



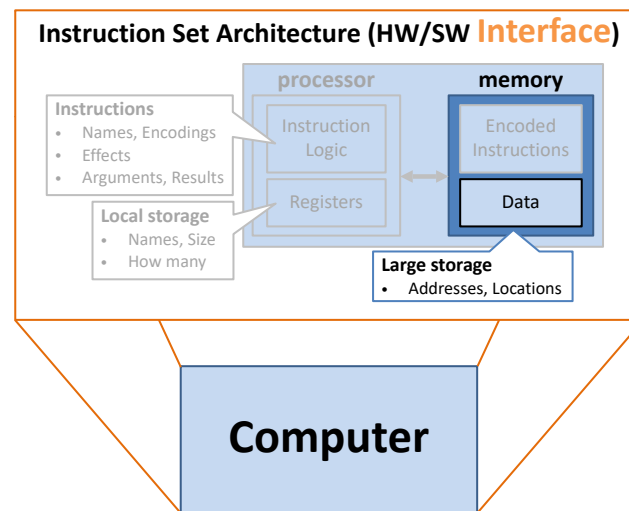
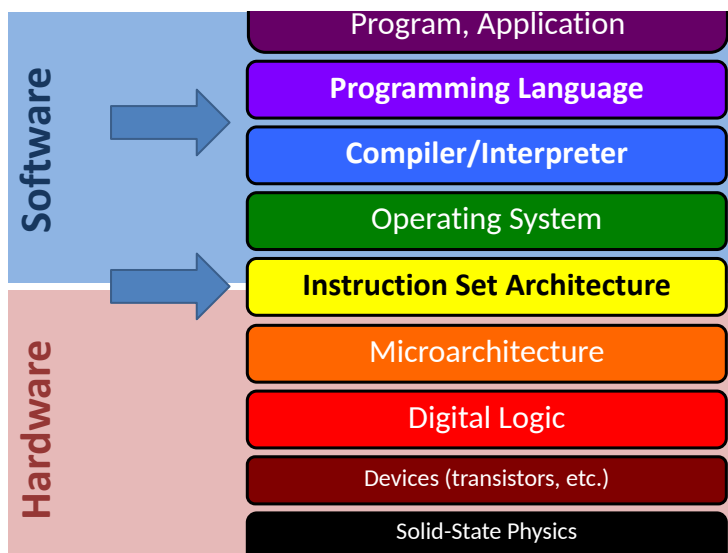
CS 240 Stage 2! Hardware-Software Interface

Memory addressing, C language, pointers
Assertions, debugging
Machine code, assembly language, program translation
Control flow
Procedures, stacks
Data layout, security, linking and loading



Programming with Memory


the memory model
pointers and arrays in C



Byte-addressable memory = mutable byte array

Location / cell = element

- Identified by unique numerical address
- Holds one byte (8 bits)



Address = index

- Unsigned number
- Represented by one word
- Computable and storable as a value

Operations:

- **Load:** read contents at given address
- **Store:** write contents at given address

5

Multi-byte values in memory

Store across contiguous byte locations.


Example: 8 byte (64 bit) values

Alignment

Multi-byte values start at addresses that are multiples of their size

Bit order within byte always same.



Recall: byte ordering within larger value?



6

Endianness: details

In what order are the individual bytes of a multi-byte value stored in memory?




Address	Contents	Little Endian: least significant byte first
03	2A	• low order byte at low address
02	B6	• high order byte at high address
01	00	• used by x86, ... and CS240!
00	0B	

Address	Contents	Big Endian: most significant byte first
03	0B	• high order byte at low address
02	00	• low order byte at high address
01	B6	• used by networks, SPARC, ...
00	2A	

7

Data, addresses, and pointers

For these slides, we'll draw the bytes in this reverse order so that multi-byte values can be read directly



memory drawn as 32-bit values, little endian order

8

Data, addresses, and pointers

address = index of a location in memory

pointer = a reference to a location in memory, represented as an address stored as data

Let's store the number 240 at address **0x20**.

$$240_{10} = F0_{16} = 0x00\ 00\ 00\ F0$$

At address **0x08** we store a pointer to the contents at address **0x20**.


At address **0x00**, we store a pointer to a pointer.

The number 12 is stored at address **0x10**.

Is it a pointer?

How do we know if values are pointers or not?

How do we manage use of memory?



