

CS107, Lecture 7

C Strings, Buffer Overflows and Security

Reading: K&R (1.6, 5.5, Appendix B3) or Essential
C section 3

This document is copyright (C) Stanford Computer Science and Nick Troccoli, licensed under
Creative Commons Attribution 2.5 License. All rights reserved.

Based on slides created by Cynthia Lee, Chris Gregg, Jerry Cain, Lisa Yan and others.

NOTICE RE UPLOADING TO WEBSITES: This content is protected and may not be shared,
uploaded, or distributed. (without expressed written permission)

CS107 Topic 2

How can a computer represent and manipulate more complex data like text?

Why is answering this question important?

- Shows us how strings are represented in C and other languages (last time)
- Helps us better understand buffer overflows, a common bug (this time)
- Introduces us to pointers, because strings can be pointers (next time)

assign2: implement 2 functions a 1 program using those functions to find the location of different built-in commands in the filesystem. You'll write functions to extract a list of possible locations and tokenize that list of locations.

Learning Goals

- Understand how to use the built-in string functions for common string tasks
- Learn more about the risks of buffer overflows and how to mitigate them

Lecture Plan

- **Recap:** Strings so far
- Searching in Strings
- **Practice:** Password Verification
- Buffer Overflows and Security

```
cp -r /afs/ir/class/cs107/lecture-code/lect7 .
```

Lecture Plan

- **Recap: Strings so far**
- Searching in Strings
- **Practice:** Password Verification
- Buffer Overflows and Security

```
cp -r /afs/ir/class/cs107/lecture-code/lect7 .
```

C Strings

C strings are arrays of characters ending with a **null-terminating character** `'\0'`.

<i>index</i>	0	1	2	3	4	5	6	7	8	9	10	11	12	13
<i>value</i>	'H'	'e'	'l'	'l'	'o'	','	' '	'w'	'o'	'r'	'l'	'd'	'!'	'\0'

String operations such as `strlen` use the null-terminating character to find the end of the string.

Side note: use `strlen` to get the length of a string. Don't use `sizeof`!

Common string.h Functions

Function	Description
<code>strlen(<i>str</i>)</code>	returns the # of chars in a C string (before null-terminating character).
<code>strcmp(<i>str1</i>, <i>str2</i>)</code> , <code>strncmp(<i>str1</i>, <i>str2</i>, <i>n</i>)</code>	compares two strings; returns 0 if identical, <0 if <i>str1</i> comes before <i>str2</i> in alphabet, >0 if <i>str1</i> comes after <i>str2</i> in alphabet. <i>strncmp</i> stops comparing after at most <i>n</i> characters.
<code>strchr(<i>str</i>, <i>ch</i>)</code> <code>strrchr(<i>str</i>, <i>ch</i>)</code>	character search: returns a pointer to the first occurrence of <i>ch</i> in <i>str</i> , or <i>NULL</i> if <i>ch</i> was not found in <i>str</i> . <code>strrchr</code> find the last occurrence.
<code>strstr(<i>haystack</i>, <i>needle</i>)</code>	string search: returns a pointer to the start of the first occurrence of <i>needle</i> in <i>haystack</i> , or <i>NULL</i> if <i>needle</i> was not found in <i>haystack</i> .
<code>strcpy(<i>dst</i>, <i>src</i>)</code> , <code>strncpy(<i>dst</i>, <i>src</i>, <i>n</i>)</code>	copies characters in <i>src</i> to <i>dst</i> , including null-terminating character. Assumes enough space in <i>dst</i> . Strings must not overlap. <i>strncpy</i> stops after at most <i>n</i> chars, and <u>does not</u> add null-terminating char.
<code>strcat(<i>dst</i>, <i>src</i>)</code> , <code>strncat(<i>dst</i>, <i>src</i>, <i>n</i>)</code>	concatenate <i>src</i> onto the end of <i>dst</i> . <i>strncat</i> stops concatenating after at most <i>n</i> characters. <u>Always</u> adds a null-terminating character.
<code>strspn(<i>str</i>, <i>accept</i>)</code> , <code>strcspn(<i>str</i>, <i>reject</i>)</code>	<i>strspn</i> returns the length of the initial part of <i>str</i> which contains <u>only</u> characters in <i>accept</i> . <i>strcspn</i> returns the length of the initial part of <i>str</i> which does <u>not</u> contain any characters in <i>reject</i> .

