

Enterprise Systems Connection (ESCON) Architecture— System overview

by S. A. Calta
J. A. deVeer
E. Loizides
R. N. Strangways

This paper serves as an introduction to a wholly new IBM data processing interconnection system called Enterprise Systems Connection (ESCON™) Architecture™. Utilizing state-of-the-art fiber optic technology, the ESCON system introduces a unique concept to computer interconnection topology, the dynamic switched point-to-point connection. A comprehensive solution to the interconnection of data processing equipment and systems, the ESCON system offers superior connectivity, bandwidth, distance, and ease of installation. The ESCON architecture is directed toward the structuring of large dispersed multisystem data processing centers with campus distributed user communities, but it is equally well suited to the needs of small processing configurations. The paper first reviews the objectives of the ESCON development and

then gives a more detailed discussion of the system design alternatives and choices which were made. Topics discussed are the fiber optic technology, the interconnection topology, the ESCON architecture, and the design of the major system elements.

Introduction

Since its introduction in 1964, only incremental enhancements have been made to the IBM System/360™ I/O interface. Faster processors and tightly coupled multiprocessing eventually outstripped I/O capability. Performance projections indicated that the number of channels could exceed practical limits. Clearly, a new basis for I/O interconnection was needed. The Enterprise Systems Connection (ESCON™) Architecture™ is a comprehensive interconnection system, using fiber optic technology, that embodies a synergy of technology, topology, and architecture. This paper describes the

©Copyright 1992 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.




Figure 1

Elements of the ESCON system. Single lines represent optical fiber connections. T1 is a telephone line.

system design decisions that went into its development, how its elements evolved, and how these elements can be utilized to operate as a cohesive I/O interconnection network.

Elements of the ESCON system

The ESCON system makes possible the structuring of a high-speed backbone network in a data processing center, to which gateways to networks of lesser speed may be attached. The primary system elements from which an ESCON network can be configured are the fiber optic link, the ESCON channels, the ESCON Director™, and the ESCON control units.

Software functions supporting the operation of the ESCON system are

- The ESCON Manager™ program for ESCON Director configuration control.
- The ESCON Dynamic Reconfiguration Management.

Figure 1 illustrates some of the interconnection capability of the ESCON elements. See [1] for an overview of the features of the ESCON system. Discussions of

various ESCON elements can be found in the companion papers in this issue [2-6] and elsewhere [7].

Objectives of the ESCON system

The major customer concerns with respect to the interconnection of systems, control units, and channels are system disruption, complexity, cable bulk, cable distance, and the need for increased data rates.

This system was developed with the following objectives:

1. To permit operating systems and applications to run unchanged on computers installing the ESCON system.
2. To permit the insertion of additional control units and systems into running configurations without having to turn off the power ("hot pluggability"), thus avoiding scheduled and unscheduled outages for installation and maintenance.
3. To improve interconnection capability within data processing centers, for both increased intersystem connections and increased device sharing between systems.

4. To allow an increased number of devices to be accessible by channels.
5. To extend the distance for direct attachment of control units and direct system-to-system interconnection in the campus environment.
6. To provide significantly higher instantaneous data rates, and develop more efficient control and data transfer protocols; to provide an order of magnitude increase in I/O channel throughput.
7. To reduce significantly the bulk and number of cables required to interconnect the system elements of a data processing complex.
8. To develop availability approaches for ESCON-configured systems that will support continuous operation.
9. To provide a base for the total systems solution of interconnection, suitable for both current and future generation systems and I/O.
10. To support the orderly migration from today's systems to ESCON-configured systems.

The scope of the objectives, and of the ESCON system, was limited to the extended machine room and campus environments. With this system serving as the hub or backbone interconnection network within a data processing complex, remote users can access, by means of gateways (for example, via the IBM 3174 Establishment Controller), any of the data processing applications.

The machine room or extended machine room environment is defined as the area containing one or more systems of processors and channel-attached I/O, where I/O devices are possibly shared between systems. This environment, which may be spread out over different physical locations, excludes equipment attached via common carrier facilities. It is arbitrarily distinguished from the campus environment only by the assumption of distances under a kilometer (km) between the system (processors and directors) and control units. This environment includes data centers where equipment is installed on several levels of a high-rise building.

A campus is defined as an area, including multiple buildings, which is wholly under the control of an enterprise. Interconnections may extend 2-3 km or more. It is an area where dedicated cables can be installed without interfering with a "right of way." The ESCON objectives included the capability to interconnect systems in data centers in widely dispersed buildings, to support database backup in a remote secure building, and to support database sharing, if desired, between systems in separate buildings. The objectives also included the support of local area network (LAN) connections (both high- and low-speed) between and within campus buildings [6].

ESCON system design considerations

To meet the design objectives discussed above, a new technology, a new architecture, and a new interconnection topology were required. The first decision made, after studying the capabilities and potential of various transmission media, was to use serial fiber optic links. This was followed by the selection of the technology and the development of appropriate transmission techniques. A major effort was spent on the analysis and selection of an interconnection topology that would provide the most suitable characteristics in the data processing environment. Finally, the architecture for the control of what was envisioned as a campus backbone network had to be defined in painstaking detail.

● *Selection of link technology*

Computer complexes are configured with the ESCON system by connecting the system elements with serial point-to-point fiber optic links. At the outset, it was decided that serial transmission would be used for our new interconnection links, because, at distances greater than 400 feet, the signal skew of parallel cables would become intolerable. This meant, however, that a much higher transmission rate would be required to achieve the desired performance. Serial transmission cables would, however, result in a significant saving in cable bulk.

After reviewing laboratory experiments with copper coaxial lines and fiber optic cables, it was concluded that, without resorting to exotic and bulky cables, copper coaxial lines were probably too limited in bandwidth and distance. Fiber optic cables, on the other hand, showed promise of higher bandwidth at much greater distances. At the time, fiber optics suitable for our application was still highly experimental, requiring the development of reliable devices and connectors as well as suitable fiber and jacket material. Fiber optics also did not appear to be capable of supporting a multidrop technology, in which several units share the same interface. This implied that our parallel OEMI (original equipment manufacturers' information [8]) topology would be unsuitable.

As the technology study commenced, short-wave light-emitting diodes (SW LEDs) were becoming quite stable as devices, but at limited output power which restricted transmission distances. Experimental long-wave (LW) laser and long-wave LED devices were also being developed that could potentially answer the distance problem. LW LEDs were selected for their reliability and lower cost. A target transmission rate of 200 Mb/s was established, and from the fiber loss information available, it was concluded that transmission distances of 2 km could be achieved. This has since been extended to 3 km, including allowance for a reasonable number of connectors and splices. The LW LED also had the benefit of not producing dangerous levels of light intensity, as did lasers. On the negative side,

the detection mechanism for light signals is a low-level analog amplifier, which could be subject to intermittent errors. The link architecture had to include a retry mechanism to account for this possibility.

