

CENG 3420

Computer Organization & Design



Lecture 05: Arithmetic and Logic Unit – 1

Bei Yu

CSE Department, CUHK


byu@cse.cuhk.edu.hk

(Textbook: Chapters 3.2 & A.5)

Spring 2022



Overview





Where we've been: **abstractions**

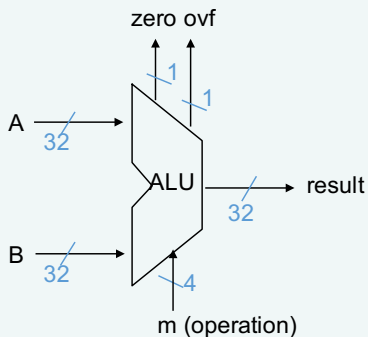
- Instruction Set Architecture (ISA)
- Assembly and machine language



Where we've been: **abstractions**

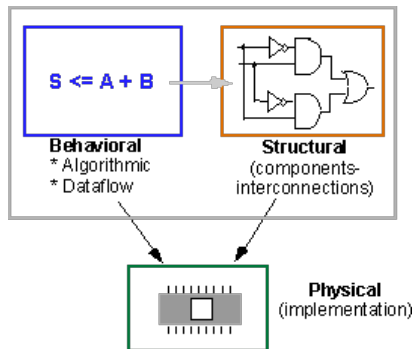
- Instruction Set Architecture (ISA)
- Assembly and machine language

What's up ahead: Implementing the ALU architecture






- Supports design, documentation, simulation & verification, and synthesis of hardware
- Allows integrated design at **behavioral** & **structural** levels





Basic Structure

- Design **entity-architecture** descriptions
- Time-based execution (discrete event simulation) model





Entity

defines externally visible characteristics

- Ports: channels of communication
 - signal names for inputs, outputs, clocks, control
- Generic parameters: define class of components
 - timing characteristics, size (fan-in), fan-out



Architecture

defines the internal behavior or structure of the circuit

- Declaration of internal signals
- Description of behavior
 - collection of Concurrent Signal Assignment (CSA) statements (indicated by \leq);
 - can also model temporal behavior with the delay annotation
 - one or more processes containing CSAs and (sequential) variable assignment statements (indicated by $:=$)
- Description of structure
 - interconnections of components; underlying behavioral models of each component must be specified



```
entity ALU is
  port(A, B:  in std_logic_vector (31 downto 0);
        m:  in std_logic_vector (3 downto 0);
        result: out std_logic_vector (31 downto 0);
        zero: out std_logic;
        ovf: out std_logic)
end ALU;

architecture process_behavior of ALU is
  . . .
begin
  ALU: process(A, B, m)
  begin
    . . .
    result := A + B;
    . . .
  end process ALU;
end process_behavior;
```



- Bits are just bits (have no inherent meaning)¹
- Binary numbers (base 2) – **integers**

Of course, it gets more complicated:

- **storage locations** (e.g., register file words) are finite, so have to worry about overflow (i.e., when the number is too big to fit into 32 bits)
- have to be able to represent **negative numbers**, e.g., how do we specify -8 in

```
addi    $sp, $sp, -8    #$sp = $sp - 8
```

- in real systems have to provide for more than just integers, e.g., fractions and real numbers (and **floating point**) and alphanumeric (**characters**)

¹conventions define the relationships between bits and numbers



32-bit signed numbers (2's complement):

```

0000 0000 0000 0000 0000 0000 0000 0000two = 0ten
0000 0000 0000 0000 0000 0000 0000 0001two = + 1ten
0000 0000 0000 0000 0000 0000 0000 0010two = + 2ten
...

0111 1111 1111 1111 1111 1111 1111 1110two = + 2,147,483,646ten
0111 1111 1111 1111 1111 1111 1111 1111two = + 2,147,483,647ten
1000 0000 0000 0000 0000 0000 0000 0000two = - 2,147,483,648ten
1000 0000 0000 0000 0000 0000 0000 0001two = - 2,147,483,647ten
1000 0000 0000 0000 0000 0000 0000 0010two = - 2,147,483,646ten
...

1111 1111 1111 1111 1111 1111 1111 1101two = - 3ten
1111 1111 1111 1111 1111 1111 1111 1110two = - 2ten
1111 1111 1111 1111 1111 1111 1111 1111two = - 1ten

```

What if the bit string represented addresses?