9

C: Variables are locations

The compiler creates a map from variable name → location.

Declarations do not initialize!


```
int x; // x @ 0x20
int y; // y @ 0x0C
```

```
x = 0; // store 0 @ 0x20
```

```
// store 0x3CD02700 @ 0x0C
y = 0x3CD02700;
```

```
// 1. load the contents @ 0x0C
// 2. add 3
// 3. store sum @ 0x20
```

```
x = y + 3;
```



10

C: Variables are locations

The compiler creates a map from variable name → location.


Declarations do not initialize!

```
int x; // x @ 0x20
int y; // y @ 0x0C
```

```
x = 0; // store 0 @ 0x20
```

```
// store 0x3CD02700 @ 0x0C
y = 0x3CD02700;
```

```
// 1. load the contents @ 0x0C
// 2. add 3
// 3. store sum @ 0x20
x = y + 3;
```



11

C: Pointer operations and types

address = index of a location in memory

pointer = a *reference* to a location in memory, an address stored as data

Expressions using addresses and pointers:

& ___ **address of** the memory location representing ___
a.k.a. "reference to ___"

* ___ **contents at** the memory address given by ___
a.k.a. "dereference ___"

Pointer types:

___* **address of a memory location holding a** ___
a.k.a. "a reference to a ___"

12

C: Types determine sizes

Sizes of data types (in bytes)

Java Data Type	C Data Type
boolean	<i>bool</i>
byte	char
char	
short	short int
int	int
float	float
	long int
double	double
long	long long
	long double
(reference)	(pointer) *

Used by CS Linux, most modern machines

32-bit word	64-bit word
1	1
1	1
2	2
2	2
4	4
4	4
4	8
8	8
8	16
4	8

address size = word size

13

C: Pointer example

& = address of
* = contents at

```
int* p;
```

Declare a variable, p that will hold the address of a memory location holding an int

```
int x = 5;
int y = 2;
```

Declare two variables, x and y, that hold ints, and store 5 and 2 in them, respectively.

```
p = &x;
```

Take the address of the memory representing x ... and store it in the memory location representing p. Now, "p points to x."

```
y = 1 + *p;
```

Add 1 to the contents of memory at the address given by the contents of the memory location representing p ... and store it in the memory location representing y.

14

C: Pointer example

C assignment:

location \leftarrow Left-hand-side = right-hand-side; value

```
int* p; // p @ 0x04
int x = 5; // x @ 0x14, store 5 @ 0x14
int y = 2; // y @ 0x24, store 2 @ 0x24
p = &x; // store 0x14 @ 0x04
```

```
// 1. load the contents @ 0x04 (=0x14)
// 2. load the contents @ 0x14 (=0x5)
// 3. add 1
// 4. store sum as contents @ 0x24
y = 1 + *p;
```

```
// 1. load the contents @ 0x04 (=0x14)
// 2. store 0xF0 as contents @ 0x14
*p = 240;
```

& = address of
* = contents at

What is the type of *p?
What is the type of &x?
What is *(&y) ?

00	00	00	08	0x24	Y
				0x20	
				0x1C	
				0x18	
00	00	00	05	0x14	X
				0x10	
				0x0C	
				0x08	
00	00	00	14	0x04	P
				0x00	
x ₃	x ₂	x ₁	x ₀		

15

C: Pointer type syntax

Spaces between base type, *, and variable name mostly do not matter. The following are equivalent:

```
int* ptr;
int * ptr;
```

I see: "The variable ptr holds an address of an int in memory."

```
int *ptr;
```

more common C style

Looks like: "Dereferencing the variable ptr will yield an int."
Or "The memory location where the variable ptr points holds an int."

Caveat: do not declare multiple variables unless using the last form.
`int* a, b;` means `int *a, b;` means `int* a; int b;`

16

C: Arrays

Declaration: `int a[6];`


element type

name

number of elements

Arrays are adjacent memory locations storing the same type of data.

`a` is a name for the array's base address, can be used as an *immutable* pointer.



17

C: Arrays


Declaration: `int a[6];`

Indexing: `a[0] = 0xf0;`

Arrays are adjacent memory locations storing the same type of data.

`a` is a name for the array's base address, can be used as an *immutable* pointer.

Address of `a[i]` is base address `a` plus `i` times element size in bytes.



