

THE UNIVERSITY OF AUCKLAND

SECOND SEMESTER, 2014

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 60 questions in this exam, worth 60 marks in total. Allocate your time appropriately.

CONTINUED

Question 1

[1 mark] Which of the following is NOT usually considered a requirement in a language used for operating system implementation?

- (a) It is easy to produce fast and efficient code with the language.
- (b) The language is compiled.
- (c) The language allows access to memory locations.
- (d) The language is dynamic and weakly typed.

Question 2

[1 mark] How did a time-sharing operating system differ from a batch operating system?

- (a) Time-sharing systems allowed multiple processes to be running simultaneously.
- (b) Security was simpler in time-sharing systems.
- (c) Many of the scheduling decisions could no longer be made by the operating system.
- (d) The response time was slower on a time-sharing system.

Question 3

[1 mark] Which of the following is an advantage of virtual machines?

- (a) Different operating systems can run simultaneously on the same real machine.
- (b) It is easy to copy a virtual machine when a new one is required.
- (c) Errors within a server running in one virtual machine are unlikely to affect other virtual machines on the same real machine.
- (d) All of the above.

Question 4

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

- (a) Paravirtualization is used to provide a number of virtual machines all sharing the same underlying kernel.
- (b) Application virtualization allows applications written for one operating system to run on a different operating system.
- (c) Type 2 hypervisors are special purpose operating systems optimised to run guest operating systems for enterprise virtualization.
- (d) OS level virtualization requires more resources and is less efficient than traditional trap and emulate virtualization.

Question 5

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

- (a) A thread blocking in the kernel doesn't stop all other threads in the same process.
- (b) Different threads can be scheduled on different processors in a multicore machine.
- (c) Each thread can be scheduled separately, rather than using the timeslice of one process over many threads.
- (d) All of the above.

Question 6

[1 mark] Which of the following best explains system-level thread implementations.

- (a) The operating system allocates and controls the threads. Information about the state of each thread is maintained inside the processes. The operating system sends regular interrupts to the processes which determine when thread scheduling changes should be allowed.
- (b) The operating system sees one thread per process. This makes it possible to run threads on the widest range of operating systems and can schedule each thread individually on multiple processors if available.
- (c) The operating system sees one thread per process. Code in the process, usually a library, allocates the threads and controls their scheduling, enabling greater flexibility. The operating system sends regular interrupts to the process which can be used to schedule a different thread.
- (d) The operating system allocates and controls the threads. It maintains information about the state of each thread and can schedule each thread individually on multiple processors if available.

Question 7

[1 mark] Which of the following is NOT a standard thread state?

- (a) Being created or initialized
- (b) Runnable, running or ready to run
- (c) Resting or recuperating
- (d) Waiting or blocked

Question 8

[1 mark] With original Linux threads the clone system call created a new process which shared all of its memory with the original process. The new process was then used as a thread of the original process. Which of the following statements was TRUE about original Linux threads?

- (a) If one thread in the process finished, all threads finished.
- (b) If a thread made a blocking system call, the other threads in the same process could still continue.
- (c) If one thread called exec to run another program, all threads in the same process would be killed.
- (d) All of the above

Question 9

[1 mark] Given three processes A, B, and C with corresponding burst times 10, 5, and 1, what is the average waiting time with a round-robin scheduler and a time slice of 5? All processes are ready to run at the same time but they arrived in the order A, B, then C.

- (a) 10
- (b) 5
- (c) 5.33...
- (d) 7

Question 10

[1 mark] Here is a simple Python program to run on a Linux system.

```
import os

i=0
while i < 2:
    print(i)
    os.fork()
    i += 1
```

How many times does the `print` function get called by the code above?

- (a) 2
- (b) 3
- (c) 6
- (d) 4

Question 11

[1 mark] Which of the following is a significant difference between the Windows `create process` function call and the way Unix creates a new process running a particular program?

