

CS 521: Systems Programming

Arrays and Strings

Today's Schedule

- Set up and Discuss Lab 3
- Arrays
- Strings

Today's Schedule

- **Set up and Discuss Lab 3**
- Arrays
- Strings

Setting up...

...and touring the code

Make

- All 521 projects will include a `Makefile`
 - This tells the `make` utility what to do
 - Essentially just a recipe for building your program
 - All our projects are composed of several C files
- Hints (applies to Lab 3 as well):
 - `make` – compile and produce executable
 - `make test` – run the test cases
 - `make clean` – clean up all build artifacts

How make Works

Makefiles are a recipe composed of instructions like this:

```
target: dependency
    instructions
```

(Note the tabs)

- You provide a target, like ‘array’ – the name of the file that gcc will generate
- The dependency tells us what files you need to build your program. In this case, it’s ‘array.c’

Today's Schedule

- Set up and Discuss Lab 3
- **Arrays**
- Strings

Arrays

- In C, arrays let us store a collection of values of the same type
 - `int list[10];`
 - `double dlist[15];`
- Internally, they are represented as a chunk of memory large enough to fit all the required elements
- Note that the arrays must be dimensioned when they're declared
 - In older versions of C the dimension had to be a constant

Accessing Array Elements

- Setting/retrieving the values of an array is the same as it is in Java:
 - `list[2] = 7;`
 - `list[1] = list[2] + 3;`
- However, one interesting note about C is there is **no boundary checking**, so:
 - `list[500] = 7;`
 - ...may work just fine.
 - 😬

Experiment: When will it Break?

- We can try modifying out-of-bounds array elements
- We can even do it in a loop to test the limits
 - Different operating systems / architectures may react differently
 - Let's try it now. Open your editor, create an array, and write a loop that iterates beyond its boundaries.
 - When does it segfault? How big was your initial array?
- At this point, you might be wondering:
 - What is wrong with C?!
 - What is the meaning of life?

Out-of-bounds Access

- So we can do things like this in C:
 - `int list[5];`
 - `list[10] = 7;`
- Your program may work fine... or crash!
- It's never a good idea to do this
- So why does C let us do it anyway?

Safety vs. Performance

- C favors performance over safety
 - Compare: C program vs Python equivalent
 - Helpful: time command
- Especially in the glory days of C, adding lots of extra checks meant poor performance
 - Additional instructions for those checks
 - If you don't want/need them, then the language shouldn't force it on you!
- This can lead to dangerous bugs

Initializing an Array [1/2]

- Let's create our list of integers:
 - `int list[10];`
- When we do this, C sets aside a place in memory for the array
 - It doesn't clear the memory **unless we ask it to**
 - Another common cause of subtle bugs
- Creating a list of integers initialized to zero:
 - `int list[10] = { 0 };`

Initializing an Array [2/2]

Thus far we've always specified the array size. There is a shorthand for doing this if you already know the contents of the array:

```
// Will auto-size to 5:  
int nums[] = { 1, 82, 9, -3, 26 };
```

Here, the compiler will fill in the size for you.

Memory Access


- What happens when we retrieve the value of `list[5]` ?
- Find the location of `list` in memory
- Move to the proper offset: $5 * 4 = \text{byte } 20$
 - Assuming `sizeof(int) = 4`
- Access the value
- Accessing, say, `list[500]` is just moving to a position in memory and retrieving whatever is there

Visualizing Arrays in Memory

```
/* Note: calculating the array  
 * dimensions automatically! */
```

```
int list[] = {  
    1,  
    2,  
    15,  
    2001  
};
```

```
sizeof(int) = 4
```



The sizeof Operator

- We can use the sizeof operator in C to determine how big things are
 - Somewhat like:
 - len() in python
 - .length in Java, or
 - .size() in Java
- Much more low-level
 - `size_t sz = sizeof(int);`
 - `printf("%zd\n", sz); // Prints 4 (on my machine)`

Array Size [1/2]

- Let's try this out:
 - `int list[10];`
 - `size_t list_sz = sizeof(list);`
- Any guesses on the output?
- On my machine, it's 40:
 - 40 bytes (10 integers at 4 bytes each)
 - This can be different depending on architecture
- In C, `sizeof(char)` is guaranteed to be 1.

Array Size [2/2]

- Knowing the number of bytes in the array can be useful, but not that useful
- Usually we want to know how many elements there are in an array
- To do this, we'll divide by the array **type** (int - 4 bytes):
 - `int list[10];`
 - `size_t list_sz = sizeof(list) / sizeof(list[0]);`
 - `printf("%zd\n", list_sz); /* 10 (for me) */`

Behind the Scenes

- Arrays in C are actually (constant) pointers
 - `int list[5];`
 - `list` is the same as `&list[0];`
- You can't change what they point at, but otherwise they work the same
- So accessing `list[2]` is really just dereferencing a pointer that points two memory addresses from the start of the array
 - ...one reason we have 0-based arrays

We can make this more “fun...”

