WDM Optical Communication Networks: Progress and Challenges

Biswanath Mukherjee, *Member, IEEE*

Invited Paper

***Abstract*—While optical-transmission techniques have been researched for quite some time, optical "networking" studies have been conducted only over the past dozen years or so. The field has matured enormously over this time: many papers and Ph.D. dissertations have been produced, a number of prototypes and testbeds have been built, several books have been written, a large number of startups have been formed, and optical WDM technology is being deployed in the marketplace at a very rapid rate. The objective of this paper is to summarize the basic optical-networking approaches, briefly report on the WDM deployment strategies of two major U.S. carriers, and outline the current research and development trends on WDM optical networks.**

***Index Terms*—Carrier strategies, optical communication network, research trends, wavelength routing, WDM.**

I. INTRODUCTION

WE ARE moving toward a society which requires that we have access to information at our fingertips *when we need it, where we need it, and in whatever format we need it*. The information is provided to us through our global mesh of communication networks, whose current implementations, e.g., today's Internet and asynchronous transfer mode (ATM) networks, do not have the capacity to support the foreseeable bandwidth demands.

Fiber-optic technology can be considered our saviour for meeting our above-mentioned need because of its potentially limitless capabilities [1], [2]: huge bandwidth [nearly 50 terabits per second (Tb/s)], low signal attenuation (as low as 0.2 dB/km), low signal distortion, low power requirement, low material usage, small space requirement, and low cost. Our challenge is to turn the promise of fiber optics to reality to meet our information networking demands of the next decade (and well into the 21st century!).

Thus, the basic premise of the subject on optical wavelength-division multiplexing (WDM) networks is that, as more and more users start to use our data networks, and as their usage patterns evolve to include more and more bandwidth-intensive networking applications such as data browsing on the worldwide web (WWW), java applications, video conferencing, etc., there emerges an acute need for very high-bandwidth transport network facilities, whose capabilities are much beyond those that current high-speed (ATM) networks can provide. There is just not enough bandwidth in our networks today to support the exponential growth in user traffic!

Fig. 1 shows the past and projected future growth of data and voice traffic reported by most telecom carriers [3]. Although voice traffic continues to experience a healthy growth of approximately 7% per year (which would be considered to be very strong growth in most business sectors), it is the data-traffic growth that is drawing people's attention. Most carriers report that data traffic has just recently overtaken or will soon overtake the voice traffic in their fiber links and networks [4].

Given that a single-mode fiber's potential bandwidth is nearly 50 Tb/s, which is nearly four orders of magnitude higher than electronic data rates of a few gigabits per second (Gb/s), every effort should be made to tap into this huge opto-electronic bandwidth mismatch. Realizing that the maximum rate at which an end-user—which can be a workstation or a gateway that interfaces with lower-speed subnetworks—can access the network is limited by electronic speed (to a few Gb/s), the key in designing optical communication networks in order to exploit the fiber's huge bandwidth is to introduce concurrency among multiple user transmissions into the network architectures and protocols. In an optical communication network, this concurrency may be provided according to either wavelength or frequency [wavelength-division multiplexing (WDM)], time slots [time-division multiplexing (TDM)], or wave shape [spread spectrum, code-division multiplexing (CDM)].

Optical TDM and CDM are somewhat futuristic technologies today. Under (optical) TDM, each end-user should be able to synchronize to within one time slot. The optical TDM bit rate is the aggregate rate over all TDM channels in the system, while the optical CDM chip rate may be much higher than each user's data rate. As a result, both the TDM bit rate and the CDM chip rate may be much higher than electronic processing speed, i.e., some part of an end user's network interface must operate at a rate higher than electronic speed. Thus, TDM and CDM are relatively less attractive than WDM, since WDM—unlike TDM or CDM—has no such requirement.

Specifically, WDM is the current favorite multiplexing technology for long-haul communications in optical communication networks since all of the end-user equipment needs to operate

Manuscript received November 15, 1999; revised June 1, 2000. This work was supported in part by the NSF under Grants NCR-9508239, ECS-9521249, and ANI-9805285.

The author is with the Department of Computer Science, University of California, Davis, CA 95616 USA (e-mail: mukherjee@cs.ucdavis.edu).

Publisher Item Identifier S 0733-8716(00)08999-X.

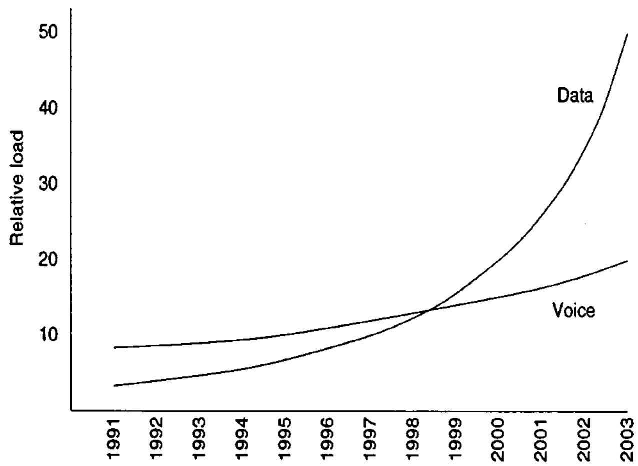

Fig. 1. Past and projected future growth of data and voice traffic.

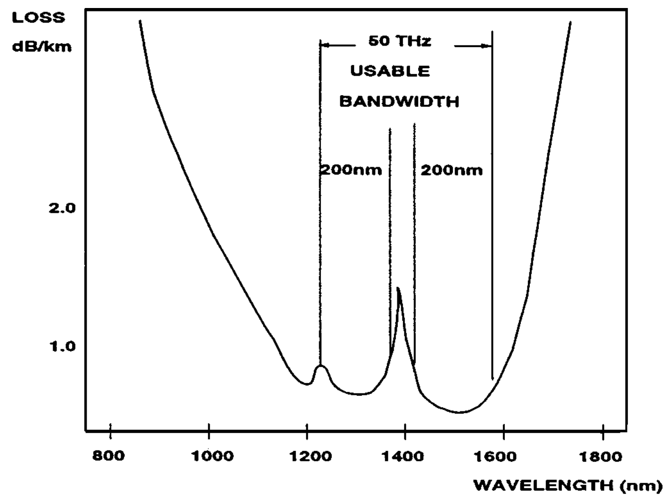

Fig. 2. The low-attenuation regions of an optical fiber.

only at the bit rate of a WDM channel, which can be chosen arbitrarily, e.g., peak electronic processing speed. Hence, the major carriers today all devote significant effort to developing and applying WDM technologies in their businesses. In particular, market analysts indicate that the North American WDM transport market (for optical backbone networks) is expected to grow from $1.9 billion in 1998 to over $3 billion in 2002.

II. Wavelength-Division Multiplexing (WDM)

Wavelength-division multiplexing (WDM) is an approach that can exploit the huge opto-electronic bandwidth mismatch by requiring that each end-user's equipment operate only at electronic rate, but multiple WDM channels from different end-users may be multiplexed on the same fiber. Under WDM, the optical transmission spectrum (see Fig. 2) is carved up into a number of nonoverlapping wavelength (or frequency) bands, with each wavelength supporting a single communication channel operating at whatever rate one desires, e.g., peak electronic speed. Thus, by allowing multiple WDM channels to coexist on a single fiber, one can tap into the huge fiber bandwidth, with the corresponding challenges being the design and development of appropriate network architectures, protocols, and algorithms. Also, WDM devices are easier to implement since, generally, all components in a WDM device need to operate only at electronic speed; as a result, several WDM devices are available in the marketplace today, and more are emerging.

Research and development on optical WDM networks have matured considerably over the past few years, and they seem to have suddenly taken on an explosive form, as evidenced by recent publications [1], [5]–[10] on this topic as well as new conferences, workshops, journals, and magazines devoted to ths topic. A number of experimental prototypes have been and are currently being developed, deployed, and tested mainly by telecommunication providers including a plethora of startup companies. It is anticipated that the next generation of the Internet will employ WDM-based optical backbones.

Current development activities indicate that this sort of WDM network will be deployed mainly as a backbone network for large regions, e.g., for nationwide or global coverage. (However, optical WDM networks for local and metropolitan applications are also being developed and prototyped.) End-users—to whom the architecture and operation of the backbone will be transparent except for significantly improved response times—will attach to the network through a wavelength-sensitive switching/routing node. An end-user in this context need not necessarily be a terminal equipment, but the aggregate activity from a collection of terminals—including those that may possibly be feeding in from other regional and/or local subnetworks—so that the end-user's aggregate activity on any of its transmitters is close to the peak electronic transmission rate.

We should also remark that an important WDM market segment that is emerging but that has not received much attention in the academic research literature is the "access network," one potential solution to which is the passive optical network (PON). In the absence of much research literature, we are unable to do justice to this topic in this paper.

III. WDM Networking Evolution

A. Point-to-Point WDM Systems

WDM technology is being deployed by several telecommunication companies for point-to-point communications. This deployment is being driven by the increasing demands on communication bandwidth. When the demand exceeds the capacity in existing fibers, WDM is turning out to be a more cost-effective alternative compared to laying more fibers. A study [11] compared the relative costs of upgrading the transmission capacity of a point-to-point transmission link from OC-48 (2.5 Gb/s) to OC-192 (10 Gb/s) via the following three possible solutions. (The terminology OC-n is a widely used telecommunications jargon. "OC" stands for "optical channel" and it specifies electronic data rates. "OC-n" stands for a data rate of $n \times 51.84$ megabits per second (Mb/s) approximately; so OC-48 and OC-192 correspond to approximate data rates of 2.5 Gb/s and 10 Gb/s, respectively. OC-768 (40 Gb/s) is the next milestone in highest achievable electronic communication speed.)

