IBM Enterprise Systems multimode fiber optic technology

by N. R. Aulet
D. W. Boerstler
G. DeMario
F. D. Ferraiolo
C. E. Hayward
C. D. Heath
A. L. Huffman
W. R. Kelly
G. W. Peterson
D. J. Stigliani, Jr.

This paper describes the first implementation of optical fiber technology for the I/O channel connections of the IBM Enterprise Systems Connection (ESCON™) Architecture™. The ESCON optical link line rate is 200 megabits per second and is capable of transmission over distances of 3 km. The link is composed of a serializer, electro-optic transmitter, duplex fiber optic cable, electro-optic receiver, and deserializer. The serializer and deserializer respectively perform the conversions from parallel to serial and serial to parallel formats. The clock which is used to retime the serial data in the deserializer is extracted from the encoded serial signal using a phase-locked loop (PLL) technique. The optical link technology selected to achieve the data processing system requirements is InGaAsP/InP 1300-nm LED, InGaAsP/InP PIN photodiode, and multimode optical fiber. A duplex fiber jumper cable is designed with a rugged, low-profile, polarized connector, with a unique protective cap which recedes as it is mated. The optical link loss budget is determined by dividing the link into two major categories: available optical power and cable plant loss. The link design ensures that the minimum available power is greater than the maximum cable plant loss. The design parameters and trade-offs of the optical link are discussed in this paper. Unique measurement techniques and tools to ensure reliable and consistent link performance are described.

Introduction

To meet future information and data processing needs, the large processor complex must be able to respond to new and varied data communications requirements. To achieve this goal, the data center interconnection method must be able to move data at higher speeds and over longer distances than are possible with today's parallel OEMI (original equipment manufacturers' information) interface [1] (approximately 122 m). For example, the interconnection technique must allow processors and control units to communicate over a campus environment (2–3 km) and achieve high data rate performance while allowing for growth with minimal disruption to existing

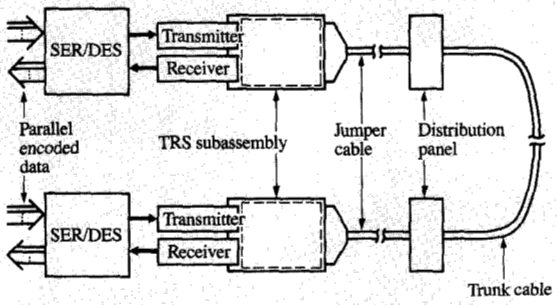

Figure 1

Block diagram of fiber optic link elements.

users and ensuring a high level of processor availability during a failure.

Clearly, any solution must be easier to use, lighter in weight, smaller in size, and more reliable than the existing OEMI interface. It must also allow for data processing growth throughout the 1990s. To achieve this objective IBM has developed the Enterprise Systems Connection (ESCON™) Architecture™. The ESCON architecture meets the needs of the future by providing full interconnectivity for processor-to-processor and processor-to-control-unit communication [2, 3]. The topology used to achieve this communication network is a switched point-to-point star with the ESCON Director™ [4] as the switch. Point-to-point communication is achieved with the fiber optic link described in this paper. The approach allows interconnectivity among all units in the network as well as dual-path (with two directors) capability to ensure high system availability.

The development of the ESCON link represents a significant advance in the utilization of fiber optic technology for high-end processor interconnection. Several technical advances and innovations were required to achieve the product objectives:

- 200 Mb/s LED/multimode link at 3 km distance.
- Duplex fiber optic connector.
- Statistical optical loss analysis.
- 8B/10B data encoding technique [5].
- PLL/serializer/deserializer implementation.
- Low bit error rate (BER).

IBM has worked jointly with other companies to develop the transmitter–receiver subassembly and the duplex cable assembly.

Figure 1 shows the key link elements, outlined below:

1. *Serialization–deserialization (SER/DES)*
   The serializer converts the encoded 10-bit parallel data to serial data for transmission over the link. The deserializer converts the serial data to 10-bit parallel data. Both of these functions are performed in a single hybrid module called the SER/DES module.
2. *Transmitter–receiver subassembly (TRS)*
   The TRS contains a transducer to convert the high-speed serial electrical signal to an optical signal; it contains a receiver which performs the optical-to-electrical conversion. The TRS is designed to provide a mechanical interface with the duplex connector. The electro-optic technology consists of a light-emitting diode (LED) operating at 1300 nm wavelength in the transmitter and a PIN photodiode in the receiver.
3. *IBM jumper cable assembly*
   The jumper cable assembly consists of a duplex connector and a dual fiber optic cable. The IBM duplex connector is a low-profile, polarized connector which makes two physical contact connections with a single plug. The fiber is a 62.5-μm graded-index multimode fiber. The connector system and optical cable are designed for low loss.
4. *Trunk*
   The ESCON link is designed to support trunk fibers with core diameters of 62.5 or 50 μm. The connection loss is minimized, achieving the longest possible distance with 62.5-μm fiber.

Transmitter–receiver subassembly (TRS)

The TRS forms the interface between the optical fiber transmission medium and the electronic environment of the system. It establishes the physical connections and performs the necessary signal transformations. In the link topology chosen for the ESCON architecture, an optical fiber transfers data in only one direction. Therefore, two separate optical interfaces are required: a transmitter, to launch an optical signal into the fiber, and a receiver, to capture this signal at the other end. The ESCON TRS consists of independent transmitter and receiver modules joined by a duplex connector receptacle. The transmitter and receiver [6] are optimized for 40–200-Mb/s operation.

- *Transmitter*

The light source of the transmitter is an InGaAsP surface-emitting light-emitting diode (LED). This material set has been chosen for its long-wavelength (1300-nm) spectral characteristics, which are well matched to the attenuation and dispersion characteristics of the optical fiber. The LED is a planar structure comprising doped InP, InAs, and InGaAsP epitaxial layers on an InP substrate. The optical signal is generated in an "active" region within the

InGaAsP layer. The diameter of the active region is about 20 μm, which provides the best balance of speed, power, and reliability. This is realized through appropriate geometries on the top (epitaxial) and bottom (substrate) electrical contacts, which control the current flow pattern through the active layer. Emission from this region takes place in all directions, with relative maxima in the directions perpendicular to the junction plane. Because of the close proximity of the active layer, optical emission through the top surface is largely blocked by the top electrical contact. However, much farther away by comparison, the bottom electrical contact can be "windowed" or offset from the optical axis, thus allowing emission through the transparent (at 1300 nm) InP substrate. To partially recover "top" emission, a reflecting plane is implemented via the top contact metallization to redirect this power toward the fiber.

The optical emission pattern from the LED is approximately Lambertian (light intensity varies with the cosine of the emission angle measured from the normal to the LED surface), with a half-power angle of the order of 100 degrees. To maximize the coupling of this light to the optical fiber core, additional optics are required in the form of a double lens system. One lens is mounted as close to the LED as possible in order to capture much of the optical energy distribution and reduce the half-power angle. A second, larger lens is placed at a specific distance to capture the modified optical distribution and focus it on the fiber end face of the connector. An active alignment of the optical system is employed in manufacturing to optimize optical coupling from the LED to the fiber. The precision and stability of this alignment are critical to meeting specification requirements.

Since the optical fiber in the connector is fixed at the center of a precision cylindrical body (known as the connector ferrule), the transmitter provides a precision cylindrical alignment sleeve which accepts this ferrule and aligns its center with the optical axis of the LED and lens system. The stability and repeatability of the optical connection are greatly influenced by the clearance between the ferrule and the sleeve. These characteristics are optimized when the clearance is minimized; however, insufficient clearance can cause problems with module/connector mating. To satisfy stability and repeatability objectives and facilitate mating, it is necessary to ensure that ferrule and sleeve diameter tolerances are less than 1 μm. Additional features are also required on the sleeve entry to facilitate ferrule engagement. These features include a larger-diameter conical sleeve opening with a gradual transition between the conical and cylindrical regions. This allows for correction of slight angular misalignments between the sleeve and ferrule axes. This transition region is critically important for eliminating ferrule jamming.

The transmission scheme is digital amplitude modulation of light, realized through LED drive current modulation. To implement the required circuit functions, high-speed silicon bipolar integrated circuit (IC) technology operating from a single +5 V supply is used. This IC, consisting of dc-coupled differential current switching amplifiers, converts differential ECL input signals into LED drive current signals (>100 mA). The IC is mounted on a substrate with ceramic chip bypass capacitors, resulting in minimal current switching noise generated on the power supply lines.

All of these elements are integrated into a single module package which provides all necessary electrical connections. This package also provides EMI/RFI shielding by using a grounded, conductive (metal) housing to effectively create a Faraday cage around the active circuitry. This protects the internal circuitry from external noise sources and also minimizes emissions from the transmitter into the environment.

- *Receiver*

The photodetector of the receiver is an InGaAsP PIN photodiode. This technology was chosen for its 1300-nm compatibility, high-speed operation at 5 V bias, and excellent reliability.

This PIN photodiode is a planar structure consisting of InGaAs and InGaAsP epitaxial layers on an InP crystal substrate. The conversion of optical energy to electrical energy takes place in an "active" region within the InGaAs layer, where incoming photons are absorbed, generating electron–hole pairs. The device is reverse-biased during operation, establishing an electric field across the active region which depletes it of free carriers. The absence of free carriers maximizes sensitivity to photogenerated carriers, which flow out of the device under the influence of the electric field, forming a current signal proportional to the optical input.

It is possible to introduce optical power to this active region through either the top or bottom surface of the diode structure. However, much higher quantum efficiency can be realized if the light enters through the substrate surface. With this arrangement, photogenerated carriers are typically developed away from the p junction (contrary to top entry), which significantly reduces the probability of their recombining within the device, thus increasing "signal" carriers exiting the device. Additionally, substrate illumination allows for smaller top contact geometry, since an optical entry window is not required. This allows junction capacitance to be minimized and facilitates electric field optimization in the active region.

To maximize the amount of optical power which is coupled to the PIN photodiode active region from the optical fiber, the fiber end face should ideally be placed very close to the photodiode surface. As with the

Table 1 Nominal transmitter and receiver parameters.

Transmitter	
Optical coupled power* (dBm)	−17
Extinction ratio (optical) (dB)	11
Center wavelength (nm)	1325
Spectral width (nm)	150
Optical rise/fall time (ns)	1.2
Receiver	
Sensitivity (@ 10^{-12}) (dBm)	−35
Saturation (dBm)	−13

* Coupled into 62.5-μm core fiber.

transmitter, practical considerations relative to photodiode mounting, connection, and protection, combined with the removable nature of the fiber/connector, make it impossible to place the fiber end face close enough to adequately capture the optical power. Consequently, additional optics are required to channel and focus the optical power from the incoming fiber connector to the photodiode. This is accomplished by using a spherical lens to focus the light on the active area of the photodiode or a very short length of optical fiber to channel the light to the photodiode. As with the transmitter, an active alignment of the optical system is used in manufacturing to optimize optical coupling from the fiber to the photodiode.

To implement the circuit functions required, high-speed silicon bipolar IC technology is used in the receiver. These functions include a preamplifier, postamplifiers, a comparator, and ECL differential output drivers. Also included is a signal detect function which serves to indicate the presence or absence of a valid optical input signal. This circuitry also operates from a single 5-V power supply.

The PIN photodiode behaves as a very high impedance source while typically operating with very small (>1 μA) signal currents. To effectively reduce signal impedance, a transimpedance amplifier circuit is used in which the current signal from the photodiode passes through a circuit impedance, producing a voltage signal representation of the optical input. This configuration minimizes voltage swings across the photodiode, thus minimizing the bandwidth-limiting effects of junction capacitance.

The very low signal amplitude generated by the photodiode precludes amplification with respect to an independent dc reference, since drift and noise could then dominate the signal. Instead, a reference derived from the signal itself must be used. This is the signal "average" level obtained through integration and is realized by ac-coupling the transimpedance voltage signal to the input of the preamplifier.

The optical sensitivity of the receiver is directly related to the signal-to-noise ratio (SNR) in the preamplifier. To minimize the SNR degradation resulting from this ac coupling scheme, the data transmission code must maintain dc balance close to 50%. This requirement is met using the 8B/10B code with a maximum run length of 5 bits, allowing full function at data rates from 40 to 200 Mb/s.

