Guest Editorial
Multiwavelength Optical Technology and Networks

A BASIC property of single mode optical fiber is its enormous low-loss bandwidth of many terahertz (THz). Unfortunately, single channel transmission is limited in speed to much less than the fiber capacity due to limitations in optoelectronic component speed and dispersive effects. However, a popular and straightforward method of more fully utilizing the fiber bandwidth is to transmit several channels simultaneously on a single fiber, with each channel located on a different wavelength. Such wavelength division multiplexing (WDM) or multiwavelength networks not only enable significant capacity enhancements, but will also enable new networks in which the routing path is wavelength dependent. These networks offer enormous aggregate capacity and greater flexibility. The research challenges in WDM networks involve novel approaches to 1) enabling device technologies, 2) physical layer system implementations, and 3) network architectures and infrastructure issues.

Since the mid-1980's, considerable research efforts worldwide have been focused in these three areas. This initial flurry of work culminated in August 1990 with the first Special Issue on Wavelength Division Multiplexing, which was a joint publication of the IEEE/OSA JOURNAL OF LIGHTWAVE TECHNOLOGY (J-LT) and the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS (residing primarily with J-SAC). Since that first Special Issue, truly astounding progress has been made in multiwavelength optical communications. It is not uncommon to encounter results demonstrating the transmission of as many as 20 wavelengths, or WDM systems with channels each modulated at >10 Gb/s. There are now several possible multiple wavelength sources, each with its own advantages and disadvantages depending on the specific system's requirements. Furthermore, packet-switched WDM architectures and circuit-switched WDM network demonstrations have been highlighted at many conferences. This Joint J-LT/J-SAC Special Issue is the next attempt to coalesce many state-of-the-art contributions concerning WDM in one publication. Whereas the first Special Issue was under the J-SAC cover, this Special Issue is under the J-LT cover.

Recent widespread interest in information infrastructures has heightened interest in the high performance achievable in multiwavelength optical networks. WDM networks offer potential advantages, including higher aggregate bandwidth per fiber, new flexibility for automated network management and control, noise immunity, transparency to different data formats and protocols, low bit-error rates, and better network configurability and survivability—all leading to more cost effective networks. Several large consortia in Optical Networks are being funded around the world. In the U.S. they are the ARPA-funded consortia and the RACE program in Europe, which rely upon WDM as one of their base technologies. Due to the importance of optical networks, this special issue is a companion to the Special Joint Issue on Optical Networks of J-SAC/J-LT, which will be published in June 1996.

This special issue of J-LT includes a Foreword written by Dr. Charles A. Brackett of Bellcore, who provides a special perspective on the growth and evolution of WDM networks. In addition, we have eight invited papers from leaders in different aspects of multiwavelength networking. We received an enormous number of contributed papers in response to the Call for Papers, and we have included approximately 40 papers spanning all aspects of WDM networking.

The organization of this special issue is divided into three broad sections encompassing the enabling component technologies, wavelength routing: systems and protocols, and networking demonstrations and infrastructure issues. There is an approximately even representation of papers among these areas. A brief description of some of the areas which are covered by the papers is as follows. 1) Enabling technologies for multiwavelength networks: novel wavelength tunable and selectable sources and receivers, wavelength selective components, wavelength routing components, wavelength translation techniques, and optical amplifiers for multiwavelength applications. 2) Wavelength routing: WDM crossconnects, packet switching approaches, network configuration issues, WDM network architectures, wavelength assignment, and management and control protocols for the WDM layer. 3) WDM networking demonstrations and infrastructure issues: performance of WDM networking testbeds, advantages of format transparency, reference wavelength and wavelength registration techniques for networking, performance limits and bandwidth management in WDM networks, potential use of WDM technology in personal communications systems, network control and management and operational issues associated with configuration at the WDM layer, and WDM networking within the National Information Infrastructure.

It is a pleasure to acknowledge the efforts of numerous reviewers for this issue who did a superb job in reviewing the contributed papers and maintaining a high standard in the selection process. Special thanks are due to Laura VanSavage

Publisher Item Identifier S 0733-8724(96)05619-8.

of the IEEE office for her efforts in producing this issue. We also thank the Editors of J-LT and J-SAC for their willingness to support this project.

We have attempted to provide a special issue of lasting value to the optics, communications, networking, and quantum electronics communities. We hope that the quality and quantity of the papers reflect the rapid development of this field. With this intense activity, it is almost certain that another special issue on WDM will appear in the next five years. Now, read and enjoy this stimulating issue!