The fiber optic link was defined to consist of a pair of single fibers, capable of supporting the full-duplex transmission of data. The two fibers comprising a link are packaged in a single multilayer cable jacket with a diameter of 4.8 mm. This functionally replaces two parallel I/O interface cables, each about one inch in diameter, resulting in a significant saving of both weight and physical bulk. The ends of a cable are terminated with IBM duplex connectors, which are specially designed to prevent misinsertion and to ensure proper fiber alignment. All of the ESCON system elements are provided with matching connectors for link cable attachment. The link cables can be connected and disconnected while the elements remain in an operational state. Other advantages afforded by fiber optic technology include improved security, because the cable does not radiate signals, and an insensitivity to external electrical noise. See [2] for a detailed discussion of the ESCON fiber optic technology.

• Selection of link architecture

Serial transmission has long been used in communications media. However, the higher speeds contemplated, and the system application, required new data transmission techniques. A character transmission code was needed that would assist the detection and clocking process, but not be an overhead burden. The "bit-stuffing" of Synchronous Data Link Control (SDLC) [9] was felt to be impractical at high speeds, and the Manchester code at 18 bits per byte was considered to be too much overhead. Also, some additional code points (one element of a code set) for character and frame demarcation were required. The "8 out of 10" (8/10) code [10] was selected for this reason. Data are transmitted in the form of packets of characters called frames. Because clock recovery was felt to take too long, "idle" characters are transmitted continuously between frames so that the photodiode receivers can maintain synchronization with the transmitter. Each character contains 10 bits when use is made of the 8/10 code, which is a dc balanced code with good code point separation characteristics and space provided for additional control codes. The control codes are unique 10-bit characters. For transmitted data characters, the 10-bit code provides a more robust checking than does the conventional 8-bit plus parity code.

It was believed that fiber optics might be subject to intermittent bit errors, not from external electrical interference as is copper cable, but from the low signal levels which could be expected at the detector. Therefore, in addition to the 8/10 character checking, a cyclic redundancy check (CRC) field was specified to be included

at the end of each transmission frame, to decrease the probability of undetected bit errors. The bit-error rates projected for fiber optics were of the order of 10^{-12} to 10^{-15} . It was felt that this was low enough to obviate the need for frame-level recovery, and that a hardware-initiated command retry could be utilized for recovery.

It was next decided that the link architecture should be defined with two distinct layers to separate transmission functions from the I/O operations which the link was to convey. A symmetric "link" layer was defined, containing routing and transfer control information, and a "device" layer was defined to contain the System/370™ Extended Architecture (XA) I/O information. The Enterprise Systems Architecture/390™ (ESA/390™) device layer obviously could not be symmetric because of the master/slave nature of the System/370-XA [11] architecture. Because they were somewhat dependent on the topology, the final link architecture definitions had to wait until a topology was selected.

Maximizing cable distance formed the basic rationale of the architecture. This required minimization of the number of "handshakes" (message exchanges), streamlining the control protocol, and designing data transfer pacing to allow full data rate transmission. To address the control exchange, control frames were designed to contain an entire channel command word (CCW). Data transfer control was tied to the availability of buffers at the data receiver. The receiver sends to the sender "data request" frames which indicate the number of bytes to be sent. By implementing sufficient buffers for the round-trip distance, a receiver can keep the sender supplied with enough data requests to support continuous data transfer. The data transfer protocol could be characterized as a "go ahead n (bytes)" type of protocol.

• Selection of topology

The term "topology" is used to refer to the physical or natural features of objects and their structural relationships. In the ESCON case, the term is used to refer to the network structure which is used to interconnect the elements of a data processing complex. It utilizes a topology which is called "switched point-to-point"; that is, all elements are connected by point-to-point links, in a star fashion, to a central switching element. By way of contrast, the IBM parallel OEMI system [8] uses a multidrop topology of "daisy-chained" cables.

Having decided that the new interconnection system was to be software-compatible, one could have chosen to implement a serial channel extender. The IBM 3044 Fiber Optic Channel Extender Link was just such a product; it not only offered distance extension, but was the initial IBM fiber optic product. It was soon decided, however, that this approach would be too limiting for a new

interconnection system. It offered no additional connectivity and did not reduce control unit adapters. A new topology approach was clearly required.

For many reasons, selection of the topology was the most difficult task, next to the architecture itself, of the system design. Many topology alternatives had to be defined, analyzed, and modeled with no preconceptions before the final approach was selected. Simulation models were required to determine the effective performance of the collision bus and various ring topologies. An atmosphere of competition prevailed during this period, with proposals being received from all quarters. It was known that something similar to an interconnection network would be desirable, since with the parallel OEMI multidrop system too many control unit adapters were required as the number of systems increased. Some of the topologies studied are described in the following sections.

Collision bus

The collision bus [12], basically a large OR circuit, derives from the Aloha radio network protocol using a common frequency. All users have equal access to the medium and compete for its use. A user attempting to send something to another user must determine that the medium is "quiet" before starting to send. Even then, a collision with another user is possible, and a retry algorithm must be defined so that the second try is treated as a reservation. Usually priorities are assigned among users to resolve conflicts.

The bandwidth of such a shared facility was determined to be unsatisfactory, and the fiber optic technology does not easily lend itself to the implementation of a collision bus. Another undesirable characteristic of the collision bus was the rapid degradation of response time as utilization of the facility increased.

Token ring

A ring structure, in our opinion, would be one of the easiest to cable. As with the multidrop, one would simply daisy-chain devices and channels using point-to-point links. To add a device, one could open the ring at the nearest point and plug in the new device with one additional link. Somewhat like the collision bus, the token ring is a contention system which is resolved by allowing only the possessor of the token to use the facility.

Like the collision bus, the token ring provided too limited a bandwidth for the System/390™ I/O application, and exhibited poor response time characteristics as well. Furthermore, unless a dual path was provided, the reliability was less than desired.

Host star

A topology consisting entirely of point-to-point connections, the host star was perhaps the easiest to understand. Both channels and control units required

multiple tails, and links were simply connected to any units with which a channel was to communicate. There was sufficient bandwidth for all units to communicate simultaneously, and thus response time was optimal. The host star, however, was by far the most expensive topology studied. It also was the most difficult to install, and resulted in a large number of cables in a large establishment.

