K i n g S a u d U n i v e r s i t y
College of Computer and Information Sciences
Computer Science Department

Course Code:	CSC 227
Course Title:	Operating Systems
Semester:	Spring 2015
Exercises Cover Sheet:	**Final Exam**

Duration: 3hours (180 min)

Student Name:	
Student ID:	
Student Section No.	

Tick the Relevant	Computer Science B.Sc. Program ABET Student Outcomes	Question No. Relevant Is Hyperlinked	Covering %
	a) Apply knowledge of computing and mathematics appropriate to the discipline;		
	b) Analyze a problem, and identify and define the computing requirements appropriate to its solution		
	c) Design, implement and evaluate a computer-based system, process, component, or program to meet desired needs;		
	d) Function effectively on teams to accomplish a common goal;		
	e) Understanding of professional, ethical, legal, security, and social issues and responsibilities;		
	f) Communicate effectively with a range of audiences;		
	g) Analyze the local and global impact of computing on individuals, organizations and society;		
	h) Recognition of the need for, and an ability to engage in, continuing professional development;		
	i) Use current techniques, skills, and tools necessary for computing practices.		
	j) Apply mathematical foundations, algorithmic principles, and computer science theory in the modeling and design of computer-based systems in a way that demonstrates comprehension of the tradeoffs involved in design choices;		
	k) Apply design and development principles in the construction of software systems of varying complexity;		

King Saud University
College Of Computer and Information Sciences
Computer Science Department
CSC 227: Operating Systems

Total Marks: 40 **Time: 08:00am – 10:00am (120 minutes)**
Spring 2015-16 **Name:** ……
Final Exam **ID#:** ……
Date: 19-May-2016 **Section#:**…… **or Teacher Name**: ……

Instructions:

- This exam has 11 pages (including title page) and has 6 questions.
- Do not use pencil.
- Write clearly and neatly.

Question 1.[10 marks] Circle your answer and copy it to the table ONLY ONE ANSWER (the best answer).
Copy your answer for question 1-1 to 1-15 in the table on page2. ONLY THAT TABLE WILL BE GRADED.

1.	If a thread attempts a wait() operation when a counting semaphore has been decremented to zero, the thread
a.	**must wait until a resource is returned to the pool by a signal() operation**
b.	obtains a resource from the pool and the semaphore remains equal to 0
c.	returns a resource to the pool and increments the semaphore to 1
d.	obtains a resource from the pool and decrements the semaphore to 1

2.	A process is a program in execution, it exists in main memory and it may be:
a.	Either OS process or User process
b.	Either I/O bound process or CPU bound process
c.	Either Independent process or Cooperating process
d.	**All the above**

3.	Which of the following statements about semaphores is true?
a.	If several threads attempt a wait(S) operation simultaneously, only one thread should be allowed to proceed.
b.	wait() and signal() operations should be indivisible operations.
c.	A semaphore implementation should guarantee that threads do not suffer indefinite postponement.
d.	**all of the above**

4.	A critical section is a program segment
a.	which should run in a certain specified amount of time.
b.	which avoids deadlocks.
c.	**where shared resources are accessed**
d.	which must be enclosed by a pair of semaphore operations, P and V.

5.	A scheduling method in which a processes can be interrupted whether they have completed their current task or not is
a.	dynamic
b.	**preemptive**
c.	non-preemptive
d.	static

6.	A frame is a portion of
a.	registers memory
b.	virtual Memory
c.	logical Memory
d.	**physical memory**

7.	Which of the following statements about critical sections is false?
a.	Only one thread at a time can execute the instructions in its critical section for a particular resource.
b.	If one thread is already in its critical section, another thread must wait for the executing thread to exit its critical section before continuing.
c.	Once a thread has exited its critical section, a waiting thread may enter its critical section.
d.	**All threads must wait whenever any critical section is occupied.**

8.	Semaphores can be used for each of the following purposes except:
a.	**to notify the process that an event has occurred**
b.	to prevent shared variables from getting corrupted
c.	to synchronize two or more concurrent threads
d.	to synchronize two or more concurrent threads

9.	A Foreground processes is:
a.	Detached from the terminal it was started
b.	Like a background process runs without user interaction
c.	**Able to receive input and return output from/to the user**
d.	All the above

10.	When a process creates a new process using the fork() operation, which of the following state is shared between the parent process and the child process?
a.	Stack
b.	Heap
c.	**Shared memory segments**
d.	All the above