18

C: Arrays


Declaration: `int a[6];`

Indexing: `a[0] = 0xf0;`
`a[5] = a[0];`

Arrays are adjacent memory locations storing the same type of data.

`a` is a name for the array's base address, can be used as an *immutable* pointer.

Address of `a[i]` is base address `a` plus `i` times element size in bytes.



19

C: Arrays

Declaration: `int a[6];`


Indexing: `a[0] = 0xf0;`
`a[5] = a[0];`

No bounds check: `a[6] = 0xBAD;`

Arrays are adjacent memory locations storing the same type of data.

`a` is a name for the array's base address, can be used as an *immutable* pointer.

Address of `a[i]` is base address `a` plus `i` times element size in bytes.



20


C: Arrays

Declaration: `int a[6];`
 Indexing: `a[0] = 0xf0;`
`a[5] = a[0];`
 No bounds check: `a[6] = 0xBAD;`
`a[-1] = 0xBAD;`

Arrays are adjacent memory locations storing the same type of data.

`a` is a name for the array's base address, can be used as an *immutable* pointer.

Address of `a[i]` is base address `a` plus `i` times element size in bytes.



C: Arrays


Declaration: `int a[6];`
 Indexing: `a[0] = 0xf0;`
`a[5] = a[0];`
 No bounds check: `a[6] = 0xBAD;`
`a[-1] = 0xBAD;`

Arrays are adjacent memory locations storing the same type of data.

`a` is a name for the array's base address, can be used as an *immutable* pointer.

Address of `a[i]` is base address `a` plus `i` times element size in bytes.

Pointers: `int* p;`
 equivalent `{ p = a;`
`p = &a[0];`



C: Arrays


Declaration: `int a[6];`
 Indexing: `a[0] = 0xf0;`
`a[5] = a[0];`
 No bounds check: `a[6] = 0xBAD;`
`a[-1] = 0xBAD;`

Arrays are adjacent memory locations storing the same type of data.

`a` is a name for the array's base address, can be used as an *immutable* pointer.

Address of `a[i]` is base address `a` plus `i` times element size in bytes.

Pointers: `int* p;`
 equivalent `{ p = a;`
`p = &a[0];`
`*p = 0xA;`



C: Arrays


Declaration: `int a[6];`
 Indexing: `a[0] = 0xf0;`
`a[5] = a[0];`
 No bounds check: `a[6] = 0xBAD;`
`a[-1] = 0xBAD;`

Arrays are adjacent memory locations storing the same type of data.

`a` is a name for the array's base address, can be used as an *immutable* pointer.

Address of `a[i]` is base address `a` plus `i` times element size in bytes.

Pointers: `int* p;`
 equivalent `{ p = a;`
`p = &a[0];`
`*p = 0xA;`



C: Arrays

Declaration: `int a[6];`
 Indexing: `a[0] = 0xf0;`
`a[5] = a[0];`
 No bounds check: `a[6] = 0xBAD;`
`a[-1] = 0xBAD;`

Pointers: `int* p;`
 equivalent $\left\{ \begin{array}{l} p = a; \\ p = \&a[0]; \\ *p = 0xA; \end{array} \right.$


equivalent $\left\{ \begin{array}{l} p[1] = 0xB; \\ *(p + 1) = 0xB; \end{array} \right.$

array indexing = address arithmetic
 Both are scaled by the size of the type.

Arrays are adjacent memory locations storing the same type of data.

`a` is a name for the array's base address, can be used as an *immutable* pointer.

Address of `a[i]` is base address `a` plus `i` times element size in bytes.



25

C: Arrays

Declaration: `int a[6];`
 Indexing: `a[0] = 0xf0;`
`a[5] = a[0];`
 No bounds check: `a[6] = 0xBAD;`
`a[-1] = 0xBAD;`

Pointers: `int* p;`
 equivalent $\left\{ \begin{array}{l} p = a; \\ p = \&a[0]; \\ *p = 0xA; \end{array} \right.$


equivalent $\left\{ \begin{array}{l} p[1] = 0xB; \\ *(p + 1) = 0xB; \\ p = p + 2; \end{array} \right.$

array indexing = address arithmetic
 Both are scaled by the size of the type.

Arrays are adjacent memory locations storing the same type of data.

`a` is a name for the array's base address, can be used as an *immutable* pointer.

Address of `a[i]` is base address `a` plus `i` times element size in bytes.



