

CS 3330: More C

30 August 2016

Layers of Abstraction

x += y

"Higher-level" language: C

add %rbx, %rax

Assembly: X86-64

60 03

Machine code: Y86

(we'll talk later)

Logic and Registers

1

2

Last time

compile / assemble / link

C data types, lack of bool

arrays and pointers

command-line tips:

man, chmod, tar, tab completion, history

printf and <stdio.h>

printf

```
int custNo = 1000;
const char *name = "Jane\u00a0Smith"
printf("Customer\u00a0#\u00d3:\u00a0%s\n",
       custNo, name);
// "Customer #1000: Jane Smith"
// same as:
cout << "Customer#\u00a0" << custNo
     << ":"&u00a0" << name << endl;
```

format string must **match types** of argument

3

4

printf formats quick reference

Specifier	Argument Type	Example(s)
%s	char *	Hello, World!
%p	any pointer	0x4005d4
%d	int/short/char	42
%u	unsigned int/short/char	42
%x	unsigned int/short/char	2a
%ld	long	42
%f	double/float	42.000000 0.000000 4.200000e+01 4.200000e-19
%e	double/float	42, 4.2e-19
%g	double/float	
%%	(no argument)	%

5

goto

```
for (...) {  
    for (...) {  
        if (thingAt(i, j) == found) {  
            goto found;  
        }  
    }  
}  
printf("not found!\n");  
return;  
found:  
printf("found!\n");
```

assembly:
jmp foundassembly:
found:

6

struct

```
struct rational {  
    int numerator;  
    int denominator;  
};  
// ...  
struct rational two_and_a_half;  
two_and_a_half.numerator = 5;  
two_and_a_half.denominator = 2;  
struct rational *pointer = &two_and_a_half;  
printf("%d/%d\n",  
    pointer->numerator,  
    pointer->denominator);
```

typedef struct (1)

```
struct other_name_for_rational {  
    int numerator;  
    int denominator;  
};  
typedef struct other_name_for_rational rational;  
// ...  
rational two_and_a_half;  
two_and_a_half.numerator = 5;  
two_and_a_half.denominator = 2;  
rational *pointer = &two_and_a_half;  
printf("%d/%d\n",  
    pointer->numerator,  
    pointer->denominator);
```

7

8

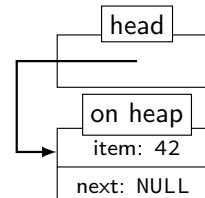
typedef struct (2)

```
struct other_name_for_rational {  
    int numerator;  
    int denominator;  
};  
typedef struct other_name_for_rational rational;  
// same as:  
typedef struct other_name_for_rational {  
    int numerator;  
    int denominator;  
} rational;  
// almost the same as:  
typedef struct {  
    int numerator;  
    int denominator;  
} rational;
```

9

linked lists / dynamic allocation

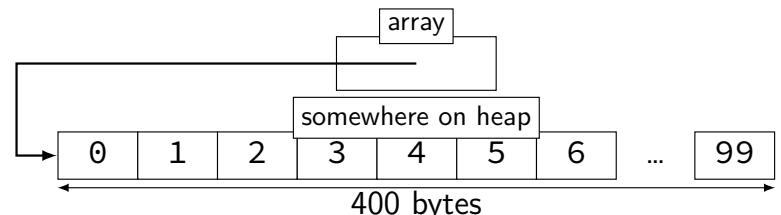
```
typedef struct list_t {  
    int item;  
    struct list_t *next;  
} list;  
// ...  
  
list* head = malloc(sizeof(list));  
/* C++: new list; */  
head->item = 42;  
head->next = NULL;  
// ...  
free(head);  
/* C++: delete list */
```



10

dynamic arrays

```
int *array = malloc(sizeof(int)*100);  
// C++: new int[100]  
for (i = 0; i < 100; ++i) {  
    array[i] = i;  
}  
// ...  
free(array); // C++: delete[] array
```



Miss vector? (1)

```
typedef struct range_t {  
    int size;  
    int *data;  
} range;  
  
range vec;  
vec.size = 100;  
vec.data = malloc(sizeof(int) * 100);  
// like: vector<int> vec(100);
```

