國立中央大學資訊工程學系博士班 114 學年度第一學期資格考試題紙

科目：**作 業 系 統** (Operating System)

1. (10%) A disk driver must schedule disk requests for tracks 10, 22, 20, 2, 40, 6, and 38. A seek takes 6 msec per track moved. Assume that the disk arm is initially at track 20, and the driver follows the SSTF (Shortest seek time first) scheduling algorithm. How much seek time is needed to handle all of these requests?

2. (20%) Given the following stream of page references by an application, calculate the number of page faults the application would incur for (a) FIFO page replacement and (b) LRU page replacement, respectively. Assume that all pages are initially free, and 3 physical pages are available.
   Reference Stream: A, B, C, D, A, B, E, A, B, C, D, E, B, A, B

3. (20%) Are the following statements true or false? For each statement, you will get 4 points for correct answer, zero point for blank, or -2 point for incorrect answer.
   (a) A worm is a process that uses the spawn mechanism to duplicate itself.
   (b) Breach of confidentiality involves unauthorized modification of data.
   (c) Little's Law can be applied to a system with multiple queues
   (d) A deadlock-free solution is guaranteed to be starvation-free.
   (e) A context switch always involves updating the page table, whether it is between processes or threads

4. (15%) A deep-learning inference engine runs on a server with 8 GB DRAM and a PCIe-attached GPU (6 GB memory). The neural-network weights (5 GB total) are stored in virtual memory and accessed through page-mapped host memory that the GPU can DMA directly.
   (a) Each page is 4 KB. When the model first runs, 3 GB of weights must be paged in from disk into DRAM before the GPU can access them. How many page faults occur during the first pass? Show the formula and numeric result.
   (b) Assume the average page-fault service time (including disk I/O) is 2 ms and DRAM access time is 200 ns. Compute the effective access time (EAT) if the page-fault rate $p = 10^{-4}$. Give the result in ns, rounded to two significant digits.
   (c) The system is upgraded to huge pages (2 MB) for weight tensors. Assuming total weight size and total memory stay the same, how does this change the number of TLB entries required and the probability of TLB misses (qualitatively)?

背面還有 Please Turn Over

5. (20%) A system has 32-bit virtual addresses, 4 KB pages, and 64 MB of physical memory. It uses demand paging and a single-level page table.
   (a) How many page frames are in physical memory? How many entries are in the page table? If each entry is 4 bytes, compute the page-table size in bytes.
   (b) The system experiences an average page-fault service time of 8 ms. If effective access time must stay below 250 ns, find the maximum allowable page-fault rate p.

6. (15%) Modern SSDs reorder or merge writes inside their Flash Translation Layer (FTL) and perform erase-before-write on flash blocks.
   (a) Explain how these SSD characteristics can interfere with a file system's crash-recovery guarantee.
   (b) Compare how a journaling file system (e.g., ext4) and a copy-on-write file system (e.g., btrfs / F2FS) maintain metadata consistency after a crash when running on SSDs.
   (c) State one performance and one endurance trade-off introduced by each approach under flash storage.