

Princeton University

Computer Science 217: Introduction to Programming Systems



Machine Language

1

Machine Language



This lecture (Monday) is about

- Machine language (in general)
- A motivating example from computer security
(And, therefore, Assignment 5: Buffer Overrun)
- AARCH64 machine language (in particular)

These slides also contain

- The assembly and linking processes (Wednesday)

2

Instruction Set Architecture (ISA)



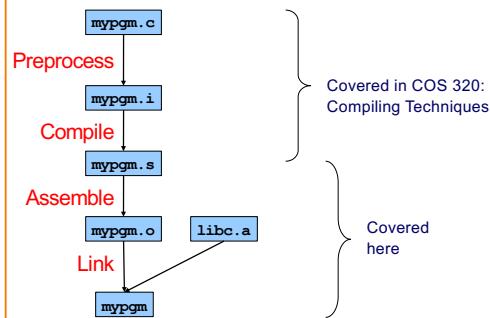
There are many kinds of computer chips out there:

ARM
Intel x86 series
IBM PowerPC
RISC-V
MIPS
(and, in the old days, dozens more)

Each of these different
“machine architectures”
understands a different
machine language

3

The Build Process



Covered here

Covered in COS 320:
Compiling Techniques

4

A Program



```

#include <stdio.h>
int main(void)
{
    char name[12], c;
    int i = 0, magic = 42;
    printf("What is your name?\n");
    while ((c = getchar()) != '\n')
        name[i++] = c;
    name[i] = '\0';
    printf("Thank you, %s.\n", name);
    printf("The answer to life, the universe, "
          "and everything is %d\n", magic);
    return 0;
}

$ ./a.out
What is your name?
John Smith
Thank you, John Smith.
The answer to life, the universe, and everything is 42
  
```

5

Why People With Long Names Have Problems



```

#include <stdio.h>
int main(void)
{
    char name[12], c;
    int i = 0, magic = 42;
    printf("What is your name?\n");
    while ((c = getchar()) != '\n')
        name[i++] = c;
    name[i] = '\0';
    printf("Thank you, %s.\n", name);
    printf("The answer to life, the universe, "
          "and everything is %d\n", magic);
    return 0;
}

$ ./a.out
What is your name?
Christopher Moretti
Thank you, Christopher Mor ← ?
tti.
The answer to life, the universe, and everything is 6911092
  
```

??? (!)
(Depending on the records, this might be an interesting phone number, but probably not one you should call for the answer to life, the universe, and everything.)

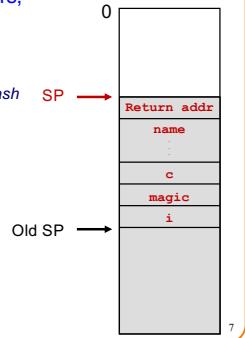
6

Explanation: Stack Frame Layout

When there are too many characters, program carelessly writes beyond space "belonging" to name.

- Overwrites other variables
- This is a *buffer overrun*, or *stack smash*
- The program has a security bug!

```
#include <stdio.h>
int main(void)
{
    char name[12], c;
    int i = 0, magic = 42;
    printf("What is your name?\n");
    while ((c = getchar()) != '\n')
        name[i++] = c;
    name[i] = '\0';
    printf("Thank you, %s.\n", name);
    printf("The answer to life, the universe,
        "and everything is %d\n", magic);
    return 0;
}
```

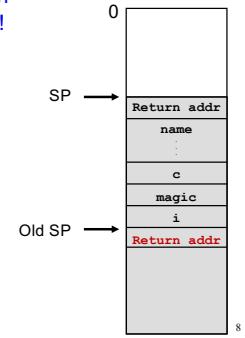


7

It Gets Worse...

Buffer overrun can overwrite return address of a previous stack frame!

```
#include <stdio.h>
int main(void)
{
    char name[12], c;
    int i = 0, magic = 42;
    printf("What is your name?\n");
    while ((c = getchar()) != '\n')
        name[i++] = c;
    name[i] = '\0';
    printf("Thank you, %s.\n", name);
    printf("The answer to life, the universe,
        "and everything is %d\n", magic);
    return 0;
}
```



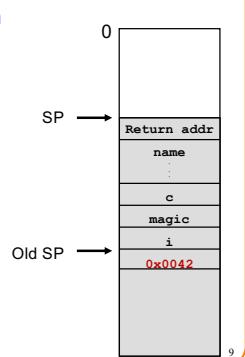
8

It Gets Worse...

Buffer overrun can overwrite return address of a previous stack frame!

- Value can be an invalid address, leading to a segfault,...

```
#include <stdio.h>
int main(void)
{
    char name[12], c;
    int i = 0, magic = 42;
    printf("What is your name?\n");
    while ((c = getchar()) != '\n')
        name[i++] = c;
    name[i] = '\0';
    printf("Thank you, %s.\n", name);
    printf("The answer to life, the universe,
        "and everything is %d\n", magic);
    return 0;
}
```



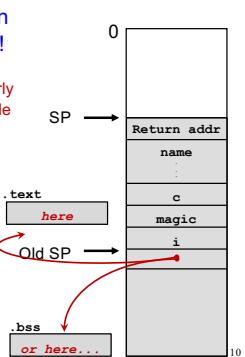
9

It Gets Much, Much Worse...

Buffer overrun can overwrite return address of a previous stack frame!

