

University of California, Berkeley
College of Engineering
Computer Science Division — EECS

Fall 1999

John Kubiawicz

Midterm I
October 6, 1999
CS152 Computer Architecture and Engineering

Your Name:	
SID Number:	
Discussion Section:	

Problem	Possible	Score
1	20	
2	15	
3	35	
4	30	
Total		

Problem 1: Performance

Problem 1a:

Name the three principle components of runtime that we discussed in class. How do they combine to yield runtime?

Problem 1b:

What is Amdahl's law for speedup? State as a formula which includes a factor for clock rate.

Let us suppose that you have been running an important program on your company's 300MHz Acme II processor. By running a detailed simulator, you were able to collect the following instruction mix and breakdown of costs for each instruction type:

Instruction Class	Frequency (%)	Cycles
Integer arithmetic and logical	40	1
Load	20	1
Store	10	2
Branches	20	3
Floating Point	10	5

Problem 1c:

What is the CPI and MIPS rating of the Acme II for this program?

Problem 1d:

Suppose that you turn on the optimizer and it eliminates 30% of the arithmetic/logic instructions (i.e. 12% of the *total* instructions), 30% of load instructions, and 20% of the floating-point instructions. None of the other instructions are effected. What is the speedup of the optimized program? (Be sure to state the formula that you are using for speedup and show your work)

Problem 1e:

What is the CPI and MIPS rating with the optimized version of the program? Compare your result to that of (1c) and explain the difference:

Problem 1f:

Now, suppose that the Acme III has just been introduced with a faster clock rate (450 MHz). However, in order to make the clock rate faster, the Acme engineers had to increase the CPI for arithmetic, logical, and load instructions to 2 cycles and floating point instructions to 6 cycles. What is the speedup of the Acme III over the Acme II on the *unoptimized* program? Show work

Problem 1g:

The engineers for Acme Inc are currently working on the Acme IV. Instead of increasing the clock rate again, they are working on reducing the time for the floating-point instructions. Use Amdahl's law to show the *maximum* speedup that you could expect between the Acme III and Acme IV on the *unoptimized* program (if the clock rates are both 450 MHz)?

Problem 2: Propagation Delay

Problem 2a:


Assume the following characteristics for NAND gates:

Input load: 150fF,

Internal delay: $T_{Plh}=0.2\text{ns}$, $T_{Phl}=0.5\text{ns}$,


Load-Dependent delay: $T_{Plhf}=.0020\text{ns}$, $T_{Phlf}=.0021\text{ns}$

For the circuit below, assume that inputs $X_0 - X_5$ are all set to 1. What are the propagation delays from A to Y (for rising and falling-edges of Y)?



Problem 2b:

Suppose that we construct a new gate, XOR, as follows:



Compute the standard parameters for the linear delay models for this complex gate, assuming the parameters given above for the NAND gate:

A Input Capacitance:

B Input Capacitance:

Load-dependent Delays:

TPAYlhf:

TPAYhlf:

TPBYlhf:

TPBYhlf:

Internal delays for $A \Rightarrow Y$, assuming that B is set to 1 (worst case delays):


TPAYlh:

TPAYhl:

Problem 2c:

Now, suppose we use our new XOR gate in the circuit below. Let $X_0 - X_5$ be set to 1.

Compute the propagation delays from $A \Rightarrow Y$ (both rising and falling edges):



Problem 3: Square Root

Suppose that you have a 32-bit value, M , and you wish to find the closest integer, S , less than its square-root, \sqrt{M} . Let's call S the "integer square root of M ". Since you are forcing S to be an integer, you will end up with a remainder, $R = M - S^2$. In this problem, we will come up with an iterative mechanism to compute S one bit at a time.

