

Comp 104: Operating Systems Concepts

Revision Lectures
(separate questions and answers)

1

Today

- Here are a sample of questions that could appear in the exam
 - Please LET ME KNOW if there are particular subjects you want to know about???

2

Notes on the Exam

- Exam will be a multiple choice paper
 - Duration: 2 hours
 - Number of Questions: 40
- Each question will have five statements which involve identifying:
 - the **one correct** statement
 - the **correct combination** of valid statements

3

Question 1

- If you type 'cat prog.c' at a UNIX command prompt, which of the following sequences of system calls would be invoked?
 - a) The shell makes an exec() call
 - b) The shell calls fork(); the child process calls exec() and the parent calls wait()
 - c) The shell calls fork(); the child calls wait() and the parent calls exec()
 - d) The shell calls exec() and then wait() and then fork()
 - e) The shell calls wait() then fork(), creating a child which calls exec()

4

Question 2

- A running process makes a system call to read data from a file. Which process state should it enter next?
 - New
 - Ready
 - Running
 - Blocked
 - Terminated

5

Question 3

- In calculating the formula $ut + \frac{1}{2}at^2$ using maximal concurrency, which of the operations might be computed in parallel?
 - $u*t$; $a/2$; $t*t$
 - $u*t$; $t+\frac{1}{2}$; $a*t$
 - $u+a$; $t*t$
 - $u+a$; $t*t$; $\frac{1}{2}$
 - no parallelism is possible for this formula

6

Question 4

- A Java object called 'helper' contains the two methods opposite, where num is an integer variable that is private to helper. Its value is initially 100.
- One thread makes the call
 - helper.addone();
- At the same time, another thread makes the call
 - helper.subone();
- What value will num have afterwards?

```
public void addone() {
    num = num + 1;
}

public void subone() {
    num = num - 1;
}
```

- 100
- 99
- 101
- either 99 or 101, but not 100
- the value of num is undefined

7

Question 5

- Consider the following situation regarding two processes (A and B), and two resources (X and Y):
 - Process A is granted resource X and then requests resource Y.
 - Process B is granted resource Y and then requests resource X.
- Which of the following is (are) true about the potential for deadlock?
 - Deadlock can be avoided by sharing resource Y between the two processes
 - Deadlock can be avoided by taking resource X away from process A
 - Deadlock can be avoided by process B voluntarily giving up its control of resource Y
 - I only
 - I and II only
 - I and III only
 - II and III only
 - I, II and III

8

Question 6

- A starvation-free job-scheduling policy guarantees that no job waits indefinitely for service. Which of the following job-scheduling policies is starvation-free?
 - a) Round-robin
 - b) Priority queuing
 - c) Shortest job first
 - d) Youngest job first
 - e) None of the above

9

Question 7

- A program is split into 3 segments. The segment table contains the following information:

segment	datum	limit
0	1700	5500
1	5600	8100
2	8300	9985

- where 'limit' is the physical address following the end of the segment, and instructions take the form opcode segment, offset
- If the program executes
 - **LOAD 1, 135**
- what physical address is accessed?
 - a) 1835
 - b) 5735
 - c) 8435
 - d) 8235
 - e) 5635

10

Question 8

- Process A and process B both share the same code segment S. Which of the following statements is (are) true?
 - I. An entry for S appears in both segment tables
 - II. The segment code must be re-entrant
 - III. The segment code must be recursive
 - a) I only
 - b) II only
 - c) I and II
 - d) I and III
 - e) I, II and III

11

Question 9

- In a paged memory system, why are page sizes invariably a power of 2?
 - a) Because computer memory is usually a multiple of 1K, which is a power of 2.
 - b) Because pages have to begin at address boundaries that are even.
 - c) Because virtual address spaces are usually a power of 2 in size
 - d) Because it simplifies indexing of the page table and the calculation of the page offset.
 - e) Because most data types occupy even numbers of bytes.

12

Question 10

- A computer uses 16-bit addressing. Its page size is 512 bytes. What is the maximum number of entries that the page table must be capable of holding?

