RL-TR-91-132
Final Technical Report
July 1991

MULTI-GIGABIT FIBER OPTIC WIDE AREA NETWORK DEVELOPMENT

North Carolina State University

J. Ben O'Neal, Jr.; Michael R. Slawson; J. Keith Townsend

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

Rome Laboratory
Air Force Systems Command
Griffiss Air Force Base, NY 13441-5700

This report has been reviewed by the Rome Laboratory Public Affairs Office (PA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

RL-TR-91-132 has been reviewed and is approved for publication.

APPROVED:

JOHN B. EVANOWSKY
Project Engineer

FOR THE COMMANDER:

JOHN A. GRANIERO
Technical Director
Communications Directorate

If your address has changed or if you wish to be removed from the Rome Laboratory mailing list, or if the addressee is no longer employed by your organization, please notify RL(C3D) Griffiss AFB NY 13441-5700. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document require that it be returned.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	July 1991	Final Mar 89 - Aug 90

4. TITLE AND SUBTITLE
MULTI-GIGABIT FIBER OPTIC WIDE AREA NETWORK DEVELOPMENT

5. FUNDING NUMBERS
C - F30602-88-D-0025
- Task - C-9-2015
PE - 62702F
PR - 4519
TA - 22
WU - PG

6. AUTHOR(S)
J. Ben O'Neal, Jr.; Michael R. Slawson; J. Keith Townsend

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
North Carolina State University
Box 7914
Raleigh NC 27695-7914

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Rome Laboratory (C3D)
Griffiss AFB NY 13441-5700

10. SPONSORING/MONITORING AGENCY REPORT NUMBER
RL-TR-91-132

11. SUPPLEMENTARY NOTES
Rome Laboratory Project Engineer: John B. Evanowsky/C3D/(315) 330-7667

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

The development of the physical layer for a high speed (> 10 Gigabits/sec) fiber optic wide area network, implementable by the Air Force in a five to ten year timeframe is addressed. First, to gain a generally acceptable degree of survivability while maintaining efficiency, it is concluded that a counter-rotating ring physical topology is most appropriate. To obtain high link capacity, links will consist of multiple channels, each operating at data rates manageable using switching. Then, for efficient information transfer in a multi-service environment, each channel will operate with Asynchronous Transfer Mode (ATM), a fast packet switching technique. Many common logical network topologies are then overlayed onto this physical multichannel ATM ring foundation. Based on comparisons relative to several important criteria, a logical multi-ring architecture, with wavelength division multiplexing, is concluded to be the most promising alternative. An investigation of the media level implementation of this network, using the example of an eight channel link where each channel operates at 2.49 Gbps, shows that symmetrical, single-mode grating-based wavelength division demultiplexers and traveling wave optical amplifiers can be used to create bidirectional multichannel link spans exceeding 2000 Km without electronic regeneration. A simple register insertion protocol is proposed to serve as the media control mechanism for the network.

14. SUBJECT TERMS	15. NUMBER OF PAGES
Communications Network, Fiber Optic Communications Network, Optical Communications, Multi-Gigabit Networks	182
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	U/L

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF FIGURES, continued

1. Executive Summary

The coming revolution in optical-based communication networks will present two challenges to military organizations like the Air Force. The lesser challenge is one of engineering and network management. This lesser challenge is the subject of this report. It involves making the right strategic choices which will permit the Air Force to acquire the new photonic technologies for its communications systems.

The greater challenge involves the evolution of the structure of the command and control function of the Air Force to accommodate itself to the new environment provided by photonic networks. Organizations structured to take advantage of these new photonic communications networks will be superior to organizations structured without these networks in mind. In other words, the photonic revolution is so pervasive that its capabilities cannot be *fully* utilized by simply overlaying these networks on top of an existing organization. The full power of these new technologies will be harnessed only by restructuring the organization to conform to the possibilities that they provide. Peter Drucker and others[1] have outlined the characteristics of organizations structured to take full advantage of the new photonic communications systems. Such organizations are decentralized into autonomous units, have fewer levels of middle management, are less hierarchical, and rely upon decisions making processes that involve more people at many levels of the organization. These traits seem antithetical to military organizations. Nevertheless, the proper response to the new technology of optical-based networks is not to just acquire the new

---

1 See the last two references in Section 10.

technology, but rather use it along with other factors to shape the organization. This is not a new idea--technologies have always shaped organizations.

The purpose of this report is to provide information which can be used by the Air Force to acquire the new optical technologies in communications systems. Following is a list of recommendations:

Network Design

1. There will be an increase of several orders of magnitude in the communications capacity or bit rate available and the cost of this capacity will continue to decrease. The Air Force must determine how to best use this capacity to fulfill its mission.

2. There will be a host of new protocols in the commercial communications world. The first thing to do is learn these new protocols and how to use them, and determine how well they satisfy Air Force requirements. These new protocols include SONET, ISDN, DQDB, ATM, and FDDI. Other protocols will be developed. It will be wise to stay in touch with standards organizations and committees in order to be informed of new developments.

3. There will be a marked improvement in the survivability of leased facilities. This must be factored into the design of military networks that contain a mix of leased and owned facilities.

4. Self-healing rings will be widely available in common carriers and for military owned networks. The appropriate use of these rings to improve overall survivability and reliability should have high priority.

5. The commercial networks will provide security for their customers. Military networks should incorporate these commercial security features to achieve improved security at reduced cost.

Military Communications Equipment Design

1. The primary technology bottlenecks in photonic communications systems are, and will remain for the next five years, electronic speed and optical power limitations. The speed of the electronic interface components is significantly below the speed of the optical transmission medium. The relatively low optical power available in optical devices places a limitation on the capacity of optical communications systems. These two constraints will govern the performance of links and nodes. Encouraging research in these two areas should be of high priority in the Air Force.

2. Due to the high speed of transmission in modern photonic systems, the ratio of transmission time to propagation time is quite small. This means that a long transmission line may contain many messages which are in transient from transmitter to receiver. This severely complicates error control. This fact along with the extremely low error rate possible on optical facilities implies that current protocols for error control (ARQ) are inefficient. These links are somewhat like satellite links with the exception that the photonic links operate at speeds which make the forward error control coding used in satellite systems very difficult to implement. New protocols for error control should be developed for these long, high speed links.

3. The next step in the future of photonic transmission technology will be more extensive use of wave division multiplex. This will be followed by coherent communications systems. These developments should be encouraged.

When it was first created in 1969, ARPANET, conceptualized by the Advanced Research Projects Agency (ARPA), was a revolutionary high speed nationwide backbone network linking research, government, and military agencies [1]. The transmission rate between ARPANET nodes was around 50 Kbps. At that time, only the Department of Defense's slower AUTODIN (AUTOmatic Digital

Network) network existed for nationwide packetized data transmission. ARPANET has been highly successful in providing electronic mail and file transfer services.

ARPANET is now being phased out due to its low data rate and now plays a lesser role as the unclassified research and development segment of the Defense Data Network (DDN). It is part of a consortium of national networks, including classified and unclassified networks, referred to as the Internet. While currently operating at 448 Kilobits/sec (Kbps) speeds, many of the Internet networks are in the process of being upgraded to DS1 (1.544 Megabits/sec) rates. At 448 Kbps, delays in transfer of a one Megabyte file can still last five minutes to one hour, during peak traffic periods [2]. However, due to increasing usage, the conversion to DS1 rates may not dramatically affect the performance.

The Air Force will soon convert to an Integrated Services Digital Network (ISDN), based on commercial standards but altered to meet Air Force needs. Then, as high speed applications like teleconferencing (for example, for setting up tactical planning sessions) and real-time graphics (for example, for use in distributed flight or battle simulator systems) come of age; and high speed local area networks such as ANSI X3T9.5 (Fiber Distributed Data Interface) and IEEE 802.6 (DQDB) become widespread at Air Force bases; network rates will have to be dramatically increased, probably to several

Gigabits/sec (Gbps). Toward this goal, agencies like NSF and DoD are initiating research in implementing wide area networks with speeds in excess of one Gbps. The development of a three Gigabit national network has been called for by members of Congress ("National High-Performance Computer Technology Act of 1989," Senate Bill no. 1067), and the President of the United States.

This report contains our assessment of the network technology, protocols, and architectures that will be necessary to construct a multi-Gbps wide area network, that efficiently meets basic survivability requirements, within the next ten years. We conclude that fast packet switching (using Asynchronous Transfer Mode, being standardized for Broadband-ISDN) will be the preferred technique to transfer information in a multi-service environment. To efficiently obtain immunity from a single point of failure or destruction, we also conclude that a counter-rotating ring offers the best alternative for use as a physical topology. Given the physical limitations on electronic and opto-electronic switching, we believe that to obtain the multi-Gbps network rates some form of optical multiplexing will have to be used to achieve parallelism over the links connecting nodes. Thus, we begin with a fast packet-switched, multi-channel counter-rotating physical ring as our foundation upon which to build future networks.

We then define several new architectures by overlaying conventional logical topologies - like the ring, bus, perfect shuffle (or banyan), star, and fixed assignment - upon this foundation. These new architectures, with provisions incorporated to insure that the survivability features of the physical ring are maintained, are compared relative to several important network criteria, including reliability and complexity. Based on this evaluation, we conclude that, for a wide area network serving as a high speed backbone, a multi-channel logical ring architecture is the most appropriate. It allows for the smooth addition of both nodes and capacity, and is robust in the event of transmission component failures.

In similar fashion we also describe optical multiplexing techniques available for creation of the multi-channel links. We conclude that wavelength division multiplexing can most effectively be used to provide a long distance link with capacity of at least twenty Gbps within the next ten years. Cascaded traveling wave optical amplifiers and symmetrical diffraction grating-based wavelength division multiplexers can be used to construct a bidirectional link with 40 Gbps capacity, having a spacing of over 2000 kilometers between active regeneration and a maximum electronic switching rate of about 2.5 Gbps.

Regarding protocol issues above the physical layer we conclude that a register insertion protocol, using the age or urgency of the fast packet to determine its transmission priority, will provide a simple but fair network media access control protocol. At the network and transport layers, the use of a protocol like the Express Transfer Protocol, being developed in the high speed local area network arena, is recommended due to the reduction in handshaking and connection time that it allows. An overview of telecommunications standardization activities, current practices for commercial telecommunications security and a simulation study of three newly defined network architectures are provided in appendices.

2. Introduction

This report addresses the question of how a survivable high speed military wide area network can be designed while still taking advantage of the high bandwidth capability of the optical fiber medium, and keeping the significant transmission costs at a reasonable level. As such, the primary concentration will be on considerations related to the network's physical layer, as defined in the Open Systems Interconnect (OSI) model (see Section 8). A multitude of possible broadband network architectures, protocols, and topologies have been proposed for the physical layer of high speed networks. A categorization of many of these architectures has been developed by Daddis and Torng in [3]. In [4], Skov presents an overview and analysis of many of the proposed protocols, including physical layer protocols, for high speed local and metropolitan area network applications. Kaminov, in [5] assesses several fiber optic-based architectures, presumably for applications where survivability is not a major concern. Many of the architectures reviewed in these references are clearly inappropriate for a high speed survivable wide area network, generally because they have a single point of network failure and/or are not designed for a multi-service environment.

Physical layer considerations include the physical and logical topologies of the network. The first step in developing a survivable high performance wide area network is to choose the most efficient physical topology meeting the survivability requirements. While the tradeoff between survivability and efficiency (or cost) varies from application to application, a physical counter-rotating ring topology most efficiently provides a level of survivability acceptable for most applications. In Section 3 we discuss this conclusion in greater detail with a practical example involving the connection of several Air Force bases.

Since the wide area network will be implemented, at least in part, using single-mode fiber optics as the transmission media, we next proceed, in Sections 4, with an overview of fiber optic technology as it applies to high speed long-haul single-mode systems. These systems, due to their extent, are often either optical power-limited or bandwidth-limited. Important characteristics of the transmitter, fiber, and receiver are addressed from the point of view of how they limit link performance. Assessments of the current state of fiber optic technology, and research to expand the capability of the medium, are interspersed throughout the section.

Returning to physical layer architectural issues, we investigate multiplexing techniques that can be used to create multiple channels on a link without requiring an increase in electronic switching speeds in Section 5. "Multi-channel" links will become necessary as electronic multiplexing techniques become incapable of keeping pace with network capacity requirements. The optical multiplexing techniques addressed include space, code, optical time, and wavelength (or frequency) division multiplexing. These techniques are then compared relative to several important criteria such as complexity and reliability.

Section 6 begins with the important assertion that Asynchronous Transfer Mode (ATM), being standardized by the CCITT international standards body, will serve as the foundation for high-speed broadband networks of the future. Then, we overlay conventional logical network configurations upon the physical multi-channel counter-rotating ATM ring foundation. The result is that several new architectures most appropriate for future wide area network are defined. It is important to note that the physical topology of a network does not have to be equivalent to its logical topology. As a primary contribution of this report, Survivability, however, must be preserved in both the physical and logical topology. These newly conceived architectures, having equal survivability, are then compared relative to other important network criteria, including reliability, efficiency, complexity, fairness, ease of growth, and security.

Based on evaluations in Sections 5 and 6, a multi-channel logical ring network, employing wavelength division multiplexing and ATM, was chosen to be superior to the other alternatives for use as the physical layer backbone of future survivable broadband networks. In Section 7, we address the implementation of the media dependent level of the physical layer of the multi-channel ring network using fiber optic technology that has been experimentally proven. We conclude that a ring-based wide area network can be implemented within the next ten years over multi-channel, bidirectional fiber links. Key technology enablers include narrow linewidth lasers, symmetric single-mode dense wavelength division demultiplexers, and traveling wave optical amplifiers.

With many physical layer protocol issues being discussed in detail in Sections 3 through 7, we briefly introduce higher layer protocol issues applicable to the multi-channel ATM ring network in Section 8. The use of a register insertion protocol for the media access control function is introduced.

Conclusions are reiterated, and areas for future research identified, in Section 9. Section 11 is Appendix A, covering telecommunications standardization efforts in the Synchronous Optical Network (SONET), ISDN, and Broadband ISDN areas. Section 12, Appendix B, presents the results of a simulation study of the performance of three contending networks, including the multi-channel ring network.

3. The Physical Topology of Survivable Wide Area Networks

Many military network applications will require a high degree of survivability and well as high network capacity. As an example of such an application, consider the interconnection of the sixteen Air Force bases shown in Figure 3.1. The interconnection of these bases must occur in such a way so that the network provides high capacity while maintaining a high

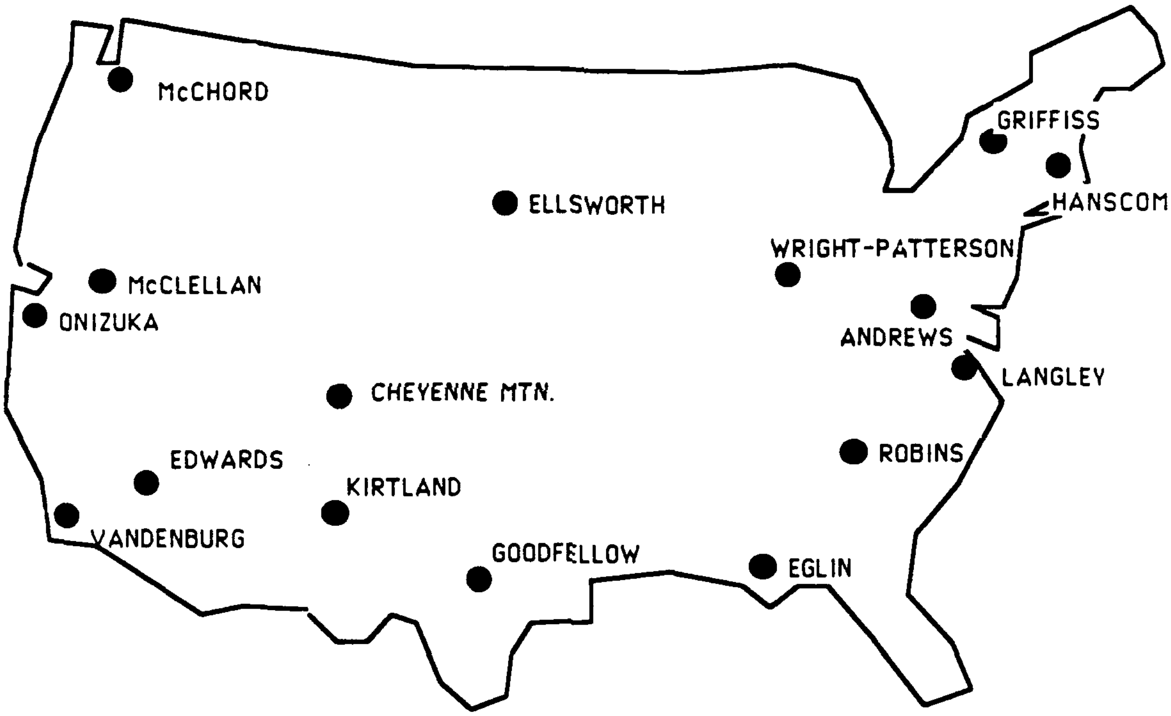

Figure 3.1: Selected U.S. Air Force Bases

degree of survivability and reliability, without being grossly inefficient or costly. For private wide area networks where transmission facilities, such as T1 links, SONET OC-3 channels, or whole fibers, are leased from telecommunications companies, there is obviously a tradeoff between network survivability and cost. This is because the leased transmission facility is likely to be a significant cost in the network.

A fully-connected mesh network in which there is a (bidirectional) link between any two nodes, would obviously be the most survivable and reliable, but the minimum spanning tree network, such as the one obtained using Kruskal's algorithm [6] and shown in Figure 3.2, is likely to be the most cost effective, assuming links have adequate capacity and that

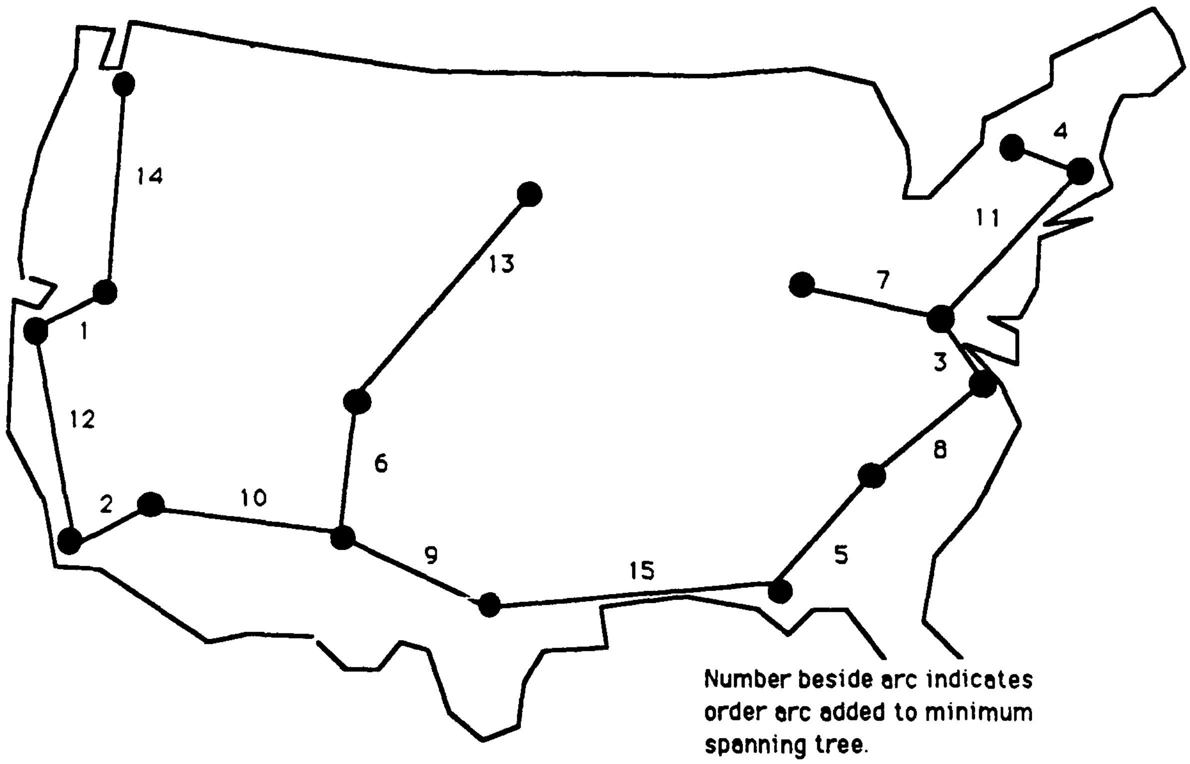

Figure 3.2: Minimum Spanning Tree Connection of Bases using Kruskal's Algorithm

link costs throughout the network are approximately uniformly proportional to their length. Somewhere between these two extremes lies the appropriate degree of connectivity for a given application. This decision can only be made after quantifying survivability requirements and transmission and switching costs, and using optimization techniques.

However, if we define the "acceptable level of survivability" to be that there are no single points of network failure or partition then the network must minimally be a two-connected network. In [7], Wu, Kolar, and Cardwell define network survivability as "the capability for the network to recover from a single failure." An additional stipulation of two-connectedness is that a group of one or more stations should never be removed from the rest of the network by any single failure in the network. Networks which are clearly not acceptable include the following:

- Those that consist of a single "master" node and several "slave" nodes, where the functions of the master node are not transferrable to any and all slave nodes. The loss of the master node would completely disable the network.

- Those for which timing or synchronization is derived from a single central clock source. The loss of the clock source would disable the network.

- Those physically configured in a star topology. Even if the star is passive and therefore does not pose a reliability problem, it could be the target of sabotage, which would completely disconnect the network.

A minimum cost two-connected network, where link cost is assumed to be roughly uniformly proportional to link length and links have adequate capacity, will be a ring topology, such as that shown in Figure 3.3. Some links must be added to the minimum spanning tree to form the ring; occasionally a few can be removed. For the example here, there is approximately a 33% increase in transmission media required to move from the minimum spanning tree to the ring topology. For even greater survivability a three-

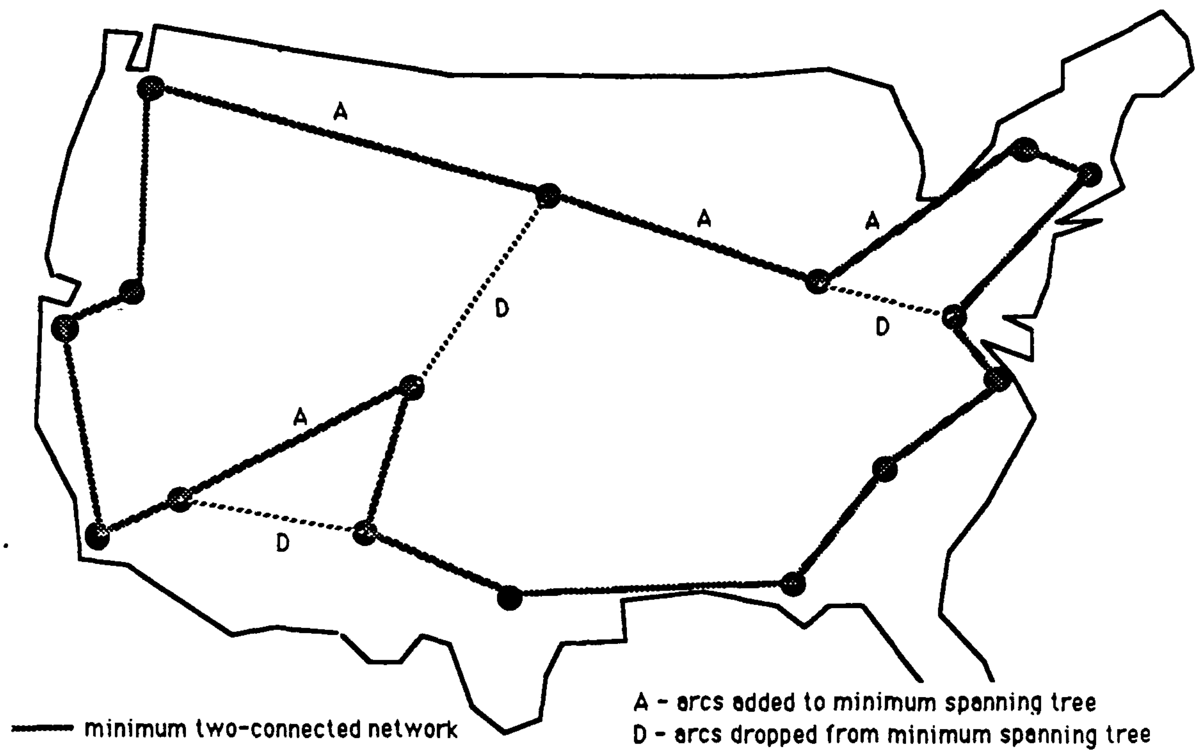

Figure 3.3: Minimum Two-Connected Network

connected network can be implemented, as shown in Figure 3.4. However, as can be seen, for this example the marginal increase in transmission media requirements is much greater when moving from two- to three-connectivity than when moving to the two-connected network from the minimum spanning tree. While this cannot be said to hold true in general, for situations where nodes are relatively uniformly distributed throughout a geographic region it is often the case.

To obtain two-connectedness, the links connecting nodes in the physical ring must be bidirectional, i.e., they must support traffic flow in each direction. These ring networks can be thought of as consisting of two rings with traffic flowing in opposite directions, and are therefore called "counter-rotating" rings. A counter-rotating ring is shown in Figure

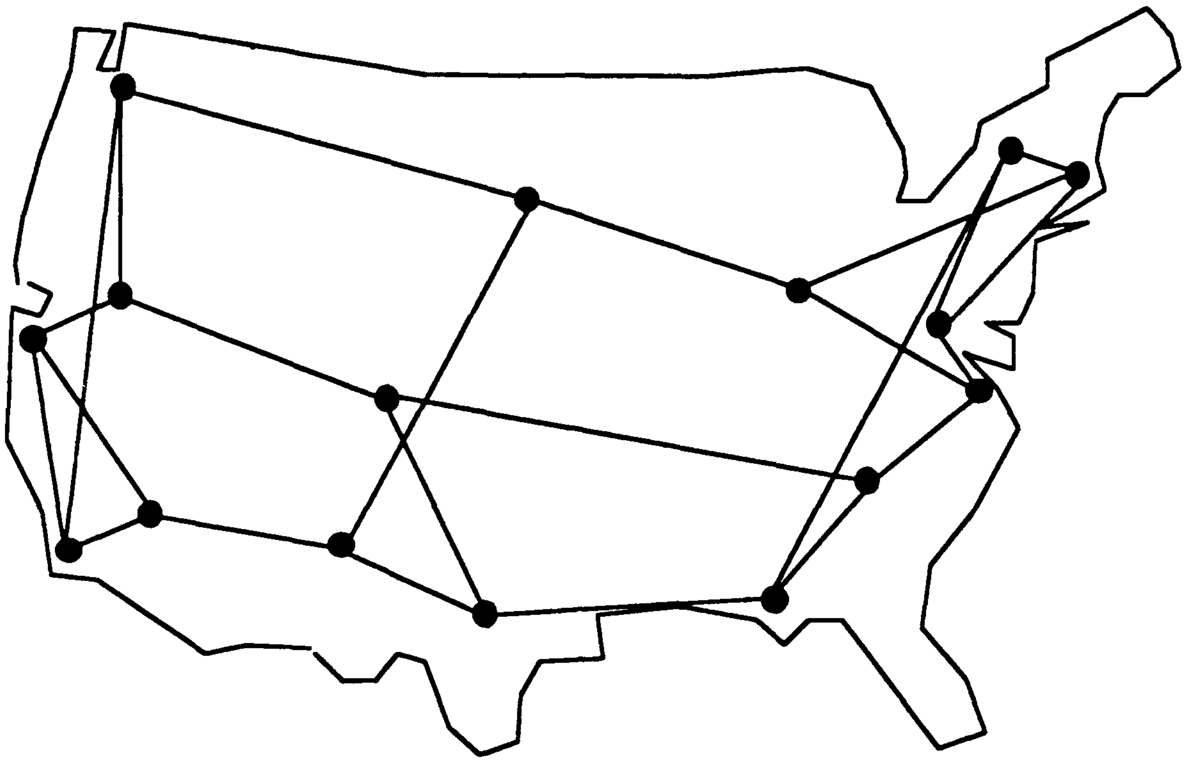

Figure 3.4: Three-Connected Network

3.5a. If a link failure occurs the ring takes on the configuration shown in Figure 3.5b. When equipped with automatic fault detection and isolation capabilities, the counter-rotating ring configuration is often called a self-healing ring.

The rings are often designated as primary ring and secondary ring, even in cases where they are peers. Use of the secondary ring under normal operation obviously increases the efficiency of the network, in that all available transmission resources can be utilized, and information can take the shorter route around the ring to the destination. However, complexity is increased due to the routing decisions that must be made at the node and the increase in management functions necessary to prevent and recover from failures. In many local networks where the cost of the unused media is not dominant, such as in the Fiber

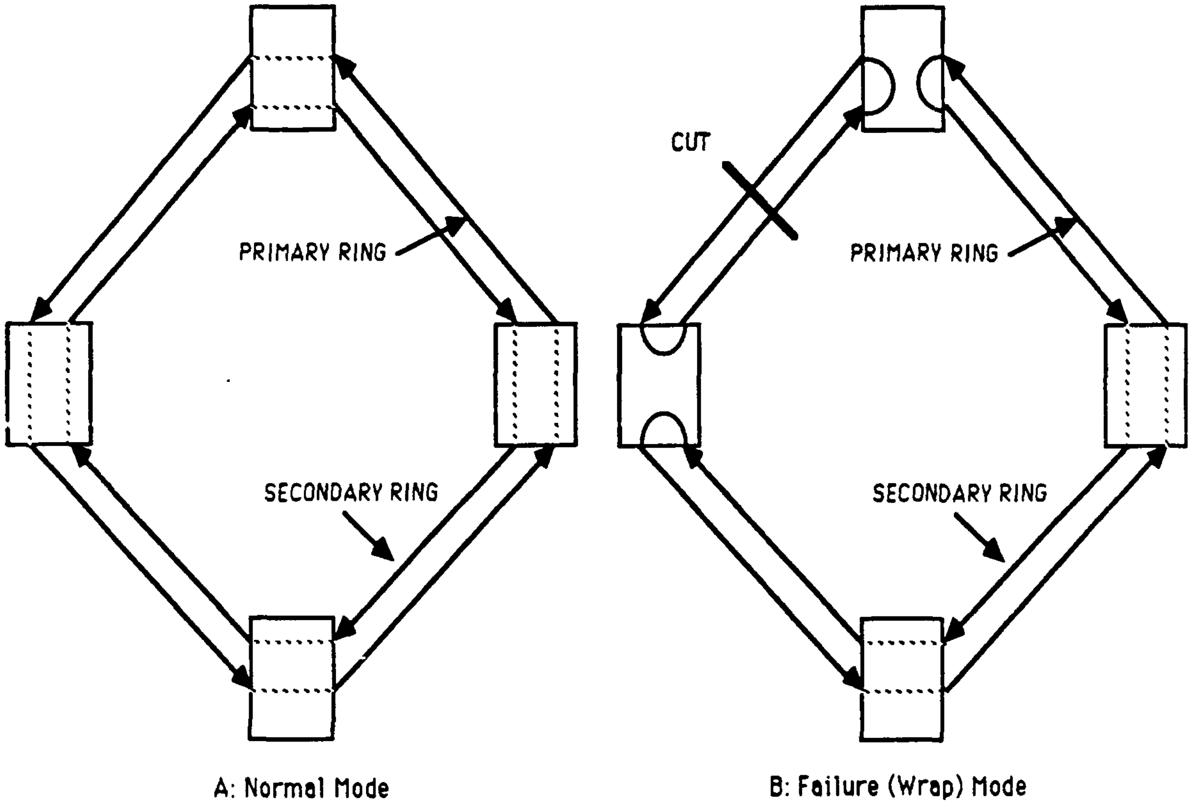

Figure 3.5: Counter-Rotating Ring in Normal and Failure (Wrap) Modes

Distributed Data Interface (FDDI) standard [8], the secondary ring is basically used only for ring initialization and in the event of a failure which requires reconfiguration. System delays are longer for this approach, but the network is easier to manage. In wide area networks where media costs are high and delays are nonnegligible, the secondary ring should be fully used.

In this section the rationale for use of the counter-rotating ring physical topology was presented. Now we continue toward the complete determination of the physical layer of a survivable wide area network architecture by looking at characteristics of long-haul fiber optic link technology.

4. Fiber Optics as a Transmission Media for Wide Area Networks

We assume that leased single-mode optical fiber will be the preferred medium for implementation of most of the private wide area networks, although microwave links will perhaps be used on particularly mountainous tracts. This is a realistic assumption, given the extent of the currently installed U.S. optical fiber base, and its continued growth. According to [9] AT&T planned to have a 25,000 mile fiber network installed by 1989. Fiber routes geographically cover most of the continental U.S. Other carriers, such as MCI and the seven company consortium known as the National Telecommunications Network, also have substantial national fiber coverage. Fiber optic cables generally contain a large number of fibers, often as many as 144.

Single-mode optical fiber is considered attractive for a number of reasons, including its extremely high potential bandwidth. The low attenuation region from 1500 nm to 1600 nm represents, when converted to the frequency domain, a 12.5 TeraHertz band (i.e. 1 nm = 125 GHz). This is enough bandwidth to carry the total voice traffic of the United States at any given time. A similar band exists at the low dispersion region around 1300 nm. Other characteristics of optical fiber that make it attractive are its immunity to electromagnetic interference and its resistance to tapping.

Organizations, like the Air Force, wishing to create wide area networks can lease part or all, if necessary, of the optical fiber capacity. They can also provide their own switching systems, or rely on the carrier for switching as well. If switching is done privately, and an administrator has access to the whole single-mode fiber, then an understanding of the characteristics of a long haul fiber optics transmission system is required. This section provides an overview of issues especially important to long-haul fiber systems. It is assumed that the reader has a basic knowledge of fiber optics.

4.1 The Single-Mode Fiber.

An optical fiber is a dielectric waveguide, which, according to Maxwell's equations, supports the transmission of light energy at certain discrete angles, or modes. These angles of propagation are separated by λ/D radians, or $(180\ \lambda)/(\pi D)$ degrees, where D is the core diameter and λ is the wavelength of the light source in the optical fiber, equal to the wavelength of the light source in air divided by the index of refraction of the fiber core, typically about 1.5. If the core diameter of an optical fiber is reduced to approximately an order of magnitude greater than the wavelength of light in the medium, *then only one mode*, the axial mode, can exist in the fiber. Total internal reflection occurs when the propagation angle relative to the fiber axis is less than $\cos^{-1}(n_{cladding}/n_{core})$, where $n_{cladding}$ is the index of refraction of the cladding and n_{core} is the index of refraction of the core. For single-mode fibers $n_{cladding}/n_{core}$ is typically 0.995, yielding a critical angle of 5.7° [10]. For 1300 nm operation, the first angle of propagation is greater than the critical angle of 5.7° if the core diameter is less than about 8.7 µm. Therefore, no non-axial modes are propagated, and single-mode operation exists, if the core diameter is less than 8.7 µm. For 850 nm operation the core must be less than about 5.7 µm, and for 1550 nm operation it must be less than about 10 µm.

There are actually two axial modes propagating in "standard" single-mode fiber, having orthogonal polarization. If the fiber were perfectly cylindrical and isotropic, then power coupled into one of the two modes would not couple into the other, orthogonal mode. However, since the fiber has small variations in its density and geometry, there is a random mixing of the power in the two orthogonal modes and light at the output of the fiber is in an arbitrarily polarized state. This is acceptable for most direct detection applications but is generally unacceptable for coherent detection applications, discussed later.

Eliminating all non-axial modes is advantageous because it eliminates one of the sources of dispersion in multimode optical fibers, modal dispersion. Modal dispersion occurs due to differences in propagation delays between modes, and obviously if only one mode is permitted to propagate, no modal dispersion can occur. In multimode fiber, a parabolic index profile across the core is often used so that mode travel times are roughly, although never completely, equalized. For single-mode fiber, a step index profile, where the index of refraction is constant across the core, can be used. Figure 4.1 shows the index of refraction profiles of conventional graded and step index optical fiber. The other type of dispersion, known as chromatic dispersion, is still present in single-mode systems. Chromatic dispersion is caused by differences in the index of refraction, and thus the speed of propagation, of the fiber as a function of wavelength. We will discuss chromatic dispersion in greater detail in a later subsection.

It is difficult to couple power into a single-mode fiber because of its small core diameter. Laser sources, with their highly directed emission distribution, relative to light emitting diodes, are required for long distance operation, and, even then, microlenses are needed to

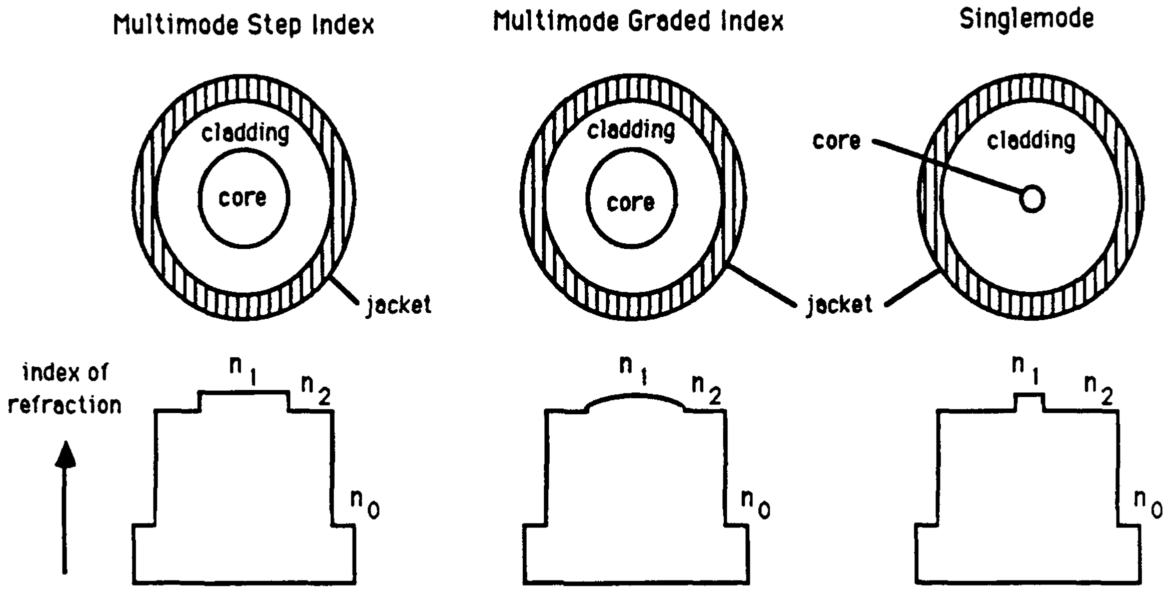

Figure 4.1: Index of Refraction Profiles for Conventional Fibers

further reduce the beam spread. A high level of mechanical precision must be maintained at the source-fiber junction and at any fiber-to-fiber junctions in the link. However, precision at the fiber-receiver junction can be relaxed somewhat due to the fact that the detector active area is generally large compared to the core diameter.

4.2 The Fiber Optic Transmission Link.

A fiber optic transmission link, shown in Figure 4.2, is here defined to stretch from the point where a coded electrical signal is converted to an optical signal, to the point where the optical signal is converted back to a digital signal and retimed. The link includes the transmitter, with the light source; the optical fiber, with any passive (or optically active) components such as optical amplifiers, switches, wavelength division multiplexers, and splitters (hereafter referred to as the fiber plant); and the receiver, which includes the detector, amplifiers, and retiming circuitry. Now, for the sake of clarity, the link will be considered unidirectional, and will consist of one serial bit stream from transmitter to receiver. As we shall see, several links can be operated over the same physical medium, using multiplexing, or several links can be logically grouped to form a parallel, or byte-wide, channel. A single fiber can also be used for bidirectional transmission.

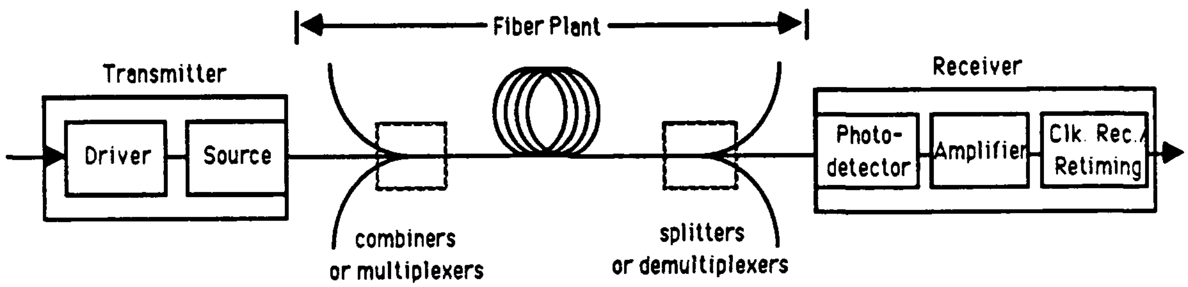

Figure 4.2: Fiber Optic Transmission Link

Various random and systematic noise sources, contributed from many different points in the link, affect the *bit error rate at the threshold point in the receiver* retiming and decision circuitry. The magnitude of the noise at the receiver can be roughly evaluated by looking at the "eye pattern" of the received signal prior to retiming, and measuring the degree of "eye opening," as shown in Figure 4.3a. The eye pattern is generated by superimposing in time the relative positions of all possible received bit patterns. If the eye pattern is closed predominately by encroachment at the top and bottom of the eye, as in Figure 4.3b, then the link is said to be power-limited, in that an increase in optical power at the receiver would dramatically improve the signal-to-noise ratio and therefore the bit error rate. If the eye pattern is closed by encroachment from the edges of adjacent symbols, as in Figure 4.3c, the link is said to be dispersion-limited, or bandwidth-limited. A reduction in the symbol rate would dramatically improve the bit error rate.

