

Optical data communication: fundamentals and future directions

Casimer DeCusatis

IBM Corporation

522 South Road MS P343

Poughkeepsie, New York 12601

E-mail: casimer@vnet.ibm.com

Abstract. An overview of optical data communications is provided, beginning with a brief history and discussion of the unique requirements that distinguish this subfield from related areas such as telecommunications. Each of the major datacom standards is then discussed, including the physical layer specification, distances and data rates, fiber and connector types, data frame structures, and network considerations. These standards can be categorized by their prevailing applications, either storage [Enterprise System Connection (ESCON), Fiber Channel Connection (FICON), and Fibre Channel], coupling (Fibre Channel), or networking [Fiber Distributed Data Interface (FDDI), Gigabit Ethernet, and asynchronous transfer mode/synchronous optical network (ATM/SONET)]. We also present some emerging technologies and their applications, including parallel optical interconnects, plastic optical fiber, wavelength multiplexing, and free-space optical links. We conclude with some cost/performance trade-offs and predictions of future bandwidth trends. © 1998 Society of Photo-Optical Instrumentation Engineers. [S0091-3286(98)01412-3]

Subject terms: optical data communication; fiber; Enterprise System Connection; fiber channel; DeCusatis' law; Fiber Distributed Data Interface; asynchronous transfer mode; synchronous optical network.

Paper FIB-14 received Mar. 15, 1998; revised manuscript received Apr. 1, 1998; accepted for publication May 22, 1998.

1 Introduction

Fiber optic communication systems have seen widespread commercial use for some time. However, in the past 10 yr many new applications have emerged for optical fiber in computer and data communication networks, as compared with their applications in voice and telephone communications (telecom). Despite recent efforts toward convergence of voice, data, video, and image processing, which attempt to combine all elements of multimedia on a common hardware and software platform, the design of data communication (datacom) links remains a subdiscipline with its own distinguishing characteristics. In this paper, we describe the unique requirements of optical data communications and the applications that are driving this subfield. We then discuss the prevailing standards in the datacom industry, including Enterprise System Connection (ESCON), Fiber Distributed Data Interface (FDDI), Fibre Channel, Gigabit Ethernet and asynchronous transfer mode/synchronous optical network (ATM/SONET). The physical layer of each standard is described, including requirements for distance, data rates, physical media connectors, and other distinguishing features. Finally, we present several emerging technologies that are expected to have a significant effect on this field in coming years.

To provide a perspective on the scope of optical datacom, we can represent the field in a tree structure, as shown in Fig. 1. The roots of the optical datacom tree lie in the use of copper data links in computer systems and telecommunications. These systems were the basis for most of our communication needs until the development of fundamen-

tal optical fiber technologies, including light-emitting diodes and semiconductor lasers capable of high-speed modulation, solid state optical detectors and receiver circuitry, and low-loss glass fiber cables. It is interesting to note that each of these fundamentals developments became mature enough for commercial applications at about the same time, in the early 1970s; this greatly facilitated the adoption of fiber optic communications by the telecom industry, and later by the datacom industry. Today, the branches of this tree are the many applications of fiber optic datacom, which we discuss in detail:

1. ESCON and Serial Byte Connection (SBCON)
2. Fiber Channel Connection (FICON)
3. FDDI for local and wide area networks (LAN/WAN)
4. Fibre Channel (FC) standard
5. ATM/SONET
6. Gigabit Ethernet
7. emerging technologies, including
 - a. parallel optical interconnects
 - b. small-form factor fiber optic connectors
 - c. plastic optical fiber
 - d. free-space data links
 - e. wavelength-division multiplexing
 - f. other applications [serial High Performance Parallel Interface (HIPPI) and HIPPI 6400, multifiber trunking systems, and more].




Fig. 1 Optical data communications tree (after Ref. 1).

It is convenient to briefly review the history of fiber optic datacom by organizing the many computer applications into a few common platforms: personal computers, workstations, mainframes/large servers, and supercomputers. At the low end of functionality and processing power, personal computers (PCs) using modems to transfer data across public carrier telephone lines, which often contain a fiber optic backbone, were perhaps among the first to indirectly use optical fiber for data transfer. While this enabled the telecom carriers to multiplex many low-speed copper lines onto a single fiber, the presence of fiber in the system had minimal impact on the end user because the high-bandwidth telecommunications fiber was bottlenecked by much lower speed copper coax or twisted pair wiring at either end of the communication link. As the functionality of the PC increased and a second higher performance platform of workstations emerged, dedicated networks to share data became increasingly important. This began the proliferation of peer-to-peer LANs and the need emerged to scale into larger platforms including mainframes and supercomputers. This has led to the emergence of client-server-based networks, employing either circuit switching or packet switching. Large systems now offer native fiber optic attachments, which can reduce or eliminate the bottleneck at either end of a link. Further, high-speed fiber optic links

have enabled new computer architectures, including Parallel Sysplex for mainframes and clustered parallel supercomputers. Some have characterized the history of computing into three eras: the age of centralized computing using mainframes, the personal computer revolution, and the current age of network-centric computing in which a powerful datacom network between all types of platforms is the most important resource.

When discussing the role of fiber optics in data communication, it is natural to ask how this technology compares with alternatives such as copper data links, and which applications are best served by considering an optical solution. This is, of course, an ongoing and vigorous debate; rather than attempt to settle it here, we simply provide a framework for discussions in the rest of the paper. Traditionally, optical communication solutions have been successful in those areas that require higher data rates and longer unrepeated distances than can be provided by copper solutions; this is reflected in Table 1, which suggests that copper solutions are most cost effective for distances less than a few meters, including card-to-card, chip-to-chip, and intrachip applications.² Although there is a great deal of promising research being conducted with optics in these areas that may lead to innovative solutions, it is beyond the scope of this paper. Rather, we consider those areas in which fiber has already found commercial success, namely, distances ranging from tens of meters to kilometers. Table 1 also suggests a cost point in dollars per gigabit per second that is required for a successful solution at these distances today. This information is presented as a rough guide only, but it is interesting that the typical cost seems to scale by roughly the same order of magnitude as the distance.

A number of unique features distinguish the subfield of optical data communications as compared with telecommunications. First, datacom systems must maintain a much lower bit error rate (BER), defined as the number of errors per second in the communication link. This is so because the consequences of even a single bit error are potentially very serious; critical data such as bank account balances or medical records could be corrupted, and in extreme cases, the computer or even the entire network could be shut down. Even in noncritical applications, any errors impact the performance of the attached computer systems. In voice communications, the ultimate receiver is the human ear, and voice signals have a bandwidth of only about 4 kHz; transmission errors often manifest themselves as excessive background noise or static, such as encountered on a long distance international call or one from a mobile phone, and most users can ignore or tolerate this level of fidelity. Typical low-end datacom links operating at 10 to 15 Mbytes/s must support a BER of 10^{-12} or better (about 1 error every

Table 1 Copper versus optical data link solutions and distance and cost considerations;² costs are estimated market cost required to compete at a given distance.

Distance	Kilometers+	10–100 m	1–10 m	0.1–1 m	<10 cm
Cost per gigabit per second	\$1000	\$100	\$10	\$1	\$0.1
Copper	Coax/repeated	Coax/twisted pair	Coax TP	Backplane	Chip-on-board
Optics	Single mode	Multimode	Multimode	Waveguide	Free space

hour for a 200 Mbits/s channel), as compared with about 10^{-9} for telecom systems (about 1 error every 5 s). Some datacom applications must support a BER as low as 10^{-15} ; even testing such a link becomes challenging when an error occurs only once every 57 days (higher data rates are helping to reduce the time required to test very low BERs).

Another unique requirement of datacom systems is eye safety versus distance trade-offs. Before the introduction of fiber optics in the data center, parallel copper coax was limited to a maximum distance of perhaps a few hundred meters by noise considerations. The introduction of optical fiber has driven a trend toward higher data rates and longer distances; new applications such as distributed computing and remote backup of data for disaster recovery have pushed fiber optic data links to relatively long unrepeated distances (10 to 20 km). At the same time that these applications have been expanding the data center, the computer networks are often serviced by the users themselves or by a staff that is not trained in optical safety or handling of fiber cables. In contrast, most telephone users would never think of crawling into a manhole and reconnecting cables under the street if their phone service failed to work; the telecom providers completely control the equipment installed in their network, and ensure that it is serviced by only trained personnel. This means that all datacom links must meet international class 1 laser safety standards,* which limit the optical transmitter output power. For datacom links, this implies a trade-off between the requirements for long-distance, high-fidelity transmission and limitations on the amount of optical power that can be safely launched into the fiber. This makes optical link budget analysis and link planning even more important for the data link designer (some standards attempt to circumvent this limitation using a technique called open fiber control, which we discuss later). Because most datacom end users are not trained or experienced with fiber optics, datacom equipment must be designed to withstand large numbers of connections and disconnections, high pulling forces on fiber cables, and other user-inflicted stress. The equipment should be easy to use and maintain; for example, connectors are keyed to prevent misplugging, and receptacles can be designed to minimize the necessity for frequent cleaning. Fiber optic connectors must exhibit the same degree of reliability as common two-prong electrical wall outlets; they must plug into the receptacle and work properly the first time, every time. Despite the requirement for rugged, high-reliability components, the datacom market is very sensitive to the cost of optoelectronics, because there are a much larger number of transceivers per kilometer of fiber than in telecom applications. Telecom links may stretch hundreds of kilometers, even around the world, by using many low-jitter repeaters and signal regenerators, and make more extensive use of multiplexing techniques to effectively utilize the available bandwidth of the fiber. Typical datacom links have not used multiplexing as extensively (until very recently), and use at most only a few repeaters in a long link.

*U.S. laser safety standards are regulated by the Department of Health and Human Services (DHHS), the Occupational Safety and Health Administration (OSHA), the Food and Drug Administration (FDA) Code of Radiological Health (CDRH) 21 Code of Federal Regulations (CFR) subchapter J; the relevant standards are given in Ref. 3. Elsewhere in the world, the relevant standard is as given in Ref. 4.

