

- 1) (POINTS 10/20) Consider a bus-based multicore that support the DRAGON cache coherence protocol. The cost of a read/write operation is 1 cycle, the cost of a BusRd (or BusRdX) transaction is 90 cycles; all caches are write-back, write-allocate and initially empty. For the DRAGON protocol the BusUpgr is not used. The cost of a BusUpdate is 60 cycles. The cost of Flush and Flush\* is 0. Evaluate the total cost of executing the following streams by completing the table below:

Stream1: R1, W1, R1, W1, R2, W2, R2, W2, R3, W3, R3, W3

Stream2: R1, R2, R3, W1, W2, W3, R1, R2, R3, W3, W1

Stream3: R1, R2, R3, R3, W1, W1, W1, W1, W2, W3

stream-1 DRAGON

Core Operation	C1	C2	C3	Bus Transaction	Data from	Cycles
PrRd1						
PrWr1						
PrRd1						
PrWr1						
PrRd2						
PrWr2						
PrRd2						
PrWr2						
PrRd3						
PrWr3						
PrRd3						
PrWr3						
TOTAL						

stream-2 DRAGON

Core Operation	C1	C2	C3	Bus Transaction	Data from	Cycles
PrRd1						
PrRd2						
PrRd3						
PrWr1						
PrWr2						
PrWr3						
PrRd1						
PrRd2						
PrRd3						
PrWr3						
PrWr1						
TOTAL						

stream-3 DRAGON

Core Operation	C1	C2	C3	Bus Transaction	Data from	Cycles
PrRd1						
PrRd2						
PrRd3						
PrRd3						
PrWr1						
PrWr1						
PrWr1						
PrWr1						
PrWr2						
PrWr3						
TOTAL						

- 2) (10/20) For the following CUDA kernel, check if it will cause branch divergence for warp 0 and for warp 1 in thread block 0 assuming that each thread block contains more than 2 warps (explain in detail the result).

```

gid= threadIdx.x + blockIdx.x * blockDim.x;
If (gid < 8) { //this the branch of interest
    //TAKEN_PATH
} else {
    //NOT_TAKEN_PATH
}

```