

CS 453/698: Software and Systems Security

Module: Background

Lecture: Abstractions in OS, PL, and SE

Meng Xu (*University of Waterloo*)

Spring 2025

Outline

- 1 Introduction
- 2 Abstractions Done by Compilers

Layered abstraction

Modern computing systems are among the most complex systems ever built.

One of the key engineering techniques that enables the construction of such complex systems is the use of **layered abstractions**:

Layered abstraction

Modern computing systems are among the most complex systems ever built.

One of the key engineering techniques that enables the construction of such complex systems is the use of **layered abstractions**:

- the system is designed as a stack of layers, where
- each layer hides implementation details of lower layers.

The hello-world example

```
1 #include <stdio.h>
2
3 int main(void) {
4     printf("Hello World");
5     return 0;
6 }
```

The hello-world example

```
1 #include <stdio.h>
2
3 int main(void) {
4     printf("Hello World");
5     return 0;
6 }
```

Compile:

cc hello-world.c

The hello-world example

```
1 #include <stdio.h>
2
3 int main(void) {
4     printf("Hello World");
5     return 0;
6 }
```

Compile:

cc hello-world.c

Execute:

./a.out

The hello-world example

```
1 #include <stdio.h>
2
3 int main(void) {
4     printf("Hello World");
5     return 0;
6 }
```

Compile:

cc hello-world.c

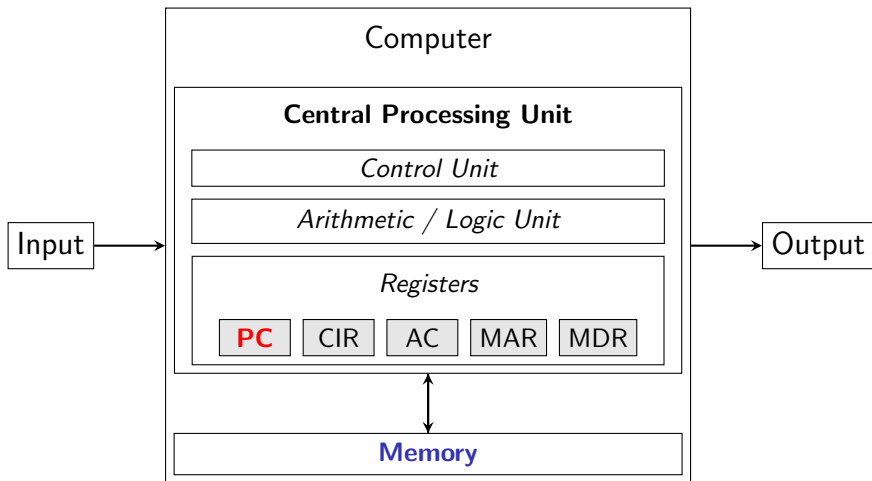
Execute:

./a.out

Q: What happens behind the scenes exactly?

Von Neumann architecture

Von Neumann architecture



Program the low-level machine

Suppose there is a CPU instruction called **output** **<char>**, with opcode **0B** **<char>**, which sends a single character **<char>** to the output device.

Q: How to display “Hello World” in the output device?

Program the low-level machine

Suppose there is a CPU instruction called **output** **<char>**, with opcode **0B** **<char>**, which sends a single character **<char>** to the output device.

Q: How to display “Hello World” in the output device?

A: This is a multi-step process:

Step 1: Find a suitable memory location (e.g., address **0x0010**)

Step 2: Put the following bytes into this memory location

0B 48 // ASCII code for 'H'

0B 65 // ASCII code for 'e'

...

0B 64 // ASCII code for 'd'

Step 3: Put value **0x0010** into the **PC** register.

A simplified view of compilation and loading

In this overly simplified example, we consider

- getting the bytes 0B 48 0B 65 ... 0B 64 from source code as **compilation**, and

A simplified view of compilation and loading

In this overly simplified example, we consider

- getting the bytes 0B 48 0B 65 ... 0B 64 from source code as **compilation**, and
- the rest as **loading**, including
 - ① Find a suitable memory location (e.g., address 0x0010)
 - ② Put the bytes 0B 48 0B 65 ... 0B 64 into this memory location
 - ③ Put value 0x0010 into the PC register.

