



The Essentials of Computer Organization and Architecture

Linda Null and Julia Lobur
Jones and Bartlett Publishers, 2003

Chapter 7 Instructor's Manual

Chapter Objectives

Chapter 7, Input/Output and Storage Systems, provides a detailed overview of I/O fundamentals, bus communication and protocols, and typical external storage devices, such as magnetic and optical disks, as well as the various formats available for each. DMA, programmed I/O, and interrupts are covered as well. In addition, various techniques for exchanging information between devices are introduced. RAID architectures are covered in detail, and various data compression formats are introduced.

This is a very large chapter, but instructors need not lecture on all topics. This was written with the intention of students reading the material. Many of the topics (such as disk layouts and logical formats) can easily be read and understood by the student. In addition, some material is more FYI in nature. However, there are some topics which should be covered in lecture. Lectures over the material in Chapter 7 should focus on the following points:

- **Amdahls's Law.** This law characterizes the interrelationship of the components within the system. Understanding the speedup potential in a system is particularly important for the concepts in Chapter 10, Performance.
- **I/O architectures.** I/O control methods (programmed I/O, interrupt-driven I/O, DMA, and channel I/O) and I/O bus operation are important topics in helping students understand the system as a whole.
- **RAID.** RAID devices allow for redundancy (in different ways) in storing data, thus offering improved performance and increased availability for systems employing these devices. Levels 0 through 6, in addition to some hybrid systems, are introduced.
- **Data compression.** Data compression not only helps economize on disk usage, but it reduces transmission time in data communications as well. This section contains some mathematically challenging concepts and should probably not be left for the student to read on his or her own.

An instructor could also lecture over sections 7.4 (Magnetic Disk Technology), 7.5 (Optical Disks), and 7.6 (Magnetic Tape). However, these sections are easily read by the student.

Required Lecture Time

The important concepts in Chapter 7 can typically be covered in 2 lecture hours. However, if a teacher wants the students to have a mastery of all topics in Chapter 7, 6 lecture hours are more reasonable.

Lecture Tips

Students often have difficulty understanding the importance and application of Amdahl's Law. Going over several examples (perhaps looking ahead to Chapter 10) can help motivate the study of this law.

It is important to be specific about the differences among the I/O control methods. In particular, instructors should focus on interrupt-driven I/O and how it differs from DMA, as both involve interrupts.

Some of the RAID levels are theoretical and have not been deployed, and this should be pointed out in lecture. While covering these levels, the advantages and disadvantages of each should be covered, in addition to situations in which one level is better or worse than another.

Answers to Exercises

- ◆ 1. Your friend has just bought a new personal computer. She tells you that her new system runs at 1GHz, which makes it over three times faster than her old 300 MHz system. What would you tell her?

Ans.

Hint: Consider how Amdahl's Law applies. You would explain Amdahl's Law. The processor is only one component contributing to the overall performance of a system.

-
2. Suppose the daytime processing load consists of 60% CPU activity and 40% disk activity. Your customers are complaining that the system is slow. After doing some research, you have learned that you can upgrade your disks for \$8,000 to make them 2.5 times as fast as they are currently. You have also learned that you can upgrade your CPU to make it 1.4 as fast for \$5,000.
- Which would you choose to yield the best performance improvement for the least amount of money?
 - Which option would you choose if you don't care about the money, but want a faster system?
 - What is the break-even point for the upgrades? That is, what price would be charged for both upgrades to make their cost and performance improvement equal?

Ans.

Fraction of work: 60% CPU, 40% disk.

$$S_{\text{CPU}} = 1 / ((1-f) + (f/k)) = 1 / ((1-0.60) + (0.60/1.4)) = 120.69\%$$

$$S_{\text{DISK}} = 1 / ((1-f) + (f/k)) = 1 / ((1-0.40) + (0.40/2.5)) = 131.58\%$$

- Choose the CPU upgrade:
\$5000/120.69% = \$41.43

$$\$8000/131.58\% = \$38.00$$

- b. The disk option gives a better performance improvement.
 - c. All things are equal when you pay \$5451.16 for the disk upgrade.]
-

- ◆ 3. How would you answer Question 2 if the system activity consists of 55% processor time and 45% disk activity?

Ans.

