

All questions are multiple-choice and the answer can be empty, i.e., no choice is correct.

1. (5 points) Suppose we have the following state in Banker's Algorithm for deadlock avoidance.

processes	Allocation			Max			Available		
	A	B	C	A	B	C	A	B	C
P1	0	0	1	4	0	1	3	5	2
P2	1	0	0	1	7	5			
P3	1	3	5	2	4	5			
P4	0	6	3	1	6	5			

Which of the following is/are correct?

- (a) After granting the request (2,0,0) by P1, the system is still safe.
- (b) After granting the request (0,3,2) by P2, the system is still safe.

2. (5 points) Continuing from the last question, which of the following is/are correct?

- (a) After granting the request (0,4,1) by P2, the system becomes unsafe.
- (b) After granting the request (1,0,0) by P3, the system becomes unsafe.
- (c) After granting the request (0,0,2) by P4, the system becomes unsafe.

3. (5 points) The following is the critical section solution by Leslie Lamport in 1974.

```

(1) do {
(2)     choosing[i]=true;
(3)     number[i]=max(number[0], ...number[n-1]);
(4)     choosing[i]=false;
(5)     for (j=0; j < n; j++) {
(6)         while choosing[j] ;
(7)         while (number[j] != 0 && (number[j],j)≤(number[i],i)) ;
(8)     }
(9)     critical section
(10)    number[i]=0;
(11)    remainder section
(12) } while (1);
    
```

The parenthesized numbers on the left are the line numbers. Please answer which of the following statements is/are correct!

- (a) Line (3) is incorrect!
- (b) Line (7) is incorrect!
- (c) Line (10) is incorrect!

4. (5points) Continuing from the last question, for the correct version of Lamport's Bakery algorithm, please answer which of the following statements is/are correct ?

- (a) The solution relies on OS or hardware support for the max operation in line (3).
- (b) The solution is incorrect for distributed implementation across internet.
- (c) The entry section complexity is $O(3n+2)$ queries to peer processes.

見背面

5. (5 points) Please answer which of the following statements about disk array is/are correct!
- (a) RAID 1 consists of exact copies of the same data on two or more disks.
 - (b) RAID 2 uses block-level stripe data.
 - (c) RAID 2 uses parity bit for error detection.
6. (5 points) Continuing from the last question, please answer which of the following statements is/are correct!
- (a) RAID 3 uses byte-level striping with a dedicated parity disk.
 - (b) RAID 4 consists of bit-level striping with a dedicated parity disk.
 - (c) RAID 5 consists of bit-level striping with distributed parity.
7. (5 points) Consider the following page reference sequence in a demand paging system. Assume that initially all the frames of the process are empty.
- 1, 2, 3, 4, 2, 1, 5, 6, 8, 1, 2, 3, 7, 6, 3, 2, 4, 2, 3, 6.
- Please answer which of the following statements is/are correct!
- (a) If the process is allocated 4 frames, the LRU algorithm incurs 10 page faults.
 - (b) If the process is allocated 6 frames, the optimal algorithm incurs 8 page faults.
 - (c) If the process is allocated 5 frames, the FIFO algorithm incurs 13 page faults.
8. (5 points) Continuing from the last question, please answer which of the following statements is/are correct!
- (a) If the process is allocated 4 frames, the reference-bit algorithm incurs 15 page faults. Assume that we start searching from the frame next to the frame with the last brought in page. If the page last brought in is in the last frame, then we start searching from the first frame.
 - (b) If the process is allocated 5 frames, the 2nd-chance algorithm incurs 14 page faults. Assume that we start searching from the frame next to the frame with the last brought in page. If the page last brought in is in the last frame, then we start searching from the first frame.
 - (c) If the process is allocated 6 frames, the LFU counting algorithm incurs 10 page faults. Assume that we start searching from the frame next to the frame with the last brought in page. If the page last brought in is in the last frame, then we start searching from the first frame.
9. (5 points) Which of the following statements regarding UNIX and Linux is/are correct ?
- (a) *Root* is the first process created in system booting.
 - (b) Daemon is a program running as a background process.
 - (c) A sysV-style init process can switch between different runlevels.
10. (5 points) Continuing from the last question, which of the following statements is/are correct?
- (a) The *init* process automatically adopts all orphaned processes.
 - (b) In System V, runlevel 1 is for single user mode.
 - (c) The */etc/inittab* file defines what each configured runlevel does in a given system.

