

THE UNIVERSITY OF AUCKLAND

SECOND SEMESTER, 2013

Campus: City

COMPUTER SCIENCE

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.

CONTINUED

Question 1

[1 mark] Which of the following models was NOT discussed as a model for operating system design and implementation?

- (a) the manager model
- (b) the onion model
- (c) the ogre model
- (d) the resource allocator model

Question 2

[1 mark] Which of the following lists of the development of operating systems is in chronological order?

- (a) resident monitors, time-sharing systems, batch systems, personal computers
- (b) resident monitors, batch systems, time-sharing systems, personal computers
- (c) resident monitors, time-sharing systems, personal computer, batch systems
- (d) time-sharing systems, resident monitors, batch systems, personal computers

Question 3

[1 mark] Which of the following statements about tightly coupled and loosely coupled multiprocessor systems is FALSE?

- (a) A loosely coupled system requires locks on shared resources, a tightly coupled system doesn't.
- (b) The operating system on a tightly coupled system is the same for each processor; on a loosely coupled system it could be different.
- (c) Communication between processes is handled differently; tightly coupled systems can share memory, loosely coupled systems generally use a network connection.
- (d) Processes share the same real-time clock on a tightly coupled system but not on a loosely coupled system.

Question 4

[1 mark] Which of the following statements about real-time systems is FALSE?

- (a) Virtual memory makes it more difficult to implement a hard real-time system.
- (b) A soft real-time system is adequate to playback videos.
- (c) Hard real-time systems must satisfy requests within definite time periods.
- (d) Most modern operating systems can handle hard real-time constraints.

Question 5

[1 mark] Which of the following statements about mobile operating systems is FALSE?

- (a) Android and iOS don't use virtual memory.
- (b) iOS is derived from MacOS X.
- (c) Early mobile operating systems were forms of resident monitors.
- (d) Android is built on top of a version of Linux.

Question 6

[1 mark] Which of the following is NOT an advantage of developing an operating system as a series of layers?

- (a) It simplifies verifying that the system is correct.
- (b) It can help with security by including checks between the layers.
- (c) It makes debugging of the system easier.
- (d) It helps get the design correct for future changes.

Question 7

[1 mark] Unix style operating systems try to make many resources look like files. Which of the following resources appear as files on Linux?

- (a) named pipes
- (b) processes
- (c) devices
- (d) All of the above.

Question 8

[1 mark] Which of the following best explains what happens when a damaged C program comes to an end but doesn't call the exit routine?

- (a) The operating system takes control when the program tries to execute an illegal instruction or attempts to access unallocated memory.
- (b) The operating system creates a new process and restarts the damaged program in that process so that it gets another chance to complete.
- (c) The C standard library takes control when the program fails to return to the code which called the main function.
- (d) The damaged program can corrupt memory used by other processes and cause them to crash or perform illegal instructions.

Question 9

[1 mark] Which of the following statements explain why threads are sometimes called lightweight processes?

- (a) Changing from one thread to another in the same process is a simple operation.
- (b) Threads are easier to create than processes.
- (c) Threads in the same process share most of the same resources.
- (d) All of the above.

Question 10

[1 mark] Which of the following statements is an advantage of system-level threads over user-level threads?

- (a) System-level threads work even if the operating system doesn't support threads.
- (b) When a system-level thread blocks in the kernel it doesn't stop all other threads in the same process.
- (c) It is easier to switch between system-level threads than user-level threads.
- (d) It is easier to create system-level threads than user-level threads.

Question 11

[1 mark] Older versions of Unix essentially treated the kernel as a monitor (the concurrency construct). Which of the following was the Unix kernel equivalent to a monitor's publicly accessible procedures?

- (a) the system wide open file table entries
- (b) page table entries
- (c) system calls
- (d) the entry queue

Question 12

[1 mark] What is the principal reason that the kernel in modern Unixes cannot be seen as a monitor?