Fraction of work: 55% CPU, 45% disk.

$$S_{\text{CPU}} = 1/((1-f)+(f/k)) = 1/((1-0.55)+(0.55/1.4)) = 118.64\%$$

$$S_{\text{DISK}} = 1/((1-f)+(f/k)) = 1/((1-0.45)+(0.45/2.5)) = 136.99\%$$

- a. Choose the CPU upgrade. This will cost \$42.14 per 1% improvement versus \$58.40 for the disk.
 - b. The disk upgrade gives the greater improvement: 136.99% versus 118.64% for the processor.
 - c. The break-even point would be with a disk upgrade costing \$5772.59.
-

- 4. Name the four types of I/O architectures. Where are each of these typically used and why are they used there?

Ans.

- 1. Programmed I/O – good for control applications, e.g., sensors
 - 2. Interrupt-driven I/O – used in personal systems
 - 3. DMA – good for small to medium-sized systems
 - 4. Channel I/O – used in the largest enterprise-class systems
-

- ◆ 5. A CPU with interrupt-driven I/O is busy servicing a disk request. While the CPU is midway through the disk-service routine, another I/O interrupt occurs.

- ◆ a. What happens next?
- ◆ b. Is it a problem?
- ◆ c. If not, why not? If so, what can be done about it?

Ans.

- a. A CPU should disable all interrupts before it enters an interrupt service routine, so the interrupt shouldn't happen in the first place.
- b. Not a problem.
- c. If interrupts are disabled, the second interrupt would never happen so it is not a problem.

Alternative good answer:

- a. The CPU could save all information from the interrupt service routine that it was working on, just as if it were a user program.
 - b. Could be a problem.
 - c. Information coming from multiple sources at the same time could get garbled. The best thing to do is to disable all maskable interrupts while I/O is occurring.)
-

- 6. Why are I/O buses provided with clock signals?

Ans.

Bus clocks maintain synchronization and define bit cell boundaries.

7. If an address bus needs to be able to address eight devices, how many conductors will be required? What if each of those devices also needs to be able to talk back to the I/O control device?

Ans.

Three conductors will be needed. To provide bidirectionality, either we add three more conductors for the other direction or add a single control line that indicated the "direction" of the signal on the other three.

8. We pointed out that I/O buses do not need separate address lines. Construct a timing diagram similar to Figure 7.7 that describes the handshake between an I/O controller and a disk controller for a write operation. (Hint: You will need to add a control signal.)

Ans.

Note: The control signal can be either "error" or "ack."

- * 9. If each interval shown in Figure 7.7 is 50 nanoseconds, how long would it take to transfer 10 bytes of data? Devise a bus protocol, using as many control lines as you need, that would reduce the time required for this transfer to take place? What happens if the address lines are eliminated and the data bus is used for addressing instead? (An additional control line may be needed.)

Ans.

Four time slots to set up the transfer => $50 * 4 = 200\text{ns}$

Ten timeslots to transfer the data => $50 * 10 = 500\text{ns}$

One timeslot to tear down transfer => 50ns .

Total time = $200\text{ns} + 500\text{ns} + 50\text{ns} = 750\text{ns}$.

10. Define the terms seek time, rotational delay, and transfer time. Explain their relationship.

Ans.

Seek time is the time it takes for a disk arm to position itself over a requested track. Rotational delay is the time that it takes for the required sector to position itself under a read/write head. The sum of the rotational delay and seek time is known as the access time. If we add to the access time the time that it takes to actually read the data from the disk, we get transfer time.

- ◆ 11. Why do you think the term random access device something of a misnomer for disk drives?

Ans.

Some people think that retrieving specific data from a particular disk is not a "random" act.

12. Why do differing systems place disk directories in different track locations on the disk? What are the advantages of using each location that you cited?

Ans.

Some systems place the directory in the center of the disk for faster access time. Others place it on the outer tracks because the medium has fewer errors there. Outer track directory placement provides greater data integrity (less error recovery), but does not, in general, perform as well as center disk placement.

- ◆ 13. Verify the average latency rate cited in the disk specification of Figure 7.11. Why is the calculation divided by 2?

Ans.

$[(60 \text{ seconds}/4,464 \text{ rpm}) * (1000\text{ms}/\text{second})]/2 = 6.72 \text{ ms}$. The answer is divided by 2 because, on average, the required sector will be 180 degrees from the current position of the read/write head.]

14. By inspection of the disk specification in Figure 7.11, what can you say about whether the disk drive uses zoned-bit recording?

Ans.

The disk does not provide zoned-bit recording because the number of sectors per track is fixed.

