Computer Architecture and Engineering
Lecture 8: Divide, Floating Point, Pentium Bug

February 12, 1995

Dave Patterson (patterson@cs) and
Shing Kong (shing.kong@eng.sun.com)

Slides available on http://http.cs.berkeley.edu/~patterson

Review: MULTIPLY HARDWARE Version 1

- 64-bit Multiplicand reg, 64-bit ALU, 64-bit Product reg, 32-bit multiplier reg

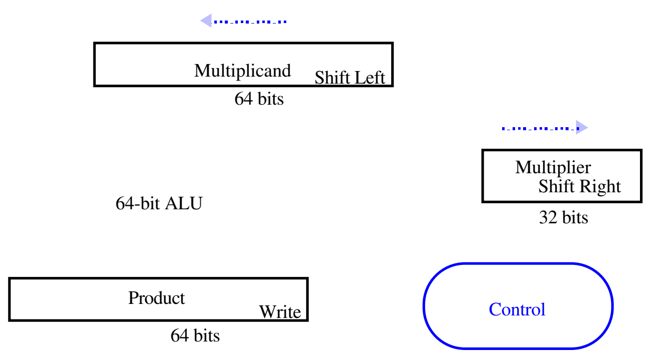

Review: MULTIPLY HARDWARE Version 3

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

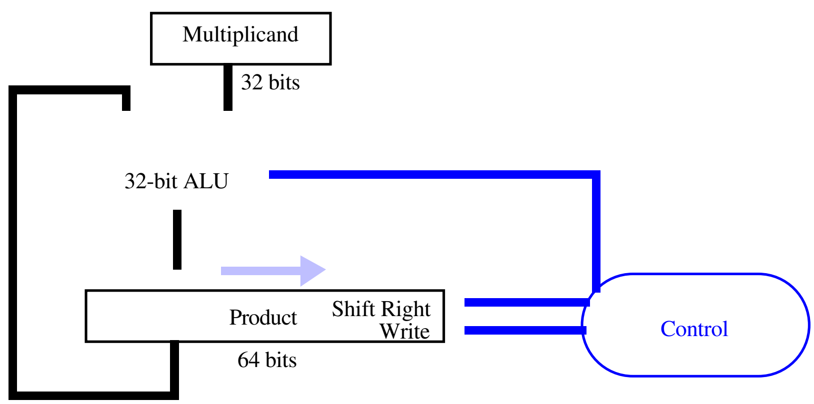

Review: Booth's Algorithm Insight

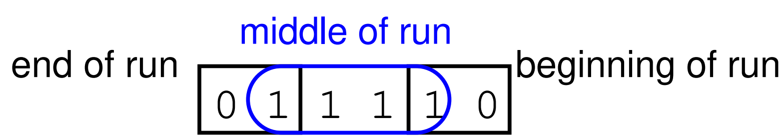

Current Bit	Bit to the Right	Explanation	Example
1	0	Beginning of a run of 1s	0001111000
1	1	Middle of a run of 1s	0001111000
0	1	End of a run of 1s	0001111000
0	0	Middle of a run of 0s	0001111000

Originally for Speed since shift faster than add for his machine

Review: Booth's Algorithm

1. Depending on the current and previous bits, do one of the following:
   - 00: a. Middle of a string of 0s, so no arithmetic operations.
   - 01: b. End of a string of 1s, so add the multiplicand to the left half of the product.
   - 10: c. Beginning of a string of 1s, so subtract the multiplicand from the left half of the product.
   - 11: d. Middle of a string of 1s, so no arithmetic operation.

2. As in the previous algorithm, shift the Product register right (arith) 1 bit.

Multiplicand	Product (2 x 3)	Multiplicand	Product (2 x -3)
`0010`	`0000 0011 0`	`0010`	`0000 1101 0`

Recap of Last Lecture

- Instruction Set drives the ALU design
- Shifter: success refinement from 1/bit at a time shift register to barrel shifter
- 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)

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.)
- Hands On Divide & FP (30 min.)
- Questions and Break (5 min.)
- Pentium Bug (25 min.)

