

# THE UNIVERSITY OF AUCKLAND

---

SECOND SEMESTER, 2012

Campus: City

---

**Computer Science and Software Engineering**

**Operating Systems**

**(Time Allowed: TWO HOURS)**

**NOTE:**

- The use of calculators is NOT permitted.
- Compare the exam version number on the Teleform sheet supplied with the version number above. If they do not match, ask the supervisor for a new sheet.
- Enter your name and student ID on the Teleform sheet. Your name should be entered left aligned. If your name is longer than the number of boxes provided, truncate it.
- Answer all **Multiple-choice** questions on the Teleform answer sheet provided. Attempt all questions. There are no negative marks.
- Use a dark pencil to mark your answers in the multiple choice answer boxes on the Teleform sheet. If you spoil your sheet, ask the supervisor for a replacement.
- Choose the BEST answer for each question based on the material covered in the course.
- There are 57 questions in this exam, worth 60 marks in total. Allocate your time appropriately.

CONTINUED

**Question 1**

[1 mark] Which of the following is NOT a good reason why operating systems are written in C?

- (a) The mapping from C to machine language is straightforward.
- (b) C has sophisticated memory management capabilities.
- (c) C provides direct access to memory locations.
- (d) C programs are very fast.

**Question 2**

[1 mark] Which of the following statements about the earliest digital computers is FALSE?

- (a) They were very expensive.
- (b) The CPUs were idle most of the time.
- (c) Programmers had complete access to the machine.
- (d) The operating system on the computers included simple memory management.

**Question 3**

[1 mark] Instructions which query the running state of a processor (i.e. kernel or user mode) make virtualization difficult. Which of the following best explains why?

- (a) An instruction running in user mode of a guest operating system would report that it was running in kernel mode.
- (b) An instruction running in the kernel of the host operating system would report that it was running in user mode.
- (c) An instruction running in the kernel of the host operating system would report that it was running in kernel mode.
- (d) An instruction running in the kernel of a guest operating system would report that it was running in user mode.

**Question 4**

[1 mark] What advantages are there in developing an operating system as a series of layers?

- (a) It makes debugging easier and helps get the design correct for future expansion.
- (b) It is very efficient and uses the client/server model.
- (c) It simplifies verifying that the system is correct and makes debugging easier.
- (d) It simplifies verifying that the system is correct and is more efficient than alternative approaches.

**Question 5**

[1 mark] In order to allow safe multiprogramming what hardware advances needed to be made?

- (a) Hardware memory protection, interrupt processing, privileged instructions.
- (b) Hardware memory protection, time-sharing terminals, real-time processing.
- (c) Interrupt processing, privileged instructions, time-sharing terminals.
- (d) Interrupt processing, time-sharing terminals, hardware memory protection.

**Question 6**

[1 mark] Which of the following groups of information could sensibly be found associated with a PCB?

- (a) memory, devices, process state, which processor
- (b) access rights, links to other processes, cpu usage, owner
- (c) open files, condition handlers, resource limits, priority
- (d) All of the above.

Here is an attempt at a solution to the producer/consumer problem.

```
exclusive_access = Semaphore.new(1)
number_deposited = Semaphore.new(0)
shared_buffer = 0

producer = Thread.new do
  while true
    next_result = whatever
    exclusive_access.wait()
    shared_buffer = next_result
    number_deposited.signal()
    exclusive_access.signal()
  end
end

consumer = Thread.new do
  while true
    exclusive_access.wait()
    number_deposited.wait()
    next_result = shared_buffer
    exclusive_access.signal()
    puts next_result
  end
end
```

### Question 7

[1 mark] What could possibly go wrong with the above semaphore solution to the producer/consumer problem?

- (a) The consumer could grab the `exclusive_access` semaphore and stop the producer from being able to store its data in the `shared_buffer`.
- (b) A fast producer could overwrite the `shared_buffer` variable with another result before the previous result was retrieved by the consumer.
- (c) If the consumer is not running, the producer continues to produce all of its data.
- (d) All of the above.