11

12

Miss vector? (2)

```
typedef struct range_t {  
    int size;  
    int *data;  
} range;  
  
range vec2;  
vec2.size = vec.size;  
vec2.data = malloc(sizeof(int) * vec.size);  
for (int i = 0; i < vec.size; ++i) {  
    vec2.data[i] = vec.data[i];  
}  
// like: vector<int> vec2 = vec;  
Why not range vec2 = vec?
```

unsigned and signed types

type	min	max
signed int = signed = int	-2^{31}	$2^{31} - 1$
unsigned int = unsigned	0	$2^{32} - 1$
signed long = long	-2^{63}	$2^{63} - 1$
unsigned long	0	$2^{64} - 1$

⋮

unsigned/signed comparison trap

```
int x = -1;  
unsigned int y = 0;  
printf("%d\n", x < y);
```

result is 0

short solution: don't compare signed to unsigned:

```
(long) x < (long) y
```

```
int x = -1;  
unsigned int y = 0;  
printf("%d\n", x < y);
```

compiler converts both to **same type** first
int if all possible values fit

otherwise: first operand (x, y) type from this list:
 unsigned long
 long
 unsigned int
 int

C evolution and standards

1978: Kernighan and Ritchie publish *The C Programming Language* — “K&R C”
very different from modern C

1989: ANSI standardizes C — C89/C90/-ansi
compiler option: -ansi, -std=c90
looks mostly like modern C

1999: ISO (and ANSI) update C standard — C99
compiler option: -std=c99
adds: declare variables in middle of block
adds: // comments

2011: Second ISO update — C11

Middle of blocks?

Examples of things not allowed in 1989 ANSI C:

```
printf("Before\u002c calling\u002c malloc()\n");
int *pointer = malloc(sizeof(int) * 100);

for (int x = 0; x < 10; ++x)
```

17

18

Undefined behavior example (1)

```
#include <stdio.h>
#include <limits.h>
int test(int number) {
    return (number + 1) > number;
}

int main(void) {
    printf("%d\n", test(INT_MAX));
}
```

without optimizations: 0
with optimizations: 1

Undefined behavior example (2)

```
int test(int number) {
    return (number + 1) > number;
}
```

Optimized:

```
test:
    movl $1, %eax      ; eax <- 1
    ret
```

Less optimized:

```
test:
    leal 1(%rdi), %eax ; eax <- rdi + 1
    cmpl %eax, %edi
    setl %al           ; al <- eax < edi
    movzbl %al, %eax   ; eax <- al
    ret
```

19

20

Undefined behavior

compilers can do **whatever they want**

what you expect
crash your program

...

common types:

signed integer overflow/underflow
out-of-bounds pointers
integer divide-by-zero
writing read-only data
out-of-bounds shift (later)

Bit-twiddling

21

22

some truth tables

AND	0	1
0	0	0
1	0	1

OR	0	1
0	0	1
1	1	1

XOR	0	1
0	0	1
1	1	0

&&, &

||, |

(nothing), ^

Logical Operators

Treat value as true (1) or false (0)

Recall: false = 0 (only)

1	&&	1	==	1	1		1	==	1
2	&&	4	==	1	2		4	==	1
1	&&	0	==	0	1		0	==	1
0	&&	0	==	0	0		0	==	0
-1	&&	-2	==	1	-1		-2	==	1
""	&&	""	==	1	""		""	==	1

23

24

Short-Circuit (`&&`)

```
1 #include <stdio.h>
2 int zero() { printf("zero()\n"); return 0; }
3 int one() { printf("one()\n"); return 1; }
4 int main() {
5     printf(">%d\n", zero() && one());
6     printf(">%d\n", one() && zero());
7     return 0;
8 }
```

zero()

> 0

one()

zero()

> 0

	0	1
0	0	0
1	0	1

25

Short-Circuit (`||`)

```
1 #include <stdio.h>
2 int zero() { printf("zero()\n"); return 0; }
3 int one() { printf("one()\n"); return 1; }
4 int main() {
5     printf(">%d\n", zero() || one());
6     printf(">%d\n", one() || zero());
7     return 0;
8 }
```

zero()

one()

> 1

one()

> 1

	0	1
0	0	1
1	1	1

26

Bitwise AND — `&`

Treat value as **array of bits**

`1 & 1 == 1`

$$\begin{array}{r} \dots 0 0 1 0 \\ \& \dots 0 1 0 0 \\ \hline \dots 0 0 0 0 \end{array}$$

