

The exam is worth 100 points. There are 13 questions. Answer all questions.

You are to use ONLY the space provided, NOT the back of the page.

Please be brief, succinct, and to-the-point in your answers.

You will lose points for incorrect statements--even when your answer is correct!

(33 pts) **1) Circle the correct statements.**

- (1) Turnaround time of a process p is the actual clock time interval that p exists.
- (2) CPU utilization percentage of a process p is the percentage of time that CPU executes interrupts for p .
- (3) In the Variable Size Partitioning memory management scheme, compaction is used to eliminate external fragmentation.
- (4) Segmentation provides several virtual address spaces for different procedures, tables, library modules, and data structures.
- (5) MULTICS operating system used segmentation with paging.
- (6) A thread in a multi-threaded environment shares code, data, and stack space with other threads.
- (7) Threads may destroy each other's work if they use the same reentrant library procedure.
- (8) Implementing threads in the kernel space is slower than implementing them in the user address space because system calls by threads are time consuming.
- (9) SOLARIS operating system provides user level thread libraries.
- (10) Distributed systems allow incremental growth since computing power can be added in small increments.
- (11) CPU scheduling for user-level threads is problematic since timer interrupts are not available.
- (12) The client stub of RPCs is in the user address space.
- (13) For RPC executions, the client stub locates the server stub using a technique called static binding.
- (14) In UNIX and MS-DOS, files are unstructured byte sequences.
- (15) Unix uses fork, wait, exit, join for process creation, waiting, joining, and termination primitives, respectively.
- (16) Unix commands are simply filenames containing executable binaries.
- (17) Demand-paging based memory management algorithms occasionally use compile-time knowledge acquired by compilers to prefetch pages that may be needed during execution.
- (18) Computer systems use either interrupt-driven I/O or polling-based I/O.
- (19) Text segments in Unix are always read-only.
- (20) Shortest-CPU-burst-process-first CPU scheduling algorithm gives the best throughput, but also results in bad turnaround times for long processes.
- (21) Binding of code and data to memory can be done at compile-time, load-time, or execution-time.
- (22) In today's computers, paging algorithms are always implemented in hardware.

(30 pts) 2) Define **in few words** the terms

(1) page fault

Requested page not in main memory; has to be brought from disk

(2) locality phenomenon

For any process, at any given time, only a few pages (not necessarily adjacent in address space) are in active use.

(3) problems with the FCFS disk-head scheduling

Thrashing (frequent disk head seek movements over different cylinders).

(4) thrashing

(i) frequent page faults; (ii) Continuous disk head seek movements.

(5) two-level paging

Two levels of page tables. Given the address $(p, d1, d2)$, $Table1[p]$ points to $Table2$; $Table2[d1]$ points to a frame f . And, $(f, d2)$ is the physical address.

(6) reentrant program

Read-only code.

(7) trap

Software-generated interrupt.

(8) zero-capacity communication link

Communication link with no ability to buffer messages.

(9) Unix line discipline

A filter in the Unix I/O system that takes a raw character stream and produces a cooked character stream (for local line editing).

(10) "tight coupling" in parallel systems

Parallel systems where CPUs share clock and main memory.

(11) data migration in distributed systems

Migrating data transparently to speed up task completions.

(12) computation migration in distributed systems

Migrating processes transparently to speed up task completions.

(13) spooling

While executing one job, the OS (i) reads the next job into the job queue, and (ii) outputs the printout of completed jobs to printers.

(14) transport layer

ISO/OSI layer responsible from reliable end-to-end communication.

(15) network layer

ISO/OSI layer responsible from routing and congestion control.

(3 pts) **3)** List four conditions under which a process enters a ready queue.

(3 pts) **4)** What is disk-head scheduling and why do we need it?

(3 pts) **5)** State three necessary conditions for deadlocks.

(3 pts) **6)** What is the difference between interrupt-driven data transfer and direct-memory-access data transfer?

(3 pts) **7)** How does an operating system prevent a user program from getting stuck in an infinite loop, and never returning control to the operating system?

(3 pts) **8)** In UNIX, due to symbolic links, directory loops of files/directories exist. How does UNIX avoid an infinite loop while searching a directory for a file?

(3 pts) **10)** List three items in a Unix Process' swappable process image.

(4 pts) **11)** What are the main difficulties in designing an OS for a real-time environment?

(6 pts) **12)** You are the president of Dot.Cheapo.com Computronics, and your star hardware designer has suggested a brilliant idea: Implement segmentation, but let the least significant m bits of a virtual address be used to select the segment, and let the other bits determine the offset. Should the designer retire or get a raise? Why?

(6 pts) **13)** Prove that, in an environment where there is a single resource instance of any resource type, the resource allocation graph contains a cycle if and only if there is a deadlock in the system.