

15-213

"The course that gives CMU its Zip!"

Machine-Level Programming V: Miscellaneous Topics

Feb 10, 2004

Topics

- Linux Memory Layout
- Understanding Pointers
- Buffer Overflow
- Floating Point Code

class09.ppt

Linux Memory Layout

Red Hat v. 6.2 ~1920MB memory limit

Upper 2 hex digits of address

- Stack**
 - Runtime stack (8MB limit)
- Heap**
 - Dynamically allocated storage
 - When call malloc, calloc, new
- DLLs**
 - Dynamically Linked Libraries
 - Library routines (e.g., printf, malloc)
 - Linked into object code when loaded
- Data**
 - Statically allocated data
 - E.g., arrays & strings declared in code
- Text**
 - Executable machine instructions
 - Read-only

- 2 - 15-213, S'04

Linux Memory Allocation

Initially Linked Some Heap More Heap

- 3 - 15-213, S'04

Text & Stack Example

```
(gdb) break main
(gdb) run
Breakpoint 1, 0x804856f in main ()
(gdb) print $esp
$3 = (void *) 0xbffffc78
```

Main

- Address 0x804856f should be read 0x0804856f

Stack

- Address 0xbffffc78

Initially

- 4 - 15-213, S'04

Dynamic Linking Example

```
(gdb) print malloc
$1 = {<text variable, no debug info>
0x8048454 <malloc>}
(gdb) run
Program exited normally.
(gdb) print malloc
$2 = {void *(unsigned int)}
0x40006240 <malloc>}
```

Initially

- Code in text segment that invokes dynamic linker
- Address 0x8048454 should be read 0x08048454

Final

- Code in DLL region

Linked

BF
↓
Stack
80
7F
40
3F
DLLs
08
Data
00
Text

- 5 - 15-213, S'04

Memory Allocation Example

```
char big_array[1<<24]; /* 16 MB */
char huge_array[1<<28]; /* 256 MB */

int beyond;
char *p1, *p2, *p3, *p4;

int useless() { return 0; }

int main()
{
  p1 = malloc(1 << 28); /* 256 MB */
  p2 = malloc(1 << 8); /* 256 B */
  p3 = malloc(1 << 28); /* 256 MB */
  p4 = malloc(1 << 8); /* 256 B */
  /* Some print statements ... */
}
```

- 6 - 15-213, S'04

Example Addresses

\$esp	0xbffffc78	→ BF
p3	0x500b5008	→ 80
p1	0x400b4008	→ 7F
Final malloc	0x40006240	→ 40
p4	0x1904a640	→ 3F
p2	0x1904a538	→ 08
beyond	0x1904a524	→ 00
big_array	0x1804a520	→ 00
huge_array	0x0804a510	→ 00
main()	0x0804856f	→ 08
useless()	0x08048560	→ 08
Initial malloc	0x08048454	→ 00

BF
Stack
↓
80
7F
Heap
40
3F
DLLs
08
Heap
Data
00
Text

- 7 - 15-213, S'04

C operators

Operators	Associativity
() [] -> .	left to right
! ~ ++ -- + - * & (type) sizeof	right to left
* / %	left to right
+ -	left to right
<< >>	left to right
< <= > >=	left to right
== !=	left to right
&	left to right
^	left to right
	left to right
&&	left to right
	left to right
?:	right to left
= += -= *= /= %= &= ^= != <<= >>=	right to left
,	left to right

Note: Unary +, -, and * have higher precedence than binary forms

- 8 - 15-213, S'04

C pointer declarations

<code>int *p</code>	p is a pointer to int
<code>int *p[13]</code>	p is an array[13] of pointer to int
<code>int *(p[13])</code>	p is an array[13] of pointer to int
<code>int **p</code>	p is a pointer to a pointer to an int
<code>int (*p)[13]</code>	p is a pointer to an array[13] of int
<code>int *f()</code>	f is a function returning a pointer to int
<code>int (*f)()</code>	f is a pointer to a function returning int
<code>int ((*f())[13])()</code>	f is a function returning ptr to an array[13] of pointers to functions returning int
<code>int ((*x[3])())[5]</code>	x is an array[3] of pointers to functions returning pointers to array[5] of ints

- 9 -

15-213, S'04

Avoiding Complex Declarations

Use Typedef to build up the decl

Instead of `int ((*x[3])())[5]`:

```
typedef int fiveints[5];
typedef fiveints* p5i;
typedef p5i (*f_of_p5is)();
f_of_p5is x[3];
```

X is an array of 3 elements, each of which is a pointer to a function returning an array of 5 ints.

