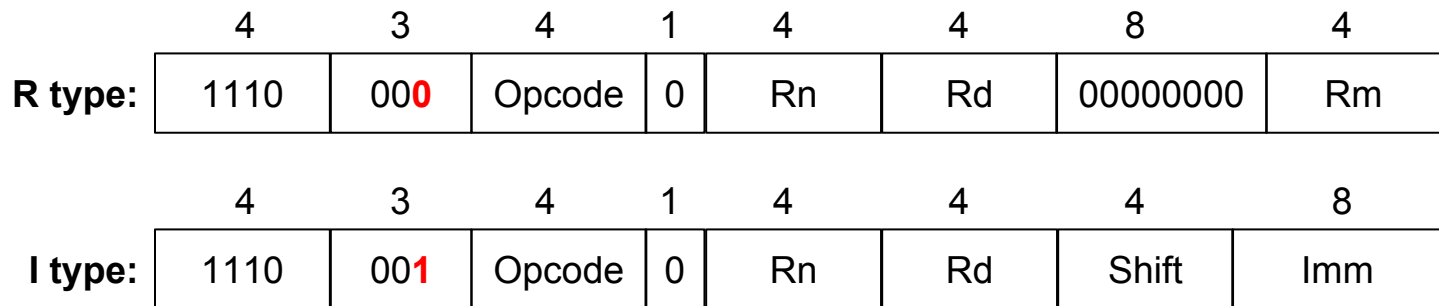




# BASIC ARM INSTRUCTIONS

- Instructions include various "fields" that encode combinations of *Opcodes* and *arguments*
- special fields enable extended functions (more in a minute)
- several 4-bit *OPERAND* fields, for specifying the sources and destination of the operation, usually one of the 16 registers
- Embedded constants ("immediate" values) of various sizes,

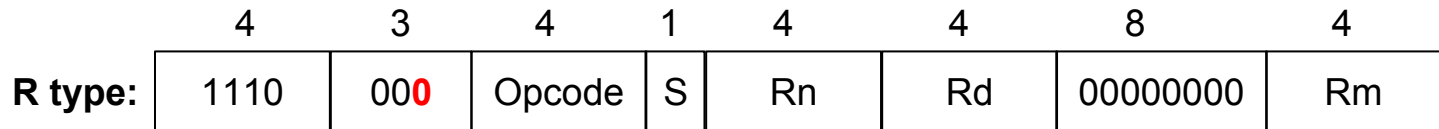
The "basic" data-processing instruction formats:





# R-TYPE DATA PROCESSING

Instructions that process three-register arguments:



Simple R-type instructions follow the following template:

OP      Rd, Rn, Rm

Later on we'll introduce more complex variants of these 'simple' R-type instructions.



- 0000 - AND
- 0001 - EOR
- 0010 - SUB
- 0011 - RSB
- 0100 - ADD
- 0101 - ADC
- 0110 - SBC
- 0111 - RSC
- 1000 - TST
- 1001 - TEQ
- 1010 - CMP
- 1011 - CMN
- 1100 - ORR
- 1101 - MOV
- 1110 - BIC
- 1111 - MVN

ADDR0, R1, R3

Is encoded as:

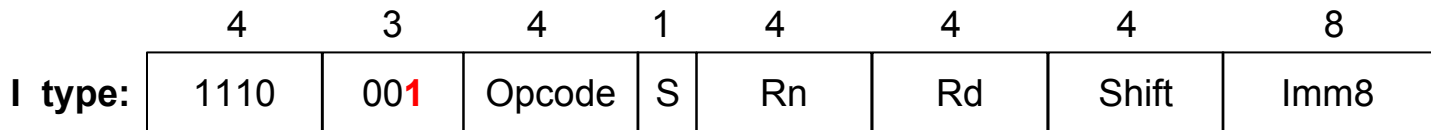
1110 0000 1000 0001 0000 0000 0000 0011

0xE0810003



# I-TYPE DATA PROCESSING

Instructions that process two registers and a constant:



