

1. [30pts] Caching is very important in modern computer systems, especially when the performance gap between a processor and the DRAM memory keeps increasing. Please answer the following questions: You must provide explanation to receive any point.
  - A. [12pts] Which of the following layer in the memory hierarchy is managed by the operating system: Registers, Cache, Main-Memory, Disk.
  - B. [8pts] When multicore processors are considered, there could be a problem called "Memory Stall." Please explain what "Memory Stall" is.
  - C. [10pts] The ideas of caching and buffering also happen to I/O-subsystem designs. Please explain the difference between the ideas of buffering and caching.
  
2. [25pts] Process scheduling determines when to allocate CPU to which process. There must be a criterion for process scheduling to fulfill the system goals. Please answer the following questions: You must provide explanation to receive any point.
  - A. [5pts] There are different ways to classify process scheduling algorithms. Please explain the difference between the short-term scheduler and the mid-term scheduler.
  - B. [10pts] When virtualization is considered, the virtualization software presents one or more virtual CPUs to each of the virtual machines running on the system and then schedules the use of the physical CPUs among the machines. Please tell us the potential impacts of virtualization on process scheduling.
  - C. [10pts] Processor affinity might help to prevent process from migration among processors. Why we like to have processor affinity to avoid process migration?
  
3. [30pts] Memory management tries to keep more processes in the main memory for multiprogramming. Please answer the following questions: You must provide explanation to receive any point.
  - A. [10pts] Paging avoids external fragmentation and allows a process to occupy non-consecutive memory regions. Logical addresses are mapped to physical addresses by a page table. Is the number of bits of a logical address always larger than that of a physical address?

- B. [8pts] Why the idea of inverted page table works?
  - C. [12pts] Based on the Locality Model, a process executes by moving from a locality to another locality, where a locality is defined as a set of pages. Please explain the relationship between the idea of Working Set and the Locality Model.
4. [15pts] In a distributed systems, processes might run on different computers and acquire services from another computer. One well-known example is a network file system. There are two major approaches for a client to access a remote file: The stateful service and stateless service. As a hint, the stateless service makes each request self-contained, e.g., a file read with file identification and the position to read. For the stateful service, the server might track each file being accessed by each client. Please give me two advantages and disadvantages in choosing the stateless service.