- 10 -

15-213, S'04


Internet Worm and IM War

November, 1988

- Internet Worm attacks thousands of Internet hosts.
- How did it happen?

July, 1999

- Microsoft launches MSN Messenger (instant messaging system).
- Messenger clients can access popular AOL Instant Messaging Service (AIM) servers



- 11 -

15-213, S'04

Internet Worm and IM War (cont.)

August 1999

- Mysteriously, Messenger clients can no longer access AIM servers.
- Microsoft and AOL begin the IM war:
 - AOL changes server to disallow Messenger clients
 - Microsoft makes changes to clients to defeat AOL changes.
 - At least 13 such skirmishes.
- How did it happen?

The Internet Worm and AOL/Microsoft War were both based on *stack buffer overflow* exploits!

- many Unix functions do not check argument sizes.
- allows target buffers to overflow.

- 12 -

15-213, S'04

String Library Code

- Implementation of Unix function gets
 - No way to specify limit on number of characters to read

```

/* Get string from stdin */
char *gets(char *dest)
{
    int c = getc();
    char *p = dest;
    while (c != EOF && c != '\n') {
        *p++ = c;
        c = getc();
    }
    *p = '\0';
    return dest;
}
    
```

- Similar problems with other Unix functions
 - strcpy: Copies string of arbitrary length
 - scanf, fscanf, sscanf, when given %s conversion specification

- 13 -

15-213, S'04

Vulnerable Buffer Code

```

/* Echo Line */
void echo()
{
    char buf[4]; /* Way too small! */
    gets(buf);
    puts(buf);
}
    
```

```

int main()
{
    printf("Type a string:");
    echo();
    return 0;
}
    
```

- 14 -

15-213, S'04

Buffer Overflow Executions

```

unix> ./bufdemo
Type a string:123
123
    
```

```

unix> ./bufdemo
Type a string:12345
Segmentation Fault
    
```

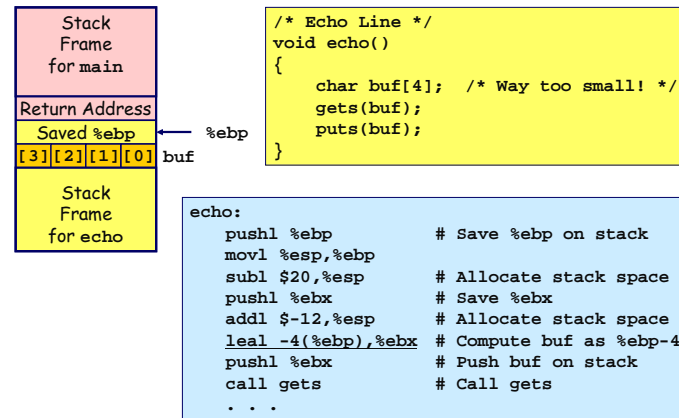
```

unix> ./bufdemo
Type a string:12345678
Segmentation Fault
    
```

- 15 -

15-213, S'04

Buffer Overflow Stack



- 16 -

15-213, S'04

Buffer Overflow Stack Example

```

unix> gdb bufdemo
(gdb) break echo
Breakpoint 1 at 0x8048583
(gdb) run
Breakpoint 1, 0x8048583 in echo ()
(gdb) print /x *(unsigned *)$ebp
$1 = 0xbffff8f8
(gdb) print /x *((unsigned *)$ebp + 1)
$3 = 0x804864d
    
```

Before call to gets

Stack Frame for main	Return Address	08 04 86 4d
	Saved %ebp	bf ff f8 f8
	[3][2][1][0] buf	xx xx xx xx
Stack Frame for echo		

8048648: call 804857c <echo>
804864d: mov 0xfffffe8(%ebp),%ebx # Return Point

- 17 - 15-213, S'04

Buffer Overflow Example #1

Before Call to gets Input = "123"

