3 Instruction pipelining

1. The performance of a pipelined processor suffers if:
   (A) the pipeline stages have different delays
   (B) consecutive instructions are dependent on each other
   (C) the pipeline stages share hardware resources
   (D) All of the above

2. Consider following situations for a pipelined CPU having a single processing unit:
   I. The $j + 1_{th}$ instruction uses the result of the j_{th} instruction as an operand.
   II. The execution of a conditional jump instruction
   III. The j_{th} and $j + 1_{th}$ instructions require the ALU at the same time.
   Which of the above can cause a hazard
   A. I and II only
   B. II and III only
   C. III only
   D. All the three

3. Which of the following statements about a superscalar processor is always true?
   (i) It executes more than one instruction per clock cycle.
   (ii) It executes two instructions in one clock cycle.
   (iii) It consists of more than one pipeline.
   (iv) It allows instructions without dependencies on one another to proceed in parallel through the separate pipelines.
   (A) (i) and (ii)
   (B) (i) and (iii)
   (C) (ii) and (iii)
   (D) (ii) and (iv)
   (E) (i), (iii) and (iv)
   (F) (i), (ii), (iii) and (iv)

4. Consider an instruction pipeline with four stages with the stage delays 5 nsec, 6 nsec, 11 nsec, and 8 nsec respectively. The delay of an inter-stage register stage of the pipeline is 1 nsec. What is the approximate speedup of the pipeline in the steady state under ideal conditions as compared to the corresponding non-pipelined implementation?
   (A) 4.0
   (B) 2.5
   (C) 1.1
   (D) 3.0

5. The stage delays in a five stage pipeline are 950,500,1150,600 and 790 picoseconds. The third stage (with delay 1050 picoseconds) is replaced with an equivalent split design involving two stages with respective delays of 800 and 550 picoseconds. The throughput increase of the pipeline is:
(A) 21%
(B) 27%
(C) 33%
(D) 47%

6. An instruction pipeline consists of 4 stages – Instruction Fetch (IF), Instruction Decode field (ID), Instruction Execute (IE) and Write Back (WB). The 5 instructions in a certain instruction sequence need these stages for the different number of clock cycles as shown by the table below:

Instruction No.	**IF**	**ID**	**IE**	**WB**
1	1	3	3	1
2	1	2	2	1
3	2	3	5	1
4	1	3	3	1
5	2	3	2	1

The number of clock cycles needed to perform the above five instructions is:
(A) 20
(B) 21
(C) 22
(D) 23

7. A branch history table is:
(A) It table which gives a better guess on whether or not an unconditional branch will be taken.
(B) It stores first K previously taken branch-targets with a taken/not taken bit for a smaller program, where K denotes the number of possible entries in the table.
(C) It stores last K previously taken branch-targets with a taken/not taken bit, where K denotes the number of possible entries in the table.
(D) Both (A) and (C).

8. A 5-stage pipeline has the stage delays as 155, 162, 143, 168 and 130 nanoseconds respectively. Registers that are used between the stages have a delay of 5,8 and 7 nanoseconds respectively. Assuming constant clocking

rate, the total time taken to process 1000 data items on this pipeline will be:
(A) 180.4 microseconds
(B) 160.5 microseconds
(C) 165.5 microseconds
(D) 176.7 microseconds

9. Comparing the time T_1 taken for a single instruction on a pipelined CPU with time T_2 taken on a non-pipelined but identical CPU, we can say that
(A) $T_1 \leq T_2$
(B) $T_1 \geq T_2$
(C) $T_1 < T_2$
(D) $T_1 = T_2$

10. A processor takes 20 cycles to complete an instruction. The corresponding pipelined processor uses 6 stages with the execution times of 3, 4, 5, 4, 5 and 3 cycles respectively with an inter-stage register delay of 1 cycle. What is the asymptotic speedup assuming that a very large number of instructions are to be executed?
(A) 3.33
(B) 2
(C) 3
(D) 4
(E) 5

11. Consider an instruction pipeline for the MIPS32 processor where data references constitute 42% of the instructions, and the ideal CPI ignoring memory structural hazards is 1.25. How much faster is the ideal machine without the memory structural hazard versus the machine with the hazard?
(A) 1.34
(B) 1.26
(C) 1.38
(D) 1.45

