
5DV008 Computer Architecture Umeå University Department of Computing Science

Stephen J. Hegner

Topic 3: Arithmetic

These slides are mostly taken verbatim, or with minor changes, from those prepared by

Mary Jane Irwin (www.cse.psu.edu/~mji)

of The Pennsylvania State University

[Adapted from *Computer Organization and Design, 4th Edition*,
Patterson & Hennessy, © 2008, MK]

11/16/10

1

5DV008 20101611 t:3 sl:1

Hegner UU

Key to the Slides

□ The source of each slide is coded in the footer on the right side:

- Irwin CSE331 = slide by Mary Jane Irwin from the course CSE331 (Computer Organization and Design) at Pennsylvania State University.
- Irwin CSE431 = slide by Mary Jane Irwin from the course CSE431 (Computer Architecture) at Pennsylvania State University.
- Hegner UU = slide by Stephen J. Hegner at Umeå University.

11/16/10

2

5DV008 20101611 t:3 sl:2

Hegner UU

Review: MIPS (RISC) Design Principles

□ Simplicity favors regularity

- fixed size instructions
- small number of instruction formats
- opcode always the first 6 bits

□ Smaller is faster

- limited instruction set
- limited number of registers in register file
- limited number of addressing modes

□ Make the common case fast

- arithmetic operands from the register file (load-store machine)
- allow instructions to contain immediate operands

□ Good design demands good compromises

- three instruction formats

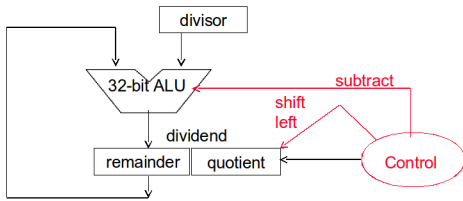
11/16/10

3

5DV008 20101611 t:3 sl:3

Irwin CSE431 PSU

Left Shift and Subtract Division Hardware



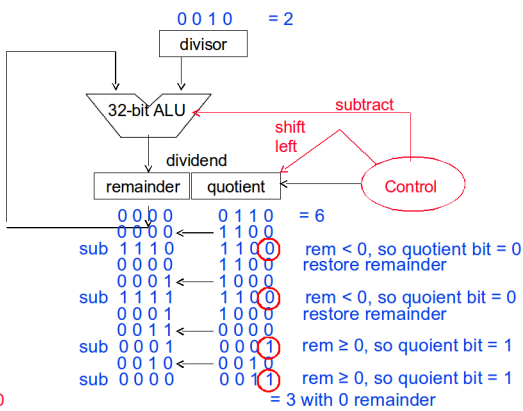
11/16/10

SDV008 20101611 t:3 sl:16

16

Irwin CSE431 PSU

Left Shift and Subtract Division Hardware



11/16/10

SDV008 20101611 t:3 sl:17

17

Irwin CSE431 PSU

MIPS Divide Instruction

- Divide (`div` and `divu`) generates the remainder in `hi` and the quotient in `lo`

```
div    $s0, $s1    # lo = $s0 / $s1
                    # hi = $s0 mod $s1
```

0	16	17	0	0	0x1A
---	----	----	---	---	------

- Instructions `mghi rd` and `mflo rd` are provided to move the quotient and remainder to (user accessible) registers in the register file
- As with multiply, divide ignores overflow so software must determine if the quotient is too large. Software must also check the divisor to avoid division by 0.

11/16/10

SDV008 20101611 t:3 sl:18

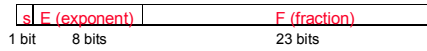
18

Irwin CSE431 PSU

Representing Big (and Small) Numbers

- What if we want to encode the approx. age of the earth?
4,600,000,000 or 4.6×10^9
or the weight in kg of one a.m.u. (atomic mass unit)
0.000000000000000000000000000000166 or 1.6×10^{-27}
There is no way we can encode either of the above in a 32-bit integer.

- Floating point representation $(-1)^{\text{sign}} \times F \times 2^E$
 - Still have to fit everything in 32 bits (single precision)



- The base (2, not 10) is hardwired in the design of the FPALU
- More bits in the fraction (F) or the exponent (E) is a trade-off between **precision** (accuracy of the number) and **range** (size of the number)

11/16/10

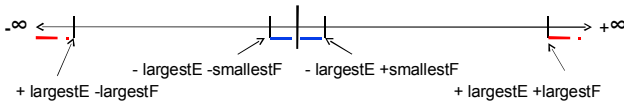
SDV008 20101611 t:3 sl:19

19

Irwin CSE431 PSU

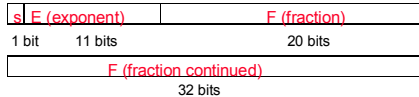
Exception Events in Floating Point

- Overflow** (floating point) happens when a positive exponent becomes too large to fit in the exponent field
- Underflow** (floating point) happens when a negative exponent becomes too large to fit in the exponent field



- One way to reduce the chance of underflow or overflow is to offer another format that has a larger exponent field

- Double precision – takes two MIPS words



11/16/10

SDV008 20101611 t:3 sl:20

20

Irwin CSE431 PSU

IEEE 754 FP Standard

- Most (all?) computers these days conform to the IEEE 754 floating point standard $(-1)^{\text{sign}} \times (1+F) \times 2^{E-\text{bias}}$