Simple I-type instructions follow the following template:

OP      Rd, Rn, #constant

In the I-type instructions the second register operand is replaced by a constant that is encoded in the instruction



- 0000 - AND
- 0001 - EOR
- 0010 - SUB
- 0011 - RSB
- 0100 - ADD
- 0101 - ADC
- 0110 - SBC
- 0111 - RSC
- 1000 - TST
- 1001 - TEQ
- 1010 - CMP
- 1011 - CMN
- 1100 - ORR
- 1101 - MOV
- 1110 - BIC
- 1111 - MVN

RSBR7, R10, #49

Is encoded as:



0xE26A7031



# I-TYPE CONSTANTS

ARM7 provides only 8-bits for specifying an immediate constant value. Given that ARM7 is a 32-bit architecture, this may appear to be a severe limitation. However, by allowing for a shift (actually a "right rotation") to be applied to the constant.

$$\text{imm32} = (\text{imm8} \gg (2 * \text{shift})) | (\text{imm8} \ll (32 - (2 * \text{shift})))$$

Example: 1920 is encoded as:

Shift	Imm8
1101	00011110

$$\begin{aligned} &= (30 \gg (2 * 13)) | (30 \ll (32 - (2 * 13))) \\ &= \quad 0 \quad | \quad 30 * 64 \\ &= 1920 \end{aligned}$$

How would 256 be encoded?