26

C: Arrays

Declaration: `int a[6];`
 Indexing: `a[0] = 0xf0;`
`a[5] = a[0];`
 No bounds check: `a[6] = 0xBAD;`
`a[-1] = 0xBAD;`

Pointers: `int* p;`
 equivalent $\left\{ \begin{array}{l} p = a; \\ p = \&a[0]; \\ *p = 0xA; \end{array} \right.$


equivalent $\left\{ \begin{array}{l} p[1] = 0xB; \\ *(p + 1) = 0xB; \\ p = p + 2; \end{array} \right.$

array indexing = address arithmetic
 Both are scaled by the size of the type.

Arrays are adjacent memory locations storing the same type of data.

`a` is a name for the array's base address, can be used as an *immutable* pointer.

Address of `a[i]` is base address `a` plus `i` times element size in bytes.



27

C: Arrays

Declaration: `int a[6];`
 Indexing: `a[0] = 0xf0;`
`a[5] = a[0];`
 No bounds check: `a[6] = 0xBAD;`
`a[-1] = 0xBAD;`

Pointers: `int* p;`
 equivalent $\left\{ \begin{array}{l} p = a; \\ p = \&a[0]; \\ *p = 0xA; \end{array} \right.$

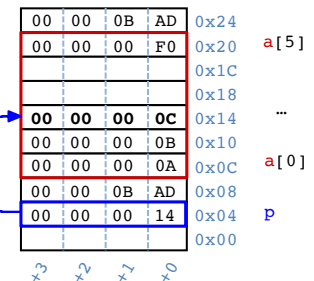
equivalent $\left\{ \begin{array}{l} p[1] = 0xB; \\ *(p + 1) = 0xB; \\ p = p + 2; \end{array} \right.$

array indexing = address arithmetic
 Both are scaled by the size of the type.

Arrays are adjacent memory locations storing the same type of data.

`a` is a name for the array's base address, can be used as an *immutable* pointer.

Address of `a[i]` is base address `a` plus `i` times element size in bytes.



`*p = a[1] + 1;`

28

C: * and []

ex

C programmers often use * where you might expect []:

e.g., `char*`:

- pointer to a `char`
- pointer to the first `char` in a string of unknown length

```
int strcmp(char* a, char* b);
```

33

C: 0 vs. '\0' vs. NULL

0
 Name: zero
 Type: `int`
 Size: 4 bytes
 Value: `0x00000000`
 Usage: The integer zero.

'\0'
 Name: null character
 Type: `char`
 Size: 1 byte
 Value: `0x00`
 Usage: Terminator for C strings.

NULL
 Name: null pointer / null reference / null address
 Type: `void*`
 Size: 1 word (= 8 bytes on a 64-bit architecture)
 Value: `0x0000000000000000`
 Usage: The absence of a pointer where one is expected.
 Address 0 is inaccessible, so *NULL is invalid; it crashes.

Is it important/necessary to encode the null character or the null pointer as 0x0?

What happens if a programmer mixes up these "zeroey" values?

34


Memory address-space layout

Addr	Perm	Contents	Managed by	Initialized
2 ^N -1 ↓ Stack	RW	Procedure context	Compiler	Run time
↑ Heap	RW	Dynamic data structures	Programmer, malloc/free, new/GC	Run time
Statics	RW	Global variables/ static data structures	Compiler/ Assembler/Linker	Startup
Literals	R	String literals	Compiler/ Assembler/Linker	Startup
Text	X	Instructions	Compiler/ Assembler/Linker	Startup
0				

35

C: Dynamic memory allocation in the heap

Heap:



Managed by memory allocator:

pointer to newly allocated block
of at least that size

number of contiguous bytes required

```
void* malloc(size_t size);
```

```
void free(void* ptr);
```

pointer to allocated block to free

36

C: standard memory allocator

```
#include <stdlib.h> // include C standard library
void* malloc(size_t size)
Allocates a memory block of at least size bytes and returns its address.
If memory error (e.g., allocator has no space left), returns NULL.
```

Rules:

- Check for error result.
- Cast result to relevant pointer type.
- Use sizeof(...) to determine size.

```
void free(void* ptr)
```

Deallocates the block referenced by ptr, making its space available for new allocations.

ptr must be a malloc result that has not yet been freed.