2 Datacom Standards

As implied earlier, the datacom network is a multivendor environment in which the end user assembles equipment from many different sources into a single network (traditionally, telecom providers have not faced this restriction, although this situation is changing somewhat in the United States). This has placed an increasing emphasis on interoperability through standardization. The goal of standardization is to document a minimum set of requirements (to encourage innovation and multiple solutions to the same problem) that achieve a maximum amount of interoperability. It is important to distinguish a standard from a specification; while multiple products may claim compliance with a given standard, they may have very different component or functional specifications. Standards may be set by the *de facto* agreement of many different companies, but are more formally administered by professional organizations such as the IEEE, ISO, and ANSI. Because datacom standards are evolving rapidly, the latest revisions should be consulted frequently; with the advent of the World Wide Web, many resources are now available online related to fiber optic datacom (see Table 2). These specifications are consistent with the applicable European and U.S. industry standards for optical fiber cables.⁵⁻⁷ In the following sections, we review the prevailing international datacom standards for fiber optics; ESCON, FDDI, FC, Gigabit Ethernet, and ATM over SONET.

We can categorize the applications that are driving optical datacom into three areas: input/output (I/O) device attachment for data storage, coupling, and networking. Storage applications typically employ ESCON, FICON, or FC; coupling applications utilize FC or some variant; and networking applications tend to use FDDI, Gigabit Ethernet, or ATM/SONET. Large servers require dedicated links to interconnect with storage devices, such as magnetic disk and tape drives, printers and other direct access storage device (DASD) equipment. This has been the driving application for ESCON, FICON, and FC links, and has addressed the need to backup data in real time at a remote location for disaster recovery. It has also played a role in storage-intensive applications such as data mining. Mainframes and large servers contain an estimated 70% of the world's data. The recent trend toward clustered, parallel processors to enhance performance has driven the second requirement for high-bandwidth coupling links. For example, large water-cooled mainframe computers are being replaced by smaller, air-cooled complementary metal-oxide semiconductor (CMOS)-based servers that couple together many processors (central processing units, or CPUs) in parallel. Currently, leading-edge applications are using FC-type data links for this application, or copper if the distances are very short. High-performance applications have traditionally been the area in which new technologies are first introduced, however, and these coupling links may be supplanted by other technologies such as parallel optical interconnects, as we discuss later (indeed, the area of intramachine communications may develop into a fourth application driver, encompassing optical data buses, clocks, and other links within a single machine). Finally, the third and perhaps the most visible application has been in networking, including LANs, WANs, and metropolitan area networks (MANs). This has been the domain of FDDI, Gigabit

Table 2 Optical datacom information on the World Wide Web.

ATM Forum	http://www.atmforum.com
Desktop ATM25 Alliance	http://222.atm25.com
Fibre Channel Association	http://www.amdahl.com/ext/CARP/FCA/FCA.html
ANSI FDDI X3T12 Home Page	http://www.sholeh.nswc.navy.mil/x3t12
Gigabit Ethernet Alliance	http://www.gigabit-ethernet.org
HIPPI standards	http://www.esscom.com/hnf
Infrared Data Association	http://www.irda.com
American National Standards Institute (ANSI)	http://www.ansi.org
ANSI T1 Committee	http://www.t1.org
International Telecommunications Union (ITU, formerly CCITT)	http://info.itu.ch
Optoelectronic Industry Association (OIDA)	http://www.oida.com
Network Computing Reference Profile	http://www.nc.ihost.com
IBM Corporation (ESCON-related information)	http://www.ibm.com
IEEE Standards	http://www.ieee.org
Optical Society of America (OSA) [also Technical Home for Photonics and homepage for Institute of Optical Data Communication (IODC)]	http://www.osa.org

Note that this list is not comprehensive.

Ethernet, and SONET-based ATM as well as other non-fiber-based network standards including ethernet and token ring.

2.1 ESCON/SBCON

From the introduction of the IBM System/360 mainframe in 1964 until the first use of ESCON 26 yr later, only incremental upgrades had been made to the mainframe computer I/O interface. The basic interconnection technology for attachment to DASD such as magnetic tape or disk drives and storage control units consisted of a parallel copper interface with a peak data rate of 4.5 Mbytes/s and a maximum distance of 144 m (200 ft). Indeed, this distance limitation is one reason why most mainframe complexes were contained in a single room or “glass house” with storage devices and peripherals clustered next to the mainframe CPU. This technology remained unchanged until the introduction of the ESCON[†] architecture on the IBM System/390 family in 1990 as an alternative high speed I/O channel attachment.^{8–10} The ESCON interface became the *de facto* industry standard for large system attachments; the specifications were subsequently published and adopted by other original equipment manufacturers^{11–13} (OEMs), and in 1996 they were adopted by the ANSI X3T1 committee as the SBCON standard;¹⁴ the name change was required because ESCON is a trademark of IBM. The architecture and protocols have been adopted by a number of mainframe and large server vendors, and several companies currently manufacture ESCON-compatible fiber optic transceivers. Recently, IBM has authorized use of the ESCON trademark by selected fiber optic cable vendors whose products are demonstrated to meet the quality, functionality, and performance specifications equivalent to components obtained directly from IBM.

[†]ESCON System/390 and S/390 are registered trademarks of IBM Corporation.

The ESCON/SBCON channel is a bidirectional, point-to-point 1300-nm fiber optic data link with a maximum data rate of 17 Mbytes/s (200 Mbits/s minus overhead). The physical layer specifications are given in Tables 3–7. ESCON supports a maximum unrepeated distance of 3 km using 62.5- μ m multimode fiber and LED transmitters with an 8-dB link budget, or a maximum unrepeated distance of 20 km using single mode fiber and laser transmitters with a 14-dB link budget. The laser channels are also known as the ESCON extended distance feature (XDF). Physical connection is provided by an ESCON duplex connector (illustrated in Fig. 2); the connector features a unique spring-loaded retractable dust cover, which protects the ferrules when the connector is not in use and retracts on insertion into an ESCON receptacle. The connector was designed specifically to provide a low profile; ease of use; consistently low loss; and a rugged, repeatable connection. Initially, the ESCON duplex connector was offered in both multimode and single-mode versions, color coded with black connectors and orange fiber jacket for multimode and gray connectors and yellow fiber jacket for single mode. Recently, the single-mode ESCON links have adopted the SC duplex connector as standardized by FC (Fig. 3).

The specifications in Tables 3–7 assume typical values of the link budget; by trading off elements of the link design, customized ESCON links have been installed that can

Table 3 ESCON allowed maximum link loss (at 1300 nm wavelength).¹¹

Maximum Length (km)	Maximum Loss (dB)	Fiber Type
2.0	8.0	62.5 μ m, 500 MHz km
2.0	8.0	50.0 μ m, 800 MHz km
2.0 to 3.0	8.0	62.5 μ m, 800 MHz km
20.0	14.0	9 μ m (single mode)

Table 4 Multimode ESCON output interface specification.¹¹

Parameter	Minimum	Maximum	Unit
Average power ^{1,2}	-20.5	-15.0	dBm
Center wavelength	1280	1380	nm
Spectral width (FWHM)		175.0	nm
Rise time (t_r) (20 to 80%) ^{1,3}		1.7	ns
Fall time (t_f) (80 to 20%) ^{1,3}		1.7	ns
Eye window ¹	3.4		ns
Optical output jitter ⁴		0.8	ns
Extinction ratio ^{1,5}	8		dB
t_r and t_f at optical path output ^{3,6}		2.8	ns

¹Based on any valid 8B/10B code. The length of jumper cable between the output interface and the instrumentation is 3 m.

²The output power shall be greater than -29 dBm through a worst-case link, as specified in Table 3. Higher order mode loss (HOML) is the difference in link loss measured using the device transmitter as compared to the loss measured with a source conditioned to achieve an equilibrium mode distribution in the fiber. The transmitter shall compensate for any excess HOML occurring in the link (e.g., HOML in excess of 1 dB for a 62.5- μ m link).

³The minimum frequency response bandwidth range of the optical waveform detector shall be 100 kHz to 1 GHz.

⁴The optical output jitter includes both deterministic and random jitter. It is defined as the peak-to-peak time-histogram oscilloscope value (minimum of 3000 samples) using a 2⁷-1 pseudorandom pattern or worst-case 8B/10B code pattern. The transmitter output light is coupled to a *p-i-n* photodiode optical-to-electrical (O/E) converter (e.g., Tektronix P6703A or equivalent) via a 3-m cable and jitter measured with a digital sampling oscilloscope.

⁵Measurement shall be made with a dc-coupled optical waveform detector that has a minimum bandwidth of 600 MHz and whose gain flatness and linearity over the range of optical power being measured provide an accurate measurement of the high and low optical power levels.

⁶The maximum rise or fall time (from, e.g., chromatic, modal dispersion, etc.) at the output of a worst-case link as specified in Table 3. The 0 and 100% levels are set where the optical signal has at least 10 ns to settle. The spectral width of the transmitter shall be controlled to meet this specification.

operate at 10 to 20% or more beyond these distance limitations.⁹ Using repeaters or switches, an ESCON link can be extended up to three to five times these distances; however, performance of the attached devices typically falls off quickly at longer distances due to the longer round-trip latency of the link, making this approach suitable only for applications that can tolerate a lower effective throughput, such as remote backup of data for disaster recovery. ESCON devices and CPUs may communicate directly through a channel-to-channel attachment, but more commonly attach to a central nonblocking dynamic cross-point switch. The resulting network topology is similar to a star-wired ring, which provides both efficient bandwidth utilization and reduced cabling requirements. This structure also makes it easy to add or reconfigure equipment while maintaining the ability to isolate and partition subdivisions of the network. The switching function is provided by an ESCON director, a nonblocking circuit switch. Although ESCON uses 8 B/10 B encoded data, it is not a packet-switching network; instead, the data frame header includes a request for connection that is established by the director for the duration of the data transfer. An ESCON data frame

Table 5 Multimode ESCON input interface specification.¹¹

Parameter	Minimum	Maximum	Unit
Sensitivity ^{1,2}		-29.0	dBm
Saturation level ¹	-14.0		dBm
Acquisition time ³		100	ns
Loss of light (LOL) threshold ⁴	-45	-36	dB
LOL hysteresis ^{4,5}	0.5		dB
Reaction time for LOL state change	3	500	μ s

¹Based on any valid 8B/10B code pattern measured at, or extrapolated to, 10⁻¹⁵ BER measured at center of eye. This specification shall be met with worst case conditions as specified in Table 4 for the output interface and Table 3 for the fiber optic link. This value allows for a 0.5-dB retiming penalty.

