

## COMPUTER SCIENCE TRIPOS Part IB – 2024 – Paper 5

### 8 Introduction to Computer Architecture (swm11)

- (a) Why do multicore systems-on-chip use hierarchical caches rather than one large shared cache? [3 marks]
- (b) Why are page tables hierarchical rather than flat? Illustrate your answer by considering the memory required to store the page table for:
- an application with a contiguous 6 MiB memory footprint starting at virtual address zero,
  - running on a 32-bit processor, with a 32-bit physical address space and 4 KiB pages. [5 marks]
- (c) What is a TLB and why do we need one per core? [3 marks]
- (d) What is the difference between inclusive, exclusive and non-inclusive non-exclusive (NINE) cache inclusion policies? [3 marks]
- (e) In multicore system with private level-1 and level-2 caches and one shared level-3 cache, why might the level-2 cache be inclusive but the level-3 cache be NINE? [3 marks]
- (f) What is the difference between cache coherence and consistency? Give examples of two consistency models. [3 marks]