The packaging of the receiver is similar to that of the transmitter, but with special consideration given to EMI/RFI shielding. Given the very high impedance and high gain of the amplifier circuit, this shielding is essential for proper receiver performance in the application environment. The most significant external noise source affecting receiver performance is an electrostatic discharge (ESD) event, in which the receiver is subjected to broadband, high-amplitude electric fields. Testing shows that a complete Faraday shield alone is not sufficient to protect the internal circuits from ESD noise. Theoretically, the Faraday shield prevents an external electric field from establishing an electric field within the shield. The currents generated within the shield induce noise currents in the circuit elements within the shield via magnetic fields. Therefore, it is necessary to minimize this magnetic coupling. This is accomplished by minimizing the area of the critical circuit network and by using a highly permeable material to contain the magnetic flux, preventing it from coupling to the critical circuit network.

- *TRS hardware evaluation*

The high-frequency operation of the TRS and the sensitivity of the optical alignment between the fiber and the component optics required development of new definitions and techniques for hardware evaluation. The electro-optic parameters were measured and correlated to system performance. **Table 1** summarizes the key parameters of the transmitter and receiver portions. This section discusses their definition and dependencies, the evaluation techniques developed, and results.

Transmitter measurements

Optical coupled power is the average optical power coupled from the LED into the fiber core. The coupled power varies not only for different fibers, but also from connection to connection of the same fiber. This variation results from several factors:

- *Fiber core diameter and numerical aperture* Process variations during manufacturing of optical fiber produce variations of the fiber core diameter and numerical aperture. Typical fiber specifications for core diameter and numerical aperture are 62.5 ± 3 μm and 0.275 ± 0.015 μm, respectively. The cable end face polishing also affects the coupled power. The resulting optical power variation is ±0.5 dB.
- *Cladding modes* The optical power in a short fiber contains modes launched into the cladding which are

very lossy. In a data link, this light does not contribute to the signal strength when it reaches the receiver because these modes are not guided and radiate out of the fiber. In measuring the coupled power of a transmitter, these modes can be eliminated either by using a long-fiber cable (>3 meters) or by adding a mandrel wrap in the test fiber, i.e., bending the fiber at a radius that causes the cladding modes to be eliminated within a short distance.

- *Alignment* The *x*/*y* alignment of the focused LED light and the fiber core strongly affects the optical power. This alignment is controlled by several factors:
  - Clearance between the connector ferrule and transmitter sleeve.
  - On-axis alignment of the light spot in the sleeve.
  - Centrality of the fiber core in the connector ferrule.

The coupled power is sensitive to less than 1-μm variations of any of these factors. The technique developed to minimize the sensitivity to misalignment is to overfill the fiber. This approach is being adopted more universally throughout the industry. The axial alignment is less critical, but it is defined by the seating plane of the transmitter and the geometry and polishing of the fiber end face.

- *Transmitter lensing* If the lensing optics and fiber are very close, optical coupling can be maximized; however, it cannot be placed so close as to be in the Fresnel regime of the optical beam, possibly causing unstable results.

Given these considerations, the transmitter is designed to minimize the effects of these variations, and to ensure a 3-km link capability. For hardware evaluation, any fiber used is measured for the mentioned characteristics and these values accounted for in the results. The result is transmitters which nominally couple −17 dBm into a 62.5-μm fiber with variations on any one transmitter less than ±1 dB.

The transmitter *extinction ratio* (*ER*) is defined as

$$ER = 10 \log \left(\frac{\text{optical amplitude of a } 1 \text{ bit}}{\text{optical amplitude of a } 0 \text{ bit}} \right) = 10 \log \left(\frac{P_H}{P_L} \right). \quad (1)$$

Traditionally, for ease of measurement the extinction ratio is done as a dc measurement of the high and low levels, but this does not correlate to the extinction ratio of the transmitter under normal modulated conditions. In the past, ac techniques (representing normal modulated conditions) were avoided because long-wavelength detectors were not stable for the very low *0* power levels due to high dark currents and temperature variations. A technique was developed for measuring the ac extinction ratio by measuring the total optical power at various duty cycles, then extrapolating the 100% (P_H) and 0% (P_L) power levels. These values are used to calculate the extinction ratio (**Figure 2**). If the duty cycle deviates far from 50%, the optical peak power as a function of duty cycle becomes nonlinear because of thermal effects of the LED and bandwidth limitations of both the LED and drive circuitry; for this reason, the dc technique did not correlate to the ac extinction ratio. The 8B-10B code ensures that the duty cycle remains close to 50% and the linear behavior is valid. Under ac conditions, the transmitters have a nominal extinction ratio of 11 dB.

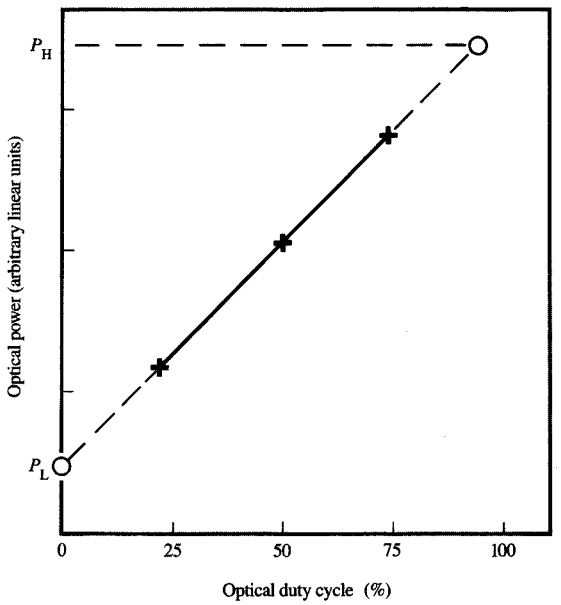

Figure 2

Average optical power as a function of duty cycle: ✚ = data points; ○ = extrapolated power levels.

The *wavelength* and *spectral width* of the LED in the transmitter strongly affect the dispersion characteristics of the signal through the fiber. Instead of the peak wavelength being used as the criteria for the transmitter, the center wavelength between the full width half maximum of the wavelength spectrum is measured, because this is where the energy is centered.

Receiver measurements

The receiver *sensitivity* is the average power of the weakest optical signal the receiver can detect and maintain a specific BER. Several days are required to accumulate

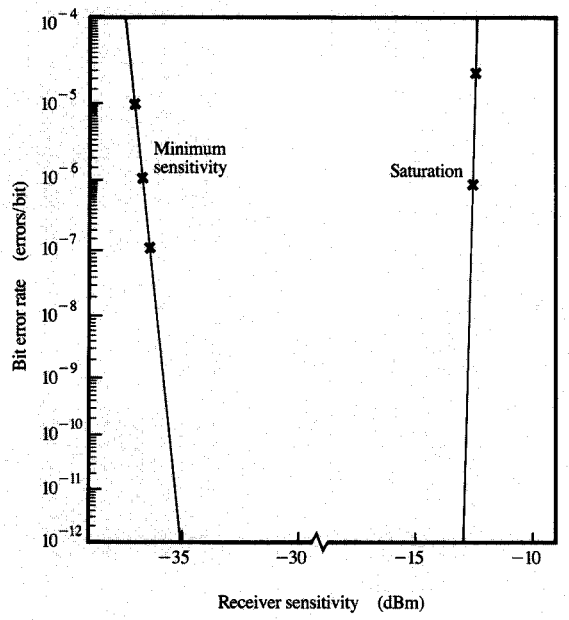

Figure 3

BER vs. receiver optical input power.

enough errors to determine a sensitivity for a very low BER. For example, to accumulate six errors with a BER of 10^{-12}, it takes 8.3 hours. However, if the noise in the receiver is Gaussian, a technique can be used in which data can be taken at high error rates, which takes seconds, and a calculation can be done to determine how much additional signal power would be needed to achieve another level of BER performance. For Gaussian noise, a term Q as defined in Equation (2) can be used to determine the relationship of the input signal strength to the BER. This term is a statistical measure of how distinct the signal levels are from the receiver decision threshold level and is directly proportional to the input signal strength:

$$Q = \frac{|\text{ decision threshold level } - \text{ expected signal level }|}{\text{standard deviation of signal level}}. \quad (2)$$

The BER is related to Q through the complementary error function:

$$BER = \frac{1}{\sqrt{2\pi}} \int_{Q}^{\infty} e^{-x^2/2}\, dx. \quad (3)$$

The relationship between the BER and receiver input optical power can be determined using Equation (3) [7]. A typical plot of BER versus receiver input optical power signal for an ESCON receiver is shown in **Figure 3**.

Theoretically, if the noise is purely Gaussian the slope should be 1.8 dB from 10^{-6} to 10^{-12} BER. If the BER slope differs from the above, it means that the receiver noise is no longer Gaussian. For example, a crosstalk problem was identified and corrected by noticing a very steep BER slope on early prototype hardware. On the other hand, if the BER slope flattens toward the horizontal instead of continuing down toward improved BER, it can mean that the receiver is reaching a noise floor and the noise is no longer Gaussian.

The receiver *saturation* is the maximum average optical signal power the receiver can detect while maintaining the specified BER. Saturation is usually a result of saturating the front-end amplifier of the receiver. In comparison to minimum sensitivity, saturation occurs suddenly as the power increases (Figure 3). Saturation is strongly dependent on the data rate and the spectral content of the data code. The faster the data rate, or the more spectral components in the data code, the easier it is to saturate the receiver. The final receiver design demonstrated approximately a 1-dB difference in saturation from 100 to 200 Mb/s, and a 2-dB difference between a square-wave signal and a $2^7 - 1$ pseudorandom signal. This code, available in most test equipment, simulates the 8B/10B code with regard to dc balance and transition density.

Another aspect of the receiver that is important to the link BER performance is the characteristic of the random and data-dependent jitter, and its effect on the *eye-opening* (data valid time window) of the receiver output signal. The larger the eye-opening, the more accurately the phase-locked loop (PLL) will retime the data (see the section on the serializer/deserializer module). The eye-opening of the receiver is measured by delaying the phase of the sampling clock used to measure the BER and profiling how it changes (**Figure 4**). Typical eye-opening at an optical power of −34 dBm is 1.9 ns for a BER of less than 10^{-12}.

A link status indicator function is included in the receiver module to provide diagnostic support. This circuit detects whether the input signal is too weak (e.g., a connector is loose) or whether the link is inoperative because the signal is latched (i.e., stuck high or low). An intentional hysteresis is designed in this circuit to aid evaluation of the threshold power level so that there is no chattering on this line if the input signal is near the threshold level.

As outlined in **Table 2**, the transmitter has several temperature- and voltage-sensitive parameters. The receiver, however, is relatively insensitive to temperature and voltage variations. All of the temperature variations in the transmitter are due to the LED. As the temperature increases, the LED radiative bandgap energy levels spread, which allows lower-energy bandgap

recombinations (longer wavelengths) and also offers nonradiative recombination paths, causing the LED to generate less light. This is the source of the shifts in the power output, extinction ratio, wavelength, and spectral widths. The only voltage-sensitive parameter is the transmitter transition time, which is smaller at higher voltages.

• *Assembly*

The design of the composite TRS assembly (transmitter–receiver and receptacle) is dictated by the duplex jumper/connector. The intent is to have a duplex connector which connects one channel (composed of a transmit and receive link) in a single operation. The duplex connector is composed of two floating ceramic ferrules (each with a single fiber) in a common housing. The primary function of the TRS receptacle is to guide and retain the duplex connector and ensure proper ferrule seating within the transmitter and receiver. **Figure 5** illustrates the dimensions and tolerances required of the receptacle to ensure connector pluggability and proper ferrule seating.

To accomplish the assembly of the transmitter, receiver, and receptacle, features are incorporated enabling the three components to be pressed together and properly aligned in a single process step. The optical port of each module is orthogonal to the module front body surface within 0.5 degrees. When fully pressed together, the optical ports are precisely located with respect to one another and to the inner top and bottom surfaces of the receptacle. This alignment includes precise lateral and angular tolerances which take into account the amount of movement of the connector ferrule (called connector float). The TRS alignment a is expressed by

$$a = x + l \sin \alpha, \tag{4}$$

where x is the lateral offset of the optical ports from the true position of the connector, l is the effective ferrule length, and α is the angular offset of the optical ports. These factors make up the total optical port offset. To ensure connector pluggability, the ferrule float (see the section on the ESCON jumper cable assembly) must be greater than a.

The module optical port location in the TRS assembly is measured using a CMM (coordinate measurement machine) by placing a 28-mm-long pin with the diameter of a connector ferrule (2.5 mm −0.0012/−0.0020 mm) into the TRS module optical port. The end of the 28-mm pin should fall within a circle of specified radius (Figure 5, Section B-B) with respect to the true center, which is related to the inside receptacle reference surfaces.

The interface of the TRS assembly and the duplex connector is key to the overall functionality of that optical interconnection. Therefore, considerable attention was given to this area during qualification and reliability testing. The major factors addressed during the testing were

- TRS/connector compatibility.
- Stability and repeatability of the optical coupling loss.
- Alignment/dimensional stability.
- Effects of external forces on the TRS/connector interface.
- Environmental compatibility.

