

Lecture 0: Introduction to Computer Systems

CS 105

Fall 2023

<https://cs.pomona.edu/classes/cs105/>



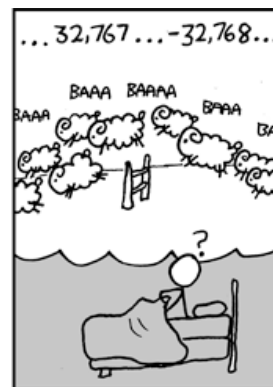
Abstraction



Correctness

- **Example 1: Is $x^2 \geq 0$?**

- Floats: Yes!



- Ints:

- $40000 * 40000 \rightarrow 1600000000$
- $50000 * 50000 \rightarrow ??$

- **Example 2: Is $(x + y) + z = x + (y + z)$?**

- Ints: Yes!

- Floats:

- $(2^{30} + -2^{30}) + 3.14 \rightarrow 3.14$
- $2^{30} + (-2^{30} + 3.14) \rightarrow ??$

Performance

```
void copyij(int src[2048][2048],
            int dst[2048][2048]){
    int i,j;
    for (i = 0; i < 2048; i++){
        for (j = 0; j < 2048; j++){
            dst[i][j] = src[i][j];
        }
    }
}
```

4.3ms

```
void copyji(int src[2048][2048],
            int dst[2048][2048]){
    int i,j;
    for (j = 0; j < 2048; j++){
        for (i = 0; i < 2048; i++){
            dst[i][j] = src[i][j];
        }
    }
}
```