MASAHIKO FUJIWARA, *Guest Editor*
NEC Corporation
MATTHEW S. GOODMAN, *Guest Editor*
Bellcore
MICHAEL J. O'MAHONY, *Guest Editor*
University of Essex
OZAN K. TONGUZ, *Guest Editor*
SUNY at Buffalo
ALAN E. WILLNER, *Guest Editor*
University of Southern California

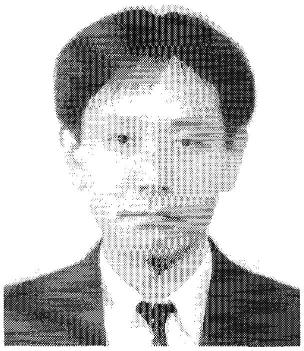

Masahiko Fujiwara (M'91) was born in Tokyo, Japan, on February 18, 1953. He received the B.S., M.S., and Ph.D. degrees in electrical engineering all from Yokohama National University, Kanagawa, Japan, in 1975, 1977, and 1993, respectively.

From 1987 to 1994, he was with Opto-Electronics research Laboratories, NEC Corporation, Kawasaki, Kanagawa, Japan. His major technical activities there were in the research and development of optical information processing systems using laser diodes, semiconductor optical switches/integrated optics, and photonic switching networks mainly using photonic space division and wavelength division switching technologies. In 1994, he moved to Video Development Division, NEC Corporation, Fuchu, Tokyo, Japan, and is now a senior manager of the Third Development Department. He is currently working on the development of commercial WDM networks for a TV broadcast center application.

Dr. Fujiwara is a member of the Laser and Electro-Optic Society (LEOS) and the Institute of Electronics, Information and Communication Engineers Japan.

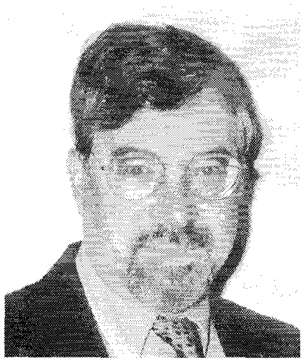

Matthew S. Goodman (M'90) received the B.S. degree in physics (cum laude) from Indiana University in Bloomington in 1969. His graduate studies were at the Johns Hopkins University in Baltimore, MD, where he received the M.A. and Ph.D. degrees in physics in 1971 and 1977, respectively.

From 1977 to 1985, he was a member of the faculty of the Department of Physics at Harvard University. During this period, his research was in the area of elementary particle physics and he performed experiments including work at the SLAC, Stanford, CA, Fermilab, Batavia, IL, Oak Ridge National Laboratory, Oak Ridge, TN, and the CERN Laboratories, Geneva, Switzerland. Since joining Bellcore in 1985, he has been working in the areas of multiwavelength optical network research, multiwavelength optical network architectures, optical networks for local access, and technology assessment. His research work has included some of the earliest multiwavelength network experiments, multiwavelength network designs and dense multiwavelength applications, including applications of optical networks to computer interconnects and optical packet switching. He was the coordinator of the Phase II ONTC (Optical Networks Technology Consortium) program on reconfigurable multiwavelength optical networks, partially funded by ARPA. He is currently working on the MONET (Multiwavelength Optical Networking) program. He holds four U.S. patents and is the author of more than 65 research papers on physics and technology and several invited book chapters.

Dr. Goodman is the recipient of numerous awards including a Harvard Graduate Society Award and the Thomas T. Hoopes III Teaching Award in 1984. In 1994, he was a recipient of the Bellcore President's Award. He is a member of Sigma Pi Sigma, the Society of Sigma Xi and the IEEE Communications Society (COMSOC), and is currently the Area Editor for Optical Communications of the IEEE TRANSACTIONS ON COMMUNICATIONS. He has chaired and served on numerous professional organizations and committees and is currently secretary of the IEEE COMSOC Technical Affairs Committee.

Michael J. O'Mahony (SM'95) was born on August 28, 1944. He graduated from Essex University, Colchester, England, in 1974 and received the Ph.D. degree for research on digital transmission systems in 1977.

In 1979, he joined the Optical System Research Division of British Telecom working on research into fiber-optic systems for undersea systems; in particular, experimental and theoretical studies of receiver and transmitter design. In 1984, he became a Group Leader responsible for the study and application of optical amplifiers to transmission systems. In 1988, he became a Head of Section responsible for inland long haul optical systems and networks. His areas of interest included optical amplifiers, coherent optics, picosecond pulse systems and optical networks. In 1991, he joined the Department of Electronic Systems Engineering at the University of Essex as Professor of Communication Networks and Director of the Network Research Centre. He is also Head of the Photonic Networks Laboratory, which is focused on the study and design of photonic networks, the main current area of interest is in understanding the physical limitations associated with the realization of optical transport and access networks. He is the author of more than 120 papers relating to optical transmission.

Dr. O'Mahony is a member of the IEE.