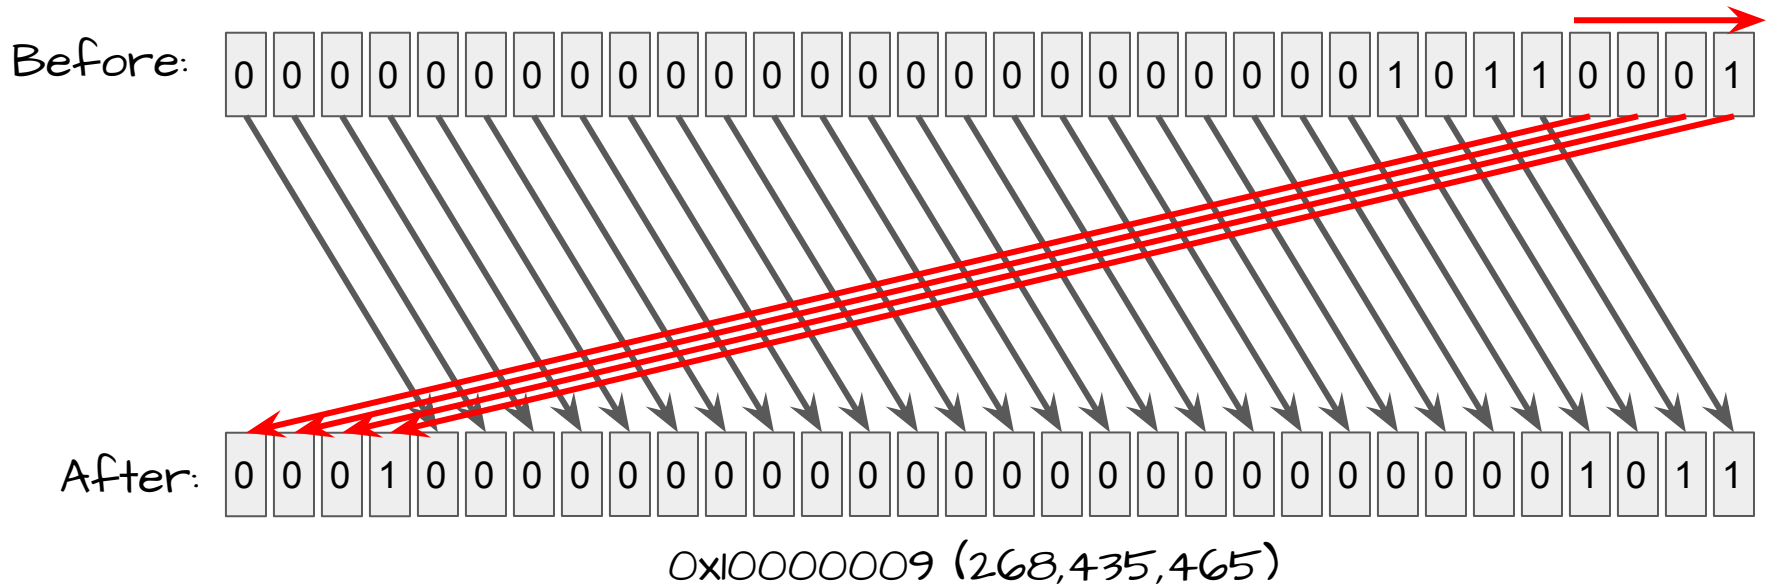
Shift	Imm8
1100	00000001



# ILLUSTRATING A RIGHT ROTATION

The "*shift*" field of the I-type instruction is a misnomer. It is actually a "*rotate-right*". What's a rotate right?

0xB1 (209) rotated 4 positions to the right





# ARM IMMEDIATE CONSTANTS

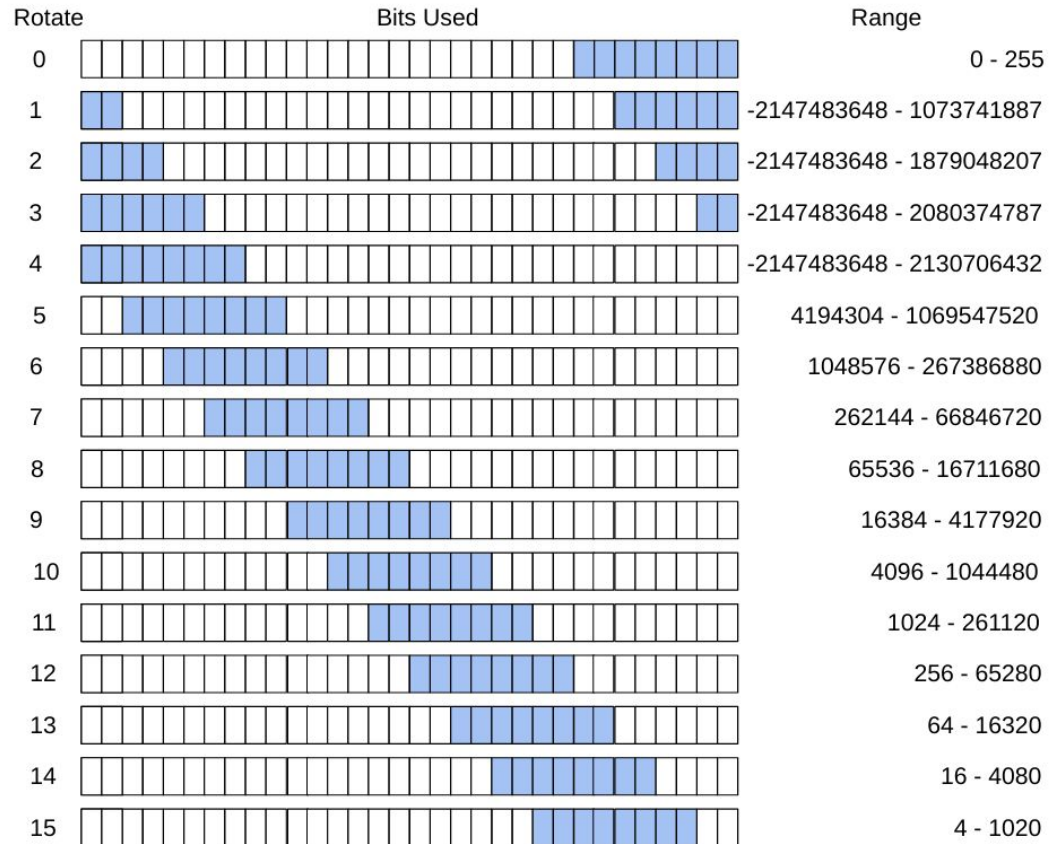
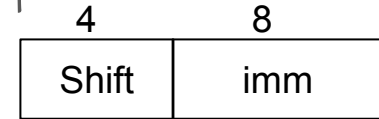
Recall that immediate constants are encoded in two parts:

Some constants can be encoded in multiple ways.