- (a) threads can communicate with each other when in the kernel
- (b) multiple threads can run concurrently in the kernel
- (c) threads can be blocked waiting on resources in the kernel
- (d) there is no equivalent construct in the kernel to condition variables

Question 13

[1 mark] Which of the following best describes a semaphore?

- (a) A lock which allocates a resource to only one thread at a time.
- (b) An integer count which signals processes to continue when the count reaches the same number as the number of threads waiting on the semaphore.
- (c) A region of code in which only one thread is guaranteed to be executing at a time.
- (d) An integer count with atomic operations to increase and decrease the count along with conditions on when the operations should return.

Question 14

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

- (a) Condition variables are commonly used to block threads until the state of the program means they can continue.
- (b) Calling wait on a condition variable only puts a thread to sleep if the thread should not continue at that time.
- (c) Condition variables only come into existence when a particular condition is true.
- (d) Condition variables keep a count of how many times they have been signalled.

Question 15

[1 mark] Which of the following statements about the readers-writers problem is FALSE?

- (a) A readers preferred solution is appropriate when the most recent data is not required.
- (b) Multiple readers may be executing in the critical section at a time.
- (c) A writers preferred solution is always appropriate when there are many writers and hardly any readers.
- (d) Only one writer may be executing in the critical section at a time.

Question 16

[1 mark] Which of the following best describes a lock-free algorithm?

- (a) A lock-free algorithm ensures that no thread waits for an indefinite time when accessing a shared resource without using a lock.
- (b) A lock-free algorithm ensures that a shared resource is kept consistent even when accessed by multiple threads without using a lock.
- (c) A lock-free algorithm frees up locks whenever multiple threads access the shared resource.
- (d) A lock-free algorithm uses a busy wait to ensure that only one thread at a time has access to the shared resource.

Question 17

[1 mark] Which of the following is NOT one of Havender's conditions for deadlock?

- (a) Resources can only be shared by passing them around the circular list of processes.
- (b) There is a circular list of processes each wanting a resource owned by another in the list.
- (c) A process can hold a resource while requesting another.
- (d) Only the owner can release the resource.

Question 18

[1 mark] Which of the following is not sufficient to prevent deadlock?

- (a) Allocate all resources in a specific order. Resources earlier in the order cannot be requested after resources later in the order have been obtained.
- (b) Allow a process to wait for a resource only if the process has a better priority than the process holding the resource.
- (c) Allocate all resources a program needs before the program starts running.
- (d) Inserting time stamps on resource requests in a distributed environment.

Question 19

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

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

Question 20

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

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

Question 21

[1 mark] `mv` is the Unix command to rename (or move) a file. `echo` is used here to put a string into a file.

Given the following Unix commands which create a file, make a soft-link to the file, rename the file and the soft-link and then create another file, what is the result of reading file “ddd”?

```
$ echo 'something' > aaa
$ ln -s aaa bbb
$ mv aaa ccc
$ mv bbb ddd
$ echo 'else' > aaa
$ cat ddd
```

- (a) somethingelse
- (b) else
- (c) something
- (d) No such file error

Question 22

[1 mark] If the link created in the previous question was a hard link what would the result of reading file “ddd” have been?

- (a) somethingelse
- (b) else
- (c) something
- (d) No such file error

Question 23

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

Question 24

[1 mark] Which of the following is NOT a technique which can be used to keep track of the blocks allocated to a file?

- (a) Linked blocks
- (b) Extents
- (c) Clusters
- (d) Index blocks

Here is the current version of a file and the log of changes maintained by a file versioning system as in the lecture notes.

Current version

Dear Mum, I can't tell if this exam is terrible. I hope you are well.

changes from version 1 to version 2:

11i20

changes from version 2 to version 3:

24d
swell
24i8

changes from version 3 to version 4:

11i16

Question 25

[2 marks] Which of the following strings shows version 2 of the file? The first version was the original version of the file. Indexes in the string start at position 1.

- (a) Dear Mum, I hope you are well.
- (b) Dear Mum, this exam is swell. I hope you are well.
- (c) Dear Mum, this exam is terrible. I hope you are well.
- (d) Dear Mum, I can't tell if this exam is terrible. I hope you are well.