Ozan K. Tonguz (S'86–M'90) was born in Nicosia, Cyprus, in May 1960. He received the B.Sc. degree in electronic engineering from the University of Essex, Essex, England, in 1980; and the M.Sc. and Ph.D. degrees in electrical engineering from Rutgers University, New Brunswick, NJ, in 1986 and 1990, respectively.

In 1981, he returned to Cyprus. After two years of mandatory military service, he was an Assistant Lecturer at the Eastern Mediterranean University of Northern Cyprus, during the academic year of 1983–1984. In September 1984, he joined the Department of Electrical and Computer Engineering, Rutgers University. Between January 1988–May 1990, he was a visiting doctoral student at the Advanced Lightwave Systems Division of Bell Communications Research, Red Bank, NJ, where he conducted research on coherent lightwave technology. From May–August 1990, he was a Member of Technical Staff at Bellcore, and continued to do work on coherent lightwave systems technology and optical amplifiers. He joined the Department of Electrical and Computer Engineering, State University of New York at Buffalo (SUNY/Buffalo), Buffalo, NY, as an Assistant Professor in September 1990, where he was granted early tenure and promoted to the rank of Associate Professor in June 1995. At SUNY/Buffalo, he leads substantial research activity in the broad area of telecommunications. His current research interests are in mobile radio and personal communication systems (PCS), fiber optic communication systems and networks, and high-speed networking. He has published in the areas of optical communication systems and networks, mobile radio and personal communication systems and networks, and is the author or co-author of more than 50 technical papers in leading technical journals and conference proceedings. While his research on optical networks has been sponsored by National Science Foundation via a Research Initiation Award, his research in wireless communications is currently being funded by several companies active in the mobile radio and PCS industry.

Dr. Tonguz frequently acts as a reviewer for various IEEE and IEE TRANSACTIONS and JOURNALS, has served on the Technical Program Committees of the IEEE Lasers and Electro-Optics Society (LEOS), and has chaired technical sessions in the IEEE International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC'95), IEEE 46th Vehicular Technology Conference (VTC'96), and IEEE LEOS Annual Meeting (LEOS'95). He is a member of the Optical Society of America and Eta Kappa Nu.

Alan E. Willner (S'87–M'92–SM'93) received the B.A. degeee in physics from Yeshiva University in 1982 and the M.S. and Ph.D. degrees in electrical engineering from Columbia University in 1984 and 1988, respectively.

He was a Postdoctoral Member of the Technical Staff at AT&T Bell Laboratories (Crawford Hill) and a Member of Technical Staff at Bellcore. He is currently an Associate Professor in the Department of Electrical Engineering—Systems at the University of Southern California. His research is in high-capacity optical communication systems, specifically wavelength division multiplexed optical systems and networks. He has coauthored more than 125 journal and conference publications, including 28 invited talks, two book chapters, and two short courses.

Dr. Willner has received the NSF Presidential Faculty Fellows Award from the White House, the Packard Foundation Fellowship in Science and Engineering, the NSF Young Investigator Award, the USC Outstanding Junior Engineering Faculty Research Award, the USC Outstanding Engineering Teacher Award, and the Armstrong Foundation Memorial Prize for the highest-ranked EE graduate student at Columbia University. He was also an NSF Alan T. Waterman Award Finalist. He is a Fellow of the Semiconductor Research Corporation. He is the Vice-President for Technical Affairs for the IEEE Lasers and Electro-Optics Society (LEOS), Chair of the Optical Communications Technical Committee for IEEE LEOS, a member of the Board of Governors for IEEE LEOS, and the Vice-Chair of the Optical Communications Group for the Optical Society of America (OSA). He was a General Co-Chair of the IEEE LEOS '95 Summer Topical Meeting on Technologies for a Global Information Infrastructure and Chair of the Optical Communications Subcommittee for the 1994 and 1996 LEOS Annual Meeting. He is a member of the Program Committee for the 1996 and 1997 Optical Fiber Communications Conference, and is a Program Committee Member for the Optical Amplifier Conference. He is an Associate Editor of the IEEE/OSA JOURNAL OF LIGHTWAVE TECHNOLOGY (J-LT), and a Guest Editor for a Special Issue on Multi-Wavelength Technologies and Networks of the JOURNAL OF LIGHTWAVE TECHNOLOGY. He is in *Marquis' Who's Who in America.*

Foreword
Is There an Emerging Consensus on WDM Networking?