Divide: Paper & Pencil

```
                 1001     Quotient
              ----------
Divisor 1000 |  1001010     Dividend
              -1000
                 10
                 101
                 1010
                -1000
                   10     Remainder
```

- 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
- 3 versions of divide, successive refinement

DIVIDE HARDWARE Version 1

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

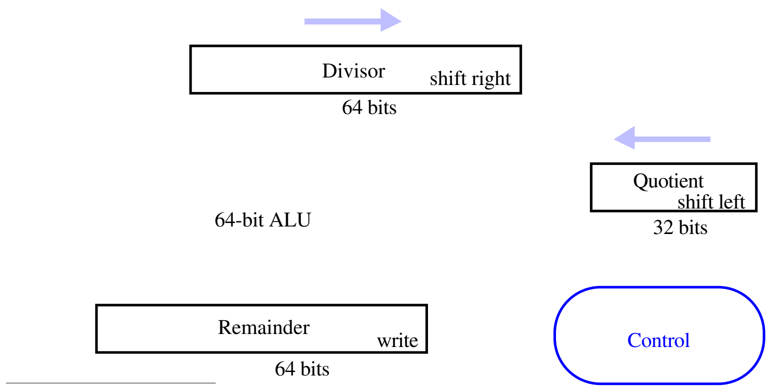

Divide Algorithm Version 1

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

Quotient Divisor Remainder

0000 0010 0000 0000 0111

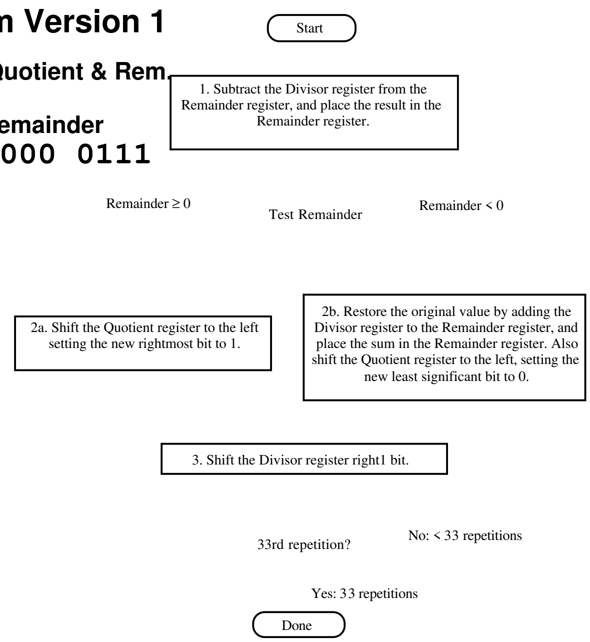

Observations on Divide Version 1

- 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

DIVIDE HARDWARE Version 2

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

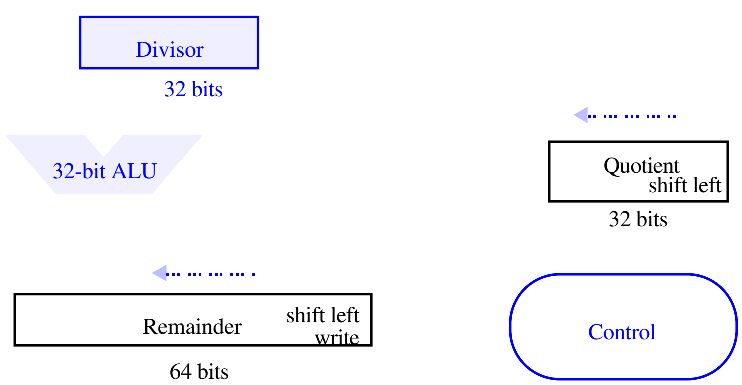

Divide Algorithm Version 2

Quotient	Divisor	Remainder
0000	0010	0000 0111

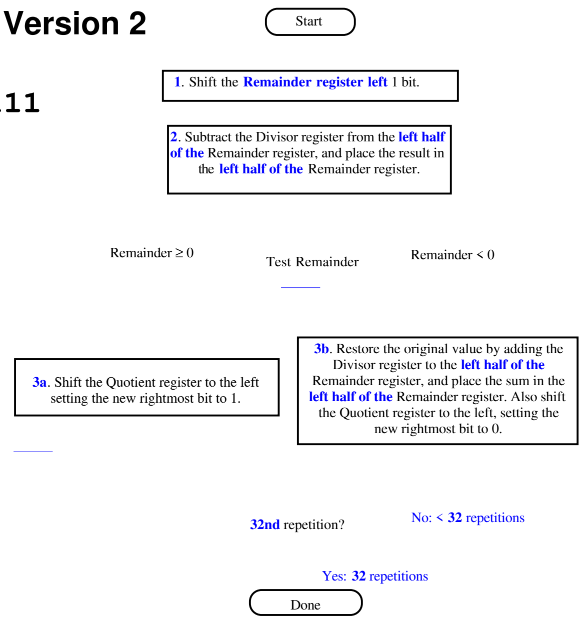

Administrative Matters

° Video tapes of lectures available for viewing in 205 McLaughlin,
