

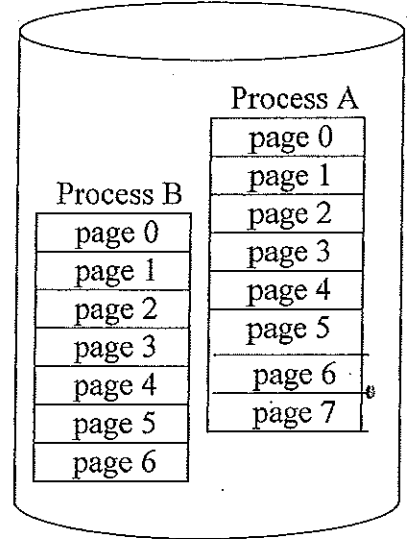
Question 1. (10 points) Consider a demand paging system with 1024-byte pages.

Process B Page Table

CPU	Frame #	Valid bit (Loaded In Memory)
0	2	1
1	6	1
2	-	0
3	4	1
4	0	1
5	-	0
6	-	0
7	-	0

Physical Memory

Frame Number	Memory
0	page 4 of B
1	page 2 of A
2	page 0 of B
3	page 1 of A
4	page 3 of B
5	page 4 of A
6	page 1 of B



5 a) Complete the above page table for Process B.

5 b) If process B is currently running and the CPU generates a logical/virtual address of 2032₁₀, then what would be the corresponding physical address?

physical addr. = 7152₁₀

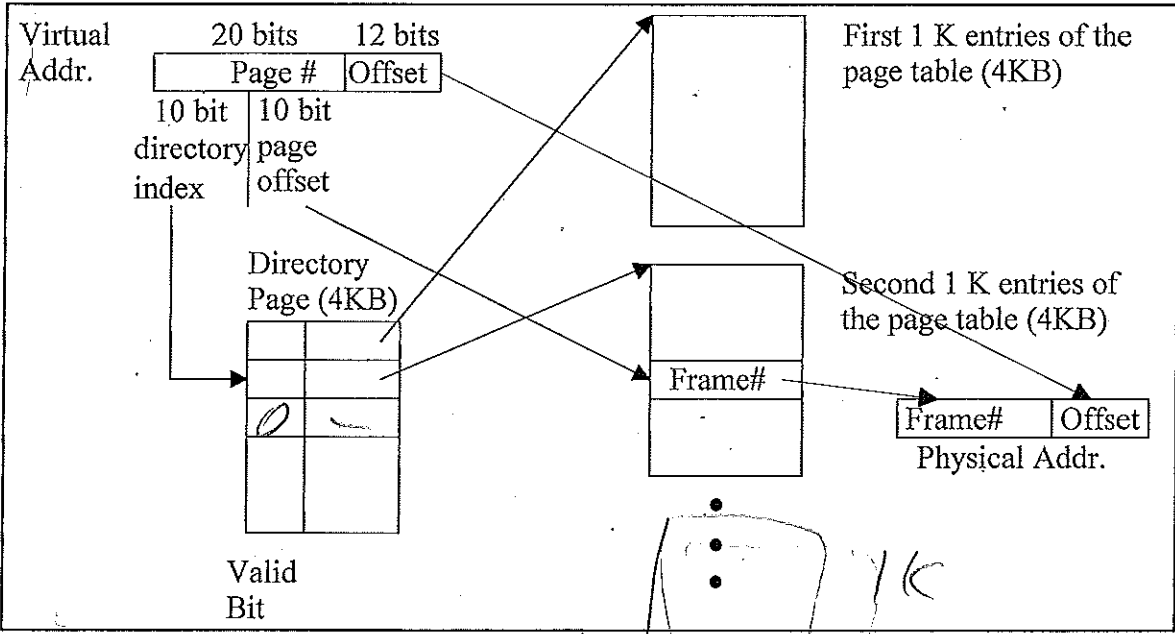
Handwritten calculation for physical address:

frame	offset
6	1008

1024 offset

$$\begin{array}{r} 1024 \\ \times 6 \\ \hline 6144 \\ + 1008 \\ \hline 7152 \end{array}$$

10 Question 2. (10 points) One way to handle large page tables is to use two-level (or more) page tables where the first level (the "directory") acts as an index into the page table which is scattered across several pages. Consider a 32-bit virtual addresses with 4KB pages and 4 byte page table entries.



a) If large sections of the virtual address space are unused, how can sections of the page table be eliminated?

5 In the directory page, the valid bit = 0 means that the corresponding section of page table (1024 entries) do not exist.

b) In a computer using two-level page tables, what would be stored in the TLB?

5/10 PT entries like before from the second level of the PT.

