CHAPTER 1:
DEFINING ALL-OPTICAL NETWORKS

As the name suggests, All-Optical Networks (AONs) form a class of networks where optical technology (rather than electronics) plays the major role in network functionality. Before defining the various types of AONs, it is instructive to recall that today's "traditional" networks consist, for the most part, of a collection of electronic switches interconnected by *point-to-point* optical fiber links, and spanning local to national (or even global) distances. To accommodate continuously increasing demand for bandwidth and flexibility, such networks are being enhanced by adding more fibers and switches, by increasing the bit rate per fiber (e.g., to 2.5- or 10-Gb/s SONET rates), and by upgrading the switches' size, throughput and functionality (e.g., from circuit switching to various forms of packet switching such as ATM and FDDI). It is anticipated that such enhancements will eventually lead to very large and complex networks that are difficult and expensive to construct, operate and maintain. Recent advances and envisioned future capabilities of optical technology promise novel all-optical networks capable of providing improved economy, flexibility and robustness, while still capable of making use of the embedded fiber plant.

There are three general types of AONs, which, from a functionality point of view, can be classified as follows: (1) Passive Optical Networks (PONs); (2) Transparent AONs; and (3) Ultra-High-Speed AONs, each of which will now be discussed in more detail.

PONs, traditionally, refer to all-optical networks utilizing only *passive* optical components (e.g., fibers, directional couplers, star couplers, wavelength routers and multiplexers, and filters). The emphasis in such networks is on low cost and low (or no) maintenance, and the intended applications are fiber-in-the-loop and fiber-to-the-home. Thus, economy and the "bury it and forget it" philosophy are prime objectives, even with a sacrifice in performance. Occasionally, optical amplifiers, particularly those of the erbium-doped fiber types, are employed to overcome propagation and distribution losses in PONs, even though they are clearly not passive. The optical signaling formats in PONs can employ wavelength-division multiplexing (WDM), subcarrier multiplexing, time-division multiplexing (TDM) or any combination of these.

Transparent AONs resemble PONs in many ways, but, rather than economy, they stress flexibility, high performance and local to global coverage. The idea is to provide optical transparency enabling each of a large number of optical WDM channels to propagate from source to destination without optical-to-electronic data conversion within the network. The signaling format of each of the WDM channels can be virtually arbitrary, providing tremendous potential for flexibility and economic upgradability. Besides the aforementioned passive optical components used in PONs, transparent WDM AONs can employ devices that control the flow of the optical signals, such as optical switches and configurable wavelength routers and wavelength converters. The control of these devices can be done electronically, as long as the path of the signal itself remains optical. Such control can be slow if it is used mainly to reconfigure the network to accommodate various connection requests and varying traffic patterns.

The terminal equipment in transparent AONs would employ tunable-laser transmitters and/or tunable-filter (or heterodyne) receivers. Two users can establish a communications path through the network by picking the appropriate transmitter and receiver wavelengths, perhaps after making a prior request from a network controller. After providing the connection, the network should be transparent to the modulation format employed. More interestingly, any group of users can pick the right set of wavelengths and reconfigure the AON so as to create their own transparent virtual network. Any user should be able to take part in multiple simultaneous sessions. Moreover, the AON can, in principle, be used to provide packet-switching service optically. In this case, the network may have to be reconfigurable within packet time scales, and optical packet storage may be needed to mitigate contention. Alternatively, electronic packet switches could be placed throughout the network to provide store-and-forward service which, when used in conjunction with signal regeneration and wavelength translation, can permit full connectivity at the packet level among access stations which are only sparsely connected at the optical level. This approach does not require the aforementioned fast tunable transmitters/receivers, fast optical reconfiguration, or optical packet storage. Thus, a transparent AON can potentially provide a very rich variety of services to a large community of users separated by local to global distances.

Ultra-high-speed AONs take advantage of the tremendous speed of various optical phenomena to realize peak channel data rates that are much higher than can be generated, controlled or switched electronically. In such networks, communication is effected by on-off modulated, ultra-high-rate (e.g., 100 Gb/s) time-division multiplexed (TDM) streams of short optical pulses or solitons. Virtually all network controls and signal switching functions are accomplished optically; i.e., these are *truly* all-optical networks. However, they are *not* optically transparent.

It appears that such ultra-high speed TDM AONs may be limited to local to metropolitan geographical extent. This is caused by the combination of single-mode fiber dispersion and nonlinearity (e.g., four-wave mixing and, particularly, stimulated Raman scattering), which seems to present an insurmountable barrier to the propagation of pulse streams at more than a few tens of gigabits per second over transcontinental or transoceanic distances. These efforts may be mitigated by low-gain fiber amplifiers and dispersion-compensated fiber. At present, however, such networks may be more natural at facilitating ultra-fast packet switching.

CHAPTER 2:
THE IMPORTANCE OF OPTICAL NETWORKS

Optical networking has emerged from its infancy, which tended to focus on photonic devices and early demonstration of physical level potential, and has now matured to the extent that a focused interdisciplinary effort can bring forth major knowledge advances, systems and device-oriented innovation, education opportunities at all degree levels, and a solid technological platform for subsequent commercialization. Involved disciplines would include photonic devices, high speed electronics, telecommunication networking, high-speed protocols, and computer engineering.

The great (and timely) importance of optical networks arises from their potential to offer unprecedented low cost telecommunications capacity. It is estimated that a single strand of single mode optical fiber possesses a potentially accessible bandwidth approaching 100,000 Gigahertz! Let us put this number into perspective. The entire radio spectrum managed by the Federal Communications Commission, containing all AM, FM, TV, cellular telephony, and satellite communication signals, among others, consists of approximately 60 Gigahertz. One strand of single mode optical fiber contains a potentially addressable spectrum over three orders-of-magnitude greater.

There is a need to develop the system, network, and device concepts required to expand our utilization of the enormous bandwidth of optical fibers in a manner which will provide networks which will support a variety of advanced applications. Fundamental research is needed to produce novel system architectures and enabling device capabilities which allow the collective spectrum of networks containing many such fibers to be captured and shared among a plurality of users. Further research is needed to permit each network access port to provide as high a peak and average data rate to the application being supported through that port as can be produced by state-of-the-art electronic, photonic, and electro-optic technologies. Finally, research is needed to discover the types of high speed applications which might be enabled by such a network and the requirements which such applications might present to the network. Such research advances will help to maintain American pre-eminence in telecommunications and computing systems, which is beginning to buckle under the pressures of intense international competition and the industry's attempt to shore-up short-term profitability by reducing R&D expenditures.

Potential applications for optical networks include (1) ultra-high-speed computer interconnects for parallel processing, (2) Local and Metropolitan Area networks service business, academic, and medical needs for voice, data, image, and video communications, and (3) wide area networks for multimedia communications enabling a much more natural type of people-to-people communication, access to vast electronic libraries, and creation of superlative distributed computing environments. Direct beneficiaries of the research agenda described in this report include telecommunication carriers, and computer and telecommunication equipment vendors. Indirect beneficiaries include the manufacturing and service segments of the American economy, both critically dependent upon advanced telecommunication and computing services for access to information, access to customers, inventory control, and improved worker productivity. In short, optical networks hold forth the potential to emerge as a key element of an American infrastructure designed for economic competitiveness and improved quality-of-life as we approach the dawn of the 21st century.

The need for an interdisciplinary focus is best demonstrated by the tight coupling which exists between network architectures and device capabilities. In essence, the early accomplishments of researchers in this field have taught us that research on new optical network architectures, in the absence of a sound knowledge of device capabilities and limitations, can unproductively lead to approaches which are either naively unrealizable or unnecessarily short-sighted. Similarly, research on new active and passive optical devices, conducted without some system context, can unproductively lead to proof-of-concept of technologically sophisticated but otherwise useless components. Interdisciplinary research is essential to the effective identification of useful device needs and appropriate system directions.

Research directed toward optical networks is clearly gaining momentum, as evidenced by the growing number of researchers and students working in this field and the number of papers being published reporting on new system innovations and device proof-of-concept demonstrations. Now is the time for a major new research effort to harness these diverse activities and produce the research/education advances needed to fully exploit the aforementioned potential.

At its core, an optical network taps into the vast optical communication spectrum to create and effectively utilize many parallel broadband channels which serve to interconnect the various network access stations or interfaces. In general, these channels may be created by means of wavelength division multiplexing and/or time division multiplexing. Both approaches require the type of high-risk, high-reward research appropriate in the university environment, with the potential to dramatically reduce the cost per unit bandwidth, both on a collective aggregate basis and on a per-user basis, relative to conventional network approaches wherein switching and multiplexing equipment is simply interconnected by fixed-point transmission links. With Time Division Multiplexing, each channel corresponds to a time slot allocation from a Time Multiplexed frame, and the actual physical link data rate is a high multiple of the electro-optic user interface speed. With WDM, each channel corresponds to a unique wavelength allocation. Both approaches allow for spectral re-use (the same wavelengths or time slots can be re-used on a non-interfering basis in geographically remote parts of the network). In addition, WDM also affords transmission transparency: the channels can carry information in *any* format (including analog, if appropriate), and it is easy to upgrade their capability by deploying more sophisticated access station technology without requiring a commensurate upgrading of the possibly wide-area optical medium. Furthermore, each access station can use technology with a complexity commensurate with the application supported through that interface; access station complexity is not driven by the needs of the most sophisticated application. Finally, and very importantly, both WDM and TDM permit dynamic channel re-assignment: by altering the wavelength or time slot assignments, access station interconnecting channels can easily be redeployed to change the connection pattern among the stations. This feature can be used to “match” the connection pattern to prevailing traffic demands and to re-route traffic around failed network elements. In effect, an optical network creates a dynamically re-allocable “pool” of high speed channels which can be deployed to appropriately interconnect network apparatus and applications. Reconfiguration can be effected in response to demographic changes, time-of-day variation in traffic patterns, or real-time demand for end-user connections, and highly reliable connections are produced.

A further advantage of optical networks is the ability to assign/reassign channels to different types of traffic (voice, data, image, video). In effect, multiple virtual networks can share a common physical medium. By creating logically segregated networks with the ability to rapidly redeploy resources from one to the other, it is expected that the many traffic control issues that traditionally arise in the integrated multimedia traffic environment will be much more readily addressed, and the great difficulty in managing performance to guarantee quality-of-service will be substantially simplified without significant sacrifice of aggregate network capacity or end-user bandwidth.

