THE UNIVERSITY OF NEW SOUTH WALES

Final Examination

SAMPLE

COMP3231/COMP9201

Operating Systems

- Time allowed: **2 hours**

- Reading time: **10 minutes**

- Total number of questions: **6**

- Answer **all** questions

- The questions are **not** of equal value

- This paper may **not** be retained by the candidate

- **Answers must be written in ink**, with the exception of graphs and multiple-choice answers sheets.

- No examination materials may be used.

- Electronic calculators provided by the University may be used.

- **Provide answers to Question 1 in the space provided on the examination paper itself.**

- **Use a *separate* answer book for *each* of the other questions.**

- **Marks will be subtracted if the above rules are not followed.**

Question 1 [16 Marks]

Answer this question *on the examination paper* itself.

Classify each statement as true, by crossing the ☒T☒ box, or false, by crossing the ☒F☒ box. You will receive one mark for each correct classification, and lose one mark for each incorrect classification. You gain zero marks for each answer left unclassified. The overall mark for this question will not be negative, i.e. the minimum mark is zero.

☐T ☐F   A smaller page size leads to smaller page tables.

☐T ☐F   A smaller page size leads to more TLB misses.

☐T ☐F   A smaller page size leads to fewer page faults.

☐T ☐F   A smaller page size reduces paging I/O throughput.

☐T ☐F   Threads are cheaper to create than processes.

☐T ☐F   Kernel-implemented threads are cheaper to create than user-level-implemented threads.

☐T ☐F   A blocking kernel-implemented thread blocks all threads in the process.

☐T ☐F   Threads are cheaper to context switch than processes.

☐T ☐F   A blocking user-level-implemented thread blocks the process.

☐T ☐F   Different user-level threads of the same process can have different scheduling priorities in the kernel.

☐T ☐F   All kernel-scheduled threads of a process share the same virtual address space.

☐T ☐F   The *optimal* page replacement algorithm is the best choice in practice.

☐T ☐F   The operating system is not responsible for resource allocation between competing processes.

☐T ☐F   System calls do not change the privilege mode of the processor.

☐T ☐F   A scheduler favouring I/O-bound processes usually does not significantly delay the completion of CPU-bound processes.

Question 2 [12 Marks]

Answer this question in a *separate* book

A) **[6 marks]**

Consider a demand-paging system with a paging disk that has an average access and transfer time of 5 milliseconds for a single page. Addresses are translated through a page table in main memory, with an access time of 100 nanoseconds per memory access. Thus, each memory reference through the page table takes two accesses. The system has a 48-entry TLB to speed up memory accesses. Assume that 99% of memory accesses result in a TLB hit, and of the remaining 1%, 5 percent (or 0.05% of the total) cause page faults. What is the effective memory access time?

B) **[6 marks]**

Some versions of UNIX store the first part of each file in the same disk block as the inode. Discuss why this might be advantageous in practice.

Question 3 [14 Marks]

Answer this question in a *separate* book

A) **[8 marks]**

Suppose that the head of a moving-head disk with 192 tracks, numbered 0 to 191, is currently serving a request at track 80 and has just finished a request at track 62. The queue of requests is kept in the FIFO order: 119, 58, 114, 28, 111, 55, 103, 30, 75. What is the total number of tracks traversed by head movements needed to satisfy these requests for the following disk-scheduling algorithms?

   i) FCFS.

   ii) SSTF.

  iii) Elevator (SCAN).

  iv) Modified Elevator (C-SCAN).

B) **[6 marks]**

User-level threads packages generally implement cooperative scheduling. What is cooperative scheduling, and why is it the common method used to schedule user-level threads?

Question 4 [10 Marks]

Answer this question in a *separate* book

Explain how a 32-bit virtual address is translated into a physical address on a system using a two-level page table and 4kb pages. Your explanation, where it refers to parts of an address, **must** specifically state which bits of the address you are talking about. (Ignore any TLB). **State any assumptions you make!**

Question 5 [14 Marks]

Answer this question in a *separate* book

A) **[6 marks]**

Describe the difference between external and internal fragmentation. Indicate which of the two are most likely to be an issues on a) a simple memory memory mangement machine using base limit registers and static partitioning, and b) a similar machine using dynamic partitioning.

B) **[8 marks]**

Describe the difference between a normal test-and-set spinlock and a read-before-test-and-set spinlock. Why would the latter be advantageous over the former on multiprocessor systems.

Question 6 [8 Marks]

Answer this question in a *separate* book

Describe the four conditions required for deadlock to occur. Describe a common method for deadlock prevention that prevents one of the conditions occurring.