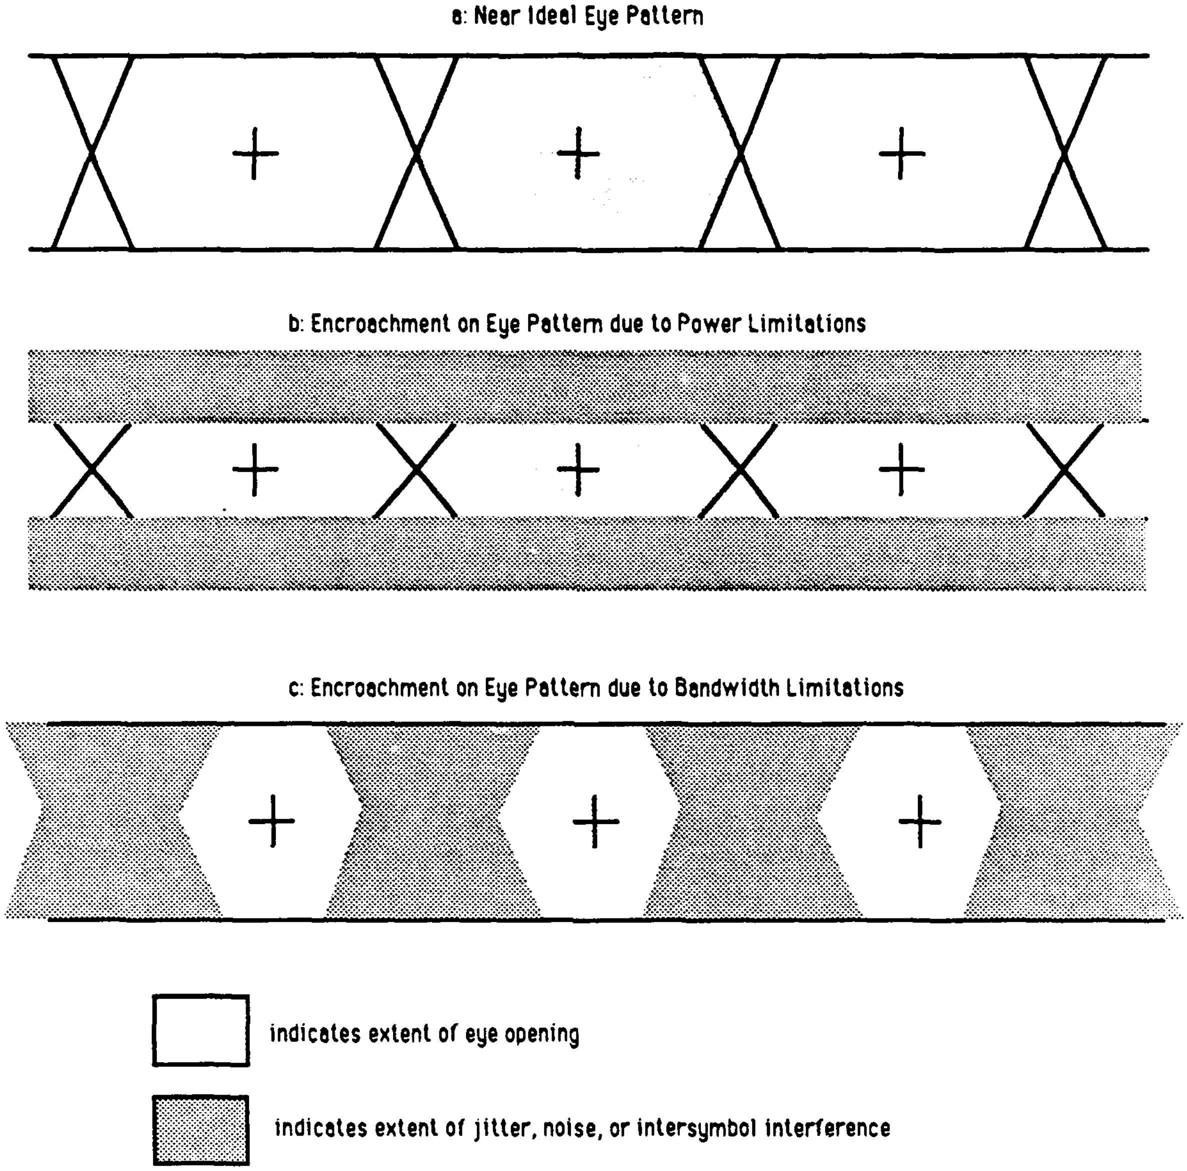

Figure 4.3: Eye Patterns showing Effects of Power and Bandwidth Limitations

There is a relationship between the power and dispersion limitations, in that a link which is dispersion-limited, or limited in symbol rate, can be somewhat improved through an increase in power; and the bit error rate in a power-limited link can be improved by a reduction in dispersion effects or symbol rate. This relationship is illustrated in Figure 4.4, which shows the worst case reduction in receiver sensitivity as a function of the ratio of the risetime of an incoming symbol to the symbol duration [11]. This relation is approximated by the following equation:

$$\Delta P_{in} \text{ (dB)} = 20\ ((0.4\ t_r / T)^2), \qquad \text{(Equation 4.1)}$$

where ΔP_{in} is the sensitivity degradation, T is the symbol duration, and t_r is the risetime (or falltime).

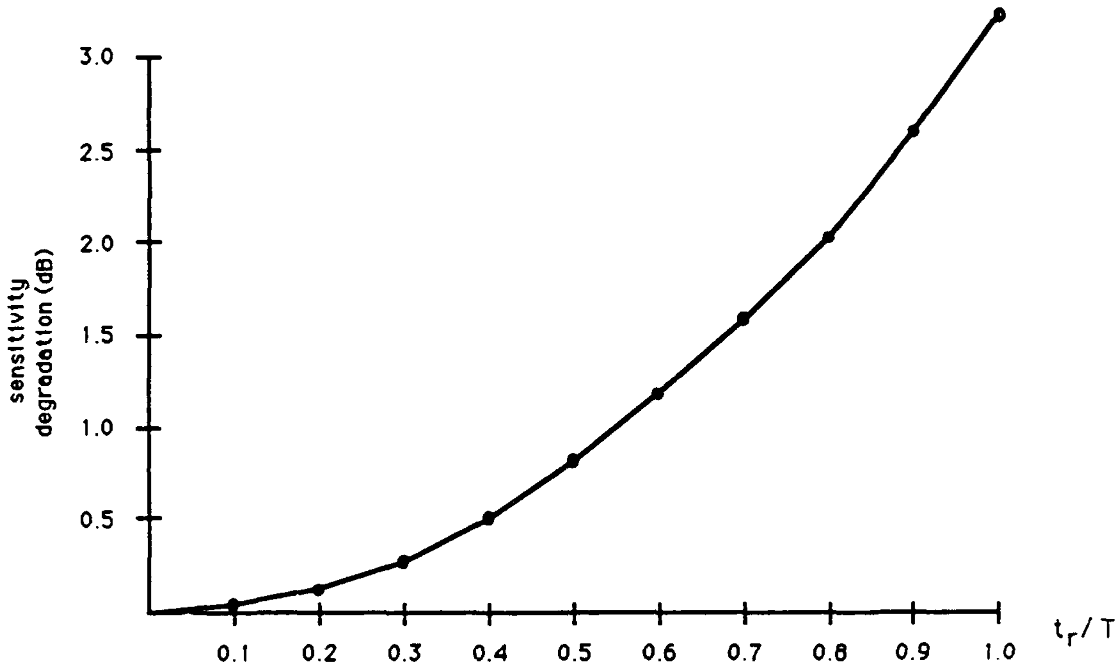

Figure 4.4: Sensitivity Degradation as a Function of Symbol Risetime

The slope of the curve becomes steep as the ratio nears one, with the results becoming increasingly unpredictable. Therefore, the maximum symbol rate which can be transmitted through a link is often said to be that rate which causes the receiver sensitivity to be reduced by one deciBel (dB). (In other words, one dB more power is necessary at the receiver to preserve the bit error rate attained at a low data rate). An intuitive explanation for the increasing steepness of the curves is that, after the rise- and falltime increases past a certain fraction of the symbol rate, the symbol amplitude is limited from attaining its full range. Palais states that maintaining a limit on the ratio of risetime (or falltime) to symbol duration of less than 0.7 will result in a sensitivity degradation of less than one dB [12].

We will now examine characteristics of the components of the fiber optic link that affect the "power budget" and the "dispersion budget" of the link.

4.3 The Optical Power Budget

In this section we will discuss characteristics of the source, the fiber plant, and the receiver that have impact on the top and bottom levels of the eye pattern at the decision threshold. A sufficient amount of signal power must be present at the receiver so that when this power is compared to the noise equivalent power (NEP) errors occur with low probability (on the order of 10^{-9} to 10^{-12}). Therefore, a sufficient amount of power must be introduced into the fiber by the source to insure the appropriate power level at the receiver. The allocation of optical power to link components, and for a safety margin, is known as the optical power budget.

4.3.1 Source Characteristics

The more optical power that can be coupled into the axial mode in the fiber core the more power will be available at the receiver, and - providing the power level does not overdrive (i.e. saturate) the receiver or cause nonlinear behavior in the fiber - the better the bit error rate. Also important is the extinction ratio of the optical source, measured as the ratio between optical high and optical low power levels. If the extinction ratio is small, resulting in the optical low power level being above the receiver noise floor, then the low level from a high powered source could be mistaken as the high level from a low power source. Furthermore, if the receiver is ac-coupled, then it makes decisions based on the difference between optical high and optical low levels, rather than on their absolute value. For applications where high sensitivity is required ac-coupled receivers are generally used because low frequency noise, such as photodetector "dark current," can be filtered out between amplifier stages in the receiver. A tradeoff therefore exists, because laser sources are generally operated with as small an extinction ratio as possible so that risetimes and falltimes are minimized, chirping is minimized, and the laser is operated in its linear region. Lasers are also generally very sensitive to changes in temperature, often caused by changes in the average current driving them, and either temperature regulation, perhaps using Peltier cooling, or current regulation via feedback is necessary to control the optical output power.

The shape of the optical pulse in the time domain can also be important to the maintenance of the eye pattern. Excessive overshoot or undershoot, or droop (see Figure 4.5) can reduce the effective eye pattern in ac-coupled receivers, especially if thresholds are set based on peak detection. The coding scheme used on the link is also important. Coding schemes where the average power level can deviate over time periods comparable to the time constant of the ac-coupling in the receiver can affect the threshold level. This effect is called "baseline wander."

There is an upper limit to the improvement of the power budget through increasing power

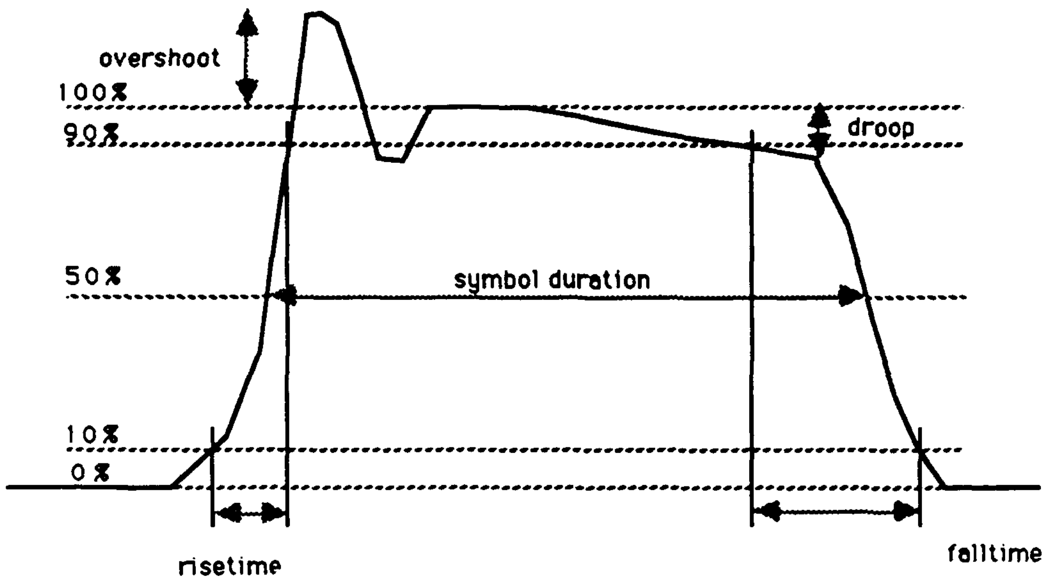

Figure 4.5: Waveform Parameters

coupled into the fiber, set not only by receiver saturation levels but by laser reliability and by the level where the optical fiber begins to operate in a nonlinear fashion. Stimulated Brillouin Scattering (SBS), whereby photons are scattered back toward the source at a wavelength slightly greater than the original wavelength, can be detrimental to the system by not only resulting in higher attenuation than expected, but by resulting in feedback to the laser source, causing instability, and crosstalk in optical channels counterpropagating on the same fiber. SBS can affect performance in long-haul single-mode systems at incident power levels beginning at around 10 dBm (10 mW) [13]. For multiwavelength systems, the optical power in each channel should be kept under 10 dBm to avoid SBS.

Other nonlinear phenomena, like Stimulated Raman Scattering (SRS) can also cause crosstalk in copropagating multiwavelength systems at higher power levels than those necessary for SBS. These nonlinear phenomena, and others, if used correctly, can enhance link performance as will be discussed later. This technology is still in the experimental stage, and could require dramatic changes in the installed fiber plant to be commercially useful.

4.3.2 Fiber Characteristics

Naturally, the loss in the optical cable, equal to the attenuation in dB/km multiplied by the distance in kilometers, is important. The solid line in Figure 4.6 shows the current practically achievable attenuation versus wavelength characteristic of common single-mode fiber. The dotted line represents the theoretically lowest achievable attenuation, set by Rayleigh scattering and absorption in silica glass devoid of impurities.

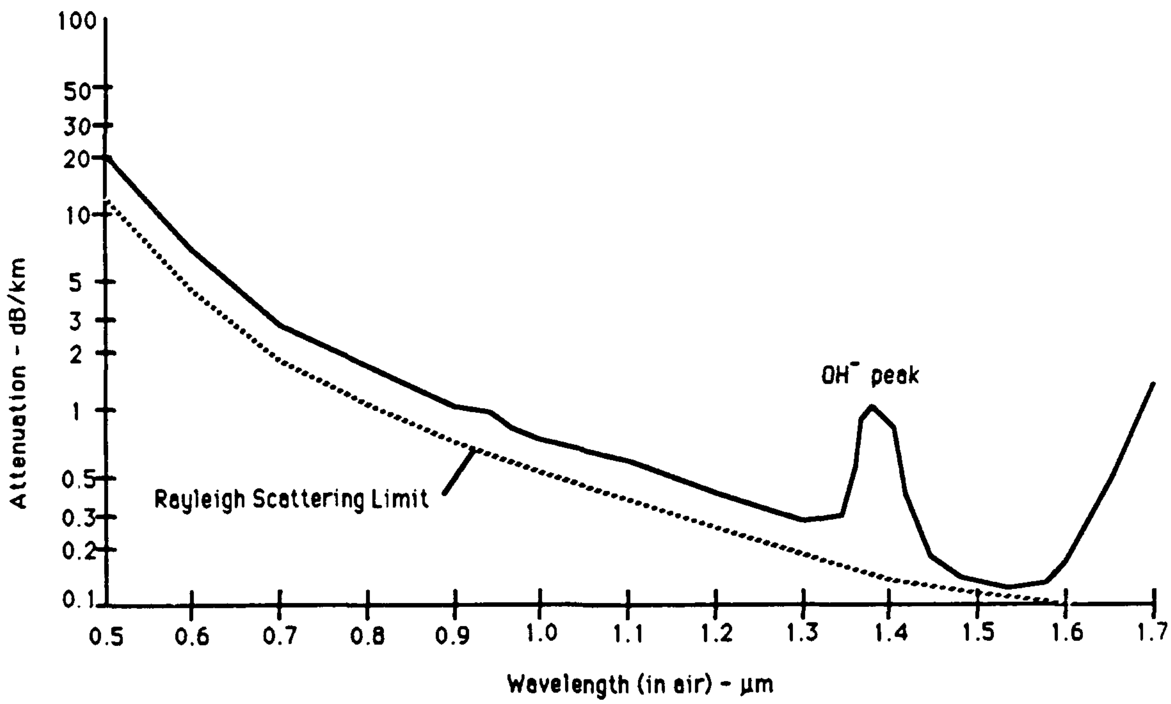

Figure 4.6: Fiber Attenuation as a function of Wavelength

The attenuation peak at about 1370 nm, due to the presence of OH^- ion impurities at the level of less than one part per million, is called the "water peak."

Three wavelength "windows," where attenuation is relatively low and where sources and detectors can be fabricated, are used for optical fiber communication. The "first" window, around 850 nm, is generally used for short distance applications, such as local area networks, due to the higher attenuation of this window (approximately 5-8 dB/km) and the lower cost of the optoelectronic technology. Designers have great experience with the Gallium Arsenide (GaAs), Silicon (Si), and Aluminum Gallium Arsenide (AlGaAs) materials used to construct sources and detectors at this wavelength, due to their previous use in visible optoelectronic applications.

The two "long wavelength" windows are at 1300 nm and 1550 nm. 1300 nm is used because it provides relatively low attenuation (0.8-1.5 dB/km) and is the region of zero chromatic dispersion in the fiber. 1550 nm is the region of lowest attenuation in the fiber, at around 0.2-0.5 dB. Techniques used to construct optoelectronic devices at these wavelengths are less mature, resulting in lower reliability and higher cost. Optoelectronic devices are fabricated from the quaternary compound Indium Gallium Arsenide Phosphorus (InGaAsP), where the compound fractions are altered to change the wavelength. Devices are also made from InGaAs or Germanium (Ge).

Fixed losses due to components in the fiber plant, such as connectors, splices couplers, splitters, switches, tunable filters, wavelength division multiplexers and demultiplexers, and couplers also affect the power budget. Reflections at interfaces can affect the stability of the laser source and must be prevented using anti-reflection coatings and isolators at the laser. Splice technology has become very mature, and splice losses of only 0.01 dB can be achieved. Connector losses are much higher, at around 0.5 dB, and are usually avoided in long-haul systems.

Some fiber plant components have a filtering effect on the signal with respect to wavelength, causing signal attenuation to change. Also, optical components can alter the mode distribution in the fiber core, such that some light is transferred out of the axial mode, or to cladding modes. Signal attenuation in the fiber immediately after the component will be higher than when measured under equilibrium conditions. Additionally, crosstalk from other wavelength multiplexed channels that may be present on the link can cause degradation in the signal-to-noise ratio at the receiver.

Recently, optical amplifiers have been developed that permit the linear amplification of an optical signal without the conversion to and from an electronic signal, stretching the capabilities of a power-limited system. Applications for optical amplifiers lie in three areas; *first*, as a post-amplifier directly after a laser (perhaps where power had to be sacrificed to achieve high modulation rates or narrow linewidths), second, as a mid-span repeater requiring less power and complexity than a conventional repeater requiring electro-optic conversion, and third, as a low-noise pre-amplifier to boost the signal prior to a photodetector.

These amplifiers can be classified into three categories, Fabry-Perot amplifiers, injection-locked amplifiers, and traveling wave amplifiers. The first two types are essentially conventional lasers that utilize externally introduced optical power to generate stimulated emission. Fabry-Perot amplifiers employ the conventional Fabry-Perot resonance between facets to amplify the incoming light. The amplifier is biased just below laser threshold. Because of the Fabry-Perot resonance this amplifier has an extremely narrow bandwidth (approx. 2-10 GHz) [14], and requires high wavelength and polarization stability in the incoming light. In traveling wave amplifiers, the Fabry-Perot resonance is suppressed due to the introduction of anti-reflection coating at the facets. These amplifiers are "single-pass" amplifiers, and can amplify light over a wide wavelength range of up to 70 nm [15]. Traveling wave amplifiers can be created out of semiconductor material, as for conventional lasers, or by using optical fiber doped with rare earth ions like erbium or neodymium [16]. When fiber amplifiers are pumped, a population inversion occurs resulting in optical gain at a wavelength associated with the dopant. Interfaces between the fiber amplifier and the fiber plant can be made through conventional splices, avoiding the problem of interfacing to and from the semiconductor material from the optical fiber. The cylindrical cross-section of the fiber amplifier also results in insensitivity to polarization, relative to the semiconductor laser amplifier. Fiber amplifiers, however, have the disadvantage of requiring an optical pump laser.

Currently, researchers can obtain a net gain of about 19 dB for semiconductor amplifiers, with saturation occurring at amplifier output powers of about 5 dBm. The application of optical amplifiers to a bidirectional, wavelength division multiplexed link will be discussed in greater detail in Section 7.

4.3.3 Receiver Characteristics

The receiver sensitivity, a measure of the lowest level of average (or peak) optical power that a signal can have and still be detected at a specified bit error rate, is the most important characteristic for power budget calculations at the receiver. Factors at the receiver that affect the sensitivity include the responsivity of the photodetector to the wavelength of the signal, the amount of dark current generated by the photodetector, the gain of the photodetector, the signal and noise amplification of the pre- and post-amplifiers, and the shot and thermal noise generated in the amplifiers.

There is a lower limit on the sensitivity of a receiver, called the "quantum limit." Even for an ideal receiver having no dark current, and a transmitter with infinite extinction ratio, when photons reach the photodetector they randomly, not deterministically, create hole-electron pairs according to a Poisson distribution given by:

$$p(n) = \Lambda^n e^{-\Lambda}/n!, \qquad \text{(Equation 4.2)}$$

where n is the number of hole-electron pairs generated and Λ is the average number of photons per received pulse. This randomness in the amount of hole-electron pairs created for a fixed number of incident photons is known as "quantum noise."

The threshold of the ideal receiver could be set so that if at least one hole-electron pair is detected then the receiver indicates a pulse was sent. Then, in order for the probability that no hole-electron pairs are generated when photons are sent (i.e. an error) to be less than 10^{-9}, the number of photons sent per pulse must be at least 21 [17]. At a certain pulse rate B, the average power required at an ideal receiver for a bit error rate of 10^{-9} is:

$$P_r = 1/2 \times 21 \times hf \times B, \qquad \text{(Equation 4.3)}$$

where h is Planck's constant (6.67×10^{-34} J/Hz), f is the frequency of the light, and the factor 1/2 is used for codes having a 50% duty factor (i.e. the average number of photons per bit, if there are an equal number of optical ones and zeroes sent, is 10.5). For example, an ideal system operating at 1550 nm using a balanced NRZ code at 2.49 Gbps/s would have a receiver sensitivity of about 3.4×10^{-9} W, or -54.7 dBm.

In practice, direct detection receivers (where a signal current is created based solely on the intensity of the incident light) come within 10 to 27 dB of the quantum limit. Two types of photodetectors are used. Avalanche photodiodes (APDs) achieve a varying amount of gain with a high reverse bias voltage of ten to one hundred volts. This gain is induced by the collision of holes or electrons with other molecules in the active region, creating more holes and electrons. Because the number of ions created by collisions is also a random variable, quantum noise is in effect also amplified in an APD. When the amplified random noise from the APD equals the noise of the following preamplifier circuit there can be no further increase in signal-to-noise ratio with increasing APD gain. Typical APDs require about 1000 photons per bit for 10^{-9} operation. APD noise starts to increase faster than signal gain at bandwidths in excess of two to five Gbps, representing the upper limit on practical APD use. PIN (Positive-Intrinsic-Negative) photodiodes do not have gain but require a lower reverse bias voltage, typically around five volts, and a simpler structure, lowering their cost and improving their reliability. PIN photodiodes are often combined with field-effect transistor (FET) amplifiers because of their low noise. Detectors at long wavelengths are fabricated with InGaAs or Ge, both of which have good responsivity over wavelengths ranging from 1000 to 1600 nm. Figure 4.7 shows the average sensitivity

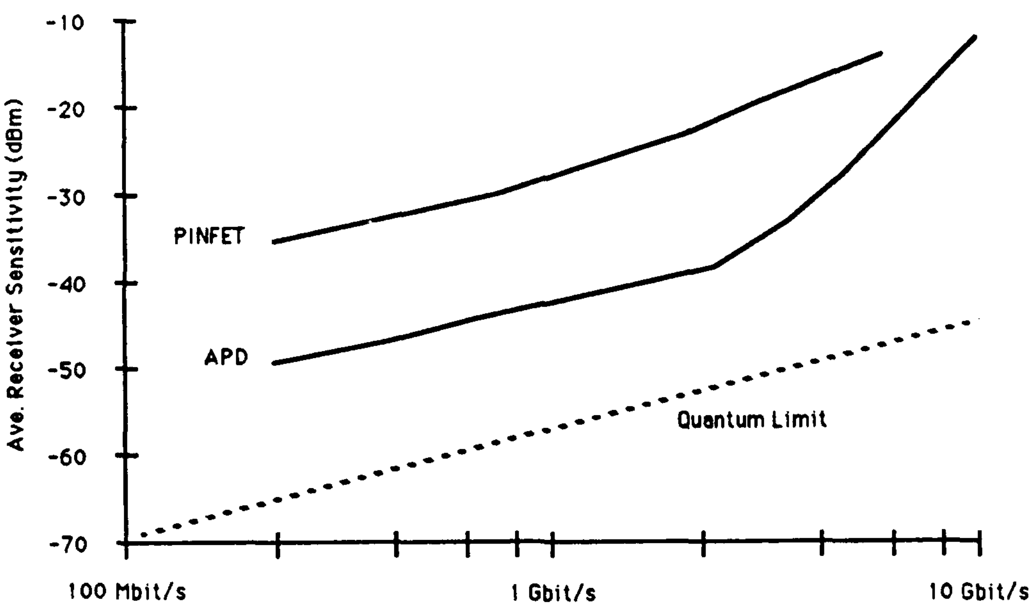

Figure 4.7: Average Receiver Sensitivity vs. Bit Rate for APD and PINFET Receivers [18]

(for 10^{-9} bit error rate) for a typical PIN-FET combination and a typical APD, relative to bit rate [18]. Noise sources in the receiver include the dark current, i.e. the residual current flowing in a photodetector even when there is no incident light, as well as shot noise and thermal noise. Receivers are generally designed to be limited by shot noise rather than thermal noise, improving their dynamic range.

Integration of the photodetector and amplifier circuitry onto the same integrated circuit can improve the sensitivity by reducing the inductance and stray capacitance of the interconnection between the two components. In long wavelength applications requiring the use of Indium Phosphate (InP) materials, there is difficulty in manufacturing semiconductors because of the fragility of the necessary InP oxide insulating layers. If manufacturing problems can be overcome, optoelectronic integrated circuits based on InP offer the potential for a factor of two improvement in speed over GaAs circuits, which are themselves five times faster than Silicon circuits.

The dynamic range of the receiver can also be important if the receiver has the potential to be used in both short distance, low attenuation applications and long distance, high attenuation applications. Such a situation could occur in a network where nodes which are inactive or failed can be optically bypassed. If dynamic range is poor then a receiver with good sensitivity may not perform well when subjected to a high power optical signal. Bit error rates can increase due to poor response times caused by device saturation.

If coherent detection is used, receiver sensitivities can be dramatically improved over those attained in practical direct detection systems. In contrast to direct detection techniques, where photons are directly converted into hole-electron pairs, coherent detection receivers first optically add a "local oscillator" signal, generated by a resident laser, and then perform demodulation or processing on the sum. In homodyne detection, a form of coherent detection, the optical frequency and phase of the local oscillator are equal to that of the incoming signal. The sum of the two signals is then a demodulated baseband signal. In heterodyne detection, analogous to the techniques used in radio communications, the local oscillator has an optical frequency slightly different from that of the incoming signal, resulting in the sum having a modulation frequency in the radio frequency region.

Theoretically, the sensitivity achievable with homodyne detection, assuming the amplitude of the local oscillator is much larger than that of the incoming signal, is equal to the quantum limit discussed above, even in the presence of noise in the amplifier. If it were possible to match the amplitude of the local oscillator with that of the incoming signal then the theoretical sensitivity would be 3 dB better than the quantum limit. With heterodyne techniques the theoretically achievable sensitivity is 3 dB worse than that for homodyne detection. [19]

Laser phase noise, caused by to the unavoidable spontaneous emission of photons which accompanies the stimulated emission, is a major cause for degradation in coherent systems. In the wavelength (frequency) domain this noise shows up in the nonzero spectral linewidth of the laser. With lasers having long external cavities and high output powers the linewidth can be reduced to tens of kilohertz, but with practical semiconductor lasers this value is typically around 10 to 100 MHz. The extent of the degradation in coherent systems is inversely proportional to the ratio of signal modulation rate to laser linewidth. Of the types of modulation used, including phase shift keying (PSK), differential phase shift keying (DPSK), and frequency shift keying (FSK), DPSK can operate with a modulation rate to linewidth ratio as low as 300, giving this form of coherent detection great promise for practical systems [19].

Another obstacle to the implementation of coherent detection systems is the requirement that the incoming signal and the local oscillator be equivalently polarized so that their electromagnetic fields will add correctly. Since with conventional fiber the optical signal will exit the fiber with random polarization, some technique must be used to overcome polarization differences. Possible techniques include the use of either polarization-maintaining fiber, an active optical device to alter the polarization of the incoming signal or the local oscillator, or a polarization diversity receiver which demodulates signals at orthogonal polarizations.

With the quantum limit on receiver sensitivity at around -55 dBm for 2.49 Gbps operation, even for coherent detection techniques, and the maximum optical power that can be launched into the fiber while avoiding nonlinear effects at around 10 dBm, we see that there is a maximum power budget of around 65 dB. This is less than the power budgets currently achievable with microwave and coax systems [20]. However, given the low attenuation of the fiber, particularly at 1550 nm, very long links can still be supported. If the attenuation is 0.2 dB/km, a power-limited repeaterless link length of 325 km can be achieved. However, these systems can be dispersion-limited much earlier than they are power-limited at the 2.49 Gbps rate, as we shall discuss in the next section.

4.4 The Bandwidth or Dispersion Budget

We will now discuss characteristics of the components of an optical system which can contribute to intersymbol interference or edge jitter, resulting in encroachment of the bit edges into the eye pattern and preventing the link from attaining the performance promised from power budget calculations.

4.4.1 Source Characteristics

Characteristics in both the time domain and the wavelength (or frequency) domain of the transmitter and optical source can contribute to intersymbol interference at the receiver. In the time domain, the rise- and falltimes (i.e. the bandwidth) of the driver-source combination could limit the link bandwidth. This is because the link bandwidth consists of the convolution of the individual bandwidths of the link components. If the propagation delay from optical low to optical high is not equal to the delay from optical high to optical low, then duty cycle distortion exists which will impact the eye pattern. Furthermore, if either of the propagation delays change depending on the bit pattern, then the resulting data dependent jitter will close the eye.

In the frequency domain, the interaction of the spectral characteristics of the source with those of the fiber will impact the temporal shape of pulses leaving the fiber. An optical source is usually defined spectrally using two parameters, the spectral linewidth, or wavelength spread, of the source, generally measured as a full-width, half-maximum value (see Figure 4.8), and the center wavelength, defined as the midpoint of the half-maximum points. Generally, the wider the spectral linewidth of the source, the greater will be the spreading of the optical pulse as it travels down the fiber, due to the fact that the index of refraction in the fiber is not constant relative to wavelength. Thus, different chromatic components of the pulse travel at different rates. This issue will be discussed in greater detail in the next section.

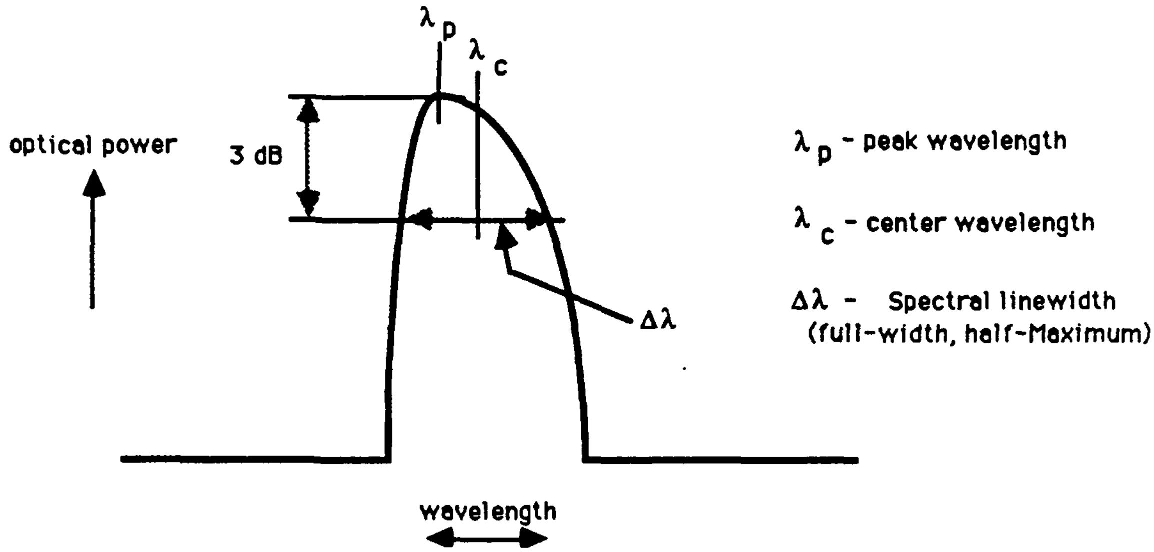

Figure 4.8: Optical Source Spectral Parameters

Stimulated emission in lasers is created by the reflection of spontaneously emitted photons back and forth within the laser material, imparting energy that causes the further emission of photons. Constructive and destructive interference occurs because of these reflections within the laser, resulting in higher coherence than for light emitting diodes. However, some spontaneously emitted radiation still escapes from the laser because reflection cannot be total. This spontaneously emitted radiation results in the nonideal (i.e. not infinitesimally narrow) spectral linewidth. In Fabry-Perot (F-P) lasers, the reflection, or feedback, of light energy occurs due to the presence of polished facets at each end of the lasing cavity. Constructive interference can occur at a number of slightly different wavelengths, where the resonance length is equal to an integer multiple of the wavelength; and the particular wavelength of operation can jump somewhat randomly among the wavelengths of highest gain. The spectral characteristics of the Fabry-Perot laser are shown in Figure 4.9. The spectral linewidth, $\Delta\lambda$, ranges from about 1-5 nm for these lasers. The spectral peaks are known as modes. Random jumping of the optical output

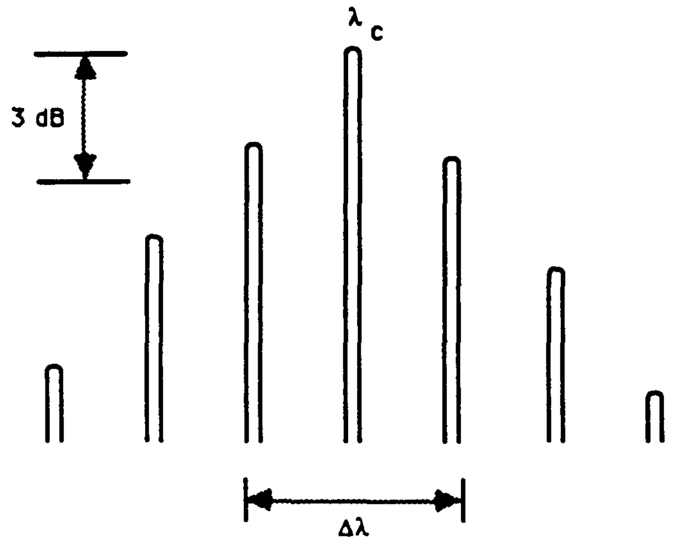

Figure 4.9: Fabry-Perot Laser Optical Spectrum Showing Multiple Modes

from mode to mode, a phenomena known as mode partition noise, can cause pulse spreading.

Distributed Feedback (DFB) lasers have been developed with the feedback mechanism incorporated throughout the extent of the laser cavity, using a corrugated structure positioned below the active region as shown in Figure 4.10. The period of the corrugated structure allows for a tighter control on wavelengths experiencing constructive interference, because of the shorter resonance lengths. These types of lasers are included in a family of lasers known as single mode lasers, because all modes except the main mode are effectively suppressed.

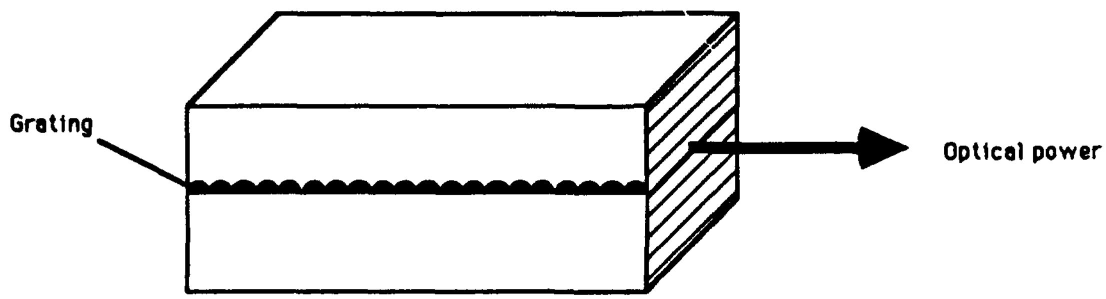

Figure 4.10: Distributed Feedback Laser

Another source of degradation is laser chirp. The spectral characteristics of the laser will generally vary linearly with the charge in the laser cavity, due to its effects on the index of refraction. The effect on a pulse is that the spectral components comprising the beginning of the pulse will differ from those at the end. In most 1550 nm lasers the wavelength increases (i.e. the chirp rate is negative) with pulse duration. This is detrimental since shorter wavelength components of the pulse travel faster than longer wavelength components in the 1550 nm region of the fiber, meaning that the pulse spreads as it travels.

External modulation of a laser is seen as one method to eliminate laser chirping. In this technique the laser is driven with a constant current. The continuous wave optical signal then passes through an electro-optic or acousto-optic modulator, where the transmittance of the modulator varies with the data signal to be sent. Lithium Niobate switches are often used as electro-optic modulators. They are capable of being modulated at rates over ten Gbps.

4.4.2 Fiber Characteristics

The silica glass material used to make optical fibers does not have a constant index of refraction over the range of wavelengths used in optical communications. Also, the propagation speed in a single-mode dielectric waveguide is affected by the wavelength of optical power. The geometry of the waveguide can be altered to change the propagation characteristics of the fiber. These two effects, known respectively as material and waveguide dispersion, combine to create the chromatic dispersion characteristic of an optical fiber. The relative delay between wavelength components in a single-mode fiber is shown in Figure 4.11. The slope of this curve is zero at around 1310 nm, meaning that wavelengths near 1310 nm will propagate at the same rate through the fiber. This wavelength is known as the zero dispersion wavelength. When the derivative of this curve is taken relative to wavelength it yields the curve for dispersion, in ps/nm-km, as a function of wavelength. This curve, shown in Figure 4.12, is approximated by the Sellmeier Equation as follows:

$$D = -S_0 \lambda [1-(\lambda_0/\lambda)^4]/4, \qquad \text{(Equation 4.4)}$$

where S_0 is the slope of the dispersion curve at the zero dispersion wavelength, λ_0 is the zero dispersion wavelength, and λ is the wavelength of operation [21].

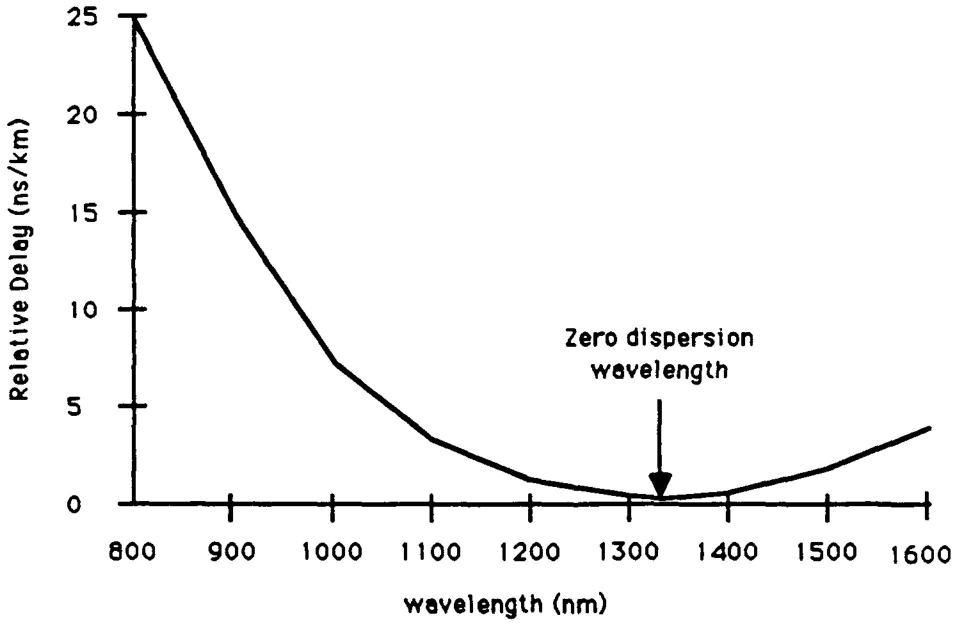

Figure 4.11: Relative Fiber Delay Characteristics as a Function of Wavelength [21]

As an example of the effect of chromatic dispersion, consider a 1550 nm laser source having an effective spectral linewidth of 1 nm, transmitting over a distance of 50 km at one Gbps. The chromatic dispersion at 1550 nm is about 20 ps/nm-km. Due to the laser spectral linewidth this yields a dispersion of 20 ps/km, and over the 50 km distance the total pulse spread is 1000 ps, or one ns. This pulse spread is equal to the total bit duration at one Gbps, meaning that the signal would be completely distorted at the receiver.

Near the zero dispersion wavelength of 1310 nm the Sellmeier equation is not completely accurate in predicting performance because of the possibility that the source spectrum falls into both the positive and negative regions of the dispersion curve, perhaps leading to some pulse compression. Computer models are often used to predict exact performance in this region, if required.

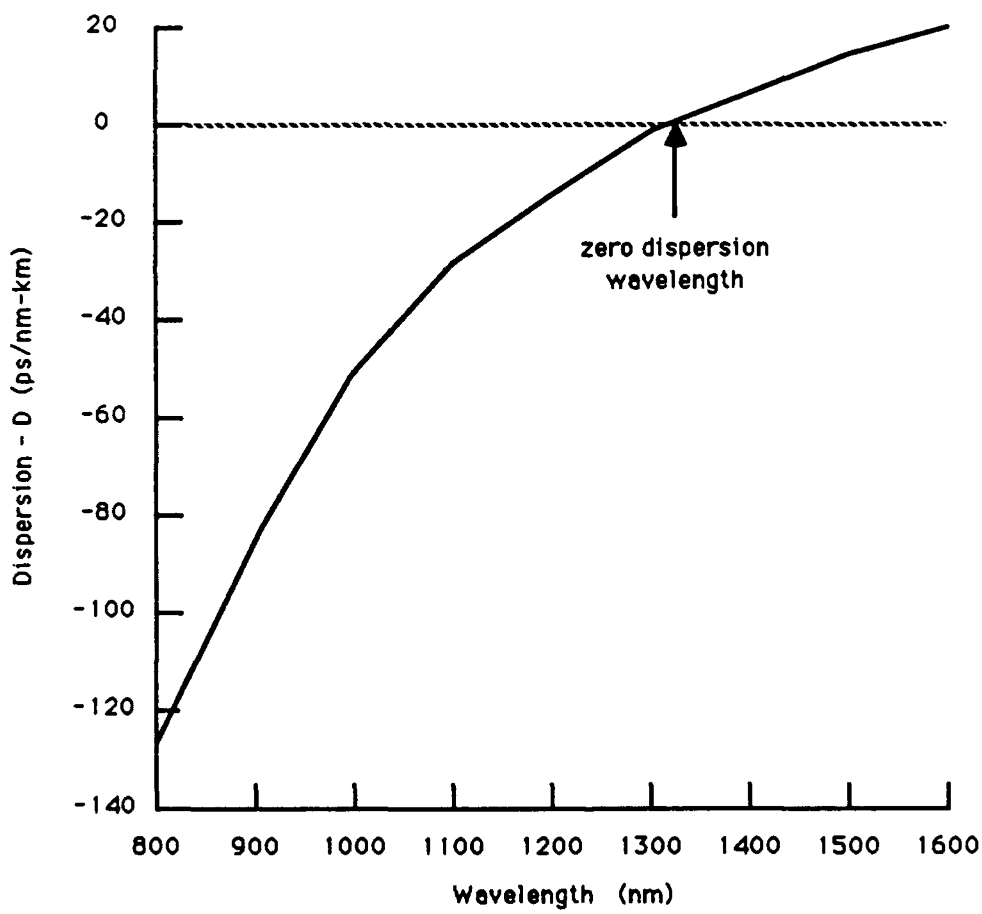

Figure 4.12: Fiber Dispersion as a Function of Wavelength [21]

The geometry of the fiber can be altered to flatten the dispersion characteristics of the optical fiber, so that dispersion is low over a wavelength range from 1300 to 1550 nm, or to shift the zero dispersion wavelength to 1550 nm, where fiber attenuation is lowest. Index of refraction profiles for fiber designs with flattened dispersion characteristic are shown in Figure 4.13a [22]. Profiles for designs which shift the dispersion characteristics to the higher wavelength are shown in Figure 4.13b. These profiles can be contrasted with those for conventional fibers shown in Figure 4.1. The dispersion curves for the new fiber designs are shown in Figure 4.14.

