UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA

Department of Computer Science
Faculty of Engineering, Built Environment & IT
University of Pretoria

COS222
Operating Systems

Semester Test 1

25th August 2016

Examiners

Stefan Gruner and Vreda Pieterse

Instructions

1. Read the question paper carefully and answer all the questions that follows.
2. The examination opportunity comprises of **54** questions on **11** pages.
3. You have **120** minutes to complete the paper.
4. This is a closed book paper. You may therefore **not** have any study material with you.
5. Please **switch off** your cell phone, and keep it off for the duration of the paper.
6. You may **not** use a pocket calculator or any other electronic devices. All calculaters/cellphones and other electronic devices – including hearing aids – should be stored away.
7. Answer Section 1 **on side 2 of the answer sheet** and Section 2 in the answer book.
8. Please write the question numbers 46 to 53 in the appropriate column on the cover page of your answer book
9. Good Luck!

SECTION 1: Multiple choice questions
Answer on side 2 of the answer sheet

Each of the questions in this section counts 1 mark, 45 marks total for this section. For each question there is only one correct answer. Some questions may seem to have more than one correct answer. In such a case, you should select the option that is most correct or more convincing.

1. The four main structural elements of a computer system are: (1)
   A. Processor, Main Memory, I/O Modules and System Bus
   B. Processor, I/O Modules, System Bus and Secondary Memory
   C. Processor, Registers, Main Memory and System Bus
   D. Processor, Registers, I/O Modules and Main Memory
   E. Registers, Main Memory, System Bus, I/O Modules

2. Instruction processing consists of two steps: (1)
   A. start and halt
   B. fetch and execute
   C. instruction and execute
   D. instruction and halt
   E. fetch and instruction

3. The _____ determines the nature of the interrupt and performs whatever actions are needed. (1)
   A. execute instruction
   B. instruction signal
   C. program handler
   D. interrupt signal
   E. interrupt handler

4. Small, fast memory located between the processor and main memory is called: (1)
   A. Block memory
   B. Warm memory
   C. Cache memory
   D. Direct memory
   E. WORM memory

5. In a uniprocessor system, multiprogramming increases processor efficiency by: (1)
   A. Taking advantage of time wasted by long wait interrupt handling
   B. Disabling all interrupts except those of highest priority
   C. Eliminating all idle processor cycles
   D. Underclocking
   E. Overclocking

6. Operating systems must evolve over time because: (1)
   A. advances in software require upgrades
   B. hardware must be replaced when it fails
   C. hardware is hierarchical
   D. new hardware is designed and implemented in the computer system
   E. users will only purchase software that has a current copyright date

7. A user program executes in a _____ , in which certain areas of memory are protected from the user's use, and in which certain instructions may not be executed. (1)
    A. kernel mode
    B. user mode
    C. task mode
    D. inner mode
    E. batch mode

8. The technique where a system clock generates interrupts, and at each clock interrupt the OS regains control and assigns the processor to another user, is _____ . (1)
    A. time slicing
    B. multi-threading
    C. round robin
    D. clock cycle
    E. interrupt handler

9. _____ is concerned with the proper verification of the identity of users and the validity of messages or data. (1)
    A. Availability
    B. Confidentiality
    C. Privacy
    D. Authenticity
    E. Data integrity

10. The Windows OS services, the environment subsystems, and the applications are structured using the _____ computing module. (1)
    A. monolithic
    B. hierarchical
    C. layered
    D. object oriented
    E. client/server

11. Which is the operating system's most important data structure for process management? (1)
    A. The process status flag
    B. The process control flag
    C. The process control block
    D. The process status block
    E. The process control list

12. Indicate the state which does *not* belong to the classical five-state process model (1)
    A. Blocked
    B. Exit
    C. Running
    D. Kernel
    E. New
    F. Ready

13. Which unit of the operating system gets active immediately before a process starts running? (1)

    A. The CPU
    B. The RAM Driver
    C. The Dispatcher
    D. The Program Counter
    E. The Kernel Interface
    F. The I/O Manager
    G. The Interrupt
    H. The Memory Table
    I. The Virtual Machine

