Universidad Carlos III de Madrid
Computer Science and Engineering Department
Operating Systems Course

RULES:

- The final grades and the review dates will be announced in Aula Global. The exam duration is **two hours and a half.** Books and notes are not allowed. A valid ID document will be necessary to submit the exam.

NAME: **GROUP:**

Exercise 1 (20 points). Quiz (20 minutes)

Answer the Quiz questions writing the letters of each correct answer. Three wrong answers deducts one correct answer. Non-answered questions don't reduce the grade (are ignored).

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15

1.- Select the correct answer about PCB:
A) All information referred to a process is stored in its PCB.
B) When a process execution is finalized, the PCB is released.
C) Data can be shared in the PCB.
D) All the previous answers are incorrect.

2.- What is SWAP used for in the lifecycle of a process?
A) To backup a process in secondary storage
B) To improve performance by sending some processes to swap area
C) To make the threads of a process easily configurable by means of an STD
D)To make the system boot faster

3.- Select the incorrect process mode
A) User
B) Privileged
C) Register
D) Kernel

4.- Which of the following is not an advantage of concurrent execution of processes ?
A) Concurrent execution improves application interactivity
B) Concurrent execution may require communication between processes
C) Concurrent execution improves CPU utilization
D) Concurrent execution accelerates computation execution

5.- Nonpreemptive kernels are not subject to race conditions, but preemptive kernels are. Why do we choose preemptive kernels then?
A) They can be more responsive than nonpreemptive.
B) There can be a smaller chance of being blocked in kernel mode for a long time.
C) It is more practical for real-time programming.
D) All of the above.

6.- Which memory management method/s avoid/s the need for external fragmentation and compaction?

A) Segmentation
B) Paging
C) Segmentation and paging
D) None of them

7.- Which of the following information is not stored in an i-node
A) Owner ID
B) Pointer to the current file read position
C) The creation date
D) The RWX bits

8.- Select the correct answer regarding the readers-writers problem.
A) Readers can have concurrency while only one writer can perform any operation when it is being executed
B) A reader can only be executed when there are no other readers that are waiting for execution.
C) If a reader performs a requests while a writer is being executed, it will be executed the first after the write operation is finished.
D) When readers have priority, read operations can enter the critical section as long as there is another read in the critical section and no writer requests, because as soon as a writer appears it will be executed before any other read.

9.- What is one problem with asynchronous cancellation?
A) The thread may not check often enough to see if it has been cancelled.
B) It is slower than synchronous cancellation.
C) The thread may be cancelled at an inopportune time (like when updating its data)
D) There are no risks with asynchronous cancellation.

10.- If in Round-Robin scheduling "q" equals infinity, we could say that this would be the same as:
A) Shortest job first
B) First to come first to serve
C) Both are correct
D) None of the above

11- For the fork() system call in Linux, which of the following is true?
A) The fork() process takes in an argument and returns a process ID.
B) The process ID of the existing parent process is 0.
C) A child process is created but it does not execute the fork() function.
D) When fork() is executed, an identical copy of its address space is created as well as its PCB.

12.- Which one is the correct order for computer structure?
A) hardware-services-kernel-system calls-system software
B) hardware-kernel-services-system calls-system software
C) hardware-kernel-system calls-services-system software
D) hardware-kernel-system calls-system software-services

13.- In case of mutual exclusion, how many processes will be allowed to access information?
A) 1
B) 0
C)It is impossible to know
D) All the processes that take part in the mutual exclusion

14.- The use of a linked list for managing file blocks …
A) Generates a lot of disk fragmentation
B) Reduces the disk fragmentation

C) It is the most efficient approach
D) It is based on a b-tree structure

15.- On which situations Operating System is activated ?
A) A service request from a process.
B) An interrupt (peripheral or clock).
C) Hardware exception.
D) All of the above

Exercise 2 (20 points).

Given a hard disk with a Unix-like filesystem using an i-node structure with a data block size of 1024Bytes. The filesystems uses 10 direct pointers to data blocks, 1 single indirect pointer, 1 double indirect pointer and 1 triple indirect pointer.

The filesystem contains the information shown next:

I-node table:

I-node number	2	3	4	5	6
Type	Directory	Directory	File	Symbolic link	
Link counter	3	2	2	1	
Data block address	11	12	13	14	
…….					