Dual insertion ring

Unlike the token ring, where one transmission occupies the ring until the possessor of the token releases it, the insertion ring operates on a frame multiplex basis. Any station, seeing a gap after a frame, can insert a frame of its own. It must, of course, buffer any frame which arrives during its insertion process, and insert the buffered frame immediately following its own frame. As each frame reaches its target station, it is removed, creating a gap which can be used for a new frame. The dual form utilized two fibers per link to form counter-rotating rings. Upon detecting an open link, a station was defined to wrap the remaining link onto itself to form a single ring. It could thus survive a failed link or station, and new stations could be dynamically inserted.

The dual insertion ring exhibited much better bandwidth and response time than even a dual token ring. With uniform station activity, the resulting bandwidth was four times that of a single fiber, because each frame would travel an average of halfway around the ring. The elastic nature of the insertion ring gave it very good response-time characteristics. Logically a distributed "any-to-any" switch, the dual insertion ring was the most cost-effective topology studied. It was decided, however, that at four times the link rate, the bandwidth available was insufficient for the application.

Switched point-to-point topology

The switched point-to-point topology utilizes a central dynamic crossbar switching element, to which the elements to be interconnected attach by means of point-to-point links. For example, an I/O channel so connected can be connected to any other element also connected to the switch.

The topology eventually selected for the ESCON system, switched point-to-point, except for the time to complete switch connections, provides the bandwidth of the host star at a fraction of the cost. Switched point-to-point is logically equivalent to a ring, where a ring represents a distributed switch. Using the same number of cables as the ring, switched point-to-point substitutes the switching element for the more complex adapters required for a ring. These switching elements provide more connectivity paths, yielding much higher bandwidth. Switched point-to-point is also more convenient for adding

and reconfiguring equipment than the ring. Also, switched point-to-point adds the ability to provide isolation and partitioning to the network.

The switched point-to-point topology met all criteria and provided some additional useful features as well. The next problem was to determine whether a suitable switch could be economically designed and built, and to develop a definitive supporting architecture.

With a properly designed switching element, the switched point-to-point topology forms a peer network; that is, the basic symmetry allows any attached element to connect to any other attached element. With a fully dynamic switch of N ports, $N/2$ concurrent connections can be made. The term "dynamic" in reference to the switch means that its connections are automatically established and broken for each transmission sequence. The switched point-to-point topology is also an economical network, since only a single cable is required to join any element in the network. Thus, the ESCON Director switching element, while providing an extremely high connection bandwidth, at the same time supports an economical form of interconnection network.

While data are transferred in the form of packets called frames, the ESCON switched point-to-point network is not a packet-switching network. The lead frame of a transmission includes the request for a connection, and then the ESCON Director establishes and maintains that connection, in the manner of a "circuit switch," for the duration of the conversation between the two connected elements. Once a connection has been established, data do not experience a "store and forward" delay while passing through the switch path. The connected path is "full duplex"; that is, a two-way conversation is established.

• *Functional attributes of the switched point-to-point topology*

The switched point-to-point topology satisfies the following functional requirements:

1. *"Multidrop" function* The switched point-to-point topology gives system I/O channels the ability to connect to many control units, thus effectively supplying the function provided by multidrop of the System/370 parallel OEMI channels. With the 60-port ESCON Director design, many more than the current limit of eight control units can be accessed by one channel.
2. *Multitail or shared control unit function* Conversely, once a control unit has been attached to an ESCON Director, its devices can be shared by any system's ESCON I/O channels that are also attached to that director. For example, with System/370-XA, a direct access storage device (DASD) controller, which is to be shared by four systems and requires attachment to

two channels per system (for performance reasons), requires eight parallel OEMI adapters. With the ESCON switched point-to-point topology, two ESCON link adapters, the interface to the ESCON links, preferably connected to two different directors for availability, provide the equivalent connectivity. In addition, the ESCON I/O channels of more than four systems can be attached to those same directors to achieve a much higher degree of sharability at no additional cost to the control unit.

3. *Alternate pathing function* With the switched point-to-point topology, it is recommended that control units be connected to two ESCON Directors, each of which is attached to an ESCON I/O channel of the using system. This approach provides two available paths to the control unit devices, which can be used interchangeably as alternates under busy or failure conditions. While it is possible for a system to use either of two ESCON I/O channels attached to the same director as alternates, it would not afford the degree of availability of the two-director approach.
4. *Channel-to-channel (CTC) function* With the switched point-to-point topology, any ESCON I/O channel of one system can connect to an ESCON I/O channel of any other system which is attached to the same director. For CTC operation, however, the ESCON I/O channel of one of the communicating systems must have been initialized with the ESCON CTC microcode. This restriction has been applied in order to support current System/370-XA "channel-to-channel" protocols. However, the ESCON switched point-to-point topology effectively replaces the CTC interconnection function provided by the IBM 3088 Multisystem Channel Communication Unit. With the announced ESCON I/O channels, an ESCON CTC transfer operates at up to 17 MB/s, and significantly more than eight systems can be so connected.
5. *Install function* The switched point-to-point topology simplifies the addition of equipment to (or its deletion from) a given configuration. With one jumper cable, an ESCON I/O channel or control unit can be connected to a director in an existing network. Furthermore, because of the use of fiber optics, connections can be made while the installed equipment continues to be operational.
6. *Fencing function* The term "fencing" is used to describe the isolation or disabling of a director port. When a port is fenced off, all connection requests to that port are rejected, and the attached element cannot contact the director as well. Fencing is used primarily to isolate an element during maintenance. The ability to fence off the director port to a failing element allows the remainder of the complex to remain operational while repairs are made to that element. Fencing also

permits new equipment to be added and tested without affecting the operational elements of a complex. Any element can be added to an existing fiber optic network without disruption.

7. *Dedicated path function* Dedicated connections are established by the controller and are insensitive to controls contained in frame link headers. Such connections can be used for configuring the serial side of the interface converters (both those that allow parallel control units to communicate with a serial channel and those that allow parallel channels to communicate with a serial control unit).
8. *Cascaded Director operation* The design of the ESCON Director supports only one level of dynamic director path selection. However, two directors can be installed in a tandem (series) configuration if the path through one of the directors is a dedicated connection. This "cascade" may be useful for certain configurations. Since the director ports provide a link repeater or redrive function, greater distances (up to 9 km with three 3-km cables) can be covered by a cascade configuration.
9. *Partitioning function* The ability to partition director ports, that is, to restrict the connectivity of a given port to a subset of all other ports, permits portions of the system to be isolated as required for test configurations.
10. *Point-to-point operation* The ESCON architecture does support direct point-to-point connections. For very small system configurations, each ESCON I/O channel of the system can be cabled directly to one ESCON control unit without the use of an ESCON Director, using only point-to-point connections.