Let us suppose that we have an estimate, S_i , for the square root of M . We will assume that S_i is less than the desired value S , i.e. $S_i \leq S$. Next, assume that we add a small increment to this estimate to make a better estimate, S_{i+1} . Call this increment N_{i+1} :

$$S_{i+1} = (S_i + N_{i+1}) \leq S$$

Now, the remainder after the first estimate is: $R_i = M - S_i^2$,
 while after the second estimate is: $R_{i+1} = M - S_{i+1}^2$
 $= M - (S_i + N_{i+1})^2$
 $= M - S_i^2 - N_{i+1} \times (2 \times S_i + N_{i+1})$

So, each time we pass through the algorithm, we subtract the following from the remainder:

$$R_i - R_{i+1} = N_{i+1} \times (2 \times S_i + N_{i+1})$$

In binary, the increment values (the N 's) are single bits. Thus, each iteration through the algorithm, we multiply the previous estimate by 2, add in the new bit (N_{i+1}), then shift by the number of zeros in N_{i+1} before subtracting from our remainder. This is very much like a divide in which the divisor keeps changing. For example, consider finding the 4-bit square root of 118:

Starting:	$M = R_0 =$	01110110	and	$S_0 =$	0000
Try:	$N_1 =$	1000	-	1000	$\Leftarrow (2 \times S_0 + 1000) \times 1000$
		$R_1 =$		00110110	$\Rightarrow S_1 = S_0 + 1000 = 1000$
Try:	$N_2 =$	0100	-	10100	$\Leftarrow (2 \times S_1 + 0100) \times 0100$
		Result < 0			$\Rightarrow S_2 = S_1 = 1000$
		$R_2 =$		00110110	(unchanged)
Try:	$N_3 =$	0010	-	10010	$\Leftarrow (2 \times S_2 + 0010) \times 0010$
		$R_3 =$		00010010	$\Rightarrow S_3 = S_2 + 0010 = 1010$
Try:	$N_4 =$	0001	-	10101	$\Leftarrow (2 \times S_3 + 0001) \times 0001$
		Result < 0			$\Rightarrow S_4 = S_3 = 1010$
		$R_4 =$		00010010	(unchanged)

Final result: $\sqrt{01110110_2} = 1010_2$ with 10010_2 remainder
 or: $\sqrt{118} = 10$ with 18 remainder!

Problem 3a:

The above example showed unsigned M . Is it easy extend the algorithm for a signed M ?

Problem 3b:

From this point on, assume M is unsigned. For a 64-bit, unsigned-value M , what is the largest possible integer square-root, S_{\max} ? How many bits would it take to represent? Explain without using a calculator. (*hint: Start by finding the smallest integer that is bigger than S_{\max} .*)

Problem 3c:

Also for a 64-bit unsigned-value M , what is the largest possible remainder, R_{\max} ? How many bits would it take to represent? Explain without using a calculator. (*Use the same hint as above.*)

Here is pseudo-code for a square root algorithm. Assume that the input value of M has been restricted so that S_{max} is no more than 31 bits in size and R_{max} is no more than 32 bits in size. Let **Result** and **Remain** be 32-bit global values which will store the square root and remainder respectively. Inputs **M_{hi}** and **M_{low}** are 32-bit arguments that give the upper and lower 32-bits of the input. **This code is modeled after version 3 of the divider from class:**

```

isqrt(Mlow,Mhi) ⇒ (Result, Remainder)
{ /* All temporaries are 32-bit values */
  int nextbit, temp, topbit, lowerbits;

  /* missing initialization instructions */

  while (nextbit > 0) {
    ROL96(topbits,Remainder,lowerbits);

    /* Above restrictions on M ensure temp only 32 bits. */
    temp = (2 * Result) | nextbit;
    if (topbits > 0 || Remainder ≥ temp) {
      Result = Result | nextbit;
      SUBcarry(topbits, Remainder, temp);
    }
    nextbit = nextbit >> 1;
  }
}

```

The ROL96(*hi*, *low*, *extra*) pseudo-instruction takes three 32-bit registers and treats them as a combined 96-bit register. It shifts the combined value left by one position, inserting a zero at the far right (of the extra register).

