

CSSE132

Introduction to Computer Systems

4 : Integer Arithmetic

March 7, 2013

Today: Integer arithmetic

- **Data type ranges**
- **Addition (and subtraction)**
 - Overflow & modularity
 - Unsigned
 - Signed (Two's complement)
- **Multiplication**
 - Unsigned
 - Signed (Two's complement)
- **Division**

Data Representations (byte count)

C Data Type	Typical 32-bit	Intel IA32	x86-64
char	1	1	1
short	2	2	2
int	4	4	4
long	4	4	8
long long	8	8	8
float	4	4	4
double	8	8	8
long double	8	10/12	10/16
pointer	4	4	8

Values for Different Word Sizes

	W			
	8	16	32	64
UMax	255	65,535	4,294,967,295	18,446,744,073,709,551,615
TMax	127	32,767	2,147,483,647	9,223,372,036,854,775,807
TMin	-128	-32,768	-2,147,483,648	-9,223,372,036,854,775,808

- Sums require more digits than the inputs:

$$9+9 = 18$$

$$99+99 = 198$$

- Same issue occurs when adding binary numbers

Today: Integer arithmetic


- **Data type ranges**
- **Addition (and subtraction)**
 - Overflow & modularity
 - Unsigned
 - Signed (Two's complement)
- **Multiplication**
 - Unsigned
 - Signed (Two's complement)
- **Division**

Binary addition

- **Start with 4 basic cases:**
 - $0 + 0$
 - $0 + 1$
 - $1 + 0$
 - $1 + 1$
- **These 4 cases form the 'truth table'**
 - Similar to the boolean truth tables from yesterday
- **May result in carry out (1+1)!**
 - Add another output to truth table

Unsigned Addition


Operands: w bits



True Sum: $w+1$ bits



Discard Carry: w bits



■ Standard Addition Function

- Ignores carry output

■ Implements Modular Arithmetic


$$s = \text{UAdd}_w(u, v) = u + v \bmod 2^w$$

$$\text{UAdd}_w(u, v) = \begin{cases} u + v & u + v < 2^w \\ u + v - 2^w & u + v \geq 2^w \end{cases}$$

Visualizing (Mathematical) Integer Addition

■ Integer Addition

- 4-bit integers u, v
- Compute true sum $\text{Add}_4(u, v)$
- Values increase linearly with u and v
- Forms planar surface



Binary addition


- **Start with 4 basic cases:**
 - $0 + 0$
 - $0 + 1$
 - $1 + 0$
 - $1 + 1$
- **Basic table may result in carry out (1+1)**
 - Refine table with 3 inputs: A, B, C
 - 2 outputs: R, C

Visualizing Unsigned Addition


■ Wraps Around

- If true sum $\geq 2^w$
- At most once

True Sum



Modular Sum



Negation: Complement & Increment

- Claim: Following Holds for 2's Complement

$$\sim x + 1 == -x$$

- Complement

- Observation: $\sim x + x == 1111\dots111 == -1$

x	1	0	0	1	1	1	0	1	
+	$\sim x$	0	1	1	0	0	0	1	0
-1	1	1	1	1	1	1	1	1	1

Complement & Increment Examples

x = 15213

	Decimal	Hex	Binary
x	15213	3B 6D	00111011 01101101
~x	-15214	C4 92	11000100 10010010
~x+1	-15213	C4 93	11000100 10010011
y	-15213	C4 93	11000100 10010011

x = 0


	Decimal	Hex	Binary
0	0	00 00	00000000 00000000
~0	-1	FF FF	11111111 11111111
~0+1	0	00 00	00000000 00000000

Two's Complement Addition

Operands: w bits



True Sum: $w+1$ bits



Discard Carry: w bits



■ TAdd and UAdd have Identical Bit-Level Behavior

- Signed vs. unsigned addition in C:

```
int s, t, u, v;
```

```
s = (int) ((unsigned) u + (unsigned) v);
```


```
t = u + v
```