The majority of the ESCON objectives were directed toward improving system and subsystem interconnectivity in a data center. The selection of a switched point-to-point ESCON topology offers the potential to reduce the number of cables required. At the same time, it supports better connectivity, and allows a greater number of devices per channel. Switched point-to-point also simplifies installation and isolation. By incorporating the CTC function into the ESCON channels, and flexible switching controls into the ESCON Director, the number of interconnection products required is reduced.

Figures 2 and 3 illustrate the difference in the basic cabling concepts between the parallel OEMI multidrop and the ESCON switched point-to-point topologies. In comparable four-system configurations, the difference in cables and control unit adapters can readily be seen. The ESCON configuration illustrates the recommended availability structure, which includes two paths, through different directors, from each system to any control unit.

Figure 4 illustrates the striking advantage of the ESCON system over the parallel OEMI system in the number of cables required to effect an interconnection network. Where the number of cables with the multidrop parallel OEMI system increases with the square of the number of systems, with the ESCON system the number of cables increases linearly.

• *ESCON architecture*

The ESCON architecture can be divided into two basic parts:

1. That which defines how the network transmission functions operate over a link.
2. That which defines how the ESCON elements (control units, channels, etc.) operate.

The first of these is referred to as the link function or transport level, and the second is referred to as the device function or level. It was decided at the outset that a strict distinction would be maintained between the levels, and the formats of the transmission frames are defined accordingly. Each frame contains a link-level portion, and some frame types also include a device-level portion as well. The link level consists of the link header and the link trailer, and is concerned with the proper delivery of the frame to its intended destination. The device level is concerned with the control information which is exchanged by channels and control units, and the data to be transferred between them.

The architecture defines the exact frame formats to be used, and how each field of the frames is to be interpreted. It also defines the control sequences used to operate I/O devices. The basic frame types employed in the ESCON system are link control, command, status, and device frames. Device frames include both frames for devices and frames for data transfer.

Link operations

All information is transported on the ESCON fiber optic network in the form of frames or packets of data. Frames can be of various lengths from 7 bytes to 1035 bytes. Supporting the formation, clocking, and receipt of frames are the 10-bit mode (8/10 code) idle and frame start characters. The remaining data within frames are translated from the 8-bit mode (8-bit plus parity) to the 10-bit mode for the actual transfer. The frames are then translated back to the 8-bit mode before being interpreted by the receiver of the frame.

The link header contains a network destination address in the form of an ESCON Director target link address, plus a link control field, and the source link address of the sender. When a frame is received at a director input port, its header is examined by the switch matrix controller.




Figure 2

Parallel OEMI system configuration example.

If a connection is requested, the controller checks its connection tables for the status of the requested target link address. If the corresponding target port is available, the controller sets the switch connection matrix path and passes the frame through. If the target port is already connected to a different port, the controller sends a port busy frame to the sender on its inbound fiber. Upon receiving a busy signal, the sender must restart the transmission sequence.

The link trailer of each frame contains a 16-bit cyclic redundancy check field (CRC) which is computed by the sender on the remainder of the frame contents. All adapters, except the director ports, are responsible for checking the CRC of received frames, as well as for checking the validity of each of the 10-bit coded characters in that frame. An adapter discards any bad frames it receives. If the first frame an adapter receives is bad, it returns an error message frame. If a transfer is in progress when a bad frame is received, an I/O command retry may

be initiated. The ESCON Director does not establish a connection when it detects a bad frame, but returns an error message frame to the sender.

The link level of the ESCON architecture is basically symmetric in nature. The new function of path selection through the ESCON Director has been included in the link level to support the switched point-to-point topology. Potentially, any attached element can request a path connection, and once a connection has been established, the director has no further function except handling the clock jitter between those elements and waiting for a disconnect request.

Device operation

At the outset, it was established that the device level of the ESCON architecture had to support device operations that were compatible with System/370-XA I/O architecture. That is, existing applications software, operating systems, and channel programs had to work with minimal change




Figure 3

ESA/390 ESCON system configuration example.

when running on systems equipped with ESCON channels, Directors, and control units.

The ESCON device layer, consequently, is used to transfer System/370-XA operational I/O sequences between a channel and control unit, and to establish any included data transfer. The architecture defines the exact protocol to be used under all conditions, and also defines the protocol and pacing of data transfer. The protocol has been implemented in the ESCON channel and in the ESCON control units, in addition to the normal functions that these elements perform to operate I/O devices. The device protocol basically duplicates the sequences implemented by parallel channels and described in [8], but is substantially modified to suit the nature of the serial environment. A more efficient protocol has resulted from the adoption of the serial media, particularly in the data transfer protocol.

Data transfer is accomplished without "handshakes." The receiving element paces the transfer by sending "data requests" to the sender, specifying an amount of data the

receiver is willing to accept. The receiver is permitted to send a subsequent request before the sender has satisfied the previous request, and typically requests are sent early enough to avoid any suspension of data transfer.

The ESCON device-level architecture faithfully reproduces all responses expected by System/370-XA software. Additional function and enhancements have been included, however, to support the new dynamic capabilities of the ESCON system. A detailed discussion of the ESCON architecture is given in [3], and the interface architecture is described in [13].

● *ESCON software support*

While the ESCON architecture gave one the ability to operate without changes to channel-program-generating software, it was decided that new software was needed to complement and enhance the ESCON systems aspect. The key software support for the ESCON system addressed the areas of configuration management, dynamic configuration support, and MVS ESCON-related enhancements.




Figure 4

Number of interconnection cables vs. number of systems for the parallel OEMI and ESCON systems.

Configuration management

Given that the switched point-to-point topology had been selected, while the switches could be transparent in normal operations, they were not transparent in error reporting and switch management. A local director console was included to support configuration control and maintenance of the ESCON Director when used with current software. It was planned to also provide the means to support centralized control through new software. The switch controller of the ESCON Director therefore included the ability to receive switch configuration commands from system channels. The software which utilizes this ability is the ESCON Manager.

The ESCON Manager is a host software product which automates many director configuration management functions in a system. It assists the system operator to manage the migration of equipment to the ESCON system, to establish the configuration of equipment for maintenance, to reconfigure for failed paths, and to direct any connectivity changes of director ports.

