

Programming with Memory

via C, pointers, and arrays

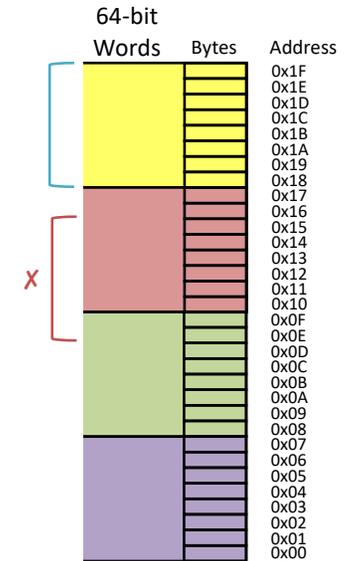
Why not just registers?

- Represent larger structures
- Computable addressing
- Indirection

multi-byte values in memory

Store across contiguous byte locations.

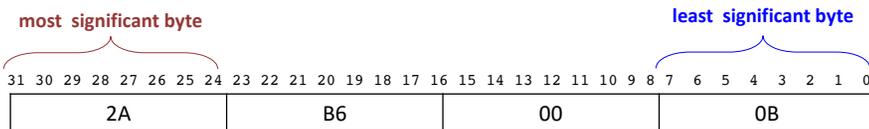
Alignment (Why?)



Bit order within byte always same.

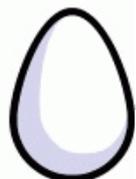
Byte ordering within larger value?

Endianness: To store a multi-byte value in memory, which byte is stored first (at a lower address)?



Address	Contents
03	2A
02	B6
01	00
00	0B

Address	Contents
03	0B
02	00
01	B6
00	2A



Little Endian: least significant byte first

- low order byte at low address, high order byte at high address
- used by x86, ...

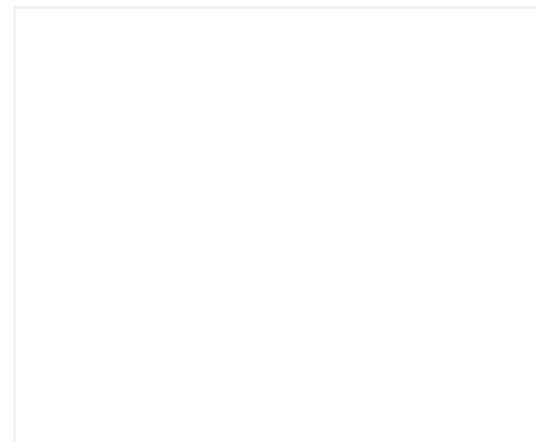
Big Endian: most significant byte first

- high order byte at low address, low order byte at high address
- used by networks, SPARC, ...

Data, Addresses, and Pointers

address = index of a cell in memory

pointer = address represented as data



				0x24
00	00	00	F0	0x20
				0x1C
				0x18
				0x14
00	00	00	0C	0x10
				0x0C
00	00	00	20	0x08
				0x04
				0x00

0x03 0x02 0x01 0x00

memory drawn as 32-bit values, little endian order

C: variables are memory locations (for now)

Compiler maps variable → memory location.

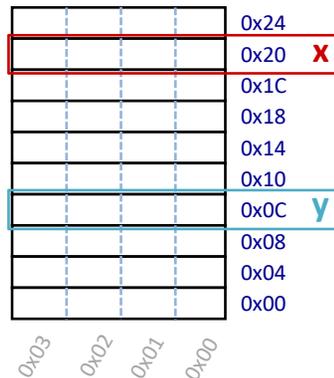
Declarations do not initialize!

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

x = 0; // store 0 at 0x20

// store 0x3CD02700 at 0x0C
y = 0x3CD02700;

// load the contents at 0x0C,
// add 3, and store sum at 0x20
x = y + 3;
```



14

C: Address and Pointer Primitives

address = index of a cell/location in memory

pointer = address represented as data

Expressions using addresses and pointers:

& ___ address of the memory location representing ___
***** ___ contents at the memory address given by ___
 a.k.a. "dereference ___"

Pointer types:

___* address of a memory location holding a ___

C: Address and 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.

Get the address of the memory location
`p = &x;` representing x
... and store it in p. Now, "p points to x."

Add 1 to the contents of memory at the address
`y = 1 + *p;` stored in p
... and store it in the memory location representing y.

19

C: Address and Pointer Example

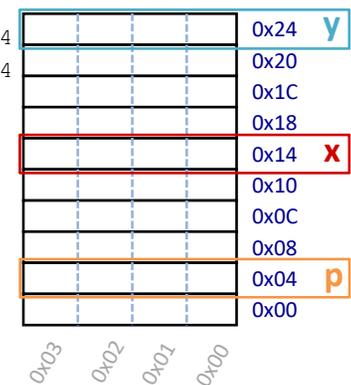
& = address of
*** = contents at

C assignment:

Left-hand-side = right-hand-side;

location value

