

## ECE 2049 LECTURE 3

### TODAY

- MORE C EXAMPLES,
- MEMORY LAYOUT, BYTE ORDER

### OFFICE HOURS

- Today, Thurs: 4-6PM

### ADMINISTRIVIA

- HW1 - DUE TODAY ON CANVAS
- LAB 0: SIGNOFF DUE THURSDAY BY  
END OF OFFICE HOURS (6PM)
- HW2: ONLINE AFTER CLASS, DUE NEXT  
TUESDAY
- LAB 1: STARTS THURSDAY
- ED DISCUSSION: ASK QUESTIONS  
ONLINE HERE!

# ECE2049: Homework 1

---

**Material:** Lecture 1

**Due:** Start of Lecture 3: Tuesday, 24 May 2022 by 2pm EDT

**Submission notes:**

- For full credit, please show your work and denote your answers with a circle or a box.
  - Always write and draw your diagrams neatly! We cannot be expected to GUESS what you meant to write!
  - Please see the submission guidelines on the homework page of the course website for details.
1. (5 pts) Please do the following logistical tasks to help you get started with the course:
    - a. Register for the course discussion board (EdStem) using your WPI email address:  
<https://edstem.org/us/join/VBSzQc>
    - b. Complete the course background survey to provide some information about prior courses you have taken. This will help me calibrate course content to accommodate everyone. You can find the survey here: [https://wpi.qualtrics.com/jfe/form/SV\\_3ZV5flgTa8swdOm](https://wpi.qualtrics.com/jfe/form/SV_3ZV5flgTa8swdOm)
  2. (5 pts) You are given three **16-bit** values shown below. Each of these values can be interpreted as:
    - An unsigned number
    - A sign-magnitude number
    - A two's complement number

Provide the decimal (base 10) equivalent of each value for each of these interpretations. Show your work.

- a. 0x4048
- b. 0x448C
- c. 0xDEED

$\Rightarrow$  0x4048

2. (5 pts) A hardware device is responsible for reading the state of 8 relays that control a manufacturing process. The device represents the state of each relay ( $R_0$ — $R_7$ ) in an 8-bit value  $v$ , with the state of relay  $R_0$  is stored in the least significant bit, and the state of  $R_7$  in the most significant bit.

If the device returns the value  $v = 0x5B$ , which relays are on?

QW1

Part 1a

0x404F  $\Rightarrow$  0100 0000 0100 1000

- UNSIGNED  $2^{19} + 2^6 + 2^3 = \boxed{16456}$

- SIGN-MAG: SIGN BIT IS 0  $\Rightarrow +$

$$2^{19} + 2^6 + 2^3 = \boxed{16456}$$

- 2'S-COMP

$$2^{19} + 2^6 + 2^3 = \boxed{16456}$$

Part 1c

0xDEED  $\Rightarrow$  1101 1110 1110 1101

- SIGN-MAG: NUMBER IS NEGATIVE

- 2'S COMP: NUMBER IS NEGATIVE.

$\hookrightarrow$

$$\begin{array}{r} \underline{1101 \ 1110 \ 1110 \ 1101} \\ 0010 \ 0001 \ 0001 \ 0010 \text{ COMP} \\ + \phantom{0010 \ 0001 \ 0001 \ 0010} 1 \\ \hline \end{array}$$

$$\begin{array}{r} 0010 \ 0001 \ 0001 \ 0010 \\ - (2^{13} + 2^8 + 2^4 + 2^1 + 2^0) = \boxed{-8467} \end{array}$$

0x4048

- WRITE AS
  - UNSIGNED
  - SIGN-MAGNITUDE
  - 2'S COMP

FIRST: WRITE 0x4048 AS BINARY

0x4048  $\Rightarrow$   $\overset{15}{0100} \overset{14}{0000} \overset{13}{0100} \overset{12}{1000}$

UNSIGNED:  $2^{14} + 2^6 + 2^3 = \boxed{16456}$

SIGN-MAGNITUDE:  $\begin{matrix} \text{S} & \text{M} & \text{M} & \text{M} & \text{M} & \text{M} & \text{M} & \text{M} & \text{M} & \text{M} \\ \text{S} & \text{M} & \text{M} & \text{M} & \text{M} & \text{M} & \text{M} & \text{M} & \text{M} & \text{M} \end{matrix}$

- MSB IS SIGN BIT = 0 = POSITIVE

$$2^{14} + 2^6 + 2^3 = \boxed{16456}$$

2'S COMP

- POSITIVE, SO SAME.