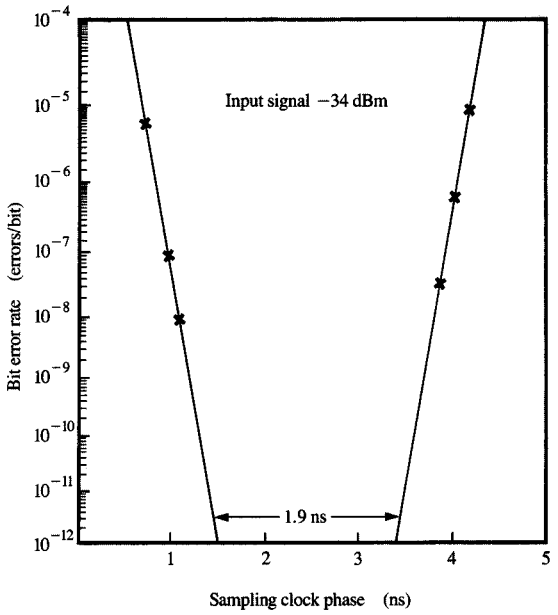

Figure 4

BER vs. sampling clock phase.

Table 2 Temperature and voltage sensitivity values.

Temperature (25 to 60°C)	
Optical power (dB)	−0.5
Extinction ratio (optical) (dB)	−1
Center wavelength (nm)	+3
Spectral width (nm)	+3
Optical rise/fall time (ns)	+0.2
Voltage (4.75 to 5.25 V)	
Optical rise/fall time (ns)	−0.2

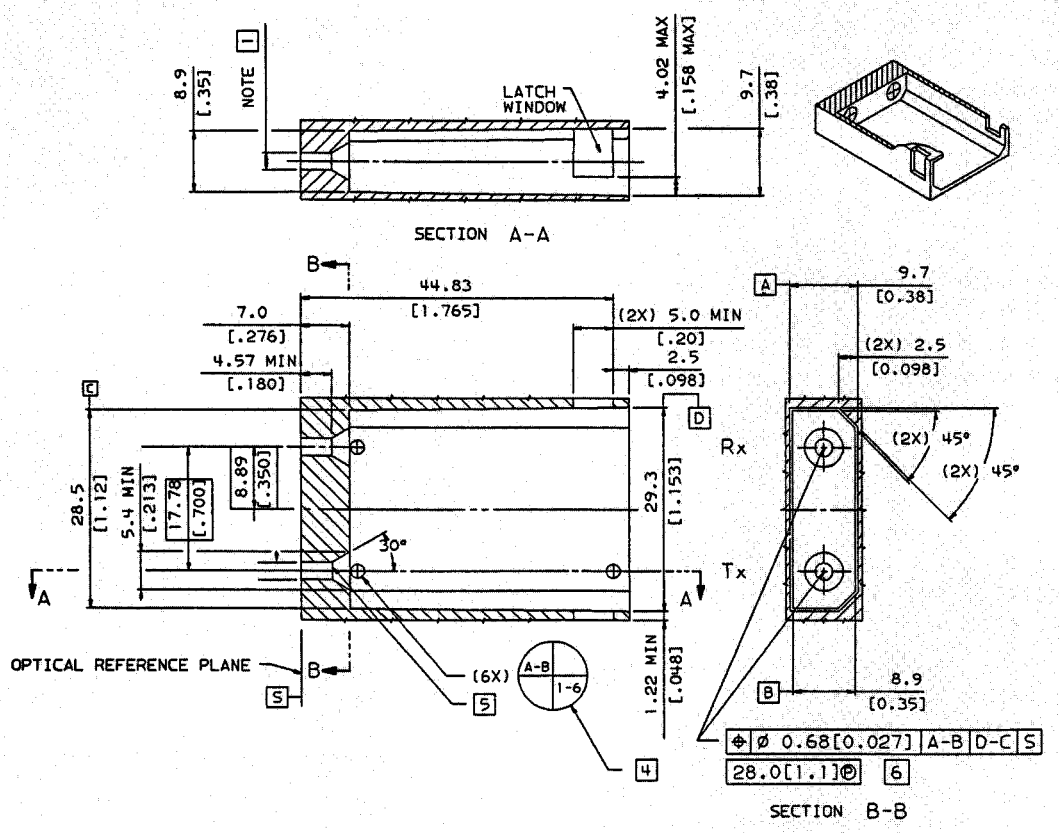

Figure 5

TRS receptacle specification.

Serializer/deserializer (SER/DES) module

The SER/DES is a 28-mm-square module designed and manufactured by IBM. The module contains two 3 × 3-mm silicon chips, 12 chip capacitors, two high-temperature chip resistors, and two trimmable inductors. These components are mounted on a 77-pin metallized ceramic substrate carrier (see **Figure 6**). The module operates off a single 5-V power supply and typically dissipates 1.9 W. The module is capable of operating in an air-cooled environment.

The silicon technology has 2.5-μm ground rules, a unity gain frequency of 3 GHz, 12-V breakdown voltage, and approximately 1300 active devices per chip. The module carrier is a metallized ceramic polyimide (MCP) substrate

with two layers of metal. The first metal layer is predominantly used as a ground plane. The second layer of metal is used for signal distribution. The active laser-trimmable inductors are made from substrate metal using the extra-fine line spacing available in the MCP substrate technology.

In the two-chip SER/DES design, one chip is used for the clock generation and serialization function and the other chip is used for the clock recovery and deserialization function. The circuitry on both chips comprises custom analog and digital gates. The predominant circuit choice for the digital logic is collector-coupled, two-level differential cascode current switch logic. The predominant circuit choice for the analog circuitry is differential current-switch emitter-follower. This circuit family was chosen to minimize power dissipation, maintain pulse symmetry, and reduce switching noise. Typical gate delays are 800 ps.

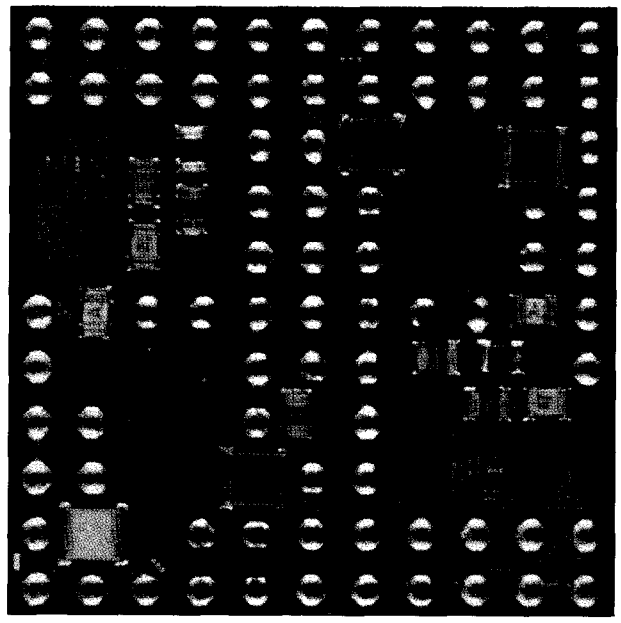

Figure 6

Photograph of SER/DES module.

- *Data encoding*

An 8B/10B code [5] has been chosen for the ESCON link design. A few of the significant advantages of this code are

- High coding efficiency (80%).
- dc balance.
- High transition density.
- Ease of implementation in byte-oriented design.
- Excellent error detection properties.

On the transmitter side the encoder converts eight bits of raw data (byte denoted by *ABCDE–FGH*) into a bit format (denoted by *abcdei–fghj*). The 8B/10B code is broken down into two smaller coding schemes, a 5B/6B bit code and a 3B/4B bit code. The 8B/10B code guarantees a maximum of five contiguous identical bits. This feature greatly simplifies the clock recovery and receiver gain control circuits. Coded bits are always transmitted while the link is in operation. While data are not being transmitted, an idle sequence is sent. This idle sequence is used for both bit and byte synchronization. On the receiver side, the decoder converts the ten bits of coded data back to the original eight bits of raw data. Properties of the code are monitored by the decoder to ensure that the data have been successfully transmitted, received, and regenerated.

- *Serializer design*

A functional block diagram of the SER/DES module is shown in **Figure 7**. The serializer portion of the SER/DES module converts parallel encoded data to a serial NRZ (non-return-to-zero) data stream. The parallel data bus is a 10-bit-wide CMOS TTL interface with each bit typically operating at 20 Mb/s. The parallel data are latched into the serializer in two groups. The first group consists of the six bits a, b, c, d, e, and i. The second group consists of the four bits f, g, h, and j. A parallel word is converted into a serial data stream with the a bit transmitted first and the j bit transmitted last. The serial data interface to the fiber optic transmitter is raised differential emitter-coupled logic (ECL), with typical voltage swings of 3.4 to 4.2 V. The nominal serial data rate is 200 Mb/s.

The transmit clock is generated in two ways. The first method uses a 200-MHz on-card quartz crystal with an on-chip oscillator. This feature is useful for applications when no 20-MHz local clock exists, or where the application requires a nonsynchronous parallel interface. The second method of generating the transmit clock uses a phase-locked loop (PLL) to synthesize a 200-MHz transmit clock from a 20-MHz reference clock. This feature allows the serializer to be synchronous with the local system, eliminating the asynchronous parallel interface inherent in the crystal oscillator design. This significantly reduces complexity in the parallel logic of the ESCON directors. ESCON directors have many outbound (transmit) ports. Synthesizing the transmit clock from the local system clock eliminates the need to buffer data and implement synchronization logic on each transmit port.

The transmit clock drives two flip-flops, each clocked on positive edges of the transmit clock. The first flip-flop is used to clock data to the serial output drivers. The second flip-flop is used to divide the transmit clock by two. The divide-by-two clock is used to drive a ring counter which