Stack Frame for main	Return Address	08 04 86 4d
	Saved %ebp	bf ff f8 f8
	[3][2][1][0] buf	00 33 32 31
Stack Frame for echo		

No Problem

- 18 - 15-213, S'04

Buffer Overflow Stack Example #2

Input = "12345"

Stack Frame for main	Return Address	08 04 86 4d
	Saved %ebp	bf ff 00 35
	[3][2][1][0] buf	34 33 32 31
Stack Frame for echo		

Saved value of %ebp set to 0xbffff0035
Bad news when later attempt to restore %ebp

```

echo code:
8048592: push %ebx
8048593: call 80483e4 <_init+0x50> # gets
8048598: mov 0xfffffe8(%ebp),%ebx
804859b: mov %ebp,%esp
804859d: pop %ebp # %ebp gets set to invalid value
804859e: ret
    
```

- 19 - 15-213, S'04

Buffer Overflow Stack Example #3

Input = "12345678"

Stack Frame for main	Return Address	08 04 86 00
	Saved %ebp	38 37 36 35
	[3][2][1][0] buf	34 33 32 31
Stack Frame for echo		

Invalid address
No longer pointing to desired return point
%ebp and return address corrupted

```

8048648: call 804857c <echo>
804864d: mov 0xfffffe8(%ebp),%ebx # Return Point
    
```

- 20 - 15-213, S'04

Malicious Use of Buffer Overflow

Stack
after call to `gets()`

```

void foo(){
  bar();
  ...
}

void bar() {
  char buf[64];
  gets(buf);
  ...
}

```

- Input string contains byte representation of executable code
- Overwrite return address with address of buffer
- When `bar()` executes `ret`, will jump to exploit code

- 21 - 15-213, S'04

Exploits Based on Buffer Overflows

Buffer overflow bugs allow remote machines to execute arbitrary code on victim machines.

Internet worm

- Early versions of the finger server (`fingerd`) used `gets()` to read the argument sent by the client:
 - `finger droh@cs.cmu.edu`
- Worm attacked `fingerd` server by sending phony argument:
 - `finger "exploit-code padding new-return-address"`
 - exploit code: executed a root shell on the victim machine with a direct TCP connection to the attacker.

- 22 -

15-213, S'04

The Internet Worm

11/2	18:24	first west coast computer infected
	19:04	ucb gateway infected
	20:00	mit attacked
	20:49	cs.utah.edu infected
	21:21	load avg reaches 5 on cs.utah.edu
	21:41	load avg reaches 7
	22:01	load avg reaches 16
	22:20	worm killed on cs.utah.edu
	22:41	cs.utah.edu reinfected, load avg 27
	22:49	cs.utah.edu shut down
	23:31	reinfected, load reaches 37

- 23 -

15-213, S'04

Exploits Based on Buffer Overflows

Buffer overflow bugs allow remote machines to execute arbitrary code on victim machines.

IM War

- AOL exploited existing buffer overflow bug in AIM clients
- exploit code: returned 4-byte signature (the bytes at some location in the AIM client) to server.
- When Microsoft changed code to match signature, AOL changed signature location.

- 24 -

15-213, S'04

Avoiding Overflow Vulnerability

```

/* Echo Line */
void echo()
{
    char buf[4]; /* Way too small! */
    fgets(buf, 4, stdin);
    puts(buf);
}
    
```

Use Library Routines that Limit String Lengths

- fgets instead of gets
- strncpy instead of strcpy
- Don't use scanf with %s conversion specification
 - Use fgets to read the string

IA32 Floating Point

History


- 8086: first computer to implement IEEE FP
 - separate 8087 FPU (floating point unit)
- 486: merged FPU and Integer Unit onto one chip

Summary

- Hardware to add, multiply, and divide
- Floating point data registers
- Various control & status registers

Floating Point Formats

- single precision (C float): 32 bits
- double precision (C double): 64 bits
- extended precision (C long double): 80 bits




FPU Data Register Stack

FPU register format (extended precision)



FPU registers

- 8 registers
- Logically forms shallow stack
- Top called %st(0)
- When push too many, bottom values disappear



FPU instructions

Large number of fp instructions and formats

- ~50 basic instruction types
- load, store, add, multiply
- sin, cos, tan, arctan, and log!

