

Introduction to Computer Architecture: Supervision 4

Lectures covered by the supervision: <https://www.cl.cam.ac.uk/teaching/2526/IntComArch/>

- Lecture 13: Multicore processors – part II
- Lecture 14: Graphics processing units (GPUs) – part I
- Lecture 15: Graphics processing units (GPUs) – part II

Past exam questions:

<https://www.cl.cam.ac.uk/teaching/exams/pastpapers/t-ComputerDesign.html>

Supervision questions:

1. [2015 Paper 5 Question 3](#) - part b.
 - a. With load linked / store conditional, why does the store alter the source register containing the value to write?
 - b. Can you link these operations used for synchronisation and ensuring cache coherence to the architectural improvements (number of transistors) used for boosting performance of CPUs?
2. [2016 Paper 5 Question 3](#) - part b.
3. Why do we place a memory barrier after taking a lock and before releasing the lock, but not the other way round (i.e., before taking and after releasing)?
4. Assuming a single instruction, xchg, that can do an atomic exchange of a register and memory location, how would you implement a naive spin lock?
5. Explain relaxed consistency model. What are the alternatives?
6. [2019 Paper 5 Question 3](#) - part e.
7. Describe SIMT, one-way GPUs exploit parallelism.
8. How does GPU processor pipeline differ from a general-purpose CPU to allow SIMT execution? What are the implications for other parts of processor for providing SIMT (e.g. the data cache)?
9. How do GPUs allow branching within the different threads?
 - a. What exactly happens during the branching?
 - b. Can you discuss the limit of branching at which GPU parallelism stops being useful?
10. How do we calculate efficiency of code execution on GPU? Sketch an example.
11. Explain warp parallelism and SIMT parallelism.
12. What are the types of memory available within a GPU and what are their relative advantages or disadvantages?
13. Summarize the main message from lecture 13 in 1-3 sentences?
14. Summarize the main message from lecture 14 in 1-3 sentences?
15. Summarize the main message from lecture 15: GPU memory in 1 sentence?

Save your answers into MS Teams or email them to me. Please use the following naming pattern:

ICA_Supervision_4_Answers_<last name>_<first name>_Michaelmas_2025

Send your answers as a pdf, doc, image, or any other format of a document for which there exists an easily available software to open.

Jasmin JAHIC

jj542@cam.ac.uk

<https://www.cl.cam.ac.uk/~jj542/>