COE 501: Computer Architecture

Problem Set 3: Memory Hierarchy

1) (10 pts) The transpose of a matrix interchanges its rows and columns. Here is the code:

```
for (i=0; i<N; i++)
  for (j=0; j<N; j++)
    output[j][i] = input[i][j];
```

Both the input and output matrices are stored in row-major order. Assume that you are executing *N*×*N* double-precision (8 bytes per element) matrix transpose on a processor with 16 KB D-Cache, which is 2-way set-associative, and 64-byte blocks. The D-Cache is a write-back with write-allocate policy on a write miss.

a) (4 pts) Assume each set in the D-Cache stores one block from the input matrix and a second block from the output matrix. How many sets exist in the D-Cache? What is the maximum value of *N* such that both the input and output matrices can fit in the 16-KB D-Cache?

b) (6 pts) A compulsory cache miss occurs when a block is referenced for the first time. Given that *N*=16, how many cache misses are caused in the 16 KB 2-way set associative cache? If each cache miss stalls the processor for 8 cycles (assuming hit in L2 cache) then what is the total number of stall cycles for matrix transpose when *N*=16? What is the impact on the CPI if the execution CPI = 1.1 (excluding cache misses)? Assume six instructions are fetched and executed per inner loop iterate plus 2 instructions per outer loop iterate.

2) (10 pts) A sample DDR SDRAM timing diagram is shown below. The cache block size is 64 bytes and each transfer consists of 8 bytes. Two transfers = 16 bytes are done per bus cycle.

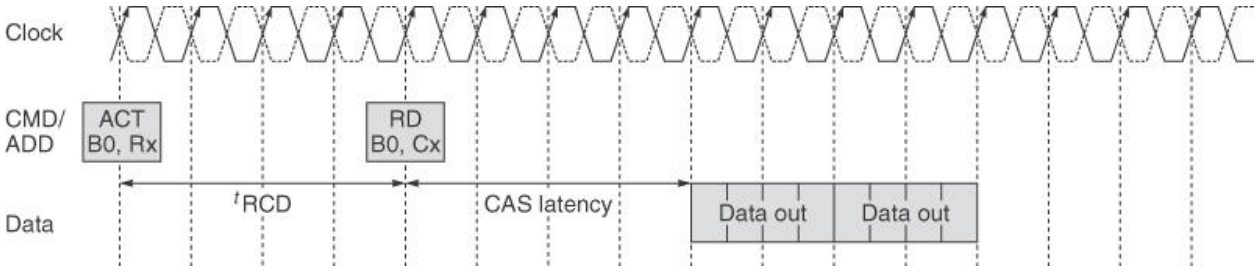

a) (5 pts) Assume that tRCD = 5 and CL = tCAS = 5 bus cycles. In addition, for each requested block (64 bytes) we automatically prefetch a second adjacent block (another 64 bytes). How much time is required from the presentation of the ACTIVE command until the two blocks are read from memory? Assume that a DDR2-800 DIMM is used (800 Millions of transfers per second at double data rate). Show your answer in bus clock cycles as well as in nanoseconds.

b) (5 pts) What is the latency when using the DDR2-800 DIMM of a burst read = 64 bytes requiring a bank ACTIVE command versus one to an already open row (no ACTIVE command)? The on-chip latency of a cache miss through levels 1 and 2 and back, not including the DRAM access is 20 ns. Add the time required to process a cache miss through L1 and L2 caches.

3) (12 pts) A processor with in-order execution runs at 1 GHz and has an execution CPI of 1.2 without counting memory stall cycles. The only instructions that access memory are loads (20% of all instructions) and stores (6% of all instructions). The memory system is composed of an I-cache and a D-cache that each has a hit time of 1 clock cycle.

The I-cache has a 3% miss rate. The D-cache is write-back with a 5% miss rate for reads and a 2% miss rate for writes. It takes 15 ns on average to access and transfer a block from the unified L2 cache into the I-cache or D-cache. Of all memory references sent to the L2 cache in the system, 20% miss in the L2 cache and require main memory access. It takes 50 ns on average to access and transfer a block from the main memory into the I-cache or D-cache.

a) (3 pts) What is the average memory access time for instruction fetching?
b) (3 pts) What is the average memory access time for data reads?
c) (3 pts) What is the average memory access time for data writes?
d) (3 pts) What is the overall CPI, including memory stall cycles?

4) (9 pts) Cache organization is often influenced by the desire to reduce the cache's energy consumption. For that purpose we assume that the cache is physically distributed into a data array (holding the data), tag array (holding the tags), and replacement array (holding information needed by replacement policy). In addition, each one of these arrays is physically distributed into multiple sub-arrays (one per way) that can be individually accessed. For example, a four-way set associative LRU cache has four data sub-arrays, four tag sub-arrays, and four replacement sub-arrays. The replacement sub-arrays are accessed once every access when the LRU replacement is used. However, it is not needed when Random replacement is used. For a specific cache, it was determined that the accesses to the different arrays have the following energy consumption:

Array	**Energy Consumption Per Way Accessed**
Data Array	20 energy units
Tag Array	5 energy units
Replacement Array	1 energy unit

Estimate the energy consumption for the following configurations. The cache is 4-way set associative. Main memory access and cache refill (although important) are not considered here. Provide answers for the LRU and Random replacement policies.

a) (3 pts) A cache read hit. All arrays are read simultaneously.

b) (3 pts) The cache access is now split across two cycles. In the first cycle, all tag sub-arrays are accessed. In the second cycle, only the sub-array whose tag matched will be accessed. Repeat part (a) for a cache read hit.

c) (3 pts) Repeat part (b) for a cache read miss, assuming that no data array access in the second cycle for a cache miss.

5) (9 pts) The following table shows parameters of a virtual memory system:

Virtual Address	Maximum Physical Memory	Page Size	Page Table Entry
50 bits	64 GB	16 KB	4 bytes

a) (2 pts) For a single-level page table, how many page table entries are needed? How much physical memory is needed for storing the page table?

b) (3 pts) Using a multilevel page table can reduce the physical memory allocation of page tables. How many levels of page tables will be needed, given the size of the page table at any level is the size of a page?

c) (2 pts) The architect wants to support a large page size, what will be the best choice for a large page size and why?

d) (2 pts) The architect wants to design a 64 KB cache that should be indexed in parallel with the TLB (address translation). The cache should be physically tagged and should not have any aliasing problem. What should be the minimum associativity of the cache to avoid aliases?