81.8ms

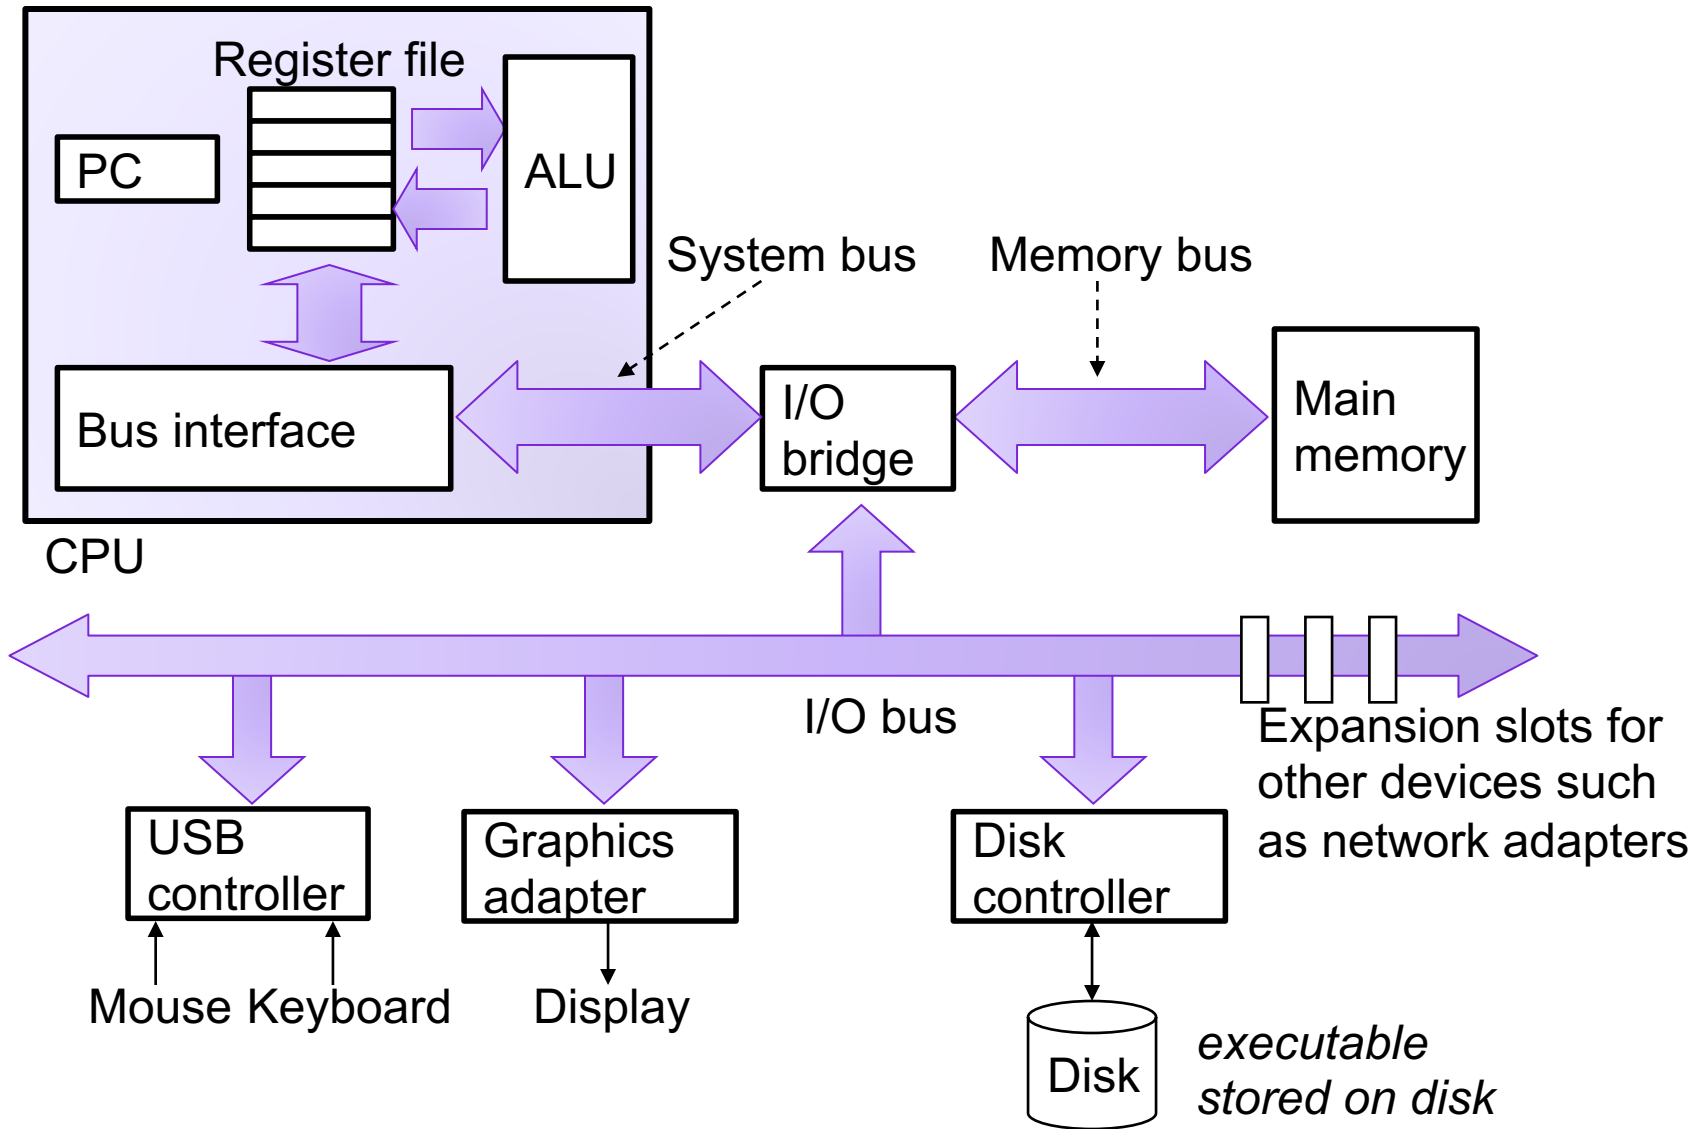
- Hierarchical memory organization
- Performance depends on access patterns
 - Including how step through multi-dimensional array

Security

```
void admin_stuff(int authenticated){
    if(authenticated){
        // do admin stuff
        // should only happen if user is authenticated
        printf("The answer is 42\n");
    }
}

int dontTryThisAtHome(char* user_input, int size) {
    char data[size];
    int ret = memcpy(*user_input, data);
    return ret;
}
```