²A minimum receiver output eye opening of 1.4 ns at 10⁻¹² should be achieved with a penalty not exceeding 1 dB.

³The acquisition time is the time to reach synchronization after the removal of the condition that caused the loss of synchronization. The pattern sent for synchronization is either the idle character or an alternating sequence of idle and data characters.

⁴In the direction of decreasing power; if power > -36 dBm, the LOL state is inactive; if power < -45 dBm, LOL state is active. In the direction of increasing power, if power < -44.5 dBm, the LOL state is active; if power > -35.5 dBm, LOL state is inactive.

⁵Required to avoid random transitions between LOL being active and inactive when input power is near threshold level.

includes a header, a payload of up to 1028 bytes of data, and a trailer. The header consists of a two-character start-of-frame delimiter, a 2-byte destination address, a 2-byte source address, and 1 byte of link control information. The trailer is a 2-byte cyclic redundancy check (CRC) for errors and a three-character end-of-frame delimiter.¹⁰ ESCON uses a dc-balanced 8 B/10 B coding scheme developed¹⁴ by IBM.

2.2 FDDI

The FDDI was among the first open networking standards to specify optical fiber. It was an outgrowth of the ANSI X3T9.5 committee proposal in 1982 for a high-speed token-passing ring as a back-end interface for storage devices. While interest in this application waned because of the development of higher speed (800 Mbits/s) proprietary interfaces, FDDI found new applications as the backbone for LANs; compatible hardware is available from many different companies. The FDDI standard was approved in 1992 as ISO standards IS 9314/1-2 and DIS 9314-3; it follows the architectural concepts of IEEE standard 802 (although it is controlled by ANSI, not IEEE, and therefore has a different numbering sequence) and is among the family of standards including token ring and ethernet, which are compatible with a common IEEE 802.2 interface.^{16,17} FDDI is a family of four specifications, namely the physical layer (PHY), physical-media dependence (PMD), media access control (MAC), and station management (SMT). These four specifications correspond to sublayers of the data link and physical layer of the open systems interconnect (OSI) reference model; as before, we will concentrate on the physical layer implementation.

Table 6 Single-mode ESCON output interface specification.¹¹

Parameter	Minimum	Maximum	Unit
Average ¹ power into single-mode fiber (SMF)	-8.0	-3.0	dBm
Center wavelength ¹	1261	1360	nm
Spectral width [root mean square (rms)]		7.7	nm
Rise time (t_r) (20 to 80%) ^{1,2}		1.5	ns
Fall time (t_f) (80 to 20%) ^{1,2}		1.5	ns
Eye window ¹	3.5		ns
Optical output jitter ³		0.8	ns
Extinction ratio ^{1,4}	8		dB
Relative intensity noise (RIN_{12}) ⁵		-112.0	dB/Hz
t_r and t_f at optical path output ^{3,6}		2.8	ns

¹Based on any valid 8B/10B code pattern. The measurement is made using a 4-m single-mode jumper cable and includes only the power in the lowest order fundamental mode.

²The minimum frequency response bandwidth range of the optical waveform detector shall be 100 kHz to 1 GHz.

³The optical output jitter includes both deterministic and random jitter. It is defined as the peak-to-peak time-histogram oscilloscope value (minimum of 3000 samples) using a 2⁷-1 pseudorandom pattern or worst-case 8B/10B code pattern. The transmitter output light is coupled to a *p-i-n* photodiode O/E converter (e.g., Tektronix P6703A or equivalent) via a 3 m cable and jitter measured with a digital-sampling oscilloscope.

⁴Measurement shall be made with a dc-coupled optical waveform detector that has a minimum bandwidth of 600 MHz and whose gain flatness and linearity over the range of optical power being measured provide an accurate measurement of the high and low optical power levels.

⁵The relative intensity noise is measured with a 12-dB optical return loss into the output interface.

⁶The maximum degradation in input interface sensitivity (from e.g., jitter, mode hopping, intersymbol interference) that can occur by using a worst case link as specified in Table 3. The spectral width of the transmitter shall be controlled to meet this specification.

The FDDI network is a 100 Mbits/s token-passing ring, with dual counterrotating rings for fault tolerance. The dual rings are independent fiber optic cables; the primary ring is used for data transmission, and the secondary ring is a backup in case a node or link on the primary ring fails. Bypass switches are also supported to reroute traffic around a damaged area of the network and prevent the ring from fragmenting in case of multiple node failures. The actual data rate is 125 Mbits/s, but this is reduced to an effective data rate of 100 Mbits/s by using a 4 B/5 B coding scheme.¹⁶ This high speed enables FDDI to be used as a backbone to encapsulate lower speed 4, 10, and 16 Mbits/s LAN protocols; existing ethernet, token ring, or other LANs can be linked to an FDDI network via a bridge or router. Although FDDI data flows in a logical ring, a more typical physical layout is a star configuration with all nodes connected to a central hub or concentrator rather than to the backbone itself. There are two types of FDDI nodes, either dual attach (connected to both rings) or single attach; a network supports up to 500 dual-attached nodes, 1000 single-attached nodes, or an equivalent mix of the two types. FDDI specifies 1300-nm LED transmitters operating over 62.5- μ m multimode fiber as the reference media, al-

Table 7 Single-mode ESCON input interface specification.¹¹

Parameter	Minimum	Maximum	Unit
Sensitivity ^{1,2}		-28.0	dBm
Saturation level ¹	-3.0		dBm
Return loss ³	12.5		dB
Acquisition time ³		100	ns
LOL threshold ⁴	-45	-32	dB
LOL hysteresis ⁵	1.5		dB
Reaction time for LOL state change	0.25	5000	μ s

¹Based on any valid 8B/10B code pattern measured at, or extrapolated to, 10⁻¹⁵ BER measured at the center of the eye. This specification shall be met with worst-case conditions as specified in Table 6 for the output interface and Table 3 for the fiber optic link. This value allows for a 0.5-dB retiming penalty.

²A minimum receiver output eye opening of 1.4 ns at 10⁻¹² should be achieved with a penalty not exceeding 1 dB.


³The measurement is made using a 4-m single-mode jumper cable and includes only the power in the lowest order fundamental mode.

⁴The acquisition time is the time to reach synchronization after the removal of the condition that caused the loss of synchronization. The pattern sent for synchronization is either the idle character or an alternating sequence of idle and data characters.

⁵Required to avoid random transitions between LOL being active and inactive when input power is near threshold level.

though the standard also provides for the attachment of 50-, 100-, 140-, and 85- μ m fiber (see Table 8). Using 62.5- μ m fiber, a maximum distance of 2 km between nodes is supported with an 11-dB link budget; since each node acts like a repeater with its own phase lock loop to prevent jitter accumulation, the entire FDDI ring can be as large as 100 km. However, an FDDI link can fail due to either excessive attenuation or dispersion; for example, insertion of a bypass switch increases the link length and may cause dispersion errors even if the loss budget is within specifications. For most other applications, this does not occur because the dispersion penalty is included in the link budget calculations or the receiver sensitivity measurements. The physical interface is provided by a special media interface connector (MIC), illustrated in Fig. 4. The connector has a set of three color-coded keys that are interchangeable, depending on the type of network connection; Table 9 shows the allowed attachments and corresponding keys. This is intended to prevent installation errors and assist in cable management.

An FDDI data frame is variable length, and contains up to 4500 8-bit bytes, or octets, including a preamble, start of frame, frame control, destination address, data payload, CRC error check, and frame status/end of frame. Each node

**Fig. 2** ESCON duplex fiber optic connector.




Fig. 3 Single mode SC duplex fiber optic connector, per ANSI FC Standard specifications, with one narrow key and one wide key. Multimode SC duplex connectors use two wide keys.




Fig. 4 FDDI duplex fiber optic connector.

has a MAC sublayer that reviews all the data frames looking for its own destination address. When it finds a packet destined for its node, that frame is copied into local memory, a copy bit is turned on in the packet, and it is then sent on to the next node on the ring. When the packet returns to the station that originally sent it, the originator assumes that the packet was received if the copy bit is on; the originator will then delete the packet from the ring. As in the IEEE 802.5 token ring protocol, a special type of packet called a token circulates in one direction around the ring and a node can transmit data only when it holds the token. Each node observes a token retention time limit, and also keeps track of the elapsed time since it last received the token; nodes may be given the token in equal turns, or they can be given priority by receiving it more often or holding it longer after they receive it. This enables devices having different data requirements to be served appropriately.