12. A machine has the following five pipeline stages, with their respective time requirements in nanoseconds (ns):
IF-stage — instruction fetch (9 ns),
ID-stage — instruction decode and register fetch (3 ns),
EX-stage — execute/address calculation (7 ns),
MM-stage — memory access (9 ns),
WB-stage — write back to a register (2 ns).
Assume that for each stage, the pipeline overhead is 1 ns. A program having 100 machine instructions runs on , where every 3rd instruction needs

a 2-cycle stall before the EX-stage. Calculate the CPU time in seconds for completing .
(A) 1200
(B) 1370
(C) 3000
(D) 1700
(E) 1470

13. Consider the following four organizations of pipelined processors.
P_1: Four-stage pipeline with stage latencies 1 ns, 2 ns, 2 ns, 1 ns with inter-stage-register-latency as 2 ns.
P_2: Four-stage pipeline with stage latencies 1 ns, 1.5 ns, 1.5 ns, 1.5 ns with inter-stage-register-latency as 1 ns.
P_3: Four-stage pipeline with stage latencies 0.5 ns, 1 ns, 1 ns, 0.6 ns with inter-stage-register-latency as 1 ns.
P_4: Four-stage pipeline with stage latencies 0.5 ns, 0.5 ns, 1 ns, 1 ns with inter-stage-register-latency as 1.5 ns.
Arrange these processors as a tuple (a,b,c,d) where ordering is opposite to the NON-INCREASING order of their peak clock frequency (in case of equal frequencies select a processor organization with smaller subscript)?
(A) P_1, P_2, P_4, P_3
(B) P_1, P_2, P_3, P_4
(C) P_3, P_4, P_2, P_1
(D) P_3, P_2, P_4, P_1

14. An instruction pipeline has six stages where each stage take 3 nanoseconds and all instruction use all six stages.
If Branch instructions are not overlapped and there are only 36% Branch instructions then which of the following represents the average instruction execution time (in ns), (ignoring conditional branch instructions)
?
(A) 3.6
(B) 4.4
(C) 5.6
(D) 6.4
(E) 8.4

15. In an MIPS32 pipeline,consider following measures that can be taken to reduce the impact of data hazards:
(i) Splitting the memory into separate Instruction and Data memories.
(ii) Implement data forwarding in the datapath.
(iii) Allow split register write and read during the two halves of the same clock cycle.
(iv) Replicate the register bank.

Which of the following correctly solves the problem:
(A) (i) and (ii) only
(B) (i) and (iii) only
(C) (ii) and (iv) only
(D) (ii) and (iii) only

16. Consider a pipeline processor with four stages S_1 to S_4. We want to execute the following loop:
$for\ (i = 0;\ i < 3;\ i++)$
I_1, I_2, I_3, I_4, I_5
where the time taken (in ns) by instructions I_1 to I_5 for stages S_1 to S_4 are given below:

Instruction No.	**S1**	**S2**	**S3**	**S4**
I1	1	2	1	2
I2	1	2	2	1
I3	2	3	2	1
I4	1	3	3	1
I5	2	1	2	1

The time required to execute above code segment is:
(A) 33
(B) 34
(C) 39
(D) 43

17. Instruction execution in a processor is divided into 5 stages, Instruction Fetch (IF), Instruction Decode (ID), Operand fetch (OF), Execute (EX), and Write Back (WB). These stages take 5, 4, 20, 18 and 3 nanoseconds (ns) respectively. A pipelined implementation of the processor requires buffering between each pair of consecutive stages with a delay of 2 ns. Two pipelined implementation of the processor are shown below:
(i) a naive pipeline implementation (NP) with 5 stages and
(ii) an efficient pipeline (EP) where the OF stage is divided into stages OF1, OF2 and OF3 with execution times of 10 ns, 6 ns and 8 ns respectively and Execution stage is divided into EX1 and EX2 with execution times of 11 ns and 9 ns respectively with advanced inter-stage-registers having delay of 1 ns. The speedup (correct to four decimal places) achieved by EP over NP in executing 100 independent instructions if there are 30% stalls of 2 cycles and 20% stalls of 3 cycles in (i) and 40% stalls of 2 cycles and 10% stalls of 3 cycles in (ii) :
(A) 1.7275
(B) 1.8925

