

CS2100 Computer Organization A Quick Introduction to C

1. Introduction

This document gives a very quick introduction to the C Programming Language. It is assumed that the reader is already reasonably proficient in programming methodology, and hence this document does not explain basic programming concepts like data-types and functions. It will however explain concepts that are unique to C, like pointers. The objective of this document is not to teach how to program in C, but to make the relevant sections of the CS2100 lecture notes comprehensible to programmers who are unfamiliar with the C programming language.

2. Data Types in C

Unlike languages like Python and JavaScript, C is a strictly typed language. C also strictly requires that all variables are declared before being used. The code fragment below shows how to declare a signed integer and a floating point number:

```
int x; // x is a signed integer
float y; // y is a floating point number (i.e. it holds real numbers)
```

2.1 Numerical Data Types

The common C numerical data types are:

Type	Meaning	Size and Encoding	Range
char	Signed byte	8 bits, 2's complement	-128 to 127
unsigned char	Unsigned byte	8 bits, unsigned	0 to 255
short	Signed integers	16 bits, 2's complement	-32768 to 32767
unsigned short	Unsigned integers	16 bits, unsigned	0 to 65535
int	Signed integers	32 bits, 2's complement	-2,147,483,648 to 2,147,483,647
unsigned	Unsigned integers	32 bits, unsigned	0 to 4,294,967,295
long	Signed integers	64 bits, 2's complement	-9,223,372,036,854,775,808 to 9,223,372,036,854,775,807
unsigned long		64 bits, unsigned	0 to 18,446,744,073,709,551,615
float	Signed floating point	32-bit IEEE 754	+/- 1.2e-38 to 3.4e+38
double	Signed float	64-bit IEEE 754	+/- 2.3e-308 to 1.7e+308

2.2 Alphanumeric Data

All data in C is (or can be) represented as integers. Characters are represented by 8-bit “char” integers (based on the ASCII table) and strings are represented as an array of char:

```
char mystr[128]; // This is a string of 127 characters.
```

Notice that an array of 128 char stores a string of 127 characters. This is because all strings in C must be terminated by a ‘\0’ (ASCII 0) character (called a NULL character).

2.3 Boolean Values

C does not have Boolean values; True and false values are represented by integers, with 0 being False, and any non-zero value being True. Hence:

```
if(0) {  
    // The statement here will never be executed since 0 is false  
}  
  
if(-1) {  
    // The statement here will always be executed since -1 is true.  
}
```

In C we do not assume that the “True” value is always 1. It can be any non-zero value. This can cause problems; The statement below shows the correct way to do comparisons:

```
if(x == 5) {  
    // Executed if x is 5  
}
```

But:

```
if(x = 5) {  
    // Always executed. C assigns 5 to x, then takes the result 5 in the  
    // if statement. Since 5 is true, this block is always executed.  
}
```


b. Iterations

C supports several types of iterations with slight differences:

i) The For Loop

The most basic C iteration is the for loop. It consists of 3 parts: An initializer, a continuation condition, and a modification operation, separated by semi-colons. To count from 0 to 9 we would do:

```
for(i=0; i<10; i++) {  
    ... C statements ...  
}
```

Due to its flexibility, the for statement is very powerful; we can for example count downwards from 9 to 0:

```
for(i=9; i>=0; i--) {  
    ... C statements ...  
}
```

Each part is optional; If we don't want to initialize i, we can do:

```
for(; i <=9; i++) {  
    ... C statements ...  
}
```

You can also implement a loop that counts infinitely from any initial value:

```
i=5; // Some statement that sets i to 5  
  
for(;;i++) {  
    // i will start from 5 and increment indefinitely.  
}
```

If we had a string "mystr" we can count how many characters are in the string using:

```
for(ctr=0; mystr[ctr] != 0; ctr++);
```

When this for loop ends ctr will contain the number of characters in mystr.

Lastly the for loop can also be used for infinite loops, simply by leaving out every part:

```
for(;;) {  
    // This is an infinite loop  
}
```

ii) The While Loop

The while loop works similarly to the for loop, except that the “while” statement itself contains only the continuation condition; initialization and update are done separately. Our “string count” operation would be written as:

```
ctr = 0;

while(mystr[ctr] != 0) {
    ctr++; // Increment ctr
}
```

Since any non-zero value is true, we can do an infinite loop using while by doing:

```
while(1) {
    // This body executes infinitely
}
```

iii) The Do..While Loop

The do..while loop is fairly unique to C. Since the continuation condition is tested only at the end of the body, this means that the body will always be executed at least once:

```
do {
    ctr++;
} while(mystr[ctr] != 0)
```

Here ctr will be incremented as least once, even if mystr[0] is 0.

c. Conditional Statements

i) The if..else statement

The if..else statement in C is fairly straightforward:

```
if( .. condition.. ) {
    // This part is executed if the ..condition.. is true (non-zero)
} else {
    // Otherwise this part is executed.
}
```

C however does not have an explicit “elif” (else if) statement, and this must be handled with:

```
if(condition1) {  
    ...  
} else  
    if(condition2) {  
        ...  
    }  
    else  
    {  
        ... Executed if condition2 is false...  
    }  
}
```

ii) The ternary if..else statement

Just as in Javascript, we can use the ternary if..else operation:

(<condition> ? <if part> : <else part>)