Because of the flexibility built into the FDDI standard, many changes to the base standard have been proposed to enable interoperability with other standards, reduce costs, or extend FDDI into the MANs or WANs. These include a single-mode PMD layer for channel extensions up to 20 to 50 km. An alternative PMD provides for FDDI transmission over copper wire, either shielded or unshielded twisted pairs; this is known as copper distributed data interface (CDDI). A new PMD is also being developed to adapt FDDI data packets for transfer over a SONET link by stuffing approximately 30 Mbits/s into each frame to make up for the data rate mismatch (we discuss SONET as an ATM physical layer in a later section). An enhancement called FDDI-II uses time-division multiplexing to divide the

bandwidth between voice and data; it would accommodate isochronous, circuit-switched traffic as well as existing packet traffic. Recently, an option known as low-cost (LC) FDDI has been adopted. This specification uses the more common SC duplex connector (Fig. 2) instead of the expensive MIC connectors, and a lower cost transceiver with a nine-pin footprint similar to the single-mode ESCON parts (Fig. 3). This is an attempt to provide lower cost transceivers and cables for the FDDI market. Most recent installations, comprising about 10% of the total LAN types including ethernet and token ring, use the LC FDDI instead of the MIC interface. Applications include the recent availability of LC FDDI adapters for large “super-servers” on the World Wide Web or private corporate intranets.¹⁸

2.3 FC Standard

Development of the ANSI FC Standard (FCS) began in 1988 under the X3T9.3 Working Group, as an outgrowth of the Intelligent Physical Protocol Enhanced Physical Project. The motivation for this work was to develop a scalable standard for the attachment of both networking and I/O devices using the same drivers, ports, and adapters over a single channel at the highest speeds currently achievable. The standard applies to both copper and fiber optic media, and uses the English spelling “fibre” to denote both types of physical layers. In an effort to simplify equipment design, FC provides the means for a large number of existing upper-level protocols (ULPs), such as internet protocol (IP), Small Computer Interface (SCI), and HIPPI, to operate over a variety of physical media. Different ULPs are mapped to FC constructs, encapsulated in FC frames, and transported across a network; this process remains transparent to the attached devices. The standard

Table 8 FDDI PMD specification.¹⁷

Transmitter output (min)	−20.0 dBm
Transmitter type	LED
Transmitter center wavelength	1300 nm
Link loss, without switch	9.0 dB
Maximum Loss, FDDI bypass switch	2.0 dB
Maximum distance, unrepeated	2.0 km
Fiber type	62.5/125 μm multimode preferred, options include 50/125, 85/125, and 100/140 μm ,
Fiber bandwidth	500 MHz km (recommended at 62.5 μm)
BER (min)	10^{-9}
Data rate	125 Mbits/s (reduced to 100 Mbits/s with 4B/5B coding)

Table 9 FDDI MIC connector keys color coding.¹⁷

There are four types of keys/labels in the FDDI MIC standard:	
Port A (A)=red key/label	
Port B (B)=blue key/label	
Master (M)=green key/label	
Slave (S)=no key/white label	
FDDI Connection Type	Key Type
Workstation to wall	S/S
Distribution panel to concentrator M port	M/M
Distribution panel to concentrator A port	A/A
Distribution panel to concentrator B port	B/B
Concentrator A to concentrator M port	A/M
Concentrator B to concentrator M port	B/M
Concentrator 1A to concentrator 2B port	A/B
Concentrator 1B to concentrator 2A port	B/A
Workstation to concentrator M port	S/M

consists of five hierarchical layers,^{19–21} namely a physical layer; an encode/decode layer, which has adopted the dc-balanced 8B/10B code,¹⁵ a framing protocol layer; a common services layer (at this time, no functions have been formally defined for this layer); and a protocol mapping

layer to encapsulate ULPs into FC. The second layer defines the FC data frame; frame size depends on the implementation, and is variable up to 2148 bytes long. Each frame consists of a 4-byte start of frame delimiter, a 24-byte header, a 2112-byte payload containing from 0 to 64 bytes of optional headers and 0 to 2048 bytes of data, a 4-byte CRC, and a 4-byte end-of-frame delimiter. In October 1994, the FC physical and signaling interface standard FC-PH was approved as ANSI standard X3.230-1994.

Logically, FC is a bidirectional point-to-point serial data link. Physically, there are many different media options (see Table 10) and three basic network topologies. The simplest, default topology is a point-to-point direct link between two devices, such as a CPU and a device controller. The second, FC Arbitrated Loop (FC-AL), connects between 2 and 126 devices in a loop configuration. Hubs or switches are not required, and there is no dedicated loop controller; all nodes on the loop share the bandwidth and arbitrate for temporary control of the loop at any given time. Each node has equal opportunity to gain control of the loop and establish a communications path; once the node relinquishes control, a fairness algorithm ensures that the same node cannot win control of the loop again until all other nodes have had a turn. As networks become larger, they may grow into the third topology, an interconnected switchable network or fabric in which all network management functions are taken over by a switching point, rather

Table 10 FC standard physical layer.¹⁹

Media Type	Data Rate (Mbytes/s)	Maximum Distance	Signaling Rate (Mbaud)	Transmitter
SMF	100	10 km	1062.5	LW laser
	50	10 km	1062.5	LW laser
	25	10 km	1062.5	LW laser
50- μ m multimode fiber	100	500 m	1062.5	SW laser
	50	1 km	531.25	SW laser
	25	2 km	265.625	SW laser
	12.5	10 km	132.8125	LW LED
62.5- μ m multimode fiber	100	300 m	1062.5	SW laser
	50	600 m	531.25	SW laser
	25	1 km	265.625	LW LED
	12.5	2 km	132.8125	LW LED
105- Ω type 1 shielded twisted pair electrical	25	50 m	265.125	ECL
	12.5	100 m	132.8125	ECL
75 Ω mini coax	100	10 m	1062.5	ECL
	50	20 m	531.25	ECL
	25	30 m	265.625	ECL
	12.5	40 m	132.8125	ECL
75 Ω video coax	100	25 m	1062.5	ECL
	50	50 m	531.25	ECL
	25	75 m	265.625	ECL
	12.5	100 m	132.8125	ECL
150 Ω twinax or STP	100	30 m	1062.5	ECL
	50	60 m	531.25	ECL
	25	100 m	265.625	ECL

Here LW is long wavelength, SW is short wavelength, and ECL is emitter-coupled logic.




Fig. 5 ANSI FC Standard-compatible optical transceiver module, also known as Optical Link Module (OLM) or Gigabit Link Module (GLM).

than each node. An analogy for a switched fabric is the telephone network; users specify an address (phone number) for a device with which they want to communicate, and the network provides them with an interconnection path. In theory, there is no limit to the number of nodes in a fabric; practically, there are only about 16 million unique addresses. FC also defines three classes of connection service, which offer options such as guaranteed delivery of messages in the order they were sent and acknowledgment of received messages.

As shown in Table 10, FC provides for both single-mode and multimode fiber optic data links using long wavelength (LW) (1300-nm) lasers and LEDs as well as short wavelength (SW) (780 to 850-nm) lasers. The physical connection is provided by an SC duplex connector defined in the standard (see Fig. 3), which is keyed to prevent misplugging of a multimode cable into a single-mode receptacle. This connector design has since been adopted by other standards, including ATM, LC FDDI, Gigabit Ethernet, and single-mode ESCON. There have been efforts from the hub and switch providers to develop a lower cost optical connector with a smaller form factor than the SC duplex (see Section 3.3). Recently, the subscriber group (SG) connector was adopted as an alternative physical media for FC links. A popular vendor enhancement to FC is the use of a pluggable 80-pin electrical interface, rather than requiring the transceiver to be soldered to the host card (Fig. 5). The requirement for international class 1 laser safety is addressed using open fiber control (OFC) on some types of multimode links with SW lasers. This technique automatically senses when a full duplex link is interrupted, and turns off the laser transmitters on both ends to preserve laser safety. The lasers then transmit low-duty-cycle optical pulses until the link is reestablished; a handshake sequence then automatically reactivates the transmitters. Today, many products offer FC attachments for both storage and networking applications; in particular, leading workstation vendors including IBM, Sun, and Hewlett Packard have formed the Common Optics Group to develop a common standard for workstation communications based on FC. Many workstation products use either the 266 Mbits/s or 1.0625 Gbits/s fiber optic attachment options from Table 10. Other applications are also possible; in 1994, IBM announced the use of links based on the FC physical layer as part of its Parallel Sysplex[‡] architecture for large servers.^{22,23} These IBM InterSystem Channels* (ISCs) or coupling links are used to couple processors in parallel to

increase performance. Two options were originally available, an SW, 531 Mbits/s link on 50- μ m multimode fiber with a maximum distance of 1 km (which has recently been discontinued) and an LW, 1.0625 Gbits/s link on SMF with a maximum distance of 10 km (recently increased from 3 km). The gigabit ISC links deviate slightly from the FC standard; although OFC is not defined for LW links, they use OFC timings corresponding to the 266 Mbits/s SW links to support a maximum distance 1 km longer than the standard. Both options have been used on five generations of large CMOS-based servers to date, including the Generation 5 (G5)[†] servers announced in May 1998.

2.4 FICON

Introduced in May 1998 on the IBM G5 servers, FICON[‡] is IBM's implementation of FC; it is a high-speed I/O link with provisions to facilitate limited use of the ESCON physical cable plant. All FICON links use LW (1300-nm) laser sources. The FICON physical layer supports data rates of 1 Gbits/s on SMF up to 10 km with a 7-dB link budget, compatible with FC physical layer specifications. In addition, FICON enables the use of both 50- and 62.5- μ m multimode fiber with a mode conditioning patch cable, similar to the emerging gigabit ethernet standard, at distances up to 550 m with a 5-dB link budget. A FICON adapter card will be available for ESCON directors (model 5) which will enable up to 8 ESCON channels to be time division multiplexed on a FICON link, and FICON directors are expected to be available in the near future. Several other companies, including Hitachi Data Systems, have announced their intention to support FICON interfaces as an *ad hoc* industry standard; IBM has announced its intention to seek approval of FICON as an ANSI standard. Further documentation of FICON is available from IBM.²³

2.5 ATM/SONET

The development of ATM networks has received a great deal of recent attention. Developed by the ATM Forum, this protocol has promised to provide a common transport media for voice, data, video, and other types of multimedia. To date, it has been deployed mostly as a data communications network, which is how we treat it in this section. Actually, ATM is a high-level protocol that can run over many different physical layers including copper; part of ATM's promise to merge voice and data traffic on a single network comes from plans to run ATM over the SONET transmission hierarchy developed for the telecommunications industry. In this section, we begin with a description of SONET followed by an overview of ATM; the relationship between an "asynchronous" protocol running on top of a "synchronous" physical layer will become apparent.