CHAPTER 3:
RELATIONSHIP OF AN ENHANCED RESEARCH AGENDA IN OPTICAL NETWORKING TO OTHER MAJOR PROGRAMS

3.1 High Performance Computing and Computer Communications Program

NSF requested the development of a research agenda as an outcome of this workshop. The workshop participants believe that there are significant research opportunties in all-optical networking. Furthermore, such research would support the goals of HPCC, namely by providing the types of scalable, high speed, high performance, cost-effective networks needed for parallel-processing computer interconnects, and Local Area, Metropolitan Area, and National-scale multimedia services. As described in the Grand Challenges 1993 - HPCC "teal book," the overriding objectives are to:

- Extend U.S. technological leadership in high performance computing and computer communications.
- Provide wide dissemination and application of the technologies both to speed the pace of innovation and to serve the national economy, national security, education, and the global environment.
- Spur gains in U.S. productivity and industrial competitiveness by making high performance computing and networking technologies an integral part of the design and production process.

Since HPCC is to support high-performance networking as well as high-performance computing, all-optical networks will be essential to both. Furthermore, knowledge from work on all-optical interconnects will help us understand how to deal with communication systems in general.

Many future engineering and scientific research problems are computation and memory intensive. Today's supercomputers, based on a single high speed computer, are incapable of supporting the speed required by these applications. The future generation of high performance computers will consist of many hundreds or thousands of simpler and less expensive interconnected processors, and can achieve speeds radically different from single processor based computers. One of the primary thrusts of the HPCC initiative is the development of a teraflop, general-purpose processing system by 1996.

Over the last decade, processors and related peripheral modules have generally advanced in speed by two orders of magnitude, but interconnects by only one. Compensating adjustments have been made in architectures and algorithms, but these strategies seem stretched to near their limits, and are unlikely to satisfy the requirements of the next generation teraflop systems. Traditional electronic interconnection networks for such massively parallel systems might require a prohibitively large wire density and may be unable to provide the required speed of communication. Optical interconnects are a promising candidate for future interconnection networks due to their high bandwidth, the high density they provide and their low power requirements. In addition, the volume of a single fiber is considerably less than a massively parallel (100+ wide) wiring on stripline board. Fiber's immunity to electromagnetic interference and ground loops is well known and would greatly ease the interconnection of a distributed processor. The flow-through electronic switching protocols already used map naturally onto photonic implementations, with little significant change. If single wideband optical devices perform the word-parallel operations of switching, power regeneration, and retiming, savings will also be realized in complexity and perhaps even in volume where it is at a premium at the switching points.

On a related note, a generic computer-oriented km-scale interconnect with considerably higher performance could be profitably retrofitted into currently heterogeneous, large-scale computer centers. Such centers, with a few diverse supercomputers at the top of the hierarchy, new high bandwidth-devouring systems such as frame buffers, and an ever-growing horde of workstations at the bottom, have already strained their interconnection networks past reasonable limits. Centers such as that at NCSA in Illinois consider the problem so serious that they are developing their own new systems to replace the mixture of ethernet, FDDI, ULTRA, and HIPPI hastily pressed into service.

All-optical networks should be viewed as a potential candidate for a national telecommunications infrastructure supporting low-cost *universal* access to supercomputer centers and vast multimedia data bases, along with substantially enhanced, more natural people-to-people communications. As such, a virtual subnetwork of such an all-optical telecommunications infrastructure might serve as the National Research and Education Network (NREN), providing user access speeds measured in gigabits per second. Other virtual subnetworks created from the same physical infrastructure could provide additional non-NREN services including, for example, general-purpose public multimedia telecommunications and delivery of High Definition TV signals for recreational, cultural, and educational needs.

Finally, advanced optical network technologies will also accelerate the research process in all disciplines and enable educators to integrate new knowledge and methodologies directly into course curricula. Students at all levels will be drawn into learning and participating in a wide variety of research experiments in all components of an expanded research effort.

3.2 ARPA's Optical Networking Initiatives

Recently, ARPA started several new initiatives in all-optical networks that resulted in the formation of industry-led, precompetitive consortia of industrial organizations and universities to foster technology and systems research and development in these very important areas. The optical network efforts address scalable optical networks that can support a broad set of services over a local geographic area as well as metropolitan and wide geographic areas.

Each consortium uses a testbed as the main point-of-focus of the effort. These testbeds will be built with devices that are commercially available as well as state-of-the-art devices supplied by research laboratories. The testbeds are expected to be fully functional in two years. Although some exploratory research is included in these efforts, these ARPA programs emphasize the construction of testbed hardware in the near term.

The central role of NSF in providing supportive university-based research should be to address optical network research of a longer time-horizon as well as pursuing near-term alternative approaches that look promising but are not specifically included in the plans of these consortia. The longer-time-horizon role of NSF is extremely important because optical network technology is still in its infancy. There are many budding innovative ideas that are worthy of support and a substantial fraction of these ideas will not be funded by the ARPA program. For example, both network testbeds are wavelength-division-multiplexed based. A possible alternative, namely, a time-division-multiplexed-short-pulse based network, was not funded, perhaps because of the lack of mature devices ready for a near-term testbed construction. Also, many aspects and opportunities of wavelength-multiplexed networks are not being addressed, such as the inherent broadcast and multicast capabilities over wide service areas, all-optical crossbar switching networks, applicability of devices which are rapidly tunable over small wavebands, and full exploitation of the "clear channel" to enable new applications, perhaps due to a combination of inadequate funding and the relative immaturity of network concepts and requisite devices.

In addition, since the field of optical networks will likely be an important component of the telecommunication industry in the future, it requires that the training of students articulate in the system and technology of optical networks begin now. NSF can adequately fulfill this role through a broad-based optical network research program. Thus, the NSF supported research will have a complimentary role to the ARPA effort in that it will be centered mostly in the academic environment rather than being industry-led, and diverse research paths of longer time horizons will be a primary part of the program.

CHAPTER 4:
RESEARCH FOCUS AND RESEARCH MODES

To bring research in optical networks to fruition, a close interaction between device, system, and network researchers is needed because of the intimate relationship between device, system, and network issues.

Both theoretical and experimental efforts are needed including systems-level proof-of-concept demonstration experiments. Full-scale testbeds involving multiple interfaces and application demonstrations are very expensive and may preclude the use of risky, experimental devices. Furthermore, full-scale optical network testbeds based on nearer-term, relatively mature technology are being developed under ARPA funding.

The much-needed interaction between device, system and network researchers, along with the required broad scope of theoretical and experimental work, makes collaborative interdisciplinary research groups (perhaps consisting of two, three or four professors working in different fields, along with their students) highly desirable. Such groups will combine the best talent to address the foregoing important research problems while crossing institutional, geographical and disciplinary borders. The research groups are expected to require a substantial amount of test and measurement equipment.

The focus of NSF-sponsored research groups should be on long-term, high-risk research areas. Its time scale should expand beyond that of ARPA and other mission agencies to support the needs and the infrastructure of the American society of the 21st century.

The support of research groups must not come at the expense of individual (single professor) research efforts. Individual research may yield innovative ideas hard to generate in any other way. Individual research efforts are, therefore, encouraged, along with research group efforts, and the continued support of such efforts is urged.

The focused effort in the field of optical networks should be directed toward WDM, TDM (short pulse) and hybrid networks for LAN, MAN and WAN applications. It should also address the issues of interconnecting potentially heterogeneous optical networks.

There are many examples of challenging research topics which are of fundamental importance to eventual practical implementation. These approaches involve support of broadband integrated services in a manner employing the unique capabilities of optical networks. Such research topics might be addressed by a cooperative research effort conducted by device, system and network researchers.

An interdisciplinary group of three to four professors (and their graduate students) might need a budget of \$500K - \$1M per year (including equipment). We estimate that up to ten such research groups will be needed to address the broad range of the foregoing research issues.

APPENDIX A:
DEVICE ISSUES

Successful implementation of future optical networks requires many novel optical and electro-optical devices. Key issues in device research pertinent to optical networks are discussed in this appendix.

A.1 Devices For Wavelength Tuning

For future WDM applications, rapidly tunable lasers will be necessary. Stable operation at any one of a number of different, controlled wavelengths will be the goal. Some 100 or more separate wavelengths will be required to utilize more of the inherent bandwidth of the fibers. The wavelengths must be far enough apart, and stable enough, to allow for the bandwidth of the information being sent at each wavelength. The lasers could be operated continuously at any particular wavelength, with integrated modulators to put the information on the optical beam, or the laser could be directly modulated. Several or even many lasers, each capable of stable operation at a number of wavelengths, can be used to achieve the desired number of wavelengths for the utilization of the fiber bandwidth. Each laser can be designed, through the composition and thickness of the quantum wells, and through the period of its distributed Bragg reflection feedback, to achieve the particular wavelengths. The means of altering the wavelength of the lasers can be the applied bias to one or more terminals of the laser, or possibly the bias applied to a tunable external resonant structure.

Some specific laser structures being used to electrically change wavelength include: dual-waveguide lasers with a single active layer; and a one-dimensional array of lasers, separately biased, with either a graded-period grating, or a graded-period distributed Bragg reflection structure behind the lasers. To select an operating wavelength fast enough to allow useful WDM, rapid wavelength tuning is necessary. Bias tuning should be accomplished in 100 picoseconds, which will allow an appropriate settling time at the new wavelength. Because matching-wavelength tunable filters might be needed to select the signals at the different wavelengths, there will eventually need to be correlation of standard wavelengths. This will allow the laser transmissions to be received on the proper channel.

For TDM operation, short pulses at particular wavelengths will be necessary. These light pulses will need to have the proper shape to allow soliton operation. The short optical pulses can be formed by pulsed bias operation, with bias applied to one or more electrodes on the laser, or with bias applied to an external modulator. Study of the laser non-linear coupling between modes, and the related effects of the multiple-electrode bias, will be necessary in order to achieve the desired results.

Finally, both the optical source and the optical receiver will need to have integration of electronic and optical devices on the same chip. Such OEIC structures will be necessary for the high electronic bandwidth.

A.2 Star Couplers and Wavelength Changers

Star couplers

N x N star couplers are the main building block of broadcast-and-select networks. Good star couplers should have low excess loss and uniform transmissivity among input/output ports and over the wavelength band of interest.

Star couplers have been made out of 2 x 2 fiber couplers and on planar integrated optics waveguide. The latter could be mass produced, but have higher loss and need to be pigtailed. There is no intrinsic size limit for star couplers, but their cost and bulk increase faster than linearly with N.

Broadcast star networks can also be built out of N x 1 and 1 x N couplers. This configuration has the advantage of allowing growth simply by cascading couplers. However its larger attenuation makes the use of optical amplifiers mandatory.

Frequency-selective couplers

If special care is taken to make the coupling frequency-selective rather than frequency-insensitive, then the couplers can be used either as wavelength routers or as filters. As routers, they are highly desirable from a system point of view since they can largely eliminate splitting losses and they allow frequency reuse. Devices that discard some wavelengths are more limited in their applications, and are usually used as filters only.