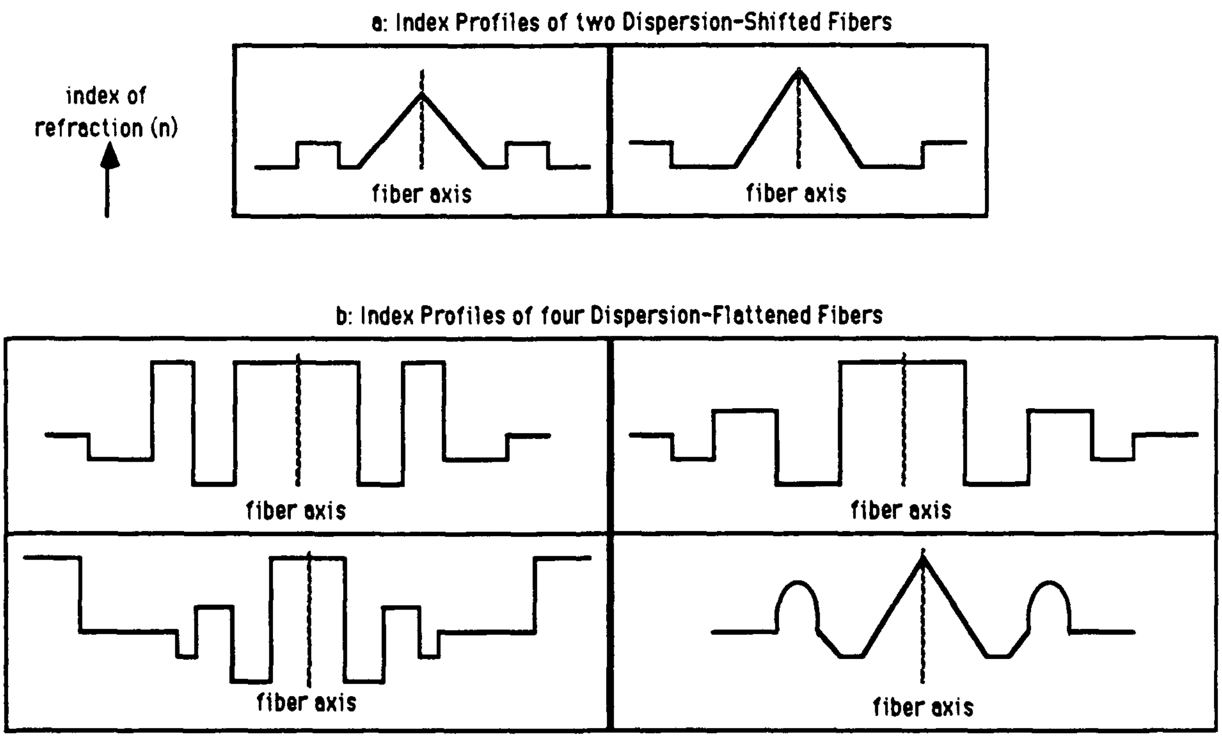

Figure 4.13: Index of Refraction Profiles for Dispersion Shifted and Flattened Fiber [22]

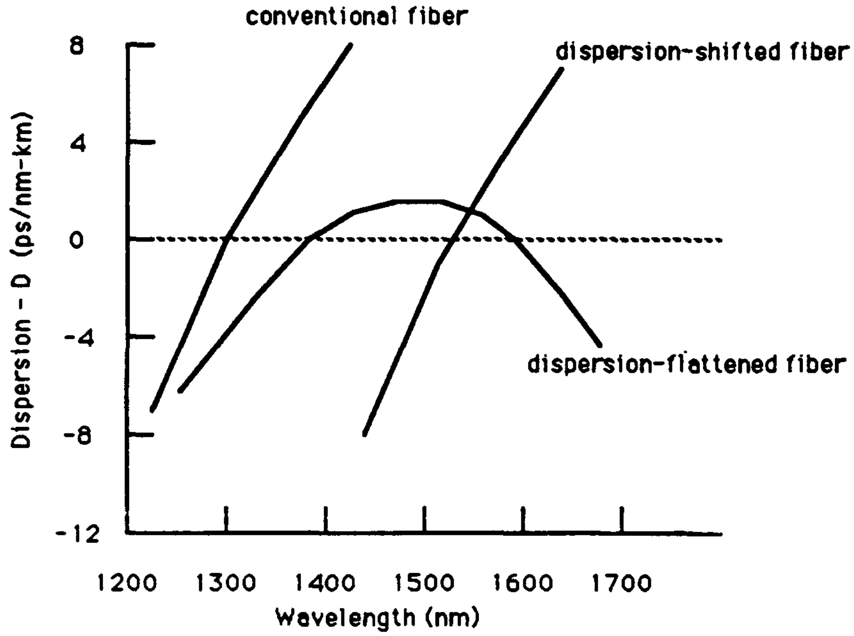

Figure 4.14: Dispersion in Dispersion-Shifted and Dispersion-Flattened Fibers

4.4.3 Receiver Characteristics

The bandwidth of the receiver, like the transmitter, affects the overall bandwidth of the link. Just as the random jitter associated with the shot and thermal noise closes the height of the eye pattern, it can close the width. The systematic jitter components of duty cycle distortion and data dependent jitter in the receiver can also lead to intersymbol interference.

Assuming the data link is digital, the information must be retimed and converted to digital logic levels at the receiver. The position of the clock pulse relative to the center of the information bit (i.e. the position where the eye opening is the widest) will, along with the other factors discussed in this section, affect the bit error rate of this conversion. If the clock position deviates from the bit center, the bit error rate will increase. Many clock regeneration mechanisms are based on deriving energy from edges in the data stream. If these edges are infrequent, the clock frequency may drift away from the intended frequency to the "natural" frequency of the retiming circuit, often based on a phase locked loop or a surface acoustic wave filter. If data edges are positioned incorrectly due to jitter, the clock position can also drift from center. Other sources of drift of the clock edge from bit center include temperature and ageing effects.

4.5 Nonlinear Fiber Optics

Because optical fibers closely confine optical power and exhibit low loss, allowing for long distances over which interactions can occur, nonlinear effects are evident in optical fibers at lower power levels than for most other media. On a quantum mechanics level, nonlinear performance occurs when photons are modulated by vibrating molecules in the fiber, or when the characteristics of the material are altered causing a change in the index of refraction or polarization. Nonlinear effects may detrimentally affect performance by causing unexpectedly higher attenuation, crosstalk, or even permanent alteration of the fiber, and therefore should be accounted for in long-haul fiber systems. Nonlinear effects can, however, be used to enhance the performance of the system. In this section we will discuss the important nonlinear effects of Stimulated Raman Scattering (SRS), Stimulated Brillouin Scattering (SBS), Self-Phase Modulation (SPM), and the optical soliton.

4.5.1 Stimulated Raman Scattering (SRS)

SRS occurs when photons traveling at a certain frequency are modulated at the frequency of the vibrating molecules in the fiber, causing sidebands to occur in the spectrum at frequencies separated from the initial frequency by the molecular vibration frequency. The molecule then makes a transition to a lower vibrational state. If the intensity of the original wave (called the pump) is high enough most of the energy will be

transferred to waves of lower frequency, called Stokes waves. There is a relatively wide band of frequencies at which these Stokes waves can occur, due to the non-crystalline nature of the silica fiber, although for silica there is a maximum transfer to frequencies shifted 13.2 THz down from the pump frequency. Obviously Raman scattering is detrimental in multi-channel systems, but can have important applications in optical broadband amplifiers. Raman scattering begins to build up exponentially in single-mode optical fibers at incident power levels starting at about 3.3 W [23].
SRS can be used to amplify a signal if the pump and the signal (referred to as the probe) copropagate in the fiber, with the probe having a frequency differing from that of the pump by about 13.2 THz. Over the length of the fiber, energy will be transferred to the probe from the pump, provided the pump has sufficient initial power.

4.5.2 Stimulated Brillouin Scattering (SBS)

SBS is similar to Raman scattering in that the frequency of photons in a wave acting as a pump is downshifted to a new frequency, and a Stokes wave is created. The pump wave modulates the index of refraction of the material, creating, in effect, a diffraction grating moving in the material at an acoustic velocity. Diffraction is maximum in the reverse direction from the direction of propagation of the pump, different from the Raman case. As previously mentioned, SBS effects also become evident at a lower power level than SRS, typically around 10 mW in conventional single-mode fiber. Also,the SBS frequency shift is in the acoustic band at around 10 GHz, whereby the frequency shift for Raman scattering is in the optical band. Due to the small frequency shift the forward diffracted light usually falls within the spectral linewidth of the source and is dominated by it. Therefore, in the forward direction of propagation, SRS is the dominant nonlinear process.

Brillouin scattering could be detrimental to full duplex fiber systems, in that energy from a wave could be transferred to a counterpropagating wave of slightly lower frequency, causing crosstalk. SBS can also be beneficially used to amplify a signal from the opposite end of the fiber from the signal input. The problem arises as to how the pump can be synchronized with the signal so that backscattering does not occur in the absence of signal.

4.5.3 Self-Phase Modulation

Self-Phase Modulation occurs when the index of refraction is dependent on the intensity of an optical pulse traveling through it. Although the index change is very small, on the order of 10^{-13}, the long interaction length and confinement of the fiber permits noticeable effects at input optical power levels exceeding about 200 mW in single-mode fiber [24]. The index of refraction becomes greater with intensity, causing the phase of the pulse to vary directly with the intensity. This implies that the instantaneous frequency of the pulse varies in proportion to its slope in the following manner:

$$\delta\omega(T) = -(\partial\phi_{NL}/\partial T) = -(\partial|U(0,T)|^2/\partial T)(1-\exp(-\alpha z))\gamma P_0/\alpha, \quad \text{(Equation 4.5)}$$

where $\delta\omega(T)$ is the change in frequency with respect to time, ϕ_{NL} is the phase, $U(0,T)$ is the pulse shape, z is the distance the pulse has traveled in the fiber, α is the fiber attenuation, P_0 is the peak input optical power, and γ is a measure of the nonlinearity of the index of refraction relative to intensity [25]. Figure 4.15 shows the change in frequency for a Gaussian shaped pulse. Note that there is a portion of the pulse over which the frequency varies linearly. This effect therefore presents the opportunity to create pulses having a linear variation in frequency over time. A linearly dispersive delay line, often formed by a grating pair, can then be used to perform pulse compression, resulting in extremely fast, high energy pulses. Figure 4.16 shows a dual-stage pulse compressor twice employing this technique, resulting in a pulse having 90 fsec duration. The 5.9 psec pulse, having a peak power of 2.4 kW, enters the three meter optical fiber, where self-phase modulation spreads the pulse to 10 psec, with the wavelength varying linearly across the pulse. The first linearly dispersive delay line, consisting of the grating and the prism, compress the pulse to 200 fsec. The grating first separates the pulse into its

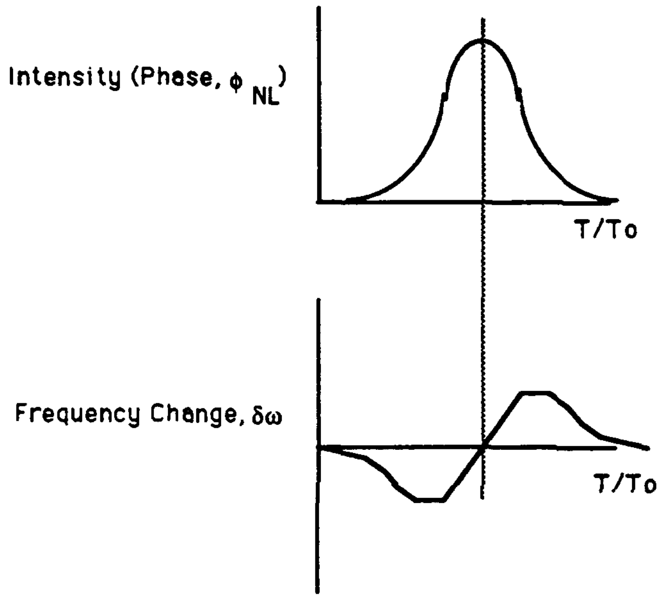

Figure 4.15: Self-Phase Modulation of a Gaussian Pulse

wavelength components. These components travel a linearly varying distance, dependent on wavelength, between the grating and the prism and back to the grating. The components are recombined again at the grating. Then, self-phase modulation again occurs in the 55 cm fiber and the pulse compression process is repeated with the second grating/prism pair. The end result is a pulse of 90 fsec duration and 20 kW peak power, which could be used in the code or optical time division multiplexing techniques discussed in Section 5.

Since self-phase modulation causes the creation of new frequencies as the pulse travels down the fiber it can have a detrimental effect on optical communication system performance, due to dispersion. As can be seen from Figure 4.15, the frequency is lower at the beginning of the pulse and higher at the end. Positive dispersion (occurring at wavelengths higher than the zero dispersion wavelength) will therefore initially compress the pulse before spreading occurs. This point is important for the understanding of soliton creation, discussed next. Negative dispersion spreads the pulse from the onset.

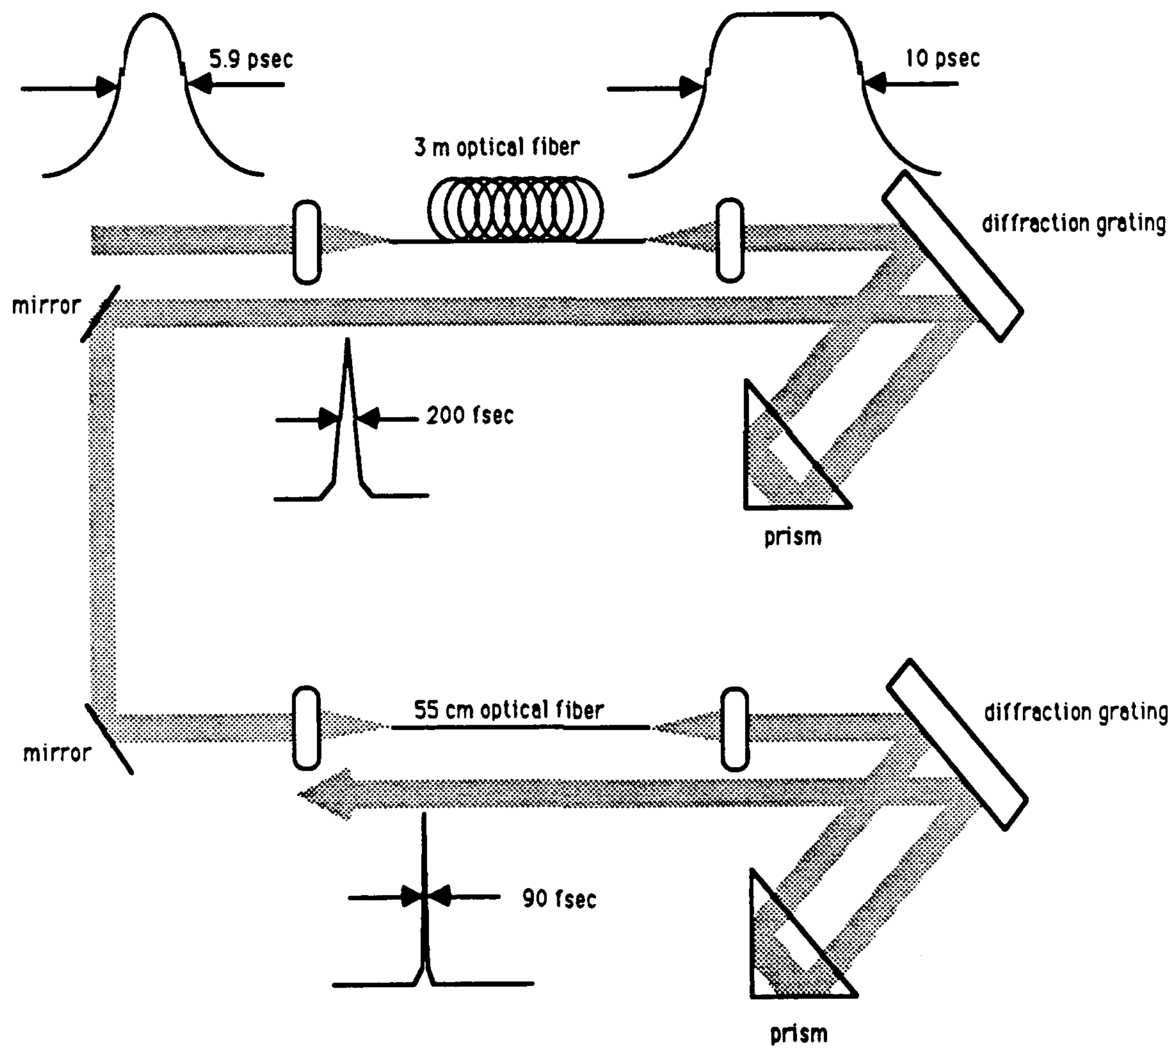

Figure 4.16: Dual Stage Optical Pulse Compression System [26]

4.5.4 Optical Solitons

It can be shown that, if we are operating in the region of positive dispersion, pulse shapes can be created having a form such that effects due to self-phase modulation and to chromatic dispersion exactly cancel each other out, allowing pulses to travel undistorted through the medium. If the amplitude of the pulses are periodically refreshed to compensate for fiber attenuation, possibly through the use of Raman amplifiers, then solitons could travel well over 1000 km at bit rates of 100 Gbps/s before breaking down due to spontaneous emission in the amplifiers [27]. Transmission of a 55 ps soliton over 4000 km has been experimentally demonstrated using dispersion-shifted fiber and Raman amplification every 42 km by Mollenauer and Smith [28].

If the laser used to generate the soliton has chirp, then the soliton degrades faster than for the unchirped case due to the fact that the interaction between the self-phase modulation phenomena and the fiber dispersion is disrupted. If the chirp is great the soliton may not form at all. Different solutions for the soliton pulse shape could theoretically be developed to account for chirping. Another design consideration is that the solitons propagating in the fiber should be separated by four times their pulse width, or interaction between adjacent solitons can occur, negatively affecting the performance. [29]

The pulse shapes that provide exact cancellation of dispersion and self-phase modulation in the absence of chirping are given by:

$$a = (A/A_0)(\mathrm{sech}(t/T)), \qquad \text{(Equation 4.6)}$$

where A/A_0 is an integer (N), and A_0 is the area of the pulse, dependent on system parameters. If A/A_0 is one, then the pulse will maintain its shape all throughout its travel. If A/A_0 is greater than one then the pulse will periodically attain a much more pronounced impulse shape as it travels down the fiber, as shown in Figure 4.17.

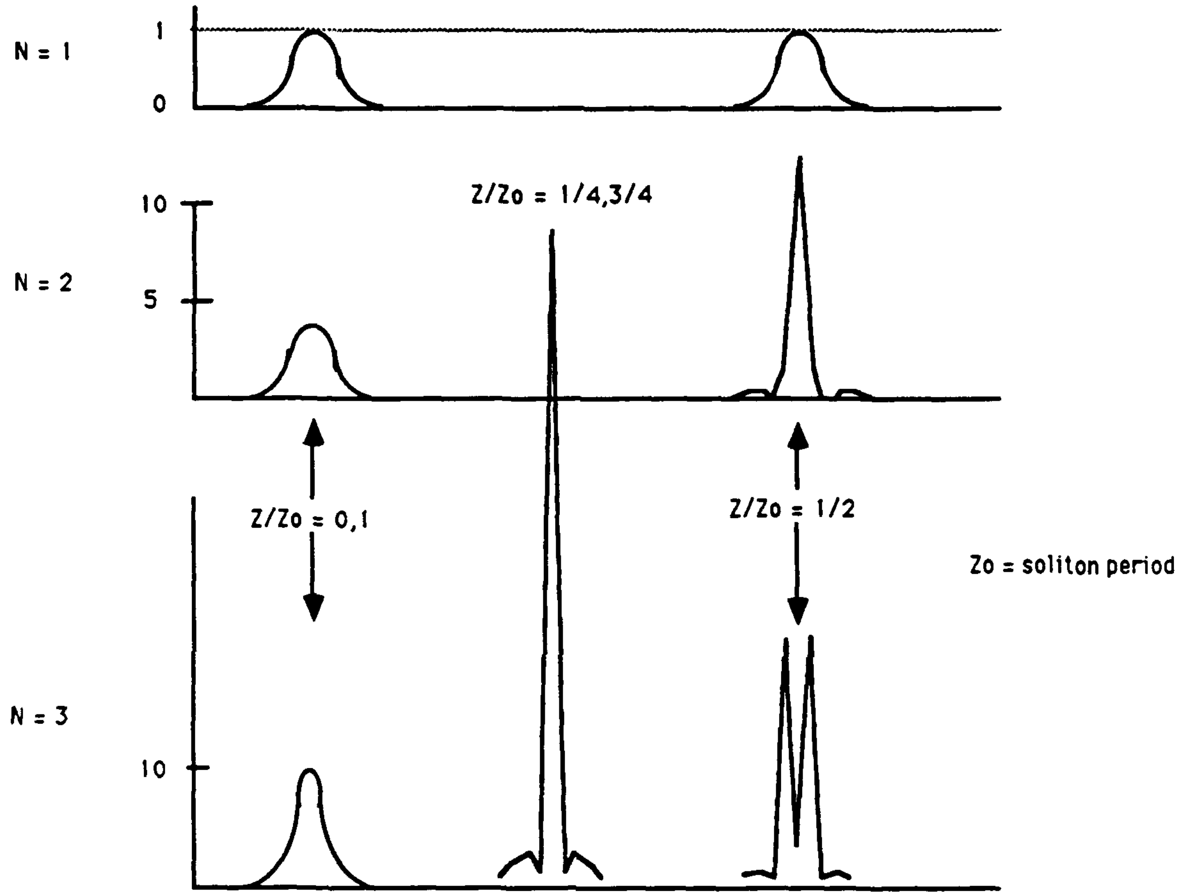

Figure 4.17: Fundamental, Secon and Third Order Solitons [30]

Now that an overview of fiber optic technology as it applies to long-haul transmission systems has been presented, we will turn to an investigation of how these techniques can be used in optical multiplexing techniques to increase the capacity of a fiber link.

5. Optical Multiplexing Techniques

The current state-of-the-art in high speed receiver pre-amplifier, post-amplifier, decision, and retiming circuit development using Gallium Arsenide (GaAs) is around 2.4 Gbps [31]. Furthermore, laser performance begins to degrade due to reflections and mode partition noise at bit rates of over one Gbps [32]. Although the top switching speeds for Silicon (Si) and Gallium Arsenide could possibly reach 20 Gbps within the foreseeable future [33], schemes other than strict electronic time domain multiplexing should be evaluated due to the fact that traffic requirements may soon even exceed this limit.

Using other multiplexing techniques, "multi-channel" networks can be developed, where the overall traffic rate is very high, but the rate on individual channels remains low enough to be manageable with conventional *electronics*. The multiple channels may be realized through one or a combination of different techniques on the fiber optic medium, which include wavelength, frequency, code, optical time, and space division multiplexing.

5.1 Space Division Multiplexing

Space division multiplexing simply involves the use of separate fibers for each channel in the physical link. As previously mentioned, optical cables can include a high number of optical fibers. In [34], Jacobson proposes a multi-channel token bus using space division multiplexing. For a high speed space division multiplexed network each channel could be made fully compatible with SONET and ATM requirements, whereby wavelength and code division multiplexing would not allow for compatibility with (current) standards. This allows for the potential use of commercial, off-the-shelf switching and transmission equipment. The chance of crosstalk between channels would be almost zero for space division multiplexing. The drawback to this approach is, of course, the requirement for several fibers for each physical link instead of one, and the inefficient use of their potential bandwidth. The cost of adding channels once existing capacity is exhausted can be extremely high. An illustration of space division multiplexing is shown in Figure 5.1.

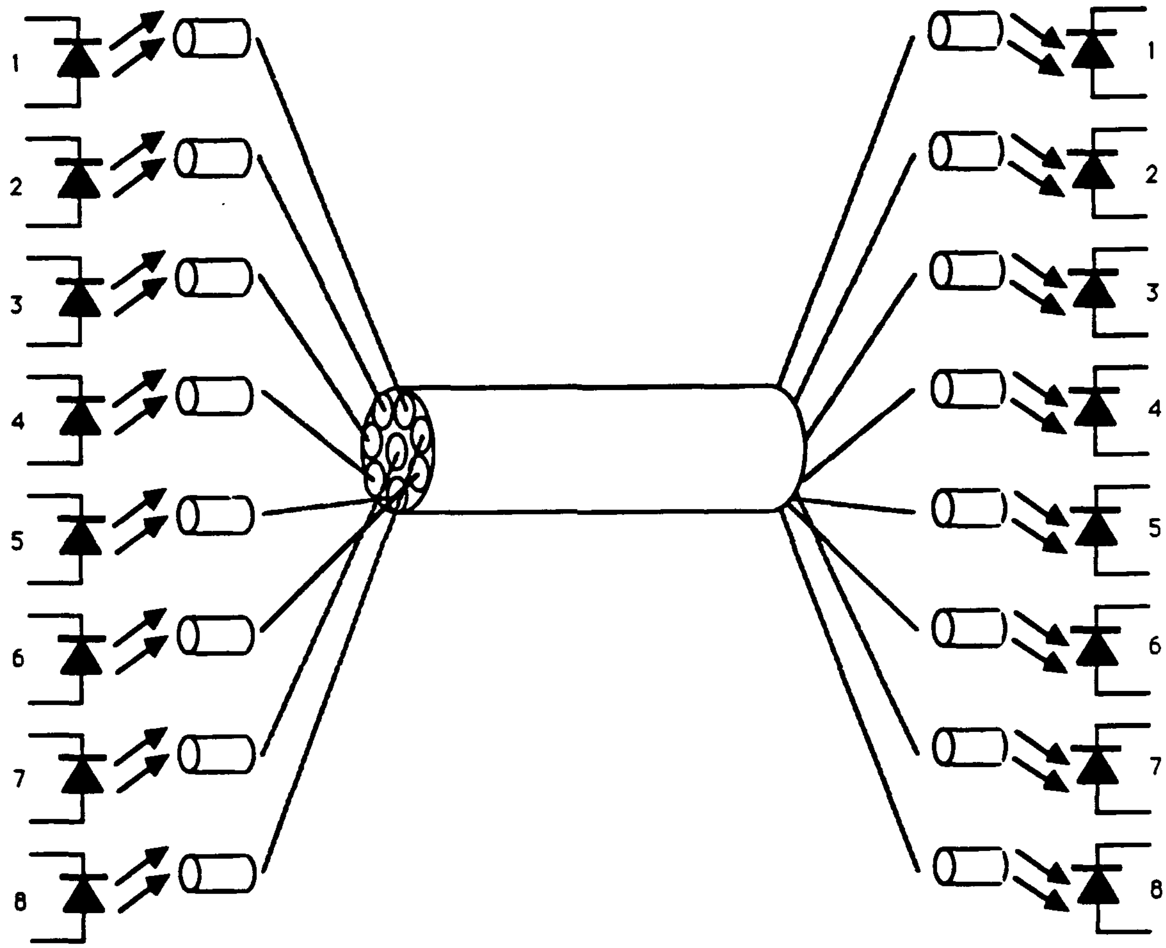

Figure 5.1: Eight Channel Space Division Multiplexed Link

5.2 Code Division Multiplexing

In code division multiplexing each channel is distinguished through the use of orthogonal codes and the assignment of a particular orthogonal code to each channel. Codes are chosen so that the peak autocorrelation of a code sequence is much higher than the sum of any other value of its autocorrelation function and the peak crosscorrelations of the code sequence with all other code sequences. Optical direct detection systems are "positive" systems [35], in that there is no way for the optical signals at the receiver to sum to zero, as there would be if ± 1 could be sent instead of 0 and +1. Therefore code sequences are different than those that may be used in spread spectrum systems over other media. An algorithm for generating optical orthogonal codes is specified by Salehi in [35].

When the transmitter for a particular channel wishes to send a logical one to the receiver it transmits its assigned code sequence. Signals from all transmitters at a node are coupled into the same fiber and sent to the receiver. At the receiving node the signal is split among the receivers. Once split, the signal enters a tapped delay line in each receiver, where the relative position of the taps corresponds to the code sequence from one of the transmitters. Thus, a channel is formed. Outputs from the taps are added, and when the sum exceeds a threshold, set just below the expected level for autocorrelation, a logical "1" is assumed. Channels can be completely asynchronous with respect to one another, i.e. bit boundaries do not have to be aligned. The code division multiplexing concept is illustrated in Figure 5.2.

To perform code division multiplexing the pulse, or "chip," rate must obviously be much higher than the data rate. Specifications for an orthogonal code usually include its length,

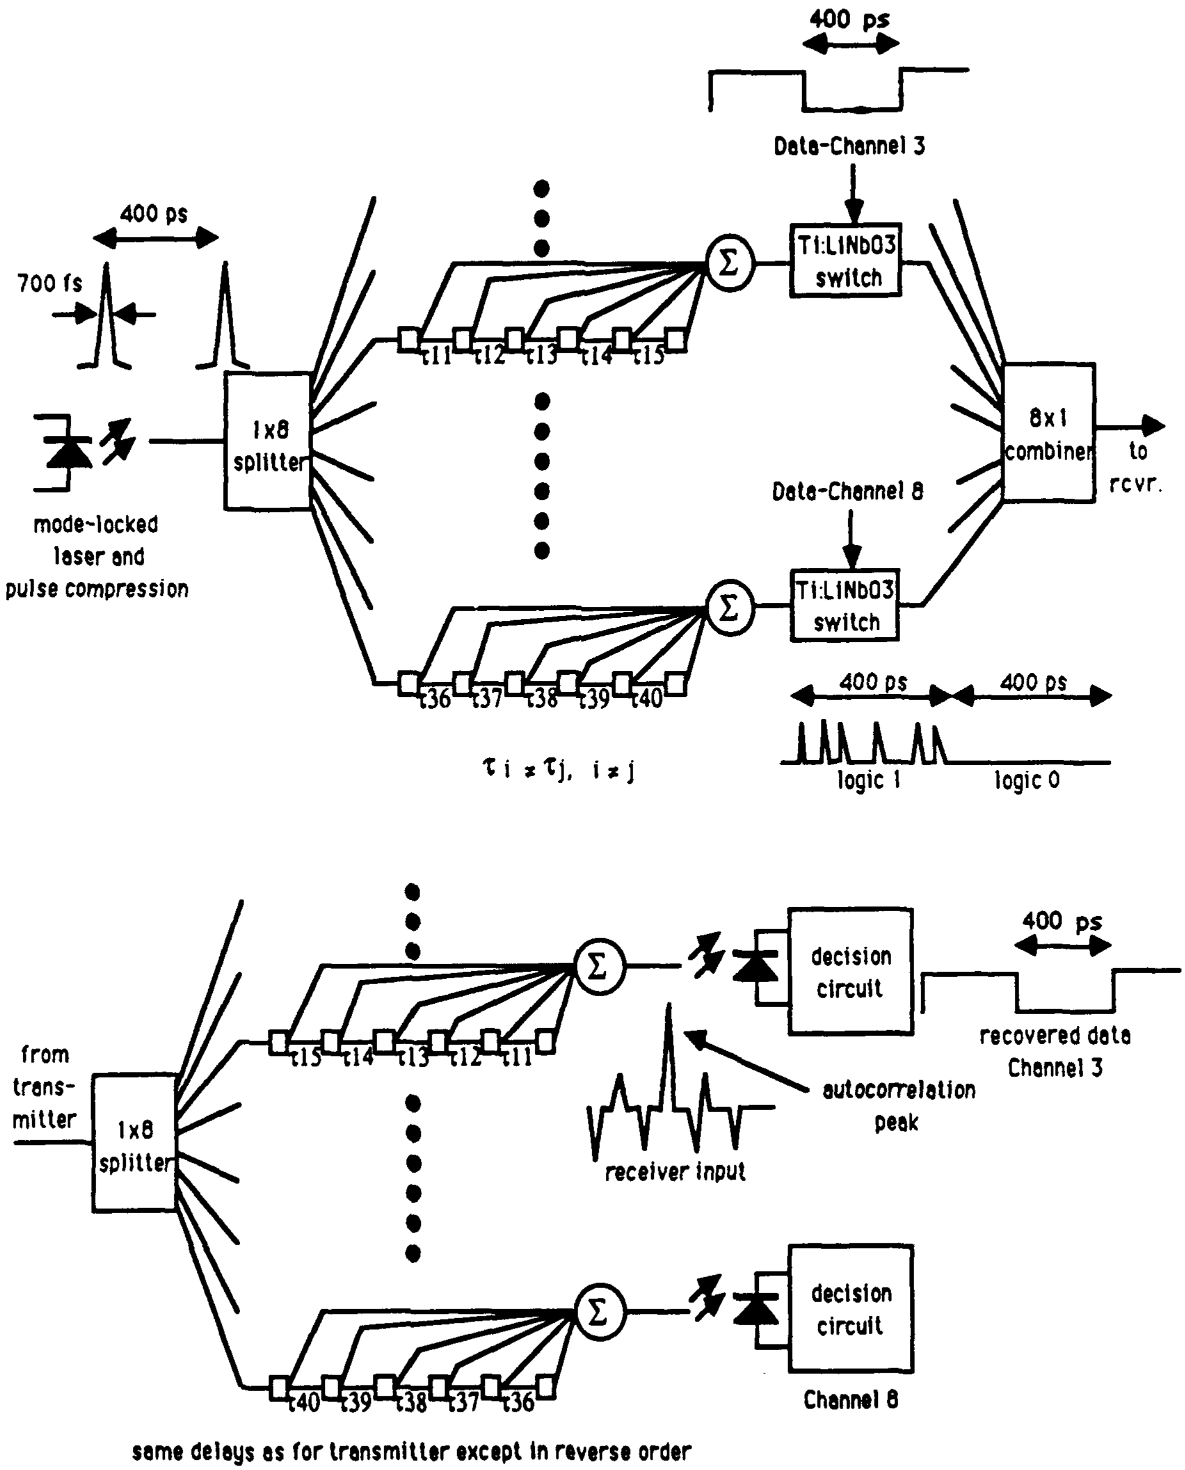

Figure 5.2: Eight Channel Code Division Multiplexed Link (channels 3 and 8 shown)

or number of chips, F, and its weight, or number of pulses sent per code sequence, K. For a code sequence where the autocorrelation peak is K times higher than the rest of the autocorrelation function and the cross-correlation functions, the number of orthogonal codes, N, available is:

$$N \leq \lfloor (F-1)/(K(K-1)) \rfloor, \qquad \text{(Equation 5.1)}$$

where the operation $\lfloor\ \rfloor$ is here taken to represent the greatest integer function [35]. Therefore, for eight discrete channels and a weight of nine, a code length of 576 is needed. If it proved necessary to increase the capacity from eight channels to nine and the weight from nine to ten, the length of the code would have to be increased to 810, and different code sequences would have to be used. The weight of nine was initially chosen because it could be possible for the sum of the cross-correlation functions of a given channel with the remaining seven channels, plus the shifted value of its own autocorrelation function, to equal eight. If other noise sources were negligible and the threshold were above eight, then the bit error rate would be zero in this case.

The weight of the code can actually be reduced below eight with a bit error rate less than 10^{-9}, due to the low probability of the cross-correlations summing in the fashion described above. An upper bound on the probability of error is given in [36] as:

$$P_{error} = (1/2) \sum_{i=Th}^{N-1} \binom{N-1}{i} \left(\frac{K^2}{2F}\right)^{i} \left(1 - \frac{K^2}{2F}\right)^{N-1-i}, \qquad \text{(Equation 5.2)}$$

where Th is the threshold level of the decision circuit. For example, a ten channel system with a length of 1000 was found to require a weight of seven and a threshold of just below seven, for a bit error rate of 10^{-9} [36].

For a bit rate of 2.49 Gbps, the chip duration for F = 576 is about 700 fsec. Pulses of the chip duration can be generated through the use of a mode-locked laser, in combination with the optical pulse compression techniques discussed in section 3.5.3. The correct code sequence can be obtained for a given 2.49 Gbps channel by creating these pulses approximately every 400 psec, then splitting each pulse into four parallel delay lines, which are delayed by differing amounts using short lengths of optical fiber to create the code sequence for the channel. The code sequence is then gated by the channel data using a $Ti:LiNbO_3$ (Lithium Niobate) optical switch. For a "1" a code sequence is sent, while no sequence is sent for a "0." This approach is used for a code division multiple access network, where nodes rather than channels are assigned code sequences, by Prucnal, Santoro, and Sehgal in [37].

If the optical output power is high enough, and nonlinear effects are not an issue, then a single mode-locked laser can serve as the optical source for all channels. However, splitting losses alone for the above-described eight channel link would be 23 dB in this case, as opposed to 14 dB if a separate laser were used for each channel. Additional excess losses, especially those for the optical switch, could be high, resulting in a channel that is severely power-limited. The use of optical solitons could be used if the soliton width was much narrower than the chip duration.

5.3 Optical Time Division Multiplexing

Optical time division multiplexing differs from conventional time division multiplexing only in that the mechanism for generating slots and relegating channels to the slots is done optically. As with the code division multiplexing system just described, a single mode-locked laser can be used to generate short pulses at the channel rate. For this technique the pulse stream is split among all channels. For an eight channel link having channel data rates of 2.49 Gbps, the pulse duration should be 50 psec, repeated every

400 psec. Optical solitons of narrower pulse duration could conceivably be used as well. The pulses are then gated by the individual channel data streams using a Lithium Niobate switch, and then delayed by differing amounts depending on the channel number using a fiber optic delay line. The pulse streams are then combined. Channels must be bit synchronous with respect to each other. The system is represented in Figure 5.3.

At the receiver end there are at least two proposed methods to demultiplex the data without requiring bandwidths in the receiver demodulation electronics higher than the channel data rate. One approach, similar to that proposed in [37], involves sending the optical pulse stream generated by the mode-locked laser on a separate fiber to the receiving node. This pulse train could be split according to the number of channels, and delayed so that clock pulses are aligned with the appropriate bit position for a given channel. This delay would have to be variable, to compensate for phase changes that may randomly occur in the optical clock due to temperature effects on the propagation delay of the fiber and random polarization changes. The clock pulse is then added with the data stream at each receiver, in effect placing the data for the receiver's channel upon an optical "pedestal." Figure 5.4 shows idealized representations of the optical clock, the optical time division multiplexed data, the sum of the optical clock and data stream, and the recovered data for channel five of an eight channel link.

