

## CS4760 Homework 2: Caches, Disk Arrays, Multiprocessors

**Due Tuesday, November , 1998 in class (12:15pm) (No late homework accepted)**

---

For this assignment and for all other homework assignments, you may collaborate with other students, but the work you turn in must be your own. If you work with other students, list their names along with your own.

Show all work to receive full credit on answers.

### 1. Memory Hierarchies

Consider the following cache:

total size 8 KBytes, word size 64bits, block size 128 bits

- (a) How many total blocks are in the cache?
- (b) How many bits are required to specify the offset portion of the address?
- (c) Fill in the following table for the different cache types:

**TABLE 1.**

	<b>Direct-Mapped Cache</b>	<b>4-Way Set-Associative</b>	<b>Fully Associative</b>
<b>Number of blocks per set</b>			
<b>Number of sets</b>			
<b>Number of bits to specify the index</b>			
<b>Number of bits in tag</b>			

- (d) In a few words, name the most important optimization to improve the performance of write-through caches.

### 2. 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) 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)
  - A cache miss takes 10 cycles.
  - The TLB miss rate is 2%
  - The cache miss rate is 5%.
  - (Assume no page faults occur.)

**What is the average memory access time for the cache?**

(b) 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?**

3. Consider an 8-way set-associative cache with the following properties:

32-bit addresses

32-bit data words

block size: 32 bytes

total cache size: 32 kilobytes

Write back cache (does write allocate)

50% of block in cache are dirty

Page table lookup on TLB miss: 20 cycles

TLB hits take 0 cycles (can be combined with cache hit time).

TLB miss rate is 0.2%

Cache hits take 1 cycle

Cache miss rate is 1%

Assume there is **no write buffer**

Physically-addressed cache

Assume a simple memory model: Memory access time is 40 cycles, transfer rate is 4 bytes per cycle (In other words, don't worry about access time vs. cycle time, etc.)

**(a) What is the average time (in cycles) to do a read operation?**

(b) Now consider a standard 5-stage L/S pipeline (IF ID EX MEM WB). If the CPI (cycles per instruction) for this pipeline is 1.5 assuming a perfect memory system (no TLB misses, no cache misses), then **what is the additional contribution to the CPI** for the non-ideal memory system we have described, assuming that 20% of instructions are loads or stores?

4. Disk Arrays

(a) What are the advantages and disadvantages of **fine-grained interleaving** in disk arrays?

(b) What is the **“small write problem”** in disk arrays and **in what specific types of disk arrays does it occur?**

(c) **Describe** two redundancy schemes used to provide higher reliability in disk arrays.

5. Multiprocessor Cache Coherency Protocols

(a) Consider a multiprocessor with centrally shared memory. If the multiprocessor uses a write-invalidate snoop cache protocol, explain how the protocol works when a data item X is initially cached by several processors and processor P decides to write the data.

(b) How does the protocol work when another processor Q wants to read data item X after P writes it?