(C) 2.0225
(D) 1.4578

18. An instruction pipeline has six stages where each stage take 3 nanoseconds and all instruction use all six stages. Branch instructions, which are 36% of total instructions, are not overlapped.
If a branch instruction is a conditional branch instruction, the branch need not be taken. If the branch is not taken the following instructions can be overlapped. Note that 75% of all branch instructions are conditional branch instructions, and (approximately) 60% of the conditional branch instructions are such that the branch is not taken.
Which of the following represents the average instruction execution time for 100 instructions (in nanoseconds).(Choose the value closest to the right answer)
(A) 360
(B) 400
(C) 560
(D) 600

19. A CPU has a eight-stage pipeline and runs at 4 GHz frequency. Instruction fetch happens in the first stage of the pipeline. A conditional branch instruction computes the target address and evaluates the condition in the sixth stage of the pipeline. The processor stops fetching new instructions following a conditional branch until the branch outcome is known. A program executes 10^8 instructions out of which 20% are conditional branches. If each instruction takes two clock cycles to complete on an average and the total execution time of the program is T seconds, then the value of $57 * T$ is:
(A) 5.7
(B) 123.33
(C) 189.67
(D) 57.57

20. Instruction scheduling can be used to eliminate data and control hazard by:
(A) Scheduling the execution of the instruction only if there is no hazard.
(B) Allowing the compiler to move instructions around to fill the LOAD/BRANCH delay slot(s) with meaningful instructions.
(C) Using a special hardware to check for hazard and issue instructions only when possible.
(D) None of the above.

21. Consider a pipelined processor with the following five stages:
IF: Instruction Fetch

ID: Instruction Decode and Operand Fetch
EX: Execute
MM: Memory access
WB: Write Back
The IF, ID and WB stages take one clock cycle each to complete the operation. The MM stage takes 4 clock cycles to load the word while it takes 6 clock cycles to store the word. The number of clock cycles for the EX stage depends on the instruction. The ADD instruction needs 1 clock cycle, SUB instruction needs 2 clock cycle, the MUL and Memory instructions need 3 clock cycles in the EX stage. Operand forwarding is used in the pipelined processor. An instruction not needing any resource may use bypassing. What is the number of clock cycles taken to complete the following sequence of instructions?

$ADD\ R_2, R_1, R_0;$	$R_2 \leftarrow R_1 + R_0$
$MUL\ R_4, R_3, R_2;$	$R_4 \leftarrow R_3 * R_2$
$SUB\ R_6, R_5, R_4;$	$R_6 \leftarrow R_5 - R_4$
$STORE\ R_6, (2019);$	$M(2019) \leftarrow R_6$

(A) 17
(B) 18
(C) 20
(D) 22
(E) 27

22. Consider a non-pipelined processor with a clock rate of 2.5 gigahertz and average cycles per instruction of five. The same processor is upgraded to a pipelined processor with five stages; but due to the internal pipeline delay, the clock speed is reduced to 2 gigahertz. Assume that there are 60% instruction which do not cause STALLS in the pipeline while remaining do introduce a stall of 1.1 ns. The speedup achieved in this pipelined processor is:
(A) 1.57
(B) 1.82
(C) 2.067
(D) 2.2
(E) 2.75

23. Consider a 5-stage pipeline:
IF (Instruction Fetch), ID (Instruction Decode and register read), EX (Execute), MEM (memory), and WB (Write Back). All (memory or register) reads take place in the second phase of a clock cycle and all writes occur in the first phase. Consider the execution of the following instruction sequence:
I1: sub r2, r3, r4; /*r2 ← r3 - r4 */
I2: sub r4, r2, r3; /*r4 ← r2 - r3 */
I3: sw r2, 100(r1); /*M[r1 + 100] ← r2 */

