

Final
Aug 2, 2004 16.30–18.30

Part A. Each question is 5 points

1 Why is the separation of mechanism and policy a desirable principle?

2 What is a **race condition**?

3 What is starvation, give an example?

4 i. What are seek time and latency time? (in disk storage)

ii. What characteristics determine transfer time?

5 What is meant by the term *context switch*? What might cause a context switch to occur?

6 Consider the following sets of processes, with the length of the CPU-burst time given in milliseconds. The processes are assumed to have arrived in the order P1, P2, P3, P4, P5 all at time 0.

Process	Burst Time	Priority
P1	1	1
P2	2	3
P3	8	3
P4	1	4
P5	5	2

i. Draw charts illustrating the execution of these processes using FCFS, SJF (non-preemptive), RR (quantum=1) and a non-preemptive priority (a smaller priority number implies a higher priority) scheduling.

ii. What is the average turnaround time of the processes for each of the scheduling algorithms.

Process	RR	PS	FCFS	SJF
P1				
P2				
P3				
P4				
P5				
Average				

7 Describe internal and external fragmentation and explain the difference between them. Which one occurs in paging systems? Which one occurs in systems using pure segmentation?

8 Explain how, under the basic segmentation scheme, the logical address `<segment-number,offset>` is translated into a physical address.

9 What is the memory management unit (MMU) and what does it do?

10 What causes a page fault? What actions may be taken by the O/S when servicing a page fault?

11 Someone has written new memory allocator to replace the standard `malloc()/free()` implementation. It works as follows: one half of available memory is divided into fixed-sized units of 4KB, and the other half is managed by a best-fit free list. If an allocation request is less than or equal to 4KB and there is space in the fixed-sized half, a 4KB unit is allocated from the fixed-sized half; otherwise, the best-fit algorithm is used over the other half of memory, and the requested size is returned (if space is available).

a) Assuming 32KB of total memory is available, what series of allocation requests will most quickly lead to all of memory getting allocated, all while requesting the least total amount of memory?

b) What type(s) of fragmentation occurs with this new allocator?

12 Name and describe two examples of where we saw the OS needed support from the hardware in order to implement some needed functionality (hint: think about process management and memory management).

13 Explain the UNIX index node (inode) structure.

Part C. Each question is 20 points

14 The Sleeping-Barber Problem. A barbershop consists of a waiting room with n chairs and the barber room containing the barber chair. If there are no customers to be served, the barber goes to sleep. If a customer enters the barbershop and all chairs are occupied, then the customer leaves the shop. If the barber is busy but chairs are available, then the customer sits in one of the free chairs. If the barber is asleep, the customer wakes up the barber. Write pseudo code to coordinate the barber and the customers.

Hints: Need a global counter for the number of customers that are currently waiting in a chair. Use 3 semaphores - one for the customers (barber waits til a customer signals they are ready), one for the barber (customer waits til barber signals he is ready) and one for the critical section (where the global counter is modified).