`1 & 0 == 0`

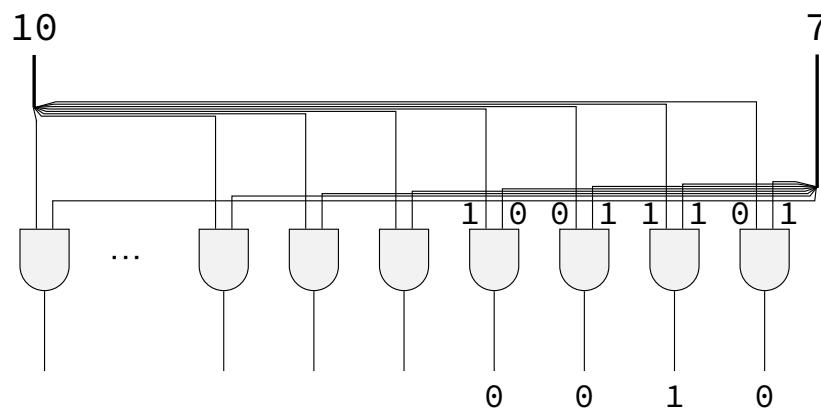
`0 & 0 == 0`

`2 & 4 == 0`

$$\begin{array}{r} \dots 1 0 1 0 \\ \& \dots 0 1 1 1 \\ \hline \dots 0 0 1 0 \end{array}$$

`10 & 7 == 2`

Bitwise hardware (`10 & 7 == 2`)



27

28

Bitwise OR — |

$$1 \quad | \quad 1 \quad == \quad 1$$

$$1 \mid 0 == 1$$

θ | θ ≡≡ θ

3 | 4 == 6

10 | 7 == 15

$$\begin{array}{r|ccccc} & \dots & 0 & 0 & 1 & 0 \\ | & \dots & 0 & 1 & 0 & 0 \\ \hline & \dots & 0 & 1 & 1 & 0 \end{array}$$

$$\begin{array}{r|ccccc} & \dots & 1 & 0 & 1 & 0 \\ | & \dots & 0 & 1 & 1 & 1 \\ \hline & \dots & 1 & 1 & 1 & 1 \end{array}$$

31

Bitwise xor — ^

$$1 \wedge 1 == 0$$

$$1 \wedge 0 == 1$$

$$\theta \wedge \theta = \theta$$

3 \wedge 4 = 6

$$10^7 = 13$$

$$\begin{array}{r}
 & \dots & 0 & 0 & 1 & 0 \\
 \wedge & \dots & 0 & 1 & 0 & 0 \\
 \hline
 & \dots & 0 & 1 & 1 & 0
 \end{array}$$

$$\begin{array}{r}
 & \dots & 1 & 0 & 1 & 0 \\
 \wedge & \dots & 0 & 1 & 1 & 1 \\
 \hline
 & 1 & 1 & 0 & 1
 \end{array}$$

30

Negation / Not — ~

`~` ('complement') is bitwise version of `!`:

$$! \theta == 1$$

`!notZero == 0`

`~0 == (int) 0xFFFFFFFF`

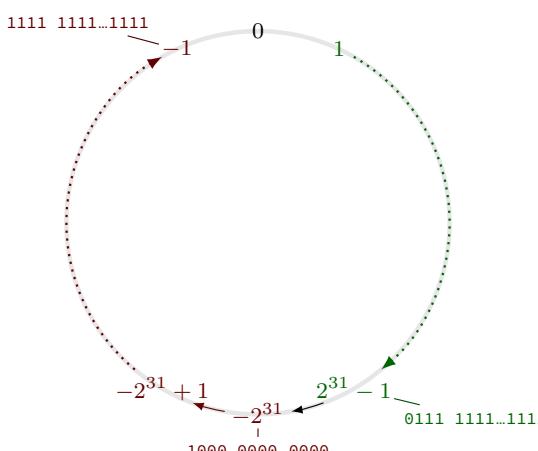
$$\sim 2 == -3$$

$$\sim -7 == 6$$

```
~((unsigned) 2) == 0xFFFFFFFFED
```

Two's complement refresher

$$-1 = \begin{matrix} -2^{31} & +2^{30} & +2^{29} \\ 1 & 1 & 1 \end{matrix} \quad \begin{matrix} +2^2 & +2^1 & +2^0 \\ 1 & 1 & 1 \end{matrix}$$



32

Two's Complement and ~

$-x == \sim x + 1$ (flip the bits and add one)

