

CS152
Computer Architecture and Engineering
Lecture 6

Divide, Floating Point, Pentium Bug

Feb 10, 1999

John Kubiatowicz (<http://cs.berkeley.edu/~kubitron>)

lecture slides: <http://www-inst.eecs.berkeley.edu/~cs152/>

2/10/99

©UCB Spring 1999

CS152 / Kubiatowicz
Lec6.1

Outline of Today's Lecture

- Recap of Last Lecture and Introduction of Today's Lecture (4 min.)
- Divide (20 min.)
- Questions and Administrative Matters (2 min.)
- Floating-Point (25 min.)
- Questions and Break (5 min.)
- Pentium Bug (25min.)

2/10/99

©UCB Spring 1999

CS152 / Kubiatowicz
Lec6.2

Recap of Last Lecture: Summary

- Intro to VHDL
 - entity = symbol, architecture ~ schematic, signals = wires
 - behavior can be higher level
 - `x <= boolean_expression(A,B,C,D);`
- On-line Design Notebook
 - Open a window with editor, or our tool, => cut&paste
- Multiply: successive refinement to see final design
 - 32-bit Adder, 64-bit shift register, 32-bit Multiplicand Register
 - Booth's algorithm to handle signed multiplies
 - There are algorithms that calculate many bits of multiply per cycle (see exercises 4.36 to 4.39 in COD)
- Shifter: Best implemented with technology-specific methodologies
- What's Missing from MIPS is Divide & Floating Point Arithmetic: Next time the Pentium Bug

2/10/99

©UCB Spring 1999

CS152 / Kubiatowicz
Lec6.3

Recap: VHDL combinational example

```
ENTITY nandnor is
  GENERIC (delay: TIME := 1ns);
  PORT (a,b: IN VLBIT; x,y: OUT VLBIT)
END nandnore

ARCHITECTURE behavioral OF nandnor is
BEGIN
  x <= a NOR b AFTER delay;
  y <= a NAND x AFTER delay;
END behavioral;
```

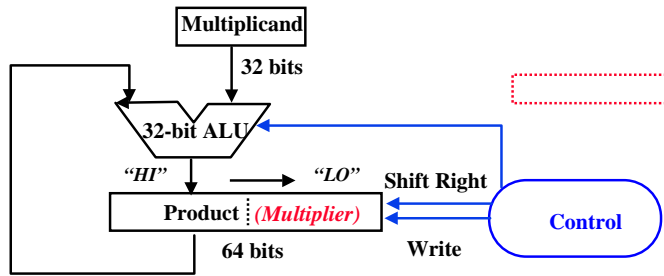
2/10/99

©UCB Spring 1999

CS152 / Kubiatowicz
Lec6.4

Review: MULTIPLY HARDWARE Version 3

- **32-bit** Multiplicand reg, **32-bit** ALU,
64-bit Product reg (**shift right**), (**0-bit** Multiplier reg)



2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.5

Review: Booth's Algorithm Insight

end of run middle of run beginning of run
 $0 \boxed{1} 1 1 \boxed{1} 0$

Current Bit	Bit to the Right	Explanation	Example	Op
1	0	Begins run of 1s	000111 1 00	sub
1	1	Middle of run of 1s	00011 1 1000	none
0	1	End of run of 1s	00 0 1111000	add
0	0	Middle of run of 0s	0 0 01111000	none

Originally for Speed (when shift was faster than add)

- Replace a string of 1s in multiplier with an initial subtract when we first see a one and then later add for the bit after the last one

$$\begin{array}{r}
 -1 \\
 + 10000 \\
 \hline
 01111
 \end{array}$$

2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.6

Radix-4 Modified Booth's Algorithm

Current Bits	Bit to the Right	Explanation	Example	Recode
00	0	Middle of zeros	00 00 00 00 00	0
01	0	Single one	00 00 00 01 00	1
10	0	Begins run of 1s	00 01 11 10 00	-2
11	0	Begins run of 1s	00 01 11 11 00	-1
00	1	Ends run of 1s	00 00 11 11 00	1
01	1	Ends run of 1s	00 01 11 11 00	2
10	1	Isolated 0	00 11 10 11 00	-1
11	1	Middle of run	00 11 11 11 00	0

- Same insight as one-bit Booth's, simply adjust for alignment of 2 bits.
- Allows multiplication 2-bits at a time.


$$\begin{array}{r}
 -1 \\
 + 10000 \\
 \hline
 01111
 \end{array}$$

2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.7

Review: Combinational Shifter from MUXes



- What comes in the MSBs?
- How many levels for 32-bit shifter?
- What if we use 4-1 Muxes ?