The ESCON Manager provides host control of the ESCON Directors in an ESCON I/O configuration. Some of the functions of the ESCON Manager are 1) to support

the migration from parallel OEMI to the ESCON system by establishing dedicated switching paths from ESCON channels for ESCON converters; 2) to manage configuring of devices for scheduled maintenance through port isolation; and 3) to maintain consistency of logical operating system pathing with physical pathing in the director.

Dynamic configuration support

Fiber optic technology made practical "hot pluggability," the possibility of dynamically adding I/O equipment to and removing it from data processing configurations without the need to power down all of the systems. To make this a reality, however, the means for software and the channel subsystem to dynamically add or remove the control blocks needed for device communications had to be developed.

This resulted in the development of Dynamic Reconfiguration Management, which allows an installation to add, delete, or modify the definitions of channel paths, control units, and I/O devices in both the software and hardware I/O configuration representations. Changes are performed without requiring either a hardware or software outage (i.e., power-on-reset of the hardware or initial program load of the software).

This function is complementary to the nondisruptive installation capability of ESCON control units and devices that fiber optic technology makes possible. It increases system availability by allowing changes to the I/O configuration while systems are in production. It eliminates the disruption of power-on-reset and initial program load caused by I/O configuration changes, as well as the subsequent restart of subsystems and networks. Also, it allows installations to make changes to the I/O configuration as required rather than scheduling around a planned outage. With planned outages becoming less frequent, with an implicit objective of minimizing duration in order to accommodate international operations, this has a secondary benefit in that time previously allocated for I/O configuration changes can now be used for other activities.

Dynamic Reconfiguration Management builds on the existing Multiple Virtual Storage/Enterprise Systems Architecture (MVS/ESA™) Hardware Configuration Definition (HCD) function. HCD is an interactive process that allows the definition of an I/O configuration to both the hardware and the control program software. Specified input is verified, at entry, for syntactic and semantic consistency with hardware and control program software definitional requirements.

At control program initialization, an instruction in the interface program allows the control program to determine whether its representation of the I/O configuration is consistent with that understood by the hardware. Once consistency has been verified, the installation can replace

the current I/O configuration definition with a new definition by using an HCD interactive panel or an MVS operator command to invoke the Dynamic Reconfiguration Management function. The control program determines the changes that must be made to the existing definition (additions, deletions, and modifications) and performs the necessary changes to the software and, through the interface instructions, to the hardware. Changes are synchronized with existing I/O activity to minimize or eliminate disruption. The resultant hardware definition can be maintained for use during subsequent initialization so as not to lose the changes that were made dynamically.

The control program provides services that allow installation applications and program and vendor products to be notified of a planned or completed configuration change. Applications which are sensitive to the I/O configuration (e.g., ESCON Manager and RMF, the Resource Management Facility of MVS) can upon notification make appropriate changes to either use the added resources or adjust their understanding of the configuration.

Dynamic Reconfiguration Management is supported both in basic mode and logically partitioned mode. This function, invoked in a logical partition, changes the hardware definition of I/O resources across all affected partitions. It provides a single point of control for Dynamic Reconfiguration Management related to hardware definitions. Installations must coordinate the planned I/O configuration change, across all affected partitions, before changing the hardware definition.

MVS ESCON-related enhancements

A number of system enhancements were included in the ESCON architecture to improve the reliability, availability, and serviceability of these products.

Configuration validation It was recognized that the optical fiber technology allowed for ease of recabling. MVS exploits features of the ESCON system to provide protection against accidental miscabling and switching actions. During system initialization, MVS constructs a table that contains all of the self-description data from all of the devices that support this feature (all ESCON-attached devices are required to support self-description). While constructing this table, MVS validates that the physical configuration is consistent. This means that all the device paths that are defined to go to the same physical device actually go to this device. Independent means are provided for notifying the system whenever a resetting event condition occurs. Upon notification of a resetting event condition through the System/390 architecture, MVS initiates its device-independent configuration validation function to ensure that the device that was in use is still the same device, before any application I/O is permitted.

Automatic recovery from I/O failures With the dynamic pathing feature of I/O devices came the requirement that MVS synchronize the actual physical state of the device with the logical state maintained by MVS. Failure to do this leads to integrity exposures during I/O recovery scenarios and system degradation in terms of missing interrupts and lower device utilization. Whenever MVS detects a symptom that the physical state of the device may have deviated from the MVS logical state, MVS synchronizes the two states. Some of the symptoms include missing interrupt conditions, resetting events, and "not operational" conditions. If MVS cannot determine the state of a device path or I/O device, it takes the resource off-line. If the errors are transient or if the failing I/O resource was repaired, manual intervention is required to bring the resource on-line in the parallel OEMI environment.

With the ESCON system, MVS is provided with notification of previously available resources again becoming available. Such an indication was not available in the parallel OEMI environment. Upon receiving the notification, MVS automatically reconfigures to take advantage of the available resource and revalidate the I/O configuration for consistency.

Exploitation of ESCON topology In the parallel OEMI environment, certain device errors require that the operating system perform I/O system resets to recover the device. For example, if a hot-I/O condition is detected (that is, if a detected error condition appears to be permanent rather than intermittent), MVS recovery may issue a "Reset Channel Path" instruction trying to clear the condition. With the "daisy chaining" of control units in the parallel OEMI environment, all of the devices on the channel path are affected by the I/O system reset, even though they were not experiencing a problem.

With the ESCON switched point-to-point topology, the same recovery action by MVS would needlessly affect devices not experiencing a failure. A new function was defined that provided MVS with the ability to reset only those devices that are on the control unit causing the hot-I/O condition. MVS maintains tables that describe the topology of the ESCON devices and uses the new capability to make its hot-I/O recovery more granular, allowing reset of selected devices rather than resetting all devices attached to a channel.

Additionally, with parallel OEMI devices, resetting events that are detected by MVS cause every device on the channel path to be recovered, since there is no way for MVS to determine the actual scope of the reset on a parallel channel path. With ESCON attached control units, the reset event recovery actions are linked to those devices on the same control unit as the device presenting the reset event.

Supplemental status In System/390, various device and channel program errors are reported to the system with an I/O interrupt that includes unit check status. This status indicates that the program must retrieve additional data from the device in order to determine the exact cause of the failure. The additional sense data, retrieved via the Sense command, are device-state-dependent, and require that no other activity occur at the device from any other interface before retrieval of the sense data, which would invalidate the data. In the parallel OEMI environment, the devices would establish a contingent allegiance and reject any I/O attempts to the device until the sense data were retrieved from the interface that reported the unit check. During the period of time between the host's acceptance of the unit check and retrieval of the sense data (or clearing the contingent allegiance in other ways), the device is inaccessible to other hosts. If the system that received the unit check fails and does not retrieve the sense data or break the contingent connection, the device is lost until some manual action is performed to reset the interface.