Thus fewer than 4096

32-bit numbers can be represented.

There are actually only 3073 distinct constants. There are 16, "0s" and 4 ways to represent all powers 2. How might you encode 256?



1100	00000001	1101	00000100
------	----------	------	----------

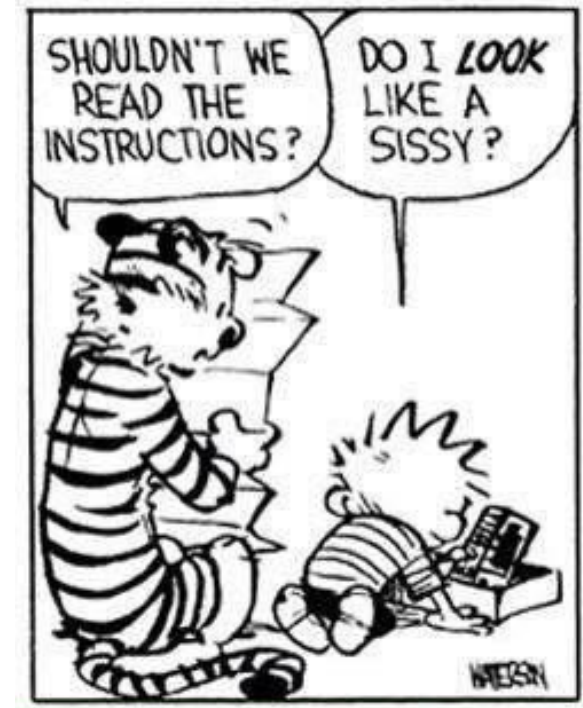
1110	00010000	1111	01000000
------	----------	------	----------



# READ THE INSTRUCTIONS

.. when all else fails

- What do instructions do?
- How are instructions decoded?
- Uniformity and Symmetry
- Cramming stuff in
- CPU state
  - Condition codes
  - Program Status Register (PSR)



# A CLOSER LOOK AT THE OPCODES



The Opcode field is common to both of the basic instruction types

	4	3	4	1	4	4	8	4
R type:	1110	000	Opcode	S	Rn	Rd	00000000	Rm
I type:	1110	001	Opcode	S	Rn	Rd	Rotate	Imm8

ARM data processing instructions can be broken into four basic groups:

- Arithmetic (6)
- Logic (4)
- Comparison (4)
- Register transfer (2)



- 0000 - AND
- 0001 - EOR
- 0010 - SUB
- 0011 - RSB
- 0100 - ADD
- 0101 - ADC
- 0110 - SBC
- 0111 - RSC
- 1000 - TST
- 1001 - TEQ
- 1010 - CMP
- 1011 - CMN
- 1100 - ORR
- 1101 - MOV
- 1110 - BIC
- 1111 - MVN

We haven't discussed the "S" field yet. If set, it tells the processor to retain some "state" after the instruction has executed.

This "state" is in the form of 5-flags.



Many instructions (all we've seen thus far) have a special variant that sets the state flags. In these variants the opcode has an "S" appended.





# ARITHMETIC INSTRUCTIONS

ADD R3, R2, R12



$$R3 \leftarrow R2 + R12$$

Registers can contain either 32-bit unsigned values or 32-bit 2's-complement signed values.

SUB R0, R4, R6



$$R0 \leftarrow R4 - R6$$

Once more, either 32-bit unsigned values or 32-bit 2's-complement signed values.

