

COMP1521 26T1

Week 4 Lecture 2

Bitwise Operators

Announcements

- **Week 3 Test Due Tomorrow:** Thursday 9pm
 - Many people have not done it yet!
 - Topic: MIPS Basics and MIPS Control
- **Week 4 Test Out Tomorrow**
 - Topic: MIPS Basics, Control and 1D Arrays
- **Census Date** : Tomorrow Thursday **12th March**
- **Assignment 1 Due:** Week 5 Friday (next week) at 6pm
- **See Help Sessions Schedule**
 - It is starting to get busy so get in now...

Plagiarism

Get help from the right places

- Staff in lectures, tuts, labs
- Forum, Help sessions
- Do not get 'help' or submit code from external sources like:
 - ChatGPT, external tutors, other people's code etc
- We run plagiarism checking on all submissions

**POV: YOU GET 90% OUT OF 100%
AT THE UNIVERSITY**

**BUT IT'S YOUR PLAGIARISM
SIMILARITY REPORT**



student.unsw.edu.au/plagiarism

Today's Lecture

- Integers Recap Exercises
- End of last lecture
 - Loading in MIPS
- Bitwise Operators

I'll drink $5 \wedge 6$
beers today

Software developers:



Mathematicians:



Recap: Bit and Bytes

What does this represent in C?

10110110111110001110110101110110

Recap: Bit and Bytes

What does this represent in C?

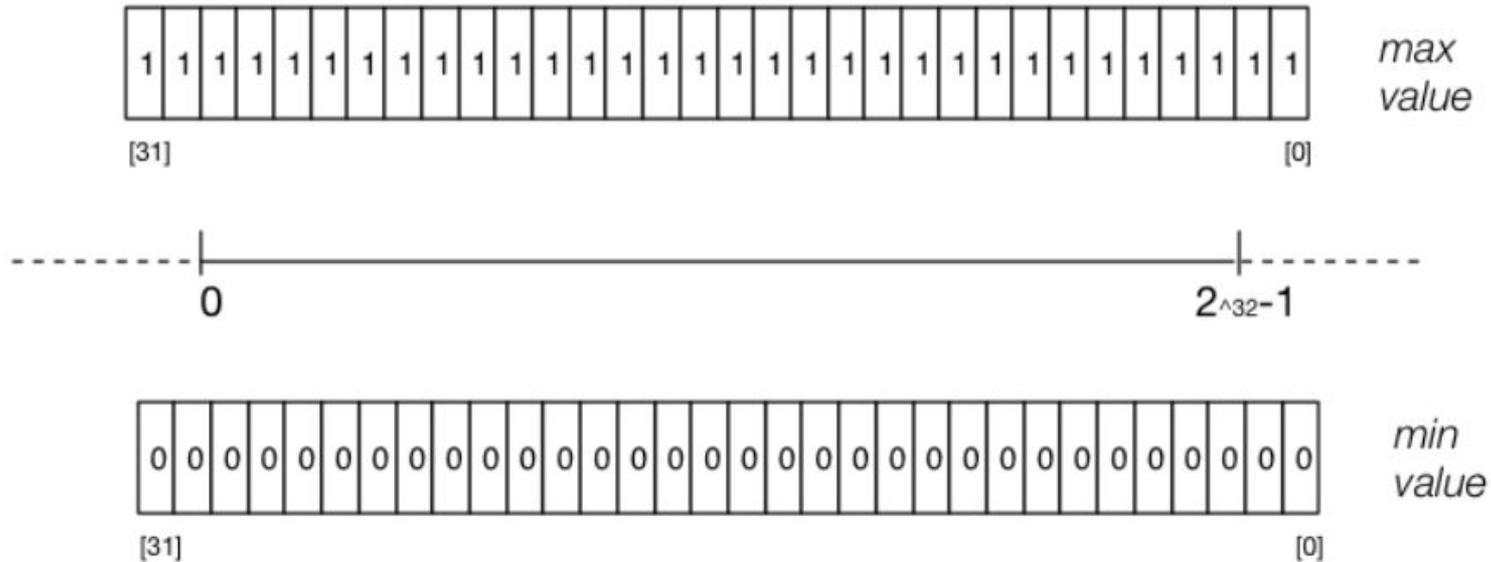
10110110111110001110110101110110

We can't know without knowing its type!

Is it: int, unsigned int, float, unicode character, MIPS instruction?

Unsigned integers

- In C the `unsigned int` data type is 4 bytes on our system
 - means we can store values from the range $0 \dots 2^{32}-1$