14. As there are typically very many processes active in a large computer system, the operating system typically holds tables of lists as its internal data structures for process managements, whereby different lists correspond to the various states in which processes can be. Assume a classical hardware architecture with a mono-processor. Now an operating system design engineer is thinking about the optimal length of the list data structure for the running processes. What advice about the optimal list-lenght would you give to the designer? (1)

    A. List length 1
    B. List length 10
    C. List length 100
    D. List length 1‘000
    E. List length 10‘000
    F. List length 100‘000
    G. List length 1‘000‘000
    H. List length 10‘000‘000

15. What is the thread:process relation in the operating system TRIX? (1)

    A. 1:1
    B. N:1
    C. 1:M
    D. N:M

16. What is the thread:process relation in the operating system Solaris? (1)

    A. 1:1
    B. N:1
    C. 1:M
    D. N:M

17. What is the thread:process relation in a traditional UNIX system? (1)

    A. 1:1
    B. N:1
    C. 1:M
    D. N:M

18. What is the thread:process relation the operating system Emerald? (1)

    A. 1:1
    B. N:1
    C. 1:M
    D. N:M

19. What is the main advantage of using user-level threads (ULT) instead of kernel-level threads (KLT)? (1)

    A. Scheduling can be omitted
    B. Scheduling can be application-specific
    C. Thread-scheduling can be done in hardware, which is fastest
    D. Interrupts are no longer needed, which leads to higher system-reliability
    E. The operating system does not need any kernel any more

20. What is the main disadvantage of a pure ULT design (without any KLT)? (1)

    A. The user needs a Micosoft Certificate, which not many users actually have
    B. The system can then easily come into a dangerous state of running too fast (a.k.a. over-clocking)
    C. The scheduler will then need too many different priority-queues for the various ULT priorities
    D. The system gets in danger of wrongly mixing up the various threads from different users
    E. The user must manually design the data structures for the thread control blocks (TCBs)
    F. Multi-threaded applications cannot effectively utilise multiprocessing

21. In which operating system is a ‘fiber’ a unit of execution which must be scheduled by an application itself? (1)

    A. Solaris
    B. Windows 8
    C. Linux
    D. Android
    E. Mac OS

22. In which operating system are the two states ‘Interruptible’ and ‘Uninterruptible’ distinguished from each other in the state-transition model? (1)

    A. Solaris
    B. Windows 8
    C. Linux
    D. Android
    E. Mac OS

23. Which operating system has the concept of ‘Lightweight Processes’? (1)

    A. Solaris
    B. Windows 8
    C. Linux
    D. Android
    E. Mac OS

24. With reference to the previous question: Where in the system architecture are those ‘Lightweight Processes’ functionally located? (1)

    A. Between the processes and the user threads
    B. Between the user threads and the kernel threads
    C. Between the OS kernel and the hardware
    D. Between the processes and the compiler
    E. Between the compiler and the virtual machine
    F. Between the kernel threads and the system calls
    G. Between the virtual machine and the Internet

25. In which operating systems do namespaces enable a process to have a view of the system that differs from the views which other processes have of the system? (1)

    A. Android
    B. MySpace Web-OS
    C. Mac OS
    D. Linux
    E. Windows 4

26. What is ‘thread jacketing’? (1)

    A. It does not exist at all; that is only a silly trick by Professor Gruner to confuse the students
    B. It is a programming trick to convert blocking system calls into non-blocking system calls
    C. It is a programming trick to convert non-blocking system calls into blocking system calls
    D. It is a programming trick to prevent insecure thread-access on the vulnerable hardware
    E. It is a programming trick to wrap kernel-level threads into user-level threads

27. The management of multiple processes within a uniprocessor system is _____ . (1)

    A. multiprogramming
    B. structured applications
    C. uniprocessing
    D. distributed processing
    E. multiprocessing

28. A situation in which a runnable process is overlooked indefinitely by the scheduler, although it is able to proceed, is _____ . (1)