IN the last several years, there has been a growing excitement among those working in the area of wavelength division multiplexing (WDM) and its related applications, such as WDM networking. This was particularly evident this year at the Optical Fiber Communication Conference, (OFC'96) held in San Jose, CA, at the end of this February, the Photonics in Switching Topical Meeting held in Salt Lake City, UT, in the middle of March 1995, and the European Conference on Optical Communication (ECOC '95) held in Brussels, Belgium, in October 1995. The excitement is that this field may actually be approaching some degree of commercial applicability as central telephone administrations and network operators around the world appear to recognize their need for tremendous network capacity growth due to the expected traffic demands for video and multimedia services. WDM transport technology is expected by many to play a significant role in helping to achieve the needed bandwidth, starting in the very near future. Indeed, WDM amplified network systems products are being marketed by at least one major telecommunications equipment supplier, and more are expected soon.There are, however, other projected applications of WDM technology beyond the point-to-point transmission link, for example, in transparent optical networking and in all-optical switch fabrics to name two, and it is this writer's purpose to propose here that a consensus is emerging as to the likely and appropriate applications for these approaches. In order to make this consensus more obvious, I will first try to trace the historical development of the concepts of multiwavelength networking and then discuss what I see as the probable course of development of this field, along with the reasons why. This is a highly personal view (for which I offer no apology), but the reader should be aware that other views do exist which may be more accurate in the last analysis. I also restrict myself to the single case of multiwavelength networks rather than try to discuss the much broader class of all-optical networks in general, which include such things as very high-speed time domain techniques, all-optical time-domain switching, optical code-division multiplexed networking, etc. Each of these may find their niche, but I believe the time is almost here for the multiwavelength case, and that it will achieve significant commercial importance and is therefore of special interest.

I. CONCEPTUAL DEVELOPMENT OF OPTICAL NETWORKING

Optical networking began with a few very simple and basic concepts, and has evolved dramatically toward solving the real world problems of building large-scale networks that are robust against failure and traffic surges and that will evolve smoothly with time and growth. The following is this author's view of the principle conceptual stepping stones that have been at the heart of the progress made to date. Clearly, there are others, not to mention the technological steps that have been achieved which make this all possible. The purpose here is to remind ourselves how little we knew just a short time ago, and to suggest that there are probably as many more steps to go before we have finished the job.

A. *Point-to-Point Multiwavelength Transmission (WDM)*

This is the traditional wavelength division multiplexing, WDM, that was introduced in the early 1970's, in which several optical signals at different wavelengths are multiplexed together onto a single fiber in order to increase the capacity of a link between two distinct points. Early WDM systems did not achieve much commercial significance because it turned out that it was less costly to increase the speed of time-division multiplexing to achieve the same capacity.

That was before the erbium-doped fiber amplifier (EDFA) became practical in providing efficient, low noise, and broadband gain in the 1.55 μm low-loss fiber band. Once the EDFA became practical, it fundamentally altered the perceived economics of transmission systems because the EDFA could amplify more than one wavelength at a time, thereby replacing one regenerator for each channel. The primary applications of this point-to-point WDM technology are expected now to include increasing the capacity of long-distance transmission systems and solving route-exhaust problems in metropolitan networks.

There is still much work remaining on this problem and commercialization is just beginning.

B. *Broadcast-and-Select Networks*

Point-to-point WDM was interesting, but there soon evolved proposals to try to implement in the optical domain some of the functions that had previously been done in the electronic domain. Early (*single-wavelength*) optical networks based on passive broadcast stars had been proposed in which each user on the network transmitted its signals into a broadcast star coupler which was used to distribute those signals passively to all other nodes on the network. A media-access protocol was required to control the transmissions of the various network nodes to avoid collisions and manage contention for the network bandwidth.

Publisher Item Identifier S 0733-8724(96)04673-7.

The advent of the distributed feedback (DFB) laser provided the ability to generate well-defined single-wavelength optical signals that could be directly modulated at high speeds, and several network demonstrations showed the feasibility of multiwavelength broadcast-and-select networks. The essential feature was that if there were N wavelengths available, then N simultaneous transmissions could be carried out in a single broadcast-star network, with each end station selecting the wavelength destined for it. This use of WDM would produce an N-fold concurrency and thereby dramatically increase the total capacity of the network. If the number of nodes on the network were less than or equal to the number of available wavelength channels, then a completely connected network could be achieved on a simple star network. With enough wavelengths, the use of WDM even eliminated the contention problem.

It became obvious that one of the potential attractions of these networks was their "transparency" to signals of different modulation formats. These were "single-hop" networks, in which the signals did not go through any intermediate electronic functions on their way from source node to destination node. The format of the signal was determined only by the transmitting and receiving equipment of the various nodes, with the optical routing being accomplished simply by broadcasting all signals everywhere, and each receiving node using an optical filter in the receiving process. Different transmission speeds and formats could exist simultaneously in the same star network.

The principal limitations of these networks were that they were not scalable to large numbers of nodes because there was a linear relationship between the number of nodes and the number of wavelengths. Proposals were made for using very large numbers of wavelengths to create very large interconnection networks, but practical limitations on crosstalk and filter resolution, as well as the impracticality of administering a network with a very large number of differing lasers and the general requirements on high-speed tunable filters, restricted these WDM broadcast-and-select networks to numbers of nodes on the order of 16 to 32.