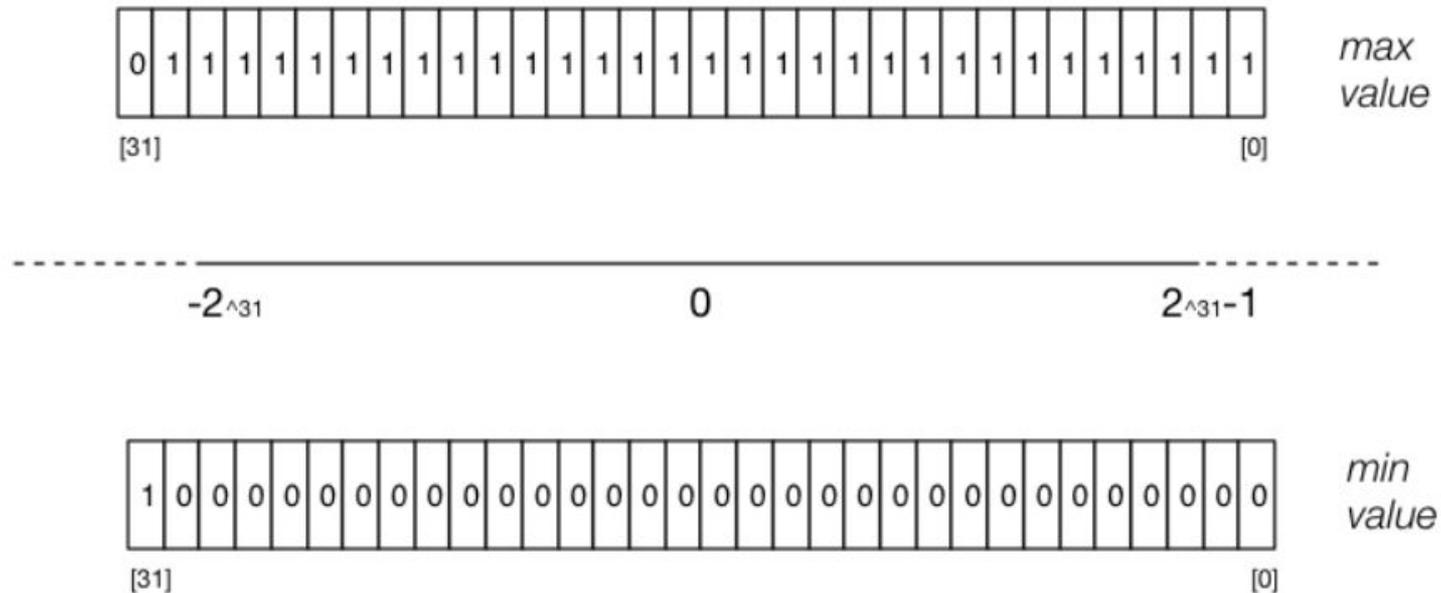


What do signed binary numbers look like?

- Modern computers use **two's complement** for integers
- Positive integers and zero represented as normal
- Negative integers represented in a way to make maths ✨easy✨ for the computer (not humans)
 - For an n -bit binary number, the number $-b$ is $2^n - b$
 - E.g. 8-bit number “-5” is represented as $2^8 - 5 = 251 = 1111\ 1011_2$
 - Conversely, $1111\ 1011_2$ represents “-5” because $251 - 2^8 = -5$.

Signed integers

- In C the `int` data type is 4 bytes on our system
 - we can store values from the range $-2^{31}.. 2^{31} - 1$



Signed vs Unsigned types

What do these bit patterns represent:

01000011

11000011

in decimal if the type is:

- `uint8_t`
- `int8_t`

Signed vs Unsigned types

What do these bit patterns represent:

01000011

11000011

in decimal if the type is:

- `uint8_t`
 - $01000011 = 2^6 + 2^1 + 2^0 = 64 + 2 + 1 = 67$
- `int8_t`

Signed vs Unsigned types

What do these bit patterns represent:

01000011

11000011

in decimal if the type is:

- `uint8_t`
 - $01000011 = 2^6 + 2^1 + 2^0 = 64 + 2 + 1 = 67$
 - $11000011 = 2^7 + 2^6 + 2^1 + 2^0 = 128 + 64 + 2 + 1 = \mathbf{195}$
- `int8_t`

Signed vs Unsigned types

What do these bit patterns represent:

01000011

11000011

in decimal if the type is:

- `uint8_t`
 - $01000011 = 2^6 + 2^1 + 2^0 = 64 + 2 + 1 = 67$
 - $11000011 = 2^7 + 2^6 + 2^1 + 2^0 = 128 + 64 + 2 + 1 = \mathbf{195}$
- `int8_t`
 - $01000011 = 2^6 + 2^1 + 2^0 = 64 + 2 + 1 = 67$

Signed vs Unsigned types

What do these bit patterns represent:

01000011

11000011

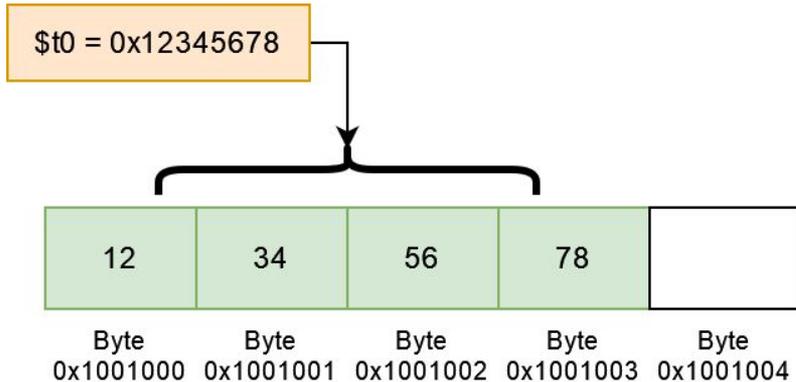
in decimal if the type is:

- `uint8_t`
 - $01000011 = 2^6 + 2^1 + 2^0 = 64 + 2 + 1 = 67$
 - $11000011 = 2^7 + 2^6 + 2^1 + 2^0 = 128 + 64 + 2 + 1 = \mathbf{195}$
- `int8_t`
 - $01000011 = 2^6 + 2^1 + 2^0 = 64 + 2 + 1 = 67$
 - $11000011 = \mathbf{195} - 2^8 = 195 - 256 = -61$

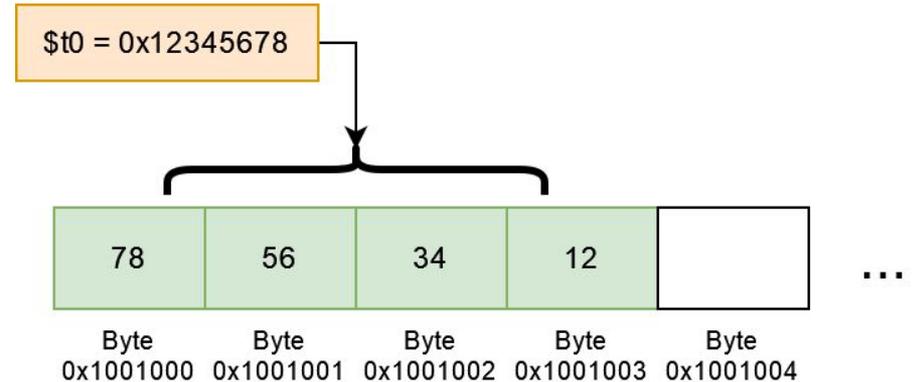
New concept: Endian-ness

- “What order to put things in” is a hard question to answer
- Two schools of thought:
 - **Big-endian**: MSB at the “low address” - big bits “first!”
 - **Little-endian**: LSB at the “low address” - little bits “first!”

BIG:



LITTLE:



Code example

- Mipsy-web is **little-endian**

```
.text
```

```
main:
```

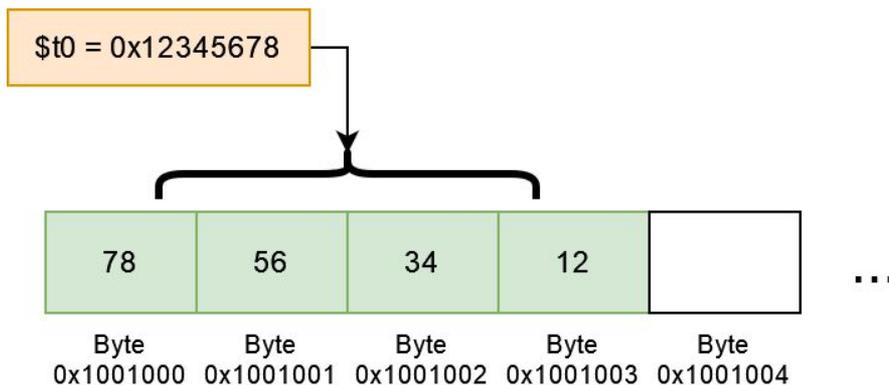
```
li $t0, 0x12345678
```

```
sw $t0, my_word
```

```
.data
```

```
my_word:
```

```
.space 4
```



Loading bytes, half-words

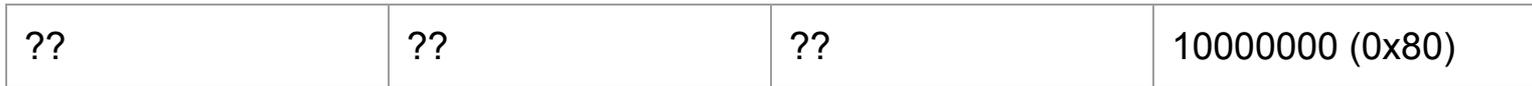
The results of these will depend on endianness:

- `lh/lb` assume the loaded byte/halfword is signed
- `lhu/lbu` for doing the same thing, but unsigned

- What might we need to do differently if we are loading a signed vs unsigned number?
- Why did we not need 2 versions for `lw`?

Loading bytes, half-words

- Consider storing 0x80 in a 4 byte (32 bit register) register
- This only takes up 1 byte (8 bits)
- What do we store in the rest of the register?



Loading bytes, half-words

- Consider storing 0x80 in a 4 byte (32 bit register) register
- This only takes up 1 byte (8 bits)
- What do we store in the rest of the register?

00000000 (0x00)	00000000 (0x00)	00000000 (0x00)	10000000 (0x80)
-----------------	-----------------	-----------------	-----------------

If we set them all the 0s, what happens if 0x80 is meant to be treated as a signed value?

Loading bytes, half-words

- Consider storing 0x80 in a 4 byte (32 bit register) register
- This only takes up 1 byte (8 bits)
- What do we store in the rest of the register?