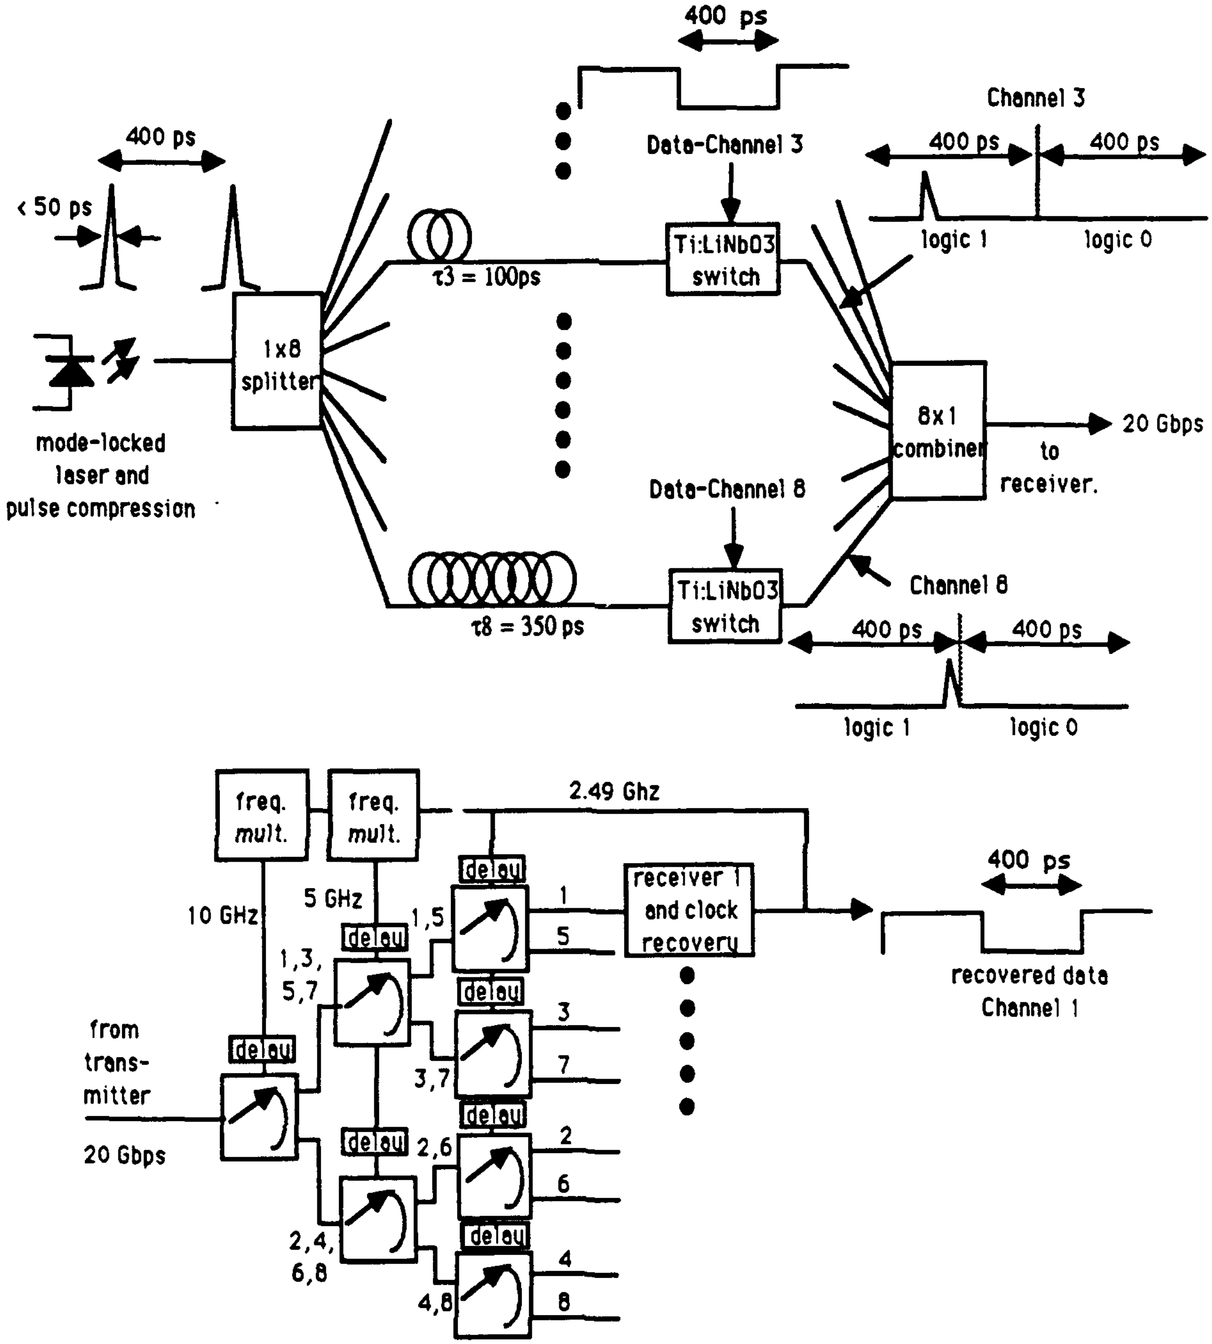

Figure 5.3: Eight Channel Optical Time Division Multiplexed Link with Optical Switches at Receiver (as in [38])

Another technique, which is more reliable and does not require sending an optical clock from transmitting node to receiving node, is illustrated in the receiver shown in Figure

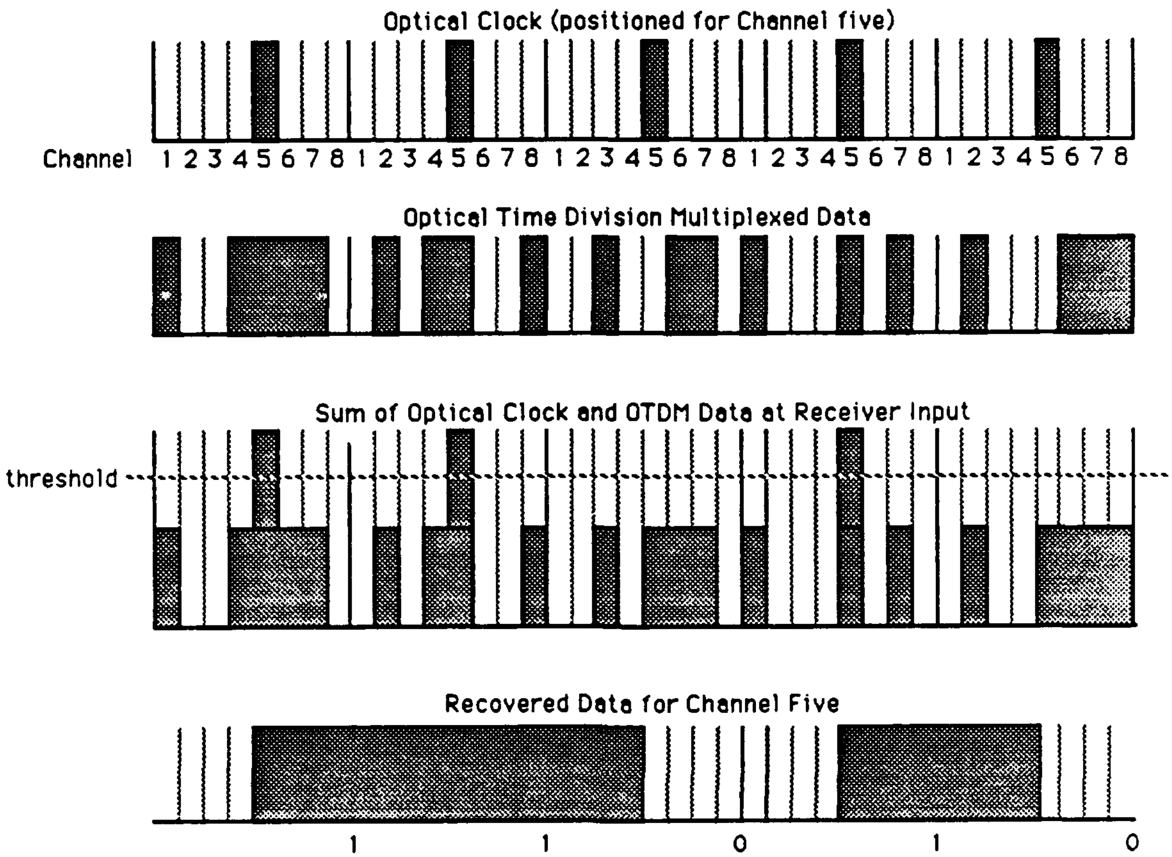

Figure 5.4: Idealized Optical Waveforms for Optical Time Division Multiplexed Link

5.3 [38]. Upon synchronization, involving manipulation of the delays, the recovered clock signal, a sinusoid at the channel frequency, is multiplied up by factors of two to the multiplexed bit rate (in three stages for an eight channel example). The sinusoids drive Ti:LiNbO3 switches which demultiplex the bit stream as shown. Although the electrical clock signals are at the higher speed, they are relatively easy to generate due to their sinusoidal nature. Demodulation in the receiver still occurs at the channel rate. Using this technique, four 4 Gbps channels have been time division multiplexed to 16 Gbps and transmitted over 8 km of single-mode fiber. [38]

5.4 Wavelength (or Frequency) Division Multiplexing

Analogous to the separation of radio and television stations into different transmission frequencies, different wavelengths, or colors, can be used for channels on the fiber. This technique is called wavelength division multiplexing (WDM) if the channels used are widely separated, and frequency division multiplexing if the channels are narrowly separated and coherent detection is used. Modulated optical signals are combined and demultiplexed using wavelength differentiating devices such as diffraction gratings, dichroic filters, or prisms. Because of the wide range of useful wavelengths, the decreasing spectral linewidths of the laser sources, and the promise of coherent detection, the potential number of channels that can be supported on a wavelength division multiplexed link is very large and growing. Another advantage of this approach is that channels can be added at unused wavelengths without affecting the other channels Channel specifications, such as laser linewidth, can vary from channel to channel as long as spectrums do not overlap. A wavelength multiplexed link is shown in Figure 5.5. The filters at the receivers are not needed if the filter response of the wavelength division demultiplexer is sharp enough to sufficiently reduce crosstalk from other channels.

Wavelength division demultiplexing of sixteen 1.5 Gbps channels, spaced by two nanometers in the 1550 nanometer wavelength window, has been demonstrated using conventional distributed feedback (DFB) lasers [39]. Channel spacings can be reduced to under one nanometer using tunable etalon filters [40]. Laser spectral widths are being reduced to less than 100 MHz (one nm =125 GHz at 1550 nm) using newly developed DFB lasers, and coherent detection techniques are progressing [41], meaning that the number of channels that can be implemented in the 1550 nm window of the optical fiber will increase, along with the maximum transmission distance.

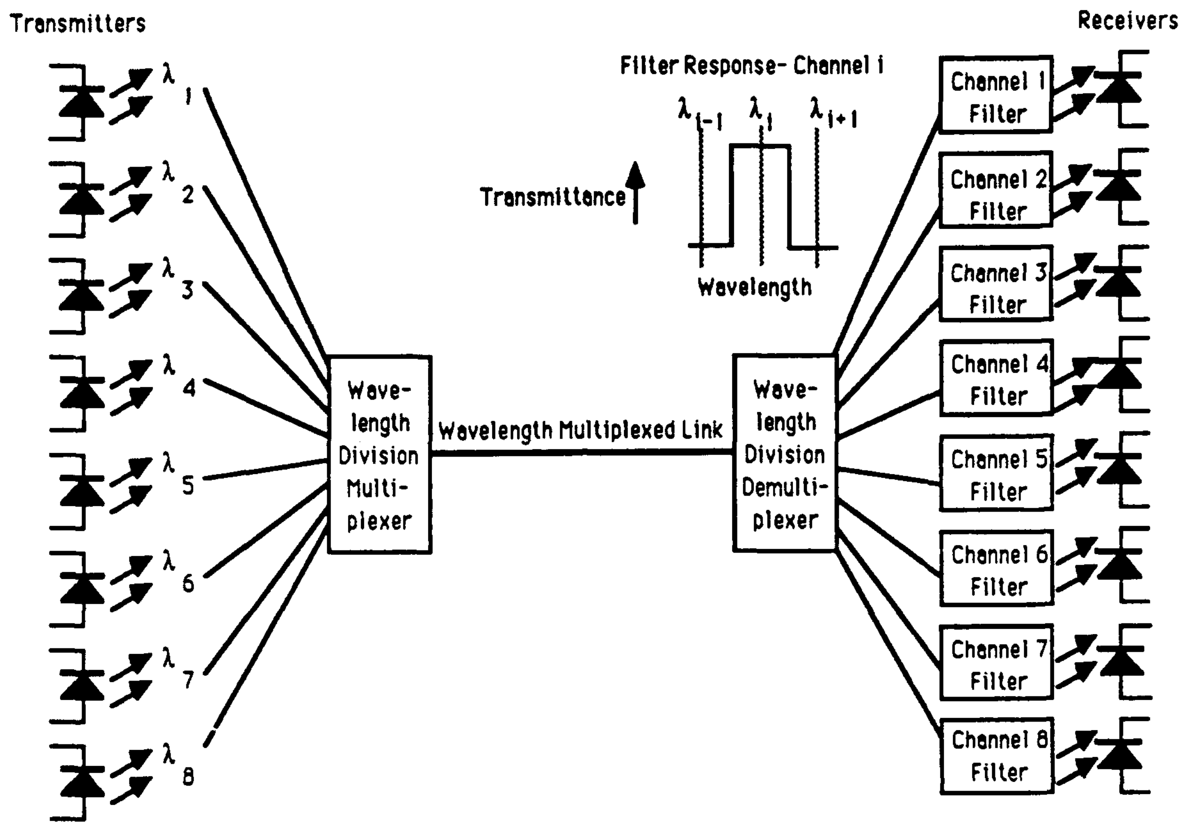

Figure 5.5: Eight Channel Wavelength Division Multiplexed Link

Technical issues in the development of dense wavelength multiplexed systems are discussed in detail in Section 7. Current research areas include expanding the tunability ranges of lasers to more than a few nanometers, which would allow a laser to be set to a particular channel; and controlling drift in the spectral characteristics of a laser over time and temperature, which could cause it to move out of its allotted channel. Reducing device sizes is also being investigated. The quaternary compounds based on Indium Phosphate, necessary for operation in the low dispersion and low attenuation windows at 1300 and 1550 nm, are difficult to fashion into opto-electronic integrated circuits (OEICs). OEICs would reduce the size and cost of the many electro-optic and photonic devices required to implement wavelength division multiplexing. Another interesting theoretical and practical challenge is the use of optical solitons in a multi-wavelength environment.

5.5 Comparison of the Optical Multiplexing Techniques

Comparison of the multiplexing techniques in a general sense will be somewhat arbitrary given the influences of application specific variables like geographic span, number of channels required, data rate, and time frame of implementation. However, if we assume that a physical link connecting two nodes on the ring will be comprised of eight channels each operating at a high data rate, for example the SONET OC-48 rate of 2.49 Gbps, then we can make rough comparisons of the four multiplexing techniques based on several important criteria, including reliability, cost effectiveness, complexity, and growth potential. In Table 5.1 the multiplexing techniques are rated on a scale of one to ten, with ten representing the best performer relative to a given criteria. Assuming that the identified criteria are all equally important, the figures are totalled and a conclusion is reached.

Reliability - If a single mode-locked laser is used to generate the optical pulses for the code or optical time division multiplexed techniques, then its failure would obviously cripple the link. If, in the optical time division multiplexed system an optical clock was sent to the receiver, failure of that channel would also cause the loss of the whole link. Space division multiplexing is obviously the most reliable multiplexing technique, because the link would not be eliminated by a single fiber cut. Single fiber cuts, rather that complete cable cuts, occur in some failure situations, including those where the cable is damaged by burrowing animals. It would take a complete cable cut to eliminate the link. Also, crosstalk from malfunctioning lasers would not affect correctly operating channels.

Table 5.1: Comparison of Optical Multiplexing Techniques

MUX TECHNIQUE / CRITERION	SPACE DIVISION MULTIPLEX	CODE DIVISION MULTIPLEX	OPTICAL TIME DIVISION MULTIPLEX	WVLNGTH/ FREQUENCY DIVISION MULTIPLEX
RELIABILITY	10	2	6	6
COST EFFECTIVENESS	3	5	8	8
COMPLEXITY	10	2	5	7
GROWTH POTENTIAL	2	3	6	10
TOTAL	25	12	25	31

Cost - Regarding the cost of the multiplexing technique, the choice between space division and wavelength division multiplexing would be based on balancing the extra fiber cost of the space division approach with the extra optical component cost of the wavelength division approach. Optical bandwidth is utilized more efficiently with wavelength division multiplexing. Although the incremental cost of extra optical fibers is small compared to the cost of installing multi-fiber optical cables, the extreme increase in the number of fibers required for the space division approach could likely result in very high cost, given the large geographic extent of the network and the probability that the fibers will be leased. This cost would probably outweigh the cost of the wavelength division demultiplexer, the optical combiner, and optical isolators that would be required for the wavelength multiplexing approach. For this reason wavelength division multiplexing will likely be the most efficient multiplexing technique.

If use of a single laser source is possible, the optical time or code division multiplexing techniques could allow for a cost effective use of optical source capabilities. However, the bandwidths of transmitters used in these systems would have to be much higher than for the other two techniques, and these two techniques require complicated optical systems for multiplexing and demultiplexing. Their total cost will likely be higher than that of the wavelength multiplexed system. The bit rate over the optical fiber is much lower for the optical time division technique than for the code division technique. A single source could be used for space division multiplexing as well, with a much simpler multiplexing system and the elimination of the need for demultiplexing optics.

Complexity - Although the number of sources and detectors for the space and wavelength multiplexing approaches is the same for a given number of channels, the complexity is not. Each channel in a space division system is exactly identical, which would allow for the use of standard equipment for all links, and the easy replacement of failed components. The wavelength multiplexing technique would generally require either individual selection of sources or tunability of sources, and laser stabilization could be required depending on the channel spacing. However, the optical bypassing of failed stations is much more difficult, and the cost and complexity of repeaters is much greater, if space division techniques are used. The complexity of the higher layers is unaffected by this choice. The precise delays and the narrow optical pulses required for both code and optical time division multiplexing make these approaches relatively complex, with code division multiplexing being the most complex.

Growth Potential - the multiplexing technique used also affects the growability of the network. Networks using optical time division multiplexing have to increase optical switching speeds and alter delays at the transmitter and receiver to add capacity, affecting the network availability during cutover and therefore not allowing for smooth growth. Code division multiplexing techniques limit the number of channels due to the limited number of orthogonal codes that exist in a given code length. With wavelength or space division multiplexing, network capacity can be increased by adding new wavelengths or fibers parallel to the existing channels, generally without affecting in-place physical layer electronics, or availability. However, with space division multiplexing, as previously mentioned, the cost of adding additional fibers to a link may be very costly. With wavelength (frequency) division multiplexing, there is the potential for smooth growth to multi-terabit per second capacity through the use of coherent techniques.

Given the significantly higher total in Table 5.1 for wavelength division multiplexing we can conclude that it would be preferable for future Air Force applications. This is largely due to its much greater growth potential relative to the other multiplexing techniques. Code division multiplexing proved to be the least desirable choice due to its low score for all criteria used here.

6. Logical Architectures for the Physical Counter-Rotating Ring

Due to the sensitive nature of many of the military applications that would be operating over it, survivability and availability are the two most important attributes of the network. Beyond these requirements the network must be able to provide the complete range of services required by the user, both upon installation and throughout the lifetime of the network, in a reliable and efficient manner. Consistent with current trends in government and military applications, the network must also follow as closely as possible the standards emerging from the commercial sector. In this section, we will define potential network configurations and look at how these configurations measure up against the above ideals.

This section first discusses a basic assumption that will likely be used in hypothesizing alternative logical network topologies and architectures for the future. This assumption is that the network will use Asynchronous Transfer Mode (ATM), a fast packet switching technique. We are also assuming that the network will be comprised of "multi-channel" links connecting the nodes. Four multiplexing techniques for constructing the multi-channel point-to-point links were discussed in the previous section.

Then, the various logical architectures that can be overlayed upon a physical counter-rotating ring, concluded in Section 3 to be the most efficient physical topology meeting minimum survivability requirements, will be defined. In this section it is assumed that all network nodes are peer entities (probably gateways), likely to be operating at the top of a network hierarchy.

Next, the logical architectures discussed in this section are assessed relative to their reliability, complexity, efficiency, fairness, growth potential, cost, and security.

Conclusions are drawn as to which configurations hold the most promise for use in survivable high speed wide area networks.

Obviously, different logical topologies, such as the dual hubbing approach discussed by Wu, Kolar, and Cardwell [7], would be appropriate if the nodes were not all peers, as in the region between local area networks or central offices and the gateways at the top of the hierarchy. Survivable hierarchical architectures will be covered briefly at the conclusion of this section.

6.1 Fast-Packet Switching

The wide area networks of the future will be integrated services networks, i.e. they will support a variety of services, including voice, video (video telephone, video conferencing, standard NTSC TV, and HDTV) and data, with the same physical layer common to to all services. Physical layer commonality is one of the major steps called for by the Defense Communications Agency as it upgrades the Defense Communications System for the twenty-first century [42]. Having one physical layer for voice, data, and video leads to a greater efficiency in providing services when two or more service types are required for an application (such as video telephone). It also leads to lower costs due to economies of scale, and to greater availability due to greater ease in adding stations or capacity. These advantages apply even more strongly in a wide area network environment, where distances between switching points are likely to be great.

Most private government and military networks do and will likely continue to lease the transmission (and possibly the switching) facilities of the public network. For example, the Department of Defense's Automatic Voice Network (AUTOVON) leases both the dedicated switching facilities and the transmission facilities from telecommunications companies [43], while the Defense Data Network (DDN) leases the transmission facilities with the switching performed by Bolt, Beranek, and Newman (BBN) C/30 Packet Switching Nodes (PSNs) located at DDN secure facilities [44]. Network compatibility with emerging standards in commercial telecommunications (i.e. ANSI T1 and CCITT), to the extent that these standards meet military requirements for security and survivability, is therefore very important [45]. Special attention should be paid to the Integrated Services Digital Network (ISDN) standards [46], the Synchronous Optical Network (SONET) standards, and the Asynchronous Transfer Mode (ATM) work being done for the Broadband-ISDN (B-ISDN) standards. These telecommunications standardization efforts are summarized in Appendix A.

Consistent with government trends toward greater standardization in networking, we assume that ATM cells will become the unit of information on the network of the future. The use of ATM means that the time domain on a link is divided into 53-byte fixed intervals (i.e. cells, slots, or fast packets). Each cell contains a 5 octet header containing minimal information including 8 to 12 bits for the cell's virtual channel identifier (VCI), an eight-bit header error check, and other information like a grade-of-service field.

There is no fixed assignment of slots to particular nodes in ATM. A node can use any vacant slot and can harness as much of the available bandwidth as necessary as long as it does not "hog" the network. This permits the network to efficiently handle traffic from many different types of sources. Synchronization cells are used to fill vacant slots. The distinction between ATM and conventional synchronous transfer mode, or synchronous time division multiplexing, can best be understood through Figure 6.1 [47]. In conventional time division multiplexing, each channel is assigned a slot position within a repeating frame structure. In ATM any unused slot can be accessed by a channel, with the requirement of the short header.

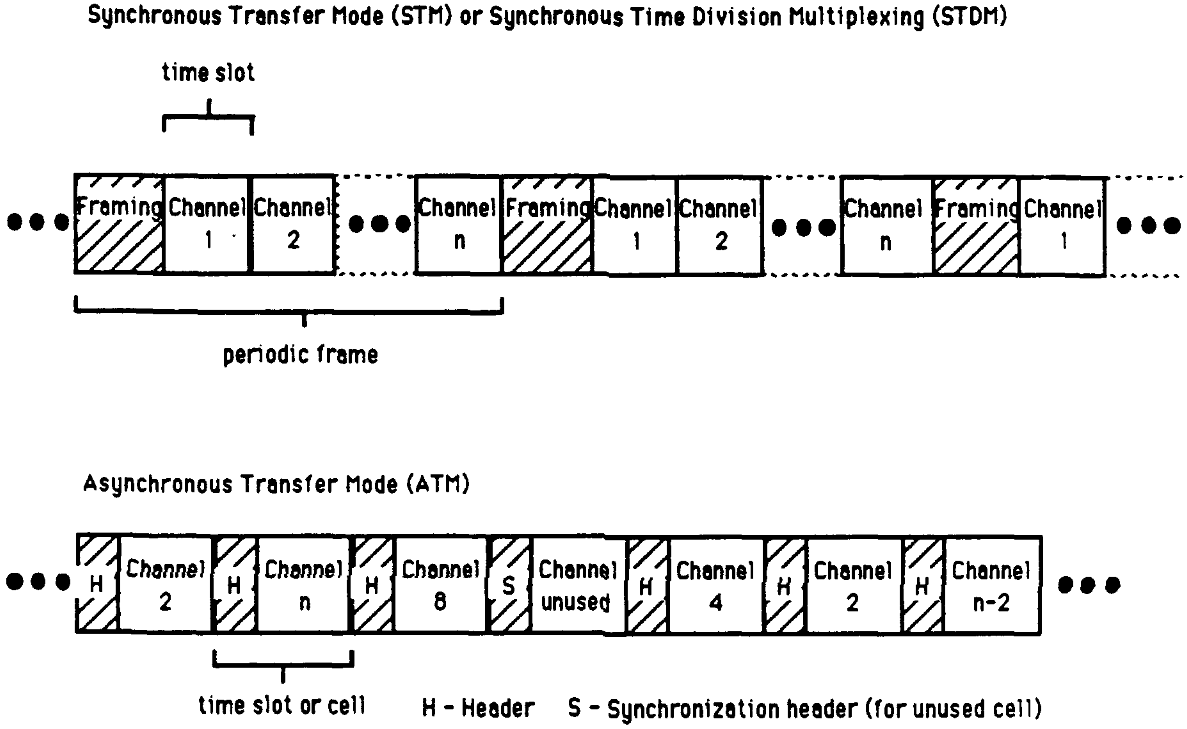

Figure 6.1: Synchronous Transfer Mode (STM) and Asynchronous Transfer Mode (ATM) [47]

The possibility of B-ISDN and ATM has led to the development of a number of self-routing, fast packet switching fabrics [48-53]. Many of these approaches internally employ Batcher sorting and Banyan routing techniques (for example, see ref. [48]), with various mechanisms to avoid the collision of packets at the output links. The approaches generally can accommodate multi-cast and broadcast functions (necessary for video conferencing and network management), and can incorporate priority levels to speed urgent packets through the network.

6.2 Logical Topologies

In this section we discuss six topologies that can potentially be overlayed onto the physical counter-rotating ring topology chosen in Section 3. We assume that all nodes on the ring are peer entities. This condition would arise if the physical ring serves as the top layer of a wide area network hierarchy and all nodes are gateways, perhaps between the network and a number of local area and metropolitan area networks or digital central offices, or the communications port of a supercomputer.

In the networks described here we assume that ATM will be used in the time domain, and that multiple channels will exist between nodes, using one of the techniques discussed in the previous section. Furthermore, it is assumed that the logical architecture will retain the level of survivability of the physical ring. If the network is partitioned due to multiple failures the remaining nodes forming subnetworks must be capable of resuming communication within themselves, implying that the network must have distributed control and management. Each node must have the capability of restarting the management process after a failure or error is detected.

For survivability, there should be no centralized clocking mechanism that could be a targetable single point of failure. A plesiochronous clocking system would be preferable. In *plesiochronous clocking* each node has a clock for transmitting information onto a link. This clock may differ by a small unknown amount from other node clocks in the network due to temperature, ageing, and manufacturing tolerances. The maximum deviation that can exist between clocks is fixed and specified in the design.

In plesiochronous clocking, receivers recover a clock from the incoming bit stream and use this recovered clock to retime the incoming data. An elasticity buffer is used to convert the data rate of the incoming bit stream to the operating rate of the resident node. Data enters the buffer at one rate and leaves it at the other. For a given buffer size and specified maximum deviation between node clocks, the buffer must be periodically reset to prevent overflow or underflow. For example, an elasticity buffer of ten bits operating in a network where the maximum clock deviation was controlled to within ± 50 parts per million (ppm) of some nominal frequency would require resetting of the buffer every 50,000 bits to assure that overflow or underflow would not occur. Resetting of the buffer could occur during the periodic transmission of synchronization cells in an ATM environment. A synchronization cell would have to be sent at least once for every 117 cells sent.

We also assume that the secondary ring in the counter-rotating ring is fully utilized, in the same fashion as the primary ring. While this assumption may not be valid in local area networks, such as FDDI, where transmission resources are less expensive relative to switching resources, it is valid in a wide area network. Even though maintaining both rings as peers increases the complexity of the nodes, it would be extremely inefficient to maintain the secondary ring only for standby purposes, due to the generally higher cost of transmission resources relative to switching resources in a wide area network environment.

The use of the secondary ring as a peer to the primary ring greatly improves the efficiency of the ring. Average shortest path propagation delays between nodes are reduced from one half the ring transition time to one quarter the ring transition time, assuming a balanced traffic pattern. The overall capacity is doubled. Additional routing information would have to be maintained at each node to determine which ring should be used to send information to a particular destination. This routing information could be either static, or dynamically updated to reflect changing traffic patterns on the network.

6.2.1. The Multi-Channel Ring Network

Obviously, a logical ring topology could be overlayed on a physical ring topology. We first discuss a multi-channel version of the ring.

Multi-channel configurations have been proposed for token-passing ring and CSMA/CD networks [54], and also token-passing bus networks [55]. Here, we propose a multi-channel configuration for a slotted ring. Each channel in the physical link connecting two adjacent nodes serves as a serial channel for ATM cells. ATM cells could be transmitted side-by-side on the same physical link. When cells reach a node, the virtual channel and path identifier is examined to see if the cell is to be removed at that point. If not, then the node sends the cell onward over <u>any</u> available logical channel. The distinction that cells can use any unused channel between nodes turns this concept into a single multi-channel network, rather than several single channel networks operating in parallel, and allows for more efficient use of the transmission media. With this approach, physical links are not even required to support the same number of channels throughout the network. A link between two nodes sharing a high volume of traffic could have more channels than a remote link.

A unidirectional eight-channel physical link in the multi-channel slotted ring is shown in Figure 6.2. In the system shown here, *slot boundaries are asynchronous over the eight* channels. Due to differences in propagation times between the various logical links in the media, skewing of cell positions could occur naturally over long transmission distances,

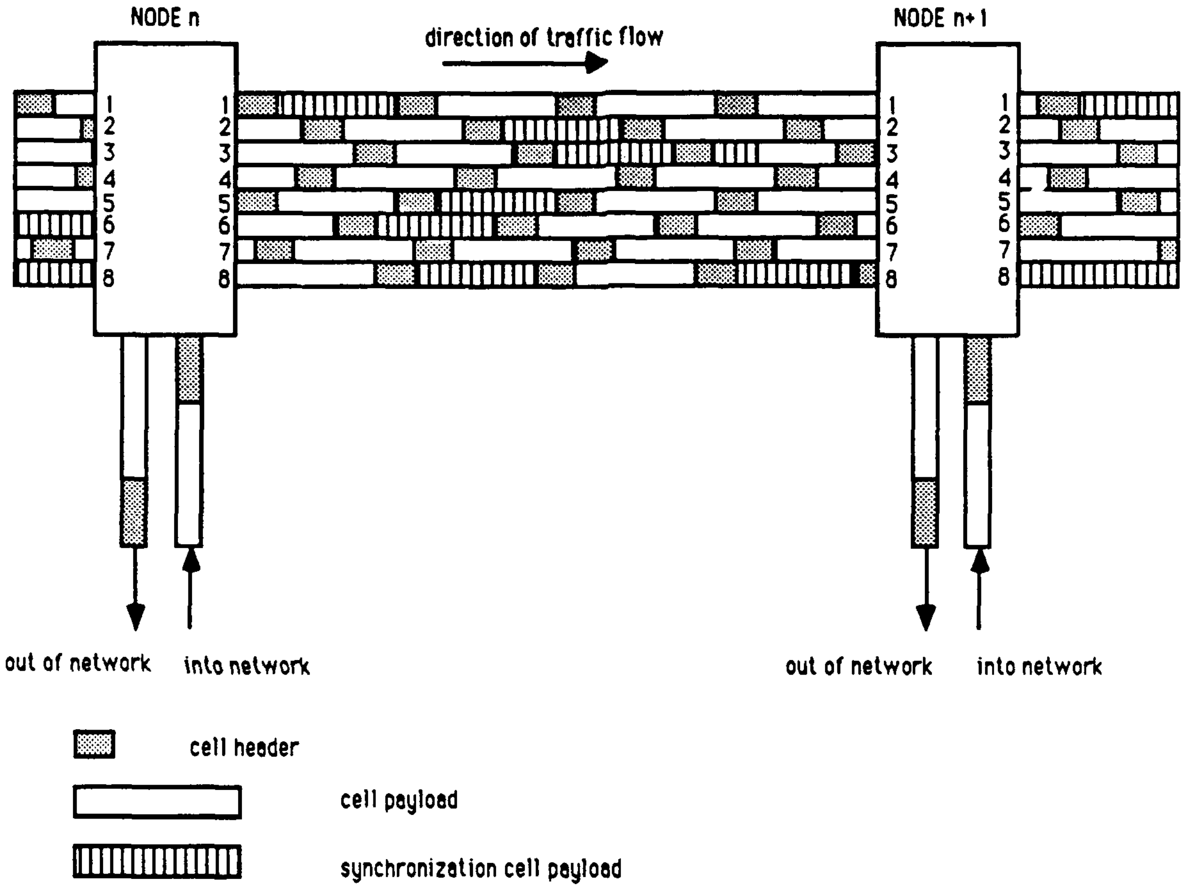

Figure 6.2: Multi-Channel Slotted ATM Ring

particularly if wavelength division multiplexing were used to create channels on a fiber optic link. Skewing may also be induced to allow for sharing of the processing capability needed to handle cell header information. It may also result in the need for larger buffers at the nodes, because it would make it more difficult for a node to immediately match an incoming cell with an outgoing channel. In a lightly loaded network this skewing may allow for quicker matching of a cell arriving at the node from outside the network with an outgoing link, reducing buffer requirements.

Propagation delay differences between channels in the same physical link must be controlled so that a cell starting out behind another cell from the same message can not overtake the first cell later in the network, causing misordering of information at the destination. If cell boundaries are not resynchronized at each node, then this delay difference accumulates through the network.

Problems of fairness can arise when a node or link failure occurs and the ring goes into wrap mode (as in Figure 3.5). In wrap mode the primary and secondary rings are joined to form a single ring, but nodes which are not immediately adjacent to the point of failure have two access points to the reconstructed single ring. The two nodes which perform the wrap function around the failure have only one access point. Nodes on the opposite side of the ring from the failure have shorter transmission distances to other nodes, on the average, than nodes nearer the point of failure.

In the event of ring wrap, the protocol could include a mechanism to give highest priority to cells from the nodes which performed the wrap, with cell priority decreasing away from the wrap nodes. Organization of priority sublevels according to the age of a cell, and the use of a protocol which sends the oldest cells first, independent of their origin, could also reduce problems of fairness. Another possible solution is to begin operating the network as if it were a dual bus, with the nodes adjacent to the failure serving as cell origination and termination nodes. It is possible, through increased complexity, to achieve satisfactory fairness with a dual bus protocol, as has been done in the Distributed Queue Dual Bus (DQDB) IEEE 802.6 standard [56].

6.2.2 The Single Channel Byte-Wide Ring Network

Networks having multiple channels between nodes are not necessarily multi-channel networks. In the single channel byte-wide network, channels are tied together so that cells are sent in a parallel, or byte-wide, fashion, instead of the bit serial fashion of the first approach. Therefore, only one cell is transmitted at a time on a unidirectional physical link, but the transmission rate per cell is eight times faster (assuming eight parallel channels) than for the multi-channel network, assuming bits are sent at the same rate in each network. The overall cell throughput is therefore the same. Figure 6.3 shows how cells travel between nodes for the single channel byte-wide network. This figure should be contrasted with Figure 6.2.

The protocol for the single channel byte-wide slotted ring network is similar to that for an individual channel of the multi-channel network. The virtual channel or path indicator of the cell header for each cell is checked to see if the current node is the cell's exit point. If not, the cell is sent onward, upon availability of the output link.

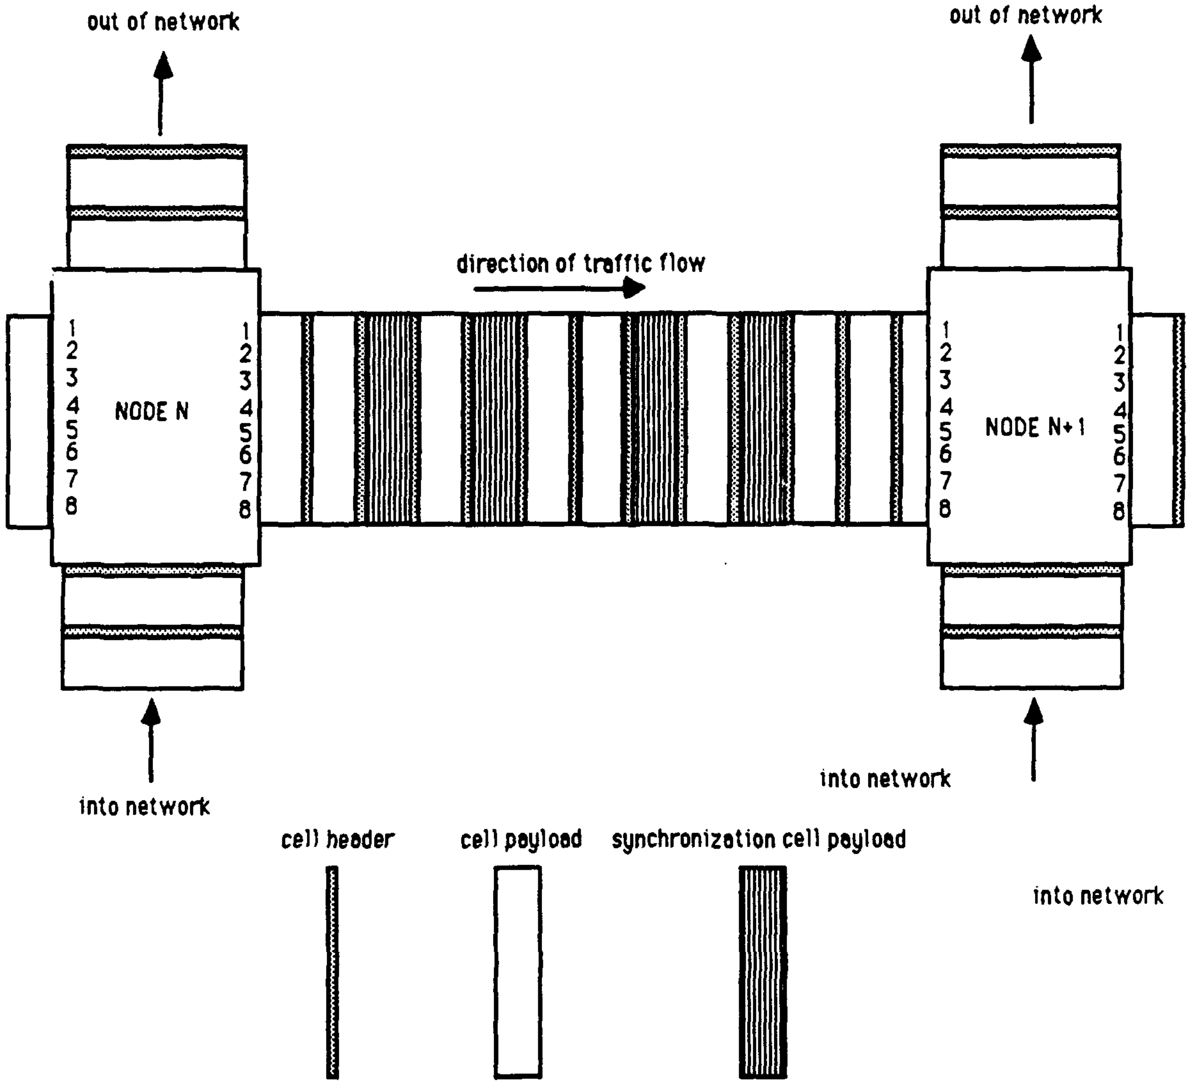

Figure 6.3: Single Channel Byte-Wide ATM Network

The concept of a byte-wide channel using wavelength division multiplexing on optical fiber has been proposed by Loeb [57] as a means to take advantage of the fact that information is

communicated in a parallel fashion within computers. The parallel-to-serial and serial-to-parallel conversions can be avoided with the byte-wide network. The issue of fairness when the ring goes to wrap mode is the same as for the multi-channel network.

Propagation delay differences between the different logical links must be more closely controlled in this network than in the multi-channel network so that it can be determined to which byte a received bit belongs. Loeb and Stilwell [58] have developed a technique to compensate for bit skews when operating in the 1300 nm region. The dispersion-limited data rate for sending data in eight-bit byte format is still less than that which could be attained on eight independent serial channels. If dispersion-flattened or shifted fiber were used the algorithm could be translated to 1550 nm, where attenuation is lower.

6.2.3 The Perfect Shuffle Network

The perfect shuffle network [59], shown for an eight node example in Figure 6.4a, has been developed as a means for expanding the capacity of networks having multi-channel capability in the transmission links, without requiring an increase in the transmission or processing capability of the nodes. Nodes are logically arranged in columns and each node is logically connected to two or more nodes in the adjacent column in the "perfect shuffle," or banyan pattern. The number of nodes is related to the number of columns by the following equation:

$$N = kp^k, \qquad \text{(Equation 6.1)}$$

where N is the number of nodes, k is the number of columns, p^k is the number of nodes per column, and p is the degree of the network, or the number of logical inputs (and outputs) each nodes has. Nodes are connected to all other nodes after an appropriate number of "hops." Routing information must be available at each node to determine which of the two logical output channels is appropriate for the given cell.

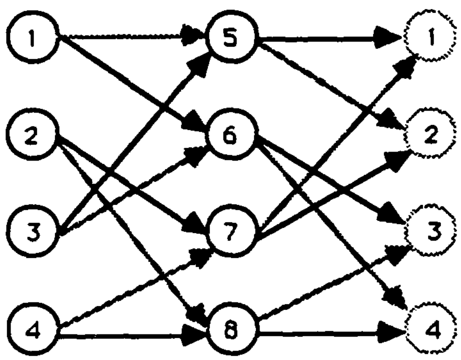

a. Perfect Shuffle configuration

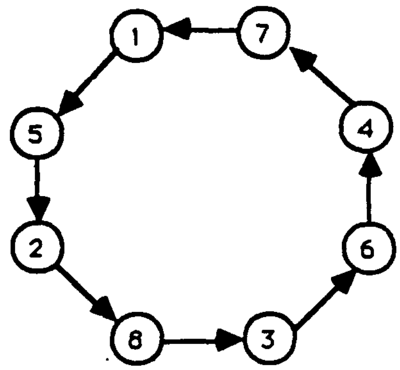

b. Correct node connection on physical ring

Current Node \ Destination	1	2	3	4	5	6	7	8
1	X	5	5	6	5	6	6	5
2	7	X	8	8	7	8	7	8
3	6	5	X	6	5	6	6	5
4	7	7	8	X	7	8	7	8
5	1	2	2	2	X	2	2	2
6	4	4	3	4	4	X	4	4
7	1	2	1	1	1	1	X	1
8	3	3	3	4	3	3	3	X

c. Next Node Routing Table

Figure 6.4: Perfect Shuffle Network

The implementation of the perfect shuffle network on a physical ring is shown in Figure 6.4b. The correct configuration for converting from the perfect shuffle pattern in Figure 6.4a to the physical ring can be found by weaving a path through the pattern without passing through a node twice until all nodes are included (and a cycle is formed). This path is referred to as a "Hamiltonian" and for the eight node network in Figure 6.4a is represented by the lightly shaded arrows. With this configuration the average propagation delay for a cell is one half the ring transition time, equivalent to that for the ring networks. Many channels physically pass through each node in the ring, but only two logically attach to each node.

The routing table is shown in Figure 6.4c. One row of the routing table is resident at each node in the network. The table determines on which of two possible output channels a cell should be placed in order to reach its destination in a minimum time. Because of the Hamiltonian ordering of the nodes and the routing table, cell propagation time is always less than one ring transition. The routing tables can be set up so that, even in the event of the failure of one of the logical links between nodes, any cell can still reach its destination, perhaps with an increase in the number of hops required. The routing table also attempts to balance the amount of traffic logically passing through each node (evident by the fact that each node appears on the routing table an equal number of times).

As an example of the operation of perfect shuffle routing we will describe the transfer of a cell from node one to node seven. Row one of the routing table, corresponding to the information that would be resident at node one, indicates that if the final destination is to be node seven, then the cell should be placed on the logical channel connecting to node six. Note that in the physical ring topology, the logical connection that exists between nodes one and six physically passes through nodes five, two, eight, and three. These intermediate nodes do not logically affect the connection. Once at node six, the routing table indicates that the cell should be placed on the channel to node four, where it will then be logically connected to node seven.

The advantage of the perfect shuffle approach can best be understood through Figure 6.5, showing as an example the internal configuration of node five from Figure 6.4. Due to the fact that the network is configured on a physical ring, several channels physically pass through node five for regeneration. Only two of these channels are logically connected to node five. Therefore, cell processing is only required for two of the six channels passing through node five.

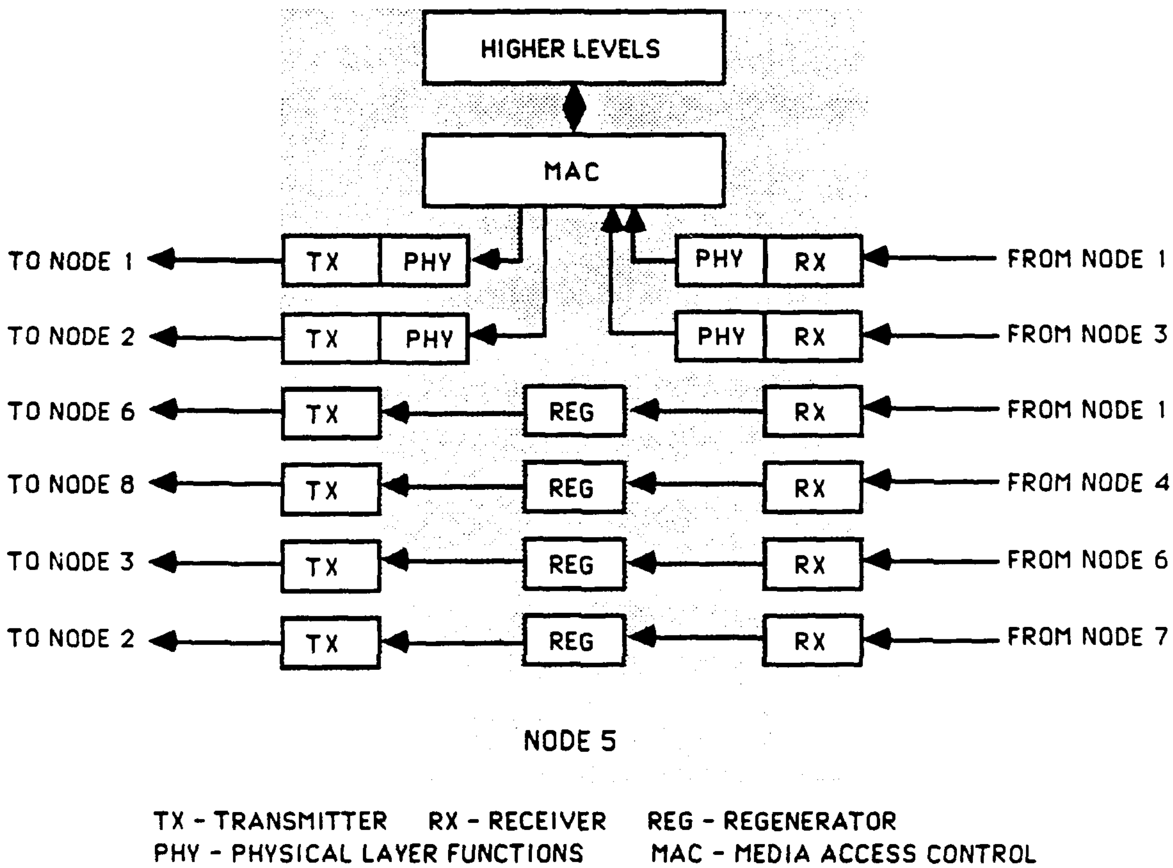

Figure 6.5: Internal Configuration of Perfect Shuffle Node Five

The secondary ring would most effectively be used to implement a second perfect shuffle network, similar to the perfect shuffle network on the primary ring except that propagation occurs in the opposite direction. Physical nodes would have different logical node numbers in the two networks. As for the logical ring networks, additional routing information would have to be available at each node to determine which of the two physical rings should handle each cell. Average propagation delays would then be reduced to one quarter of the ring transition time.

