

B.E. Computer Science and Engineering SECOND YEAR SECOND SEMESTER-2023

COMPUTER ARCHITECTURE

Time: 3 hours

Full Marks: 100

GROUP-A

Answer 40 questions

40 × 2 = 80

Choose the unique correct answer

1. Message-passing achieves

- (a) communication
- (b) synchronization
- (c) both (a) and (b)
- (d) none of the above

2. A NUMA machine is a

- (a) shared memory system
- (b) distributed memory system
- (c) distributed shared memory system
- (d) none of the above

3. In a single-bus multiprocessor with centralized arbitration and separate request and grant lines, the arbiter accepts a request

- (a) when the bus busy line is passive
- (b) when the bus busy line is active
- (c) as soon as it arrives
- (d) by deactivating the bus busy line

4. In the MSI protocol for cache coherence, if a cache controller observes a BusRdX transaction while its cache is in the M-state, it changes its state to

- (a) I
- (b) M
- (c) S
- (d) none of the above

5. In the MESI protocol for cache coherence, if a cache controller receives a PrWr request while the cache is in state E, it changes the state to

- (a) E
- (b) I
- (c) M
- (d) S

[Turn over

6. In directory-based cache coherence, a node incurring a cache miss
- (a) obtains the required block from its local memory immediately
 - (b) communicates with the directory entry of the missing block
 - (c) asks the Communication-Assist to fetch the data
 - (d) none of the above
7. Circuit-switching is a form of
- (a) topology
 - (b) routing
 - (c) flow control
 - (d) none of the above
8. A 4×5 crossbar is implemented with 5 $4:1$ multiplexers. Each input is connected to
- (a) 5 multiplexers
 - (b) 4 multiplexers
 - (c) 3 multiplexers
 - (d) none of the above
9. Consider a symmetric Clos network with $m=3$, $n=3$, $r=4$. Each middle-stage switch has
- (a) 1 input
 - (b) 2 inputs
 - (c) 3 inputs
 - (d) 4 inputs
10. the number of crossbar switch nodes in each stage of a k -ary n -fly butterfly network is
- (a) k^{n-2}
 - (b) k^{n-1}
 - (c) k^n
 - (d) k^{n+1}
11. The efficient allocation to a packet of a network's resources, e.g. channel bandwidth, buffer capacity, and control state, is determined by
- (a) flow control
 - (b) routing
 - (c) topology
 - (d) none of the above
12. The unit of information that is transferred across a channel in a single clock cycle is called a
- (a) flit
 - (b) message
 - (c) packet
 - (d) phit

13. Dropping flow control is a technique of

- (a) buffered flow control
- (b) bufferless flow control

14. In Store-and-Forward flow control, a node

- (a) waits for a complete packet to be received before forwarding it
- (b) doesn't wait for the complete packet

15. Cut-through flow control

- (a) waits for a complete packet to be received before forwarding it
- (b) forwards a packet as soon as the header is received

16. The shared signal (**S**) in the MESI protocol

- (a) enables a cache controller to determine on a BusRd whether any other cache currently holds the data
- (b) is activated by the main memory whenever any block is written through
- (c) is activated by DMA controllers whenever data transfers take place via DMA
- (d) none of the above

For Q17-Q18: Consider a directory-based cache coherence scheme with bit-vector organization. Node **I** encounters a read or write miss. Suppose the home node of the block is a remote node.

17. If the dirty bit is ON

- (a) the communication assist at the home node sends to **I** (requestor) the identity of the owner (i.e. dirty) node
- (b) the requestor sends a request to the owner node
- (c) the owner supplies the block to the requestor
- (d) all of the above

18. If the dirty bit is OFF, the communication assist at node **I**

- (a) obtains the block from the main memory
- (b) obtains the block from the exclusive node
- (c) obtains the block from the dirty node
- (d) none of the above

For Q19-Q24: Consider a 5-stage instruction pipeline consisting of the functional units IM, Reg, ALU, DM and the stages IF, ID, EX, MEM, WB. Successive instructions entering the pipeline are designated i_1, i_2, i_3 , etc. The clock cycles are designated CC1, CC2, etc. Forwarding is not employed unless explicitly stated.