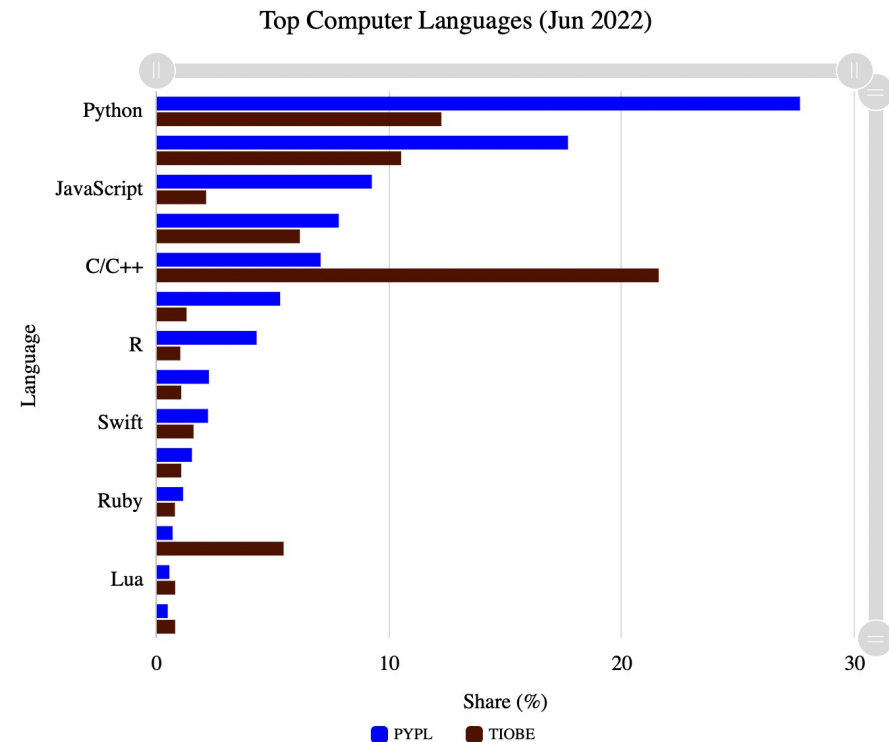
A Computer System



C

- compiled, imperative language that provides low-level access to memory
- low overhead, high performance

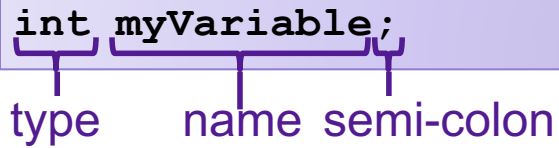
- developed at Bell labs in the 1970s
- C (and related languages) still today



Variables

- Declaration

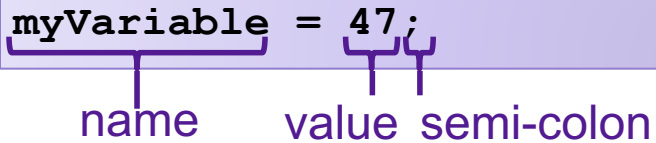
```
int myVariable;
```



A diagram showing the components of a variable declaration. The code `int myVariable;` is displayed in a light purple box. Brackets and lines identify the parts: `int` is labeled 'type', `myVariable` is labeled 'name', and `;` is labeled 'semi-colon'.

- Assignment

```
myVariable = 47;
```



A diagram showing the components of a variable assignment. The code `myVariable = 47;` is displayed in a light purple box. Brackets and lines identify the parts: `myVariable` is labeled 'name', `47` is labeled 'value', and `;` is labeled 'semi-colon'.

