

## DISC INTERFACES FOR MPE-XL SYSTEMS

Gary Vogelsberg  
Hewlett-Packard Company  
Disc Memory Division  
P.O. Box 39  
Boise, Idaho 83707


### INTRODUCTION

With the introduction of HP-FL, Hewlett-Packard's new fiber-optic link, there are now two disc interface alternatives for HP 3000 MPE-XL systems (they are also supported on Series 800 HP-UX systems). This paper is intended to give the HP 3000 computer user greater insight into those disc interface alternatives and how they work. It will also provide some guidelines for deciding which interface is appropriate for the user's system.

### WHAT IS A DISC INTERFACE?

The disc drive is a storage device, used to store data so it can be accessed by the CPU. Internal to the CPU, data is stored in memory. There are various hierarchies of memory within the CPU, allowing the CPU to be designed for optimum price/performance. Because CPU memory is expensive and volatile, disc drives are used to complement CPU memory. They provide a place where large amounts of data can be stored in a non-volatile form, and be easily accessible to the system.

The disc interface is the connection between the disc drive and the CPU. The interface provides a means of moving data between disc mechanisms and memory within the CPU. Examples of disc interfaces on HP systems include the parallel differential interface on the 7906/20/25 disc drives; HP-IB, which is used on many current products; and HP-FL, which was introduced earlier this year and will be shipping soon on MPE-XL systems. The components of a typical disc interface are pictured below.



## Driver/Interface Card

Within the CPU, hardware and software work together to manage disc I/O's. The disc driver is the software portion, while the interface card is the hardware portion. The interface card plugs into the internal CPU bus. The purpose of the interface card is to receive I/O requests issued across the CPU bus by the driver and to communicate those requests to the disc controllers. The interface card then manages the communications necessary to complete the I/O request. This includes managing the transmission or reception of the data as packets across the cable.

The intelligence of the interface card varies. Microprocessors can be placed on the interface, effectively downloading some of the functions off the main CPU. This can result in less CPU overhead and improved performance.

## Cable/Protocol

The cable is the physical connection between the interface card and the disc controller. The cable traditionally consists of multiple wires, and is used to send commands and data between the interface card and the disc controller.

Commands are sent across the cable using a protocol. The protocol can be defined as the language used by the interface card and the controller as they communicate across the cable.

## Controller

The disc controller's function is to provide the intelligence required to execute commands issued across the interface. The main functions of the disc controller are to decode commands, execute the commands, manage error recovery when necessary, manage the transmission of data across the cable, and report the execution status of commands back to the host. The controller can also be designed to provide diagnostics to aid in troubleshooting a disc drive.

## HOW THE COMPONENTS WORK TOGETHER - THE TYPICAL DISC I/O

Now that the pieces of a disc interface have been defined, let's take a look at the execution of a typical disc input/output (I/O) request. The I/O request is initiated by the CPU when it discovers that it needs data that does not reside in memory, or that some data needs to be posted to disc. The CPU decides what data it needs to transmit/receive, and then issues a command across the CPU

bus. This portion of the I/O is referred to as CPU overhead.

Upon receipt of the I/O request, the interface card looks for idle time on the bus and sends the command down to the disc controller. The disc controller then acknowledges receipt of the I/O request. The time it takes the interface to receive and then issue these commands is referred to as interface overhead.

After the controller receives the command, it decodes the command and issues a command to the disc to seek to the correct location. The time required to do this is referred to as controller overhead.

The disc mechanism must then execute the requested command. It does this by means of a mechanical movement of the actuator, placing the heads in the disc drive over the correct track. Once the heads are over the correct track, additional time is required for the disc surface to rotate so that the correct data is under the head. The time required to move the heads to the correct position is referred to as seek time, and the time spent waiting for the correct data to rotate under the head is called latency.

Once all of this is complete, data can be transferred. Data is sent to or from the disc mechanism through the interface, cable, and controller. Once the data transmission is complete, the controller sends a message that the transaction has completed properly. The time spent transmitting data to/from disc is known as transfer time.

In a normal disc transaction on current HP systems, the majority of the time is spent in seek and latency. CPU overhead, interface overhead, and controller overhead are generally small components of the I/O execution time. Transfer time is variable, and is dependent on the amount of data being transferred.

