國立成功大學

113學年度碩士班招生考試試題

編　號：196

系　所：電機資訊學院-資訊聯招

科　目：計算機組織與系統

日　期：0201

節　次：第 1 節

備　註：不可使用計算機

國立成功大學 113 學年度碩士班招生考試試題

系　　所：電機資訊學院-資訊聯招

考試科目：計算機組織與系統

考試日期：0201，節次：1

※ 考生請注意：本試題不可使用計算機。 請於答案卷(卡)作答，於本試題紙上作答者，不予計分。

1. [20%] Please select a correct answer for each of the following questions.

(1) [2%] What statement concerning privileged instructions is considered false?

(a). They may cause harm to the system.

(b). They can only be executed in kernel mode.

(c). They cannot be attempted from user mode.

(d). They are used to manage interrupts.

(2) [2%] A message-passing model is ____.

(a). easier to implement than a shared memory model for intercomputer communication

(b). faster than the shared memory model

(c). a network protocol, and does not apply to operating systems

(d). only useful for small simple operating systems

(3) [2%] The major difficulty in designing a layered operating system approach is ____.

(a). appropriately defining the various layers

(b). making sure that each layer hides certain data structures, hardware, and operations from higher-level layers

(c). debugging a particular layer

(d). making sure each layer is easily converted to modules

(4) [2%] A boot block ____.

(a). typically only knows the location and length of the rest of the bootstrap program

(b). typically is sophisticated enough to load the operating system and begin its execution

(c). is composed of multiple disk blocks

(d). is composed of multiple disk cylinders

(5) [2%] _____ allow operating system services to be loaded dynamically.

(a). Virtual machines

(b). Modules

(c). File systems

(d). Graphical user interfaces

(6) [2%] A cycle in a resource-allocation graph is ____.

(a). a necessary and sufficient condition for deadlock in the case that each resource has more than one instance

(b). a sufficient condition for a deadlock in the case that each resource has more than once instance

(c). is neither necessary nor sufficient for indicating deadlock in the case that each resource has exactly one instances

(d). a necessary and sufficient condition for a deadlock in the case that each resource has exactly one instance

(7) [2%] Absolute code can be generated for ____.

(a). compile-time binding

(b). load-time binding

(c). execution-time binding

(d). interrupt binding

(8) [2%] In a dynamically linked library, _____.
(a). loading is postponed until execution time
(b). system language libraries are treated like any other object module
(c). more disk space is used than in a statically linked library
(d). a stub is included in the image for each library-routine reference

(9) [2%] The vfork() system call in UNIX _____.
(a). allows the child process to use the address space of the parent
(b). uses copy-on-write with the fork() call
(c). is not intended to be used when the child process calls exec() immediately after creation
(d). duplicates all pages that are modified by the child process

(10) [2%] Belady's anomaly states that _____.
(a). giving more memory to a process will improve its performance
(b). as the number of allocated frames increases, the page-fault rate may decrease for all page replacement algorithms
(c). for some page replacement algorithms, the page-fault rate may decrease as the number of allocated frames increases
(d). for some page replacement algorithms, the page-fault rate may increase as the number of allocated frames increases

2. [10%] Please answer the following True/False questions.
(1) [1%] Solid state disks are considered volatile storage.
(2) [1%] The process runs with effective UID all the time until the process is stopped.
(3) [1%] Mac OS X is a hybrid system consisting of both the Mach microkernel and BSD UNIX.
(4) [1%] Many operating system merge I/O devices and files into a combined file because of the similarity of system calls for each.
(5) [1%] The communication with named pipes requires parent-child relationship.
(6) [1%] All access to POSIX shared memory requires a system call.
(7) [1%] init is the very first process for a typical Linux system.
(8) [1%] One-to-one model provides more concurrency than the many-to-one model by allowing another thread to run when a thread makes a blocking system call.
(9) [1%] The trend in developing parallel applications is to use implicit threading.
(10) [1%] Grand Central Dispatch requires multiple threads.

3. [5%] Given the following snapshot of the system: Please determine whether there exists a safe sequence. (Hint: You need to create the Need array before you examine it)

	Allocation			Max			Available		
	A	B	C	A	B	C	A	B	C
P0	1	1	2	6	5	3	2	1	2
P1	0	0	2	3	4	2			
P2	4	0	0	8	5	2			
P3	0	3	0	1	4	1			
P4	1	3	3	4	3	3			

4. [15%] Please answer the following questions for the demand paging:
   (1) [5%] What is the maximum number of instruction re-starts when the following instruction is executed: Add A, B, C (which means C = A+B). Suppose the page which contains the instruction is already in the DRAM.
   (2) [5%] If the indirect addressing is considered, what is the minimum number of frames to execute an instruction ADD (A), (B), (C)? Note that (X) is an indirect address.
   (3) [5%] Modern systems tend to use much larger page sizes since the CPU speed and the memory capacity grow much faster than the disk speed. Please show three advantages and two disadvantages of adopting the larger page size in the systems. Brief explanation is required.

5. [20%] Please answer the following True/False questions.
   (1) [2%] In order to alleviate the memory wall problem, a small, fast memory called cache memory, which is built with the dynamic random access memory (DRAM) technology, is used in the processor and acts as a DRAM buffer.
   (2) [2%] There are many metrics and they are good for measuring and comparing the performance delivered by different processors, including execution time, CPU time, million instructions per second (MIPS), and CPU clock cycles.
   (3) [2%] A general-purpose register (GPR) architecture provides the GPRs for storing addresses or data for instructions and it can be found on many processor architectures, such as Intel 8086, MIPS, and RISC-V.
   (4) [2%] Parallel execution strategies that work great for integer data also work for floating-point data. This means that computer arithmetic generates the same results when going from sequential to parallel.
   (5) [2%] Pipelining techniques can be implemented independent of process technology. For instance, the delayed branch is a simple solution to alleviate the overhead caused by control hazards, and it is very useful for processors with longer pipelines.
   (6) [2%] The high-level abstraction provided by programming languages, such as C and C++, allows the programmers and compilers to produce a high-performance code without considering memory system behavior of a computer.
   (7) [2%] Operating Systems (OS) is responsible for serving the disk accesses requested by user-space programs. This makes OS the best place to schedule the disk accesses for improving the performance.
   (8) [2%] The instruction number of a program and the clock cycles required for each instruction are independent.
   (9) [2%] The performance delivered by a processor architecture is usually improved by lowering the clock rate of the processor.

(10) [2%] The performance of a compute-intensive program is dominated by the operating frequency of a processor. That is, the program will execute faster on a machine with a processor running at 1 GHz clock than on a machine with another processor running at 800 MHz clock.

6. [30%] Please evaluate the performance of the loop below on different machines. The nested loops perform the operations on the integer arrays, A and B, where each array element is a 32-bit integer.

```
for (i=0; i<1000; i++)
    for (j=0; j<1000; j++)
        A[i][j] = B[j][i] * 100 + 20;
```

(1) [10%] What is the common optimization technique that can be used by a compiler to generate a more efficient version of the code to alleviate branch penalties during the execution? Please provide the optimized code generated by the compiler.

(2) [10%] We assume a machine α with a 64-bit dual-core processor. What is the ideal speedup of the above code running in parallel on the machine α, compared with the sequential execution version? Why? (Hint: There is no control hazard, and memory miss overhead is negligible.)

(3) [10%] We assume another machine β with a single-core processor with a SIMD engine. The SIMD engine is equipped with 128-bit data registers for data processing. What is the performance ratio of the above code running with and without the support of the SIMD engine? Why? (Hint: There is no control hazard, and memory miss overhead is negligible.)