### Question 8

[1 mark] Which of the following is not a necessary component of a monitor (the concurrency construct)?

- (a) publicly accessible entry points
- (b) a scheduler
- (c) a readers/writers lock
- (d) a shared resource which is protected by the monitor

**Question 9**

[1 mark] Which of the following statements is FALSE about condition variables?

- (a) Condition variables are commonly used to block threads until the state of the program means they can continue.
- (b) Condition variables are associated with a queue that can hold threads.
- (c) Calling wait on a condition variable only puts a thread to sleep if the thread should not continue at that time.
- (d) Condition variables do not maintain any state about how many times they have been signaled.

**Question 10**

[1 mark] There is a lock variable associated with each Java object, on entering a synchronized method the lock is checked to see if it is free, already held by the current thread, or held by another thread. Which of the following statements is FALSE?

- (a) Once a thread acquires the lock, it releases the lock when it has finally left the synchronized method or it calls wait on the lock object.
- (b) Only one thread can be holding the lock at a time.
- (c) On entering a synchronized method the thread only progresses if the lock is free.
- (d) A thread can recursively call the same synchronized method safely.

**Question 11**

[1 mark] The Dining Philosophers can be used to illustrate what types of problem?

- (a) Priority inversion
- (b) Deadlock
- (c) Deadlock and livelock
- (d) Livelock

**Question 12**

[1 mark] Which of the following statements is TRUE?

- (a) Semaphores are at least as powerful as monitors.
- (b) Monitors are at least as powerful as semaphores.
- (c) Simple locks are at least as powerful as semaphores.
- (d) All of the above.

Here is some code with a compare and swap function “cas”.

```
add_to_balance(increase):  
    previous_amount = balance  
    while (!cas(&balance,  
               previous_amount,  
               previous_amount + increase)):  
        previous_amount = balance
```

**Question 13**

[1 mark] What is the code above doing?

- (a) It safely increments `balance` by `increase` using a lock-free algorithm.
- (b) It increments `balance` by `increase` using a condition variable.
- (c) It repeatedly increments `balance` by `increase` until `balance` overflows.
- (d) It safely swaps the values of `balance` with `previous_amount + increase` using a wait-free algorithm.

**Question 14**

[1 mark] Which of the following statements about deadlock is TRUE?

- (a) A wait-for graph can be used to detect deadlock.
- (b) The best way to solve the problem of deadlock is to use Dijkstra’s Banker’s algorithm.
- (c) It is impossible for deadlock to occur in a system with only one process.
- (d) Havender described the five necessary conditions for deadlock to occur.

**Question 15**

[1 mark] Which of the following answers best explains the Wound-Wait deadlock prevention scheme?

- (a) Processes can take resources from younger processes. Processes must restart rather than waiting for resources held by older processes.
- (b) Processes can take resources from younger processes. Processes may wait for resources held by older processes.
- (c) Processes can take resources from older processes. Processes may wait for resources held by younger processes.
- (d) Processes must restart rather than waiting for resources held by older processes. Processes may wait for resources held by younger processes.

**Question 16**

[1 mark] Which of the following is a problem with deadlock in distributed systems?

- (a) False cycle detection.
- (b) How to maintain the wait-for graph.
- (c) The Banker's algorithm is more complicated in a distributed system.
- (d) All of the above.

**Question 17**

[1 mark] Which of the following statements about files is FALSE?

- (a) All Unix files are stored on secondary storage.
- (b) With sparse files it is possible to have the size of a file larger than the device the file is stored on.
- (c) Executable files have different structures specific to their particular operating systems.
- (d) Moving a file does not necessarily mean that all of the data needs to be copied and then the original file deleted.

**Question 18**

[1 mark] File types are useful for which of the following reasons?

