Ex/CSE/PC/B/T/226/2022

B.E. Computer Science and Engineering SECOND YEAR SECOND SEMESTER-2022

COMPUTER ARCHITECTURE

Time: 3 hours **Full Marks: 100**

GROUP-A

Answer 40 questions $40 \times 2 = 80$

Choose the unique correct answer

For Q1-Q2

Consider a 5-stage pipeline (IF, ID, EX, MEM, WB). The i-th instruction entering the pipeline in cycle-1 is a branch which is detected in the ID-stage. In order to prevent a control hazard, a bubble is inserted in the pipeline. The (i+1)-th instruction is the instruction that lexically follows the branch.

1.In cycle-2, instruction i+1 is in
(a) IF
(b) ID
(c) EX
(d) none of the above

2. In cycle-3, instruction i+1 is in
(a) IF
(b) ID
(c) EX
(d) none of the above

For the rest of this Question Paper:

Unless otherwise mentioned, we will assume a 5-stage instruction pipeline: IF, ID, EX, MEM, WB.

For Q3-Q5

Consider a scheme for reducing branch penalty in a 5-stage pipeline. Every branch is initially ASSUMED to be NOT TAKEN, i.e. the instruction i+1 lexically following the i-th instruction, a branch, is fetched. Whether the branch (instruction-i) is actually taken is known only in its decode stage ID. If it is found that the branch will be taken, then the instruction at the branch target is fetched and NOT instruction i+2.

Now consider the i-th instruction, a branch, doing its IF (instruction fetch) in cycle-1, and the branch IS TAKEN.

3. In cycle-2, the IF stage is occupied by
(a) instruction-i
(b) instruction i+1
(c) instruction at the branch target
(d) none of the above

4. In cycle-3, the ID (instruction decode) stage is occupied by
(a) instruction i
(b) instruction i+1
(c) instruction at the branch target
(d) none of the above

5. The branch target instruction is fetched in
(a) cycle-2
(b) cycle-3
(c) cycle-4
(d) none of the above

For Q6-Q7:

Consider a 5-stage pipeline. Instruction-i is a Load instruction and does IF (fetch) in cycle-1. Both instructions and data are accessed through a single memory port, i.e. in a particular cycle, only one instruction can access memory.

6. In cycle-4, the IF stage is occupied by
(a) instruction i+3
(b) instruction i+4
(c) instruction i+5
(d) none of the above

7. Instruction i+3 occupies EX stage in
(a) cycle-6
(b) cycle-7
(c) cycle-8
(d) none of the above

For Q8-Q10:

Consider the following latencies of FP (floating-point) operations:

Sl No	Instruction producing result	Instruction using result	Latency in clock-cycles
1	FP ALU op	Another FP ALU op	3
2	FP ALU op	Store double	2
3	Load double	FP ALU op	1
4	Load double	Store double	0

Now consider the following code:

```
-----------------------------------------------
Loop:   L.D       F0, 0(R1)
        ADD.D     F4, F0, F2
        S.D       F4, 0(R1)
        DADDUI    R1, R1, #-8
        BNE       R1, R2, Loop
-----------------------------------------------
```

Suppose L.D is issued in cycle-1. Instruction scheduling is NOT employed.

8. ADD.D is issued in cycle-x and the applicable Rule-number y (Sl No of the table given above) is
(a) x=2, y = 4
(b) x =3, y = 3
(c) x =4, y = 2
(d) x = 5, y = 1

9. Cycle x in which S.D is issued and the applicable Rule number y is
(a) x = 3, y = 3
(b) x = 4, y =3
(c) x = 5, y = 2
(d) x = 6, y = 2

10. x and y for BNE is
(a) x = 6, y = 3
(b) x = 7, y = 2
(c) x = 8, y : NOT APPLICABLE
(d) x = 9, y: NOT APPLICABLE

For Q11-14: Consider the following program for the MIPS floating-point unit:

```
-----------------------------------
1) L.D        F6,32(R2)
2) L.D        F2, 44(R3)
3) MUL.D      F0,F2,F4
4) SUB.D      F8, F2, F6
5) DIV.D      F10,F0,F6
6) ADD.D      F6, F8, F2
-----------------------------------
```

Consider the moment when
(i) all instructions have issued;
(ii) all instructions except DIV.D have finished execution;
(iii) all instrutions except MUL.D and DIV.D have written their result; and
(iv) MUL.D is ready to write its result.

Using the conventional notation, answer the following questions.

11. Vj for MUL.D is
(a) Mem[32 + Regs[R2]]
(b) Mem[44 + Regs[R3]]
(c) Regs[F2]
(d) Regs[F4]

12. Vk for MUL.D is
(a) Regs[F2]
(b) Regs[F0]
(c) Mem[44 + Regs[R3]]

13. Vj for DIV.D is
(a) Regs[F0]
(b) Regs[F6]
(c) Mem[32 + REgfs[R2]]
(d) none of the above

14. Qj for DIV.D is
(a) L.D F6, 32(R2)
(b) SUB.D
(c) MUL.D
(d) ADD.D

For Q15-Q20: Answer the following questions with reference to the MSI protocol for cache coherence.