```
int* p; // p: 0x04
int x = 5; // x: 0x14, store 5 at 0x14
int y = 2; // y: 0x24, store 2 at 0x24
p = &x; // store 0x14 at 0x04
// load the contents at 0x04 (0x14)
// load the contents at 0x14 (0x5)
// add 1 and store sum at 0x24
y = 1 + *p;
// load the contents at 0x04 (0x14)
// store 0xF0 (240) at 0x14
*p = 240;
```



C: Arrays

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

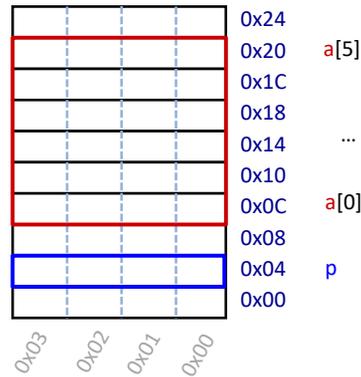
Pointers:
 equivalent `int* p;`
`p = a;`
`p = &a[0];`
`*p = 0xA;`

equivalent `p[1] = 0xB;`
`*(p + 1) = 0xB;`
`p = p + 2;`

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

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

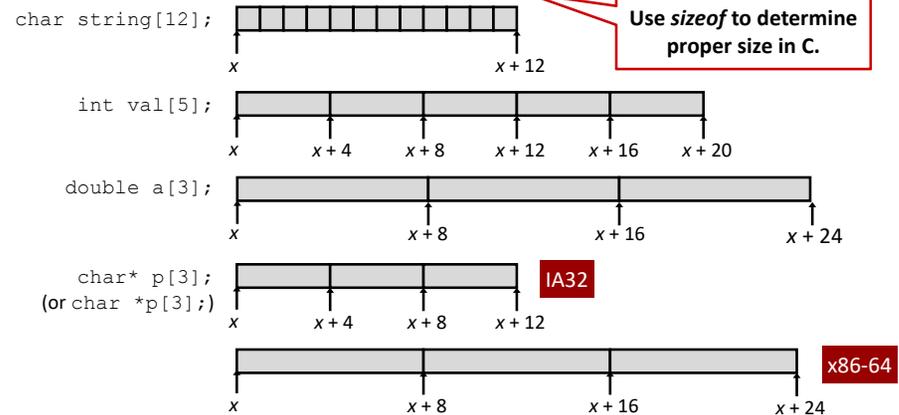
Arrays are adjacent memory locations storing the same type of data.
a is a name for the array's base address, can be used as an *immutable* pointer.
 Address of **a[i]** is base address **a** plus **i** times element size in bytes.



C: Array Allocation

Basic Principle

`T A[N];`
 Array of length *N* with elements of type *T* and name *A*
 Contiguous block of $N * \text{sizeof}(T)$ bytes of memory

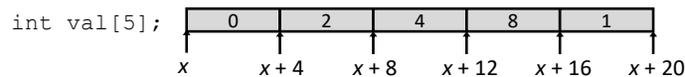


C: Array Access

ex

Basic Principle

`T A[N];`
 Array of length *N* with elements of type *T* and name *A*
 Identifier *A* has type

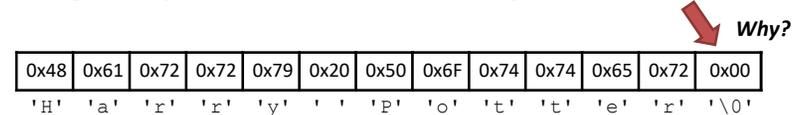


Reference	Type	Value
<code>val[4]</code>	<code>int</code>	1
<code>val</code>	<code>int*</code>	
<code>val+1</code>	<code>int*</code>	
<code>&val[2]</code>	<code>int*</code>	
<code>val[5]</code>	<code>int</code>	
<code>*(val+1)</code>	<code>int</code>	
<code>val + i</code>	<code>int*</code>	

C: Null-terminated strings

ex

C strings: arrays of ASCII characters ending with *null* character.



Does Endianness matter for strings?