11. (5 points) One critical factor in powering a server farm is cooling. If heat is not removed from the computer efficiently, the fans will blow hot air back onto the computer, not cold air. We will look at how different design decisions affect the necessary cooling, and thus the price, of a system. Please use the components and devices in the following for your power calculations.

Component type	Product	Performance	Power
Processor	P1	4GHz, 16-core	102-130 W peak
	P2	6GHz	49-60 W
DRAM	M1	184-pin, 4GB	2.3 W
	M2	240-pin, 2GB	1.8 W
Hard drive	D1	8000rpm	8.2 W read/seek, 3 W idle
	D2	9200rpm	9.5 W read/seek, 3.9 W idle

Please answer which of the following statements is/are correct.

- (a) A cooling door for a rack costs \$4000 and dissipates 16 KW (into the room; additional cost is required to get it out of the room). The cooling door can cool 223 servers with a P2 processor, an M2 DRAM, and a single D2 hard drive.
- (b) You are considering providing fault tolerance with a P2 processor, an M2 DRAM, and seven D1 hard drives running RAID 2. Now 134 systems can be cooled down on a single rack with a single cooler.

12. (5 points) Continuing from the last question, please answer which of the following statements is/are correct. Suppose that our server farm can dissipate a maximum of 300 W per square foot and a server rack requires 11 square feet (including front and back clearance).

- (a) A cooling door for a rack costs \$6000 and dissipates 16 KW (into the room; additional cost is required to get it out of the room). The servers are all with a P1 processor, an M1 DRAM, and two D1 hard drives with RAID 1. In the farm, a single rack can host up to 28 servers.
- (b) From part (a), two cooling doors are required for a single rack.
- (c) Suppose a cooling door for a rack costs \$3500 and dissipates 28KW (into the room; additional cost is required to get it out of the room). The servers are all with a P2 processor, an M1 DRAM, and five D2 hard drives with RAID 6. In the farm, a single rack can host up to 30 servers.

13. (5 points) About cache optimization, which of the following statements is/are correct ?

- (a) Larger block sizes reduce compulsory misses, increase the miss penalty, and can increase capacity or conflict misses.
- (b) Larger cache capacities reduces miss rate but may incur longer hit time and higher cost and power.
- (c) Greater associativity reduces conflict misses. Greater associativity can come at the cost of increased hit time.

14. (5 points) About cache optimization, which of the following statements is/are correct ?

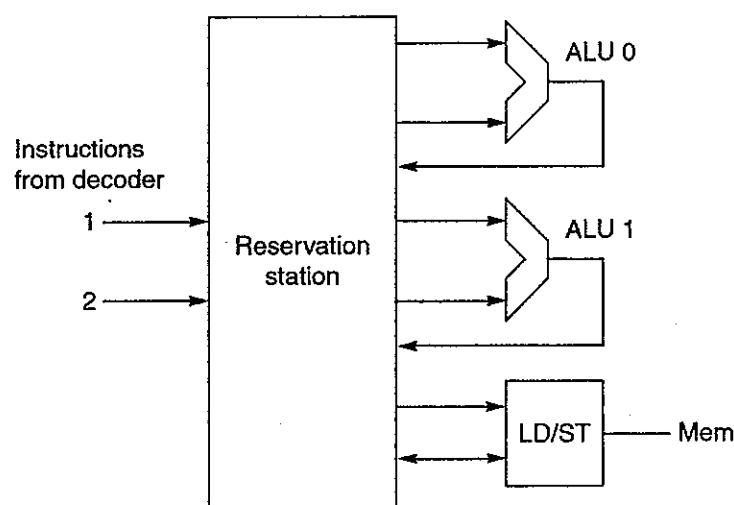
- (a) In a multilevel cache system, the first-level cache can be small enough to match a fast clock cycle time, yet the second-level (or third-level) cache can be large enough to capture many accesses that would go to main memory.
- (b) Most processors give cache write misses priority over read misses to reduce miss penalty.
- (c) A common optimization is to use the page offset to index the cache in order to remove the translation lookaside buffer (TLB) access from the critical path.

