

CENG 3420

Computer Organization & Design



Lecture 06: Arithmetic and Logic Unit

Bei Yu

CSE Department, CUHK

byu@cse.cuhk.edu.hk

(Textbook: Chapters 3.2 & A.5)


2024 Spring



- ① Overview
- ② Addition Unit
- ③ Multiplication & Division
- ④ Shifter



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





- 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





I-Type


Type	opcode	funct	imm[11:5]
ADDI	0010011	000	xx (any)
SLLI	0010011	001	0000000
SLTI	0010011	010	xx
SLTIU	0010011	011	xx
SRLI	0010011	101	0000000
SRAI	0010011	101	0100000
ORI	0010011	110	xx
ANDI	0010011	111	xx

R-Type

Type	opcode	funct
ADD	0110011	0000000 000
SUB	0110011	0100000 000
SLL	0110011	0000000 001
SLT	0110011	0000000 010
SLTU	0110011	0000000 011
XOR	0110011	0000000 100
SRL	0110011	0000000 101
SRA	0110011	0100000 101



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)

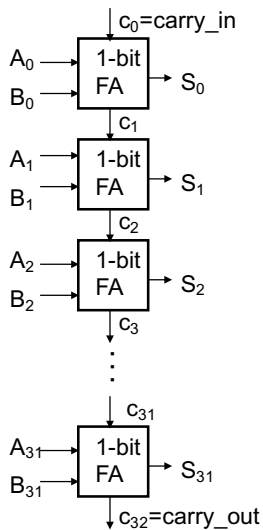


Glitch

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

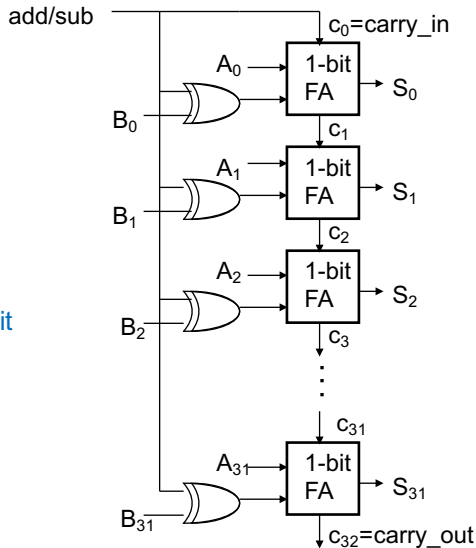


- complement all the bits



- add a 1 in the least significant bit


A	0111	->	0111
B	- 0110	->	+ 1001
	<u>0001</u>		<u>1</u>