Sample instructions:

Instruction	Effect	Description
fildz	push 0.0	Load zero
flds Addr	push M[Addr]	Load single precision real
fmulx Addr	%st(0) ← %st(0)*M[Addr]	Multiply
faddp	%st(1) ← %st(0)+%st(1);pop	Add and pop

Floating Point Code Example

Compute Inner Product of Two Vectors

- Single precision arithmetic
- Common computation

```
float ipf (float x[],
          float y[],
          int n)
{
    int i;
    float result = 0.0;

    for (i = 0; i < n; i++)
    {
        result += x[i]*y[i];
    }
    return result;
}
```

```

pushl %ebp          # setup
movl %esp,%ebp
pushl %ebx

movl 8(%ebp),%ebx  # %ebx=&x
movl 12(%ebp),%ecx # %ecx=&y
movl 16(%ebp),%edx # %edx=n
fldz               # push +0.0
xorl %eax,%eax    # i=0
cmpl %edx,%eax    # if i>=n done
jge .L3

.L5:
flds (%ebx,%eax,4) # push x[i]
fmuls (%ecx,%eax,4) # st(0)*=y[i]
faddp          # st(1)+=st(0); pop
incl %eax      # i++
cmpl %edx,%eax # if i<n repeat
jl .L5

.L3:
movl -4(%ebp),%ebx # finish
movl %ebp, %esp
popl %ebp
ret                # st(0) = result
    
```

Floating Po

```

pushl %ebp          # setup
movl %esp,%ebp
pushl %ebx

movl 8(%ebp),%ebx  # %ebx=&x
movl 12(%ebp),%ecx # %ecx=&y
movl 16(%ebp),%edx # %edx=n
fldz               # push +0.0
xorl %eax,%eax    # i=0
cmpl %edx,%eax    # if i>=n done
jge .L3

.L5:
flds (%ebx,%eax,4) # push x[i]
fmuls (%ecx,%eax,4) # st(0)*=y[i]
faddp          # st(1)+=st(0)
incl %eax      # pop
cmpl %edx,%eax # i++
jl .L5

.L3:
movl -4(%ebp),%ebx # finish
movl %ebp, %esp
popl %ebp
ret                # st(0) = result
    
```

Floating Po

```

pushl %ebp          # setup
movl %esp,%ebp
pushl %ebx

movl 8(%ebp),%ebx  # %ebx=&x
movl 12(%ebp),%ecx # %ecx=&y
movl 16(%ebp),%edx # %edx=n
fldz               # push +0.0
xorl %eax,%eax    # i=0
cmpl %edx,%eax    # if i>=n done
jge .L3

.L5:
flds (%ebx,%eax,4) # push x[i]
fmuls (%ecx,%eax,4) # st(0)*=y[i]
faddp          # st(1)+=st(0)
incl %eax      # pop
cmpl %edx,%eax # i++
jl .L5

.L3:
movl -4(%ebp),%ebx # finish
movl %ebp, %esp
popl %ebp
ret                # st(0) = result
    
```

Inner Product Stack Trace

Initialization

1. fldz	0.0	%st(0)
---------	-----	--------

Iteration 0

2. flds (%ebx,%eax,4)	0.0	%st(1)
	x[0]	%st(0)
3. fmuls (%ecx,%eax,4)	0.0	%st(1)
	x[0]*y[0]	%st(0)
4. faddp	0.0+x[0]*y[0]	%st(0)

Iteration 1

5. flds (%ebx,%eax,4)	x[0]*y[0]	%st(1)
	x[1]	%st(0)
6. fmuls (%ecx,%eax,4)	x[0]*y[0]	%st(1)
	x[1]*y[1]	%st(0)
7. faddp	x[0]*y[0]+x[1]*y[1]	%st(0)

Final Observations

Memory Layout

- OS/machine dependent (including kernel version)
- Basic partitioning: stack/data/text/heap/DLL found in most machines

Type Declarations in C

- Notation obscure, but very systematic

Working with Strange Code

- Important to analyze nonstandard cases
 - E.g., what happens when stack corrupted due to buffer overflow
- Helps to step through with GDB

IA32 Floating Point

- Strange "shallow stack" architecture

- 37 -

15-213, S'04

Final Recommendation

Sign your mail.

Google "pgp"

- 38 -

15-213, S'04