

University of Bath

**DEPARTMENT OF COMPUTER SCIENCE
EXAMINATION**

CM30225: PARALLEL COMPUTING

Wednesday, 25 January 2023, 9:30–11:30

No calculators may be brought in or used.

Full marks will be given for correct answers to THREE questions. If you opt to answer more than the specified number of questions, you should clearly identify on the cover sheet which of your answers you wish to have marked. In cases you have failed to identify the correct number of answers the marker is only obliged to consider the answers in the order they appear up to the number of questions required.

DURING THIS EXAM YOU ARE NOT PERMITTED TO COMMUNICATE WITH ANY PERSON(S) EXCEPT AN INVIGILATOR OR AN ASSIGNED SUPPORT WORKER.

YOU MUST NOT HAVE ANY UNAUTHORISED DEVICES OR MATERIALS WITH YOU.

YOU MUST KEEP YOUR LIBRARY CARD ON YOUR DESK AT ALL TIMES.

PLEASE FILL IN THE DETAILS ON THE FRONT OF YOUR ANSWER BOOK/COVER AND SIGN IN THE SECTION ON THE RIGHT OF YOUR ANSWER BOOK/COVER, PEEL AWAY ADHESIVE STRIP AND SEAL.

TAKE CARE TO ENTER THE CORRECT CANDIDATE NUMBER AS DETAILED ON YOUR DESK LABEL.

DO NOT TURN OVER YOUR QUESTION PAPER UNTIL INSTRUCTED TO BY THE CHIEF INVIGILATOR.

1. Describe each of the following primitives, explaining their main purpose and giving an example of a typical use:
 - (a) locks (mutexes)
 - (b) semaphores
 - (c) barriers
 - (d) atomics

[4 each]

Primitives like these are typically *not fair* in real implementations. Explain what this means, and give some reasons why fairness is not enforced. [4]

2. There are several kinds of shared memory systems. Write notes on each of the following architectures, paying particular attention to scalability and memory bottlenecks:
 - (a) symmetric shared
 - (b) NUMA
 - (c) virtual shared
 - (d) vector processor

[4 each]

Which of the above architectures scales the best to large numbers of processors? Explain. [2]

Why is virtual shared memory not a popular architecture? [2]

3. What does it mean when we say a computation is *small grain* or *large grain*? [2]
 What is *load balancing*, and why is it important in parallel systems? [3]
 Describe each of the following work distribution models:

- (a) Provider/Consumer
- (b) Manager/Worker

[2 each]

For both, explain what they need to do for load balancing. [2 + 2]

A *thread pool* can be used to implement Manager/Worker. Describe this. [3]

Grand Central Dispatch (GCD) is a thread pool mechanism. What is it and what are its advantages and disadvantages? [4]

4. There have been many different approaches attempting to provide support for writing parallel programs. Write an essay describing these approaches, giving their advantages and disadvantages. You should cover *at least*: library-based approaches, modifications of existing programming languages, and designing new programming languages. [20]