#### HP-IB COMPONENTS

HP-IB disc drives have now been available on HP systems for over eight years. This interface, offering one of the first intelligent controllers, is based on the IEEE 488 electrical specification and the CS-80 protocol. This section will give a short overview of the hardware components of the HP-IB interface.

There are three basic hardware components of an HP-IB connection: the HP-IB interface card, the HP-IB cable, and the HP-IB disc controller. The HP-IB interface card plugs

into the CPU bus. The HP-IB interface card provides a means of attaching multiple types of devices to the CPU. Supported peripherals include disc drives, tape drives, printers, and plotters. There is an HP-IB connector provided with the interface card, allowing the connecting of the HP-IB cable.

The HP-IB cable is a multi-wire copper cable which plugs into the HP-IB connector on the interface card. HP-IB provides 8 device addresses, so up to 8 peripherals can be daisy chained together off one HP-IB interface card. The open back of the HP-IB connector provides for easy daisy-chaining of cables.

The HP-IB disc controller resides in the disc drive. The HP-IB cable plugs into the controller. As I/O requests are received by the addressed controller across the HP-IB cable, the controller accepts and decodes the requests, and gives commands to the disc mechanism to execute the requests.

The CS-80 command set is the protocol used to communicate with discs across the interface. All hardware components of the disc interface are designed to communicate using this protocol.

#### HP-FL COMPONENTS

HP-FL, Hewlett-Packard's fiber-optic-link, was introduced earlier this year, and will be shipping on MPE-XL systems in the near future. Since the new interface is only recently introduced, this section will give a more in-depth view of the components of the HP-FL interface, outlining how they work together.

#### Cabling

There are two methods of cabling incorporated into the new HP-FL architecture. A fiber-optic cable is used to connect the CPU to a group of disc drives, while a multi-wire PBus cable is used to daisy chain a group of disc drives together.

The fiber-optic cable is a duplex cable of glass fiber; there are actually two fiber-optic strands in the cable. One strand is used to transmit data in the direction of the disc drives, while the other is used to send data back to the CPU. The fiber-optic cable has a burst transfer rate of 5 megabytes per second, and is supported in lengths of up to 500 metres.

The PBus cable is a 64-pin copper cable. It is used to daisy-chain up to eight drives to the fiber-optic cable. The PBus cable is designed to operate at a burst transfer rate of 5 megabytes per second, the same transfer rate achieved across the fiber-optic cable.

#### The HP-FL Interface Card

The HP-FL interface is a CPU resident card, providing an interface between peripherals and the CIO backplane on HP PA systems. Each HP-FL interface card has an improved CIO backplane interface circuit, protocol controller, high-speed parallel/serial/parallel converter and encoder/decoder, and high-performance microprocessor. A pair of fiber-optic connectors is also included with each HP-FL interface card.

The fiber-optic connectors provide a means of attaching the fiber-optic cable to the interface card. As previously noted, the fiber-optic cable consists of two fiber-optic strands. One of the connectors, equipped with a fiber-optic transmitter, is used to transmit data from the interface card to the group of disc drives. The other connector, equipped with a fiber-optic receiver, is used to receive data transmitted from the group of disc drives.

The function of the electronics is to convert electronic signals into a format allowing them to be communicated across the fiber-optic cable. When data is to be sent from the CPU to the disc drive, the interface card accepts the data issued in parallel across the CIO backplane and converts it into serial signals with the proper protocol. An LED in the transmitter then sends the data, in the form of light pulses, across the fiber-optic cable. The reverse is done when signals are sent back to the CPU from the disc drive; the interface card accepts the light pulses transmitted across the fiber-optic cable, converts them into parallel signals, and forwards them on to the CPU across the CIO backplane.

The high-performance microprocessor on the board allows the downloading of I/O processing from the CPU, reducing the amount of CPU time spent executing I/O's. The microprocessor also provides the interface card with the "horsepower" required to meet the critical timing required by the interface's high transfer rate. It should be noted that a large amount of custom VLSI work (for which HP is recognized as an industry leader) was required to provide the performance and functionality of the new interface at a reasonable cost.

