

GEORGIA INSTITUTE OF TECHNOLOGY

College of Computing

CS6290/CS4290 — High-Performance Computer Architecture Fall 2001

CS6290/CS4290
Homework 2

Issued: August 31, 2001
Due: September 14, 2001

Purpose: This homework covers caching.

Reading: H&P Chapter 5 for caches
[Jouppi90] to be discussed in class Friday, September 7th.
[Dennis65] to be discussed in class Friday, September 14th.

Problems:

1. Cache Organization.
2. Cache Performance.
3. Cache Analysis.
3. Cache Mechanics.
5. Read [Jouppi90] and [Dennis65].

Collaboration: (*As in the syllabus*) collaboration on projects and homework in **pairs** is encouraged. If you work in a pair, turn in one write-up with the names of both collaborators. You're welcome to discuss high-level concepts with other groups, but all homework solutions must be worked out and written up separately.

Problem 1: Cache Organization

A cache has a total size of N data words, with B words per block and A -way set associativity. Assume 32-bit data words, 32-bit addresses and that N , B and A are all powers of two.

A: What is the layout of an address in terms of the number and positions of bits used for the tag, index and offset?

B: What is the total number of bits of storage required (data + tags) in the cache hardware?

C: Why are the index bits taken from the “middle” of the address and not from the least significant or most significant bits of the address?

Problem 2: Cache Performance

A: Problem 5.4 in the book.

Problem 3: Cache Analysis

A: Problem 5.3 in the book.

Problem 4: Cache Hacking

Adapted from a problem in Ward & Halstead, Computation Structures, McGraw-Hill, 1990.

You are hired by companies *A*, *B* and *C* to recommend cache designs for their respective machines. Each company's product is a clone of the DLX from the text and executes an identical instruction set. For the sake of variety, you have recommended a different cache to each manufacturer as follows:

A: A one-way (direct-mapped) cache, 4K bytes total size, 16 data bytes per line.

B: A two-way set associative cache, 4K bytes total size, 16 data bytes per line.

C: A four-way set associative cache, 4K bytes total size, 16 data bytes per line.

Note that (for the sake of consistency) you have recommended a cache with a total size of 4K bytes and 16 bytes per cache line. The write policy is write-through and the set-associative caches use an LRU replacement policy among the lines in a set.

Each of the companies is outraged to learn that you have been consulting for the others and you are threatened with multiple lawsuits. To placate your clients, you need to assure each company that its machine now outperforms the others in certain applications. Your challenge is to find three appropriate programs.

Describe three small program fragments. Each should be a loop that would run forever and each should perform best on one of the three machines. For each of your three program fragments, give the following information (example format below):

1. The memory reference string; i.e., the sequence of addresses referenced during the execution of the loop. The sequence is, of course, an infinite cycle; use ellipses as appropriate.
2. For each cache/machine, an indication of whether each reference hits or misses in the cache.
3. The hit ratio for each cache/machine.

Here is an example format (the patterns are nonsensical):

```
ref string:  1  37  42 1025  1  37  42 1025  ...
cache A:    -  H   H   -   -  H   H   -   ... (HR = 0.5)
cache B:    -  H   -   -   -  H   -   -   ... (HR = 0.25)
cache C:    -  -   -   -   -  -   -   -   ... (HR = 0)
```

A: Write three answers in the format above, one that performs best for each cache/machine.

Problem 5: Reading

A: Read Jouppi's paper, [Jouppi90] on victim caches and stream buffers in order to be prepared to discuss it in class on Friday, September 7th. What performance metric does Jouppi use to judge his proposed hardware? What are the drawbacks of this metric?

B: Read [Dennis65] on virtual memory in order to be prepared to discuss it in class on Friday, September 14th. This is quite an old paper and virtual memory systems have generally become *simpler* since 1965. In particular, segmentation is rarely used. What is the key benefit of segmentation? What do current operating systems do to get this benefit without segmentation hardware?