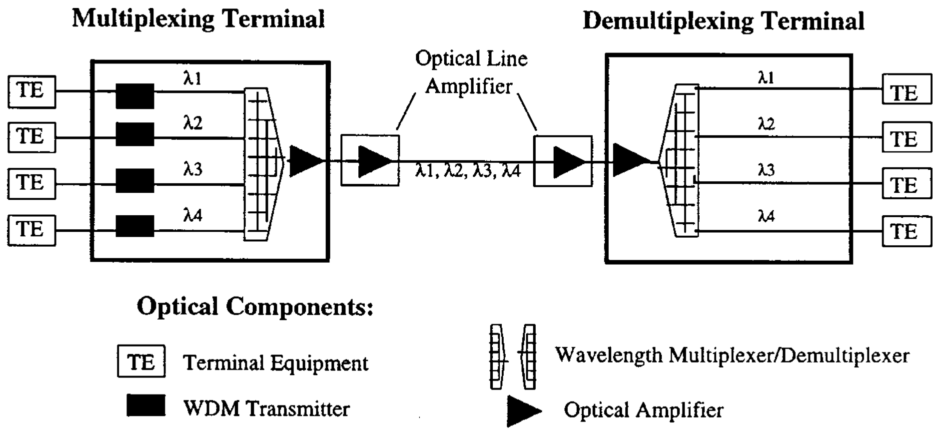

Fig. 3. A four-channel point-to-point WDM transmission system with amplifiers.

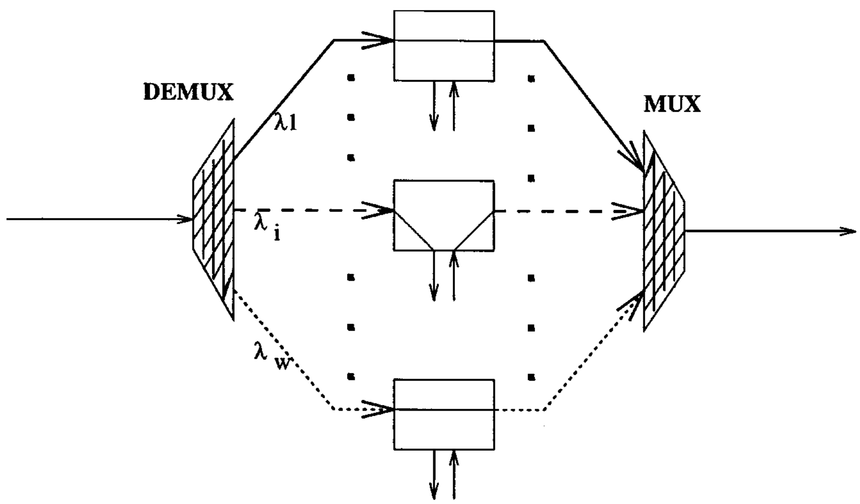

Fig. 4. A wavelength add/drop multiplexer (WADM).

1) Installation/burial of additional fibers and terminating equipment (the "multifiber" solution).
2) A four-channel "WDM solution" (see Fig. 3) where a WDM multiplexer (mux) combines four independent data streams, each on a unique wavelength, and sends them on a fiber; and a demultiplexer (demux) at the fiber's receiving end separates out these data streams.
3) OC-192, a "higher-electronic-speed" solution.

The analysis in [11] shows that, for distances lower than 50 km for the transmission link, the "multifiber" solution is the least expensive; but for distances longer than 50 km, the "WDM" solution's cost is the least with the cost of the "higher-electronic-speed" solution not that far behind.

WDM mux/demux in point-to-point links is now available in product form from several vendors such as IBM, Pirelli, and AT&T [12]. Among these products, the maximum number of channels is 64 today, but this number is expected to increase soon.

B. Wavelength Add/Drop Multiplexer (WADM)

One form of a wavelength add/drop multiplexer (WADM) is shown in Fig. 4. It consists of a demux, followed by a set of 2×2 switches—one switch per wavelength—followed by a mux. The WADM can be essentially "inserted" on a physical fiber link. If all of the 2×2 switches are in the "bar" state, then all of the wavelengths flow through the WADM "undisturbed." However, if one of the 2×2 switches is configured into the "cross" state (as is the case for the λ_i switch in Fig. 4) via electronic control (not shown in Fig. 4), then the signal on the corresponding wavelength is "dropped" locally, and a new data stream can be "added" on to the same wavelength at this WADM location. More than one wavelength can be "dropped and added" if the WADM interface has the necessary hardware and processing capability.

The above "functionality" of add/drop multiplexing can be enhanced for practical implementation of an optical add/drop multiplexer (OADM). For example, an OADM may drop part of the traffic on a wavelength or it may drop a band of wavelengths.

C. Fiber and Wavelength Crossconnects—Passive Star, Passive Router, and Active Switch

In order to have a "network" of multiwavelength fiber links, we need appropriate fiber interconnection devices. These devices fall under the following three broad categories:

- passive star (see Fig. 5),
- passive router (see Fig. 6), and
- active switch (see Fig. 7).

The *passive star* is a "broadcast" device, so a signal that is inserted on a given wavelength from an input fiber port will have

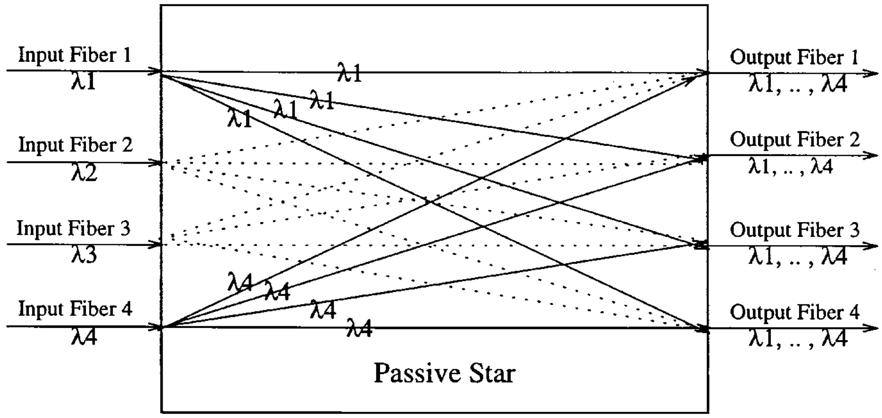

Fig. 5. A 4 × 4 passive star.

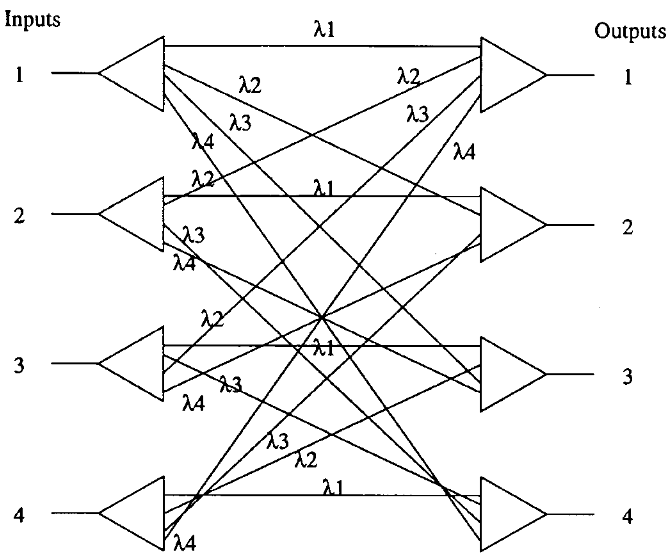

Fig. 6. A 4 × 4 passive router (four wavelengths).

its power equally divided among (and appear on the same wavelength on) all output ports. As an example, in Fig. 5, a signal on wavelength λ_1 from Input Fiber 1 and another on wavelength λ_4 from Input Fiber 4 are broadcast to all output ports. A "collision" will occur when two or more signals from the input fibers are simultaneously launched into the star on the same wavelength. Assuming as many wavelengths as there are fiber ports, an $N \times N$ passive star can route N simultaneous connections through itself.

A *passive router* can separately route each of several wavelengths incident on an input fiber to the same wavelength on separate output fibers, e.g., wavelengths λ_1, λ_2, λ_3, and λ_4 incident on Input Fiber 1 are routed to the same corresponding wavelengths to Output Fibers 1, 2, 3, and 4, respectively, in Fig. 6. Observe that this device allows *wavelength reuse*, i.e., the same wavelength may be spatially reused to carry multiple connections through the router. The wavelength on which an input port gets routed to an output port depends on a "routing matrix" characterizing the router; this matrix is determined by the internal "connections" between the demux and mux stages inside the router (see Fig. 6). The routing matrix in a *passive router* is "fixed" and cannot be changed. Such routers are commercially available, and are also known as Latin routers, waveguide grating routers (WGRs), wavelength routers (WRs), etc. Again, assuming as many wavelengths as there are fiber ports, a $N \times N$ passive router can route N^2 simultaneous connections through itself (compared to only N for the passive star); however, it lacks the broadcast capability of the star.

The *active switch* also allows *wavelength reuse*, and it can support N^2 simultaneous connections through itself (like the passive router). But the active star has a further enhancement over the passive router in that its "routing matrix" can be *reconfigured* on demand, under electronic control. However the "active switch" needs to be powered and is not as fault-tolerant as the passive star and the passive router which don't need to be powered. The active switch is also referred to as a *wavelength-routing switch (WRS)*, *wavelength selective crossconnect (WSXC)*, or just crossconnect for short. (We will refer to it as a WRS in the remainder of this paper.)