- (a) Unix requires two system calls, a call to copy the current process and then a call to change the program running in that copied process.
- (b) Unix checks to see if the particular program is executable whereas Windows does not.
- (c) Windows creates an initial thread for the new process whereas Unix can create several threads in the new process.
- (d) Windows releases all memory currently held in the creating process whereas Unix maintains that memory until after the new process has been created.

Question 12

[1 mark] Which of the following problems can be helped with the use of a copy on write (COW) or similar mechanism?

- (a) Minimising the copying of data in a Unix process when the process is forked before a call to `exec`.
- (b) Minimising the down time of a process which is being migrated to another machine.
- (c) Minimising data movement between machines when implementing a distributed shared memory system.
- (d) All of the above.

This information is for the next 4 questions. Here is a portion of process information on a running Linux machine:

PID	PPID	STAT	SIZE	RSS	WCHAN	COMMAND
3366	2459	S+	600	1524	wait	man
3375	3366	S+	336	960	n_tty_	pager
3521	2398	S+	2992	5492	wait	python3
3522	3521	S+	328	652	wait	sh
3523	3522	R+	928	876	-	ps

Question 13

[1 mark] Which process is currently running?

- (a) ps
- (b) man
- (c) python3
- (d) sh

Question 14

[1 mark] Which process created the sh process?

- (a) man
- (b) ps
- (c) python3
- (d) pager

Question 15

[1 mark] Which process is currently occupying the smallest amount of real memory?

- (a) ps
- (b) sh
- (c) man
- (d) pager

Question 16

[1 mark] What does the WCHAN column represent?

- (a) The wait condition handler for the process.
- (b) The system call the process was executing when it last ran.
- (c) The address within the kernel where the process is waiting.
- (d) The owner of the process.

Question 17

[1 mark] Which of the following statements about the Unix fork system call is FALSE?

- (a) Most fork implementations no longer make an actual copy of the code of the parent process.
- (b) Fork stops the parent process running until the child process finishes.
- (c) The open file table of the parent process is copied into the child process by fork.
- (d) Fork makes a copy of the data of the parent process.

Question 18

[1 mark] Why are spin locks particularly bad when run on a single core processor?

- (a) As only one thread can run on a single core processor there is no need for a lock of any kind, including spin locks.
- (b) Single core processors get overheated more easily if a thread is running in a spin lock.
- (c) Because time-slice preemption is the only way another thread can run on a single core, it is possible for the spin lock to fail.
- (d) Running in the spin lock is pointless as the thread holding the lock cannot run and therefore release the lock at the same time.

Question 19

[1 mark] How does the Bakery algorithm guarantee a unique ordering of process requests to use a shared resource?

- (a) The algorithm uses the scheduling priority of each process as its unique ordering identifier.
- (b) The process id of the process is appended to the ticket number distributed by the algorithm.
- (c) The ticket number distributed by the algorithm is kept unique by the use of a spin lock.
- (d) A hash value is generated from the process id of the process and the ticket number distributed by the algorithm.

Question 20

[1 mark] Which of the following statements about semaphores is correct?

- (a) A binary semaphore can be used in the same way as a simple lock.
- (b) A semaphore is an integer count with some indivisible operations and an initialization.
- (c) Returning a resource when no process is waiting causes the semaphore value to increase.
- (d) All of the above.

Question 21

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

- (a) A process can hold a resource while requesting another.
- (b) Only the owner can release the resource.
- (c) Resources cannot be shared.
- (d) There is a circular list of resources each allocated to a process.

Question 22

[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 23

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

- (a) Deadlock prevention is impossible.
- (b) The Bakery algorithm is more complicated in a distributed system.
- (c) Deadlock detection is harder.
- (d) All of the above.