2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.8

Funnel Shifter

Instead Extract 32 bits of 64.



- Shift A by i bits (sa= shift right amount)
- Logical: Y = 0, X=A, sa=i
- Arithmetic? Y = Sign, X=A, sa=i
- Rotate? Y = A, X=A, sa=i
- Left shifts? Y = A, X=0, sa=32-i

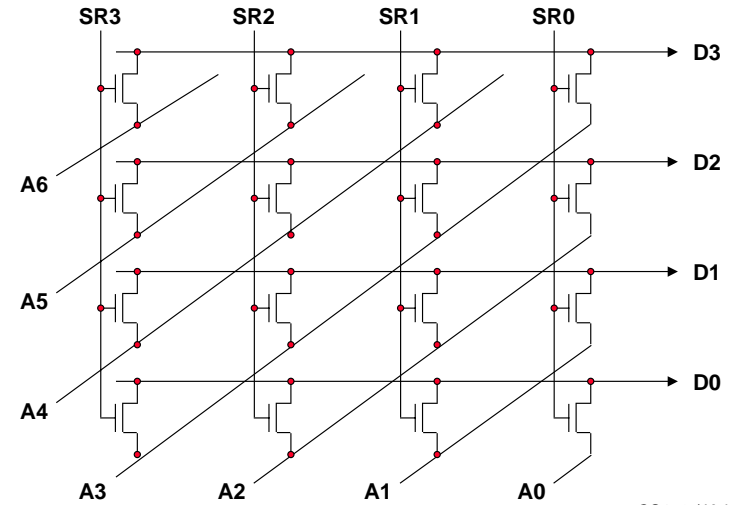
2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.9

Barrel Shifter

Technology-dependent solutions: transistor per switch



2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.10

Divide: Paper & Pencil

```

      1001      Quotient
  Divisor 1000 | 1001010      Dividend
                -1000
                ---
                 10
                 101
                 1010
                 -1000
                 ---
                  10      Remainder (or Modulo result)
    
```

See how big a number can be subtracted, creating quotient bit on each step

Binary => 1 * divisor or 0 * divisor

Dividend = Quotient x Divisor + Remainder
=> | Dividend | = | Quotient | + | Divisor |

3 versions of divide, successive refinement

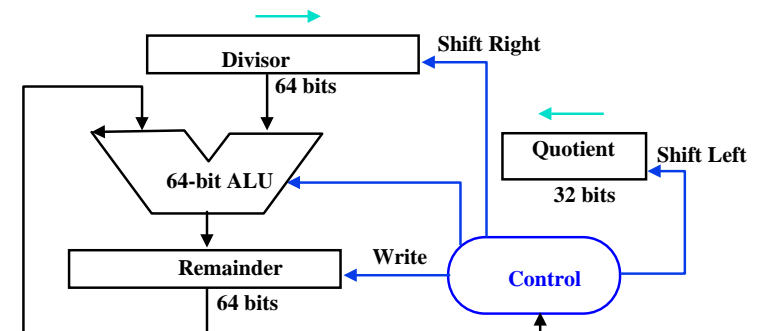
2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.11