Ideally, each individual passband should be broad and flat in order to ease channel accuracy requirements and accommodate high bandwidth signals, yet there needs to be high isolation of the rejected bands. These requirements may be in conflict.

As for all devices, low loss and low polarization sensitivity are desirable. In particular, the latter tends to become more difficult to achieve as the couplers are made frequency selective. The wavelength-selective routers and filters in any system must be matched to each other and also to all the sources used in the system. Even these nominally static devices may require some initial or slow adjustment.

Possible desired configurations for wavelength-selective couplers depend on the number of inputs and outputs. For example, all the wavelengths may enter on a single fiber and the device is required to route some wavelengths to one output and all others to a second output. Or, each wavelength may be sent to a separate output. A harder requirement to meet asks that signals be routed such that all signals of a given wavelength be sent to the same output regardless of input port.

Most technologies to accomplish wavelength selection are based on interference effects, although some filtering can be accomplished by using material absorption characteristics. The well-known technology of diffraction gratings has long been used to make N x 1 (1 x N) devices known as wavelength (de)multiplexers. Also, N x N frequency-selective couplers have been fabricated in planar integrated optics, using interference between waveguides of different path-lengths or reflection gratings, with N in the range 20-100. All these implementations need to be perfected. In particular, flexibility in specifying what sets of frequencies should be split off needs to be improved.

Dynamic frequency-selective couplers

The couplers just discussed are largely static, although they can be adjusted by mechanical pressure or temperature change. More sophisticated devices whose characteristics can be easily varied are known as tunable filters, or dynamic wavelength routers, depending on their use. In addition to the previous requirements for low loss, uniformity, and polarization insensitivity, they should also have reliable and accurate tuning characteristics, and enjoy the stability required of all wavelength-sensitive devices.

Tunable filters are extremely useful from a system point of view, as they can switch optical signals without O/E conversion. Their utility increases greatly with their size and their tuning speed, and very fast tuning will be required for packet switching. Unfortunately, the state of development of such devices is rather limited.

As in the case of static routers, the ideal device can provide arbitrary routing. However, in practice, most technologies do not do this easily. For example, Fabry-Perot filters with piezo-electric tuning are restricted to selecting a comb of evenly-spaced frequencies. Furthermore, such methods as electro-mechanical tuning are slow. The 1 x N demultiplexers formed by a cascade of Mach-Zehnder filters (temperature-tuned) cannot be used easily to route an arbitrary subset of wavelengths to the same output. The 1 x 2 (or 2 x 2) acousto-optic filters (tuning in about a microsec) have the opposite restriction. Despite better speed and the ability to separate the wavelengths into two arbitrary subsets, separation into multiple outputs requires cascading devices.

Larger N x N structures could be realized by using a stage of static wavelength demultiplexers, a stage of wavelength insensitive lithium niobate matrix switches (ns response time), and a stage of static wavelength multiplexers, but this arrangement is far from ideal.

A critical issue with any device is control. If the device is continuously tunable, then drift is a potentially serious problem. On the other hand, if the device is discretely tunable, short-term operation is simplified but initial matching to other components in the system becomes more difficult.

Frequency changers

The next level of device complexity appears in frequency changers, which could be static or dynamic. When used together with static demultiplexers they can also function as wavelength routers, and they thus enjoy a dual role. Various designs are known to change the frequency of a single channel, e.g. demodulation followed modulation at another frequency. This is neither elegant nor very practical. More sophisticated designs are based on 4 wave mixing (third order intermodulation products) between a band of signals and a pump signal in a nonlinear medium such as a semiconductor amplifier. Such devices can translate a band of signals by THz in ns, but the demonstrations to date have suffered from low efficiencies, and much more research and development efforts are required to make them practical.

A.3 Optical Amplifiers

Optical amplifiers (OA) are key components in lightwave communication networks. They provide the means to directly regenerate optical signals without need for electronic conversion. Additionally, several optical WDM channels can be amplified simultaneously, without need for demultiplexing. The operating bandwidth of OAs ranges from DC to several Terahertz, which considerably surpasses the performance of electronics.

The technology of rare-earth doped fiber amplifiers has now reached a mature stage. In particular, praseodymium-doped and erbium-doped fibers pumped with compact laser diodes provide THz-wide gain bandwidths centered near $\lambda = 1.3$ μm, and $\lambda = 1.5$ μm, which corresponds to the two lightwave transmission windows.

Erbium-doped amplifiers (EDFA), which operate at 1.5 μm, have already revolutionized the field of long-distance communications. Indeed, by 1995 EDFAs will be deployed in a new generation of transoceanic links. In such systems, signals will be transmitted without any form of electronic regeneration over distances near 10,000 km, and at bit rates up to 5 Gbit/s.

Praseodymium-doped fiber amplifiers (PDFA), which operate at 1.3 μm, also have a tremendous potential for long-distance communications, but in the field of terrestrial systems. The very large fiber plant already existing in the US territory, which represents more than 1 million km of standard single-mode fiber, could be upgraded in the near future by the substitution of PDFAs in place of conventional electronic repeaters for higher operating bit rates.

In the field of lightwave networks, it is anticipated that OAs based on EDFA and PDFA devices will likewise have considerable impact. A first category of applications of OAs in lightwave networks include power boosters, in-line repeaters and optical preamplifiers. A plethora of recently demonstrated OA-based devices can also be implemented in network systems, forming a second category of potential applications. The two types of applications of OAs in lightwave networks are discussed below.

The performance of unrepeated communication systems, as expressed in bit-rate x transmission distance product, is determined primarily by the system's power budget, and secondarily, by fiber dispersion and nonlinearities. Given the transmitter power and receiver sensitivity, the maximum possible transmission distance is determined by system loss. In the case of lightwave networks, signal loss is due to many factors, which include fiber transmission loss, splice loss, and excess loss of various optical components. In particular, active optical components such as 2 x 2 switches are characterized by significant insertion loss (i.e., 3-5 dB), which has limited so far the user capacity of all-optical networks. Such a perspective is radically changed by the implementation of OAs as loss compensators. The direct result of loss compensation by OAs is a dramatic increase of the network capacity in number of end-users. This number cannot be indefinitely increased, however, as each OA contributes for some amount of amplified spontaneous emission noise (ASE). The accumulation of this background ASE noise throughout the transmission medium eventually results in the degradation of signal-to-noise ratio (SNR) at the receiver end.

One of the key research issues in the implementation of OAs in lightwave networks concerns the analysis of ASE and SNR as these parameters evolve throughout the system. Such analysis should help define optimal network architectures in which SNR degradation is minimized, and hence, network capacity is enhanced. Given the network architecture, optimal locations of OAs for minimal ASE accumulation must be determined. In lightwave networks based on WDM, relatively large amounts of optical power can be concentrated in different nodes, which causes amplifier gain saturation. In such nodes, power OAs with high output saturation power characteristics must be implemented. In passive lightwave networks (whether WDM or TDM), OAs can be used as in-line repeaters. As a result, a virtually transparent transmission medium from transmitting to receiving ends can be realized. Optimum locations of these optical repeaters must then be determined for minimizing ASE noise accumulation in the system.

We have discussed the potential of OAs as loss compensators. In addition, OAs offer the possibility of many other all-optical functions. All-optical OA-based devices include time- and wavelength- domain switches (nonlinear loop mirrors), optically-controlled 2 x 2 space switches (tandem amplifier gates), optical memories (recirculating delay lines), time-domain demultiplexers (4-photon mixing), and various signal-processing devices (e.g., power equalizers, time-slot interchangers). The main attractive feature of these OA-based devices is that they are bit-rate and modulation-format transparent. The functions that such devices can offer open new perspectives in the field of all-optical lightwave networks and related system concepts. Research should focus on the demonstration of practical and compact OA-based devices (e.g., store-and-forward all-optical switches) and, in parallel, analyze with realistic experimental parameters how such devices could impact the system performance.

In summary, OAs can be used in lightwave networks either as a means of overall loss compensation, or as all-optical and high-performance components. As OAs are available at both 1.3 µm and 1.5 µm, new network architectures and protocols that would make full use of the corresponding transmission windows could be investigated.

In the field of fiber amplifiers, several research issues remain to be addressed at the device level. For EDFAs these include:

- Determination of optimal pump bands - such an optimization consists in finding trade-offs between high power conversion efficiency (1480nm band), minimal noise figure (980nm band) and cost effectiveness (810nm band, available from inexpensive compact-disk laser diodes);

- Equalization of gain spectrum, which is relevant to WDM systems. Gain equalization can be achieved by various methods (e.g. acousto-optic filtering, gain saturation) whose performance in network implementation must be compared. Novel equalization techniques yielding improved gain flatness and power dynamic range must also be investigated;

- Automatic gain control (AGC). Implementation could be required to compensate possible effects of transient gain saturation in the amplifiers. Such transients can be generated by accidental bursts of traffic or packet collisions, resulting in a temporary increase of OA signal input. The relative merits of several possible AGC techniques, based on pump or signal feedback, must be tested experimentally.

A.4 TDM Devices

While much current research is directed at wavelength-division multiplexed systems, time-division multiplexing can also be used in optical networks. Ultrafast all-optical switching may enable operation beyond electronic speeds (> 50 Gb/sec.). Such switches have almost instantaneous response and recovery times because they rely on virtual transitions (i.e., deformation of electron clouds) in nonlinear media. Since most virtual transitions are nonresonant interactions, they may dissipate less power than high-speed electronics. For ultrafast time-domain systems using picosecond or sub-picosecond pulses in optical fibers, solitons are crucial for avoiding the deleterious effects of nonlinearity and group-velocity dispersion.

Achieving optical control in networks requires development of several key enabling technologies. First, all-optical logic gates are required for processing the header on a packet. Second, compact, synchronizable, short pulse lasers are required as the optical power supplies for the network. Third, synchronization circuits and techniques for reducing timing jitter are required for the accurate alignment with only picosecond bit periods. Finally, we must show that bit-rate and packet-rate components can be integrated to self-route packets through the network interchanges. Below we discuss research required in each of these topics.

A.4.1 <u>All-optical logic gates and routing devices</u>

All-optical logic gates rely on third-order nonlinearities and virtual optical transitions to have an almost instantaneous response time. The logic gates must be cascadable, have fan-out, and be capable of performing all Boolean operations with little or no gate delay. It is also desirable that the switches be three terminal devices with input/output isolation and logic-level restoration (e.g., regenerative gates). For example, soliton-dragging logic gates have been demonstrated to satisfy most of the above requirements for digital logic gates, to operate as fast as 200 Gbit/sec., and to have switching energies approaching a picojoule. In addition, all-optical routing devices such as nonlinear optical loop mirrors and four-wave-mixing gates have also been demonstrated in optical fibers. Routing devices are particularly important for random, non-periodic demultiplexing of bit streams.