The SUBcarry(*hi*, *low*, *subvalue*) pseudo-instruction takes three 32-bit registers. It treats the first two as a combined 64-bit register. It subtracts the 32-bit subvalue from this 64-bit register.

Problem 3d:

The pseudo-code is missing some initialization instructions. What should be there? (*hint: look at the example square root again and try to figure out what the various arguments to ROL96 must be. Also, make sure that every variable has an initial value!*):

Problem 3e:

Assume that you have a MIPS processor that is missing the `isqrt()` instruction. Implement `isqrt()` as a procedure. Assume that M_{low} and M_{high} are in the `$a0` and `$a1` registers respectively, and that the `Result` and `Remain` values are returned in registers `$v0` and `$v1` respectively. You can use `ROL96` and `SUBcarry` pseudo-instructions, but don't use any other pseudo-instructions. Make sure to adhere to all MIPS calling conventions!

Problem 3f:

Implement the `ROL96($t0,$t1,$t2)` pseudo-instruction in 7 MIPS instructions. Assume that `$t0`, `$t1`, and `$t2` are the three input registers (with `$t0` the most significant).

(hint: what happens if you use signed `slt` on unsigned numbers?)

Problem 3g:

Implement the `SUBcarry($t0,$t1,$t2)` pseudo-instruction in 3 MIPS instructions.


Problem 3h:

What is the maximum “CPI” of your `isqrt()` procedure? (i.e. what is the total number of cycles to perform an `isqrt`)? Assume that each real MIPS instruction takes 1 cycle, and pseudo-instructions `ROL96` and `SUBcarry` take 7 and 3 cycles respectively:

EXTRA CREDIT [5pts => Save until last!]:

Draw the data path for a hardware square-root engine that does 64-bit square-roots. Explain what you are doing and how this will be controlled.

Problem 4: New instructions for a multi-cycle data path



The Multi-Cycle datapath developed in class and the book is shown above. In class, we developed an assembly language for microcode. It is included here for reference:

Field Name	Values For Field	Function of Field
ALU	Add	ALU Adds
	Sub	ALU subtracts
	Func	ALU does function code (Inst[5:0])
	Or	ALU does logical OR
SRC1	PC	PC \Rightarrow 1 st ALU input
	rs	R[rs] \Rightarrow 1 st ALU input
SRC2	4	4 \Rightarrow 2 nd ALU input
	rt	R[rt] \Rightarrow 2 nd ALU input
	Extend	sign ext imm16 (Inst[15:0]) \Rightarrow 2 nd ALU input
	Extend0	zero ext imm16 (Inst[15:0]) \Rightarrow 2 nd ALU input
	ExtShft	2 nd ALU input = sign extended imm16 \ll 2
ALU Dest	rd-ALU	ALUout \Rightarrow R[rd]
	rt-ALU	ALUout \Rightarrow R[rt]
	rt-Mem	Mem input \Rightarrow R[rt]
Memory	Read-PC	Read Memory using the PC for the address
	Read-ALU	Read Memory using the ALUout register for the address
	Write-ALU	Write Memory using the ALUout register for the address
MemReg	IR	Mem input \Rightarrow IR
PC Write	ALU	ALU value \Rightarrow PCibm
	ALUoutCond	If ALU Zero is true, then ALUout \Rightarrow PC
Sequence	Seq	Go to next sequential microinstruction
	Fetch	Go to the first microinstruction
	Dispatch	Dispatch using ROM

In class, we made our multicycle machine support the following six MIPS instructions:

```
op | rs | rt | rd | shamt | funct = MEM[PC]
op | rs | rt |      Imm16      = MEM[PC]
```