- (a) They specify whether a file is executable or not.
- (b) They provide information to the file system about how the contents of a file are to be treated.
- (c) They specify which programs can open which files.
- (d) All of the above.

**Question 19**

[1 mark] Which of the following statements about file access is TRUE?

- (a) The only way to read from a particular position in a sequential access file is to first seek to that position.
- (b) The file system must maintain the current position pointer every time a direct (random) access file is read.
- (c) Writing to a file is usually faster than reading from it.
- (d) The file system does not always have to update the file length when a file is written to.

**Question 20**

[1 mark] Which of the following statements about MacOS resource forks or NTFS alternate data streams is FALSE?

- (a) Files with alternate data streams or resource forks can easily be handled by a wide range of operating systems.
- (b) The main use of NTFS alternate data streams is to hide extra data on the device.
- (c) The text in a word processing document would normally be stored in the resource fork.
- (d) All of the above.

The following commands are input to a correct solution of assignment 2. The next two questions are about the state of the system after all of these commands have been executed.

```
mkdir d1
mkdir d1/d1
link d2 d1/d1
create d2/f1
append "hello-" d2/f1
create d1/f1
append "world" d1/f1
enter d1
append "again" d1/f1
listfiles
```

And here is some possible output.

Possible output 1:

```
=== / ===
d1          d          6
d2          d          3
```

Possible output 2:

```
=== /d1/ ===
d1          d          3
f1          d          5
```

Possible output 3:

```
=== /d1/ ===
d1          d          3
f1          d          11
```

Possible output 4:

```
=== / ===
d1          d          6
d2          d          3
f1          d          5
```



**Question 21**

[2 marks] Which of the possible outputs on the previous page is the correct output?

- (a) Possible output 4
- (b) Possible output 1
- (c) Possible output 2
- (d) Possible output 3

**Question 22**

[1 mark] What is the content of `/d2/f1` after running the commands on the previous page?

- (a) hello-
- (b) hello-again
- (c) worldagain
- (d) hello-world

**Question 23**

[1 mark] What was the main reason given in lectures that file systems usually do not allow hard links to directories?

- (a) Directories can be on different disk devices and this stops hard links from working.
- (b) Hard links don't allow simultaneous access to files and this is particularly bad if the files are directories.
- (c) Hard links are slower than soft links and directory access has to be fast.
- (d) It is possible to create cycles in the directory tree.

**Question 24**

[1 mark] For contiguous allocation of disk space to files which of the following algorithms leaves the smallest amount of external fragmentation when a single file allocation is made?

- (a) Best fit
- (b) Last fit
- (c) First fit
- (d) Worst fit

**Question 25**

[1 mark] Using contiguous allocation for files if empty disk space is maintained in a linked list with random ordering, which of the following algorithms would be fastest to allocate space to a file.

- (a) Worst fit
- (b) First fit
- (c) Last fit
- (d) Best fit

**Question 26**

[1 mark] Which of the following is TRUE regarding linked allocation of files?

- (a) Linked allocation decreases the amount of internal fragmentation.
- (b) Linked allocation increases the amount of external fragmentation.
- (c) Linked allocation is slower than contiguous allocation for direct (random) access of file data.
- (d) Damage to one block never causes the loss of the rest of the file.

**Question 27**

[1 mark] Which of the following is TRUE regarding indexed allocation of files?

- (a) Damage to one block never causes the loss of multiple blocks of the file.
- (b) Indexed allocation is considerably faster than linked allocation for sequential access.
- (c) All indexed allocation schemes have no limits on the file size.
- (d) Indexed allocation means there is no external fragmentation.


**Question 28**

[1 mark] NTFS uses extents to keep track of the data of a file. Which of the following is the best description of an extent?

- (a) Extents are the current versions of File Allocation Table (FAT) linked data structures.
- (b) An extent holds the first 1K of a file followed by links to further extents.
- (c) An extent consists of a start cluster number and a length.
- (d) An extent is a pointer to a Master File Table entry.