The next 5 questions are associated with the flat file system of assignment 2. If required, assume that the current output of the `ls -l A2dir` command before each question is:

```
total 0
-rw-rw-r-- 1 robert robert 0 Oct  5 10:05 -aa
-rw-rw-r-- 1 robert robert 0 Oct  5 10:05 -bb-cc
-rw-rw-r-- 1 robert robert 0 Oct  5 10:05 -bb-dd-ee
-rw-rw-r-- 1 robert robert 0 Oct  5 10:05 -bb-gg
-rw-rw-r-- 1 robert robert 0 Oct  5 10:05 -ff
-rw-rw-r-- 1 robert robert 0 Oct  5 10:05 -hh-ii
-rw-rw-r-- 1 robert robert 0 Oct  5 10:05 -hh-jj-kk
```

Question 24

[1 mark] How many directories would be seen by a user of this flat file system (include the root directory)?

- (a) 7
- (b) 5
- (c) 4
- (d) 6

Question 25

[1 mark] The following commands get executed:

```
ffs> cd
ffs> ls
```

Here are some possible outputs from these commands:

1	2	3
f: aa d: bb f: bb-cc d: bb-dd f: bb-dd-ee f: bb-gg f: ff d: hh f: hh-ii d: hh-jj f: hh-jj-kk	f: aa d: bb d: hh f: ff	f: aa f: bb-cc f: bb-dd-ee f: bb-gg f: ff f: hh-ii f: hh-jj-kk

Which of the outputs above is the correct one?

- (a) 2
- (b) Unknown as we don't know what the current directory is.
- (c) 3
- (d) 1

Question 26

[1 mark] The following commands get executed:

```
ffs> cd -bb-dd
ffs> cd ..
ffs> pwd
```

What is the output from the sequence of commands above:

- (a) -
- (b) -dd-
- (c) -bb-
- (d) -bb-dd-

Question 27

[1 mark] Several minutes later the following commands are executed:

```
ffs> cd -
ffs> add -bb-cc someInputText
ffs> rls
```

What differences would be seen between the output from this sequence of commands and the original output from the `ls -l A2dir` command?

- (a) There would be a new file in the listing called `-bb-cc-someInputText`.
- (b) The total info, the size of the `-bb-cc` file, the modified time of the `-bb-cc` file.
- (c) The total info, the size of `-bb-cc` file, the modified times of the `-bb-cc` and `-bb-` files.
- (d) The size of the `-bb-cc` file only.

Question 28

[1 mark] Which of the following statements about the flat file system of assignment 2 is TRUE?

- (a) Adding a new file embedded in several directories which don't exist yet is simpler in the flat file system than in a traditional Unix file system.
- (b) It is not possible to have a file with the same name in different directories of the flat file system.
- (c) Searching for a file with a long pathname in the flat file system is always easier than searching for a file with a long pathname in a traditional Unix file system.
- (d) None of the above.

Question 29

[1 mark] Where can file type information be stored?

- (a) In the file name.
- (b) Inside the file system but external to the file.
- (c) Inside the file as a magic number.
- (d) All of the above.

Question 30

[1 mark] What is the job of an open file system call?

- (a) To verify the rights the process has to access the file.
- (b) To record information as to which process is accessing the file.
- (c) To make the connection between a process file object or data structure and the external file.
- (d) All of the above.

Question 31

[1 mark] Which of the following statements relating to block size is FALSE?

- (a) Most files are small and hence can be completely stored in blocks of 8KB or less.
- (b) Large block sizes are more efficient for large files, especially on rotating disk devices.
- (c) The use of a small block size means a small amount of space is lost to internal fragmentation.
- (d) The use of a large block size means a large amount of space is lost to external fragmentation.

Question 32

[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 cannot easily be handled by a wide range of operating systems.
- (b) The text in a word processing document would normally be stored in the data fork.
- (c) The main use of NTFS alternate data streams is to hide extra data on the device.
- (d) All of the above.

Question 33

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

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

Question 34

[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) An extent holds the first 1K of a file followed by links to further extents.
- (b) Extents are the current versions of File Allocation Table (FAT) linked data structures.
- (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 s


C: itting_

D: tudyng for_

E: my Operating System exam at the moment

F: have finished my degree

G: ._Yeah!