0xDEC0

UNSIGNED: ✓

SIGN-MAG: ✓

2's COMP: SIGN BIT IS 1, SO  
NEGATIVE.

$$\begin{array}{r} \text{DEC0} \Rightarrow \begin{array}{cccc} 1101 & 1110 & 1110 & 1101 \\ \hline 0010 & 0001 & 0001 & 0010 \end{array} \quad \begin{array}{l} \text{BIN} \\ \text{COMP} \end{array} \\ + \quad \begin{array}{cccc} \hline 0010 & 0001 & 0001 & 0010 \\ \hline \end{array} \quad \begin{array}{l} \text{ADD } 1 \\ \hline \end{array} \end{array}$$

$$-(2^{13} + 2^6 + 2^4 + 2^1 + 2^0) = -[2467]$$

## Module 3. Of Integers and Endians & Floating Point Representations

### Topics

- Memory organization and endianness
- More data representations: overview of floating point

### Last Time

- C programming basics
- Data representations for characters

### Warmup: try the following...

```
int z = 0x4007;
```

```
// a. What is the size of z (in bytes)?
```

```
// b. In C, how is z stored (unsigned, sign-magnitude, 2's comp)?
```

```
if (z & 0x8000) {
    alpha();
} else {
    beta();
}
```

```
// c. Based on the value of z, which function would get called?
```

$\Rightarrow \text{INT} = 2 \text{ BYTES} = 16 \text{ BITS.}$

X	Y	X	Y
0	0	0	0
0	1	0	0
1	0	0	0
1	1	1	1

$z$  0100 0000 0000 0111

1000 0000 0000 0000

0000 0000 0000 0000  $\rightarrow \text{BETA()}'$

X	Y	X	Y
0	0	0	0
0	1	0	0
1	0	0	0
1	1	1	1

1100 0011 ✓

0000 0110 0+06

0000 0010 "MASKING"

LET'S TEST IF

A BIT OR BITS ARE SET

IS POSITIVE ( )

— TEST IF A NUMBER  $x$   
IS POSITIVE W/O  $<$   $>$   
 $\leq$   $\geq$

INT IS POSITIVE (INT  $x$ )?

// CHECK IF SIGN BIT IS SET,  
RETURN TRUE IF IT IS.

RETURN  $!(x \& 0x\text{f000})$

~

$x$  0 - - - - -  
1000 0000 0000 0000  
0x f000

Ex.  $-1 \Rightarrow$  1111 1111 1111 1111  
& 1000 0000 0000 0000  
-----  
1000 0000 0000 0000

RETURN  $(0x\text{ffff} \& 0x\text{f000}) \Rightarrow$  (TRUE)

$-1 \Rightarrow$  FALSE

$> 0 \Rightarrow$  TRUE

?

## Memory organization

INT a = 42;

What does it mean to type "int a" in C? This is called variable declaration, which allocates space in the program's memory to store an int.

What do we mean by memory? You can think of memory as a big table of "addresses" that each map to a certain piece of data. This data could be a variable (as above), or it could be a piece of code, a portion of the hardware, etc., but for now let's focus on variables.

On the MSP430, addresses are 16-bits long, and each address refers to one byte.

Recall that the MSP430 is a 16-bit architecture,

$(2^{16} \text{ POSSIBLE ADDRESSES})(1 \text{ BYTE})$   
 $= 65536 \text{ BYTES}$

$= 16 \text{ KiB}$   
 TOTAL MEMORY

$(1 \text{ KiB} = 1024 \text{ BYTES})$

x86: 32 BIT ADDRESSES  
 $2^{32} \approx 4 \text{ GiB}$

x86-64  
 $2^{64} \approx 16 \text{ EiB}$

ADDRESS	VALUES
0x0000	12h
0x0001	AAh
⋮	⋮
0xFFFFE	0xFF
0xFFFF	0x10

$\downarrow 2^{16} - 1$   
 16 BITS  $\Rightarrow$  1 BYTE (8 BITS)

Unfortunately, this is no longer completely true! Newer MSP430 variants (like ours MSP430F5529) utilize 20-bit addresses. Why?

$2^{20} = 1 \text{ MiB}$   
 (MSP430X)

$\rightarrow$  EXTEND ADDRESS SPACE  
 w/ HARDWARE CHANGE.