<u>INST</u>	<u>Register Transfers</u>	
ADDU	$R[rd] \leftarrow R[rs] + R[rt];$	$PC \leftarrow PC + 4$
SUBU	$R[rd] \leftarrow R[rs] - R[rt];$	$PC \leftarrow PC + 4$
ORI	$R[rt] \leftarrow R[rs] + \text{zero_ext}(\text{Imm16});$	$PC \leftarrow PC + 4$
LW	$R[rt] \leftarrow \text{MEM}[R[rs] + \text{sign_ext}(\text{Imm16})];$	$PC \leftarrow PC + 4$
SW	$\text{MEM}[R[rs] + \text{sign_ext}(\text{Imm16})] \leftarrow R[rs];$	$PC \leftarrow PC + 4$
BEQ	if ($R[rs] == R[rt]$) then $PC \leftarrow PC + 4 + \text{sign_ext}(\text{Imm16}) \parallel 00$ else $PC \leftarrow PC + 4$	

For your reference, here is the microcode for two of the 6 MIPS instructions:

Label	ALU	SRC1	SRC2	ALUDest	Memory	MemReg	PCWrite	Sequence
Fetch	Add	PC	4		ReadPC	IR	ALU	Seq
Dispatch	Add	PC	ExtShft					Dispatch
RType	Func	rs	rt	rd-ALU				Seq
BEQ	Sub	rs	rt				ALUoutCond	Fetch

In this problem, we are going to add three new instructions to this data path:

```
lui      $rd, <const>      =>  $R[rd] \leftarrow \text{Imm16} \parallel 0000000000000000$ 
multacc $rd, $rs, $rt      =>  $R[rd] \leftarrow (R[rs] \times R[rt]) + R[rd]$ 
bltual  $rs, $rt <offset> => if ( $R[rs] < R[rt]$ ) then
                                 $PC \leftarrow PC + 4 + \text{sign\_ext}(\text{Imm16}) \parallel 00$ 
                                 $R[31] \leftarrow PC + 4$ 
                                else
                                 $PC \leftarrow PC + 4$ 
```

1. The `lui` instruction is familiar to you from the normal MIPS instruction set. It places the 16 bit immediate field into the upper 16 bits of $R[rd]$, filling the lower 16 bits of $R[rd]$ with zeros. *Important note: the encoding for the `lui` instruction has a zero in the `rs` field.*
2. The `multacc` instruction (multiply-accumulate) uses register $R[rd]$ as both a source and a destination register. It multiplies the values $R[rs]$ and $R[rt]$, adds the result to register $R[rd]$, then places the result back into register $R[rd]$. Assume that this instruction does not overflow.
3. The `bltual` instruction (branch on less than unsigned and link) checks to see if $R[rs]$ is less than $R[rt]$. If it is, it will save the PC in $\$ra$ (like `jal`), then branch to the offset.

Problem 4a:

How *wide* are microinstructions in the original datapath (answer in bits and show some work!)?

Problem 4b:

Draw a block diagram of a microcontroller for the unmodified datapath. Include sequencing hardware, the dispatch ROM, the microcode ROM, and decode blocks to turn the fields of the microcode into control signals. Make sure to show all of the control signals coming from somewhere. (*hint: The PCWr, PCWrCond, and PCSrc signals must come out of a block connected to the PCWrite field of the microinstruction.*)

Problem 4c:

Come up with a binary encoding for the ALUDest field of the microinstruction (rd-ALU, rt-ALU, rt-Mem, or blank). Construct logic which maps this binary field to the appropriate control signals from problem 4b.

Problem 4d:

Describe/sketch the modifications needed to the datapath for the new instructions (`lui`, `multacc`, and `bltval`). Assume that the original datapath had only enough functionality to implement the original 6 instructions. Try to add as little additional hardware as possible. Make sure that you are very clear about your changes.

Problem 4e:

Describe changes to the microinstruction assembly language for these new instructions. How wide are your microinstructions now?

Problem 4f:

Write complete microcode for the three new instructions. Include the Fetch and Dispatch microinstructions. If any of the microcode for the original instructions must change, explain how (*Hint: since the original instructions did not use $R[rd]$ as a register input, you must make sure that your changes do not mess up the original instructions*).

Problem 4g:

What are the CPI values for each of the three new instructions?