The ESCON and ESA architectures allow MVS to receive the device sense data along with the unit check status. The sense data arrive at the host in the form of supplemental status. The ESA/390 architecture then allows the data to be presented to the software with the I/O interrupt. This feature minimizes the length of time that devices remain unavailable in a multisystem environment because of device errors. It also eliminates the exposure of devices being left unavailable because a system failed to break a contingent allegiance.

◆ *ESCON reliability, availability, and serviceability*

The ESCON design improves overall system reliability simply by requiring fewer components to make the I/O and system interconnection. Just by the reduction in the cables required, and by the use of serial cables instead of parallel, the number of potential points of failure of connector contacts is reduced by orders of magnitude. In addition, the ESCON system displaces the IBM 3088 Multisystem Channel Communications Units, the IBM 3814 Switch Management System, and the IBM 3044 Fiber Optic Channel Extender Link of the System/370.

With the ESCON system, availability is increased in several ways:

- ◆ "Hot pluggability" of equipment returns previous scheduled and unscheduled maintenance and installation periods to productive operating time.
- ◆ Most intermittent bit errors on the fiber optic link are corrected transparently to the user and the application.
- ◆ The ESCON Director can continue to operate with several failed elements because of the redundancy incorporated in its design.

- ◆ With the recommended dual-path configuration strategy, systems can continue to operate while sustaining failures of channels, directors, and control units.

The ESCON facilities also are planned to improve serviceability; they have the ability to "fence off," or isolate, director ports, permitting maintenance service to be performed on equipment without affecting the remainder of the system. They also have the ability to utilize logical partitioning of systems, and the ability to partition the director, which allows repaired equipment to be thoroughly tested without affecting production before it is returned to service. The ESCON Monitor [1] supports more effective service through its ability to detect and analyze exceptions in the installed equipment.

ESCON element design considerations

◆ *ESCON channel*

The ESCON fiber network is accessible by the ESA/390 systems through the ESCON channels. Each ESCON channel can operate any of the multiple I/O devices which it can access through the fiber network. In addition, ESCON channels can operate with any channel-to-channel-configured ESCON channel to form an extensive system-to-system fiber network. The ESCON channels operate asynchronously to the system processors. In response to I/O instructions executed by the central processor, the channels are initiated by the system I/O processor (IOP). The channels then execute I/O programs to operate various I/O devices. Each channel has independent access to the main memory of the system to fetch its channel program, and to fetch and store data for I/O operations. The ESCON channel is designed around a high-speed processor, which executes programs from the writable control store (WCS). The channels include a considerable number of working data buffers implemented both with hardware registers and RAM, used to prepare and receive I/O data. Each channel also contains the logic which communicates with the ESCON fiber optic link.

With the switched point-to-point topology, each ESCON channel drives only a single ESCON link, which can be attached to an ESCON Director port or to an ESCON control unit adapter. The link interface of the ESCON channel consists of independent transmit and receive sections. These include the hardware logic to serialize and deserialize the optical signals of the link [2]. The channel acquires and maintains character synchronization on the 10-bit characters received, and translates these characters to the 8-bit plus parity form used internally.

◆ *ESCON channel-to-channel (CTC) adapter and Converter Model 1 implementation*

By personalizing the ESCON channel with microcode, an ESCON channel can be converted to an inboard CTC

control unit, including the local channel functions. The ESCON inboard CTC can be used to communicate with any ESCON channel on its fiber network. This facility provides a high-bandwidth intersystem backbone network, suitable for the rapid transfer of either control information or large data blocks between cooperating ESA/390 systems. Additional microcode is provided to permit the channel platform to operate with the ESCON Converter Model 1. See [4] for a detailed discussion of the ESCON channel design.

● *Design of the ESCON Director*

The heart of the ESCON fiber optic network is the ESCON Director. The ESCON Director is a dynamic switching unit which can potentially establish connections between any of its attached ESCON elements. The 60-port IBM 9032 Director supports up to 30 concurrent full-duplex 200Mb/s connection paths. Switched in microseconds, the ESCON Director reacts to control signals included in the headers of arriving frames in order to establish a requested connection. Connections are full-duplex, and are maintained until one of the connected elements signals a disconnect. Once connected, the connection is transparent to the elements communicating on that path. Multiple paths operating through the ESCON Director cannot interfere with one another, since a port cannot be connected simultaneously to multiple ports, or to any port already connected.

There were many questions to answer, and many decisions to make in arriving at a design approach for the ESCON Director. Should it be a circuit switch (crossbar) or packet switch (store-and-forward)? Transparent or not? Any-to-any? How should busy and error conditions be handled?

A crossbar design was selected because it would support more bandwidth than would a store-and-forward design. It was then decided that the switch should be handled entirely by the link or transmission layer, and thus be transparent to the device layer and software. The link layer was planned to be a completely symmetric, any-to-any design, but with the ability to partition and fence off ports. Logically, the switch had to be managed by a single high-speed controller.

The first frame of a sequence was planned to carry the target link address of the required connection. It was specified that each port would contain a receive buffer in which to accept that first frame, and where its target address could be examined by the switch controller. The switch controller was to maintain a table of active connections against which each new request could be tested to determine the availability of its target. Upon encountering a busy target, the switch controller would send a port busy frame to the originator so that the sequence could be restarted.

A basic function performed by the ESCON Director is the rationalization of the individual link clocks of the ESCON elements in its network. The clocking for the ESCON Director link ports is derived individually from the attached channels or control units. The receive side of each link port acquires character synchronization individually with the bit stream being received. The transmit side bit stream of all ports is driven from the director's 20MB/s system clock. Each individual port therefore operates with a slightly different receive clock. This results in a slight difference in the clocks used when data are transferred through a switch connection. The data are received at one link port as clocked by the sender, and then clocked out of the target link port at a bit rate based on the director's system clock. To accommodate the clock difference, two additional idle characters are inserted between frames by the sender. The director then has the option of deleting or adding an idle character, between frames, as the frames pass through. Whenever a character's difference between the clocks builds up, the ESCON Director makes a dynamic adjustment by adding or deleting an idle character.