The active switch can be enhanced with an additional capability, viz., a wavelength may be converted to another wavelength just before it enters the mux stage before the output fiber (see Fig. 7). A switch equipped with such a wavelength-conversion facility is more capable than a WRS, and it is referred to as a *wavelength-convertible switch*, *wavelength interchanging crossconnect (WIXC)*, etc.

Although the active switch in Fig. 7 is assumed to include the wavelength demux/mux devices, note that the mux/demux devices could be part of the (multiwavelength) fiber transmission system. Thus, the mux/demux devices are not included in the optical crossconnect products being developed by most vendors today.

The passive star is used to build local WDM networks, while the active switch is used for constructing wide-area wavelength-routed networks. The passive router has mainly found application as a mux/demux device. The WADM device and its variants are mainly used to build optical WDM ring networks which are expected to deployed mainly in the metropolitan area market.

IV. New and Interesting WDM Technologies

Readers interested in the enabling fiber-optic technologies are referred to [13] for detailed discussions and to [1] for an overview. However, newer approaches and technologies are

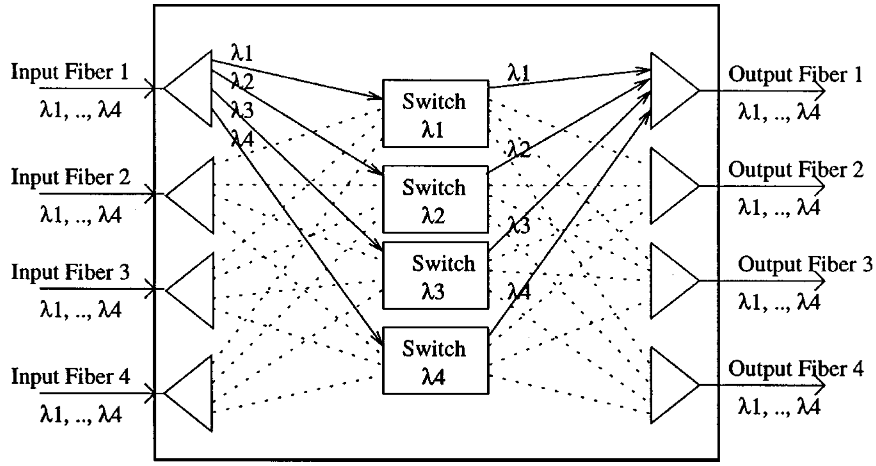

Fig. 7. A 4×4 active switch (four wavelengths).

constantly under development. While the objective of this paper is to abstract the device-layer functionalities to create useful network architectures, the author wishes to highlight two new and exciting technologies that can potentially further revolutionize the design and effectiveness of WDM networks. These two approaches are briefly reviewed in this section so that this exposure will enable a number of our readership to exploit these features and create novel network architectures.

Recent developments in fiber optics have further expanded the usable fiber bandwidth. A new type of fiber, called the *allwave fiber*, does not have the 1385 nm "water-peak window" (which the conventional fiber has) and hence it provides a more usable optical spectrum. Fig. 8 shows the low-attenuation region of the allwave fiber in comparison to that of the conventional fiber.

Another new and interesting technology is a new type of amplifier device which uses the erbium-doped fiber amplifier (EDFA) as a building block. The normal EDFA has a gain spectrum of 30–40 nm (typically in the 1530–1560 nm range). Thus, fibers employing EDFAs for long-haul communication would normally require that their carried signals be confined to the EDFA gain spectrum. However, by using the EDFA as an elementary gain building block, we can build a "circuit" of EDFAs configured such that the "circuit designer" has full control of the gain as well as gain bandwidth of the composite "circuit" (see Fig. 9 for an example) [14]. This "amplifier circuit" is referred to as an *ultrawide-band EDFA*, which can fully exploit the expanded low-attenuation region of the "allwave fiber." As an analogy, if we think of the EDFA as equivalent to a "transistor," then the *ultrawide-band EDFA* device would be analogous to an "integrated circuit" [14].

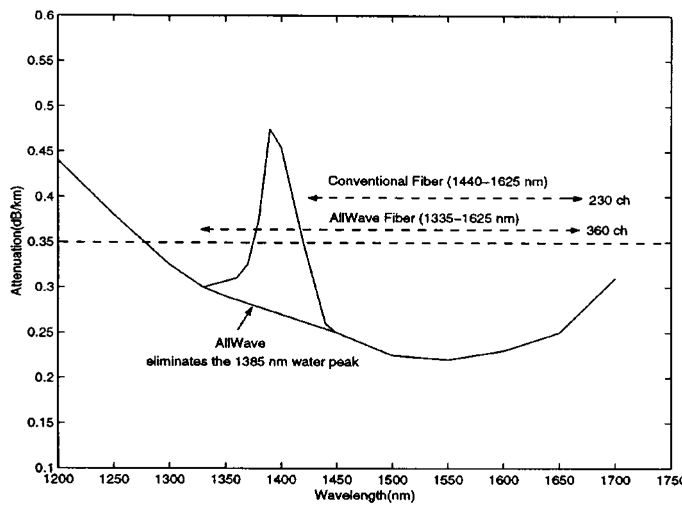

Fig. 8. Low-attenuation region of allwave fiber versus conventional fiber.

V. WDM Network Constructions

A. Broadcast-and-Select (Local) Optical WDM Network

A local WDM optical network may be constructed by connecting network nodes via two-way fibers to a passive star, as shown in Fig. 10. A node sends its transmission to the star on one available wavelength, using a laser which produces an optical information stream. The information streams from multiple sources are optically combined by the star and the signal power of each stream is equally split and forwarded to all of the nodes on their receive fibers. A node's receiver, using an optical filter, is tuned to only one of the wavelengths; hence it can receive the information stream. Communication between sources and receivers may follow one of two methods: 1) *single-hop* [15], or 2) *multihop* [16]. Also, note that, when a source transmits on a particular wavelength λ_1, more than one receiver can be tuned to wavelength λ_1, and all such receivers may pick up the information stream. Thus, the passive-star can support "*multicast*" services.

Detailed, well-established discussions on these network architectures can be found elsewhere, e.g., [15], [16], [1]. Hence, local WDM network architectures will not be discussed here any further; instead, the remainder of this paper will concentrate mainly on wide-area (wavelength-routed) optical networks.

B. Wavelength-Routed (Wide-Area) Optical Network

A wavelength-routed (wide-area) optical WDM network is shown in Fig. 11. The network consists of a *photonic switching fabric*, comprising "active switches" connected by fiber links to form an arbitrary *physical topology*. Each end-user is connected to an active switch via a fiber link. The combination of an end-user and its corresponding switch is referred to as a network *node*.

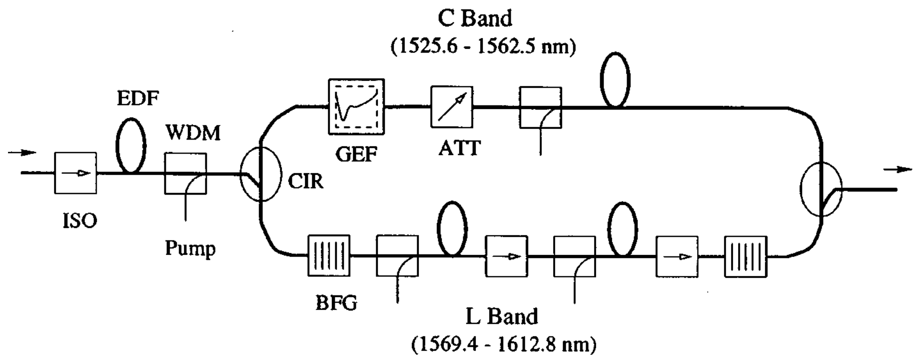

Fig. 9. Ultrawide-band EDFAs.

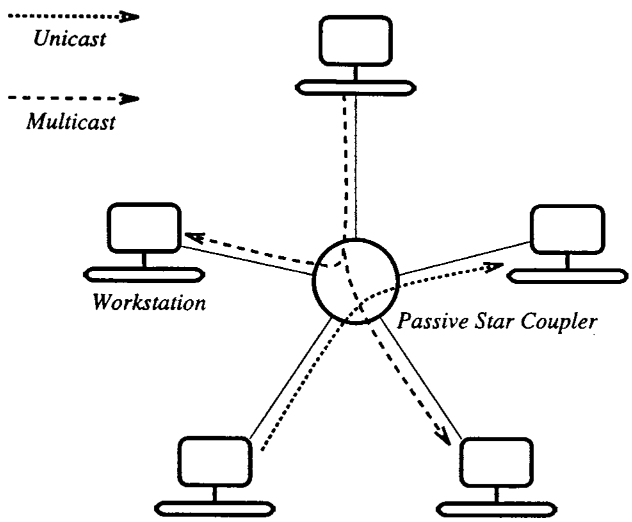

Fig. 10. A passive-star-based local optical WDM network.

Each node (at its access station) is equipped with a set of transmitters and receivers, both of which may be wavelength tunable. A transmitter at a node sends data into the network and a receiver receives data from the network.

The basic mechanism of communication in a wavelength-routed network is a *lightpath*. A *lightpath* is an all-optical communication channel between two nodes in the network, and it may span more than one fiber link. The intermediate nodes in the fiber path route the lightpath in the optical domain using their active switches. The end-nodes of the lightpath access the lightpath with transmitters and receivers that are tuned to the wavelength on which the lightpath operates. For example, in Fig. 11, lightpaths are established between nodes A and C on wavelength channel λ_1, between B and F on wavelength channel λ_2, and between H and G on wavelength channel λ_1. The lightpath between nodes A and C is routed via active switches 1, 6, and 7. (Note the wavelength reuse for λ_1.)