15. If the state of a cache block in processor Pi is M (Modified, i.e. dirty),
(a) the main memory contains a valid copy
(b) other caches do NOT contain a valid copy
(c) the cache of Pi doesn't contain valid data
(d) none of the above

16. If the state of a cache block in processor Pi is S (shared),
(a) main memory contains stale data
(b) every other cache contains a valid copy of the block
(c) main memory contains a valid copy of the block
(d) none of the above

17. A Bus Write Back (BusWB) transaction
(a) is generated by a cache controller
(b) is generated by a processor
(c) does not involve main memory
(d) none of the above

18. A cache changes its state from S (shared) to M (modified) on a
(a) processor write (PrWr)
(b) processor read (PrRd)
(c) Bus Read (BusRd)
(d) Bus Read Exclusive (BusRdX)

19. On a cache miss following a processor read, a BusRd transaction is generated. After this,
(a) the main memory supplies the data
(b) a different cache supplies the data
(c) the requesting cache, after obtaining the data, changes the cache state to S (shared), even if no other cache is sharing the data
(d) none of the above

20. On a cache miss following a processor write (PrWr), a bus read exclusive(BusRdX) transaction is generated. After this, all **other** cache states become
(a) modified (M)
(b) invalid (I)
(c) shared (S)
(d) none of the above

For Q21-23: Answer the following questions with reference to the MESI protocol for cache coherence.

21. If the state of a cache-block of processor Pi is E (exclusive),
(a) main memory does not have a valid copy of the block
(b) no other cache has a (valid) copy of the block
(c) no other cache has a (valid) copy of the block
(d) none of the above

22. On a cache read miss, the state changes to
(a) S (shared)
(b) E (exclusive)
(c) M (modified)
(d) S if the shared signal (***S***) is asserted, E if shared signal(**S**) is NOT asserted

23. If a cache-block is in E (exclusive) state, and it observes a BusRd (Bus Read) transaction, it changes its state to
(a) E (exclusive)
(b) I (invalid)
(c) M (modified)
(d) S (shared)

For Q24-Q26: Consider a directory-based cache coherence scheme in a scalable multiprocessor. The directory contains a bit-vector of presence bits, one bit for each node of the multiprocessor.

24. On a read miss for cache block at node i, and if the home node j is a remote node,
(a) node j immediately supplies data to node i
(b) node i sends a request to node j to look up the directory.
(c) if the block is dirty, node j supplies to i the identity k of the owner node
(d) both (b) and (c)

25. When a dirty node provides data to a requestor node on a read miss, the state of the block
(a) changes to "shared"
(b) remains "dirty"

26. On a write-miss at node i for a cache block that is not dirty,
(a) the state of the cache block remains shared
(b) the home node sends the presence bit-vector to node-i
(c) node-i sends invalidation requests to all caches sharing that block
(d) both (b) and (c)

---

27. An interconnection network is a "programmable system". It is a **system** because
(a) it is a supercomputer
(b) it connects many different buildings or geographically separated workstations
(c) it is composed of many components: buffers, channels, switches and controls that work together to deliver data
(d) none of the above

28. In an n × m crossbar,
(a) there are n number of m:1 multiplexers
(b) there are m number of n:1 multiplexers
(c) each multiplexer is connected to one input
(d) none of the above

29. In a 2-ary 3-fly butterfly network, there are
(a) 3^2 source terminal nodes
(b) 3 crossbar nodes in each stage
(c) each switch node is a 2 × 2 crossbar
(d) none of the above

30. The basic unit of routing and sequencing is a
(a) phit
(b) flit
(c) packet
(d) message

31. The basic unit of bandwidth and storage allocation is a
(a) phit
(b) flit
(c) packet
(d) message

32. The typical bit-length of a flit is
(a) 512K
(b) 128
(c) 8
(d) 64

For Q33-Q35: Consider a dropping flow control mechanism with explicit negative acknowledgement. A 5-flit packet consists of a head flit H, three successive body flits B1, B2, and B3, and a tail-flit T. The packet is routed through a 4-hop route consisting of channels c0, c1, c2, and c3. Each channel has two links, a forward link F, and a reverse link R. The head flit enters channels c0 (F-link) in cycle-0. In cycle-3, the packet is unable to proceed to channel c3. A negative acknowledgement flit 'N' is triggered on the R-link of c2. The N-flit travels backwards along the reverse links (R) from c2 to c0 in successive cycles and reaches c0 in cycle-5. In cycle-6, a retransmission of the packet is initiated.