  Mon. to Fri. 9 AM to 5 PM; Wed & Fri. 6 PM to 10 PM

° 1st Midterm 2 weeks, Feb. 22; 5 PM to 8 PM for 1.5 hour test in Sibley Auditorium; (on hnadout 1st day if class)
  2 hand drawn pieces of paper
  (will be able to bring to second midterm too)

° LaVal's Afterwards to meet students, TAs, profs

° Other topics?

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 shifted left one time too many.**
- **Thus the final correction step must shift back only the remainder in the left half of the register**

DIVIDE HARDWARE Version 3

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

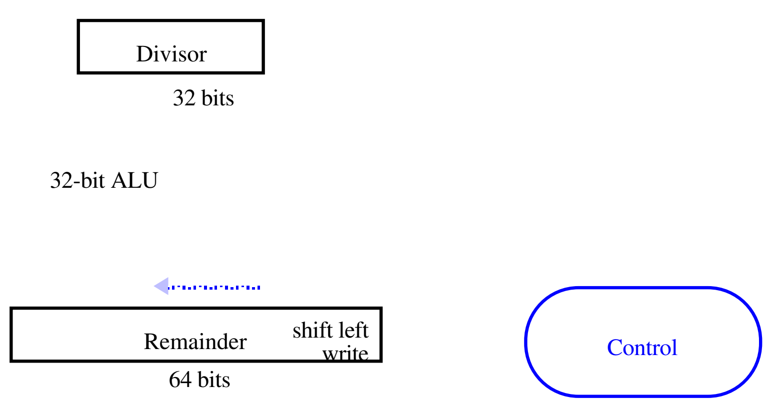

Divide Algorithm Version 3

Divisor Remainder

0010 0000 0111

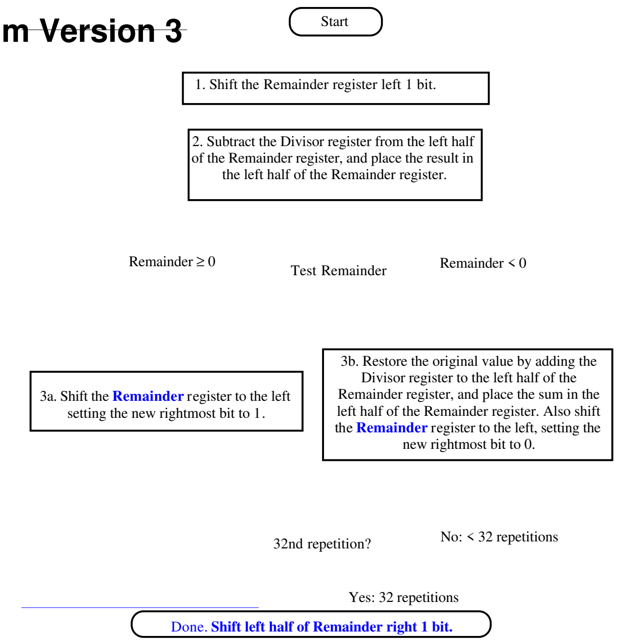

Observations on Divide Version 3

- 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

Recall Scientific Notation

decimal point

exponent — Sign, magnitude

6.02×10^{23} $\qquad$ 1.673×10^{-24}

Mantissa — Sign, magnitude $\qquad$ *radix (base)*

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

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 $A - B \neq 0$)

