

Computer Architecture Homework #5

Due: 3/30/12 (F)

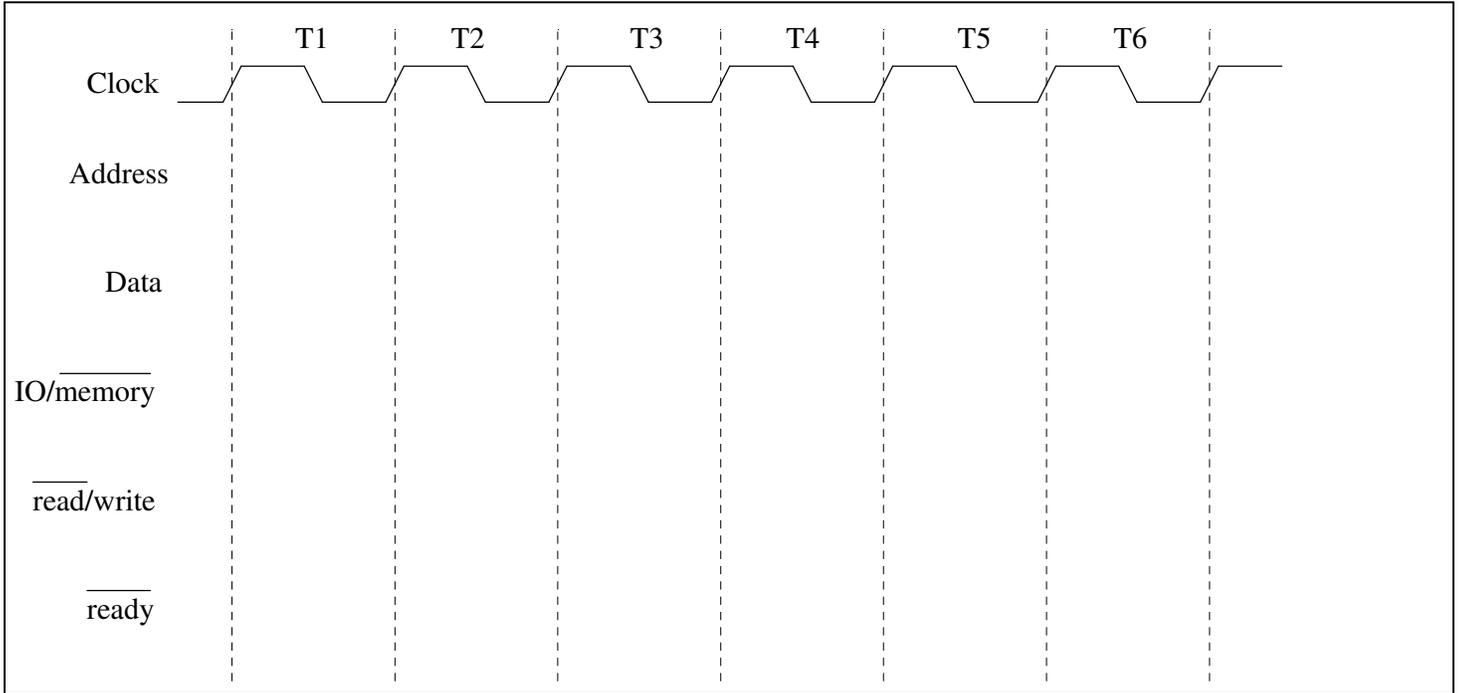
Chapter 7: Exercises: 22, 27, and the following problems: (You might want to refer to lecture 16 notes at:

http://www.cs.uni.edu/~fienup/cs2420s12/lectures/lec16_questions.pdf)

Extra Credit: Exercise: 40

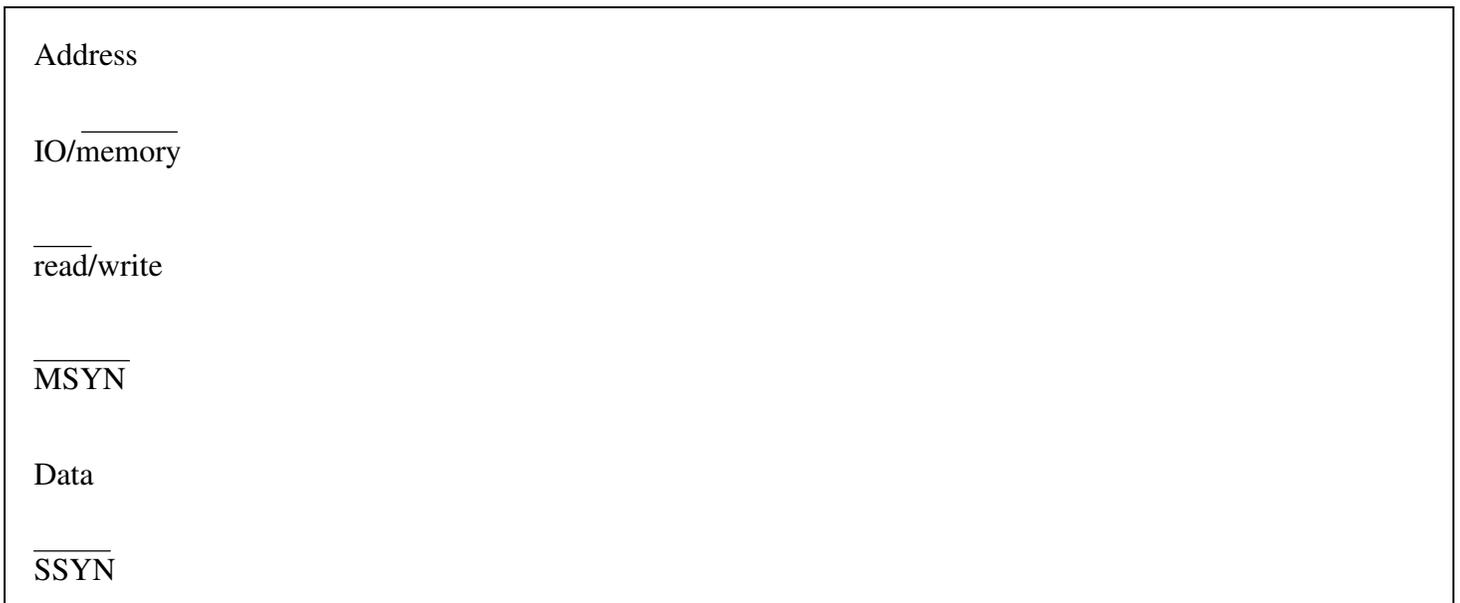
A) Draw a timing diagram for a synchronous, memory-write operation with two wait states (memory not ready to receive the data for two-clock cycles). Clearly label:

- which device (CPU or Memory) puts a value on each line, and
- when the data is read off the Data lines



B) Draw a timing diagram for an asynchronous, memory-write operation. Clearly label:

- which device (CPU or Memory) puts a value on each line,
- when the data is read off the Data lines, and
- where delays for bus skew are occurring.



C) Suppose we have an 8 disk RAID array with each disk having a 100 MB/sec data transfer rate. Complete the following table **assuming none of the disks are faulty**.

RAID Level	Maximum number of concurrent, independent READs	Maximum number of concurrent, independent WRITEs	Data Transfer Rate for a single large READ
RAID 0 (no redundancy with large strips)			
RAID 1 (Mirroring with large strips)			
RAID 3 (bit-interleaved parity)			
RAID 5 (block-level distributed parity)			
RAID DP (block-level double parity)			

D) Suppose we have an 8 disk RAID array with each disk having a 100 MB/sec data transfer rate. Complete the following table **assuming ONE of the disks is faulty**.

RAID Level	Maximum number of concurrent, independent READs	Maximum number of concurrent, independent WRITEs	Data Transfer Rate for a single large READ
RAID 0 (no redundancy with large strips)			
RAID 1 (Mirroring with large strips)			
RAID 3 (bit-interleaved parity)			
RAID 5 (block-level distributed parity)			
RAID DP (block-level double parity)			

CONFIGURATION:		RELIABILITY AND MAINTENANCE:	
Formatted Capacity, GB	1500	MTTF	300,000 hours
Integrated Controller	SATA	Start/Stop Cycles	50,000
Encoding Method	EPRML	Design Life	5 years (minimum)
Buffer Size	32MB	Data Errors	(nonrecoverable) <1 per 10 ¹⁵ bits read
Platters	8	PERFORMANCE:	
Date Surfaces	16	Seek times	
Tracks per Surface	16,383	Track to Track	
Track Density	190,000 tpi	Read	0.3 ms
Recording Density	1482 KbpI	Write	0.5 ms
Bytes per Sector	512	Average	
Sectors per Track	63	Read	4.5 ms
PHYSICAL:		Write	5.0 ms
Height	26.1mm	Average Latency	4.17 ms
Length	147.0mm	Rotational Speed	
Width	101.6mm	(+/-0.20%)	7,200 rpm
Weight	720g	Data Transfer Rate:	
Temperature (C°)		To/from Disk	1.2MB/Sec
Operating	5°C to 55°C	To/from Host	3GB/Sec
Non-operating/Storage	-40°C to 71°C	Start Time	
Relative Humidity	5% to 95%	(0 – Drive Ready)	9 sec
Acoustic Noise	33dBA, Idle		

POWER REQUIREMENTS		
Mode	+5VDC +5% - 10%	Power +5.0VDC
Spin-up	500mA	16.5W
Read/write	1080mA	14.4W
Idle	730mA	9.77W
Standby	270mA	1.7W
Sleep	250mA	1.6W

FIGURE 7.15 A Typical Rigid Disk Specification as Provided by a Disk Drive Manufacturer

22. The disk specification in Figure 7.15 gives a data transfer rate of 60MB per second when reading from the disk, and 320MB per second when writing to the disk. Why are these numbers different? (Hint: Think about buffering.)
26. Suppose a disk drive has the following characteristics:
- 6 surfaces
 - 16,383 tracks per surface
 - 63 sectors per track
 - 512 bytes/sector
 - Track-to-track seek time of 8.5 milliseconds
 - Rotational speed of 7,200 RPM.
- a) What is the capacity of the drive?
 - b) What is the access time?
27. Suppose a disk drive has the following characteristics:
- 6 surfaces
 - 953 tracks per surface
 - 256 sectors per track
 - 512 bytes/sector
 - Track-to-track seek time of 6.5 milliseconds
 - Rotational speed of 5,400 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 26? Explain.

EXTRA CREDIT. Exercise 40:

*40. 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 slowed 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 \$2,000 for each 10 gigabyte spindle. The average access time of these new disks is 15ms with an MTTF of 20,000 hours and an 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 four 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.