Work in this area is still underway for interconnection of computers in local and metropolitan area networks, and such systems tend to operate at the moment in a circuit switched or burst mode. Application to large-scale networks is however not feasible due to their lack of graceful scaling.

C. Broadcast-Star Switching Fabrics

The ease and simplicity of broadcast-and-select WDM networks brought about several proposals for various kinds of broadband switching fabrics, as would be found within a high-speed digital switch. One goal has been to design very high-capacity asynchronous transfer mode (ATM) switches that utilized the natural broadcast connectivity of optical star couplers. This would simplify the switch-fabric interconnection network, which is one of the principal limitations in building large switches.

Another, more recent direction has been to try to utilize the wavelength domain as the third dimension in space-time-wavelength switching fabrics. These designs are aimed at reducing the complexity in the time and space domains (number of equivalent ports) while extending the overall capacity of the switch toward the Terabit/s range.

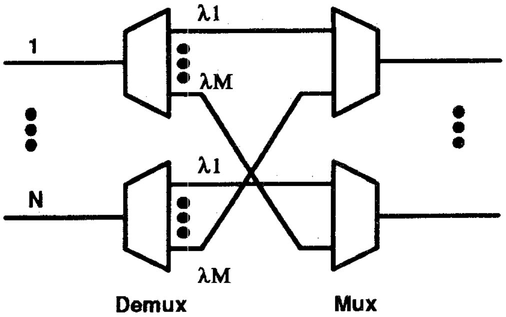

Fig. 1. Schematic of a fixed WDM cross-connect with N input-output port pairs and M wavelengths.

These "all-optical ATM switches" suffer from the problem that all inputs and outputs need to be synchronized, and the logical operations inherent in ATM switching such as header replacement need to be implemented. Heroic attempts have been made to achieve these functions with multiwavelength delay-line switching. However, lacking a compact and low power optical memory technology, such logical and memory operations in switching have a very difficult time in competing with electronics.

D. Wavelength Routing

Wavelength routing is defined to be the selective routing of optical signals according to their wavelengths as they travel through the network elements between source and destination. There are two salient features of wavelength routing in optical networks.

First, wavelength routing determines the path taken by the optical signal, and if multiple signals are launched from a given node, each may go to a separate distinct destination. The number of such destinations is equal to the number of wavelengths generated at each node.

The second feature is that because each signal is restricted to a particular path, it is possible to have each wavelength reused many times on different paths throughout the network as long as these other paths do not try to coexist on the same fiber link.

Wavelength routing is achieved by implementing some form of wavelength-selective elements at the nodes of the fiber network. Fixed wavelength routing would most likely use WDM multiplexers in a back-to-back configuration to allow interchange of wavelengths between input and output fibers in a prearranged pattern. This configuration (shown in Fig. 1) has been called a WDM cross-connect, and in its simplest form does not have any automated rearrangeability.

Rearrangeability is introduced by adding space division switches, as shown in Fig. 2. Using such WDM cross-connects, each wavelength on any input fiber can be interconnected to any output fiber providing that output fiber is not already using that wavelength. The total cross-connect

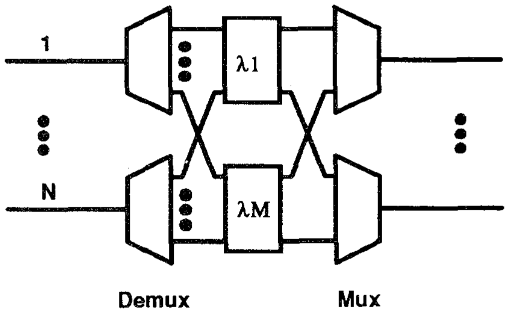

Fig. 2. Schematic of a rearrangeable WDM cross-connect using space division switches.

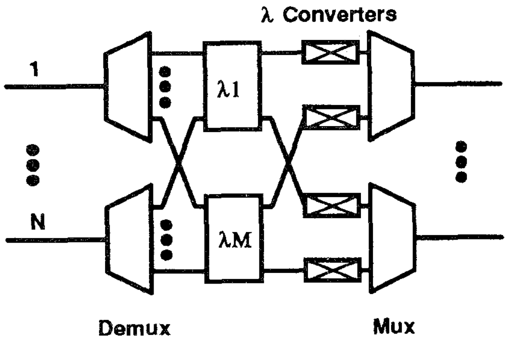

Fig. 3. A schematic of a wavelength interchanging cross-connect using wavelength converters with fixed-wavelength outputs and variable inputs.

bandwidth is proportional to $N \bullet M \bullet B$, where N is the number of input fibers, M is the number of wavelengths, and B is the bit rate per wavelength.