11111111 (0xFF)	11111111 (0xFF)	11111111 (0xFF)	10000000 (0x80)
-----------------	-----------------	-----------------	-----------------

If our value is meant to be signed we extend the sign bit. For a negative value this means extending with 1

Loading Examples: lb

```
.text
```

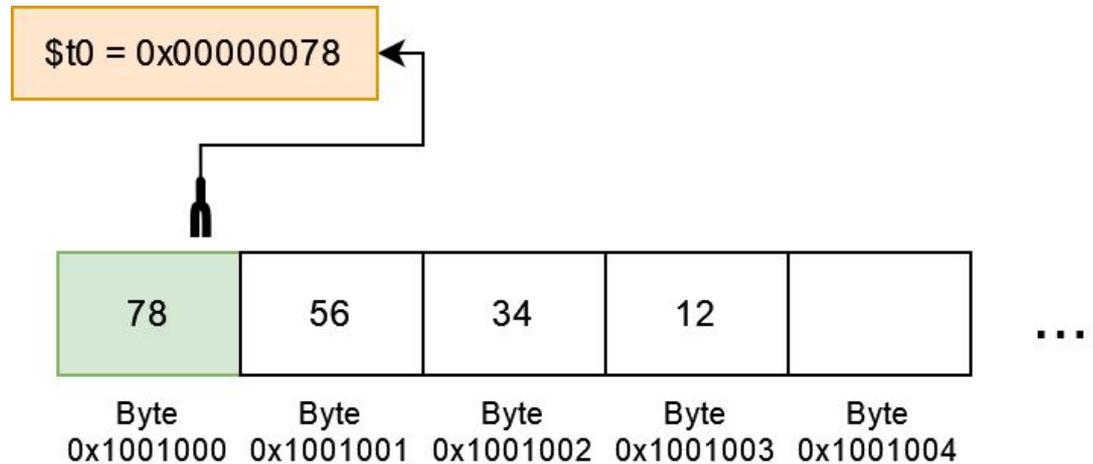
```
main:
```

```
    lb $t0, my_label
```

```
.data
```

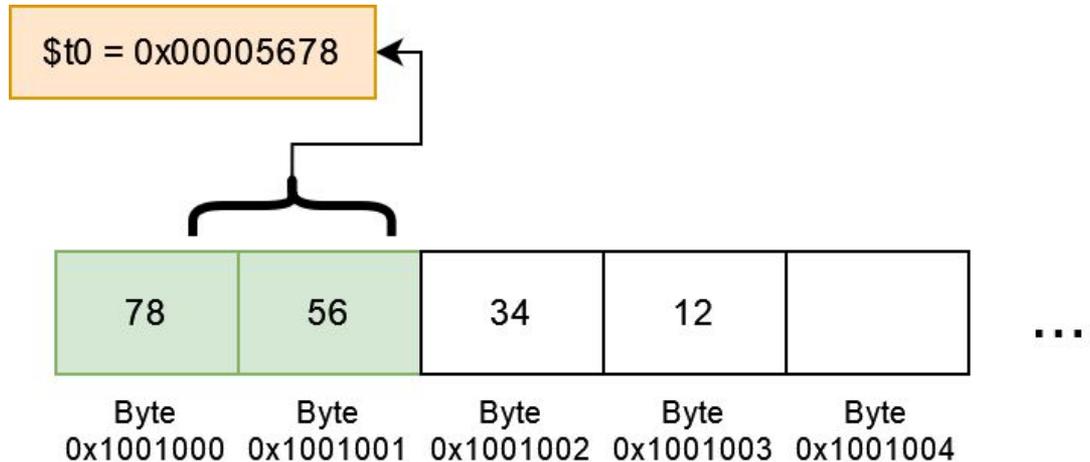
```
my_label:
```

```
    .word 0x12345678
```



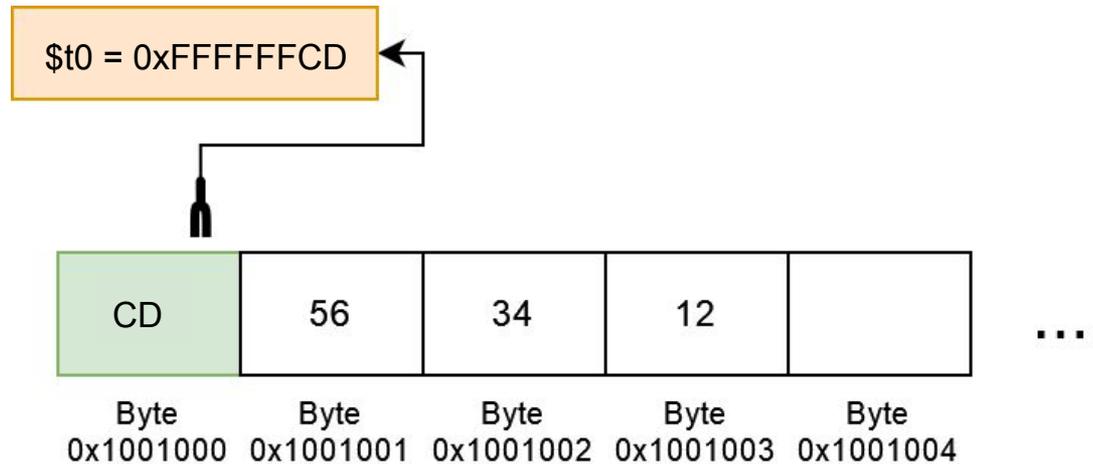
Loading Examples: lh

```
.text  
main:  
    lh $t0, my_label  
.data  
my_label:  
    .word 0x12345678
```