(WE WON'T DEAL w/  
 IT MUCH)



## Laying out variables in memory

When you declare variables in your program, they are arranged in memory starting at a certain address. For now, it is sufficient to know that variables in `main` start at address `0x4400`. We will discuss why in an upcoming lecture.

When variables are declared, they are (usually) arranged in order from this starting address.

For example:

*MAIN()*

```
char a = 0x11;
char b = 0x22;
```

...can be arranged in memory as follows:


*BY CONVENTION  
START FROM BOTTOM,  
GROW UP*

Address	Data	Variable
<i>0x4401</i>	<i>22h</i>	<i>B</i>
<i>0x4400</i>	<i>11h</i>	<i>A</i>

In our class, we will arrange memory in a table like the one above, with the starting address at the bottom. We use this convention because we are typically representing variables on the program stack, which starts at a fixed base address and grows up.

*MAIN ( ) {  
 CHAR A, B;  
 F();  
 F();  
}*

*F ( ) {  
 CHAR C;  
 G();  
}*



## Endianness: Ordering bytes

In the previous example, we have left out an important detail. How do you store variables that are larger than a byte?

As declared on the MSP430, a long is has a size of four bytes:

```
long v = 0xAABBCCDD; // AAh is the most significant byte (MSB), and
                    // DDh is the least significant byte (LSB)
```

For multi-byte variables, we have a choice—do we arrange the data with the least significant byte first, or with the most significant byte first? Which is correct? Does it matter?

This concept is known as *endianness*, which governs how a processor orders bytes in memory. There are two forms of endianness:

### Little Endian (LE)

Little Endian stores the least significant byte first, meaning that the memory in this example would be arranged as follows:

$V = 0xAABBCCDD;$

↑ MSB                      ↑ LSB

Address	Data	Variable
0x4403	AA	v
0x4402	BB	
0x4401	CC	
0x4400	DDh	

← MSB (pointing to 0x4403)  
← LSB (pointing to 0x4400)

LE LOOKS 'OUT  
OF ORDER' WHEN  
WE READ LEFT → RIGHT

### Big Endian (BE)

Big Endian stores the *most significant byte first*, as follows:


Address	Data	Variable
0x4403	DD	v
0x4402	CC	
0x4401	BB	
0x4400	AAh	

← MSB (pointing to 0x4403)  
← LSB (pointing to 0x4400)


BE LOOKS 'IN ORDER'  
WHEN WE READ  
LEFT TO RIGHT.

Here is an example of being *Endian-ed*!

A nice plot of a file of unsigned integers as created on a little endian machine.



Below is a plot of the same data file having being read in as unsigned integers on a big endian machine. The data is good! It is the same as above! All that has changed is the endianness of the machine that read the data.



"GARBAGE"

The table below shows the first few unsigned integer values from the data file created on the little endian computer as read by both of the machines. The byte swap is evident in the hexadecimal values.

Read as Little Endian		Read as Big Endian	
8178	1FF2 h	61983	F21F h
8193	2001 h	288	0120 h
8194	2002 h	544	0220 h
8182	1FF6 h	63007	F61F h
8201	2009 h	2336	0920 h
8201	2009 h	2336	0920 h

ENDIAN-NESS is Part of Microprocessor Design!  
Not a function of the OS!!

## Important points on endianness

- Endianness is a fundamental part of the architecture's design. When a processor is designed, it is designed to use a specific byte order—you cannot change this with a compiler setting.
- Is one endianness better than the other? No, they simply reflect different design choices.
- Big endian is read "left to right", which is intuitively easier to read for those accustomed to languages written left to right
- Little endian makes it easier to slice out small portions of a variable (eg, what if you only want the first byte of a long?)

## When will you deal with endianness?

Endianness becomes especially important when you need to transfer data between different architectures. Examples include any stored data format or network protocol.

LE: MSP430, x86, M1 <sup>APPLE</sup>  
 BE: PowerPC, DSP CHIPS,  
 "NETWORK BYTE ORDER"  
 (INTERNET TRAFFIC)

ENDIANNESS MATTERS FOR ALL METHODS  
 BY WHICH DATA IS TRANSFERRED  
 BETWEEN SYSTEMS  
 (FILE FORMATS, NETWORK...)