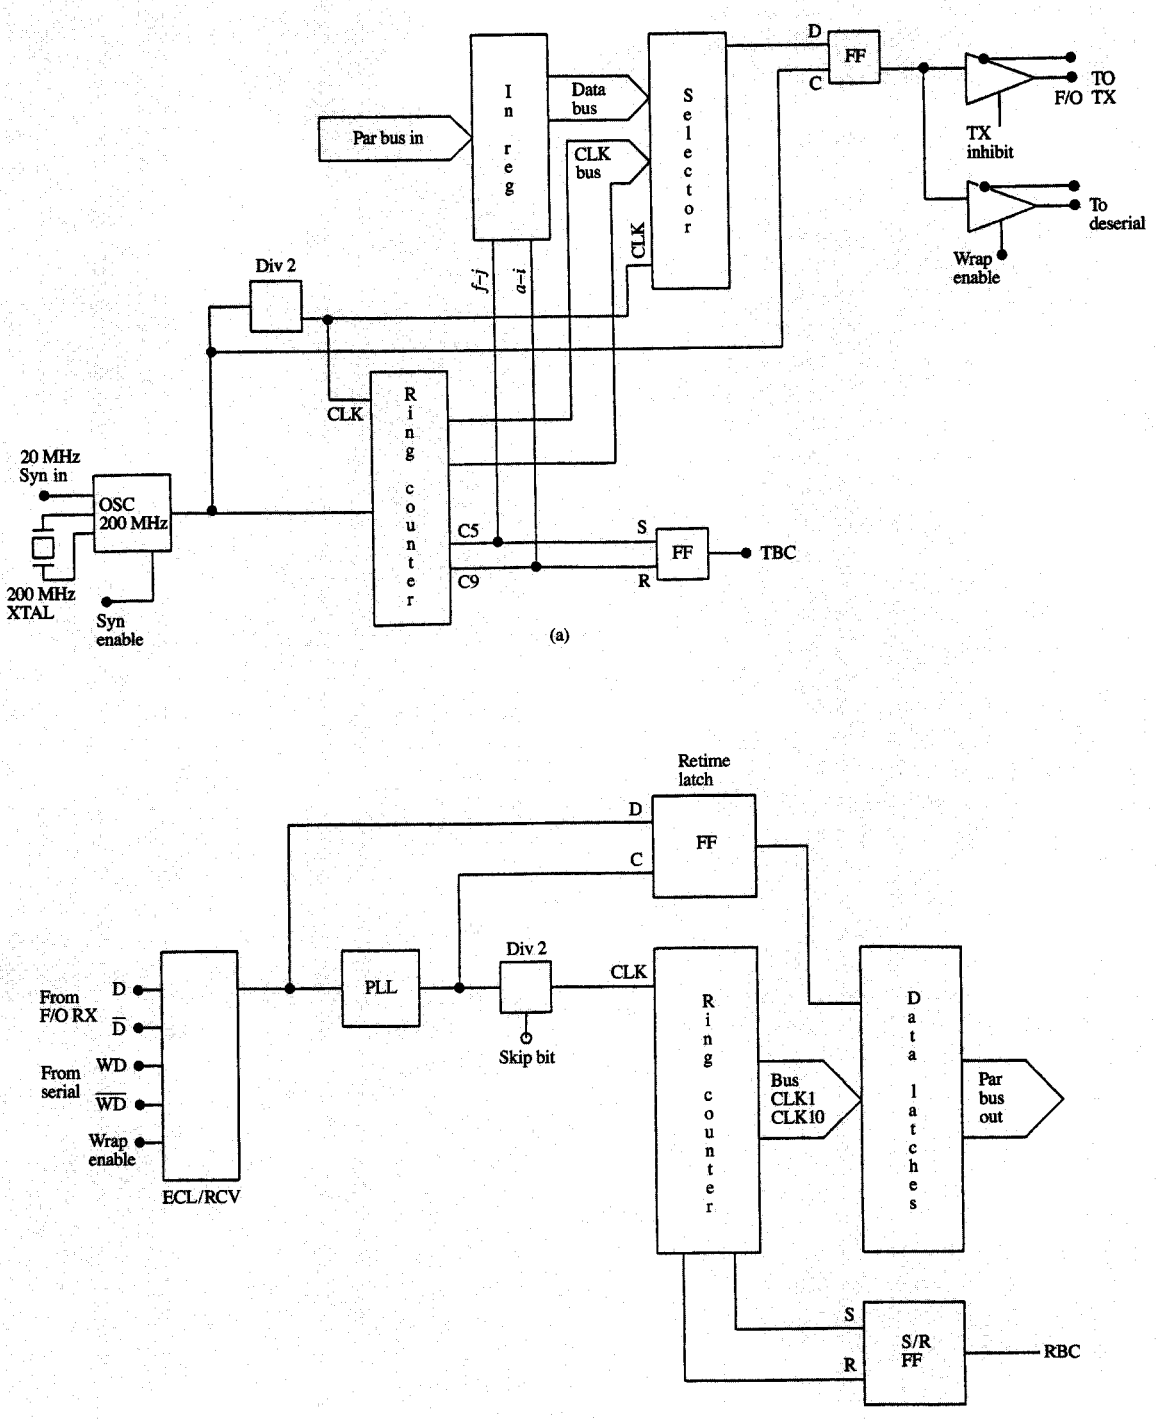

Figure 7

Functional block diagram of SER/DES module: overview of (a) serial chip; (b) deserial chip.

is the heart of both the serializer and deserializer logic. The ring counter consists of five pairs of level-sensitive latches. Each of these pairs has a positive level-sensitive latch that drives a negative level-sensitive latch. All five latch pairs are serially connected and are clocked by a common divide-by-two transmit clock. Two logic one levels are circulated in the ring counter to produce ten synchronous clock pulses which are 10 ns wide and separated in time by 5 ns. These pulses are used to address the data, each pulse causing a data bit to be moved from a parallel data latch to the input of a clocked flip-flop. The clocked flip-flop is used to gate data to the serial output driver synchronously.

There are two differential serial output drivers on the serializer chip. During normal operation one set of drivers is used to drive the fiber optic transmitter. The second set of drivers, when enabled, is directly connected to the adjacent deserializer inputs, thereby achieving an electrical wrap (Figure 7) which is used for local diagnostics. The fiber optic transmitter inputs can be forced to a logical *0* at any time by asserting the inhibit line on the serializer chip.

• *Deserializer design*

Digital data are passed from the fiber optic receiver to the deserializer. The serial data signals are raised ECL differential signals. The deserializer (Figure 7) uses a PLL to derive a clock signal from the incoming data. This synthesized clock is used for retiming the data (sampling the data at the optimum point), thereby removing signal jitter. The PLL clock signal is divided by ten and driven off-chip to be used as a timing reference for the outbound parallel data. The parallel data are passed off-chip in two groups of five bits each. The deserializer does not establish the byte boundaries. However, the bits appear in the order in which they are received. The byte boundary may be aligned with the "skip bit" control line.

The "skip bit" function is a means for establishing byte synchronization. The function is based on error statistics accumulated in the parallel logic as a result of receiving the output from the deserializer. If the byte boundary is not properly established, code violations will occur. In this case the parallel logic will send a pulse to the deserializer, causing the deserializer logic to shift the byte boundary by one bit time. The byte boundary is advanced one bit time by inserting an additional clock edge into the ring counter. This is done synchronously to avoid disruption of the ring counter. This process is repeated until the byte boundary has been properly established. The skip bit function can be disabled by grounding a module input pin so that other forms of byte synchronization can be used.

The output clock from the PLL is used to drive two flip-flops. One flip-flop is clocked in phase, the other out of phase with the PLL clock. The first flip-flop is used to sample (retime) serial data. The second flip-flop is a divide-by-two flip-flop; it drives a ring counter. The ring counter generates ten synchronous clocks, each 10 ns long, separated by 5 ns in time. Each serial bit is addressed by a ring counter output clock pulse. Each pulse causes a serial data bit out of the retime data flip-flop to be moved to a parallel output latch. Parallel addressing significantly reduces circuit speed and critical timing compared with a conventional shift register approach. The data are placed directly on the parallel output interface in two groups of five bits each. The receive byte clock (RBC) is also derived from the ring counter. Hence, there is a fixed relationship between the RBC clock and the parallel data. A falling edge on the RBC output signals that the first five bits are valid at the deserializer interface. A rising edge on the RBC output signals that the second five bits are valid at the deserializer interface.

The deserializer has a two-to-one multiplexer on its serial data input port. During normal operation the data from the fiber optic receiver are manipulated by the deserializer. During diagnostics the serial data input path can be altered to allow data from the adjacent serializer to be passed to the deserializer (electrical wrap) to verify local operation.

• *Clock generation*

The serializer chip contains the serializer logic and the clock generation circuitry. The clock generation may be controlled either by a local 200-MHz crystal or by synthesizing a 200-MHz clock from a 20-MHz external reference clock.

The crystal is contained in a TO-5 package and operates on the seventh overtone with excellent frequency stability and reliability; it is used when an external reference is not available. An LRC tank circuit is used with the crystal, its Q chosen so that it provides sufficient bandwidth to eliminate the problems associated with component drift and aging while being narrow enough to allow only the desired crystal harmonic mode to dominate. The tank inductor is made from substrate metal, and is active-laser-trimmed so that circuit process tolerances may be fully compensated. The tank capacitor is fixed, and is mounted on the ceramic substrate.

The oscillator signal, a sine wave, is clipped differentially to develop a 50% duty cycle, and is gated through a clock selector circuit. The clock selector allows an external clock to be used during chip and module testing. The output of the clock selector is the 200-MHz serializer clock, which is distributed throughout the chip. If an external 20-MHz reference clock is available, the clock generation may be performed by using the synthesizer within the serializer chip. This synthesizer is very similar to the clock recovery circuits described below, except for a decade counter placed in the feedback loop of the phase-locked loop to force synchronism

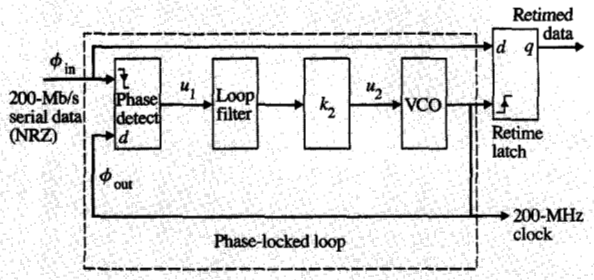

Figure 8

Block diagram of clock recovery design.

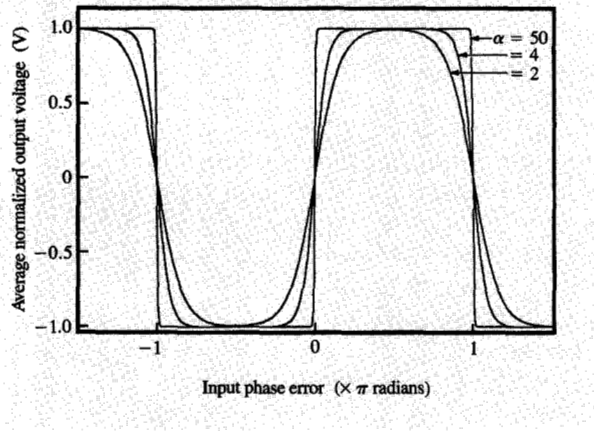

Figure 9

D-type phase detector characteristic.

between the external clock and the character boundary of the data.

- *Clock recovery*

The two major techniques in use for clock recovery in fiber optic transmission systems use either surface acoustic wave (SAW) filters or PLLs. Although PLLs have not traditionally been popular in systems operating above approximately 50 MHz because of design difficulties [8], the PLL has overwhelming advantages over the SAW in flexibility, integrability, and cost. Furthermore, many advances in circuit simulation tools and techniques have been developed in recent years which have significantly increased our ability to analyze and design complex systems, including PLLs (see [9–11] and footnote 1).

The clock recovery system resides on the deserializer chip with the retiming latch and the deserializer logic. A PLL is used to generate a 200-MHz clock which is in phase with the incoming serial data stream. This clock is used for retiming the data, which may have large amounts of jitter contributed by preceding portions of the link, and for subsequent use as a deserializer clock. The degree to which the jitter can be removed from the data is directly related to the retiming penalty, or the optical equivalent link power penalty between an ideal clock (i.e., the transmit clock) and the recovered clock. The retiming penalty is one of many parameters which is included in the link loss budget (see the section on link system design), since it limits the capability of the link to some extent, and is designed to be less than 1.0 dB in the worst case. A block diagram of the clock recovery system is shown in **Figure 8**.

Phase detector

A D-type flip-flop is used as the phase detector for the PLL, which operates as an early-late detector between the input serial data and the output from the voltage-controlled oscillator (VCO); for a steady-state locked condition, the negative transitions of the VCO will align with the serial data transitions. A differential cascode current switch structure is used in the phase detector to exploit the low setup and hold times, the low power, and the relative simplicity that this circuit topology provides in contrast to other techniques. Since the 200-Mb/s NRZ data have a maximum spectral component at 100 MHz, the PLL must effectively double this component to create a 200-MHz clock signal. Since the D-type phase detector is sensitive to harmonics of the input data, there is no need for data preprocessing or feedback counters to double the frequency.

Since the D-type phase detector is essentially sampling the state of the VCO at every negative transition of the serial input data, the output of the phase detector is constant during intervals when no data transitions are occurring. The 8B/10B encoding ensures that there will be a high density of transitions because of its dc balancing features, limiting the run length to five and the maximum interval between like transitions to ten-bit intervals.

With reference to Figure 8, the loop phase error $\phi(t) = \phi_{in}(t) - \phi_{out}(t)$ is simply the difference between the phase of the input signal ϕ_{in} and the phase of the output from the VCO ϕ_{out}. The output of the phase detector is $u_1(t) = k_1 f_1[\phi(t)]$, where $f_1[\phi(t)]$ represents the normalized phase detector characteristic and k_1 is the output amplitude limit. We can also define the loop

[1] Marco L. Padula and Terrence R. Scott, "Performing Convolution in ASTAP," IBM internal memo, March 14, 1989.

frequency error $\dot{\phi} = d\phi/dt$ and the loop acceleration error $\ddot{\phi} = d^2\phi/dt^2$ to be the first and second derivatives of the phase error, respectively.

Figure 9 shows the average normalized output of a D-type phase detector as a function of phase error. This response may be modeled as the sum of three hyperbolic tangents,

$$u_1(t) = k_1\{\tanh[\alpha\phi(t)] - \tanh[\alpha(\phi(t) - \pi)] - \tanh[\alpha(\phi(t) + \pi)]\}, \quad (5)$$

where the constant α, defined as the phase detector slope factor, controls the slope of the characteristic at phase errors approaching zero and $\pm\pi$ radians.

Filter

The output of the phase detector is connected to the PLL filter, which is a first-order differential lag-lead with the transfer function

$$F_2(s) = \frac{(1 + s\tau_2)}{(1 + s\tau_1)}, \quad (6)$$

where τ_1 and τ_2 represent the time constants corresponding respectively to the pole and zero frequencies of the filter. The filter capacitor is mounted on the MCP substrate near the deserializer chip.

Referring again to Figure 8, the input to the VCO u_2 is

$$u_2(t) = k_2 u_1(t) \otimes f_2(t), \quad (7)$$

where $\otimes$ denotes convolution, $f_2(t)$ represents the impulse response of the filter determined by the inverse Laplace transform of (6), and k_2 is a constant which represents an additional linear gain stage within the loop.