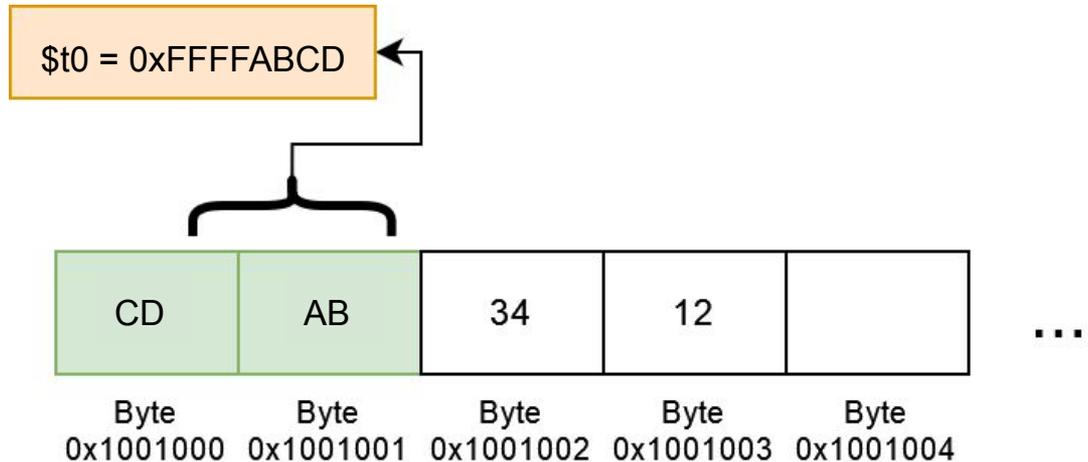
Loading Examples Negative: lb

```
.text  
main:  
    lb $t0, my_label  
  
.data  
my_label:  
    .word 0x1234ABCD
```



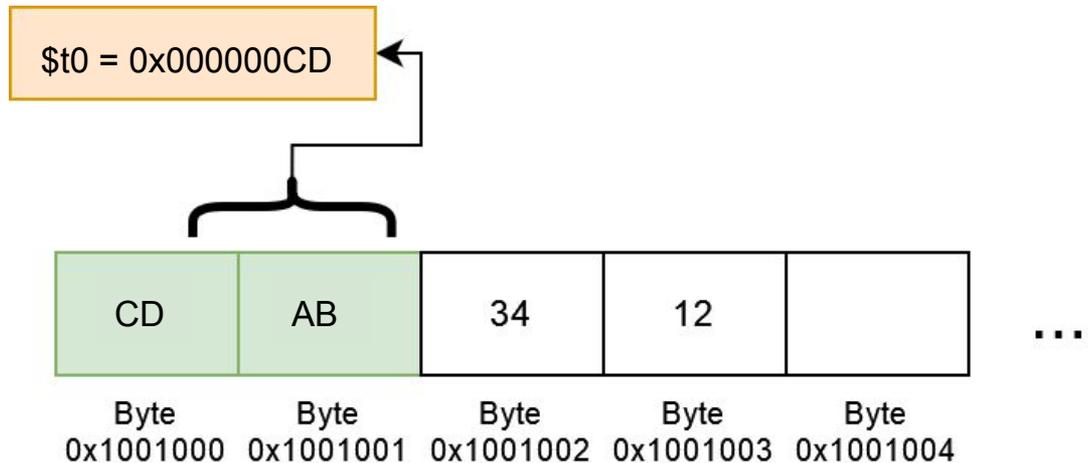
Loading Examples Negative: lh

```
.text  
main:  
    lh $t0, my_label  
.data  
my_label:  
    .word 0x1234ABCD
```



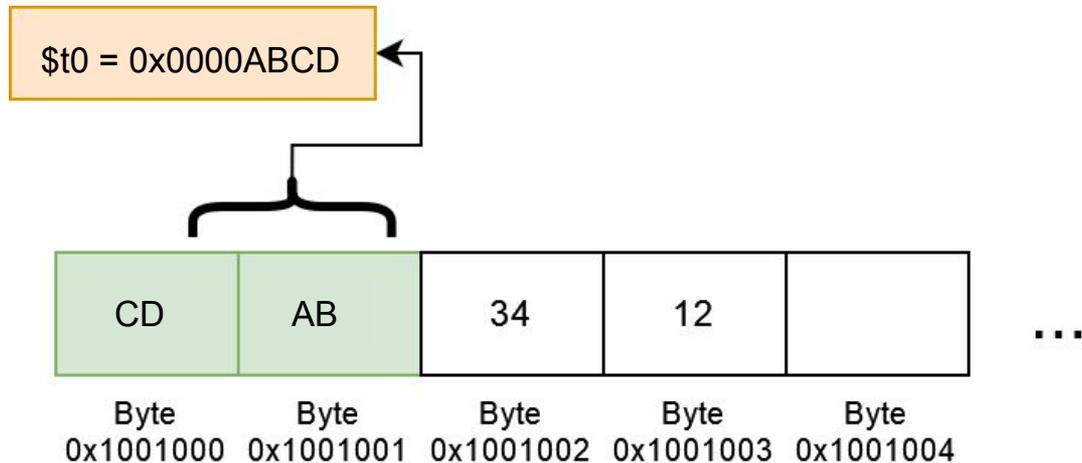
Loading Examples: lbu

```
.text  
main:  
    lbu $t0, my_label  
.data  
my_label:  
    .word 0x1234ABCD
```



Loading Examples Negative: lhu

```
.text  
main:  
    lhu $t0, my_label  
  
.data  
my_label:  
    .word 0x1234ABCD
```



What MIPS instruction is this?