Substrings

We can combine pointer arithmetic and copying to make any substrings we'd like.

```
// Want just "ace"
char str1[8];
strcpy(str1, "racecar");

char str2[4];
strncpy(str2, str1 + 1, 3);
str2[3] = '\0';
printf("%s\n", str1);           // racecar
printf("%s\n", str2);          // ace
```

char * vs. char[]

What kind string do I want to create?

- To create a string with space for new characters -> create a **char[]**
 - space on the stack for actual characters
- To create a string without making stack space for new characters -> **char ***
 - Space for 8 byte pointer – not making space for new characters on the stack

```
char str1[8]; // Space for new chars
strcpy(str1, "racecar");
char *otherstr = str1 + 4; // points to existing chars
```

C Strings As Parameters

When we pass a string as a parameter, it is passed as a **char *** (pointer to its first character). We can still operate on the string the same way as with a `char[]`.

```
int doSomething(char *str) {  
    char secondChar = str[1];  
    ...  
}
```

```
// can also write this, but it is really a pointer  
int doSomething(char str[]) { ...
```

char * vs. char[]

You can reassign char * pointers like any other variable type:

```
char myString[6];  
strcpy(myString, "Hello");  
char *myOtherString = myString;  
myOtherString = myString + 2;           // ok
```

You cannot reassign arrays (but you can change their contents):

```
char myString[6];  
char otherString[6];  
strcpy(myString, "Hello");  
otherString = myString;           // not allowed  
strcpy(otherString, myString);       // ok
```

Lecture Plan

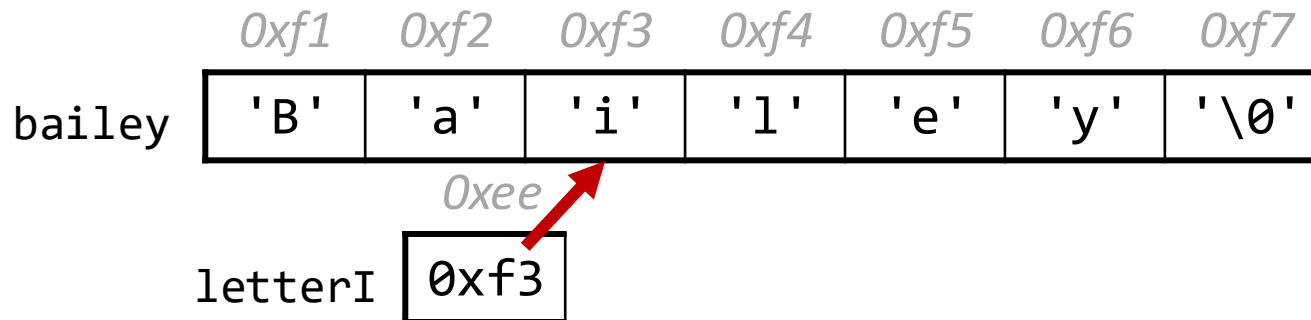
- **Recap:** Strings so far
- **Searching in Strings**
- **Practice:** Password Verification
- Buffer Overflows and Security

```
cp -r /afs/ir/class/cs107/lecture-code/lect7 .
```

Searching For Letters

`strchr` returns a pointer to the first occurrence of a character in a string, or `NULL` if the character is not in the string.

```
char bailey[7];  
strcpy(bailey, "Bailey");  
char *letterI = strchr(bailey, 'i');  
printf("%s\n", bailey);           // Bailey  
printf("%s\n", letterI);         // iley
```