11.	In a paging system, a frame
a.	has no relation to the page size
b.	**is of the same size as that of the page**
c.	is always smaller than the page
d.	is always larger than a page

12.	In segmentation, each address is specified by:
a.	**segment number and offset**
b.	page number and offset
c.	segment number
d.	page number

13.	A multilevel page table results in:
a.	Less number of frames
b.	**reduced size of page table**
c.	it is required by the translation look aside buffer
d.	reduced number of pages

14.	After synchronous I/O starts, control returns to user program
a.	Immediately after I/O start
b.	After the end of its quantum
c.	after I/O completion
d.	According to DMA speed

15.	Which of these activities is NOT accomplished by the Operating Systems
a.	Creating and deleting processes
b.	**Generating interrupts**
c.	Providing mechanisms for deadlock handling
d.	Mapping files onto secondary storage

16.	When the operating system is divided into a number of levels, each one built on top of lower other, this is called:
a.	Modular approach
b.	Monolithic approach
c.	**Layered approach**
d.	Micro-kernel approach

17.	Which of the following instructions should be privileged?
a.	Set value of timer
b.	**Switch from user to kernel mode**
c.	Read the clock
d.	Create a new thread

18.	Dynamic linking is the linking method done at:
a.	Compilation time
b.	Loading time
c.	**Execution time**
d.	Proper time

19.	Process execution consists of a cycle of:
a.	IO burst then CPU burst
b.	**CPU burst then I/O burst**
c.	CPU burst then CPU burst then IO burst
d.	None of the above

20.	CPU scheduling decisions may take place when a process
a.	Switches from running to waiting state
b.	Switches from running to ready state
c.	Switches from waiting to ready
d.	**All of the above**

1.	2.	3.	4.	5.	6.	7.	8.	9.	10.
a	**d**	**d**	**c**	**b**	**d**	**d**	**a**	**c**	**c**

11.	12.	13.	14.	15.	16	17	18	19	20
b	**a**	**b**	**c**	**b**	**c**	**b**	**c**	**b**	**d**

Question 2. [5 marks]

2-a) [1 mark] How does aging solve the starvation problem in priority based scheduling system?

It works by increasing the priority of a process as time passes and will eventually be the highest priority request (after it has waited long enough)

2-b) [1 mark] Assuming special hardware is available to support swapping, when is it appropriate to use execution-time binding?

If the process can be moved during its execution from one memory segment to another, then binding must be delayed until execution time.

2-c) [1 mark] How do base register and limit register help in protecting against illegal memory access?

Any access to memory is checked against these two registers. The address should be greater or equal than the base address and smaller than (Base register + limit register).

2-d) [1 mark] How does shared-pages help in saving memory?

By having only one copy of the code in physical memory.

2-e) [1 mark] What is the use of Translation Look-aside Buffer (TLB)?

With standard paging scheme two memory accesses are needed to access a byte (one for the page-table entry, one for the byte). TLB is small, fast lookup hardware cache that stores recent translations of virtual memory to physical addresses for faster retrieval.

Question 3. [7 marks]

3-a) [1 mark] Considering the following diagram of a process state:

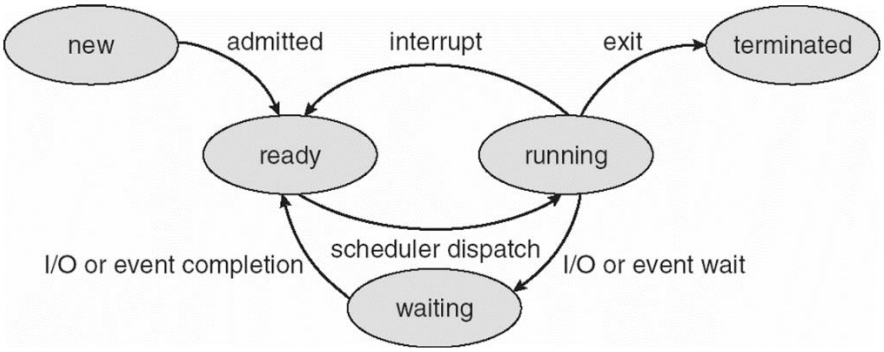

Give three reasons that cause a process to move from the waiting/blocked state.

