference" due to a node's signal propagating around the closed circular loop and interfering - albeit greatly attenuated - with its own later transmissions. In addition, this attenuation also permits wavelengths to be reused in different portions of the network. In other words, there can be many simultaneous transmissions on each wavelength in different portions of the network. This spatial reuse greatly reduces the number of wavelengths that need be multiplexed on an optical fiber to yield a large, high-capacity, multichannel network with many concurrent transmissions.

We now present bounds on the power penalty due to interference caused by wavelength reuse (including the "self interference").[22] Assuming the interference signals do not add coherently, the power penalty due to this interference is

$$10\log_{10}\left(\frac{1-A}{1-2A}\right)\text{dB}, \qquad (1)$$

where A is the total attenuation along the ring between nodes that transmit on the same wavelength. An upper bound on the power penalty is obtained by assuming all the interference adds coherently. Then, the power penalty is

$$20\log_{10}\left(\frac{1-\sqrt{A}}{1-2\sqrt{A}}\right)\text{dB}. \qquad (2)$$

Equations (1) and (2) assume all passive taps are wavelength-insensitive and all nodes transmit at the same power level.

Figure 9 shows the upper bound (1) and lower bound (2) on the power penalty for a ring-shaped passive bus. Experimental results show that the penalty falls within these bounds for various cases of laser linewidth, ring length, polarization state, and modulation scheme.[23] In all cases, the penalty is below 1.2 dB with 18-dB attenuation between reused wavelengths.

6. WDM STAR-COUPLER-BASED NETWORKS

In a WDM system, one way to allow communication between any pair of users is to provide all users with a transmitter and a receiver for each wavelength. For most applications, though, this becomes prohibitively expensive. So, although many wavelength channels can be multiplexed to provide the concurrency needed in a high-capacity network, it is desirable to keep the interface costs low by providing each user with only a small number of optical transmitters and receivers (e.g., one or two). If all the transmit and receive frequencies are fixed, then there are fixed point-to-point links between nodes and packets may need to be relayed through intermediate nodes to reach their destinations (as covered in the previous two sections). Such a multihop lightwave network requires an O/E and an E/O conversion each hop (using the same access equipment needed at each node in any event).

Alternatively, an optically-transparent packet network can be designed using a passive star coupler and tunable transmitters or tunable receivers, or both. Naturally, the optical components need be rapidly tunable over a broad spectrum for high-speed packet communications. Also, nodes somehow must know which frequency channels to listen to (or transmit on) at what times to yield the necessary connections between users. The simple solution of a common control channel will restrict network throughput.[1] Suppose each user has fixed-tuned transmitters and tunable receivers. Then each user must learn what frequency channels to listen to at what times. Alternatively, suppose each user has tunable transmitters and fixed-tuned receivers. Then the contention for channels need be resolved and pretransmission coordination is necessary to avoid collisions. Later in this section we describe some contention resolution procedures for WDM star-coupler-based networks. To evaluate and compare their performances, though, we first present a few basic queueing-theory results.

In high-performance packet networks, queueing is necessary since two or more packets with the same destination may arrive simultaneously. At most one packet can be transmitted per frequency channel (otherwise there will be a collision), so other packets needing the same frequency channel must be queued for later transmission.

In a star-coupler system, packets can be queued either at the inputs or the outputs of the system. For example, packets are queued at the inputs in the HYPASS[24] and Star-Track[25] systems (these systems are described below), whereas packets are queued at the output ports in the Photonic Knockout Switch.[26] The performance limitations of input queueing were first quantified in [27]. Specifically, it was shown that the capacity (i.e., the maximum achievable throughput) of a large input-queued system is about 59% of the capacity of an output-queued system. This assumes a time-slotted, synchronous system with a uniform traffic pattern and First-In First-Out (FIFO) buffers at each input. Head-of-Line (HOL) Blocking causes this throughput reduction. Of all the packets contending for an output, one is successful (i.e., selected for transmission) and the others are blocked. Those time slots on the blocked inputs are wasted, even if other packets in the FIFO buffers are awaiting transmission to idle output ports. That is, a packet at the head of the FIFO blocks access to idle output ports.

In contrast, output queueing allows all packets to (at least conceptually) be routed to their appropriate output ports, where they are served one-at-a-time. Figure 10 shows the additional (normalized) capacity available to each user in an output-queued system with the number of users N greater than approximately 20.