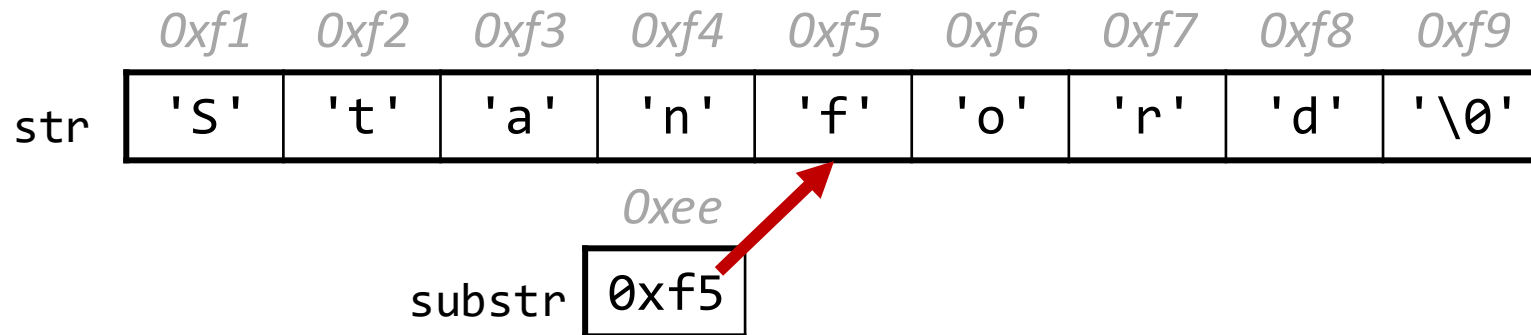


For multiple occurrences, `strchr` returns a pointer to the *first* one. Use `strrchr` to obtain a pointer to the *last* occurrence.

Searching For Strings

`strstr` returns a pointer to the first occurrence of the second string in the first, or `NULL` if it cannot be found.

```
char str[9];  
strcpy(str, "Stanford");  
char *substr = strstr(str, "for");  
printf("%s\n", str);           // Stanford  
printf("%s\n", substr);       // ford
```



If there are multiple occurrences of the string, `strstr` returns a pointer to the *first* one.

String Spans

`strspn` returns the *length* of the initial part of the first string which contains only characters in the second string. (useful for chopping strings up, checking characters used).

```
char bailey[11];  
strcpy(bailey, "Bailey Dog");  
int spanLength = strspn(bailey, "aBeoi");           // 3
```

“How many places can we go in the first string before I encounter a character not in the second string?”

String Spans

`strcspn` (c = “complement”) returns the *length* of the initial part of the first string which contains only characters not in the second string.

```
char bailey[11];  
strcpy(bailey, "Bailey Dog");  
int spanLength = strcspn(bailey, "driso"); // 2
```

“How many places can we go in the first string before I encounter a character in the second string?”

str[c]spn vs. strstr

strspn/strcspn can't search for substrings because it does not pay attention to the order of the characters in the second string. **strstr** allows us to search for substrings within another string.

// these are all equivalent

```
int spanLength = strcspn(bailey, "driso");
```

```
int spanLength = strcspn(bailey, "sirdo");
```

```
int spanLength = strcspn(bailey, "odris");
```

Arrays of Strings

We can make an array of strings to group multiple strings together:

```
char *stringArray[5]; // space to store 5 char *s
```

We can also use the following shorthand to initialize a string array:

```
char *stringArray[] = {  
    "Hello",  
    "Hi",  
    "Hey there"  
};
```

Arrays of Strings

We can access each string using bracket syntax:

```
printf("%s\n", stringArray[0]); // print out first string
```

When an array of any type is passed as a parameter in C, C passes a *pointer to the first element of the array*. This is what **argv** is in **main**!

```
void myFunction(char *stringArray[]) {  
// is really this (both are equivalent):  
void myFunction(char **stringArray) {
```

Practice: Password Verification

Write a function **verifyPassword** that accepts a candidate password and certain password criteria and returns whether the password is valid.

```
bool verifyPassword(char *password, char *validChars, char  
*badSubstrings[], int numBadSubstrings);
```

password is valid if it contains only letters in **validChars**, and does not contain any substrings in **badSubstrings**.