1.	**I/O completion**
2.	**Termination of a child process**
3.	**A signal/message from another process**

3-b) [2.5 marks] Give at least five components of the process's PCB.

1.	**Process state**
2.	**PID (Process ID)**
3.	**PC (Program counter)**
4.	**Contents of the processor's registers**
5.	**Memory limits**
6.	**List of I/O devices allocated to the process**
7.	**List of open files** **etc. (e.g., priority of the process, amount of CPU time used, time limits, …)**

3-c) [1 mark] Why should a web server not run as a single-threaded process?

For a web server that runs as a single-threaded process, only one client can be serviced at a time. This could result in potentially enormous wait times for a busy server.

3-d) [1.5 marks] What is a semaphore? Explain a binary semaphore with the help of an example? (4)

A semaphore is a synchronization tool that provides a general-purpose solution to controlling access to critical sections.

A semaphore is an abstract data type (ADT) that defines a nonnegative integer variable which, apart from initialization, is accessed only through two standard operations: wait and signal. The classical definition of wait in pseudo code is

```
wait(S){
while(S<=0) ; // do nothing
```

```
S--;
}
The classical definitions of signal in pseudocode is
signal(S){
S++;
}
```

A binary semaphore is one that only takes the values 0 and 1. These semaphores are used to implement mutual exclusion.

3-e) [1 mark] Mention at least two parameters needed to define a multilevel-feedback-queue scheduler.

Number of queues

Scheduling algorithms for each queue

Method used to determine when to upgrade a process

Method used to determine when to demote a process

Method used to determine which queue a process will enter when that process needs service

Question 4.[8 marks]

4-a) Consider the following set of processes, with the length of the CPU burst given in milliseconds:

Process	Burst	Priority	Arrival Time
P1	8	4	0
P2	6	1	2
P3	1	2	2
P4	9	2	1
P5	3	3	3

i) [2 marks] Find out turnaround time for each process using the Non-Preemptive Priority Scheduling (a smaller priority number implies a higher priority)? **Two possible answers.**

TT1	**8**
TT2	**14**
TT3	**15**
TT4	**24**
TT5	**27**

TT1	**8**
TT2	**14**
TT3	**24**
TT4	**23**
TT5	**27**

ii) [2 marks] Find out turnaround time for each process using Preemptive Priority Scheduling (a smaller priority number implies a higher priority)? **Two possible answers.**

TT1	**27**
TT2	**8**
TT3	**9**
TT4	**17**
TT5	**20**

TT1	**27**
TT2	**8**
TT3	**17**
TT4	**16**
TT5	**20**

iii) [2 marks] Find out turnaround time for each process using Non-Preemptive Shortest Job First Time?

TT1	**8**
TT2	**18**
TT3	**9**
TT4	**27**
TT5	**12**

iv) [2 marks] Find out turnaround time for each processes using Preemptive Shortest Job First Time (or Shortest-Remaining-Time-First)? **Two possible answers.**

TT1	**12**
TT2	**18**
TT3	**3**
TT4	**27**
TT5	**6**

TT1	**18**
TT2	**12**
TT3	**3**
TT4	**27**
TT5	**6**

Question 5. [6 marks]

5-a) [3 marks] Given eight memory holes of 10KB, 4KB, 20KB, 18KB, 7KB, 9KB, 12KB, and 15KB (in order), how would the first-fit, best-fit, and worst-fit algorithms will place P1 (requires 12KB), P2 (requires 10KB), P3 (requires 7KB) and P4 (requires 9KB). Use memory maps given below to answer the question.

