

CS202-002: Lecture 2

Processes 1

- Admin:

Quiz timing: Last 10 minutes of class!
First one on Tuesday

- PROCESSES

- WHAT

- WHY

- ABSTRACT MACHINE - EXEC ENVIRONMENT

- ↳ REGISTERS

- MEMORY LAYOUT

- EXECUTION

- STACK

- FUNC. CALL & RETURN

- SYSTEM CALLS

- OTHER WAYS TO TRANSFER CONTROL

Some of this
might be
familiar
from 201!

PROCESS

- Instance of a running program

BROWSER

VS CODE

Zoom

← PROCESSES

GOAL: RUN MULTIPLE PROCESSES ON ONE COMPUTER

WHY: Discussed this last time, but

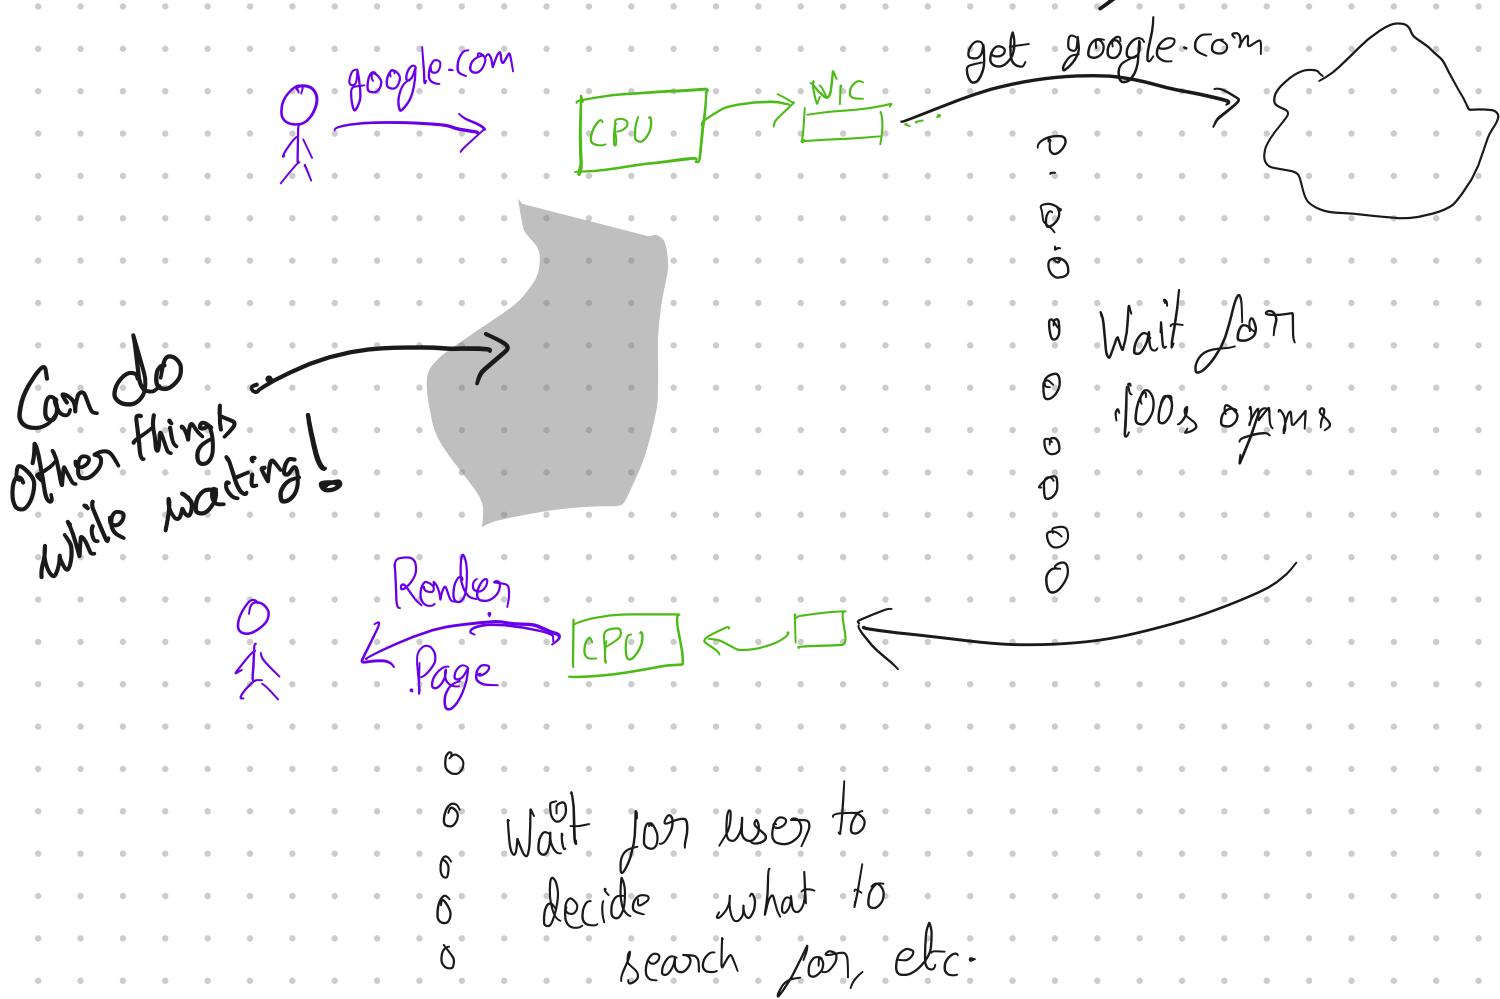
(a) Users want to

- Listen to music while working on HW

