

EECS 361 Homework 4
Fall 2006
Due: 11/23/06

1. [10] Here is a series of address references given as word addresses: 1, 4, 8, 5, 20, 17, 19, 56, 9, 11, 4, 43, 5, 6, 9, 17. Assuming a direct mapped cache with 16 one-word blocks that is initially empty, label each reference in the list as a hit or a miss and show the final contents of the cache.
2. [10] Using the series of references given in Problem 1, show the hits and misses and final cache contents for a direct-mapped cache with four-word blocks and a *total size* of 16 words.

Average Memory Access Time

To capture the fact that the time to access data for both hits and misses affects performance, designers often use average memory access time (AMAT) as a way to examine alternative cache designs. Average memory access time is the average time to access memory considering both hits and misses and the frequency of different accesses; it is equal to the following:

$$\text{AMAT} = \text{Time for a hit} + \text{Miss rate} \times \text{Miss penalty}$$

AMAT is useful as a figure of merit for different cache systems.

3. [5] Find the AMAT for a machine with a 2-ns clock, a miss penalty of 20 clock cycles, a miss rate of 0.05 misses per instruction, and a cache access time (including hit detection) of 1 clock cycle. Assume that the read and write miss penalties are the same and ignore other write stalls.
4. [5] Suppose we can improve the miss rate to 0.03 misses per reference by doubling the cache size. This causes the cache access time to increase to 1.2 clock cycles. Using the AMAT as a metric, determine if this is a good trade-off.
5. [10] If the cache access time determines the processor's clock cycle time, which is often the case, AMAT may not correctly indicate whether one cache organization is better than another. If the machine's clock cycle time must be changed to match that of a cache, is this a good trade-off? Assume the machines are identical except for the clock rate and the number of cache miss cycles; assume 1.5 references per instruction and a CPI without cache misses of 2. The miss penalty is 20 cycles for both machines.

6. [10] You have been given 18 32K x 8-bit SRAMs to build an instruction cache for a processor with a 32-bit address. What is the largest size (i.e., the largest size of the data storage area in bytes) direct-mapped instruction cache that you can build with one-word (32-bit) blocks? Show the breakdown of the address into its cache access components (for an example, see Figure 1) and describe how the various SRAM chips will be used. (Hint: You may not need all of them.)




Figure 1 The caches in the DECStation 3100 each contain 16K blocks with one word per block. This means that the index is 14 bits and that the tag contains 16 bits.

7. [10] This exercise is similar to Problem 6, except that this time you decide to build a direct-mapped cache with four-word blocks as in Figure 2. Once again show the breakdown of the address and describe how the chips are used.