19. If i_1, i_2, \dots, i_5 enter the pipeline in consecutive cycles CC_1, CC_2, \dots, CC_5 , two simultaneous memory accesses occur in
- (a) CC_3
 - (b) CC_4
 - (c) CC_5
 - (d) both (b) and (c)
20. Suppose i_4 enters the pipeline in CC_5 instead of CC_4 . This measure
- (a) delays the pipeline without reason
 - (b) avoids memory conflict if there is only one memory port for both instructions and data
 - (c) prevents register-memory conflicts
 - (d) none of the above

For Q21-QQ22: Consider the following program fragment:

```

DADD    R1, R2, R3
DSUB    R4, R1, R5
AND     R6, R1, R7
OR      R8, R1, R9
XOR     R10, R1, R11

```

The instructions are allowed to enter the pipeline in successive cycles $CC_1.. CC_5$.

21. $DSUB$ reads R_1 in CC_3 and this value is
- (a) INCORRECT
 - (b) CORRECT
22. If forwarding hardware is present, $DSUB$ can obtain the output of $DADD$ in
- (a) CC_1
 - (b) CC_2
 - (c) CC_3
 - (d) CC_4

For Q23: Consider the following program fragment:

```

LD      R1, 0(R2)
DSUB    R4, R1, R5
AND     R6, R1, R7
OR      R8, R1, R9

```

Assume that forwarding hardware IS PRESENT.

23. DSUB

- (a) obtains the correct value of R1 by means of forwarding
- (b) cannot obtain the correct value of R1, no matter what means are employed
- (c) can obtain the correct value of R1 if it enters ALU in CC5
- (d) none of the above

24. Suppose i1 is a branch instruction. For correct execution, i2 should be fetched in

- (a) CC2
- (b) CC3
- (c) CC4
- (d) CC5

-----END-Q19-Q24-----

For Q25-Q27: Consider the following table of latencies of Floating Point (FP) operations for MIPS. Latency is defined to be the number of **intervening** clock cycles needed to avoid a stall.

Instruction producing result	Instruction using result	Latency in clock cycles
FP ALU op	Another FP ALU op	3
FP ALU op	Store double	2
Load double	FP ALU op	1
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
    