Although fibers have proven to be an almost ideal nonlinear material for proof-of-principle demonstrations of all-optical devices, most of the devices require fiber lengths of 300m to several kilometers because of the weak nonlinearity in fibers. Therefore, further research is required on novel nonlinear materials to reduce the latency (reflected in gate delay) of all-optical devices. For example, nonlinear spectroscopy of semiconductor waveguides or organics is required to assess their use in all-optical switches. In addition, all-optical devices need further improvements in terms of switching energy, tolerance to timing jitter, and sensitivity to environmental variations.

A.4.2 Short pulse lasers

High-speed, time-domain networks will require low-noise, stable, compact, synchronizable short pulse sources. The femtosecond pulses, which may be solitons, must be transform-limited, operate in the wavelength range compatible with erbium-doped fiber amplifiers, have high repetition rates, and provide sufficient energy per pulse. For example, passively mode-locked, self-starting, all-polarization maintaining erbium-doped fiber lasers have been demonstrated that generate nearly transform-limited pulses as short as 320 fsec with pulse energies up to 40 pJ. In addition, recently, arrays of semiconductor lasers have been mode-locked to generate picosecond pulses. The optical power supplies for any network application will have to be either semiconductor lasers or diode-pumped lasers with relatively high wall-plug efficiency.

A.4.3 Synchronization circuits and reducing timing jitter

As the bit-rate increases and the bit-periods become shorter, the timing restrictions and synchronization issues become increasingly difficult. Devices and circuits must be demonstrated for temporarily aligning two bit streams, for synchronizing two lasers, and for synchronizing a laser to an incoming bit stream. For example, all-optical phase locked loops are required for master-slaving two lasers. All-optical “elastic stores” are required to reduce the timing jitter in a bit stream. Techniques for reducing jitter in long-haul transmission systems--such as spectral and temporal filters--may also be valuable for short pulse networks.

A.4.4 Integration of bit and packet rate devices

All-optical devices are still at an early stage of research, and, for a while at least, the devices can be expected to be expensive. Therefore, it may be necessary to restrict the use of high-speed, bit-rate switches to key areas where they bring a value-added service to the network, and we want to use more conventional packet-rate switches in other areas. In other words, one may use a handful of all-optical devices in carefully chosen positions within the network. As an example, a soliton ring network has been designed that uses four logic gates to decode the packet header while electro-optic switches are used to route the packet. Even trying to demonstrate header processing in a 100 Gbit/sec. soliton ring network will raise important issues such as power requirements, synchronization, thermal stability and compatibility and cascadability of logic gates.

A.5 Device Integration and Packaging

Optical devices such as wavelength agile laser diodes, tunable filter, semiconductor optical amplifiers, waveguides and high-speed logic circuits are important basic components in high-speed optical networks. However, as the size of the network and the port bandwidth increase, the optimized integration of optical components on a common substrate and its compact packaging become increasingly important in view of network performance, reliability, and cost. The assembly of discrete optical components or circuits is not only difficult due to alignment problems, but also causes significant power loss and speed bottleneck problems in the interconnects. For many multiplexing technologies such as WDM, integrated arrays of optical sources, photoreceivers, optical waveguides, and control circuits are in high demand.

The demand for advanced compact packages for optical devices and circuits is as critical as integrated devices. The current art of packaging provides bulky packages which take too much space in the system, thus enlarging the physical dimension of the overall system. Devices should also be designed such that they can be readily packaged for systems insertion.

For some time, the cost of optical components has been high due to manufacturing difficulties and the lack of volume demand. As a result, optical components have not been able to replace their electronic counterparts, thus limiting the proliferation of optoelectronic ICs. At the same time, the user's volume demand has not been generated due to high cost. With development of mature device integration technology and its continued evolution, the industry's catch-22 dilemma can be resolved. Such resolution will in turn expedite the proliferation of photonics technology.

MOSIS-like service for opto-electronic integrated circuits fabrication and packaging, tentatively called here, Optoelectronic Circuits Integration Service (OCIS), to both industry and universities are much needed for fast prototyping and development of important classes of subsystems and for education of future engineers in this field. Adequate CAD tools for optical circuits and systems should be in place for design analysis and optimization before submission of masks. When such service is in place, users of OCIS can submit their designs in CIF files electronically and receive packaged optical circuits for systems insertion or testbed experiments for systems applications.

A.6 Limitations on WDM Systems Caused By Fiber Nonlinearities, Dispersion, and Jitter

Nonlinear effects in optical fibers impose serious limitations on optical WDM communication systems by setting bounds on the optical power, bit-rate, and channel separations. The advent of erbium-doped fiber amplifiers (EDFAs) makes the WDM technology very attractive due to the large gain bandwidth of the EDFA in the 1.55μm wavelength region. However, limitations due to nonlinear effects caused by the EDFA itself occur as well.

The nonlinear effects in an optical fiber impose limitations by depleting the power (stimulated Raman and Brillouin scatterings), creating light at new wavelengths (four-photon mixing), or disrupting the phase (carrier-induced phase modulation).

- Stimulated Raman scattering (SRS): This effect is caused by inelastic light scattering due to vibrational excitation modes of SiO_2. The scattering results in light pumping from the high-frequency part of the spectrum to the low-frequency part. The Raman gain spectrum in silica fibers extends to about 30 THz, with a peak at the order of 10^{-11} cm.W^{-1} occurring at a Stokes shift of about 13 THz. The typical critical power P_α for the onset of SRS is 1 W. For an *N*-channel WDM system with a channel spacing of Δf and an average power per channel *P*, this would translate into a limitation on the product of the total optical bandwidth by the total optical power, or $[N - 1) \Delta f] [NP]$. The maximum value of this product set by SRS is about 500 GHz.W. In terms of allowable number of channels, taking $\Delta f = 10$ GHz, the value of P_α starts decreasing from its 1-W value for about 10 channels, to about 1-mW for 200 channels, to a few tens of μW for 1,000 channels. One way to reduce SRS effects is to reduce the channel spacing.

- Stimulated Brillouin scattering (SBS): This effect is similar in principle to SRS, except that the nonlinear interaction involves an acoustic phonon in SBS as compared to an optical phonon in SRS. The bandwidth of the Brillouin gain is much narrower than the Raman gain, and is approximately 10 MHz, with a peak at the order of 6×10^{-9} cm.W^{-1} occurring at a Stokes shift of about 10 GHz. The gain is only significant in the backward direction and can therefore be greatly reduced by the use of optical isolators. The typical critical power for the onset of Brillouin effects is less than 10 Mw and does not depend on the number of channels for a given channel separation. To reduce the effects of SBS, high bit-rates or carrier suppression can be used.

- Self phase and cross-phase modulations (SPM and XPM): By contrast with SRS and SBS, this is an elastic effect, since there is no energy exchange between the nonlinear medium and the electromagnetic field. Both SPM and XPM are governed by the Kerr effect, which involves the third-order susceptibility $\chi^{(3)}$ of the optical fiber, resulting in a change of refraction that is linear with the intensity. In the case of SPM, this intensity is the one carried by the channel itself, while in XPM the nonlinear index varies with the intensity carried by adjacent channels. This nonlinear index of refraction causes a nonlinear phase shift that sets a limitation on the channel capacity. For a channel separation of $\Delta f = 10$ GHz, the onset for those carrier-induced phase-modulation effects decreases from ~ 1 W for two channels to about 1 mW for 1,000 channels. Those effects can be reduced by reducing the amplitude modulation on the data channels, and by using quiet light sources.

- Four-photon mixing (FPM): This is also a $\chi(3)$ effect, in which photons from one or more waves are annihilated and new photons at different frequencies are created in such a way that the total energy and momentum are conserved during the interaction. This process does not occur unless a phase-matching condition--implying a specific relation between the frequencies and the refracted indices--is satisfied. This is fundamentally different from the case of SRS and SBS in which the phase-matching condition is automatically satisfied by the participation of the medium in the nonlinear process. In step-index fibers, the phase-matching efficiency is close to unity for channel separations not exceeding 15 GHz. The power in the "created" frequency is close to 1/10 of the power P_0 in the channel for $P_0 \simeq 0.1$ mW. FPM induces a strong limitation on WDM channels: the maximum power per channel is a few mW for a channel separation $\Delta f = 100$ GHz, and it drops to only a few tens of µW for more than 10 channels if $\Delta f =$ 1 GHz. The effects of FPM can be greatly reduced by increasing Δf or by using dispersive fibers.

Among the limitations from nonlinear effects introduced by EDFAs is one due to the non-uniform wavelength-dependent gain profile and saturation characteristics of the EDFA. Each channel in a WDM system will experience different optical gain. This leads to a degradation of the bit-error rate performance of some channels. For equal intensities launched in a seven-channel WDM with 2-nm channel spacing, after an 840-km propagation and a 70-km amplifier spacing, the output power varies by ~ 245 dB, from $\lambda = 155.0$ nm to $\lambda = 1562.0$ nm. A way to circumvent this is by "gain-equalization" that consists of adjusting the powers of the different channels at the terminals with variable attenuators, while keeping the total power in the fiber constant, following an algorithm that allows to equalize the output signal-to-noise ratios.

APPENDIX B:
NETWORKING ISSUES

The telecommunication networking aspects of optical networks represent fertile grounds for the pursuit of fundamental scientific and engineering research of the highest quality. Much new knowledge must be generated, new innovative architectural approaches proposed, and new analytical methodologies developed to fulfill the potential of optical networks by harnessing the capabilities of optical device technologies while respecting their fundamental limitations. Pursuit of new knowledge and technical feasibility demonstration of new approaches further provide a unique opportunity to involve students at all degree levels in the program of research.

Optical networks truly present unique opportunities but are subject to unique constraints. We have already mentioned the potential benefits arising from the availability of (1) universally affordable access to a nationwide optical network, (2) local and metropolitan Area optical networks meeting industrial and business needs, and (3) optical networks for high performance parallel processing supercomputers. Commercial realization of these potential long-term benefits to enhance economic competitiveness and improve quality of life requires that the appropriate program of research be defined and undertaken, with any delay in program initiation resulting in, at best, a similar delay in realization of the benefits and, at worst, a total loss of the opportunity to other more willing or better able to take the associated risks. The networking issues which must be addressed by such a program of research are described in the remainder of this chapter.