## More memory layout: Arrays

### How do arrays work, anyway?

In C, we can declare arrays and use them as follows:

```
// Declare an array of 5 bytes
char arr[5];
// Declare an array of 5 bytes, and initialize it (set it with some initial values)
char arr[5] = { 0xAA, 0xBB, 0xCC, 0xDD, 0xEE };

// You can access elements of an array by "indexing" into it
// In C, array indexes start at 0
char c = arr[0]; // The first element
char d = arr[4]; // The last element (arr has size of 5, so last index is 5 - 1 = 4)
```

You can think of the elements of the array laid out like this:

Index	0	1	2	3	4
Element	<code>ARR[0]</code>	<code>ARR[1]</code>	<code>ARR[2]</code>	<code>ARR[3]</code>	<code>ARR[4]</code>
Value	AAh	BB	CC	DD	EE

0x400 0x401 0x402 0x403 0x404

Why is it important that array elements are contiguous? (And must contain elements of the same type?)

WANT TO MAKE IT "EASY" TO FIND THE  $i^{\text{TH}}$  ELEMENT

$$A_i = A_0 + (i * \text{SIZEOF}(X))$$

↑  
BASE ADDRESS

↑  
SIZE OF ONE ELEMENT  
(INT, CHAR, ETC)

"EASY" MEANS  
CONSTANT TIME  
(NO SEARCHING)  
 $O(1)$

What would happen if we tried to get the 6<sup>th</sup> element of arr?

→ `ARR[1000] = ?`

NO ERROR CHECKING!

↳ THIS WILL COMPILE, BUT PROGRAM WILL READ WHATEVER DATA IS IN THIS SPACE, BUT IT'S NOT PART OF ARR (NOT WHAT YOU WANT)

BUFFER OVERFLOW

CLASSIC SECURITY PROBLEM.

## How endianness affects arrays (or rather, how it does not)

A fundamental property of arrays is that their elements are stored contiguously in memory in order of their index (as discussed above).

**Endianness does not change the order of array elements.**

For example, if we laid out the array from the above example on a Big Endian (BE) and a Little Endian (LE) system, it would look like this:

Address	BE	LE
0x4404	EE	EE
0x4403	DD	DD
0x4402	CC	CC
0x4401	BB	BB
0x4400	AA	AA

ARR[4]

ARR[0]

However, endianness *does* affect the ordering of the bytes in each element of the array! In the previous example, the elements were just 1 byte each!

**Example:** an array of ints:

`int iarr[2] = {0x1122, 0x3344};`

ELEMENTS ARE TYPE INT

EACH ELEMENT IS 2 BYTES

Here, the memory would be organized as follows:

Address	BE	LE
0x4403	44	33
0x4402	33	44
0x4401	22	11
0x4400	11	22

1ARR[1] = 0x4402 →

1ARR[0] = 0x4400 →

1ARR[1]

1ARR[0]

DOES CHANGE

DOES NOT CHANGE  
w/ DIFF. BYTE  
ORDER

## Using addresses as data

We can also have variables that contain memory addresses. These are called pointers.

*You work w/ pointers all the time!*

You can get the address of a variable with the “address-of” operator (&):

```
long v = 0x11223344;
long *pv = &v;
```

In this example, we say that `pv` is declared as the type “pointer to long,” which is indicated by the “\*” before the name `pv`.

How big is `pv`?

*pv holds one memory address  
on MSP430  $\Rightarrow$  2 bytes (16 bits)*

What is the value of `pv`?

*— starting address of v*

*$PV = \&V = 0x4400$*

We can lay out these variables in memory as follows:

Address	BE	LE
0x4405	00	44
0x4404	44	00
0x4403	44	11
0x4402	23	22
0x4401	22	33
0x4400	11	44



*} PV*  
*} v*

*PV CONTAINS  
v's ADDRESS  
NOT ITS  
CONTENT.*

## How big is a pointer?

A pointer is the size of a memory address for a given architecture. On the MSP430, an address has a size of 2 bytes (16 bits).

Type	Size (bytes)
int	2
long	4
char	1
long long	8

Type	Size (bytes)
int *	2
long *	2
char *	2
long long *	2

ALL HAVE THE SAME SIZE!  
(ON SAME SYSTEM)

This is one way in which pointers are powerful: a pointer can represent a larger data structure in the program—by passing around the pointer, we can avoid copying or moving the larger data structure.

## How are pointers used with arrays?

Whenever you use arrays, you use pointers. Consider the following example:

```
int iarr[10];
int i = iarr[5];
```

When you index into an array, the program actually does the following:

```
int i = *(iarr + 5); // Equivalent to writing iarr[5]
```

Here, the `*` is the *dereference operator*, which **gets the value at the given address**. This is called *dereferencing* the pointer—it is the opposite of the *address-of* (`&`) operator.