The following text snippets with their labels are associated with the version tree below.

- A: Dear Mum, I\_
- B: am
- C: thought I was
- D: \_doing really well in my Operating System course
- E: \_until I sat the test
- F: .\_I\_
- G: hope you are well.



**Question 29**

[1 mark] Using the above tree what is the file data for version 2 of the file?

- (a) Dear Mum, I\_am\_doing really well in my Operating System course.\_I\_ hope you are well.
- (b) am\_doing really well in my Operating System course.\_I\_
- (c) Dear Mum, I\_amthought I was\_doing really well in my Operating System course\_until I sat the test.\_I\_ hope you are well.
- (d) Dear Mum, I\_ hope you are well.

Here is some output from Assignment 3, it was generated repeatedly by the same algorithm.

```
ref : 1 : 2 : 3 : 4 : 1 : 2 : 5 : 1 : 2 : 3 : 4 : 5
=====
      : 1 : = : = : = : = : = : = : = : = : = : = : = : 5
      : 0 : 2 : = : = : = : = : = : = : = : = : = : = : =
      : 0 : = : 3 : = : = : = : 5 : = : = : = : = : 4 : =
      : 0 : = : = : 4 : = : = : = : = : = : = : 3 : = : =
page faults: 8
```

**Question 30**

[2 marks] Which of the following replacement algorithms was most likely used to produce the output above?

- (a) Least recently used (LRU)
- (b) First in first out (FIFO)
- (c) Random
- (d) Optimal

Here is some more output from Assignment 3, it was generated repeatedly by the same algorithm.

```
ref : 1 : 2 : 3 : 4 : 1 : 2 : 5 : 1 : 2 : 3 : 4 : 5
=====
      : 1 : = : = : = : = : = : = : = : = : = : = : 4 : =
      : 0 : 2 : = : = : = : = : = : = : = : = : = : = : =
      : 0 : = : 3 : = : = : = : = : = : = : = : = : = : =
      : 0 : = : = : 4 : = : = : 5 : = : = : = : = : = : =
page faults: 6
```

### Question 31

[2 marks] Which of the following replacement algorithms was most likely used to produce the output above?

- (a) Random
- (b) Least recently used (LRU)
- (c) Optimal
- (d) First in first out (FIFO)

### Question 32

[1 mark] The algorithm for writing to a Unix file was given in class. This included information about accessing inode information. On such a system two processes write to an empty file at the same time. One writes “hello”, the other writes “world”. Which of the following contents of the file could NOT occur?

- (a) hellwoorld
- (b) helloworld
- (c) worldhello
- (d) world

**Question 33**

[1 mark] Which of the following statements best describes locality of reference and why it is important?

- (a) Over a long period of time most memory accesses are concentrated in small areas. This is important because it can be used to minimise the amount of memory that needs to be available in the system.
- (b) Memory or file accesses over a short period of time are largely predictable. This is important because it means that we can design our programs without having to take access behaviour into account.
- (c) The names of files used in a small period of time are closely related. Parts of the pathnames are the same and this is important in both local and distributed file systems.
- (d) Many memory or file accesses in a small period of time are close together. This is important for the efficiency of virtual memory, read-ahead in file data and in distributed file systems.

**Question 34**

[1 mark] Which of the following statements about transparency in distributed file systems is FALSE?

- (a) Both NFS and AFS provide location transparency.
- (b) Migration transparency means that files can be moved without programs which use those files needing to be changed.
- (c) Location transparency means that there is no visible connection between a file name and its position.
- (d) Location transparency requires migration transparency.

**Question 35**

[1 mark] Which of the following is NOT usually regarded as an advantage of a distributed file system?

- (a) It provides uniform access times to all files.
- (b) It provides greater flexibility for sharing data.
- (c) It provides a greater amount of storage.
- (d) It enables easy replication of file resources.

**Question 36**

[1 mark] In which of the following areas would caching of file data be better than using remote service?