- Watch YouTube while compiling

- Monitor e-mail while working

⑤ Efficiency (people who manage computers want this)



What is a process provided



Book (Ch 4)
The OS



virtualizes

physical resources
to create an
abstract machine.

Two things we could look at

PROCESS

- ABSTRACT MACHINE
 - Memory
 - CPU
 - ...
 - Running exactly one program
- Interface to manipulate abstract machine & shared resources

KERNEL

- How to virtualize?

What does the abstract machine look like & how does it execute programs.

Today, Next class

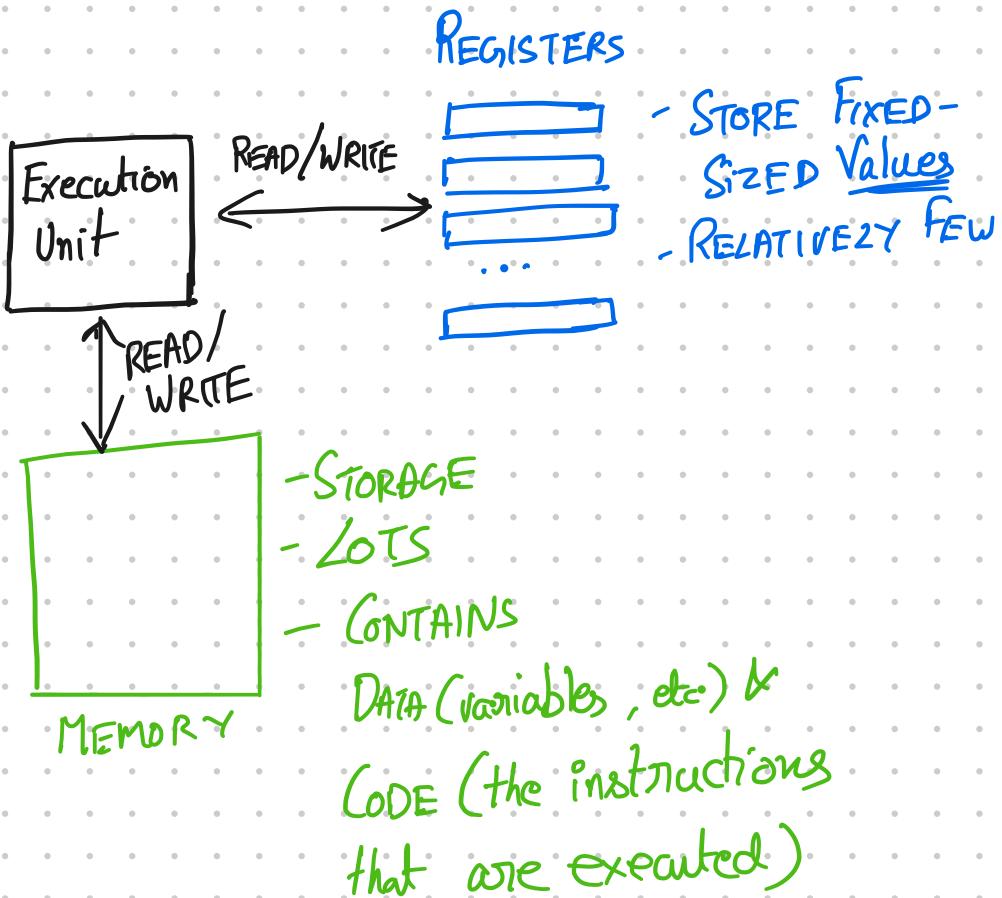
Later
in
semester 2.