Both of the above cross-connects can be termed wavelength-selective cross-connects, because they select wavelengths and rearrange them in the spatial domain.

A third type of WDM cross-connect has recently been defined to allow cross-connection in the wavelength domain as well as the space domain, and has been called a wavelength-interchanging cross-connect, one form of which is shown in Fig. 3.

Such a wavelength interchanging cross-connect allows any input wavelength on any input fiber to be cross-connected with any output wavelength on any output fiber.

The importance of the WDM cross-connect switch, and the closely related WDM add-drop multiplexer, is that they allow the optical network to be reconfigured on a wavelength-by-wavelength basis to optimize traffic, congestion, network growth, and survivability. They also permit the configuration of special circuits for transmission of alternative format signals. The WDM cross-connect switch and the WDM add-drop multiplexer are the essential wavelength-selective format transparent elements upon which multiwavelength networks will be built.

E. Scalability

As has been noted above, networks based on the broadcast-and-select principle are not scalable. The reason for this is that while a wavelength is being used for communication between one pair of nodes, that wavelength may not be used for communication between any other nodes without causing interference.

What was needed was a way of reusing wavelengths throughout a network, similar to frequency reuse in a cellular radio network. Wavelength reuse through the use of wavelength routing was introduced to solve that problem, but there is currently an active debate as to whether wavelength routing is sufficient to create a scalable network.

Scalability needs a definition. In electronic switching fabrics, scalability is generally considered to mean that the number of crosspoints increases with the number of input ports according to something like an $N \log N$ dependence. This is sufficient for a space-division switching fabric because large switches can be built and because one never really runs out of physical space.

In WDM networks, however, the problem is different. There are only a finite number of wavelengths, no matter how closely spaced in wavelength they are. Original optimistic estimates were that thousands of independent wavelengths could be used. In reality, it is very likely that both practical and fundamental limits will restrict the number of wavelengths to very small numbers such as 8, 16, or 32, depending on the specific application.

Scalability then, has a completely different interpretation for WDM networks. One definition is that one must always be able to add more nodes to the network, no matter how large it already is. If the number of wavelengths is eventually limited, the implication is that ultimately the number of nodes on the network must be completely independent of the number of wavelengths. Wavelength routing alone is not sufficient to allow that in an all-optical network.

However, that level of scalability can be achieved in multihop networks. In single-hop networks, an optical signal travels from the point of origination to the destination without encountering electronic regeneration. However, in a wavelength routed network, each node can only transmit in a single hop to as many nodes as it has wavelengths. To go further requires that the signal be switched to a new wavelength-path, which may be on a different wavelength, thereby requiring wavelength conversion. This effectively requires both space and wavelength switching, in which not only are signals switched to different output space ports, but where they are also switched from one wavelength to another. This can be performed on a circuit switched basis, or on a packet switched basis, depending on the type of switching protocol being used. If the switching from one wavelength to another can be done, and controlled, all-optically, then a transparent all-optical network can be created that achieves this "true" scalability. At the present time, the technology for doing such switching on time scales associated with ATM cells is not practical and the only achievable cell-switched scalable WDM network is one that utilizes electronic ATM switches at each network access node. This requires detection of each incoming wavelength, switching on a cell-by-cell basis in the ATM switch, and retransmission on the appropriate output wavelengths.

Such a network has been demonstrated using multiwavelength array lasers and WDM cross-connect switches. Such networks have a mixture of single-hop and multihop optical paths. By configuring the WDM switches, the single-hop paths can be utilized to serve those connections which require single-hop transparent connections. Other traffic must be carried on multihop connections, and network optimization can be achieved by the appropriate choices for the single-hop paths.

The view that is evolving for these types of networks is that the all-optical, WDM portion of the network needs to be rearrangeable but not on a call-by-call, or cell-by-cell basis. The inner all-optical portion of the network will be reconfigured to meet overall network traffic demand and growth needs, which occurs much less frequently. This is much easier to do than to reconfigure rapidly and frequently because of the difficulty of synchronization and management across an extended network.

The above scenario is not the only view of scalable WDM networks. A second view is that all network connections must be done on a single-hop basis. The implication of this is that only a limited number of connections can be established at a given time and the network must be somehow reconfigured several times to accept a new set of connections. In this scheme, one can theoretically bound the number of wavelengths required and it is roughly proportional to the square root of the number of nodes. The difficulty with this type of design, besides the limitations to the number of wavelengths, is the required synchronization of the reconfiguration of the network, or the tuning states of the lasers and receivers, across the entire network. It would seem (to this author) that such rapid reconfiguration would be limited to networks of modest size.

Another critical remark about this form of network is that although the stated objective is to achieve transparency, the transmissions may need to be scheduled into time slots in frames. This would not appear to achieve the original goal of transparency any more than the need to go through a multihop arrangement.