DIVIDE HARDWARE Version 1

- 64-bit Divisor reg, 64-bit ALU, 64-bit Remainder reg, 32-bit Quotient reg



2/10/99

©UCB Spring 1999

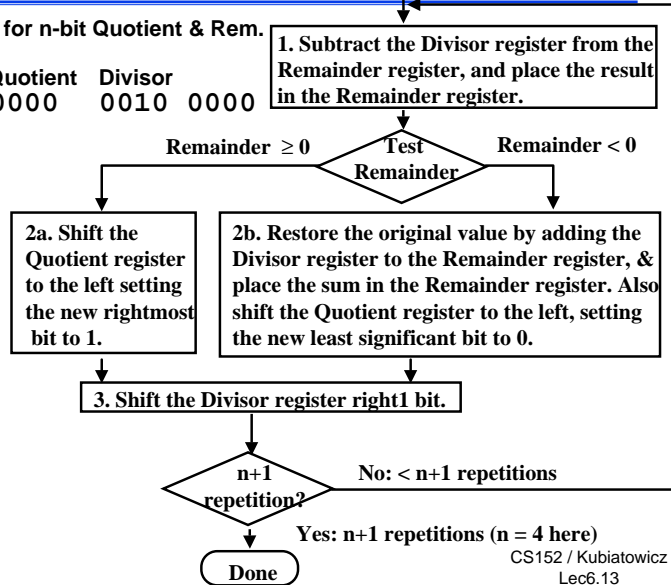
CS152 / Kubiawicz
Lec6.12

Divide Algorithm Version 1

Start: Place Dividend in Remainder

° Takes $n+1$ steps for n -bit Quotient & Rem.

Remainder Quotient Divisor
0000 0111 0000 0010 0000



2/10/99

CS152 / Kubiawicz
Lec6.13

Observations on Divide Version 1

- ° 1/2 bits in divisor always 0
=> 1/2 of 64-bit adder is wasted
=> 1/2 of divisor is wasted
- ° Instead of shifting divisor to right, shift remainder to left?
- ° 1st step cannot produce a 1 in quotient bit (otherwise too big)
=> switch order to shift first and then subtract, can save 1 iteration

2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.14

Divide: Paper & Pencil

```

      01010   Quotient
  Divisor 0001 | 00001010   Dividend
                00001
                -0001
                0000
                 0001
                 -0001
                 0
                  00   Remainder (or Modulo result)
  
```

- Notice that there is no way to get a 1 in leading digit! (this would be an overflow, since quotient would have $n+1$ bits)

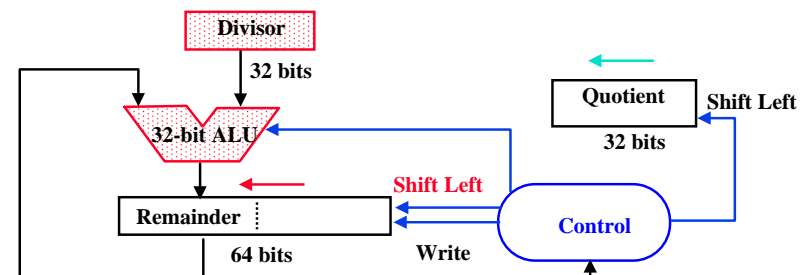
2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.15

DIVIDE HARDWARE Version 2

- ° 32-bit Divisor reg, 32-bit ALU, 64-bit Remainder reg, 32-bit Quotient reg

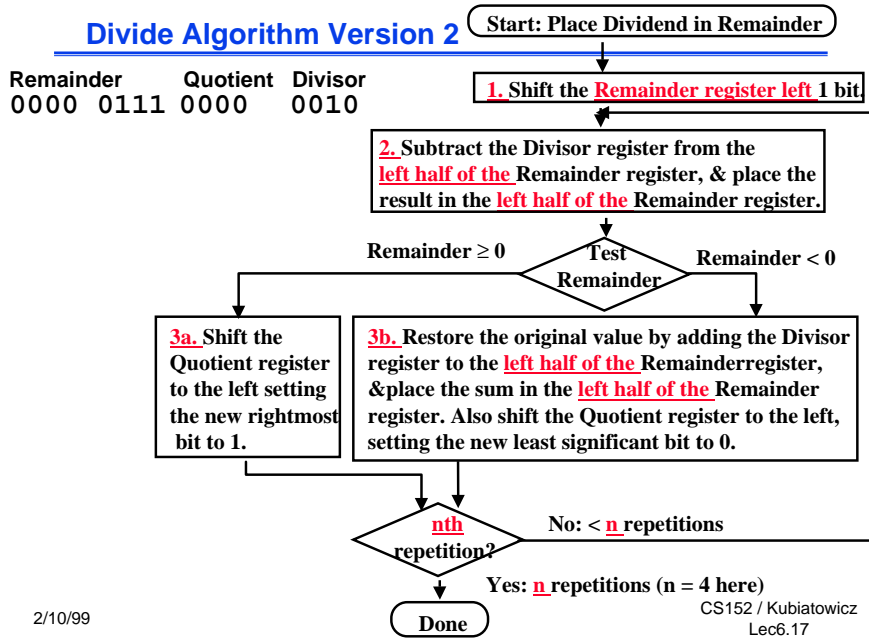


2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.16

Divide Algorithm Version 2



Observations on Divide Version 2

- Eliminate Quotient register by combining with Remainder as shifted left
 - Start by shifting the Remainder left as before.
 - Thereafter loop contains only two steps because the shifting of the Remainder register shifts both the remainder in the left half and the quotient in the right half
 - The consequence of combining the two registers together and the new order of the operations in the loop is that the remainder will be shifted left one time too many.
 - Thus the final correction step must shift back only the remainder in the left half of the register