- 16
- 64
- 128
- 256
- 512

13

Answers 1-10

- Answer: b** - The shell calls `fork()`, the child process calls `exec()` and the parent calls `wait()`. The shell is just another process that can take a string as standard input, look for the program referenced by the string, and then run this program. Unless the program is put in the background, the shell will wait until the program has finished.
- Answer: d** - Blocked; it may take some time before the file system can read the file (e.g. on a networked file store), so the process is blocked until the data is available.
- Answer: a** - u^2 , a^2 ; and t^4 - i.e. only those parts of the formula that have no dependencies on other parts of the formula can be run concurrently. Think how the formula could be written in 3-code...
- Answer: d** - either 99 or 101, but not 100 - if the two threads are run simultaneously, then it depends on the order in which the threads are executed by the ready queue. However, as "num" is not protected by a semaphore, its final value could be either value.
- Answer: a** - I, II and III - as all three options will avoid exclusive ownership of the resources.
- Answer: a** - Round Robin - this gives all processes equal access to the processor. The other techniques each select some "types" of processes to others (e.g. short processes, high priority processes etc).
- Answer: b** - 5600 (from segment 1) + 135 (offset)
- Answer: c** - I and II - as the segment is shared, both processes need to index it (i.e. include an entry in their respective segment tables, and, the code must not be changed by its use (i.e. it should be re-entrant). Recursion is irrelevant to this issue.
- Answer: d** - Because it simplifies indexing of the page table and the calculation of the page offset.
- Answer: c** - 128; 9 bits are used for addressing the 512 bytes in each page, so the remaining 7 bits are used to address the pages.

14

Question 11

- Which of the following programming constructs tend to contribute to the phenomenon expressed in the Principle of Locality?

- Iteration (e.g. FOR and WHILE loops)
- Selection (e.g. IF-statements)
- Recursion

- I only
- III only
- I and III only
- II and III only
- I and II only

15

Question 12

- Consider the following sequence of page references in a paged memory management system:

page	p	q	r	q	q	q	p	r	r	q	
time	0	1	2	3	4	5	6	7	8	9	10

- What is the working set expressed as $W(3,4)$?

- q
- r
- qr
- pq
- pqr

16

Question 13

- Consider the following sequence of page references in a paged memory management system:

page	p q r q q q p r r q
time	0 1 2 3 4 5 6 7 8 9 10

- What would be the predicted working set expressed as $W(10,3)$?
 - q
 - r
 - qr
 - pq
 - pqr

17

Question 14

- Consider the following sequence of page references in a paged memory management system:

page	p q r q q q p r r q
time	0 1 2 3 4 5 6 7 8 9 10

- Page s arrives at time 10. Which of the following policies suggests we should throw out page p to make room for s?
 - LRU
 - LFU
 - FIFO
 - I only
 - III only
 - I and II
 - I and III
 - I, II and III

18

Question 15

- When should a deleted file *not* be garbage collected?
 - When there are multiple links to it.
 - When the file contains program code.
 - When there is only one copy of the file.
 - When the file is a system file rather than a user file.
 - When the file contains encrypted data.

19

Question 16

- Which of the following is true about linked filestore allocation versus contiguous allocation?
 - Linked is slower for both sequential and direct access.
 - Linked is faster for both sequential and direct access.
 - Linked is faster for sequential access, but slower for direct access.
 - Linked is slower for sequential access, but faster for direct access.
 - The performance for linked and contiguous is roughly the same for both forms of access.

20

Question 17

- Consider the following grammar, where S, A and B are non-terminals, and a and b are terminals:

S ::= AB
 A ::= a
 A ::= BaB
 B ::= bbA

- Which of the following is FALSE?
 - The length of every string derived from S is even.
 - No string derived from S has an odd number of consecutive b's.
 - No string derived from S has three consecutive a's.
 - No string derived from S has four consecutive b's.
 - Every string derived from S has at least as many b's as a's.

21

Question 18

- A BNF grammar includes the following statement:

```
<statement> ::= <iden> := ( <expr> );
```

- What kind of message would be produced by the following line of code?

```
a := (2 + b);
```

- A Syntax Error.
- A Static Semantic Error.
- A Dynamic Semantic Error.
- A Warning, rather than an error.
- None of the above.

22

Question 19

- If the array x contains 20 ints, as defined by the following declaration:

```
int x[] = new int[20];
```

- What kind of message would be generated by the following line of code?

```
a := 22;  
val := x[a];
```

- A Syntax Error.
- A Static Semantic Error.
- A Dynamic Semantic Error.
- A Warning, rather than an error.
- None of the above.

23

Question 20

- Which of the following is usually NOT represented in a subroutine's activation record frame for a stack-based programming language?