Star-Track[25] is a multiwavelength packet-switching fabric using fixed-wavelength transmitters and tunable receivers (see Fig. 11). It uses a token reservation scheme to schedule transmissions for each of the wavelength channels, thereby resolving the output-port contention. Because of its input-queueing architecture, HOL Blocking limits its (normalized) capacity to approximately 0.59 per channel. Furthermore, the use of a token reservation scheme restricts the use of Star-Track to small physical distances (i.e., to centralized switch fabrics) and small numbers of nodes.

The HYPASS system,[24] illustrated in Fig. 12, not only suffers from HOL Blocking, it also has some additional inefficiency due to a multiaccess arbitration protocol with collisions. A tree-polling algorithm resolves the contention, yielding a (normalized) capacity of 0.32 per channel for a uniform traffic model.

Glance recently proposed using a simple Protection-Against-Collision (PAC) circuit to solve this media access problem and achieve full optical connectivity.[28] The PAC system has the potential to interconnect a thousand ports transmitting at one gigabit per second.

The PAC Optical Packet Network (see Fig. 13) is an $N \times N$ star-coupler-based FDM system with tunable transmitters and fixed-tuned receivers. Each node can direct its packet to any other node by tuning its carrier frequency to the destination frequency. Packet collisions are avoided by allowing access to the network only if the addressed channel is available. Once a user "acquires" a channel, other users are temporarily blocked from accessing it. Also, packets simultaneously addressing the same available channel are denied access to the star. These two properties result from controlling each input of the network by a PAC circuit. Each node has such a PAC circuit, located at the central hub. The PAC circuit prevents collisions in the star coupler by blocking users from transmitting on already-occupied channels. Protection-Against-Collision (PAC) circuits provide rapid "contention resolution" of all input/output pairs, even for large systems.

A PAC circuit probes the state of the addressed channel (i.e., the energy present in the addressed channel) using an n-bit carrier burst that precedes the packet. The carrier burst is switched through a second $N \times N$ control star coupler, where it is combined with a fraction of all the packets coming out of the network star plus all the carrier bursts trying to gain access to this star. The resulting electrical signal (derived from simple optical power measurements) controls an optical switch (shown in Fig. 13) that connects the input transmission fiber to the network. The switch is closed, allowing access to the star coupler, only if there isn't any energy on the addressed channel from other users (that either are attempting to or have access to the channel). If two or more users simultaneously attempt to access an idle channel, all those users detect energy on the channel and are blocked. "Collisions" in the control star coupler only occur over the length of a packet's carrier burst. A packet that "collides" with another packet in the control star coupler is denied access to the network and is reflected back to its input port, which is thus informed of the packet status. The detection of a reflected signal triggers retransmission of the packet (stored in an electronic buffer at the input) at a later time. The absence of a reflected signal implies successful transmission and allows erasure of the stored data. Because noise can lead to a wrong estimation of the channel state, there is a 10^{-14} probability of collision in the network star when the carrier burst is 20 bits long.[28] Furthermore, if the "feedback" propagation delay from a user to the hub (indicated in Fig. 13) is small, then there is the possibility of many contention processes per packet length. That is, as soon as a user receives its reflected packet, the user knows that the channel is occupied (or that another user is simultaneously accessing the same channel), and it can abort its transmission and attempt to access another frequency channel.

We will assume users transmit fixed-length packets. Let T represent the switching time (normalized to the packet transmission time): the amount of time it takes to switch a transmitter from one frequency channel to another and begin transmitting a packet (including the tuning time and the locking time). Next, let B denote the transmission time of the carrier burst (i.e., the probing delay), including the bits needed for sensing the arrival of the carrier burst and controlling the optical switch. Finally, let F represent the feedback propagation delay (from the transmitter to the central hub, and back), normalized to the packet length.

Suppose ρ ($0 \leq \rho \leq 1$) represents the utilization (i.e., throughput) of a frequency channel in the PAC Network with a uniform traffic pattern. Figure 14 shows the maximum achievable throughput ρ_{max} as a function of F for various switching times T with a probing delay $B = 0$. In geographically-distributed (local and metropolitan area network) applications (i.e., $F \geq 1$) the capacity (normalized to the transmission rate) falls between 0.4 and 0.5 per channel for switching times less than one-half the packet length.[29] For example, the maximum achievable throughput is between 400 and 500 Mbit/s for a 1000-channel system with a 1-Gbit/s transmission rate. Also, for ρ between 0.4 and 0.5 it takes on average between 2 and 2.5 attempts to access each channel. In a centralized switch where the inputs and outputs are all at the same location (i.e., $F < 1$), the (normalized) capacity can exceed 0.8 per channel for small switching times; "collision detection" leads to the improved performance. At $\rho = 0.8$, it takes an average of five attempts to access each channel.