- Declaration and assignment

```
int myVariable = 47;
```


Operations

- Arithmetic Operations: +, -, *, /, %

```
int x = 47;  
int y = x + 13;  
y = (x * y) % 5;
```

- Boolean Operators: ==, !=, >, >=, <, <=

```
int x = (13 == 47);
```

- Logical Operations: &&, ||, !

```
int x = 47;  
int y = !x;  
y = x && y;
```

- Bitwise Operations: &, |, ~, ^

```
int x = 47;  
int y = ~x;  
y = x & y;
```

Control Flow

Conditionals

```
int x = 13;
int y;
if (x == 47) {
    y = 1;
} else {
    y = 0;
}
```

Do-While Loops

```
int x = 47;
do {
    x = x - 1;
} while (x > 0);
```

While Loops

```
int x = 47;

while (x > 0) {
    x = x - 1;
}
```

For Loops

```
int x = 0;
for (int i=0; i < 47; i++){
    x = x + i;
}
```

Functions

Declaring a Function

```
int myFunction(int x, int y){  
  
    int z = x - 2*y;  
    return z * x;  
  
}
```

Calling a Function

```
int a;  
  
a = myFunction(47, 13);
```

Exercise 1

- Define a function `sum_interval` that takes two integers and returns an integer. If the second integer argument is greater than (or equal to) the first, it returns the sum of the integer values between those two numbers (inclusive). Otherwise it returns -1.

Main Functions

- By convention, main functions in C take two arguments:
 1. `int argc`
 2. `char** argv`
- By convention, main functions in C return an int
 - 0 if program exited successfully

```
int main(int argc, char** argv){  
    // do stuff  
  
    return 0;  
}
```

Aside: Printing

```
printf("Hello world!\n");
```

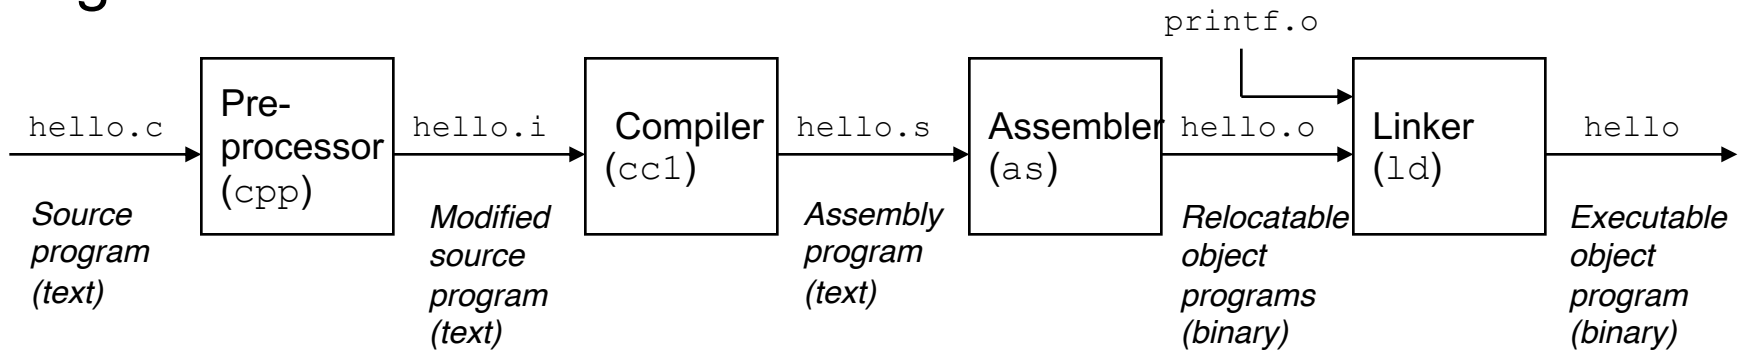
```
printf("%d is a number\n", 13);
```

```
printf("%d is a number greater than %f\n", 47, 3.14);
```

Compilation

compiler output name filename

- gcc -o hello hello.c



```
#include<stdio.h>

int main(int argc,
         char ** argv){

    printf("Hello
           world!\n");

    return 0;
}
```

```
...
int printf(const char *
           restrict,
           ...)
    __attribute__((__format__
                  (__printf__, 1, 2)));
...
int main(int argc,
         char ** argv){

    printf("Hello
           world!\n");

    return 0;
}
```

```
pushq   %rbp
movq    %rsp, %rbp
subq    $32, %rsp
leaq   L_.str(%rip), %rax
movl   $0, -4(%rbp)
movl   %edi, -8(%rbp)
movq   %rsi, -16(%rbp)
movq   %rax, %rdi
movb   $0, %al
callq  _printf
xorl   %ecx, %ecx
movl   %eax, -20(%rbp)
movl   %ecx, %eax
addq   $32, %rsp
popq   %rbp
retq
```