SONET is really a family of standards defined by ANSI T1.105-1988 and T1.106-1988, as well as by several Consultative Committee for International Telegraphy and Telephony (CCITT) recommendations.²⁴⁻²⁷ Several different data rates are defined as multiples of 51.84 Mbits/s, known

*InterSystem Channel and ISC are registered trademarks of IBM Corporation.

[†]Generation 5 and G5 are registered trademarks of IBM Corporation.

[‡]FICON is a registered trademark of IBM Corporation.

[‡]Parallel Sysplex is a registered trademark of IBM Corporation.

Table 11 SONET physical layer specifications.²⁸

Parameter	Line Rate								Unit
	OC-1	OC-3	OC-9	OC-12	OC-18	OC-24	OC-36	OC-48	
Transmitter type	MLM/LED	MLM/LED	MLM/LED	MLM/LED	MLM	MLM	MLM	MLM	
$\lambda_{T_{\min}}$ to $\lambda_{T_{\max}}$	1260–1360	1260–1360	1260–1360	1260–1360	1260–1360	1260–1360	1260–1360	1265–1360	nm
$\Delta\lambda_{\text{rms}}$	80	40/80	19/45	14.5/35	9.5	7	4.8	4	nm
$P_{T_{\max}}$	−14	−8	−8	−8	−8	−5	−3	−3	dBm
$P_{T_{\min}}$	−23	−15	−15	−15	−15	−12	−10	−10	dBm
$r_{e \text{ min}}$	8.2	8.2	8.2	8.2	8.2	8.2	8.2	8.2	dB
Optical path									
System ORL _{max}	NA	NA	NA	NA	20	20	24	24	dB
Attenuation	0–7	0–7	0–7	0–7	0–7	0–7	0–7	0–7	dB
Max. reflectance between <i>S</i> and <i>R</i>	NA	NA	NA	NA	−25	−25	−27	−27	dB
Receiver									
$P_{R \text{ max}}$	−14	−8	−8	−8	−8	−5	−3	−3	dBm
$P_{R \text{ min}}$	−31	−23	−23	−23	−23	−20	−18	−18	dBm

NA stands for not applicable. MLM stands for multi-longitudinal laser. $\lambda_{T_{\min}}$ and $\lambda_{T_{\max}}$ are minimum and maximum values of transmitter center wavelength. $\Delta\lambda_{\text{rms}}$ is the transmitter root mean square spectral width. $P_{T_{\max}}$ and $P_{T_{\min}}$ are the maximum and minimum transmitter optical output powers. r_e is the extinction ratio. ORL stands for optical return loss. *S* and *R* stand for source and receiver, respectively. $P_{R_{\max}}$ and $P_{R_{\min}}$ are the maximum and minimum receiver input optical powers.

as optical carrier level 1 (OC-1). The numerical part of the OC-level designation indicates a multiple of this fundamental data rate, thus, 155 Mbits/s is called OC-3 (see Table 11). The standard provides for seven incremental data rates, OC-3, OC-9, OC-12, OC-18, OC-24, OC-36, and OC-48 (2.48832 Gbits/s). Both single-mode links with laser sources and multimode links with LED sources are defined for OC-1 through OC-12; only single-mode laser links are defined for OC-18 and beyond. In addition to the specifications in Table 11, SONET also contains provisions to carry sub-OC-1 data rates, called virtual tributaries, which support telecom data rates including digital signal level 1 (DS-1) (1.544 Mbits/s), DS-2 (6.312 Mbits/s), and 3.152 Mbits/s (DS1C). The basic SONET data frame is an array of nine rows with 90 bytes per row, known as a synchronous-transport-signal level 1 (STS-1) frame. In an OC-1 system, an STS-1 frame is transmitted once every 125 μs (810 bytes per 125 μs yields 51.84 Mbits/s). The first three columns provide overhead functions such as identification, framing, error checking, and a pointer which identifies the start of the 87-byte data payload. The payload floats in the STS-1 frame, and may be split across two consecutive frames. Higher speeds can be obtained either by concatenation of *N* frames into an STS-Nc frame (the *c* stands for concatenated) or by byte interleaved multiplexing of *N* frames into an STS-N frame.^{9,28}

ATM technology has been standardized by the concurrence of the ATM Forum; the most recent version is available online (Table 2). It incorporates elements of both circuit and packet switching. All data is broken down into a 53-byte cell, which may be viewed as a short, fixed-length packet. Five bytes make up the header, providing a 48-byte payload. The header information contains routing information (cell addresses) in the form of virtual path and channel identifiers, a field to identify the payload type, an error

check on the header information, and other flow control information. Cells are generated asynchronously; as the data source provides enough information to fill a cell, it is placed in the next available cell slot. There is no fixed relationship between the cells and a master clock, as in conventional time-division multiplexing schemes; the flow of cells is driven by the bandwidth needs of the source. ATM provides bandwidth on demand; for example, in a client-server application the data may come in bursts, and several data sources could share a common link by multiplexing during the idle intervals. Thus, the ATM adaptation layer allows for both constant- and variable-bit-rate services; this is an effective method that does not require that bandwidth be allocated in fixed multiples. The combination of transmission options is sometimes described as a pleisio-synchronous network, meaning that it combines some features of multiplexing operations without requiring a fully synchronous implementation. Note that the fixed cell length enables the use of synchronous multiplexing and switching techniques, while the generation of cells on demand enables flexible use of the link bandwidth for different types of data, characteristic of packet switching. Higher level protocols may be required in an ATM network to ensure that multiplexed cells arrive in the correct order, or to check the data payload for errors (given the typical high reliability and low BER of modern fiber optic technology, it was considered unnecessary overhead to replicate data error checks at each node of an ATM network). If an intermediate node in an ATM network detects an error in the cell header, cells may be discarded without notification to either end user. Although cell loss priority may be defined in the ATM header, for some applications the adoption of unacknowledged transmission may be a concern.

ATM data rates were intended to match SONET rates of 51, 155, and 622 Mbits/s; an FDDI compliant data rate of

Table 12 Gigabit ethernet physical layer specifications.

Fiber Type	Distance
SM 1000BaseLX	5 km
MM 62.5	220 meters/160 MHz-km
1000BaseSX	275 meters/200 MHz-km
MM 62.5 1000BaseLX (with mode conditioner)	550 meters/500 MHz-km
MM 50.0 1000BaseSX	550 meters/500 MHz-km
MM 50.0 1000BaseLX (with mode conditioner)	550 meters/500 MHz-km

SX indicates short wavelength lasers, LX indicates long wavelength lasers.

100 Mbits/s was added to facilitate emulation of different types of LAN traffic over ATM. To provide an LC copper option and compatibility with 16 Mbits/s token ring LANs to the desktop, a 25 Mbits/s speed is also approved. For premises wiring applications, ATM specifies the SC duplex connector, color coded beige for multimode links and blue for single-mode links (Fig. 3). At 155 Mbits/s, multimode (MM) ATM links support a maximum distance of 3 km while single-mode (SM) links support up to 20 km. Because of its ability to emulate different LAN protocols, ATM has sometimes been cast as a competitor to FC; however, in most applications the two standards actually complement one another, since FC is applied to high-speed dedicated connections, while ATM is employed in a networked environment.

2.6 Gigabit Ethernet

With the commercial success of 10 Mbit/s ethernet and 100 Mbit/s fast ethernet standards, a new proposal for 1000 Mbit/s or gigabit ethernet was finalized this past year.²⁹ Using a common ethernet frame format (IEEE 802.3z), this widely anticipated standard is being supported by many datacom equipment manufacturers. Many view the large installed base of ethernet products as an indication that this new interface will be widely accepted. Since the link topologies and data formats for ethernet are well established, we will not cover them in detail here. The fiber optic physical layer specifications for gigabit ethernet are given in Table 12, the standard supports a family of interfaces, including short distances on copper, hundreds of meters on multimode fiber, and kilometers on singlemode fiber. All of the fiber optic interfaces use the SC duplex connector. Both short wavelength (SX) and long wavelength (LX) laser options are supported in the standard, operating at a data rate of 1.25 Gbit/s. A new feature in this standard is a provision for operating the same long wavelength optical transmitter over both singlemode and multimode fiber types. If a long wavelength laser transmitter is optimized for a launch into singlemode fiber, it will typically not be optimal for multimode fiber; modal noise problems such as differential mode delay can result, making the link inoperable over any useful distance.³⁰⁻³¹ The standard provides for the use of optical mode conditioners to facilitate the launch into multimode fiber; these mode conditioners typically provide an offset

launch of light into the edge of multimode fiber, and avoid exciting mode groups which are prone to causing differential mode delay.

3 Emerging Technologies

There are many developing technologies that may play an important role in datacom systems within the next 5 to 10 yr. In the following sections, we provide a brief overview of some of the more promising emerging technologies. It is not possible to make this list all-inclusive; rather, we discuss some of the major industry trends and directions. As before, three dominant applications tend to drive emerging technologies and standards, but these applications are not exclusive. For example, future storage systems may benefit from wavelength multiplexing and free-space interconnects at long distances, or from parallel optical interconnects at short distances; future coupling systems will probably take increasing advantage of parallel optical interconnects; and future networking systems will use gigabit ethernet, plastic fiber, and wavelength-multiplexing technologies.