CORE IDEA: EACH PROCESS APPEARS TO EXECUTE ON
ITS OWN ABSTRACT MACHINE.

That is, it is isolated from other processes.

What does the abstract machine have

- Can execute an instruction!
- No state!



IMPORTANT: All state is in registers or memory.

```
22
23     x = f(&arg);
24
25     printf("x: %lu\n", x);
26     printf("dereference q: %lu\n", *q);
27
28     return 0;
29 }
30
31 uint64_t f(uint64_t* ptr)
32 {
33     uint64_t x = 0;
34     x = g(*ptr);
35     return x + 1;
```

→ The location of the next instruction.

→ The next instruction

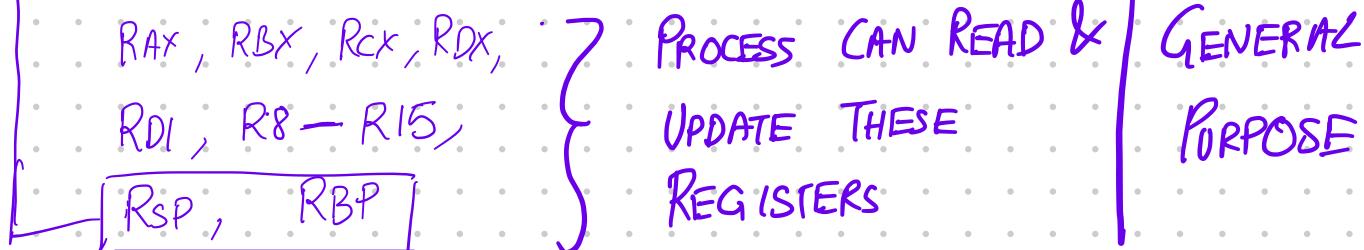
The value of X

• • •

Registers:

- This class x86-64 / AMD64.
- Registers in the processes abstract machine

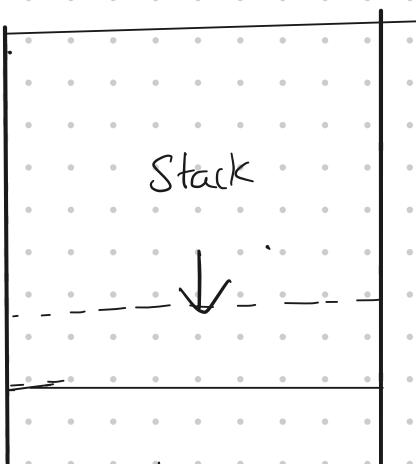
~~special significance~~



RIP

POINTS TO NEXT INSTRUCTION TO EXECUTE.
READ / MODIFIED BY SPECIAL INSTRUCTIONS CALL/RET/JMP/-

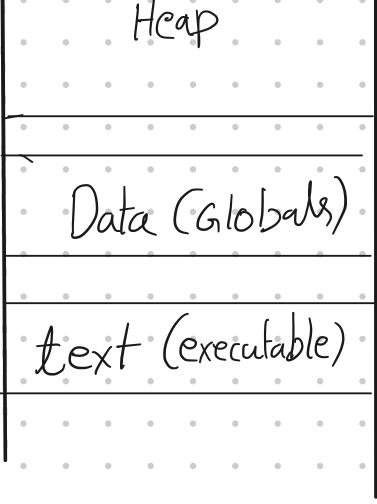
MEMORY



fffff..f ← High

- Layout (what goes where) determined by CONVENTION.

