

1. (9pts) Three Requirements must be satisfied while designing a solution to the "critical-section" problem. Please briefly describe the three requirements.
2. Consider the following set of processes, assumed to have arrived at time 0, in order P1, P2, ..., P5, with the length of the CPU-burst time given in milliseconds.

Process	Burst Time	Priority
P1	9	3
P2	1	1
P3	2	3
P4	1	4
P5	5	2

Assume that we use "priority scheduling" to schedule these processes. Please answer the following two questions:

- (a) (3pts) Plot the Gantt chart according to which we would schedule these processes.
 - (b) (3pts) Compute the average waiting time for these processes.
3. (10pts) As a process executes, it changes *state*. The state of a process is defined in part by the current activity of that process. There are five states in one of which each process may be. Please depict the transition diagram of process state. (Hint: you have to write down the five states first, and then specify the transition between each two states.)
 4. (5pts) List in order of speed of access (fastest = 1; for instance, if magnetic tape is the 2nd, disk is 3rd, register is 4th, main memory is 5th, and CPU cache is 1st, your answer will be "2,3,4,5,1"): magnetic tape, disk, register, main memory, CPU cache
 5. Consider a system with an average memory access time of 50 nanoseconds, a three level page table (meta-directory, directory, and page table). For full credit, your answer must be a single number and not a formula.
 - (a) (10pts) If the system had an average page fault rate of 0.01% for any page accessed (data or page table related), and an average page fault took 1 millisecond to service, what is the effective memory access time (assume no TLB or memory cache)?
 - (b) (10pts) Now assume the system has no page faults, we are considering adding a TLB that will take 1 nanosecond to lookup an address translation. What hit rate in the TLB is required to reduce the effective access time to

memory by a factor of 2.5?

6. In this problem set, show your answers in the following format:

<a> ? CPU cycles

Derive your answer.

 $CPI = ?$

Derive your answer.

<c> Machine ? is ?% faster than ?

Derive your answer.

<d> ? CPU cycles

Derive your answer.

Both machine A and B contain one-level on-chip caches. The CPU clock rates and cache configurations for these two machines are shown in Table 1. The respective instruction/data cache miss rates in executing program P are also shown in Table 1. The frequency of load/store instructions in program P is 20%. On a cache miss, the CPU stalls until the whole cache block is fetched from the main memory. The memory and bus system have the following characteristics:

1. the bus and memory support 16-byte block transfer;
2. a 32-bit synchronous bus clocked at 200 MHz, with each 32-bit transfer taking 1 bus clock cycle, and 1 bus clock cycle required to send an address to memory (assuming shared address and data lines);
3. assuming there is no cycle needed between each bus operation;
4. a memory access time for the first 4 words (16 bytes) is 250 ns, each additional set of four words can be read in 25 ns. Assume that a bus transfer of the most recently read data and a read of the next four words can be overlapped.

	Machine A	Machine B
CPU clock rate	800 MHz	400 MHz
I-cache configuration	Direct-mapped, 32-byte block, 8K	2-way, 32-byte block, 128K
D-cache configuration	2-way, 32-byte block, 16K	4-way, 32-byte block, 256K
I-cache miss rate	6%	1%
D-cache miss rate	15%	4%

Table 1

To answer the following questions, you don't need to consider the time required for

writing data to the main memory:

- (a) (10pts) What is the data cache miss penalty (in CPU cycles) for machine A?
 - (b) (5pts) What is the average CPI (Cycle per Instruction) for machine A in executing program P? The CPI (Cycle per Instruction) is 1 without cache misses.
 - (c) (5pts) Which machine is faster in executing program P and by how much? The CPI (Cycle per Instruction) is 1 without cache misses for both machine A and B.
 - (d) (5pts) What is the data cache miss penalty (in CPU cycles) for machine A if the bus and memory system support 32-byte block transfer? All the other memory/bus parameters remain the same as defined above.
7. (4pts) Given the bit pattern 10010011, What does it represent assuming
- (a) It's a two's complement integer?
 - (b) It's an unsigned integer?
- Write down your answer in decimal format.
8. (6pts) Draw the schematic of a 4-bit 2's complement adder / subtractor that produces $A+B$ if $K=1$, $A-B$ if $K=0$. In your design try to use minimum number of the following basic logic gates (1-bit adders, AND, OR, INV, and XOR).
9. (6pts) We want to add four 4-bit numbers, $A[3:0]$, $B[3:0]$, $C[3:0]$, $D[3:0]$ together using carry-save addition. Draw the schematic using 1-bit full adders.
10. (9pts) We have an 8-bit carry ripple adder that is too slow. We want to speed it up by adding one pipeline stages. Draw the schematic of the resulting pipeline adder. How many 1-bit pipeline register do you need? Assuming the delay of 1-bit adder is 1 ns, what's the maximum clock frequency the resulting pipelined adder can operate?