3.1 Parallel Optical Interconnect

In the quest for 2.5 Gbits/s data rates and beyond, one approach is to develop high-speed serial transceivers similar to those currently in use for ESCON, ATM, etc. Another approach that may be more cost effective for short distances involves striping several lower speed fibers in parallel; this space-division multiplexed approach is the motivation behind parallel optical interconnects (POIs). Recent advances in vertical-cavity surface-emitting laser (VCSEL) arrays, optical fiber-ribbon connectors, and receiver arrays have contributed to the commercialization of this technology. The first commercially available POI transceiver was the Motorola Optobus, announced³² in 1995; it provides a 10-fiber ribbon between transmitter and receiver arrays, with an aggregate data rate of 2.4 Gbits/s over 300 m (this product was recently discontinued). Since then, there have been many vendors who have demonstrated prototype POIs (see Table 13), and more commercial products will be available in the near future. Although there is not yet a common standard for these devices, most are based on SW VCSELs and are limited to a few tens or hundreds of meters. There is no standardization on the ribbon connectors; although many products are based on the MTP/MPO 12-fiber ferrule (Fig. 6), which has been accepted as part of various industry standards,^{*} different manufacturers have developed at least a half dozen proprietary implementations of optical connectors based on this ferrule that are not compatible (see Table 14). Applications for POI include highly parallel supercomputers or clustered multiprocessors for large servers, large telecom or datacom switches, channel extensions for existing parallel data buses such as the SCI, and options for emerging parallel bus standards such as HIPPI 6400.[†] With the potential for significant cost reductions in packaging and testing, POIs could also supplement

*International industry standards based on the MT ferrule include those in Ref. 33. Additional work is currently under discussion by the ANSI HIPPI 6400 committee and the EIA/TIA FOCIS group for draft proposal.

[†]As preceding footnote.

Table 13 POI transceivers.³²

Company	Number of Channels	Maximum Mb/s per channel	Maximum Distance	Fiber	Source/Wavelength
Siemens	12	1000	300	MM	LD/850
Gore	12	1000	300	MM	LD/850
Fujitsu	6	622	400	SM	LD/1300
Hitachi	12	200	100	SM	LD/1300
NEC	6	1100	—	SM/MM	LD/1300
NTT	5	2800	250	MM	LD/1300
Oki	12	311	200	MM	LD/1310
Motorola ¹	10	400	300	MM	VCSEL/860
OETC ²	32	500	100	MM	VCSEL/850
POLO ³	10	622	400	MM	VCSEL/980
Vixel	4	1250	300	MM	VCSEL/840
Optobahn 1	10	650	300	SM/MM	LD/1301
Optobahn 2	10	1000	300	MM	LD/1310
Jitney ⁴	20	1500	300	MM	VCSEL/850
SPIEBOC ⁵	12	2500	—	SM	VCSEL/850

¹The Motorola POI was commercially known as Optobus.

²OETC is the Optoelectronics Technology Consortium, an Advanced Research Projects Agency (ARPA)-funded consortium consisting of IBM, AT&T, Honeywell, and Lockheed Martin.

³POLO is the Parallel Optical Link Organization, an ARPA-funded consortium led by Hewlett Packard.

⁴Jitney is the product of an ARPA-funded consortium consisting of IBM, 3M, and Lexmark.

⁵SPIEBOC is the Standardized Packaging and Interconnection for Inter- and Intra-Board Optical Communications, a consortium funded by the Commission for the European Communities under the ESPRIT II OLIVES program, whose members include British Telecom, Bull SA, ETH Zurich, Europ-tics Ltd., Graseby Microsystems Ltd., IMC Stockholm, Ericsson, Philips Optoelectronics Centre, Telefonica de Espana, and White Cross Advanced Products Ltd.

existing duplex transceivers in applications that require multiple duplex channels per card and are typically short distance.

Emerging highly parallel computer architectures may find applications for POIs; for example, ARPA has recently granted several contracts to build the world's largest super-computer by 1998 under the Advanced Supercomputing Initiative (ASCI). To meet the requirements of large simulations for the U.S. Department of Energy, including modeling of nuclear explosions that will be heavily relied on since live atomic testing is no longer permitted, several companies are pursuing different paths to a teraflops machine. These machines are also applicable to so-called "grand challenge" problems, such as mapping the human genetic structure or some types of astronomical modeling of our universe. Cray is exploring a nonuniform memory architecture, Intel is developing a massively parallel machine with 9000 Pentiums, and IBM is developing coupled RISC PowerParallel[‡] workstations.³⁴ As applications such as these become more commonplace, the practical need for POIs will continue to grow.

3.2 Small Form Factor Optical Connectors

Conventional duplex fiber optic connectors mount a pair of fibers about 12.5 mm apart; since the optical fiber is only about 125 μ m in diameter, there is potential for significantly smaller optical connectors. Recently, several new

connector types have emerged that attempt to mount a pair of fibers in the same space as a standard RJ-45 electrical jack; this allows for increased density of fiber optic ports on networking equipment. There are many different technical approaches to this problem, and there is no clear industry standard yet for small form factor connectors. One option is to use a pair of standard 2.5-mm-diam ferrules placed closer together; this is the approach used by the OPTI-JACK* connector developed by Panduit Corporation. Another way to achieve smaller connectors is by using smaller ferrules. The LC connector, developed by Lucent, is commonly used in telecom applications and standardized for data rates exceeding 800 Mb/s by IEEE 1394. The LC uses 1.25-mm-diam ferrules for connectors, which are available as either simplex or duplex. Another option is to use existing multifiber ferrules populated with only two fibers. The mini-MT connector developed by U.S. Conec (under discussion as a possible IEC standard) is based on the 12-fiber MT ferrule, which mounts fibers on a 750- μ m pitch and uses a pair of precision metal pins for alignment. Similarly, the MT-RJ connector, developed by an industry consortium including Amp, Hewlett Packard, Siecor, Siemens, and Fujikura, uses an MT ferrule and RJ-45 connector body. Another option is the SC-DC connector[†] developed by Siecor, which features a new type of ferrule designed to accommodate either two or four fibers; the con-

*OPTI-JACK is a registered trademark of Panduit Corporation.

†SC-DC is a registered trademark of Siecor Corporation.

[‡]RISC and PowerParallel are registered trademarks of IBM Corporation.




Fig. 6 (a) MTP style multifiber connector and (b) multifiber ferrule with alignment pins.

connector body is based on the conventional SC, with different keying and ferrule alignment features. Finally, it may be possible to eliminate ferrules altogether; the SG connector (also known as Volition or VF-45[‡]), developed by 3M has been adopted as an alternative for FC. It cantilevers a pair of optical fibers in free space and uses plastic V-grooves to facilitate alignment with a receptacle. Important issues remain to be resolved for many of these small connector types, including reliability, cost, and support of both MM and SM fiber.³⁵ The various small form factor connectors are illustrated in Fig. 7.

3.3 Plastic Fiber

There has been recent research into the use of large-core plastic fiber as an optical data link, and some vendors offer commercially available plastic fiber link components. Because of the much higher volumes of glass fiber manufactured today, the cost of plastic optical fiber remains higher than glass. Plastic fibers are currently limited by their high loss (120 to 150 dB/km) to fairly short distances (around 50 m). Despite this, there is ongoing interest in this technology because plastic fiber links can use very low cost visible light sources; a 680-nm (red) LED transceiver could be as inexpensive as \$5 to \$10, compared with over \$100 for a

1300-nm transceiver. Plastic fiber links have already found some applications in laser printers and microprocessor control systems for automobiles. Emerging applications include the consumer market and small office/home office (SOHO) environment, where there is a need for links between computers, fax machines, printers, and other equipment that can withstand the high levels of electromagnetic noise present in a typical home. Many home appliances, such as microwave ovens, emit broad spectrum noise; existing home wiring, such as stereo audio cables, may be insufficient to prevent interference with relatively delicate computer equipment. Plastic fiber is potentially less expensive than unshielded category 5 wiring, and may be easy enough for a typical homeowner to install without special training, since it does not require precise polishing when terminating a new connector. Several plastic fiber connectors have been proposed, including variations on the F07, SMA 905 and 906, EIAJ Digital Audio, ST, and FC connectors. The ATM Forum is investigating new standards for this home area network (HAN), including a 50-m plastic fiber link at 51.8 Mbits/s and a 17-dB link budget. Some have even proposed that red VCSEL lasers could operate up to 1 Gbit/s over short distances, providing a solution for cost-sensitive applications such as fiber to the desktop and PC/television convergence. Higher bandwidth applications are anticipated in this area, coincident with the development of universal serial bus attachments to the PC at 1.5 or 12 Mbits/s [sufficient to carry broadband Integrated Services Digital Network (ISDN) traffic] as well as IEEE P1394 (better known by the Apple trademark, Firewire) at 100 or 400 Mbits/s. Such applications are expected to far exceed the 2.5 Mbits/s capacity of a typical parallel computer port, and may provide additional incentive for plastic fiber solutions.

3.4 Free-Space Data Links

In 1995, the Infrared Data Association (IrDA) finalized a standard for wireless IR communications; over 100 vendors including IBM, Microsoft, Motorola, and others have endorsed this new standard and announced plans to build compliant equipment.³⁶ While wireless IR data links have been available for some time in various forms, this new standard is expected to accelerate their widespread acceptance, especially in areas where it is too expensive or impractical to retrofit optical fiber, such as large warehouses using computerized data management. This is an emerging market, and most applications do not compete with fiber since they are limited to a few hundred kilobits per second and a few meters distance.

There is a different application for free-space optical communications, which involves the use of 4- to 8-in. telescope optics to provide unguided optical communication across several miles in clear weather. One application is wireless optical data links between office buildings in large cities, where the cost to run optical fiber across the street is prohibitive because it involves leasing “dark fiber” from the local telecom providers. For example, AstroTerra Corporation currently offers 155 Mbits/s ATM services on an unguided line-of-sight laser beam connection, as well as other types of datacom links. This technology is best suited to areas with stable weather conditions, such as the deserts

[‡]Volition and VF-45 are registered trademarks of 3M Corporation.

Table 14 Multifiber connector options.

