

資格考試科目：高等作業系統

Instructions: There are **five** questions each of which counts 20 points. Each question may have several sub-questions. Please read the questions carefully before answering the questions.

1. (Concurrency – 20 pts) - Threads and Processes :
 - A. Describe the reason why increasing the number of active processes or threads (i.e., the degree of concurrency) can improve system throughput.
 - B. Describe the reason why increasing the number of active processes or threads (i.e., the degree of concurrency) can reduce system throughput.
 - C. While developing a new internet banking system which accesses highly confidential data and which accepts requests from Internet users, Alice would like to a multi-process/thread system so as to shorten the user waiting time. Should Alice use a multi-process or multi-thread system to implement her new internet banking system? Please explain your answer.
 - D. While developing the next generation Google search engine which also accepts the requests from Internet users, Bob also like to a multi-process/thread system so as to shorten the user waiting time. Should Bob use a multi-process or multi-thread system to implement his new system? Please explain your answer.

2. (20 pts Deadlock) There are two general approaches to deal with deadlocks. One general approach is called deadlock prevention, which prevents deadlock from occurring. (1) What is the other general approach to deal with deadlocks? (2) How do you ensure that starvation will not occur when dealing with deadlocks (first describe the starvation problem)?

3. (20 pts Storage) A RAID 50 (RAID 5+0) combines the block-level striping with distributed parity of RAID 5, with the straight block-level striping of RAID 0. Below is an example of RAID 50.

RAID 50

Disk 1	Disk 2	Disk 3	Disk 4	Disk 5	Disk 6
Block 1	Block 2	Parity 1-2	Block 3	Block 4	Parity 3-4
Block 5	Parity 5-6	Block 6	Block 7	Parity 7-8	Block 8
Parity 9-10	Block 9	Block 10	Parity 11-12	Block 11	Block 12

To see how RAID 50 differs from RAID 5, below is an example of RAID 5

RAID 5

Disk 1	Disk 2	Disk 3	Disk 4	Disk 5
Block 1	Block 2	Block 3	Block 4	Parity 1-4
Block 5	Block 6	Block 7	Parity 5-8	Block 8
Block 9	Block 10	Parity 9-12	Block 11	Block 12
Block 13	Parity 13-16	Block 14	Block 15	Block 16
Parity 17-20	Block 17	Block 18	Block 19	Block 20

Please describe one main advantage and one main disadvantage of RAID 50 over RAID 5 (in terms of performance and reliability)?

4. (20 pts Real-time Scheduling) Say if you want to build a real-time scheduler at the user-level in Linux, which provides a priority-based scheduler. By a real-time scheduler, it means that user-level processes (e.g., a multimedia process) can reserve (execution time, period), and your user-level real-time scheduler can take this reservation and provide some kind of guarantees to them. How would you build such a user-level real-time scheduler (hint – your scheduling process can change priority of processes)? Please provide a brief high-level description. To help you explain your solution, you can assume that you have two processes X and Y. Process X needs 20

msec of CPU execution time to every 100 msec. Process Y needs 40 msec of CPU execution time to every 250 msec.

5. (20 points) Vector time stamps protocol extend Lamport's logic clock for ordering multicasting messages in distributed systems. In the following, there are three processes: P1, P2, and P3, each of which may broadcast messages to the others. Answer the following questions.
- (18 Points) Please label the vector time stamps for each event (including sending event and receiving event)
 - (6 Points) What are the message orders on Process P1, P2, and P3?
 - (6 Points) Please describe the possible causes for the out-of-order messages.