- THE ABSTRACT MACHINE



MAKES IT SO EACH
PROCESS ONLY SEES ITS
OWN MEMORY.

- PROGRAM EXECUTION

① Assembly crash course

Registers $\%rax, \%rbx, \dots \%r8$

Constants \$42 \$0x8
(Immediates)

Pointer deref $(\%r8)$ $(\%rbp)$

Pointer with $-8(\%rbp)$ $\star(\cdot.\%rbp - 8)$

$\star(\cdot.\%r8) =$

⑥ Break for stack



Orfff... - -

$\leftarrow \%rbp$: Base pointer
Beginning of running
function's stack frame

$\leftarrow \%rsp$: Stack pointer.
Last thing in the stack

- STACK FRAME
- For each function
- Contains
 - Where to return.
 - Local variables
 - ...

② Back to assembly

`movq %src, dst`: 'move' 64-bits from `src` to `dst`

`src` can be register, memory or immediate

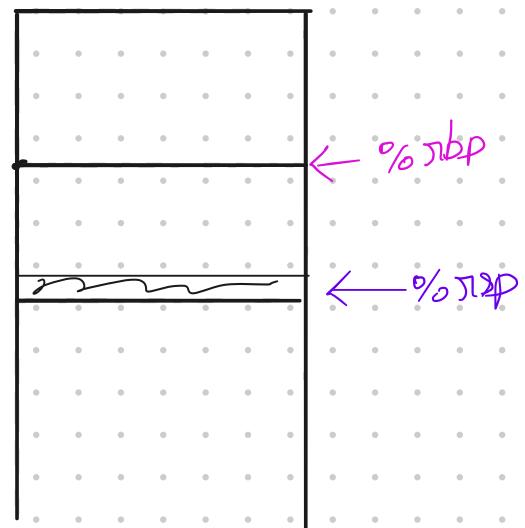
`dst` can be register or memory.

`addq/subq op1, op2` $op2 = op2 \pm op1$

`pushq %rax`: Push rax to stack

`subq $8, %rsp`

`movq %rax, (%rsp)`



`popq %rax`: Pop rax from stack
→ Do the opposite of above

`movq (%rsp), %rax`

`addq $8, %rsp`

③ Calling & returning

call f

Note: In reality $f \equiv f'$'s address.

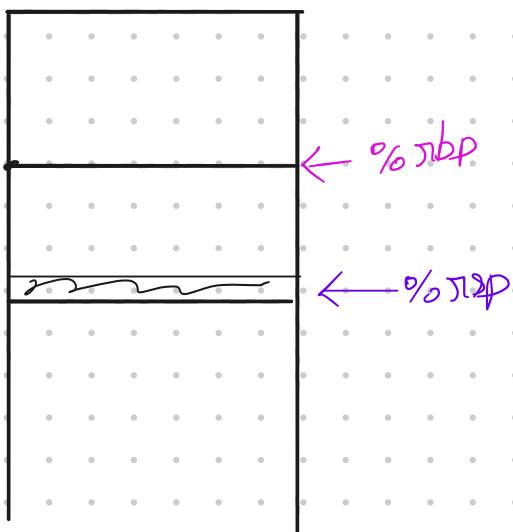
Transfer control to f

f();

=

pushq %rip
movq f, %rip

IMPORTANT: THIS IS
PSEUDOCODE! YOU
CANNOT RUN THESE
INSTRUCTIONS.



ret ← Return to the function
that called us

= popq %rip

OBSERVE: PROGRAM'S RESPONSIBILITY TO

@ ENSURE %RSP POINTS TO
THE RIGHT ADDRESS WHEN

RET IS CALLED

- CREATE STACK FRAME ON ENTRY INTO FUNC (PROLOG)
- DESTROY STACK FRAME BEFORE RET WHEN EXITING FUNC (EPILOG)

Calling convention

(b)

FIGURE OUT HOW TO PASS ARGUMENTS

(c)

MANAGE REGISTERS ACROSS FUNCTION CALLS.