Company	Alignment mechanism	Fiber Count	Notes
U.S. Conec (MTP)	MT ferrule/pins	2–12	Known as MPO in Europe
Amp	MT ferrule/pins	2–12	
Alcoa-Fujikura	MT ferrule/pins	10	Applications on Optobus
Methode (MP)	AT&T MAC, external pins	2–16	
Berg (MiniMAC)	AT&T MAC, external pins	2–32	
3M	Precision ball-in-socket	2–18	

Note that not all connectors are commercially available at this time; none are compatible with any other connector on list. MTP has been adopted as a Japanese telecom standard.

of Las Vegas or the Middle East; potential problems such as interference from rain, clouds, and bad weather are being addressed by second-generation devices.

3.5 Wavelength Multiplexing

Multiplexing wavelengths is a way to take advantage of the high bandwidth of fiber optic cables without requiring extremely high modulation rates at the transceiver. With an available bandwidth of about 25 THz, a single optical fiber could carry all the telephone traffic in the United States on the busiest day of the year (recently, Mother's Day has been slightly exceeded by Valentine's Day). Traditionally, optical wavelength-division multiplexing (WDM) has been widely used in telecom applications, but has found only limited applications in datacom. This may change, as a number of companies are now offering multiplexing alternatives to datacom networks that need to make more efficient use of their existing bandwidth. One example is the IBM 9729 Optical Wavelength Division Multiplexer.^{37,38} The first WDM product developed specifically for the datacom industry, it enables the transmission of up to 20 independent data streams over a single fiber. Using different adapter cards, a mixture of ESCON, FC, FDDI, and ATM links can be plugged into the device, which is protocol independent. The data is remodulated using distributed feedback laser diodes with wavelengths spaced 1 nm apart near the 1550-nm region. The optical signals are then combined using a diffraction grating with embedded fiber pig-tails, and coupled into a singlemode fiber; another unit at the far end of the link demultiplexes the signals. The maximum unrepeat distance is 50 km for data in the 200 Mbits/s range; higher speeds are possible at shorter distances (for example, 16 bit/s at 20 to 40 km). This type of product is a cost effective way to utilize leased fiber optic lines, which can cost \$150 to \$300 per mile per month or more (where available). This technology is being applied to geographically dispersed parallel sysplex architectures,³⁹ and may be the first step toward development of all-optical networks. Other commercial WDM products include the MultiWave WDM terminal from Ciena Corp., which accommodates up to 16 OC-48 channels and will soon be upgradable to 40 channels, and the WaveMux 3200 from Pirelli, which handles up to 10 OC-192 channels or 32 OC-48 channels. Several other corporations are also investigating applications of WDM, including the MONET con-

sortium led by AT&T, the ESPRIT program in Europe, the Optical Network Technology Consortium, the All Optical Network Consortium, and others.³⁸

3.6 Other Datacom Applications

Given the rapid pace of developments in the datacom industry, it is difficult to predict the next big technology breakthrough; there is a great deal of ongoing research in smart pixels, optical multichip modules, planar waveguides, all-optical switching, and other areas that might contribute to advancing the state of the art. For example, the first commercial product with an optical backplane connection, using the MTP connector, was introduced this past year by DCS Corporation; the use of optics in the backplane may become increasingly important. Since some datacom technologies derive from telecom applications, such as WDM, perhaps future datacom systems will incorporate advances such as erbium-doped optical amplifiers or multigigabit transmission. Fiber optics have also frequently been employed to provide a channel extender for existing datacom links. For example, some token ring applications have made use of optical fiber channels.⁴⁰ In large computing environments, the 800 Mbits/s parallel copper interface of HIPPI has been extended to several kilometers by a serial fiber optic channel extender, known as⁴¹ serial HIPPI. There is also interest in new package types for optical transceivers, such as the Gigabit Interface Converter (GBIC), which enables the transceiver to be pluggable so that different physical layers can be interchanged easily. Many of these areas could represent important markets for fiber optic datacom. The recent introduction of network computers ("thin clients" with limited internal storage that rely on network connections to a server) represent a large potential expansion in LAN bandwidth requirements (see Table 2 for the network computer specification online).

Another area that has received increased attention is cable management in large data centers. Many companies now offer preterminated fiber optic trunks intended to simplify the installation of fiber under the raised floor or in overhead conveyances of a large computer center. Many different options exist, including under-the-floor junction boxes, rack-mounted units, and even "intelligent" patch panels that can literally tell the user where to plug cables by means of flashing colored lights. Other approaches include the use of multifiber optical connectors to manage cable




Fig. 7 Small form factor fiber optic connectors: (a) SC/DC, (b) SG or VF-45, (c) LC, (d) mini-MT, (e) MT-RJ, and (f) OPTI-JACK.

congestion and make it easier to reroute or change fiber cables. An example is the Fiber Transport Services (FTS) offering developed⁴² by IBM; using 12-fiber MTP connectors on trunk cables and hybrid cable harnesses with six duplex connectors fanning into an MTP connector, this service provides direct attachment of trunk cables to an IBM parallel server or large datacom switch (128 duplex jumper cables can be condensed into only four multifiber trunk attachments). Similar cable management solutions are offered by many companies, including Amp and Siecor.

4 Conclusions

A comparison of the commercial datacom standards currently in use from Table 15 reveals that most datacom applications are presently using data rates around 200 Mbits/s, with typical MM distances up to 3 km and SM distances up to 20 km. Note that most applications get by with much shorter distances, and only a small percentage of the market is currently using the upper distance limits. However, the long-distance applications are typically asso-

Table 15 Comparative performance of datacom standards for MM fiber.

	ESCON	ATM/SONET	FCS	FDDI	POI
Preferred cable	62.5	62.5	50	62.5	100
Type (μm)					
Maximum distance (km)	3	3	1–3	2	0.10
Data rate (Mbits/s)	200	155–622	266–1062	100	200/channel
Connector	ESCON	SC duplex	SC duplex	MIC or SC duplex	MTP or ?
Connector color	Black	Beige	None	None	None
Wavelength (nm)	1300	1300–1550	850–1300	1300	850
Application	I/O	Network	I/O, network, or coupling	Network	Coupling
Other comments	—	—	OFC safety option	Bypass switch option	Not standardized

ciated with very large data centers and mission critical applications, which makes them an important factor in the datacom market. Most datacom systems are operating at a wavelength of 1300 nm; there has not been a significant cost differentiation between 1300- and 850-nm laser diodes to make up for the distance limitations of shorter wavelength links. The cost of a 200 Mbits/s LED transceiver has fallen dramatically in the past 5 yr, from several hundred dollars to under \$100; this has helped fiber to surpass copper links in some applications, for example, the displacement of parallel copper cables by ESCON. As users become more familiar with the benefits of optical fiber (noise immunity, ease of installation, high reliability, etc.) they have become more comfortable installing fiber in the data center. It takes a long time to completely displace any existing technology, however, and parallel copper will probably continue to exist in the data center for many years.

A high-performance computer system is composed of many elements: a fast, powerful CPU; memory; a power supply; I/O devices such as a monitor, keyboard, or mouse; and the means to network with other computers, servers, or DASD for mass storage. Since the advent of the microprocessor in the 1960s, the growth of processing power has been governed by improvements in the CPU; Moore's law predicted the exponential growth in devices per chip, which continues to drive the semiconductor industry today.⁴³ With modern very large scale integration (VLSI) chips measuring feature sizes in fractions of a micrometer, we have achieved a 100 million times reduction over vacuum tubes—coincidentally, this is about the same as the order of magnitude increase in speed between voice grade phone wire to the 20-THz available bandwidth of optical fiber. With the emergence of efficient coupling and parallel processing schemes, the need for higher bandwidth interconnects may become just as important in determining future generations of computer performance. This is reflected to some extent in the current roadmaps of industry associations such as the OIDA, which project future bandwidth requirements in datacom networks (see Table 2 for the OIDA Web site). Many applications are driving this trend, including the growth of the World Wide Web, data mining in large corporate databases, digital libraries, video on de-

mand electronic commerce and others. Terabyte and petabyte databases are becoming more common, and the hardware required to effectively mine these databases for useful information is demanding increasingly high bandwidth data links. As an example of the growing need, IBM has recently established a TeraPlex Complex to facilitate customers interested in using large, parallel coupled transaction servers to solve database management issues.

While different applications may have very different requirements, we can draw some conclusions about future bandwidth requirements. Because there are so many variations on LAN implementations, future technology requirements in this area are very difficult to predict, instead, we consider somewhat more stable applications. In the area of I/O attachments, the bandwidth of an individual I/O channel tends to increase in step increments, always keeping just ahead of DASD requirements. Parallel copper channels offered a net bandwidth of 3.0 Mbytes in the early 1960s; this did not increase until 1981, when the data rate increased to 4.5 Mbytes. With the introduction of optical fiber in 1990, the data rate jumped to 17 Mbytes and has been on a sharply upward trend ever since. This trend will probably continue, with the next step to 1 Gbit/s I/O being widespread use of FICON. If we consider the total bandwidth requirements for mainframes and large server I/O (number of channels multiplied by bandwidth per channel) for high-end applications over time, then we see an exponentially increasing trend since the introduction of optical fiber links (Fig. 8). If we extrapolate following Moore's law—type behavior, we can estimate the total I/O bandwidth required on a high-end server[‡] by the year 2000.

For other applications such as coupling, the future is more difficult to predict. Parallel coupled architectures have only been available for about 3 yr, and first-generation systems introduced a relatively high total coupling bandwidth (around 2 Gbits/s for two coupled CPUs). This has proven to be adequate bandwidth for two to three generations of CMOS processors, which are now available in configura-

[‡]When this chart was first presented at the 1996 Optical Society of America annual meeting by the author, one participant was kind enough to refer to this trend as "DeCusatis's law."




Fig. 8 Trend in large server total I/O bandwidth versus time based on IBM and competitive products including S/360 (1964), S/370 (1970), S/390 (1990), and G3 and Multiprise Servers (1996); and the largest available configuration for each year. Future points are extrapolations only.