Voltage-controlled oscillator

The VCO is designed to operate over a small range about the 200-MHz center frequency. The VCO circuit uses differential current switches and an LRC tank circuit with an active-laser-trimmed inductor to adjust the center frequency. The differential voltage at the input to the VCO changes the bias current through a diode connected in series with a small capacitor, which together are shunted with the tank circuit. An increase in the diode bias current increases the effective capacitance of the tank circuit, causing a lowering of the resonance and of the operating frequency.

Referring once again to Figure 8,

$$\frac{d\phi_{out}(t)}{dt} = k_3 u_2(t). \quad (8)$$

In the implementation of the VCO design, its output frequency f_{out} is controlled by a current switch characteristic,

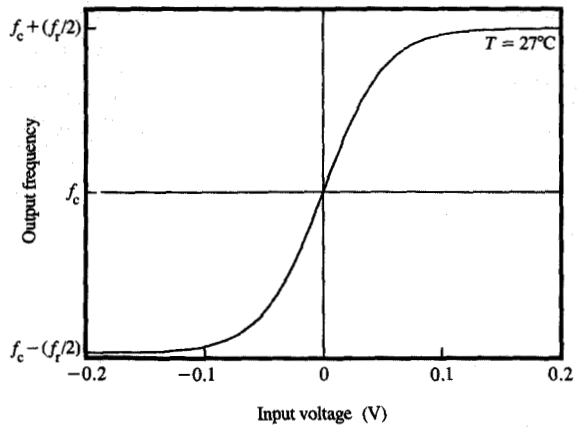

Figure 10

VCO frequency vs. voltage characteristic.

$$f_{out} = \frac{\omega_{out}}{2\pi} = \frac{1}{2\pi}\frac{d\phi_{out}}{dt} = f_c + \frac{f_r}{2}\tanh\left[\frac{u_2}{2V_t}\right], \quad (9)$$

where f_c is the VCO center frequency, f_r is the total VCO range, and $V_t = kT/q$ is the threshold voltage. **Figure 10** shows the frequency versus the voltage characteristic of the VCO. Note that the VCO has a frequency-limiting feature controlling its range of operation and that the VCO characteristic becomes progressively more nonlinear as the operating point moves away from its center frequency.

The output of the VCO goes to a clock selector circuit (not shown), which can be used for testing the circuit with an external clock signal. The output of the clock selector is fed back to the phase detector and is also used an an input by the retime latch (with the inputs complemented to trigger on positive transitions) and the deserializer logic.

PLL dynamics

Nonlinear model Finally, it can be shown [12] that

$$\ddot{\phi} + \frac{k_1k_2k_3\tau_2\alpha\{\text{sech}^2[\alpha\phi] - \text{sech}^2[\alpha(\phi - \pi)] - \text{sech}^2[\alpha(\phi + \pi)]\} + 1}{\tau_1}\dot{\phi} + \frac{k_1k_2k_3\{\tanh[\alpha\phi] - \tanh[\alpha(\phi - \pi)] - \tanh[\alpha(\phi + \pi)]\} - \Omega_0}{\tau_1} = 0, \quad (10)$$

where the initial frequency difference Ω_0 is

$$\Omega_0 = \tau_1\frac{d^2\phi_{in}}{dt^2} + \frac{d\phi_{in}}{dt} \quad (11)$$

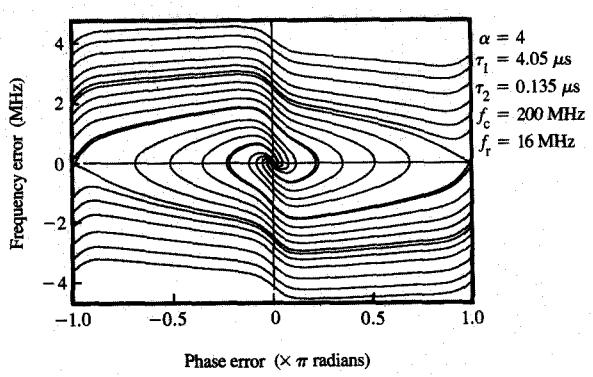

Figure 11

PLL phase-plane portrait solution.

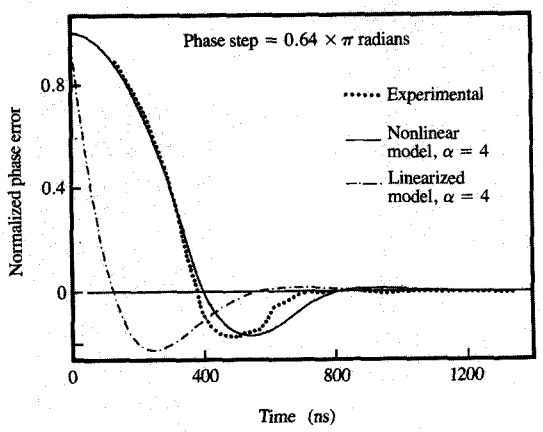

Figure 12

Transient phase error due to large step change in phase.

and

$$k_3 = \left[\frac{\pi f_r}{2V_t}\left\{1 - \frac{1}{\pi^2 f_r^2}[\dot{\phi}]^2\right\}\right] \text{ for } -1 < \frac{\dot{\phi}}{\pi f_r} < 1, \tag{12}$$

which together describe the general time domain loop response. Although this second-order nonlinear differential equation has no closed-form analytic solution, it may be solved numerically with the ASTAP (Advanced Statistical Analysis Program) [12] circuit analysis program using circuit analogies [11], or with some other simulation language. The solution to the loop equation for a specific design point may be presented as a phase-plane portrait, which shows a series of trajectories of the frequency error $\dot{\phi}$ as a function of phase error ϕ (**Figure 11**). This type of diagram is useful for showing the dynamic behavior of a PLL in response to transient phase and frequency error [11].

Linear model Under conditions of small phase error, the general time-domain loop response equation may be reduced to a linear second-order differential equation and solved, producing a linearized model of the PLL [11]. The linearized transfer function of the loop becomes

$$H(s) = \frac{\phi_{out}(s)}{\phi_{in}(s)} = \frac{\alpha k_0}{\tau_1} \frac{1 + \tau_2 s}{s^2 + \frac{1 + \alpha k_0 \tau_2}{\tau_1} s + \frac{\alpha k_0}{\tau_1}}, \tag{13}$$

where peak loop gain $k_0 = k_1 k_2 \frac{\pi f_r}{2V_t}$,

natural frequency $\omega_n = \sqrt{\frac{\alpha k_0}{\tau_1}}$,

and damping $\zeta = \frac{1}{2}\omega_n\left(\frac{1}{\alpha k_0} + \tau_2\right)$.

The loop error due to a unit step change in phase is

$$\phi(s) = \frac{1 - H(s)}{s} = \frac{s + \frac{\omega_n^2}{\alpha k_0}}{s^2 + 2\zeta\omega_n s + \omega_n^2}, \tag{14}$$

which in the time domain for $\zeta < 1$ becomes

$$\phi(t) = e^{-\zeta\omega_n t}\left(\cos\omega_n\sqrt{1-\zeta^2}\,t + \frac{\frac{\omega_n}{\alpha k_0} - \zeta}{\sqrt{1-\zeta^2}}\sin\omega_n\sqrt{1-\zeta^2}\,t\right). \tag{15}$$

Figure 12 shows the linearized solution, the nonlinear solution, and experimental data for the transient phase error response due to a large (0.64π radians) step change in phase. This example clearly shows the limitations in using a linearized approach, although for small phase shifts the agreement between the methods improves significantly.

PLL transfer function model

Although conventional methods may be used for simulating the individual components of the PLL, some special methods are required for analysis at a system level. The preceding analysis of the PLL using the loop equation gives us great insight into the dynamics of the PLL, but it is difficult to extend these techniques into all of the areas which warrant study. The PLL presents a unique problem

for traditional circuit simulation programs because as a detailed circuit model it generally requires both very small step sizes for convergence because of the oscillator characteristics and very long transient charge-up periods due to the low-pass nature of the filter and the loop itself; it is not uncommon for several hundred thousand to well over a million passes to be required for a single nominal PLL simulation. By modeling the PLL with transfer function descriptions of the components rather than complete circuit models, a thousandfold efficiency improvement may be realized in ASTAP. This enormous efficiency improvement comes largely as a result of the exploitation of the built-in FORTRAN and FORTRAN-like function libraries within ASTAP [10].

Figures 13 and **14** show examples of some of the simulation results of a PLL acquisition process using transfer function models. In this example the VCO generates a signal with zero rise and fall time which can easily be distinguished from the serial data idle sequence. The loop is initially unlocked and has an initial frequency error of −5 MHz. Figure 13 shows the serial data and the recovered clock as a function of time early in the acquisition process. As time advances, the VCO frequency and phase are adjusted until they are matched, with some small but finite steady-state error, to the input serial data. Figure 14 shows the phase error as a function of time throughout the acquisition process. Performing simulations of this type provides both the transient and steady-state response of the PLL to a virtually limitless set of initial conditions, allowing full parametric characterization.

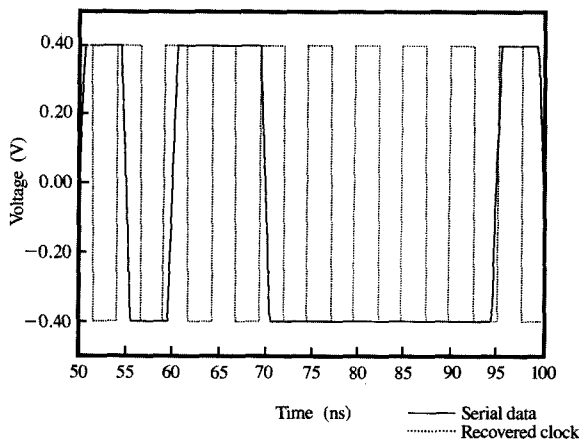

Figure 13

Simulated PLL acquisition process (PLL initially not locked).

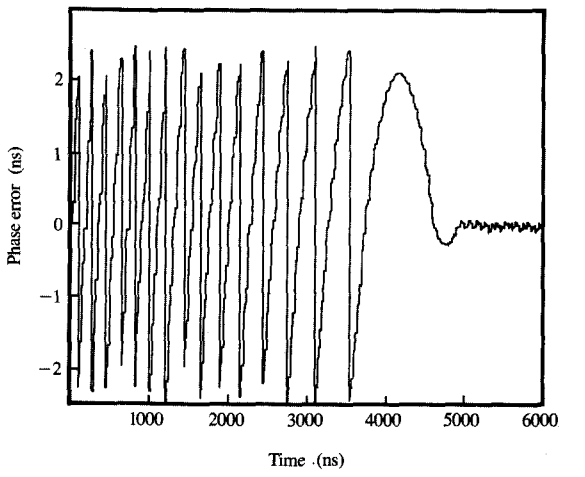

Figure 14

Simulated PLL phase error as a function of time.

- *Characterization of the SER/DES module*

Functional performance testing of the module was done in the laboratory in several different ways. A significant amount of testing determined what portion of the optical loss budget should be allocated for the SER/DES function. IBM developed a custom tester to aid in the evaluation of the SER/DES and other fiber optic link components. Two of the significant features of the IBM fiber optic data link tester are the following:

- The tester can identify the exact bit in error and its exact location within the data block. This is helpful in the study of systematic errors and in verifying hypotheses concerning which bits in a sequence are most likely to produce errors.
- The tester is capable of generating unique patterns to analyze data-dependency effects.

The SER/DES was fully tested using this custom tester. The BER is measured at the parallel interface of the deserializer.

