

A Brief C Primer for CS 350

Benjamin Cassell

University of Waterloo
becassel@uwaterloo.ca

May 6, 2014

Overview

- 1 Introduction
- 2 Data Types
 - ▶ Structs
 - ▶ Enums
 - ▶ Pointers
 - ▶ Arrays
 - ▶ Strings
- 3 Volatile Variables
- 4 File Manipulation
 - ▶ Writing Files
 - ▶ Reading Files
- 5 Memory
 - ▶ Stack
 - ▶ Heap
 - ▶ Endianness
- 6 Conclusion

Why C?



(a) Using C to code.



(b) Using other languages to code.

A scientific and unbiased comparison of C with other languages.

Pointers

Pointers are C's method of allowing the programmer to directly manipulate memory:

- Pointers are variables that contain a memory address.
- The data contained at the address is inferred by the pointer type:
 - `int* int_ptr`: A pointer to an `int`.
 - `foo* foo_ptr`: A pointer to a `foo`.
 - `void* ptr`: A generic pointer with no particular type.
 - `int** int_ptr_ptr`: A pointer to a pointer to an `int`.

Pointers

Warning!

The position of the `*` in a pointer declaration is irrelevant, but you may start a minor war based on your decision. Although `foo*`
`my_foo`, `foo * my_foo` and `foo *my_foo` are all equivalent, the suggested style (and the one used by the Linux kernel) is `foo *my_foo`. Just be sure to stay consistent!

Warning!

Be **VERY** careful when declaring multiple pointers on the same line. The `*` in a declaration only applies to the particular variable it is attached to.

Pointers

A bad declaration of multiple `int*`'s:

```
int* foo, bar, foobar;
```

Properly declaring multiple `int*`'s:

```
int* foo, * bar, * foobar;
```

Pointers

Pointers can be assigned in a variety of ways:

- Grabbing the address of a variable with the '&' operator:

```
int foo = 1;
int* bar = &foo;
```

- Casting other numerical values:

```
int* foo = (int*)0xFFFF0000;
unsigned int bar = 0xFFFF0000;
int* foobar = (int*)bar;
```

- Using a predefined macro like NULL (which is equal to 0):

```
int* foo = NULL;
```

Pointers

Remember, a pointer contains an address of a value, not the value itself. We can examine the address itself if desired:

```
printf("0x%x\n", foo); // Will print "0xFFFF0000"
```

To examine the value that resides at the address a pointer points to, use the dereference operator '*':

```
printf("%d\n", *bar); // Will print "1"
```

Pointers

Pointers are particularly useful for “passing by reference”:

- Parameters are normally passed by value.
- Passing a pointer allows the programmer to manipulate the original variable from within a function.

```
void foo(int* bar) {
    *bar = *bar + 1;
}

int main(int argc, char** argv) {
    int bar = 0;
    foo(&bar);
    printf("%d\n", *bar); // Will print "1"
    return 0;
}
```

Pointers

Pointers are very useful for passing structs around:

- Structs are passed by value by default.
- Passing by reference is more efficient and allows you to modify a struct from anywhere!

```
foo my_foo;
foo* foo_ptr = &my_foo;
do_something(&my_foo);
```

- When accessing the members of a struct via a pointer, the '.' operator annoyingly takes precedence over the '*' operator:

```
*foo_ptr.x = 0; // Wrong!
(*foo_ptr).x = 0; // Right!
```

Pointers

- To avoid issues with operator precedence and members of structs passed by reference, C has the '->' operator.

```
(*foo_ptr).x = 0; // Right!  
foo_ptr->x = 0; // Also right!
```

- Remember, pointers can be made for any type of memory, including enums and other pointers!
- The `void` pointer is a pointer that represents a memory containing untyped data. Often `void` pointers are used to work with multiple different types of data or memory buffers.
- Because `char*` points to memory divided into bytes, it too is often used to work with memory buffers.

Arrays

Arrays are collections of elements of a single type:

```
int foo[32];
```

- This example creates an array of 32 integers.
- The total size occupied by an array is `sizeof (type) * num_elements`.
- Elements are accessed using the '[]' operator.
- The elements of an array are laid out contiguously in memory.
 - If the address of `foo[0]` is `0x10`, then the address of `foo[1]` would be `0x14`.