In the absence of any wavelength-conversion device, a lightpath is required to be on the same wavelength channel throughout its path in the network; this requirement is referred to as the *wavelength-continuity* property of the lightpath. This requirement may not be necessary if we also have wavelength converters in the network. For example, in Fig. 11, the lightpath between nodes D and E traverses the fiber link from node D to switch 10 on wavelength λ_1, gets converted to wavelength λ_2 at switch 10, traverses the fiber link between switch 10 and switch 9 on wavelength λ_2, gets converted back to wavelength λ_1 at switch 9, and traverses the fiber link from switch 9 to node E on wavelength λ_1.

A fundamental requirement in a wavelength-routed optical network is that two or more lightpaths traversing the same fiber link must be on different wavelength channels so that they do not interfere with one another.

C. A Sample WDM Networking Problem

As we have described in Subsection B, end-users in a fiber-based WDM backbone network may communicate with one another via *all-optical (WDM) channels*, which are referred to as *lightpaths*. A *lightpath* may span multiple fiber links, e.g., to provide a "circuit-switched" interconnection between two nodes which may have a heavy traffic flow between them and which may be located "far" from each other in the physical fiber network topology. Each intermediate node in the lightpath essentially provides an all-optical bypass facility to support the lightpath.

In an N-node network, if each node is equipped with $N-1$ transceivers [transmitters (lasers) and receivers (filters)] and if there are enough wavelengths on all fiber links, then every node pair could be connected by an all-optical lightpath, and there is no networking problem to solve. However, it should be noted that the network size (N) should be scalable, transceivers are expensive so that each node may be equipped with only a few of them, and technological constraints dictate that the number of WDM channels that can be supported in a fiber be limited to W (whose value is a few tens today, but is expected to improve with time and technological breakthroughs). Thus, only a limited number of lightpaths may be set up on the network.

Under such a network setting, a challenging networking problem is, given a set of lightpaths that need to be established on the network, and given a constraint on the number of wavelengths, to determine the routes over which these lightpaths should be set up and also determine the wavelengths that should be assigned to these lightpaths so that the maximum number of lightpaths may be established [17]. While shortest-path routes may be most preferable from the individual point of view of each lightpath, note that this choice may have to be sometimes sacrificed, in order to allow more lightpaths to be set up. Thus, one may allow several alternate routes for lightpaths to be established [18]. Lightpaths that cannot be set up due to constraints on routes and wavelengths are said to be

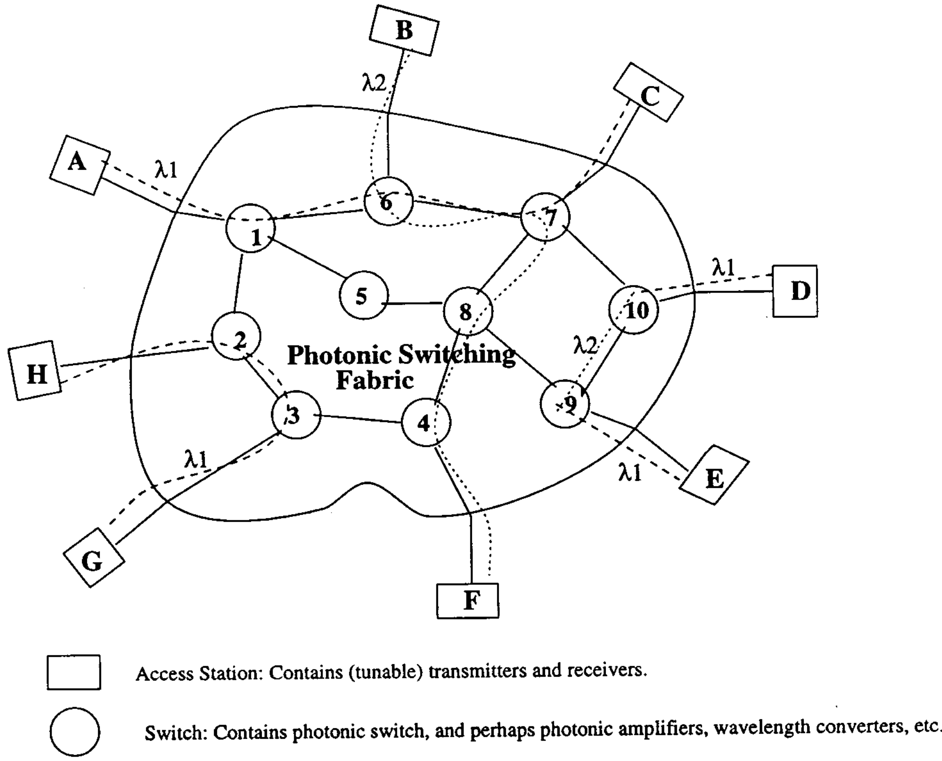

Fig. 11. A wavelength-routed (wide-area) optical WDM network.

blocked, so the corresponding network optimization problem is to minimize this blocking probability.

In this regard, note that, normally, a lightpath operates on the same wavelength across all fiber links that it traverses, in which case the lightpath is said to satisfy the *wavelength-continuity constraint*. Thus, two lightpaths that share a common fiber link should not be assigned the same wavelength. However, if a switching/routing node is also equipped with a *wavelength-converter facility*, then the *wavelength-continuity constraints* disappear, and a lightpath may switch between different wavelengths on its route from its origin to its termination.

This particular problem, referred to as the routing and wavelength assignment (RWA) problem, has been examined in detail in [19], while the general topic of *wavelength-routed networks* has been studied in [1].

Returning to our sample networking problem, note that designers of next-generation lightwave networks must be aware of the properties and limitations of optical fiber and devices in order for their corresponding protocols and algorithms to take advantage of the full potential of WDM. Often a network designer may approach the WDM architectures and protocols from an overly simplified, ideal, or traditional networking point of view. Unfortunately, this may lead an individual to make unrealistic assumptions about the properties of fiber and optical components, and hence may result in an unrealizable or impractical design.

VI. Carrier Strategies

In this section, we examine two major carriers' strategies in developing WDM networks based on published information [20], [21]. Several other carriers are also developing similar strategies; however, the objective of this section is not to provide an exhaustive list of all carriers' strategies. Instead, we focus on two major carriers whose approaches are representative of the industry trend and whose strategies were available to the author.

A. AT&T

The North American WDM transport market is expected to grow from \$1.9 billion in 1998 to over \$3 billion in 2002 [3]. AT&T is experiencing annual traffic growth of 5–15% in voice and 20–70% in data in its various markets, which leads to its total traffic (voice + data) to increase threefold in two years. AT&T is fully committed to WDM optical networking to solve its bandwidth problems [20].

The AT&T WDM approach is driven by the MONET vision. (MONET is a DARPA-sponsored project called Multiwavelength Optical Network [22].) MONET has the following three-layered approach.

1) Application layers will include voice, data, video, multimedia, image, LAN traffic, and optical access.
2) Electronic-switching layers will include ATM/SONET, legacy switches, and future formats.
3) Configurable WDM layers will include WDM Add/Drop Multiplexers (WADM), Wavelength-Routed Stars (R), and WDM Cross-Connects (XC).

The local-exchange network will have WADMs connected in ring fashion. Rings are interconnected through XCs. Stars can be connected to XCs or WADMs. The long-distance network consists of XCs interconnected in arbitrary mesh topology and it connects the local-exchange networks. Electronic equipment can interface directly with WADMs, Stars, or XCs in either the local-exchange network or the long-distance network.

AT&T believes that large-scale transparency in WDM networks has many problems. Transparency limits scalability in capacity or network size because 1) impairments accumulate (dispersion, nonlinearity, crosstalk, noise, amplifier gain flatness, etc.), 2) network engineering becomes a nightmare, 3) fault location, performance monitoring, multivendor interworking, and adoption of new technology (new wavelength plans, new modulation formats, etc.) become difficult. Hence, AT&T is currently following the “opaque” optical-networking approach where each crossconnect will perform optoelectronic conversion. This approach will aid in performance monitoring, fault location, arresting of impairments, etc. As a result, the backbone network becomes transponder-based.

AT&T also believes that there will be “islands of transparency” in its otherwise opaque network. The opaque backbone will carry most traffic at high SONET rates.

Crossconnect size plays an active role in the development and the employment of WDM networks. Currently, a large central office in AT&T’s network has tens of OC-48 signals coming in, so that a 64 $\times$ 64 switch (using mechanical switching technology) is sufficient. With WDM, in the future, a typical node may have degree 3–4, a few 100 fibers in, with 8/16/32/$\cdots$ wavelengths, so that a switch size could grow well beyond 1000 $\times$ 1000, and require multi-Tb/s switching capacity! Prospects for very large switches are encouraging due to advances in integration of switch fabrics using waveguides, semiconductor gate arrays, and MEMS.

Provisioning of a private line can be done in seconds. AT&T believes that most applications would be satisfied with subsecond restoration, so that switches need not be very fast.

B. MCI Worldcom

Just like AT&T, MCI Worldcom is experiencing strong growth in its data traffic. From 3Q97 to 3Q98, its voice traffic grew by 15%, while its data traffic grew 40–70% in various markets [21].

MCIs investigation indicates that transport speed increased from OC-12 to OC-48 in 1993 and reached OC-192 in 1998. It is expected to be OC-768 in 2001. As to data switch port speeds, it has also been growing fast during the past years. It increased from DS3 to OC-3 in 1993, was OC-12 in 1996, OC-48 in 1999, and is expected to be OC-192 in 2000 and reach OC-768 in 2002.

MCI believes the following are technology enablers for an all-optical network: fixed lasers/receivers on ITU grid, tunable lasers, next-generation fiber, optical cross-connect switch (OCCS), all-optical regenerators, broadband amplifier, grating technologies, optical ADM, wavelength routers, all-optical signal processing, and end-to-end network management.

