Tutorial Week of November 15^{th}

1. Consider a cache with the following attributes:

   - Associativity - k
   - Block Size - b
   - Cache Size - s

   What is the size of the tag for a n bit address space?

2. When do multiple threads in a fine grained multi-threaded machine require separate TLBs?

3. Assume the following parameters for a computer:

   (a) CPI assuming a perfect memory system (1 cycle latency) : 1.2
   (b) L1 local miss rate : 10%
   (c) L2 local miss rate : 20%
   (d) L1 latency : 1 cycle
   (e) L1 + L2 latency : 10 cycles
   (f) L1 + L2 + Main mem Latency : 300 cycles

   What is the CPI for this memory system?

4. Why is it necessary to write back data to main memory upon a $M \rightarrow S$ transition in the snoopy protocol?

5. Assume that two nodes desire to transition from the S state to the M state at exactly the same point of time. How will the Snoopy protocol ensure that only one of these nodes will enter the M state? What happens to the other node?

6. The snoopy protocol clearly has an issue with scalability. If we have 64 cores, we cannot broadcast messages to everybody. Can you propose solutions to circumvent this problem.

7. Assume a hard disk with the following parameters:

   - Seek time : 2 ms
   - Rotational Speed : 600 RPM
   - Bandwidth : 100 MB/s

   We observe that on an average the disk needs to move by 120° till the head is over the correct sector. How long will it take to transfer 25 MBs of data?

8. What is the bisection bandwidth of a 2d-torus, 3d-mesh, and a hypercube? *Bisection Bandwidth* is defined as the maximum of the minimum number of links that need to be cut to divide the network into two equal partitions.

9. What is the diameter of a 2d-torus, 3d-mesh, and a hypercube? *Diameter* is defined as the largest of all shortest paths between any two nodes in a network.