Once the assumption is made about whether a multihop or single-hop network is desired, or allowable, the choice between these two approaches will already have been made.

F. Wavelength Translation

There is a significant debate in progress about the necessity of translating the wavelength of a signal within a network. A distinction has been made in the literature between a wavelength path and a virtual wavelength path, depending on whether the signal stays on the same wavelength or is converted to another in midstream. Various calculations of the blocking probability have been produced and the conclusions regarding the necessity of wavelength conversion depend significantly upon the assumptions made (but not always stated).

What does seem clear is that there is a significant difference between what is necessary in a small local network, and what is necessary in a large-scale regional or national network. What is theoretically possible may not in fact be manageable.

Current telecommunications networks are divided up into regional administrative domains with simplified network interfaces in order to solve the problem of complexity. It is not considered practical to have current and complete network knowledge in a centralized location in a large scale network. Instead, each domain interacts with its neighbors to request call setup, for example, or to isolate faults, without knowing the details of its neighbors' network connectivity, etc.

In completely analogous reasoning, it is not likely to be feasible to set up end-to-end transparent paths on a single wavelength across multiple administrative domains in an efficient and robust manner. Because of this, it is very likely that wavelength conversion will take place, at least at the boundaries of administrative domains, just to lower the complexity of the network connection setup, if nothing else.

G. Transparency

Transparency is one of those concepts that is elusive. Every time one considers a simple definition, another limitation arises. The original idea was to have a network into which one could launch any optical signal and nothing in the network would interfere with it before it reached its destination, and therefore only the terminal equipment would determine the limitations on signal format, etc.

However, there are limits to almost everything, many of which depend on the physical properties of the optical signal being transported and which therefore represent a limitation away from strict transparency.

As a practical matter, the most transparent optical system imaginable of any usefulness is a very short piece of optical fiber. It has essentially zero loss, it has essentially zero total dispersion, it has essentially no power limitations or modulation format restrictions, and it introduces no additional noise to the signal. It does require the signal to have the correct beam-shape for effective launching.

However, for longer lengths of fiber, one encounters loss which decreases the signal-to-noise ratio at the detector, and that SNR impact is a direct function of the bit rate of a digital signal and of the modulation format of an analog signal. So, after some critical transmission distance, some bit rates will not be detectable whereas others may, and strict transparency is lost.

The same can be said about the effects of dispersion with regard to high speed signals. And in amplified WDM systems, fiber nonlinearities and power levels also limit the achievable transparency and independence from signal parameters. And crosstalk in the wavelength selective components imposes severe restrictions on wavelength spacings allowable.

Finally, in a WDM system, it will be necessary to specify the standard wavelengths of the system in order to achieve interoperability and volume production. Transparency loses again.

Therefore, when asking for a transparent optical network, one must take into account the very real limitations of the physical medium, some of which are of a fundamental nature and therefore not easily avoided. What will work in a local network environment will not necessarily work in a national scale network. It is also to understand why transparency is desired in particular applications. Transparency imposes certain costs; is transparency really needed?

The final assessment will be to place a value on transparency. The question will be: is the value of end-to-end transparency worth the cost?

H. Network-Layering

As more and more thought was given to actually constructing working models of real WDM networks, it became apparent that WDM was really adding another layer to present network abstractions: the optical connectivity layer.

Imagine that current networks have a physical fiber layer at the bottom of the system, a transport layer which defines the actual transport of information, and a switching layer which organizes the flow of that information from each source to its destination. An example of those three layers would be fiber, SONET, and ATM. Now we are adding a new layer. The physical fiber has a topology which is determined by where it lies in the ground and where it begins and ends.

But, by adding the possibility of WDM cross-connects at the physical nodes of the network, we can introduce an optical connectivity which is different than the fiber connectivity. This becomes a new layer between the fiber and transport layers. The transport layer (SONET) thinks it is connected to fiber, but what it is really connected to is an optical connection layer that can be configured in a wide variety of ways.

This concept allows the extension of current networking directions and trends without having to displace what is already there. There is one important exception and that is the problem of network management and operations in transparent and rearrangeable systems.

I. Network Management, Control, and Operations in Transparent Systems

There are two problems introduced by the use of transparent networks. The first is that means must be developed for monitoring the state of the network since in a transparent network the normal digital information about network performance is not available. The second problem is that current transport systems, such as SONET and SDH, have well defined internal means for dealing with faults and performance monitoring, and whatever is done at the optical layer must work together with the transport layers. If there is a fault, for example, both the SONET layer and the optical layer will know about it, and means must be implemented to coordinate their response. This is a topic of current work and includes the subject of how to carry the network control information.