According to MCI, initially, WDM can be used for increasing capacity. Then, WDM can be used for optical networking, when combined with O/E/O (like AT&T’s opaque vision). Finally, WDM can be used to implement all-optical networking, when no O/E/O is necessary in such networks, i.e., with optical 3R signal generation (which provides signal regeneration/amplification as well as reshaping and retiming for digital signals). (Please see Section VII-D-2 below for more discussion on 1R, 2R, and 3R signal regeneration.)

MCIs hierarchical network has three components:

- an access network is a network in a metro area,
- an edge network is a local or state network, and
- a core network is a national network. The core is targeted for optical networking initially.

The WDM network will be run as rings, with ATM and IP support. Electronic interfaces will run at OC-48 and OC-192. Protection/restoration will be provided. Quite a bit of analysis of these various options have been conducted.

MCI believes an all-optical network architecture should be as follows.

- Optical rings interconnected with one another to form optical meshes (like AT&T’s approach). This belief might come from the fact that MCIs economic analysis indicates that optical WDM ring is the cheapest (capital) long-term data architecture.
- Some wavelengths will carry ATM traffic, others will carry IP traffic.
- Electronic digital cross connects (DXCs) will allow the ATM and IP traffic to interface with the network.
- The network will contain OADMs and OCCSs to interconnect the rings.
- Optical regenerators will include optical amplification and dispersion compensation, and will provide transparency to bit rate, wavelength, and protocol.

Besides the all-optical network, optical OAM&P will be run in each electronic terminal (DXC), all of which collectively form the operation support system (OSS).

MCI classifies the implementation of OCCS into three phases.

- *Phase I:* The architecture can either be opaque (as in AT&T’s vision) or transparent. In an opaque architecture, mesh or ring restoration is enabled and it is a medium-sized large-loss optical matrix. In a transparent architecture, small node rings exist for restoration and it is a dedicated-functionality optical matrix.
- *Phase II:* Larger $N \times N$ optical matrix with low loss is used to provide restoration, plus elementary management of interring traffic.
- *Phase III:* Very large (1024 $\times$ 1024) optical matrix with low loss is used, and this phase has the functionality of bandwidth provisioning and management.

To summarize, it appears from [21] that MCI plans to implement an all-optical network and their analysis indicates feasibility and economic advantages. MCI believes that today’s technology is workable but still somewhat immature and the challenge is management of an all-optical network (management for performance, fault, configuration, and restoration). The requirement for the development of WDM networks is standards for vendor equipment interoperability, e.g., if a fault (fiber cut) occurs, is restoration controlled by IP, or ATM, or OADM?

VII. Recent Trends in WDM Research

Recent research interests in WDM networks include network control and management, fault management, multicasting, physical-layer issues, IP over WDM, traffic grooming, and optical packet switching, just to name a few. We briefly examine these topics below.

A. Network Control and Management

In a wavelength-routed WDM network, a control mechanism is needed to set up and take down all-optical connections (i.e., lightpaths) [19], [23], [24]. Upon the arrival of a connection request, this mechanism must be able to select a route, assign a wavelength to the connection, and configure the appropriate optical switches in the network. The mechanism must also be able to provide updates to reflect which wavelengths are currently being used on each fiber link so that nodes may make informed routing decisions. This control mechanism can either be centralized or distributed. Distributed systems are usually more robust than centralized systems; so they are generally more preferred. The objectives of various research efforts on this subject are to minimize 1) the blocking probability of connection requests, 2) the connection setup delays, and 3) the bandwidth used for control messages; as well as to maximize the scalability of such networks [23].

There are two distributed network control management schemes which have been examined in the literature. The first approach is proposed in [24], and we refer to it as the "link-state approach" because it routes connections in a link-state fashion. The second approach is proposed by us in [23], and we refer to it as the "distributed routing approach" because it utilizes the distributed Bellman–Ford routing algorithm. We describe the two approaches below.

In the link-state approach (which may be implemented using the open shortest path first (OSPF) algorithm), each node maintains the complete network topology, including information on which wavelengths are in use on each fiber link. Upon the arrival of a connection request, a node utilizes the topology information to select a route and a wavelength. Once the route and wavelength are selected, the node attempts to reserve the selected wavelength along each fiber link on the route by sending reservation requests to each node in the route. If an intermediate node is able to reserve the wavelength on the appropriate link, it sends an acknowledgment directly back to the source node. If all of the reservations are successful, then the source sends a SETUP message to each of the nodes. The appropriate switches are then configured at each node, and the connection is established. If even one of the reservations is not successful, then the call is blocked and the source node sends a TAKEDOWN message to each node on the route in order to release the reserved resources. When a connection is established or torn down, each node involved in the connection broadcasts a topology-update message which indicates any changes in the status of wavelengths being used on the node's outgoing links.

In the distributed-routing approach, routes are selected in a distributed fashion without knowledge of the overall network topology. Each node maintains a routing table which specifies the next hop and the cost associated with the shortest path to each destination on a given wavelength. The cost may reflect hop counts or actual fiber-link distances. The routing table is established by employing a distributed Bellman–Ford algorithm [25]. In the distributed-routing approach, upon receiving a connection request, a node will choose the wavelength which results in the shortest distance to the destination, and it will forward the connection request to the next node in the path. Thus, the connection request is routed one hop at a time, with each node along the route independently selecting the next hop based on routing information, and reserving the appropriate wavelength on the selected link. Once the request reaches the destination node, the destination node sends an acknowledgment back to the source node along the reverse path (i.e., it sends the acknowledgment back to the node from which it received the connection request). Upon receiving the ACK, each node along the reverse path conFigs. its wavelength-routing switch. The source node begins transmitting data after it receives the acknowledgment. If a node along the path is unable to reserve the desired wavelength on a link, it will send a negative acknowledgment back to the source along the reverse path. The nodes on the reverse path will release the reserved wavelengths as they receive the negative acknowledgment. The source node may then reattempt the connection on a different wavelength. If the source node is unable to establish the connection on any wavelength, the call is blocked. Once a connection is established, each node along the route sends to each of its neighbors an update message reflecting the status of the newly occupied link and wavelength. Each node receiving an update message may then update its routing table. Similar updates occur when a connection is taken down.

Both schemes can be implemented using a separate channel, namely, a control channel in the WDM network. Alternately, the control information can be carried in-band, as in multiprotocol label switching (MPLS), whose wavelength-routing version is referred to as MP lambda switching. Besides supporting the signaling protocol and the network-topology and status update protocol, the control channel should also have the ability to discover and recover from faults, which we explain below.

B. Fault Management

In a wavelength-routed WDM network (as well as in other networks), the failure of a network element (e.g., fiber link, cross-connect, etc.) may cause the failure of several optical channels, thereby leading to large data (and revenue) losses. Studies have been conducted to examine different approaches to protect WDM optical networks from single fiber-link failures. Earlier studies were focused on ring topologies while recently mesh-based networks have been considered. (See, for example, [26]–[29] for a discussion of these problems. We also remark that our discussion on fault management will be very brief because this is just one of several topics that are covered in this paper. Also, we won't be able to do justice to the large and increasing amount of literature on this subject.)

There are several approaches to ensure fiber-network survivability. Survivable network architectures are based either on dedicating backup resources in advance, or on dynamic restoration. In dedicated-resource survivability (which includes automatic protection switching and self-healing rings [30], [31]), the disrupted network service is restored by utilizing the dedicated network resources. In dynamic restoration, the spare-capacity available within the network is utilized for restoring services affected by a failure. Generally, dynamic restoration schemes are more efficient in utilizing capacity due to the multiplexing of the spare-capacity requirements, and they provide resilience against different kinds of failures, while dedicated restoration schemes have a faster restoration time and provide guarantees on the restoration ability.

We will examine three approaches to protecting against fiber-link failures in an optical network: a dedicated-resource survivability approach called $1 + 1$ protection, and *two* dynamic approaches called link restoration and path restoration. (Again, for detailed discussions on this subject, please consult the literature, e.g., [28] and references therein, e.g., link versus path protection, link versus path restoration, $1 + 1$ versus $1:1$ versus $\mathrm{M}:\mathrm{N}$ versus $0:1$ protection, etc.) Each of these approaches has different wavelength-capacity requirements and blocking performance. So far, research studies have focused on single-link failures, which are the most common form of failures in optical networks.

- *$1 + 1$ Protection:* For each connection that needs to be $1+1$ protected, a dedicated link-disjoint backup route and wavelength is set up in advance (at the time of connection setup). Upon a link failure on the primary path, the end-nodes of the connection start utilizing the backup route and wavelength.
- *Link Restoration:* In link restoration, all the connections that traverse the failed link are rerouted around that link. The source and destination nodes of the connections traversing the failed link are oblivious to the link failure. The end-nodes of the failed link dynamically discover a route around the link, for each wavelength that traverses the link. Upon a failure, the end-nodes of the failed link may participate in a distributed procedure to hunt for new paths for each active wavelength on the failed link. When a new route is discovered around the failed link for a wavelength channel, the end-nodes of the failed link reconFigure their cross-connects to reroute that channel onto the new route. If no new routes are discovered for a wavelength channel, the connection that utilizes that wavelength is blocked.
- *Path Restoration:* In path restoration, when a link fails, the source node and the destination node of each connection that traverses the failed link are informed about the failure (possibly via messages from the nodes adjacent to the link failure). The source and the destination nodes of each connection independently discover a backup route on an end-to-end basis (such a backup path could be on a different wavelength channel). When a new route and wavelength channel is discovered for a connection, network elements such as wavelength cross-connects are reconfigured appropriately, and the connection switches to the new path. If no new routes (and associated wavelength) are discovered for a broken connection, that connection is blocked.