In the event of a failure requiring reconfiguration, issues of fairness arise that are similar to those for the multi-channel ring network.

To permit greater sharing of link capacity and to increase efficiency of the perfect shuffle network when implemented on a physical ring, we can add the variation that, if the logical channels conventionally available to a given node are congested, the node can use one of the other physical channels passing through it, provided it is unused. This variation requires that each physical channel either be logically, not just physically, terminated at each node, and that the cell contain information defining its logical channel number. The latter technique is used in the simulation study, included as Appendix B, comparing this network to the multi-channel and single channel byte-wide networks previously discussed. With this alteration the perfect shuffle network throughput becomes equal to the throughput for the other two networks, except that an extra field is required in the header to define the logical channel. The simulation study shows that this network also has similar delay and number-in-queue performance as the other two. With the variation the network becomes very similar to the multi-channel ring network, except for the extra header field.

Although we are here applying the use of the perfect shuffle topology on a physical ring, the perfect shuffle has been more commonly applied to other physical topologies. The use of the perfect shuffle logical topology was first proposed for a physical star or bus topology by Acampora, Karol, and Hluchyj in [60]. The use of the perfect shuffle on a physical ring topology as a potential upgrade to the FDDI or IEEE 802.6 standards was discussed by Karol and Gitlin in [59]. Their technique does not, however, maintain the immunity to single points of failure realized in those standards, because it does not utilize the secondary ring to establish a counter-rotating, peer network to that implemented on the primary ring.

A network using a physical and logical perfect shuffle topology and all-optical switching has also been proposed by Sauer in [61]. Sauer's network concept uses extremely short packet lengths (the payload is the length of a single computer word), where each bit is simultaneously transmitted at a different wavelength. Optical pulse compression and dispersion-flattened fiber is proposed to lower the probability of collisions between asynchronously transmitted packets. At the node receiver, the wavelengths representing the control information are separated and used to route the packet or remove it from the network. Lithium Niobate switches are proposed to perform this routing, with optical amplifiers used to regenerate the signal after the lossy switch. No buffering of the optical packets is performed; if the correct output for a packet is in use then the packet is simply sent over the other output, possibly causing its arrival at its destination to be delayed by one or more network transition times. The use of priorities with ageing is proposed to prevent indeterminant packet circulation.

6.2.4 Fixed Assignment Networks

In a fixed assignment network, nodes are allocated a portion of the available bandwidth on which to send or receive information. An example is the TDM Ring [62], where nodes are assigned slots in which they receive data. In an ATM environment, fixed assignment in the time domain would negate the advantages of ATM, but node assignments could be made in the wavelength, frequency, or space domain, or even through code division multiplexing using optical orthogonal codes [35]. The most likely technique for fiber optic networks is to assign each node a fixed wavelength (or frequency) on which it can receive.

An example of a fixed assignment network currently being proposed is LambdaNet™. In LambdaNet™ (a registered trademark of Bell Communications Research), the transmitter at each station is assigned a fixed optical wavelength. The optical signals from all stations are mixed in a star coupler and passed to the receiver portion of all stations. Demultiplexing occurs and each station identifies packets that were intended for it [63]. Another use of wavelength division multiplexing in a fixed assignment network involves assigning a particular wavelength to a receiver (through the use of filters), and using tunable transmitters to direct messages to the correct receiver.

Code division multiple access has also been used in [37] to create a fixed assignment network, as previously mentioned. Each node is assigned a code sequence from a group of orthogonal sequences. Transmitters use the code sequence for a given receiver when sending a logical "1" to that receiver. This approach also involves the use of a star coupler. Optical time division multiple access has also been used to create fixed assignment networks in much the same way as for conventional time division multiple access. Each node is assigned a periodically recurring slot and transmitters wishing to send information to a given node place information in the node's assigned slot.

Figure 6.6 shows the internal configuration of a node employed in a fixed assignment ring network. This node has been arbitrarily designated as node five in an eight node network. All information received on channel five is destined for node five, and is removed. Information on the other channels is simply regenerated by node five and sent on around the ring. If node five wishes to send information to another node, it monitors the channel for that node until it is free. Protocols are needed to prevent "hogging" of a channel.

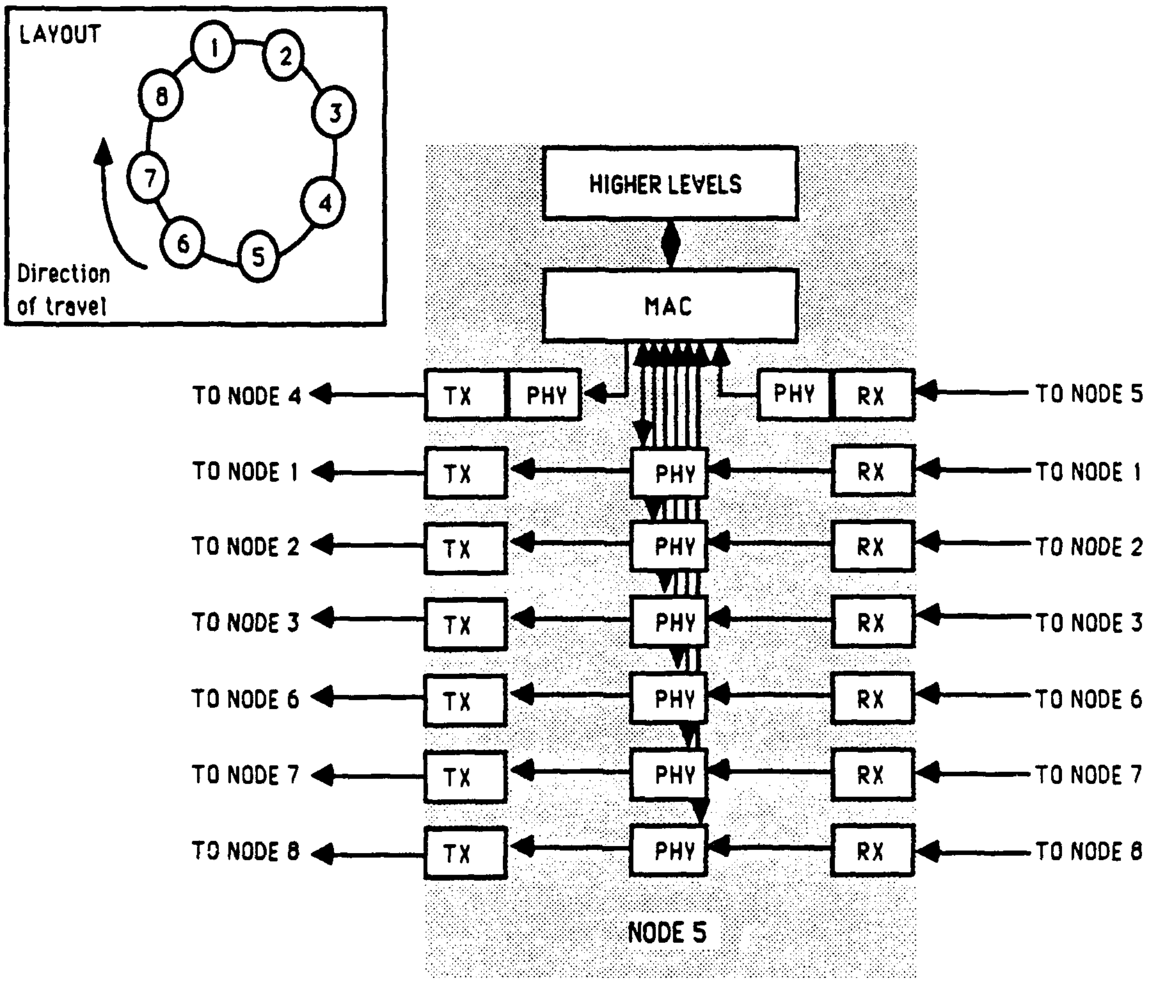

TX - TRANSMITTER   RX - RECEIVER   PHY - PHYSICAL LAYER FUNCTIONS   MAC - MEDIA ACCESS CONTROL

Figure 6.6: Configuration of Node Five in a Fixed Assignment Ring Network

When a fixed assignment network is implemented on the primary ring, the configuration resembles a web of unidirectional buses, each node being the terminating point of its own bus, or channel. For a given node, the initial point of the bus is the node directly downstream. The bus flows around the ring until it reaches the terminating node. Obviously, those nodes close to the origination of a particular node's bus will have a better chance of finding the bus idle than those near the termination point. Protocols would have to be implemented to achieve fairness.

The secondary ring would be used to implement another set of unidirectional buses flowing in the opposite direction from those on the primary ring. In the event of a link or node failure, buses from both rings would be employed to construct the network.

6.2.5 Star Network

A logical star network can be implemented on a physical ring. One node on the ring can serve as the hub point with the other nodes logically connected to it. To maintain the survivability aspects of the ring, each node must be able to serve as the hub point. Therefore a high level of complexity must reside in each node. Another problem is that links near the distribution point can become very congested relative to links opposite the distribution point. Since all links must have the sufficient capacity to allow any node to become the hub point, efficiency is low. Furthermore, nodes farther away from the distribution point have greater difficulty transferring information due to the longer propagation delays to and from the distribution point, so protocols would have to be implemented to assure fairness.

6.2.6 The Multi-Channel Bus Network

As shown in Figure 6.7, a dual bus can be implemented on a physical counter-rotating ring as has been done in IEEE 802.6. Both primary and secondary rings are peers in this

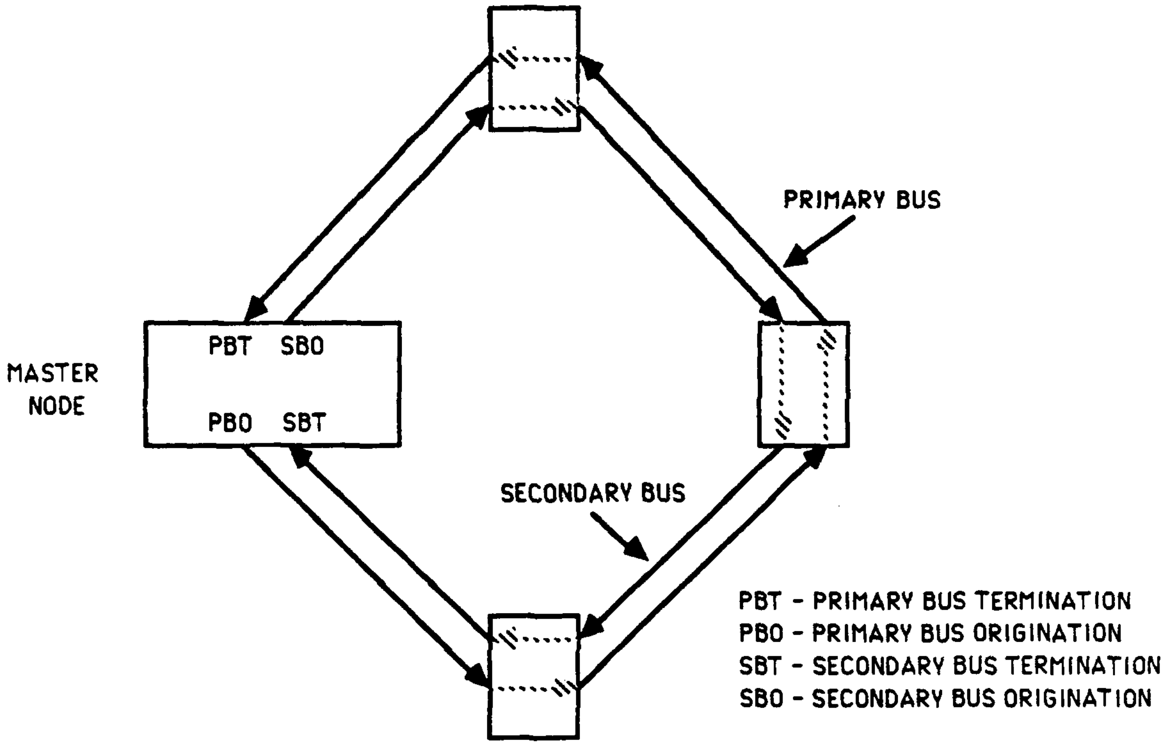

Figure 6.7: Dual Bus Implemented on Physical Counter-Rotating Ring

standard. One node serves as the bus origin and termination point, and, in the event of a node or link failure, this task is transferrable to any other node. The dual bus concept can be expanded to a multi-channel environment as for the multi-channel ring. Multi-channel versions of the Fasnet [64] and Expressnet [65] bus architectures have been proposed and analyzed by Camarda, Castagnola, and Leaci in [66].

The multi-channel bus could allow a cell to use any logical link to travel from node to node. As for the ring, this distinction would make the network a single multi-channel network rather that a set of single channel networks. The network could also be implemented as a single channel byte-wide dual bus similar to the byte-wide single channel ring previously discussed.

Fairness can be approached if a small portion of the bandwidth available to the secondary bus is used to hold reservations for positions on the primary bus, and vice versa. Routing information must be present to allow a node to know to which ring it should send a cell headed for a given destination.

6.3 Network Evaluation

As long as the logical network overlayed upon the physical counter-rotating ring does not contain a single point of failure, such as a central clock, a non-transferrable distribution point in a logical star, or a non-transferrable bus origination point in a logical bus, then all networks discussed here will continue operating and will avoid partition in the event of one failure of any nature. Therefore they can be said to be equally survivable. They are not, however, equally reliable, efficient, complex, cost effective, fair, growable, or secure. The networks will be evaluated according to these criteria in this section.

As for the multiplexing techniques, any attempt to evaluate the merits of the above discussed logical networks relative to the listed criteria will be somewhat arbitrary, due of the strong dependence on application specific issues such as geographic span and bit rate. However, the networks can be compared in a rough sense, as has been done in Table 6.1. On a scale of one to ten, characteristics for each network are rated relative to those for the other proposed networks. Ten is awarded to the "best" performer in a given category. For example, a ten is awarded to the single channel byte-wide network in the complexity category because it is the least complex. In networks where fairness can be achieved through added complexity (such as in the fixed assignment network), the additional complexity negatively affects the fairness rating of the network and not the complexity rating. For the perfect shuffle network, values are assigned assuming the use of the baseline configuration, i.e. without the proposed addition of a "logical channel" field suggested on page 79.

Table 6.1: Comparison of Six Logical Topologies on a Physical Ring Network
(Survivability of each topology assumed to be equal)

CRITERION \ NETWORK		MULTI-CHANNEL RING	SINGLE CHANNEL RING	PERFECT SHUFFLE	RECEIVER FIXED ASSIGN.	STAR	MULTI-CHANNEL BUS
RELIABILITY		10	1	4	4	3	10
EFFICIENCY		10	10	3	1	1	10
COMPLEXITY		4	10	6	7	5	3
FAIRNESS	NORMAL	5	5	5	1	1	1
	FAILURE	1	1	1	1	1	1
GROWTH POTENTIAL		10	3	6	2	2	8
SECURITY		10	7	7	1	1	10
COST		6	10	9	7	5	6
TOTAL		56	46	40	24	19	49

Reliability - The networks discussed will be implemented with components, such as laser transmitters, that fail on occasion. While any single failure will not bring down the whole network in any case, it can have varying degrees of impact. With this criteria we assess how the network would be affected by the failure of a laser transmitter.

The multi-channel ring and bus networks will be only negligibly affected by a single transmitter failure. One channel, out of several connecting two nodes, would be eliminated, but cells could be transmitted on the physical link over any of the remaining channels. In contrast, the single-channel byte-wide network would be severely affected by a transmitter failure in that the whole physical link between two nodes is eliminated. The ring would have to be wrapped and half the available bandwidth would be lost. For the perfect shuffle network, the loss of a transmitter means that one of two logical connections at a node is eliminated and either more hops will have to be taken by cells or the perfect shuffle network on the alternate ring will have to be used. The other networks are also more adversely affected by transmitter failures than the multi-channel networks.

Efficiency - Because all networks discussed here use ATM techniques, their efficiency is going to be higher than if other techniques, such as conventional time division multiplexing, were used. However, they are not equally efficient in other ways. Efficiency is here defined as the ability of the network to employ unused transmission resources in the event of congestion in other transmission resources. The multi-channel networks which permit sharing of all logical links on a physical link are the most efficient. The single channel byte-wide is as efficient as the multi-channel networks if each have comparable transmission resources. The other approaches are inefficient, although the perfect shuffle network is more efficient than the fixed assignment or star networks.

Complexity - The single channel byte-wide approach is the least complex network due to the elimination of the parallel-to-serial conversions and the fact that it has only one channel to manage. The receiver fixed assignment protocol is also reasonably simple if fairness is not considered. The star network has high complexity in that each node must be able to serve as a distribution point. The multi-channel ring is complex but less so than the multi-channel bus under normal operation, because no node needs to serve as bus arbitrator and cell originator.

Fairness - In Table 6.1, the fairness criteria is divided into two columns, one measuring fairness under normal operating conditions and one measuring fairness under failure conditions. The maximum rating for each column is five, giving a maximum possible rating for the fairness criteria of ten. The ring topologies and the perfect shuffle topology are naturally fair under normal conditions, while the other configurations require complexity to achieve fairness. In a wrap mode, all networks considered require increased complexity to achieve fairness.

Growth Potential - This criteria measures the ease with which additional transmission resources or nodes can be added to the network, and the impact of this addition on other nodes. Ideally, the network should be able to accommodate the addition of a reasonably large number of nodes to the network without necessitating hardware or software changes at existing nodes. The worst impact an entering node should have on the operation of the network is to possibly cause a disruption in currently propagating (low priority) information. Within reason, there should be no "hard limit" set on the number of nodes that can be supported, or the amount of capacity that can be added.

For the multi-channel ring, both additional capacity, in the form of additional logical links, and additional nodes can be easily added to the network with little impact. For the multi-channel bus, additional capacity can be easily added but the addition of nodes can require fine tuning of bus reservation protocols already in place. Many of these bus protocols allow a given node to reserve a portion of the capacity dependent on the number of nodes downstream from the given node. With the perfect shuffle network, both nodes and capacity can be added, although care must be taken to preserve the logical pattern. With the star and the fixed assignment networks the addition of a node requires the addition of capacity throughout the whole network, severely impacting all other nodes. In networks such as LambdaNet™, the number of nodes on the network is limited by the number of discrete wavelengths that can be supported by the network, because there must be one separate wavelength for each station.

The single-channel byte-wide network permits the smooth insertion of nodes, but the insertion of capacity has a major impact on every link and every node throughout the network. The upper bit rate can also be limited by bit skew in the byte-wide network if dispersion-flattened fiber is not used.

Security - Data security is also important to the users of wide area networks. For reasons including data security, network growability, and the variability of traffic due to the many sources, a shared allocation of network transmission resources (or bandwidth) is generally preferable to a system where allocation is fixed. In a fixed allocation system, if an eavesdropper can identify the transmission resources used by a particular network user, he can then determine how much traffic that user is creating. (The contents are still likely to remain a secret due to the probable use of encryption.)

Although the assignment of values in Table 6.1 is inexact, the large differences in the total scores for the network suggest that conclusions drawn from them are still valid. Assuming that the criteria used to evaluate the networks are all given equal weight, we can conclude that the multi-channel ring network is superior to the other networks for use as an ATM wide area network overlaying a physical ring. This is especially due to its high reliability, which is often second only to survivability as a necessary attribute of wide area networks, as well as its high efficiency, security, and growth potential. The network achieves high ratings in these three areas primarily because of the division of the physical links connecting nodes into independent serial channels which can all be accessed by a given node.

6.4 Hierarchical Architectures

In section 6.2, a number of topologies for connecting peer nodes at the highest level of a wide area network were discussed. In this section we discuss architectures where nodes may not be of the same level. This would occur, for example, in a situation where central offices serving various local military installations or government complexes were to be interconnected, and also connected to a gateway to the national network at the top of the hierarchy.

The public switched telephone network has traditionally been arranged hierarchically with five classes of switching centers [67]. Offices within each class are located in what can best be described as an irregular mesh topology across the country, with the office density becoming sparser as we move from Class 5 to Class 1. The mesh topology is extremely survivable due to its multi-connected nature.

With the greater introduction of single-mode fiber as the interconnection media between central offices, the topology of the local, or intra-LATA (local access transport area) network is fundamentally changing from the previously widespread mesh architecture to a "hub" architecture [68,69]. The hub topology improves the economic efficiency of the system by increasing the utilization of the single mode fiber transmission facilities. Instead of distributing information directly between central offices at DS1 rates, information would be multiplexed at each central office in the LATA and sent at SONET OC-1 or OC-3 rates to a central "facility" hub, as shown in Figure 6.8. The hub, which would

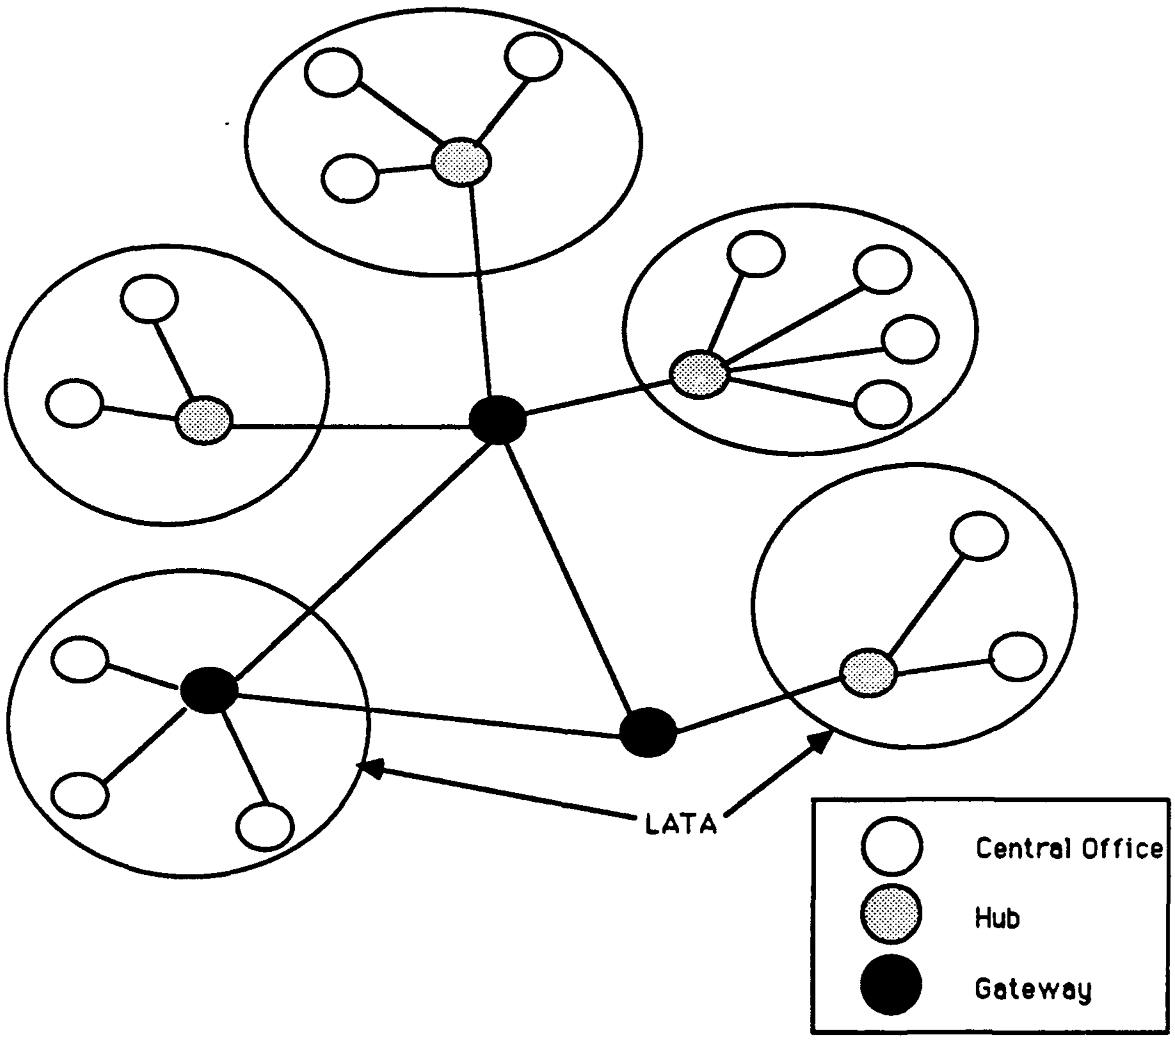

Figure 6.8: Basic Hubbing Topology (with Gateways Connected by Ring)

have to switch at the higher rate, could also serve as the input source for various broadcast video channels. Hubs would be connected to gateways on the predescribed ring network for the transfer of inter-LATA information [70].

The hub architecture can be made survivable through the use of a combination of 1:N protection, 1:N diverse protection, dual homing, and self-healing rings [7]. These are summarized as follows:

1:N protection: For every N fibers along an office-to-hub span, one extra is added to serve as a replacement in the event of a failure to one of the active fiber links. Since in almost all cases fibers connecting two central offices would reside in the same cable, this technique would not prevent against failures due to cut cables.

1:N diverse protection: The replacement fiber is physically separated from the active fibers, resulting in survivability under cable failure conditions. If the protection is 1:1 diverse (N=1) then the ability to survive single failures in the transmission system between central office and hub is 100%.

Dual homing: A central office is connected to two hubs instead of one, resulting in increased survivability in the event of failure or destruction of the cross connect system, or any other major system, in the hub. The "home" hub refers to the hub to which the central offices in a local switching area are connected. In dual homing, fiber spans from central offices connect not only to the "home" hub, but to a "foreign" hub, as well. Traffic is distributed between both routes. In the event of a hub failure, critical circuits are rearranged so that they can be connected through the remaining hub.

Self-healing rings: A physical ring similar to that used to connect gateways can also be used in the hierarchical portion of the network. Transmission pathways are available in both directions around the ring, and offices on the ring can have the capability to transmit and receive in both directions and to perform ring wraps. The self-healing ring approach can save fiber at no loss in survivability over a 1:1 diverse protection technique [7]. Dual homing with self-healing rings can be achieved if both the home and foreign hubs are included on the ring. The self-healing ring approach can be applied not only in connecting central offices to hubs, but in connecting hubs to gateways.

It is likely that a major application of this intermediate hierarchical network will be to connect local area networks and metropolitan area networks that comply with the IEEE 802.6 or American National Standards Institute (ANSI) X3T9.5 standards. Transparent connections between the local networks will be easier to achieve if this connecting network is as compatible as possible with these emerging standards in computer communications.

The ANSI X3T9.5 standard, better known as the Fiber Distributed Data Interface, or FDDI, is a 100 Mbps token-passing network employed on a counter-rotating ring topology using optical fiber. If the multimode fiber version of the specification is used, then the internode spacing limit is set at two kilometers, with 500 nodes allowed on the network, and a total ring circumference of 200 km. Frames, or packets, in FDDI can be of variable length with the maximum limitation of 4500 bytes. When a node has control of the token, it may transmit more than one frame, limited by a token-holding time determined through negotiation at network initialization. Frames completely traverse the ring and are removed by the source so that correct delivery of the information is assured. The network employs automatic reconfiguration, with a provision for optical bypass switches at nodes to allow them to be bypassed in the event of node failure or removal. A single-mode version of the standard allows for a maximum spacing of 60 kilometers [71]. An alteration to conventional FDDI, known as FDDI-II, allows for synchronous, isochronous, and asynchronous traffic [72]. In both versions the secondary ring is used only in the event of a failure that requires reconfiguration, and for ring initialization.

The IEEE 802.6 Metropolitan Area Network standardization effort, referred to as the Distributed Queue Dual Bus (DQDB), is a slotted dual bidirectional bus physically configured as a counter-rotating ring. Due to the subdivision of (part of) its bandwidth into unassigned cells and its use of the SONET OC-3 physical layer, 802.6 has the best chance of being compliant with ATM. Two main priority levels are specified, with the lower priority level being further divided into four other priority levels. Frames, or slots, are generated at eight KHz intervals on each of the buses, which operate at 155.5 Mbps, equal to the SONET OC-3 rate. Nodes can place reservations for slots on one bus by setting request bits in slots traveling on the other bus. If the master, or slot generating, functions can be handled by any station, and fault recovery capability can be distributed throughout the network, there would be no single points of failure in the network. [73]

7. Fiber Optic Specifications for the Multi-Channel ATM Ring

As we shall show in the next section, the physical layer of the OSI model, as it applies to an ATM network, consists of two sublayers, the ATM sublayer and the Physical Layer Media Dependent (PMD) sublayer. The PMD sits below the ATM layer, and provides the actual transmission capabilities. The goal of this section is to address issues leading to the creation of a target specification, implementable within the next ten years, for the PMD sublayer of the multi-channel ATM ring network. Given this implementation timeframe, we will restrict our proposed PMD to that which has been experimentally, rather than just theoretically, proven.

Based on the conclusions from Section 5, a reasonable approach is to use wavelength division multiplexing. For the sake of illustration, we will assume that each link is composed of eight wavelength multiplexed channels, each operating at the SONET OC-48 rate of 2.49 Gbps.

The choice of 1550 nm as the operating wavelength region for a long-haul network is easily made due to a number of developments, some of which were discussed in Section 4. The lowest attenuation possible on the optical fiber occurs in this region, even though dispersion at 1550 nm in conventional fiber is around 18 ps/nm-km. Reductions in laser linewidths due to the development of distributed feedback lasers have made Gbps transmission at this wavelength possible. Further chirp-reducing improvements due to separation of biasing and modulation regions in the lasers (see Figure 7.1) have enabled direct modulation laser systems to operate at multi-gigabit speeds. A distributed feedback

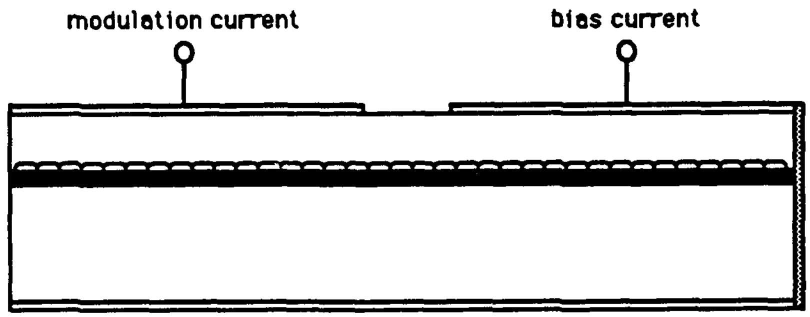

Figure 7.1: Monolithic Tunable Distributed Feedback Laser

laser has recently been directly modulated at eight Gbps [74]. Indirect modulation techniques using Lithium Niobate directional couplers have been used to externally modulate a continuously operating laser at a 20 GHz bandwidth [75]. Figure 7.2 [76] shows transmission distances as a function of bit rate over conventional fiber at 1550 nm that can be obtained using lasers having differing linewidths. The dark line at the top of the

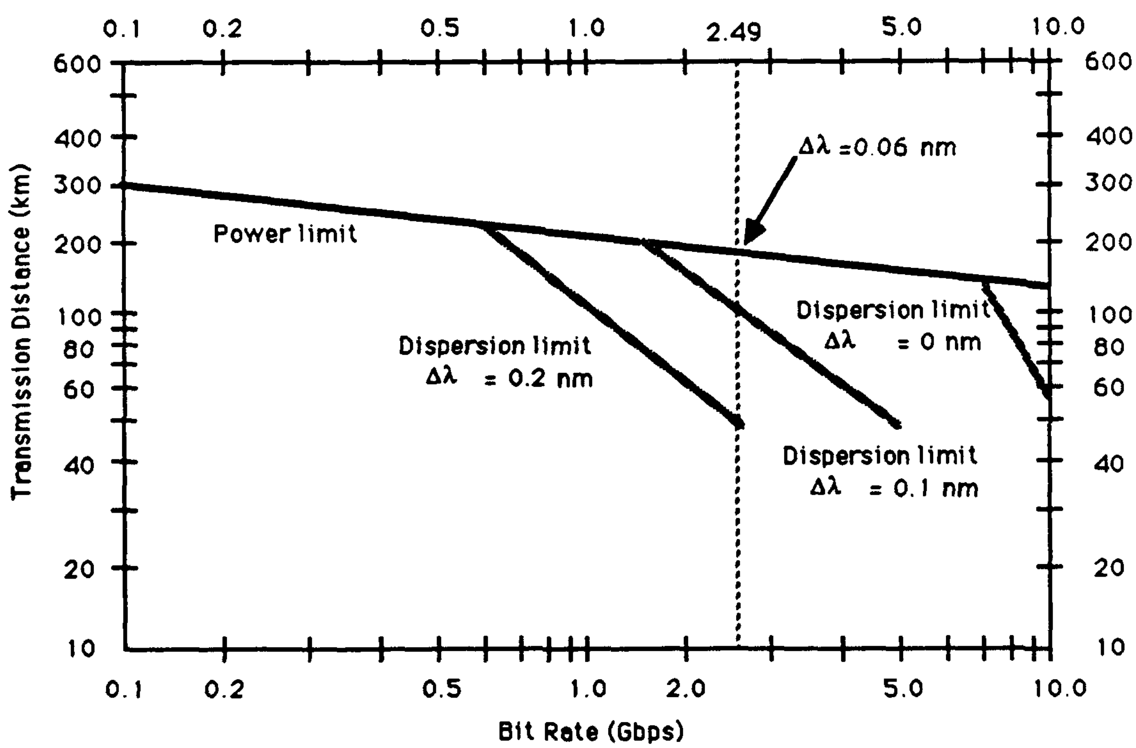

Figure 7.2: Transmission Distance versus Bit Rate for 1550 nm Laser Systems [76]

figure represents the power limit, assuming the use of state-of-the-art APDs, 0.2 dB/km fiber, and an input optical power of 1 mW. This limit would be reduced by more than 50% if the 1300 nm wavelength region were used. The more lightly shaded family of lines represent the dispersion limits for lasers with various linewidths. The fact that there is a dispersion limit for the zero linewidth case is due to the bandwidth of the information itself.

At our chosen bit rate of 2.49 Gbps, where state-of-the-art APD sensitivity is about -36 dBm, Figure 7.2 indicates that an unrepeatered and unamplified system would be power-limited at about 180 km. If the system is dispersion-limited when the pulse spreading equals one-half the bit time, then, assuming dispersion is 18 ps/nm-km, we will remain power-limited for laser linewidths up to about 0.06 nm, or 7.7 GHz. A reduction in linewidth below 0.06 nm will not increase the transmission distance of the system. Since monolithic tunable distributed feedback lasers have been made with linewidths well under 100 MHz [77], the laser linewidth requirement is not stringent. The transmission distance could be increased to 230 km if the laser power were increased to 10 mW, which is still below the level where significant nonlinear effects begin, and the laser linewidth were reduced to below 6 GHz. However, low linewidth semiconductor lasers are increasingly difficult to fabricate as output power requirements are increased.

Using dispersion-shifted fiber with index profiles as shown in Figure 4.13, the zero dispersion wavelength has been shifted to around 1550 nm with an experimentally measured attenuation of 0.17 dB/km [22]. Dispersion-flattened fibers, with dispersion over the wavelength range of 1300 and 1500 nm limited to within ±2.5 ps/nm-km and attenuation of around 0.2 dB/km, are entering production [78]. Their widespread installation into the national fiber network would assure the full-scale use of 1550 nm in new systems. The laser linewidth requirement could be significantly relaxed, permitting Fabry-Perot lasers to be used in 2.49 Gbps systems.

The next step in specifying the link will be to determine the characteristics of the wavelength division multiplexing operation. It is assumed that the system, in order to be implementable within the next ten years, will start out as a direct detection system, and may evolve into a coherent frequency division multiplexed system as that technology progresses. Two additional assumptions, that each channel of the link is to be relatively equivalent in performance, and all channels are to be simultaneously receivable, are also made. A "dense" wavelength multiplexed system, where all wavelengths used are from the 1550 nm window, will therefore be employed. The detectors used could then be identical, since the responsivity of the InGaAs APDs is relatively flat over a wide range in the 1550 nm region. The detectors could be incorporated onto a single array, with the possible inclusion of amplification circuitry. Detector arrays with eight to 16 channels have been built on a common substrate [79].

In [63] Kobrinski, et al. report the development of an eighteen channel wavelength multiple access network (known as LambdaNet™ and previously discussed in Section 6.2.4) with wavelengths ranging from 1527 to 1561 nm, spaced by 2 nm. The demultiplexing device employs a diffraction grating with a single-mode fiber used at the input, and eighteen multimode fibers used at the output. Channel widths were 0.85 nm, with an insertion loss of less than 3.5 dB and a crosstalk of less than -27 dB reported. Grating-based devices can be employed in many ways, including those shown in Figure 7.3 [80]. Figure 7.3a shows a conventional Littrow type of grating-based wavelength demultiplexer, while Figure 7.3b shows a GRIN (GRaded INdex)-rod lens Littrow type of demultiplexer. Littrow type wavelength demultiplexers are characterized by their use of a single lens. Another type of demultiplexer, where no lens is used, is represented by Figure 7.3c. This type is known as the slab-waveguide type and is characterized by a concave-shaped grating and a slab waveguide. Also in this family of demultiplexers are those employing chirped gratings to separate wavelengths [80].

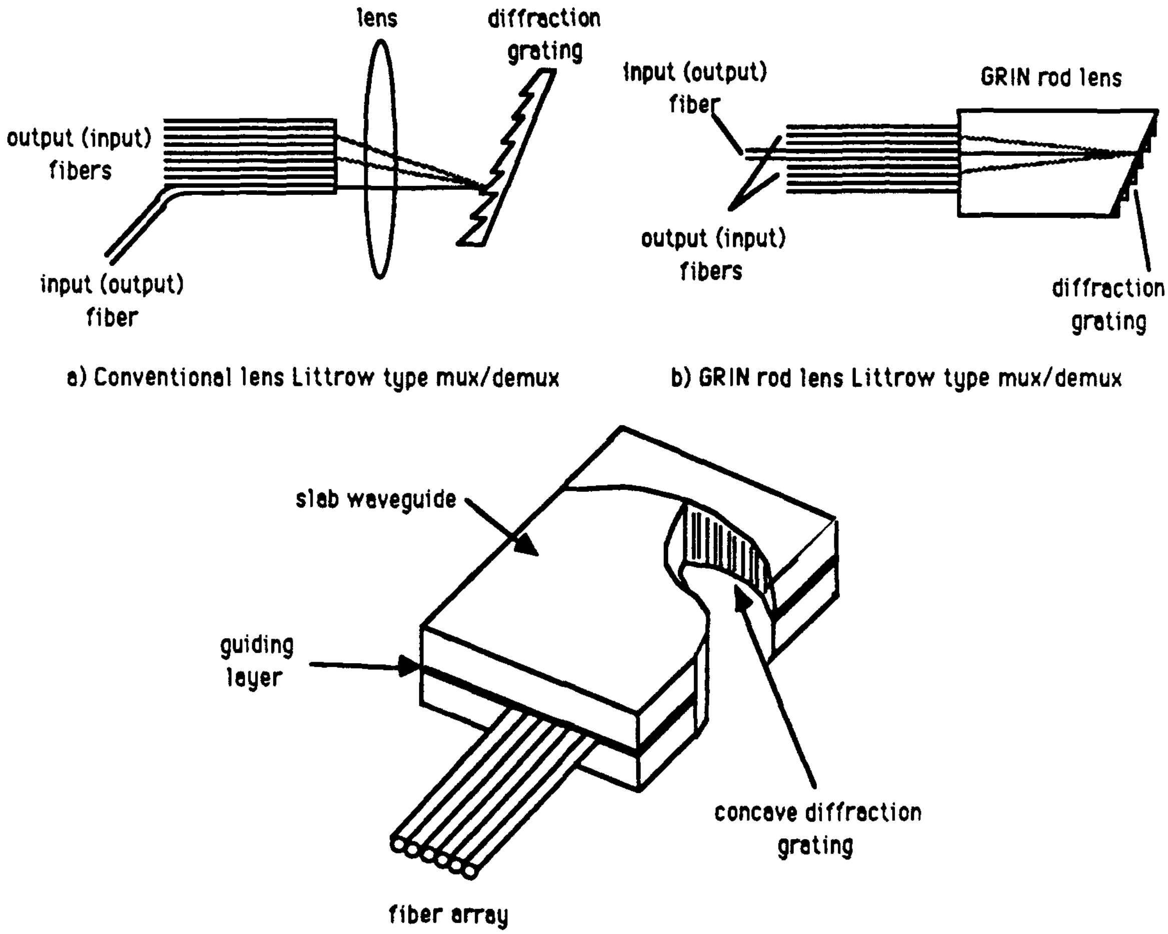

a) Conventional lens Littrow type mux/demux

b) GRIN rod lens Littrow type mux/demux

c) concave grating-slab waveguide type mux/demux

Figure 7.3: Diffraction Grating-based Wavelength Multiplexers/Demultiplexers [80]

In order to maintain adequate channel spacing between lasers under operating conditions, and to keep the lasers within the narrow wavelength channel defined for them, some sort of regulation is required. In [81] the center wavelength of distributed feedback lasers is altered by temperature effects and then stabilized in the correct channel using temperature control to within ±0.01°C. Tunable lasers could be controlled through feedback information which could then be used to alter the bias current of the device. The monolithic tunable laser shown in Figure 7.1 has a continuous tuning range of 2.1 nm [77]. Other monolithic tunable lasers, using the distributed Bragg reflector (DBR) structure, have achieved continuous tuning ranges of 4 nm [77]. Manufacturing tolerances on the center wavelength of distributed feedback lasers can currently be controlled to within a 5 nm window [82].

The multiplexing technique at the transmitter end of the link can be a simple star combiner, or it can be one of the grating devices previously discussed, operating in a reverse fashion (i.e. the output fibers to the receivers become input fibers from lasers, and the input fiber from the transmitter becomes the output fiber to the receiver). In this case the multiplexer would have to employ single-mode fiber on both sides of the grating, unlike the device used in the LambdaNet™ experiments. The reciprocal, or symmetric, nature of the all single-mode wavelength multiplexer/demultiplexer also allows for important possibilities for bidirectional transmission on the same fiber.

