COMPUTER ORGANIZATION AND DESIGN The Hardware/Software Interface

## Chapter 3

## Arithmetic for Computers

## Arithmetic for Computers

Operations on integers

- Addition and subtraction
- Multiplication and division
- Dealing with overflow

Floating-point real numbers

- Representation and operations


## Integer Addition

## Example: $7+6$



Overflow if result out of range

- Adding +ve and -ve operands, no overflow
- Adding two +ve operands
- Overflow if result sign is 1
- Adding two -ve operands
- Overflow if result sign is 0


## Integer Subtraction

Add negation of second operand
Example: $7-6=7+(-6)$

| +7: | $00000000 \ldots 00000111$ |
| :--- | :--- | :--- |
| -6: | $11111111 \ldots 11111010$ |
| $+1:$ | $00000000 \ldots 00000001$ |

Overflow if result out of range

- Subtracting two +ve or two -ve operands, no overflow
- Subtracting +ve from -ve operand
- Overflow if result sign is 0
- Subtracting -ve from +ve operand
- Overflow if result sign is 1


## Dealing with Overflow

Some languages (e.g., C) ignore overflow

- Use MIPS addu, addui, subu instructions

Other languages (e.g., Ada, Fortran) require raising an exception

- Use MIPS add, addi, sub instructions
- On overflow, invoke exception handler Save PC in exception program counter (EPC) register
Jump to predefined handler address mfc 0 (move from coprocessor reg) instruction can retrieve EPC value, to return after corrective action


## Arithmetic for Multimedia

Graphics and media processing operates on vectors of 8 -bit and 16-bit data

- Use 64-bit adder, with partitioned carry chain Operate on $8 \times 8$-bit, $4 \times 16$-bit, or $2 \times 32$-bit vectors
- SIMD (single-instruction, multiple-data)

Saturating operations

- On overflow, result is largest representable value
c.f. 2 s -complement modulo arithmetic
- E.g., clipping in audio, saturation in video


## Multiplication

Start with long-multiplication approach


Length of product is the sum of operand lengths


Chapter 3 - Arithmetic for Computers - 7

## Multiplication Hardware



Chapter 3 - Arithmetic for Computers - 8

## Optimized Multiplier

Perform steps in parallel: add/shift


## One cycle per partial-product addition

- That's ok, if frequency of multiplications is low


## Faster Multiplier

## Uses multiple adders

- Cost/performance tradeoff


Can be pipelined

- Several multiplication performed in parallel


## MIPS Multiplication

Two 32-bit registers for product

- HI: most-significant 32 bits
- LO: least-significant 32-bits

Instructions

- mult rs, rt / multu rs, rt
- 64-bit product in HI/LO
- mfhi rd / mflo rd
- Move from HI/LO to rd
- Can test HI value to see if product overflows 32 bits
- mul rd, rs, rt
- Least-significant 32 bits of product -> rd


## Division



## Check for 0 divisor

Long division approach

- If divisor $\leq$ dividend bits

1 bit in quotient, subtract

- Otherwise
- 0 bit in quotient, bring down next dividend bit
Restoring division
- Do the subtract, and if remainder goes < 0, add divisor back
Signed division
- Divide using absolute values
- Adjust sign of quotient and remainder as required


## Division Hardware



## Optimized Divider



One cycle per partial-remainder subtraction Looks a lot like a multiplier!

- Same hardware can be used for both


## Faster Division

Can't use parallel hardware as in multiplier

- Subtraction is conditional on sign of remainder

Faster dividers (e.g. SRT devision) generate multiple quotient bits per step

- Still require multiple steps


## MIPS Division

Use HI/LO registers for result

- HI: 32-bit remainder
- LO: 32-bit quotient

Instructions

- div rs, rt / divu rs, rt
- No overflow or divide-by-0 checking Software must perform checks if required
- Use mfhi, mflo to access result


## Floating Point

Representation for non-integral numbers

- Including very small and very large numbers

Like scientific notation
$=-2.34 \times 10^{56}$
normalized

- $+0.002 \times 10^{-4}$
$-+987.02 \times 10^{9}$
In binary
- $\pm 1 . x x x x x x x_{2} \times 2 y y y$