I4: sub r3, r4, r2; /*r3 ← r4 - r2 */
If number of true-data dependency is denoted by T, number of anti-data dependency is denoted by A and number of output-data dependency is denoted by P, which of the followings represents the value of the expression: $T + A * 2 - 2 * P$
(A) 18
(B) 34
(C) 6
(D) 10

24. A processor designer is provided with two designs D_1 and D_2 for a pipelined processor. D_1 has 6 pipeline stages with execution times of 2 nsec, 3 nsec, 4 nsec, 3 nsec, 2 nsec and 6 nsec while the design D_2 has 8 pipeline stages each with 4 nsec execution time. How much time can be saved using design D_2 over design D_1 for executing 100 instructions?
(A) 416 nsec
(B) 318 nsec
(C) 220 nsec
(D) 202 nsec
(E) 196 nsec

25. A non pipelined single cycle processor operating at 2.5 GHz is converted into a synchronous pipelined processor with five stages requiring 150 ps, 20 ps, 170 ps, 200 ps and 25 ps respectively. The speedup of the pipeline processor for a large number of instructions is if the delay of the latches is 10 ps.
(A) 2.8
(B) 2.5
(C) 3.33
(D) 1.9

26. A pipelined processor uses a 4-stage instruction pipeline with the following stages: Instruction fetch (IF), Instruction decode (ID), Execute (EX) and Writeback (WB). The arithmetic operations as well as the load and store operations are carried out in the EX stage.
The sequence of instructions corresponding to the statement:
$X = ((R - S) * (P + Q * S))/T$
,is given below. The values of variables P, Q, R, S and T are available in the registers R0, R1, R2, R3 and R4 respectively, before the execution of the instruction sequence.
$MUL\ R_1,\ R_1,\ R_3\ \ ;\ R_1 \leftarrow R_1 * R_3$
$ADD\ R_0,\ R_0,\ R_1\ \ ;\ R_0 \leftarrow R_0 + R_1$
$SUB\ R_2,\ R_2,\ R_3\ \ ;\ R_2 \leftarrow R_2 - R_3$
$MUL\ R_0,\ R_0,\ R_2\ \ ;\ R_0 \leftarrow R_0 * R_2$

$DIV\ R_0,\ R_0,\ R_4 \quad ;\ R_0 \leftarrow R_0/R_4$
$STORE\ R_0,\ X \quad ;\ X \leftarrow R_0$
The number of Read-After-Write (RAW) dependencies, Write-After-Read(WAR) dependencies, and Write-After-Write (WAW) dependencies in the sequence of instructions are, respectively,
(A) 4, 3, 2
(B) 5, 2, 3
(C) 4, 2, 2
(D) 6, 3, 3
(E) 8, 8, 3

27. Consider the sequence of machine instruction given below:
$MUL \quad R_5, R_0, R_1$
$DIV \quad R_6, R_2, R_3$
$ADD \quad R_7, R_5, R_6$
$SUB \quad R_8, R_7, R_4$
$STR \quad 0(R_2), R_8$
In the above sequence, R_0 to R_8 are general purpose registers. In the instructions shown, the first register shows the result of the operation performed on the second and the third registers. This sequence of instructions is to be executed in a pipelined instruction processor with the following 5 stages: (1) Instruction Fetch and Decode (IF), (2) Operand Fetch (OF), (3) Perform Operation (PO), (4) Memory access (MA) and (5) Write back the result (WB). The IF, OF and WB stages take 1 clock cycle each for any instruction. The PO stage takes 1 clock cycle for ADD and SUB instruction, 3 clock cycles for MUL instruction and 5 clock cycles for DIV instruction. MA stage takes 4 clock cycles to read and 5 clock cycles to write to memory. The pipelined processor uses operand forwarding from the PO stage to the OF stage. The number of clock cycles taken for the execution of the above sequence of instruction is:
(A) 25
(B) 21
(C) 27
(D) 19

28. Consider following question which share exactly the same context with the above question.
The IF and ID stages take 1 clock cycle each while WB stage takes 2 cycles. The EX stage takes 1 clock cycle for the ADD, 2 clock cycles for SUB, 3 clock cycles for STORE operations and 4 clock cycles each for MUL and DIV operations. Operand forwarding from the EX stage to the ID stage is used. The number of clock cycles required to complete the sequence of instructions is
(A) 24
(B) 26