```

Assume that L.D is issued in cycle-1.

25. ADD.D is issued in

- (a) cycle-2
- (b) cycle-3
- (c) cycle-4
- (d) cycle-5

26. S.D is issued in

- (a) cycle-3
- (b) cycle-4
- (c) cycle-5
- (d) cycle-6

[Turn over

27. Suppose the loop in the given code is unrolled by executing L.D, ADD.D, S.D four times in a new loop, with a proper assignment of registers. Then each iteration of the loop requires (including a stall after BNE)

- (a) 4 stalls
- (b) 9 stalls
- (c) 14 stalls
- (d) 19 stalls

-----END Q25-Q27-----

For Q28-Q30: Consider the MIPS floating-point unit which employs Dynamic Scheduling.

28. Register Renaming is provided by

- (a) Common Data Bus (CDB)
- (b) Reservation Stations
- (c) Load buffers
- (d) Store buffers

For Q29-Q30. Consider the following program fragment:

(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

All instructions have been issued but only instruction (1) has completed and written its result to the CDB.

29. V_k for MULT is

- (a) $44 + \text{Regs}[R3]$
- (b) $\text{Mem}[32 + \text{Regs}[R2]]$
- (c) $\text{Regs}[F4]$
- (d) none of the above

30. Q_j for instruction (6) is

- (a) instruction-1
- (b) instruction-3
- (c) instruction-2
- (d) instruction-4

-----END Q28-Q30-----

For Q31-Q36: Answer the following with reference to the CRAY X-MP/ Model 24 Architecture.

31. During the second cycle of the instruction issue phase of a 1-parcel instruction, the parcel is moved to the

- (a) LIP
- (b) NIP
- (c) CIP
- (d) none of the above

32. During the second cycle of the instruction issue phase of a 2-parcel instruction, the second parcel is transferred from the instruction buffer to the

- (a) LIP
- (b) NIP
- (c) CIP
- (d) none of the above

33. The RECIPROCAL APPROXIMATION vector functional unit has

- (a) 2 stages
- (b) 4 stages
- (c) 7 stages
- (d) 14 stages

34. The scalar instruction

S2 S1 + S7

reserves

- (a) S1
- (b) S2
- (c) S7
- (d) none of the above

35. Consider the following program fragment:

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

The destination register V4 becomes available after

- (a) 52 cycles
- (b) 56 cycles
- (c) 57 cycles
- (d) 61 cycles

36. Consider the following program fragment:

Line	Instruction
1	A1 10
2	VL A1
3	V4 V3 + FV2
4	V5 V4 * FV7

- Instruction-1 is issued in cycle-1.
 Then, instruction-4 begins **chained execution** in
- (a) cycle-12
 - (b) cycle-13
 - (c) cycle-14
 - (d) cycle-15

-----END Q31-Q36-----

37. Consider the following bar-graphs showing the effect of increasing the number of cores from 1 to 2 in the i7 and i5 processors.

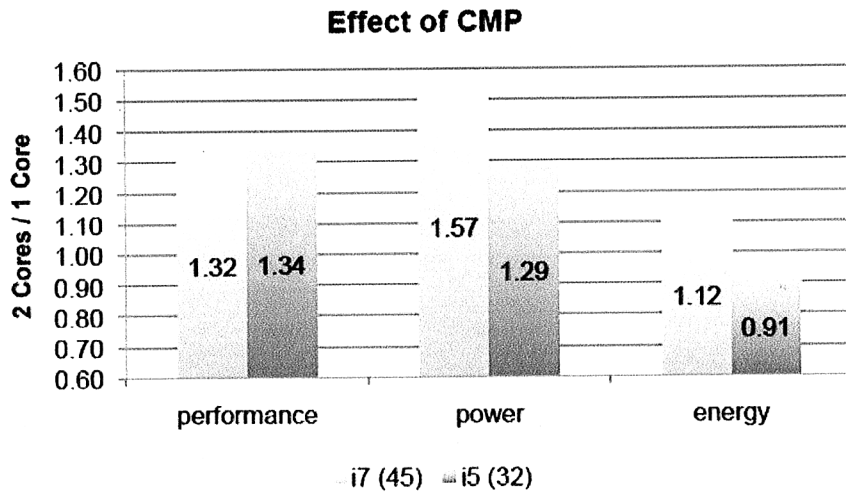


Fig Q37: Average impact of doubling cores.

While the performance improvement is roughly the same, 32 % for the i7 and 34 % for the i5, the i7 spends 12 % more energy while the i5 spends 9% less energy ! This is because

- (a) the power consumption of i7 goes up by 57% while that of i5 goes up by merely 29%, i.e half of that for i7
- (b) the clock rates of the i7 and the i5 are 2.7 GHz and 3.4 GHz, respectively
- (c) the TDP for the i7 and the i5 are 130W and 73W, respectively
- (d) i7 uses a 45nm technology while i5 uses 32nm technology.

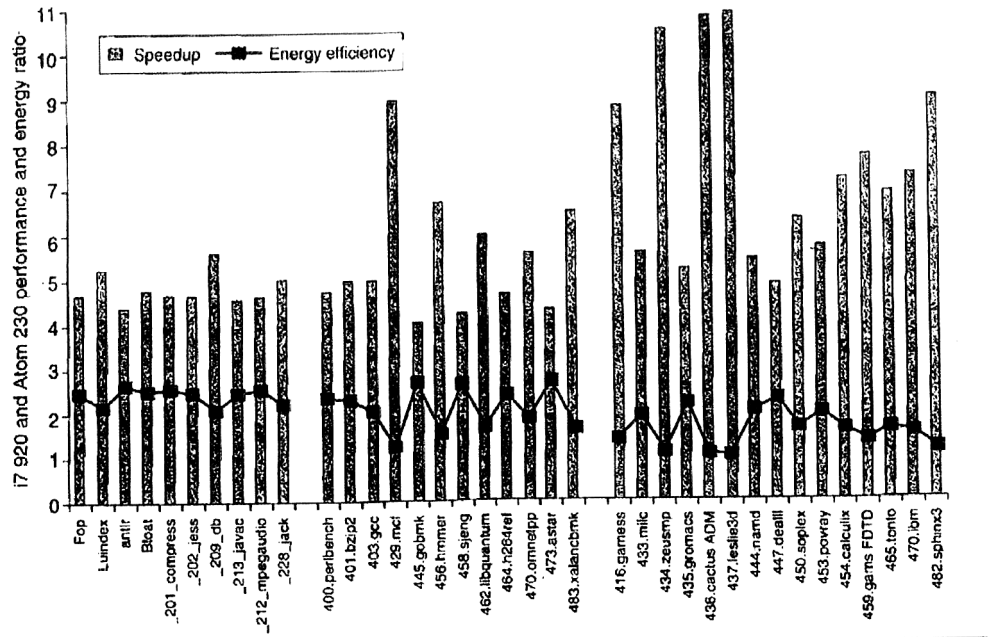


Figure Q38: Relative performance and energy-efficiency of the Intel i7 920 AND the Intel ATOM 230