Types float and double in C

## Floating Point Standard

Defined by IEEE Std 754-1985
Developed in response to divergence of representations

- Portability issues for scientific code

Now almost universally adopted
Two representations

- Single precision (32-bit)
- Double precision (64-bit)


## IEEE Floating-Point Format

single: 8 bits double: 11 bits

| $S$ | Exponent | Fraction |
| :--- | :--- | :--- |

$$
x=(-1)^{S} \times(1+\text { Fraction }) \times 2^{(\text {Exponent-Bias })}
$$

S: sign bit ( $0 \Rightarrow$ non-negative, $1 \Rightarrow$ negative) Normalize significand: $1.0 \leq \mid$ significand $\mid<2.0$

- Always has a leading pre-binary-point 1 bit, so no need to represent it explicitly (hidden bit)
- Significand is Fraction with the "1." restored

Exponent: excess representation: actual exponent + Bias

- Ensures exponent is unsigned
- Single: Bias = 127; Double: Bias = 1203


## Single-Precision Range

Exponents 00000000 and 11111111 reserved Smallest value

- Exponent: 00000001
$\Rightarrow$ actual exponent $=1-127=-126$
- Fraction: $000 . . .00 \Rightarrow$ significand $=1.0$
- $\pm 1.0 \times 2^{-126} \approx \pm 1.2 \times 10^{-38}$

Largest value

- exponent: 11111110
$\Rightarrow$ actual exponent $=254-127=+127$
- Fraction: $111 . .11 \Rightarrow$ significand $\approx 2.0$
$\pm \pm 2.0 \times 2^{+127} \approx \pm 3.4 \times 10^{+38}$


## Double-Precision Range

Exponents 0000... 00 and $1111 \ldots 11$ reserved Smallest value

- Exponent: 00000000001
$\Rightarrow$ actual exponent $=1-1023=-1022$
- Fraction: $000 \ldots 00 \Rightarrow$ significand $=1.0$
$\pm \pm 1.0 \times 2^{-1022} \approx \pm 2.2 \times 10^{-308}$
Largest value
- Exponent: 11111111110
$\Rightarrow$ actual exponent $=2046-1023=+1023$
- Fraction: 111... $11 \Rightarrow$ significand $\approx 2.0$
$- \pm 2.0 \times 2^{+1023} \approx \pm 1.8 \times 10^{+308}$


## Floating-Point Precision

Relative precision

- all fraction bits are significant
- Single: approx $2^{-23}$

Equivalent to $23 \times \log _{10} 2 \approx 23 \times 0.3 \approx 6$ decimal digits of precision

- Double: approx $2^{-52}$

Equivalent to $52 \times \log _{10} 2 \approx 52 \times 0.3 \approx 16$ decimal digits of precision

## Floating-Point Example

Represent -0.75

- $-0.75=(-1)^{1} \times 1.1_{2} \times 2^{-1}$
- $S=1$
- Fraction = 1000 $\ldots 0_{2}$
- Exponent $=-1+$ Bias
- Single: $-1+127=126=01111110_{2}$

Double: $-1+1023=1022=01111111110_{2}$
Single: 1011111101000... 00
Double: 1011111111101000... 00

## Floating-Point Example

What number is represented by the singleprecision float
11000000101000... 00

- $\mathrm{S}=1$
- Fraction $=01000 \ldots 00_{2}$
- Fxponent $=10000001_{2}=129$

$$
\begin{aligned}
\mathrm{X} & =(-1)^{1} \times\left(1+01_{2}\right) \times 2^{(129-127)} \\
& =(-1) \times 1.25 \times 2^{2} \\
& =-5.0
\end{aligned}
$$