15. The disk specification in Figure 7.11 gives a data transfer rate of 6.0 MB per second when reading from the disk, and 11.1 MB per second when writing to the disk. Why are these numbers different?

Ans.

Disk writes are buffered.

16. Do you trust disk drive MTTF figures? Explain.

Ans.

MTTF is a theoretical, statistical calculation with a tentative relationship to reality. MTTF methodologies can vary from vendor to vendor, so they are of limited use for making comparisons unless the methodology is mandated—such as in a government contract.]

- ◆ 17. Suppose a disk drive has the following characteristics:

- 4 surfaces
- 1,024 tracks per surface
- 128 sectors per track
- 512 bytes/sector
- Tract-to-track seek time of 5 milliseconds
- Rotational speed of 5,000 RPM.

- ◆ a. What is the capacity of the drive?
- ◆ b. What is the access time?

Ans.

- a. $4 \text{ surfaces} \times 1,024 \text{ tracks per surface} \times 128 \text{ sectors per track} \times 512 \text{ bytes/sector} = 4 * 1024 * 128 * 512 / (2^{20} \text{ bytes/MB}) = 256\text{MB}$.
 - b. Rotational delay = $(60 \text{ seconds}/5000 \text{ rpm}) * (1000\text{ms}/\text{second})/2 = 6 \text{ ms}$ + 5ms seek time = 11 ms.
-

18. Suppose a disk drive has the following characteristics:

- 5 surfaces
- 1,024 tracks per surface
- 256 sectors per track
- 512 bytes/sector

- Track-to-track seek time of 8 milliseconds
- Rotational speed of 7,500 RPM.
-
- a. What is the capacity of the drive?
- b. What is the access time?
- c. Is this disk faster than the one described in question 17? Explain.

Ans.

- a. $5 \text{ surfaces} \times 1,024 \text{ tracks per surface} \times 256 \text{ sectors per track} \times 512 \text{ bytes/sector} = 5 \times 1024 \times 256 \times 512 / (2^{20} \text{ bytes/MB}) = 640\text{MB}.$
 - b. $\text{Rotational delay} = (60 \text{ seconds}/7500 \text{ rpm}) \times (1000\text{ms}/\text{second}) / 2 = 4 \text{ ms} + 8\text{ms seek time} = 12 \text{ ms}.$
 - c. The track-to-track seek time makes it slower.
-

19. What are the advantages and disadvantages of having a small number of sectors per disk cluster?

Ans.

The advantage with a smaller number of sectors per cluster, is that you get more efficient usage (less wasted space) on the disk. The disadvantage is that the disk directory (or FAT) gets very large and may slow things down.

* 20. Suggest some ways in which the performance of a 1.44 MB floppy disk could be improved.

Ans.

The chaining could take place through special entries in the data sectors (e.g. header or trailer fields) that would take the place of continually re-reading the FAT. Of course, the FAT would need to be updated on writes so that the sectors are marked as "in use."

21. What is the maximum number of root directory entries on a 1.44 MB floppy? Why?

Ans.

Each root directory entry is 32 bytes long and occupies 14 sectors, so each sector contains $512/32=16$ root directory entries. There are 14 root directory sectors, so there is a maximum of $14 \times 16=224$ entries.

22. How does the organization of an optical disk differ from the organization of a magnetic disk?

Ans.

Optical disks are fundamentally sequential media, whereas magnetic disks are direct (or random). Also, a magnetic disk directory is in a fixed location on the medium, on CDs, directory entries can "float."

23. Discuss the difference between how DLT and DAT record data. Why would you say that one is better than the other?

Ans.

DLT uses serpentine recording while DAT uses helical scan. DAT is more complicated to implement than DLT. DAT puts more wear on the tape and is generally less reliable than DLT. Many implementations of DLT record data faster than DAT.

24. How would the error-correction requirements of an optical document storage system differ from the error-correction requirements of the same information stored in textual form? What are the advantages offered by having different levels of error correction for optical storage devices?

Ans.

A document image will tolerate single-bit errors. Text storage will not. By eliminating error-correcting codes from the stored data, we can store more data on a disk and access it faster.

25. You have a need to archive a large amount of data. You are trying to decide whether to use tape or optical storage methods. What are the characteristics of this data and how it is used that will influence your decision?

Ans.

How near-line does it have to be? (What kind of access times do the users require?) How long and where will the archive volumes be stored? Is the data in some coded form like ASCII or EBCDIC, or is it binary like scanned images? Does backup time matter? Price?

