

Name: \_\_\_\_\_ Username: \_\_\_\_\_

You have 20 minutes for this quiz. The quiz is closed notes/closed book. If you should finish early, you are welcome to turn in your quiz and step out of the room until the lecture starts. **There are 2 pages to this quiz.**

I have neither given nor received aid on this quiz, nor observed anyone else doing so.

---

1. Consider the following access pattern: A, B, C, D, C, A. Assume that A, B, C, and D are memory addresses each of which are in a different block of memory. Further, assume A, B, C and D are generated in a uniformly random way.
  - a. What is the probability that the second instance of “A” will be a hit on a 4-line direct-mapped cache? Briefly show your work. [15]

$$\left(\frac{3}{4}\right)^3 \sim 42.2\%$$

- b. What is the probability that the second instance of “A” will be a hit on a 4-line 2-way associative cache? Again, briefly show your work. [10]

There are 2 sets. Each of the 3 could go either in the same set as A or the other set. So 8 possibilities, each equal in probability. Half of those have 2 or more of the accesses going to the same set as A. So 50%.

Another way: 0 go with A:  $(\frac{1}{2})^3$ . Exactly 1 with A:  $3 * (\frac{1}{2}) * (\frac{1}{2})^2$ .  $\frac{1}{8} + \frac{3}{8} = \frac{1}{2}$

2. Answer the following questions about caches and TLBs.
  - a. Physically-addressed caches tend to be slower than virtually addressed caches. Explain why. [10]

You have to do the TLB access before you can do the cache access in a physically addressed cache. You don't have to in a virtually-addressed cache.

- b. In a virtually-addressed, physically-tagged (VIPT) cache, the TLB and tag store of the cache can be accessed in parallel. Using phrases like “page offset”, “page number”, “block offset”, and “set index” what are the restrictions on the VIPT's organization? [15]

The set index and block offset of the cache must come entirely from the address bits that make up the page offset.

3. Consider a case of having 3 processors using a snoopy MESI protocol where the memories can snarf data. All three have a 2-line direct-mapped cache with ***each line consisting of 16 bytes***. The caches begin with all lines marked as invalid. Fill in the following tables indicating
- If the processor gets a hit or a miss in its cache. *Note: If you need to go to the bus for the data or ownership, it is considered a miss.*
  - If a HIT or HITM (or nothing) occurs on the bus during snoop.
  - What bus transaction(s) (if any) the processor performs (BRL, BWL, BRIL, BIL)
  - For misses only, indicate if the miss is compulsory, capacity, conflict, or coherence. A coherence miss is one where there would have been a hit, had some other processor not interfered.

Finally, indicate the state of the processor after all of these memory operations have completed. The operations occur in the order shown. **[50 points, -2 per wrong or blank, minimum of 0]**

Processor	Address	Read/Write	Cache Hit/Miss	Bus transaction(s)	HIT/HITM	“4C” miss type (if any)
1	0x100	Write	Miss	BRIL	--	Compulsory
1	0x200	Read	Miss	BRL/BWL	--	Compulsory
1	0x124	Write	Miss	BRIL	--	Compulsory
1	0x100	Write	Miss	BRIL/BWL	--	Capacity
2	0x109	Read	Miss	BRL	HITM	Compulsory
2	0x110	Read	Miss	BRL	--	Compulsory
2	0x100	Read	Hit	--	--	--
1	0x100	Read	Hit	--	HIT	--
2	0x110	Write	Hit	--	--	--
1	0x100	Write	Miss	BIL	HIT	Coherence

Final state:

<u>Proc 1</u>			<u>Proc 2</u>		
	Address	State		Address	State
Set 0	0x100	M	Set 0		
Set 1			Set 1	0x110	M

*For the address, give the address of the start of the cache line.*