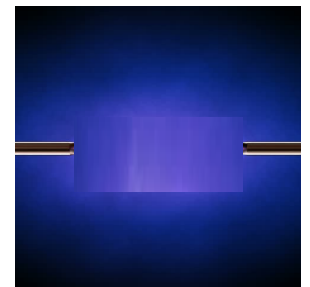
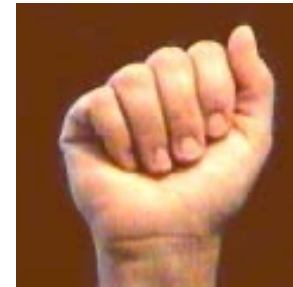
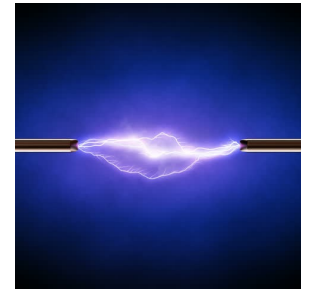
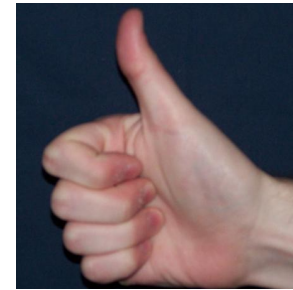
```
55
48 89 e5
48 83 ec 20
48 8d 05 25 00 00 00
c7 45 fc 00 00 00 00
89 7d f8
48 89 75 f0
48 89 c7
b0 00
e8 00 00 00 00
31 c9
89 45 ec
89 c8
48 83 c4 20
5d
c3
```

Running a Program

- `./hello`

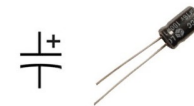
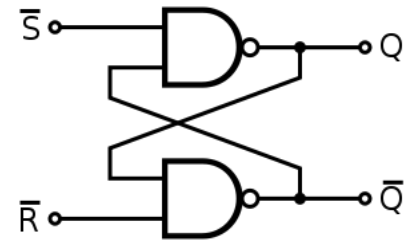
Bits

- a **bit** is a binary digit that can have two possible values
- can be physically represented with a two state device



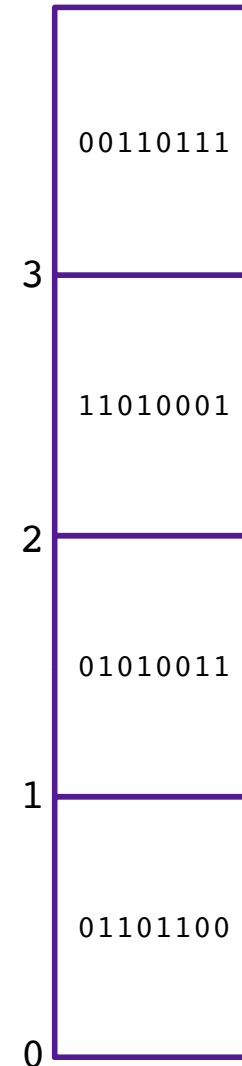
Storing bits

- Static random access memory (SRAM): stores each bit of data in a flip-flop, a circuit with two stable states
- Dynamic Memory (DRAM): stores each bit of data in a capacitor, which stores energy in an electric field (or not)
- Magnetic Disk: regions of the platter are magnetized with either N-S polarity or S-N polarity
- Optical Disk: stores bits as tiny indentations (pits) or not (lands) that reflect light differently
- Flash Disk: electrons are stored in one of two gates separated by oxide layers



Bytes and Memory

- **Memory** is an array of ~~bits~~^{bytes}
- A **byte** is a unit of eight bits
- An index into the array of memory is an **address**, **location**, or **pointer**
 - Often expressed in hexadecimal
- We speak of the *value* in memory at an address
 - The value may be a single byte ...
 - ... or a multi-byte quantity starting at that address