B.1 Reconfigurability

Optical networks offer the flexibility of adapting their topology in response to changing network traffic patterns. Given some physical topology, it is possible to construct a “virtual” or “logical” topology over it by proper assignments of wavelengths between source-destination pairs. The logical topology can be adapted, for example, by simply retuning transmitters and receivers at the end-nodes in the network. The reconfiguration can be performed over different scales, for example, on a connection-by-connection basis, or somewhat more infrequently in response to gross changes in traffic patterns.

At one end of the spectrum, one could set up logical links on a call-by-call basis. However, the real-time implementation compromises associated with such an approach may result in a highly non-optimal network topology. Thus, we need metrics to determine whether a new call should be routed on the existing logical topology, or whether a new logical link should be created for this purpose.

At the other end of the spectrum, one could set up all the desired logical links and route all calls over these links until some gross change in traffic pattern is observed. Then, an appropriate algorithm could be invoked to reconfigure the network from the current logical topology to a more optimal logical topology while causing minimal disruption to calls already in progress.

Given that arbitrary virtual topologies can be embedded on a given physical topology, the possibilities are enormous. For a given traffic demand, which topology is the optimal one? What are the objective functions to be optimized (e.g., delay, carried traffic, etc.)? Can we prove that an optimal solution can be found in polynomial time? If not, heuristic algorithms which provide sub-optimal solutions need to be found, and the quality of these solutions, relative to optimal, must be analyzed, if possible. In addition, the interplay, if any, between physical topology optimization and virtual topology optimization must be investigated.

When the prevailing traffic pattern changes, mechanisms by which network nodes learn about such changes should be investigated. How can network nodes use this information to adapt (reconfigure) the virtual topology in real time to maintain optimality? Moreover, how can we reconfigure the topology so that existing connections need not be re-routed? How can such systems accommodate both packet and circuit traffic, e.g., how can "clear-channel" circuits be set up on demand?

B.2 Integrated Packet and Circuit Switched Networks

Switched communication networks fall within two main classes: circuit switched and packet switched. While telecommunication networks have evolved along circuit-switching lines, the data processing world, with the desire to share expensive computer resources, prompted the establishment of packet-switched data communication networks. The current move is towards integration, with the goal to build a single network that will provide voice, video, image and data services to everybody. ISDN and Broadband ISDN (B-ISDN) are electronically-based switched network architectures targeted to provide services consisting of the transfer of voice, data, image and video information in a single system.

In circuit switching, or virtual circuit switching, a complete end-to-end path between the communicating entities is established for the duration of the connection. Resources on this path can be made available for the duration of the connection. The end-to-end path thus typically can provide a constant bit rate, and guarantee a continuous flow of information. Telephone networks have traditionally been circuit switched; each telephone call is given a fixed 4 KHz or 64 Kbit/sec. channel. Since data users (computers) send bursts of information, it was recognized early on that sharing resources between several end users at all system and protocol levels is a much more efficient and appropriate way to manage a data communication network. Packet switching realizes the notion of statistical multiplexing where information is collected into smaller units, packets, which can be switched inside the network as necessary resources become available.

In an optical networking scenario, the situation grows more complex, as the allocation of optical resources differs significantly from the situation in the electronic world. The optical network is basically a circuit-switched domain in which each circuit affords enormous bandwidth, but the optical analogy to electronic packet switching is almost totally absent. In many of the underlying optical technologies, the mode in which different channels can be accessed further limits system flexibility. In a WDM system, for instance, packet switching has been considered by means of addressing each packet with a specific wavelength. This requires rapid wavelength agility, but the tuning speed may impose large overhead on switching between channels. Virtual circuit switching also faces potentially severe constraints. For instance, the number of wavelengths necessary to connect all pairs of users makes a passive switched optical network impractical for all but very small systems. The number of transceivers at intermediate nodes poses an additional limitation on the number of virtual circuits that can be established concurrently. Multihop networks superimpose an electronic packet-switched architecture over the inherent circuit-switched optical infrastructure, enabling full connectivity at the virtual circuit level but not exploiting the advent of rapid, packet-by-packet wavelength agility over a limited spectral range. Integration of reservation-based time-multiplexed circuit switching with virtual connection-oriented packet switching has not yet been considered for multihop networks which today can, at best, integrate clear-channel circuit switching of enormous bandwidth with virtual connection packet switching (see, for example, Section 5.1), but cannot accommodate circuit switching of sub-rate channels requiring less bandwidth than that provided by a clear optical path. Furthermore, reprovisioning of a circuit-switched clear channel requires total network reconfiguration which true circuit/packet switching integration should seek to avoid.

The need to integrate different applications dictates a parallel integration of packet and circuit switching on the same network. Except for the clear channel possibility mentioned above, the integration of circuit and packet switching in an optical networking context has been virtually unexplored.

B.3 Scalability and Modularity

For practical applications, any network architecture must be scalable and modular. By scalable, we mean that we must always be able to add and remove nodes from the network. By modular, we mean that we must be able to add or remove just one node at a time. This is important because networks may begin to operate with a few nodes and grow with time to very large configurations. For example, many regular topologies require that the number of nodes be always an integral power of a constant number. These topologies are not modular. However, regular topologies offer certain advantages, such as easy routing and high throughputs under uniform loading, and an interesting area of research is the design of regular structures that can be reasonably scalable and modular. In general, architectures must be scalable until physical resources are fully utilized; for example, existing tunable devices may limit the number of accessible channels to a few 10s but, when better tunable devices come along, it will be desirable if they can simply replace existing ones without requiring that other network elements also be replaced. However, it will not be possible to add more channels beyond the constrained physical bandwidth limit. Alternatively, optical channel re-use techniques afforded by the inclusion of electronically-controlled wavelength-selective optical switches will reduce the effort of this physical bandwidth limit and enhance network scalability.

B.4 Local, Metropolitan and Wide Area Optical Communication Networks

Communication networks are categorized in terms of their geographical span into local, metropolitan and wide area networks. Local area networks (LANs) are typically operated by a single organization and span distances of a few kilometers. Metropolitan area networks (MANs) span distances of around 50-100 km and serve as backbone networks interconnecting LANs and other high-speed information sources and sinks over a regional domain. MANs also serve as access networks into wide area backbone networks. Wide area networks (WANs) serve to interconnect stations, LANs and MANs over long distances. Local area networks provide for the high-speed interconnection of computer communications bursty packet streams as well as for integrated support of multimedia isochronous oriented streams. Metropolitan and wide area networks provide for the integrated support of synchronous and asynchronous services while also accommodating high-intensity multiplexed streams.

Optical LANs, MANs and WANs should be designed to provide high throughput and acceptable latencies for multimedia integrated services applications. Services to be supported should include communications transport for evolving applications for telecommunications and computer communications; distributed computing, operating systems and data base applications; medical imaging and radiology-based applications; massively parallel computing, supercomputing and processing, and processor interconnections; and distributed C^3 applications.

The key components of a communications network include the transmission, switching and processing systems. Optical communication networks provide opportunities for effectively tapping into the vast bandwidth embedded within the optical fiber link. The design of an effective optical communications network must thus provide realizable trade-offs between the efficiency attained from the sharing of the network's optical communications bandwidth resources with the associated complexity required of the (electronically and optically based) switching, process and transmission resources. A key role in this trade-off is played by the optical channel architectures used. The effectiveness and feasibility of these structures jointly depend upon the efficiency and responsiveness of the selected network system architecture and protocols and on the efficiency and implementability of the associated optical, or hybrid electrical/optical, devices.

Another aspect of optical LANs, MANs and WANS is that of public versus private networks. A key issue that will determine the geographic scope of private optical networks is the availability of “dark” fiber. Dark fiber refers to fiber provided to a user, where the user is allowed to send signals using any format (bit rate, modulation format, total bandwidth) he/she chooses. While 50 percent of the installed fiber base today is currently dark, most of it is owned by the common carriers, who may be unwilling to provide it to private customers. Without dark fiber, private optical networks will be limited to LANs and MANs.

Any optical network architecture must also consider power budget constraints. The power available at a transmitter is limited, and the optical signal undergoes attenuation and splitting losses within the network. Reliable optical reception requires a minimum received optical power, which then limits the total allowable losses within the network. Optical amplifiers will be required to boost the power budget but their use must be studied carefully. Optimal placement of amplifiers will have a major impact on performance. Cascading amplifiers with non-flat spectral shapes will result in a significant reduction in the available optical bandwidth, and we must find ways to overcome these potential problems. Other physical layer considerations include factors such as channel bandwidth and spacing for WDM systems, and time slot durations and guard bands for TDM systems.

B.5 Multiplexing for Optical Networks

The sharing of optical fiber (and, as a result, the consequent sharing of other transmission, switching and processing resources) can be accomplished along the frequency, time, space and code dimensions.

1. In sharing the frequency resource of the optical link, we employ Wavelength Division Multiplexing (WDM) or Multiple-Access (WDMA) schemes. Multiple streams are transmitted across the link, simultaneously in time, over distinct wavelength channels. Subcarrier multiplexed (SCM) methods can also be employed in multiplexing multiple RF carriers over a single optical carrier.

2. Time Division Multiplexing (TDM) or Multiple-Access (TDMA) methods are used to share the optical links (and other transmission and processing resources) on a time division basis. Packet and stream transmissions across a TDM optical channel proceed in a time sequential manner.

3. Space Division Multiplexing (SDM) and Multiple-Access (SDMA) techniques are employed to share multiple fiber link groups, as well as to provide for spatial re-use along a single optical link (or succession of links).

4. Code Division Multiplexing (CDM) or Multiple-Access (CDMA) schemes are used for the time simultaneous sharing of a single optical channel through the application of encoding/decoding operations to each traffic stream.

5. Hybrid multiplexing and multiple-access methods are employed to devise an access-control procedure which integrates the sharing of multiple link resources. For example, hybrid wavelength/time/space multiple access methods combine the use of multiple wavelengths, time division and spatial diversity techniques.

A few WDM architectures have been explored recently; the limitations and possibilities of these are not yet fully understood. There is a lot of potential in developing new innovative architectures using WDM, and WDM in combination with other forms of multiplexing, and evaluating their capabilities.

The broadcast-and-select architecture is among the simplest considered so far. Here, a transmission from a node is broadcast to all other network nodes by a passive (except for optical amplification) optical fabric and some form of wavelength selectivity (tuning) is required either at the transmitters, receivers, or both. Thus, this falls into the category of passive optical networks.

At the systems level, efficient multi-access protocols are required to resolve contention, which occurs when two or more nodes transmit to a single destination on different channels concurrently, as well as collision, which occurs when two or more nodes transmit on the same channel concurrently. Efficient schemes to synchronize all the network nodes to a common frame clock will greatly aid the design of efficient multi-access protocols.