Reality is more complicated

However, in reality, things are way more complicated. But the operating system, compiler, and software engineering practices **abstract** the complications away.

Outline

- 1 Introduction
- 2 Abstractions Done by Compilers

A simple C program

```
1  #include <stdio.h>
2  #include <string.h>
3
4  int main(void) {
5      int pass = 0;
6      char buff[8];
7
8      printf("Enter the password: ");
9      gets(buff);
10
11     if(strcmp(buff, "warriors")) {
12         printf("Wrong password\n");
13     } else {
14         printf("Correct password\n");
15         pass = 1;
16     }
17
18     if(pass) {
19         printf ("Root privileges granted\n");
20     }
21     return 0;
22 }
```

A simple C program

```
1  #include <stdio.h>
2  #include <string.h>
3
4  int main(void) {
5      int pass = 0;
6      char buff[8];
7
8      printf("Enter the password: ");
9      gets(buff);
10
11     if(strcmp(buff, "warriors")) {
12         printf("Wrong password\n");
13     } else {
14         printf("Correct password\n");
15         pass = 1;
16     }
17
18     if(pass) {
19         printf ("Root privileges granted\n");
20     }
21     return 0;
22 }
```

Try with

gcc -m64 -fno-stack-protector

And password "golden-hawks"

Stack layout (Linux x86-64 convention)

```
1 long foo(  
2     long a, long b, long c,  
3     long d, long e, long f,  
4     long g, long h)  
5 {  
6     long xx = a * b * c;  
7     long yy = d + e + f;  
8     long zz = bar(xx, yy, g + h);  
9     return zz + 20;  
10 }
```

High address

RBP + 24

h

RBP + 16

g

RBP + 8

return address

RBP

saved rbp

RBP - 8

xx

RBP - 16

yy

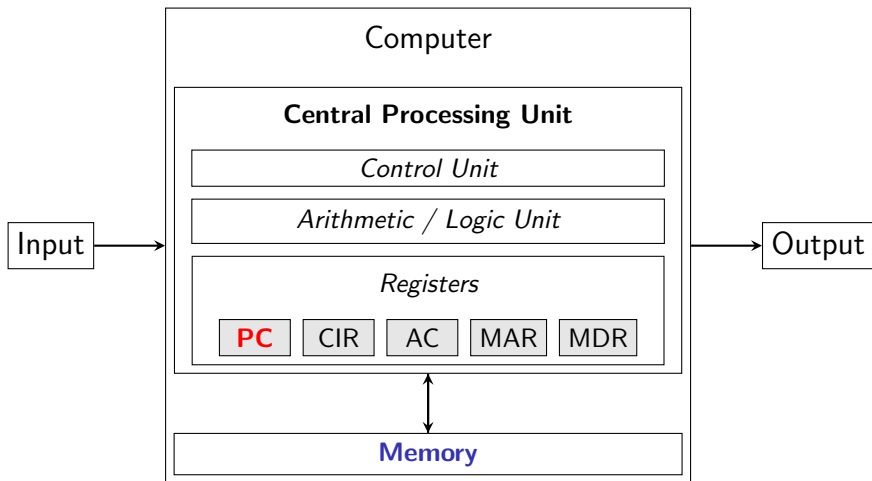
RBP - 24

zz

Low address

Argument a to f passed by registers.

Von Neumann architecture



Implications of the Von Neumann architecture

- Code and data reside in the same memory space and can be addressed in a unified way
 - If you manage to get the PC register to point to a memory address contains your logic, you have effectively hijacked the control flow.

Implications of the Von Neumann architecture

- Code and data reside in the same memory space and can be addressed in a unified way
 - If you manage to get the PC register to point to a memory address contains your logic, you have effectively hijacked the control flow.
- There is only one unified memory. It is the job of the compiler / programming language / runtime to find a way to utilize the memory efficiently.
 - Variables declared in a program (e.g., `int i = 0;`) need to be mapped to an address in the memory, and the mapping logic needs to be (ideally) consistent on the same architecture.

Definition: memory

Q: What is a conventional way of dividing up the “memory”?