So if we did:

```
y = (x > 0 ? 3 : 2);
```

The “y” variable would be set to 3 if x is greater than 0, and 2 otherwise.

iii) The switch statement:

The C Programming Language has the switch statement, which can handle multiple choices without having to deeply nest if..else statements. The statement takes the form:

```
switch(<integer variable>) {
    case 1:
        ... something ...
        break;

    case 2:
        ... something else ...
        break;

    case 3:
        .. something else again ...
        break;

    default:
        .. Every other case is handled here ..
}
```

The switch statement only works with integer variables (including “char”, which if you recall is a 1-byte integer). You also need a “break” statement after each case, to force execution to exit from the switch statement, instead of falling through all the other cases. The “default” keyword is used to catch all other cases that are not explicitly stated. You can use switch to implement a menu system:

```
choice = readMenuOption(); // Get menu option from user

switch(choice) {

    case 'a': callOptionA();
        break;

    case 'b': callOptionB();
        break;

    case 'c': callOptionC();
        break;

    default:
        printf("Don't understand. Please choose a, b or c.");
}
```

4. Pointers

The idea of pointer variables is arguably the most difficult concept in C to grasp. Every memory location in a computer is indexed with an address (also called a “memory location”). All variables in C must be stored in memory, and a pointer variable simply stores the address of another variable.

So if we had:

```
int *x;
```

This declares a pointer variable *x*, and it will point to another variable of type *int*.

Let’s suppose variable *y* is stored in location 0x5004. The figure below shows memory contents from addresses 0x5000 to 0x5008.

Address	Value	Description
0x5000	55	Variable z
0x5001		
0x5002		
0x5003		
0x5004	72	Variable y
0x5005		
0x5006		
0x5007		
0x5008		
...

Notice firstly that every memory location has an “address”, and stores “values”. The Description column tells us what the memory holds, but is not actually stored anywhere – it is there just for our information. Also notice that each variable occupies four memory locations; this is because each *int* is 32-bits long, which is four bytes, and each address refers to an individual byte in memory (byte-addressable memory).

Now if we did:

```
x = &y;
```

The “&” operator, called the “address-of” operator, returns the address of *y*. This is stored into *x*, which is a pointer variable and stores addresses of other variables. Now if we did:

```
printf(“x is %x”, x);
```

The program will tell us 0x5004, which is the address of *y*.

Now what can we do with this? We could make two address variables point to exactly the same variable:

```
int *a;  
  
a = x; // Now both a and x point to y.
```

This is very useful in “call by reference” parameter passing which we will see shortly.

Now that x (and a) point to y, can we use them to access (and change) y’s value? Yes, using the “de-referencing” operation.

In a fit of ingenuity, the designers of C decided to use the same “*” operation (the one used to declare pointer variables) to also access the value pointed to.

So if we did:

```
printf("The value pointed to by x is %d.", *x);
```

We would get:

The value pointed to by x is 72.

This is the value that is in y.

Now remember that a is also pointing to y. If we do:

```
*a = 123;
```

Our memory will now look like this:

Address	Value	Description
0x5000	55	Variable z
0x5001		
0x5002		
0x5003		
0x5004	123	Variable y
0x5005		
0x5006		
0x5007		
0x5008		
...

This is because C uses the “*” de-referencing operator to access a, getting the address 0x5004, then going to that address and storing 123.

5. Functions

As in other languages, a function in C is a unit that takes inputs, performs some sort of transformation on the input and produces an output. However as with other languages, a function may not necessary take inputs (called arguments) and may not necessary produce a value.

A C function is declared as follows:

```
<return type> <function name>(<param1 type><param1>, <param2 type>
    <param2>, ...) {
    .. function body ..
}
```

The function below returns the sum of two integers:

```
int sum(int a, int b) {
    return a + b;
}
```

The function below halves the argument:

```
float half(int a) {
    return a / 2.0;
}
```

The function below doesn't take any arguments nor does it return a value:

```
void say() {
    printf("Say what?\n");
}
```

5.1 Call by Value vs. Call by Pointer


All arguments are “passed by value” to functions (i.e. in C, all function arguments are “call by value”). To understand what this means, let’s look at our earlier sum function:

```
int sum(int a, int b) {  
    return a + b;  
}
```

Now let’s declare two variables x and y and call sum with them:


```
int x = 5, y = 6;  
  
z = sum(x, y);
```

The diagram below shows what happens:




Here 5 is copied from the argument x into the parameter a, and 6 is copied from the argument y into the parameter b. This means that a and b are second copies of x and y.

This has implications. Supposed we modify a within the body:



The picture below shows what happens:



Notice that while a is changed to 10, x remains as 5. Therefore we cannot use this as a means of passing back values through the parameters.

This is where call-by-pointer comes in. We rewrite our function as:

```
int sum(int *a, int *b) {  
    *a = 10;  
    return *a + * b;  
}
```

Now we call it with:

```
sum(&x, &y);
```

What happens now is that the ADDRESSES of x and y are copied “by value” into a and b. When we de-reference a (which points to x) in this statement:

```
*a = 10;    // a contains address of x
```

C will look at the address in a, go to that address and write in 10. Since a is pointing to x, this causes x to be modified to 10.

6. Conclusion

This is a very short brief introduction to the C programming language. It is not sufficient to get you started in programming C, but should be enough for you to understand the lectures.