A System/370 I/O operation is usually originated by a channel, which attempts to send a command, derived from a CCW, to a control unit. The channel then expects a response from the control unit indicating the action it is taking. Depending on the CCW program, several more control exchanges may take place. It was decided that it would be most efficient to maintain a switch path connection, in both directions, for such exchanges. If a control unit were to perform a long operation, such as a DASD seek, the control unit could terminate the connection upon the completion of a command sequence, in order to free the channel for another operation. The control unit then would become responsible for initiating the reconnection when it was ready to resume the operation. To implement these protocol exchanges in the switch environment, both channels and control units must not only be aware of the switch, but must manage their connections and task switching to suit the switch path availability. For example, if a control unit receives a "port busy" signal in attempting to reconnect an operation, the control unit has the option of trying an alternate path to a system and completing the operation with a different channel. Thus, the application view of I/O is oblivious of the switch network, but the I/O subsystem is intimately involved in its operations.

The director was planned to support I/O configuration management functions. The switch controller was defined to be addressable as a control unit via a unique target port address, so that it could be interrogated for status by software. The director was also designed to collect error information, which could also be obtained, on request, by software. The director also furnishes the port identifier, on

request, to a connected channel or control unit. Additionally, the switch performs the tie-breaking when two requests for the same target port are received. And, of course, the director can be instructed to fence off ports, create dedicated connections, or create partitions, by means of commands to the switch controller. Alternatively, the director can be configured manually at its control console.

◆ *ESCON control units*

Complementing the ESCON channels on the fiber optic network are the ESCON control units, which are provided with one or more ESCON fiber link adapters and continue to serve as the systems interface for multiple I/O devices. For example, certain models of the 3990 Storage Controller are supplied with up to 16 ESCON adapters, and thus support multiple access paths to any ESCON channel attached to the fiber network. Like the channels, the ESCON control units utilize the ESCON architecture protocols to communicate on the fiber network. Because it was required that the ESCON control units be implemented with record buffers, they are therefore potentially be able to transmit or receive data on the fiber network at channel speed, rather than the device speed of the parallel OEMI connections.

To minimize the connect time for operations, on receipt of a read request from the device to the channel, many of the ESCON control units disconnect from the channel (and disconnect the path), while data are being acquired from a device. After the data have been moved to the read buffer of the control unit, reconnect is requested. If dynamic pathing has been invoked, any available path can be used to establish the reconnect with the channel. On the reconnect operation, the control unit indicates to the channel that it is ready to transfer the requested data. Since the channel is to be the receiver of data in this case, the channel paces the transfer by sending "data request" control frames as it clears its data input buffers.

Similarly, on write operations, the channel transfers the data to a control unit buffer, and then the control unit initiates a disconnect while it is disposing of the data. When the write is complete, the control unit reconnects and informs the channel that the command has been completed.

While disconnected during one I/O operation, the channel can proceed with another I/O operation on a different subchannel, and the control unit can accept additional commands from the same or a different channel. Most ESCON control units, as well, are capable of concurrently processing multiple commands from different channels. For example, the IBM 3990 ESCON control units incorporate two controllers, each of which can be provided with four or eight ESCON link adapters. The attached DASD devices are provided with two input paths

to allow for access by either controller. Thus, I/O operations can be initiated by either controller, and reconnect may be made through either controller.

The IBM 3990 ESCON control unit [7], in order to assist migration, offers a special adapter microcode, which supports operation with the ESCON Converter Model 2. Any serial ESCON adapter can be so modified. The following section discusses the migration aids available.

◆ *ESCON migration aids*

To meet the objective to support an orderly migration to ESCON, two units were proposed to allow interconnections between installed System/370 and new System/390 devices. The first of these, the ESCON Converter Model 1, is intended to permit ESCON channels to perform I/O operations with installed System/370 control units. This allows the new System/390 to share I/O devices with the installed systems. The second migration aid, the ESCON Converter Model 2, was planned to allow the channels of the installed System/370 to attach to new ESCON control units, so that the new devices could be shared by the installed systems. The problem faced was that the control content of the ESCON architecture is greater than that available on the parallel OEMI channels; thus, while conceptually it is relatively easy to develop the OEMI controls as a subset of the ESCON controls, the reverse is not possible. It was therefore not possible to support dynamic Director pathing on the intervening fiber links.

The Converter Model 1 unit attaches by fiber link to an ESCON channel containing special microcode which prepares and receives the information exchanged with the Converter unit. This fiber link can be connected through a dedicated path of an ESCON Director, or be connected as a direct point-to-point link. The Converter Model 1 unit, driving parallel OEMI cables, emulates the System/370 parallel channel to operate System/370 devices. The parallel OEMI cables can attach to multiple System/370 control units in the normal multidrop fashion. The operation of the Converter Model 1 at a distance from the System/390 system is only possible because of the design decision to make the System/370 channel function remotely.

The ESCON Converter Model 2 unit attaches to a System/370 channel as a normal parallel OEMI control unit. The Converter unit then produces and interprets fiber optic transmission over a fiber optic link to a special adapter available on certain IBM 3990 models. In this case the Converter unit is designed to emulate an ESCON channel, but with the exception of the dynamic pathing information. For this reason, the Converter can only be connected to a cooperating ESCON 3990 converter adapter. The fiber link, however, can be connected by way of a dedicated ESCON Director path, and can function at




Figure 5

ESCON migration aids.

the normal ESCON link distance. **Figure 5** illustrates the possible configurations of the two ESCON Converters.

Summary

The ESCON facilities offer a substantially new approach to the attachment of control units to large systems channels. The first major I/O attachment architectural change since the announcement of System/360 in 1964, ESCON widens the scope of local channel attachment to campus distances, eliminates cable bulk, and provides dramatically higher data rates. The ESCON facilities improve system availability and simplify the management of the complex, while providing the basis for support of multiple systems with substantially greater processing power.

Since the introduction of ESCON in 1990, IBM has introduced a number of enhancements to the initial product set. The ESCON Extended Distance Feature (XDF) extended the maximum distance of an individual fiber optic

link to 20 km using laser technology and single-mode fiber. IBM introduced two new ESCON-attachable control units: the 9343 DASD Storage Subsystem and the 3495 Tape Library Subsystem. Maximum distances for some ESCON control units have been extended well beyond the initial 9 km (i.e., the 3990 to 15 km, the 3490 to 23 km, and CTC to 60 km). For customers that use the Processor Resource/Systems Manager™ (PR/SM™) capabilities of the ES/9000 product family [14], IBM introduced the ESCON Multiple Image Facility (EMIF), allowing the sharing of physical channels by any number of logical partitions (LPARs). With PR/SM installed, the resources of the processor are distributed among multiple system control programs, each of which has an independent set of resources called a logical partition. Additionally, IBM ADSTAR™ introduced three performance enhancements to the IBM 3990 DASD Subsystem that are available only on ESCON-attachable control units.