- Value can be an invalid address, leading to a segfault, or it can cleverly point to unintended or malicious code

```
#include <stdio.h>
int main(void)
{
    char name[12], c;
    int i = 0, magic = 42;
    printf("What is your name?\n");
    while ((c = getchar()) != '\n')
        name[i++] = c;
    name[i] = '\0';
    printf("Thank you, %s.\n", name);
    printf("The answer to life, the universe,
        "and everything is %d\n", magic);
    return 0;
}
```



10

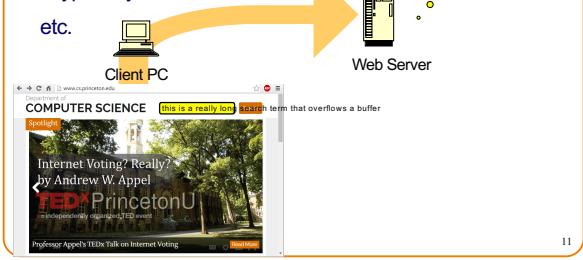
Attacking a Web Server

URLs

Input in web forms

Crypto keys for SSL

etc.



11

11

Attacking a Web Browser

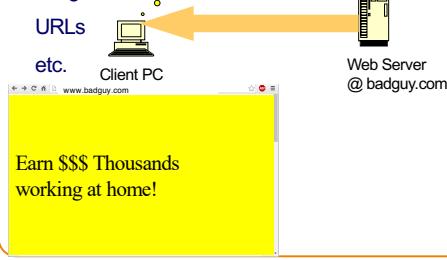
HTML keywords

Images

Image names

URLs

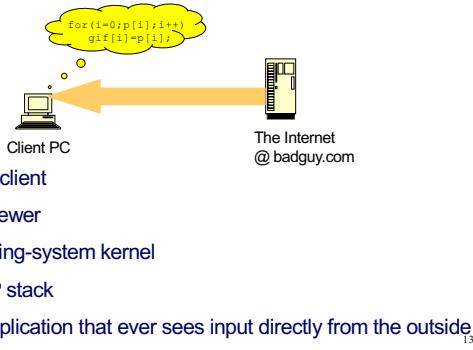
etc.



12

2

Attacking Everything in Sight



13

Defenses Against This Attack

Best: program in languages that make array-out-of-bounds impossible (Java, C#, ML, python,)

None of these would have prevented the “Heartbleed” attack

If you must program in C: use discipline
and software analysis tools to check
bounds of array subscripts

Otherwise, stopgap security patches:

- Operating system randomizes initial stack pointer
 - “No-execute” memory permission
 - “Canaries” at end of stack frames

14

14

Asgt. 5: Attack the “Grader” Program

```
$ ./grader
What is your name?
Bob
D is your grade.
Thank you, Bob.

$ ./grader
What is your name?
Andrew Appel
B is your grade.
Thank you, Andrew A
```

15

15

Asgt. 5: Attack the “Grader” Program

```

/* Read a string into name */
void readstring() {
    char buf[BUFSIZE];
    int i = 0; int c;

    /* Read string into buf[] */
    for (; ;) {
        c = fgetchar(stdin);
        if (c == EOF || c == '\n')
            break;
        buf[i] = c;
        i++;
    }
    buf[i] = '\0';

    /* Copy buf[] to name[] */
    for (i = 0; i < BUFSIZE; i++)
        name[i] = buf[i];
}

```

```
/* Prompt for name and read it */
void getName() {
    printf("What is your name?\n");
    readString();
}
```

Unchecked
write to
buffer!

16

16

Asgt. 5: Attack the “Grader” Program

```
int main(void) {
    getname();
    if (strcmp(name, "Bob") == 0)
        grade = "B";
    printf("%c is your grade\n", grade);
    printf("Thank you.\n");
    return 0;
}
```

17

17

Asgt. 5: Attack the “Grader” Program

```
int main(void) {
    getname();
    if (strcmp(name, "Andrew Appel") == 0)
        grade = 'B';
    printf("%c is your grade.\n", grade);
    printf("Thank you, %s.\n", name);
    return 0;
}
```

```
$ ./grader  
What is your name?  
Susan@0?!*??????!!*!%?!?(!*%(*^^?  
A is your grade.  
Thank you, Susan.
```

18

18

Agenda

Buffer overrun vulnerabilities

AARCH64 Machine Language

AARCH64 Machine Language after Assembly

AARCH64 Machine Language after Linking



38

An Example Program

A simple (nonsensical) program, in C and assembly:

```
#include <stdio.h>
int main(void)
{ printf("Type a char: ");
  if (getchar() == 'A')
    printf("Hi\n");
  return 0;
}
```

```
.section .rodata
msg1: .string "Type a char: "
msg2: .string "Hi\n"
.section .text
.global main
main:
    sub    sp, sp, 16
    str   x30, [sp]
    adr   x0, msg1
    bl    printf
    bl    getchar
    cmp   w0, 'A'
    bne  skip
    adr   x0, msg2
    bl    printf
skip:
    mov   w0, 0
    ldr   x30, [sp]
    add   sp, sp, 16
    ret
```



Let's consider the machine language equivalent...