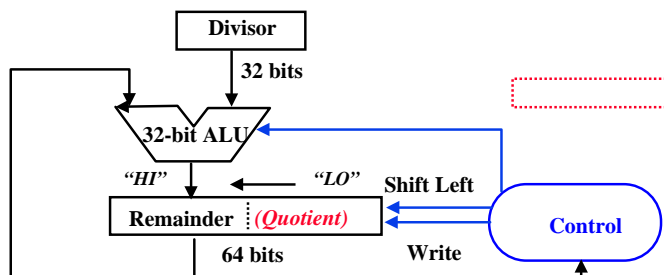
2/10/99

©UCB Spring 1999

CS152 / Kubiawicz Lec6.18

DIVIDE HARDWARE Version 3

- 32-bit Divisor reg, 32-bit ALU, 64-bit Remainder reg, (0-bit Quotient reg)

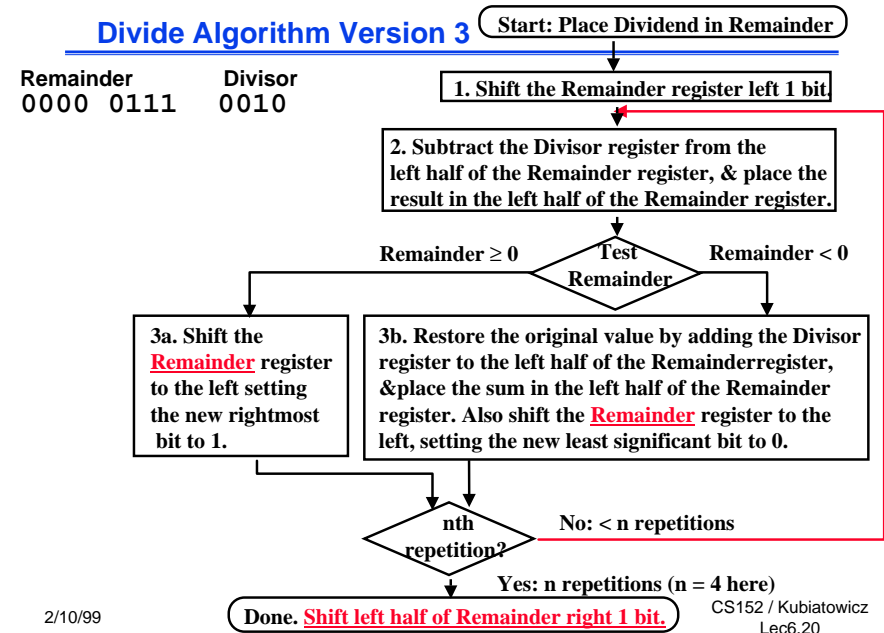


2/10/99

©UCB Spring 1999

CS152 / Kubiawicz Lec6.19

Divide Algorithm Version 3



Observations on Divide Version 3

- Same Hardware as Multiply: just need ALU to add or subtract, and 63-bit register to shift left or shift right
- Hi and Lo registers in MIPS combine to act as 64-bit register for multiply and divide
- Signed Divides: Simplest is to remember signs, make positive, and complement quotient and remainder if necessary
 - Note: Dividend and Remainder must have same sign
 - Note: Quotient negated if Divisor sign & Dividend sign disagree e.g., $-7 \div 2 = -3$, remainder = -1
- Possible for quotient to be too large: if divide 64-bit integer by 1, quotient is 64 bits ("called saturation")

2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.21

Administrative Matters

- Midterm Wednesday 3/3
- Lab 2 due on Friday. How are people doing?
- Next Monday is a Holiday (so, homework/quiz due on Wednesday)
- Finishing Chapter 4 today, moving to Chapter 5 next Wednesday.

2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.22

Floating-Point

- What can be represented in N bits?