GET THE DATA AT THIS MEMORY ADDRESS

$\text{IARR} = \text{STARTING ADDRESS OF} = \&\text{IARR}[0]$   
IARR



## Working with Pointers

**Pointer math:** When performing arithmetic operations on pointers, the address changes in increments based on the type of the pointer.

$$A_i = A_0 + i \cdot \text{sizeof}(A)$$

```
// Example 1: array of char
char carr[4];
// How big is the array?
// Say the starting address is 0x4400, what is the address of carr[3]?
```

$$(4 \text{ ELEMENTS})(1 \text{ BYTE/ELEMENT}) = 4 \text{ BYTES}$$

$$0x4400 + 3 * \text{sizeof}(\text{CHAR})$$

$$= 0x4403$$

```
// Example 1: array of int
int carr[4];
// How big is the array?
```

$$(4 \text{ ELEMENTS})(2 \text{ BYTES/INT}) = 8 \text{ BYTES}$$

```
// Say the starting address is 0x4400, what is the address of iarr[3]?
```

$$0x4400 + 3 * \text{sizeof}(\text{INT})$$

2

$$= 0x4400 + 6 = 0x4406$$

**Passing arrays:** Further, when you use the name of an array (either to store or pass to a function), you are *passing a pointer to the first element of the array*. This is the “starting point” of the array used as input to calculate the index.

```
int *ptr = iarr;    // Could also write &iarr[0]
do_thing(iarr, 10); // Same here

void do_thing(int* arr, int size) { // Function takes pointer to array (+ size)
    // ...
}
```

ARR[0] ⇒ SAME AS IARR[0]

How do you know the size of an array in C?

No way to tell from just the pointer.

Up to the programmer to have a convention

E.G.

- ARGUMENT

- NULL TERMINATED STRINGS

## Memory organization example

Here's a larger example of memory organization. How would we organize the following variables?

```
unsigned int a = 0x00FF;
long int b[2] = { 65540, -5 };
char c = 'c'; // 'c' = 0x43
```

→ 0x4400 is start

- VARIABLES ALLOCATED IN ORDER.

STEP 1. WRITE EVERYTHING IN HEX,

B, EACH ELEMENT IS 4 BYTES (TYPE IS LONG INT)

$$B[0] = 65540 \\ = 65536 + 4 = 2^{16} + 2^2$$

$$\begin{array}{cccc} 31 & 16 & 15 & 7 \\ 0000 & 0000 & 0000 & 0000 \\ 00 & 01 & 00 & 04 \end{array}$$

$$B[0] = 00010004h$$

B[1] = -5 (NEGATIVE SO DO 2'S COMP)

0... 0000 0000 0000 0101 MAGNITUDE

1... 1111 1111 1111 1010 COMP

$$\begin{array}{cccc} 31 & 15 & 7 & 1 \\ 1111 & 1111 & 1111 & 1011 \\ FF & FF & FF & FB \end{array}$$

How many bytes of memory do we require?

$$B[1] = FFFF FFBh$$

A: 2 BYTES

B: (2 ELEMENTS)(4 BYTES/ELEMENT) = 8 BYTES

C: 1 BYTE ⇒ 11 BYTES TOTAL.

So using the above information, we can make our table:

$$A = 0x000F$$


$$B[0] = 00010004h$$

$$B[1] = FFFFFFFBh$$

$$C = 43h$$

Address	BE	LE
0x440C		
0x440B		
0x440A	43	43
0x4409	FB	FF
0x4408	FF	FF
0x4407	FF	FF
0x4406	FF	FB
0x4405	04	00
0x4404	00	01
0x4403	01	00
0x4402	00	04
0x4401	0F	00
0x4400	00	0F

(A | B[0] | B[1] | C)



ORDER OF VARIABLES ALWAYS THE SAME, BUT ORDER OF BYTES WITHIN A VARIABLE CHANGES w/ ENDIANNESS.

CHAR \*STR = "HELLO";

## STRINGS IN C:

CHAR \*S = "HELLO"

	ASCII	HEX	
4405	'\0'	00h	S[5]
4404	'O'	4F	S[4]
4403	'L'	4C	S[3]
4402	'L'	4C	S[2]
4401	'E'	45h	S[1]
0x4400	'H'	48h	S[0]

INT STR LENGTH(CHAR \*S) ;