In [83] an all single-mode grating device, similar to that shown in Figure 7.3a, is described. This device has an insertion loss of about 6 dB and a crosstalk of less than -34 dB. However, only five channels, spaced by about five nm, where incorporated in the design. In [82], Payne and Hill report the development of a single-mode to single-mode device of similar construction having 20 channels spaced by about 3.6 nm, in the range from 1484 to 1554 nm. Channel widths were 0.3 nm, with insertion loss reported to be typically less than 3 dB. Payne and Hill state that the use of phase gratings produced in volume holograms are most capable of allowing for reduced channel spacing to channel width ratios.

The use of wavelength multiplexer/demultiplexers at each end of the link will likely reduce the available optical power in the system by about 6 dB, reducing the power limited transmission distance from 180 km to 150 km. The linewidth requirement of the lasers can therefore be relaxed to 0.074 nm, or 9.26 GHz.

Traveling wave optical amplifiers can be used to increase the power limited distance of the link because these amplifiers have sufficient optical bandwidth to pass all the channels in the wavelength multiplexed system. Jopson and Darcie report in [84] that as long as channels are spaced by at least one or two GHz, then intermodulation distortion effects from other channels will be 20 dB below the signal level of a given channel, assuming channels have equal power. This effect would not be a problem at 2.49 Gbps operation, because channel spacing would have to be wider than one GHz to accommodate the information bandwidth.

Important parameters of the optical amplifier are not only its wavelength range, but its unsaturated gain, its saturated gain, and the optical power level at which gain begins to fall into the saturated region. Koga and Matsumoto [15] report that, for a semiconductor traveling wave amplifier with facet reflectivity reduced to 0.02%, an unsaturated gain of 19 dB has been attained. This value includes eight dB loss associated with coupling between optical fibers and the semiconductor. This gain decreases as input optical power increases until it is reduced by 3 dB (to 16 dB gain) when input optical power reaches -12 dBm for single channel operation. Gain saturation eventually limits the number of channels that can be supported by the amplifier at about 100 [85].

The level at which saturation begins to occur is reduced as more channels are added to the link, because changes in gain as a function of the number of channels in the link will lead to crosstalk between channels. In order to keep the change in the channel's signal-to-noise ratio below one dB, crosstalk, or the change in gain of one channel as a function of the use of the other channels, must be less than -10 dB [15]. If there are eight channels simultaneously operating, then this requirement means that the maximum input optical power into the amplifier is further reduced by about 19 dB, to -31 dBm, according to calculations in [15]. There is about a twelve dB range of input powers (i.e from -31 dBm to about -43 dBm) over which the amplifier will produce a useful gain, enabling freedom in the placement of the amplifier in the span.

Optical amplifiers can be cascaded in a link to increase its distance substantially. Each amplifier would contribute spontaneously emitted light to the signal, however, effectively reducing the sensitivity of the receiver. The filtering effect of the wavelength division demultiplexers lessens the effects of the spontaneously emitted noise due to its filtering of the signal. In [86], Brain reports that four semiconductor traveling wave amplifiers were cascaded to achieve an additional gain of 73 dB in a coherent system, with a receiver sensitivity degradation of 6.5 dB due to spontaneous emission from the amplifier. The effective increase in power of 66.5 dB would allow for an additional transmission distance of 332.5 km (assuming 0.2 dB/km attenuation), for a total transmission distance of 482.5 km counting the initial power available in the system. Of course, if the laser linewidth is not reduced to below 0.023 nm, or 2.88 GHz, the link would be dispersion-limited before this point, since the optical amplifiers do not regenerate the signal. The limit on the number of amplifiers that can be cascaded is determined at the point where gain saturation due to the amplified spontaneously emitted light occurs [85]. Dietrich, et al, estimate that about twenty semiconductor amplifiers of the type typically available at this time could be cascaded, with more possible if spontaneous emissions were filtered at intermediate points. Clearly, the possibility for multi-channel links spanning 2000 km without electronic regeneration exists in the near future.

The availability of symmetrical, or reciprocal, wavelength multiplexer/demultiplexers , along with wide bandwidth, "single pass" traveling wave amplifiers have created the potential for bidirectional multi-channel transmission between nodes over the same fiber, resulting in large savings in long-haul systems [87]. For example, the eighteen channel grating-based multiplexer/demultiplexer previously described could be used to support nine channels in each direction. The need for optical isolators, to prevent interchannel crosstalk from disturbing laser stability, positioned between the multiplexer/ demultiplexer and the laser would increase, although isolators may already be present to prevent the negative effects of reflections. The increase in the number of channels passing through the optical amplifier would naturally reduce the maximum input power level for each channel, so that gain saturation and its crosstalk effect can be avoided. In moving from a unidirectional eight channel link to a bidirectional sixteen channel link (eight channels per direction) the maximum input power for each channel would be reduced by about 4 dB [15], reducing the dynamic range to about eight dB. To avoid system degradation when multiple optical amplifiers are cascaded, reflections at facets would have to be suppressed even more than is required for unidirectional operation, perhaps to levels on the order of 10^{-5} for twenty stages [85].

The effect of the differences in propagation rates of the different channels in a wavelength multiplexed multi-channel network must also be considered. As an example, we assume an eight channel network operating in the 1550 nm region with channels spaced by 2 nm. If dispersion is 18 ps/nm-km, then signals on the "fastest" channel will arrive 252 ps/km ahead of signals on the "slowest" channel. An ATM cell is 53 octets, or 424 bits, in duration, which, at a transmission rate of 2.49 Gbps, is 170.3 ns. Therefore, a cell transmitted on the fastest channel immediately after the completion of transmission of a cell on the slowest channel would catch the first cell after about 675 km. If channels were operated asynchronously relative to each other, this skew could accumulate through the network. Several possible solutions to this problem exist, including the resynchronization of channels prior to the 675 km limit, the use of protocols which send cells from the same message over the same wavelength, or the use of protocols that can reorder cells at the destination. Of course, tighter channel spacing would allow a longer distance before cells became misordered, although, as discussed, this could cause other problems to arise.

To summarize, our recommendations for the specification of a PMD sublayer that can be implemented within the next ten years include the use of wavelength division multiplexing in the 1550 nm region to create eight 2.49 Gbps channels spaced by about 2 nm. A symmetrical multiplexer/demultiplexer using a diffraction grating, and having sixteen channels, is recommended so that bidirectional transmission can be implemented over the same fiber. Using conventional fiber having an attenuation of 0.2 dB/km, widely available distributed feedback lasers having linewidths under modulation of less than 0.06 nm, and state-of-the-art APDs, an unrepeatered and unamplified link span of 150 km can be attained. Lasers can be stabilized in a channel either by tuning the bias current or through temperature control.

Traveling wave optical amplifiers, initially implemented with semiconductor material and later with rare earth ion-doped fiber, could be cascaded to significantly stretch the power limitations of the system, allowing for link spans of 2000 km without electronic regeneration. These amplifiers have wide bandwidths permitting the simultaneous amplification of as many as 100 channels, provided that input power levels are low enough to prevent gain saturation and its associated crosstalk. Back reflections at the facets of the amplifiers would have to be almost completely suppressed, and narrow laser linewidth lasers (not outside the capability of current technology) or dispersion-shifted fiber would be required to prevent the link from being dispersion-limited. As coherent technology matures, direct detection channels can be gradually replaced with coherent detection channels. Inline traveling wave amplifiers can continue to be used in the coherent system.

8. Higher Layer Protocols for the Multi-Channel ATM Ring

If information is to be conveyed between distinct entities, there must be rules to govern the format of the information. For example, in speech, the frequency, duration, and volume of sounds produced by a person must form certain patterns in order to be recognized as words by another human. These words must be in a language common to the listener. Words must be grouped together according to conventions or rules in order to form sentences, or complete thoughts. On a higher level, sentences must be organized to allow the conveyance of complex ideas. The need for rules, or protocols, for the transfer of information is universal. This section will discuss protocols for an ATM ring network at levels above the Physical Media Dependent sublayer, discussed in the previous section.

As the example of human speech shows, there are many levels of protocol that must be present in order for complex entities to communicate. To provide a framework for the layering of network protocols, two layered protocol stacks have been widely promoted; the U. S. Department of Defense (DoD) four-layer protocol stack, developed initially to control ARPANET; and the International Standards Organization's (ISO) Open Systems Interconnect (OSI) seven-layer model. A communications network consisting of completely homogeneous entities would be relatively easy to effectively connect. However, most networks include products from different companies or from different lines within the same company, creating the need for detailed specification and standardization of protocols. Since companies do not collaborate completely on these specifications, a compromise has been reached in which standards have been developed which fit into one of the two frameworks. These standards define the functions required in the layer and the type of information shared between layers. The set of rules governing the transfer of information at layer boundaries is called the interface [88].

An illustration of how the current Department of Defense protocol stack, based on the Transmission Control Protocol/Internet Protocol (TCP/IP), is organized relative to the OSI layers is shown in Figure 8.1. The Transmission Control Protocol/Internet Protocol (TCP/IP) has been in wide use since its inception in 1973. The function of the Internet Protocol is to divide messages from the TCP into small packets with appropriate headers for transmission in a packet switched network like the Defense Data Network. It also reassembles the packets at the destination. IP cannot guarantee error free delivery, so TCP

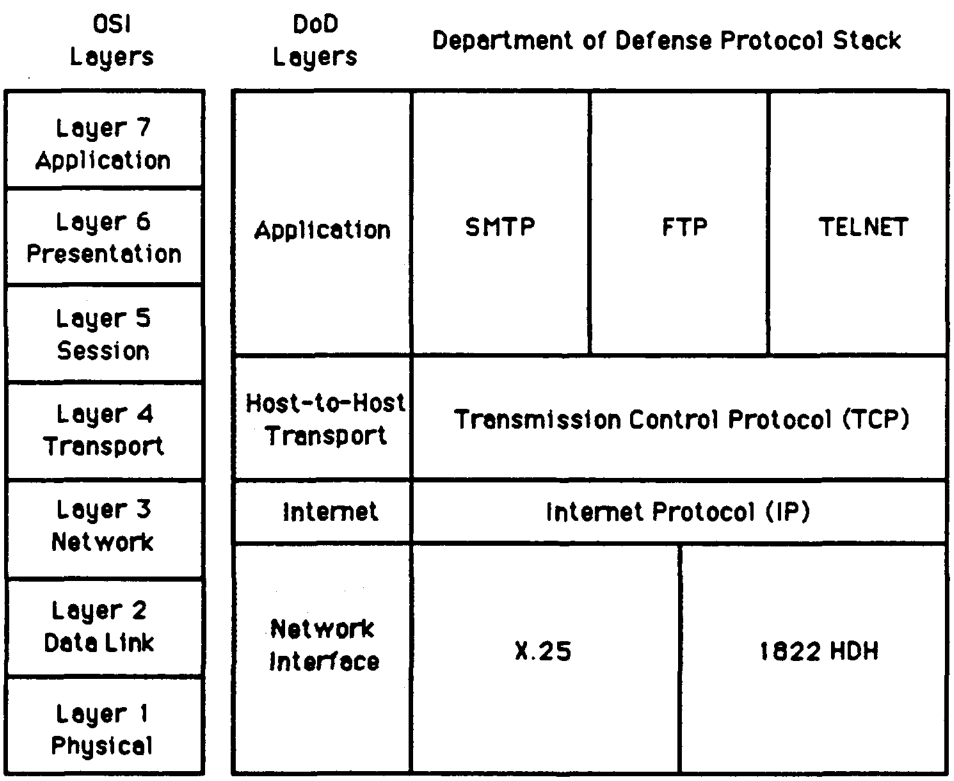

Figure 8.1: Comparison of the Department of Defense protocol stack to the OSI layers

performs error detection and correction functions with the whole message. TCP uses acknowledgments, packet numbering, and checksums to guard against packet loss, packet errors, and misordering of packets [89]. A large portion of the current OSI protocol was based on knowledge gained from the use of the TCP/IP-based stack.

Although the Department of Defense protocols are well established, and reasons, such as a large installed base, exist to use them in place of the OSI model, the U.S. government has adopted a policy for transferring to the OSI model. The specifications for government conversion to use of a pure OSI stack are presented in the Government Open Systems Interconnect Profile (GOSIP) policy. Proponents of the shift state that adhering to the ISO standards will result in lower cost, and greater network availability, due to greater availability of replacement parts and service, and will provide a higher possibility that government communication equipment will be compatible with that of international allies. The application layer standards within OSI model are also expected to allow a greater level of functionality when compared with the current application layer protocols. The OSI application layer currently has standards under development in remote file access, synchronous terminal support, bit-mapped graphics, multi-media electronic mail, and real-time process control. The OSI protocol stack which is called for in GOSIP is shown in Figure 8.2 [90].

Within the two protocol stacks, the Simple Mail Transfer Protocol (SMTP) and the X.400 provide electronic mail, the File Transfer Protocol (FTP) and the File Transfer Access and Management (FTAM) provide file transfer functions, and the TELNET and Virtual Terminal

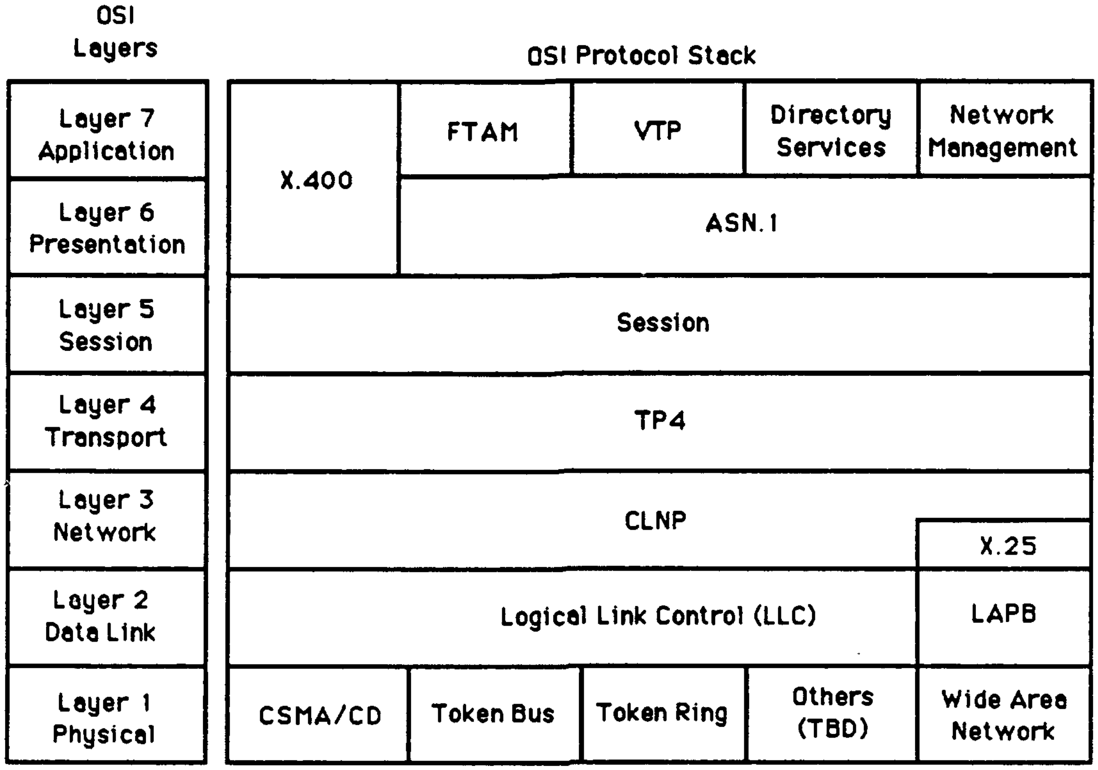

Figure 8.2: OSI Protocol Stack [90]

Protocol (VTP) provide remote login capabilities. As in the Department of Defense stack, the X.25 protocol extends down through the physical layer, and is used for the transmission of conventional, variable length packet information over a private or public wide area network. Much of X.25 would be replaced in an ATM network. The LAPB (link access procedure-balanced) is for data link control in networks exchanging X.25 packets on a peer-to-peer basis, where each entity takes equal responsibility for the success of the communication. The LLC data link control standard (IEEE 802.2) is intended for local area network communication. A basic difference between it and the LAPB standard (standardized by CCITT) is that the LLC simply discards packets received in error, while the LAPB protocol requests a retransmission of the packet.

As a first step in the movement to pure OSI, the Defense Communications Agency is developing dual protocol hosts that can run on either the DoD or the OSI stacks and can serve as application layer gateways between the stacks [90]. A new OSI-compatible transport protocol, called TP-4 and standardized as ISO 8073, has been developed to replace TCP in the protocol stack. The two protocols do not differ significantly, except in the checksum algorithm, according to experiments referred to in [89]. Both TCP and TP-4 are end-to-end connection-oriented protocols, in that a "connection" is established between to communicating entities through the transmission of connection-request and acknowledgment packets. The connection is maintained and data is transferred until connection-termination and acknowledgment packets are sent. End user resources are dedicated to the connection as long as it is established.

The tasks of each of the layers in the seven layer OSI model are as follows [91]:

Application - performs tasks relevant to a particular user application, such as file transfer and electronic mail.

Presentation - performs data compression (of voice or video signal, for example), code conversion, and data encryption. Data encryption is performed after compression because encrypted data would be difficult to correctly compress.

Session - assists with the matching of users to the correct computing resources or services to which they have access rights. This process is called the establishment of a session. The session layer also tabulates use of the network for accounting purposes.

Transport - responsible for the segmentation and reassembly of messages into units of information that are to be transmitted discretely over the network. It would be responsible for the reordering of ATM cells from the same voice message. High-speed sessions might be broken up into several easier-to-handle sessions by the transport layer, or low rate sessions might be multiplexed together into a single session. In many cases the transport layer is also the location of gateway functions linking networks that are incompatible below the transport layer. Transport layer gateways would be required to link a packet-based local area network like FDDI or IEEE 802.3 to an ATM wide area network, for example.

The transport protocol can be either connection-oriented, connectionless, or both. In connectionless service, also known as datagram service, single packets are sent without prior knowledge that the destination can or will accept them, and no acknowledgment is expected. ISO is standardizing a connectionless protocol known as ISO 8602 [92].

Network - responsible for the implementation of routing and flow control in the network. This layer would decide on which of two counter-rotating rings to place a cell, for example. Many network control cells would be generated at the network layer. The network layer generally provides for the buffering of cells entering or leaving the network. Network bridges could be used to link networks having compatible network layers, but incompatible lower layers, such as IEEE 802.3 (Ethernet) and IEEE 802.5 (token ring). The network layer protocol can also be connectionless or connection-oriented.

Data Link Control - The purpose of this layer is to convert the unreliable communications link established by the Physical layer to a reliable link through the introduction of error checks, flags, or framing structures. In some cases cell reordering would occur in the data link control. A major function of the data link layer in multi-access applications is to provide media access control, so that anarchy does not exist on the network media. Common media access protocols include Carrier Sense Multiple Access with Collision Detection (CSMA/CD), token-passing, register insertion, and time division multiple access. Later in this section we will discuss the register insertion protocol as a promising protocol for use in the multi-channel ATM ring network, due to its low complexity and high level of fairness.

Physical - This layer is often referred to as a virtual bit pipe, because it is responsible for the transfer of bits between nodes. In many cases this layer is divided into two sublayers, the Physical Media Dependent (PMD) sublayer and the physical protocol sublayer. PMD deals with signal transmission and mechanical issues pertinent to the media, as was discussed in the previous section. Many issues of physical layer protocol have already been introduced in Sections 5 and 6. The physical protocol would accomplish coding and synchronization, and is responsible for the formation of ATM cells through the addition of the header information. The division of the ring into slots of ATM cell length is a physical protocol decision, as is the assignment of header fields according to ATM standards. The use of plesiochronous clocking and the generation of synchronization cells, needed to periodically allow for node elasticity buffers to reinitialize themselves, are also issues of physical protocol.

The ISDN and ATM standards are being developed in conjunction with international standards bodies including CCITT (Comite Consultatif International Telegraphique et Telephonique) and ISO, and will be compliant with the OSI model. The ATM standard fits into the OSI model as shown in Figure 8.3. The ATM layer sits at the top of the physical layer,

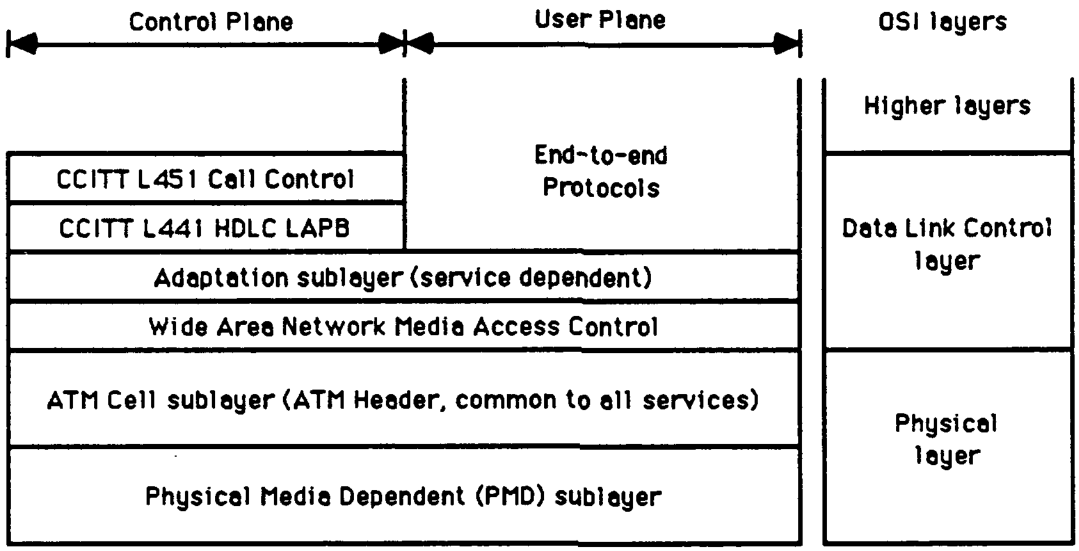

Figure 8.3: ATM within the OSI Model

above the Physical Media Dependent (PMD) sublayer, and is responsible for the cell header information. The adaptation sublayer is responsible for the conveyance of service dependent information not needed by the general switching and multiplexing structure of the network. This information would include service type, payload interior channel information, payload error checks, acknowledgment requirements, the number of bytes present in an incompletely filled cell, or cell sequencing information.

As previously mentioned, ATM cells are 53 octets in duration, with five octets being used for the header. Depending on whether the cell is passing through the User-Network Interface or through the Network-Node Interface, there are two formats for the header of the ATM cell, as is shown in Figure 8.4 [93]. The generic flow control field is used at the

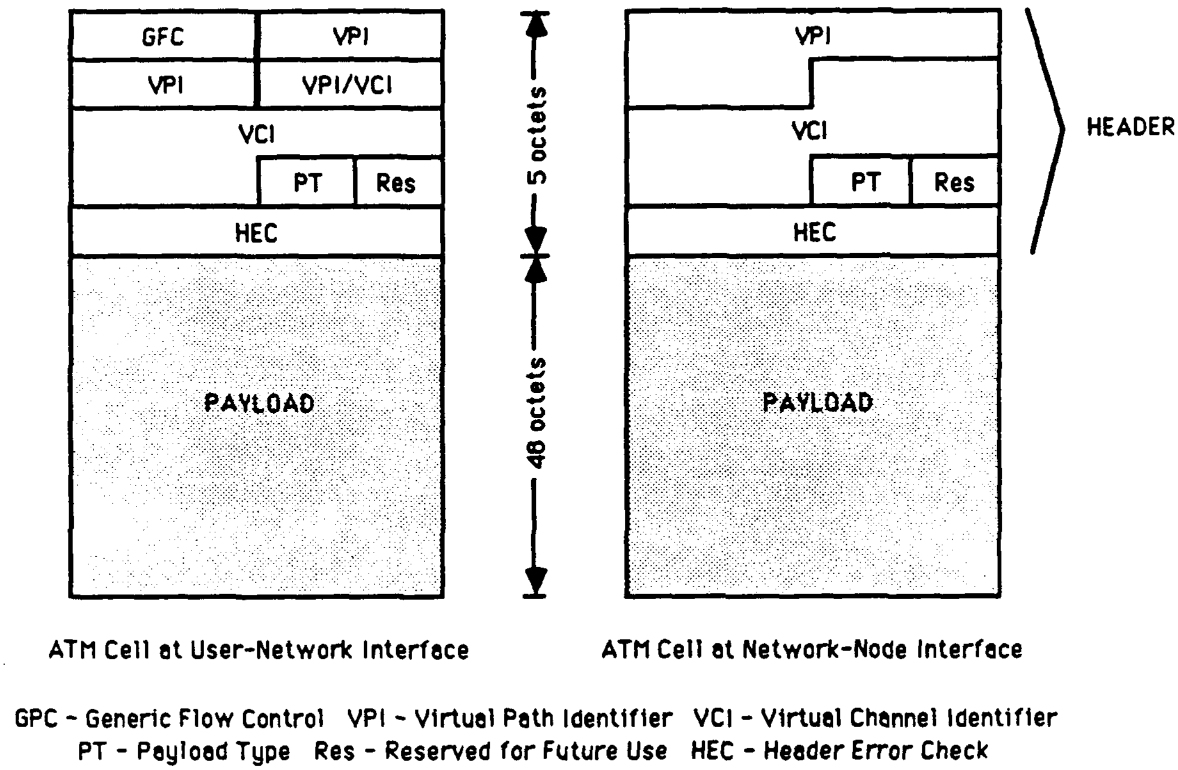

Figure 8.4: ATM Cell Formats

customer premises to control the flow of traffic of different service types at the premises. This field would only be present in cells at the user-network interface. The Virtual Channel Identifier (VCI) and the Virtual Path Identifier (VPI) fields are used, in combination with call set-up information (if necessaary), for packet addressing and routing. The header error check is an eight-bit cyclic redundancy check, capable of correcting one error in the cell header, and then switching to a mode where it has increased error detection capabilities. The type of error checks performed on the payload depend on the type of information in the payload and are done at higher layers.

One two-bit field of the cell header has been reserved for future study. For the sensitive network applications addressed in this report it is recommended that that field, combined with the Payload Type field, be used to prioritize packets. Priority levels could be used, for example, to handle confidential, secret, and top secret traffic in different ways in the network, to indicate the urgency of the transmission, or to organize traffic into groups requiring different levels of network performance. For example, "real-time" traffic, like voice and video, could be given higher priority in routing, without requiring an acknowledgment of correct reception to the source from the destination. Data traffic, like file transfers, which must be transmitted with a high degree of reliability but does not require timeliness, could have a lower priority. The priority of a cell could be changed as it travels through the network depending on its "age," or time spent in the network. Different types of packets could be "aged" at different rates, depending on the degree of timeliness required.

The choice of the Media Access Control protocol is very important to network performance. For the multi-channel ATM ring network the use of the simple, mature register insertion protocol has many benefits. Register insertion protocols increase the efficiency of the transmission media by allowing any node on the network to take advantage of unused capacity on the transmission link connecting it to another node. The protocol can best be understood through Figure 8.5. Each node has one or more queues for holding cells traveling to other nodes in the network, and for cells entering the network at that node. The node transmits cells from the queue on outgoing transmission links. A wide range of queueing disciplines are conceivable, including first-in-first-out (FIFO) queues, priority queues, and even more complicated queuing protocols involving the age of the cell. Cells already on

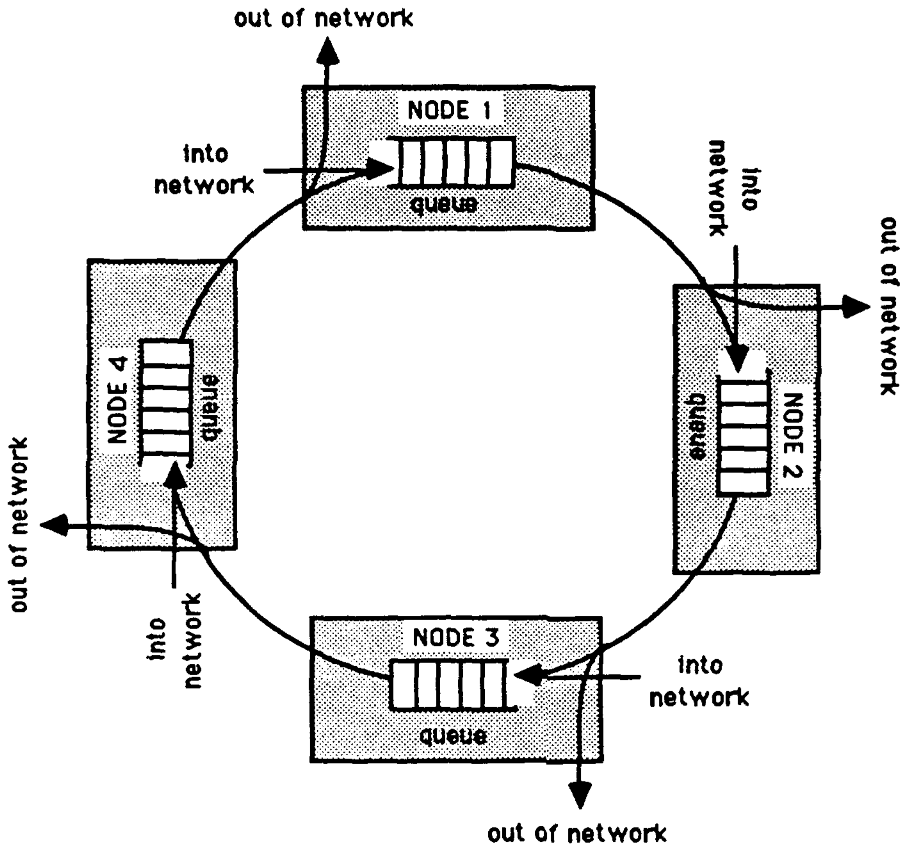

Figure 8.5: Register Insertion Protocol on a Four Node Ring

the network and "just passing through" a node would either circumvent the queue of that node, reducing buffer size requirements, or be required to join the queue as for cells just entering the network. In the first case the protocol must include a mechanism to prevent "hogging" of the network.

The use of a queueing policy where all cells entering the node enter the queue, and where primary ordering is determined by cell priority, and secondary ordering by cell age, would be inherently fair. This fairness would apply even in ring reconfiguration situations, where nodes implementing the wrap have limited ring access, such as were discussed in Section 6. This queueing scheme is used in the simulation project discussed in Appendix B. There, an example of the difference in performance between this priority scheme and a FIFO ordering scheme is shown. An obvious disadvantage of this choice is the requirement for greater buffer space and more complexity in buffer management. The ATM cell header would also have to contain a field that in some fashion represented the cell age.

Another data link layer protocol that holds promise for use in wide area networks is based on the use of the optical fiber for storage of cells to reduce buffer requirements at nodes. If the outgoing channels of a node were all busy, then incoming cells would be redirected back out on links going in the opposite direction. Cells would circulate on the fiber between nodes in this fashion either until capacity opened up on the outgoing link, or until the cell achieved a high enough priority to postpone the transmission of another cell. The priority mechanism would have to be incorporated to insure that cells would not perpetually circulate between two heavily congested nodes, and to make sure that cells eventually continued in the correct direction around the ring. This concept is derived from BlazeNet [94], which employs fiber loops, instead of simple point-to-point links, to connect nodes. BlazeNet nodes are assumed to be more generally configured than in the counter-rotating ring topology assumed here.

Also extremely important to the performance of the network is the speed at which the higher layer protocols can be operated. Recently, the high data rate capability and high reliability of fiber optic links have resulted in a communications bottleneck in the higher layers, due to their general unwieldiness and the constraints of software. For this reason there have been efforts to develop fast "protocol engines" on very large scale integrated (VLSI) circuits. One proposal, developed by Protocol Engines Incorporated, and known as the Express Transfer Protocol (XTP) [95], has been widely supported by groups involved in the standardization of military local area networks [96]. This protocol encompasses the network and transport layers, often grouped together and referred to as the Transfer layer. XTP's major gain in speed is due to the fact that it performs many of the connection set-up and termination functions implicitly, rather than explicitly as is done in TP-4. A comparison of the information exchanged between two nodes to accomplish the transmission of one data packet in a connection-oriented environment is shown for the two protocols in Figure 8.6. The clear gain in efficiency, and transfer speed, can be seen for the case of single packet transmission. Since many network resources are dedicated to a connection as long as it is established, the quicker that connections can be set up, used, and torn down, the quicker network resources can be allocated to other connections. The efficiency of TP-4 improves with multiple packet transfers, but, for tactical military local area networks, where single packet communication is common, the XTP can greatly improve performance [96]. The implicit connection setup aspects of XTP can also be useful in a wide area ATM network, especially if real-time information makes up a large portion of the traffic.

Other higher level protocol issues, including the use of the secondary ring as a peer to the primary ring, have been previously mentioned. Routing decisions, an OSI network layer function, could be made using either dynamic or static algorithms. Algorithms requiring only local information are more stable and are likely to be adequate for the simple routing requirements of a counter-rotating ring. Flow control, or the decision as to whether or not to permit cells onto the network, should be incorporated in a distributed fashion among the nodes. Also, flow control should operate using an algorithm requiring only local, rather than global information. The priority level of a cell, perhaps alterable with cell age, would be an important element in the flow control algorithm.

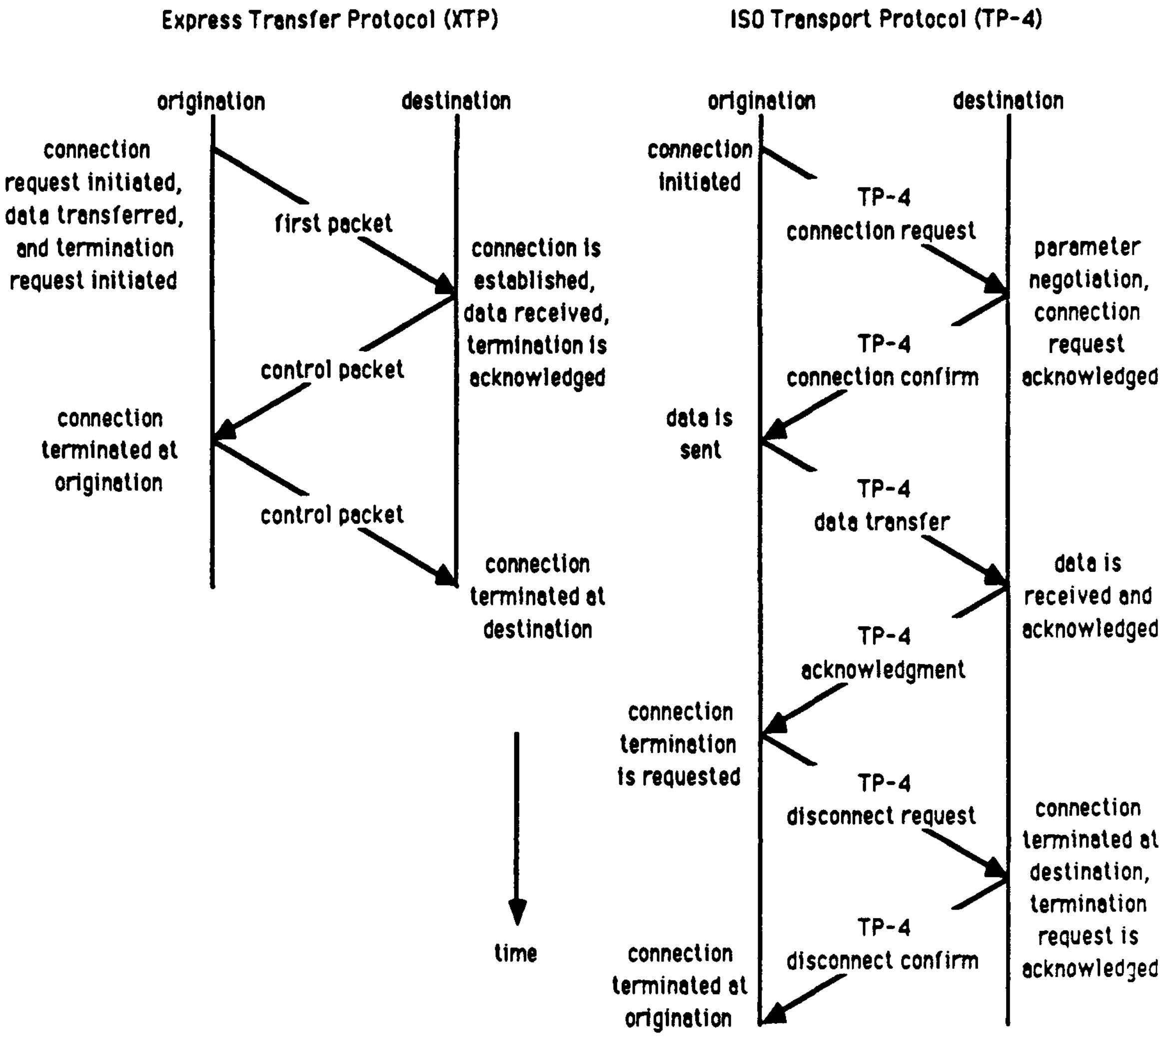

Figure 8.6: Comparison of XTP and TP-4 for Single Packet Transmission [96]

9. Conclusions and Recommendations for Future Work

This report presents a wide-ranging investigation into the physical layer technology of a survivable optical fiber-based wide area network intended for use as a high capacity (> ten Gbps) communications backbone. Issues addressed include the characteristics of the single-mode fiber media, the effects of component parameters on link performance, possible multiplexing techniques to achieve high capacity, the physical and logical topology of the network, and network protocols.

Three important conclusions are reached early in the report. The most appropriate physical topology of the network is determined to be a counter-rotating ring, because, without application-specific knowledge of the exact survivability-versus-cost tradeoffs, the ring is the topology most efficiently offering a generally acceptable level of survivability. Secondly, consistent with commercial standardization efforts moving toward providing a common fast-packet switched environment for multi-service traffic, Asynchronous Transfer Mode (ATM) is chosen as the baseline for operation of the network physical layer. Thirdly, conventional electronic time division multiplexing will not be able to keep pace with the capacity requirements of future networks, so we assume that links connecting nodes will be implemented as multiple channels in parallel.

A major contribution of this report is the extrapolation of the logical topologies of several conventional network architectures - including the ring, the perfect shuffle, the fixed assignment, the star, and the bus - into this new multi-channel ATM ring-based wide area environment. Once these logical topologies are applied to the new environment, the resulting architectures are compared relative to their reliability, growth potential, cost, complexity, fairness, and security. The multi-channel ATM ring configuration was concluded to be superior, largely due to its high reliability and growth potential. Using similar criteria, the various multiplexing techniques for realization of the multi-channel links, including space, code, optical time, and wavelength (or frequency) division multiplexing, were also compared. Wavelength division multiplexing was determined to hold the most promise due to its smooth growth potential and relatively low complexity.

The implementation of wavelength division multiplexing on a multi-channel ATM ring wide area network was then studied in greater detail, with an eye toward the implementation of a practical system within ten years. It was assumed that each channel in the multi-channel links between nodes would operate at the Synchronous Optical Network (SONET) OC-48 rate of 2.49 Gbps, and that there would initially be eight channels. Due to the narrow laser linewidths of distributed feedback lasers, we conclude that the low attenuation region of the fiber at 1550 nm will serve as the transmission window. Within this window, direct detection dense wavelength division multiplexing systems, where channels are spaced by two to five nanometers, will soon become practical. A smooth transition to the subsequent generation, frequency multiplexed coherent systems, will likely be possible soon thereafter. The symmetric nature of the grating-based wavelength demultiplexers, along with the wide bandwidth of the "single pass" traveling wave optical amplifiers, present a great opportunity for the development of multi-wavelength bidirectional links within the next ten years. Traveling wave amplifiers also will allow for the creation of links spanning over 2000 kilometers without electronic regeneration, matching the distance performance of optical solitons with less complexity.

There are, of course, many areas in the field of high speed wide area networking that present opportunities for further research. Only a few, which are directly related to this report, will be mentioned here. These range from component to protocol issues and include the following:

1) The register insertion protocol is promising for use as the media access control protocol for the multi-channel ATM ring network. The questions of what queueing discipline to use, whether or not cells on the network passing through a node should enter the queue, what buffer sizes are needed for various throughputs and grades of service, and how to "age" the cells on the network so that fairness can be maintained, remain to be answered.

2) The statistics of the combined traffic in a multi-service network, necessary network sizing and performance evaluation, need to be better understood.

3) The amount of "processing power" required to process cell headers at the nodes in the multi-channel ATM ring network as a function of the relative positions of cells in the channels presents an interesting optimization problem.

4) Most of the questions in the development of transport layer gateways for converting between a conventional local area network, based on an IEEE 802 or FDDI standard, and an ATM-based wide area network, have yet to be decided.

5) New, low overhead, hardware-based transport protocols, such as the Express Transfer Protocol (XTP), have yet to be applied to the wide area ATM environment.

6) With wavelength division multiplexing, different wavelengths, and therefore different channels, will travel at different rates through the network. Therefore, the potential exists that cells originating in the same message could become misordered prior to

reaching their destination. The question of where in the protocol hierarchy this misordering should be handled should be addressed, as should the possibilities for resolving this problem in the physical layer.

7) Optical solitons, which retain their shape over long transmission distances due to interactions between dispersion and nonlinear effects, should be studied relative to their use in a multi-channel (i.e. multi-wavelength) environment.

8) Traveling wave amplifiers constructed from fibers doped with rare earth ions are advantageous due to the ease with which they can be coupled to the optical fiber system, and their low dependence on polarization. However, there are difficulties involved in pumping the amplifier with an external laser which must be addressed. An alternative area of research is in the development of semiconductor laser amplifiers that can be more easily coupled to the fiber, and that are independent of input light polarization. The development of coatings that suppress reflections down to 10^{-5} will also be necessary so that amplifiers can be cascaded to their full potential in a single fiber bidirectional system.

9) The development of 1550 nm lasers with chirping characteristics which, when combined with the dispersion characteristics of the fiber at 1550 nm, initially resulted in pulse compression rather than pulse spreading presents an opportunity in device physics.

10. List of References

[1] B. Schultz, "The Evolution of ARPANET," *Datamation*, vol. 34, no. 15, pp. 71-74, Aug. 1, 1988.

[2] W. Schatz and M. J. Foley, "Users Welcome New NSF Network with Glee and Caution," *Datamation*, vol. 34, no. 17, pp. 33-34, Sept. 1, 1988.

[3] G. E. Daddis and H. C. Torng, "A Taxonomy of Broadband Integrated Switching Architectures," *IEEE Communications Magazine*, vol. 27, no. 5, pp. 32-42, May 1989.

[4] M. Skov, "Implementation of Physical and Media Access Protocols for High-Speed Networks," *IEEE Communications Magazine*, vol. 27, no. 6, pp 45-53, June 1989.

[5] I. P. Kaminow, "Photonic Multiple-Access Networks: Topologies," *AT&T Technical Journal*, vol. 68, no. 2, pp.61-71, March/April 1989.

[6] D. Bertsekas and R. Gallager, *Data Networks*, Prentice Hall, Inc., Englewood Cliffs, NJ, 1987, p. 313.