Question 35

[1 mark] Using the tree above what is the file data for version 3 of the file?

- (a) have finished my degree
- (b) Dear Mum, I_am studying for_my Operating System exam at the moment._Yeah!
- (c) Dear Mum, I_have finished my degree._Yeah!
- (d) Dear Mum, I_am sitting_my Operating System exam at the moment._Yeah!

Question 36

[1 mark] Which of the following answers is a list of operations or data NONE of which are required to be implemented or kept at the file server in a pure stateless distributed file system?

- (a) A block write operation, and a block read operation.
- (b) An open file operation, file locking information, and stored file position information.
- (c) An open file operation, a block read operation, and stored file position information.
- (d) A block read operation, file locking information, and access control information.

Question 37

[1 mark] What is the detection window in a self-securing storage system?

- (a) The period of time during which all modifications to the system are recorded. Earlier changes are lost due to the pruning of storage.
- (b) An area of data which if modified by a possible attack on the system causes an alarm to be sent to the system administrators notifying them of the attack.
- (c) An indicator that the file system has failed and some of the files have been lost.
- (d) None of the above.

Question 38

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

- (a) Access to files under NFS is mostly done with caching.
- (b) NFS has difficulty working in a heterogeneous environment.
- (c) NFS file location information is stored in volume location databases.
- (d) Directory name lookups are expensive when the mount point to a remote directory is crossed.

Here is a figure from the textbook section on NFS. Use it for the next 2 questions.

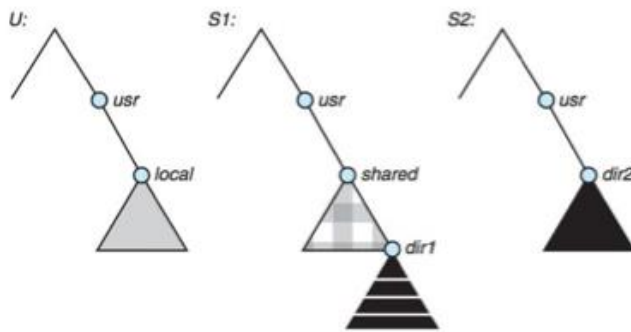


Figure 12.13 Three independent file systems.

Question 39

[1 mark] If $S1:/usr/shared$ is mounted over $U:/usr/local$ what is the pathname on U to files in $S1:/usr/shared/dir1$?

- (a) $/usr/local/dir1$
- (b) $/usr/shared/dir1$
- (c) $/usr/dir1$
- (d) $/usr/local/shared/dir1$

Question 40

[1 mark] If $S2:/usr/dir2$ is then mounted over $U:/usr/local/dir1$ what is the pathname on U to files in $S1:/usr/shared/dir1$?

- (a) $/usr/shared/dir1$
- (b) $/usr/local/dir1$
- (c) $S1:/usr/shared/dir1$ is now inaccessible from U .
- (d) $/usr/local/dir2$

Question 41

[1 mark] There were two different formulas for working out the access time of memory in lectures.

(1) The average access time of memory using page tables, given a TLB with all pages in real memory and (2) The effective access time of memory given that some of the pages are swapped out and the cost of swapping them back in again. Which of the following statements about these formulas is FALSE?

- (a) The access time of (1) is far smaller than the access time of (2).
- (b) The average access time of (1) is included as a factor in (2).
- (c) The most important factor in (1) is the TLB hit ratio.
- (d) The time of swapping a page in once a free frame is found is negligible in (2).


Question 42

[1 mark] Why are page sizes powers of 2?

- (a) The page number is determined by the n most significant bits of the virtual address. So the page size is 2^n .
- (b) It doesn't have to be but manufacturers only make memory in those sizes.
- (c) The displacement within the page is determined by the n least significant bits of the virtual address. So the page size is 2^n .
- (d) Pages need to be repeatedly able to be halved.