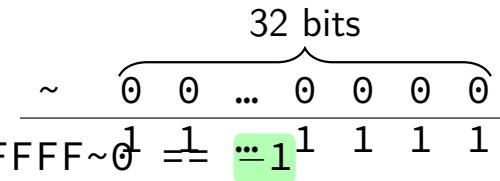
$-x - 1 == \sim x$

Negation / Not — ~

~ ('complement') is bitwise version of !:

$\text{!}0 == 1$

$\text{!notZero} == 0$

\sim 

$\sim 0 == (\text{int}) 0xFFFFFFFF$

$\sim 2 == -3$

$\sim -7 == 6$

$\sim ((\text{unsigned}) 2) == 0xFFFFFFFF$

Left shift

$1 \ll 0 == 1$ 0000 0001

$1 \ll 1 == 2$ 0000 0010

$1 \ll 2 == 4$ 0000 0100

$10 \ll 0 == 10$ 0000 1010

$10 \ll 1 == 20$ 0001 0100

$10 \ll 2 == 40$ 0010 1000

$$x \ll y = x \times 2^y$$

Right shift

Undefined: ~~$x \ll -1$~~ Instead: $x >> 1$

$1 >> 0 == 1$ 0000 0001

$1 >> 1 == 0$ 0000 0000

$1 >> 2 == 0$ 0000 0000

$10 >> 0 == 10$ 0000 1010

$10 >> 1 == 5$ 0000 0101

$10 >> 2 == 2$ 0000 0010

$$x >> y = \lfloor x \times 2^{-y} \rfloor$$

Shifts and negative numbers

`-10 >> 1 == ??? (-10 = 1111 ... 1111 0110)`

binary ?111 ... 1111 1011

Option 1: binary 1111 ... 1011 =
 $-5 = -10 \times 2^{-k}$

copy sign bit
arithmetic

Option 2: binary 0111 ... 1011 = $2^{31} - 5$

always use zero
logical

Aside: implementation-defined behavior

C standard: negative `>> 1` is
implementation-defined

compiler chooses

37

38

Arithmetic shift

`-1 >> 0 == -1`

1111 1111

`-1 >> 1 == -1`

1111 1111

`-1 >> 2 == -1`

1111 1111

`-10 >> 0 == -10`

1111 0110

`-10 >> 1 == -5`

1111 1011

`-10 >> 2 == -3`

1111 1101

`10 >> 2 == 2`

0000 0010

$$x >> y = \lfloor x \times 2^{-y} \rfloor$$

signed versus unsigned types

```
/*signed*/ int x = -10;  
/* arithmetic: */  
x >> 1 == -5  
x >> 4 == -1
```

```
unsigned int y = 0xFFFFFFFF6;  
/* logical */  
y >> 1 == 0x7FFFFFFB  
y >> 4 == 0xFFFFFFFF
```

39

40

Sign-extension vs. zero-extension

```
signed char x = -10;      // 1111 0110
int y = x;                // 1111.. 1111 0110
printf("%d\n", y);        // outputs "-10"

unsigned char x = 0xF6;    // 1111 0110
int y = x;                // 0000.. 1111 0110
printf("%d\n", y);        // outputs "246"
```

Aside: integer promotions

```
unsigned short number = 1;
unsigned short offset = 20;
printf("0x%x\n", number << offset);
```

Outputs (on lab machines)?

- A. 0x100000 (2^{20})
- B. 0x0
- C. Undefined behavior — varies

integers types smaller than **int** converted to **int**

41

42

Shifts: undefined behavior

`0 >> 32` is undefined behavior

`0 << 32` is undefined behavior

`(long) 0 << 32` is okay

`(long) 0 << 64` is undefined behavior

Summary

struct — functionless classes

typedef struct or write **struct typeName**
malloc, free instead of new/delete.

