Ref. No.:EX/CSE/PC/B/T/226/2025

B.E. COMPUTER SCIENCE AND ENGINEERING SECOND YEAR SECOND SEMESTER – 2025

Subject: COMPUTER ARCHITECTURE Time: Three Hours Full Marks:100

Instructions: Answer **a total of 41** questions, 40 from Group A through Group F, and 1 from GROUP-G, choosing **a minimum number from each group as specified below:**

Group	Minimum number of questions to be answered from group	Total number of questions to be answered
A	5	40
B	3	
C	14	
D	3	
E	4	
F	4	
G	1 (no choice)	01

GROUPs A-F 40×2=80

Choose the unique correct answer

GROUP-A

Answer at least 5 questions

For Q1-Q3: The MIPS Floating-Point Unit has latencies given in the table below. 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

The following MIPS code adds a scalar in register F2 to every element of a vector V whose left boundary (array slot preceding the starting element) is stored in R2. Initially, R1 points to the last element.

CODE-C1

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

The compiler unrolls the loop to generate a **code fragment CODE-C2** in which 4 iterations of the loop in C1 constitute 1 iteration of a loop in C2. The four successive elements of the vector V loaded by C1 in four successive iterations correspond to F0, F6, F10, and F14, respectively. The sums are computed in F4, F8, F12, ad F16, respectively. Each element of the vector occupies 8 bytes. The instructions in C2 are referred to as C2.j, where j is the line-number. Thus, C2.1 implies the first instruction of C2. Now answer the following questions.

1. C2.9 is
(a) L.D F14, -24(R1)
(b) ADD.D F8, F6, F2
(c) S.D F12, -16(R1)
(d) none of the above

2. C2.4 is
(a) S.D F16, -24(R1)
(b) L.D F6, -8(R1)
(c) S.D F8, -8(R1)
(d) none of the above

3. C2.13 is
(a) BEQ R1, R2, Loop
(b) BNE R1, R2, Loop
(c) DADDUI R1, R1, #32
(d) DADDUI R1, R1, #-32

For Q4-Q6: Consider the following MIPS Floating Point Unit Code employing Dynamic Scheduling:

```
L.D     F6, 32(R2)
L.D     F2, 44(R3)
MUL.D   F0, F2, F4
SUB.D   F8, F2, F6
DIV.D   F10, F0, F6
ADD.D   F6, F8, F2
```

Consider the situation when MUL.D is ready to write its result. All instructions except DIV.D have completed execution (need not imply register write-back).
[**N.B.** MUL.D is loaded in Reservation Station Mult1. DIV.D is loaded in Reservation Station Mult2.]

4. For Mult1, Vj is
(a) Mem[32+Regs[R2]]
(b) Mem[44+Regs[R3]]
(c) Regs[F2]
(d) Regs[F4]

5. For Mult2, Qj is
(a) Regs[F0]
(b) Regs[F6]
(c) Add2
(d) Mult1

6. For Mult2, Vk is
(a) Add1
(b) Mult1
(c) Mem[32+Regs[R2]]
(d) Mem[44+Regs[R3]]

For Q7-Q8: Consider the following code fragment

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

The code is executed in a 5-stage pipeline (IF, ID, EX, MEM, WB). LD enters IF stage in cycle-1.

7. The content of the memory location accessed by LD reaches the pipeline
(a) at the beginning of cycle-4
(b) at the end of cycle-4
(c) at the end of cycle-5
(d) none of the above

8.The data hazards present in the above code can be resolved by
(a) bypassing
(b) stalling
(c) loop unrolling
(d) none of the above

GROUP-B

Answer at least 3 questions.

Answer the following with respect to the CRAY X-MP:

9. For a two-parcel instruction, when the first parcel moves from the NIP to the CIP, the second parcel is transferred to the
(a) Program Address Register
(b) NIP
(c) LIP
(d) Vector Length Register

10. For a scalar operation,
(a) input operands are reserved
(b) output operands are reserved
(c) both input and output operands are reserved
(d) neither input, nor output operands are reserved

For Q11-Q12: Consider the following code fragment:

Line	Instruction	
1	A1	10
2	VL	A1
3	V4	V3+FV2
4	V6	V3*FV7

Line-1 is issued in cycle-1.

11. Line-4 starts execution in
(a) cycle-5
(b) cycle-25
(c) cycle-14
(d) cycle-17

12. The first element of V4 is available at the end of
(a) cycle-12
(b) cycle-21
(c) cycle-26
(d) cycle-35

GROUP-C

Answer at least 14 questions

13. Consider a 2-ary 4-fly (Butterfly) Network. Each crossbar switch node of this network is labelled i.j, where $0 \leq i \leq 3$ (i is the stage-number)and $0 \leq j \leq 7$ (j is the switch-number in stage-i). Terminal number are in the range 0..15.