First fit

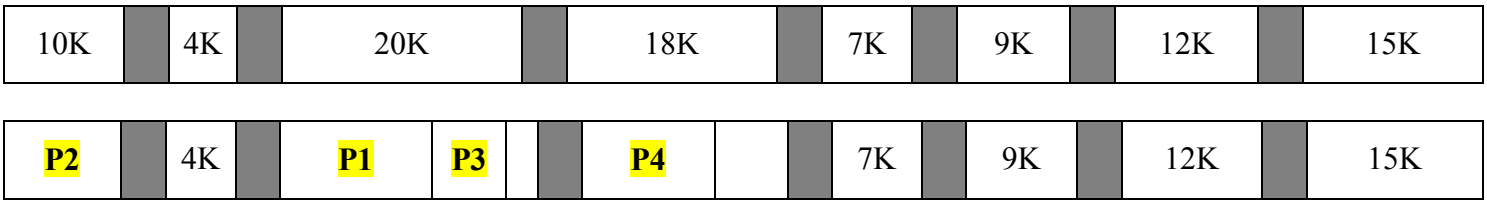

Best fit

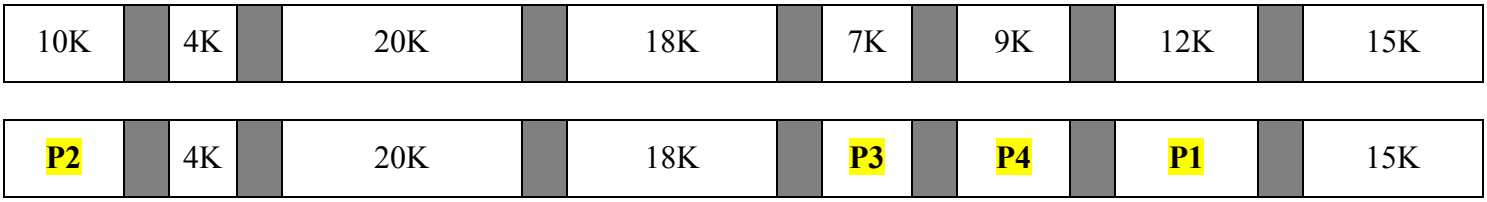

Worst fit

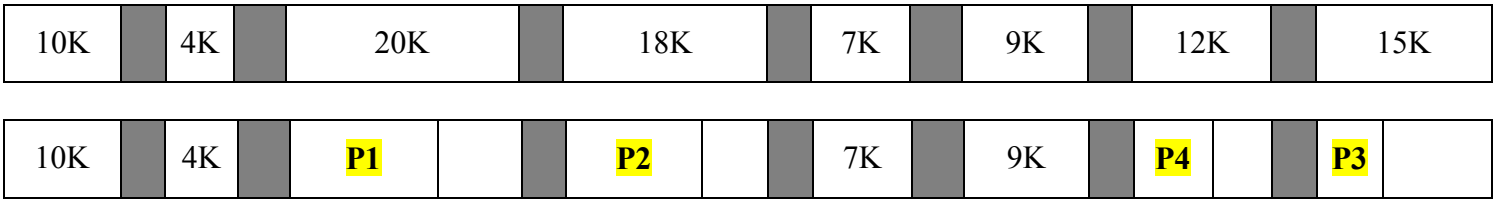

5-b) [3 marks] Given a logical address filed with the following format:

Page #	Page Offset
16-bits	8-bits

a) What is the maximum number of pages?
b) What is the maximum size of page table (PT)? Assume each page entry requires 4 bytes.
c) Consider that PT is stored in physical memory, the memory access time is 100ns and the access time to associative memory is 10ns. If the hit ratio is 80%, what is the effective access time (EAT)?

a) 2^{16} pages
b) $2^{16} \times 4 = 2^{16} \times 2^2 = 2^{18}$
c) EAT = (access time for associative memory for page# + access time for memory) × hit ratio + (access time for associative memory for page# + access time for memory for PT + access time for memory) × (1 - hit ratio)

$$(110 \times 0.80) + (10 + 100 + 100) \times (0.20) = 130$$

Question 6.[4 marks]

6-a) [1 mark] List two different definitions of the operating system.

i) **OS is a resource allocator**

a) **Manages all resources**

b) **Decides between conflicting requests for efficient and fair resource use**

ii) **OS is a control program**

a) **Controls execution of programs to prevent errors and improper use of the computer**

iii) **"Everything a vendor ships when you order an operating system" is good approximation**

6-b) [1 mark] What is Direct Memory Access (DMA)?

A hardware device controller that can transfer a block of data data to/from any device buffer its own buffer storage without the intervention of the CPU. Only one interrupt is generated when the block transfer is finished rather than one interrupt by byte generated for low-speed devices.

6-c) [1 mark] Micro-kernel moves most of the kernel into user space. List three advantages of Micro-kernels:

Easier to extend a microkernel

Easier to port the operating system to new architectures

More reliable (less code is running in kernel mode)

More secure

6-d) [1 mark] Security and Protection are some of the major roles of an operating system. What is the difference between Security and protection?

- **Protection involves ensuring that all access to system resources is controlled**
- **Security of the system from outsiders requires user authentication, extends to defending external I/O devices from invalid access attempts**