(C) 28
(D) 29

29. An instruction pipeline has five stages, namely, instruction fetch (IF), instruction decode and register fetch (ID/RF), instruction execution (EX), memory access (MEM), and register writeback (WB) with stage latencies 1 ns, 3 ns, 2 ns, 1 ns, and 0.75 ns, respectively (ns stands for nanoseconds). To gain in terms of frequency, the designers have decided to split the ID/RF stage into three stages (ID, RF1, RF2) each of latency in the ratio of 1 : 2 : 3 in ns respectively. Also, the EX stage is split into two stages (EX1, EX2) each of latency 1 ns. The new design has a total of eight pipeline stages. A program has 30% branch instructions which execute in the EX stage and produce the next instruction pointer at the end of the EX stage in the old design and at the end of the EX2 stage in the new design. The IF stage stalls after fetching a branch instruction until the next instruction pointer is computed. All instructions other than the branch instruction have an average CPI of one in both the designs. The execution times of this program on the old and the new design are P and Q nanoseconds, respectively. The value of $\frac{P}{Q}$ is:
(A) 1.28
(B) 3.12
(C) 2.50
(D) 2.05

30. Consider a RISC processor with all of its instructions are of type register-register arithmetic instructions that have the format $R_1 \leftarrow R_2$ *op* R_3. The pipeline for these instructions runs with a 1 GHz clock with the following stages: instruction fetch = 1 clocks, instruction decode = 2 clock, fetch operands = 3 clock, execute = 4 clocks, and store result = 1 clock. Which of the following rates represents the rate (in MIPS) at which processor can execute register-register instructions that have no data dependencies with other instructions?
(A) 125
(B) 250
(C) 350
(D) 450
(E) 566.67

31. In context with the above question, at what rate (in MIPS) can processor execute the instructions when every instruction depends on the results of the previous instruction?
(A) 125
(B) 250
(C) 350

(D) 450
(E) 566.67

32. In context with the above question, consider we implement internal forwarding. At what rate (in MIPS) can processor now execute the instructions when every instruction depends on the results of the previous instruction?
(A) 125
(B) 250
(C) 350
(D) 450
(E) 566.67

33. You are asked to design a pipeline for a new instruction architecture. The unpipelined implementation of each instruction in this architecture has 6-cycle execution. The names for the stages in the pipeline are the same as those used for the cycles in the unpipelined implementation:
IF = instruction fetch,
RF = instruction decode and register fetch,
ALU1 = effective address calculation for memory references and branches,
MEM = memory access,
ALU2 = ALU operations and branches comparison,
WB = write back.
Assume all stages in this pipeline are perfectly balanced and the instruction is started every clock cycle. 10 instructions, out of which third and sixth instructions are branch instructions (branches never taken) with two delay slots, are executed on this pipeline. The number of cycles needed to complete this sequence of instructions is: —?

34. Delayed branching can help in the handling of control hazards. For all delayed conditional branch instructions, irrespective of whether the condition evaluates to true or false,
(A) The instruction following the conditional branch instruction in memory is executed
(B) The first instruction in the fall through path is executed
(C) The first instruction in the taken path is executed
(D) The branch takes longer to execute than any other instruction

35. Consider a non-pipelined architecture of a processor, which when redesigned as a 6 stage pipelined counter-part the clock gets slower by 25%. When an application is executing on this 6-stage pipeline, the speedup achieved with respect to non-pipelined execution if 25% of the instructions incur 2 pipeline stall cycles is:

(A) 2.5
(B) 3.2
(C) 4.0
(D) 4.5

36. Consider the following code in a 5-stage MIPS pipeline with forwarding. Assume the branch is resolved in the ID stage, and there is only one memory port to handle instruction and data. The initial value for R1 is arrayBase and for R8 is arrayBase+40.
Loop: LW R2, 0(R1)
SUB R4, R2, R3
SW R4, 0(R1)
LW R5, 4(R1)
SUB R6, R5, R3
SW R6, 4(R1)
ADDI R1, R1, 8
BNE R1, R8, Loop

(a) Show all possible pipeline hazards in the code. Draw a pipeline chart and solve those hazards using stalls.
(b) Suppose there are enough memory and you are allowed to schedule the code by changing the order of instructions. How does the scheduled code compare to the original unscheduled code in clock cycles?