			1 0001






- 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

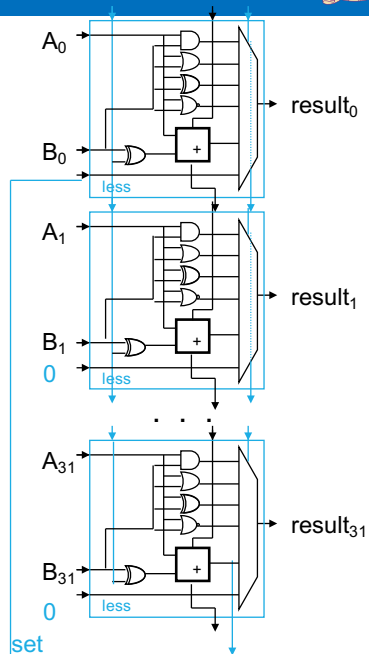


Modifying the ALU Cell for `slt`

Modifying the ALU for `slt`



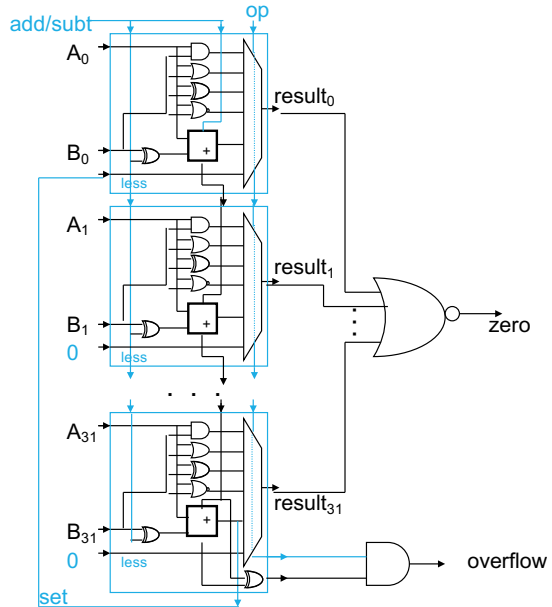
- First perform a subtraction
- Make the result 1 if the subtraction yields a negative result
- Make the result 0 if the subtraction yields a positive result
- Tie the most significant sum bit (sign bit) to the low order `less` input





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.





Multiplication & Division




- More complicated than addition
- Can be accomplished via **shifting** and **adding**

$$\begin{array}{r} 0010 \quad (\text{multiplicand}) \\ \times 1011 \quad (\text{multiplier}) \\ \hline 0010 \\ 0010 \\ 0000 \\ 0010 \\ \hline 00010110 \quad (\text{product}) \end{array}$$


The partial products (0010, 0010, 0000, 0010) are grouped by a bracket and labeled "(partial product array)". The final product (00010110) is highlighted with a blue box.

- Double precision product produced
- More time and more area to compute

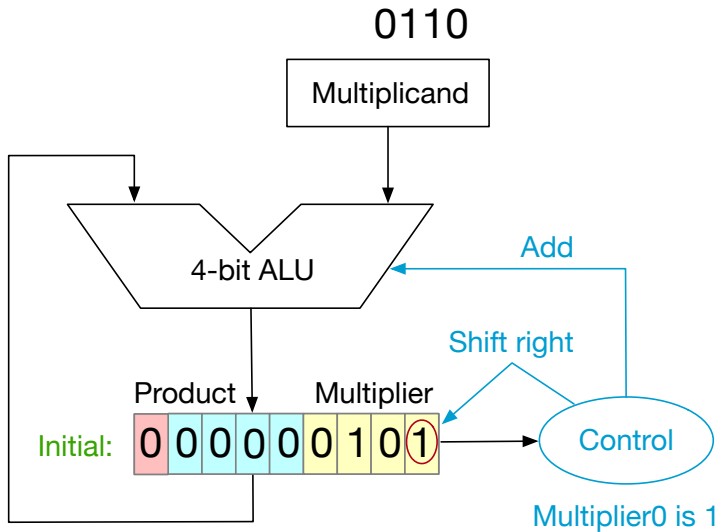
First Version of Multiplication Hardware

