

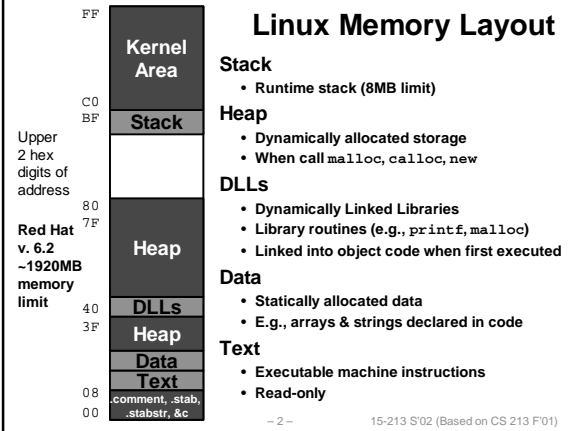
15-213 Machine Language V: Miscellaneous Topics February 12, 2002

Topics

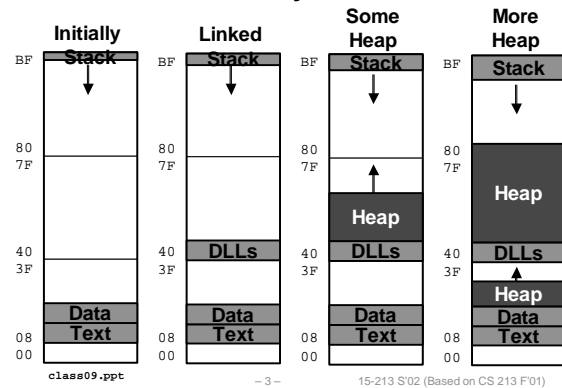
- Linux Memory Layout
- Understanding Pointers
- Buffer Overflow
- Floating Point Code
- Reading: 3.12 – 3.16
- Problems: 3.24 and 3.39

class09.ppt

Linux Memory Layout



Linux Memory Allocation



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 ... */
}
```

class09.ppt

– 4 –

15-213 S'02 (Based on CS 213 F'01)

Finding the Stack and Heap

```
(gdb) break main
(gdb) run
Breakpoint 1, 0x80483fc in main ()
(gdb) print $esp
$3 = (void *) 0xbffffb00
```

Main

- Address 0x80483fc should be read 0x80483fc

Stack

- Address 0xbffffb00

class09.ppt

- 5 -

15-213 S'02 (Based on CS 213 F'01)

Loading of A Dynamic Library

```
(gdb) disassemble 0x400734e0
No function contains specified address.
(gdb) print * (char *) 0x400734e0
Cannot access memory at address 0x400734e0
(gdb) break malloc
Breakpoint 1, main () at q.c:11
11      pl = malloc(1 <<28); /* 256 MB */
Dump of assembler code for function __libc_malloc:
0x400734e0 <__libc_malloc>:      push    %ebp
0x400734e1 <__libc_malloc+1>:      mov     %esp,%ebp
0x400734e3 <__libc_malloc+3>:      sub     $0x8,%esp
...
```

Main

- Address 0x80483fc should be read 0x80483fc

Stack

- Address 0xbffffb00

malloc()

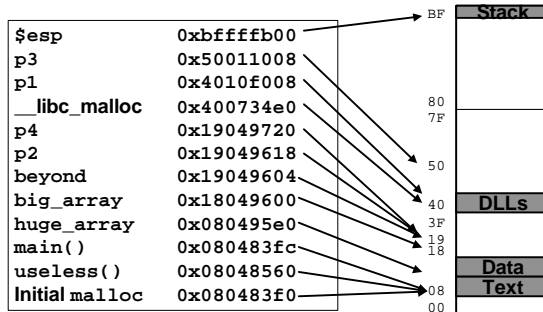
- Implemented by __libc_malloc
- Linked into memory at runtime - not present before program starts

class09.ppt

- 6 -

15-213 S'02 (Based on CS 213 F'01)

Example Addresses



class09.ppt

- 7 -

15-213 S'02 (Based on CS 213 F'01)

C operators

Operators

() [] -> .	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

class09.ppt

- 8 -

15-213 S'02 (Based on CS 213 F'01)

C pointer declarations

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

class09.ppt

- 9 -

15-213 S'02 (Based on CS 213 F'01)


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



class09.ppt

- 10 -

15-213 S'02 (Based on CS 213 F'01)

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, such as `gets()` and `strcpy()`, do not check argument sizes.
- allows target buffers to overflow.

class09.ppt

- 11 -

15-213 S'02 (Based on CS 213 F'01)

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;
}
  
```

class09.ppt

- 12 -

15-213 S'02 (Based on CS 213 F'01)

Buffer Overflow Executions

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

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

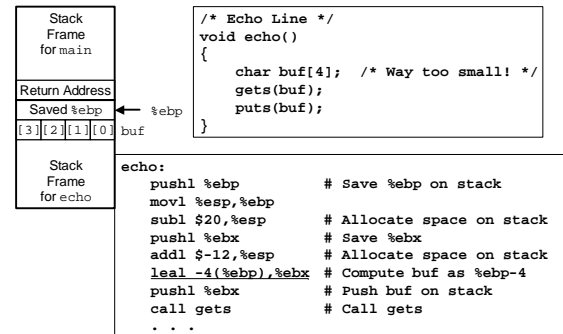
```
unix> ./bufdemo
Type a string:12345678
Segmentation Fault
```

class09.ppt

- 13 -

15-213 S'02 (Based on CS 213 F'01)