Rules:

- ptr must be a malloc result that has not yet been freed.
- Do not use *ptr after freeing.

37



C: Dynamic array allocation

```
#define ZIP_LENGTH 5
int* zip = (int*)malloc(sizeof(int)*ZIP_LENGTH);
if (zip == NULL) { // if error occurred
    perror("malloc"); // print error message
    exit(0); // end the program
}

zip[0] = 0;
zip[1] = 2;
zip[2] = 4;
zip[3] = 8;
zip[4] = 1;

printf("zip is");
for (int i = 0; i < ZIP_LENGTH; i++) {
    printf(" %d", zip[i]);
}
printf("\n");

free(zip);
```



38

C: Array of pointers to arrays of ints

```
int** zips = (int**)malloc(sizeof(int*) * 3);


zips[0] = (int*)malloc(sizeof(int)*5);
int* zip0 = zips[0];
zip0[0] = 0;
zips[0][1] = 2;
zips[0][2] = 4;
zips[0][3] = 8;
zips[0][4] = 1;

zips[1] = (int*)malloc(sizeof(int)*5);
zips[1][0] = 2;
zips[1][1] = 1;
zips[1][2] = 0;
zips[1][3] = 4;
zips[1][4] = 4;

zips[2] = NULL;
```


Why
no NULL?

Why terminate
with NULL?



39

Zip code

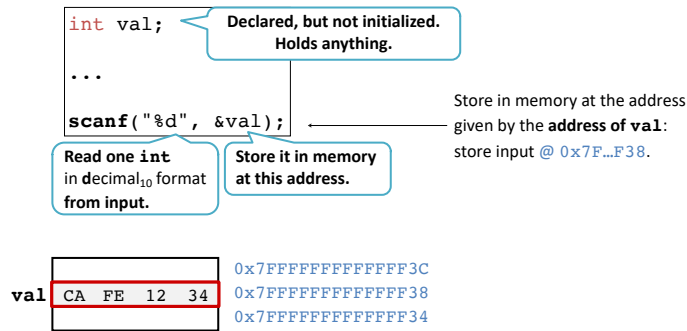


```
// return a count of all zips that end with digit endNum
int zipCount(int* zips[], int endNum) {
```

```
}
```

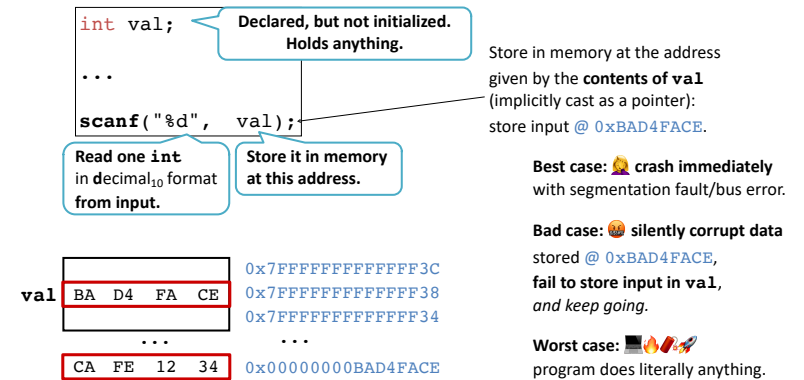
40

scanf reads formatted input



41

C: Classic bug using scanf



42

C: Memory error messages

- 11: **segmentation fault** ("segfault", SIGSEGV)
accessing address outside legal area of memory
- 10: **bus error** (SIGBUS)
accessing misaligned or other problematic address

More to come on debugging!



<http://xkcd.com/371/>

43

C: Why?

Why learn C?

- Think like actual computer (abstraction close to machine level) without dealing with machine code.
- Understand just how much Your Favorite Language provides.
- Understand just how much Your Favorite Language might cost.
- Classic.
- Still (more) widely used (than it should be).
- Pitfalls still fuel devastating reliability and security failures today.


Why not use C?

- Probably not the right language for your next personal project.
- It "gets out of the programmer's way" ... even when the programmer is unwittingly running toward a cliff.
- Advances in programming language design since the 70's have produced languages that fix C's problems while keeping strengths.

44

Group example: longest string starts with

ex



```
// Return the starting character of the longest string in the
// null-terminated strings array.
// You can use: int strlen(char *str)
char longest_string_starts_with(char ** strings) {
```