Networks using this architecture are critically dependent on tunable devices. Transmitters (lasers) and receivers (filters) capable of rapidly tuning over a large number of channels are essential to realize packet-switched networks. Optical amplifiers are required to recover from splitting and attenuation losses occurring when each transmission is broadcast to every node in the network. The passive optical fabric is commonly realized using star couplers, which are currently bulk devices made from cascading 2 x 2 elements. Cheap integrated-optic star couplers must be developed.

The two main limitations of the broadcast-and-select architecture are the lack of wavelength re-use and the splitting loss. These two factors are likely to limit this architecture to LANs and MANs. WAN architectures must find ways to re-use wavelengths and avoid splitting loss problems. A potential solution is to use wavelength routers in combination with switching inside the network. A number of important topics in this area are yet to be addressed, starting with a fundamental understanding of the network complexity (switches, wavelengths, wavelength converters--wavelength converters are devices that can convert an input data stream at a wavelength to an output data stream at a different wavelength) required to support a given traffic requirement among a given number of users, and how different dimensions (switches, wavelengths) trade off against each other.

To start with, the optical switches within the network may be electronically controlled and slowly reconfigurable, perhaps on a circuit-by-circuit basis, analogous to the cross-connect switches in the telephone network today. This is because the technology required to decode optical headers and dynamically switch packets is yet to be developed. On a longer-term basis, it may be possible for the switches in the network to decode packet headers optically and switch packets optically. The lack of optical buffering may require us to use deflection routing, rather than store-and-forward routing, in such networks. The relative performance of these schemes must be compared.

Call admission and routing policies must be developed. For example, a new call may be routed on the existing network, or routed on a new logical link created for this purpose by setting up a new path across the network using a certain wavelength. Distributed algorithms for routing and topology maintenance must be developed.

TDM has been used in computer networks since their inception. However, trying to use conventional electronic TDM results in an electronic “bottleneck” in high-speed networks. The basic problem is that all the electronic processing has to be done at the aggregate bit rate of the network, and not at the individual per-node bit rates. However, optical TDM using narrow light pulses offers the promise of being able to achieve very high bit-rate transmission; bit durations can be as small as picoseconds or lower.

In combination with appropriate devices, optical TDM can be used to build very high capacity networks in the long term. Several important problems must be solved before this potential can be realized. First, compact sources for generating very short pulses must be developed. Second, all-optical ways of multiplexing and demultiplexing data by individual nodes from a TDM stream at such high speeds must be developed, since electronics just simply cannot do the job. Moreover, the high transmission rates may require very accurate network-wide synchronization.

Multiplexers might be implemented by accurate optical delay lines. Demultiplexers might also be implemented using delay lines along with optical correlators. Experiments have already demonstrated the feasibility of these devices but their true potential remains to be explored; also, most of these devices are just beginning to be demonstrated in research laboratories. Finally, it is likely that even in the long term, data will be generated and received electronically, and that will limit per-node bit rates to the order of a few gigabits per second. In order to make use of the very high optical transmission speeds, nodes must be able to compress their data to match the optical transmission rate. For example, if data is generated at 1 Gb/s for a 10 Gb/s optical link, then a compression ratio of 10 is required.

For time-division local and metropolitan area networks, broadcast-and-select or spatial-reuse broadcast-and-select multiple-access techniques can be used. Such schemes utilize the regularity of an implemented network topology to reduce the complexity of the routing and switching functions. Regular topologies include stars, buses, rings, trees, structured grids and their interconnections.

Access of stations to the network for the transmission of their packets and streams are determined by multiple-access architectures which correspond to the link-sharing multiplexing architectures described above, noting that under the spatial-reuse broadcast-and-select structure the regular topology can be regarded as a shared logical link. Hence, included are the architectures (see the discussion above for their corresponding system and device features):

1. Centralized slotting structures with fixed or demand-assigned slot allocations. An in-band or out-of-band signaling channel is established for the transmission of request/reservation packets.

2. Distributed slotting architectures.

3. Hybrid centralized and distributed slotting structures, implementing integrated circuit and packet switching disciplines.

4. Buffer-insertion networks.

5. Limited-buffering switching architectures over regular topologies.

6. Hybrid time/wavelength/space division networks. Time division networking methods are employed over allocated wavelength sets and space segments.

Other multiplexing and multiple-access techniques, such as those based on modulation formats and power splitting configurations, and their use within a hybrid architecture, are also of interest.

For time-division based wide-area networks, mesh topologies are implemented. The topological overlay can correspond to a regular structure or involve an arbitrary configuration. As a result, while the sharing of the optical links (or logical link subnets) is carried out through the use of the multiplexing (or multiple access) techniques (1) - (6) mentioned above, switching modules are required to provide for transport of information among links. Circuit and packet switching systems, and their hybrids, should be considered.

1. Optical circuit switches are required for the implementation of a time-division circuit-switching optical network. Such a switch operates to establish end-to-end paths across the {time x space} domain, or across the {time x space x wavelength} domain when hybrid WDMA/TDMA or WDMA/TDMA/SDMA architectures are implemented. A separate signaling network, electronically or optically based, is used to provide for the transmission of the control information used to establish and take-off the required circuit connections. Circuit allocation algorithms must be devised to utilize efficiently the network's transmission, switching and processing resources.

2. For the implementation of time-division packet-switched mesh optical wide-area networks, optical packet switches are required. All-optical self-routing fabrics, or combined electrical/optical fabrics which employ electrically-based control subsystems should be considered.

For the implementation of wide-area networks which assume a regular topology, the switching function can be simplified:

1. Simplified routing and switching methods can be employed (such as deflection routing, or other topological-based routing and switching strategies), leading to a simplified packet switch fabric.

2. The optical wide-area network can be constructed through the interconnection of optical local/metropolitan networks using optical bridges and routers. (See section on Optical Network Interconnection.)

In addition to medium access-control, routing and switching methods, the design of an optical communications network transport system requires the incorporation of flow, congestion, error and fairness control procedures and algorithms. Such schemes must be efficiently implemented to ensure timely delivery of information across the network at desired throughput, error-rate, blocking, availability, reliability, and other quality-of-service levels. Such control algorithms should provide for the integration of multimedia services.

Network management architectures and procedures must be implemented to manage the network's allocation of resources at acceptable performance levels to the accommodated service classes.

B.6 Time-Division Channelization and Medium Access

The time-division based channelization structure can consist of the following approaches:

1. Synchronous (or isochronous) slotting structure (local link based, or global network based). Slot boundaries are defined for each link, using centralized or distributed synchronization approaches;

2. Asynchronous temporal structure. Time boundaries are associated with packet transmissions and are derived from packet delimiter fields.

Under asynchronous (or isochronous) time slotting structure, stations gain access for the transmission of their information packets and streams in accordance with the following architectures:

I. Centralized Slotting Structure. A central or regional controller generates and allocates time slots for designated source-destination stations. We observe:

(a) a fixed-assigned TDM (synchronous TDM) access-control scheme, under which slots are allocated to stations on a fixed (or quasi-fixed basis);

(b) demand-assigned TDM access control scheme, under which slots are allocated to a requesting station upon demand.

For such systems, the station's module which provides interface and access to the optical channel must be able to recognize time slot boundaries and allocations. Stations can employ fixed wavelength optical transmitters and receivers.

II. Distributed Slotting Structure. A central or regional controller generates time slots, but does not allocate them. Slot headers are used to designate their status.

(a) idle slots are captured by a ready station for the insertion and transmission of a packet, with the access rule determined through the use of a distributed access control algorithm;

(b) busy slots, or designated subsets of busy slots, and their payloads are examined by each station's interface and access module to determine if they contain packets destined to the underlying station.

For such systems, the station's module which provides interface and access to the optical channel must be able to recognize time slot boundaries as well as read, delay and process slot and packet headers and subsequently use a logic module to control the packet removal, packet cut-through or packet insertion operations at the station's interface to the optical link. Stations can employ fixed wavelength optical transmitters and receivers. Multiple fixed sets, or agile, optical transmitters and receivers are used when such architectures are combined with wavelength-division multiplexing systems.

III. <u>Hybrid Slotting Structure</u>. This structure combines the operations used in I and II. A central or regional controller generates time slots, allocates a subset of them to designated connections, may allocate other subsets to designated stations on a receive- or transmit-only basis, and leaves the rest of the slots unallocated, to be used for distributed access control.

For such systems, the station's module which provides interface and access to the optical channel must be able to recognize time slot boundaries and allocations, as well as read, delay and process slot and packet headers and subsequently use a logic module to control packet removal, packet cut-through or packet insertion operations at the station's interface to the optical link. Multiple fixed sets, or agile, optical transmitters and receivers, are used when such architectures are combined with wavelength-division multiplexing systems.

Implementations may also effectively combine optical transmission, delay and header processing with electrical intelligent control, header processing, buffering and switching. In either case, circuit-switched, packet-switched and hybrid circuit and packet switched, under an asynchronous temporal structure, the station's optical medium interface and access module must be able to multiplex the station's packets with other station packets traversing the shared link. Methodologies include the utilization of:

IV. Buffer-insertion systems which can lead to architectures similar to those discussed above for slotted systems (II).

V. Limited-buffering switching systems, which capitalize on unique topological features to minimize packet buffering.

VI. Buffered switching systems.

VII. Hybrid time/wavelength/space multiplexing.

Implementations may use all-optical operations or effectively combine optical transmission, delay, switching and header processing with electrical intelligent control, header processing, buffering and switching. When wavelength division methods are incorporated, related devices (including multiple or agile transmitters and receivers and optical filters and wavelength routers) are required. Optical loop delay lines and switches have been proposed and need to be explored further.

B.7 Single-Hop vs. Multihop Systems

A local optical network may be constructed by employing a broadcast-and-select architecture in which the end-users are connected to a broadcast hub (e.g., a passive star coupler) via two-way fibers. Two major approaches have been proposed based on transceiver tunability speeds. In the first, fast tunable transceivers are employed and the wavelength assignments are performed on a per-packet basis, resulting in a single-hop network. In the second, fixed or slowly-tunable transceivers are employed, resulting in a multihop virtual topology (with switching between hops).

An important question that needs to be investigated is the following: Are single-hop and multihop architectures comparable in a meaningful fashion? In a single-hop network, the nodes compete for two resources--a wavelength channel and the potential destination. The corresponding protocols have to account for the transceiver tuning time penalties, limited transceiver ranges, transmitter-receiver coordination overhead, and the fact that channel feedback may be delayed due to large channel propagation delays. The design of multihop virtual topologies, on the other hand, has to consider the trade-off between the system's average end-to-end delay (which is a function of the average number of hops and the switching time required at each intermediate hop) and the system cost in terms of the number of transceivers at each node. These designs must accommodate the fact that the offered traffic may be asymmetric and time-varying.