[7] T.-H. Wu, D. J. Kolar, and R. H. Cardwell, "Survivable Network Architectures for Broad-Band Fiber Optic Networks: Model and Performance Comparison," *IEEE Journal of Lightwave Technology*, vol. 6, no. 11, pp. 1698-1709, Nov. 1988.

[8] F. Ross, "FDDI - A Tutorial," IEEE Communications Magazine, vol. 24, no. 5, pp. 10-18, 1986.

[9] C. D. Chaffee, *The Rewiring of America*, Academic Press, Boston, 1988, p.93.

[10] S.D. Personick, *Optical Fiber Transmission Systems*, Plenum Press, New York, 1981, p.18.

[11] ibid, p. 132.

[12] J. C. Palais, *Fiber Optic Communications*, Prentice Hall, Inc., Englewood Cliffs, NJ, 1984, p. 256.

[13] G. P. Agrawal, *Nonlinear Fiber Optics*, Academic Press, Inc., San Diego, 1989, p. 284.

[14] H. J. Westlake and M. J. O'Mahony, "Gain Characteristics of a 1.5 µm DCPBH InGaAsP Resonant Optical Amplifier," *Electronics Letters*, vol. 21, no 1, pp. 33-35.

[15] M. Koga and T. Matsumoto, "The Performance of a Traveling-Wave-Type Semiconductor Laser Amplifier as a Booster in Multiwavelength Simultaneous Amplification," *IEEE Journal of Lightwave Technology*, vol. LT-8, no. 1, pp. 105-112, 1990.

[16] M. J. O'Mahony, "Semiconductor Laser Optical Amplifiers for use in Future Fiber Systems," *IEEE Journal of Lightwave Technology*, vol. LT-6, no. 5, pp. 531-544, 1987.

[17] Personick, pp. 84-86.

[18] P. A. Kirkby, "Multichannel Wavelength-Switched Transmitters and Receivers-New Component Concepts for Broad-Band Networks and Distributed Switching Systems," *IEEE Journal of Lightwave Technology*, vol. LT-8, no. 2, pp. 202-211, 1990.

[19] J. Salz, "Coherent Lightwave Communications," *AT&T Technical Journal*, vol. 64, no. 10, pp. 2153-2210, Dec. 1985.

[20] J.R. Jones, "A Comparison of Lightwave, Microwave, and Coaxial Transmission Techniques," *IEEE Journal of Quantum Electronics*, vol. QE-18, no. 10, pp. 1524-1536, 1982.

[21] J. J. Refi and H.-T. Shang, "Determining the Zero Dispersion Wavelength and Dispersion Slope of Fibers Used for the Fiber Distributed Data Interface (FDDI)," *Proceedings of SPIE*, August 17-19, 1987, vol. 841, pp. 66-69.

[22] V. A. Bhagavatula, "Single-Mode Fiber Designs," *Optical Fiber Communications Conference Digest*, Tutorial Sessions," Jan. 25-28, 1988, New Orleans, pp. 125-154.

[23] R. H. Stolen, "Nonlinear Properties of Optical Fibers," in *Optical Fiber Telecommunications*, S. E. Miller and A. G. Chynoweth, editors, Academic Press, New York, 1979, pp. 125-150.

[24] ibid.

[25] Agrawal, p. 76.

[26] Y. R. Shen, *The Principles of Nonlinear Optics*, John Wiley and Sons, New York, 1984, p. 524.

[27] Agrawal, p. 133.

[28] L.F. Mollenauer and K. Smith, "Demonstration of Soliton Transmission over more than 4000 km in Fiber with Loss Periodically Compensated by Raman Gain," *Optics Letters*, vol. 13, no. 8, pp. 675-677, Aug. 1988.

[29] Agrawal, p. 132.

[30] Y. R. Shen, p. 520.

[31] W. I. Way, et al., "High Speed Circuit Technology for Multi-gigabit/sec Optical Communication Systems," *International Communications Conference '88*, pp. 313-318, 1988.

[32] J. L. Gimlett, et al., "Degradations in Practical High Speed Transmission Systems using Distributed Feedback Lasers," *International Communications Conference '88*, pp. 319-324, 1988.

[33] R. G. Swartz, "Electronics for High Bit Rate Systems," *Optical Fiber Communication Conference Digest*, Optical Society of America, Houston, p. 82, Feb. 6, 1989.

[34] D. Jacobson, "High Performance Multiple Channel Token Bus Protocol," *IEEE Pacific Rim Conference on Communications, Computers, and Signal Processing*, June 4-5, 1987.

[35] J. A. Salehi, "Code Division Multiple-Access Techniques in Optical Fiber Networks-Part I: Fundamental Principles," *IEEE Transactions on Communications*, vol. 37, no. 8, pp. 824-833, Aug. 1989.

[36] J. A. Salehi and C. A. Brackett, "Code Division Multiple-Access Techniques in Optical Fiber Networks-Part II: Systems Performance Analysis," *IEEE Transactions on Communications*, vol. 37, no. 8, pp. 834-842, Aug. 1989.

[37] P. R. Prucnal, M. A. Santoro, and S. K. Sehgal, "Ultrafast All-Optical Synchronous Multiple Access Fiber Networks," *IEEE Journal of Selected Areas in Communications*, vol. SAC-4, no. 9, pp. 1484-1493, Dec. 1986.

[38] R. S. Tucker, G. Eisenstein, and S. K. Korotky, "Optical Time-Division Multiplexing for Very High Bit-Rate Transmission," *IEEE Journal of Lightwave Technology*, vol. LT-6, no. 11, pp. 1737-1749, Nov. 1988.

[39] M. P. Vecchi, et al., "High-bit-rate measurements in the LAMBDANET wavelength optical star network," *Optical Fiber Communication Conference Digest*, Optical Society of America, New Orleans, p. 95, Jan. 27. 1988.

[40] A. Frenkel and C. Lin, "Tunable Etalon Filter for Optical Channel Selection in High-density Wavelength-division-multiplexing Systems," *Optical Fiber Communication Conference Digest*, Optical Society of America, New Orleans, p. 96, Jan. 27. 1988.

[41] I. W. Stanley, G. R. Hill, and D. W. Smith, "The Application of Coherent Optical Techniques to Wide-Band Networks," *IEEE Journal of Lightwave Technology*, vol. LT-5, no. 4, pp. 439-450, Apr. 1987.

[42] "The 1995 DCS - A Step to Vision 21," Defense Communications Agency white paper, April 1988.

[43] R. F. Rey, technical editor, *Engineering and Operations in the Bell System*; 2nd Edition, p. 91, 1983.

[44] E. S. Elsam, "The Defense Data Network Hits Its Stride," *Telecommunications*, vol. 20, no. 5, p. 121, May 1986.

[45] R. A. Mercer and W. L. Edwards, "Issues in the Migration of Military Communications to an ISDN," *IEEE MILCOM '86* conf. proc., pp. 50.5.1-50.5.5, 1986.

[46] ---, "A Users Guide to ISDN Standards," *Telecommunications*, vol. 22, no. 5, p. 34, May 1988.

[47] R. Haendel, "Evolution of ISDN Towards Broadband ISDN," *IEEE Network*, vol. 3, no. 1, pp. 7-13, Jan. 1989.

[48] B. Bingham and H. Bussey, "Reservation-based Contention Resolution Mechanism for Batcher-Banyan Packet Switches," *Electronics Letters*, vol. 24, no. 13, pp. 772-773, June 23, 1988.

[49] T. T. Lee, R. Boorstyn, and E. Arthurs, "The Architecture of a Multicast Broadband Packet Switch," *INFOCOM '88*, pp. 1-8, 1988.

[50] H. Suzuki, T. Takeuchi, and T. Yamaguchi, "Very High Speed and High Capacity Packet Switching for Broadband ISDN," *GLOBECOM '86*, pp. 24.3.1-24.3.6, 1986.

[51] J. S. Turner, "Design of a Broadcast Packet Network," *INFOCOM '86*, pp. 667-675, 1986.

[52] A. Huang and S. Knauer, "Starlite: A Wideband Digital Switch," *GLOBECOM '84*, pp. 5.3.1-5.3.5, 1984.

[53] K. Y. Eng, M. G. Hluchyj, and Y.-S. Yeh, "A Knockout Switch for Variable-Length Packets," *IEEE Journal of Selected Areas in Communications*, vol. SAC-5, no. 9, pp. 1426-1435, Dec. 1987.

[54] J. M. Senior, S. D. Cusworth, and A. Ryley, "Multichannel Access Strategies for WDM Optical Fibre LANs," *SPIE Fibre Optics '88*, vol. 949, pp. 19-27, 1988.

[55] D. Jacobson, "High Performance Multiple Channel Token Bus Protocol," *IEEE Pacific Rim Conference on Communications, Computers, and Signal Processing*, June 4-5, 1987.

[56] IEEE Computer Society, Draft of Proposed IEEE Standard 802.6 - Distributed Queue Dual Bus (DQDB) Metropolitan Area Network (MAN), Draft D. 0., June 1988.

[57] M. L. Loeb, "Parallel Data Transmission through an Optical Fiber using Wavelength Division Multiplexing," doctoral thesis, North Carolina State Univ., 1985.

[58] M. L. Loeb and G. R. Stilwell, Jr., "An Algorithm for Bit-Skew Correction in Byte-Wide WDM Optical Fiber Systems," *IEEE Journal of Lightwave Technology*, vol. LT-8, no. 2, pp. 239-242, Feb. 1990.

[59] M. J. Karol and R. D. Gitlin, "High Performance Optical Local and Metropolitan Area Networks: Enhancements of FDDI and IEEE 802.6 DQDB," *GLOBECOM '89*, pp. 28.3.1-28.3.8, 1989.

[60] A. S. Acampora, M. J. Karol, and M. G. Hluchyj, "Terabit Lightwave Networks: The Multihop Approach," *AT&T Technical Journal*, vol. 66, no. 6, pp. 21-34, Nov./Dec. 1987.

[61] J. R. Sauer, "Multi-Gb/s Optoelectronic Interconnection System," printed by Center for Optoelectronic Computing Systems, Univ. of Colorado at Boulder, 1989.

[62] J. W. Reedy, "The TDM Ring - A Fiber Optic Transport System for Campus or Metropolitan Networks," *IEEE Journal of Selected Areas in Communications*, vol. SAC-4, no. 9, pp. 1474-1483, Dec. 1986.

[63] H. Kobrinski, et al., "Demonstration of High Capacity in the LAMBDANET Architecture: A Multiwavelength Optical Network," *Electronics Letters*, vol. 23, no. 16, pp. 824-826, July 30, 1987.

[64] J. O. Limb and C. Flores, "Description of Fasnet - A Unidirectional Local-Area Communications Network," *Bell System Technical Journal*, vol. 61, no. 7, pp. 1413-1440, Sept. 1982.

[65] F. A. Tobagi, F. Borgonovo, and L. Fratta, "Expressnet: A High-Performance Integrated-Services Local Area Network," *IEEE Journal of Selected Areas in Communications*, vol. SAC-1, no. 5, pp. 898-912, Nov. 1983.

[66] P. Camarda, B. Castagnolo, and G. Leaci, "Optical Fiber Multichannel Local Area Networks," *International Communications Conference '88 Proc.*, pp. 46.7.1-16.7.5, 1988.

[67] R. F. Rey, technical editor, pp.104-110.

[68] S. S. Cheng and H. L. Lemberg, "Single-mode Optical Fiber Networks for Intra-LATA Communications," *International Communications Conference '85 Proc.*, pp. 1.1.1-1.1.4, 1985.

[69] L. T. Wu and N.-C. Huang, "Synchronous Wideband Network - An Interoffice Facility Hubbing Network," *1986 Zurich Seminar on Digital Switching*, pp. 33-39, 1986.

[70] T.-H. Wu and R. H. Cardwell, "Optimum Routing for Fiber Network Design: Model and Applications," *International Communications Conference '88 Proc.*, pp. 2.5.1-2.5.6, 1988.

[71] *FDDI Single-Mode Physical Layer Medium Dependent* (SMF-PMD), Draft Proposed American National Standard, X3T9.5/88 - 155, Rev. 1.3, Nov. 23, 1988.

[72] FDDI-II (ANSC X3T9.5) Working Paper, Feb. 1, 1987.

[73] R. M. Newman, et al., "The QPSX Man," *IEEE Communications Magazine*, vol. 26, no. 4, pp. 20-28, Apr. 1988.

[74] K. Kamite et al., "DFB Laser with Bandwidth Larger than 9 GHz," *10th IEEE International Semiconductor Laser Conference Proceedings*, Kanazawa, Japan, 1986, pp. 180-181.

[75] S. K. Korotky et al., "Optical Intensity Modulation to 40 GHz using a Waveguide Electro-Optic Switch," *Applied Physics Letters*, vol. 50, June 1987, pp. 1631-1633.

[76] T. Li and R. A. Linke, "Multigigabit-per-second Lightwave Systems Research for Long-Haul Applications," *IEEE Communications Magazine*, vol. 26, no. 4, pp. 29-35, Apr. 1988.

[77] I. Mito, "Recent Advances in Frequency Tunable DFB-DBR Lasers," *Optical Fiber Communication Conference Digest*, Optical Society of America, New Orleans, p. 176, Jan. 27. 1988.

[78] P. Creittner and H. Lydtin, "Dispersion-Modified Single-mode Fibers," *Optical Fiber Communication Conference Digest*, Optical Society of America, New Orleans, pp. 72-73, Jan. 27. 1988.

[79] Y. Kanabar and N. Baker, "Demonstration of a Novel Optical Multichannel Grating Demultiplexer Receiver (MGD) for HDWDM Systems," *Electronics Letters*, vol. 25, pp. 817-819, June 1989.

[80] H. Ishio, J. Minowa, and K. Nosu, "Review and Status of Wavelength-Division Multiplexing Technology and Its Applications," *IEEE Journal of Lightwave Technology*, vol. LT-2, no. 4, pp. 448-463, Aug. 1984.

[81] M.P. Vecchi et al., "High-bit-rate Measurements in the LambdaNet Multiwavelength Optical Star Network," *Optical Fiber Communication Conference Digest*, Optical Society of America, New Orleans, p. 95, Jan. 27. 1988..

[82] D. B. Payne and A. M. Hill, "Multiwavelength Technology for Local Loop Applications," *Globecom '89 Proceedings*, Nov. 27-30, Dallas, pp. 1.5.1-1.5.5.

[83] A. Lord and J. M. Boggis, "Novel Single-Mode, Grating WDM Device Based Network," *Electronics Letters*, vol. 24, no. 11, pp. 672-674, May 26, 1988.

[84] R. M. Jopson and T. E. Darcie, "Calculation of Multicarrier Intermodulation Distortion in Semiconductor Optical Amplifiers," *Electronics Letters*, vol. 24, no. 22, pp. 1372-1374.

[85] E. Dietrich et al., "Semiconductor Laser Optical Amplifiers for Multichannel Coherent Optical Transmission," *IEEE Journal of Lightwave Technology*, vol. LT-7, no. 12, pp. 1941-1955, Dec. 1989.

[86] M. C. Brain, "DPSK Coherent Field Demonstration System, and Use of Optical Amplifier Repeaters," *Globecom 1989 Proceedings*, Nov. 27-30, Dallas, pp. 10.4.1-10.4.6.

[87] M. Koga, T. Matsumoto, and J. Minowa, "Multiwavelength Bidirectional Repeater System using a Semiconductor Laser Amplifier," *CLEO '89 Tech. Digest*, 1989.

[88] L. Pouzin and H. Zimmermann, "A Tutorial on Protocols," *Proceedings of the IEEE*, vol. 66, no. 11, pp. 1346-1368, Nov. 1978.

[89] F. L. Maybaum, "Defense Data Network - An Overview," *IEEE MILCOM '86* conf. proc., pp. 15.1.1-15.1.7, 1986.

[90] P. Gross and G. Reichlen, "DoD Support for OSI Protocols," *IEEE MILCOM '88 Conf. Proc.*, pp. 343-347, 1988.

[91] Bertsekas and Gallagher, pp. 14-26.

[92] J. E. Israel and A. J. Weissberger, "Communicating Between Heterogeneous Networks," *Data Communications*, March 1987, pp. 215-235.

[93] CCITT Study Group XVIII, WP XVIII/8-Report of Meeting," Temporary Document No. 16 (Plenary), Geneva, Switzerland, June 1989.

[94] Z. Haas and D. R. Cheriton, "BLAZENET: A High Performance Wide-Area Packet-Switched Network using Optical Fibers," *IEEE Pacific Rim Conference on Communications, Computers, and Signal Processing*, June 4-5, 1987.

[95] G. Chesson et al., "Express Transfer Protocol Definition", Rev. 3.2, Protocol Engines Inc., June 1988.

[96] M. Cohn, "A High-Performance Transfer protocol for Real-Time Local Area Networks," *IEEE MILCOM '88 Conf. Proc.*, pp. 1041-1047, 1988.

[97] R. Ballart and Y.-C. Ching, "SONET: Now It's the Standard Optical Network," *IEEE Communications Magazine*, vol. 29, no. 3, pp. 8-15, Mar. 1989.

[98] S. Fleming, "To Know SONET, Know Your VTs," *Telephone Engineer & Management*, June 15, 1989, p. 62.

[99] R. S. Bloomfield, "Architectural Considerations for Accessing the Post-1995 CONUS DCS/DSN: Target MILDEP Local Information Transfer Networks," *IEEE MILCOM '88 Conf. Proc.*, pp. 361-368, 1988.

[100] CCITT Recommendation I.121, Geneva, June 1988 (TD 140).

[101] M. J. Rider, "Protocols for ATM Access Networks," *IEEE Network*, vol. 3, no. 1, pp. 17-22, Jan. 1989.

[102] W. R. Byrne, et al.,"Broadband ISDN Technology and Architecture," *IEEE Network*, vol. 3, no. 1, pp. 23-28, Jan. 1989.

[103] J. Gechter and P. O'Reilly, "Conceptual Issues for ATM," *IEEE Network*, vol. 3, no. 1, pp. 14-16, Jan. 1989.

[104] A. A. B. Pritscker, *Introduction to Simulation and SLAM-II*, 3rd Edition, John Wiley & Sons, New York, 1986.

[105] P. F. Drucker, "The Coming of the New Organization," *Harvard Business Review*, January-February, pp. 45-53, 1988.

[106] L. M. Applegate, J. I. Cash, Jr. and D. Q. Mills, "Information Technology and Tomorrow's Manager," *Harvard Business Review*, November-December, pp. 128-136, 1988.

11. Appendix A-Summary of Telecommunications Standardization Activities

Within the last five years telecommunications standards work has been done, both in international and national bodies, which will likely be of great significance to the development of future high-speed WANs, including those for military applications. Standards efforts of the most significance include ISDN, SONET, and B-ISDN. Much of the work, with the exception of SONET, originated in CCITT. Within the United States, the T1 committee, accredited by the American National Standards Institute (ANSI), began work on applying and expanding the CCITT efforts for U.S. use.

11.1 Synchronous Optical Network (SONET)

The Synchronous Optical Network (SONET) standards are being defined by the T1X1 Technical committee within T1. Standardization efforts originated in the United States and spread to the CCITT Study Group XVIII. SONET essentially specifies the physical and data link layer for high speed single mode fiber transmission between telecommunications network facilities, such as central offices and hubs. The SONET standards specify a family of rates, which are multiples of a basic rate at 51.840 Mbps. The basic rate was chosen so that video payloads could be handled, and so DS3 (44.736 Mbps) blocks could be transferred without much inefficiency. It was also chosen because of its ability to efficiently carry multiplexed channels at the American DS1 rate of 1.544 Mbps and at the European (CEPT) rate of 2.048 Mbps. Because of its compatible with both rates, SONET will serve as an international standard.

The SONET payload is organized into frames of duration 125 μs, containing 810 eight-bit bytes (resulting in a 51.840 Mbps stream). The basic framing block is called the Synchronous Transport Signal-Level 1 (STS-1) at the electrical interface (not standardized), and the Optical Carrier-Level 1 (OC-1) at the optical interface, which is standardized in SONET. This basic frame is organized into nine rows of ninety bytes each, with four rows carrying (section, line, and path layer) overhead and embedded operations channels (EOC), as shown in Figure 11.1 [97]. The portion of the STS-1 that is not reserved for overhead or for data communications channels is referred to as the synchronous payload envelope (SPE). STS-1 frames can be easily multiplexed together using byte interleaving to form standardized optical channels at rates currently specified up to 2.49 Gbps. Rates higher than the STS-1 rate are exact multiples of the STS-1 rate and are designated by their multiple; for example, 2.49 Gbps is designated as OC-48 [97]. The OC-3 and OC-12 interfaces are likely to be important in the near future as basic rates for the transport of broadband services to end users.

The SPE can be further divided in a number of fashions to accommodate lower bit rate signals, such as DS1, CEPT1, and DS2. For example, the SPE can be subdivided into a number of three-column sections each containing a DS1 stream. The standard specifies two modes, locked and floating, for positioning blocks of information such as DS1's within the SPE. Locked mode is generally used when network equipment is run from the same clock. The payload is located starting at a fixed position in the SPE. On the other hand, in floating mode blocks of information can start at any point in the SPE, and are located through the use of a virtual tributary (VT) pointer [98]. This allows SONET to operate in a plesiochronous environment.

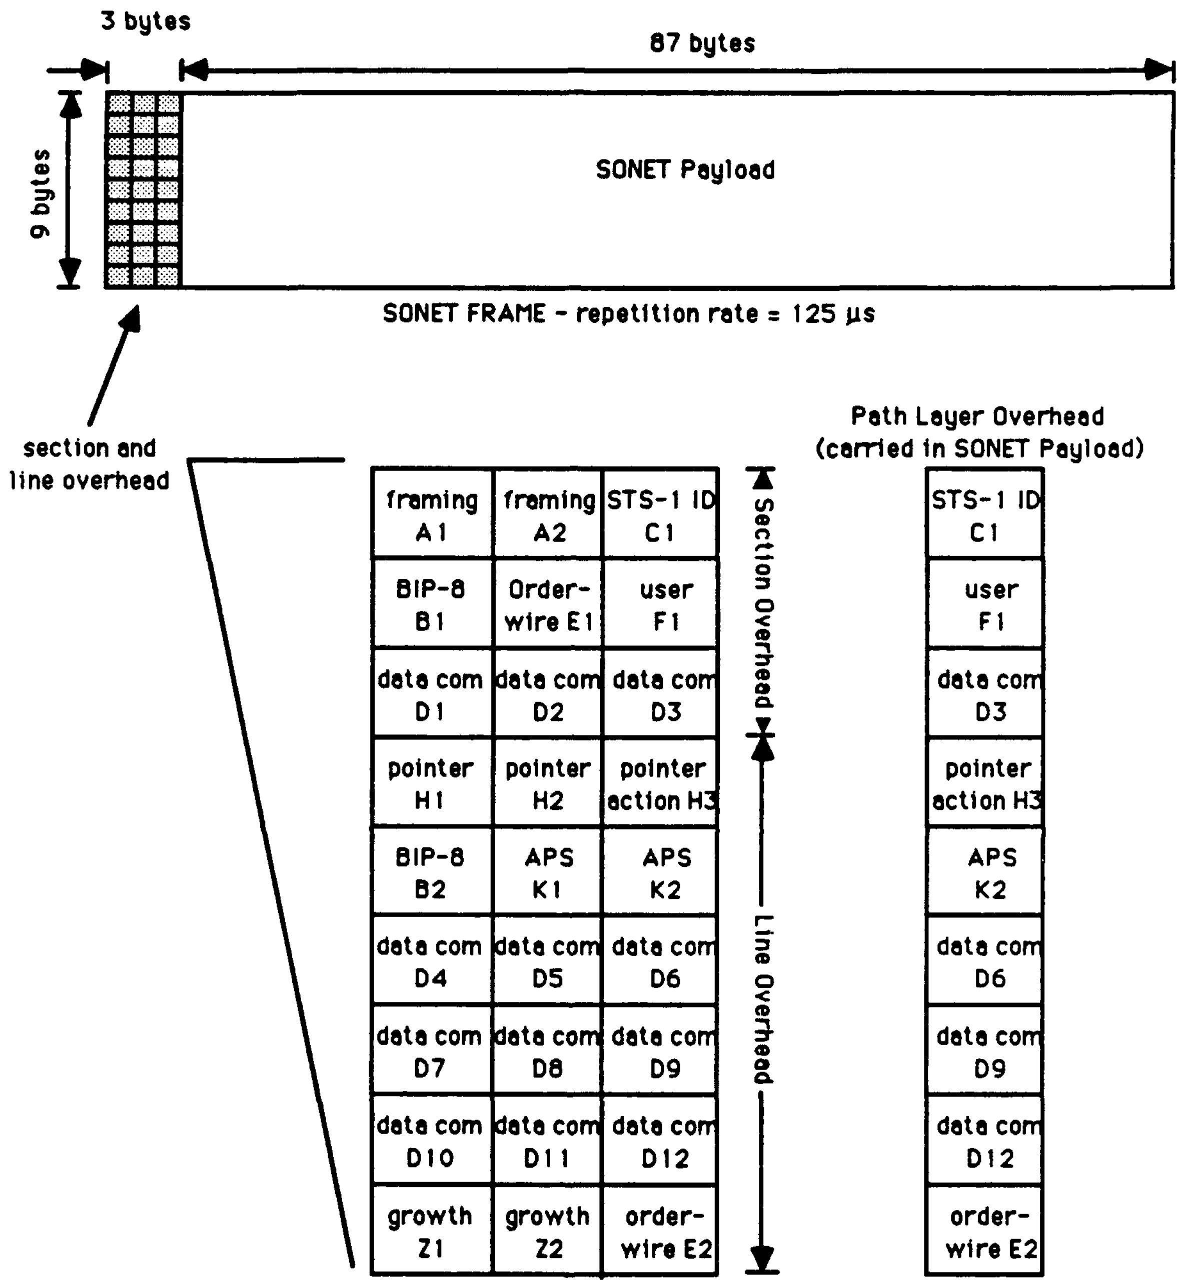

Figure 11.1: SONET Framing Structure [97]

The standardized, high-speed fiber optic interface specified in SONET is an important step to the realization of ISDN networks capable of transporting broadband traffic at ccmpetitive cost in an increasingly multi-vendor environment.

11.2 Integrated Services Digital Network (ISDN)

The basic goal of the Integrated Services Digital Network (ISDN) standards efforts is to provide generic digital connectivity all the way to the end user. The Basic Rate Interface (BRI) at each telephone provides two Bearer (B) channels at 64 Kbps and one Delta (D) channel at 16 Kbps, for an overall rate of 192 Kbps to the user. Another interface, the Primary Rate Interface (PRI), provides a 1.544 Mbps interface (2.048 Mbps in Europe), which can be allocated in a number of ways, including having 23 B channels and one 64 Kbps D channel. Each of the B channels can be used to transfer digitized voice or (packet or circuit switched) data, with the D channel controlling the usage of the channels and also providing another packetized data channel.

The various military departments like the Air Force, have plans to convert to ISDN switches in the next ten years [99], and the Defense Communications Agency has plans to insure that the military departments can achieve base-to-base ISDN connectivity on a national level [89], likely by leasing facilities from the commercial long haul carriers. As an example of the possible use of ISDN in a military application, a user could be involved in routine training exercises via a data base at another location, connected through the B channels, and could be alerted to the presence of incoming tactical information requiring immediate attention by the D channel [45].

ISDN standard efforts are being carried out by the T1S1 group within T1, in cooperation with the X3T5 technical committee within X3 (Information Processing). X3T5 work centers on the standardization of the OSI upper layers for ISDN, assuring that ISDN will be compatible with the OSI model [46].

11.3 Broadband-ISDN, Asynchronous Transfer Mode (ATM)

Standards for Broadband ISDN (B-ISDN) are being developed within ANSI T1S1 and CCITT Group XVIII for the purpose of determining a generic architecture capable of delivering a multiplicity of traffic types, including various types of video, to and from the user. Traffic which must be handled is likely to include digitized voice and data from ISDN terminals, high-speed data and file transfers from supercomputers and workstations, video telephone, standard NTSC TV, Enhanced Quality TV (EQTV), High Definition TV (HDTV), video conferencing, and video on demand.

As an initial implementation of B-ISDN, it is likely that a symmetric SONET OC-3 (155.52 Mbps) interface will be provided to the end user from the central office. This interface would be capable of providing two or three standard video channels, video telephone, voice, and/or high-speed data services to a residential or business end user. It would also be capable of connecting a high-speed LAN or MAN, like FDDI or IEEE 802.6, or an ISDN PBX, to the national network. As technology progresses and standards and markets solidify, an asymmetric interface might be defined, where a 600 Mbps signal (probably SONET OC-12) is passed from the network to the user, and an OC-3 is passed from the user to the network. The reason for the asymmetry is that the video service to the end user would likely progress to HDTV, requiring anywhere from 150 to 600 Mbps per channel. HDTV would not be transferred from the user to the network. Single mode fiber optics would have to be installed to the user interface or broadband terminal from the central office. Communications between central offices or to hubs would have to take place at a much higher rate, possibly 2.49 Gbps.

To handle the varying types of traffic over a common physical channel, CCITT has established as a baseline the use of a fast packet switching technique call Asynchronous Transfer Mode (ATM) [100]. Fast packet switching has been proposed for the multi-service environment because of its ability to efficiently handle bursty traffic and to multiplex traffic patterns of varying rates and periodicity onto a single physical channel. ATM takes advantage of the improved physical layer reliability of single mode fiber to reduce the latency of a packet as it passes through a network, and to reduce the amount of header overhead required. An end-to-end check on errors can be performed, instead of the traditional store-and-forward, or link-by-link error checking that has been used with packet networks in the past.

Fast packet switching is a hybrid between circuit switching, the technique currently used in the *telecommunications network*, and packet switching, the technique used in most local area networks exchanging data. With circuit switching, a fixed connection, or "circuit" is set up between the source and destination, and this circuit remains in place until it is dismantled at the end of the call. Packet switching is analogous to mailing portions of a large manuscript through the U.S. Postal System using separate envelopes for each page. Fast packet switching is similar to conventional packet switching in that information is packetized, however, packet sizes are very small in comparison, to allow for time sensitive information to be packetized, travel through the network, and be de-packetized in time. Packet lengths are 53 bytes, compared to around 2000 or more bytes for conventional packet networks. Fast packet switching, in conjunction with coding mechanisms, enables higher network efficiency (silent time in phone conversations does not tie up network resources) and permits greater sharing of the network among many different types of users [101,102].

First implementations of ATM will likely use the 155.52 Mbps SONET interface, including the framing structure previously discussed. ATM cells will be placed into the envelope of the frame, as shown in Figure 11.2 [102].

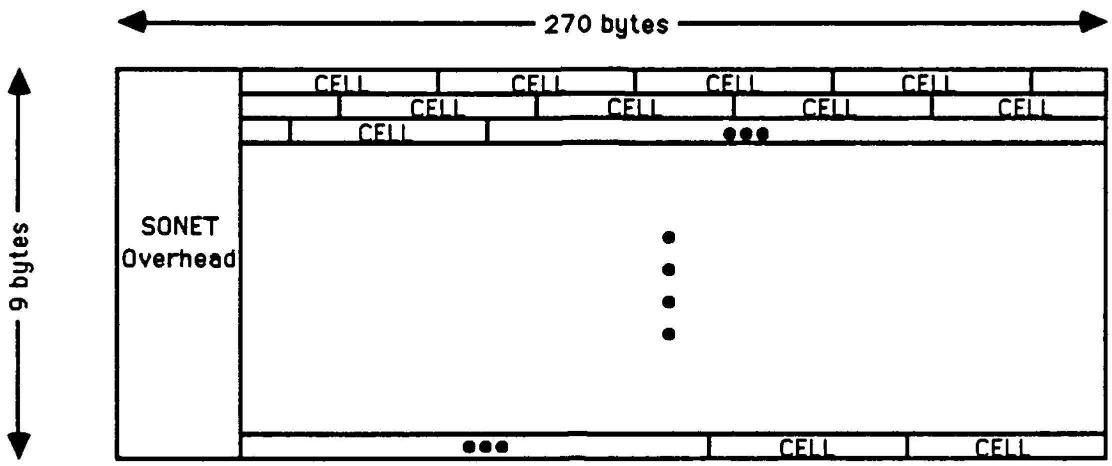

Figure 11.2: ATM Cells in SONET OC-3 Frame [102]

A pure ATM protocol (i.e without the use of the SONET frame), could be used at 600 Mbps, although it is not known if packet switching technology will be capable of switching ATM packets in a public network at the higher rate. This, along with the problem of echo suppression in voice transmission using ATM, could possible impede its complete acceptance [103]. The simpler switching required in the physical ring topology could possibly permit the use of ATM at higher rates in private, ring-based networks earlier than in the public network.

12. Appendix B - Simulation of Three Network Architectures

Three logical network topologies discussed in Section 6, the multi-channel ATM ring network, the single channel byte-wide ATM ring network, and the perfect shuffle, or banyan, ATM network, were modeled using the SLAM-II simulation language [104]. Each network was assumed to have eight nodes and to use a register insertion protocol. Section 12.1, 12.2, and 12.3 contain the SLAM-II code for the three networks. Comments describing particular aspects of the code are included at appropriate places. The SLAM-II network models representing one of the eight nodes in each network are shown in Figures 12.1, 12.2. and 12.3, respectively.

There are several common features in the SLAM-II network implementation of the three models which allow for a fair comparison of the network architectures being studied. These include the technique for synchronization of the network, the priority scheme for ordering of packets in the queues at each node, the requirement for an acknowledgment of data packets, the distributions of the random input sequences controlling the cell interarrival times and choice of destination node, the queue capacity at each node and the methodology for handling lost cells, the overall node transmission capability, and the link-to-link delay. These points will now be discussed in greater detail.

Each node has the same overall transmission resources at its disposal. As is seen in the SLAM-II code or network model, nodes in two multi-channel networks have a service activity following the priority queues which has eight servers with a fixed service time of eight time units. The single channel network has one server with a fixed service time of one

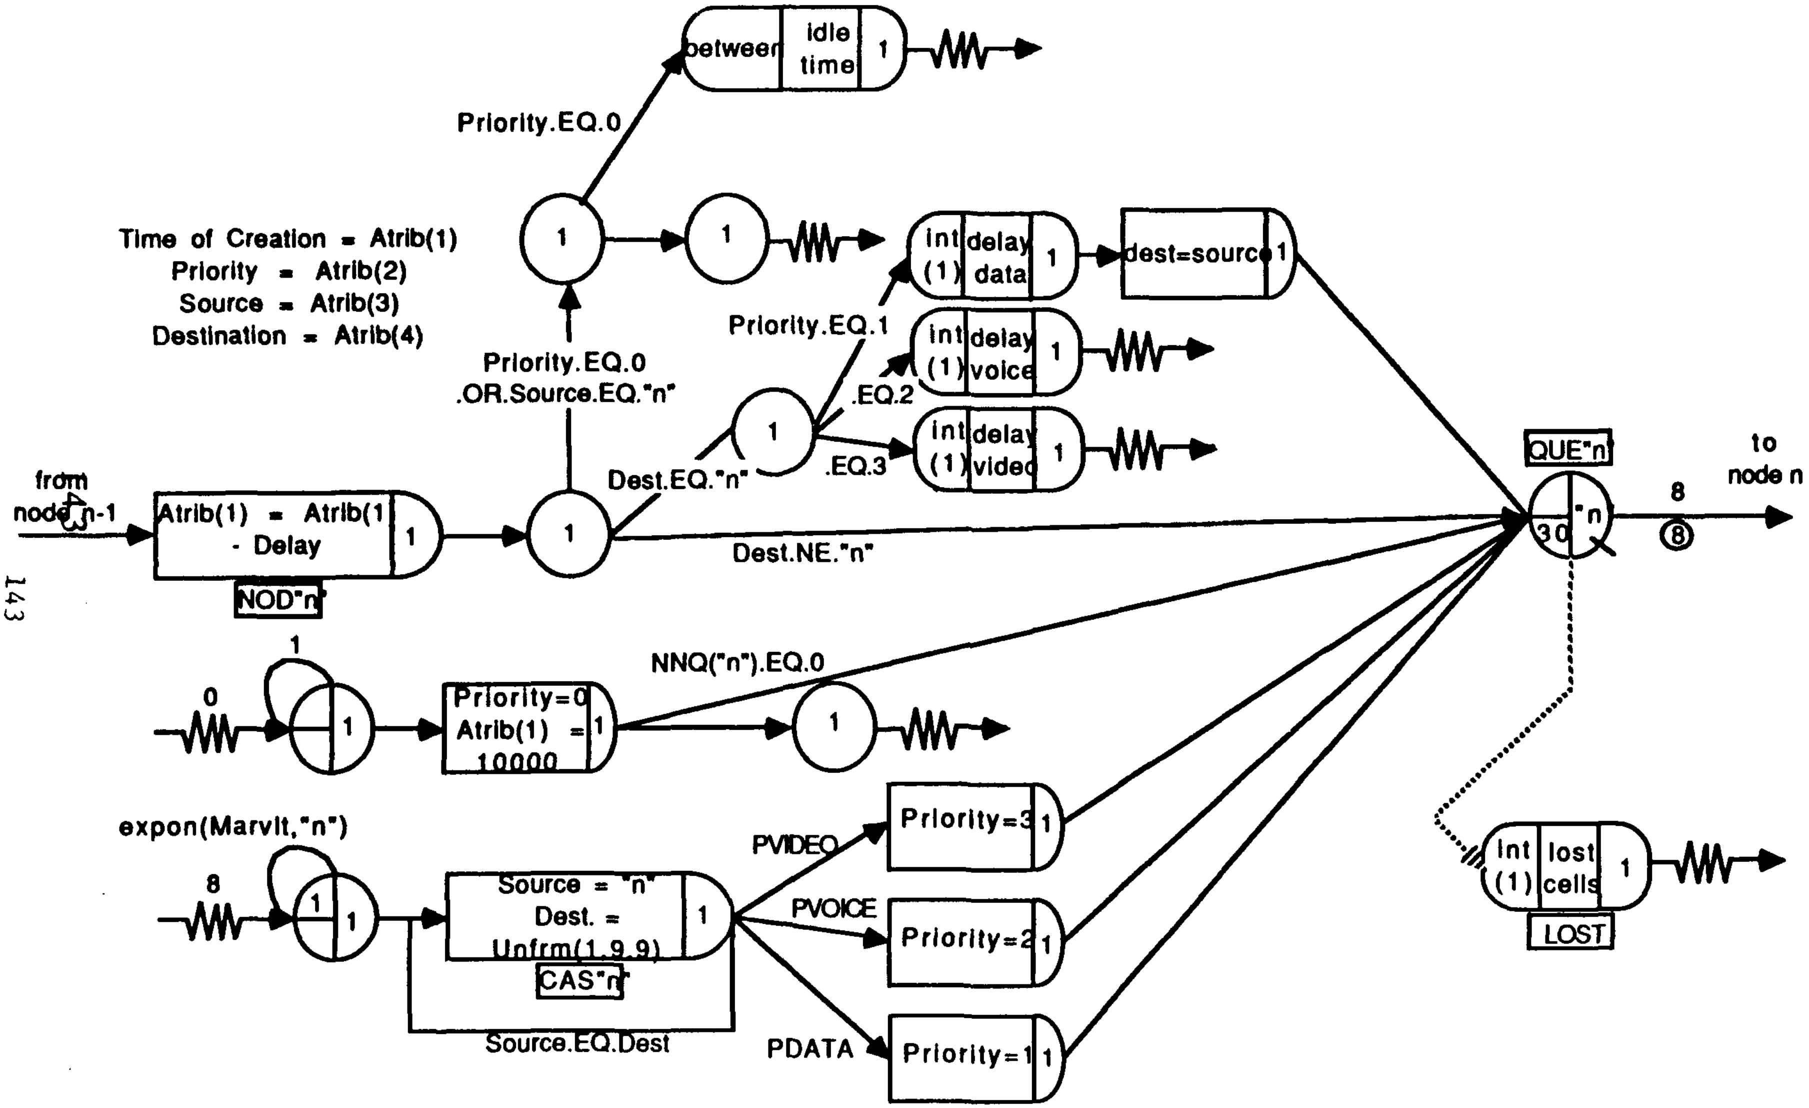

Figure 11.1: SLAM-II Network Model for Node "n" of Multi-Channel ATM Ring Network

