CS6760 Final, December 14, 1996 Name:

Total points:
Read all parts of the questions. Questions have different point values.
One sheet of notes allowed, calculators ok
You must show your work to receive partial credit

1. Coinsider a cache with 32 blocks (0, 1, ..., 31) that holds blocks from a memory of total size 128 blocks (0, 1, ..., 127). Determine **in which block(s) in each cache the memory block 50 may reside.** Make your reasoning clear by **drawing a picture of the cache** and identifying blocks and sets.

   (a) (5 pts) A direct-mapped cache.

   (b) (5 pts) A fully associative cache.

   (c) (5 pts) A 2-way set-associative cache. (Use the notation we described in class, where an n-way set-associative cache is divided into sets of size n, and assume the blocks within each set are numbered sequentially in the cache.)

4. Disk Arrays
   (a) (5 pts) What are the advantages and disadvantages of **fine-grained interleaving** in disk arrays?

   (b) (5 pts) What is the **"small write problem"** in disk arrays and i**n what specific types of disk arrays does it occur?**

   (c) **(5 pts) Describe** two redundancy schemes used to provide higher reliability in disk arrays.

CS6760 Final, December 14, 1996 Name:

3. Virtual Memory

(a) **(10 pts) Draw the components and data paths** of a virtual memory system that includes a CPU, a physically-addressed cache, a TLB, main memory and disk.

(b) Assuming the folowing cache configurations, **describe step-by-step how the following operations are handled.** (Use high-level descriptions, such as "data read from disk" or "cache hit detected".)

(1) (10 pts) A **read miss** in a **write-back, physically addressed cache** where the block to be replaced in the cache is **dirty**, there is a **TLB hit** and a **page fault on the read data**. (No page fault on the write data.)

Step 1: CPU generates a virtual address for the read access.
Step 2:
Step 3:
Step 4:
Step 5:
Step 6:
Step 7:
Step 8:
Step 9:
Step 10:

(2) (10 pts) A **write miss** in a **write-through, virtually-addressed cache** that does **write-allocate** on a miss. Also assume a **TLB mis**s and that data resides in main memory (**no page fault)**.

Step 1: CPU generates a virtual address for the write access.
Step 2:
Step 3:
Step 4:
Step 5:
Step 6:
Step 7:
Step 8:
Step 9:
Step 10:

5. 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

(a) **(3 pts) How many sets** are in the cache (again using the notation that an n-way set associative cache contains sets of size n)?

(b) (10 pts) Label the parts of the address that are used to decide whether there is a cache hit or miss. **For each portion of the address, calculate the number of bits**. Briefly describe **how the different parts** of the address are used to decide the hit or miss.

(c) (20 pts) Assume the following additional information about the cache described above:
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.)

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

(d) (10 pts) 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?

CS6760 Final, December 14, 1996 Name:

6. Briefly Answer the following questions
   (a) (5 pts) How does a victim cache work?

   (b) (5 pts) How are WAR hazards handled in a scoreboard? in a Tomasulo's scheme?

   (c) (5 pts) At what stage in a pipeline can we use a branch target buffer? a branch prediction buffer? Explain.

   (d) (5 pts) Describe data forwarding in pipelines. How does it eliminate pipeline stall cycles?

   (e) (5 pts) What are some instruction types and/or addressing modes that are available in a R/M machine that are not available in an L/S machine?

   (f) (5 pts) What is the average memory access time equation for a single-level cache? When we add a second level to the cache, which componenets of the equation are affected?

(g) (5 pts) Consider a multiprocessor with centrally shared memory. If the multiprocessor uses a write-invalidate snoopy 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.

(h) (5 pts) How does the protocol work when another processor Q wants to read the data after P writes it?

(i) (5 pts) What are some of the limits to speedup in a vector processor?

(j) (5 pts) Write the access time equation for reading 4 sequential words from a word-wide nibble-mode DRAM.

Scratch Paper: