multiple choice

The memory access time is _____.
a) The time from initiation to completion of a word or byte transfer from RAM
N the cache memory that uses write-through protocol, if a write miss occurs _____.
b) The data is written directly to the main RAM.
In pipelined RISC processors, data dependencies can be handled.
d) All of the above.
A block-direct-mapped cache consists of a total of 32 blocks. The main memory contains
2K blocks, each consisting of 8 words. Each word is 4 bytes. Assuming a 16-bit byte-addressable
Address space, how many bits are there in each of the Tag, Block, and Word fields?
c) Tag= 6 bits. Block=5 bits, Word= 5 bits.

multiple choice (cont)

A block-associative-mapped cache consists of a total of 32 blocks. The main memory contains
2K blocks, each consisting of 8 words. Each word is 4 bytes. Assuming a 16-bit byte-addressable
Address space, how many bits are there in each of the Tag, Block, and Word fields?
A) Tag= 11 bits. Block=0 bits, Word= 5 bits.
A block-4-way associative-mapped cache consists of a total of 32 blocks. The main memory
Contains 2K blocks, each consisting of 8 words. Each word is 4 bytes. Assuming a
16-bit byte-addressable address space, how many bits are there in each of the Tag, Set, and Word fields?
B) Tag= 9 bits. Set=2 bits, Word= 5 bits.
By using the Booth algorithm recoding technique, the binary multiplier

multiple choice (cont)

0 1 0 0 1 1 0 1 1 1 0 will be recoded as:
A) +1 -1 0 +1 0 -1 +1 0 0 -1 0
The processor uses the Instruction Register (IR) to:
Keep the instruction to be executed until it is finished.
Q1: How many chips you need to build a memory module of size
4M 8 bits, if you only have
Chips of size 256K 1 bits.
a) 128
Q2: A block-direct-associative cache consists of a total of 64 blocks. The main memory
Contains 1K blocks, each consisting of 8 words. Each word is 4 bytes. Assuming a
16-bit byte-addressable address space, how many bits are there in each of the Tag, Set,
And Word fields?
a) Tag= 5 bits. Block=6 bits, Word= 5 bits.

multiple choice (cont)

Q3: The memory access time is:
a) The time from initiation to completion of a word or byte transfer.
Q4: Compared to the static RAM, the dynamic RAM is:
a) Slower, Cheaper, Simpler, and always needs refreshing.
Q5: The Double-Data-Rate (DDR) SDRAM:
a) Transfers data on the rising and falling edges of its own clock cycle.
Q6: The hierarchy of the memory in a computer is (starting from the processor):
a) Registers, Cache, RAM, and Hard disk.
Q7: The flash memory is made of:
a) EEPROM cells.
Q1: How many chips you need to build a memory module of size
4M 8 bits, if you only
Have chips of size 256K 2 bits.
a) 64

multiple choice (cont)

Q2: A block-direct- cache consists of a total of 32 blocks. The main memory
Contains 4K blocks, each consisting of 4 words. Each word is 4 bytes. Assuming
A 16-bit byte-addressable address space, how many bits are there in each of the Tag,
Set, and Word fields?
a) Tag= 6 bits. Block=5 bits, Word= 4 bits.
Q8: To solve the problem of data dependency in pipelining architectures, ----------:
a) Processors use data forwarding.
b) The compiler insert NOP operation between the instructions.
c) The pipeline has to stall until the operands are ready.
d) Any of the other options.

multiple choice (cont)

Q9: To solve the problem of branch penalty in pipelining architectures, ----------:
a) Processors use delayed branching technique.
b) Processors use branch prediction technique.
c) Processors use branch target buffer.
d) All of the other options.
Q10: The pipeline may stall because of:
a) Data dependency.
b) Cache miss.
c) Limited hardware resources.
d) All of the other options.
Q11: Although the throughput increases by increasing the number of stages in
The pipeline, the number of
a) The probability of data dependency will increase.
b) The branch penalty will increase.

multiple choice (cont)

c) The cache misses may increase.
d) All of the other options
Q2: The processor's control signals are generated by:
a) Only hardware and called Hardwired approach
b) Only software and called Microprogrammed approach
c) Either (a) or (b)
Q3: In pipelining, to alleviate the problem of branch delayed slot problem:
a) The compiler tries to find a suitable instruction that precedes the branch and move it after the branch instruction.
Q5: The Double-Data-Rate (DDR) SDRAM:
a) Transfers data on the rising and falling edges of its own clock cycle.
Q1: By using the Booth algorithm recoding technique, the binary multiplier 0 1 0 1 1 1 0 1 0 1 1 will be:

multiple choice (cont)

b) a) +1 -1 +1 0 0 -1 +1 - 1 +1 0 -1
Q2: Which statement from the following is true?
a) The ripple-carry adder is slower than the carry-lookahead adder.
Q3: Which statement from the following is true?
a) The 2-dimensional combinational array is faster than sequential circuit multiplier.
Q8: How many chips you need to build a memory module of size 1G 32 bits, if you only have chips of size 128M 64 bits.
a) 4
Q9: A block direct-mapped cache consists of a total of 128 blocks. The main memory contains 64K blocks, each consisting of 16 words. Each word is 4 bytes. Assuming a 16-bit byte-addressable address space, how many bits are there in each of the Tag, Block, and Word fields?

multiple choice (cont)

a) Tag= 3 bits. Block=7 bits, Word= 6 bits.
Q10: A block associative-mapped cache consists of a total of 64 blocks. The main memory contains 64K blocks, each consisting of 8 words. Each word is 4 bytes. Assuming a 16-bit byte-addressable address space, how many bits are there in each of the Tag, Block, and Word fields?
a) Tag= 11 bits. Block=0 bits, Word= 5 bits.
Q13: The access time for the hard disk is:
a) The sum of the seek time and rotational delay.
Q14: The Digital Versatile Disk stores up to 17GB because:
a) It uses two-layered two-sided disks, and red-light laser.
Q15: The CD-ReWritables (CD-RW) uses:
a) Three different laser powers with an organic dye and alloy of materials in the recording layer.

multiple choice (cont)

Q16: The magnetic hard disk is formatted as:
a) Tracks and sectors in each platter.
Q17: In virtual memory, a page fault occurs:
a) When a virtual address has no corresponding physical address.
Q18: Which statement is true?
a) The processor always issues a virtual address and the MMU translates it to a physical address.
Q19: In the cache memory, the no hit occurs when:
a) The tag match occurs and the valid bit is 0.
Q20: In the cache memory that uses write-back protocol, if a write miss occurs:
a) First transfer block containing the addressed word into the cache and then overwrite specific word in cached block.
Q21: The cache memory makes the RAM to appear to the processor as much faster because of:

multiple choice (cont)

a) The locality of reference of computer programs.
Q22: The choice of a RAM chip for a given application depends on:
a) The cost, speed, power dissipation, and size of the chip.
Q23: The main problem in DRAMs is:
a) They are slow and need refreshing of their contents.
Q24: The main problem in static RAMs is:
a) They use six transistors to build each cell.
Q25: Pipelining is used in RISC processors to:
a) Increase the processor's throughput.
Q26: The 5-satges in the pipelined RISC processors is in the following order:
a) Fetch, Decode, Compute, Memory, Write.
Q27: In pipelined RISC processors, data dependencies can be handled by:

multiple choice (cont)

a) Any one of the other options.
Q28: In pipelined RISC processors, the ideal time (throuput) needed to execute one instruction is:
a) One clock cycle.
Q29: In pipelined RISC processors, what makes the actual throughput less than the ideal one is:
a) All of the other options.
b) The stall times due to data dependencies.
c) The branch penalties.
d) The caches misses.
Q30: The processor uses the program counter (PC) to:
a) Keep track of the address of the next instruction to be fetched and executed.
Q31: In pipelined RISC processors, executing the instruction (Add R3, R4, R5) requires the following steps in order:
a) Fetch, Decode, Compute, No-action, and Write result into register.

multiple choice (cont)

Q32: The processor's control signals can be generated by using:
a) Either hardwired or microprogramming control

rate of instruction execution

1-GHz clock. Instruction statistics in a large program are as follows: Branch 20% Load 20% Store 10% Computational instructions 50%

90% of instruction fetch operations are completed in one clock cycle and 10% are completed in 4 clock cycles. On average, access to the data operands of a Load or Store, instruction is completed in 3 clock cycles

> On average, instruction fetch takes 0.9+0.1x4 = 1.3 cycles. All instructions, except Load and Store, take four more cycles to complete. Load and Store instructions take two additional cycles, on average. Average completion time = 1.3 + (0.2 + 0.5) x 4 + (0.2 + 0.1) x 6 = 5.9 cycles Instruction rate = 109 /5.9 = 169.5 million instructions per second

a) Access to the memory is always completed in 1 clock cycle.

> Execution rate = 1x109 /5 = 200 million instructions per second

chapter 5 example

Call_register R9

• Calls a subroutine whose address is in register R9:

1. Memory address <- [PC], Read memory, IR <-Memory data, PC <-[PC] <- 4
2. Decode instruction, RA <-[R9]
3. PC-Temp<-[PC], PC <-[RA]
4. RY <- [PC-Temp]
5. Register LINK <- [RY]

chapter 5 example (cont)

Q1: Assume that all memory access operations are completed in one clock cycle in a processor that has a 1-GHz clock. What is the frequency of memory access operations if Load and Store instructions constitute 20 percent of the dynamic instruction count in a program? (The dynamic count is the number of instruction executions, including the effect of program loops, which may cause some instructions to be executed more than once.) Assume that all instructions are executed in 5 clock cycles

There is one memory access to fetch each instruction. Then, 20 percent of the instructions have a second memory access to read or write a memory

chapter 5 example (cont)

operand. On average, each instruction has 1.2 memory accesses in 5 clock cycles. Therefore, the frequency of memory accesses is (1.2/5) × 10^9 , or 240 million accesses per second.(1 Ghz= 10^9 hz)

Give the sequence of actions for a Return-from-subroutine instruction in a RISC processor. Assume that the address LINK of the general-purpose register in which the subroutine return address is stored is given in the instruction field connected to address A of the register file (IR31–27).

Whenever an instruction is loaded into the IR, the contents of the generalpurpose register whose address is given in bits IR31–27 are read and placed

chapter 5 example (cont)

into register RA. Hence, a Return-from-subroutine instruction will cause the contents of register LINK to be read and placed in register RA. Execution proceeds as follows:
1. Memory address←[PC], Read memory, Wait for MFC,
IR←Memory data, PC←[PC] + 4 2. Decode instruction, RA←[LINK] 3. PC←[RA] 4. No action 5. No action
At the time the instruction Load R6, 1000(R9) is fetched, R6 and R9 contain the values 4200 and 85320, respectively. Memory location 86320 contains 75900.
step 1 and 2: the values are determined by the previous instructions

chapter 5 example (cont)

stept 3: RA = 85320, RB = 4200. Step 4: RA = 85320, RB = 4200, RZ = 86320, RM= 4200. Step 5: RA = 85320, RB = 4200, RZ = 86320, RM= 4200 RY = 75900.

16M × 32 memory using 1M × 4 memory chips

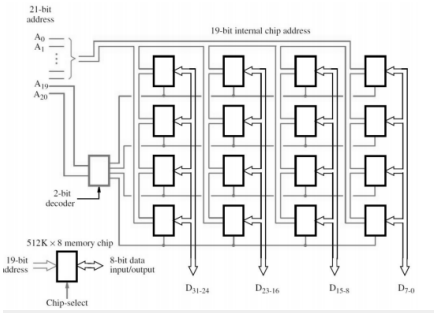

A 16M module can be structured as 16 rows, each containing eight 1M x 4 chips. A 24-bit address is required. Address lines A19-0 should be connected to all chips. Address lines A23-20 should be connected to a 4-bit decoder to select one of the 16 rows.

A block-set-associative cache

A block-set-associative cache consists of a total of 64 blocks, divided into 4-block sets. The main memory contains 4096 blocks, each consisting of 32 words. Assuming a 32-bit byte-addressable address space, how many bits are there in each of the Tag, Set, and Word fields?

A block-set-associative cache (cont)

Each block contains 128 bytes, thus requiring a 7-bit Word field. There are 16 sets, requiring a 4-bit Set field. The remaining 21 bits of the address constitute the tag field.

problem set Mult

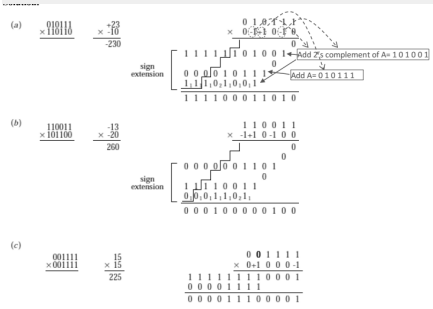

A = 010111 and B = 110110
b) A = 110011 and B = 101100
c) A = 001111 and B = 001111

Question about the sequential circuit

how to implement multiplication of 2's-complement n-bit numbers using the Booth algorithm, by clearly specifying inputs and outputs for the Control sequencer and any other changes needed around the adder and register A.
Both the A and M registers are augmented by one bit to the left to hold a sign extension bit.

Question about the sequential circuit (cont)

The adder is changed to an n + 1-bit adder. A bit is added to the right end of the Q register to implement the Booth multiplier recoding operation. It is initially set to zero. The control logic decodes the two bits at the right end of the Q register according to the Booth algorithm. The right shift is an arithmetic right shift as indicated by the repetition of the extended sign bit at the left end of the A register..

R2, R3, R4, R5, R6, and R7

```
Load R2, #A_VEC
Load R3, #B_VEC
Load R4, #3
And R5, R5, R0
 LOOP: Load R6, (R2)
Load R7, (R3)
Multiply R6, R6, R7
Add R5, R5, R6
Add R2, R2, #4
Add R3, R3, #4
Subtract R4, #1
Branch>0 LOOP
```

R2, R3, R4, R5, R6, and R7 (cont)

```
Load R7, # RESULT
Store R5. (R7)
End
ORIGIN 500
A_VEC: DATAWORD 05, -20,
10
B_VEC: DATAWORD 09, 04, 07
RESULT: RESERVE 4
R2=512, R3=524, R4=0,
R5=25 R6=7 R7=524
```

pipeline provides forwarding paths

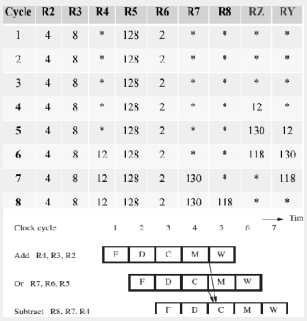

Cycle	R2	R3	R4	R5	R6	R7	R8	RZ	RY
1	4	8	*	128	2	*	*	*	*
2	4	8	*	128	2	*	*	*	*
3	4	8	*	128	2	*	*	*	*
4	4	8	*	128	2	*	*	12	*
5	4	8	*	128	2	*	*	130	12
6	4	8	12	128	2	*	*	118	130
7	4	8	12	128	2	130	*	*	118
8	4	8	12	128	2	130	118	*	*

The result from theALU is 130 – 12 = 118. This result is available in register RZ during cycle 6. The result of the Or instruction, 130, is in register RY during in cycle 6. In cycle 6, the Subtract instruction is in the Memory stage. The unspecified instruction following the Subtract instruction is generating a result in the Compute stage. In cycle 7, the result of the unspecified instruction is in register RZ, and the result of the Subtract instruction is in register RY.

Chapter 6 execution time and speed up

Assume that 20% of the dynamic count of the instructions executed for a program are branch instructions. There are no pipeline stalls due to data dependencies. • Static branch prediction is used with a not-taken assumption. a) Determine the execution times for two cases: when 30 percent of the branches are taken, and when 70 percent of the branches are taken. b) Determine the speedup for one case relative to the other. Express the speedup as a percentage relative to 1

In first case, 30% of branches are taken but we assumed not-taken, so they are mispredicted (one cycle penalty 30% from 20% branch instructions): The value of δbranch_penalty = 0.20 × 0.30 ×1 = 0.06

Chapter 6 execution time and speed up (cont)

In second case, 70% of branches are taken but we assumed not-taken, so they are mispredicted (one cycle penalty 70% from 20% branchs): The value of δbranch_penalty = 0.20 × 0.70 ×1 = 0.14

Using S = 1 + δbranch_penalty, the execution time:

in first case is (1.06 × N)/R and (1.14 × N)/R for the second case.

b) Because the execution time for the first case is smaller, the performance improvement as a speedup percentage is:(1.14/1.06 - 1)*100=7.5%

Chapter 9

Overflow =cn ⊕ cn−1 or xn−1 yn−1(‾sn−1) +(‾xn−1)(‾yn−1) sn−1

For the subtraction operation X−Y on 2's-complement numbers X and Y • We form the 2's-complement of Y and add it to X . • Set Add/Sub =0 and c0= 0 for addition. • Set Add/Sub =1 and c0 = 1 for subtraction

Chapter 9 (cont)

4-bit adder has four carry-out signals: c1 = G0 + P0 c0, c2 = G1 + P1G0 + P1P0 c0, c3 = G2 + P2G1 + P2P1G0 + P2P1P0 c0, c4 = G3 + P3G2 + P3P2G1 + P3P2P1G0 + P3P2P1P0 c0

A sequence of n addition cycles generates a 2n-bit product Delay = n * (adder + control delays)

For n = 32, delay is approximately 32 14 = 448 gate delays

Registers A and Q are shift registers, together, they hold partial product PPi while multiplier bit qi generates the signal Add/Noadd.At the end of each cycle, C, A, and Q are shifted right one bit position to allow for growth of the partial product as the multiplier is shifted out of register Q

• Non-Restoring Division:

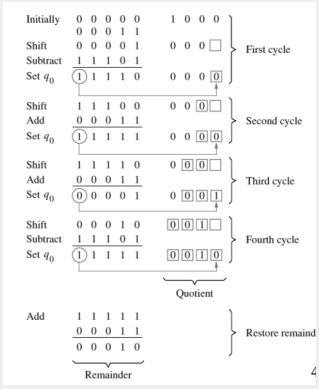

Stage 1: Do the following two steps n times:

1. If the sign of A is 0, shift A and Q left one bit position and subtract M from A; otherwise, shift A and Q left and add M to A.
2. Now, if the sign of A is 0, set q0 to 1; otherwise, set q0 to 0.

Stage 2: If the sign of A is 1, add M to A.

example

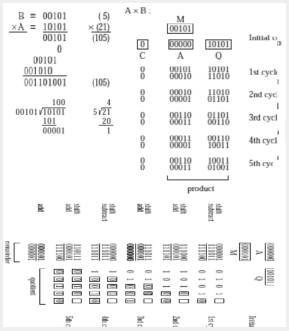

Store R6, X(R8)

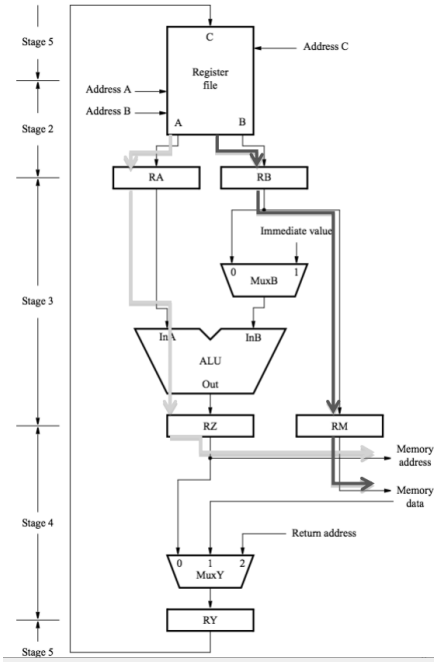

Memory address [PC],
Read memory,
IR Memory data,
PC [PC] 4
2. Decode instruction,
RA [R8], RB [R6]
3. RZ [RA] Immediate value X,
RM [RB]
4. Memory address [RZ],
Memory data [RM],
Write memory
5. No action

Add R5, R6

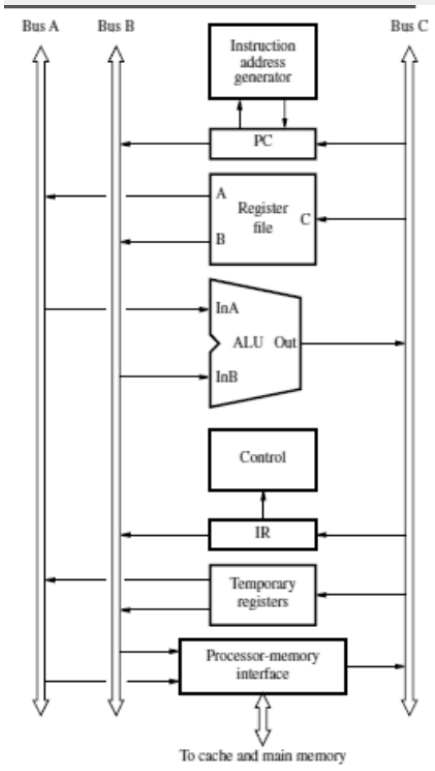

(on Bus B) Memory address ← [PC],
Read memory, Wait for MFC, (on Bus C) IR ← Memory data,
PC ←[PC] + 4
2. Decode instruction
3. The contents of R5 and R6 are sent to ALU
using buses A and B,
R5 ←[R5] + [R6],
the sum is written to R5 using bus C

sequential mult UNSIGNED

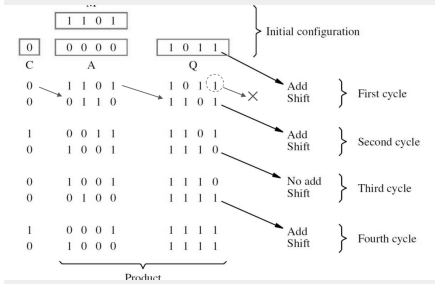

13*11

Booth algorithm

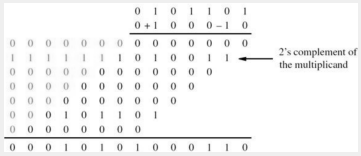

booth mult example

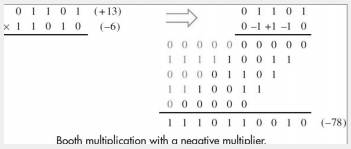

Restoring Division

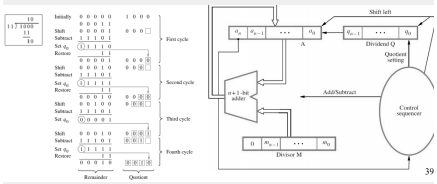

• Do the following three steps n times:
1. Shift A and Q left one bit position.
2. Subtract M from A, and place the answer back in A.
3. If the sign of A is 1, set q0 to 0 and add M back to A (that is, restore A); otherwise, set q0 to 1.