The sensitivity of the fiber optic receiver is first characterized by plotting a BER curve, as discussed previously. The receiver is then mounted on a card with the SER/DES module. An output test point is designed into the SER/DES module that allows the serial data to be brought out of the module after being sampled by the PLL but before being passed to the logic. This test point is then used to generate a BER curve. This new BER curve is compared to the receiver-sensitivity BER curve. The

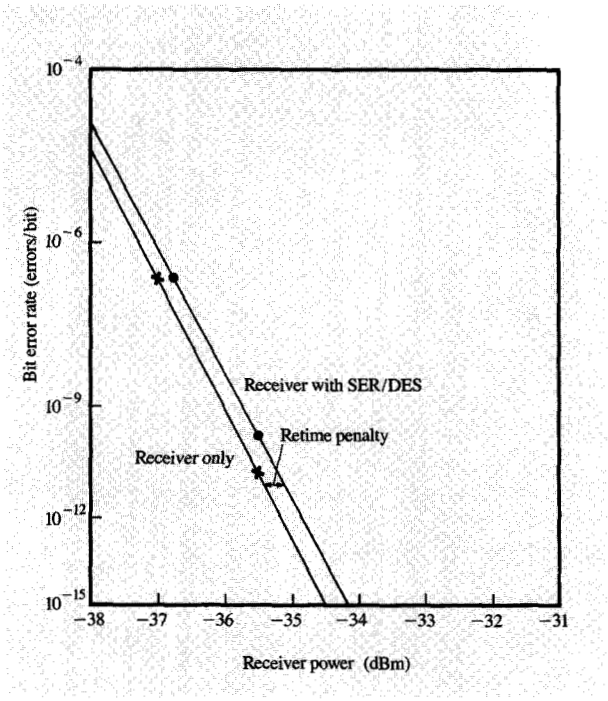

Figure 15

Link BER graph with/without SER/DES module.

difference in optical power between these two curves is the retime penalty (see **Figure 15**). This testing yields a figure of merit in optical power for the SER/DES function that is accounted for in the optical link budget.

Extensive pattern testing was also conducted to ensure that no significant pattern sensitivities existed in the fiber optic receiver, PLL, or SER/DES logic and to quantify the data dependencies that do exist. The fiber optic receiver was tested by generating families of BER curves as a function of the data pattern. The difference in optical power required to obtain a specified BER for various data patterns is used to define a term denoted "data-dependent penalty" which is also included in the link loss budget. The PLL was tested in a similar manner by generating these same families of BER curves as a function of the data pattern. These curves are compared to the BER curves of the receiver. This test yielded a figure of merit in optical power for the PLL performance as a function of data patterns. Families of BER curves are generated at the parallel interface of the deserializer to ensure that no additional penalty exists. These test results are incorporated as parameters in the optical link budget. Finally, the average VCO input voltage of the PLL is monitored to examine the effects of interchip noise (digital-to-analog interaction) and data dependency on the individual PLL circuits.

ESCON jumper cable assembly

The data processing environment places unique demands on the ESCON jumper cable assembly (duplex fiber optic connector and cable):

- Ease of use.
- Ease of cable reconfiguration.
- "Under the machine room floor" environment.
- Ruggedness.
- Tolerance of environmental dust.
- Ease of cleaning.
- Direct card connection.
- High reliability.
- Reduction in size from that of current parallel copper cable.

Special design characteristics are included in both the connector and the cable to meet these requirements:

- The ESCON connector is a polarized, duplex connector which mates to the TRS in a single plug operation. This is particularly important with large systems, since there could be many jumper cable assemblies at the channel and director interfaces.
- The ESCON connector is a low-profile connector allowing high-density, direct-card plugging.
- The connector latches and ferrules are protected during installation and handling by designing the latches into the connector body and including a retracting protective cap for the ferrules.
- The fiber optic cable is a "unitary strength member" cable with two tightly buffered fibers. This construction is rugged enough to be routed under the computer room floor without requiring special conduits or raceways.
- The small diameter, light weight, and pliable properties of the cable allow increased flexibility for layout and installation.
- The two ferrules within the connector are designed to float independently. This alleviates the TRS alignment requirements and eases the connector plugging force.

- *ESCON duplex cable assembly*

Figure 16 shows a comparison of the parallel copper cable (bus and tag) [1] with the ESCON jumper cable assembly. For each channel, the fiber optic cable provides a reduction of one cable and two connectors when compared with the bus and tag cable, and the weight reduction ratio is 74 to 1. **Figure 17** shows the cross section of the ESCON fiber optic cable and its optical and physical characteristics. It also offers high data rate and distance advantages, as discussed in the Introduction. There are two multimode glass fibers that are protected by a buffer coating; the two buffered fibers are surrounded by aramid fiber strength members, and together are encased in an

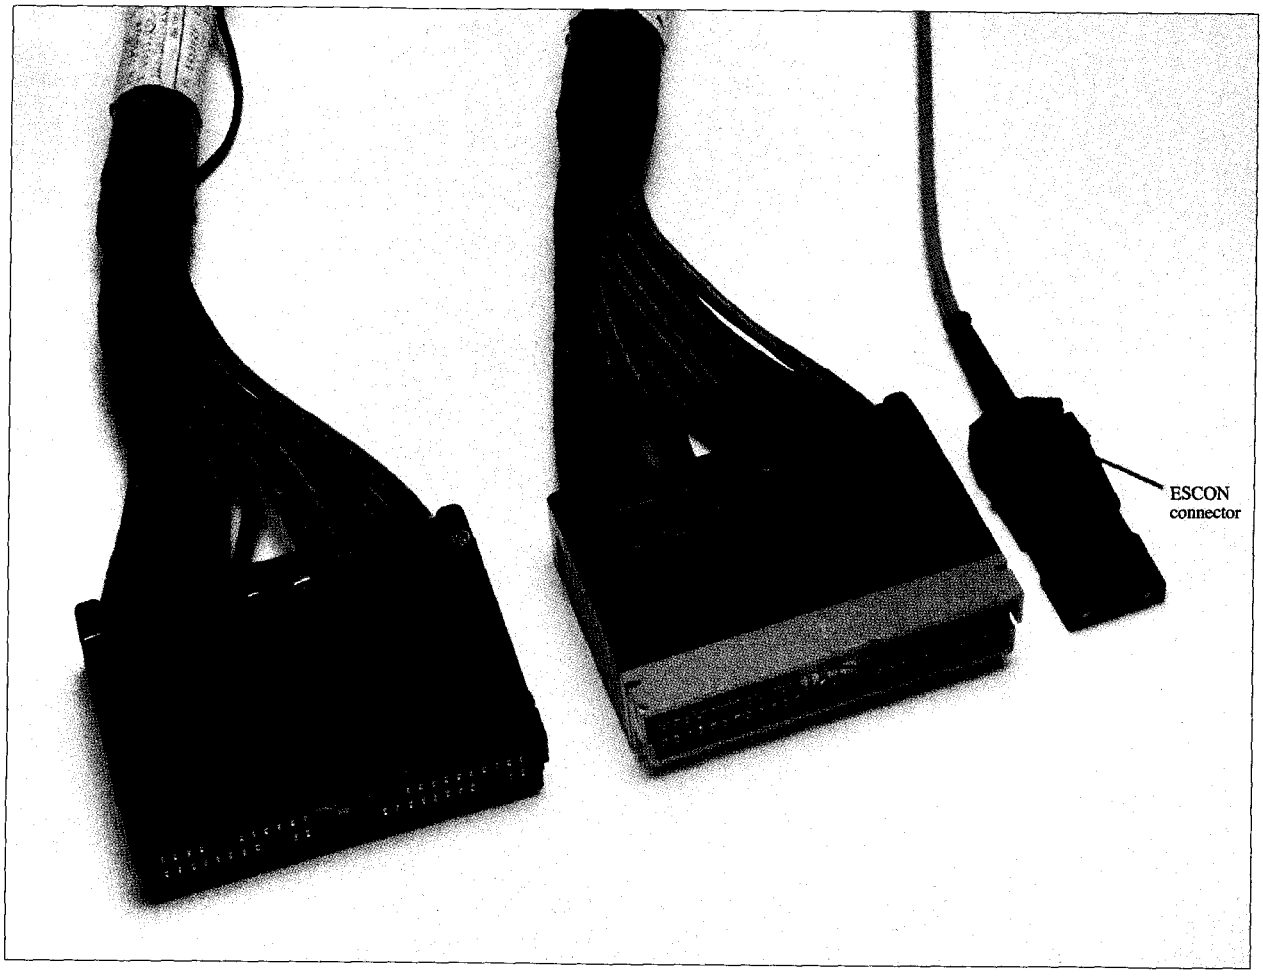

Figure 16

Comparison of parallel copper cable and ESCON jumper cable assembly.

outer jacket material. This construction results in a package capable of withstanding tensile loads of 1000 newtons during installation, as well as impact resistance of ten impacts on a hammer radius of 12.5 mm with a potential energy of 4.5 newton-meters.

Figure 18 shows the uncovered duplex connector illustrating the fiber optic cable termination. The outer jacket material is stripped away, and the aramid fiber is secured to a metal strain-relief member of the connector. This ensures that any mechanical force applied to the cable is translated to the connector body and not the glass fibers, providing a rugged connector cable interface. The two coated fibers are routed through the connector body in an S-bend to provide enough extra fiber in the connector body to prevent the bend radius of glass fiber from being too small. The connector body length is primarily determined by the requirement for the S-bend limitation. The fibers are then terminated in their respective zirconia ceramic ferrule.

In the fabrication process, the buffer coating is removed, and the bare glass is inserted into the small hole (microhole) of the ceramic ferrule and secured with epoxy. The end face of the ceramic ferrule is then polished to a precise radius to ensure glass-to-glass physical contact in a connection. This reduces the optical loss and results in a lower and more repeatable link loss. The final assembly of the connector is completed without tools by pressing the plastic cover and case together. A spring-loaded protective cap is placed on the end of the connector to provide mechanical protection for the ceramic ferrules; it can also be moved back to allow the ferrules to be cleaned before being plugged into a TRS or coupler.

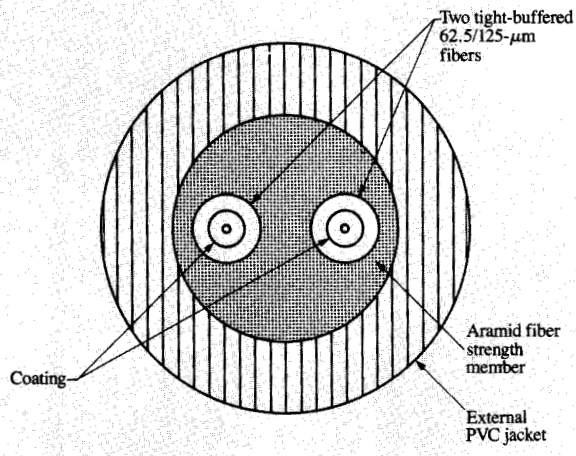

Bulk cable cross section

Cable characteristics

Core diameter — 62.5 μm
Cladding diameter — 125 μm
Numerical aperture — 0.275
Maximum attenuation — 1.75 dB/km
Minimum bandwidth — 1300 nm, 500 MHz·km
Maximum tensile load at 75-mm bend radius — 1000 N

Figure 17

Physical and optical characteristics of the ESCON cable.

- *TRS/connector pluggability*

A significant requirement of the duplex connector is the ability of the connector to be easily plugged into a TRS. This is accomplished in several different ways, the first of which is the tapered design of the TRS receptacle. The connector is designed to match the taper, resulting in ease of plugging. Once the connector is plugged, any movement of the cable will not cause significant optical loss variations in the link.

The second design consideration for ease of plugging is the allowance for the ferrule assemblies to float in the connector plastic housing. The float is the physical distance between the ferrule assembly and the plastic housing; it allows both lateral and angular movement. It is important that the connector float be greater than the allowed mechanical alignment tolerance of the TRS optical subassembly (described in the section on the transmitter/receiver subassembly).

- *Connector optical characteristics*

The duplex connector assembly provides an optical interface to two different plugging options.

First, it can be plugged into a TRS module. The ferrule end face is plugged into a precision bore and comes to rest against a mechanical stop in the module. In this type of plugging, it is important for the fiber to be centered in the opening so that the transmitted light can be focused on the fiber for maximum coupling of the light into the fiber. The coupling repeatability of the TRS-to-duplex optical connection is primarily a function of the tolerances between the connector ferrule and the transmitter and receiver alignment sleeves (2.7-μm nominal spacing), the sleeve alignment to the LED or photodiode, the fiber core alignment to the ferrule, and the residual lateral force the connector ferrule exerts on the sleeve. **Figure 19** illustrates the optical coupling variation for a transmitter-to-duplex connection. The receiver variation is significantly better, typically ±0.2 dB max.