In addition, in transparent optical networks, it is likely that diverse transport systems may share the same medium (FDDI and SONET, for example). Each of these different systems will have some internal means for dealing with faults, etc. An additional challenge in network management for transparent systems is to have an integrated system which works across different systems, as well as across the different system layers.

Both of these problems have solutions but they tend to complicate matters over what would be required for smaller-scale networks.

Network management and control is probably a greater challenge at this point than are the technologies because the implementation of WDM and optical networking will not take place until satisfactory management systems are in place.

II. AN EMERGING CONSENSUS?

Is WDM going to happen? Where in the network will it have its most important impact? How does the cost effectiveness of WDM compare with that of higher speed TDM? Is there a cost effectiveness associated with the flexibility of transparent networking? What is the evolutionary path from today's electronic networks to tomorrow's optical ones? What needs to be done to bring this about?

Is there a consensus emerging about the future of WDM and optical networking?

These and many more questions are prevalent in everyone's minds as the progress of this technology is followed and evaluated. While not having answers to many of these questions, there does seem to be some consensus that is emerging which are summarized in a few brief comments below.

- WDM for point-to-point applications is progressing much faster than many have estimated. Several commercial suppliers are selling and installing systems now (late 1995), and several more are expected to begin in 1996. Early systems have been four-wavelength systems running OC-48 for an OC-192 equivalent. These systems are expected to be extended to eight-wavelengths in 1996.
  WDM is effectively here now and growing rapidly.
- The first applications appear quite naturally in long-distance routes where the savings in regenerators by using amplifiers on multiwavelength signals has a pronounced effect. In addition, there are many administrations where there are routes that are exhausted in both fiber capacity and in duct space, and WDM systems on a point-to-point basis for capacity upgrade may be cost effective.
- Various multimedia and video services are expected to drive the need for increased network capacity dramatically. WDM is likely to be one of the principle ways of accommodating that increase in need.
- Many optical networking programs around the world have made tremendous progress in identifying the issues and finding solutions. In particular, several demonstrations of networking on a small scale have shown that the technology is available and capable, except for the network management functions which require further development.
- Laser arrays, wavelength selective switching elements, cascaded WDM amplifiers, and power management are the technologies that have been investigated the most and that have been proven feasible for the scale of networks investigated to date.
  Wavelength conversion is under intensive investigation but a clear winner has not yet emerged.
- WDM networking for telecommunications has universally arrived at the layered architecture approach, where the optical connectivity layer is a sublayer of the physical layer and contributes to the overall functionality of additional transport and switching layers riding on top of it.

- Not all connections need to be "single-hop." This is a performance and service issue, but the multihop architecture does give a much larger degree of scalability in which the number of network access nodes and the number of wavelengths are independent of each other.
- Optical networking is still being pursued mostly as a core backbone network capability. Whether it will find application in the subscriber loop will depend on the benefit-to-cost ratio. Most of the systems and components developed for telecommunications optical networks are probably far too complicated to have low enough costs for subscriber loop application.
- One of the most pressing needs to help stimulate commercial success of WDM and optical networks in general is the need for standards. It is essential that standards be introduced in a timely manner, and that they be structured in such a way that the standards for WDM do not force commercial directions which are not advantageous to more general optical networking directions, and in a way that encourages a fundamental reduction in the cost of bandwidth.

III. CONCLUDING COMMENTS

The promise of optical communications is to provide so much bandwidth that the cost of bandwidth decreases significantly, with benefits to the end users and society in general. Whoever captures the lead in driving down the cost of user bandwidth will be in a dominant position to profit from the coming of the information age. Network bandwidth is increasingly becoming an extension of our computers, our entertainment, our medicine, and our commerce. Network bandwidth will soon become as important to the progress of computing as was the progress in random access memory of a few years ago. The pressure to increase the available bandwidth to end-users is being driven by the ubiquitous availability of personal computers and the public's recognition that connecting these computers together and to information resources is an essential proposition.

The major challenge to all those working in optical communications, communications in general, and in optical networking in particular, is to find the most practical way of driving down the cost of that bandwidth.

CHARLES A. BRACKETT
Bellcore
Red Bank, NJ 07701 USA

Charles A. Brackett (S'60–M'64–SM'89–F'91) received the Ph.D. degree in electrical engineering from the University of Michigan in 1968.

He is Executive Director for Optical Networking Research at Bellcore, where he has led the optical networking activities since 1984. He organized and led the Optical Networks Technology Consortium (ONTC) which demonstrated a multihop transparent WDM network using advanced laser array technology and wavelength selective cross-connect switches. He has also been the Bellcore program manager for the MONET Consortium and acted as Chair of its Technical Management Committee. Prior to Bellcore, he was responsible for optical receiver development at AT&T Bell Laboratories.

Dr. Brackett is a member of the Optical Society of America.