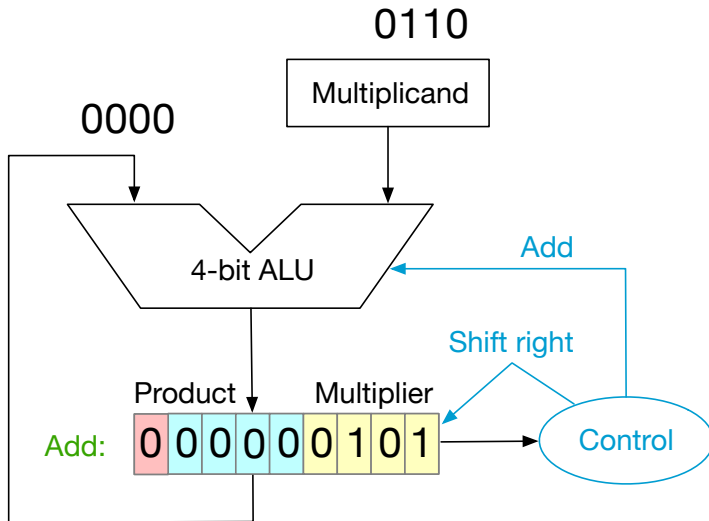



Note: $n\text{-bit} \times n\text{-bit}$ needs $2n\text{-bit}$ adder

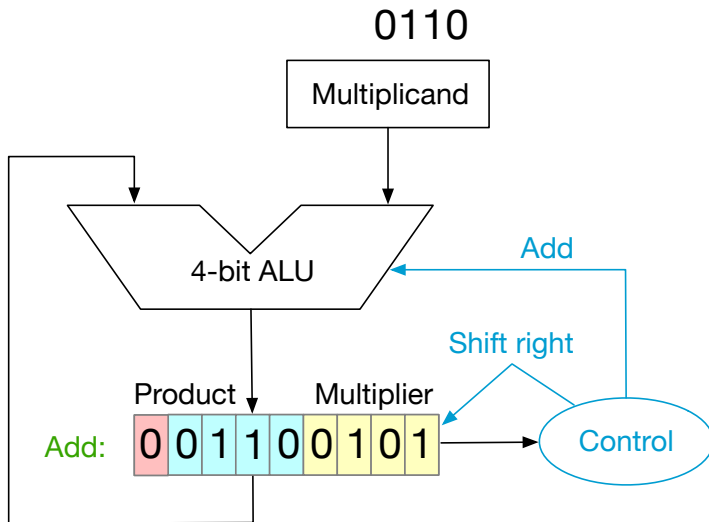



Note: $n\text{-bit} \times n\text{-bit}$ needs only $n\text{-bit}$ adder

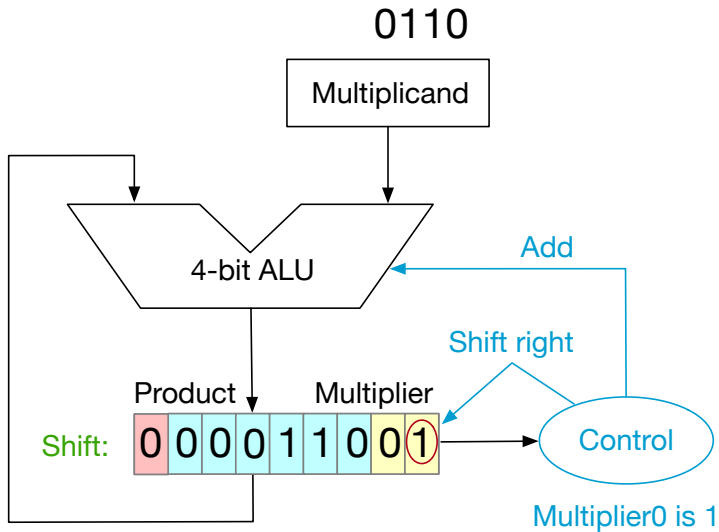


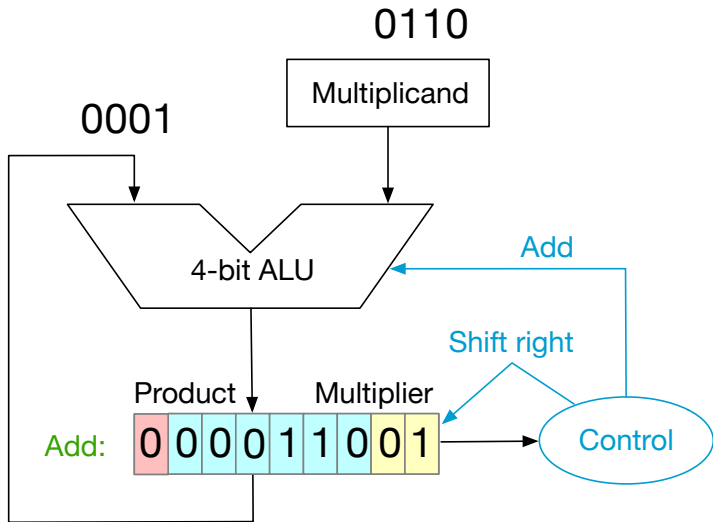


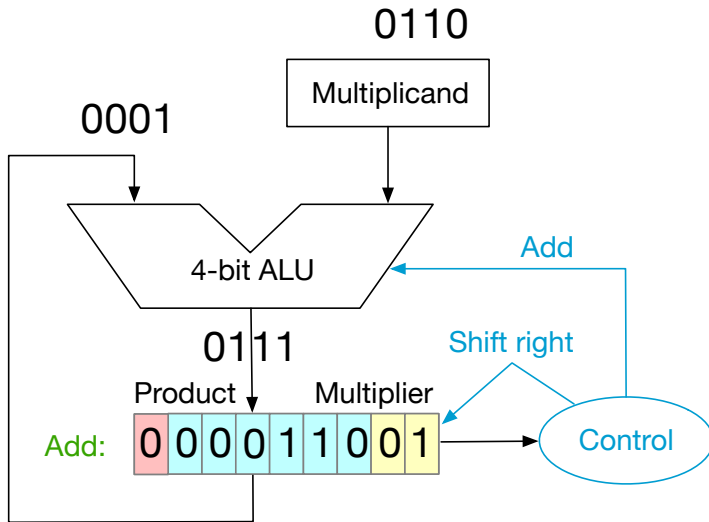


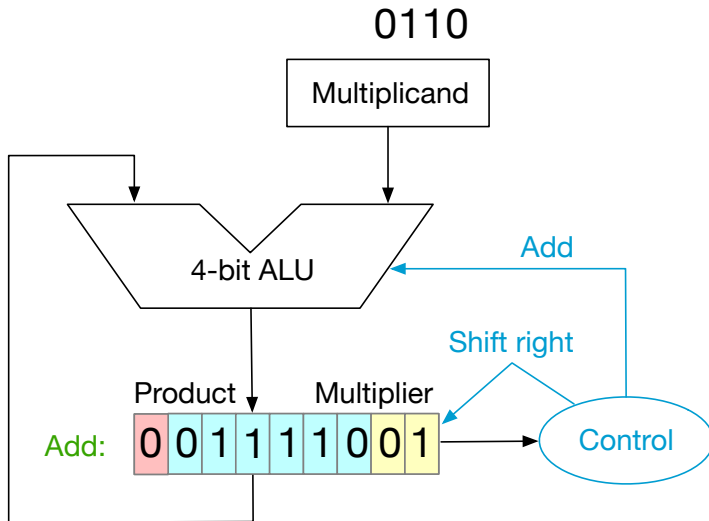


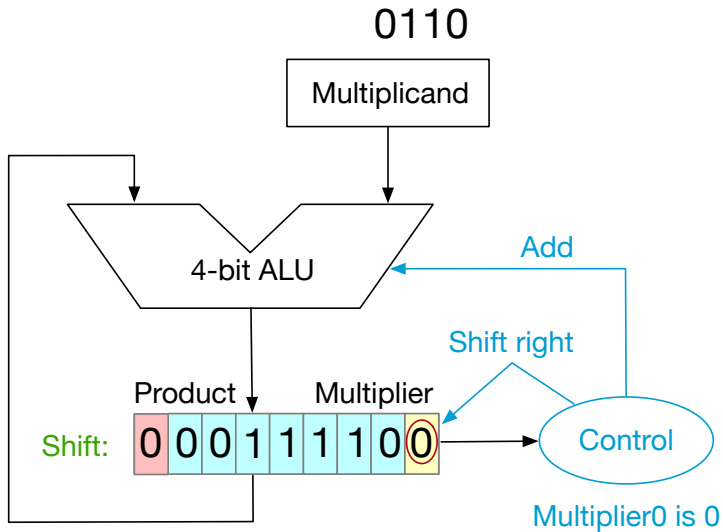


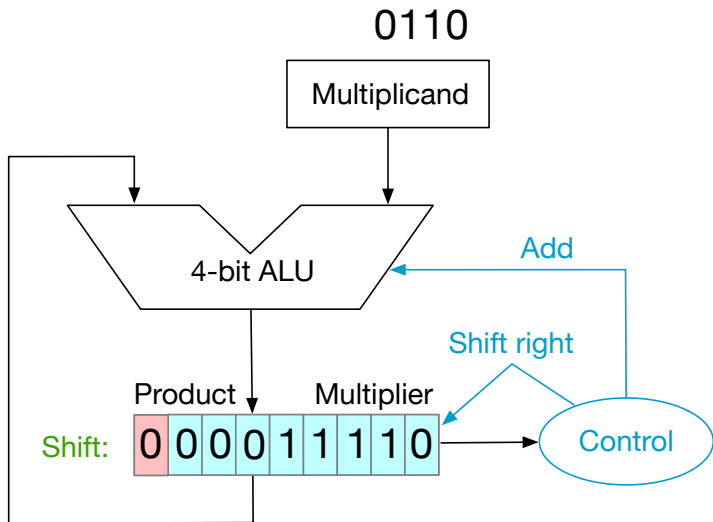













Final Result: 00011110 = 30



- `mul` performs an 32-bit \times 32-bit multiplication and places the lower 32 bits in the destination register.


```
mul rd, rs1, rs2
```



- `mulh`, `mulhu`, and `mulhsu` perform the same multiplication but return the upper 32 bits of the full 64-bit product, for signed \times signed, unsigned \times unsigned, and signed \times unsigned multiplication respectively.