undefined behavior — who knows what'll happen

logical operators — `&&`, `||`, `!`: only care if 0/not 0

bitwise operators — `&`, `|`, `^`, `~`: all bits in parallel

bitshift — `<<`, `>>`: same as multiplying by $2^{\pm x}$

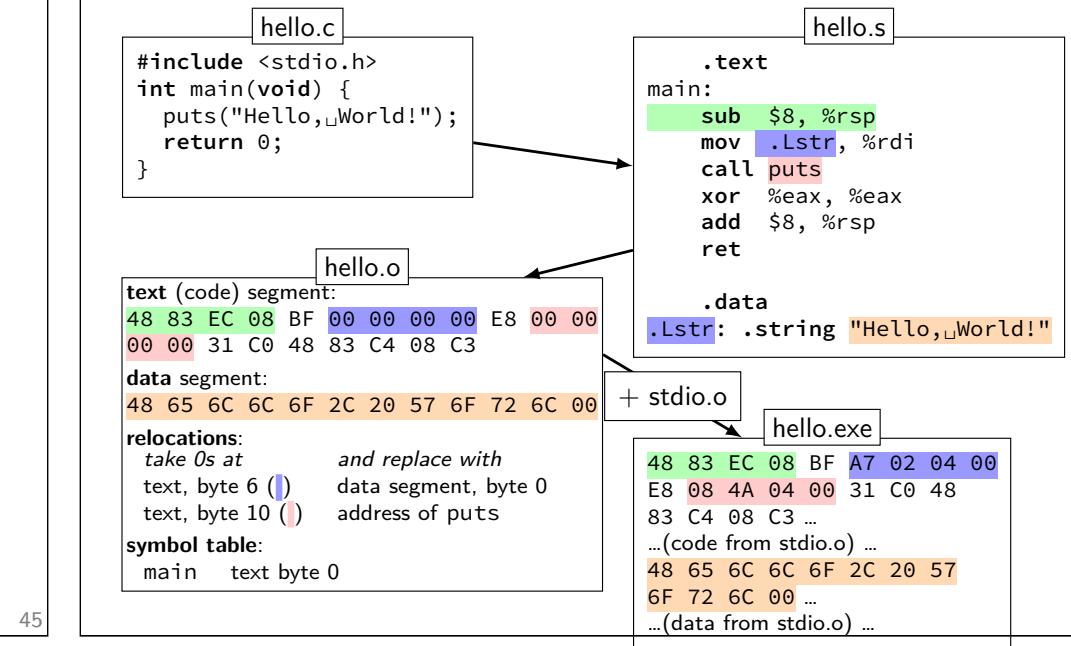
arithmetic right shift — borrow sign bit

43

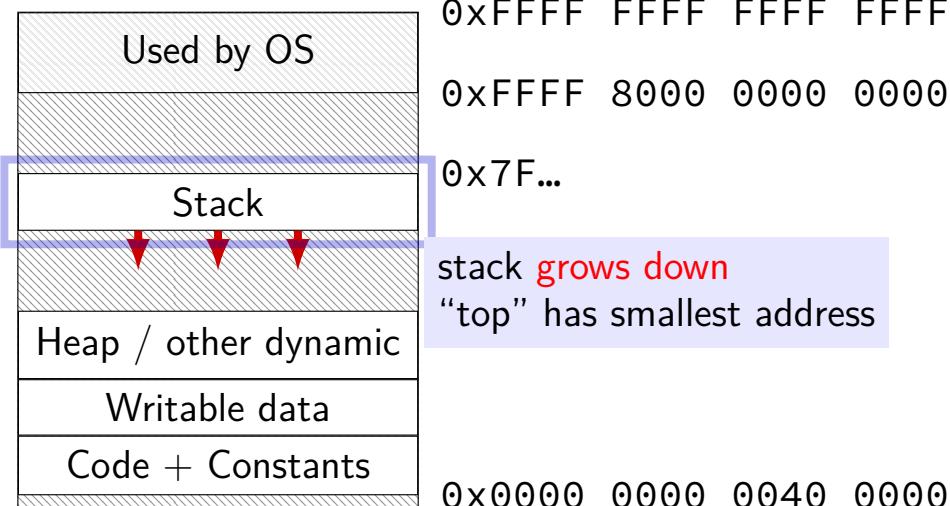
44

Backup Slides

What's in those files?



Program Memory (x86-64 Linux)



Arrays: not quite pointers (1)

```
int array[100];
int *pointer;
```

Legal: `pointer = array;`
same as `pointer = &(array[0]);`

Illegal: ~~`array = pointer;`~~

Arrays: not quite pointers (2)

```
int array[100];
int *pointer = array;

sizeof(array) == 400
    size of all elements

sizeof(pointer) == 8
    size of address

sizeof(&array[0]) == ???
    (&array[0] same as &(array[0]))
```

chmod

```
chmod --recursive og-r /home/USER
```

user (yourself) / group / others
- remove / + add
read / write / execute or search

49

50

tar

the standard Linux/Unix file archive utility

Table of contents: tar tf filename.tar

eXtract: tar xvf filename.tar

Create: tar cvf filename.tar directory

(v: verbose; f: file — default is tape)

51