Arrays may be of any type, including structs and enums:

```
struct foo bar[32];
```


Arrays

In C, arrays are managed with pointer arithmetic:

- The variable name of the array is, in reality, just the memory address of the first element of the array (in other words, `&foo[0]` is equal to `foo`).
- Using the `[]` operator actually does the following:

```
int foo[32];
int bar1 = foo[8];
int bar2 = *(foo + 8); // Identical to bar1
int bar3 = *(((int*)((unsigned int)foo) + sizeof(int) * 8));
           // Identical to bar1
```

- This is why arrays are zero-based in C.

Arrays

Warning!

An common pitfall arises from forgetting that pointers use pointer arithmetic. Also common is improper casting when manipulating pointers. Use explicit casting and brackets when in doubt.

Warning!

Pointers and arrays are similar, but **NOT** identical. Declaring a pointer allocates memory to store an address. Declaring an array only allocates for the elements, meaning the array variable itself cannot be mutated. Using '&' on a pointer results in a pointer to a pointer. Using '&' on an array simply results in the same value as the array variable!

Arrays

The similarities between pointers and arrays, and C's deference of memory management responsibility to the programmer enables some cool tricks:

```
typedef struct {
    int x;
    int y;
} foo;

int main(int argc, char** argv) {
    char bar[sizeof(foo) * 32]; // bar is an array of char
    foo* foobar = (foo*)bar; // Treat bar as an array of foo
    int* barfoo = (int*)bar; // Treat bar as an array of int
    return 0;
}
```

Arrays

Warning!

Be very cautious when using `sizeof` with arrays. Running `sizeof` on an array declared locally inside a function will return the total byte size of all elements contained in an array. Running `sizeof` on a pointer to the same array, however, will return the size of a pointer! Even more confusing, passing an array as an argument to a function only actually passes a pointer to the array. It is almost **ALWAYS** more useful to write functions that accept pointers as arguments instead of arrays, as this is more faithful to what is actually happening behind the scenes.

Strings

Remember that a pointer can be directed to point at any memory address at all, making something like the following possible:

```
char foo[] = "Hello World!";
char* bar = foo;
// foo[i] and bar[i] are now the same for any i
// sizeof(foo) and sizeof(bar) are still different
```

You can even declare strings to be larger than their contents:

```
char foo[] = "Hello World!";
char bar[13] = "Hello World!"; // foo and bar are identical
char foobar[20] = "Hello World!"; // Has 7 uninitialized bytes
```


Volatile Variables

Consider a more urgent example:

```
unsigned int foo; // foo is memory-mapped
void wait_foo() {
    foo = 0;
    while (foo != 255);
}
```

This could be optimized into the following code:

```
unsigned int foo; // foo is memory-mapped
void wait_foo() {
    foo = 0;
    while (1);
}
```

Volatile Variables

Again, the fix is to make our variable volatile:

```
volatile unsigned int foo; // foo is memory-mapped
void wait_foo() {
    foo = 0;
    while (foo != 255);
}
```

File Manipulation

C allows for direct interaction with and manipulation of files:

- People are often intimidated by file manipulation in C, but once you're used to using it, it's just as easy as file manipulation in any other language.
- Open files in C are dealt with through file descriptors, numbers which represent open files and allow you to interact with them.
- Your main tools when interacting with files will be the `open`, `close`, `read` and `write` system calls.
- All of these calls are very well-documented online.
- If you don't close open files you can run out of file descriptors!

Writing Files

```
#include <fcntl.h>
#include <stdio.h>
#include <stdlib.h>
#include <sys/stat.h>
#include <sys/types.h>

int main (int argc, char** argv) {
    int i, rc, fd;
    unsigned int buf[40];
    for (i = 0; i < 40; i+= 1) {
        buf[i] = i;
    }
    fd = open("test-output", O_CREAT | O_WRONLY, S_IRWXU);
```


Reading Files

```
#include <fcntl.h>
#include <stdio.h>
#include <stdlib.h>
#include <sys/stat.h>
#include <sys/types.h>

#define PER_ROW 4

int main(int argc, char** argv) {
    int i, rc, fd;
    unsigned int buf[40];
    if ((fd = open("test-output", O_RDONLY)) < 0)
        exit(1);
}
```


