Course (410 or 689):_____

This examination is closed book and closed notes **except** you are permitted a one page review sheet. The review sheet

1. Must be attached to the exam
2. Must have been prepared by you
3. Is limited to one side of 8.5 by 11 inch paper

The examination has **5** pages (including this one) and **20** questions.

Put your name on each page of the exam.

If a question appears ambiguous to you, it is to your advantage to state your interpretation.

All questions will be given equal weight in grading.

Question 1: (Ch.1) What is the difference between a *hard real-time system* and a *soft real-time system*?

A hard real-time system guarantees that critical tasks be completed on time.
In a soft real-time system a critical real-time task gets priority over other tasks and retains that priority until it completes.

Question 2: (Ch. 2) For a magnetic disk, what are the two components of its *positioning time* (sometimes called the *random-access time*)? (Give the name and a brief definition of the two components.)

The two components are seek time and rotational latency. The seek time is the time needed to move the disk arm to the desired cylinder. The rotational latency is the time for the desired sector to rotate to the disk head.

Question 3: (Ch. 2) What are the two modes in dual mode operation? (Indicate the name of the two modes and tell what the distinction is between them.)

User mode and monitor mode (also known as supervisor mode, system mode, or privileged mode). In user mode, a task is executed on behalf of the user. In monitor mode it is executed on behalf of the operating system. Privileged instructions can only be executed in monitor mode.

Question 4: (Ch. 3) What is the main function of the *command interpreter*?

The main function of the command interpreter is to get and execute the next user-specified command.

Question 5: (Ch. 4) From what process state to what process state does the dispatch operation switch between?

The dispatch operation switches from the ready state to the running state.

Question 6: (Ch. 4) When a process creates a new process, there are two possibilities in terms of the address space of the new process. What are they?

1. the child process is a duplicate of the parent process
2. the child process has a program loaded into it

Question 7: (Ch. 4) What is cascading termination?

If a process terminates (either normally or abnormally) then all of its children must also be terminated.

Question 8: (Ch. 5) Name four elements of lightweight processes that are *not* shared among each other (the same four are also distinct for heavyweight processes).

1. thread ID
2. program counter
3. register set
4. stack

Question 9: (Ch. 5) Which (generally) are faster to create and manage: user-level threads or kernel-level threads. Why?

User-level threads. With kernel-level threads, thread management is done by the operating system, so their creation and scheduling requires kernel intervention.

Question 10: (Ch 6) Draw the Gantt chart assuming the following process burst times and SRTF scheduling with quantum 3, starting at time 0. P0 arrived before P1, and P1 before P2.

Process	*Burst time*	*Arrival time*
P0	7	before 0
P1	3	0
P2	3	1
P3	3	2

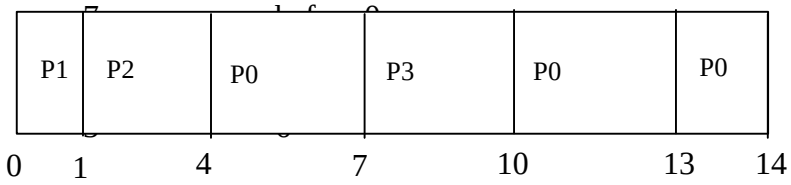

Question 11: (Ch. 6) What are the two circumstances under which CPU scheduling decisions may take place given nonpreemptive scheduling? (These involve process states and the short-term scheduler.)

1. When a process switches from the running state to the waiting state (e.g. I/O request)
2. When a process terminates

Question 12: (Ch. 7) What is the function of the V semaphore operation?

V atomically increments its counter (often, but not always, signifies release of a shared resource).

Question 13: (Ch. 8) Draw the wait-for graph corresponding to this resource allocation graph.

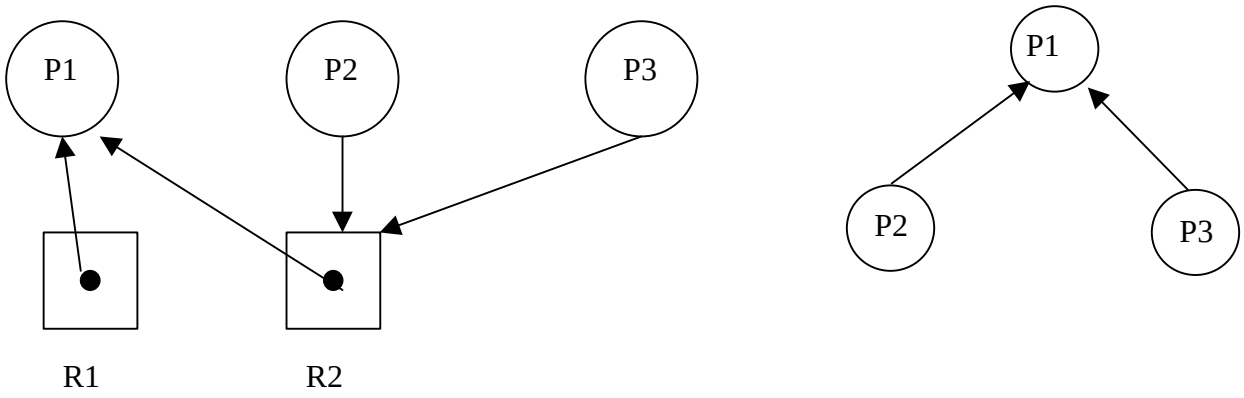

Question 14: (Ch. 8) What are the four necessary conditions for a deadlock to occur?

1. Mutual exclusion
2. Hold-and-wait
3. No preemption
4. Circular wait

Question 15: (Ch. 9) Define the term "external fragmentation".

External fragmentation exists when enough total memory space exists to satisfy a request, but it is not contiguous; storage is fragmented into a large number of small holes.

Question 16: (Ch. 9/10) What is the difference between a page and a frame?

The page is the unit of memory within the context of processes. The frame is the unit of memory within the context of the hardware. The page and the frame are the same size as each other.

Question 17: (Ch. 9) In translating to physical addresses, what is the role of the page number? What is the role of the page offset?

The page number serves as an offset into the page table. The frame's base address is found at the referenced location in the page table. The frame's base address is combined with the page offset to obtain the physical address.

Question 18: (Ch 10) What is Belady's anomaly?

For some page-replacement algorithms, the page-fault rate may increase as the number of allocated frames increases.

Question 19: (Ch. 10) Consider a demand-paging system with the following time-measured utilizations:

CPU utilization	20%
Paging disk utilization	97.7%
Other I/O devices utilization	5%

Would we expect the CPU utilization to increase if we increase the degree of multiprogramming? Why or why not?

No, we would not expect CPU utilization to increase. The system is thrashing. Increasing the degree of multiprogramming will only lead to additional thrashing.

Question 20: (Ch. 10) Given three frames, initially empty, and the following reference string

1 2 3 4 2 1 4 3 2 3 1

Show a[illegible] frames

1	2	3	4	2	1	4	3	2	3	1
1			4							1
	2						3			
		3			1			2		
y	y	y	y	n	y	n	y	y	n	y

Page fault rate: 8/11