## The HP-FL Controller

The HP-FL controller is disc resident, replacing the HP-IB controller in 7936 and 7937 disc drives. The function of the disc controller is to manage communications between the disc drives and the CPU.

Like the interface card, each HP-FL controller is equipped with two fiber-optic connectors, one for sending data and one for receiving data. The fiber-optic strand connected to the transmitter on the interface card must be connected to the receiver on the controller card, and the fiber-optic strand connected to the receiver on the interface card must be connected to the transmitter on the controller card. LEDs on the controller indicate if the fiber-optic cable is not properly connected.

The HP-FL controller also has two connectors for PBus cables. Using these connectors, up to eight disc drives can be daisy-chained together. The end drives in a chain must have PBus terminators on their unused connectors.

With two cabling methods incorporated into the HP-FL architecture, the HP-FL controller becomes the point of management of communications moving from the fiber-optic cable to the PBus cable. As serial light pulses are received across the fiber-optic cable, the controller converts them into a parallel signal that can be sent across the PBus cable to the appropriate disc drive. Signals sent across the PBus cable destined for the CPU are converted from parallel to serial so that they can be transmitted as light across the fiber-optic cable. The interface card and controller share some common electronics, since the same functions of encoding and decoding signals and converting them back and forth from parallel to serial must be executed at both ends of the fiber-optic connection.

Once the command is received at the appropriate disc drive, the controller carries out the traditional functions of decoding the commands, executing the commands, executing error correction, and reporting back to the host.

The HP-FL interface, with its 5 megabytes per second burst transfer rate, has a much higher transfer rate than the 7937 disc drive. To use the channel efficiently, the controller has been designed to allow interleaved transfers of data from multiple disc drives. This is accomplished through buffering in the controller and active management of the PBus connection. As data is read off disc, it is stored in the controller buffer. When a critical mass of data is stored, the controller arbitrates for ownership of the

interface, sends its data at the full 5 megabytes per second transfer rate of the interface, and disconnects. The controller will connect/disconnect several times when executing a large transfer. This intelligent use of the interface allows it to be used efficiently by multiple drives.

Since all components have very high burst transfer rates, the components of the controller must be designed to meet critical timing requirements. The controller was designed with precision and performance in mind, using custom VLSI extensively.

#### PERFORMANCE IMPACT OF HP-FL

The following shows the time to process a typical MPE-XL I/O with a 7937H and 7937FL disc drive once it gets to the disc controller. The HP-FL disc controller overhead is slightly larger than that of the HP-IB controller. It should be noted that a feature called "command queueing" allows the pre-processing of controller overhead, effectively masking controller overhead on busy disc drives. Since the 7937 disc mechanism is common to both products, the seek and latency times are the same.

The typical transfer is assumed to be 8 kilobytes on an MPE-XL system, as compared to 1 kilobyte on an MPE system. Although the HP-FL interface has a roughly five times higher transfer rate, transfer time is reduced by less than half. Since the interface rate is now faster than the disc transfer rate; transfer time for an I/O is determined by the disc transfer rate (1.89 megabytes per second across one track).

	7937H	7937FL
Controller overhead	1.0 ms	1.5 ms
Seek time	20.5 ms	20.5 ms
Latency	8.3 ms	8.3 ms
Transfer time	8.0 ms	4.2 ms
	-----	-----
Total access time	37.8 ms	34.5 ms
I/O index	26.5	29.0

The comparison shows that the typical 8 kilobyte I/O will execute in 9% less time using the HP-FL interface. At first glance, it might be expected that the HP-FL interface will have an immediate impact of at least 9% on HP PA system performance. However, there are some reasons that this is not the case.

Total access time is an important measure of disc performance. However, it is not important because a given disc drive allows the system to get back to a given user 3 or 4 milliseconds faster, but because it provides a relative measure of how many I/O's a disc drive is capable of executing.

#### Impact on Today's MPE-XL System

In the early design phases of precision architecture systems, conscious decisions were made to improve system level performance by reducing disc I/O requirements. The transaction look-aside buffer, memory mapped files, seek aheads, and the transaction manager were all designed to reduce the number of disc I/Os. Because of these features, MPE-XL systems will require far fewer I/Os than an MPE-V system to complete the same task.