39

Examining Machine Lang: RODATA

Assemble program; run objdump

```
$ gcc217 -c detecta.s
$ objdump --full-contents --section .rodata detecta.o

detecta.o: file format elf64-littlesearch64

Contents of section .rodata:
0000 54797065 20612063 6861723a 20004869 Type a char: .Hi
0010 0000
```

Offsets

- Assembler does not know addresses
- Assembler knows only offsets
- "Type a char: " starts at offset 0x0
- "Hi\n" starts at offset 0xe

40



40

Examining Machine Lang: TEXT

Run objdump to see instructions

```
$ objdump --disassemble --reloc detecta.o
detecta.o: file format elf64-littlesearch64

Disassembly of section .text:
0000000000000000 <main>:
  0: d10043fe  sub    sp, sp, #0x10
  4: f90003fe  str   x30, [sp]
  8: 10000000  adr   x0, 0 <main>
  c: R_AARCH64_ADR_P2B_L021    rodata
  10: 94000000  bl    0 <printf>
    R_AARCH64_CALL26
  14: 7101041f  cmp   w0, #0x1
  18: 54000061  b.ne  24 <skip>
  1c: 10000000  adr   x0, 0 <main>
  20: 94000000  bl    0 <printf>
    R_AARCH64_CALL26
  24: f28000000 <skip>:
  28: f94003fe  ldr   x30, [sp]
  2c: 910043ff  add   sp, sp, #0x10
  30: d65f03c0  ret
```



Assembly language

41

41

Examining Machine Lang: TEXT

Run objdump to see instructions

```
$ objdump --disassemble --reloc detecta.o
detecta.o: file format elf64-littlesearch64

Disassembly of section .text:
0000000000000000 <main>:
  0: d10043fe  sub    sp, sp, #0x10
  4: f90003fe  str   x30, [sp]
  8: 10000000  adr   x0, 0 <main>
  c: R_AARCH64_ADR_P2B_L021    rodata
  10: 94000000  bl    0 <printf>
    R_AARCH64_CALL26
  14: 7101041f  cmp   w0, #0x1
  18: 54000061  b.ne  24 <skip>
  1c: 10000000  adr   x0, 0 <main>
  20: 94000000  bl    0 <printf>
    R_AARCH64_CALL26
  24: f28000000 <skip>:
  28: f94003fe  ldr   x30, [sp]
  2c: 910043ff  add   sp, sp, #0x10
  30: d65f03c0  ret
```

42



42

Examining Machine Lang: TEXT

Run objdump to see instructions

```
$ objdump --disassemble --reloc detecta.o
detecta.o: file format elf64-littlesearch64

Disassembly of section .text:
0000000000000000 <main>:
  0: d10043fe  sub    sp, sp, #0x10
  4: f90003fe  str   x30, [sp]
  8: 10000000  adr   x0, 0 <main>
  c: R_AARCH64_ADR_P2B_L021    rodata
  10: 94000000  bl    0 <printf>
    R_AARCH64_CALL26
  14: 7101041f  cmp   w0, #0x1
  18: 54000061  b.ne  24 <skip>
  1c: 10000000  adr   x0, 0 <main>
  20: 94000000  bl    0 <printf>
    R_AARCH64_CALL26
  24: f28000000 <skip>:
  28: f94003fe  ldr   x30, [sp]
  2c: 910043ff  add   sp, sp, #0x10
  30: d65f03c0  ret
```



Let's examine one line at a time...

43

43

sub sp, sp, #0x10

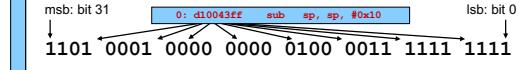
```
$ objdump --disassemble --reloc detecta.o
```

detecta.o: file format elf64-littlearched64

Disassembly of section .text:

0000000000000000 <main>:

0: d10043ff sub sp, sp, #0x10



```
1c: 10000000 ldr x0, 0 <main>
20: 94000000 bl 0 <printf>
```

```
0000000000000024 <skip>:
24: f5280000 mov w0, #0x0
28: f94003fe ldr x30, [sp]
2c: 910403ff add sp, sp, #0x10
30: d65f03c0 ret
```

44

sub sp, sp, #0x10

msb: bit 31 0: d10043ff sub sp, sp, #0x10 lsb: bit 0
↓ ↓ ↓

1101 0001 0000 0000 0100 0011 1111 1111

- opcode: subtract immediate
- Instruction width in bit 31: 1 = 64-bit
- Whether to set condition flags in bit 29: no
- Immediate value in bits 10-21: $10000_2 = 0x10 = 16$
- First source register in bits 5-9: 31 = sp
- Destination register in bits 0-4: 31 = sp
- Additional information about instruction: none

45

44

45

str x30, [sp]

```
$ objdump --disassemble --reloc detecta.o
```

detecta.o: file format elf64-littlearched64

Disassembly of section .text:

0000000000000000 <main>:

0: d10043ff sub sp, sp, #0x10

```
4: f90003fe str x30, [sp]
8: 10000000 adr x0, 0 <main>
b: R_AARCH64_ADR_PREL_LO21 .rodata
c: 94000000 bl 0 <printf>
f: R_AARCH64_CALL26 printf
10: 94000000 bl 0 <getchar>
13: R_AARCH64_ADR_PREL_LO21 .rodata+0xe
14: 7101041f cmp w0, #0x41
18: 54000061 b.ne 24 <skip>
1c: 10000000 adr x0, 0 <main>
1f: R_AARCH64_ADR_PREL_LO21 .rodata+0xe
20: 94000000 bl 0 <printf>
23: R_AARCH64_CALL26 printf
```

```
0000000000000024 <skip>:
24: f5280000 mov w0, #0x0
28: f94003fe ldr x30, [sp]
2c: 910403ff add sp, sp, #0x10
30: d65f03c0 ret
```

