Philadelphia University Faculty of Engineering Department of Computer Engineering		First Semester 2015/2016 Date:- 15/04/2015 Allowed time:- 2 Hours

Operating Systems (630422) Final Exam

Student Name: - ……..... ID: - ……

Question 1: chose the correct answer for the followings.

1- A process may transition to the Ready state by which of the following actions?
A) Completion of an I/O event B) Awaiting its turn on the CPU
C) Newly-admitted process D) All of the above

2- Which of the followings involves distributing tasks across multiple computing cores.
A) Concurrency B) Task parallelism C) Data parallelism D) Parallelism

3- One of the followings allows a thread to run on only one processor.
A) Processor affinity B) Processor set C) NUMA D) Load balancing

4- A(n) _____ refers to where a process is accessing/updating shared data.
A) critical section B) entry section C) mutex D) test-and-set

5- Which of the following data structures in the banker's algorithm is a vector of length *m*, where *m* is the number of resource types?
A) Need B) Allocation C) Max D) Available

6- One of the followings is the dynamic storage-allocation algorithm which results in the smallest leftover hole in memory.
A) First fit B) Best fit C) Worst fit D) None of the above

7- Consider a 32-bit address for a two-level paging system with an 8 KB page size. The outer page table has 1024 entries. How many bits are used to represent the second-level page table?
A) 10 B) 8 C) 12 D) 9

8- In the enhanced second chance algorithm, which of the following ordered pairs represents a page that would be the best choice for replacement?
A) (0,0) B) (0,1) C) (1,0) D) (1,1)

9- The _____ allocation algorithm allocates available memory to each process according to its size.
A) equal B) global C) proportional D) slab

10- In an environment where several processes may open the same file at the same time, ____.
A) the operating system typically uses only one internal table to keep track of open files
B) the operating system typically uses two internal tables called the system-wide and per-disk tables to keep track of open files
C) the operating system typically uses three internal tables called the system-wide, per-disk, and per-partition tables to keep track of open files
D) the operating system typically uses two internal tables called the system-wide and per-process tables to keep track of open files

11- Order the following file system layers in order of lowest level to highest level.
[1] I/O control [2] logical file system [3] basic file system
[4] file-organization module [5] devices

A) 1, 3, 5, 4, 2 B) 5, 1, 3, 2, 4 C) 1, 5, 3, 4, 2 D) 5, 1, 3, 4, 2

12- Which of the following is the simplest method for implementing a directory?
A) tree data structure B) linear list C) hash table D) nonlinear list

Question 2:- Ordinarily the `exec()` system call follows the `fork()`. Explain what would happen if a programmer place the call to `exec()` before the call to `fork()`.

Ans: Because `exec()` overwrites the process `fork()` will never execute and no new processes would be created. Rather, the program specified in the parameter to `exec()` would be run instead.

Question 3:- Distinguish between data and task parallelism.

Ans: Data parallelism involves distributing subsets of the same data across multiple computing cores and performing the same operation on each core. Task parallelism involves distributing tasks across the different computing cores where each task is performing a unique operation.

Question 4:-Explain the difference between response time and turnaround time.

Ans: 1- Turnaround time is the sum of the periods that a process is spent waiting to get into memory, waiting in the ready queue, executing on the CPU, and doing I/O. Turnaround time essentially measures the amount of time it takes to execute a process.
2- Response time is a measure of the time that elapses between a request and the first response

Question 5:- Write two short methods that implement the simple semaphore `wait()` and `signal()` operations on global variable `S`.

Ans:

```
wait (S) {
  while (S <= 0);
  S--;
}

signal (S) {
  S++;
}
```

Question 6:- What is the difference between deadlock prevention and deadlock avoidance?

Ans: 1- Deadlock prevention is a set of methods for ensuring that at least one of the necessary conditions for deadlock cannot hold.
2- Deadlock avoidance requires that the operating system be given, in advance, additional information concerning which resources a process will request and use during its lifetime.

Question 7:- Explain how paging is implemented and how CPU can access data within a page.

Ans: Physical memory is broken up into fixed-sized blocks called frames while logical memory is broken into equal-sized blocks called pages. Whenever the CPU generates a logical address, the page number and offset into that page is used, in conjunction with a page table, to map the request to a location in physical memory.

Question 8:- Explain the sequence of events that happens when a page-fault occurs.

Ans: 1- the memory reference is checked for validity. In the case of an invalid request, the program will be terminated.
2- If the request was valid, a free frame is located.
3- A disk operation is then scheduled to read the page into the frame just found
4- update the page table
5- restart the instruction that was interrupted because of the page fault

Question 9:- Why do all file systems suffer from internal fragmentation?

Ans: Disk space is always allocated in fixed sized blocks. Whenever a file is written to disk, it usually does not fit exactly within an integer number of blocks so that a portion of a block is wasted when storing the file onto the device.

Question 10:- What are the problems associated with linked allocation of disk space routines?

Ans: 1- can be used effectively only for sequential-access files.
2- the space required for the pointers in each block.
3- decreased reliability due to lost or damaged pointers.

Question 11:- Suppose we have the following page accesses: 1 2 3 5 4 2 3 5 4 1 2 5 1 1 3 1 4 and that there are three frames within our system. what is the number of page faults for the given reference string using Optimal Replacement, FIFO and LRU algorithms?

Optimal Replacement:-

1	1	1	5	4			4		1					1		4
	2	2	2	2			2		2					3		3
		3	3	3			5		5					5		5

Number of page faults=9

FIFO

1	1	1	5	5	5	3	3	3	1	1	1			3	3	3
	2	2	2	4	4	4	5	5	5	2	2			2	1	1
		3	3	3	2	2	2	4	4	4	5			5	5	4

Number of page faults=15

LRU

1	1	1	5	5	5	3	3	3	1	1	1			1		1
	2	2	2	4	4	4	5	5	5	2	2			3		3
		3	3	3	2	2	2	4	4	4	5			5		4

Number of page faults=14

Question 12:- Assume a system has a TLB hit ratio of 90%. It requires 15 nanoseconds to access the TLB, and 85 nanoseconds to access main memory. What is the effective memory access time in nanoseconds for this system?

Ans: when hit the access time is 15 ns (to access TLB) + 85ns (to access the page in memory)=100ns
When miss the access time is 15ns (to access TLB)+85ns (to access page table in memory)+85ns (to access the page in memory) = 185ns
EAT=hit ratio × access time(when hit) + miss ratio × access time(when miss)
EAT=0.9×100 + 0.1 × 185 = 90 + 18.5 = 108.5 ns

Good Luck

Sultan M. Al-Rushdan