```
int string_length(char str[]) {
    // ...
}
```

C: * and []

ex

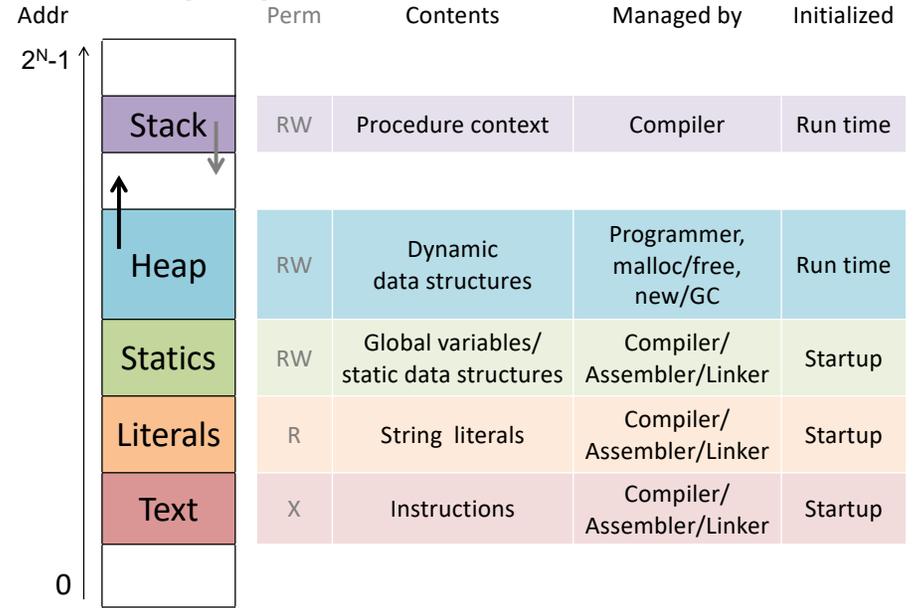
C programmers often use * where you might expect []:

e.g., char*:

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

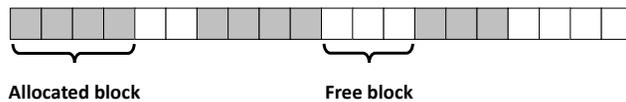
```
int strcmp(char* a, char* b);
int string_length(char* str) {
    // Try with pointer arithmetic, but no array indexing.
}
```

Memory Layout



C: Dynamic memory allocation in the heap

Heap:

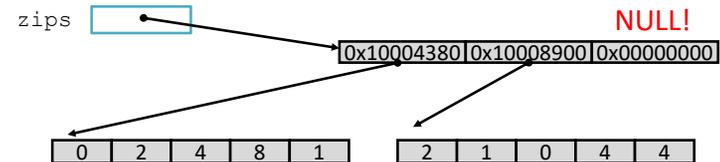


Managed by memory allocator:

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

Annotations:
- pointer to newly allocated block of at least that size (points to malloc)
- number of contiguous bytes required (points to size)
- pointer to allocated block to free (points to ptr)

Zip Cycles



// return a count of all zips the end with digit endNum

```
int zipCount(int* zips[], int endNum) {
}
```

C: scanf reads formatted input

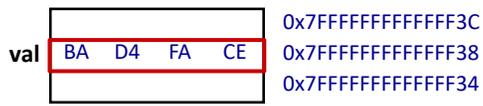
```
int val;
...
scanf("%d", &val);
```

Declared, but not initialized
– holds anything.

Read one int
from input.

Store it in memory
at this address.

i.e., store it in memory at the address
where the contents of val is stored:
store into memory at 0xFFFFFFFF38.



C: classic bug using scanf

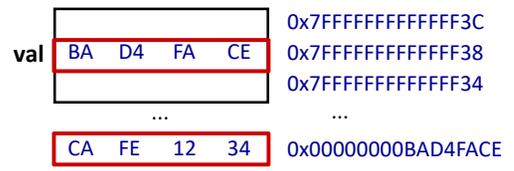
```
int val;
...
scanf("%d", val);
```

Declared, but not initialized
– holds anything.

Read one int
from input.

Store it in memory
at this address.

i.e., store it in memory at the address
given by the contents of val:
store into memory at 0xBAD4FACE.



Best case: segmentation fault,
or bus error, crash.

Bad case: silently corrupt data
stored at address 0xBAD4FACE,
and val still holds 0xBAD4FACE.

Worst case: arbitrary corruption