

C Memory

CS 2130: Computer Systems and Organization 1

Xinyao Yi Ph.D.
Assistant Professor

Announcements

- Homework 8 releasing soon, due next Monday on Gradescope

C Reference Guide

Functions Definition *We only want the function definition to appear in one place*

```
int f(int x) {  
    return 2130 * x;  
}
```

- Definition of the function

We only want this in **one** .c file

- Do not want 2 definitions
- Which one should the linker choose?

Header Files

C header files: `.h` files

- Written in C, so look like C
- Only put header information in them
 - Function headers
 - Macros
 - `typedefs`
 - `struct` definitions
- Essentially: information for the **type checker** that does not produce any actual binary
- `#include` the header files in our `.c` files *(take the entire file and paste it at the top of my code).*


Big Picture

Header files

- Things that tell the type checker how to work
- Do not generate any actual binary

C files

- Function definitions and implementation
- Include the header files (one header file can be used by multiple C files).



Including Headers

`#include "myfile.h"` *(I want to include a file that I wrote).*

- Quotes: look for a file where I'm writing code
- Our header files

`#include <string.h>`

- Angle brackets: look in the standard place for includes
- Code that came with the compiler *(a set of includes that came with the compiler).*
- Likely in `/usr/include`

Macros

`#define NAME something else`

- Object-like macro

→ It does this at the text level.

- Replaces NAME in source with **something else**

`#define NAME(a,b) something b and a`

→ whatever is after the comma until the second parentheses is the second thing.

- Function-like macro

→ Whatever is the first thing after the opening parentheses until the comma.

- Replaces NAME(X,Y) with **something Y and X**

→ text level

Lexical replacement, *not* semantic

Interesting Example

```
#define TIMES2(x)  x * 2          /* bad practice */
#define TIMES2b(x) ((x) * 2)     /* good practice */

int x = ! TIMES2(2 + 3);
```

$$x = \underbrace{!}_{0} \underbrace{2+3}_{6} * 2; \Rightarrow x = 0 + 6 \Rightarrow x = 6.$$

```
int y = ! TIMES2b(2 + 3);
```

$x = !(2+3)*2;$ Be very explicit when you're using the macro define. Wrap them in parentheses and enforce an order of operations.

Examples and More

- header example
- string.h
- variadic functions

`#ifndef` } If ... is already defined, it won't do
`#endif` } anything between the `ifndef` and `endif`

why? We want to include a header file in all of our C files. if I include the same file multiple times, it's correct but weird.

Memory

I have multiple programs running on my computer all the time. But they're all going to think they can see all of memory.

In reality, that's not true — each process only has its own isolated view of memory.

(process A and process B might both think they have memory at address 0x1000, but they're actually referring to different physical locations).

If each process thinks it has access to “all addresses” — like terabytes or exabytes of possible memory — how is that possible when your computer might only have 16 GB of RAM?

That's the “weird” part that virtual memory solves.

We're not going to go deep into memory management right now — you'll learn about address translation, page tables, and virtual-to-physical mapping later in Computer Systems II (CSO2).

Virtual memory lets a program use **virtual addresses**, which the hardware (Memory Management Unit, MMU) automatically **translates** into **physical addresses** — the real locations in RAM or on disk.

Memory

ffff...f

Kernel Space
or kernel
Memory.

reserved
for the
kernels

the upper
half of the
memory.

If we try
to read or
write from
them, you
are not
allowed.

User
Space

If you're
not in the
kernel, you
are in
user
space.

From assembly:

- ①. Generate a memory address and go to memory.
- ②. hardware is going to check if we're actually allowed to read/write from that place in memory. (rely on a piece of software).

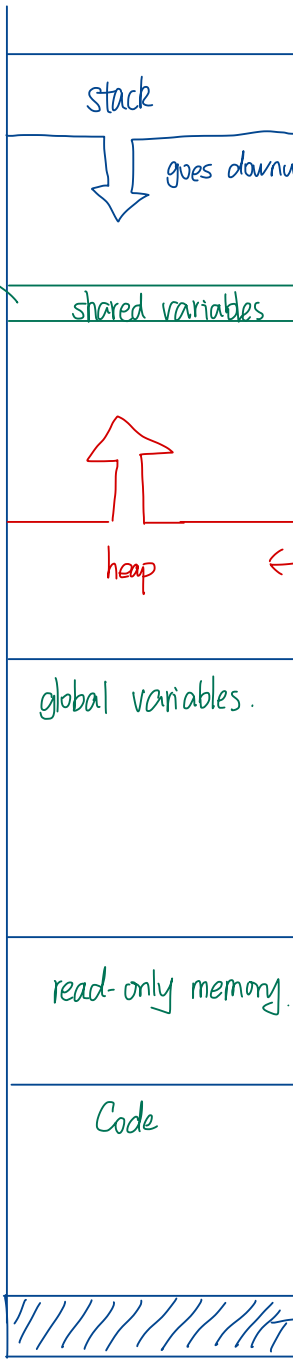
2 names for that: kernel or OS (operating system).

↳ help the hardware decide whether these addresses are good or not.

It divides memory up into a couple of different segments.

a chunk of memory

It's going to have different properties based on what segment of memory I'm in.



variables available to multiple functions

stack frame :

Anytime I call a function, there are a couple of things that get put on the stack:

- ① parameters
- ② return address.
- ③ local variables.

When I return, RSP moves away, memory does not get erased.

One thing missing: pointers, variables that we want to leave when our functions return. (pass to another function?)

In C, I can define variables outside of functions. For local variables in functions: put values in stack. But for global variables, we put them here. (The size of the variables are set at compile time, and the size cannot be changed.)

Example: string literals
Things are not code, but that were in my code that are really helpful for my program.

My code is going to be stuck in memory somewhere near the bottom.

We put the code near the bottom, but not at the bottom. (Explanation next page).

At the very bottom of memory — near address 0 — there's a special protected region. The operating system doesn't let us read or write anything there. Why? To protect us from ourselves.

If we accidentally take an integer, cast it to a pointer, and try to access that address — for example, address 5 — the hardware will stop us. That's why we get a segmentation fault: we tried to read or write from a memory segment we're not allowed to touch.

In C, the constant NULL is not a keyword but is usually defined as a pointer to address 0. This means that dereferencing a null pointer — trying to access the memory at address 0 — will also trigger a segmentation fault. And that's intentional. We want the system to complain immediately instead of silently reading some random value from the bottom of memory.

So, the “bottom of memory” is deliberately left empty to catch common pointer mistakes — for example, dereferencing NULL or casting an integer like 5 to a pointer and using it as an address. When that happens, the operating system stops the program with a segmentation fault rather than letting it continue with undefined behavior.

In short:

The low-address region exists as a safety barrier — it's the system's way of saying, “You're trying to access memory that doesn't belong to you.”