Second, the duplex connector can be mated to a corresponding duplex connector via a coupler. The connection loss of mated randomly selected cable

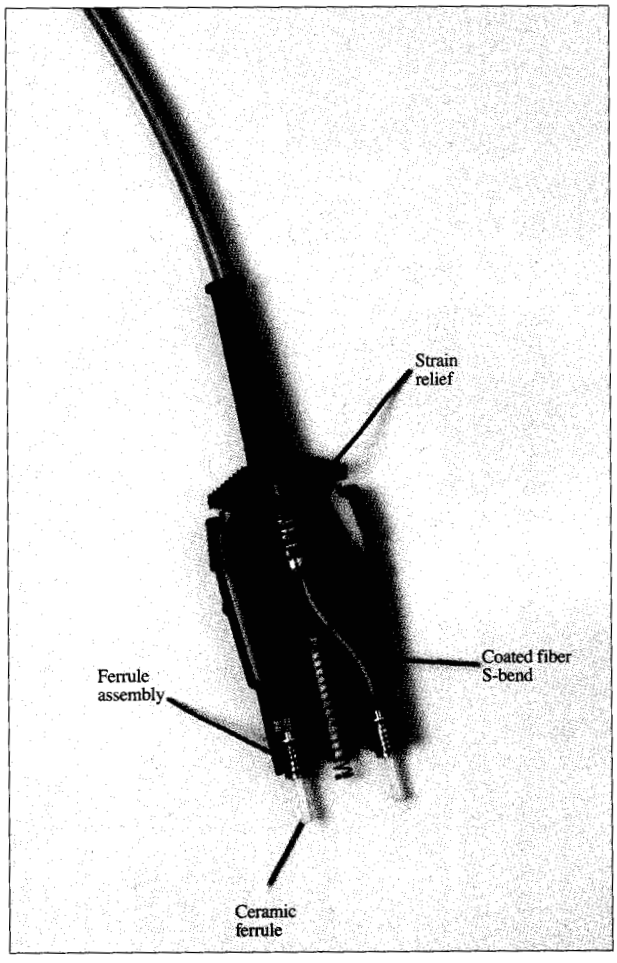

Figure 18

ESCON duplex connector.

assemblies is a function not only of the connector (both the ferrule and fiber) but also of the alignment sleeve in the coupling mechanism that keeps the ferrules aligned and in physical contact. The manufacturing tolerances of the ferrules, sleeves, centrality of fiber core within the ferrule, and fiber contribute significantly to the connection loss. A histogram of the optical connection loss for random ESCON cables is shown in **Figure 20**.

Extensive qualification and reliability testing has been conducted on the jumper cable assembly and associated optical components. A test matrix of more than 350 cable assemblies was used to ensure that the cable assembly met its reliability objectives.

Link system design

A common technique used in the industry to determine the allowable link loss is to directly add up the individual optical losses. An additional optical loss margin, generally 3 dB, is added to account for contingencies and unknowns. The total link loss must be less than the available power from the respective transmitter and receiver pair.

Another technique for determining allowable link loss is to statistically sum the individual link loss elements and select an acceptance criterion as the mean, plus some number of standard deviations of the distribution (e.g., $\mu + n\sigma$). This maximum loss must be less than the available power.

There are significant advantages in using a statistically based loss budget analysis. The main goal of the analysis is to accurately describe the link performance in the field. The statistical technique also allows the link designer more flexibility because the link performance reliability can be tailored to the application. The increased accuracy gained by using link distributions allows better focus on the design optimization. It also makes it possible to more accurately assess cost/performance trade-offs.

The approach to determining link loss in the ESCON link design is based on a statistical characterization of the link parameters. The characterization is grouped into two areas:

1. The installation loss budget includes cable attenuation, connectors, and splices.
2. The available power budget includes the transmitter power, receiver sensitivity, BER margin, coupling variation, data dependence, EOL (end of life), retiming penalty, and mode equalization.

The available power budget statistically determines the maximum allowable link loss specification (the -3σ value of the available power distribution). The installation loss budget determines whether one can build the fiber optic cable plant to meet the maximum allowable link loss specification. Assumptions are made about the jumper length and the number of connectors and splices used at the maximum distance.

The statistical approach requires the components to be described as statistical distributions. These distributions may also be truncated to reflect a selection from a larger population. These parameters are incorporated into a

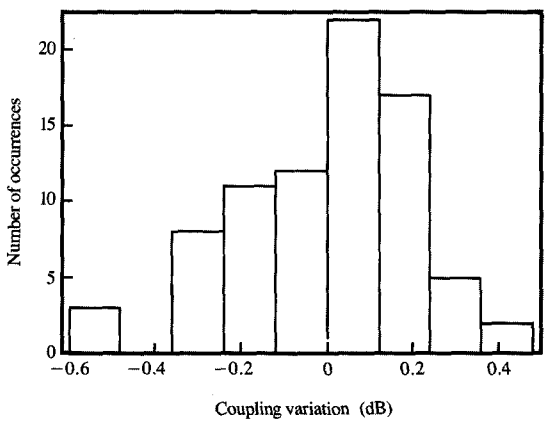

Figure 19

TRS/connector coupling variation (sample size = 100).

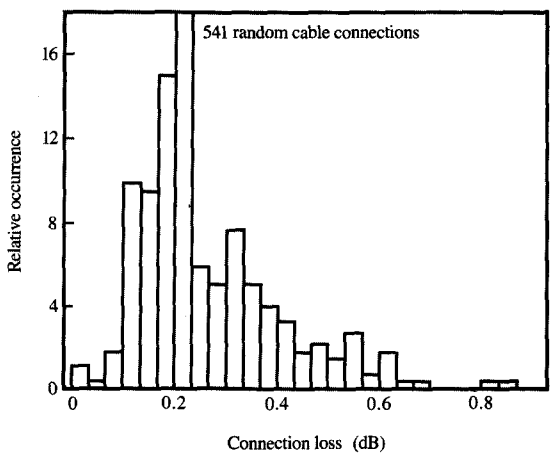

Figure 20

Mated connection loss of randomly selected ESCON cable assemblies.

Table 3 Maximum link loss (at 1300-nm wavelength).

Maximum link length	*Maximum link loss* (dB)	*Trunk size* (μm)	*Minimum trunk modal bandwidth* (MHz · km)
2	8.0	62.5	500
2	8.0	50.0	800
3	8.0	62.5	800

model which carries out a Monte Carlo simulation and combines these components to determine the link margin. The approach is based on data which reflect worst-case temperature and power supply voltage.

The loss budget parameters used in the determination of the available optical power are described below:

Transmitter power The average optical power coupled into a reference 62.5/125-μm jumper.

Receiver sensitivity Specified at a 10^{-12} BER as measured at the center of the bit (see the section on the transmitter–receiver subassembly).

TRS variability Accounts for the difference in transmitter-coupled optical power and received optical power between using a reference jumper and a normal jumper. This includes such factors as optical alignment tolerances, ferrule tolerances, different fiber parameters (NA, core size, etc.), and optical spot distortion.

Retiming penalty The effective optical power penalty of the link between an ideal clock and the recovered clock (see the section on the serializer/deserializer module).

Data dependence penalty Since the receiver sensitivity is characterized with a $2^7 - 1$ pseudorandom pattern, this parameter accounts for the possible degradation for sending any 8/10 encoded data.

Stressed jumper variation The application allows for tension to be applied to the jumper while it is in use in a TRS. This parameter accounts for the additional optical loss.

Tailgate connection loss This parameter accounts for the loss of the additional duplex connection within the product.

Link loss measurement uncertainty Link optical attenuation is measured using a substitution technique similar to FOTP-171 [13]; however, reference jumpers are not used. This parameter accounts for the possibility that a link with more than 8 dB of loss measures less than 8 dB.

Transmitter EOL degradation Accounts for the degradation in output optical power of the transmitter over its lifetime.

Source/cable variation This depends on the design of the launching optics in the transmitter. It relates the optical loss of a fiber link using the optical mode conditioner (OMC, see Section 6) tool to the link loss as measured with a product transmitter. Part of this difference is included in the installation loss budget as higher-order-mode loss; the remainder is included in the available power budget as source/cable variation.

Higher-order-mode loss (HOML) The additional optical attenuation of a fiber which occurs during the first few hundred meters of cable. This loss occurs with product transmitters; it does not occur when the OMC tool is used. Therefore, the HOML must be added to link loss measurements which require the use of the OMC tool.

Dispersion penalty This accounts for the additional optical power needed by the receiver to compensate for signal degradation due to the chromatic and modal bandwidth limitations [14] of the fiber.

The following transmitter parameters are very important to the link performance; their effects are included in the above loss budget parameters:

- Optical extinction ratio.
- Optical rise/fall time.
- Center wavelength.
- Spectral width.

Combining the loss budget parameters at each supported distance by means of a Monte Carlo simulation yields the minimum available power budget, which is equivalent to the maximum allowable link loss. The results of the simulation are used to determine the parameters and specifications for links using 62.5/125 or 50/125-μm fiber trunk cable. An analysis of trade-offs among link distance, trunk modal bandwidth, and maximum link loss was performed to yield the configuration specifications in **Table 3**. [The link length shown in the table includes both jumper and trunk cables. The maximum total jumper length is 244 meters (800 ft) when jumpers are used with a trunk.]

Table 4 presents an example of a typical installation loss budget for a common link configuration. The table shows that a 3σ design criterion can be maintained for each fiber type and link configuration at the maximum allowed distances.

Application support for ESCON links

Fiber optic serial links require application support in the areas of link loss testing, transceiver optical wrap, cable management, and cable strain relief to ensure their performance and reliability in the field.

- *Link loss testing*

It is well established that the effective modal volume (EMV) carried by a fiber affects attenuation and connection loss measurements. The telecommunications industry has established standards for the proper launch conditions for these measurements. These are an approximation of the equilibrium (steady-state) mode power distribution (EMD) existing after transmission through long lengths of fiber. However, the EMV carried in short links may never reach the EMD. It will likely have a mode power distribution that depends on the fiber mode excitation distribution from the transmitter. Therefore, the link loss measurement will be different for different transmitters.

Three techniques are specified by FOTP-50 [15] (dummy-fiber, mandrel-wrap mode filter, and limited-beam optics launch) for obtaining the required EMD launch. However, none are convenient for field application.

The IBM optical mode conditioner (OMC) hand-held tool is designed for field application, to achieve loss test results equivalent to those obtained with EMD conditions in a fiber optic link. Through spatial overfilling and angular underfilling, the IBM OMC tool achieves source isolation and correlation to EMD measurements for all link configurations and sources in a single, cost-effective package.

- *Transceiver optical wrap*

Transmitter/receiver modules used in ESCON products rely on an optical wrap for testing and diagnostics. The optical wrap tool consists of a short length of fiber to optically couple light from the transmitter to the adjacent receiver. This tool is inserted into the TRS to allow optical wrap testing. Optical wrap testing is performed before product ship, during product installation, and while isolating field failures.

The optical wrap plug uses a short piece of optical fiber terminated with a plastic ferrule at each end. This subassembly is then contained within a plastic shell assembly functionally similar to the duplex connector. When inserted, plastic ferrules extend into the transmitter and receiver ports of the transceiver module to provide precise optical alignment of fiber to the active elements housed within module ports.

The wrap plug is designed to satisfy the optical characteristics of both transmitter and receiver by attenuating the transmitter power to ensure that the receiver input power is between the minimum receiver sensitivity and saturation. This is ensured for the total product distribution. This requirement is primarily achieved by using optical fiber with a 50-μm core diameter.

Table 4 Installation losses for a typical 3-km/62.5-μm link.

Parameter (length or number)	*Loss* (dB)	
	Mean (μ)	*Std. dev.* (σ)
Trunk fiber (2.76 km)	2.48	0.0
Jumper fiber (0.24 km)	0.42	0.0
62/62 connections (2)	1.4	0.28
62/50 connections (0)	—	—
50/62 connections (0)	—	—
Splices (4)	0.6	0.2
Higher-order-mode loss	1.0	0.0
Loss subtotal	5.9	0.34
Total loss ($\mu + 3\sigma$)	6.9	

The wrap plug is engineered for minimum cost. Tolerances on the ferrule diameters are relaxed, molded plastic is used as a ferrule material, and the ferrule polishing process is eliminated by employing a unique "cleave only" fiber preparation technique. The "cleave only" technique involves recessing the the fiber slightly below the ferrule end face. The ferrule end face is then terminated with a concave transparent epoxy meniscus, which protects the fiber and eliminates polishing.

- *Cable strain relief and management*

Fiber optic cables may operate in an environment where physical handling could result in axial loads of up to 60 pounds on the cable body. Loads of this magnitude transmitted via the cable body to the cable connector, or onto the mated transmitter/receiver assembly, will cause light loss or physical damage to these components. Protection in IBM products against damage of this nature (when necessary) is provided by strain-relief devices. These devices restrain the cable to prevent the transmission of undesired forces to the sensitive cable connectors and mating components.

Strain-relief devices used in IBM products include mandrel- and groove-type designs. In mandrel-type designs (usually used with small cable counts), strain relief is accomplished by wrapping the cable approximately two and one half times around the mandrel. When pulled, the cable tightens, thereby increasing the static friction between the mandrel and the cable jacket to provide the required holding action.

IBM products with many cables use a "serpentine groove"[2] strain-relief device. In one ESCON application

[2] U.S. Patent 5,115,260.