- Values of locally declared variables
- A heap area
- The return address
- A pointer to the calling activation record
- Parameter values passed to the subroutine

24

Answers 11-20

11. **Answer:** c - I and III only, as statements in iterated loops (or called through recursion) may be called several times before other code is called. Statements in an if construct tend to be called only once within a given block (unless of course the code is also within an iterated loop).
12. **Answer:** d - pq, as these are the pages used for 4 time segments after time 3.
13. **Answer:** c - qr, as the predicted working set is approximately equivalent to the working set for the previous epoch, i.e. the 3 time segments before time t=10.
14. **Answer:** e - I, II and III.
15. **Answer:** a - When there are multiple links to it. Links allow a file to appear in multiple places around a file system, without having to maintain or synchronise several copies. Thus, if the link count is greater than one, then it should be decremented, but the file left. This way, other references to the file can still access the file.
16. **Answer:** a - Linked is slower for both sequential and direct access. With contiguous allocation, the OS knows the start address and can read forward (i.e. offset from the start address) for direct access, or use large reads from a single area of the disc for sequential access. Linked allocated files have their blocks scattered around the disc, so each block needs to be read to find the link to locate the subsequent block.
17. **Answer:** c - No string derived from S has three consecutive a's
18. **Answer:** a - A syntax error - all the tokens are valid, but the close parenthesis is missing, resulting in an error in the grammar.
19. **Answer:** c - A dynamic semantic error - the value of a would cause an array out of bounds error
20. **Answer:** b - A heap area, as this maintains memory storage for the whole program, and not for a subroutine. Instances of classes, allocated memory and container data structures are normally found in the heap.

25

Question 21

- In generating code for the following assignment

$$X := X / 4$$
- an optimising compiler might use strength reduction and generate code equivalent to which of the following?
 - $X := X \ll 4$
 - $X := X \ll 2$
 - $X := X \gg 2$
 - $X := X \gg 4$
 - None of the above.

26

Question 22

- Concerning compilation, which of the following is NOT a method for symbol table access?
 - Sequential lookup
 - Direct lookup
 - Binary chop
 - Hash addressing
 - Hash chaining

27

Question 23

- If the symbol table for a compiler is size 4096, how many comparisons on average need to be made when performing a lookup using the binary chop method?
 - 2
 - 11
 - 12
 - 16
 - 31

28

Question 24

- Which of the following postfix expressions is equivalent to the following expression?

$a*b - c/d$

- a) $a b c d * - /$
- b) $a b * - c d /$
- c) $a b c d / - *$
- d) $a b * c d / -$
- e) $a b c * - d /$

29

Question 25

- Which of the following is NOT a form of intermediate representation used by compilers?

- a) Postfix
- b) Tuples
- c) Context-free grammar
- d) Abstract syntax tree
- e) Virtual machine code

30

Question 26

- What optimisation technique could be applied in the following examples?

$a = b^2$

$a = a / 2$

- a) Constant Folding
- b) Code Deletion
- c) Common Sub-Expression Elimination
- d) Strength Reduction
- e) Global Register Allocation

31

Question 27

- Why can't we allocate data frames statically, i.e. have one fixed area for each subprogram? Which of the following are true

- I. Data Structures may be dynamically allocated
- II. Object Orientation demands the creation of Instances
- III. Recursion causes data frames to grow arbitrarily

- a) I only
- b) III only
- c) I and II only
- d) II and III only
- e) I, II and III only

32

Question 28

- Lex is a software tool that can be used to aid compiler construction. It is an example of which of the following?
 - A scanner generator
 - A parser generator
 - A code generator generator
 - A semantic analyser
 - A code debugger

33

Question 29

- The following two statements describe the performance of two programs:
 - A performs a total of 20 seconds of computation and 15 seconds of input/output.
 - B performs a total of 30 seconds of computation and 10 seconds of I/O
- Which of the following are true?
 - It will take up to 50 seconds to run A and B sequentially
 - It will take up to 75 seconds to run A and B sequentially
 - Using multiprogramming, the shortest time to execute both is 30 seconds
 - Using multiprogramming, the shortest time to execute both is 50 seconds
 - I and III
 - I and IV
 - II and III
 - II and IV
 - None of the above

34

Question 30

- Suppose two users simultaneously type the following command at the unix shell command prompt (\$):


```
$ ls -l
```
- Which of the following are true?
 - One process and one program is involved
 - Two processes and two programs are involved
 - One process and two programs are involved
 - Two processes and one program are involved
 - None of the above

35

Answers 21-30

21. Answer: c - $X = X \gg 2$ as dividing by four is equivalent to dividing by 2^2
22. Answer: b - Direct Lookup! All the other methods describe a way of searching for an ident within a dynamically sized symbol table. Direct lookup mechanisms work if the table is fixed. However, the table's size has to be dynamic as the compiler cannot know how many variables may appear within a user's code.
23. Answer: b - 11 - as there are $\log_2 N-1$ comparisons on average
24. Answer: d - With postfix expressions, operands are pushed on the stack first, and then when an operator is encountered, it pops off the values from the stack, performs the operation, and pushes the resulting value on the stack. Thus, with "a b * c d / -", the values a and b will be pushed on the stack, then popped off for the multop (*) operation, and the result returned to the stack before values c... etc.
25. Answer: c - A context-free grammar defines the language used by the compiler; the rest are intermediate representations
26. Answer: d - Both expressions can be reduced by changing the operator: $a = b \wedge 2$ can be reduced to $a = b * b$; whereas $a = a / 2$ is a right shift operation: $a = a \gg 1$
27. Answer: c - I and II only; recursion does not affect the size of data frames
28. Answer: a - Lex is responsible for identifying tokens using regular expressions. It is therefore a scanner generator
29. Answer: d - When I/O tasks are performed, the process will be held in the blocked queue, until the process has completed. This answer assumes that a single processor is used.
30. Answer: d - Only one program (ls) is involved, but this will be run as two processes.