◦ Unsigned	0	to	2^N	
◦ 2s Complement	-2^{N-1}	to	$2^{N-1} - 1$	
◦ 1s Complement	$-2^{N-1} + 1$	to	$2^{N-1} - 1$	N-1
◦ Excess M	$-M$	to	$2^M - 1$	
• (E = e + M)				
◦ BCD	0	to	$10^{N/4} - 1$	
- But, what about?
 - very large numbers? 9,349,398,989,787,762,244,859,087,678
 - very small number? 0.00000000000000000000000045691
 - rationals $\frac{2}{3}$
 - irrationals $\sqrt{2}$
 - transcendentals e,

2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.23

Recall Scientific Notation



IEEE F.P. $\pm 1.M \times 2^{e-127}$

- Issues:
 - Arithmetic (+, -, *, /)
 - Representation, Normal form
 - Range and Precision
 - Rounding
 - Exceptions (e.g., divide by zero, overflow, underflow)
 - Errors
 - Properties (negation, inversion, if $A \neq B$ then $A - B \neq 0$)


2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.24

Review from Prerequisites: Floating-Point Arithmetic

Representation of floating point numbers in IEEE 754 standard:



exponent: excess 127 binary integer
 mantissa: sign + magnitude, normalized binary significand w/ hidden integer bit: 1.M

actual exponent is $e = E - 127$

$$N = (-1)^S 2^{E-127} (1.M) \quad 0 < E < 255$$

$$0 = 0 \text{ 00000000 } 0 \dots 0 \quad -1.5 = 1 \text{ 01111111 } 10 \dots 0$$

Magnitude of numbers that can be represented is in the range:

$$2^{-126} (1.0) \quad \text{to} \quad 2^{127} (2 - 2^{-23})$$

which is approximately:

$$1.8 \times 10^{-38} \quad \text{to} \quad 3.40 \times 10^{38}$$

(integer comparison valid on IEEE Fl.Pt. numbers of same sign!)

2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.25

Basic Addition Algorithm

For addition (or subtraction) this translates into the following steps:

- (1) compute $Y_e - X_e$ (*getting ready to align binary point*)
- (2) right shift X_m that many positions to form $X_m 2^{X_e - Y_e}$
- (3) compute $X_m 2^{X_e - Y_e} + Y_m$

if representation demands normalization, then normalization step follows:

- (4) left shift result, decrement result exponent (e.g., 0.001xx...)
 right shift result, increment result exponent (e.g., 101.1xx...)
 continue until MSB of data is 1 (NOTE: Hidden bit in IEEE Standard)
- (5) doubly biased exponent must be corrected:

$X_e = 7$	$X_e = 1111$	$= 15$	$= 7 + 8$
$Y_e = -3$	$Y_e = 0101$	$= 5$	$= -3 + 8$
Excess 8	10100	20	4 + 8 + 8

extra subtraction step of the bias amount

- (6) if result is 0 mantissa, may need to zero exponent by special step

2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.26

Extra Bits for rounding

"Floating Point numbers are like piles of sand; every time you move one you lose a little sand, but you pick up a little dirt."

How many extra bits?

IEEE: As if computed the result exactly and rounded.

Addition:

1.xxxxx	1.xxxxx	1.xxxxx
+ 1.xxxxx	0.001xxxxx	0.01xxxxx
1x.xxxxxy	1.xxxxxyyy	1x.xxxxxyy

post-normalization pre-normalization pre and post

◦ **Guard Digits:** digits to the right of the first p digits of significand to guard against loss of digits – can later be shifted left into first P places during normalization.

◦ Addition: carry-out shifted in

◦ Subtraction: borrow digit and guard

◦ Multiplication: carry and guard, Division requires guard

2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.27

Rounding Digits

normalized result, but some non-zero digits to the right of the significand --> the number should be rounded

E.g., $B = 10$, $p = 3$:

	0	2	1.69	= 1.6900 * 10 ^{2-bias}
-	0	0	7.85	= - .0785 * 10 ^{2-bias}
	0	2	1.61	= 1.6115 * 10 ^{2-bias}

one round digit must be carried to the right of the guard digit so that after a normalizing left shift, the result can be rounded, according to the value of the round digit