Floating-Point Arithmetic

Representation of floating point numbers in IEEE 754 standard:

single precision

	1	8	23
sign	S	E	M

exponent: excess 127 binary integer

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

actual exponent is

$e = E - 127 = 10000001_2 - 127$

$0 < E < 255$

$N = (-1)^S \, 2^{E-127} (1.M)$

$0 = 0\ 00000000\ 0 \ldots 0$ $\qquad$ $-1.5 = 1\ 01111111\ 10 \ldots 0$

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 fp numbers!)

What numbers are represented?

rationals, p digits of precision on radix B, exponent in range

–bias to 2^l - 1 - bias

(this assumes significand is left justified, i.e., normalized)

the standard reserves the smallest and largest exponent field values

(0 and 2^l - 1) for special purposes

Layout on the number line:
ex: B = 3, p = 2

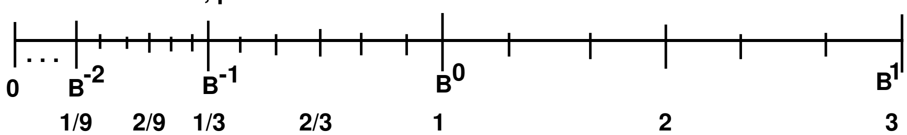

Notice that each interval $[B^i, B^{i+1})$ increases in size over the prior interval $[B^{i-1}, B^i)$, yet the number of values represented in the interval remains the same: $(B - 1)\, B^{p-1}$ - the number of non-zero, left adjusted significands for each value of the exponent. So the density of represented values decreases by a factor of B for each successive interval.

Normalized Numbers

Significand is left adjusted --> as large as possible,
exponent is as small as possible

E.g., $B = 2^4$, p = 3:

0	0110	0000 . 0110 1100	= 0.6C

denormalized

0	0101	0110 . 1100 0000	= 6.C0

normalized

In B = 2, the significand MSB is always 1 when the significand is left adjusted. So not necessary to store this "hidden" bit in memory.

0	011	1.01	1

w/o hidden bit

0	011	.011

w/ hidden bit = improved precision

Within the FPU, the hidden bit is inserted because of the denormalization step that precedes FP add/subtract.

2^{-bias}

Smallest normal #: 0 0 ... 0 "1" .0 ... 01 hidden bit *must distinguish from zero!*

0 0 ... 0 1.0 ... 00 no hidden bit

Basic Addition Algorithm

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

(1) compute Ye - Xe

(2) right shift Xm that many positions to form Xm 2^{Xe-Ye}

(3) compute Xm 2^{Xe-Ye} + Ym

if representation demands normalization, then a normalization step follows:

(4) left shift result, decrement result exponent
right shift result, increment result exponent
continue until MSB of data is 1 (NOTE: Hidden bit in IEEE Standard)

(5) doubly biased exponent must be corrected:

Xe = 7	Xe = 1111	= 15	= 7 + 8
Ye = -3	Ye = <u>0101</u>	= <u>5</u>	= <u>-3 + 8</u>
Excess 8	10100	20	4 + 8 + 8

extra subtraction step of the bias amount

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

Datapath for FPU

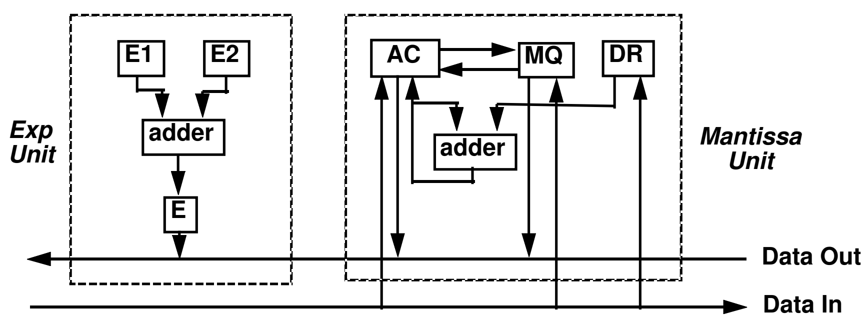

Addition Algorithm:

$AC<n_m-1:0>$, $DR<n_m-1:0>$, $E1<n_m-1:0>$, $E2<n_m-1:0>$, $E<n_m-1:0>$,

AC_OVERFLOW, ERROR

Begin: AC_OVERFLOW := 0; ERROR := 0;
Load: E1:= Xe; AC := Xm;
E2 := Ye; DR := Ym;

Extra Bits

"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.xxxxy	1.xxxxxyyy	1x.xxxxyyy
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

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

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)
= B/2 then round to nearest even digit

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

Sticky Bit

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

```
  d0 . d1 d2 d3 . . . dp-1  0  0  0   <-- Sticky bit: set to 1 if any 1 bits fall off
+  0 .  0  0  X . . .  X    X  X  S       the end of the round digit
                            X  X  S

  d0 . d1 d2 d3 . . . dp-1  0  0  0         d0 . d1 d2 d3 . . . dp-1  0  0  0
-  0 .  0  0  X . . .  X    X  X  0       -  0 .  0  0  X . . .  X    X  X  1
                            X  X  0                     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

Denormalized Numbers

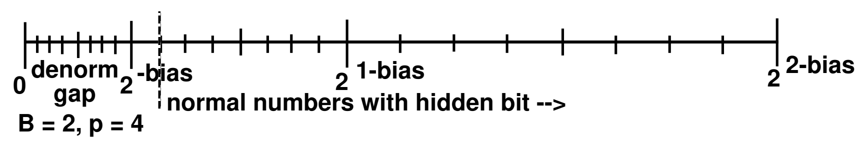

The gap between 0 and the next representable number is much larger than the gaps between nearby representable numbers.

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

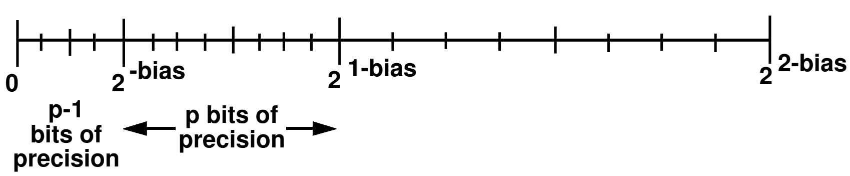

same spacing, half as many values!

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

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 | S | 1 . . . 1 | 0 . . . 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.q. sqrt(-4))
invalid operation exception (unless operation is = or ≠)

NaN | S | 1 . . . 1 | non-zero | HW decides what goes here

NaNs propagate: f(NaN) = NaN

Exceptions

Invalid operation:
result of operation is a NaN (except = or ≠)
inf. +/- inf.; 0 * inf; 0/0; inf./inf.; x remainder y where y = 0;
sqrt(x) where x < 0, x ≠ +/- inf.

Overflow:
result of operation is larger than largest representable #
flushed to +/- inf. if overflow exception is not enabled

Divide by 0:
x/0 where x = 0, +/- inf.;
flushed to +/- inf. if divide by zero exception not enabled

Underflow:
subnormal result OR non-zero result underflows to 0

Inexact:
rounded result not the actual result (rounding error ≠ 0)

IEEE Standard --> specifies defaults and allows traps to permit user to handle the exception

contrast with the more usual result of aborting the computation altogether!

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 reminder, 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

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 x 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)
  - Rare since error in 5th significant digit

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 = 4.2 million 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.

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: Aaaaaaaiiiiiiiiieeeeeeeeeeeee!
- 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

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	Craig R. Barrett	Gordon E. Moore
President /CEO	Executive Vice President &COO	Chairman of the Board

Summary

- 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
- 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?