Data block:

Block number	11	12	13	14	15
Content	. 2	. 3	File data	/tests/data	
	.. 2	.. 2			
	tests 3	data 4			
		report 4			
		chart 5			

Answer the following questions:

a) Obtain the maximum file size taking into account only i-node configuration.

b) Obtain how many disk blocks are used by a 10 Mbyte file (taking into account the overall space with both the data blocks and the addresses).

c) Show what are the results in the i-node and data block configurations after the following operations (use the tables provided)

1. rm /tests/data
2. rm /tests/report
3. rm /tests/chart
4. mkdir /tests/Dir1

Exercise 3 (30 points).

Given the following concurrent program specification:

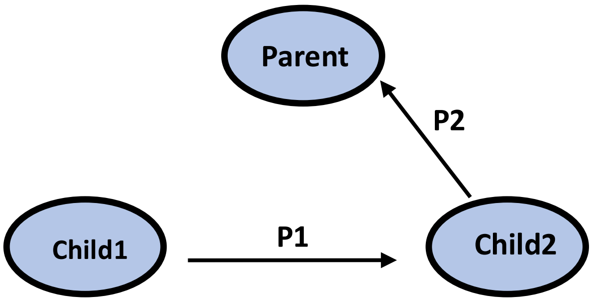

- There are three processes: a parent (called Parent) and 2 child processes called Child1 and Child2
- There are two pipes:
  - P1 that links Child1 with Child2
  - P2 that links Child2 with Parent
- Child 1 executes the Linux system command (command prompt) that lists in the standard output the files and directories existing in the current directory. The command output is sent to Child2 by means of pipe P1.
- Child 2 executes the Linux system command (command prompt) that displays in the standard output the number of characters received from the standard input. That is, given an input string, the command returns the number of characters on it. This command has to be employed using P1's output as standard input and P2's input as standard output.
- Parent process reads and displays on screen the output of P2. This output will be the number of characters (counted by Child2) of the list (created by Child1) that contains the existing files and directories in the current directory.

Complete the following tasks:

1. Write the Linux command names that are used by Child1 and Child2.

2. Implement the problem in a C-language code using the standard system calls (*fork(), pipe(), execlp()*, etc.). **Justify how the code works by means of comments in the code.**

3. Show the changes that have to be done in the code to make the parent process receive the error messages of the command executed by Child2. Note: it is only necessary to consider Child2's command (the error messages of Child1's command are not considered and have to be discarded).

Exercise 4 (30 points).

Write a program that executes in parallel three threads called A, B and C.

- A thread has three code blocks (a1, a2 and a3).
- B thread has four code blocks (b1, b2, b3 and b4)
- C thread has three code blocks (c1, c2 and c3).

a) Code a function called Write() that displays on the screen a text (received as input parameter) and after that performs a random delay between 3 seconds. Note: use rand() function that generates a random number with a value between 0 and MAX_INT.

b) Write a main program that creates the threads A, B and C. Each thread will use the Write() function to show the code block that is being executed. This operation is performed **without** any concurrency control.

c) Modify the previous program using as many semaphores as necessary to enforce the following execution order:

- c1 can only be executed after a1 completion.
- a2 can only be executed after b1 completion
- c2 can only be executed after b2 completion
- b3 can only be executed after a2 completion
- a3 can only be executed after c2 completion
- b4 has to be the last data block executed.

Note: show clearly and justify how to use the system calls that employed to manage the semaphores.

Exercise 2: Fill the following tables for section 3.

1) **I-node table after rm /tests/data:**

I-node number	2	3	4	5	6
Type					
Link counter					
Data block address					

1) **Data blocks after rm /tests/data:**

Block number	11	12	13	14	15
Content					

2) **I-node table after rm /tests/report:**

I-node number	2	3	4	5	6
Type					
Link counter					
Data block address					
…….					

2) Data blocks after rm /tests/report:

Block number					
Content					

3) I-node table after rm /tests/chart:

I-node number	2	3	4	5	6
Type					
Link counter					
Data block address					
…….					

3) Data blocks after rm /tests/charts:

Block number	11	12	13	14	15
Content					

4) I-node table after mkdir /tests/Dir1:

I-node number	2	3	4	5	6
Type					
Link counter					
Data block address					
…….					

4) Data blocks after mkdir /tests/Dir1:

Block number	11	12	13	14	15
Content					