## Floating-Point Addition

Consider a 4-digit decimal example

- $9.999 \times 10^{1}+1.610 \times 10^{-1}$

1. Align decimal points

- Shift number with smaller exponent
- $9.999 \times 10^{1}+0.016 \times 10^{1}$

2. Add significands

- $9.999 \times 10^{1}+0.016 \times 10^{1}=10.015 \times 10^{1}$

3. Normalize result \& check for over/underflow

- $1.0015 \times 10^{2}$

4. Round and renormalize if necessary

- $1.002 \times 10^{2}$


## Floating-Point Addition

Now consider a 4-digit binary example
$=1.000_{2} \times 2^{-1}+-1.110_{2} \times 2^{-2}(0.5+-0.4375)$

1. Align binary points

- Shift number with smaller exponent
- $1.000_{2} \times 2^{-1}+-0.111_{2} \times 2^{-1}$

2. Add significands

- $1.000_{2} \times 2^{-1}+-0.111_{2} \times 2^{-1}=0.001_{2} \times 2^{-1}$

3. Normalize result \& check for over/underflow

- $1.000_{2} \times 2^{-4}$, with no over/underflow

4. Round and renormalize if necessary

- $1.000_{2} \times 2^{-4}$ (no change) $=0.0625$


## FP Adder Hardware

Much more complex than integer adder Doing it in one clock cycle would take too long

- Much longer than integer operations
- Slower clock would penalize all instructions FP adder usually takes several cycles
- Can be pipelined


## FP Adder Hardware



## Floating-Point Multiplication

Consider a 4-digit decimal example

- $1.110 \times 10^{10} \times 9.200 \times 10^{-5}$

1. Add exponents

- For biased exponents, subtract bias from sum
- New exponent $=10+-5=5$

2. Multiply significands

- $1.110 \times 9.200=10.212 \Rightarrow 10.212 \times 10^{5}$

3. Normalize result \& check for over/underflow

- $1.0212 \times 10^{6}$

4. Round and renormalize if necessary

- $1.021 \times 10^{6}$

5. Determine sign of result from signs of operands

- $+1.021 \times 10^{6}$


## Floating-Point Multiplication

Now consider a 4-digit binary example

- $1.000_{2} \times 2^{-1} \times-1.110_{2} \times 2^{-2}(0.5 \times-0.4375)$

1. Add exponents

- Unbiased: $-1+-2=-3$
- Biased: $(-1+127)+(-2+127)=-3+254-127=-3+127$

2. Multiply significands

- $1.000_{2} \times 1.110_{2}=1.1102 \Rightarrow 1.110_{2} \times 2^{-3}$

3. Normalize result \& check for over/underflow

- $1.110_{2} \times 2^{-3}$ (no change) with no over/underflow

4. Round and renormalize if necessary

- $1.110_{2} \times 2^{-3}$ (no change)

5. Determine sign: +ve $\times-\mathrm{ve} \Rightarrow-\mathrm{ve}$

- $-1.110_{2} \times 2^{-3}=-0.21875$


## FP Arithmetic Hardware

## FP multiplier is of similar complexity to FP

 adder- But uses a multiplier for significands instead of an adder
FP arithmetic hardware usually does
- Addition, subtraction, multiplication, division, reciprocal, square-root
- FP $\leftrightarrow$ integer conversion

Operations usually takes several cycles

- Can be pipelined


## FP Instructions in MIPS

FP hardware is coprocessor 1

- Adjunct processor that extends the ISA Separate FP registers
- 32 single-precision: \$f0, \$f1, ... \$f31
- Paired for double-precision: \$f0/\$f1, \$f2/\$f3, ...

Release 2 of MIPs ISA supports $32 \times 64$-bit FP reg's FP instructions operate only on FP registers

- Programs generally don't do integer ops on FP data, or vice versa
- More registers with minimal code-size impact

FP load and store instructions

- 1wc1, 1dc1, swc1, sdc1
e.g., 1dc1 \$f8, 32(\$sp)


## FP Instructions in MIPS

Single-precision arithmetic

- add.s, sub.s, mu7.s, div.s
e.g., add.s \$f0, \$f1, \$f6

Double-precision arithmetic

- add.d, sub.d, mu7.d, div.d
e.g., mul.d \$f4, \$f4, \$f6

Single- and double-precision comparison

- C. $x x$.s, c. $x x . d$ ( $x x$ is eq, $7 \mathrm{t}, 7 \mathrm{e}, \ldots$ )
- Sets or clears FP condition-code bit
e e.g.c.7t.s \$f3, \$f4
Branch on FP condition code true or false
- bc1t, bc1f
- e.g., bc1t TargetLabe1