38. Consider the 433.milc BENCHMARK in Figure Q38. The speedup (performance) ratio of i7 to ATOM is approximately 5.35. The average power consumption of the i7 and the Atom are approximately 43W and 4.2W, respectively. We can define energy-efficiency of a processor as the no. of operations per joule of energy expended. The performance is, of course, measured as the number of operations per second. Then, the energy-efficiency ratio of the Atom to the i7 is approximately

- (a) 1.0
- (b) 1.5
- (c) 1.9
- (d) none of the above

39. Consider the following bar-charts showing the effect of SMT(Simultaneous multithreading)




Fig Q39: Average impact of two-way SMT.

The energy-saving due to multithreading is maximum in the

- (a) Atom
- (b) i5
- (c) i7
- (d) Pentium 4

40. The CPUFreq framework in Linux is a technique for

- (a) Software guided power management
- (b) Dynamic Frequency and Voltage Scaling
- (c) Active Power Management
- (d) all of the above

41. Domain-specific architectures (DSAs) can achieve higher performance and greater energy efficiency because

- (a) they exploit a more efficient form of parallelism for the specific domain
- (b) they can make more efficient use of the memory hierarchy
- (c) they can use less precision (number of bits of data) when it is adequate
- (d) all of the above

42. The guiding principles in defining the RISC-V Instruction Set Architecture (ISA) was

- (a) to make it suitable for nearly any computing device
- (b) to make it open and free to implement
- (c) to separate the ISA into a small base ISA and optional extensions
- (d) all of the above

43. The standard extension A (Atomic operations) of the RISC-V architecture supports
- (a) compressed instructions
 - (b) double-precision floating-point
 - (c) multiprocessor synchronization
 - (d) integer multiply/divide
44. I-Buffer, Scoreboard, Issue, and SIMT stack in a **GPU SIMT-core** are part of the
- (a) instruction fetch loop
 - (b) instruction issue loop
 - (c) register access scheduling loop
 - (d) none of the above
45. A warp (or wavefront) in a **GPU** is
- (a) a crossbar
 - (b) a group of threads that serves as a unit of scheduling
 - (c) an SIMD execution unit
 - (d) register file bank

GROUP-B

46. An out-of-order superscalar processor has two execution units. Instructions may be issued out-of-order. Each execution unit is capable of executing any instruction.

In this context, the latency of an instruction is defined as the delay between the time at which the instruction 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.

Load operations have a latency of 3 cycles, and all other operations have a latency of 2 cycles.

We assume a “greedy scheduling assumption”, i.e. if more instructions can issue in a cycle than the processor has execution units, the processor issues the instructions that occur earliest in the program, even if a different order of issue takes less time.

The instruction window of the processor is large enough to cover the entire code fragment. We define “instruction window” to be the number of instructions the processor examines to select instructions to issue in each cycle.

Now consider the following code fragment:

```

-----
(1)   LD    r4, (r5)      ; Load r4 from memory location whose address is in r5
(2)   LD    r7, (r8)
(3)   ADD   r9, r4, r7    ; r9 := r4 + r7
(4)   LD    r10, (r11)
(5)   MUL   r12, r13, r14
(6)   SUB   r2, r3, r1    ; r2 := r3 - r1
(7)   ST    (r2), r15    ;Store r15 in memory location whose address is in r2
(8)   MUL   r21, r4, r7
(9)   ST    (r22), r23
(10)  ST    (r24), r21
-----

```

How many cycles are required for issuing the above sequence ?

20

-----END-----