- Division is just a bunch of quotient digit guesses and left shifts and subtracts






FIGURE 3.8 First version of the division hardware. The Divisor register, ALU, and Remainder register are all 64 bits wide, with only the Quotient register being 32 bits. The 32-bit divisor starts in the left half of the Divisor register and is shifted right 1 bit each iteration. The remainder is initialized with the dividend. Control decides when to shift the Divisor and Quotient registers and when to write the new value into the Remainder register.




Question: Division

Dividing 1001010 by 1000



- `div` generates the remainder in `hi` and the quotient in `lo`

```
div rd, rs1, rs2
```



- `div` perform an 32 bits by 32 bits signed integer division of `rs1` by `rs2`, rounding towards zero.
- `div` and `divu` perform signed and unsigned integer division of 32 bits by 32 bits.
- `rem` and `remu` provide the remainder of the corresponding division operation.




Shifter



- Shifts by a constant are encoded as a specialization of the I-type format. The operand to be shifted is in `rs1`, and the shift amount is encoded in the lower 5 bits of the I-immediate field.


```
srl  rd, rs1, imm[4:0]
srai rd, rs1, imm[4:0]
```





- `slli` is a logical left shift; `srl` is a logical right shift; and `srai` is an arithmetic right shift.
- Logical shifts fill with **zeros**, arithmetic left shifts fill with the **sign bit**

The shift operation is implemented by hardware separate from the ALU


Using a barrel shifter, which would takes lots of gates in discrete logic, but is pretty easy to implement in VLSI









Logarithmic Shifter Structure

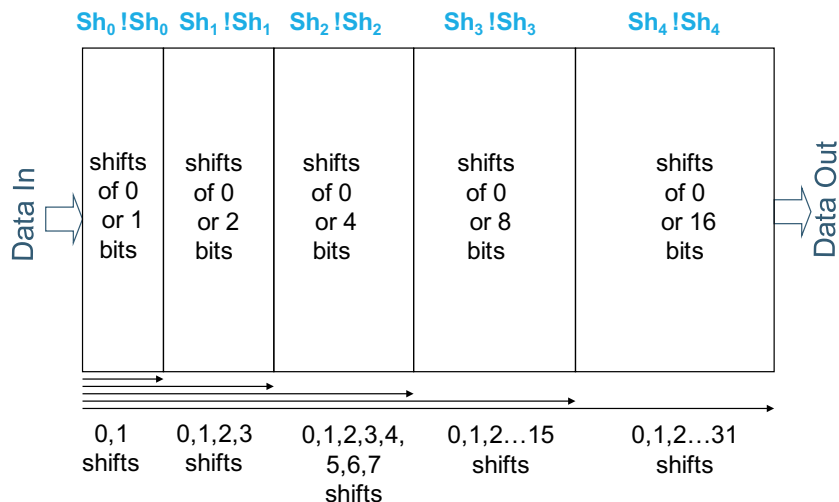




Logarithmic Shifter Structure



Logarithmic Shifter Structure



Logarithmic Shifter Structure

