CSE 5600 – Practice Problem Set 7

1. You are building a system around a processor with in-order execution that runs at 2.66 GHz and has a CPI of 0.7 excluding memory accesses. The only instructions that read or write data from memory are loads (20% of all instructions) and stores (5% of all instructions).

   The memory system for this computer is composed of a split L1 cache that imposes no penalty on hits. Both the I-cache and the D-cache are direct mapped and hold 32 KB each. The I-cache has a 2% miss rate and 32-byte blocks, and the D-cache is write through with a 5% miss rate and 16-byte blocks. There is a write buffer on the D-cache that eliminates stalls for 95% of all writes (i.e., there is no stall for a hit, even though the cache is write through).

   The 512 KB write-back, unified L2 cache has 64-byte blocks and an access time of 15 ns. Of all memory references sent to the L2 cache in this system, 80% are satisfied without going to main memory.

   The main memory has an access latency of 60 ns, after which any number of bus words may be transferred at the rate of one per cycle on the 128-bit-wide 133 MHz main memory bus.

   Assume that writes are similar to reads.

   Note, you may want to define a few symbols such as $t_{HIT,I\$}$, $t_{MISS,I\$}$, $\%miss_{I\$}$, $t_{AVG,I\$}$, etc. to handle each cache and memory in your formulation.

   (a) What is the average memory access time for instruction accesses?

   (b) What is the average memory access time for data reads?

   (c) What is the average memory access time for data writes?

   (d) What is the overall CPI, including memory accesses?

2. We are interested in assessing the performance impact of two different data cache organizations on the performance of a single-issue, in-order processor with a 5-stage pipeline that uses 32-bit addresses for memory references. Assume that the CPI with a perfect cache is 2.0, the clock cycle time is 1 ns (i.e., 1 GHz clock rate), and there are 0.2 memory references per instruction (20% of instructions are loads or stores).

The size of the two caches is 64 KB, and both have a block size of 64 bytes. One cache (cache A) is direct mapped and the other (cache B) is two-way set associative.

To accommodate the additional complexity of a set-associative cache, the processor cycle time must be stretched 1.25 times its value with a perfect cache. The direct mapped cache does not impact the processor's clock rate. The cache miss penalty is 75 ns for either cache organization. Assume the hit time for both caches is 1 cycle, the miss rate of the direct-mapped cache is 1.4%, and the miss rate of the set-associative cache is 1.0%.

(a) If the memory is byte-addressable, how many address bits represent the offset within the cache block for cache A? For cache B?

(b) How many blocks fit (total) in cache A? In cache B?

(c) How many address bits are used to select the index for cache A? For cache B?

(d) How many tag bits are required per cache block in cache A? In cache B?

(e) What is the average memory access time (in ns) for cache A? For cache B?

(f) What is the CPI for the machine with cache A? With cache B?

(g) For an application with 10^9 instructions, what is the total execution time for the machine with cache A? With cache B?

(h) Which machine completes the application faster?