33. In cycle-4, c1 (F-link) is occupied by
(a) H
(b) B1
(c) B2
(d) B3

34. In cycle-4, c0 (F-link) is occupied by
(a) T
(b) B1
(c) B2
(d) B3

35. In the retransmitted packet, B3 occupies c3 (F-link) in
(a) cycle 11
(b) cycle 12
(c) cycle 13
(d) cycle 14

For Q36-Q41: Consider a node of an interconnection network. The node is a 2×2 crossbar having two input channels I1, I2 and two output channels O1, O2. A particular output can be connected to exactly one input. If a flit f2 arriving on input channel Ij wants a connection to output channel Ow while it is connected to input Ik for transporting flit f1, then f2 MUST WAIT until f1 has been completely transmitted through the node. Transport is achieved by WORMHOLE ROUTING. Each input channel has an associated flit buffer having a capacity

of 2 flits. The states of an input channel are idle (I) when the buffer is empty, active (A) when the buffer is partially or fully occupied with flits, and waiting (W) when the flits in the buffer are waiting for transmission but cannot proceed since the desired output channel cannot be allocated.

Now consider a 4-flit packet having a head flit H, two consecutive body flits B1, B2, and a tail flit T. The packet arrives at the node in cycle-1 on input channel I1 whose buffer is empty and desires a connection to output channel O1 which is busy since it is transmitting a packet arriving at channel I2. In cycle-5, O1 becomes free and can be connected to I1.

36. In cycle-2, the number of filled buffers of I1 is
(a) 0
(b) 1
(c) 2
(d) 3

37. If the state of I1 in cycle-1 is I, the state in cycle-2 is
(a) I
(b) W
(c) A
(d) none of the above

38. In cycle-3, the buffer of I1 is occupied by
(a) H
(b) H and B1
(c) B1 and B2
(d) B2 and T

39. In cycle-4, the buffer of I1 is occupied by
(a) H
(b) H and B1
(c) B1 and B2
(d) B2 and T

40. In cycle-5, when O1 is connected to I1, the state of buffer I1 is
(a) I
(b) W
(c) A
(d0 none of the above

41. In cycle-8, the state of F1 becomes
(a) I
(b) W
(c) A
(d) none of the above

For Q42-Q44: Consider a bus arbitration scheme with a daisy-chained grant in which there are N bus masters 1,..., N, master-1 being connected to the arbiter.

42. The number of bus request lines is
(a) 1
(b) 2
(c) 3
(d) N

43.A master can access the bus if
(a) bus-busy line is active
(b) bus-busy line is passive

44. If master-i ($1 \leq i \leq N-1$) finds its input grant line active but doesn't need the bus, it
(a) deactivates its bus request
(b) activates the bus busy line
(c) activates its output grant line, i.e. the input grant line of master (i + 1)
(d) none of the above

---

45. In a distributed shared-memory architecture, the local memories are
(a) accessible only by the associated processors
(b) accessible by any processor

46. Distributed memory multicomputers
(a) use semaphores for synchronization
(b) are highly scalable
(c) do not need load-balancing
(d) none of the above

For Q47-Q48: Consider the following program for the CRAY X-MP:

Line	Instruction	
1	A1	53
2	VL	A1
3	V4	V2 + V3

47. The first result (corresponding to the first element of the 53-element vector) emerges after
(a) 3 cycles
(b) 6 cycles
(c) 9 cycles
(d) 61 cycles

48. The functional unit, i.e. the vector adder, can be reused after
(a) 52 cycles
(b) 56 cycles
(c) 57 cycles
(d) 61 cycles

---

49. The B registers (B0 .. B77) of the CRAY X-MP have a bit-length of
(a) 8
(b) 16
(c) 24
(d) 32

50. During instruction fetch of a 1-parcel instruction, the parcel is first transferred to
(a) CIP
(b) LIP
(c) NIP
(d) none of the above

GROUP-B

51. A pipelined processor has two execution units. Each execution unit is capable of executing any instruction. Each processor has a 6-stage execution pipeline. Instructions CANNOT be reordered.

The latency of an instruction type is defined as the delay between the time at which an instruction of that type issues AND the time at which a dependent instruction may issue. An instruction is said to have issued when it passes into the execution stage.

In this processor, load operations have a latency of 3 cycles, and all other operations have a latency of 2 cycles.

Now consider the following program:

```
-------------------------
LD      r4, (r5)        ; Load  r4 := (r5), i.e. r5 contains the memory address
LD      r7, (r8)
ADD     r9, r4, r7      ; r9 := r4 + r7
LD      r10, (r11)
MUL     r12, r13, r14
SUB     r2, r3, r1
ST      (r2), r15       ;Store
MUL     r21, r4, r7
ST      (r22), r23
ST      (r24), r21
-------------------------
```

How long would the program take to execute (issue-time + execution-time) ? **20**

---------- END ----------