

- 1) (26/40) Let us consider a Directory-Based Invalidation Coherence Protocol for a 64-core system. In order to implement the sharer list (S, list of cores sharing a block of data) in the HOME core, a per-block bit-vector is maintained. The list S tracks all cores that share the associated data block. The k-th bit of the bit vector indicates whether or not the corresponding k-th core has a copy of the block (bit==1 → has the block, bit==0 → the core does not have a copy the block). Assume that the block-size is 32 bytes.

Question 1a) Calculate the cost of the “bit-vector” solution as a fraction of the bits that are needed to store the data. Perform the calculation for the cases of a 64-core machine and a 4-core machine.

Question 1b) Assume that the S-list is replaced by an optimized solution (“Single-Sharer”) that uses, instead of the S-list, a single 6-bit identifier (SID) which indicates which a core (the last one) which has a copy of the block. When there is a Bus-Read Transaction involving a certain block, the copy of that block in the core indicated by SID has to be invalidated (the HOME core initiates an invalidation request, then the current sharer core sends an acknowledge message to the HOME core; the HOME core then replaces the SID content with the core ID of the new sharer core and sends an acknowledge message to the new sharer core). Compare the efficiency of the “Bit-Vector” solution with the “Single-sharer” solution by filling up the first two empty columns in the following tables with the number of invalidation requests that are generated by the coherence protocol at each step. Assume that the data block B is not initially present in any cache.

SEQUENCE 1	Bit-Vector -- # of invalidations	Single-Sharer -- # of invalidations	Single-Sharer+Global-Bit -- # of invalidations
STEP1: Core #0 reads B	0	0	
STEP2: Core #1 reads B			
STEP3: Core #0 reads B			

SEQUENCE 2	Bit-Vector -- # of invalidations	Single-Sharer -- # of invalidations	Single-Sharer+Global-Bit -- # of invalidations
STEP1: Core #0 reads B	0	0	
STEP2: Core #1 reads B			
STEP3: Core #2 reads B			

Question 1c) Assume that the previous solution (“Single-Sharer”) is improved by adding a global bit G besides the S-ID (“Single-Sharer + Global-Bit”). This global bit is set to 1 to indicate that more than one core has a copy of the block. When G==1, the HOME core does not track any more the specific sharer and it assumes that all cores could potentially share the block. Fill-up the third empty column in the previous table.

Question 1d) Describe with a state diagram the complete functioning of the three coherence protocol variants of the cases 1a, 1b, 1c above, assuming three possible status conditions for the block: U=Uncached (or invalid), EM=Exclusive-or-Modified (or owned, which means not shared and possibly modified), S=Shared (Shared, which is also clean, i.e., coherent with the memory).

The possible transactions regard a Read-Block, Invalidation and Flush (or Write-back) which implies to evict the block from the cache. Use this notation: L=Sharer-List, Cx=Previous Core doing an operation on a block, Cy=Current Core doing an operation on a block.

- 2) (13/40) Write a simple program in the CUDA programming model in order to add two vectors A and B and obtain a third vector C. Each vector has a size of 100'000 floating point elements.