Practice: Password Verification

```
bool verifyPassword(char *password, char *validChars, char  
*badSubstrings[], int numBadSubstrings);
```

Example:

```
char *invalidSubstrings[] = { "1234" };
```

```
bool valid1 = verifyPassword("1572", "0123456789",  
    invalidSubstrings, 1);    // true
```

```
bool valid2 = verifyPassword("141234", "0123456789",  
    invalidSubstrings, 1);    // false
```

Practice: Password Verification



```
verify_password.c
```

Lecture Plan


- **Recap:** Strings so far
- Searching in Strings
- **Practice:** Password Verification
- **Buffer Overflows and Security**

```
cp -r /afs/ir/class/cs107/lecture-code/lect7 .
```

Recall: Buffer Overflows

We must make sure there is enough space in the destination to hold the entire copy, *including the null-terminating character*. Writing past memory bounds is called a “buffer overflow”. It can allow for security vulnerabilities!

```
char str1[14];  
strcpy(str1, "hello, world!");  
char str2[6];  
strcpy(str2, str1); // not enough space - overwrites other memory!
```



Buffer Overflow Impacts

Buffer overflows are not merely functionality bugs; they can cause a range of unintended behavior:

- Let the user access memory they shouldn't be able to access
- Let the user modify memory they shouldn't be able to access
 - Change a value that is used later in the program
 - User changes the program to execute their own custom instructions instead
 - And more...

It's our job as programmers to find and fix buffer overflows and other bugs not just for the functional correctness of our programs, but to protect people who use and interact with our code.

Buffer Overflow Example: ./buf

```
// first argument is the name, second is the password
int main(int argc, char *argv[]) {
    char secret[4] = "123";
    // assume secret comes right after name in memory
    // (this is not always true)
    char name[4];
    strcpy(name, argv[1]);

    if (strcmp(secret, argv[2]) == 0) {
        printf("You're in!\n");
    }
    return 0;
}
```

<i>name</i>				<i>secret</i>			
?	?	?	?	'1'	'2'	'3'	'\0'

Which of these arguments would cause the program to print "You're in!"?

- A. ./buf abcdefg efg
- B. ./buf abcd abcd
- C. ./buf a a
- D. ./buf abcdefgh abcd



Respond on PollEv:
pollev.com/cs107



L7. Buffer Overflow Example

0 surveys completed



0 surveys underway

Which line in this example is problematic?

The creation of the secret string

The strcpy function call

The strcmp function call

The printf function call

Which of these arguments would cause the program to print "You're in!"?

A) `./buf abcdefg efg`

B) `./buf abcd abcd`


C) `./buf a a`

D) `./buf abcdefg abcd`

Buffer Overflow Example: ./buf

```
// first argument is the name, second is the password
int main(int argc, char *argv[]) {
    char secret[4] = "123";
    // assume secret comes right after name in memory
    // (this is not always true)
    char name[4];
    strcpy(name, argv[1]);

    if (!strcmp(secret, argv[2])) {
        printf("You're in!\n");
    }
    return 0;
}
```




Which of these arguments would cause the program to print "You're in!"?

- A. **./buf abcdefg efg**
- B. ./buf abcd abcd
- C. ./buf a a
- D. ./buf abcdefgh abcd

Buffer Overflow Example: ./buf

```
// first argument is the name, second is the password
int main(int argc, char *argv[]) {
    char secret[4] = "123";
    // assume secret comes right after name in memory
    // (this is not always true)
    char name[4];
    strcpy(name, argv[1]);

    if (!strcmp(secret, argv[2])) {
        printf("You're in!\n");
    }
    return 0;
}
```



Which of these arguments would cause the program to print "You're in!"?

- A. **./buf abcdefg efg**
- B. ./buf abcd abcd
- C. ./buf a a
- D. ./buf abcdefgh abcd

Buffer Overflow Impacts

- AOL instant messenger buffer overflow: allowed remote attackers to execute code: <https://www.cvedetails.com/cve/CVE-2002-0362/>
- Morris Worm: first internet worm to gain widespread attention; exploited buffer overflow in Unix command called "finger":
<https://www.zdnet.com/article/the-morris-worm-internet-malware-turns-25/>

How can we fix buffer overflows?

There's no single solution to fix all buffer overflows; instead, it's a combination of techniques to avoid them as much as possible:

- Constant vigilance while programming (checking arrays and where they are modified)
- Carefully reading documentation
- Thorough testing to uncover issues before release
- Thorough documentation to document assumptions in your code
- (Where possible) use of tools that reduce the possibility for buffer overflows

How can we fix buffer overflows?