- need operations that also deal with only positive (unsigned) integers



- Negating a two's complement number – complement all the bits and then add a 1
 - remember: “negate” and “invert” are quite different!
- Converting n-bit numbers into numbers with more than n bits:
 - 16-bit immediate gets converted to 32 bits for arithmetic
 - **sign extend**: copy the most significant bit (the sign bit) into the other bits

```
0010  -> 0000 0010
1010  -> 1111 1010
```

- sign extension versus zero extend (lb vs. lbu)




- Must support the Arithmetic/Logic operations of the ISA

RV 32I:



```
add, sub, mul, mulh, mulhu, mulhsu,  
div, divu, rem, li, addi, sll, srl,  
sra, or, xor, not, slt, sltu, slli,  
srli, srai, andi, ori, xori, slti,  
sltiu,
```

RV 64I:

```
addw, subw, remu, mulw, divw, divuw,  
remw, remuw, addiw, sllw, srlw, sraw,  
srliw, sraiw,
```



- With special handling for:
 - sign extend: addi, slti, sltiu
 - zero extend: andi, xori
 - Overflow detected: add, addi, sub




Type	op	funct
ADDI	001000	xx
ADDIU	001001	xx
SLTI	001010	xx
SLTIU	001011	xx
ANDI	001100	xx
ORI	001101	xx
XORI	001110	xx
LUI	001111	xx

Type	op	funct
ADD	000000	100000
ADDU	000000	100001
SUB	000000	100010
SUBU	000000	100011
AND	000000	100100
OR	000000	100101
XOR	000000	100110
NOR	000000	100111

Type	op	funct
	000000	101000
	000000	101001
SLT	000000	101010
SLTU	000000	101011
	000000	101100



Addition Unit




A	B	carry_in	carry_out	S
0	0	0	0	0
0	0	1	0	1
0	1	0	0	1
0	1	1	1	0
1	0	0	0	1
1	0	1	1	0
1	1	0	1	0
1	1	1	1	1

$$S = A \text{ xor } B \text{ xor } \text{carry_in}$$

$$\text{carry_out} = A \& B \mid A \& \text{carry_in} \mid B \& \text{carry_in}$$

(majority function)

- How can we use it to build a 32-bit adder?
- How can we modify it easily to build an adder/subtractor?




- Just connect the carry-out of the least significant bit FA to the carry-in of the next least significant bit and connect ...
- Ripple Carry Adder (RCA)
 - 😊: simple logic, so small (low cost)
 - ☹️: slow and lots of glitching (so lots of energy consumption)




Glitch

invalid and unpredicted output that can be read by the next stage and result in a wrong action

Example: Draw the propagation delay






A	B	carry_in	carry_out	S
0	0	0	0	0
0	0	1	0	1
0	1	0	0	1
0	1	1	1	0
1	0	0	0	1
1	0	1	1	0
1	1	0	1	0
1	1	1	1	1

But What about Performance?




- Critical path of n-bit ripple-carry adder is $n \times CP$
- **Design trick:** throw hardware at it (Carry Lookahead)



A 32-bit Ripple Carry Adder/Subtractor

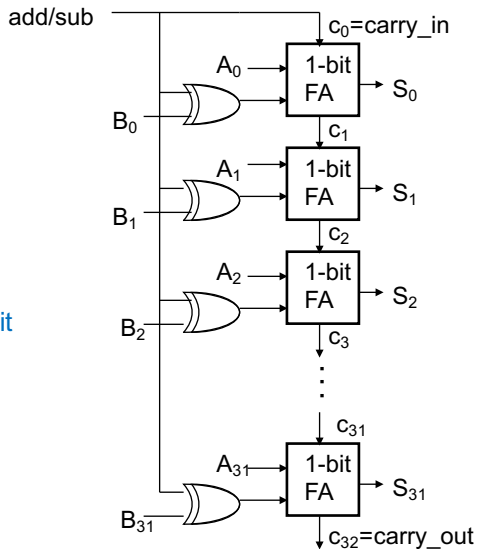


- complement all the bits



- add a 1 in the least significant bit

$$\begin{array}{r}
 A \quad 0111 \quad \rightarrow \quad 0111 \\
 B \quad -0110 \quad \rightarrow \quad + \quad 1001 \\
 \hline
 \quad \quad 0001 \\
 \hline
 \quad \quad \quad 1 \quad 0001
 \end{array}$$






Gate library: inverters, 2-input NANDs, or-and-inverters

```
architecture concurrent_behavior of full_adder is
    signal t1, t2, t3, t4, t5: std_logic;
begin
    t1 <= not A after 1 ns;
    t2 <= not cin after 1 ns;
    t4 <= not((A or cin) and B) after 2 ns;
    t3 <= not((t1 or t2) and (A or cin)) after 2 ns;
    t5 <= t3 nand B after 2 ns;
    S <= not((B or t3) and t5) after 2 ns;
    cout <= not((t1 or t2) and t4) after 2 ns;
end concurrent_behavior;
```




- Also need to support the logic operations (*and*, *nor*, *or*, *xor*)
 - Bit wise operations (no carry operation involved)
 - Need a logic gate for each function and a mux to choose the output
- Also need to support the set-on-less-than instruction (*slt*)
 - Uses subtraction to determine if $(a - b) < 0$ (implies $a < b$)
- Also need to support test for equality (*bne*, *beq*)
 - Again use subtraction: $(a - b) = 0$ implies $a = b$
- Also need to add overflow detection hardware
 - overflow detection enabled only for *add*, *addi*, *sub*
- **Immediates** are sign extended outside the ALU with wiring (i.e., no logic needed)