- Since arrays are just constant pointers, we have another way to access them:
 - `list[5]` is the same thing as: `*(list + 5)`
- Workflow:
 1. Locate the start of the array
 2. Move up 5 memory locations (4 bytes each*)
 3. Dereference the pointer to get our value


Pointer Arithmetic

- Manipulating pointers in this way is called **pointer arithmetic**
- `arr[i];` is the same thing as: `*(arr + i);`
- `arr[6] = 42;` is the same as `*(arr + 6) = 42;`

Visualizing Arrays with Pointer Arithmetic

```
int list[] = {  
    1,  
    2,  
    15,  
    2001  
};
```

```
sizeof(int) = 4
```



A Note on Pointer Arithmetic

- In general, stick with using regular array syntax
- You may see pointer arithmetic in production code, but it should only be used in situations that make the code **more** understandable
- Haphazardly showing off your knowledge of pointer arithmetic is a recipe for confusing code 😊

Arrays as Function Arguments

- When we pass an array to a function, its pointer-based underpinnings begin to show
- If we modify an array element inside a function, will the change be reflected in the calling function?
 - ...
 - ...why?
- In fact, when an array is passed to a function it **decays** to a pointer
 - The function just receives a pointer to the first element in the array. That's it!

Array Decay

- When an array decays to a pointer, we lose its dimension information
- Let's imagine someone just gives us a pointer
 - Do we know if it points to a single value?
 - Is it the start of an array?
- Functions are in the same situation: they don't know where this pointer came from or where it's been
 - `sizeof()` doesn't work as expected

Dealing with Decay

- Array dimensions are often very useful information!
 - If we don't know how many elements are in the array, then we could read/write beyond the end of it
- There are two viable strategies to deal with this:
 1. Pass the size of the array into the function as an argument
 2. Put some kind of identifier at the end of the array so we know where it ends as we iterate through

Today's Schedule

- Set up and Discuss Lab 3
- Arrays
- **Strings**

C Strings

- In C, strings are just *special* arrays of characters:
 - `char str[] = "Hello World!"; // Mutable (array)`
 - `char *str = "Hello World!"; // Immutable (str. literal)`
- You can't see it, but the reason these character arrays are special is because they end in `\0`
 - The `NUL` terminator
- As we already discussed, we need to either pass dimensions along with arrays **OR** include some way of knowing where they end
 - `\0` means “the end!”

Strings as Arrays

- Let's look at C strings:

“HELLO!” → The diagram shows the string "HELLO!" being converted into an array of characters. Each character is enclosed in a light blue square box with a darker blue border. The characters are H, E, L, L, O, !, and \0. An arrow points from the string to the first box containing 'H'.

- Note how our string contains 6 characters, but the array representation has 7 due to the `NUL` byte
- `\0` is a control character
 - Just like `\n`, etc., we write it with two characters but it is just shorthand for a single character
 - Its value also happens to be 0 (decimal)
 - C string functions assume this is present; if it's not, you only have an **array of characters** and your program will crash

Some C String Library Functions

- `#include <string.h>`
- `strcpy` – copy one string to another
- `strcat` – concatenate two strings
- `strcmp` – test for string equality
- `strlen` – returns the length of the string (ignoring `\0`)
- `strtok` – tokenize the string (split it up)
- Remember: documentation available in the `man` pages

Avoiding Buffer Overruns

- The string functions you just saw have one weakness...
 - If they lack the `\0`, they break!
 - This can lead to bugs, crashes, or even security issues
- Most C string functions also have a version that allows you to specify a fixed length
 - `strncmp`, `strncpy`, etc.
 - Notice the `n`: `strNcpy`
- If you can use one of these functions instead, it is recommended to do so

Copying a String [1/3]

- Let's say you want to copy one string into another:

- `char str1[] = "Hello World!";`

- `char *str2 = str1;`

- This doesn't make a copy; it just points to `str1`

- What about:

- `char str2[] = str1;`

- Nope:

```
error: array initializer must be an initializer
list or string literal
```

Copying a String [2/3]

- We could loop through the array and copy each character into the other, but that's a lot of work
- Better solution: strcpy
 - (let's take a quick peek at the man page)

```
char str1[] = "Hello World!";  
char str2[12];  
strcpy(str2, str1);  
printf("%s\n", str2);
```

But wait... This code has a big problem: array size

Copying a String [3/3]

Let's fix our bug:

```
char str1[] = "Hello World!";  
char str2[13];  
strcpy(str2, str1);  
printf("%s\n", str2);
```

We could also create a much larger array to copy into.

Reading a String With scanf

Let's greet the user:

```
char str[100];  
printf("Enter your name: ");  
scanf("%s", str);  
printf("Hi, %s!\n", str);
```

- Wait a minute! Where's our &?!
- Well, remember that when we see the `[]` brackets, we're grabbing the actual values (via dereference)
 - The array name only = a pointer to the first element

Getting String Lengths

We can use the `strlen` function to find out how many characters (not including the `\0`) are in a string:

```
char str[] = "Hello";  
printf("Length = %zd\n", strlen(str));
```

Concatenating a String

`strcat` (and `strncat`) concatenate strings:

```
char *strcat(char *dest, const char *src);
```

```
char x[128] = "Hello";  
char *y = "World";  
  
strcat(x, " ");  
strcat(x, y);  
strcat(x, "!");  
  
printf("%s\n", x);
```

Be careful: *dest* must be initialized before using `strcat` !

Concatenation: Another Option

You can use `printf`-style format specifiers to combine strings with `sprintf` and `snprintf`:

```
char a[128];  
  
char x[] = "Hello";  
char *y = "World";  
sprintf(x, "%s %s!", x, y);
```

Here, you're basically “printing” to a string.

More String Functions

- This is only a very light introduction to C strings.
- We'll be working with them quite a bit more this semester!