The sequence of switch-nodes traversed for transporting a message from terminal-12 to terminal-3 is
(a) 0.6, 1.1, 2.0, 3.1
(b) 0.6, 1.3, 2.1, 3.1
(c) 0.6, 1,6, 2.2, 3.1
(d) 0.6, 1.2, 2.0, 3.1

For Q14..Q18: Consider a 4-flit packet (H, B1, B2, T), arriving at input I1 of a switch-node of a network. The switch has two inputs I1, I2, two outputs O1, O2, and a buffer capacity of 2 flits. Initially, O1 is connected to I2. However, the packet has to travel to O1. Thus the packet has to be buffered until O1 becomes free.In cycle-1, flit-H arrives on I1. O1 becomes free in cycle-5. Wormhole Routing is employed.

14. In cycle-5, the buffer contains
(a) H only
(b) H and B1
(c) B2 and T
(d) B1 and B2

15. In cycle-6, the flit transferred to O1 is
(a) H
(b) B1
(c) B2
(d) T

16. In cycle-7, the flit transferred to O1 is
(a) H
(b) B1
(c) B2
(d) T

17. In cycle-4, the buffer contains the buffer contains
(a) H and B1
(b) B1 and B2
(c) B2 and T
(d) only H

18. In cycle-5, the buffer contains
(a) H only
(b) T only
(c) H and B1
(d) B1 and B2

19. The maximum length of a flit is
(a) 16 bits
(b) 64 bits
(c) 512 bits
(d) 1K bits

For Q20-Q23: Answer the following with reference to the MSI Protocol.

20. When a cache block in the shared (S) state receives a PrWr request, its new state becomes
(a) I
(b) M
(c) S
(d) none of the above

21. When a cache block in the shared(S) state observes a BusRdX event, its new state becomes
(a) I
(b) M
(c) S
(d) none of the above

22. When a cache block in the modified (M) state receives a PrRD request, its new state becomes
(a) I
(b) M
(c) S
(d) none of the abve

23. If a cache block in the modified (M) state observes a BusRd transaction, it
(a) does nothing
(b) changes to Invalid state
(c) it flushes the data on the bus and changes its state to 'shared'.
(d) none of the above

For Q24-Q26: Consider a Bit-Vector scheme for Directory-based Cache-coherence.

24. There is a write-miss at processor-i and the dirty bit os ON. The dirty node is j. Then, after processing the transaction,
(a) only presence[i] remains ON
(b) presence[i] and presence[j] are ON
(c) only presence[j] is ON
(d) none of the above

25. On a read-miss at node-i, when the dirty-bit at the home node j is OFF, the communication assist at node-j
(a) sends from its main memory the data to i and presence[i] is turned ON
(b) an exclusive node supplies the data to i
(c) a dirty node supplies data to i
(d) none of the above

26. When a node incurs a cache-miss, it first
(a) sends a broadcast on the network to find where the block resides
(b) communicates with the directory entry for the block
(c) snoops the bus
(d) none of the above

27. The idea behind Simultaneous Multithreading (SMT) is to
(a) utilize multiple cores
(b) utilize functional units with independent operations from the same or different threads
(c) simulate RISC on a CISC
(d) none of the above

For Q28-Q30: Answer the following with reference to the Intel Nehalem:
28. The unit responsible for retrieving blocks of macro-instructions from memory and translating them into micro-ops and buffering them for the downstream stages is the
(a) Execution Unit
(b) Scheduler
(c) Retirement Unit
(d) Front-End Pipeline (FEP)

29. The fact that the DRAM controller, the L3 cache, and QPI ports are all housed within the same silicon die as the four cores
(a) causes the Nehalem to run slower than chips with North-Bridge
(b) is the principal cause of the high performance of the Nehalem
(c) is not relevant to speed
(d) none of the above

30.The unit which accepts 16 bytes from the L1 instruction cache or pre-fetch buffers and prepares the Intel64 instructions found there for instruction decoding downstream is
(a) Instruction Queue
(b) Instruction Decoder
(c) Instruction Length Decoder
(d) Micro-fusion

GROUP-D

Answer at least 3 questions

For Q31-Q34:

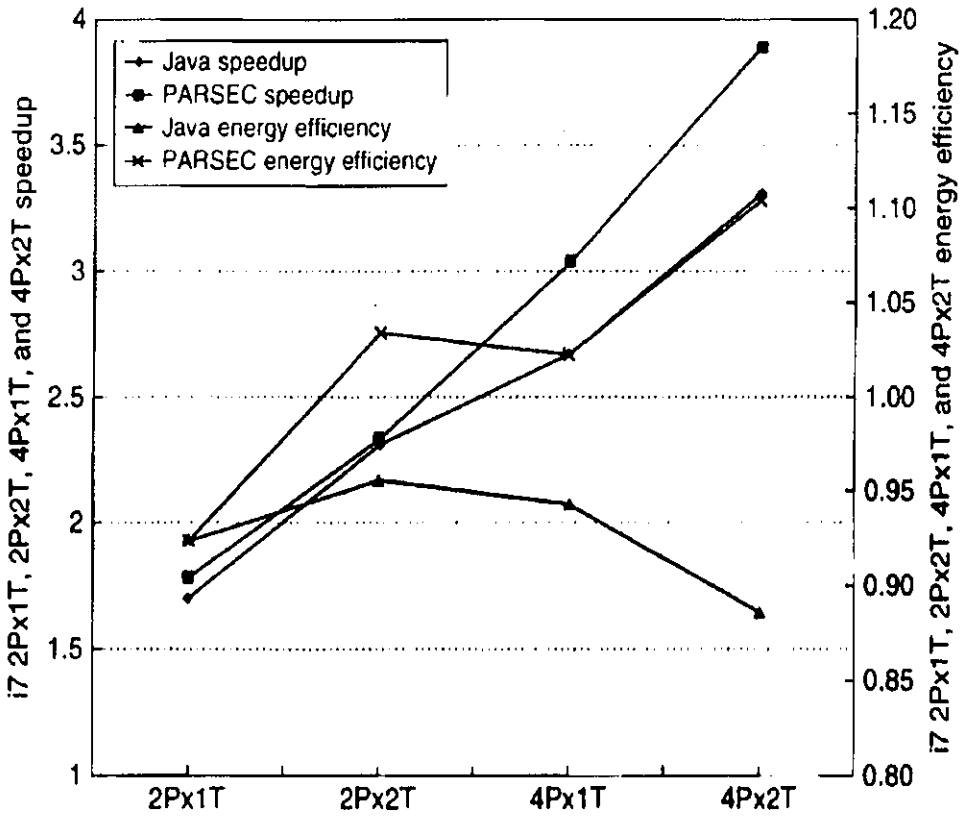

FigQ31

In Fig Q31, speedup and energy efficiency of the Intel Core i7 (single chip, 4 cores) are plotted for two benchmarks: Java and PARSEC. We consider 4 cases:

(i) 2 cores, no multithreading (2P×1T)
(ii) 2 cores with 2 threads per core (2P × 2T)
(iii) 4 cores, no multithreading (4P × 1T)
(iv) 4 cores with 2 threads per core (4P× 2T)

We define

$$\text{Speedup Efficiency} = \frac{Actual\ speedup}{Processor\ Count}$$

Now answer the following:

31. For the 4-core, no SMT, the PARSEC speedup is approximately
(a) 2.10
(b) 2.60
(c) 3.04
(d) 1.75

32. For the 4-core with SMT, the Java benchmark shows a speedup efficiency of
(a) 70%
(b) 83%
(c) 91%
(d) 97%

33. For the 4-core, with SMT, PARSEC speedup efficiency is
(a) 67%
(b) 75%
(c) 83%
(d) 97%

34. For the 4-core, no SMT, Java speedup is approximately
(a) 2.667
(b) 3.257
(c) 3.916
(d) 1.823

GROUP-E

Answer at least 4 questions

35. Direct emissions coming from fuel combustion (e.g. diesel, natural gas, and petrol), refrigerants in offices and data centres, transportation, and the use of chemicals and gases in semiconductor manufacturing are categorized by the Greenhouse Gas Protocol as
(a) Scope 1
(b) Scope 2
(c) Scope 3
(d) none of the above

36. Carbon emissions from hardware manufacturing can be reduced by
(a) using coal for powering fabs
(b) using renewable energy for semiconductor fabrication units
(c) no means, since carbion nemissions cannot be controlled
(d) none of the above

37. In Quantum Logic Circuits, operations are
(a) defined by Linear Algebra over Hilbert Space and can be represented by unitary matrices with complex elements
(b) defined by Boolean Algebra

38. The loss of information from a quantum system into the environment is called
(a) Suoperposition
(b) Entanglement
(c) Decoherence
(d) none of the above

39. The number of neuromorphic cores in the Intel Loihi chip is
(a) 32
(b) 64
(c) 128
(d) 256

GROUP-F

Answer at least 4 questions.

40. Digital Signal Processors (DSP), ASIC's (Application Specific Integrated Circuits), FPGA's (Field Programmable gate arrays) are all special purpose processors that
(a) use custom instructions
(b) optimize power by offloading processing from CPU
(c) employ energy aware scheduler
(d) none of the above

41. The technique to reduce Power Consumption by shutting off the current to blocks of the circuit that are not in use is called
(a) Dynamic Voltage Frequency Scaling
(b) Clock Gating
(c) Power Gating
(d) none of the above

42. A GPU is composed of many cores. For NVIDIA GPUs, the cores are called
(a) compute units
(b) streaming multiprocessors

43. GPU's can obtain improved performance over superscalar out-of-order CPU's on highly parallel workloads by dedicating a larger fraction of their die area to
(a) arithmetic logic units
(b) control logic

44. In a CMOS device, the supply voltage V, the switching threshold voltage, temperature, and transistor size determine
(a) dynamic power
(b) leakage power
(c) both (a) and (b)
(d) neither (a), nor (b)

45. A System level power and thermal management scheme keeps the system in its highest operating state (frequency/voltage) in order to complete the workload as fast as possible and then go to sleep at its lowest operating state (frequency/voltage). Such a scheme is called
(a) Crawl-to-halt
(b) Thermal Management
(c) Idle Power Management
(d) Race-to-idle

GROUP-G

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----------