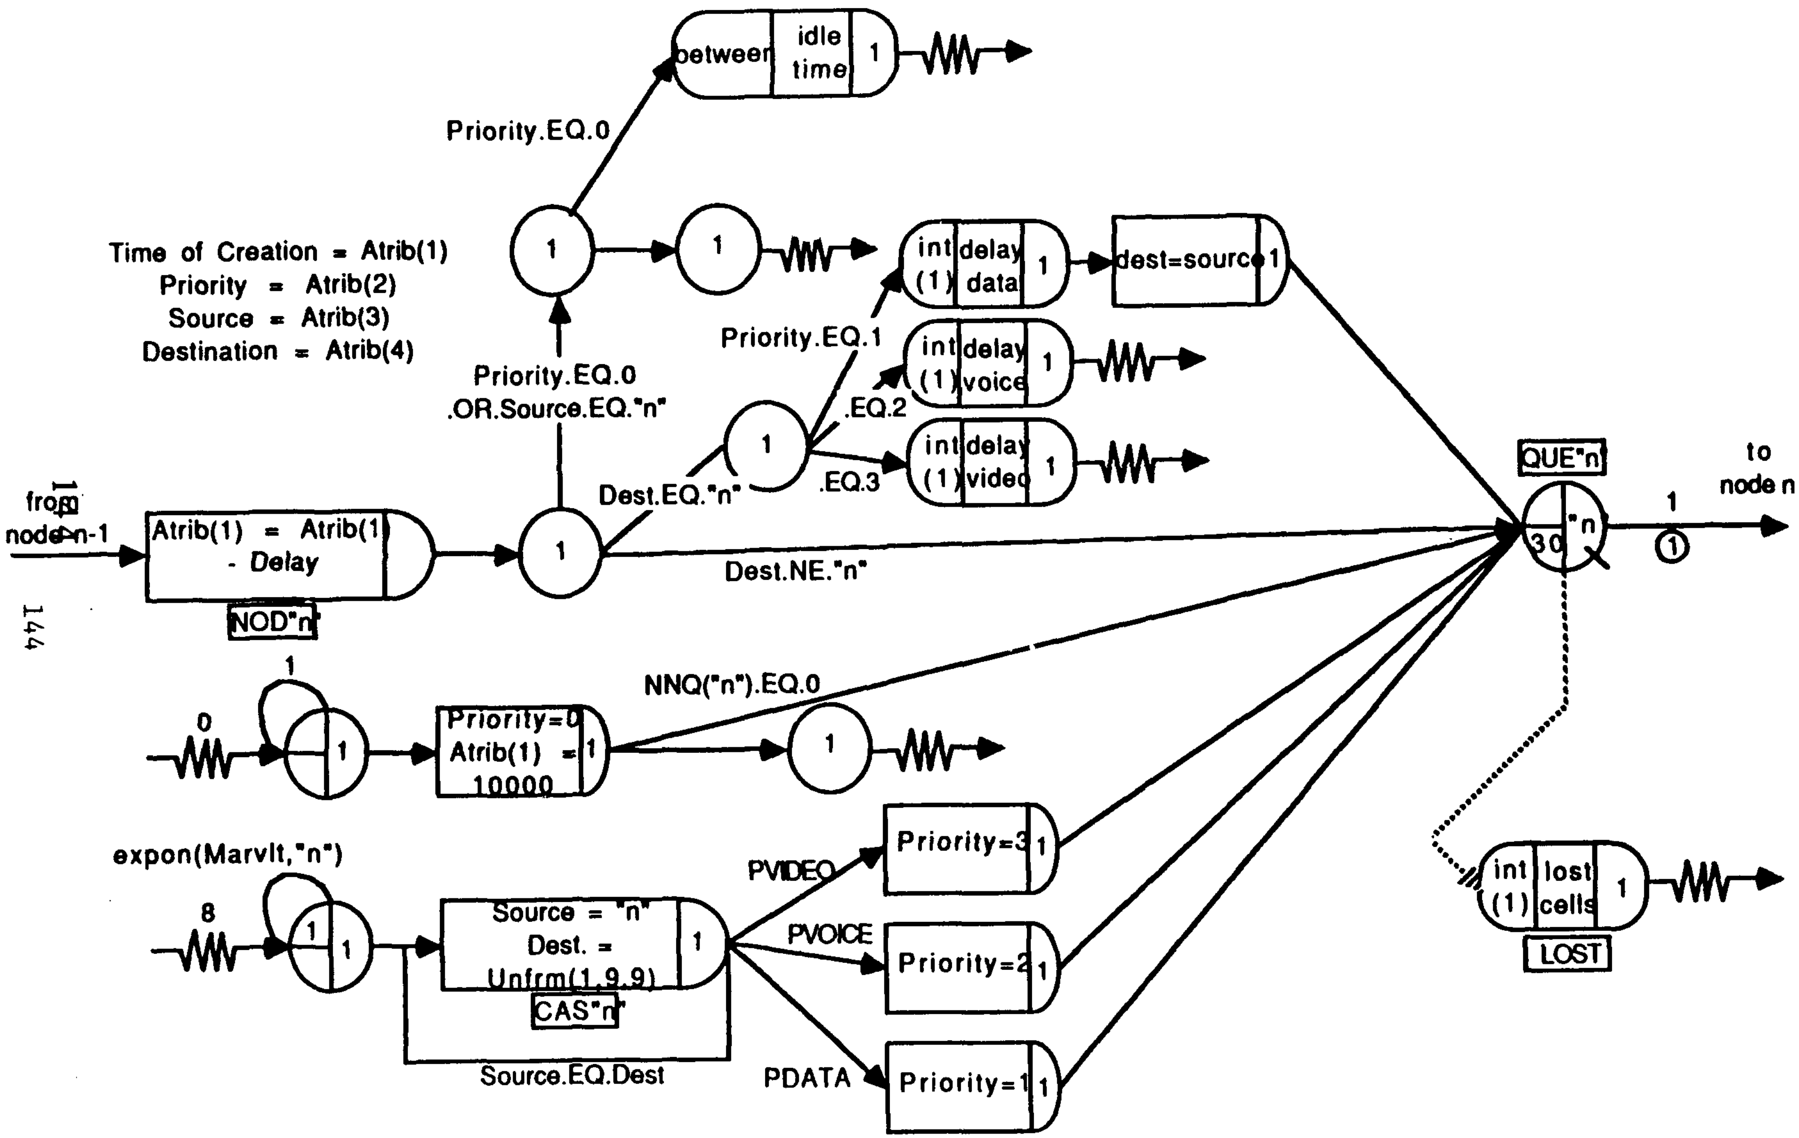

Figure 11.2: SLAM-II Network Model for Node "n" of Single Channel Byte-wide ATM Ring Network

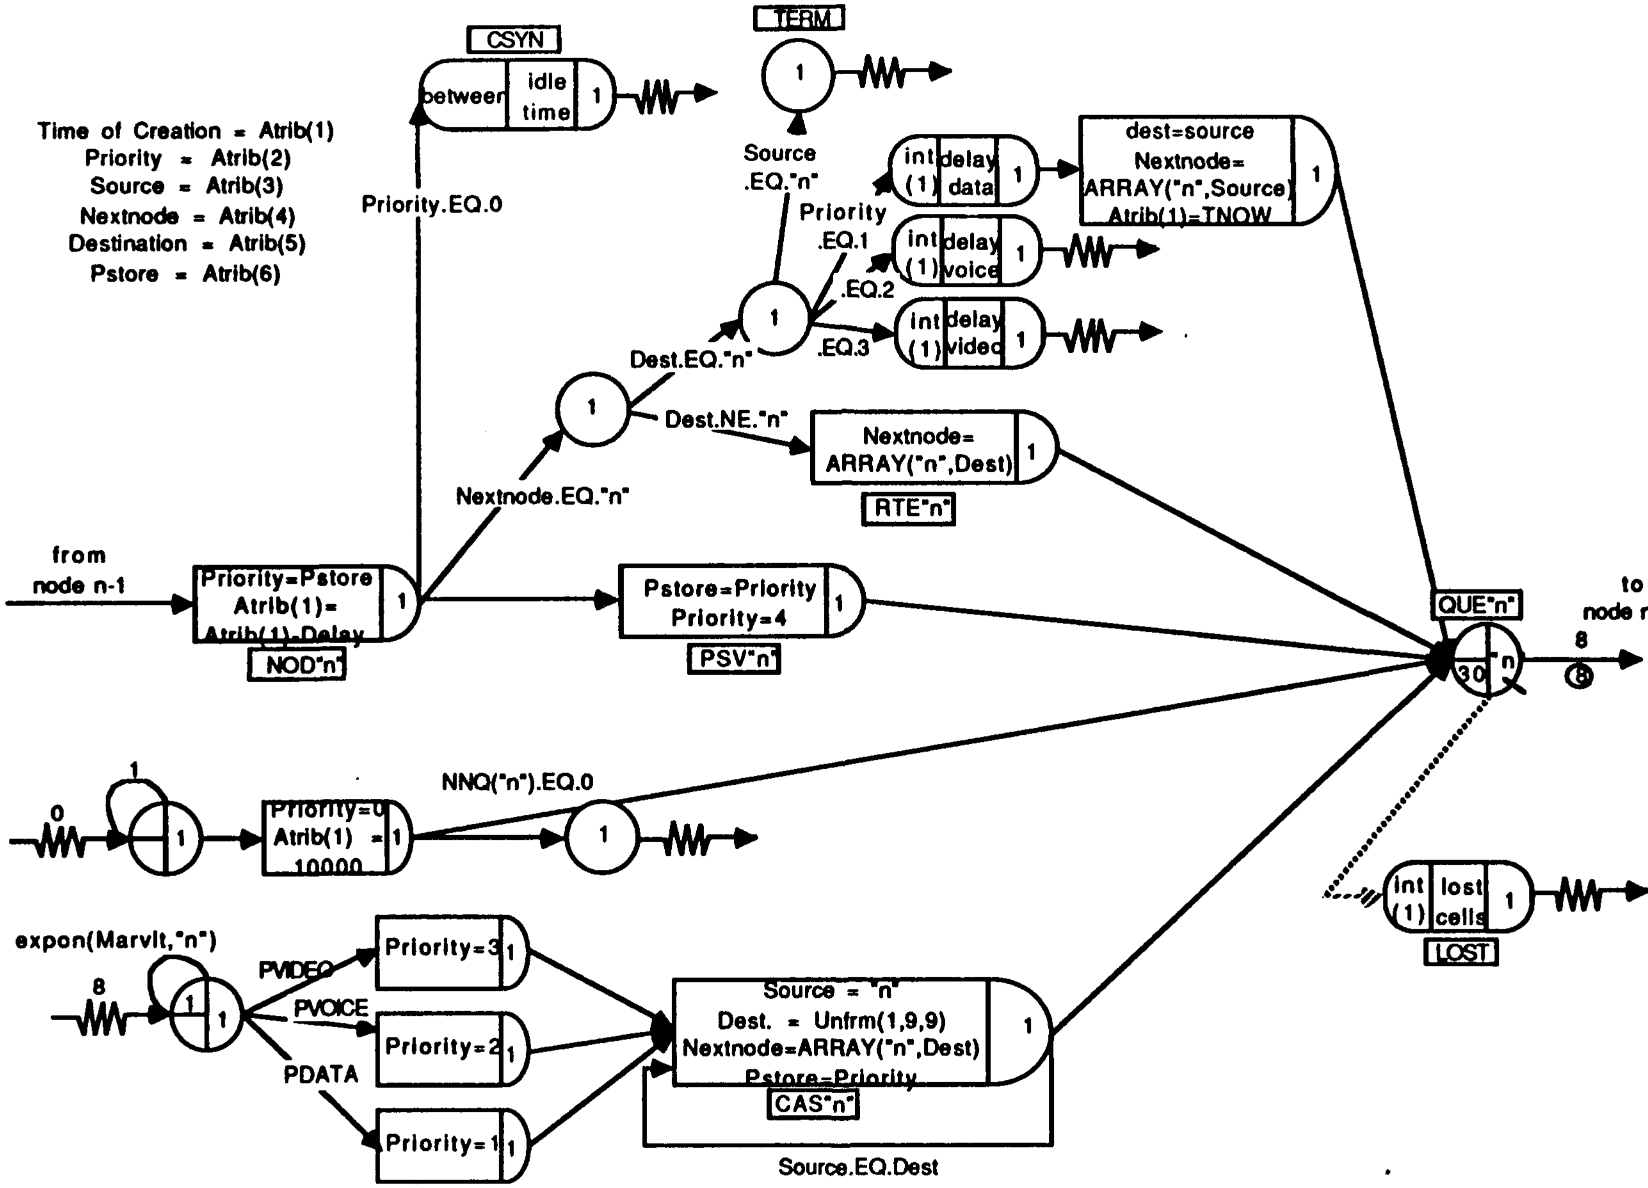

Figure 11.3: SLAM-II Network Model for Node "n" of Perfect Shuffle ATM Ring Network

time unit. Thus all networks can transmit one cell, or entity, per time unit.

Each queue is organized from high priority to low priority, with the priority of the cell being contained in attribute two. Video cells have priority three, voice cells have priority two, and data cells have priority one. In the perfect shuffle network, cells for which the node is not the "next node" in the cell's travels are immediately given a priority level of four so they can speed through the node unaffected by its current traffic. This is a temporary assignment which is reevaluated at the next node. The criterion for settling ties in the queue is the time of creation of the cell, which is contained in attribute one. The older the cell, the closer to the front of the queue it is placed. Data acknowledgment cells have their time of creation updated from that of the original data cell, and they retain their level one priority.

To keep synchronization, idle cells must be sent whenever there is a void in the traffic. This is accomplished by creating entities of priority level zero which are sent to the queues at each node if the number in queue falls to zero. These entities also have very large values in attribute one so that the event clock will not cause their transmission before another entity coming into the node from the network and destined for another node can make its way through the logic between a node's input and its output.

Synchronization entities are removed at the node immediately following their introduction, and sent to be counted before they are terminated. If necessary, new synchronization cells are generated by the next node. Synchronization cells are counted to determine the total number of idle slots for throughput calculations and to verify the correct operation of the network, as will be discussed later.

The eight nodes in the wide area networks modeled here are intended to represent the highest level in the telecommunications hierarchy. Therefore, due to the large number of sources of information coming into the network, burstiness of traffic is likely to be eliminated. The traffic can best be modeled as having exponential interarrival times at each node. The network is assumed to be uniform, so traffic distributions into each node are assumed to be identical. The distribution for choosing the destination node for information arriving at a particular source node is also assumed to be uniform. The source node and destination node are forced to be different, as it is assumed that such a cell which was already at its destination node would not enter the network. Each network is initialized to the same starting seed for use in random number generation. This reduces the variance in comparing the outputs from the three models.

The queue capacity for all queues in all network models is set to thirty, reflecting realistic implementation capabilities. When a queue is full, balking occurs and the cell is considered lost and is sent to be counted before it is terminated. The number of cells lost is a basis for comparing the three networks.

The propagation delay between nodes can be modeled by either inserting an activity of duration equal to the delay, or by reducing the effective time of creation of the entity by the delay for each delay encountered. For long delays between networks, the first technique would require that a prohibitively large number of entities exist on the network at the same time, quickly exceeding the memory capacity allocated for the simulation. Therefore the second technique is used. Variable XX(6) is used to represent the link delay, which should be normalized to the transmission time required to transmit one cell in the byte-wide format. For example, for a 2.49 Gbps fiber optic transmission link, a cell length of 53 bytes, and a link distance of 500 kilometers, the delay should be set at 117453. Therefore, 117453 cells in byte-wide format would simultaneously exist on one link between nodes.

Model differences reflect the differences in the networks as pointed out in Section 6. These differences include the routing table required for the perfect shuffle network, accompanied by the need for implementation of a correct node-to-node connection pattern; the transmission rate of the channel for the single versus multi-channel versions; and the number of attributes required. These differences will now be discussed in greater detail.

The perfect shuffle network requires the routing table, which is implemented using the ARRAY capability in SLAM-II. The routing table is closely linked to the physical connectivity of the nodes, which is different for this network than for the other two. There are only a few network configurations which yield minimum system delays for a given routing table determined from a logical perfect shuffle network.

As previously mentioned the single channel byte-wide network has a channel transmission rate which is eight times faster than for the other networks, but there is only one channel.

The number of attributes required to implement the two slotted ring networks is four in both cases. Attribute one is the time of creation, attribute two is the priority, attribute three is the source of the cell, and attribute four is the cell destination. The perfect shuffle network requires these four attributes plus two additional attributes. One of the attributes contains the next node, or logical channel, of the cell, while the other provides temporary storage of a cell's actual priority as it receives a priority level of four at nodes where it is intended to pass directly through.

The models were first verified by using input conditions for which expected performance could be easily analyzed. For low utilizations, the average cell time in system should be slightly greater than four times the sum of the transmission time plus the link delay. The minimum time in system should be slightly greater than the transmission time plus the link delay. The maximum time in system should be slightly greater than seven times the sum of the transmission time plus the link delay. The amount by which the average, minimum, and maximum times in system exceed the above described sums should be approximately equal to four, one, and seven times the average waiting times in the queues.

For example, the average time in system for the multi-channel networks with zero delay between nodes should be slightly higher than 32 time units for low utilization cases. The minimum should be near eight, and the maximum just above 56. We will later see that this is verified by the results shown in section 12.4.

The number of slots actually filled during the simulation in a time interval should be close to the expected number of slots which should be filled during the time interval. The total number of slots actually filled is equal to four times the number of voice and video cells sent (since voice and video cells travel, on the average, half way around the eight node network) plus eight times the number of data cells sent (since data cells and their acknowledgment completely circumnavigate the eight node ring) plus the number of idle, or synchronization, cells sent. The total number of slots created is equal to eight times the time interval, since there are eight nodes and each transmits in one slot per unit of time. Table 12.1 shows the number of voice, data, and idle cells sent during a simulation interval from 500 to 2000 time units. The total number of slots filled is calculated and compared with 12000, equal to eight times the time interval of 1500. Data is taken from simulation runs summarized in section 12.4. The Table shows that the slot mechanism is functioning correctly.

Table 12.1: Slot Utilization

Run Number	1	2	3	4	5	6
Data Cells	507	506	504	503	504	503
Voice Cells	476	475	478	478	476	477
Video Cells	516	516	508	508	510	510
Idle Cells	3921	3922	3952	3952	3954	3952
Total Slots	11945	11934	11928	11920	11930	11924
Expected Slots	12000	12000	12000	12000	12000	12000

Section 12.4 presents data from six runs, each with a mean arrival time at the nodes of eight. There are two runs from each of the three networks. The first run of each pair is with the queues organized according to cell priority level, while the second run is with the queue organized in a first in-first out (FIFO) manner. The effects of the queue organization on time-in-system for each of the three data types can be seen, particularly in the maximum time in system values. This verifies that the queueing procedure used is operating correctly and is in fact more efficient than a FIFO procedure. The first run of each pair (for priority queues) shows that average, maximum, and minimum times in system are as expected.

Section 12.5 shows plots for the number of entities in queue one for voice only data at a mean arrival rate of 4.5. Two seeds were used to generate the two plots, the first plot generated with a seed of 10, the second plot generated with the opposite seed of -10. The first plot attained a maximum queue length of 14 cells, while the second run attained a maximum queue length of 19. The plots show that the seed values can impact the data if run lengths are not long enough, but that the initial condition of zero entities in queue has very little impact on the number in queue after time intervals of a few hundred units. All statistical counters in the simulations were reinitialized at 500 time units. For utilizations lower than the one represented in section 12.5, initial conditions will affect network performance for a longer time, but since the networks are operating at steady state levels closer to initial condition levels the effect is lessened.

Average and maximum time in system, along with maximum number in queue, data was collected for cell interarrival times ranging from 50 to 4.1 time units for all voice traffic for the three networks, assuming no link delay. The results are shown in Table 12.2. Similar data was collected for data only traffic having interarrival times ranging from 50 to 8.1. These results are presented in Table 12.3. PS designates the perfect shuffle network, MC the multi-channel ring, and SC the single channel ring. Values in parentheses represent the number of cells lost due to queue overflow during the run.

Table 12.2: Voice Traffic Performance

Interarrival	Ave. System Time			Max. System Time			Max. No. in Queue		
Time	PS	MC	SC	PS	MC	SC	PS	MC	SC
50	33.1	33.1	4.74	58.5	59.2	9.73	3	3	3
20	34.1	33.1	4.94	60.9	59.2	12.0	4	4	4
15	33.9	33.6	5.37	62.5	62.5	14.2	4	4	4
10	34.4	34.7	5.92	63.8	65.5	16.9	4	6	4
8	35.0	35.5	6.83	69.0	68.0	24.9	8	6	8
6	36.5	37.3	8.54	74.5	74.4	24.9	10	8	8
5	39.0	39.7	11.0	82.0	85.6	33.3	12	14	12
4.5	43.6	43.9	15.8	93.7	95.5	45.8	16	18	18
4.25	55.8	52.0	24.2	140	127	71.3	30(1)	30(4)	30(7)
4.1	70.5	67.7	39.2	194	161	111	30(23)	30(12)	30(14)

Table 12.3: Data Traffic Performance

Interarrival Time	Ave. System Time			Max. System Time			Max. No. in Queue		
	PS	MC	SC	PS	MC	SC	PS	MC	SC
50	33.2	33.2	4.82	59.3	60.2	9.73	3	3	4
20	34.8	33.8	5.50	63.7	63.2	15.9	4	4	4
15	35.3	34.9	6.49	65.7	67.7	18.2	5	6	5
12	36.4	36.2	7.60	73.8	72.5	22.7	6	7	6
10	38.6	38.1	9.32	83.2	78.7	31.9	10	8	8
9	41.0	40.6	12.1	86.9	90.8	39.8	13	13	10
8.75	42.8	42.1	13.3	92.1	96.0	45.5	14	13	11
8.5	44.9	43.7	14.8	95.8	100	51.7	15	14	14
8.25	47.5	47.6	17.2	104	105	59.0	18	17	15
8.1	52.3	52.3	20.6	114	114	63.9	18	16	16

As is seen from the tables, the single channel byte-wide network has lower system times because of its lower transmission time. The maximum number of cells in queue, and the number of cells lost are very close for all three networks. Queues begin to overflow for all three networks when utilization approaches 100%. This occurs, as expected, when the interarrival time for all voice or video traffic is four time units, or the interarrival time for all data traffic is eight time units.

Section 12.6 shows the summary of two runs where there is a delay between nodes representing a propagation distance of over 400 km (for 53 byte cells at a data rate of 2.49 Gbps), typical of the wide area network applications under consideration. The first run is for the single channel network, while the second is for the multi-channel ring network. The average and maximum time in system values indicate that for high delay between nodes the transmission time of the channel has little impact. Therefore, since maximum number in queue and number of cells lost are roughly equal, the three networks modeled are very close in performance. Other factors, such as complexity, survivability, reliability as discussed in Section 6 should be considered in choosing between the three.

13. Appendix C - Security in an ATM Environment (written by Franc E. Noel)

13.1.0 Introduction

The purpose of this paper is to review the state of the art in implementable security measures for computer communications networks. This paper is intended to supplement a paper written by Michael R. Slawson titled "Survivable Broadband Wide Area Network Selection."

This paper reviews the current security practices of both commercial and military systems, as well as the network assumptions and end user importance. A close look at the OSI Security Architecture is made since it represents a comprehensive study of network security. Included in this review, will be a summary of security threats/services that are addressed by the OSI Security Architecture. Finally the current issues in security, from an OSI layering point of view, are elaborated. Since any modern discussion of communications security involves the National Bureau of Standard's data encryption standard, the algorithm is discussed in Section 13.8.0.

13.2.0 Current Security Practices

13.2.1 Commercial Networks

The utilization of data security techniques in commercial networks has grown with the level of competition in business, which becomes more competitive daily. Also, the availability of products has improved which allow the network supplier to smoothly integrate data security functions with their existing networks and provide higher function, from a corporate point of view, at minimal increase in overhead, from an end user's point of view.

The leading edge in widely installed large scale corporate computer network architectures is represented by IBM's System Network Architecture (SNA). This architecture focuses on network security on many different planes, from the controlling of physical network resources, to the higher levels of end to end communications security. As such, this paper will use SNA as a representative example of today's widely available commercial networks, and review its communications security functions, in particular those at the higher OSI/SNA layers.

The primary, and first released, security function in SNA is the encryption by the Connection Point Manager of the Request Units in the Transmission Control function [MEIJE 88]. The encryption technique used is the National Bureau of Standards' DES algorithm with block chaining (See Section 8.0 for description of the algorithm). At session establishment, the end users agree to either mandatory or optional encryption. When optional encryption is used, the Enciphered Data Indicator (EDI) bit in the Request Header indicates if an individual packet has been enciphered.

SNA provides management functions to address the protection of session keys. Each Logical Unit (end user network attachment) is provided with a Device Encryption Key, where it is securely stored. The Device Key is also stored in the System Services Control Point (SSCP) using an SSCP master key. IBM provides special hardware to provide for the physical security of the device keys. When a session is established between two Logical Units (LUs), the SSCP selects a session dependent key that is session unique. The session key is encrypted twice, once using the primary LU's device key, and once using the secondary LU's. Both encrypted keys are forwarded by the SSCP to the primary LU, one for its own decryption and one to forward to the secondary LU as part of the session initialization process. Hence each session has a unique key, and the device keys are only used for the distribution of the session keys.

To insure that the preceding process has been successful, the secondary LU generates a 64 bit random number called the session seed, encrypts it with its session key and sends it to the primary LU. The primary LU decrypts the message and inverts the first four bytes. This so called test value is then encrypted and sent to the secondary LU for verification, which responds with the appropriate acknowledgement.

The above security measures were enhanced by IBM in 1982 with the announcement [IBM 3073, IBM 3112] of Advanced Program-to-Program Communication (APPC) which included an advanced function logical unit type, LU 6.2. LU 6.2 provided new security functions in the area of verification of the identity of two logical units before a session is established, and to verify the validity of a program/user to access a secure resource, such as a file server.

13.2.2 Military Networks

The present Defense Data Network (DDN) was initiated in 1982 to replace the aging AUTODIN network and to provide a common-user multi-level secure data communications system. As of today, this network is nearly fully operational and consists of five individual subnetworks as listed in Figure 13.1 with their appropriate security level.

DISNET 3	TOP SECRET(SCI)
DISNET 2	TOP SECRET
DISNET 1	SECRET
MILNET	UNCLASSIFIED
ARPANET	EXPERIMENTAL

Figure 13.1 Defense Data Network Subnetworks

It is the goal of the Defense Communications Agency (DCA) to integrate all of these networks into a common backbone network as affordable end-to-end encryption (E^3) devices become available. By 1995, full integration of voice, video, and data is expected with implementation of ISDN if it is available.

The current networks utilize transmission facilities primarily in the range of 9.6 kbps to 56 kbps, with limited usage of 1.544 Mbps (T1) services [MAYAB 86]. The current planning for these networks shows no references to transmission facilities above T1 and clearly indicates that the world wide DDN networks will continue to utilize leased Continental U.S. (CONUS) Commercial Carriers as well as Foreign Government Carriers [VISIO 88].

The DCA has mandated the use of the protocol architecture shown in Figure 13.2, with the optional use of TELNET, FTP, SMTP, and Native Mode (i.e. Terminal emulator) being recommended, but not mandatory.

It is of interest to note that waivers have been granted for the usage of alternative protocol suites and that those who have implemented the mandatory protocols are NOT guaranteed interoperability unless the "same vendor" protocols and X.25

Transmission Control Protocol (TCP)

Internet Protocol (IP)

ARPANET OR X.25

Figure 13.2 Defense Data Network Protocols

implementations are utilized. Currently there is discussion, initiated by the National Research Council, on replacing TCP with the ISO Transport Protocol (TP) [MITRE 86].

The security features that have been implemented include end-to-end encryption, link encryption on all interswitch trunks, passwords and other procedural protocols, key distribution center, access control center, callback techniques, security clearances, and of course physical site security. Comparing these with commercial network security features shows that current commercial networks utilize all of the above techniques except for link level encryption and separation of network traffic by security classification onto separate networks, i.e. Figure 13.1. Even techniques such as callback are utilized in commercial networks. For example, callback is used in the IBM network for use in its employees' Home Terminal Program to insure the physical identity of the end user who is not actually on IBM premises.

There appears to be a strong session level similarity between the state of the art in commercial systems and military systems, probably brought on by the wide spread availability of end-to-end encryption technology and communications architecture support of such facilities.

Certainly in the defense environment there is a stronger notion of absolute security than is found in a commercial environment. Although, implementation cost is not irrelevant, there appears to be no comparisons involving worth of information vs cost of breaking the security measures in the literature reviewed.

13.3.0 Network Assumptions

13.3.1 Topology

It is clearly necessary for large networks, both commercial and military, to utilize public switch networks, albeit supplemented at times by dedicated satellite facilities [VISIO 88]. This requirement means that packet header information cannot be encrypted, since the public switching facilities must route the packets. It also implies that the network is susceptible to traffic pattern analysis. It further implies a vulnerable point for terrorist activity since telephone central and toll offices utilize a cable vault facility for cables entering the facility. On the external side of the cable vault, the trunk cables are routed through a common manhole for routing to specific cable ducts. A terrorist attack on a finite (and small) number of these cable facilities throughout a country could cripple a communications network in the short term.

On the other hand, use of the public switched network provides for a larger number of alternative network interconnection paths than might otherwise be affordable. In addition, it provides some resiliency against terrorist attack.

Given the dependency on the public switched network, the military networks will need to migrate to ISDN. However, this migration will be difficult to accomplish in light of the slow development of the required standards and the nonuniform implementation strategy of each country, telecommunications subdivision, or political subdivision.

13.3.2 Traffic

As was seen above, the opportunity to analyze the military networks' traffic pattern is possible in the currently planned networks. Hence, it would be highly recommended to provide message stuffing on the network to insure that traffic analysis is not effective. To effectively provide for message stuffing, the additional packets must be generated at or above the level doing the encryption, which is generally the Presentation layer.

However, in order to know that message stuffing is required requires knowledge only found in the Physical or Data Link Control layers. To provide the required communication in the ISO model would be cumbersome [STALL 88]. More effective implementations are possible in the Department of Defense (DOD) architectures where management functions can directly access multiple layers simultaneously.

As network traffic rates increase, one would normally assume that vital messages will be less detectable, since they will become a small percentage of a large volume of traffic. However, the standard military practice of providing a readable security label on each packet will help insure that an intruder will not have to waste time on decrypting undesired messages.

13.3.3 End User Integrity

At the boundary of each network node lies an application process/end user. It has been estimated that the reliability of a human being is on the order of 10 E - 5 [DHILL 89]. To illustrate the vulnerability of a secure network to a human's interference, I submit the following Trojan horse example [VINIO 90]. Let's assume that we have two different networks with security classifications, Top Secret and Secret, that are interconnected and are password protected. If User 1 (Clearance = Top Secret) wishes to provide a password to User 2 (Clearance = Secret), User 1 sends hourly a non encrypted message of a meaningless nature. If the message is sent (time-stamped) before the hour a 1 bit is implied, if the message is sent after the hour a 0 bit is implied. Only 64 such messages need be sent in order to transfer a 64 bit encryption key.

The above illustration can be readily supplemented with examples of human element vulnerabilities, e.g. Why spend a million dollars on a decryption engine for breaking encryption keys, when you can bribe someone for considerably less money and get the specific document you wish. The purpose here is to state a point often made, without focus on the "why", in the security literature that "There is considerable doubt that 'absolute' security is achievable at any cost" [WOOD 83, ISO/2 88].

13.4.0 OSI Security Architecture

The International Standards Organization (ISO) has made a thorough study of the issue of computer communications security and has updated its Open Systems Interconnect (OSI) Reference Model with an addendum, ISO 7498/2 devoted to the issue [ISO/2 88]. Where the original ISO 7498 document covered the area of open systems architecture, this addendum extends the Reference Model to include communications between open systems. As with the ISO 7498 document, however, it is "a framework for coordinating the development of standards not an implementation specification."

13.4.1 Security Services

The security services necessary in an computer communications environment are developed, at length, in the ISO 7498/2 document and are listed and explained in Figure 13.3.

Peer Entity Authentication	Provides corroboration that the peer entity is the claimed entity.
Data Origin Authentication	Provides corroboration that the source of the data is the claimed entity.
Access Control Service	Provides protection against unauthorized use of resources accessible via OSI.
Data Confidentiality	Provides for the protection of data from unauthorized disclosure.
Connection Confidentiality	Provides for the confidentiality of all user data on the *connection.*
Connectionless Confidentiality	Provides for the confidentiality of all user data in a single connectionless packet.
Selective Field Confidentiality	Provides for the confidentiality of selective fields within the user data.
Traffic Flow Confidentiality	Provides for the protection of the information which might be derived from observation of traffic flows.
Connection Integrity w/ Recovery	Provides for the integrity of all user data on a connection and detects any modification, insertion, deletion, or replay of any data within an entire packet sequence (with recovery attempted).
Connection Integrity w/o Recovery	As above, with no recovery attempted.
Selective Field Connection Integrity	Provides for the integrity of selective fields within the user data over a connection.
Connectionless Integrity	Provides integrity assurance of a single connectionless packet.
Selective Field Connectionless Integrity	Provides for the integrity of a single packet.
Non-repudiation, Origin	The recipient of data is provided with proof of the origin of data.
Non-repudiation, Delivery	The sender of data is provided with proof of delivery of data.

Figure 13.3 ISO Security Services

13.4.2 Layer Security Responsibility

In order to provide for the above services, mechanisms were identified by ISO 7498/2. The primary mechanism was encipherment which handled all the services except three. An access control mechanism was identified as required to provide the access control service and a notarization mechanism was required to provide the two non-repudiation services.

To guide further standards development, ISO 7498/2 also identified which services should be incorporated into the standards for a given layer as an option. Figure 13.4 tabulates the number of services per layer which have been identified in ISO 7498/2, keeping in mind their were 14 services in total.

Layer 1	2
Layer 2	2
Layer 3	8
Layer 4	8
Layer 5	0
Layer 6	3
Layer 7	14

Figure 13.4 Security Services per Layer

13.4.3 Extension To TCP/IP Environments

The merging of the security functions with the OSI model will inevitably be accomplished. Because of the strict layer principles in the model, management functions (e.g. security) are cumbersome to implement. Use of null layers and parallel stacks for data and control will provide the framework required [STALL 88]. However, in the military networks TCP/IP is the mandatory protocol suite. In this environment the security architecture of ISO 7498/2 is directly implementable because of the freedom in the DOD architecture to provide management functions through several layers simultaneously.

13.5.0 ATM Networking Issues

In an ATM environment, packets are switched in a connectionless manner [HANDE 89]. Hence, each packet must contain some identifying feature so that the switch can determine the packet's tolerance to delay. For example, real time video, and real time voice will have to be separated from and given priority over data packets. In order to accomplish this sorting function, in a common backbone network, at least portions of the packet header must be transmitted in an unenciphered form so that packets may be queued appropriately. In a similar fashion packet headers and trailers will contain routing information and other control information that must be available to the switch without encryption.

The transport of video through the secure network raises several interesting questions. First, since compressed video will most likely be transmitted, an intruder would have to try a key on many (10's) of video packets in order to determine if a given key was successful. This is because a common way of compressing video is to transmit only the frame to frame differences. Second if we model the effect of a wrong decryption key as a significant decrease in the signal to noise (S/N) ratio, what is the effect of getting a "close" key to the video image, which can withstand large increases in S/N ratio and still be intelligible (even if it is irritating to watch).

13.5.1 Physical Layer Issues

When one considers the subject of security, the most immediate thoughts are those considering physical security, hence physical layer issues are immediately brought to mind. In the early 1970's optical fibers offered the promise of a secure (i.e. non-radiating) and untappable (e.g. it was difficult enough to install a connector) transmission medium. Today it is generally known that an optical fiber is non-radiating but it is hardly untappable. In fact many devices are marketed, in conjunction with splicing equipment that allows one to carefully bend an optical fiber, forcing some of the optical energy into the higher order modes and allowing the energy to be detected as it radiates from the fiber cladding, hence optically tapping the fiber [ENGEL 86].

In the harsh light of scientific scrutiny, one can only conclude that all transmission mediums must be assumed to be tappable, severable, or jammable with little effort.

13.5.2 Network Layer Issues

As can be seen in the ISO 7498/2 document, there is a strong case for the majority of the security effort to be placed in the network layer or in the application layer. The argument for the placement in the network layer takes the position that security is a network issue and functions, such as encipherment, should be provided by the network. If only simple bulk protection of all end to end communications is required, then network layer encipherment is a strong candidate. However, rarely is bulk encipherment the desired goal. In that case, the application layer must control the destiny of its data with appropriate support from the underlying communications subnetwork.

13.5.3 Intra-Network Routing

As one reviews the various routing strategies in use today [STALL 88], it becomes apparent that the network itself, is highly vulnerable to attack. For example, suppose that an intruder sends false routing table information through the network, disguised as legitimate information from a central node (for a centralized algorithm) or from a neighboring node (in the case of a decentralized algorithm). All such transmissions should be made using an encryption technique to remove the possibility of network tampering, which could be of great consequence.

13.6.0 Higher Layer Issues

As is implied in the ISO 7498/2 document, there is strong argument for all security functions to be provided in the application layer. While this may be of interest to the ISO, it really doesn't help the user, since in essence, the application layer is the end user to the network or directly represents them (e.g. FTP). In any case, the least the communications subnetwork can do for the end user in this context is to provide them with a strong encryption technique, which is easy to utilize on a per message or per session basis. Since the presentation layer deals with syntax issues, this is the logical place for this service.

13.7.0 Conclusions

It would seem that as current networks are brought to the state of the art in security, there is little to be gained by improving on the architectures in place. The human being, at that point, will be the most vulnerable link in the chain. It appears that the only way to circumvent that vulnerability (i.e. provide for absolute security) is to provide no information flow [RANDA 89]. The only significant improvement required, will be to occasionally increase the encryption capabilities to keep pace with semiconductor technology efforts to provide more capable decryption engines. This can be done by finding more powerful encryption algorithms or (more simply) by using longer keys.

13.8.0 DES Encryption

After initial development by IBM in Kingston N.Y., an encryption algorithm was published by the National Bureau of Standards (NBS) for a block cipher in 1977 [FIPS 77, TANEN 81]. The Federal Data Encryption Standard (DES) is commonly implemented throughout the computer industry and has been validated by NBS on a wide range of machines, literally from micros to mainframes [GAITJ 77].

The DES algorithm is based on a 64 bit key consisting of 56 data bits and 8 error detection bits, with the same key being utilized for both encryption and decryption. The encryption algorithm itself, shown in Figure 13.5, is inverted for decryption and the key bits are reordered appropriately. The algorithm manipulates data in blocks of 64 bits (including padding if required). The data is initially permutated, processed by a key dependent computation of sixteen cycles and then finally permutated. The cipher function (f) and the interaction with the 64 bit key can be found in the NBS literature [FIPS 77]. It has been asserted by experts [BARAP 64] that a system which cannot be described safely in the open literature is not sufficiently secure to be used with confidence, hence there is no difficulty in obtaining technical details for the algorithm.

Much debate has occurred [DIFFW 77, MEISP 76, BRAND 77, DAVID 78] over the ability of the 64 bit key to withstand highspeed brute force attacks. To give some perspective in the matter the following should be noted: To determine the key by brute force would require, on average, .5 X 2^{56} or 3.6 X 10^{16} trials.

Each trial requires a minimum of 34 machine cycles (assuming 2 cycles for each time through the loop) for an average of 122.2 X 10^{16} machine cycles. If one were using a 100 MIP machine, it would take 1.222 X 10^{10} seconds of machine time on average, or 387 years to find the key. This computation, of course, does not include the CPU cycles required after each key is tried for an expert system to evaluate the decrypted output to see if it is garbage or legitimate data. The consensus at the time the standard was developed, was that it was at least acceptable for business applications and nonmilitary government applications.

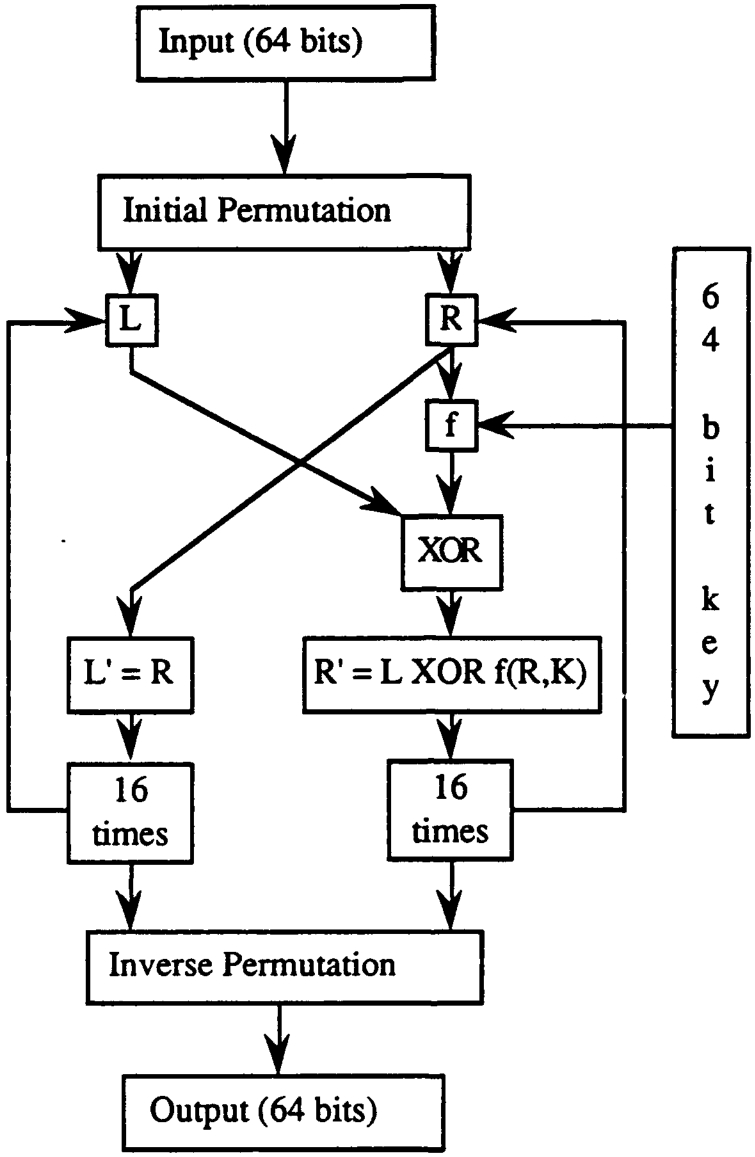

Figure 13.5 **DES Enciphering Computation**

Bibliography

BARAP 64 Barap, P., "On Distributed Communications: IX. Security, Secrecy, and Tamper-free Considerations," Rand Corporation, AD-444839, August 1964.

BRAND 78 Branstad, D.K., ed., "Computer Security and the Data Encryption Standard", National Bureau of Standards, Spec. Publ. 500-27. February 1978.

DAVID 78 Davies, D. W., and D. A. Bell, "The Protection of Data," National Physical Laboratory, Teddington ,England, NPL Report COM98, January 1978.

DIFFW 77 W. Diffie and M. Hellman "Privacy and Authentication: An Introduction to Cryptography," Proc. IEEE, Vol.67, No.3, 1979.

DHILL 89 Dhillon, B.S., "Bibliography of Literature On Safety Factors," Microelectronics Reliability, Vol. 29, No. 2, 1989 [TBV].

ENGEL 86 Engel, R., "Microprocessor-Controlled Fusion Splicer," Telcom Rep., Vol. 9, No. 5, September-October 1986.

FIPS 77 Data Encryption Standard, National Bureau of Standards, FIPS Publ. 46, January 1977.

GAITJ 77 Gait, J., "Validating the Correctness of Hardware Implementation of the NBS Data Encryption Standard," National Bureau of Standards, Spec. Publ. 500-20, November 1977.

HANDE 89 Handel, R., "Evolution of ISDN Towards Broadband ISDN," IEEE Network, January 1989.

IBM 3073 IBM Corporation, "Systems Network Architecture Technical Overview," GC30-3073, Armonk, NY.

IBM 3112 IBM Corporation, "Systems Network Architecture Format and Protocol Reference Manual: Architectural Logic," SC30-3112, Armonk, NY.

ISO/2 88 "Information Processing Systems - Open Systems Interconnection Reference Model", ISO 7498/2, 1988.

MAYBA 86 Maybaum, F.L. and Duffield, H.C., "Defense Data Network: An Overview," Proc. MILCOM 86, 1986.

MEIJE 88 Meijer, A., "Systems Network Architecture: a Tutorial," John Wiley & Sons, Inc., New York, 1988.

MEISP 76 Meissner, P., "Report of the 1976 Workshop on Estimation of Significant Advance in Computer Technology," National Bureau of Standards, August 30-31, 1976, NBS-IR 76-1189, December 1976.

MITRE 86 "ISO Transport Protocol Experiments," MTR-86W00002, The MITRE Corporation, January 1986. (Not available in the public domain)

RANDA 89 Randall, J.D., "Communications/Network Security," IEEE Eastern N.C. Section, Communications Society meeting, Raleigh NC, December 1989.

STALL 88 Stallings, W., "Data and Computer Communications", Macmillan, New York, 1988.

TANEN 81 Tanenbaum, A.S., "Computer Networks," Prentice-Hall, New York, 1981.

VINIO 90 Viniotis, I., Private conversation, January 25, 1990.

VISIO 88 "The 1995 DCS - A Step to Vision 21," Defense Communications Agency, April 1988.

WOOD 83 Wood H.W. and Cotton I.W., "Security in Computer Communications Systems," in W. Chou, Computer Communications Vol. I Principles, Prentice-Hall, New York, 1983.

MISSION

OF

ROME LABORATORY

Rome Laboratory plans and executes an interdisciplinary program in research, development, test, and technology transition in support of Air Force Command, Control, Communications and Intelligence (C^3I) activities for all Air Force platforms. It also executes selected acquisition programs in several areas of expertise. Technical and engineering support within areas of competence is provided to ESD Program Offices (POs) and other ESD elements to perform effective acquisition of C^3I systems. In addition, Rome Laboratory's technology supports other AFSC Product Divisions, the Air Force user community, and other DOD and non-DOD agencies. Rome Laboratory maintains technical competence and research programs in areas including, but not limited to, communications, command and control, battle management, intelligence information processing, computational sciences and software producibility, wide area surveillance/sensors, signal processing, solid state sciences, photonics, electromagnetic technology, superconductivity, and electronic reliability/maintainability and testability.