Buffer Overflow Stack




class09.ppt

- 14 -

15-213 S'02 (Based on CS 213 F'01)


Buffer Overflow Stack Example

Before Call to gets



```

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
        
```



```

8048648: call 804857c <echo>
804864d: mov 0xbffff8f8(%ebp),%ebx # Return Point
        
```


class09.ppt

- 15 -

15-213 S'02 (Based on CS 213 F'01)

Buffer Overflow Stack Example #1

Before Call to gets Input = "123"



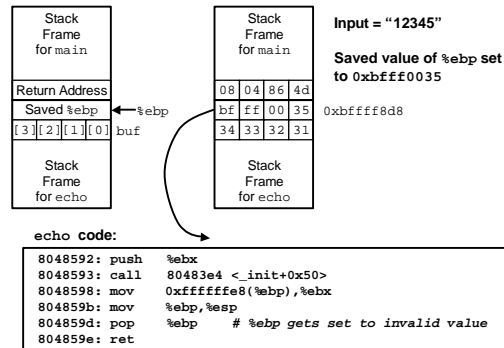
No Problem

class09.ppt

- 16 -

15-213 S'02 (Based on CS 213 F'01)

Buffer Overflow Stack Example #2

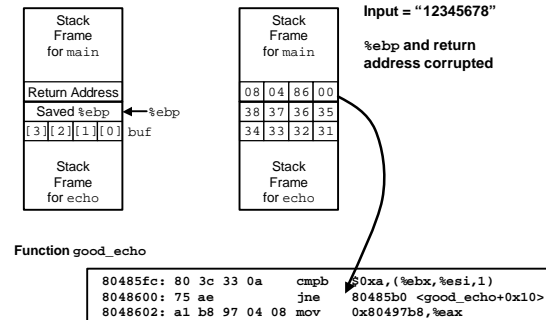


class09.ppt

- 17 -

15-213 S'02 (Based on CS 213 F'01)

Buffer Overflow Stack Example

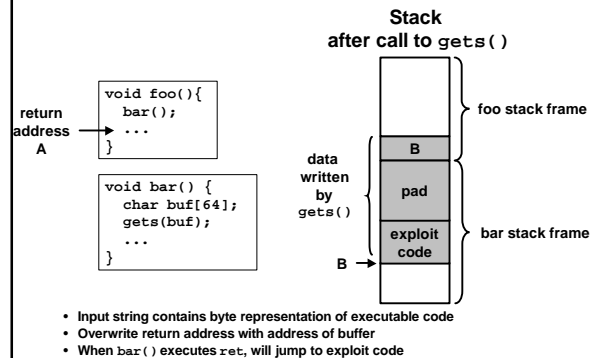


class09.ppt

- 18 -

15-213 S'02 (Based on CS 213 F'01)

Malicious Use of Buffer Overflow



class09.ppt

- 19 -

15-213 S'02 (Based on CS 213 F'01)

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.

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.

class09.ppt

- 20 -

15-213 S'02 (Based on CS 213 F'01)

Date: Wed, 11 Aug 1999 11:30:57 -0700 (PDT)
 From: Phil Bucking <philbucking@yahoo.com>
 Subject: AOL exploiting buffer overrun bug in their own software
 To: rms@pharlap.com

Mr. Smith,

I am writing you because I have discovered something that I think you might find interesting because you are an Internet security expert with experience in this area. I have also tried to contact AOL but received no response.

I am a developer who has been working on a revolutionary new instant messaging client that should be released later this year.

...
 It appears that the AIM client has a buffer overrun bug. By itself this might not be the end of the world, as MS surely has had its share. But AOL is now 'exploiting their own buffer overrun bug' to help in its efforts to block MS Instant Messenger.

...
 Since you have significant credibility with the press I hope that you can use this information to help inform people that behind AOL's friendly exterior they are nefariously compromising peoples' security.

Sincerely,
 Phil Bucking
 Founder, Bucking Consulting
 philbucking@yahoo.com

It was later determined that this email originated from within Microsoft!

class09.ppt

- 21 -

15-213 S'02 (Based on CS 213 F'01)

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




class09.ppt

- 22 -

15-213 S'02 (Based on CS 213 F'01)


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



class09.ppt

- 23 -

15-213 S'02 (Based on CS 213 F'01)

FPU instructions

Large number of floating point instructions and formats

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

Sample instructions:

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

class09.ppt

- 24 -

15-213 S'02 (Based on CS 213 F'01)

Floating Point Code Example

Compute Inner Product of Two Vectors

- Single precision arithmetic
- Scientific computing and signal processing workhorse

```
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      # 1=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
jle .L5

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

class09.ppt

- 25 -

15-213 S'02 (Based on CS 213 F'01)

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
   %st(0)
```

x[0]*y[0]+x[1]*y[1]

class09.ppt

- 26 -

15-213 S'02 (Based on CS 213 F'01)

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

class09.ppt

- 27 -

15-213 S'02 (Based on CS 213 F'01)