## FP Example: ${ }^{\circ} \mathrm{F}$ to ${ }^{\circ} \mathrm{C}$

## C code:

float f2c (float fahr) \{ return ((5.0/9.0)*(fahr - 32.0));
\}

- fahr in \$f12, result in \$f0, literals in global memory space
Compiled MIPS code:
f2c: 1wc1 \$f16, const5(\$gp)
1wc2 \$f18, const9(\$gp)
div.s \$f16, \$f16, \$f18
lwc1 \$f18, const32(\$gp)
sub.s \$f18, \$f12, \$f18
mul.s \$f0, \$f16, \$f18
jr \$ra


## FP Example: Array Multiplication

$X=X+Y \times Z$

- All $32 \times 32$ matrices, 64 -bit double-precision elements

C code:
void mm (double x[][], double y[][], double z[][]) \{ int i, j, k;
for (i $=0 ; \mathrm{i}$ ! = 32; $\mathbf{i}=\mathbf{i}+1$ )
for ( $\mathrm{j}=0 ; \mathrm{j}!=32 ; \mathrm{j}=\mathrm{j}+1$ )
for $(k=0 ; k!=32 ; k=k+1)$
$x[i][j]=x[i][j]$
$+y[i][k]$ * $z[k][j] ;$
\}

- Addresses of $\mathrm{x}, \mathrm{y}, \mathrm{z}$ in \$a0, \$a1, \$a2, and i, j, k in \$s0, \$s1, \$s2


## FP Example: Array Multiplication

## MIPS code:



## FP Example: Array Multiplication

| $\begin{array}{lc} \hline \text { s11 } & \$ t \\ \text { addu } & \$ \\ \text { s11 } & \$ \\ \text { addu } & \$ \\ 1 . d & \$ \end{array}$ | $\begin{aligned} & \text { \$t0, \$s0, } 5 \\ & \$ t 0, \$ t 0, \$ s 2 \\ & \$ t 0, \$ t 0,3 \\ & \text { \$t0, \$a1, \$t0 } \\ & \$ \mathrm{t} 18,0(\$ \mathrm{t} 0) \end{aligned}$ | $\begin{aligned} & \# \$ t 0=i * 32 \text { (size of row of } y \text { ) } \\ & \# \$ t 0=\mathrm{i} \text { *size(row) }+k \\ & \# \$ t 0=\text { byte offset of [i][k] } \\ & \# \$ t 0=\text { byte address of } y[i][k] \\ & \# \$+18=8 \text { bytes of } y[i][k] \end{aligned}$ |
| :---: | :---: | :---: |
| $\begin{aligned} & \text { mu1.d } \\ & \text { add.d } \end{aligned}$ | \$f16, \$f18, \$f16 \$f4, \$f4, \$f16 | $\begin{aligned} & \# \$ f 16=y[i][k] * z[k][j] \\ & \# \text { f4=x[i][j] }+y[i][k] * z[k][j] \end{aligned}$ |
| addiu <br> bne s.d | \$s2, \$s2, 1 \$s2, \$t1, L3 \$f4, 0 (\$t2) | $\begin{aligned} & \text { \# \$k k + 1 } \\ & \# \text { if }(k \text { ! }=32) \text { go to L3 } \\ & \# \times[i][j]=\$ f 4 \end{aligned}$ |
| addiu bne | $\begin{array}{lll} \$ \mathrm{~s} 1, & \$ \mathrm{~s} 1, & 1 \\ \$ \mathrm{~s} 1, & \$ \mathrm{t} 1, & \mathrm{~L} 2 \end{array}$ | $\begin{aligned} & \# \$ j=j+1 \\ & \# \text { if }(j!=32) \text { go to } L 2 \end{aligned}$ |
| addiu bne | $\begin{array}{lll} \hline \$ \mathrm{~s} 0, & \$ \mathrm{~s} 0, & 1 \\ \$ \mathrm{~s} 0, & \$ \mathrm{t} 1, & \mathrm{~L} 1 \end{array}$ | $\begin{aligned} & \text { \# \$i }=\text { i }+1 \\ & \text { \# if (i ! }=32 \text { ) go to L1 } \end{aligned}$ |

