

1. (20 pts) Please answer the following questions related to process management and CPU scheduling:
  - a. (5 pts) What are the differences between processes and kernel-level threads?
  - b. (5 pts) Compare the differences (and the advantages/disadvantages) between user-level threads and kernel-level threads.
  - c. (5 pts) What are the differences between preemptive and non-preemptive scheduling?
  - d. (5 pts) Compare the advantages and disadvantages of CPU scheduling algorithms First-Come-First-Served (FCFS) and Shortest-Job-First (SJF).
  
2. (15 pts) Consider the classic bounded-buffer producer-consumer problem (in which the producer writes data into a share bounded buffer and the consumer reads data from the same buffer). Assume the size of the bounded buffer is  $n$ . Please enhance/modify the following code by using binary and counting semaphores to ensure that (1) semaphores are properly initialized, (2) there is no buffer overflow and underflow, and (3) accesses to the buffer are mutual exclusive. (Suppose we have an abstract data type **semaphore** that provides three operations: **wait()**, **signal()**, and **count()**, where operation **count()** is used for initializing a semaphore. For example, assuming *sem* is a variable of type **semaphore**, *sem.count(5)* will initialize *sem* to the value of 5.)

```
// Semaphore declaration and initialization code
...

// Producer process
while (TRUE) {
    ...
    add an item to buffer;
    ...
}

// Consumer process
while (TRUE) {
    ...
    read and remove an item from buffer;
    ...
}
```

3. (15 pts) Consider a computer system with a 32-bit virtual address space where paging is used. Assuming the page size is 4K bytes and the memory is byte-addressable, please answer the following questions:
- (6 pts) How many pages can a process have at most? Suppose the maximum physical memory size is 32 GB. What is the number of bits for physical addresses? What is the maximum number of frames for the system?
  - (5 pts) Let the memory access time and TLB access time be 100ns and 20ns, respectively. If we want an effective memory access time of less than 140ns, what is the minimal TLB hit ratio that needs to be achieved?
  - (4 pts) Suppose the virtual memory of the system adopts demand paging. Assume the effective memory access time of the computer system without any page fault is 100ns, and the service time for a page fault is 15ms. If the page fault rate is 0.0000004, what is the effective access time under demand paging?
4. (20 points) Disk I/O has been a critical overhead in the operating systems.
- (10 points) Please define FCFS Scheduling, SSTF Scheduling, SCAN scheduling, C-SCAN scheduling, and LOOK scheduling algorithms.
  - (10 points) Under different system workload patterns, we may have to choose different disk scheduling algorithm. Please choose a best scheduling algorithm for the following two workload patterns.
    - Lightly loaded system that has occasional burst of disk accesses.
    - Heavily loaded system which rarely has an empty request queue.
5. (10 points) Describe the basic idea of copy-on-write and two different applications of this mechanism.
6. (20 points) As the network communication overhead becomes less and less, distributed file systems have become a more attractive mechanism in modern operating systems.
- (10 points) In a distributed file system, what are the tradeoffs between caching files (or parts of files) locally and accessing files only via remote procedure calls (RPCs)?
  - (5 points) Under what circumstances should a distributed file system perform better with RPCs only?
  - (5 points) Under what circumstances is using local caching more efficient than using RPCs only?