- Will give `s == t`

TAdd Overflow

■ Functionality

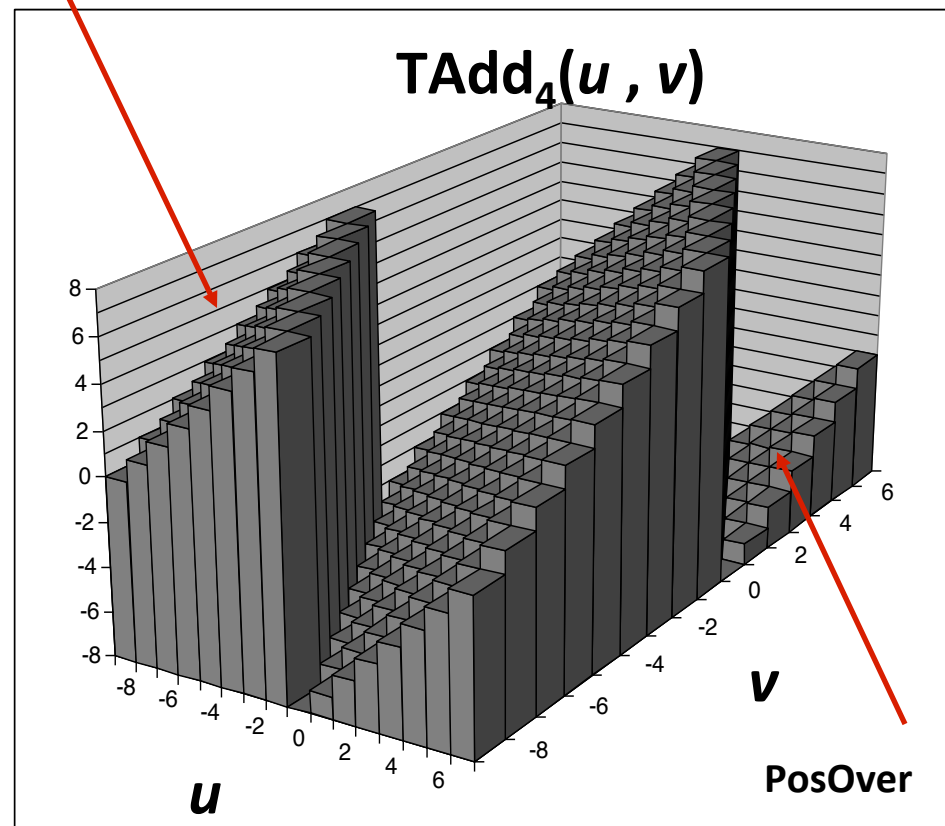
- True sum requires $w+1$ bits
- Drop off MSB
- Treat remaining bits as 2's comp. integer



Visualizing 2's Complement Addition

- **Values**
 - 4-bit two's comp.
 - Range from -8 to +7
- **Wraps Around**
 - If $\text{sum} \geq 2^{w-1}$
 - Becomes negative
 - At most once
 - If $\text{sum} < -2^{w-1}$
 - Becomes positive
 - At most once

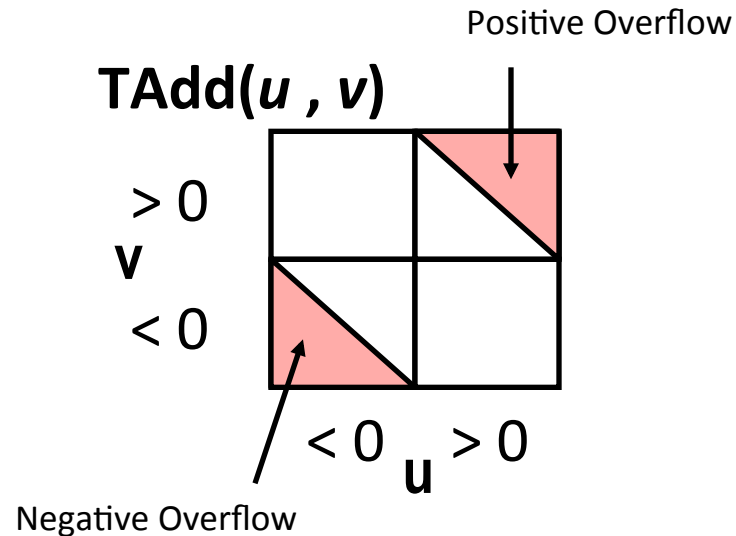
NegOver



Characterizing TAdd

■ Functionality

- True sum requires $w+1$ bits
- Drop off MSB
- Treat remaining bits as 2's comp. integer