A Simple ALU Cell with Logic Op Support



A Simple ALU Cell with Logic Op Support

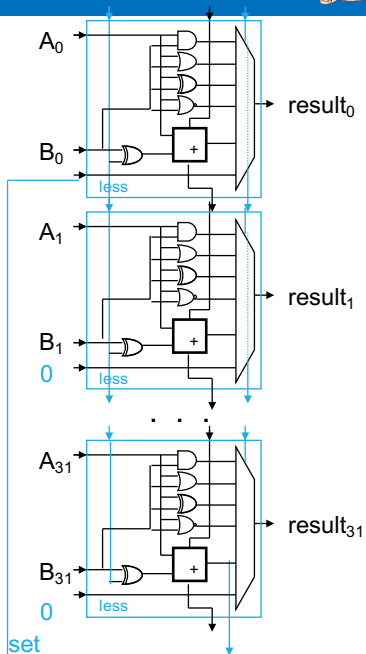


Modifying the ALU Cell for `slt`



Modifying the ALU for slt

- First perform a subtraction
- Make the result 1 if the subtraction yields a negative result
- Make the result 0 if the subtraction yields a positive result
- Tie the most significant sum bit (sign bit) to the low order **less** input







Overflow Detection

Overflow occurs when the result is too large to represent in the number of bits allocated

- adding two positives yields a negative
- or, adding two negatives gives a positive
- or, subtract a negative from a positive gives a negative
- or, subtract a positive from a negative gives a positive

Question: prove you can detect overflow by:

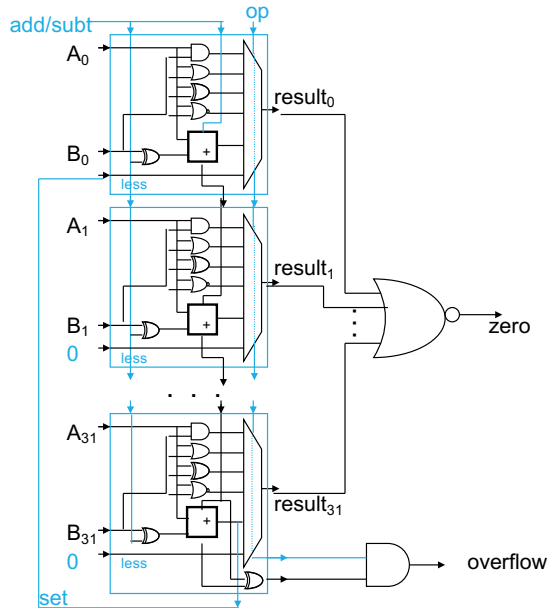
Carry into MSB xor Carry out of MSB



Modifying the ALU for Overflow



- Modify the most significant cell to determine overflow output setting
- Enable overflow bit setting for signed arithmetic (add, addi, sub)





- On overflow, an exception (interrupt) occurs
- Control jumps to predefined address for exception
- Interrupted address (address of instruction causing the overflow) is saved for possible resumption
- Don't always want to detect (interrupt on) overflow



Category	Instr	Op Code	Example	Meaning
Arithmetic (R & I format)	add unsigned	0 and 21	addu \$s1, \$s2, \$s3	\$s1 = \$s2 + \$s3
	sub unsigned	0 and 23	subu \$s1, \$s2, \$s3	\$s1 = \$s2 - \$s3
	add imm.unsigned	9	addiu \$s1, \$s2, 6	\$s1 = \$s2 + 6
Data Transfer	ld byte unsigned	24	lbu \$s1, 20(\$s2)	\$s1 = Mem(\$s2+20)
	ld half unsigned	25	lhu \$s1, 20(\$s2)	\$s1 = Mem(\$s2+20)
Cond. Branch (I & R format)	set on less than unsigned	0 and 2b	sltu \$s1, \$s2, \$s3	if (\$s2 < \$s3) \$s1=1 else \$s1=0
	set on less than imm unsigned	b	sltiu \$s1, \$s2, 6	if (\$s2 < 6) \$s1=1 else \$s1=0

- Sign extend: `addi`, `addiu`, `slti`
- Zero extend: `andi`, `ori`, `xori`
- Overflow detected: `add`, `addi`, `sub`




Binary Representations for Integers




Hey guys! How do you negate numbers?

In the early days of computing, designers made computers express numbers using **unsigned binary**.
And they were content... Until there were negative numbers.



To include negative numbers, designers came up with **sign magnitude**.
That took care of the negative numbers... But the computer had to count backwards for the negative numbers.



Plus, this introduced positive and negative zero.

Then designers created **one's complement**.
Now computers only had to count in one direction... But there were still two zeroes!



Finally, designers developed **two's complement**.
Now, there was only one zero... And they were content.