## Accurate Arithmetic

IEEE Std 754 specifies additional rounding control

- Extra bits of precision (guard, round, sticky)
- Choice of rounding modes
- Allows programmer to fine-tune numerical behavior of a computation
Not all FP units implement all options
- Most programming languages and FP libraries just use defaults
Trade-off between hardware complexity, performance, and market requirements


## Interpretation of Data

Bits have no inherent meaning

- Interpretation depends on the instructions applied
Computer representations of numbers
- Finite range and precision
- Need to account for this in programs


## Associativity

Parallel programs may interleave operations in unexpected orders

- Assumptions of associativity may fail

|  |  | $(x+y)+z$ | $x+(y+z)$ |
| ---: | ---: | ---: | ---: |
| $x$ | $-1.50 \mathrm{E}+38$ |  | $-1.50 \mathrm{E}+38$ |
| $y$ | $1.50 \mathrm{E}+38$ | $0.00 \mathrm{E}+00$ |  |
| $z$ | 1.0 | 1.0 | $1.50 \mathrm{E}+38$ |
| $z$ |  | $1.00 \mathrm{E}+00$ | $0.00 \mathrm{E}+00$ |

Need to validate parallel programs under varying degrees of parallelism

## x86 FP Architecture

Originally based on 8087 FP coprocessor

- $8 \times 80$-bit extended-precision registers
- Used as a push-down stack
- Registers indexed from TOS: ST(0), ST(1), ...

FP values are 32-bit or 64 in memory

- Converted on load/store of memory operand
- Integer operands can also be converted on load/store
Very difficult to generate and optimize code
- Result: poor FP performance


## x86 FP Instructions

| Data transfer | Arithmetic | Compare | Transcendental |
| :--- | :--- | :--- | :--- |
| FILD mem/ST(i) | FIADDP mem/ST(i) | FICOMP | FPATAN |
| FISTP mem/ST(i) | FISUBRP mem/ST(i) | FIUCOMP | F2XMI |
| FLDPI | FIMULP mem/ST(i) | FSTSW AX/mem | FCOS |
| FLD1 | FIDIVRP mem/ST(i) |  | FPTAN |
| FLDZ | FSQRT |  | FPREM |
|  | FABS |  | FPSIN |
|  | FRNDINT |  | FYL2X |

## Optional variations

- I: integer operand
- P: pop operand from stack
- R: reverse operand order
- But not all combinations allowed


## Streaming SIMD Extension 2 (SSE2)

Adds $4 \times 128$-bit registers

- Extended to 8 registers in AMD64/EM64T

Can be used for multiple FP operands

- $2 \times 64$-bit double precision
- $4 \times 32$-bit double precision
- Instructions operate on them simultaneously Single-Instruction Multiple-Data


## Right Shift and Division

Left shift by $i$ places multiplies an integer by $2^{i}$
Right shift divides by $2^{2}$ ?

- Only for unsigned integers For signed integers
- Arithmetic right shift: replicate the sign bit
- e.g., -5 / 4
$=11111011_{2} \gg 2=11111110_{2}=-2$
Rounds toward $-\infty$
- c.f. $11111011_{2} \ggg 2=00111110_{2}=+62$


## Who Cares About FP Accuracy?

Important for scientific code

- But for everyday consumer use?
"My bank balance is out by $0.0002 \phi$ !" $\cdot$
The Intel Pentium FDIV bug
- The market expects accuracy
- See Colwell, The Pentium Chronicles


## Concluding Remarks

ISAs support arithmetic

- Signed and unsigned integers
- Floating-point approximation to reals

Bounded range and precision

- Operations can overflow and underflow

MIPS ISA

- Core instructions: 54 most frequently used $100 \%$ of SPECINT, $97 \%$ of SPECFP
- Other instructions: less frequent