There's no single solution to fix all buffer overflows; instead, it's a combination of techniques to avoid them as much as possible:

- Constant vigilance while programming (checking arrays and where they are modified)
- **Carefully reading documentation**
- Thorough testing to uncover issues before release
- Thorough documentation to document assumptions in your code
- (Where possible) use of tools that reduce the possibility for buffer overflows

How Can We Fix Overflows?

Documentation & MAN Pages (Written by Others)

“The strcpy() function copies the string pointed to by src, including the terminating null byte ('\0'), to the buffer pointed to by dest. The strings may not overlap, and the destination string dest must be large enough to receive the copy. Beware of buffer overruns! (See BUGS.) ...

BUGS

If the destination string of a strcpy() is not large enough, then anything might happen. **Overflowing fixed-length string buffers is a favorite cracker technique for taking complete control of the machine.** Any time a program reads or copies data into a buffer, the program first needs to check that there's enough space. **This may be unnecessary if you can show that overflow is impossible, but be careful: programs can get changed over time, in ways that may make the impossible possible.”**

How can we fix buffer overflows?

There's no single solution to fix all buffer overflows; instead, it's a combination of techniques to avoid them as much as possible:

- Constant vigilance while programming (checking arrays and where they are modified)
- Carefully reading documentation
- **Thorough testing to uncover issues before release**
- Thorough documentation to document assumptions in your code
- (Where possible) use of tools that reduce the possibility for buffer overflows

How Can We Fix Overflows?

- **Valgrind:** Your Greatest Ally
- Write your own tests
- Consider writing tests *before* writing the main program

 cs107.stanford.edu/testing.html 

On future assignments, you'll get practice writing some tests prior to writing code, and adding additional tests as you write code.

- Makes it easier to write tests entirely based on intended behavior, without being influenced by the actual implementation.

How can we fix buffer overflows?

There's no single solution to fix all buffer overflows; instead, it's a combination of techniques to avoid them as much as possible:

- Constant vigilance while programming (checking arrays and where they are modified)
- Carefully reading documentation
- Thorough testing to uncover issues before release
- **Thorough documentation to document assumptions in your code**
- (Where possible) use of tools that reduce the possibility for buffer overflows

Code Documentation

```
/* Function: myFunction
 * This function assumes that the provided string is
 * at most length 10.
 * ...
 */
void myFunction(char *str) {
    ...
}
```

How can we fix buffer overflows?

There's no single solution to fix all buffer overflows; instead, it's a combination of techniques to avoid them as much as possible:

- Constant vigilance while programming (checking arrays and where they are modified)
- Carefully reading documentation
- Thorough testing to uncover issues before release
- Thorough documentation to document assumptions in your code
- **(Where possible) use of tools that reduce the possibility for buffer overflows**

Memory Safe Systems Programming

Idea 5: Choose your Tools & Languages Carefully


Existing code bases or requirements for a project may dictate what tools you use. Knowing C is crucial – it is and will remain widely used.

When you you are choosing tools for systems programming, consider languages that can help guard against programmer error.

- Rust (Mozilla)
- Go (Google)
- Project Verona (Microsoft)

Memory Safe Systems Programming

“Memory safety vulnerabilities are a class of vulnerability affecting how memory can be accessed, written, allocated, or deallocated in unintended ways. Experts have identified a few programming languages that both lack traits associated with memory safety and also have high proliferation across critical systems, such as C and C++. Choosing to use memory safe programming languages at the outset, as recommended by the Cybersecurity and Infrastructure Security Agency’s (CISA) Open-Source Software Security Roadmap is one example of developing software in a secure-by-design manner.”



Association for Computing Machinery (ACM) Code of Ethics

ACM Code of Ethics and Professional Conduct

ACM Code of Ethics and Professional Conduct

Preamble

Computing professionals' actions change the world. To act responsibly, they should reflect upon the wider impacts of their work, consistently supporting the public good. The ACM Code of Ethics and Professional Conduct ("the Code") expresses the conscience of the profession.