46

str x30, [sp]

msb: bit 31 4: f90003fe str x30, [sp] lsb: bit 0
↓ ↓ ↓

1111 1001 0000 0000 0000 0011 1111 1110

- opcode: store, register + offset
- Instruction width in bits 30-31: 11 = 64-bit
- Offset value in bits 12-20: 0
- “Source” (really destination) register in bits 5-9: 31 = sp
- “Destination” (really source) register in bits 0-4: 30
- Additional information about instruction: none

47

46

47

adr x0, 0 <main>

```
$ objdump --disassemble --reloc detecta.o
```

detecta.o: file format elf64-littlearched64

Disassembly of section .text:

0000000000000000 <main>:

0: d10043ff sub sp, sp, #0x10

```
4: f90003fe str x30, [sp]
8: 10000000 adr x0, 0 <main>
b: R_AARCH64_ADR_PREL_LO21 .rodata
c: 94000000 bl 0 <printf>
f: R_AARCH64_CALL26 printf
10: 94000000 bl 0 <getchar>
13: R_AARCH64_ADR_PREL_LO21 .rodata+0xe
14: 7101041f cmp w0, #0x41
18: 54000061 b.ne 24 <skip>
1c: 10000000 adr x0, 0 <main>
1f: R_AARCH64_ADR_PREL_LO21 .rodata+0xe
20: 94000000 bl 0 <printf>
23: R_AARCH64_CALL26 printf
```

```
0000000000000024 <skip>:
24: f5280000 mov w0, #0x0
28: f94003fe ldr x30, [sp]
2c: 910403ff add sp, sp, #0x10
30: d65f03c0 ret
```

48

adr x0, 0 <main>

msb: bit 31 8: 10000000 adr x0, 0 <main> lsb: bit 0
↓ ↓ ↓

0001 0000 0000 0000 0000 0000 0000 0000

- opcode: generate address
- 19 High-order bits of relative address in bits 5-23: 0
- 2 Low-order bits of relative address in bits 29-30: 0
- *Relative data location is 0 bytes after this instruction*
- Destination register in bits 0-4: 0
- Huh? That's not where **msg1** lives!
 - Assembler knew that **msg1** is a label within the RODATA section
 - But assembler didn't know address of RODATA section!
 - So, assembler couldn't generate this instruction completely, left a placeholder, and will request help from the linker

49

48

49

Examining Machine Lang: TEXT

```
$ objdump --disassemble --reloc detecta.o
detecta.o:    file format elf64-littlearm64
Disassembly of section .text:
0000000000000000 <main>:
 0: d10043ff  sub    sp, sp, #0x10
 4: f90003fe  str    x30, [sp]
 8: 10000000  adr    x0, 0 <main>
c: 94000000  bl    R_AARCH64_ADR_PREL_LO21 .rodata
 0: 0 <printf>
c: R_AARCH64_CALL26  printf
10: 94000000  bl    0 <getchar>
 0: R_AARCH64_CALL26  getchar
14: 7101041f  cmp    w0, #0x1
18: 54000061  b.ne   24 <skip>
1c: 10000000  adr    x0, 0 <main>
 0: R_AARCH64_ADR_PREL_LO21 .rodata+0xe
20: 94000000  bl    0 <printf>
 0: R_AARCH64_CALL26  printf
0000000000000024 <skip>:
24: 52800000  mov    w0, #0x0
28: f94003fe  ldr    x30, [sp]
2c: 910043ff  add    sp, sp, #0x10
30: d65f03c0  ret

Run objdump to see instructions
```

Relocation records

50

R_AARCH64_ADR_PREL_LO21 .rodata

```
$ objdump --disassemble --reloc detecta.o
detecta.o:    file format elf64-littlearm64
Disassembly of section .text:
0000000000000000 <main>:
 0: d10043ff  sub    sp, sp, #0x10
 4: f90003fe  str    x30, [sp]
 8: 10000000  adr    x0, 0 <main>
c: 94000000  bl    R_AARCH64_ADR_PREL_LO21 .rodata
 0: 0 <printf>
c: R_AARCH64_CALL26  printf
10: 94000000  bl    0 <getchar>
 0: R_AARCH64_CALL26  getchar
14: 7101041f  cmp    w0, #0x1
18: 54000061  b.ne   24 <skip>
1c: 10000000  adr    x0, 0 <main>
 0: R_AARCH64_ADR_PREL_LO21 .rodata+0xe
20: 94000000  bl    0 <printf>
 0: R_AARCH64_CALL26  printf
0000000000000024 <skip>:
24: 52800000  mov    w0, #0x0
28: f94003fe  ldr    x30, [sp]
2c: 910043ff  add    sp, sp, #0x10
30: d65f03c0  ret
```

51

Relocation Record 1

```
8: R_AARCH64_ADR_PREL_LO21 .rodata
This part is always the same, it's the name of the machine architecture!
```

Dear Linker,

Please patch the TEXT section at offset **0x8**. Patch in a **21-bit** signed offset of an **address**, relative to the **PC**, as appropriate for the instruction format. When you determine the address of **rodata**, use that to compute the offset you need to do the patch.

