

- **One sheet of notes allowed. Calculators are allowed.**
- Show your work or explain your reasoning where required to receive full credit.
- Total points: 100

1. (20 pts) Average Memory Access Time and Caches

(a) (4 pts) Give the Average Memory Access Time equation (AMAT). Label each component of the equation with words.

(b) (2 pts each) For each of the following cache optimizations, indicate which component of the AMAT equation is improved.

_____ (1) Using a second-level cache

_____ (2) Using a direct-mapped cache

_____ (3) Using a 4-way set-associative cache

_____ (4) Using a virtually-addressed cache

_____ (5) Using sub-blocks

_____ (6) Performing hardware prefetching using stream buffers

_____ (7) Using a non-blocking cache

_____ (8) Using larger blocks

2. (4 pts) To implement the cache optimization called Critical Word First, the memory system must be organized as....

3. Consider a branch target buffer for dynamic branch prediction.

(a) (2 pts) In what pipeline stage is the target buffer checked?

(b) (4 pts) Show the state machine for a two-bit branch prediction scheme.

4. (20 pts. total) Consider an 16 kilobyte 2-way set-associative cache. Words are 32 bits long. There are 512 blocks in the cache.

(a) (2 pts each) Determine the following values:

(1) Block size (in bytes):

(2) Set size (in blocks):

(3) Number of sets:

(4) Size of index field in address:

(5) Size of offset field in address (assume the memory system is byte-addressable):

(6) Size of tag field in address:

(b) (8 pts.) Draw a picture of the cache. Be sure to show how the different parts of the address are used to pick a set, check for a match, and extract the desired word from the chosen block.

5. (20 pts. total) Consider a program running on a computer system with virtual memory. Besides the CPU, you must take into account the cache, the translation lookaside buffer (TLB), and main memory.

(a) (10 pts.) First assume that the cache is physically addressed. Assume the following: a TLB hit takes zero cycles, a TLB miss takes 10 cycles, a cache hit takes one cycle (in addition to any address translation that may be required), and a cache miss takes 10 cycles. The TLB miss rate is 2%, and the cache miss rate is 5%. (Assume no page faults occur.)

What is the average memory access time for the cache?

(b) (10 pts.) Now assume that the cache is a virtually-addressed cache. Use the same values as in part (a). **What is the average memory access time for the cache?**

6. (30 pts) Tomasulo’s Algorithm and Speculative Execution

For the following example, fill in the graph for Tomasulo’s algorithm with speculative execution. Assume that reservation stations are freed up when the write-result operation is complete.

Indicate on which cycle each instruction commits; if an instruction does not commit, but is instead flushed from the reorder buffers and reservation stations, then indicate the cycle on which the flush occurs.

Assume a branch target buffer predicts the branch is NOT taken.

The branch is fully resolved after the execution stage (no write result required).

Assume the following execution times:

LW: 2 cycles ADD/SUB: 2 cycles BNEZ: 3 cycles MULT/DIV: 4 cycles

TABLE 1.

Instruction	Issue	Execution Complete	Write Result	Commit on Cycle...	Flush on Cycle...
LW R2, #1					
SUB R4, R1, R2					
BNEZ R2, DEST					
MULT R5, R2, R3					
DIV R6, R5, R4					
DEST: MULT R5, R2, R4					
ADD R6, R5, R2					

TABLE 2.

Unit	Busy	Op	Vj	Vk	Qj	Qk
MEM1						
MEM2						
ADD1						
ADD2						
BRANCH						
MULT/DIV						

Scratch Paper