C Types

C Data Type	x86-64
int	4
short	2
long	8
unsigned short	2
unsigned int	4
unsigned long	8
char	1
unsigned char	1
float	4
double	8

Pointers

- Pointers are addresses in memory (i.e., indexes into the array of bytes)
- Most pointers declare how to interpret the value at (or starting at) that address
- Example:

```
int myVariable = 47;  
int* ptr = &myVariable;
```

- Dereferencing pointers:

```
int var2 = *ptr
```

Pointer Types	x86-64
void*	8
int*	8
char*	8
:	8

& is an "address of" operator
* is a "value at" operator

& and * are inverses of one another

Exercise 2

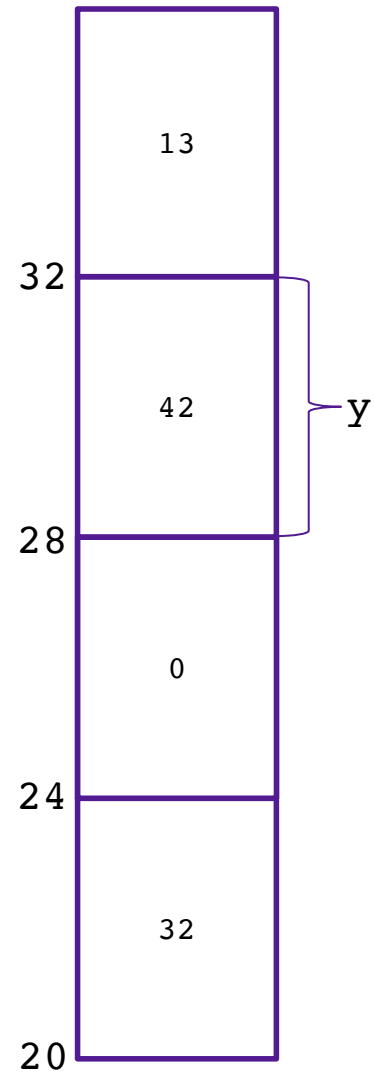
What does x evaluate to in each of the following?

1. `int* ptr = 32;`
`x = *ptr;`

2. `int y = 42; // assume allocated at address 28`
`x = &y;`

3. `int* x = 24;`
`*x = 47;`

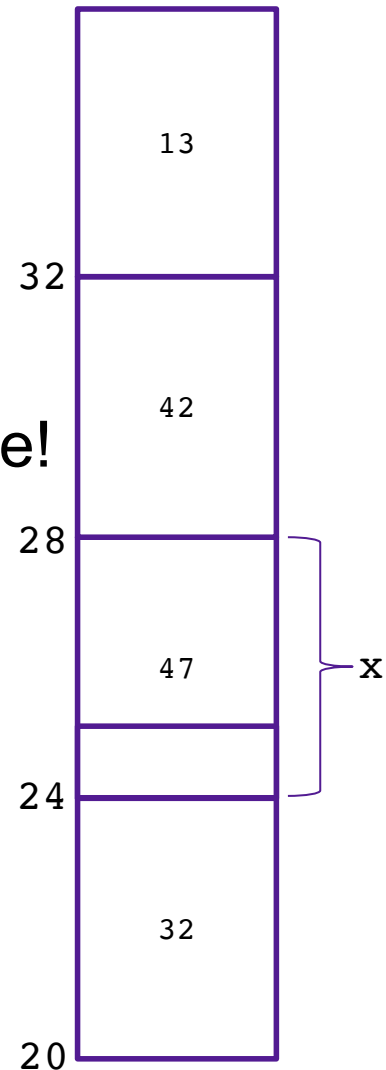
4. `int* ptr = 20;`
`x = *(*ptr);`



Casting between Pointer Types

- You can cast values between different types
- This includes between different pointer types!
- Doesn't change value of address
- Does change what you get when you dereference!
- Example:

```
int x = 47; // assume allocated at address 24
int* ptr = &x;
char* ptr2 = (char*) ptr;
int y = *ptr
char c = *ptr2;
```



Arrays

- Contiguous block of memory
- Random access by index
 - Indices start at zero

- Declaring an array:

```
int array1[5]; // array of 5 ints named array1  
  
char array2[47]; // array of 47 chars named array2  
  
int array3[7][4]; // two dimensional array named array3
```

- Accessing an array:

```
int x = array1[0];
```

- Arrays are pointers!