RSB R0, R4, R2



$$R0 \leftarrow -R4 + R2$$

The operands of the subtraction are in reversed order. It is called "Reverse Subtract". Why? The I-type version makes more sense.

ADC R1, R5, R8



$$R1 \leftarrow R5 + R8 + C$$

Where 'C' is the Carry-out from some earlier instruction (usually an ADDS or ADCS) as saved in the Program Status Register (PSR)

SBC R2, R5, R7



$$R2 \leftarrow R5 - R7 - 1 + C$$

Where 'C' is the Carry-out from some earlier instruction (usually a SUBS or SUBCS) as saved in the PSR

RSC R1, R5, R3



$$R1 \leftarrow -R5 + R3 - 1 + C$$

"Reverse Subtract" with a Carry. Usually a carry generated from a previous RSBS or RSCS instruction.

A byte-sized example:

411 =	00000001	10011011	+	00000001	10011011
-42 =	00000000	00101010	+	11111111	11010101
				1 00000001	C=1 01110001

1=1-1+C      1

= 256 + 113 = 369



# LOGIC INSTRUCTIONS

Logical operations on words operate "bitwise", that is they are applied to corresponding bits of both source operands.

**AND R0, R1, R2**

**ORR R0, R1, R2**

**EOR R0, R1, R2**

Commonly called "exclusive-or"

**BIC R0, R1, R2**

Called "Bit-clear"  
 $R0 \leftarrow R1 \& \sim(R2)$

R1:	0000 0000 0000 0000 1111 1111 0000 0000
R2:	0000 0000 0000 0000 1111 0000 1111 0000
R0:	0000 0000 0000 0000 1111 0000 0000 0000
R0:	0000 0000 0000 0000 1111 1111 1111 0000
R0:	0000 0000 0000 0000 0000 1111 1111 0000
R0:	0000 0000 0000 0000 0000 1111 0000 0000



# STATUS FLAGS

Now it is time to discuss what status flags are available. These five status flags are kept in a special register called the Program Status Register (PSR). The PSR also contains other important bits that control the processor.

- **N** - set if the result of an operation is negative (Most Significant Bit (MSB) is a 1)
- **Z** - set if the result of an operation is "0"
- **C** - set if the result of an operation has a carry out of its MSB
- **V** - set if a sum of two positive operands gives a negative result, or if the sum of two negative operands gives a positive result
- **Q** - a sticky version of overflow created by instructions that generate multiple results (more on this later on).



# COMPARISON INSTRUCTIONS

These instructions modify the status flags, but leave the contents of the registers unchanged. They are used to test register contents, and they **must** have their "S" bit set to "1". They also don't modify their Rd, and by **convention**, Rd is set to "0000".

	4	3	4	1	4	4	8	4
R type:	1110	000	Opcode	1	Rn	0000	00000000	Rm
I type:	1110	001	Opcode	1	Rn	0000	Rotate	Imm8