Here is the data from the first block of a drive (and the separator) in assignment 2.

```
1
volume exam
20
x-----x
19

** 0 **
```

Question 26

[1 mark] What does the value 20 represent above?

- (a) the block number where the root directory information starts
- (b) the number of blocks used to hold the volume meta-data
- (c) the number of blocks in the volume (or drive)
- (d) the number of times the volume has been mounted

Question 27

[1 mark] Which of the following statements would NOT cause an error in assignment 2. You may assume that the drive has been correctly created.

- (a) `Volume.format(drive, b'Volume A')`
- (b) `Volume.format(drive, b'Volume\nA')`
- (c) `Volume.format(drive, b'A' * blocks * Drive.BLK_SIZE)`
- (d) `Volume.format(drive, b'')`

Question 28

[1 mark] Which of the following approaches could have been used to store the file data block information in assignment 2?

- (a) linked blocks
- (b) multiple levels of index blocks
- (c) linked index blocks
- (d) All of the above.

Question 29

[1 mark] Which of the following was a limitation on the file system in assignment 2?

- (a) The maximum length of the volume name was 63 bytes.
- (b) The maximum length of a file name was 255 bytes.
- (c) The maximum number of directories was 1.
- (d) The maximum file size was 64 blocks.

Question 30

[1 mark] In assignment 2 some people did not write any data to the drive until the drive was unmounted. Which of the following is a reason why this technique would not be used in a real file system?

- (a) If the machine was accidentally switched off the data would not be on the drive.
- (b) There is not enough space to hold all file contents in memory.
- (c) The unmount would take too long.
- (d) All of the above.

Question 31

[1 mark] Which of the following information does NOT have to be kept by the server in a stateful distributed file system?

- (a) the current file location
- (b) how the file data is structured
- (c) the type of access the file is open for
- (d) an identifier for the process which has opened the file

auto_master

/joannemap_home

/home auto_home

map_home

joanne server1:/export/home/joanne

greg server2:/export/home/joanne

auto_home

home server3:/export/home/joanne

joanne server4:/export/home/joanne

Question 32

[1 mark] Given the NFS map files above (similar to those shown in class), on what server does /home/joanne exist?

- (a) server1
- (b) server2
- (c) server 4
- (d) server3

Question 33

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

- (a) NFS mounts remote directories in a similar way to the method Unix mounts drives in the directory tree.
- (b) NFS only works in a homogeneous environment with all clients and servers running the same operating system.
- (c) Remote file directories in NFS can only be mounted when a machine is booted.
- (d) Servers in NFS are dedicated to acting as file servers and cannot be used for general operations.

Question 34

[1 mark] Which of following best explains the use of XDR in NFS communications?

- (a) It is the protocol of function calls from clients to external servers which replicate file data.
- (b) It is used to represent data being transmitted between clients and servers in a form independent of machine type or operating system.
- (c) It is used to recreate file access requests on the server so that they appear like local file access requests.
- (d) It checks file locations to determine if the file is local or remote. If it is remote it has to pass the request over the network.

Question 35

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

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

Question 36

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

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

Question 37

[1 mark] Which of the following statements about RPC and the reference parameter problem is FALSE?

- (a) Reference parameters are problematic in RPCs because the client and server memories are different.
- (b) If a reference parameter is read but not written to we only need to copy the referenced data from the client to the server, and not back again.
- (c) Even though reference parameters are difficult to use over RPC they are faster than value parameters.
- (d) Simulating reference parameters using copy/restore semantics is a good enough solution to the problem in many cases.

Question 38

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

- (a) Tuplespaces scale very well. There are no problems with bottlenecks.
- (b) The tuplespace itself deals with synchronization. Each tuple operation is atomic.
- (c) Processes don't communicate directly with each other when using a tuplespace.
- (d) Tuplespaces make it very easy to add extra servers and clients to applications.