For a comprehensive review of the literature on the design of survivable optical networks, please consult the literature, e.g., [28], [29] and references therein.

C. Multicasting—"Light Trees"

Multicasting is the ability of a communication network to accept a single message from an application and to deliver copies of the message to multiple recipients at different locations. The connections (i.e., lightpaths) we examined earlier are point-to-point connections. However, we can extend the lightpath concept to point-to-multipoint, i.e., a multicast concept, which we call a "light tree" [29], [32]. A light tree enables a transmitter at a node to have many more logical neighbors, thereby leading to a denser virtual interconnection diagram and lower hop distance. A collection of light trees embedded on an optical WDM backbone network can improve the performance of unicast, multicast, and broadcast traffic. However, the corresponding network will require multicast-capable optical switches and more power budget to combat the effect of power losses due to signal splitting. These tradeoffs, including an integer linear program (ILP) formulation of the problem, are examined in [29], [32].

Many multicast applications exist, but their implementations are not necessarily efficient because today's networks were designed to mainly support point-to-point communication. Such applications include video conferencing, software/file distribution including file replication on mirrored sites, distributed games, Internet news distribution, e-mail mailing lists, etc. In the future, as multicast applications become more popular and bandwidth intensive, there will emerge a pressing need to provide multicast support in the underlying communication network. A multicast-capable WDM WAN should not only support efficient routing for multicast traffic, but it may also enhance routing for unicast traffic by allowing more densely connected virtual topologies. To realize multicast-capable WDM WANs, we need to develop novel multicast-capable WRS architectures and design multicast routing and wavelength assignment algorithms, as outlined below.

1) Multicast-Capable WRS Architectures: An optical splitter splits an input signal into multiple output signals, with the output powers being governed by the splitter's "splitting ratio." Fig. 12 shows a 2×2 multicast-capable wavelength-routing switch (MWRS), which can support four wavelengths on each link. The information from each incoming fiber link is first demultiplexed (DEMUX) into separate wavelengths. Then, the separate signals, each on separate wavelengths, are switched by an optical switch (OSW). Unicast signals are sent directly to OSW ports corresponding to their output links, while those signals which need to be multicast are sent to an OSW port connected to a splitter bank. (*The splitter bank may be enhanced to provide optical signal amplification, wavelength conversion, and signal regeneration for multicast as well as unicast connections.*) For example, in Fig. 12, wavelength λ_a is an unicast signal, and λ_b is a multicast signal. The output of the splitter is connected to a smaller optical switch, which routes the multicast signals to their respective output fiber links. Since an optical splitter is a passive device, the power of the signals at each output of an n-way splitter (with equal splitting ratio) is $1/n$ times the input power. To be detected, the optical signal power needs to be higher than a threshold; hence, an MWRS may have a limited multicasting capability.

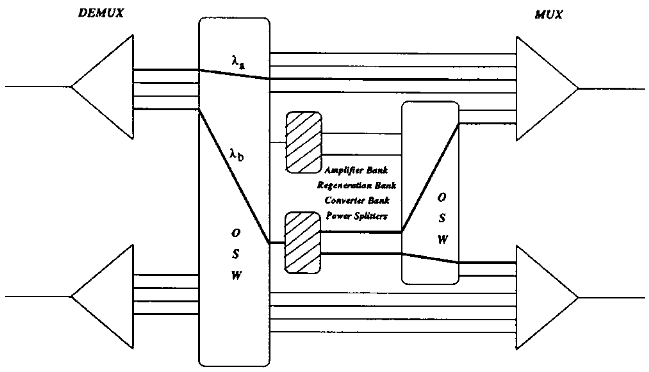

Fig. 12. A multicast-capable WRS (MWRS).

2) Multicast Routing and Wavelength Assignment: The multicasting problem in communication networks is often modeled as the Steiner tree problem in networks (SPN), which is defined as follows. Given

- a graph $G = (V, E)$,
- a cost function C: $E \to R^+$, and
- a set of nodes $D \subset V$,

find a subtree $T = (V_T, E_T)$ which spans D, such that its cost $C(T) = \sum_{e \in E_T} C(e)$ is minimized.

In [34], SPN was shown to be NP-complete. Heuristics are preferred in practice. See [33] for a tutorial on multicast routing algorithms and protocols. Mechanisms described in Section VII-A can be extended to route and assign wavelength to multicast traffic. See also [32] on how to set up "light trees" to improve the performance of unicast and broadcast traffic.

D. Physical-Layer Issues

To date, research on optical network architectures has taught us that, without a sound knowledge of device capabilities and limitations, one can produce architectures which may be unrealizable; similarly, research on new optical devices, conducted without the concept of a useful system, can lead to sophisticated technology with limited or no usefulness. Optics has many desirable characteristics, but it also possesses some not-so-desirable properties. It is beneficial to correct for several of these "mismatches" using intelligent network algorithms, two examples of which are outlined in the following subsections.

1) BER-Based Call Admission in Wavelength-Routed Optical Networks: In a wavelength-routed optical network, a new call can be admitted if an all-optical lightpath can be established between the call's source and destination nodes. Most previous networking studies have concentrated on the RWA problem to set up lightpaths while assuming an ideal physical layer, e.g., [35], [36], [17]. It should, however, be noted that a signal degrades in quality due to physical-layer impairments as it proceeds through switches (picking up crosstalk) and EDFAs (picking up amplified spontaneous emission (ASE) noise). As a result, the bit-error rate (BER) at the receiving end of a lightpath may become unacceptably high.

The work in [37] estimates the *on-line BER* on candidate routes and wavelengths before setting up a call. Thus, one approach would be to set up a call on a lightpath with minimum BER. Another approach would be to establish a call on any lightpath with a BER lower than a certain threshold, e.g., 10^{-12}; if no such lightpath is found, the call is blocked. It is feasible to develop network-layer solutions to combat the physical-layer impairments, including laser shift, dispersion in fiber, and also impairments that affect optical components such as wavelength converters, switch architectures, etc.

2) Sparse Optical Regeneration in Nationwide WDM Networks: Transparency in optical networks is a matter of intense debate. Transparency, in the strict sense, implies that the physical medium (an optical WDM channel in our case) should support end-to-end communication of data, independent of bit rates and signal formats. However, this requires that the physical medium preserve all information in the form of phase, frequency, and analog amplitude of the optical signal during its transport through the network. Such optical transparency is difficult to achieve in practice, especially for long-haul nationwide transport networks. If all future data transmission will be digital, a limited form of transparency known as *digital transparency* may be all that is needed, where digital signals of any bit rates up to a certain limit are accommodated in the network.

As has been mentioned before, the quality of an optical signal degrades as it travels "long distances" through several optical components, e.g., fiber segments, WRSs, EDFAs, etc. Thus, "long-distance" lightpaths may require signal regeneration at strategic locations in a nationwide or global WDM network. As a result, network-design algorithms that perform routing and wavelength assignment, virtual-topology embedding, wavelength conversion, etc., must also be mindful of the locations of the sparse signal regenerators in the network. Such regenerators or "refueling stations," which are placed at select locations in the network, "clean up" the optical WDM signal either entirely in the optical domain or through an opto-electronic conversion followed by an electro-optic conversion. Thus, the signal from the source travels through the network "as far as possible" before its quality (e.g., BER) drops below a certain threshold, thereby requiring it to be regenerated at an intermediate node. The same signal could be regenerated several times in the network before it reaches the destination.

The following three forms of signal regeneration are available, in practice [1].

- *1R—(Optical) Regeneration:* Simple amplification using EDFAs or other amplifiers. It provides total data transparency (since amplification process is independent of signal's modulation format).
- *2R—Regeneration and Reshaping:* In this approach, the optical signal is converted to an electronic signal which is then used to directly modulate a laser. Reshaping of the signal reproduces the original pulse shape of each bit, eliminating much of the noise. Reshaping applies primarily to digitally modulated signals, but in some cases may also be applied to analog signals.
- *3R—Regeneration, Reshaping, and Reclocking:* In today's digital networks (e.g., SONET and SDH), which use optical fiber only as a transmission medium, optical signals

are amplified by first converting the information stream into an electronic data signal, and then retransmitting the signal optically. The reclocking of the signal synchronizes the signal to its original bit timing pattern and bit rate. Reclocking applies only to digitally modulated signals.

Other networking studies which attempt to incorporate physical-layer device characteristics while attempting to solve network-layer problems include amplifier placement in wavelength-routed optical networks [38]–[41].

E. IP Over WDM

The need for high bandwidth in today's IP-based Internet, and the promise of WDM to provide this high capacity, is fueling the current investigations on IP-over-WDM networks. In an IP-over-WDM network, network nodes employ wavelength-routing switches (WRSs) and IP routers. Nodes are connected by fibers to form an arbitrary physical mesh topology. Any two IP routers in this network can be connected together by an all-optical WDM channel, called a lightpath, and the set of lightpaths that are set up form a virtual interconnection pattern.

A lightpath is a point-to-point all-optical wavelength channel that connects a transmitter at a source node to a receiver at a destination node. Using WRSs at intermediate nodes and via appropriate routing and wavelength assignment, a lightpath can create virtual (or logical) neighbors out of nodes that are geographically far apart in the network; thus, a set of lightpaths embeds a virtual (or logical) topology on the network. In the virtual topology, a lightpath carries not only the direct traffic between the nodes it interconnects, but also traffic from nodes upstream of the source (including the source) to nodes downstream of the destination (including the destination). Nodes that are not connected directly in the virtual topology can still communicate with one another using the "multihop approach," namely, by using electronic packet switching at the intermediate nodes in the virtual topology. This electronic packet-switching functionality can be provided by IP routers, ATM switches, etc., leading to an IP-over-WDM or an ATM-over-WDM network, respectively.