Figure 15 shows the capacity ρ_{max} as a function of F for various switching times T with a probing delay $B = 0.05$. This value corresponds to a probing delay of 22 bits and a small (ATM) packet length of 424 bits. Comparing Fig. 14 with Fig. 15, we see a lower capacity in Fig. 15 due to the overhead of the probing delay (equal to five percent of the packet length). Still, for typical parameters, the (normalized) capacity is between 0.4 and 0.5 per channel in geographically-distributed applications, and approaches 0.8 per channel in a centralized switch.

7. SUMMARY

In single-channel shared-medium networks the capacity available to each station falls off at least linearly as more stations are added to share the network capacity. To achieve a total throughput greater than an individual station's transmission rate, multiple stations must, on average, be able to successfully inject new packets into the network at any time. One way to provide some concurrency and increase the capacity of ring and bus networks is to remove packets from the network at their destinations, rather than let them propagate around the entire ring (as in FDDI) or to the end of the bus (as in DQDB). Destination removal of packets allows the bandwidth to be spatially reused.

Besides "destination removal," another way to increase the throughput of a single-channel network is to simply increase the transmission rate. This, however, first assumes that it is technologically feasible to operate at the higher rate. Second, even if it is possible to run the system at a higher rate, it may require the use of much more expensive components. Alternatively, suppose the network uses r fibers in parallel, and every station has a transceiver for each of r fibers. The "extra" fibers may already be installed (for reliability and network upgrades), or cost only incrementally more than a single fiber. With multiple fibers, either every station can have a transceiver for each of the r fibers, or a multihop network can be designed with each station having only two transceivers.

Rather than use multiple fibers, WDM provides a way to tap the vast bandwidth potential of optical fiber. For easy network growth and fault tolerance, suppose each node physically accesses the WDM fibers via passive taps. Then, the power-splitting losses of the fiber-optic passive taps can be exploited to yield a large, high-capacity, multi-channel network with many concurrent transmissions.

Finally, an optically-transparent packet network can be designed using a passive star coupler and tunable transmitters or tunable receivers, or both. But nodes somehow must know which frequency channels to listen to (or transmit on) at what times to yield the necessary connections between users. The PAC Optical Packet Network uses coherent optics to achieve a (normalized) capacity between 0.4 and 0.5 per channel in geographically-distributed applications, and approaching 0.8 per channel in a centralized switch.

8. REFERENCES

1. P. S. Henry, "Very-High-Capacity Lightwave Networks," *ICC'88 Conf. Rec.*, pp. 1206-1209, June 1988.
2. F. E. Ross, "FDDI - a Tutorial," *IEEE Commun. Mag.*, vol. 24, pp. 10-17, May 1986.
3. F. E. Ross, "An Overview of FDDI: The Fiber Distributed Data Interface," *IEEE J. Select. Areas Commun.*, vol. 7, pp. 1043-1051, Sept. 1989.
4. J. F. Mollenauer, "Standards for Metropolitan Area Networks," *IEEE Commun. Mag.*, vol. 26, pp. 15-19, April 1988.
5. R. M. Newman, Z. L. Budrikis, and J. L. Hullett, "The QPSX Man," *IEEE Commun. Mag.*, vol. 26, pp. 20-28, April 1988.
6. J. O. Limb, "Performance of Local Area Networks at High Speed," *IEEE Commun. Mag.*, vol. 22, pp. 41-45, Aug. 1984.
7. W. Bux, "Token-Ring Local-Area Networks and Their Performance," *Proc. IEEE*, vol. 77, pp. 238-256, Feb. 1989.
8. E. R. Hafner, Z. Nenadal, and M. Tschanz, "A Digital Loop Communications System," *IEEE Trans. Commun.*, vol. COM-22, pp. 877-881, June 1974.
9. C. C. Reames and M. T. Liu, "A Loop Network for Simultaneous Transmission of Variable Length Messages," *Proc. 2nd Annual Symp. Comput. Architecture*, pp. 7-12, 1975.