Reading Files

```
if ((rc = read(fd, buf, sizeof(buf))) < 0) {
    exit(1); // Should really perform cleanup here
}
for (i = 0; i < 40; i += 1) { // Should use rc and sizeof
    if (i % PER_ROW == 0) {
        printf("offset = %4d : ", i * sizeof(unsigned int));
    }
    printf("0x%08x ", buf[i]);
    if ((i + 1) % PER_ROW == 0) {
        printf("\n");
    }
}
close(fd);
exit(0);
}
```

Reading Files

```
offset = 0 : 0x00000000 0x00000001 0x00000002 0x00000003
offset = 16 : 0x00000004 0x00000005 0x00000006 0x00000007
offset = 32 : 0x00000008 0x00000009 0x0000000A 0x0000000B
offset = 48 : 0x0000000C 0x0000000D 0x0000000E 0x0000000F
offset = 64 : 0x00000010 0x00000011 0x00000012 0x00000013
offset = 80 : 0x00000014 0x00000015 0x00000016 0x00000017
offset = 96 : 0x00000018 0x00000019 0x0000001A 0x0000001B
offset = 112 : 0x0000001C 0x0000001D 0x0000001E 0x0000001F
offset = 128 : 0x00000020 0x00000021 0x00000022 0x00000023
offset = 144 : 0x00000024 0x00000025 0x00000026 0x00000027
```


Stack

Every variable in the following example is located on the stack:

```
void foo (int bar) {  
    int foobar;  
    return;  
}
```

Warning!

It is entirely possible to run out of space on the stack. By default, most OS's will give you "only" a few megabytes of stack space!

Stack

Warning!

One of the most common and basic mistakes you can make is to reference uninitialized variables, structs and array elements (e.g. using '+=' with an uninitialized integer, or assuming pointers are initialized to NULL). Get in the habit of initializing everything, and zero out entire structures or arrays if you have to!

Warning!

It is a critical mistake to reference stack variables from outside of their context. Passing variables downwards to functions is fine, but don't try to pass pointers to parameters or stack variables upwards! Don't even think about doing it!

Stack

Don't ever do this:

```
int* foo1 (int* bar) {
    int foobar = *bar;
    return &foobar; // This is almost certainly wrong!
}

int* foo2() {
    int bar = 0;
    int* foobar;
    foobar = foo1(&bar); // Passing bar downwards is fine
    return foobar; // Normally fine, but wrong because of foo1
}
```


Heap

Warning!

Calling `malloc` can return `NULL` if no memory is available!

Warning!

Like the stack, heap memory is uninitialized. Don't forget to initialize the memory your heap-based pointers point to!

Warning!

If you forget to return heap memory to the OS (for example in error cases), you will encounter memory leaks! This forgotten memory won't be available for use until the program terminates!

Endianness

Here are what bytes look like in little and big endian:

```
unsigned int x = 0xDEADBEEF;
```

Little endian: Least significant byte at lowest address

```
0 .. 7 8 .. 15 16 .. 23 24 .. 31
[ EF ] [ BE ] [ AD ] [ DE ]
```

Big endian: Most significant byte at lowest address

```
0 .. 7 8 .. 15 16 .. 23 24 .. 31
[ DE ] [ AD ] [ BE ] [ EF ]
```

Endianness

Here is what the sample from file reading looks like when dumped on a little endian system:

```
offset = 0 : 0x00000000 0x01000000 0x02000000 0x03000000
offset = 16 : 0x04000000 0x05000000 0x06000000 0x07000000
offset = 32 : 0x08000000 0x09000000 0x0A000000 0x0B000000
offset = 48 : 0x0C000000 0x0D000000 0x0E000000 0x0F000000
offset = 64 : 0x10000000 0x11000000 0x12000000 0x13000000
offset = 80 : 0x14000000 0x15000000 0x16000000 0x17000000
offset = 96 : 0x18000000 0x19000000 0x1A000000 0x1B000000
offset = 112 : 0x1C000000 0x1D000000 0x1E000000 0x1F000000
offset = 128 : 0x20000000 0x21000000 0x22000000 0x23000000
offset = 144 : 0x24000000 0x25000000 0x26000000 0x27000000
```


The End

Congratulations! You're on your way to becoming an expert C programmer! Everyone will want to use your OS!

Questions?