系所組別：電腦與通信工程研究所甲組

考試科目：計算機組織與作業系統 考試日期：0223，節次：1

※ 考生請注意：本試題不可使用計算機 請勿在本試題紙上作答，否則不予計分

Part I 40%

1. The Linux Scheduler is a preemptive, priority-based algorithm with two separate priority ranges: _____ and _____.
   These two ranges maps into a global priority scheme wherein numerically lower values indicate higher priority. A runnable task is considered eligible for execution on the CPU as long as it has time remaining in its time slice. When a task has exhausted its time slice, it is considered expired and is not eligible for execution again until all other tasks have also exhausted their time quanta. The kernel maintains a list of all runnable tasks in a runqueue data structure. Each runqueue contains two priority arrays:_____ and _____.
2. A semaphore S is an integer variable that, apart from initialization, is accessed only through two standard atomic operations: _____ and _____. The value of a binary semaphore can range only between _____.
   The main disadvantage of the semaphore is that it requires _____.
3. Deadlocks can be described more precisely in terms of a directed graph called _____. This graph consists of a set of vertices V and a set of edges E. The set of vertices V is partitioned into two different types of nodes: _____, and _____. If this graph does not have a cycle, then the system is not in a _____ state.
4. A thread library provides the programmer with an API for creating and managing threads. There are two primary ways of implementing a tread library. The first approach is _____.
   The second approach is _____.
5. In order to ensure that the operating system maintains control over the CPU. We cannot allow a user program to get stuck in an infinite loop or to fail to call system services and never return control to the operating system. To accomplish this goal, we can use a _____.
6. Java is a popular object-oriented programming language introduced by Sun Microsystems. The _____ is a specification provided by Java for an abstract computer. It consists of _____ and _____that executes the architecture-neutral bytecodes.
7. A socket is defined as an endpoint for communication. A pair of processes communicating over a network employ a pair of sockets-one for each process.
   A socket is identified by _____ concatenated with _____.

Part II 10%

8. What is the cause of thrashing? How does the system detect thrashing? Once it detects thrashing, what can the system do to eliminate this problem?

(背面仍有題目,請繼續作答)

※ 考生請注意：本試題不可使用計算機

Part III 50%

9. Read the following paragraph and choose the correct answers from the following multiple choice questions (10%). Each question may have more than one answer. No partial point, no penalty.

Just as CPU programmers were forced to explicitly manage CPU memories in the days before virtual memory, for almost a decade, GPU programmers directly and explicitly managed the GPU memory hierarchy. The recent release of NVIDIA's Fermi architecture, however, has brought GPUs to an inflection point: it implements a unified address space that eliminates the need for explicit memory movement to and from GPU memory structures. Yet, without demand paging, something taken for granted in the CPU space, programmers must still explicitly reason about available memory. The drawbacks of exposing physical memory size to programmers are well known.

Which of the followings are true?

a. Without virtual memory technology, CPU programmers must explicitly manage the physical memory.
b. There is no need for the programmers of the Fermi architecture to worry about the available memory to use.
c. With virtual memory, a CPU programmer implements demand paging policy in his/her program.
d. GPU stands for General Processing Unit.

10. Answer the following questions about virtual memory. (15 %)
    a. What event triggers a page fault? 5%
    b. What event triggers a TLB miss? 5%
    c. Does a process have its own page table or all the active processes share a page table? 5%

11. For CPUs, the problem of exception support was solved at a relatively early stage. This support was a key enabler to their success, and instrumental in this success was the definition of precise exception handling. In a pipelined processor, multiple exceptions may occur at the same clock cycle. Assume this is a five-stage pipeline. Write down which exception (if any) may occur at the IF, ID, EXE, MEM, WB stage.(25%)