- (a) Simpler implementation
- (b) Less local memory used
- (c) Speed of access to file data
- (d) All of the above

**Question 37**

[1 mark] Which of the following statements about consistency semantics is not TRUE?

- (a) Unix semantics is based on session semantics.
- (b) Unix semantics means that any change to a file is immediately visible to any other process.
- (c) With session semantics a process may be working on data which is out of date.
- (d) Session semantics is simpler to implement in a distributed file system than Unix semantics.

**Question 38**

[1 mark] Which of the following is FALSE about NFS as discussed in class?

- (a) NFS only uses dedicated servers.
- (b) Tables on each machine list the mount points and server information for the distributed files.
- (c) Remote directories are unmounted automatically by the automounter when unused.
- (d) Moving a collection of files is complicated.

**Question 39**

[1 mark] Which of the following is TRUE about AFS as discussed in class?

- (a) When a file is relocated all processes using the file lose access to it.
- (b) AFS is really only useful in small networks.
- (c) All AFS clients can also act as servers.
- (d) AFS uses Kerberos for system authentication.

**Question 40**

[1 mark] Which of the following statements concerning distributed services is FALSE?

- (a) A Remote Procedure Call uses stub functions on both the client and server end to connect a local procedure call on one machine to a procedure on a remote machine.
- (b) Process migration is not commonly used for load balancing because it is too expensive.
- (c) In a tuplespace the in and out primitives are atomic.
- (d) Remote Procedure Calls handle reference parameters in the same way they handle value parameters.

**Question 41**

[1 mark] Which of the following distributed computation technologies is fundamentally different from the other three?

- (a) Linda
- (b) RMI
- (c) CORBA
- (d) RPC

**Question 42**

[1 mark] Which of the following is NOT an advantage of most virtual memory systems?

- (a) Access control over areas of memory.
- (b) The ability to use more memory than exists in RAM.
- (c) The ability to easily share memory between processes.
- (d) Consistent access time to memory locations.

**Question 43**

[1 mark] Most operating systems don't use the least recently used page replacement algorithm because it is too expensive. The second chance or clock algorithm uses which of the following bits associated with page information to approximate least recently used selection?

- (a) the access bits
- (b) the mode bits
- (c) the valid bit
- (d) the referenced bit

**Question 44**

[1 mark] Which of the following is the best description of what causes thrashing?

- (a) When the foreground process has completely used up the number of frames it has been allocated.
- (b) When there is not enough contiguous memory to be allocated for all current working sets.
- (c) When all frames are currently being used.
- (d) When the sum of the pages of the current working sets exceeds the number of frames.

**Question 45**

[1 mark] Which of the following statements about virtual memory on Windows is FALSE?

- (a) Privileged processes can lock pages in real memory.
- (b) Once a page has been allocated to a process's working-set it cannot be removed.
- (c) A process can be allocated a number of frames up to its working-set maximum.
- (d) A process is guaranteed to be allocated its working-set minimum number of frames.

**Question 46**

[1 mark] Concerning protection and security which of the following definitions is incorrect?

- (a) A domain is a collection of ordered pairs <object, rights>
- (b) A setuid program is a program which can take on the privileges of the owner rather than those of the user.
- (c) A subject is a resource which can be used by a principal.
- (d) A reference monitor is a component which mediates accesses to objects.

Use the following access matrix for the next two questions.

	F1	F2	F3	D1	D2	D3	D4
D1	read* write*						switch
D2		owner read	owner	switch			
D3	read						switch
D4			read write			switch control	

CONTINUED



**Question 47**

[1 mark] Which of the following sequences of commands is the most efficient means to allow a process running in D2 to write to F3?

- (a) Add the write privilege over F3 to D2, then write to F3.
- (b) It cannot be done. A process in D2 cannot write to F3.
- (c) Switch to D1, switch to D4, then write to F3.
- (d) Switch to D1, copy the write privilege from F1 to F3, then write to F3.