LET US SEE AN EXAMPLE

```
15    movq $0, -8(%rbp)
16    movq $8, -16(%rbp)
17
18    leaq -16(%rbp), %rdi
19
20
21
22    call f
23
24    movq %rdi, -8(%rbp)
25
```

f PROLOG

```
28    f:
29        pushq %rbp
30        movq %rbp, %rsp
31
32        subq $32, %rsp
33        movq %rdi, -24(%rbp)
34
35        movq $0, -8(%rbp)
36
```

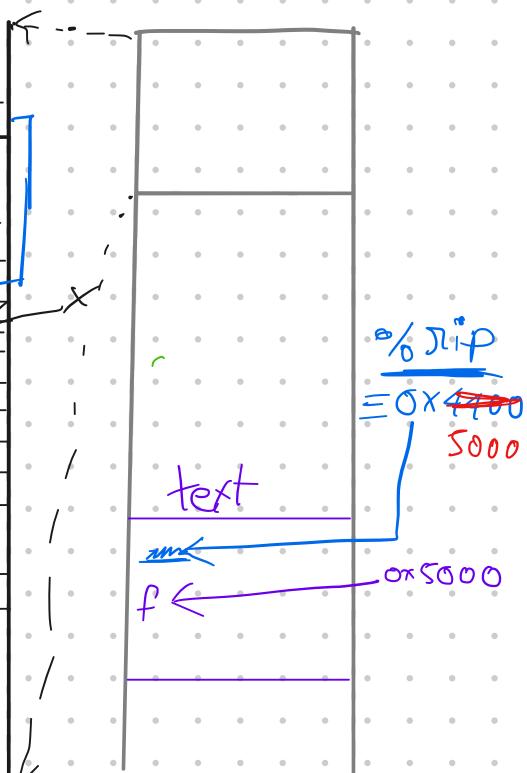
f EPILOG

```
47        movq %r10, %rax
48
49        movq %rbp, %rsp
50        popq %rbp
51        ret
```

Return Value

STACK

MEM



The actual working

```
37    movq -24(%rbp), %r8
38    movq (%r8), %r9
39    movq %r9, %rdi
40
41
```

42 call g ← g has its own
 43 prolog & epilog

f Prolog

```
28 f:
29   pushq %rbp
30   movq %rsp, %rbp
31
32   subq $32, %rsp
33   movq %rdi, -24(%rbp)
34
35   movq $0, -8(%rbp)
```

g Prolog

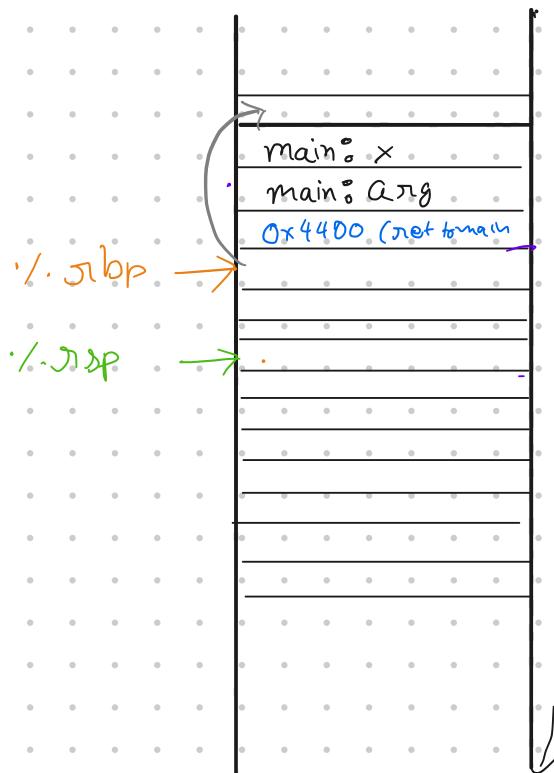
```
53 g:
54   pushq %rbp
55   movq %rsp, %rbp
56   subq $0x8, %rsp
```

g Epilog

```
59
60   movq %rbp, %rsp
61   popq %rbp
62
63   ret
```

f EPLOG

```
47   movq %r10, %rax
48
49   movq %rbp, %rsp
50   popq %rbp
51   ret
```



Calling Convention

- Arguments (in order)

%rdi, %rsi, %rdx, %rcx, %r8, %r9

- Return

%rax

- Call preserved (callee save)

$\%rbx, \%rbp, \%r12 - \%r15$

Function must restore value before
returning

- Call clobbered (caller save)

Everything else.