0x01288820 =

What MIPS instruction is this?

0x01288820 =

0000 0001 0010 1000 1000 1000 0010 0000

This is the bit pattern, but we need to know more about the MIPS instruction formats to work out what it represents.

What do MIPS instructions look like?

- 32 bits long
- Specify:
 - An operation
 - (The thing to do)
 - 0 or more operands
 - (The thing to do it over)
- For example:



R-type



I-type



J-type

0010000100001001000000000000001100

addi \$t1, \$t0, 12

What MIPS instruction is this?

0x01288820 =

000000 01001 01000 10001 00000100000

What MIPS instruction is this?

0x0128820 =

000000 01001 01000 10001 00000100000

add

What MIPS instruction is this?

0x0128820 =

000000 01001 01000 10001 00000100000

add \$17

What MIPS instruction is this?

0x0128820 =

000000 01001 01000 10001 00000100000

add \$17, \$9

What MIPS instruction is this?

0x01288820 =

000000 01001 01000 10001 00000100000

add \$17, \$9, \$8

What MIPS instruction is this?

0x0128820 =

000000 01001 01000 10001 00000100000

add \$17, \$9, \$8

add \$s1, \$t1, \$t0

Let's type it into mipsy web to check!

Bitwise Operators

Why Learn Bitwise Operators

Used extensively in this course and also:

- Optimisation
- Embedded Systems
- Data compression
- Security and Cryptography
- Graphics
- Computer Networks

Bitwise Operations

- CPUs provide instructions which implement bitwise operations
 - Provide us ways to manipulating the individual bits of a value.
 - MIPS provides 13 bit manipulation instructions
 - C provides 6 bitwise operators

& bitwise AND

| bitwise OR

^ bitwise XOR (eXclusive OR)

~ bitwise NOT

<< left shift

>> right shift

Logical AND(&&) vs Bitwise AND (&)

- && works on whole values
 - We usually use it in conditions like:
 - `if (x > 10 && x < 20)`
- & works on every individual bit in each value
 - We use it to modify and/or extract bit information from values

Bitwise AND (&)

- takes two values (eg. a & b) and performs a logical AND between pairs of corresponding bits
 - resulting bits are set to 1 if **both** the original bits in that column are 1

Example:

	128	64	32	16	8	4	2	1	
	0	0	1	0	0	1	1	1	&
&	1	1	1	0	0	0	1	1	0
	0	0	1	0	0	0	1	1	1

Used for eg. checking if a particular bit is set (that is, set to 1)

Exercise: &

For any given bit value, x what is:

$$x \& 0 = ?$$

$$x \& 1 = ?$$

Exercise: &

For any given bit value, x what is:

$$x \& 0 = 0$$

$$x \& 1 = x$$

Bit Masks

We can create bit patterns to help us isolate the bits we are interested in! We call these masks!

For example:

```
int8_t x = 0x13;           //00010011
int8_t mask = 0x7;        //00000111 &
int8_t result = x & mask;
```

Bit Masks

We can create bit patterns to help us isolate the bits we are interested in! We call these masks!

For example:

```
int8_t x = 0x13;           //00010011
int8_t mask = 0x7;        //00000111 &
int8_t result = x & mask;
```

Bit Masks

We can create bit patterns to help us isolate the bits we are interested in! We call these masks!

For example:

```
int8_t x = 0x13;           //00010011
int8_t mask = 0x7;        //00000111 &
int8_t result = x & mask; //00000011
```

bit_ops_and.c

Checking if a number is odd

The obvious way to check if a number is odd in C:

```
int is_odd(int n) {  
    return n % 2 != 0;  
}
```

Checking if a number is odd

Decimal	Binary
0	000
1	001
2	010
3	011
4	100
5	101
6	110
7	111

What pattern do you see in the binary representation of odd numbers?

Checking if a number is odd

Decimal	Binary
0	000
1	00 1
2	010
3	01 1
4	100
5	10 1
6	110
7	11 1

What pattern do you see in the binary representation of odd numbers?

They all have a 1 as the least significant bit.

We can check that bit to see if it is 1.
If it is it is odd!

Checking if a number is odd

An odd value must have a 1 bit in the 1s place:

128	64	32	16	8	4	2	1
0	0	1	0	0	1	1	1

We can use bitwise AND to check if the last bit is set .

Checking if a number is odd

```
int is_odd(int n) {  
    return n & 1;  
}
```

If the value is **ODD** (eg 39):

	128	64	32	16	8	4	2	1
	0	0	1	0	0	1	1	1
&	0	0	0	0	0	0	0	1
<hr/>								
	0	0	0	0	0	0	0	1

If the value is **EVEN** (eg 38):

	128	64	32	16	8	4	2	1
	0	0	1	0	0	1	1	0
&	0	0	0	0	0	0	0	1
<hr/>								
	0	0	0	0	0	0	0	0

Bitwise OR (|)

- takes two values (eg. $a \mid b$) and performs a logical OR between pairs of corresponding bits
 - resulting bits are set to 1 if **at least** one of the original bits are 1

Example:

	0	0	1	0	0	1	1	1
	1	1	1	0	0	0	1	1
<hr/>								
	1	1	1	0	0	1	1	1

	0	1
0	0	1
1	1	1

Used for eg. setting particular bits (ie set to 1)

Bit Masks with |

For any given bit value, x what is:

$$x | 0 = ?$$

$$x | 1 = ?$$

Bit Masks with |

For any given bit value, x what is:

$$x | 0 = x$$

$$x | 1 = 1$$

Bit Masks with |

For any given bit value, x what is:

$$x | 0 = x$$

$$x | 1 = 1$$

For example:

```
uint8_t x = 0x13;           //00010011
```

```
uint8_t mask = 0x7;        //00000111
```

```
uint8_t result = x | mask; //00010111
```

Bitwise Negation (\sim)

- takes a single value (eg. $\sim a$) and performs a logical negation on each bit

Example:

\sim	0	0	0	1	0	1	1	0
	1	1	1	0	1	0	0	1

\sim	0	1
	1	0

Note: This does NOT mean making a number negative!

Bit Masks with ~

We can have a mask we can use to both **set** bits and **unset** bits

- Example:
 - mask 0x7 with | to **set** the least significant 3 bits
 - negate that mask and use it with & to **unset** the least significant 3 bits

Bit Masks with ~

We can have a mask we can use to both **set** bits and **unset** bits

- Example:
 - mask 0x7 with | to set the least significant 3 bits
 - negate that mask and use it with & to unset the least significant 3 bits

For example:

```
uint8_t x = 0x13;           //00010011
```

```
uint8_t mask = ~0x7;       //11111000
```

```
uint8_t result = x & mask; //00010000
```

Bitwise XOR (^)

- takes two values (eg. $a \wedge b$) and performs an exclusive OR between pairs of corresponding bits
 - resulting bits are set to 1 if **exactly** one of the original bits are 1

Example:

	0	0	1	0	0	1	1	1
\wedge	1	1	1	0	0	0	1	1
	1	1	0	0	0	1	0	0

\wedge	0	1
0	0	1
1	1	0

Used for eg. cryptography, flipping a bit, checking for bits that don't match

Bit Masks with ^

For any given bit value, x what is:

$$x \wedge 0 = ?$$

$$x \wedge 1 = ?$$

Bit Masks with ^

For any given bit value, x what is:

$$x \wedge 0 = x$$

$$x \wedge 1 = \sim x \text{ (flips the bit)}$$

Bit Masks with ^

For any given bit value, x what is:

$$x \wedge 0 = x$$

$$x \wedge 1 = \sim x \text{ (flips the bit)}$$

For example:

```
uint8_t x = 0x13;           //00010011
```

```
uint8_t mask = 0x7;        //00000111
```

```
uint8_t result = x ^ mask; //00010100
```

Bit Masks with ^

For any given bit value, x what is:

$$x \wedge 0 = x$$

$$x \wedge 1 = \sim x \text{ (flips the bit)}$$

For example:

```
uint8_t x = 0x13;           //00010011
```

```
uint8_t mask = 0x7;        //00000111
```

```
uint8_t result = x ^ mask; //00010100
```

What happens if I apply the mask again?

Exercise 1:

- Evaluate the following:
 - $5 \ \&\& \ 6$
 - $5 \ \& \ 6$
- How many beers did the software developer drink? i.e. What is 5^6
- Code Example: xor.c

I'll drink 5^6
beers today

Software developers:



Mathematicians:



Left Shift (<<)

- takes a value and a small positive integer x (eg. $a \ll x$)
- shifts each bit x positions to the left
 - any bits that fall off the left vanish
 - new 0 bits are inserted on the right
 - result contains the same number of bits as the input
- Example:

$$\begin{array}{cccccccc} 1 & 1 & 1 & 0 & 0 & 0 & 1 & 1 & \ll 2 \\ \hline 1 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \end{array}$$

Implications of left shift

What does this mean mathematically?

Implications of left shift

What does this mean mathematically?

Expression	Result Binary	Result Decimal
$00000001 \ll 1$	00000010	2
$00000001 \ll 2$	00000100	4
$00000001 \ll 3$	00001000	8
$00000001 \ll 4$	00010000	16

Implications of left shift

What does this mean mathematically? Multiplies by powers of 2!

Expression	Result Binary	Result Decimal
<code>00000001 << 1</code>	<code>00000010</code>	2
<code>00000001 << 2</code>	<code>00000100</code>	4
<code>00000001 << 3</code>	<code>00001000</code>	8
<code>00000001 << 4</code>	<code>00010000</code>	16

Demo: `shift_as_multiply.c`

Right Shift (>>)

- takes a value and a small positive integer x (eg. $a \gg x$)
- shifts each bit x positions to the right
 - any bits that fall off the right vanish
 - new 0 bits are inserted on the left (for unsigned types)
 - result contains the same number of bits as the input
- Example:

1	1	1	0	0	0	1	1	>>	2
0	0	1	1	1	0	0	0		

Used for eg looping through 1 bit at a time

Implications of right shift for unsigned values

What does this mean mathematically?