10. D. E. Huber, W. Steinlin, and P. J. Wild, "SILK: An Implementation of a Buffer Insertion Ring," *IEEE J. Select. Areas Commun.*, vol. SAC-1, pp. 766-783, Nov. 1983.

11. A. M. Perdikaris and M. A. Rodrigues, "Destination Releases of Segments in the IEEE 802.6 Protocol," *Contribution 802.6-88/61 to IEEE 802.6 working group*, July 1988.

12. M. A. Rodrigues, "Erasure Node: Performance Improvements for the IEEE 802.6 MAN," *INFOCOM'90 Conf. Proc.*, June 1990.

13. C. S. Raghavendra and J. A. Silvester, "A Survey of Multi-Connected Loop Topologies for Local Computer Networks," *Computer Networks and ISDN Systems* (North-Holland), vol. 11, pp. 29-42, 1986.

14. C. S. Raghavendra, M. Gerla, and D. S. Parker, "Multi-Connected Loop Topologies for Local Computer Networks," *INFOCOM'82 Conf. Proc.*, pp. 184-190, Mar. 1982.

15. J. A. Silvester and C. S. Raghavendra, "Analysis and Simulation of a Class of Double Loop Network Architectures," *INFOCOM'84 Conf. Proc.*, pp. 30-35, April 1984.

16. K. Y. Eng, "A Multi-Fiber Ring Architecture for Distributed Lightwave Networks," *Proc. ICC'88*, pp. 1490-1496, June 1988.

17. I. Bar-David and M. J. Karol, "Nested Multi-Connected Rings for Large High-Capacity LANs and MANs," *INFOCOM'90 Conf. Proc.*, June 1990.

18. M. J. Karol, "Optical Interconnection Using ShuffleNet Multihop Networks in Multi-Connected Ring Topologies," *Proc. ACM SIGCOMM'88 Symposium*, pp. 25-34, Aug. 1988.

19. M. J. Karol and R. D. Gitlin, "High-Performance Optical Local and Metropolitan Area Networks: Enhancements of FDDI and IEEE 802.6 DQDB," *IEEE Jour. Select. Areas Commun.*, vol. 8, pp. 1439-1448, Oct. 1990.

20. P. Baran, "On Distributed Communications Networks," *IEEE Trans. Comm. Sys.*, pp. 1-9, Mar. 1964.

21. N. F. Maxemchuk, "Regular and Mesh Topologies in Local and Metropolitan Area Networks," *AT&T Tech. Jour.*, vol. 64, pp. 1659-1686, Sept. 1985.

22. M. J. Karol, "Exploiting the Attenuation of Fiber-Optic Passive Taps to Create Large High-Capacity LAN's and MAN's," *IEEE Jour. Lightwave Tech.*, vol. 9, pp. 400-408, Mar. 1991.

23. M. A. Santoro and M. J. Karol, "Experimental and Theoretical Performance of Ring-Shaped Passive-Bus Optical Networks", *IEEE Photon. Tech. Lett.*, pp. 490-492, May 1991.

24. E. Arthurs *et al.*, "HYPASS: An Optoelectronic Hybrid Package Switching System," *IEEE Jour. Select. Areas Commun.*, vol. 6, pp. 1500-1510, Dec. 1988.

25. T. T. Lee, M. S. Goodman, and E. Arthurs, "A Broadband Optical Multicast Switch," *Proc. International Switching Symposium 1990*, vol. III, pp. 7-13, May 1990.

26. K. Y. Eng, "A Photonic Knockout Switch for High-Speed Packet Networks," *IEEE Jour. Select. Areas Commun.*, vol. 6, pp. 1107-1116, Aug. 1988.

27. M. J. Karol, M. G. Hluchyj, and S. P. Morgan, "Input vs. Output Queueing on a Space-Division Packet Switch," *Globecom'86 Conf. Rec.*, pp. 659-665, Dec. 1986. Also in *IEEE Trans. Commun.*, vol. COM-35, pp. 1347-1356, Dec. 1987.

28. B. Glance, "PAC Optical Packet Network," submitted to *IEEE Jour. Lightwave Tech.*.

29. M. J. Karol and B. Glance, "Performance of the PAC Optical Packet Network," submitted to *IEEE Jour. Lightwave Tech.*.