- * 26. A particular high-performance computer system has been functioning as an e-business server on the Web. This system supports \$10,000 per hour in gross business volume. It has been estimated that the net profit per hour is \$1,200. In other words, if the system goes down, the company will lose \$1,200 every hour until repairs are made. Furthermore, any data on the damaged disk would be lost. Some of this data could be retrieved from the previous night's backups, but the rest would be gone forever. Conceivably, a poorly-timed disk crash could cost your company hundreds of thousands of dollars in immediate revenue loss, and untold thousands in permanent business loss. The fact that this system is not using any type of RAID is disturbing to you.

Although your chief concern is data integrity and system availability, others in your group are obsessed with system performance. They feel that more revenue would be lost in the long run if the system slows down after RAID is installed. They have stated specifically that a system with RAID performing at half the speed of the current system would result in gross revenue dollars per hour declining to \$5,000 per hour.

In total, 80% of the system e-business activity involves a database transaction. The database transactions consist of 60% reads and 40% writes. On average, disk access time is 20ms.

The disks on this system are nearly full and are nearing the end of their expected life, so new ones must be ordered soon. You feel that this is a good time to try to install RAID, even though you'll need to buy extra disks. The disks that are suitable for your system cost \$2000 for each 10 gigabyte spindle. The average access time of these new disks is 15ms with a MTTF of 20,000 hours and a MTTR of 4 hours. You have projected that you will need 60 gigabytes of storage to accommodate the existing data as well as the expected data growth over the next 5 years. (All of the disks will be replaced.)

- a. Are the people who are against adding RAID to the system correct in their assertion that 50% slower disks will result in revenues declining to \$5,000 per hour? Justify your answer.
- b. What would be the average disk access time on your system if you decide to use RAID-1?

- c. What would be the average disk access time on your system using a RAID-5 array with two sets of 4 disks if 25% of the database transactions must wait behind one transaction for the disk to become free?
- d. Which configuration has a better cost-justification, RAID-1 or RAID-5? Explain your answer.

Ans.

- a. No. Only 80% of the system activity involves the database. Proportionately then, doubling the disk access time would affect only 80% of the system activity. Access time would still be a lot slower according to their thinking:

(Percent of transaction on disk * half of the throughput) + (Percent of transaction in CPU * 10,000) = (0.8 * 0.5 * 10,000) + (0.2 * 10,000) = 4000 + 2000 = \$6,000 per hour.

So, using their assumptions, revenues would decline by only \$4,000 not \$5,000!

- b. In RAID-1, it takes twice as long to do a write as a read, because data has to be written twice. However, access time for a read is half of what we would expect from a system not using RAID-1, assuming that the disk arms are 180 degrees offset from one another.

Average Access Time = $0.4 * (15 \text{ ms} / 2) + 0.6 * (15 \text{ ms} * 2) = 21 \text{ ms.}$

- c. Average Access Time = $0.75 * 15\text{ms} + 0.25 * 30\text{ms} = 18.75 \text{ ms.}$

- d. Both RAID solutions will offer database response time comparable to what is currently offered by the system. The RAID-1 system will require 2*N disks while the 4-disk RAID-5 solution will require 133% of the number of disks. That is, RAID-1 will cost \$24,000 and RAID-5 will cost \$16,000. The cost of the disks isn't the big issue here, however. What matters most is system availability. With 8 disks each with a MTTF of 20,000 hours, we can expect a failure of at least two of the disks to fail within 20,000/8 hours, or 2,500 hours. So at least twice a year, we could expect a disk failure that will last 4 hours. If RAID-1 is used, the system will continue to function, while the RAID-5 system will be down, costing roughly \$4,800 in lost revenue during each outage. (No data would be lost, though!)

Cost of RAID-1: \$24,000; Cost of RAID-5: \$16,000 + \$9,600 revenue loss = \$25,600.

The RAID-1 is therefore more economical. Note: We have not included loss of goodwill and permanent business loss in the RAID-5 figure. This tilts the balance greatly in favor of the RAID-1 solution.]

- 27. a. Which of the RAID systems described in this chapter cannot tolerate a single disk failure?
- b. Which can tolerate more than one simultaneous disk failure?

Ans.

- a. RAID-0.
- b. RAID-1, RAID-2 and RAID-6. RAID-1 can tolerate multiple disk failures only if the failure does not involve a disk and its mirror image.

28. Compute the compression factors for each of the JPEG images in Figure 7.30.


Ans.