Question 3. (6 points) To approximate the LRU page-replacement algorithm, a hardware maintained reference (R) bit and history bits can be stored for each entry in the page tables. Periodically, say every 10 milliseconds, an interrupt causes the OS to shift the R-bit into the counter/history bits. Consider the following snapshot of R-bits, counter/history bits, and dirty bit (page in memory modified from disk copy).

	R	Counter/History bits	Dirty Bit (1 = modified)
page 0	1	0 0 1 1 0 1 0	0
page 1	0	1 0 1 0 0 0 0	1
page 2	0	0 1 0 1 0 1 0	1
page 3	0	0 1 0 1 0 1 0	0
page 4	1	0 0 0 0 0 1 0	1

- a) If a page fault occurs, which page should be selected for replacement? *Either page 2 or 3, but page 3 is not dirty so it would be a better choice.*
- b) How long has it been since page 1 has been referenced? (give a range) *0-20 ms.*

Question 4. (9 points) Suppose we had a block transfer from an I/O device to memory. The block consists of 1024 words and one word can be transferred to/from memory at a time. For each of the following, indicate the number of interrupts needed to transfer a block using:

- a) DMA (direct-memory access) *1*
- b) interrupt-driven I/O *1024*
- c) programmed-I/O *0*

Question 5. (5 points) What is the advantage of interrupt-driven I/O over programmed-I/O?

Interrupt-driven turns the CPU over to another process(es) while the slow I/O is performed. => increased CPU utilization.

Question 6. (5 points) What is the advantage of DMA (direct-memory access) over interrupt-driven I/O?

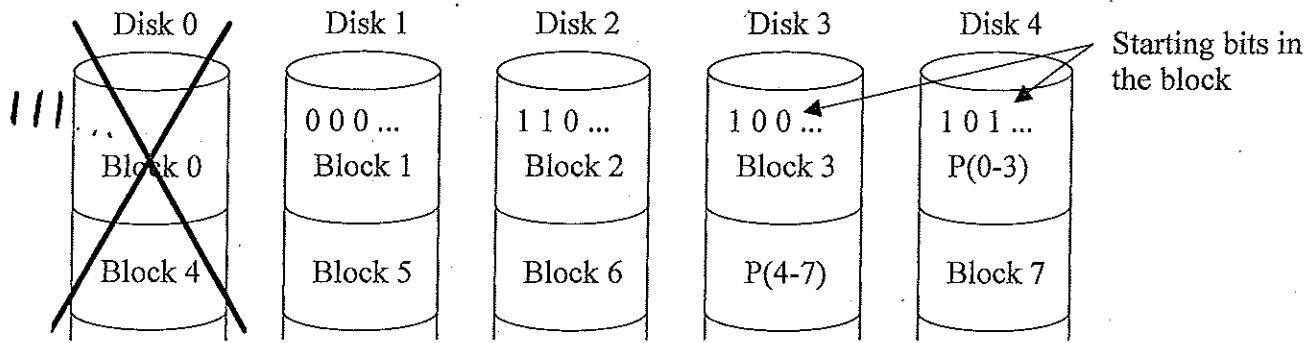
Both use interrupts, but the DMA only interrupts after whole block transferred while interrupt-driven interrupts on each word.

Question 7. (10 points) On a paged, multiprogrammed, multi-user computer system that uses memory-mapping I/O, describe what hardware support in need by the operating system to:

- a) guard against infinite loops in user programs
- CPU timer
 - dual-mode protection with CPU timer setting privileged instr.
- b) restrict a user program from accessing the disk directly thus gaining access to other users' data files
- Since it is memory-mapped I/O, the memory addresses must not be generated that corresponds to the I/O controller registers.*

Question 8. (10 points) Suppose we have a 5 disk RAID 5 (block-level distributed parity) array.

a) If Disk 0 crashes, reconstruct the first three bits of block 0. Assume even parity is being used.



b) Assume that Disk 0 has crashed. Can Block 0, Blocks 5 and Block 6 all be read at the same time? (Justify your answer)

No, Reading Block 0 from the failed disk requires reading Blocks 1, 2, 3, and P(0-3) block, so disks 1 and 2 will not be available to read Blocks 5 and 6.