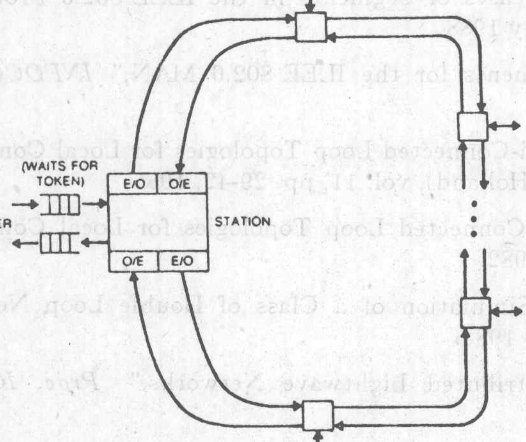

Fig. 1 Fiber Distributed Data Interface (FDDI)

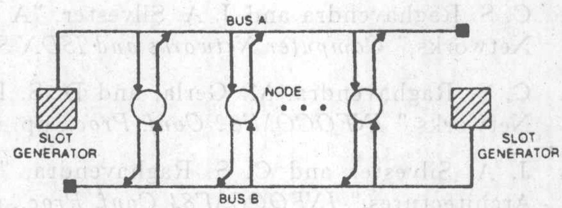

Fig. 2 Distributed Queueing Dual Bus (DQDB)