- Formats for both single and double precision
- F is stored in **normalized** format where the msb in F is 1 (so there is no need to store it!) – called the **hidden bit**
- To simplify sorting FP numbers, E comes before F in the word and E is represented in **excess** (biased) notation where the bias is -127 (-1023 for double precision) so the most negative is 00000001 = $2^{1-127} = 2^{-126}$ and the most positive is 11111110 = $2^{254-127} = 2^{+127}$

- Examples (in normalized format)

- Smallest+: 0 00000001 1.000000000000000000000000 = $1 \times 2^{1-127}$
- Zero: 0 00000000 000000000000000000000000 = true 0
- Largest+: 0 11111110 1.111111111111111111111111 = $2 \cdot 2^{23} \times 2^{254-127}$
- $1.0_2 \times 2^{-1} =$

11/16/10 0.75₁₀ × 2⁴ =

SDV008 20101611 t:3 sl:21

21

Irwin CSE431 PSU

Floating Point Addition

□ Addition (and subtraction)

$$(\pm F1 \times 2^{E1}) + (\pm F2 \times 2^{E2}) = \pm F3 \times 2^{E3}$$

- Step 0: Restore the hidden bit in F1 and in F2
- Step 1: **Align** fractions by right shifting F2 by E1 - E2 positions (assuming $E1 \geq E2$) keeping track of (three of) the bits shifted out in G R and S
- Step 2: **Add** the resulting F2 to F1 to form F3
- Step 3: **Normalize** F3 (so it is in the form 1.XXXXX ...)
 - If F1 and F2 have the same sign $\rightarrow F3 \in [1,4) \rightarrow$ 1 bit right shift F3 and increment E3 (check for overflow)
 - If F1 and F2 have different signs $\rightarrow F3$ may require *many* left shifts each time decrementing E3 (check for underflow)
- Step 4: **Round** F3 and possibly **normalize** F3 again
- Step 5: Rehide the most significant bit of F3 before storing the result

11/16/10

SDV008 20101611 t:3 sl:25

25

Irwin CSE431 PSU

Floating Point Addition Example

□ Add

$$(0.5 = 1.0000 \times 2^{-1}) + (-0.4375 = -1.1100 \times 2^{-2})$$

- Step 0:
- Step 1:
- Step 2:
- Step 3:
- Step 4:
- Step 5:

11/16/10

SDV008 20101611 t:3 sl:26

26

Irwin CSE431 PSU

Floating Point Addition Example

□ Add

$$(0.5 = 1.0000 \times 2^{-1}) + (-0.4375 = -1.1100 \times 2^{-2})$$

- Step 0: **Hidden bits restored in the representation above**
- Step 1: **Shift significand with the smaller exponent (1.1100) right until its exponent matches the larger exponent (so once)**
- Step 2: **Add significands**
 $1.0000 + (-0.111) = 1.0000 - 0.111 = 0.001$
- Step 3: **Normalize the sum, checking for exponent over/underflow**
 $0.001 \times 2^{-1} = 0.010 \times 2^{-2} = \dots = 1.000 \times 2^{-4}$
- Step 4: **The sum is already rounded, so we're done**
- Step 5: **Rehide the hidden bit before storing**

11/16/10

SDV008 20101611 t:3 sl:27

27

Irwin CSE431 PSU

Floating Point Multiplication

□ Multiplication

$$(\pm F1 \times 2^{E1}) \times (\pm F2 \times 2^{E2}) = \pm F3 \times 2^{E3}$$

- Step 0: Restore the hidden bit in F1 and in F2
- Step 1: Add the two (biased) exponents and subtract the bias from the sum, so $E1 + E2 - 127 = E3$
also determine the sign of the product (which depends on the sign of the operands (most significant bits))
- Step 2: Multiply F1 by F2 to form a double precision F3
- Step 3: Normalize F3 (so it is in the form 1.XXXXX ...)
 - Since F1 and F2 come in normalized $\rightarrow F3 \in [1,4) \rightarrow$ 1 bit right shift F3 and increment E3
 - Check for overflow/underflow
- Step 4: Round F3 and possibly normalize F3 again
- Step 5: Rehide the most significant bit of F3 before storing the result

11/16/10

SDV008 20101611 t:3 sl:28

28

Irwin CSE431 PSU

Floating Point Multiplication Example

□ Multiply

$$(0.5 = 1.0000 \times 2^{-1}) \times (-0.4375 = -1.1100 \times 2^{-2})$$

- Step 0:
- Step 1:

- Step 2:

- Step 3:

- Step 4:

- Step 5:

11/16/10

SDV008 20101611 t:3 sl:29

29

Irwin CSE431 PSU

Floating Point Multiplication Example

□ Multiply

$$(0.5 = 1.0000 \times 2^{-1}) \times (-0.4375 = -1.1100 \times 2^{-2})$$

- Step 0: Hidden bits restored in the representation above
- Step 1: Add the exponents (not in bias would be $-1 + (-2) = -3$ and in bias would be $(-1+127) + (-2+127) - 127 = (-1-2) + (127+127-127) = -3 + 127 = 124$)
- Step 2: Multiply the significands
 $1.0000 \times 1.110 = 1.110000$
- Step 3: Normalized the product, checking for exp over/underflow
 1.110000×2^{-3} is already normalized
- Step 4: The product is already rounded, so we're done
- Step 5: Rehide the hidden bit before storing

11/16/10

SDV008 20101611 t:3 sl:30

30

Irwin CSE431 PSU