見背面

15. (5 points) Assume that we have the following machine instruction program.

[1]	Loop:	LW	R1, 0(R3)
[2]		ADDI	R1, R1, #1
[3]		SW	R1, 0(R3)
[4]		LW	R1, 0(R0)
[5]		SUB	R4, R1, R2
[6]		BNZ	R4, Loop

Also assume that our processor has the following microarchitecture that the arithmetic-logical units (ALUs) can do all arithmetic ops in the above program and that the Reservation Station (RS) can simultaneously dispatch at most one operation to each functional unit per cycle (one op to each ALU plus one memory op to the LD/ST).



Suppose all of the instructions from the program in the above are present in the RS, with no renaming having been done. Please answer which of the following statements is/are correct? (Hint: An instruction with latency +2 requires two <stall> cycles to be inserted into the code sequence. Think of it this way: A one-cycle instruction has latency 1 + 0, meaning zero extra wait states. So, latency 1 + 1 implies one stall cycle; latency 1 + N has N extra stall cycles.)

- (a) Performance can be improved by renaming register R1 in instruction [1], [2], and [3].
- (b) Performance can be improved by renaming register R4 in instruction [5] and [6].
- (c) Suppose the register-renamed version of the above code is resident in the RS in clock cycle N, with latencies LW = +4, SW = +1, ADDI = +0, SUB = +0, and BNZ = +2. The number of clock cycles taken by the code sequence is 14.

16. (5 points) Continuing from the last question, please answer which of the following statements is/are correct?

- (a) Suppose that the RS is empty and the front end (decoder/register-renamer) will continue to supply two new instructions per clock cycle. In cycle 0, the first two register-renamed instructions of this sequence appear in the RS. Assume it takes one clock cycle to dispatch any op in addition to the functional unit latencies. Then 9 clock cycles are required in the first iteration of this code sequence.
- (b) Adding another ALU can save two clock cycles.
- (c) Adding another LD/ST can save four clock cycles.

17. (5 points) Consider the following code.

```
for (k=0;k<400;k++) {  
    r[2*k] = 4 / (4*k+1);  
    r[2*k+1] = -4 / (4*k+3);  
}
```

Assume that the processor runs at 700 MHz and has a maximum vector length of 128. The load/store unit has a start-up overhead of 16 cycles; the multiply/division unit, 8 cycles; and the add/subtract unit, 6 cycles. Please answer which of the following statements is/are correct.

- (a) The arithmetic intensity of this kernel is 1.
- (b) Assuming chaining and a single memory pipeline, there are 4 cycles required.

18. (5 points) Suppose that we have three classes A, B, and C of instructions with CPI 2, 3, 6 respectively. Which of the following statements is correct?

- (a) A code segment with 10 instructions in A, 8 instructions in B, and 5 instructions in C use 74 cycles.
- (b) The CPI of the code in part (a) is 3.
- (c) A code segment with 8 instructions in A, 6 instructions in B, and 8 instructions in C use 82 cycles.

19. (5 points) In the problem of cache coherence protocols, which of the following statements is/are correct ?

- (a) Directory-based protocols in general have better scalability than snooping protocols.
- (b) Snooping protocols rely on a shared bus that specifically records which caches have copies of a piece of data so that when a processor modifies its copy in its local cache, updating signals can be sent to the other caches.
- (c) Directory-based protocols is usually more complex and inefficient for small-scale systems.

20. (5 points) Assume that we have an interconnection network topology on the clock cycles per instruction (CPI) of programs running on a distributed-memory multiprocessor. The processor clock rate is 6.3 GHz and the base CPI of an application with all references hitting in the cache is 0.75. Assume that 0.2% of the instructions involve a remote communication reference. The cost of a remote communication reference is $(100 + 10h)$ ns, where h is the number of communication network hops that a remote reference has to make to the remote processor memory and back. Assume that all communication links are bidirectional. Please answer which of the following statements is/are correct.

- (a) The worst-case remote communication cost when the system is a 4×4 processor grid is 150 ns.
- (b) When the system is an $8 \times 8 \times 8$ hypercube, the average CPI of the application is smaller than 1.3.
- (c) When the system is a 32-node ring, the average CPI of the application is larger than 1.1.