I/O TECHNIQUES

Computers generate, process and delete data. However, they can only store data for very short periods. Therefore, computers move data to storage devices such as tape libraries and the disk subsystems discussed in the previous chapter for long-term storage and fetch it back from these storage media for further processing. So-called I/O techniques realise the data transfer between computers and storage devices. This chapter describes I/O techniques that are currently in use or that the authors believe will most probably be used in the coming years.

4.1 THE PHYSICAL I/O PATH FROM THE CPU TO THE STORAGE SYSTEM

In the computer, one or more CPUs process data that is stored in the CPU cache or in the random access memory (RAM). CPU cache and RAM are very fast; however, their data is lost when the power is been switched off. Furthermore, RAM is expensive in comparison to disk and tape storage. Therefore, the data is moved from the RAM to storage devices such as disk subsystems and tape libraries via system bus, host bus and I/O bus (Figure 3.1). Although storage devices are slower than CPU cache and RAM, Computers generate, process and delete data. However, they can only store data for very short periods. Therefore, computers move data to storage devices such as tape libraries and the disk subsystems discussed in the previous chapter for long-term storage and fetch it back from these storage media for further processing. So-called I/O techniques realise the data transfer between computers and storage devices. This chapter describes I/O techniques that are currently in use or that the authors believe will most probably be used in the coming years.

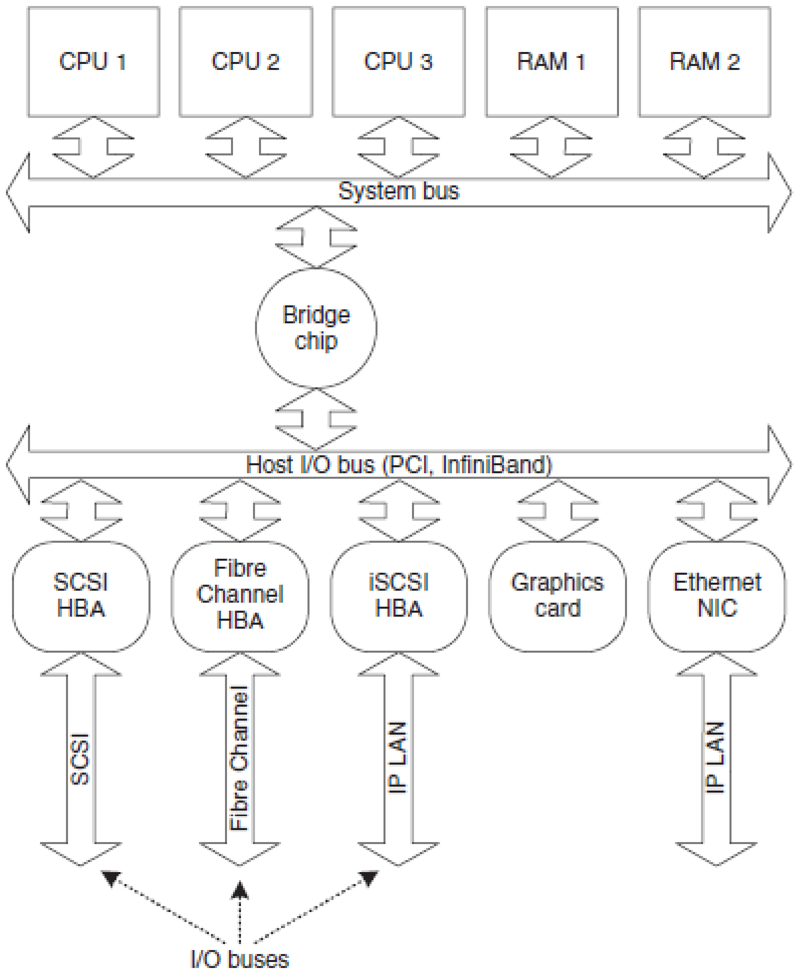

Figure 4.1 The physical I/O path from the CPU to the storage system consists of system bus, host I/O bus and I/O bus.