- a. $1 - (5.33/7.14) * 100\% = 25.35\%$.
- b. $1 - (1.96/7.14) * 100\% = 72.55\%$
- c. $1 - (1.49/7.14) * 100\% = 79.13\%$
- d. $1 - (1.11/7.14) * 100\% = 84.45\%$
- e. $1 - (0.639/7.14) * 100\% = 91.05\%$
- f. $1 - (0.556/7.14) * 100\% = 92.21\%$

29. Create a Huffman tree and assign Huffman codes for the “Star Bright” rhyme used in Section 7.8.2. Use <ws> for whitespace instead of underscores.

Ans.

One tree and its corresponding code:



Char-acter	Code	Char-acter	Code
<ws>	01	E	11111
T	100	Y	101000
H	001	N	101010
I	110	O	101011
S	1110	M	101101
G	0000	V	1010010
R	0001	L	1010011
W	10111	B	1011000
A	11110	F	1011001

30. Complete the LZ77 data compression illustrated in section 7.8.2.

Ans.

0, 0, S 0, 0, T 0, 0, A 0, 0, R 0, 0, _
 0, 0, L 0, 0, I 0, 0, G 0, 0, H 1, 1, A
 0, 5, B 3, 1, I 7, 4, F 6, 1, R 0, 2, _
 0, 5, I 7, 2, E 31, 1, _ 1, 1, O 0, 0, N
 0, 0, N 5, 1, F 5, 1, S 24, 1, _ 11, 2, M
 5, 1, Y 19, 11, I 4, 4, G 0, 0, E 0, 2, T

4, 1, E 9, 8, W 4, 4, T 0, 0, O 0, 0, N
4, 4, <EOF>

31. JPEG is a poor choice for compressing line drawings, such as the one shown in Figure 7.30. Why do you think this is the case? What other compression methods would you suggest? Give justification for your choice(s).

Ans.

The DC Coefficient gives a weighted average of the chrominance and luminance for the other 63 pixels in the block. When we “average” the pixels of a line drawing, we wind up with blocks of gray (or some other shade) instead of the sharp lines in the original graphic. Just about any other compression method would be better. LZ77 and run-length coding would be the fastest and would offer excellent compression because of the “local redundancy” (long runs of a single color) in a line drawing.

32. a. Name an advantage of Huffman coding over LZ77.
b. Name an advantage of LZ77 over Huffman coding.
c. Which is better?

Ans.

a. Huffman coding will usually result in better compression than LZ77 because it doesn't matter where in the message the characters are located.
b. LZ77 is faster and can be done in hardware.
c. It depends on what kinds of data you are dealing with and how you define “better.” If “better” means “faster,” LZ77 wins. If “better” means higher compression factor, Huffman will usually win.

33. State one feature of PNG that you could use to convince someone that PNG is a better algorithm than GIF.

Ans.

The advantages of PNG over GIF include:

- User-selectable compression modes: "Faster" or "better" on a scale of 0 to 3, respectively;
 - Improved compression ratios over GIF, typically 5% to 25% better;
 - Error detection provided by a 32-bit CRC (ISO 3309/ITU-142);
 - Faster initial presentation in progressive display mode; and, of course,
 - PNG is an open international standard, freely available and sanctioned by the World Wide Web Consortium (W3C) as well as many other organizations and businesses.]
-

Chapter 7A: Focus On Selected Disk Storage Implementations

1. Which of the types of storage architectures discussed in this section would you expect to find in a large data center or server farm? What would be the problem with using one of the other architectures in the data center environment?

Ans.

You would expect SAN or NAS in the data center or server farm. HPPI is overkill—probably too costly. Bus-based architectures would not provide sufficient throughput. There would be too much contention. USB or daisy-chained architectures are simply not fast enough.

2. How many SCSI devices can be active after the arbitration phase has completed?

Ans.

Just two. The Initiator and the Target.

3. Suppose during an asynchronous parallel SCSI data transfer someone removes a floppy disk from the drive that is the intended target of the transfer. How would the initiator know that the error has occurred during the phases:

- | | |
|--------------|----------------|
| a. Bus-free | e. Status |
| b. Selection | f. Message |
| c. Command | g. Reselection |
| d. Data | |

- a. During which of the phases is it possible that good data may be written to the floppy if the data transfer is a "write" operation?
- b. If the transfer is a "read," at which point would the system have good data in the buffer? Would the system ever acknowledge this data?

Ans.