Sincerely,
Assembler

19 High-order bits of relative address in bits 5-23: 0
2 Low-order bits of relative address in bits 29-30: 0

52

bl 0 <printf>

```
$ objdump --disassemble --reloc detecta.o
detecta.o:    file format elf64-littlearm64
Disassembly of section .text:
0000000000000000 <main>:
 0: d10043ff  sub    sp, sp, #0x10
 4: f90003fe  str    x30, [sp]
 8: 10000000  adr    x0, 0 <main>
c: 94000000  bl    R_AARCH64_ADR_PREL_LO21 .rodata
 0: 0 <printf>
c: R_AARCH64_CALL26  printf
10: 94000000  bl    0 <getchar>
 0: R_AARCH64_CALL26  getchar
14: 7101041f  cmp    w0, #0x1
18: 54000061  b.ne   24 <skip>
1c: 10000000  adr    x0, 0 <main>
 0: R_AARCH64_ADR_PREL_LO21 .rodata+0xe
20: 94000000  bl    0 <printf>
 0: R_AARCH64_CALL26  printf
0000000000000024 <skip>:
24: 52800000  mov    w0, #0x0
28: f94003fe  ldr    x30, [sp]
2c: 910043ff  add    sp, sp, #0x10
30: d65f03c0  ret
```

53

bl 0 <printf>

msb: bit 31 c: 94000000 bl 0 <printf> lsb: bit 0

↓ ↓

1001 0100 0000 0000 0000 0000 0000 0000

- opcode: branch and link
- Relative address in bits 0-25: 0

Huh? That's not where **printf** lives!

- Assembler had to calculate [addr of **printf**] - [addr of this instr]
- But assembler didn't know address of **printf** - it's off in some library (**libc.a**) and isn't present (yet)!
- So, assembler couldn't generate this instruction completely, left a placeholder, and will request help from the linker

54

R_AARCH64_CALL26 printf

```
$ objdump --disassemble --reloc detecta.o
detecta.o:    file format elf64-littlearm64
Disassembly of section .text:
0000000000000000 <main>:
 0: d10043ff  sub    sp, sp, #0x10
 4: f90003fe  str    x30, [sp]
 8: 10000000  adr    x0, 0 <main>
c: 94000000  bl    R_AARCH64_ADR_PREL_LO21 .rodata
 0: 0 <printf>
c: R_AARCH64_CALL26  printf
10: 94000000  bl    0 <getchar>
 0: R_AARCH64_CALL26  getchar
14: 7101041f  cmp    w0, #0x1
18: 54000061  b.ne   24 <skip>
1c: 10000000  adr    x0, 0 <main>
 0: R_AARCH64_ADR_PREL_LO21 .rodata+0xe
20: 94000000  bl    0 <printf>
 0: R_AARCH64_CALL26  printf
0000000000000024 <skip>:
24: 52800000  mov    w0, #0x0
28: f94003fe  ldr    x30, [sp]
2c: 910043ff  add    sp, sp, #0x10
30: d65f03c0  ret
```

55

54

9

Relocation Record 2

c: R_AARCH64_CALL26 printf

Dear Linker,

Please patch the TEXT section at offset 0xc.
Patch in a 26-bit signed offset relative to the PC,
appropriate for the function `call(bl)` instruction
format. When you determine the address of
`printf`, use that to compute the offset you need
to do the patch.

Sincerely,
Assembler

56

b1 0 <getchar>

```
$ objdump --disassemble --reloc detecta.o
detecta.o:      file format elf64-littlesearch64
Disassembly of section .text:
0000000000000000 <main>
 0: d10043ff  sub  sp, sp, #0x10
 4: f90003fe  str   x30, [sp]
 8: 10000000  adr   x0, <main>
 c: 94000000  bl    0 <printf>
10: 94000000  bl    0 <getchar>
14: 7101041f  cmp   w0, #0x41
18: 54000061  bne   x0, 0 <main>
1c: 10000000  adr   x0, 0 <main>
1d: R_AARCH64_ADR_PREL_LO21 .rodata+0xe
20: 94000000  bl    0 <printf>
24: f94003fe  ldr   x30, [sp]
28: 910043ff  add   sp, sp, #0x10
30: d65f03c0  ret

0000000000000024 <skip>
24: 52800000  mov   w0, #0x0
28: f94003fe  ldr   x30, [sp]
2c: 910043ff  add   sp, sp, #0x10
30: d65f03c0  ret
```



57

b1 0 <getchar>

msb: bit 31 10: 94000000 bl 0 <getchar> lsb: bit 0
↓ ↓ ↓
1001 0100 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000

- opcode: branch and link
- Relative address in bits 0-25: 0
- Same situation as before – relocation record coming up!

58

Relocation Record 3

10: R_AARCH64_CALL26 getchar

Dear Linker,

Please patch the TEXT section at offset 0x10.
Patch in a 26-bit signed offset relative to the PC,
appropriate for the function `call(bl)` instruction
format. When you determine the address of
`getchar`, use that to compute the offset you
need to do the patch.