B.8 Evolution

When a communication network is being deployed for services, there are five general aspects to consider with regard to how the network evolves in time. The first is how the network interworks with other networks and systems already in existence. This is the compatibility issue. The second is how the network expands itself for further capabilities such as providing higher bit-rate capacity, supporting more users or nodes, and interconnecting to one another. This is the scalability issue. The third is how the network protects itself from various kinds of faults, both accidental and deliberate. This is the reliability issue. The fourth is how the network operates, which includes provisioning and management of services, servicing, accounting and billing. These are the operational issues. The fifth is the strategy for the deployment of the network. The considerations are the capital investment, pricing of services, and market penetration. These are the deployment issues.

In the case of public commercial networks, the deployment issues can be overwhelmingly important because of the size of capital investment. All major planned public provisioning, such as fiber-in-the-loop (or fiber-to-the-home, fiber-to-the-curb, etc.) and Broadband-ISDN, have been facing this difficult deployment problem.

For optical networks, the above evolutionary issues are even more difficult than those for traditional networks that are based on electronic processing and switching. There are three reasons.

First, current optical technology is far less mature than electronic technology. This means that a direct translation of network architecture and functions of an electronic network to optics may not be feasible for a long time to come. Furthermore, the lack of optical memory and the larger size of optical devices as compared to electronic devices (from quantum mechanical arguments) mean that it may be more promising to use optical technology in ways different from those of electronic technology. How to exploit the advantages of optics versus electronics in relation to network evolution should be studied and not assumed.

Second, many optical network architectural proposals are very different from traditional networks. For example, many single-hop networks employ distributed processing without any centralized switching mechanism, and look more like small computer networks than large telecommunication networks. How can these networks evolve to become general public networks, or should these networks be confined to small sizes for specialized applications? These questions require more in-depth studies of optical network architecture and design.

Third, the technologies used in optical networks are considerably different from those of electronic networks. For example, optical fibers allow multiple broadband channels to be operated in parallel. How should the optical spectrum be allocated? How many channels should be allowed? Can more channels be added when the network evolves? How will that affect the spacing of optical amplifiers? If tunable devices are used in the network, can they support the newly-added channels without being replaced? These technological factors must be considered in a network design.

Another important development that affects network evolution is the multimedia revolution. This trend has been observed in both the computer community and the telecommunications community. For example, the Broadband-ISDN proposal, the rapid commercialization of CD-ROMs and multimedia workstations, and the proliferation of low-cost video equipment all point to the convergence of the multimedia era. One effect of this multimedia revolution is that the traditional boundary between computer networks and telecommunication networks is becoming blurred. A more serious effect is that most operational networks are incapable of satisfying real-time multimedia requirements. The affected networks include the popular LANs, the Internet, and the circuit-switched telephone networks.

How might future optical networks evolve to support real-time multimedia communications and applications? This question relates also to the issue of services integration which confronts traditional network designers as well as optical network designers. Should optical networks support ATM (the asynchronous transfer mode), which is so popularly accepted by the traditional networking community as the paradigm for services integration, or should services be provided on separate, physically distinct layers? Can the network management and congestion control problems associated with the ATM paradigm be dealt with by optical networks? Could optical networks be free of these problems? These are important questions to be addressed.

Currently, most optical network architectural studies have been focused on the performance aspects of the network, and the above equally important and practical questions concerning the evolutionary aspects have generally been ignored. It is recommended that future architectural studies should thoroughly address these important issues also.

B.9 WDM for Multiple Virtual Networks

An important advantage of optical networks is their ability to create and support several virtual subnetworks operating over the same physical network, with each virtual network occupying a given unique waveband. Wavelength-division multiplexing would be used within each waveband, and all wavebands would then be multiplexed onto a common physical medium. Examples of how this feature might be utilized are

1. to dynamically allocate the resources (i.e., subnetworks) to meet the changing traffic demands;

2. to support both circuit- and packet-switched traffic;

3. to support simultaneously a variety of applications (such as data, entertainment video, video teleconferencing, and voice) with different requirements in BER, latency, holding time, etc.;

4. to avoid the difficulties associated with traffic control and performance management arising when a multiplicity of traffic types are fully integrated onto the logically same network.

Innovative research at all levels (devices, systems, and networks) is needed to find ways to create, support, and adapt multiple virtual networks operating over a single optical network. Experimental demonstration of these ideas even on a limited scale is highly desirable.

B.10 Traffic Control and Performance Management

Traditionally, the concerns of traffic engineering have been to make efficient use of network resources and to predict the quality of service experienced by the network user. In spite of the large fiber bandwidth at their disposal, these remain areas of concern to optical networks. Optical networks must face the electronic bottleneck imposed by electrical transceivers in the system; hence, efficient use of this transceiver hardware is still a concern. If optical networks are to serve large numbers of users, switching (all-optical, fast-packet switching, or other) is needed, in which case, load management of switch contention also becomes important.

While the efficiency and performance of a network is strongly influenced by the network design, traffic control during operation will often determine the manner in which user service quality degrades as network load increases. To maintain acceptable service quality, we must understand the behavior of key performance metrics as functions of load in optical networks. This is an old network problem but there are reasons to rethink our old results and methodologies.

In an attempt to minimize the effect of electronic bottlenecks, many conceptual designs of optical networks exclude store-and-forward nodes either by minimizing storage and processing or by maintaining the signal in optical form from end to end. Thus, their end-to-end delay characteristics change accordingly. Many congestion prevention and control rules based on delay need to be revisited. The appropriate metric on which to base these rules is itself an issue.

Moreover, it is possible to provide store-and-forward capability to all-optical networks without sophisticated optical processing by using optical delay lines. Deflection routing is an example of essentially this same principle. However, the behavior of optical fiber as a storage device has different characteristics than electronic buffer queues. Also, we cannot rule out the emergence of novel optical technology suitable for network storage. So, even if optical store-and-forward principles apply, it is not clear that the current models are appropriate.

An important consideration for wide area networks or very high rate networks (networks with a high rate-delay product) is the fact that end-to-end network delay can be orders of magnitude larger than the basic data units in the system (e.g., packet duration). At the same time, some high rate computer applications are expected to exhibit bursty behavior, with temporal correlation lengths much smaller than the end-to-end delay of such networks. This means that many control feedback assumptions of traditional networks no longer hold true in this scenario since the feedback delay is much longer than the lifetime of events to be controlled. Moreover, the potentially large buffer size that would be required to operate with the feedback delays in networks of high rate-delay products threatens to complicate access interfaces to these networks.

The larger bandwidth available through optical interfaces is fostering development of new computer applications like network-distributed computing and visualization. At the moment it is not clear what the statistical behavior of traffic from such applications will be like; hence, current traffic models may also prove inadequate to characterize applications of the future.

That these and other traditional assumptions need rethinking is mainly a consequence of the dramatic changes provided by optical networks. Not surprisingly, there is no shortage of research work required to address these issues. There is ongoing work on defining performance metrics for optical networks. This includes metrics related to transceiver utilization, which serve as the basis for load balancing work in wavelength and space division systems or blocking probabilities, among others. It includes switch contention metrics not only for blocking and loss characterization but also for delay performance metrics such as hop count characterizations of multihop topologies. There are results on the access delay achievable by scheduling algorithms for networks based on reservation. Because several of these are worst-case results or topology-dependent simulation results, work is needed to enhance or complete the picture.

The performance of novel multiple access architectures also requires research, to be able to weigh the relative merits of various approaches. Besides the performance metrics mentioned above, CDMA systems suitable for optical networks call for bit error performance studies. Many versions of CDMA systems have been proposed based on pulse position, center frequency wander, and optical signal parameters. Many despreading alternatives have been proposed. Their error performance is a function of network load and, hence, a potential factor in traffic control.

These, among other examples, show the breadth of factors that are relevant to traffic control. The high-rate delay product of many optical network architectures has attracted work on bandwidth reservation and management algorithms, with emphasis on scheduling problems and global state knowledge of the system by the nodes involved. Load control in multiple access architectures has perhaps received less attention, with results borrowed from traditional multi-access analysis.

For progress to be made in the area of traffic control and performance management, it will be important to bear in mind the overall network architecture and the changing assumptions due to the new technology. It will be necessary to explicitly identify the relationship between the optical technology and the traffic performance to understand how best to use the network.

We need to understand the theoretical limitations imposed by end-to-end feedback in networks of high rate-delay product. We must also understand, from the systems and the interface hardware points of view, the problems involved in buffering large amounts of data at the edges of the network while waiting for this feedback. At the architectural level, networks with no store-and-forward nodes pose questions regarding the role of congestion control. Without storage at intermediate nodes, are the roles of congestion control and channel access control essentially the same, or should there still be two different mechanisms operating at different time scales, so that one operates without feedback and the other with feedback?

Similarly, we should understand and quantify the advantages or disadvantages that store-and-forward operation would bring to an optical network. This would help to guide research directions in optical delay line storage technology for store-and-forward node operation and research in other alternative technologies.

Much of the work on traffic control needs to rely on good characterization of the network performance parameters, as discussed before, but, equally important are good characterization of emerging traffic. Poisson traffic models may prove too simplistic for effective predictions of the statistical behavior of emerging traffic. While work is underway on models for multiplexing of circuit connections, these models are still heavily qualified with assumptions. Without better understanding of this traffic behavior, we will be unable to effectively design the control mechanisms to manage the network load and will find it hard to repeat past successes of traffic engineering. As is often the case, the research required will call for analytical work, simulations, and experimental prototypes. It would be dangerous to preclude a priori any approaches or ideas in this field.

B.12 Fault Management

In general, fault management is that area of network management that is concerned with detecting, isolating, and recovering from network faults. It consists of the provisioning of redundant, back-up facilities and the ability to detect failures and switch over to the back-up in case of failure. Optical networks pose new challenges in fault management mainly because many of the optical network implementations are "passive," that is, no processing is done on the data while it is transported inside the network. For example, while in conventional networks such as FDDI and ATM a link failure is detected and reported by the nodes at either end of the link, in passive optical networks the optical nodes have no processing capabilities. Thus, other fault detection techniques must be used. Another difficulty of fault management in passive optical networks is that they employ entirely new families of devices whose failure modes and reliability are not well characterized at this time. Thus, the well-developed fault management techniques used in digital electronic networks are not necessarily applicable to passive optical networks. For the purpose of identifying research issues in this field, it is helpful to consider and evaluate three separate classes of optical networks: passive broadcast networks, wavelength routing networks, and optical processing networks.