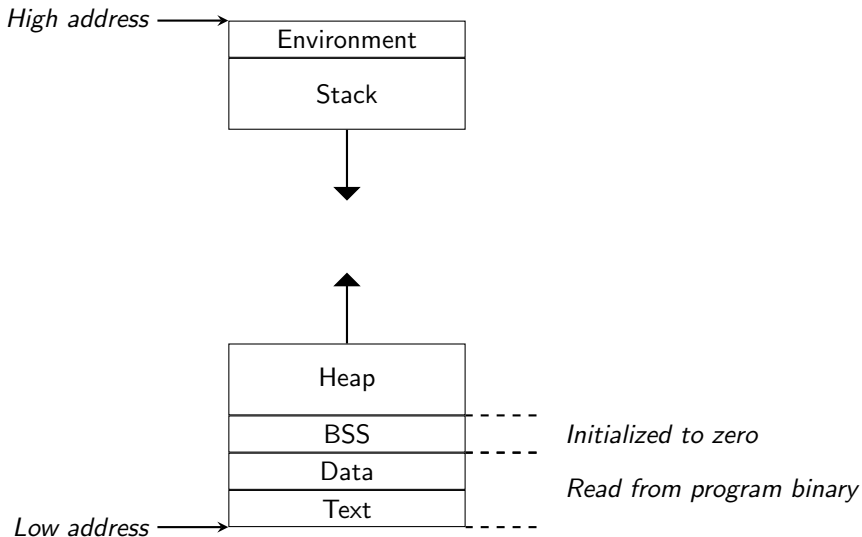
Definition: memory

Q: What is a conventional way of dividing up the “memory”?

A: Four types of memory on a conceptual level:

- Text (where program code is initially loaded to)
- Stack
- Heap
- Global (a.k.a., static)

Memory layout (Linux x86-64 convention)



Example

```
1  #include <stdlib.h>
2
3  //! where is this variable hosted?
4  const char *HELLO = "hello";
5
6  //! where is this variable hosted?
7  long counter;
8
9  void main() {
10     //! where is this variable hosted?
11     int val;
12
13     //! where is this variable hosted?
14     //! where is its content allocated?
15     char *msg = malloc(120);
16
17     //! what is freed here?
18     free(msg);
19
20     //! what is freed here (at end of function)?
21 }
22
23 //! what is freed here (at end of execution)?
```

Example (and answers)

```
1  #include <stdlib.h>
2
3  // this is in the data section
4  const char *HELLO = "hello";
5
6  // this is in the BSS section
7  long counter;
8
9  void main() {
10     // this is in the stack memory
11     int val;
12
13     // the msg pointer is in the stack memory
14     // the msg content is in the heap memory
15     char *msg = malloc(120);
16
17     // msg content is explicitly freed here
18     free(msg);
19
20     // the val and msg pointer is implicitly freed here
21 }
22
23 // the global memory is only destroyed on program exit
```

What is heap and why do we need it?

In C/C++, the **heap** is used to manually allocate (and free) **new regions of process memory** during program execution.

Heap vs stack

```
1  typedef struct Response {
2      int status;
3      char message[40];
4  } response_t;
5
6  response_t *say_hello() {
7      response_t* res =
8          malloc(sizeof(response_t));
9      if (res != NULL) {
10         res->status = 200;
11         strncpy(res->message, "hello", 6);
12     }
13     return res;
14 }
15 void send_back(response_t *res) {
16     // implementation omitted
17 }
18 void process() {
19     response_t *res = say_hello();
20     send_back(res);
21     free(res);
22 }
```

Heap vs stack

```
1 typedef struct Response {
2     int status;
3     char message[40];
4 } response_t;
5
6 response_t *say_hello() {
7     response_t* res =
8         malloc(sizeof(response_t));
9     if (res != NULL) {
10         res->status = 200;
11         strncpy(res->message, "hello", 6);
12     }
13     return res;
14 }
15 void send_back(response_t *res) {
16     // implementation omitted
17 }
18 void process() {
19     response_t *res = say_hello();
20     send_back(res);
21     free(res);
22 }
```

```
1 typedef struct Response {
2     int status;
3     char message[40];
4 } response_t;
5
6 void say_hello(response_t *res) {
7     res->status = 200;
8     strncpy(res->message, "hello", 6);
9 }
10 void send_back(response_t *res) {
11     // implementation omitted
12 }
13 void process() {
14     struct Response res;
15     say_hello(&res);
16     send_back(&res);
17 }
```

A stack-based implementation of
(roughly) the same functionality

〈 End 〉