Sincerely,
Assembler



59

cmp w0, #0x41

```
$ objdump --disassemble --reloc detecta.o
detecta.o:      file format elf64-littlesearch64
Disassembly of section .text:
0000000000000000 <main>
 0: d10043ff  sub  sp, sp, #0x10
 4: f90003fe  str   x30, [sp]
 8: 10000000  adr   x0, <main>
 c: 94000000  bl    0 <printf>
10: 94000000  bl    0 <getchar>
14: 7101041f  cmp   w0, #0x41
18: 54000061  bne   x0, 0 <main>
1c: 10000000  adr   x0, 0 <main>
1d: R_AARCH64_ADR_PREL_LO21 .rodata+0xe
20: 94000000  bl    0 <printf>

0000000000000024 <skip>
24: 52800000  mov   w0, #0x0
28: f94003fe  ldr   x30, [sp]
2c: 910043ff  add   sp, sp, #0x10
30: d65f03c0  ret
```

60

cmp w0, #0x41

msb: bit 31 14: 7101041f cmp w0, #0x41 lsb: bit 0
↓ ↓ ↓
0111 0001 0000 0001 0000 0100 0001 1111

- Recall that `cmp` is really an assembler alias:
this is the same instruction as `subs wzr, w0, 0x41`
- opcode: subtract immediate
- Instruction width in bit 31: 0 = 32-bit
- Whether to set condition flags in bit 29: yes
- Immediate value in bits 10-21: $1000001_0 = 0x41 = 'A'$
- First source register in bits 5-9: 0
- Destination register in bits 0-4: 31 = wzr
- Note that register #31 (11111_0) is used to mean either sp or xzr/wzr,
depending on the instruction



61

60

b.ne 24 <skip>

```
$ objdump --disassemble --reloc detecta.o

detecta.o:      file format elf64-littlesearch64

Disassembly of section .text:
0000000000000000 <main>:
 0: d10043ff  sub    sp, sp, #0x10
 4: f90003fe  str    x30, [sp]
 8: 10000000  adr    x0, 0 <main>
 8: R_AARCH64_ADR_PREL_LO21   .rodata
 c: 94000000  bl    R_AARCH64_CALL26   printf
10: 94000000  bl    R_AARCH64_CALL26   printf
10: R_AARCH64_CALL26   printf
14: 7101041f  cmp    w0, #0x1
18: 54000061  b.ne   24 <skip>
1c: 10000000  adr    x0, 0 <main>
1c: R_AARCH64_ADR_PREL_LO21   .rodata+0xe
20: 94000000  bl    R_AARCH64_CALL26   printf
20: R_AARCH64_CALL26   printf

0000000000000024 <skip>:
24: 52800000  mov    w0, #0x0
28: f94003fe  ldr    x30, [sp]
2c: 910043ff  add    sp, sp, #0x10
30: d65f03c0  ret
```



62

b.ne 24 <skip>

msb: bit 31 lsb: bit 0
 ↓ ↓
0101 0100 0000 0000 0000 0000 0110 0001

- This instruction is at offset 0x18, and **skip** is at offset 0x24, which is $0x24 - 0x18 = 0xc = 12$ bytes later
- opcode: conditional branch
- Relative address in bits 5-23: 11_b. Shift left by 2: 1100_b = 12
- Conditional branch type in bits 0-4: NE
- No need for relocation record!
 - Assembler had to calculate [addr of **skip**] – [addr of this instr]
 - Assembler did know offsets of **skip** and this instruction
 - So, assembler could generate this instruction completely, and does not need to request help from the linker

63

R_AARCH64_ADR_PREL_LO21 .rodata+0xe

```
$ objdump --disassemble --reloc detecta.o

detecta.o:      file format elf64-littlesearch64

Disassembly of section .text:
0000000000000000 <main>:
 0: d10043ff  sub    sp, sp, #0x10
 4: f90003fe  str    x30, [sp]
 8: 10000000  adr    x0, 0 <main>
 8: R_AARCH64_ADR_PREL_LO21   .rodata
 c: 94000000  bl    R_AARCH64_CALL26   printf
10: 94000000  bl    R_AARCH64_CALL26   printf
10: R_AARCH64_CALL26   printf
14: 7101041f  cmp    w0, #0x1
18: 54000061  b.ne   24 <skip>
1c: 10000000  adr    x0, 0 <main>
1c: R_AARCH64_ADR_PREL_LO21   .rodata+0xe
20: 94000000  bl    R_AARCH64_CALL26   printf
20: R_AARCH64_CALL26   printf

0000000000000024 <skip>:
24: 52800000  mov    w0, #0x0
28: f94003fe  ldr    x30, [sp]
2c: 910043ff  add    sp, sp, #0x10
30: d65f03c0  ret
```



64

Relocation Record 4

1c: R_AARCH64_ADR_PREL_LO21 .rodata+0xe

Dear Linker,

Please patch the TEXT section at offset 0x1c.
 Patch in a 21-bit signed offset of an address,
 relative to the PC, as appropriate for the
 instruction format. When you determine the
 address of .rodata, add 0xe and use that to
 compute the offset you need to do the patch.