37. The 5 stages of the processor have the following latencies (in ps):

TYPE	IF	ID	EX	MM	WB
a	300	400	350	550	100
b	200	150	100	190	140

Assume that each pipeline stage costs 20ps extra due to the registers between pipeline stages.
Which of the following pairs of tuples represents the latency of an instruction and the throughput for a Non-pipelined processor for both the types a and b mentioned above?
(A) (780,1/780) (1100, 1/1100)
(B) (1700, 1/1700), (1100, 1/1100)
(C) (780,1/780) , (1700, 1/1700)
(D) (1700, 1/1700), (780,1/780)

38. This question is based on the data provided in previous question.
Which of the following pairs of tuples represents the cycle time, the la-

tency of an instruction and the throughput for a pipelined processor for both the types a and b mentioned above?
(A) (570, 2800, 1/570) (220, 1100, 1/1100)
(B) (220, 1700, 1/1700), (220, 1100, 1/1100)
(C) (220, 2800 ,1/2800) , (220, 1100, 1/1100)
(D) (570, 2850, 1/570), (220, 1100, 1/220)

39. To improve the pipelined processor's design further one of the pipeline stages is split into 2 equal halves. Answer the following questions:
(a) Which one would you choose?
(b) What is the new cycle time?
(c) What is the new latency?
(d) What is the new throughput?

40. Consider following implementation of a pipelined processor:
Consider a RISC processor with all of its instructions are of type register-register arithmetic instructions that have the format $R_1 \leftarrow R_2$ *op* R_3. The pipeline for these instructions runs with a 1 GHz clock with the following stages: instruction fetch = 1 clocks, instruction decode = 2 clock, fetch operands = 3 clock, execute = 4 clocks, and store result = 1 clock.
For the RISC processor described above, we decide to implement branches by always assuming the branch will not be taken rather than implementing some form of branch prediction and we do not implement internal forwarding. We do not know that the instruction is a branch until stage 2 (decode) and we do not know the condition code (A condition which is tested in a branch instruction) also until stage 5 (operand store) is complete. Thus we can not provide the target address (of a branch taken) to stage 1 until the end of stage 5.
Assume a sequence of instructions where the instruction which modify condition-code immediately precedes the conditional branch instruction. Only three sequential instructions following the branch instructions are fetched.
Which of the following represents the penalty (in lost cycles) if the branch is not taken?
(A) 0
(B) 2
(C) 3
(D) 4
(E) 5

41. Answer this question based on the data given in previous question. If the branch is taken which of the following represents the penalty (in lost cycles)?
(A) 2

(B) 3
(C) 5
(D) 6
(E) 7

42. Answer this question based on the data given in previous question. If we implement delayed branching (with one delay slot obviously) and the conditional branch is a delayed conditional branch which of the following correctly represents the penalty (in lost cycles) for the delayed branch taken?
(A) 2
(B) 3
(C) 5
(D) 6
(E) 7

43. Answer this question based on the data given in previous question. Now we implement internal forwarding along with the delayed branch. What penalty (in lost cycles) does it incur for the delayed branch taken with internal forwarding?
(A) 2
(B) 3
(C) 5
(D) 6
(E) 7

44. Consider the MIPS32 pipeline with ideal CPI of 1. Assume that 30% of all instructions executed are branch, out of which 80% are taken branches. The penalties for predict taken and delayed branch are 1 and 0.5 respectively. The pipeline speedup for predict taken and delayed branch approaches to reduce branch penalties will be:
(A) 4.10 and 4.45
(B) 3.25 and 4.35
(C) 3.67 and 4.25
(D) 3.85 and 4.35

45. For the 1-bit prediction scheme, which of the following statements are false?
(A) There will be two mispredictions whenever the behavior of a branch instruction changes.
(B) An entry in the Branch prediction Buffer (BPB) uniquely identifies a branch instruction
(C) Does not give any advantage for the MIPS32 pipeline.