Question 43

[1 mark] The figure below is from the textbook. It shows the number of page faults generated by a process over time. Which of the following is a correct interpretation of the graph?



- (a) The peaks indicate a time when the process is demand paging a new locality.
- (b) The process has moved to three different localities, either in data or code space, then settled in for a time before moving again.
- (c) The sections with a small number of page faults after the peaks indicate that pages for the current working set have been allocated to the project.
- (d) All of the above.

The next four questions use this page reference string and a memory of 4 frames:

1,7,3,1,2,4,3,5,6,7,7,1,6,5,4,6,2,6,1,7

Each table shows the contents of the frames using a variety of different page replacement algorithms using the reference string. As in assignment 3 a zero means the frame is free, an "=" means the content of the frame is the same as in the previous step.

Question 44

[1 mark] What algorithm was used to generate the table below?

	1	7	3	1	2	4	3	5	6	7	7	1	6	5	4	6	2	6	1	7
0	1	=	=	=	=	=	=	=	=	=	=	=	=	=	=	=	=	=	=	=
0	=	7	=	=	=	=	=	=	=	=	=	=	=	=	=	=	=	=	=	=
0	=	=	3	=	=	=	=	5	=	=	=	=	=	=	4	=	2	=	=	=
0	=	=	=	=	2	4	=	=	6	=	=	=	=	=	=	=	=	=	=	=

- (a) Optimal
- (b) Least Frequently Used
- (c) First In First Out
- (d) Least Recently Used

Question 45

[1 mark] What algorithm was used to generate the table below?

	1	7	3	1	2	4	3	5	6	7	7	1	6	5	4	6	2	6	1	7	
0	1	=	=	=	=	=	=	5	=	=	=	=	=	=	=	=	=	=	=	1	=
0	=	7	=	=	=	4	=	=	=	7	=	=	=	=	4	=	=	=	=	=	7
0	=	=	3	=	=	=	=	=	=	=	=	1	=	=	=	=	2	=	=	=	=
0	=	=	=	=	2	=	=	=	6	=	=	=	=	=	=	=	=	=	=	=	=

- (a) Optimal
- (b) Least Recently Used
- (c) Least Frequently Used
- (d) First In First Out

Question 46

[1 mark] What algorithm was used to generate the table below?

	1	7	3	1	2	4	3	5	6	7	7	1	6	5	4	6	2	6	1	7	
0	1	=	=	=	=	4	=	=	=	=	=	1	=	=	=	=	=	=	=	=	7
0	=	7	=	=	=	=	=	5	=	=	=	=	=	=	4	=	=	=	=	=	=
0	=	=	3	=	=	=	=	=	6	=	=	=	=	=	=	=	2	=	=	=	=
0	=	=	=	=	2	=	=	=	=	7	=	=	=	=	=	=	=	6	=	=	=

- (a) Least Recently Used
- (b) Optimal
- (c) First In First Out
- (d) Least Frequently Used

Question 47

[1 mark] What algorithm was used to generate the table below?

	1	7	3	1	2	4	3	5	6	7	7	1	6	5	4	6	2	6	1	7
0	1	=	=	=	=	=	=	=	=	=	=	=	=	=	=	=	=	=	=	=
0	=	7	=	=	=	4	=	=	6	=	=	=	=	=	=	=	=	=	=	=
0	=	=	3	=	=	=	=	=	=	=	=	=	=	5	4	=	2	=	=	=
0	=	=	=	=	2	=	=	5	=	7	=	=	=	=	=	=	=	=	=	=

- (a) Least Frequently Used
- (b) Least Recently Used
- (c) First In First Out
- (d) Optimal

Here is a snippet from a C program with two ways to initialise an array:

```
int A[][] = new int[256][256];

for (int i = 0; i < 256; i++) /* Loop 1 */
    for (int j = 0; j < 256; j++)
        A[i][j] = 0;

for (int j = 0; j < 256; j++) /* Loop 2 */
    for (int i = 0; i < 256; i++)
        A[i][j] = 0;
```

Question 48

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

- (a) Loop 1 has smaller working sets than Loop 2.
- (b) Loop 1 has better locality of reference than Loop 2 when considering code memory access.
- (c) Loop 1 has better locality of reference than Loop 2 when considering data memory access.
- (d) Loop 1 is preferable to Loop 2 because it would cause fewer page faults.