- a. No knowledge.
- b. The target could not respond with the SEL signal because it is not ready.
- c. The target device would not raise ACK.
- d. The target device would not raise ACK.
- e. The target device would not raise ACK.
- f. The target would not respond with the "command complete" message.
- g. No knowledge.

During the Status and Message phases is it possible that good data may be written to the floppy.

During the Status, Message, and Reselection phases it is possible that the system have would have good data in the buffer.

4. Your manager has decided that the throughput of your file server can be improved by replacing your old SCSI-2 host adapter with a "fast and wide" SCSI-3 adapter. It has also been decided that the old SCSI-2 drives will be replaced with "fast and wide" SCSI-3 drives that are much larger than the old ones. After all of the files from the old SCSI-2 disks have been moved to the SCSI-3 drives, you reformat the old drives so that they can be used

again somewhere. Upon hearing that you did this, your manager tells you to leave the old SCSI-2 drives in the server, because she knows that SCSI-2 is downward compatible with SCSI-3. Being a good employee, you acquiesce to this demand.

A few days later, however, you are not surprised when your manager expresses disappointment that the SCSI-3 upgrade does not seem to be delivering the performance improvement that she expected. What happened? How can you fix it?

Ans.

A SCSI bus is as slow as the slowest device on it, so the SCSI-3 drives are working at SCSI-2 speeds. This can be fixed by reinstalling the old SCSI-2 host adapter and connecting it to the old SCSI-2 disks.

- ◆ 5. You have just upgraded your system to a "fast and wide" SCSI interface. This system has a floppy disk, a CD-ROM and five 8-gigabyte fixed disks. What is the device number of the host adapter? Why?

Ans.

The host adapter always has the highest device number so that it will always win arbitration for the bus. "Fast and wide" SCSI-3 interfaces can support as many as 32 devices; therefore, the host adapter will always be device number 32.

6. How does SCSI-2 differ from the principles behind the SCSI-3 Architecture Model?

Ans.

The SCSI-3 Architecture Model is a layered architecture, while SCSI-2 is not. SCSI-2 is part of the architecture model to maintain downward compatibility.

7. What benefits does the SCSI-3 Architecture Model provide to computer and peripheral equipment manufacturers?

Ans.

It is a layered protocol that provides backward compatibility. The interfaces at each layer are clearly defined.

8. Suppose you wish to devise a video conferencing system by connecting a number of computers and video cameras together. Which interface model would you choose? Will the protocol packet used to transfer the video be identical to the protocol packet used for data transmission? What protocol information would be in one packet and not the other?

Ans.

You would consider IEEE 1394. The protocol packet for the video will be different from the data packets because of the time-sensitivity of the transmissions. (See Fig. 7A.5.)

9. How would an SSA bus configuration recover from a single disk failure? Suppose another node fails before the first one can be fixed. How would the system recover?

Ans.

An SSA system would recover from a failure by using the ring in half-duplex mode, routing data around the failed device. (The loop is broken at the failed device.) If a second node fails, all devices situated between the two failed devices become inaccessible.

10. You have been assigned to a work group that has been given the task of placing automated controls in a chemical plant. Hundreds of sensors will be placed in tanks, vats and hoppers throughout the factory campus. All data from the sensors will be fed into a group of sufficiently high-powered computers so that plant managers and supervisors can control and monitor the various processes taking place.

What type of interface would you use between the sensors and the computers? If all computers are to have access to all of the sensor input, would you use the same type of connection to interconnect the computers among one another? Which I/O control model would you use?

Ans.

Think about IEE-1394. If data must be monitored in the split-second, isochronous protocols could be used. However, if there is little room for error, a fast-asynchronous method should be considered. Either IEEE 1394 or Fibre Channel will work, depending upon the speed and number of devices that would be daisy chained into one computer or hub. For process-control applications, programmed I/O is the only I/O control model that would be sufficiently responsive.

11. One of the engineers who works for you is proposing changes to the bus architecture of the systems that your company manufactures. She claims that if the bus is modified to support network protocols directly, the systems will have no need for network cards. She claims that you could also eliminate your SAN and connect the client computers directly to the disk array. Would you object to this approach? Explain.

Ans.

Unless the memory bus is separate from the I/O bus, this premise is flawed because memory transfers would be too slow to be useful. In fact, network cards provide functionality beyond simply negotiating protocols. (See Chapter 11.) By adding the computers directly to the SAN, you effectively have a gigantic shared bus that could collapse from contention.