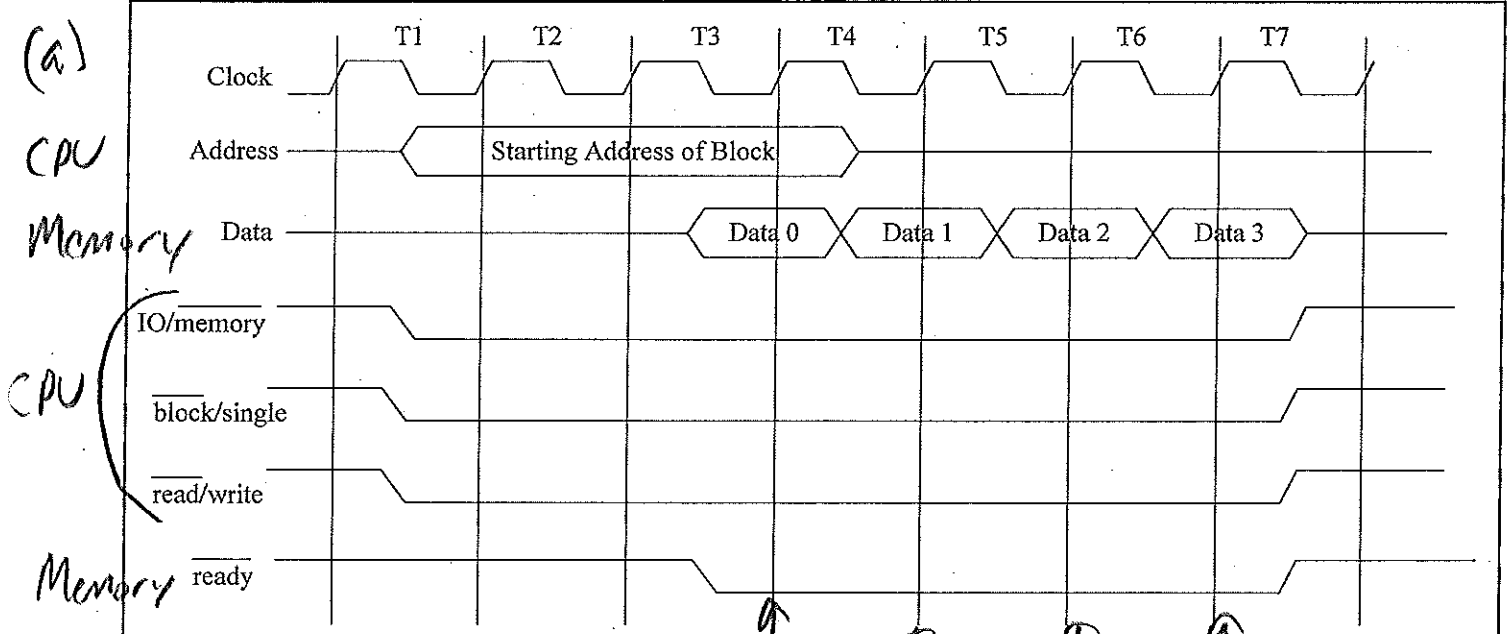
Question 9. (12 points) Suppose we have an 6 disk RAID array with each disk having a 100 MB/sec data transfer rate. Complete the following table assuming NO disks are faulty.

RAID Level	Maximum number of concurrent, independent READs	Maximum number of concurrent, independent WRITES	Data Transfer Rate for a single large READ
RAID 1 (Mirroring with large strips)	6	3	100 MB/sec
RAID 3 (bit-interleaved parity)	1	1	500 MB/sec
RAID 5 (block-level distributed parity)	6	3	100 MB/sec
RAID DP (Double parity blocks)	4	1	100 MB/sec

Question 10. (8 points) Circle all correct answers to explain why RAID level 5 (blocked with distributed parity) is good for a database server.

- a) a database read operation typically involves one disk in the RAID array
- b) a database read operation typically involves two disks in the RAID array
- c) a database read operation typically involves all of the disks in the RAID array
- d) a database write operation typically involves one disk in the RAID array
- e) a database write operation typically involves two disks in the RAID array
- f) a database write operation typically involves all of the disks in the RAID array
- g) many database I/O operations can be performed in parallel on the RAID array
- h) a single large read operation spanning several RAID disks can boost the data transfer rate of the read
- i) if one disk in the RAID array fails, then the database server can continue to operate
- j) if two disks in the RAID array fail, then the database server can continue to operate

Question 11. (15 points) On a cache miss, suppose a 4-word block of memory needs to be read to fill the cache line. Suppose our synchronous bus includes Block Memory operations via a *block/single* control wire. The below synchronous bus timing diagram shows a Block Memory Read with one wait state.



a) Indicate which wires are controlled by the CPU and which are controlled by the memory.

b) When (in clock cycles) are the 4 words of data read off of the Data wires by the CPU?

Start of: T4, T5, T6, T7

c) Why do devices change their values being sent in the middle of the clock cycles, and read the values being received at the start of the next clock cycle?

This handles bus skew

d) How does the CPU know when the first Data value is on the Data-wires?

The memory drops the ready signal to 0.

e) What advantage would a Block Memory Read operation to fill a cache-line have over 4 separate single word non-block reads?

It is faster since we only need to request + acquire the bus once and send the address once.

15