Question 49

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

- (a) Some virtual address spaces are smaller than the real address space.
- (b) The use of virtual memory slows down the average memory access time.
- (c) Some virtual address spaces can be larger than the real address space.
- (d) A virtual memory system without a TLB is faster than one with a TLB.

Question 50

[1 mark] The second chance or clock algorithm uses which of the following bits associated with page information to approximate least recently used selection?

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

object domain	F_1	F_2	F_3	laser printer	D_1	D_2	D_3	D_4
D_1	read		read			switch		
D_2				print			switch	switch control
D_3		read	execute					
D_4	write		write		switch			

This is an access matrix from the textbook, use it for the next 2 questions.

Question 51

[1 mark] Which of the following sequences of commands would allow a process running in D_4 to read from F_2 ?

- (a) write new code to F_3 (this code can read from D_3), execute F_3 (which reads F_2)
- (b) switch to D_1 , switch to D_2 , control D_4 (add read right to D_4), switch to D_4 , read F_2
- (c) switch to D_1 , switch to D_2 , switch to D_3 , read F_2
- (d) None of the above.

Question 52

[1 mark] Which domains can directly or indirectly remove the ability to write to F_1 from D_4 ?

- (a) D_1, D_2, D_4
- (b) D_1, D_2
- (c) D_4
- (d) D_2

Question 53

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

- (a) A program with privileges mistakenly prevents access to a resource which should be available.
- (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 is tricked into misusing its authority.

Question 54

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

- (a) Kerberos tickets are time expiring capabilities.
- (b) Ticket Granting Servers don't have to be trusted.
- (c) Kerberos is based on the Needham-Schroeder protocol.
- (d) If the Kerberos Authentication Server is down, no one can log in.

Question 55

[1 mark] When are dictionary attacks most likely to succeed?

- (a) If the password checking mechanism allows an unlimited number of tries.
- (b) If the password size is limited to 10 characters.
- (c) If the password must be changed every month.
- (d) If the encrypted password file can be accessed.

Question 56

[1 mark] Given the following listing from a Unix system, what access rights does the user joan who is in the staff group have to thefile?

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

- (a) no access
- (b) execute only
- (c) write only
- (d) write and execute

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

```
crw-rw-rw-   1 root root      1,   3 Oct  7 15:43 null
crw-rw-rw-   1 root tty       5,   0 Oct  7 15:43 tty
crw--w----   1 root tty       4,   0 Oct  7 15:43 tty0
crw-----   1 root root      5,   3 Oct  7 15:43 ttyprintk
crw-rw----   1 root tty       4,   5 Oct  7 15:43 tty5
crw-rw-rw-   1 root root      1,   5 Oct  7 15:43 zero
```

Question 57

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

- (a) null and zero
- (b) tty and ttyprintk,
- (c) tty0 and tty5
- (d) All of the above.

Question 58

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

- (a) Bus interfaces make it easier to attach a wide range of devices to current computers.
- (b) An IO request block (or IO request packet) is constructed by a device driver and passed to the IO system calls.
- (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 59

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

- (a) User level device drivers allow multiple drivers to be used by a device at the same time.
- (b) User level device drivers can more easily access interrupts using non-privileged code.
- (c) Problems with a user level device driver are outside the kernel, hence less likely to cause widespread damage.
- (d) User level device drivers are faster.

Question 60

[1 mark] Which of the following disk scheduling algorithms prevents new disk requests from slowing down older ones during a traversal of the disk surface?

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

Rough Working – This page will not be marked