CMP R0, R1



PSR flags set for the result R2 - R3

1000 - TST  
 1001 - TEQ  
 1010 - CMP  
 1011 - CMN

CMN R2, R3



PSR flags set for the result R2 + R3

TST R4, #8



PSR flags set for the result R4 & 8

TEQ R5, #102

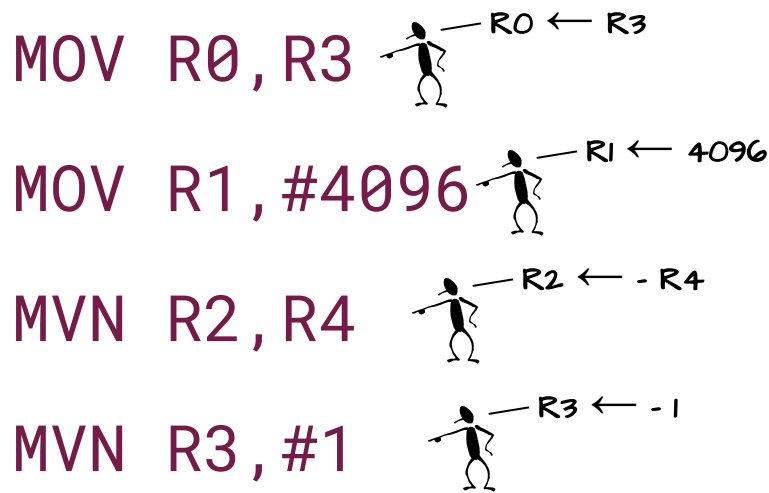


PSR flags set for the result R5 ^ 1024




# REGISTER TRANSFER

These instructions are used to transfer the contents of one register to another, or simply to initialize the contents of a register. They make use of only one operand, and, by convention, have their Rn field set to "0000".



	4	3	4	1	4	4	8	4
R type:	1110	000	Opcode	S	0000	Rd	00000000	Rm
I type:	1110	001	Opcode	S	0000	Rd	Rotate	Imm8

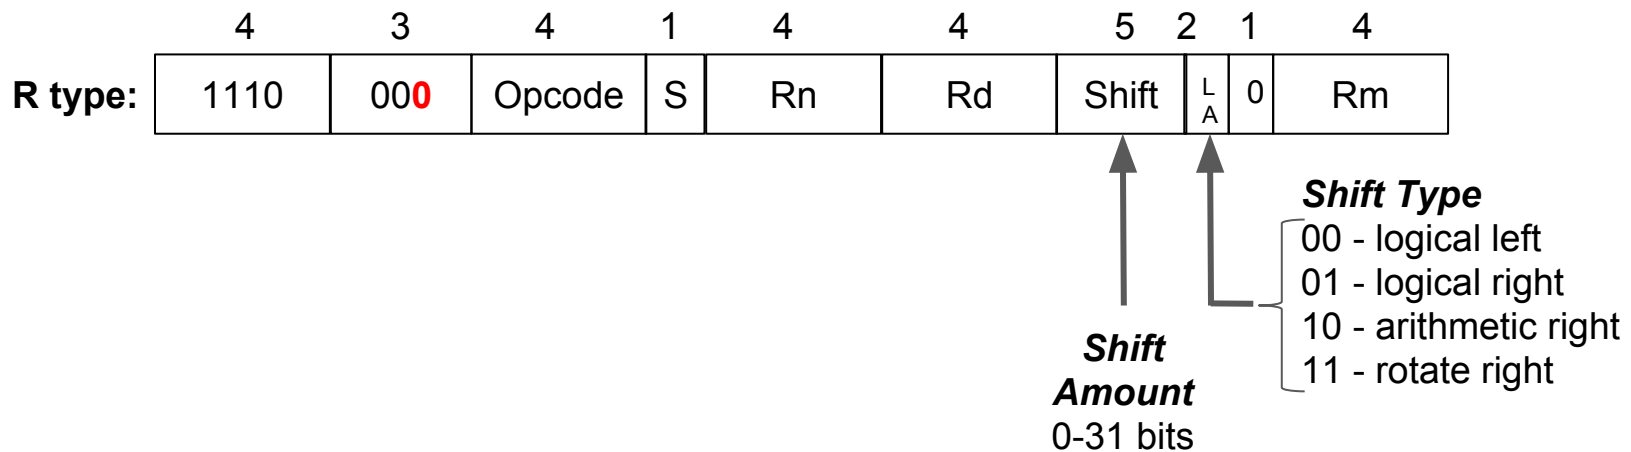


 {  
 1101 - MOV  
 1111 - MVN
 }



# ARM SHIFT OPERATIONS

A novel feature of ARM is that *all* data-processing instructions can include an optional "shift", whereas most other architectures have separate shift instructions. This is actually very useful as we will see later on. The key to shifting is that 8-bit field between Rd and Rm.





# LEFT SHIFTS

Left shifts effectively multiply the contents of a register by  $2^s$  where  $s$  is the shift amount.