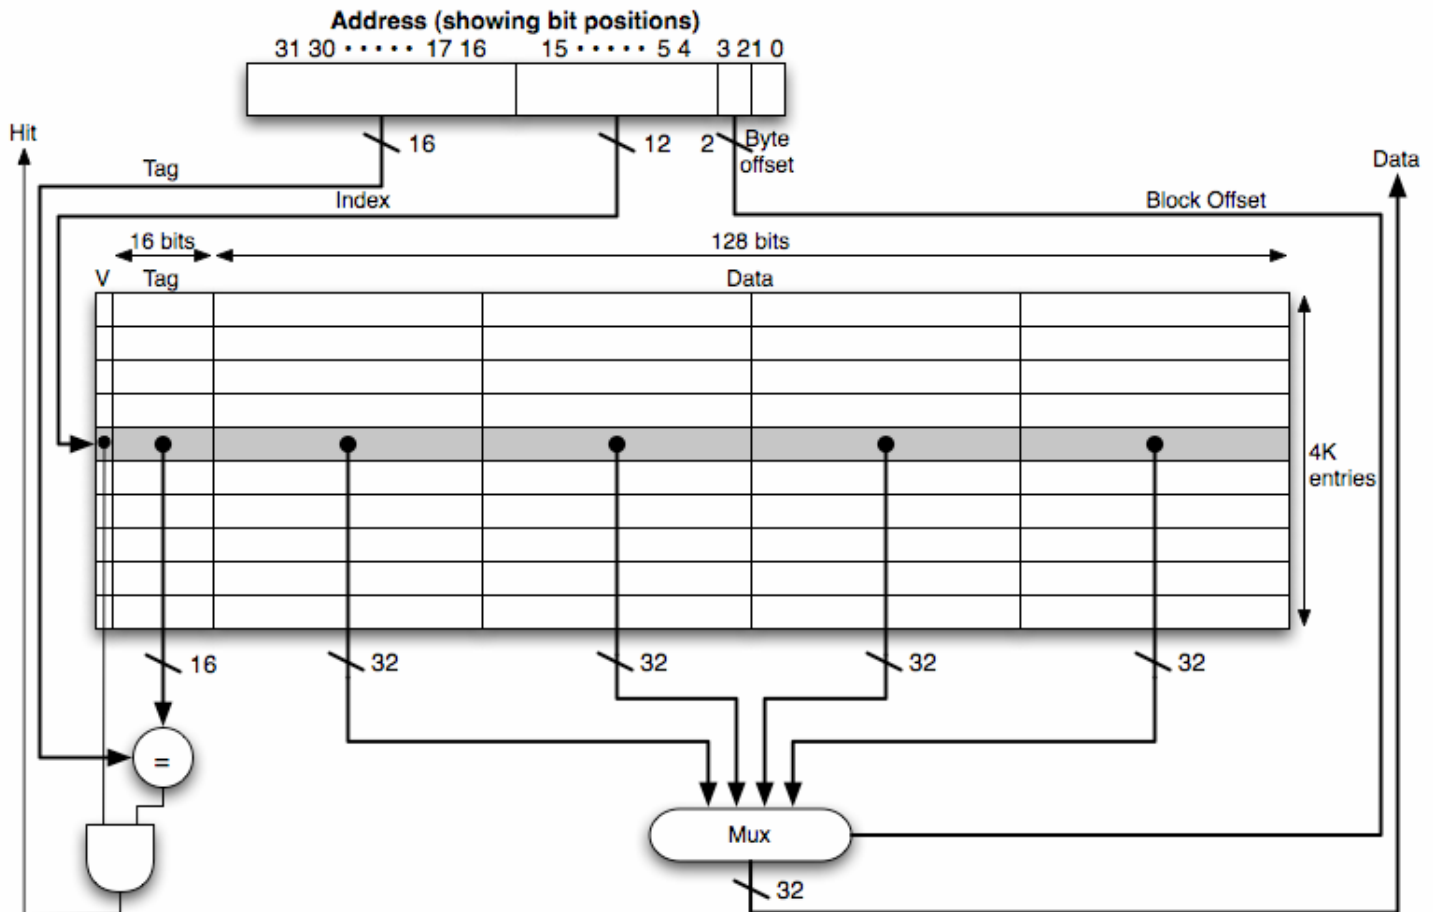


Figure 2 A 64-KB cache using four-word (16-byte) blocks. The tag field is 16 bits wide and the index field is 12 bits wide, while a 2-bit field (bits 3-2) is used to index the block and select the word from the block using a 4-to-1 multiplexor. In practice, the low-order bits of the address (bits 2 and 3 in this case) are used to enable only those RAMs that contain the desired word, eliminating the need for the multiplexor. Another way to eliminate the multiplexor is to have a large RAM for the data (with the tags stored separately) and use the block offset to supply 2 address bits for the RAM. The RAM must be 32 bits wide and have four times as many words as blocks in the cache.

8. [20] Consider three machines with different cache configurations:

- *Cache 1:* Direct-mapped with one-word blocks
- *Cache 2:* Direct-mapped with four-word blocks
- *Cache 3:* Two-way set associative with four-word blocks

The following miss rate measurements have been made:

- *Cache 1:* Instruction miss rate is 4%; data miss rate is 8%.
- *Cache 2:* Instruction miss rate is 2%; data miss rate is 5%.
- *Cache 3:* Instruction miss rate is 2%; data miss rate is 4%.

For these machines, one-half of the instructions contain a data reference. Assume that the cache miss penalty is $6 + \text{Block size}$ in words. The CPI for this workload was measured on a machine with cache 1 and was found to be 2.0. Determine which machine spends the most cycles on cache misses.

9. [5] The cycle times for the machines in Problem 8 are 2 ns for the first and second machines and 2.4 ns for the third machine. Determine which machine is the fastest and which is the slowest.