IEEE Standard:

four rounding modes: round to nearest (default)
 round towards plus infinity
 round towards minus infinity
 round towards 0

round to nearest:

round digit < B/2 then truncate
 > B/2 then round up (add 1 to ULP: unit in last place)
 = B/2 then round to nearest even digit

it can be shown that this strategy minimizes the mean error introduced by rounding

2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.28

Sticky Bit

Additional bit to the right of the round digit to better fine tune rounding

$d_0 . d_1 d_2 d_3 \dots d_{p-1} 0 0 0$
 $+ \underline{0 . 0 0 X \dots X \quad X X S}$

Sticky bit: set to 1 if any 1 bits fall off the end of the round digit

$d_0 . d_1 d_2 d_3 \dots d_{p-1} 0 0 0$
 $- \underline{0 . 0 0 X \dots X \quad X X 0}$

$d_0 . d_1 d_2 d_3 \dots d_{p-1} 0 0 0$
 $- \underline{0 . 0 0 X \dots X \quad X X 1}$

generates a borrow

Rounding Summary:

Radix 2 minimizes wobble in precision

Normal operations in +, -, *, / require one carry/borrow bit + one guard digit

One round digit needed for correct rounding

Sticky bit needed when round digit is B/2 for max accuracy


Rounding to nearest has mean error = 0 if uniform distribution of digits are assumed

2/10/99

©UCB Spring 1999


CS152 / Kubiawicz
Lec6.29

Denormalized Numbers



The gap between 0 and the next representable number is much larger than the gaps between nearby representable numbers.

IEEE standard uses denormalized numbers to fill in the gap, making the distances between numbers near 0 more alike.



same spacing, half as many values!

NOTE: PDP-11, VAX cannot represent subnormal numbers. These machines underflow to zero instead.

2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.30

Infinity and NaNs

result of operation *overflows*, i.e., is larger than the largest number that can be represented

overflow is not the same as divide by zero (raises a different exception)

+/- infinity $\boxed{S \ 1 \dots 1 \ 0 \dots 0}$

It may make sense to do further computations with infinity
e.g., $X/0 > Y$ may be a valid comparison

Not a number, but not infinity (e.g. $\sqrt{-4}$)
invalid operation exception (unless operation is = or \neq)

NaN $\boxed{S \ 1 \dots 1 \ \text{non-zero}}$
HW decides what goes here

NaNs propagate: $f(\text{NaN}) = \text{NaN}$

2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.31

Radix-4 Modified Booth's \Rightarrow Multiple representations

Once admit new symbols (i.e. $\bar{1}$), can have multiple representations of a number:

Current Bits	Bit to the Right	Explanation	Example	Recode
00	0	Middle of zeros	00 00 00 <u>00</u> 00	00 (0)
01	0	Single one	00 00 00 <u>01</u> 00	01 (1)
10	0	Begins run of 1s	00 01 11 <u>10</u> 00	$\bar{1}0$ (-2)
11	0	Begins run of 1s	00 01 11 <u>11</u> 00	$0\bar{1}$ (-1)
00	1	Ends run of 1s	00 <u>00</u> 11 11 00	01 (1)
01	1	Ends run of 1s	00 <u>01</u> 11 11 00	10 (2)
10	1	Isolated 0	00 11 <u>10</u> 11 00	$0\bar{1}$ (-1)
11	1	Middle of run	00 11 <u>11</u> 11 00	00 (0)

2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.32

Pentium Bug

- Pentium FP Divider uses algorithm to generate multiple bits per steps
 - FPU uses most significant bits of divisor & dividend/remainder to guess next 2 bits of quotient
 - Guess is taken from lookup table: -2, -1,0,+1,+2 (if previous guess too large a remainder, quotient is adjusted in subsequent pass of -2)
 - Guess is multiplied by divisor and subtracted from remainder to generate a new remainder
 - Called SRT division after 3 people who came up with idea
- Pentium table uses 7 bits of remainder + 4 bits of divisor = 2^{11} entries
- 5 entries of divisors omitted: 1.0001, 1.0100, 1.0111, 1.1010, 1.1101 from PLA (fix is just add 5 entries back into PLA: cost \$200,000)
- Self correcting nature of SRT => string of 1s must follow error
 - e.g., 1011 1111 1111 1111 1111 1011 1000 0010 0011 0111 1011 0100 (2.99999892918)
- Since indexed also by divisor/remainder bits, sometimes bug doesn't show even with dangerous divisor value