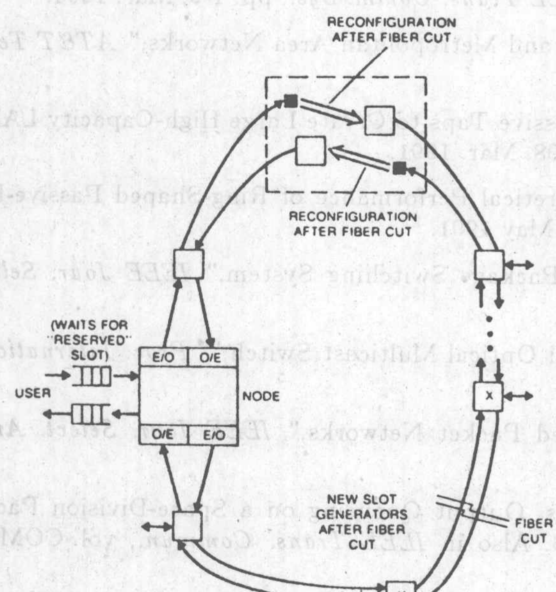

Fig. 3 Looped Bus Architecture For IEEE 802.6 DQDB

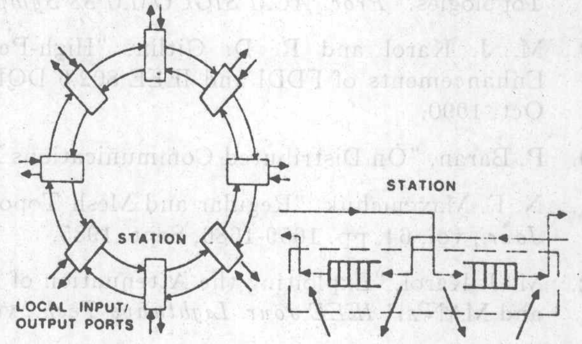

Fig. 4 Bidirectional Buffer Insertion Ring

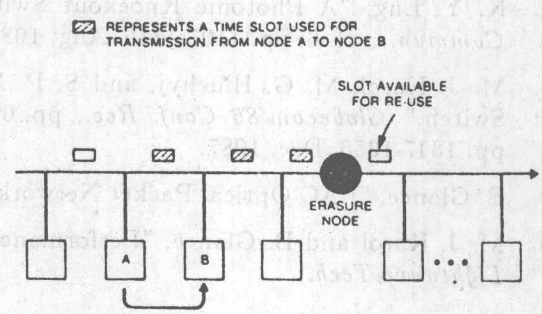

Fig. 5 Erasure Node Implementation Of Destination Release In IEEE 802.6 DQDB

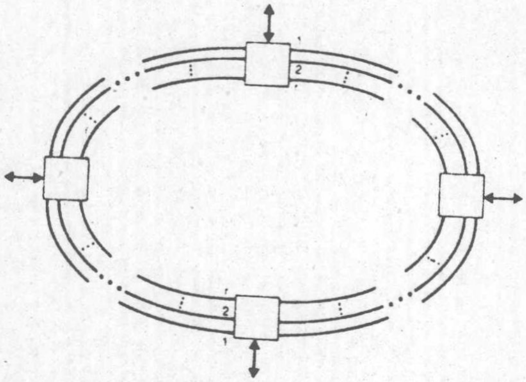

Fig. 6 Multiple Channel Ring Network

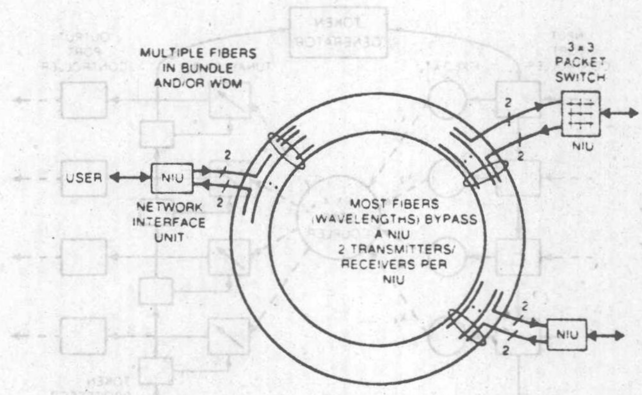

Fig. 7 Multi-Connected Ring Network

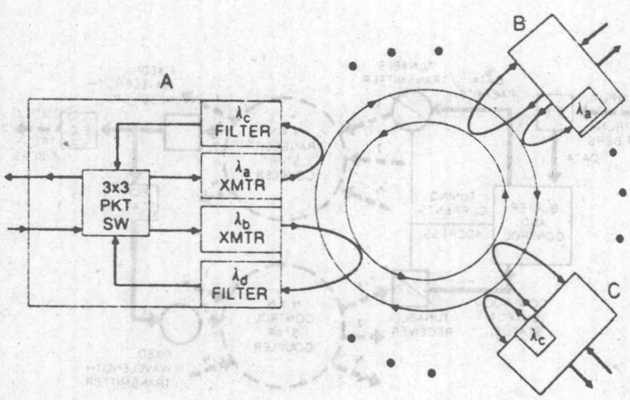

Fig. 8 A Dual Ring-Shaped Passive-Bus Optical Network

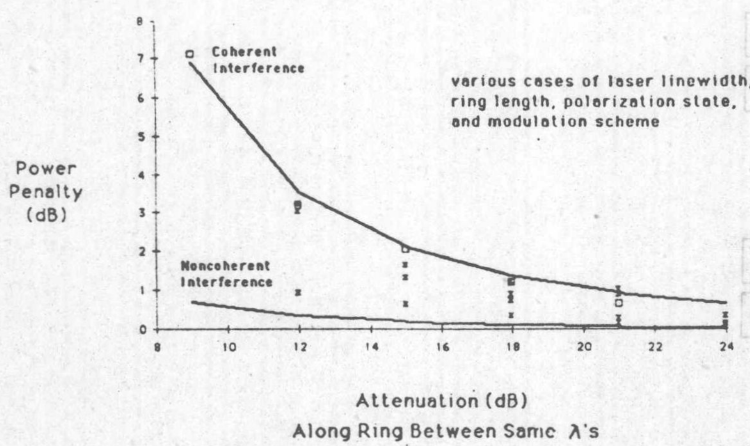

Fig. 9 Bounds On The Power Penalty For A Ring-Shaped Passive Bus (Experimental Results Also Shown)

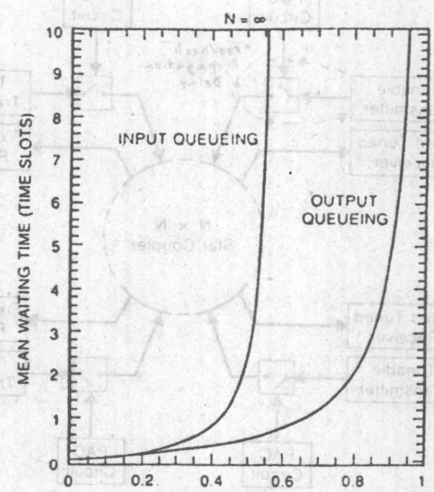

Fig. 10 Mean Waiting Time Vs. Offered Load: Input And Output Queueing