The Code is designed to inspire and guide the ethical conduct of all computing professionals, including current and aspiring practitioners, instructors, students, influencers, and anyone who uses computing technology in an impactful way. Additionally, the Code serves as a basis for remediation when violations occur. The Code includes principles formulated as statements of responsibility, based on the understanding that the public good is always the primary consideration. Each principle is supplemented by guidelines, which provide explanations to assist computing professionals in understanding and applying the principle.

<https://www.acm.org/code-of-ethics>

On This Page

Preamble

1. GENERAL

1.1 Contribute to the well-being, and respect the rights of all stakeholders.

1.2 Avoid conflicts of interest.

1.3 Be honest and trustworthy.

1.4 Be fair and avoid discrimination.

1.5 Respect the privacy and intellectual property of others, and promote new ideas.

ACM Code of Ethics on Security

2.9 Design and implement systems that are robustly and usably secure.

Breaches of computer security cause harm. Robust security should be a primary consideration when designing and implementing systems. Computing professionals should perform due diligence to ensure the system functions as intended, and take appropriate action to secure resources against accidental and intentional misuse, modification, and denial of service. As threats can arise and change after a system is deployed, computing professionals should integrate mitigation techniques and policies, such as monitoring, patching, and vulnerability reporting. Computing professionals should also take steps to ensure parties affected by data breaches are notified in a timely and clear manner, providing appropriate guidance and remediation.

To ensure the system achieves its intended purpose, security features should be designed to be as intuitive and easy to use as possible. Computing professionals should discourage security precautions that are too confusing, are situationally inappropriate, or otherwise inhibit legitimate use.

In cases where misuse or harm are predictable or unavoidable, the best option may be to not implement the system.

Buffer Overflows

- We must always ensure that memory operations we perform don't improperly read or write memory.
 - E.g. don't copy a string into a space that is too small!
 - E.g. don't ask for the string length of an uninitialized string!
- The **Valgrind** tool is a memory checker tool we'll introduce next time that can help track down (some, but not all) memory-related issues.

Recap

- **Recap:** Strings so far
- Searching in Strings
- **Practice:** Password Verification
- Buffer Overflows and Security

Lecture 7 takeaway: string searching functions allow us to search for characters, substrings, and spans. Buffer overflows can cause significant functionality and security issues!

```
cp -r /afs/ir/class/cs107/lecture-code/lect7 .
```

Extra Practice

2. Code study: strncpy

STRCPY(3)

Linux Programmer's Manual

STRCPY(3)


DESCRIPTION

The `strncpy()` function is similar, except that at most `n` bytes of `src` are copied. **Warning:** If there is no null byte among the first `n` bytes of `src`, the string placed in `dest` will not be null-terminated.

If the length of `src` is less than `n`, `strncpy()` writes additional null bytes to `dest` to ensure that a total of `n` bytes are written.

A simple implementation of `strncpy()` might be:

```
1 char *strncpy(char *dest, const char *src, size_t n) {
2     size_t i;
3     for (i = 0; i < n && src[i] != '\0'; i++)
4         dest[i] = src[i];
5     for ( ; i < n; i++)
6         dest[i] = '\0';
7     return dest;
8 }
```



What happens if we call `strncpy(buf, str, 5);`?



2. Code study: strncpy

STRCPY(3)

Linux Programmer's Manual

STRCPY(3)

DESCRIPTION

The `strncpy()` function is similar, except that at most `n` bytes of `src` are copied. **Warning:** If there is no null byte among the first `n` bytes of `src`, the string placed in `dest` will not be null-terminated.

If the length of `src` is less than `n`, `strncpy()` writes additional null bytes to `dest` to ensure that a total of `n` bytes are written.

A simple implementation of `strncpy()` might be:

```
1 char *strncpy(char *dest, const char *src, size_t n) {
2     size_t i;
3     for (i = 0; i < n && src[i] != '\0'; i++)
4         dest[i] = src[i];
5     for ( ; i < n; i++)
6         dest[i] = '\0';
7     return dest;
8 }
```

	0x60	0x61	0x62	0x63	0x64	0x65	0x66
buf	'M'	'o'	'n'	'd'	'a'	'y'	'\0'

	0x58	0x59	0x5a	0x5b
str	'F'	'r'	'i'	'\0'

dest	
src	
n	5
i	

What happens if we call `strncpy(buf, str, 5);`?