Sincerely,
 Assembler

65

Another printf, with relocation record...

```
$ objdump --disassemble --reloc detecta.o

detecta.o:      file format elf64-littlesearch64

Disassembly of section .text:
0000000000000000 <main>:
 0: d10043ff  sub    sp, sp, #0x10
 4: f90003fe  str    x30, [sp]
 8: 10000000  adr    x0, 0 <main>
 8: R_AARCH64_ADR_PREL_LO21   .rodata
 c: 94000000  bl    R_AARCH64_CALL26   printf
10: 94000000  bl    R_AARCH64_CALL26   printf
10: R_AARCH64_CALL26   printf
14: 7101041f  cmp    w0, #0x1
18: 54000061  b.ne   24 <skip>
1c: 10000000  adr    x0, 0 <main>
1c: R_AARCH64_ADR_PREL_LO21   .rodata+0xe
20: 94000000  bl    R_AARCH64_CALL26   printf
20: R_AARCH64_CALL26   printf

0000000000000024 <skip>:
24: 52800000  mov    w0, #0x0
28: f94003fe  ldr    x30, [sp]
2c: 910043ff  add    sp, sp, #0x10
30: d65f03c0  ret
```



66

Everything Else is Similar...

```
$ objdump --disassemble --reloc detecta.o

detecta.o:      file format elf64-littlesearch64

Disassembly of section .text:
0000000000000000 <main>:
 0: d10043ff  sub    sp, sp, #0x10
 4: f90003fe  str    x30, [sp]
 8: 10000000  adr    x0, 0 <main>
 8: R_AARCH64_ADR_PREL_LO21   .rodata
 c: 94000000  bl    R_AARCH64_CALL26   printf
10: 94000000  bl    R_AARCH64_CALL26   printf
10: R_AARCH64_CALL26   printf
14: 7101041f  cmp    w0, #0x1
18: 54000061  b.ne   24 <skip>
1c: 10000000  adr    x0, 0 <main>
1c: R_AARCH64_ADR_PREL_LO21   .rodata+0xe
20: 94000000  bl    R_AARCH64_CALL26   printf
20: R_AARCH64_CALL26   printf

0000000000000024 <skip>:
24: 52800000  mov    w0, #0x0
28: f94003fe  ldr    x30, [sp]
2c: 910043ff  add    sp, sp, #0x10
30: d65f03c0  ret
```



Exercise for you:
 using information
 from these slides,
 create a bitwise
 breakdown of
 these instructions,
 and convince yourself
 that the hex values
 are correct!

67

66

11

Agenda

- Buffer overrun vulnerabilities
- AARCH64 Machine Language
- AARCH64 Machine Language after Assembly
- AARCH64 Machine Language after Linking**



68

68

From Assembler to Linker



Assembler writes its data structures to .o file

Linker:

- Reads .o file
- Writes executable binary file
- Works in two phases: **resolution** and **relocation**

69

69

Linker Resolution



Resolution

- Linker resolves references
- For this program, linker:
 - Notes that labels `getchar` and `printf` are unresolved
 - Fetches machine language code defining `getchar` and `printf` from `libc.a`
 - Adds that code to TEXT section
 - Adds more code (e.g. definition of `_start`) to TEXT section too
 - Adds code to other sections too

70

70

Linker Relocation



Relocation

- Linker patches (“relocates”) code
- Linker traverses relocation records, patching code as specified

71

71

Examining Machine Lang: RODATA



Link program; run objdump

```
$ gcc217 detecta.o -o detecta
$ objdump --full-contents --section .rodata detecta

detecta:   file format elf64-littleaarch64

Contents of section .rodata:
400710 01000200 00000000 00000000 ..... .
400720 54797065 20612063 6861723a 20004869  Type a char: .Hi
400730 0400

RODATA is at 0x400710
Starts with some header info
Real start of RODATA is at 0x400720
"Type a char: " starts at 0x400720
"Hi\n" starts at 0x40072e
```

Addresses,
not offsets

72

72

Examining Machine Lang: TEXT



```
$ objdump --disassemble --reloc detecta
Run objdump to see instructions

detecta:   file format elf64-littleaarch64
...
0000000004006400 <main>:
400650: d10043ff  sub  sp, sp, #0x10
400654: f90003fe  str  x30, [sp]
400658: 10000640  adr  x0, 400720 <msg1>
40065c: 97fffffa1  bl  4004e0 <printf@plt>
400660: 97ffff9c  bl  4004d0 <getchar@plt>
400664: 7101041f  cmp  w0, #0x41
400668: 54000061  bne  w0, 400674 <skip>
40066c: 50000600  adr  x0, 40072e <msg2>
400670: 97ffff9c  bl  4004e0 <printf@plt>

000000000400674 <skip>:
400674: 52800000  mov  w0, #0x0
400678: f94003fe  ldr  x30, [sp]
40067c: 910043ff  add  sp, sp, #0x10
400680: d65f03c0  ret

Addresses,
not offsets
```

73

12

Examining Machine Lang: TEXT

```
$ objdump --disassemble --reloc detecta
detecta: file format elf64-littleaarch64

...
0000000000400650 <main>:
400650: d10043ff sub sp, sp, #0x10
400654: f90003fe str x30, [sp]
400658: 10000640 adr x0, 400720 <msg1>
40065c: 97fffffa1 bl 4004e0 <printf@plt>
400660: 97ffff9c ldr w0, [x30, #0x10]
400664: 7101041f cmp w0, #0x41
400668: 54000061 b.ne 400674 <skip>
40066c: 50000600 adr x0, 40072e <msg2>
400670: 97ffff9c bl 4004e0 <printf@plt>

0000000000400674 <skip>:
400674: 52800000 mov w0, #0x0
400678: f94003fe ldr x30, [sp]
40067c: 910043ff add sp, sp, #0x10
400680: d65f03c0 ret
```

