

Computer Architecture and Engineering
CS152 Quiz #2
March 7th, 2016
Professor George Michelogiannakis

Name: _____

This is a closed book, closed notes exam.
80 Minutes. 15 pages

Notes:

- Not all questions are of equal difficulty, so look over the entire exam and budget your time carefully.
- Please carefully state any assumptions you make.
- Please write your name on every page in the quiz.
- You must not discuss a quiz's contents with other students who have not taken the quiz. If you have inadvertently been exposed to a quiz prior to taking it, you must tell the instructor or TA.
- You will get no credit for selecting multiple-choice answers without giving explanations if the instructions ask you to explain your choice.

Writing name on each sheet	_____	1 Point
Question 1	_____	29 Points
Question 2	_____	30 Points
Question 3	_____	18 Points
Question 4	_____	22 Points
TOTAL	_____	100 Points

Question 1: Doppleganger Cache [29 points]

For this question, we will consider a new cache design that was proposed in 2015, named the doppleganger cache. The doppleganger cache is motivated by approximate computing and essentially implements lossy compression of data in cache lines, where multiple tags share the same cache line. A simplified diagram is as follows:

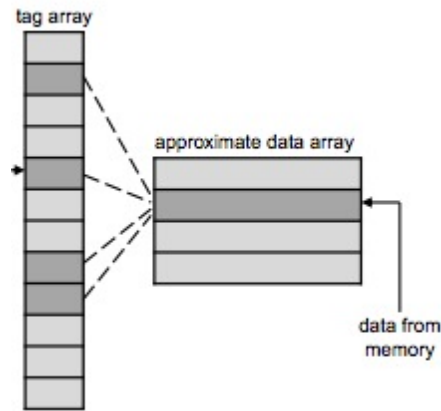


Figure 1: Doppleganger cache overview. “Data from memory” refers to incoming data in case of a cache miss.

As shown, multiple tag array entries point to the same data array line. This results in different memory addresses (different tag array entries) being serviced by the same data array line. The benefit of this design is reducing overall cache area by reducing the size of the data array, but not the coverage (entries) in the cache because the tag array preserves its number of entries.

For a more detailed view, you can use Figure A-1 in appendix A, with the modification that the data array has fewer entries than shown in the original direct-mapped cache.

For this question, we consider a doppleganger cache with 128 32-byte lines (4KB data array and blocks of 32-bytes or eight 4-byte words) and 512 tag array entries. Tag array entry I maps to data array line $I \% 128$. For example, tag arrays 0, 128, 256, and 384 all map to (use) data array entry 0. You can assume direct mapping.

When a read miss occurs, data arrives from memory into the appropriate location of the data array as determined by the tag array entry that was used. If the data array already contains a valid entry for a different tag array, the new data are “merged” with lossy compression. Therefore, the end result is that both the previous tag array and the new tag array entry now point to the same data array entry.

Q1.A Three C's of Cache Misses [3 points]

Compare the 3C's of cache misses for the doppelganger cache as provided, to a 256-entry direct-mapped cache. State which of the 3C's of cache misses increase, decrease, or remain the same for the doppelganger cache and why. How do you expect the cache miss rate to change compared to the direct mapped cache? Explain your reasoning to receive credit.

Q1.B Read Lookup [4 points]

Now lets assume that any tag array entry can map to **any** data array entry (line). That is, tag array entries are extended with a field "data array entry" to point to the data array entry the tag refers to. Now the cache needs to not only hit or miss like we have seen for each access, but also determine which data array entry to access in the case of a hit.

Name one advantage of this modified design. Name one disadvantage. Explain.

Q1.C Write Policy [5 points]

Back to the original doppleganger design in the description and Q1.A (ignore the modification in Q1.B), would you choose a write back/write allocate policy or a write through/no allocate policy in the doppleganger cache? Think of any different steps compared to the standard direct-mapped cache and how one write can affect multiple tag entries. Is there any functional (i.e., different data for the same read) difference between write back and write through in this cache?

Q1.E Cache Parameters [4 points]

Fill out the following table showing how many **bits** are in each of the caches with the assumptions in the problem description (keep ignoring Q1.B), the direct-mapped cache with 512 entries, and for a write back/allocate policy. Assume 32-bit addresses, byte addressable addresses, and 32-byte blocks.

	Unmodified Direct-mapped	Doppleganger
# tag array bits		
# data array bits		

Q1.F TLB and virtual indexing [8 points]

Now let's assume that we want to have a TLB *in parallel* with the cache. To make this happen, do we need a virtual or physical tag, and a virtual or physical index?

Using your answer above, 4KB pages, and the cache organization in this problem, is it possible to have aliasing in this cache? In other words, could there be a situation where multiple copies of the same physical page exist in the cache? Describe why not, or if there is such a case how it can happen *and how to solve it*.

Page number (virtual or physical)	Page offset
-----------------------------------	-------------

(the index uses bits starting from the least significant after the cache block offset bits)

Q1.G Two-way set associative [5 points]

Describe a modification to the doppelganger cache that makes it two-way set associative but keeps one data array. Don't forget to describe how the data array gets accessed. You can change the mapping from tag to data entries that this problem provides.

Question 2: Page Tables [30 points]

For this question, you are going to study a 32-bit machine that uses a hierarchical page table scheme. Assume that each page table segment can hold 512 entries. In other words, the first level of the page table has 512 entries. Each entry points to another segment that can hold 512 entries, etc. Virtual addresses are 32 bits long. This machine uses 4KB pages. Assume that the operating system is smart enough to only allocate the minimal needed physical memory for each program. In all questions, **explain** your answers.

Q2.A Bit indices [6 points]

Given the above information, and assuming you want to use as many bits as possible of the address, how many bits for the page offset? How many levels for the hierarchical page table? How many index bits per level?

In your above answer, it is possible that some bits cannot be assigned to anything given the constraints. What can you change from the constraints to assign all bits?

Q2.B DRAM Size [4 points]

With your original answer from Q2.A (without the modification to fit the unassigned bits), what is the maximum DRAM size we can have in this system?

Q2.C Time for a TLB Refill [6 points]

Now let's assume that the latency to access a page table segment is $\langle \text{number of entries} \rangle * 1\text{ns}$. So for 512 entries it is 512ns. Also assume that there is a fixed penalty (for example for pointer chasing) to go from one level to the next that is 10ns. Calculate the latency to walk the page table you designed in Q2.A. Can you propose a modification to this page table structure that has a lower latency but maps to the same or more amount of DRAM (i.e., no less)?

Q2.D Page Table Size [4 points]

What is the size of the page table if a program allocates the maximum it is allowed to (assume the page table structure from Q2.A. Assume a page table entry in any level is 8 bytes.

What is the size of the page table if a program allocates only one page?

Q2.E Reducing Page Walk Size [10 points]

Based on the page table structure of Q2.A, assume you want to speed up the translation of some important pages, such as pages belonging to the operating system. Let's assume that pages 0 to 127 have to be resolved with the minimum latency possible, that is smaller than all other pages. Pages from 128 and upward are to be translated with higher latency.

Describe how you would do this while not decreasing the amount of DRAM we can map in this machine (similar to Q2.C). You can modify an existing structure or add a new structure, as long as you describe what the functionality should be. However, make sure you maintain the same first-level structure (i.e., there still is one first page table segment).

For the modified structure you propose, calculate the latency to translate pages 0 to 127, and then pages 128 and upwards. You can use the assumption of $\langle \text{number of entries} \rangle * 1\text{ns}$ for any array structure you propose. Also, to go from one level of a hierarchical page table to another, add 10ns.