Question 39

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

- (a) Virtual memory gives programs their own address space, possibly starting at address zero.
- (b) Virtual memory allows programs to access more memory than is available in the real primary memory of the machine.
- (c) Virtual memory can only be provided by partitioning memory into equal sized sections.
- (d) In many cases virtual memory provides protection for the operating system.

Question 40

[1 mark] Which of the following is FALSE when comparing using pages or segments for virtual memory?

- (a) Segments have no external fragmentation whereas pages can have external fragmentation.
- (b) Segments have no internal fragmentation whereas pages can have internal fragmentation.
- (c) Segments are sized according to logical units of the program, pages are not.
- (d) Segments have limit values in the segment table whereas pages do not have limit values in the page table.

Question 41

[1 mark] Why is there a need for a translation look-aside buffer when using virtual memory?

- (a) It solves the problem of locality of reference when pages which are stored a long way from each other are accessed at the same time.
- (b) It buffers pages in real memory so that access to them is faster.
- (c) It significantly reduces the average access time to memory by caching page table information.
- (d) It is used to keep copies of the page protection information so that access can be checked before retrieving pages from secondary storage.

Question 42

[1 mark] Which of the following best explains what happens concerning the current instruction when a page fault occurs?

- (a) The current instruction is only restarted if it caused the page fault. If it didn't cause the page fault it continues normally.
- (b) After the required page has been brought into a frame the instruction must usually be restarted from its initial state before the page fault occurred.
- (c) The instruction is completed and then the required page is brought into a frame so that the program can continue correctly.
- (d) The current thread changes to another thread and the offending instruction and its thread are removed from the system.

CONTINUED

Question 43

[1 mark] Which of the following conditions have a direct influence on the effective access time of memory?

- (a) the time to retrieve page table information for the page
- (b) the time taken to provide a free frame for the page if required
- (c) the page fault rate
- (d) All of the above.

Question 44

[1 mark] Which of the following page selection algorithms normally uses the referenced bit to determine if a page should relinquish its frame to another page?

- (a) death row
- (b) first in first out
- (c) least recently used
- (d) the second chance or clock algorithm

Question 45

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

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

Here is a table showing a page reference string across the top and the corresponding contents of frames in each column. The "=" character means the frame hasn't changed from the previous step. 0 means the frame is free.

ref	1	2	3	4	1	5	5	1	2	3	4	5
0	1	=	=	=	=	=	=	=	=	=	=	5
0	0	2	=	=	=	5	=	=	=	=	4	=
0	0	=	3	=	=	=	=	=	2	=	=	=
0	0	=	=	4	=	=	=	=	=	3	=	=

Question 46

[1 mark] Which of the following replacement algorithms would repeatedly produce the result above?

- (a) first in first out (FIFO)
- (b) random
- (c) optimal
- (d) least recently used (LRU)

Question 47

[1 mark] Which of the following techniques can be used to keep the kernel area of memory safe from accidental modification by a process?

- (a) a fence register
- (b) base and limit registers
- (c) page tables entries
- (d) All of the above.

Question 48

[1 mark] Which of the following best explains what minimising the attack surface means?

- (a) Increasing the difficulty of authentication so that access to system functions is reduced.
- (b) Removing functions which can be called directly or indirectly by an unauthorised user.
- (c) Removing functions which can cause damage after an attacker has found a vulnerability in the system.
- (d) Removing error prone functions which an authorised user can call directly or indirectly.

Question 49

[1 mark] Which of the following statements about code signing enforcement is FALSE?

- (a) Code signing enforcement means that executable pages must be unwritable.
- (b) Code must be checked when it is loaded to make sure that it hasn't been modified since the code was signed.
- (c) Code signing enforcement means that apps are safe to use because they have not been modified.
- (d) Code must be checked as it is executing to make sure that it hasn't been modified since it was signed.

Question 50

[1 mark] With regards to protection what is a reference monitor?