    A. impossible
    B. mutual exclusion
    C. deadlock
    D. starvation
    E. livelock

29. The requirement that when one process is in a critical section that accesses shared resources, no other process may be in a critical section that accesses any of those shared resources is _____ . (1)

    A. critical section
    B. livelock
    C. shared access
    D. mutual exclusion
    E. atomic operation

30. A means for two processes to exchange information is with the use of _____ . (1)

    A. monitors
    B. spinlocks
    C. event flags
    D. condition variables
    E. messages

31. A semaphore that does not specify the order in which processes are removed from the queue is a _____ semaphore. (1)

A. spinning
B. weak
C. general
D. strong
E. binary

32. A _____ occurs when multiple processes or threads read and write data items so that the final result depends on the order of execution of instructions in the multiple processes. (1)

A. critical error
B. atomic operation
C. race condition
D. livelock
E. deadlock

33. A _____ is an integer value used for signaling among processes. (1)

A. boolean
B. semaphore
C. message
D. mutex
E. atomic operation

34. _____ is when the sequence of instruction is guaranteed to execute as a group, or not execute at all, having no visible effect on system state. (1)

A. Group block
B. Critical section
C. Mutual exclusion
D. Atomic operation
E. Starvation

35. _____ are memory words used as a synchronization mechanism. (1)

A. Synchronizer words
B. Semaphores
C. Event flags
D. Counting semaphores
E. Mailboxes

36. The term _____ refers to a technique in which a process can do nothing until it gets permission to enter its critical section but continues to execute an instruction or set of instructions that tests the appropriate variable to gain entrance. (1)

A. dead wait
B. spin waiting
C. general semaphore
D. critical resource
E. message passing

37. A semaphore whose definition includes the policy that the process that has been blocked the longest is released from the queue first is called a _____ semaphore. (1)

    A. binary
    B. general
    C. strong
    D. weak
    E. counting

38. Probably the most useful combination, _____ allows a process to send one or more messages to a variety of destinations as quickly as possible. (1)

    A. blocking send, blocking receive
    B. nonblocking send, blocking receive
    C. nonblocking send, nonblocking receive
    D. blocking send, nonblocking receive

39. Assume that some operating system 'X' provides the possibility of resource preemption. What is the consequence? (1)

    A. Deadlock can no longer be prevented.
    B. Deadlock is guaranteed.
    C. Starvation is possible.
    D. Starvation is guaranteed.
    E. The system becomes much faster on average.
    F. Scheduling can be prevented.
    G. Scheduling is in most cases not needed any more.
    H. The average size of the thread pool is growing.
    I. The average size of the thread pool is shrinking.
    J. Processes can no longer enter into their critical sections.

40. After a deadlock has been discovered, (1)

    A. we delete all processes, to be on the safe side (because of the zombies)
    B. we first delete the encircled process of highest priority
    C. we first delete the encircled process of lowest priority
    D. we re-install the operating system in its default settings (from factory)
    E. we delete only the encircled process control blocks, without actually doing any harm to the processes themselves.
    F. we transform the deadlock into a less harmful livelock, from which the system can continue by itself (automatic recovery)
    G. we release the spin-locks (re-set to state zero), such that the waiting processes are no longer blocked.
    H. we install the additionally needed resources on-the-fly.

NOTE: *The remaining questions, forms a* **question-cluster** *on the topic of the Banker algorithm. Answer the questions while referring to the following example. Subsequent questions should take the consqences of preceding questions into account.*

R is the only type of resource. The *Banker* holds 14 instances of R. The following table shows several processes P_i together with their *current* allocations and their *maximal* needs.

process	current allocation	maximal need
P_1	$2R$	$9R$
P_2	$3R$	$13R$
P_3	$6R$	$7R$
P_4	$2R$	$11R$

41. The configuration of above is (1)
    A. circular
    B. unsafe
    C. undefined
    D. safe
    E. semi-safe
    F. almost safe
    G. almost unsafe