ESCON supports the creation of campus-wide high-speed fiber backbone interconnection networks. The ESCON system does not simply accommodate the attachment of I/O control units to systems; it allows sysplexes (complexes of systems) in different buildings to be interconnected to form a cohesive data processing utility. By adding the IBM gateway products, extensive network hierarchies can be configured as well. Jobs can readily be moved to the most appropriate system, and users can access databases anywhere on a campus as if they were on the local system. The ESCON interconnection facility can be characterized as a design of elegant simplicity. It is certainly simple in concept, but elegant in function and utility.

Acknowledgments

The system design and implementation of the ESCON system has been a complex and lengthy effort involving many hundreds of people over a ten-year period. Many people from technology, planning, research, architecture, engineering, programming, and other disciplines and specialties, in all of the development laboratories, have contributed substantial technical effort to bring its product elements to completion. For contributions to the system design described in this paper, the authors would especially like to recognize Joseph C. Elliott, John H. Sorg, Jr., and Martin W. Sachs. The authors also extend special thanks to Richard Cwiakala and Harry Yudenfriend for their assistance on the section on ESCON software support.

Enterprise Systems Connection Architecture, ESCON, System/360, ESCON Director, ESCON Manager, Enterprise System/9000, ES/9000, System/370, Enterprise Systems Architecture/390, ESA/390, MVS/ESA, Processor Resource/Systems Manager, PR/SM, and ADSTAR are trademarks, and System/390 is a registered trademark, of International Business Machines Corporation.

References

1. *Introducing Enterprise Systems Connection*, Order No. GA23-0383-1, March 1992; available through IBM branch offices.
2. 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, and D. J. Stigliani, Jr., "IBM Enterprise Systems Multimode Fiber Optic Technology," *IBM J. Res. Develop.* **36**, 553-576 (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. J. R. Flanagan, T. A. Gregg, and D. F. Casper, "The IBM Enterprise Systems Connection (ESCON) Channel: A Versatile Building Block," *IBM J. Res. Develop.* **36**, 617-632 (1992, this issue).
5. 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).
6. J. J. Coleman, C. B. Meltzer, and J. L. Weiner, "Fiber Distributed Data Interface Attachment to System/390," *IBM J. Res. Develop.* **36**, 647-654 (1992, this issue).
7. C. P. Grossman, "Role of the DASD Storage Control in an Enterprise Systems Connection Environment," *IBM Syst. J.* **31**, 123-146 (1992).
8. *IBM System/360 and System/370 I/O Interface Channel to Control Unit Original Equipment Manufacturers' Information*, Order No. GA22-6974-09, January 1991, available through IBM branch offices.
9. *Synchronous Data Link Control General Information*, Order No. GA27-3093-03, January 1991; available through IBM branch offices.
10. A. X. Widmer and P. A. Franaszek, "A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code," *IBM J. Res. Develop.* **27**, 400-451 (1983).
11. *IBM ESA/390 Principles of Operation*, Order No. SA22-7201-00, March 1992; available through IBM branch offices.
12. A. S. Tanenbaum, *Computer Networks*, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1989.
13. *IBM Enterprise Systems Architecture/390 ESCON I/O Interface*, Order No. SA22-7202-01, March 1992; available through IBM branch offices.
14. *32389 ES/9000 PR/SM Introduction and Operations*, Order No. SR21-3674-00, February 1992; available through IBM branch offices.

Received April 18, 1991; accepted for publication April 27, 1992

Salvatore A. Calta *IBM Enterprise Systems, P.O. Box 950, Poughkeepsie, New York 12602 (retired).* Mr. Calta was a Senior Engineer in the systems design area of Large Systems Planning. He joined IBM in 1956 as an I/O channel designer on the 709 systems development program, and subsequently served as a channel designer on 7000 systems and as a storage subsystem designer on the System/360 Model 91/95 programs. Since 1968, he had specialized in the system architecture of large systems, and had held numerous system design positions, with a primary focus on I/O subsystems. In 1991, he received an IBM Outstanding Innovation Award for his contributions as one of the original developers of the ESCON system design.

John A. deVeer *IBM Enterprise Systems, P.O. Box 950, Poughkeepsie, New York 12602 (retired).* Mr. deVeer was a Senior Engineer in the I/O Architecture and System Structure Department of Large Systems Planning. He joined IBM in 1950 as a CE and served as a CE instructor and as a large systems designer on various 700 and 7000 systems. Mr. deVeer served on the System/360 design council and was technical coordinator of the Model 75. He led several FSD investment program developments before joining the group where he participated in the system design of ESCON. He received his B.S.E.E. from Swarthmore College, and his S.M. from Harvard University. Mr. deVeer holds six patents. In 1991, he received an IBM Outstanding Technical Achievement Award for ESCON system design.

Edward Loizides *IBM Enterprise Systems, P.O. Box 950, Poughkeepsie, New York 12602 (LOIZIDES at TDCSYS2).* Mr. Loizides is a Senior Engineer and manager of I/O Architecture and System Structure at the ES Mid-Hudson Valley Laboratory. He received his B.A. degree from Columbia College in 1958 and his B.S.E.E. and M.S.E.E. degrees from the Columbia University School of Engineering in 1959 and 1960, respectively. He subsequently joined IBM at the Poughkeepsie Development Laboratory, where he has worked on I/O and communications subsystems for System/360, System/370, and System/390. Mr. Loizides holds an IBM Second-Level Invention Achievement Award. In 1991 he received an IBM Outstanding Technical Achievement Award for ESCON system design.

Richard N. Strangwayes *IBM Enterprise Systems, P.O. Box 950, Poughkeepsie, New York 12602 (deceased).* Mr. Strangwayes joined IBM in Endicott, New York, as a Test Engineer in Product Assurance after receiving his B.S. degree at Fairleigh Dickinson University. In 1960, he transferred to Poughkeepsie with the 7070 Machine Program. He subsequently held several engineering and management positions and in 1982 became Program Manager of the I/O Architecture System Design Group, which, under his direction, played a major role in the origination, innovation, development, and coordination of the ESCON program. In 1990, Mr. Strangwayes was named Senior Technical Staff Member, Mid-Hudson Valley Laboratory. In this position, reporting to the Future Systems Design Manager, he had responsibility for developing future hardware and software systems interconnect applications and solutions.