**MOV R0, R0, LSL #7**

<b>R0 before:</b>	0000 0000 0000 0000 0000 0000 0000 0111	= 7
<b>R1 after:</b>	0000 0000 0000 0000 0000 0011 1000 0000	= $7 * 2^7 = 896$

A red arrow points from the right side of the R0 register to the left side of the R1 register, indicating a left shift of 7 positions.

Shifts can also be applied to the second operand of any data processing instruction

**ADD R1, R1, R0, LSL #7**



# RIGHT SHIFTS

Right Shifts behave like *dividing* the contents of a register by  $2^s$  where  $s$  is the shift amount, *if* you assume the contents of the register are *unsigned*.

MOV R0, R0, LSR 2

R0 before:	<table border="1"><tr><td>0000</td><td>0000</td><td>0000</td><td>0000</td><td>0000</td><td>0100</td><td>0000</td><td>0000</td></tr></table>	0000	0000	0000	0000	0000	0100	0000	0000	= 1024
0000	0000	0000	0000	0000	0100	0000	0000			
R1 after:	<table border="1"><tr><td>0000</td><td>0000</td><td>0000</td><td>0000</td><td>0000</td><td>0001</td><td>0000</td><td>0000</td></tr></table>	0000	0000	0000	0000	0000	0001	0000	0000	= $1024 / 2^2 = 256$
0000	0000	0000	0000	0000	0001	0000	0000			





# ARITHMETIC RIGHT SHIFTS

Arithmetic right shifts behave like *dividing* the contents of a register by  $2^s$  where  $s$  is the shift amount, *if* you assume the contents of the register are *signed*.

MOV R0, R0, ASR #2

R0 before:	<table border="1"><tr><td>1111</td><td>1111</td><td>1111</td><td>1111</td><td>1111</td><td>1100</td><td>0000</td><td>0000</td></tr></table>	1111	1111	1111	1111	1111	1100	0000	0000	= -1024
1111	1111	1111	1111	1111	1100	0000	0000			
R1 after:	<table border="1"><tr><td>1111</td><td>1111</td><td>1111</td><td>1111</td><td>1111</td><td>1111</td><td>0000</td><td>0000</td></tr></table>	1111	1111	1111	1111	1111	1111	0000	0000	= $-1024 / 2^2 = -256$
1111	1111	1111	1111	1111	1111	0000	0000			

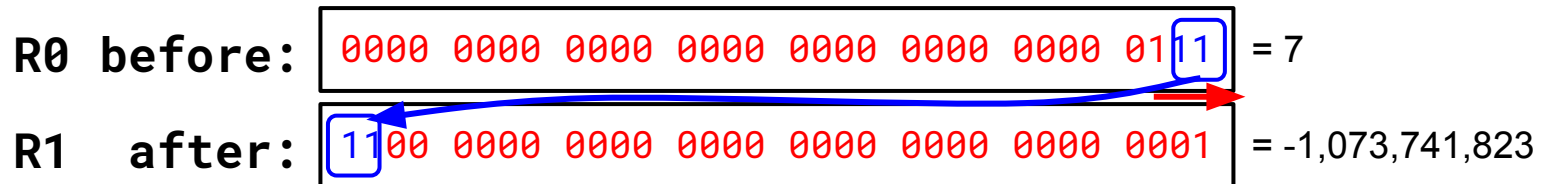
*Note: A red arrow points from the 6th bit of the R0 register to the 6th bit of the R1 register, indicating a right shift of 2 positions.*



# ROTATE RIGHT SHIFTS

Rotating shifts have no arithmetic analogy. However, they don't lose bits like both logical and arithmetic shifts. We saw rotate right shift used for the I-type "immediate" value earlier.

**MOV R0, R0, ROR #2**



Why no rotate left shift?

- Ran out of encodings?
- Almost anything Rotate lefts can do ROR can do as well!

# NEXT TIME



Instructions still missing

- Access to memory
- Branches and Calls
- Control
- Multiplication?
- Division?
- Floating point?