**Question 48**

[1 mark] Which of the following sequences of commands is the most efficient means to allow a process running in D2 to remove the ability to read F1 from D3?

- (a) Switch to D1, remove read privilege from D3.
- (b) It cannot be done. No domain has rights over the privileges of another domain.
- (c) Switch to D3, remove read privilege.
- (d) Switch to D1, switch to D4, remove read privilege from D3.

**Question 49**

[1 mark] Which of the following best describes the Confused Deputy problem?

- (a) A program with privileges is tricked into misusing its authority.
- (b) A program with privileges is tricked into giving those privileges to another program.
- (c) A program with privileges gets so many requests that it loses track of which request came from which source.
- (d) A program with privileges mistakenly prevents access to a resource which should be available.

**Question 50**

[1 mark] Which of the following techniques would NOT be suitable to keep capabilities from being modified without permission?

- (a) On a tagged memory architecture, capabilities could be tagged and only privileged instructions could modify them.
- (b) All capabilities could be stored in kernel memory and only modified by system calls.
- (c) Capabilities could be encrypted. The public key could be used to check them but the private key to produce them is held by the security system.
- (d) Each process stores its own capabilities in memory not shared by other processes and provides them only when access control is being verified.

**Question 51**

[1 mark] Which of the following protocols allows two parties to cooperatively form a secret key for communication but without authenticating themselves with each other?

- (a) Needham-Schroeder symmetric key
- (b) Diffie-Hellman
- (c) Kerberos
- (d) Menezes-Qu-Vanstone

**Question 52**

[1 mark] Given the three “c”s of security failure: change, complacency and convenience, which of the following is most likely to produce a problem based on convenience?

- (a) Making a quick patch to enable a program to work on a different system.
- (b) Allowing a superuser to access all machine state.
- (c) Not checking the bounds on an array.
- (d) Providing all users with an electronic card key.

**Question 53**

[1 mark] Which of the following statements is FALSE with regards to Kerberos authentication?

- (a) Kerberos tickets are time expiring capabilities.
- (b) Kerberos is based on the Needham-Schroeder protocol.
- (c) If the Kerberos Authentication Server is down, no one can log in.
- (d) The Ticket Granting Server maintains the private keys of all principals.

**Question 54**

[1 mark] Which of the following are commonly used to gain users' passwords?

- (a) Trojan horse
- (b) Dictionary attacks
- (c) Social engineering
- (d) All of the above

Here is a section of the `/dev` directory from a current version of Linux.

```
crw----- 1 root root      5,   1 Sep 26 17:06 console
crw-rw---- 1 root video   29,   0 Sep 26 17:06 fb0
brw-rw---- 1 root disk     8,   0 Sep 26 17:06 sda
crw-rw---- 1 root disk   21,   0 Sep 26 17:06 sg0
crw-rw-rw- 1 root tty     5,   2 Sep 26 17:06 ptmx
```

**Question 55**

[1 mark] From the information above, which of the following devices use the same device driver?

- (a) `sda` and `sg0`
- (b) `console` and `ptmx`
- (c) `fb0`, `sda` and `sg0`
- (d) `console` and `fb0`

**Question 56**

[1 mark] Which of the following is an advantage that user level device drivers have over more traditional device drivers?

- (a) If a user level device driver causes a problem it is isolated from the rest of the system.
- (b) User level device drivers can more easily access interrupts using non-privileged code.
- (c) User level device drivers are faster.
- (d) User level device drivers allow multiple drivers to be used by a device at the same time.

**Question 57**

[1 mark] Which of the following disk scheduling algorithms potentially services requests of the middle tracks up to twice as often as the innermost or outermost tracks?

- (a) N-step SCAN
- (b) SSTF - shortest seek time first
- (c) C-SCAN - circular SCAN
- (d) SCAN - elevator algorithm

**Rough Working – This page will not be marked**

---