In such an optical WDM network architecture, the failure of a network component such as a fiber can lead to the failure of all of the lightpaths that traverse the failed fiber. Since each lightpath is expected to operate at a rate of few Gb/s, a fiber failure can lead to a significant loss of bandwidth and revenue. For an IP-over-WDM network, two methods for providing protection are as follows: 1) provide protection at the WDM layer (i.e., set up a backup lightpath for every primary lightpath), or 2) provide restoration at the IP layer (i.e., overprovision the network so that, after a fiber cut, the network should still be able to carry the same amount of traffic as it was carrying before the fiber cut). (See, for example, [29] for a study of WDM protection versus IP restoration.)

We speculate some important tasks in designing IP-over-WDM networks to the following.

1) *IP/WDM Interface:* Develop a WDM interface that can work with today's IP (IPv4, IPv6, RSVP, etc.); such a solution may be suboptimal, but necessary if optical WDM is to be deployed in the Internet.
2) *Enhanced WDM Interfaces:* Develop an improved version of IP (which we call IPvWDM) that can exploit the characteristics of WDM; develop an improved version of TCP (TCPvWDM) that can allow some applications to "vertically cut-through" the IP layer if necessary; develop the appropriate NC&M interfaces including interfaces for legacy protocols such as frame relay, and for some applications to directly access the WDM layer (bypassing TCP and IP) if necessary; and support broadband local trunking, i.e., provide all-optical, "fat bit pipes" to high-capacity end-users.
3) *Interoperability:* Support multiple WDM backbones—each as an *autonomous system (AS)* with its own internal protocols—and support interoperability among them via *border gateways (BGs)*, using wavelength conversion, etc. Thus, many Internet concepts can carry over to the NGI.

F. Traffic Grooming in WDM Ring Networks

Traffic grooming is a term used to describe how different (low-speed) traffic streams are packed into higher-speed streams. In a WDM/SONET ring network, each wavelength can carry several lower-rate traffic streams in TDM fashion. The traffic demand, which is an integer multiple of the timeslot capacity, between any two nodes is established on several TDM virtual connections. A virtual connection needs to be added and dropped only at the two end nodes of the connection; as a result, the electronic add/drop multiplexors (ADMs) at intermediate nodes (if there are any) will electronically bypass this timeslot. Instead of having an ADM on every wavelength at every node, it may be possible to have some nodes on some wavelength where no add/drop is needed on any time slot; thus, the total number of ADMs in the networks (and hence the network cost) can be reduced. Under the static traffic pattern, the savings can be maximized by carefully packing the virtual connections into wavelengths, and quite a bit of work on this subject has been reported in the literature [42]–[48].

In [49], extension of the traffic-grooming problem in a single ring network is studied. Arbitrary (nonuniform) traffic is allowed and a formal mathematical definition of the problem is presented, which turns out to be an integer linear program (ILP). Then, a simulated-annealing-based heuristic algorithm is proposed for the case where all the traffic is carried on directly connected virtual connections (referred to as the "single-hop" case). Also studied is the case where a hub node is used to bridge traffic from different wavelengths (referred to as the multihop case). The simulated-annealing-based approach was found to achieve the best results in most cases relative to other comparable approaches proposed in the literature. In general, a multihop approach can achieve better equipment savings when the grooming ratio is large, but it consumes more bandwidth. In [49], traffic grooming in interconnected multi-WDM rings is examined. In typical telecommunication networks, multiple rings are interconnected together to provide large geographical coverage. The two problems of how to interconnect WDM ring networks as well as how to groom the traffic in interconnected rings are addressed in [49]. Not much research has been reported in this area so far. The work in [49] first proposes several WDM-ring interconnection strategies; then, it formulates the traffic-grooming problems for the proposed architectures. Two heuristic traffic-grooming algorithms are proposed for the cases of optically interconnected and hierarchical network architectures. It is found that the optically interconnected strategy achieves the best ADM savings; however, the hierarchical architecture is suitable for general physical network topologies and it provides moderate ADM savings. A pure digital-crossconnect scheme uses fewest wavelengths among all proposed interconnecting strategies, but it requires more network-wide ADMs.

G. *Optical Packet-Switched Networks*

As telecommunications and computer communications continue to converge, the data traffic is gradually exceeding the telephony traffic. This means that many of the existing circuit-switched networks will need to be upgraded to support packet-switched data traffic. While WDM has provided us an opportunity to multiply the network capacity, current optical switching technologies allow us to rapidly deliver the enormous bandwidth of WDM networks. Among all the switching schemes, photonic packet switching appears to be a strong candidate because of the high speed, data rate/data format transparency and configurability it offers. The goal of this area of research is to investigate critical issues involved in designing and implementing photonic packet-switched networks, and, through computer simulation and hardware experiments, find suitable solutions. Please see [50] for a tutorial on this subject.

The field of optical packet switching inevitably involves fast switch fabrics and other devices to facilitate a real-time, packet-by-packet switching functionality. However, for our purposes here, we will focus on network-level issues, such as architectures, network protocols, network control and management, etc. The fact that there is no available optical random-access memory (RAM) imposes a major difference between the design of an optical packet network and an electronic one. Instead of random-access memories, optical buffers (or delay-lines) are used to resolve contentions or adjust the packets' position in time. Research on this subject needs to search for ways to build a packet switched network which will offer significant flexibility without sacrificing the network utilization.

Contention resolution has a great impact on network performance in terms of packet-loss ratio, average packet delay, average hop distance, and network throughput. In a WDM optical packet-switched network, we have three domains to explore contention-resolution schemes: wavelength, space, and time. We can use any combination of wavelength conversion, path (space) deflection, and optical delay-lines to resolve contention. Each scheme has its own advantages and disadvantages: wavelength conversion is very efficient and able to resolve contention without introducing extra delay to the packet, but it is expensive to implement and no full-range wavelength converters are available today. Path deflection has the lowest cost since it shifts the burden of resolving contention to the whole network while lowering network overall throughput. Optical delay-lines (time buffering) has medium cost. It might introduce nonnegligible delay to the packet, depending on the packet length.

The first choice (wavelength conversion) achieves the same propagation delay and number of hops as the optimal case, and eliminates the difficulties in sequencing multiple packets. From this perspective, wavelength conversion is a very attractive solution compared to path deflection or time-buffering.

In limited delay buffering, a packet may be routed through a local fiber delay line and recirculated back into an input port of the same node. At that point, the header content will be read and routing will be attempted again. Because the packets may not be aligned, the choice of the length of the delay line may be arbitrary; however, there will be a tradeoff between the contention-resolution efficiency versus the minimum optical latency. A limited number of delay buffer lines may need to be incorporated in a node, and multiple wavelengths are accommodated in each delay buffer line.

With path deflection, it is important to prevent a packet from "wandering" in the network for a long time. The existing IP effectively utilizes the time-to-live (TTL, or maximum number of hops) field in the header to control the maximum number of hops an IP packet traverses. Similarly, deploying a TTL field in the optical packet header is an important consideration for loop-prevention and avoiding excessive path deflections.

Should optical packet-switched networks be synchronous or asynchronous? In synchronous packet-switched networks, all packets are of the same length, and all packets entering and leaving a switch are aligned. This improves the contention-resolution efficiency, but it requires additional cost in terms of hardware (packet synchronizers) and constraints (equal-length packets). These costs are absent from an asynchronous packet-switched network, which can allow packets of arbitrary lengths; however, the network's contention-resolution efficiency can be quite low. The synchronous or asynchronous tradeoff in optical packet-switched networks requires further careful investigation.

VIII. Conclusion

The objective of this paper was to provide an overview of the research and development work in the area of optical networking. Because of the large amount of activity on this subject over the past dozen years, capturing the important and most significant developments in a concise format is a challenging task; thus, I apologize if I inadvertently left out some contributions which you (the reader) would consider to be more important. This paper summarized the basic optical-networking approaches (focusing on functionalities of various devices and technologies rather than exact vendor implementations), reported on the WDM deployment strategies of two major U.S. carriers, discussed physical-layer impairments which may strongly influence network architectures, and outlined the current research and development trends on WDM optical networks, focusing mainly on the core (long-haul) wide-area network.

ACKNOWLEDGMENT

Much of the material in this paper is based on research that I have conducted over the years with my former and current Ph.D. students as well as with research scientists visiting my laboratory. Thus, I would like to acknowledge the following individuals for all their contributions: D. Banerjee, S. Banerjee, M. Borella, W. Cho, D. Datta, J. Iness, F. Jia, J. Jue, Y.-C. Kim, B. Ramamurthy, S. Ramamurthy, L. Sahasrabuddhe, S. Tridandapani, J. Wang, S. Yao, H. Zang, and K. Zhu.