(D) None of the above

46. Consider an MIPS pipeline with 8 stages, where branch target locations are known in stage 3 and branch conditions are evaluated in stage 4.
Note following abbreviations: Unconditional Branch as UB, Conditional Branch not taken as CBNT, Conditional Branch Taken as CBT, Predict Taken as PT and Predict not Taken as PNT.

	UB	**CBNT**	**CBT**
PT	2	3	2
PNT	2	0	3

Assume a benchmark with 4% unconditional branches, 6% conditional branches with 70% conditional branches taken.
Which of the following tuples represents the Predict taken penalty value and Predict not taken penalty value in (a, b) format where a and b are ten times of the corresponding penalties:
(A) (3.44, 2.06)
(B) (0.344, 0.206)
(C) (3.242, 2.66)
(D) (0.218, 0.206)

47. Delayed branching can help in the handling of control hazards The following code is to run on a pipelined processor with three branch delay slots:
$I_1: \ R_5 \leftarrow R_7 + R_8$
$I_2: \ R_3 \leftarrow R_5 - R_6$
$I_3: \ R_{10} \leftarrow R_5 * R_2$
$I_4: \ R_2 \leftarrow R_6 + R_3$
$I_5: \ R_1 \leftarrow R_2 + R_3$
$I_6: \ R_4 \leftarrow R_9 + R_{10}$
$I_7: \ Memory[R_4] \leftarrow R_1$
$Goto \ Label \ if \ R_1 == R_2$
Set of instructions among I_1, I_2, I_3, I_4, I_5, I_6 or I_7 which can legitimately occupy the delay slot without any program modification?
(A) I_1, I_2, I_7
(B) I_3, I_2, I_7
(C) I_2, I_4, I_7
(D) I_3, I_6, I_7

48. The instruction set architecture of MIPS has three 32-bit instruction formats known as I-type, Rtype, and J-type. Suppose we want to increase the number of registers in the MIPS architecture from 32 to 64 without

changing the instruction size.
(a) How many total bits should be reassigned to the register fields in each instruction type?
(b) If the bits are taken exclusively from the opcode, show the new bit layout for each instruction type. How many opcodes will this new architecture support?

49. Consider an instruction pipeline with six stages without any branch prediction: Fetch Instruction (FI), Decode Instruction (DI), Fetch Operand (FO), Execute Instruction (EI), Memory access (MA) and Write Operand (WO). The stage delays for FI, DI, FO, EI, MA and WO are 2800 ps, 1700 ps, 5500 ps, 3600 ps, 6250 ps and 2300 ps respectively. There are intermediate storage buffers after each stage and the delay of each buffer is 1400 ps. A program consisting of 20 instructions $I_1, I_2, I_3, \ldots, I_{20}$ is executed in this pipelined processor. Instruction I_4 and I_{13} are the only branch instructions and their branch targets are I_9 and I_{18} respectively. If the branch is taken during the execution of this program and Branch targets are computed in EI stage, the time (in ns) needed to complete the program is
(A) 132
(B) 165
(C) 176
(D) 328

50. Consider the following code from an if-else statement of the form:
if (A==0)
A = B;
else
A = A + 8; ,where A is at 0(R2) and B is at 0(R3):
LD R1, 0(R2)
BNEZ R1, L1
LD R1, 0(R3)
J L2
L1: DADDI R1, R1, 8
L2: SD R1, 0(R2)
Starting with a standard 5-stage MIPS pipeline with forwarding and branch resolution in the ID stage, you are asked to design a new conditional load instructions 'LDZ Rd, x(Rsl), Rs2' and 'LDNZ Rd, x(Rsl), Rs2' that do not load unless the value of Rs2 is zero or not zero, respectively.
a) Write the code using this new conditional load instruction and show possible stall cycles that would occur in the pipeline.
b) Compare the clock cycles used to that of the original code. If implementing the new conditional load instruction would increase the clock cycle by 10%, will it be worth making that change?