36

Question 31

- If the UNIX command 'head file' outputs the first 10 lines of file, the command 'tail -n file' outputs the last n lines of file, and the command 'wc -w file' counts the number of words in file, what will the following output?

```
head file | tail -1 | wc -w
```

- The number of words in the tenth line from the end of file
- The first 10 lines of file, then the last line of file, then the number of words in file
- The number of words in line 10 of file only
- The number of words in line 10, then line 9, then line 8, etc.
- The number of words in the first ten lines plus the last line of file

37

Question 32

- If a process executes a *fork()* system call, which of the following are true?

- The parent process is moved to the *blocked* state
- The child process is placed in the *running* state
- The parent process is moved to the *ready* state
- The child process is placed in the *blocked* state
- The parent process is moved to the *terminated* state

38

Question 33

- Which of the following statements about threads is FALSE?

- A Java program need not necessarily lead to the creation of any threads
- A thread is sometimes referred to as a lightweight process
- Threads share code and data access
- Threads share access to open files
- Threads are usually more efficient than conventional processes

39

Question 34

- The value of a semaphore *s* is initially 1. What could happen in the following situation?

T1	T2
V(s);	P(s);
<i>critical region</i>	<i>critical region</i>
P(s);	V(s);

- Deadlock will ensue
- T1 and T2 can both enter their critical regions simultaneously
- Neither T1 nor T2 can enter its critical region
- T1 can never enter its critical region, but T2 can enter its own
- T1 can enter its critical region, but T2 can never enter its own

40

Question 35

- In a computer memory, a 100K partition becomes available. In the ready list is a program image of size 300K, plus three others of sizes 100K, 85K and 15K.
- Assuming that our current priority is to avoid starvation of the 300K program, which of those in the list should be swapped into the available partition?
 - The 100K program
 - The 85K program
 - The 15K program
 - Both the 15K and the 85K programs
 - None of the above

41

Question 36

- A new program requires 100K of memory to run. The memory management approach adopted is a simple partitioning one, and the operating system has the following list of empty partitions:

60K, 240K, 150K, 600K, 108K, 310K
- Assuming that the 150K partition is chosen, say which of the following selection strategies is being used:
 - First fit
 - Best fit
 - Worst fit
 - All of the above
 - None of the above

42

Question 37

- The page table shown below is for a job in a paged virtual storage system with a page size of 1K:

segment	datum
0	4
1	2
2	0
3	1
- A virtual address of [1, 352] would map on to what actual address?
 - 354
 - 1376
 - 2352
 - 2400
 - 4448

43

Question 38

- To assist in locating a bug that is causing a program to crash, a programmer inserts print statements as follows:


```
begin
  print("Got to point A without crashing");
  print("Got to B without crashing");
end
```
- Too much information will probably be written to the screen to allow location of the bug.
- The very insertion of the print statements will probably alter the program's behaviour, preventing the bug from occurring.
- The new diagnostic statements will interfere with the program's existing output, introducing further bugs.
- The bug will probably make all print statements inoperable.
- The use of output buffers by the system might prevent some messages from being written.

44

Answers 31-38

31. **Answer: c** - Only line 10 will be passed to wc -w
32. **Answer: c** - The parent process will be moved back to the ready state (depending on the scheduling policy), and once the child has been admitted, will also be placed in the ready state
33. **Answer: a** - Every Java program starts as a thread! The rest of the statements are true...
34. **Answer: b** - If T1 executes first, then it acquires the semaphore, which is immediately released by T2. Both then execute the critical region. However, if T2 executes first, it releases a semaphore it does not have, which can be acquired by T1. Again, both can execute the critical region.
35. **Answer: e** - If we are avoiding the starvation of the 300K program, then we need to wait for an adjacent partition to become available to allow for the 300K program to run
36. **Answer: e** - First Fit would select 240K; Best Fit would select 108K; and Worst Fit would select 600K. As none of these select the 150K partition, then some other strategy has been used!
37. **Answer: d** - 2048 (from page 1) + 352 (offset)
38. **Answer: e** - The buffers may not be flushed, and the program may continue (and crash) giving misleading information.