Additional code



74

Examining Machine Lang: TEXT

```
$ objdump --disassemble --reloc detecta
detecta: file format elf64-littleaarch64

...
0000000000400650 <main>:
400650: d10043ff sub sp, sp, #0x10
400654: f90003fe str x30, [sp]
400658: 10000640 adr x0, 400720 <msg1>
40065c: 97fffffa1 bl 4004e0 <printf@plt>
400660: 97ffff9c ldr w0, [x30, #0x10]
400664: 7101041f cmp w0, #0x41
400668: 54000061 b.ne 400674 <skip>
40066c: 50000600 adr x0, 40072e <msg2>
400670: 97ffff9c bl 4004e0 <printf@plt>

0000000000400674 <skip>:
400674: 52800000 mov w0, #0x0
400678: f94003fe ldr x30, [sp]
40067c: 910043ff add sp, sp, #0x10
400680: d65f03c0 ret
```

No relocation records!
Let's see what the linker did with them...



75

adr x0, 400720 <msg1>

```
$ objdump --disassemble --reloc detecta
detecta: file format elf64-littleaarch64

...
0000000000400650 <main>:
400650: d10043ff sub sp, sp, #0x10
400654: f90003fe str x30, [sp]
400658: 10000640 adr x0, 400720 <msg1>
40065c: 97fffffa1 bl 4004e0 <printf@plt>
400660: 97ffff9c ldr w0, [x30, #0x10]
400664: 7101041f cmp w0, #0x41
400668: 54000061 b.ne 400674 <skip>
40066c: 50000600 adr x0, 40072e <msg2>
400670: 97ffff9c bl 4004e0 <printf@plt>

0000000000400674 <skip>:
400674: 52800000 mov w0, #0x0
400678: f94003fe ldr x30, [sp]
40067c: 910043ff add sp, sp, #0x10
400680: d65f03c0 ret
```



76

adr x0, 400720 <msg1>

msb: bit 31 | 400658: 10000640 adr x0, 400720 <msg1> lsb: bit 0
↓ 0001 0000 0000 0000 0000 0110 0100 0000

- opcode: generate address
- 19 High-order bits of offset in bits 5-23: 11001000b = 0xc8 bytes after this instruction
- 2 Low-order bits of offset in bits 29-30: 00
- *Relative data location is 11001000b = 0xc8 bytes after this instruction*
- Destination register in bits 0-4:0

• msg1 is at 0x400720; this instruction is at 0x400658 ✓



77

bl 4004e0 <printf@plt>

```
$ objdump --disassemble --reloc detecta
detecta: file format elf64-littleaarch64

...
0000000000400650 <main>:
400650: d10043ff sub sp, sp, #0x10
400654: f90003fe str x30, [sp]
400658: 10000640 adr x0, 400720 <msg1>
40065c: 97fffffa1 bl 4004e0 <printf@plt>
400660: 97ffff9c ldr w0, [x30, #0x10]
400664: 7101041f cmp w0, #0x41
400668: 54000061 b.ne 400674 <skip>
40066c: 50000600 adr x0, 40072e <msg2>
400670: 97ffff9c bl 4004e0 <printf@plt>

0000000000400674 <skip>:
400674: 52800000 mov w0, #0x0
400678: f94003fe ldr x30, [sp]
40067c: 910043ff add sp, sp, #0x10
400680: d65f03c0 ret
```



78

bl 4004e0 <printf@plt>

msb: bit 31 | 40065c: 97fffffa1 bl 4004e0 <printf@plt> lsb: bit 0
↓ 1001 0111 1111 1111 1111 1111 1010 0001

- opcode: branch and link
- *Relative address in bits 0-25: 26-bit two's complement of 1011111b. But remember to shift left by two bits (see earlier slides)! This gives -101111100b = -0x17c*
- printf is at 0x4004e0; this instruction is at 0x40065c
- 0x4004e0 - 0x40065c = -0x17c ✓



79

Everything Else is Similar...

```
$ objdump --disassemble --reloc detecta

detecta: file format elf64-littleaarch64

...
0000000000400650 <main>:
400650: d10043ff sub sp, sp, #0x10
400654: f90003fe str x30, [sp]
400658: 10000010 ldr x0, [sp] <msg1>
40065c: 91fffffa1 bl 4004e0 <printf@plt>
400660: 97ffff9c bl 4004d0 <getchar@plt>
400664: 7101041f cmp w0, #0x41
400668: 54000061 b.ne 400674 <skip>
40066c: 50000600 adr x0, 40072e <msg2>
400670: 97ffff9c bl 4004e0 <printf@plt>

0000000000400674 <skip>:
400674: 55800000 mov w0, #0x0
400678: f94003fe ldr x30, [sp]
40067c: 910043ff add sp, sp, #0x10
400680: d65f03c0 ret
```



80

Summary



AARCH64 Machine Language

- 32-bit instructions
- Formats have conventional locations for opcodes, registers, etc.

Assembler

- Reads assembly language file
- Generates TEXT, RODATA, DATA, BSS sections
 - Containing machine language code
- Generates **relocation records**
- Writes object (.o) file

Linker

- Reads object (.o) file(s)
- Does **resolution**: resolves references to make code complete
- Does **relocation**: traverses relocation records to patch code
- Writes executable binary file

81

81