Sep 04, 2023 10:32

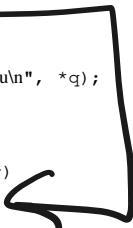
example.c

Page 1/1

```

1  /* CS202 -- handout 1
2   * compile and run this code with:
3   * $ gcc -g -Wall -o example example.c
4   * $ ./example
5   *
6   * examine its assembly with:
7   * $ gcc -O0 -S example.c
8   * $ [editor] example.s
9  */
10
11 #include <stdio.h>
12 #include <stdint.h>
13
14 uint64_t f(uint64_t* ptr);
15 uint64_t g(uint64_t a);
16 uint64_t* q;
17
18 int main(void)
19 {
20     uint64_t x = 0;
21     uint64_t arg = 8;
22
23     x = f(&arg); x = f(arg);
24     printf("x: %lu\n", x);
25     printf("dereference q: %lu\n", *q);
26
27     return 0;
28 }
29
30
31 uint64_t f(uint64_t* ptr)
32 {
33     uint64_t x = 0;
34     x = g(*ptr);
35     return x + 1;
36 }
37
38 uint64_t g(uint64_t a)
39 {
40     uint64_t x = 2*a;
41     q = &x; // <-- THIS IS AN ERROR (AKA BUG)
42     return x;
43 }

```



Sep 04, 2023 10:32

as.txt

Page 1/1

```

1  2. A look at the assembly...
2
3      To see the assembly code that the C compiler (gcc) produces:
4      $ gcc -O0 -S example.c
5      (then look at example.s)
6      NOTE: what we show below is not exactly what gcc produces. We have
7      simplified, omitted, and modified certain things.
8
9      main:
10         pushq  %rbp          # prologue: store caller's frame pointer
11         movq   %rsp, %rbp    # prologue: set frame pointer for new frame
12
13         subq   $16, %rsp     # prologue: make stack space
14
15         movq   $0, -8(%rbp)  # x = 0 (x lives at address rbp - 8)
16         movq   $8, -16(%rbp) # arg = 8 (arg lives at address rbp - 16)
17
18         leaq   -16(%rbp), %rdi # load the address of (rbp-16) into %rdi
19         # this implements "get ready to pass (&arg)
20         # to f"
21
22         call   f              # invoke f
23
24         movq   %rax, -8(%rbp) # x = (return value of f)
25
26         # eliding the rest of main()
27
28 f:
29         pushq  %rbp          # prologue: store caller's frame pointer
30         movq   %rsp, %rbp    # prologue: set frame pointer for new frame
31
32         subq   $32, %rsp     # prologue: make stack space
33         movq   %rdi, -24(%rbp) # Move ptr to the stack
34         # (ptr now lives at rbp - 24)
35         movq   $0, -8(%rbp)  # x = 0 (x's address is rbp - 8)
36
37         movq   -24(%rbp), %r8 # move 'ptr' to %r8
38         movq   (%r8), %r9    # dereference 'ptr' and save value to %r9
39         movq   %r9, %rdi     # Move the value of *ptr to rdi,
40         # so we can call g
41
42         call   g              # invoke g
43
44         movq   %rax, -8(%rbp) # x = (return value of g)
45         movq   -8(%rbp), %r10 # compute x + 1, part I
46         addq   $1, %r10        # compute x + 1, part II
47         movq   %r10, %rax     # Get ready to return x + 1
48
49         movq   %rbp, %rsp     # epilogue: undo stack frame
50         popq   %rbp          # epilogue: restore frame pointer from caller
51         ret                 # return
52
53 g:
54         pushq  %rbp          # prologue: store caller's frame pointer
55         movq   %rsp, %rbp    # prologue: set frame pointer for new frame
56         subq   $0x8, %rsp    # prologue: make stack space
57
58         ...
59
60         movq   %rbp, %rsp     # epilogue: undo stack frame
61         popq   %rbp          # epilogue: restore frame pointer from caller
62         ret                 # return

```