Key to symbols:

◆=commercial systems

○=high-End technical computing (supercomputers)

+ =IBM PowerParallel Supercomputers

FJ VPP500=Fujitsu Vector Parallel Processor Supercomputer

SGI=Silicon Graphics Supercomputers

FCS=Fiber Channel Standard

DEC=Digital Equipment Corp. proprietary "MemChannel" interconnect

SUN=proprietary "S3.mp" interconnect from SUN Microsystems

Paragon=Intel parallel processor interconnect

tions with up to 10 coupled CPUs and an aggregate system bandwidth of 64 Gbits/s (although the cost of a 1 Gbit/s transceiver has not yet fallen as dramatically as the 200 Mbits/s modules discussed earlier due to smaller volumes). Although the attached coupling bandwidth may exceed CPU capacity today, advances in CMOS processors over the next 3 to 5 yr along with more sophisticated parallel software applications may drive the need for higher bandwidth links, and the gradual replacement of full duplex 1 Gbit/s links with full duplex 2 Gbits/s links. Similar opportunities exist for massively parallel supercomputer architectures; together, these areas may help drive the introduction of parallel optical interconnects as a cost-effective way to reach higher data rates. While the use of serial or parallel links for this application is not clear, traditionally this market has been cost driven, not technology driven; if POI is to develop into its predicted market (\$0.5 B in datacom³² by

2000), then it must at least be cost competitive with serial transceiver solutions, and preferably with copper. In the past, data links were often neglected in the computer center because they fall between devices (thus are not truly owned by either end of the link) and represent a small percentage cost of a large computer system. With the increasing importance of networking, and the emerging paradigm of "networking the mainframe," this technology will merit closer attention in the years to come.

References

1. J. W. Goodman, "The Optical Data Processing Family Tree," *Opt. News* 4, 25-28 (1984).
2. M. K. Hibbs-Brenner, Y. Liu, A. Cox, D. Greenlaw, L. Galameau, J. Bristow, J. A. Lehman, T. Werner, and R. A. Morgan, "Components and packaging for low-cost optoelectronic modules," in *Proc. OSA Annual Meeting*, p. 132, Rochester, NY (1996).
3. ANSI, "Standard for the safe use of lasers," ANSI Z136.1 (1993)

- revision); ANSI, "Standard for the safe use of optical fiber communication systems utilizing laser diodes and LED sources," ANSI Z136.2 (1996–97 revision).
4. International Electrotechnical Commission (IEC/CEI), IEC/CEI 825 (1993 revision).
 5. Electronics Industry Association/Telecommunications Industry Association (EIA/TIA) "Commercial building telecommunications cabling standard," EIA/TIA-568-A, Electronics Industry Association, New York (1996).
 6. Electronics Industry Association/Telecommunications Industry Association (EIA/TIA), "Detail specification for 62.5 micron core diameter/125 micron cladding diameter class 1a multimode graded index optical waveguide fibers," EIA/TIA-492AAAA, Electronics Industry Association, New York (1996).
 7. Electronics Industry Association/Telecommunications Industry Association (EIA/TIA), "Detail specification for class IV-a dispersion unshifted single-mode optical waveguide fibers used in communications systems," EIA/TIA-492BAAA, Electronics Industry Association, New York (1996).
 8. *IBM J. Res. Develop.* special issue on IBM System/390: architecture and design, **36**(4) (1992).
 9. C. DeCusatis, "Data processing systems and optoelectronics," Chap. 6 in *Optoelectronics for Data Communications*, R. Lasky, U. Osterberg, and D. Stigliani, Eds., pp. 219–283; Academic Press, New York (1995).
 10. D. Stigliani, "Enterprise systems connection fiber optic link," Chap. 13 in *Handbook of Fiber Optic Data Communications*, C. DeCusatis, R. Lasky, D. Clement, and E. Mass, Eds., Academic Press, New York (1998).
 11. "ESCON I/O interface physical layer document," IBM document SA23-0394, 3rd. ed., IBM Corporation, Mechanicsburg, PA (1995).
 12. "Introduction to Enterprise Systems Connection," IBM document GA23-383, IBM Corporation, Mechanicsburg, PA (1991); "Enterprise Systems Architecture/390 ESCON I/O Interface," IBM document SA22-7202, IBM Corporation, Mechanicsburg, PA (1991).
 13. C. DeCusatis, A. Huffman, G. DeMario, and D. Stigliani, "ESCON fiber optic interface for System/390," *Opt. News* **2**, 33 (1991).
 14. "ANSI Single Byte Command Code Sets CONnection architecture (SBCON)," draft ANSI standard X3T11/95-469, Revision 2.2 (1996).
 15. A. X. Widmer and P. A. Franaszek, "A dc-balanced, partition-block, 8B/10B transmission code," *IBM J. Res. Develop.* **27**, 440–451 (1983).
 16. R. Thapar, "Fiber Distributed Data Interface," Chap. 14 in *Handbook of Fiber Optic Data Communications*, C. DeCusatis, R. Lasky, D. Clement, and E. Mass, Eds., Academic Press, New York (1998).
 17. ANSI "Fiber distributed data interface physical layer media dependent (PMD)," ANSI X3T9.5 (1984).
 18. "IBM Open Systems Adapter (OSA-2) announcement," *Software Server* **5**, 41 (1995).
 19. ANSI, "Fibre Channel—physical and signaling interface (FC-PH)," ANSI X3.230-1994, Revision 4.3 (1994).
 20. ANSI, "Fibre Channel—arbitrated loop (FC-AL)," ANSI X3.272-199x, Revision 4.5 (June 1995).
 21. ANSI, "Fiber Channel protocol for SCSI (FCP)," ANSI X3.269-199x, Revision 012 (May 1995).
 22. "Coupling facility channel I/O interface physical layer," IBM document SA23-0395 (1994).
 23. S. Van Doren and C. DeCusatis, "Fibre Channel standard," Chap. 12 in *Handbook of Fiber Optic Data Communications*, C. DeCusatis, R. Lasky, D. Clement, and E. Mass, Eds., Academic Press, New York (1998). Also see "S/390 fiber channel connection (FICON) I/O interface physical layer," IBM document SA24-7172 (1998).
 24. ANSI "Digital hierarchy optical rates and format specification," ANSI T1.105-1988 (1988); ANSI T1.106-1988 (1988).
 25. CCITT, "Synchronous digital hierarchy bit rates," CCITT Recommendation G.707 (1984).
 26. CCITT "Network node interfaces for the synchronous digital hierarchy," CCITT Recommendation G.708 (1984).
 27. CCITT, "Synchronous multiplexing structure," CCITT Recommendation G.709 (1984).
 28. C. Beckman, "Asynchronous transfer mode," Chap. 11 in *Handbook of Fiber Optic Data Communications*, C. DeCusatis, R. Lasky, D. Clement, and E. Mass, Eds., Academic Press, New York (1998).
 29. IEEE 802.3z, Supplement to Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications; Media Access Control (MAC) parameters, physical layer, repeater and management parameters for 1000 Mb/s operation.
 30. J. Fox, T. Giles, and A. MacGregor, "Bandwidth reduction in high speed laser communication systems over multimode fiber," (in this issue).
 31. C. DeCusatis, J. Fox, T. Giles, and A. MacGregor, "Differential mode delay in gigabit fiber optic links," *Proc. OSA Annual Meeting*, Baltimore, MD, Oct. 4–9, 1998; Paper Th004, *Opt. and Photonics News* **9**(8), 140 (August 1998).
 32. J. Crow and A. Benner, "Intramachine communications," Chap. 10 in *Handbook of Fiber Optic Data Communications*, C. DeCusatis, R. Lasky, D. Clement, and E. Mass, Eds., Academic Press, New York (1998).
 33. IEC Sectional Specification 874-16 (1994); IEC: Interface Standard 1754-5 (1996); IEC Interface Standard 1754-7 (1997); European Union Sectional Specification EN 186240 (1994); European Union Interface Standard (1997); JIS C5981 (1984); JIS C5982 (1997).
 34. A. Wolfe, "Cray enters race for teraflops computer," *EE Times*, issue 992, p. 1 (Oct. 7, 1996).
 35. C. DeCusatis, J. Trehwella, J. Fox, T. Giles, and S. Thomas, "Comparison of small form factor fiber optic connectors," *Proc. OSA Annual Meeting*, Baltimore, MD, Oct. 4–9, 1998; Paper Th005, *Opt. and Photon. News* **9**(8), 140 (August 1998).
 36. IrDA, "IrDA Serial Interface Test Suite, Link Access Protocol (IrLAP), Secondary Station Implementation Test Specification ver. 3.0" and "IrDA Serial Interface Test Suite, Link Management Protocol (IrLMP), Secondary Station Implementation Test Specification ver. 3.0," both available from Genoa Technology, Moorpark, CA (July 1995) or by anonymous FTP at ftp.gentech.com in the /pub directory.
 37. "IBM Corporation 9729 Optical Wavelength Division Multiplexer," *Photon. Spectra* special issue, **30** (June 1996).
 38. C. S. Li and F. Tong, "Emerging technology for fiber optic data communication," Chap. 21 in *Handbook of Fiber Optic Data Communications*, C. DeCusatis, R. Lasky, D. Clement, and E. Mass, Eds., Academic Press, New York (1998).
 39. C. DeCusatis, W. E. Hall, F. Janiello, and D. Petersen, "Geographically dispersed parallel sysplex architecture using optical wavelength division multiplexing," (in this issue).
 40. G. McClain, Ed., *The Handbook of International Connectivity Standards*, Van Nostrand Reinhold, New York (1992).
 41. R. J. Cypser, *Communications for Cooperating Systems*, Addison-Wesley, New York (1992).
 42. C. DeCusatis, "Wavelength and channel multiplexing for data communications," in *Proc. OSA Annu. Meeting*, p. 115, Portland, OR (1995).
 43. G. E. Moore, "Lithography and the future of Moore's law," in *Optical/Laser Microlithography VII*, *Proc. SPIE* **2440**, 2–17 (1995).

Casimer DeCusatis: Biography and photograph appear with the special section guest editorial in this issue.