$$TAdd_w(u, v) = \begin{cases} u + v + 2^w & u + v < TMin_w \text{ (NegOver)} \\ u + v & TMin_w \leq u + v \leq TMax_w \\ u + v - 2^w & TMax_w < u + v \text{ (PosOver)} \end{cases}$$

Subtraction

- **Similar to decimal**
 - $A - B$ is the same as
 - $A + (-B)$

- **Use addition and 2's complement**
 - Take 2's complement of subtrahend
 - (the number being subtracted)
 - Then add

Overflow

■ Detect using

- Operation (add or subtract)
- Sign of inputs (A and B)
- Sign of output (R)

Op	Sign of A	Sign of B	Overflow if R	Expected
+	≥ 0	≥ 0	< 0	≥ 0
+	< 0	< 0	≥ 0	< 0
-	≥ 0	< 0	< 0	> 0
-	< 0	≥ 0	≥ 0	< 0

Today: Integer arithmetic

- **Data type ranges**
- **Addition (and subtraction)**
 - Overflow & modularity
 - Unsigned
 - Signed (Two's complement)
- **Multiplication**
 - Unsigned
 - Signed (Two's complement)
- **Division**

Multiplication

■ Computing Exact Product of w -bit numbers x, y

- Either signed or unsigned


■ Ranges

- Unsigned: $0 \leq x * y \leq (2^w - 1)^2 = 2^{2w} - 2^{w+1} + 1$
 - Up to $2w$ bits
- Two's complement min: $x * y \geq (-2^{w-1}) * (2^{w-1} - 1) = -2^{2w-2} + 2^{w-1}$
 - Up to $2w-1$ bits
- Two's complement max: $x * y \leq (-2^{w-1})^2 = 2^{2w-2}$
 - Up to $2w$ bits, but only for $(TMin_w)^2$

■ Maintaining Exact Results

- Would need to keep expanding word size with each product computed
- Done in software by “arbitrary precision” arithmetic packages

Unsigned Multiplication in C




■ Standard Multiplication Function

- Ignores high order w bits

■ Implements Modular Arithmetic

$$\text{UMult}_w(u, v) = u \cdot v \bmod 2^w$$

Signed Multiplication in C



■ Standard Multiplication Function


- Ignores high order w bits
- Some of which are different for signed vs. unsigned multiplication
- Lower bits are the same

Power-of-2 Multiply with Shift


■ Operation

- $u \ll k$ gives $u * 2^k$
- Both signed and unsigned

Operands: w bits



True Product: $w+k$ bits $u \cdot 2^k$



Discard k bits: w bits

UMult_w($u, 2^k$)
TMult_w($u, 2^k$)



■ Examples

- $u \ll 3 \quad \quad \quad == \quad u * 8$
- $u \ll 5 - u \ll 3 \quad == \quad u * 24$
- Most machines shift and add faster than multiply
 - Compiler generates this automatically (strength reduction)

Compiled Multiplication Code

C Function

```
int mul12(int x)
{
    return x*12;
}
```

Compiled Arithmetic Operations

```
leal (%eax,%eax,2), %eax
sall $2, %eax
```

Explanation

```
t <- x+x*2
return t << 2;
```

- C compiler automatically generates shift/add code when multiplying by constant


Today: Integer arithmetic

- **Data type ranges**
- **Addition (and subtraction)**
 - Overflow & modularity
 - Unsigned
 - Signed (Two's complement)
- **Multiplication**
 - Unsigned
 - Signed (Two's complement)
- **Division**