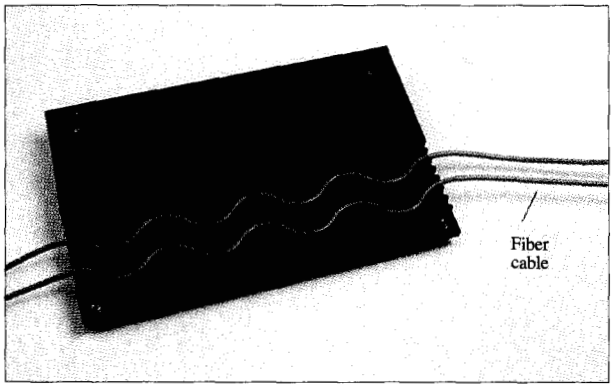

Figure 21

Cable strain relief.

more than 60 cables can be accommodated by a device of this type (**Figure 21**). Also, this device provides cable management and permits field serviceability. The cable is placed in a serpentine-shaped groove. Axial pull on the cable is translated into lateral force against the inner radii of the groove. This force increases the static friction between the cable body and the groove, thereby providing the required resistance to prevent cable movement at its termination point. These strain-relief devices control cable bend radii to satisfy the cable manufacturers' specifications. Large long-term bend radii are held to a minimum radius of 25 mm to minimize the effects of prolonged tensile stress on the glass fiber housed within the cable jacket. Short-term bend radii are held to 12 mm. The holding capability against axial pull is 267 N (60 lb) or better. These devices are also designed to prevent mechanical damage on the cable/fiber and increased optical attenuation.

Summary

The data signal connection to large processors has been primarily via copper media with optical fiber technology being used for specialized applications (e.g., distance extenders). This paper has described the first implementation of optical fiber technology for the I/O channel connections of the IBM System/390®. This link has a line rate of 200 Mb/s and is capable of transmission over 2- or 3-km distances using 50- or 62.5-μm multimode fiber, respectively. The data processing environment is significantly different from the telecommunications environment and fosters different trade-offs and design optimizations. The primary environmental requirements for data processing equipment that drive optical technology selection and link design are short distances, reliability (for both soft and hard failures), reconfigurability, ease of use, and physical environment. The above requirements, as well as improved data rate and distance as compared to copper channels, are the prime factors determining the technology and link design.

The optical link technology selected to achieve the environmental and data processing system requirements is InGaAsP/InP 1300-nm LED, InGaAs/InP PIN photodiode, and multimode optical fiber. The need for a rugged, low-profile connector offering positive engagement and excellent performance resulted in the development of a duplex optical fiber connector.

The use of fiber optics for data communication allows the data processing manager to explore new hardware configurations. The advantages of fiber optics suggest that it will become the pervasive interconnection technology between hardware units.

Acknowledgment

We wish to acknowledge the many individuals within IBM (too numerous to list here) who worked on the development of the fiber optic technology within the ESCON architecture. Their dedicated efforts made this project a success. We also wish to acknowledge the many workers within AT&T, Siemens, and Siecor corporations who contributed to the development of ESCON fiber optic technology.

References

1. *IBM System/360 and System/370 I/O Interface Channel to Control Unit Original Equipment Manufacturers' Information*, Order No. GA22-6974-5; available through IBM branch offices.
2. S. A. Calta, J. A. deVeer, E. Loizides, and R. N. Strangwayes, "Enterprise Systems Connection Architecture—System Overview," *IBM J. Res. Develop.* **36**, 535–551 (1992, this issue).
3. J. C. Elliott and M. W. Sachs, "The IBM Enterprise Systems Connection (ESCON) Architecture," *IBM J. Res. Develop.* **36**, 577–591 (1992, this issue).
4. C. J. Georgiou, T. A. Larsen, P. W. Oakhill, and B. Salimi, "The IBM Enterprise Systems Connection (ESCON) Director: A Dynamic Switch for 200Mb/s Fiber Optic Links," *IBM J. Res. Develop.* **36**, 593–616 (1992, this issue).
5. A. X. Widmer and P. A. Franaszek, "A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code," *IBM J. Res. Develop.* **27**, 440–451 (1983).
6. J. E. Morris and R. D. Brooks, "Long-Wavelength Compact Data Link Modules for 40–200 Mbit/sec Applications," *Proceedings of the Optical Fiber Communication Conference*, San Diego, February 11–13, 1985, pp. 38–39.
7. H. Kressel, *Semiconductor Devices for Optical Communications*, Springer-Verlag, New York, 1982, pp. 134–135.
8. Robert L. Rosenberg, David G. Ross, Patrick R. Trischitta, Daniel A. Fishman, and Carol B. Armitage,

"Optical Fiber Repeatered Transmission Systems Utilizing SAW Filters," *IEEE Trans. Sonics & Ultrasonics* **30**, 119–126 (1983).
9. Hoang Kim Do, "Model of a Fiber Optic Link in ASTAP," Master's Thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, May 1989.
10. David W. Boerstler and Terrence R. Scott, "Analyzing the PLL with ASTAP," *Technical Report TR-00.3333,* IBM Data Systems Division, Poughkeepsie, NY, February 7, 1985.
11. David W. Boerstler, "Dynamic Behavior of a Phase-Locked Loop Using a D-Type Phase Detector and Nonlinear Voltage-Controlled Oscillator," *Technical Report TR 21.1428*, IBM Data Systems Division, Kingston, NY, March 21, 1991.
12. W. T. Weeks, A. J. Jimenez, G. W. Mahoney, D. Mehta, H. Qassemzadeh, and T. R. Scott, "Algorithm for ASTAP—A Network Analysis Program," *IEEE Trans. Circuit Theory* **CT-20,** No. 6, 628–634 (November 1973).
13. "Attenuation by Substitution Measurement—for Short-Length Multimode Graded-Index and Single-Mode Optical Fiber Cable Assemblies," *American National Standard Document EIA-455-171*, Engineering Department, Electrical Industries Association, 2001 I St. NW, Washington, DC 20006.
14. J. J. Refi, "LED Bandwidth of Multimode Fibers as a Function of Laser Bandwidth and LED Spectral Characteristics," *IEEE J. Lightwave Technol.* **LT-4**, No. 3, 265–272 (March 1986).
15. "Light Launch Conditions for Long-Length Graded-Index Optical Fiber Spectral Attenuation Measurements," *American National Standard Document EIA-455-50A*, Engineering Department, Electrical Industries Association, 2001 I St. NW, Washington, DC 20006.

Received April 18, 1991; accepted for publication October 28, 1991

Nancy R. Aulet *IBM Enterprise Systems, P.O. Box 950, Poughkeepsie, New York 12602 (AULET at PK705VMA).* Ms. Aulet received her B.S. and M.Eng. in electrical engineering in 1983 and 1984, respectively, from Cornell University. She worked on LED-based fiber optic component development for data communications from 1984 to 1987. From 1987 to 1990 she worked on laser component development for optical storage, laser printer, and fiber optic communications applications. Subsequently she worked in a joint GTD/Research group on GaAs technology development. Ms. Aulet is currently manager of a technology development department for high-end systems.

David W. Boerstler *IBM Enterprise Systems, Neighborhood Road, P.O. Box 100, Kingston, New York 12401 (DAVEB at KGNVMA).* Mr. Boerstler received a B.S. in electrical engineering in 1978 from the University of Cincinnati, and an M.S. in electrical engineering in 1981 and an M.S. in computer engineering in 1985, both from Syracuse University. Since joining IBM in 1978 he has been involved in the design of analog and digital circuits, phase-locked loops, oscillators, and fiber optic systems. He is the recipient of an IBM Outstanding Technical Achievement Award and three IBM Invention Achievement Awards. Mr. Boerstler is an Advisory Engineer in the Advanced VLSI Technology Group.

George DeMario *IBM Enterprise Systems, Neighborhood Road, P.O. Box 100, Kingston, New York 12401 (DEMARIO at KGNVMF).* Mr. DeMario received his B.S. degree in electrical engineering from Polytechnic University, Brooklyn, New York. In 1967, he joined IBM at the Advanced Systems Development Division, Mohansic, New York, and later received his M.S. in electrical engineering from New York University. Mr. DeMario was responsible for technology applications for the IBM 8100 distributed processing systems. He later worked on the IBM PC and Workstation products and was part of the development team for the 3270/PC Engineering Workstation and the IBM Gas Panel Display. At present, Mr. DeMario is a Senior Engineer and manager of the engineering organization in IBM Enterprise Systems responsible for the development and applications of fiber optic technology in large systems. He manages the Advanced Products Technology Development Department, which was responsible for the ESCON link and the IBM duplex connector.

Frank D. Ferraiolo *IBM Enterprise Systems, Neighborhood Road, P.O. Box 100, Kingston, New York 12401 (FFRANK at KGNVMF).* Mr. Ferraiolo received the B.S. degree from Pennsylvania State University, joining IBM in 1982 at the IBM Development Laboratory in Poughkeepsie, New York, as a circuit design engineer. He is currently working as an Advisory Engineer at the IBM Kingston Development Laboratory in the Advanced Products Technology Development Department. Mr. Ferraiolo received an IBM Outstanding Technical Achievement Award in 1989 for his work on the SER/DES module for the ESCON fiber optic data link. He has received two IBM Invention Achievement Awards.

Curtis E. Hayward *IBM Enterprise Systems, Neighborhood Road, P.O. Box 100, Kingston, New York 12401.* Mr. Hayward has been involved in the development and qualification of various components in fiber optic technology over the last eight years. During that time, he has produced several technical publications and patents relating to

opto/mechanical devices. He is the recipient of an IBM Invention Achievement Award and several informal awards in recognition of his contributions to fiber optic technology. Mr. Hayward is an Advisory Engineer in the Advanced Products Technology Development Department, currently on leave of absence.

Charles D. Heath *IBM Enterprise Systems, Neighborhood Road, P.O. Box 100, Kingston, New York 12401 (retired).* Mr. Heath received his B.S. in electrical engineering in 1962 from the University of Massachusetts and his M.S. in industrial administration from Union College in 1977. He joined IBM in 1962 in Poughkeepsie. Subsequently, he was assigned to programs in the Components Division (GTD) and Data Systems Division.

Allen L. Huffman *IBM Enterprise Systems, Neighborhood Road, P.O. Box 100, Kingston, New York 12401.* Mr. Huffman received his B.S. and M.S. degrees in electrical engineering from Purdue University in 1986 and 1988, respectively. He joined IBM at the Data Systems Division in Kingston, New York, in 1988. As an engineer in the Advanced Products Technology Department, he has researched fiber modal effects, optical time domain reflectometry, phase-locked loops, and jitter in digital systems. Mr. Huffman is currently attending Yale University.

William R. Kelly *IBM Technology Products, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (WILLIAMK at POKVMCP2).* Mr. Kelly received a B.S. degree in electrical engineering from the Rochester Institute of Technology in 1981. In 1982, he joined the General Technology Division of IBM in Poughkeepsie, New York, as a Product Engineer in Corporate Component Procurement. He was responsible for electro-optical transducer design and, later, fiber optic transmitter and receiver devices. Mr. Kelly is now a Development Engineering Manager responsible for laser optoelectronic engineering and qualifications in Corporate Component Procurement.

Gerald W. Peterson *IBM Enterprise Systems, Neighborhood Road, P.O. Box 100, Kingston, New York 12401 (JERRYP at KGNVMF).* Mr. Peterson joined the IBM Components Division in Poughkeepsie, New York, where his initial technical assignment was in the Solid Logic Technology development organization. He received his B.A. in physics from Marist College in 1969, and has held engineering and management positions in the General Technology Division in East Fishkill and in the Data Systems Division in Poughkeepsie and Kingston. Mr. Peterson has been associated with the Advanced Products Technology Development Department since 1987; he is responsible for the development of the duplex connector/cable assembly for the IBM ESCON link.

Daniel J. Stigliani, Jr. *IBM Enterprise Systems, P.O. Box 950, Poughkeepsie, New York 12602 (DANST at PK705VMA).* Dr. Stigliani received the B.Engr. degree in general engineering from Stevens Institute of Technology and the M.S. and Ph.D. degrees in electrical engineering from the University of Illinois. In 1969, he joined the IBM Federal Systems Division in Owego, New York, working on a broad range of projects in optical signal processing and communications. In 1974, Dr. Stigliani transferred to the Data Systems Division, where he was responsible for the development of optical fiber communications for data processing applications. He is the recipient of a DSD Achievement Award and two IBM Invention Achievement Awards, and has published several papers in the field of optical communications. Dr. Stigliani is currently a Senior Technical Staff Member responsible for future high-speed optical communication development.