More recent technologies such as InfiniBand, Fibre Channel and Internet SCSI (iSCSI) replace individual buses with a serial network. For historic reasons the corresponding connections are still called host I/O bus or I/O bus. they compensate for this by being cheaper and by their ability to store data even when the power is switched off. Incidentally, the same I/O path also exists within a disk subsystem between the connection ports and the disk subsystem controller and between the controller and the internal hard disk (Figure 3.2).

At the heart of the computer, the system bus ensures the rapid transfer of data between

CPUs and RAM. The system bus must be clocked at a very high frequency so that it can supply the CPU with data sufficiently quickly. It is realised in the form of printed conductors on the main circuit board. Due to physical properties, high clock rates require short printed conductors. Therefore, the system bus is kept as short as possible and thus connects only CPUs and main memory.

In modern computers as many tasks as possible are moved to special purpose processors such as graphics processors in order to free up the CPU for the processing of the application. These cannot be connected to the system bus due to the physical limitations mentioned above. Therefore, most computer architectures realise a second bus, the so-called host I/O bus. So-called bridge communication chips provide the connection between system bus and host I/O bus. Peripheral Component Interconnect (PCI) is currently the most widespread technology for the realisation of host I/O buses.

Device drivers are responsible for the control of and communication with peripheral devices of all types. The device drivers for storage devices are partially realised in the form of software that is processed by the CPU. However, part of the device driver for the communication with storage devices is almost always realised by firmware that is processed by special processors (Application Specific Integrated Circuits, ASICs). These ASICs are currently partially integrated into the main circuit board, such as on-board SCSI controllers, or connected to the main board via add-on cards (PCI cards). These add-on cards are usually called network cards (Network Interface Controller, NIC) or simply controllers. Storage devices are connected to the server via the host bus adapter (HBA) or via the on-board controller. The communication connection between controller and peripheral device is called the I/O bus.

The most important technologies for I/O buses are currently SCSI and Fibre Channel. SCSI defines a parallel bus that can connect up to 16 servers and storage devices with one another. Fibre Channel, on the other hand, defines different topologies for storage networks that can connect several millions of servers and storage devices. As an alternative to Fibre Channel, the industry is currently experimenting with different options for the realisation of storage networks by means of TCP/IP and Ethernet such as IP storage and FCoE. It is worth noting that all new technologies continue to use the SCSI protocol for device communication.

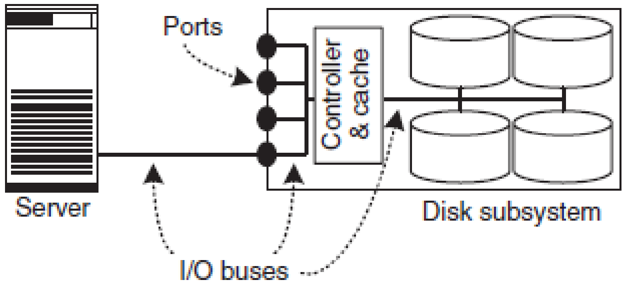

Figure 4.2 The same I/O techniques are used within a disk subsystem as those used between server and disk subsystem.

The Virtual Interface Architecture (VIA) is a further I/O protocol. VIA permits rapid and CPU-saving data exchange between two processes that run on two different servers or storage devices. In contrast to the I/O techniques discussed previously VIA defines only a protocol. As a medium it requires the existence of a powerful and low-error communication path, which is realised, for example, by means of Fibre Channel, Gigabit Ethernet or InfiniBand. VIA could become an important technology for storage networks and server clusters. There are numerous other I/O bus technologies on the market that will not be discussed further in this book, for example, Serial Storage Architecture (SSA), IEEE 1394 (Apple's Firewire, Sony's i.Link), High-Performance Parallel Interface (HIPPI), Advanced Technology Attachment (ATA)/Integrated Drive Electronics (IDE), Serial ATA (SATA), Serial Attached SCSI (SAS) and Universal Serial Bus (USB). All have in common that they are either used by very few manufacturers or are not powerful enough for the connection of servers and storage devices. Some of these technologies can form small storage networks. However, none is anywhere near as flexible and scalable as the Fibre Channel and IP storage technologies described in this book.