42. With regard to the resource allocation table of above, process P_2 asks the Banker for one more resource, whereupon which (1)
    A. the banker aborts process P_2
    B. the banker grants the request
    C. the banker denies the request
    D. the banker postpones the request until later
    E. the banker aborts process P_4
    F. the banker aborts process P_1
    G. the banker calls the scheduler
    H. the banker calls the monitor
    I. the banker triggers an interrupt
    J. the banker releases the semaphore

43. With regard to the resource allocation table of above, process P_1 asks the Banker for one more resource, whereupon which (1)
    A. the banker aborts process P_1
    B. the banker grants the request
    C. the banker denies the request
    D. the banker postpones the request until later
    E. the banker aborts process P_2
    F. the banker aborts process P_3
    G. the banker calls the scheduler
    H. the banker calls the monitor
    I. the banker triggers an interrupt
    J. the banker releases the semaphore

44. With regard to the resource allocation table of above, process P_3 asks the Banker for one more resource, whereupon which (1)

A. the banker aborts process P_4
B. the banker grants the request
C. the banker denies the request
D. the banker postpones the request until later
E. the banker aborts process P_2
F. the banker aborts process P_3
G. the banker calls the scheduler
H. the banker calls the monitor
I. the banker triggers an interrupt
J. the banker releases the semaphore

45. In any circumstances, the situation *after* the Banker has granted *any* request is (1)

A. frozen
B. unknown
C. unsafe
D. undefined
E. safe
F. semi-safe
G. possibly circular
H. preempted
I. no longer amenable to resource-sharing
J. no longer in need of further mutual exclusions

SECTION 2: Answer all questions in the answer book

Please write the question numbers 46 to 53 in the appropriate column on the cover page of your answer book.

46. Draw a diagram to illustrate the instruction cycle with interrupts (6)

47. Explain the difference between a monolithic kernel and a microkernel. (5)

48. Draw an accurate picture of a typical 6-state process model which includes a 'Suspended' state. Your picture must show the names of all 6 states, and also the correctly labeled transitions between the states. (6)

49. Explain briefly the four key benefits of multi-threading, with regard to runtime performance. (4)

50. To demonstrate the advantage of multi-threading, draw two time-line diagrams of a process P which makes two remote procedure calls (RPC) to an external server S, whereby the calling entity cannot continue while waiting for the response from the external server.
    (a) Draw the time-line diagram for a single-threaded caller. (1)
    (b) Draw the time-line diagram for a two-threaded caller (3)
    (c) Write down the mathematical formula (equation) of the speedup law. (3)

51. Consider the following implementation of the `semWait()`-function that is part of the implementation of a semaphore using the atomic `compare_and_swap` instruction:

```
semWait(s)
{
    while(compare_and_swap(s.flag, 0, 1) == 1)
        /* do nothing */;
    s.count--;
    if(s.count < 0)
    {
        /* place this process in s.queue */;
        /* block this process */
    }
    s.flag = 0;
}
```

    (a) Explain the `compare_and_swap` instruction (3)
    (b) Write code to implement a `semSignal()`-function for this semaphore using the atomic `compare_and_swap` instruction (6)

52. List the four necessary deadlock conditions (4)

53. (a) List the three main strategies of dealing with deadlock (3)
    (b) *For each* of the three strategies of above, list (6)
        - *one* major *advantage*, and
        - *one* major *disadvantage*.

54. Answering this question is voluntary. Choosing to answer or not will not have any bearing on your marks. By answering the question you consent to your data anonymously be used in reaserch about self-evlauation and time-tracking in learning. Students who answer the question qualify for a **lucky draw for R500** at the end of the semester. All time-sheets and mark estimations we recieve will be included in the lucky draw. You are encourged to keep your time-tracking sheets up to date and be prepapered to submit them with the written exam at the end of the sememester. (0)

    Redraw the following table to estimate your marks for this paper

Question:	46	47	48	49	50	51	52	53	Total
Marks (Section 2)	6	5	6	4	7	9	4	9	50
Score:									