2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.33

Pentium bug appearance

- First 11 bits to right of decimal point always correct: bits 12 to 52 where bug can occur (4th to 15th decimal digits)
- FP divisors near integers 3, 9, 15, 21, 27 are dangerous ones:
 - $3.0 > d \geq 3.0 - 36 \times 2^{-22}$, $9.0 > d \geq 9.0 - 36 \times 2^{-20}$
 - $15.0 > d \geq 15.0 - 36 \times 2^{-20}$, $21.0 > d \geq 21.0 - 36 \times 2^{-19}$
- 0.333333×9 could be problem
- In Microsoft Excel, try $(4,195,835 / 3,145,727) * 3,145,727$
 - = 4,195,835 => not a Pentium with bug
 - = 4,195,579 => Pentium with bug (assuming Excel doesn't already have SW bug patch)
 - Rarely noticed since error in 5th significant digit
 - Success of IEEE standard made discovery possible:
 - all computers should get same answer

2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.34

Pentium Bug Time line

- June 1994: Intel discovers bug in Pentium: takes months to make change, reverify, put into production: plans good chips in January 1995 4 to 5 million Pentiums produced with bug
- Scientist suspects errors and posts on Internet in September 1994
- Nov. 22 Intel Press release: "Can make errors in 9th digit ... Most engineers and financial analysts need only 4 of 5 digits. Theoretical mathematician should be concerned. ... So far only heard from one."
- Intel claims happens once in 27,000 years for typical spread sheet user:
 - 1000 divides/day x error rate assuming numbers random
- Dec 12: IBM claims happens once per 24 days: Bans Pentium sales
 - 5000 divides/second x 15 minutes = **4,200,000** divides/day
 - IBM statement: <http://www.ibm.com/Features/pentium.html>
 - Intel said it regards IBM's decision to halt shipments of its Pentium processor-based systems as unwarranted.

2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.35

Pentium jokes

- Q: What's another name for the "Intel Inside" sticker they put on Pentiums?
A: Warning label.
- Q: Have you heard the new name Intel has chosen for the Pentium?
A: the Intel Inacura.
- Q: According to Intel, the Pentium conforms to the IEEE standards for floating point arithmetic. If you fly in aircraft designed using a Pentium, what is the correct pronunciation of "IEEE"?
A: Aaaaaaaiiiiiiiiiieeeeeeeeeee!
- TWO OF TOP TEN NEW INTEL SLOGANS FOR THE PENTIUM
9.9999973251 It's a FLAW, Dammit, not a Bug
7.9999414610 Nearly 300 Correct Opcodes

2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.36

Pentium conclusion: Dec. 21, 1994 \$500M write-off

“To owners of Pentium processor-based computers and the PC community:

We at Intel wish to sincerely apologize for our handling of the recently publicized Pentium processor flaw.

The Intel Inside symbol means that your computer has a microprocessor second to none in quality and performance. Thousands of Intel employees work very hard to ensure that this is true. But no microprocessor is ever perfect.

What Intel continues to believe is technically an extremely minor problem has taken on a life of its own. Although Intel firmly stands behind the quality of the current version of the Pentium processor, we recognize that many users have concerns.

We want to resolve these concerns.

Intel will exchange the current version of the Pentium processor for an updated version, in which this floating-point divide flaw is corrected, for any owner who requests it, free of charge anytime during the life of their computer. Just call 1-800-628-8686.”

Sincerely,
Andrew S. Grove
President /CEO

Craig R. Barrett
Executive Vice President
&COO

Gordon E. Moore
Chairman of the Board

CS152 / Kubiawicz
Lec6.37

2/10/99

Summary

- Pentium: Difference between bugs that board designers must know about and bugs that potentially affect all users
 - Why not make public complete description of bugs in later category?
 - \$200,000 cost in June to repair design
 - \$500,000,000 loss in December in profits to replace bad parts
 - How much to repair Intel's reputation?
- What is technologists responsibility in disclosing bugs?
- Bits have no inherent meaning: operations determine whether they are really ASCII characters, integers, floating point numbers
- Divide can use same hardware as multiply: Hi & Lo registers in MIPS
- Floating point basically follows paper and pencil method of scientific notation using integer algorithms for multiply and divide of significands
- IEEE 754 requires good rounding; special values for NaN, Infinity

2/10/99

©UCB Spring 1999

CS152 / Kubiawicz
Lec6.38