Implications of right shift

What does this mean mathematically? $16_{10} == 00010000_2$

Expression	Result Binary	Result Decimal
<code>00010000 >> 1</code>	<code>00001000</code>	8
<code>00010000 >> 2</code>	<code>00000100</code>	4
<code>00010000 >> 3</code>	<code>00000010</code>	2
<code>00010000 >> 4</code>	<code>00000001</code>	1

Implications of right shift for unsigned values

What does this mean mathematically? $16_{10} == 00010000_2$

Expression	Result Binary	Result Decimal
<code>00010000 >> 1</code>	<code>00001000</code>	8
<code>00010000 >> 2</code>	<code>00000100</code>	4
<code>00010000 >> 3</code>	<code>00000010</code>	2
<code>00010000 >> 4</code>	<code>00000001</code>	1

Divides by powers of 2

Implications of right shift for unsigned values

But what about situations like this? We lose some bits!

$0111 \gg 1 == 0011$

This is the same as $7/2 == 3$ with integer division!

Issues with shifting

- Shifts involving negative values may not be portable, and can vary across different implementations
- Common source of bugs in COMP1521 (and elsewhere)
- Always use unsigned values/variables when shifting to be safe/portable

Demo: `shift_bug.c`

Exercise

Given the following declarations:

```
// a signed 8-bit value
uint8_t x = 0x55;
uint8_t y = 0xAA;
```

What is the value of each of these expressions?

```
uint8_t a = x & y;
uint8_t b = x ^ y;
uint8_t c = x << 1;
uint8_t d = y << 2;
```

```
uint8_t e = x >> 1;
uint8_t f = y >> 2;
uint8_t g = x | y;
```

Code Demo

- `get_nth_bit.c`

MIPS - Bit manipulation instructions

assembly	meaning	bit pattern
and r_d, r_s, r_t	$r_d = r_s \& r_t$	000000sssstttttddddd00000100100
or r_d, r_s, r_t	$r_d = r_s \mid r_t$	000000sssstttttddddd00000100101
xor r_d, r_s, r_t	$r_d = r_s \wedge r_t$	000000sssstttttddddd00000100110
nor r_d, r_s, r_t	$r_d = \sim(r_s \mid r_t)$	000000sssstttttddddd00000100111
andi r_t, r_s, I	$r_t = r_s \& I$	001100sssstttttIIIIIIIIIIIIIIIIIIII
ori r_t, r_s, I	$r_t = r_s \mid I$	001101sssstttttIIIIIIIIIIIIIIIIIIII
xori r_t, r_s, I	$r_t = r_s \wedge I$	001110sssstttttIIIIIIIIIIIIIIIIIIII
not r_d, r_s	$r_d = \sim r_s$	pseudo-instruction

MIPS - Shift instructions

assembly	meaning	bit pattern
sllv r_d, r_t, r_s	$r_d = r_t \ll r_s$	000000s s s s s t t t t t d d d d d 00000000100
srlv r_d, r_t, r_s	$r_d = r_t \gg r_s$	000000s s s s s t t t t t d d d d d 00000000110
sra r_d, r_t, r_s	$r_d = r_t \gg r_s$	000000s s s s s t t t t t d d d d d 00000000111
sll r_d, r_t, I	$r_d = r_t \ll I$	000000000000t t t t t d d d d d I I I I I 000000
srl r_d, r_t, I	$r_d = r_t \gg I$	000000000000t t t t t d d d d d I I I I I 000010
sra r_d, r_t, I	$r_d = r_t \gg I$	000000000000t t t t t d d d d d I I I I I 000011

- **srl** and **srlv** shift zeroes into most-significant bit
 - This matches shift in C of unsigned values
- **sra** and **srav** propagate most-significant bit
 - This ensures the sign is maintained

MIPS Code Demos

- `odd_even.s`
- `mips_bits.s`
- `mips_negative_shifts.s`

What did we learn today?

- Integer representation recap
- Bitwise Operators
- Next lecture:
 - Finish Bitwise
 - Floating Point Data

Feedback Please!

Your feedback is valuable!

If you have any feedback from today's lecture, please follow the link below or use the QR Code.

Please remember to keep your feedback constructive, so I can action it and improve your learning experience.



<https://forms.office.com/r/NKVTwTXixC>

Reach Out

Content Related Questions:

[Forum](#)

Admin related Questions email:

cs1521@cse.unsw.edu.au



Student Support | I Need Help With...

My Feelings and Mental Health

Managing Low Mood, Unusual Feelings & Depression



Mental Health Connect

student.unsw.edu.au/counselling
Telehealth



**In Australia Call Afterhours
UNSW Mental Health Support
Line**

1300 787 026
5pm-9am



Mind HUB

student.unsw.edu.au/mind-hub
Online Self-Help Resources



**Outside Australia
Afterhours 24-hour
Medibank Hotline**

+61 (2) 8905 0307

Uni and Life Pressures

Stress, Financial, Visas, Accommodation & More



**Student Support
Indigenous Student
Support**

student.unsw.edu.au/advisors

Reporting Sexual Assault/Harassment



**Equity Diversity and Inclusion
(EDI)**

edi.unsw.edu.au/sexual-misconduct

Educational Adjustments

To Manage my Studies and Disability / Health Condition



**Equitable Learning Service
(ELS)**

student.unsw.edu.au/els

Academic and Study Skills



**Academic Language
Skills**

student.unsw.edu.au/skills

Special Consideration

Because Life Impacts our Studies and Exams



Special Consideration

student.unsw.edu.au/special-consideration