- The array variable stores the address of the first element in the array

Pointer Arithmetic

```
int * ptr = &myVariable;  
ptr += 1;  
  
char * ptr2 = (char *) ptr;  
ptr2 += 1;
```

- Location of `ptr+k` depends on the type of `ptr`
- adding 1 to a pointer `p` adds `1*sizeof(*p)` to the address
- `array[k]` is the same as `*(array+k)`

Exercise 3

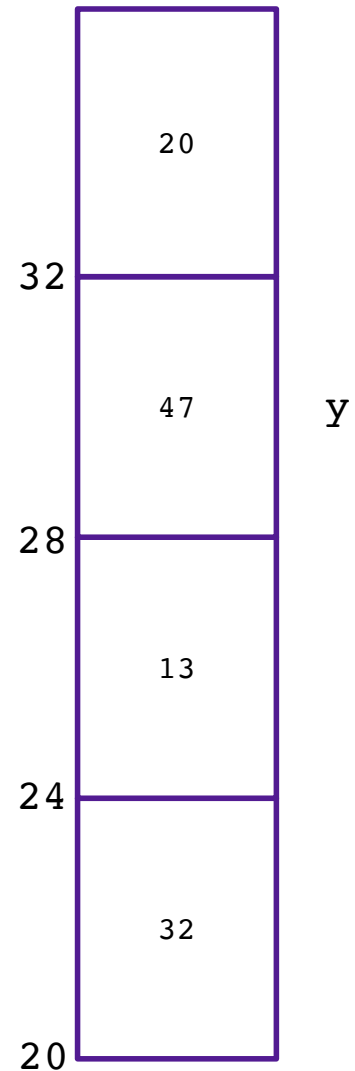
What does x evaluate to in each of the following?

1. `int* ptr = 20;`
`int* x = ptr+1;`

2. `int* ptr = 20;`
`int x = *(ptr+2)`

3. `char* ptr = 20;`
`char* x = ptr+1;`

4. `char* ptr = 20;`
`int x = *((int*)(ptr + 4));`



Strings

- Strings are just arrays of characters
 - aka strings are just pointers
- declared as type `char*`
- End of string is denoted by null byte `\0`

Structs

- Heterogeneous records, like objects

- Typical linked list declaration:

```
typedef struct cell {  
    int value;  
    struct cell *next;  
} cell_t;
```

- Usage:

```
cell_t c;  
c.value = 42;  
c.next = NULL;
```

- Usage with pointers:

```
cell_t *p;  
p->value = 42;  
p->next = NULL;
```

`p->next` is an
abbreviation for
`(*p).next`

LOGISTICS

The Course in a Nutshell

- Textbooks (not required)
 - Bryant and O'Halloran, *Computer Systems: A Programmer's Perspective*, **third edition**, Pearson, 2016
 - Arpaci-Dusseau and Arpaci-Dusseau, *Operating Systems: Three Easy Pieces*, online, 2018
- Classes
 - Monday and Wednesday, 11am – 12:15pm in Edmunds 114
- Labs
 - Wednesdays 7-8:15 in Edmunds 229/219
 - **Starts Wednesday!**
- Office Hours M 7-9pm and T 1-2:30pm
- Mentor Sessions TBA

Grading

- Assignments (10)
 - Introduced during labs, Due Tuesdays at 11:59pm
 - Tremendous fun, work in pairs
 - 10 late days
- Check-ins (5)
 - three-question quizzes
 - Sept 20, Oct 11, Nov 1, Nov 20, Dec 6
 - Can improve grade on any question(s) during "Extra Chance Check-in"
- Grades
 - Must successfully complete all the assignments
 - Beyond that, 50% assignments, 45% check-ins, 5% participation

Course website

<https://cs.pomona.edu/classes/cs105>



- All information is on the course website
- All course materials get posted on the course website
- Links from the course page:
 - Slack (#cs105-2023fa), for questions and discussion
 - Gradescope, for submitting assignments and seeing grades
 - Additional resources