Let us first consider the passive broadcast network. The typical topology is a star or tree. Internally there are no active processing elements. Optical amplifiers (which are active, but non-processing components) may be installed within the network, generally at the center of the star or at the root of the tree. Amplifier back-up strategies, with fault monitoring and switchover, must be used; these strategies can probably be implemented using conventional techniques. A more challenging problem, on the other hand, is the recovery from fiber cuts, which are a very real threat in metropolitan area networks. Here it is important to detect and locate the cut in a fiber plant which often spans several hundreds of miles so that repairs can be promptly scheduled. The location of the cut is in general not easy to determine since the internal nodes (i.e., splitters, couplers) have no link failure detection capabilities. An important research area is thus the development of topology redundancy strategies (e.g., overlaid trees in a tree topology) which permit the switchover to a back-up topology if the primary topology is damaged. The back-up strategy should provide smooth (ideally, user-transparent) switchover; it must also permit the fault manager to identify the location of the cut in order to schedule the repair.

Other types of faults which must be addressed in passive broadcast networks are those associated with the end stations (e.g., a station which “babbles,” or lets its laser drift, or gets out of step in a T/WDMA access protocol, or intentionally jams one or more WDM channels). In this case the partitioning of stations into physically separate clusters (as in a multi-level structure) or into different fiber plants (in a multifiber cable arrangement) may provide a possible solution. Another strategy that can be used effectively in passive networks is virtual topology reconfiguration. Once the faulty station/router is isolated and the wavelengths affected by the failure are identified, a new virtual topology can be deployed, excluding the faulty elements.

Wide-area wavelength routing networks present all the fault management problems already identified in passive broadcast networks. In addition, the nodal backup problem is more acute here in that the node contains much more active equipment (tunable filters, switches, amplifiers, etc.) than in broadcast networks. The problem of fiber cut detection is also more complex, due to multipath routing. Topological redundancy, on the other hand, is implicit in the mesh topology and therefore does not involve a separate back-up facility. The challenge is to develop a strategy (either centralized or distributed) which permits identification and isolation of link failure (based on end-to-end user failure reports), and to switch users affected by such failure over to an alternate path.

Optical processing networks typically perform optical, fast packet switching by virtue of header inspection at each intermediate node. If nodes have processing capability and, therefore, intelligence, they can then engage in link monitoring. For the optical networks in which this assumption is valid, conventional fault management schemes may then become feasible. However, given the very broad range of optical processing networks with widely differing characteristics and processing capabilities, it is safe to say that in most cases novel fault management schemes must be sought.

In summary, novel research in fault management of optical networks is required in order to make sure that these networks are at least as robust as conventional networks.

B.12 Optical Power Budget

The analysis of power budget plays an important role in the determination of network capacity (e.g., bitrate x number of users). The budget is determined by three factors: the power available at the transmitting ends, the sensitivity achievable at the receiving ends, and the overall maximum loss incurred over the optical transmission path. By using optical amplifiers (OAs) the power budget can be drastically enhanced. First, OAs can be used as power boosters at the transmitting ends. Second, they can be used as loss compensators whether loss is caused by transmission distance, splices, splitting or optical components. Finally, OAs can be used as preamplifiers, which considerably enhances receiver sensitivity. These three factors make it possible to increase both the number of users and the network size to extents that have not yet been fully explored.

A primary limitation, however, is the spontaneous emission noise generated at each amplification stage. As this noise accumulates throughout the optical transmission medium, a degradation of signal-to-noise ratio (SNR) results. Therefore, a full SNR analysis must be performed in addition to that of power budget analysis. Given the transmitter power and receiver sensitivity, the SNR and power budget analyses eventually define a range of bit rate and number of end-users, which correspond to the network capacity. Optimum amplifier locations must be determined in order to maximize this capacity.

Power budget is also affected by network architecture and topology. In fact, the optimal amplification strategy is influenced by the network characteristics. For the purpose of outlining different amplification strategies, it is convenient to group optical networks in three broad classes: (a) local and metropolitan passive optic networks implemented on a broadcast medium; (b) wide area networks based on the wavelength routing scheme; and (c) optical processing networks which, for example, require packet header inspection and processing at intermediate nodes.

Local and metropolitan PONs are generally based on a star or tree topology or a combination of the two. If amplification is required (as is usually the case for a large number of stations), the amplifier can be placed at the center of the star or at the root of the tree. Several stages of amplification may be required for very large networks (for example, amplifiers may be placed between the stages of a modular star coupler, or at different levels of the tree topology). An interesting option to explore is the possibility of driving the optical amplifiers with a remote (rather than local) pump in order to keep the network truly passive.

It should be pointed out that power budget in broadcast networks can be enhanced (and, therefore, the need for optical amplification reduced) by using multifiber cables, and thereby implementing an optical tree network (for example) which is the superposition of several trees, each on a separate fiber plant. By inserting and properly setting optomechanical switches and star couplers at the intermediate nodes of the tree structure one can create several "embedded" subtrees such that, given two arbitrary nodes, there is always a subtree which interconnects them. From the power budget standpoint the advantage of this scheme is that subtrees have fewer levels than the original tree and, therefore, have lower loss and may require less amplification (or no amplification at all). Power budget considerations also influence the choice of the LAN and MAN topology. From the point of view of power budget, it is clear that linear topologies are less desirable than tree or star topologies since the former require many more stages of amplification to support the same number of stations.

In wide area wavelength routing networks, the physical topology is typically that of a mesh. Thus, different from LAN and MAN optical networks, there is no single, obvious place (e.g., the root of the tree) where the amplifier can be installed. Assuming that the nodes of the wavelength routing network are passive (i.e., filters, couplers, switches), one can expect that each node introduces a substantial power loss, which is dependent on the nodal degree (i.e., number of links connected to this node). In addition, long links introduce a loss which is comparable to that of the nodes. Depending on node/link loss, amplifiers may be required at all outputs of each node, or may be installed only at some of the nodes/links. The location of amplifiers in the network must be properly chosen for optimal efficiency. The installation strategy should also account for possible link/node failures and thus should incorporate redundancy. Another area for further study is the interplay between physical topology design and amplifier allocation.

For optical processing networks, the need for power amplification depends on the particular architecture under consideration. Some systems process (i.e., inspect and regenerate) the header (or control) field of the packet at each intermediate node and transfer the data field transparently. Thus, if the network path traverses several nodes, amplification of the data field is required. In this respect the optimal allocation of amplifiers (i.e., how many to install and where to install them) is an open research issue, similar to the amplifier allocation strategy in wavelength routing networks.

Power budget enhancement strategies depend on progress in the following device technologies: flat broadband optical amplifiers, more powerful sources, and more sensitive detectors.

Further enhancement in power budget can be obtained with the use of coherent modulation/ demodulation techniques.

B.13 Physical Topology Optimization

One of the unique advantages of WDM optical networks is that of allowing the establishment of arbitrary, “virtual” subnetworks within the basic physical topology. Virtual stars, rings, trees, grids and shuffle networks, to mention a few, can be embedded in any broadcast, passive topology. Since access protocol design and performance are mostly influenced by the virtual topology, it may seem that physical topology does not play a major role in the overall design. This, however, is not true because, as we shall see, the choice of topology has a critical impact on power budget, fault tolerance and fiber layout cost. Furthermore, in mesh topologies (which are found in wavelength routing networks and optical processing networks), the physical layout also impacts on protocol performance.

In order to identify some of the research issues posed by physical topology design it behooves us to once again classify optical networks in three categories: passive broadcast networks, wavelength routing networks, and optical processing networks. We shall discuss each category separately. At the end, we also consider the use of multifiber cables and its impact on topology design.

Starting with metropolitan broadcast passive networks, the topology of choice in this case is the star or the tree. The particular choice (star or tree) and the physical layout are determined by network size, line costs, optical power budget and growth flexibility. The design must also include topology redundancy to recover from fiber cuts. There have been proposals of multi-level physical tree or star with wavelength-sensitive filters which isolate low-level clusters and allow wavelength re-use. In this case the topology design must also include cluster partitioning (based on traffic pattern and geographical proximity criteria).

In wide-area wavelength-routing networks, the physical topology is that of a mesh. Wavelength routing permits us to exploit the multiple paths offered by the mesh topology. Thus, physical topology design has direct impact on routing and, therefore, on performance. In these networks there is clearly a much closer coupling between topology and protocols than in passive broadcast networks. In addition to protocol performance, the topological design of wavelength routing networks must also consider the previously mentioned issues of power budget, fault tolerance and layout cost.

Optical processing networks can route packets at each intermediate node based on control information carried in the header. Thus, very fast optical packet-switched operations which exploit the multiple paths offered by the mesh topology can be supported. As with wavelength routing networks, in optical processing networks the physical topology has direct impact on performance. Topological optimization is driven by throughput and delay criteria as well as by power budget, fault tolerance and cost considerations.

So far we have assumed that the fiber plant consists of a single fiber on which several wavelengths can be multiplexed. It is clear, however, that since a typical cable contains dozens or even hundreds of fibers the resulting plant is a de facto superposition of several fiber plants. This multifiber arrangement can be exploited in various ways in the design of optical networks.

First, multiple fibers can make up for the scarcity of wavelengths available within a single fiber; by combining WDM and SDM (space division multiplexing) across multiple fibers we can increase the number of available channels by perhaps two orders of magnitude. It should also be pointed out that space division switching (with optomechanical or lithium niobate switches) is much cheaper, faster and less lossy than wavelength switching (with tunable lasers and filters). These multifiber advantages most directly effect the WDM networks such as passive broadcast and wavelength routing. But they can also impact optical processing network design.

Secondly, multiple fibers can be exploited to design topologies with more favorable power budgets. For instance, in large metropolitan systems with thousands of stations, a single-fiber tree topology may require multiple stages of amplification. With multifiber cables one can partition the network into clusters, each with its fiber plant (no amplification). Furthermore, intercluster communications can be supported among subsets of clusters using several small subtrees (again, no amplification required).

Finally, multiple fibers can enhance tolerance to fiber cuts by embedding diverse routed fiber plants in the same cable plant, and to fault stations by connecting different stations to different sets of fibers.

In summary, physical topology optimization is an important aspect of optical network design. It impacts not only cost but, also, throughput, fault tolerance and power budget. The availability of multifiber cables introduces opportunities which do not exist in conventional electronic network design.

Research in physical topology optimization is challenging and essential to the success of optical networking. It is also important to integrate physical topology optimization into the overall network design process. For example, choices made in physical topology design affect the options available to the virtual topology designer and the transceiver designer. Thus, an integrated network design methodology would incorporate tools for physical topology design.