- (a) a list of ordered pairs showing the objects and access rights for each subject
- (b) the mechanism which checks accesses to objects by subjects
- (c) a list of ordered pairs showing the subjects and their access rights for each object
- (d) the mechanism which checks accesses to subjects by objects

Question 51

[1 mark] In Unix operating systems what does `setuid` do and why is it considered dangerous?

- (a) It changes the permissions of a process to be the permissions of the user of the program rather than those of the owner. This is dangerous because if the process can be hijacked then the owner has all of the user's permissions.
- (b) It changes the permissions of a process to be the permissions of the owner of the program rather than those of the user. This is dangerous because the user can now freely access any resources which the owner has permission to.
- (c) It changes the permissions of a process to be the permissions of the user of the program rather than those of the owner. This is dangerous because the program can now access any of the user's permissions.
- (d) It changes the permissions of a process to be the permissions of the owner of the program rather than those of the user. This is dangerous because if the process can be hijacked then the user has all of the owner's permissions.

Here is an access matrix.

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

Question 52

[1 mark] From the matrix above which of the following requests is impossible?

- (a) a process beginning in D4 reading from F1
- (b) a process beginning in D1 reading from F3
- (c) a process beginning in D2 writing to F2
- (d) None of the above.

Question 53

[1 mark] Given the following listing from a Unix system, what access rights does the user `jane` have to `thefile`?

```
-r---w---x    1    jane  staff    thefile
```

- (a) the owner's right
- (b) read only
- (c) read, write and execute
- (d) execute only

Question 54

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

- (a) Access Control Lists make it easy to determine which subjects can access an object.
- (b) Capabilities make it easy to determine which objects can be accessed from a particular domain.
- (c) Capabilities are easier to revoke than Access Control List permissions.
- (d) Both Capabilities and Access Control Lists are now widely used.

Question 55

[1 mark] Which of the following is NOT a way to make passwords safer?

- (a) Automatically refuse the use of common passwords.
- (b) Use computer generated random passwords.
- (c) Increase the length of the passwords.
- (d) Require passwords to be changed every day.

Question 56

[1 mark] In Kerberos authentication when a ticket has been granted to principal A enabling it to talk to principal B how does B verify the identity of A?

- (a) The ticket sent to B by A has been encrypted with A's secret key and includes the key for A and B to communicate, $K_{A,B}$, which was issued by the Ticket Granting Server. The only way the ticket could have been sent to B from A was if it decoded the ticket. Only A with its secret key is able to do this, so the ticket must have come from A.
- (b) The ticket sent to B by A has been encrypted with B's secret key and includes the key for A and B to communicate, $K_{A,B}$, which was issued by the Ticket Granting Server. A could not modify the ticket. Along with the ticket, A sends an authenticator encrypted with $K_{A,B}$. $K_{A,B}$ would only be accessible to A if it was authenticated by the Ticket Granting Server.
- (c) B sends a message encrypted with its secret key to the Ticket Granting Server asking it to verify that the ticket sent to it by A is valid and requests a secret identifier to verify that the sender of the ticket must be A. It checks this secret identifier against the identifier sent to B by A along with the ticket.
- (d) None of the above.

Question 57

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

- (a) The bottom half of a device driver deals with the interrupt handling of the device.
- (b) Bus interfaces make it easier to attach a wide range of devices to current computers.
- (c) The block buffer cache is used both to hold block input/output data and to cache the data which is used frequently.
- (d) In traditional Unix devices are specified with two identifiers. The major device number refers to the device driver and the minor device number refers to the particular device.

Question 58

[1 mark] Which of the following requirements have made user-level device drivers difficult to implement?

- (a) controlling DMA from user level
- (b) accessing IO ports from user level
- (c) allocating interrupt handling resources from user level
- (d) All of the above.

Question 59

[1 mark] Which of the following disk scheduling algorithms treats all tracks equally. i.e. Outside and inside tracks are serviced as regularly as those in between.

- (a) SSTF - shortest seek time first
- (b) FCFS - first come first served
- (c) C-SCAN - circular SCAN
- (d) N-step SCAN

Rough Working – This page will not be marked

