Operating Systems, PhD Qualifying Exam, Fall 2014

- This is 2-hour closed-book test. Calculator is allowed.

Name: _____

1. [15 points] Consider the following program, which attempts to use a pair of semaphores for mutual exclusion.

```
Initially: s = 1, t = 0.

    Thread 1            Thread 2
    P(s);               P(s);
    V(s);               V(s);
    P(t);               P(t);
    V(t);               V(t);
```

   (a) [10 points] The program always causes deadlock. Explain why.

   (b) [5 points] What change to the initial semaphore values will eliminate the potential for deadlock? If any change still causes deadlock, explain.

2. [20 points] Shortest Job First (SJF) scheduler is known to give the best turnaround time compared to other schedulers. However, the length of the job is unknown when a job is submitted. It is almost impossible to construct a SJF schedule. The Multi-Level Feedback Queue (MLFQ) scheduler approximates a schedule to the ideal SJF one without a *priori* job length. Describe the MLFQ scheduling algorithm with two-level queues and explain how it minimizes the average turnaround time.

3. [15 points] Suppose that a 1 MB file consisting of 512-byte logical blocks is stored on a disk drive with the following characteristics:

Parameter	Value
Rotational speed	10000 RPM (Rotation Per Minute)
Average seek time (T_{avg_seek})	5 ms
Average number of sectors/track	1000
Surfaces	4
Sector size	512 bytes

   Suppose that a program reads the logical blocks of the file sequentially, one after the other, and that the time to position the head over the first block is average seek time (T_{avg_seek}) plus average rotation time ($T_{avg\_rotation}$). Estimate the time required to read the file given the best possible mapping of logical blocks to disk sectors (i.e., sequentially).

4. [30 points] Consider the virtual memory system using a 42-bit virtual address and a 26-bit physical address for a page size of 4K bytes.
   (a) [15 points] Page table entry (PTE) consists of one valid bit, one dirty bit, two protection bits, and reference bits for LRU replacement. The virtual memory system has the exact implementation of LRU and maintains a total order of all pages accessed. Calculate the size of page table. Note that the size of PTE is always multiple of 8 bits (byte).
   (b) [15 points] Explain the *clock* algorithm that approximates LRU, and calculate the size of page table for the clock algorithm with the same assumption in part (a). (If you do not know the clock algorithm, propose any LRU approximation algorithm, explain why it works, and answer rest of the questions.)

5. [20 points] Virtualization needs a good scheme of shared resource management for efficient resource utilization. Memory is one of resources shared by multiple *guest* operating systems (OS) and managed by virtual memory system running on *host* OS. When guest OS needs a new page, guest OS asks host OS a new page and the host OS allocates a page for guest OS in the physical memory. If different virtual machines run the same code for OS, many pages for guest OS can be shared in the physical memory. Unfortunately, host OS does not know the page requested by guest OS is same as the page existing in the memory (i.e., previously requested by other guest OS). Thus, host OS will allocate a new page in the memory, which results in the waste of memory resource. Design an efficient scheme that enables page sharing among guest operating systems.