REFERENCES

[1] B. Mukherjee, *Optical Communication Networks*. New York: McGraw-Hill, 1997.
[2] P. Cochrane, "Foreword," in *Optical Network Technology*. London: Chapman & Hall, 1995.
[3] *Ryan, Hankin, and Kent (RHK)*. http://www.rhk.com/.
[4] *WDM Forum* London, June 1998.
[5] *Special Issue on High-Capacity Optical Transport Networks, IEEE J. Select. Areas Commun.*, vol. 16, Sept. 1998.
[6] *Special Issue on Multiwavelength Optical Technology and Networks, IEEE/OSA J. Lightwave Technol.*, vol. 14, June 1996.
[7] *Special Issue on Optical Networks, IEEE J. Select. Areas Commun.*, vol. 14, June 1996.
[8] *Special Issue on Broadband Optical Networks, IEEE/OSA J. Lightwave Technol.*, vol. 11, May/June 1993.
[9] *Special Issue on Dense Wavelength Division Multiplexing Techniques for High Capacity and Multiple Access Communication Systems, IEEE J. Select. Areas Commun*, vol. 8, Aug. 1990.
[10] *Special Issue on WDM Networks, J. High-Speed Networks*, vol. 4, no. 1/2, 1995.
[11] S. Melle, C. P. Pfistner, and F. Diner, "Amplifier and multiplexing technologies expand network capacity," *Lightwave Mag.*, pp. 42–46, Dec. 1995.
[12] P. E. Green, "Optical networking update," *IEEE J. Select. Areas Commun.*, vol. 14, pp. 764–779, June 1996.
[13] ——, *Fiber Optic Networks*. Englewood Cliffs, NJ: Prentice-Hall, 1993.
[14] C. Fan, "Keynote address—Optical networking: A paradigm shift," in *WDM Forum*, London, June 1998.
[15] B. Mukherjee, "WDM-based local lightwave networks—Part I: Single-hop systems," *IEEE Network Mag.*, vol. 6, pp. 12–27, May 1992.
[16] ——, "WDM-based local lightwave networks—Part II: Multihop systems," *IEEE Network Mag.*, vol. 6, pp. 20–32, July 1992.
[17] R. Ramaswami and K. N. Sivarajan, "Optical routing and wavelength assignment in all-optical networks," *IEEE/ACM Trans. Networking*, vol. 3, pp. 489–500, Oct. 1995.
[18] S. Ramamurthy and B. Mukherjee, "Fixed alternate routing and wavelength conversion in wavelength-routed optical networks," in *Proc. IEEE Globecom'98*, Sydney, Australia, Nov. 1998.
[19] H. Zang, J. P. Jue, and B. Mukherjee, "A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks," *Opt. Networks Mag.*, vol. 1, no. 1, Jan. 2000.
[20] J. Nagel, "Prospects for optical networking," in *(D)WDM Forum*, Phoenix, AZ, Oct. 1998.
[21] S. Nathan, "Evolution to OXC based network: Impact of economic analysis," in *(D)WDM Forum*, Phoenix, AZ, Oct. 1998.
[22] R. E. Wagner and R. C. Alferness *et al.*, "MONET: Multiwavelength optical networking," *IEEE/OSA J. Lightwave Technol.*, vol. 14, pp. 1349–1355, June 1996.
[23] H. Zang, L. Sahasrabuddhe, J. P. Jue, S. Ramamurthy, and B. Mukherjee, "Connection management for wavelength-routed WDM networks," in *Proc. IEEE GLOBECOM'99*, Rio de Janeiro, Brazil, Dec. 1999, pp. 1428–1432.
[24] R. Ramaswami and A. Segall, "Distributeed network control for optical networks," *IEEE/ACM Trans. Networking*, vol. 5, pp. 936–943, Dec. 1997.
[25] J. J. Garcia-Luna-Aceves, "Distributed routing with labeled distances," in *Proc. IEEE Infocom'92*, Florence, Italy, May 1992.
[26] S. Ramamurthy and B. Mukherjee, "Survivable WDM mesh networks, Part I—Protection," in *Proc. IEEE Infocom'99*, New York, Mar. 1999, pp. 744–751.
[27] ——, "Survivable WDM mesh networks, Part II—Restoration," in *Proc. IEEE Int. Conf. Commun. (ICC'99)*, Vancouver, Canada, June 1999.
[28] S. Ramamurthy, "Optical design of WDM network architectures," Ph.D. dissertation, Univ. Calif., Davis, Dep. Computer Sci., Sept. 1998.
[29] L. H. Sahasrabuddhe, "Multicasting and fault tolerance in WDM optical networks," Ph.D. dissertation, Dep. Computer Sci., Univ. Calif., Davis, Nov. 1999.
[30] T. Wu, *Fiber Network Service Survivability*. Norwood, MA: Artech House, 1992.
[31] W. D. Grover, "The selfhealing network: A fast distributed restoration technique for networks using digital crossconnect machines," in *Proc. IEEE Globecom'87*, Tokyo, Japan, Nov. 1987, pp. 28.2.1–28.2.6.
[32] L. H. Sahasrabuddhe and B. Mukherjee, "Light-trees: Optical multicasting for improved performance in wavelength-routed networks," *IEEE Commun. Mag.*, vol. 37, pp. 67–73, Feb. 1999.
[33] ——, "Multicast routing algorithms and protocols: A tutorial," *IEEE Network*, vol. 14, pp. 90–102, Jan./Feb. 2000.
[34] R. M. Karp, "Reducibility among combinatorial problems," in *Complexity of Computer Computations*, R. M. Miller and R. M. Thatcher, Eds. New York: Plenum, 1972, pp. 85–103.
[35] R. A. Barry and P. A. Humblet, "Models of blocking probability in all-optical networks with and without wavelength changers," *IEEE J. Select. Areas Commun.*, vol. 14, pp. 858–867, June 1996.
[36] A. Birman, "Computing approximate blocking probabilities for a class of all-optical networks," *IEEE J. Select. Areas Commun.*, vol. 14, pp. 852–857, June 1996.
[37] B. Ramamurthy, D. Datta, H. Feng, J. P. Heritage, and B. Mukherjee, "Impact of transmission impairments on the teletraffic performance of wavelength-routed optical networks," *IEEE/OSA J. Lightwave Technol.*, vol. 17, pp. 1713–1723, Oct. 1999.
[38] C. S. Li *et al.*, "Gain equalization in metropolitan and wide area optical networks using optical amplifiers," in *Proc. IEEE Infocom'94*, Toronto, Canada, July 1994, pp. 130–137.
[39] B. Ramamurthy, J. Iness, and B. Mukherjee, "Optimizing amplifier placements in a multi-wavelength optical LAN/MAN: The equally-powered-wavelengths case," *IEEE/OSA J. Lightwave Technol.*, vol. 16, pp. 1560–1569, Sept. 1998.
[40] ——, "Optimizing amplifier placements in a multiwavelength optical LAN/MAN: The unequally-powered-wavelenghs case," *IEEE/ACM Trans. Networking*, vol. 6, pp. 755–767, Dec. 1998.
[41] B. Ramamurthy, D. Datta, H. Feng, J. P. Heritage, and B. Mukherjee, "Transparent vs. opaque vs. translucent wavelength-routed optical networks," in *Tech. Dig. Optical Fiber Commun. (OFC'99)*, San Diego, CA, Feb. 1999, Paper TuF2, pp. 59–61.
[42] O. Gerstel, P. Lin, and G. Sasaki, "Wavelength assignment in a WDM ring to minimize cost of embedded SONET rings," in *Proc. IEEE INFOCOM'98*, San Francisco, CA, Mar. 1998, pp. 94–101.
[43] O. Gerstel, R. Ramaswami, and G. Sasaki, "Cost effective traffic grooming in WDM rings," *Proc. IEEE INFOCOM'98*, pp. 69–77, Mar. 1998.
[44] A. L. Chiu and E. H. Modiano, "Reducing electronic multiplexing costs in unidirectional SONET/WDM ring networks vis efficient traffic grooming," in *Proc. IEEE GLOBECOM'98*, Sydney, Australia, Nov. 1998, pp. 322–327.
[45] J. M. Simmons, E. L. Goldstein, and A. A. M. Saleh, "On the value of wavelength-add/drop in WDM rings with uniform traffic," in *Proc. OFC'98*, San Jose, CA, Mar. 1998, (expanded version in J. Lightwave Technol., Jan. 1999), pp. 361–362.
[46] X. Zhang and C. Qiao, "An effective and comprehensive approach for traffic grooming and wavelength assignment in SONET/WDM rings," in *Proc. SPIE*, vol. 3531, 1998, pp. 221–232.
[47] W. Cho, J. Wang, B. Mukherjee, B. Moslehi, and R. J. Black, "Unification and extension of the traffic-grooming problem in WDM ring networks," in *Proc. OFC 2000 Conf.*, Baltimore, MD, Mar. 2000.
[48] J. Wang, W. Cho, V. R. Vermuri, and B. Mukherjee, "Improved approaches for cost-effective traffic grooming in WDM ring networks: Nonuniform-traffic case," in *Proc. IEEE ICC 2000 Conf.*, New Orleans, LA, June 2000.
[49] J. Wang, "Traffic grooming in WDM ring networks," Ph.D. dissertation, 2000.
[50] S. Yao, S. Dixit, and B. Mukherjee, "Advances in photonic packet switching: An overview," *IEEE Commun. Mag.*, vol. 38, pp. 84–94, Feb. 2000.

Biswanath Mukherjee (S'82–M'87) received the B.Tech. (Hons) degree from the Indian Institute of Technology, Kharagpur, India, in 1980 and the Ph.D. degree from the University of Washington, Seattle, in June 1987.

At the University of Washington, he held a GTE Teaching Fellowship and a General Electric Foundation Fellowship. In July 1987, he joined the University of California, Davis, where he has been a Professor of Computer Science since July 1995, and Chairman of Computer Science since September 1997. His research interests include lightwave networks, network security, and wireless networks.

Dr. Mukherjee is a corecipient of paper awards presented at the 1991 and the 1994 National Computer Security Conferences. He serves on the Editorial Boards of the IEEE/ACM TRANSACTIONS ON NETWORKING, IEEE NETWORK, *ACM/Baltzer Wireless Information Networks* (WINET), *Journal of High-Speed Networks*, *Photonic Network Communications*, and *Optical Network Magazine*. He also served as Editor-at-Large for optical networking and communications for the IEEE Communications Society. He served as the Technical Program Chair of the IEEE INFOCOM '96 Conference. He is the author of the textbook *Optical Communication Networks* (McGraw-Hill, 1997), which received the Association of American Publishers, Inc.'s 1997 Honorable Mention in Computer Science.