Because of the reduced I/O requirements of MPE-XL, the HP-IB interface is able to keep up with typical MPE-XL I/O workloads. The net result is that HP-FL will have little if any impact on system level performance on today's typical MPE-XL systems.

#### Impact on Future MPE-XL Systems

Although 'HP-FL has little impact on the performance of today's typical MPE-XL system, there are factors that will cause the disc interface to have a performance impact in the future. Some of the factors that will cause I/O loads on the interface to increase are listed below.

HP PA systems are early in their life cycle. As new systems continue to roll out, extending the family to higher performance levels, there will be corresponding increases in I/O requirements of those systems.

Along with the increases in processing power, there will also be larger mass storage configurations on those larger systems. As more gigabytes of data are stored on each interface, there will be corresponding increases in the number of I/O's executed across the interface.

With the initial release of MPE-XL, typical disc transfers are larger than they were on MPE-V systems. Since I/O loads are relatively light today, this has not yet become a factor in performance. However, as the number of I/O's increases with larger processors and disc configurations, the larger I/O's will put an added strain on the disc interface. Growing I/O requirements coupled with larger transfers will put pressure on interface transfer rates over time.



## WHICH INTERFACE IS FOR ME??

With the introduction of HP-FL, there are now two interfaces for attaching drives to MPE-XL systems. Each of the interfaces offers advantages to the customer. Those advantages, as well as guidelines for making the right decision, are outlined below.

### Benefits of HP-IB

HP-IB continues to provide an easy inexpensive method of attaching a wide variety of peripherals and instruments to HP systems. On HP systems, the same HP-IB interface can be used (within supported configuration guidelines) to attach discs, tapes, printers, and plotters to the system. Because it is inexpensive and can be used for multiple peripherals, the HP-IB interface continues to make a lot of sense for small systems.

The common interface on systems and peripherals also provides an easy migration path. Since HP-IB is supported on systems from the HP 150 to the HP 3000 Series 950, there is a great deal of flexibility in moving peripherals upward as computing needs grow.

The transfer rate of HP-IB is not a bottleneck to system performance on today's HP PA systems in most situations. HP-IB will meet the performance requirements of systems that will not be growing significantly.

### Benefits of HP-FL

Eight disc drives are supported on one HP-FL interface card. Because more disc drives are supported per HP-FL interface, HP supports larger mass storage configurations than with HP-IB. Up to 30 7937FL disc drives (17 gigabytes of mass storage) are supported on the Series 950 at first release, and supported configurations will be increased with future releases of the operating system.

A side-benefit of the larger number of discs supported per interface is that CPU I/O slot usage is reduced. The large configurations noted above can be built using a minimal number of I/O slots.

Long fiber-optic cable lengths also provides a great deal of flexibility in laying-out a data center. Some of you have probably experienced the challenges of laying-out large disc configurations while staying within supported cable length

constraints of HP-IB. With supported cable lengths of up to 500 metres, these problems will disappear.

HP-FL is the interface of the future for high-end mass storage for Series 900 HP 3000 systems. The improved transfer rate and flexibility provide a mass storage growth path for the future. HP-FL also provides a platform for designing future mass storage solutions. Future discs for high-end systems will be HP-FL compatible.

Because data is transmitted via light, the new fiber-optic cable also reduces environmental concerns. The fiber-optic cable is immune to electromagnetic interference. As a result, long cables can be run without worrying about emissions from equipment near the cabling route. The fiber-optic cable also provides electrical isolation of discs from the system, and does not emit radio frequency energy that might cause interference with other equipment.

### Conclusions

Today, Hewlett-Packard is offering interfacing alternatives on MPE-XL systems. Customers anticipating the need for large configurations, longer cables, a growth path for the future, or the environmental advantages of HP-FL should purchase new disc drives with the HP-FL interface. By purchasing FL drives today at a small price premium, the larger future costs of a field upgrade are avoided. It should be remembered, however, that there will probably be no immediate impact on system-level performance.

Those customers concerned about price and not anticipating a need for the additional features of HP-FL should continue to purchase HP-IB drives. If the decision is made to migrate to HP-FL in the future, they will have an upgrade kit available to make the transition (at a price premium).