Unsigned Power-of-2 Divide with Shift

■ Quotient of Unsigned by Power of 2

- $u \gg k$ gives $\lfloor u / 2^k \rfloor$
- Uses logical shift



	Division	Computed	Hex	Binary
x	15213	15213	3B 6D	00111011 01101101
x >> 1	7606.5	7606	1D B6	00011101 10110110
x >> 4	950.8125	950	03 B6	00000011 10110110
x >> 8	59.4257813	59	00 3B	00000000 00111011

Compiled Unsigned Division Code

C Function

```
unsigned udiv8(unsigned x)
{
    return x/8;
}
```

Compiled Arithmetic Operations

```
shrl $3, %eax
```

Explanation


```
# Logical shift
return x >> 3;
```

- Uses logical shift for unsigned
- For Java Users
 - Logical shift written as >>>

Signed Power-of-2 Divide with Shift

■ Quotient of Signed by Power of 2


- $x \gg k$ gives $\lfloor x / 2^k \rfloor$
- Uses arithmetic shift
- Rounds wrong direction when $u < 0$



	Division	Computed	Hex	Binary
y	-15213	-15213	C4 93	11000100 10010011
$y \gg 1$	-7606.5	-7607	E2 49	11100010 01001001
$y \gg 4$	-950.8125	-951	FC 49	11111100 01001001
$y \gg 8$	-59.4257813	-60	FF C4	11111111 11000100

Correcting negative shift divide

- Need to adjust result when shifting negative number
 - Add 1 to result



	Division	Computed	+1 if < 0
y	-15213	-15213	No change
$y \gg 1$	-7606.5	-7607	-7606
$y \gg 4$	-950.8125	-951	-950
$y \gg 8$	-59.4257813	-60	-59

Compiled Signed Division Code

C Function

```
int idiv8(int x)
{
    return x/8;
}
```

Compiled Arithmetic Operations

```
    testl %eax, %eax
    js    L4
L3:
    sarl $3, %eax
    ret
L4:
    addl $7, %eax
    jmp  L3
```

Explanation

```
if x < 0
    x += 7;
# Arithmetic shift
return x >> 3;
```

- Uses arithmetic shift for int
- For Java Users
 - Arith. shift written as >>

Arithmetic: Basic Rules

■ Addition:

- Unsigned/signed: Normal addition followed by truncate, same operation on bit level
- Unsigned: addition mod 2^w
 - Mathematical addition + possible subtraction of $2w$
- Signed: modified addition mod 2^w (result in proper range)
 - Mathematical addition + possible addition or subtraction of $2w$

■ Multiplication:

- Unsigned/signed: Normal multiplication followed by truncate, same operation on bit level
- Unsigned: multiplication mod 2^w
- Signed: modified multiplication mod 2^w (result in proper range)

Arithmetic: Basic Rules

- **Unsigned ints, 2's complement ints are isomorphic rings: isomorphism = casting**
- **Left shift**
 - Unsigned/signed: multiplication by 2^k
 - Always logical shift
- **Right shift**
 - Unsigned: logical shift, div (division + round to zero) by 2^k
 - Signed: arithmetic shift
 - Positive numbers: div (division + round to zero) by 2^k
 - Negative numbers: div (division + round away from zero) by 2^k
Use biasing to fix

Why Should I Use Unsigned?

■ *Don't Use Just Because Number Nonnegative*

- Easy to make mistakes

```
unsigned i;  
for (i = cnt-2; i >= 0; i--)  
    a[i] += a[i+1];
```

- Can be very subtle

```
#define DELTA sizeof(int)  
int i;  
for (i = CNT; i-DELTA >= 0; i-= DELTA)  
    . . .
```

■ *Do Use When Performing Modular Arithmetic*

- Multiprecision arithmetic

■ *Do Use When Using Bits to Represent Sets*

- Logical right shift, no sign extension

Weekly review

■ Monday

- System overview

■ Tuesday

- Bits and Bytes

■ Wednesday

- Boolean logic, signed numbers

■ Today

- Binary arithmetic