

CS202-002: OPERATING SYSTEMS

TWO DAYS AGO ... ENDED WITH HISTORY

TODAY

- PROCESSES

- PROCESS EXECUTION

↳ STACK

→ FUNCTION CALLS — STACK FRAMES

→ SYSTEM CALLS

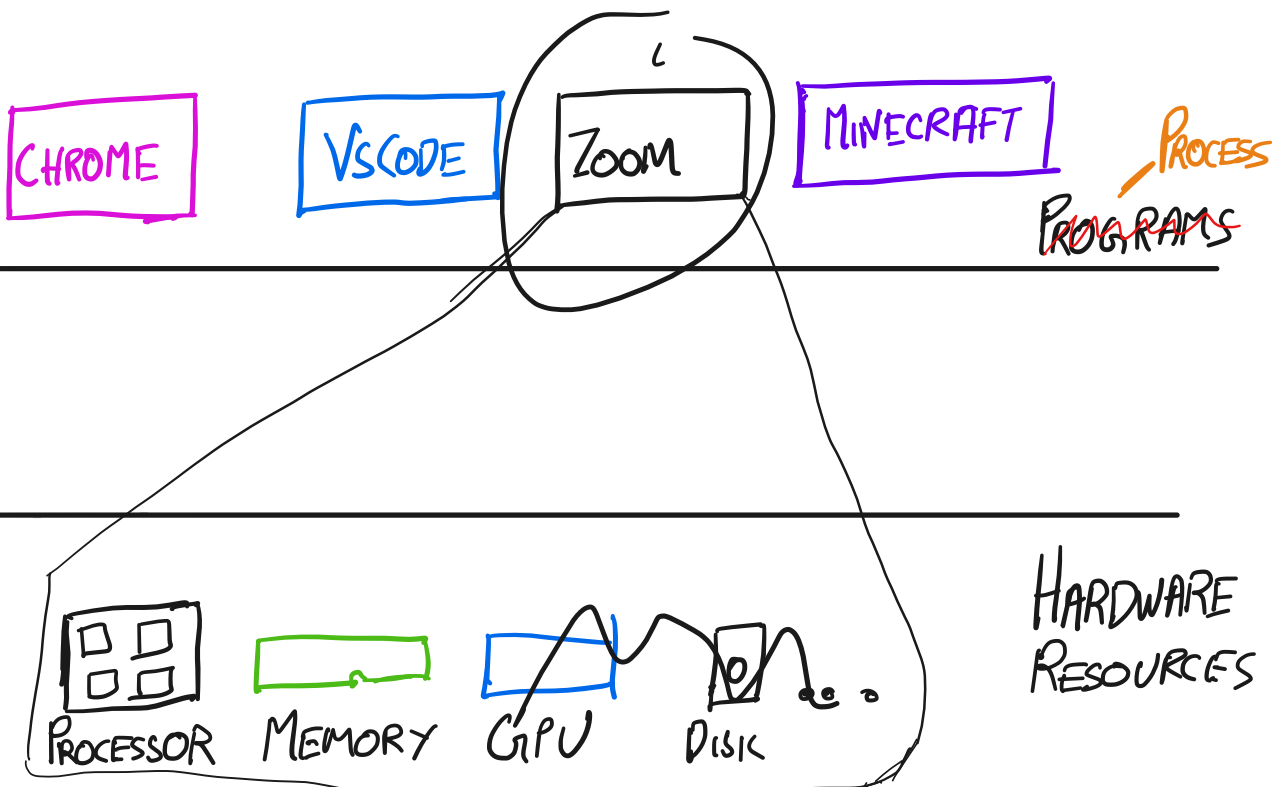
- OTHER TRAPS

} Maybe went through some of this in CS0.

- Review it again

- Agree on terminology

- Might add details



ONLY CONSIDERING PROCESSOR + MEMORY

To QUESTIONS

TWO QUESTIONS

- Q1. (a) WHAT DOES THE PROCESSOR LOOK LIKE TO A PROCESS?
- (b) WHAT DOES MEMORY LOOK LIKE TO A PROCESS?

→ TODAY!

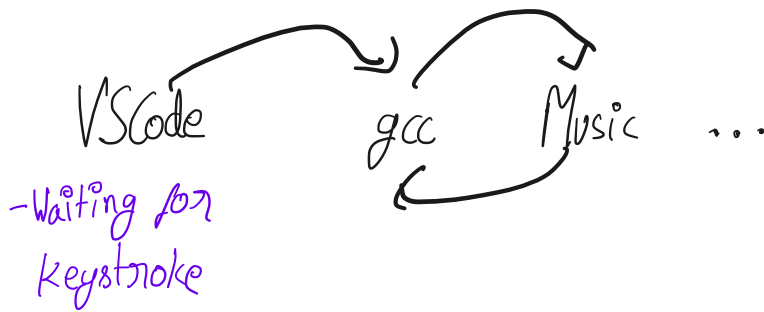
Q2. HOW DOES THE KERNEL IMPLEMENT & MANAGE PROCESSES?

→ LEAVE THIS FOR ANOTHER DAY

BUT WHY?

- CONVENIENCE

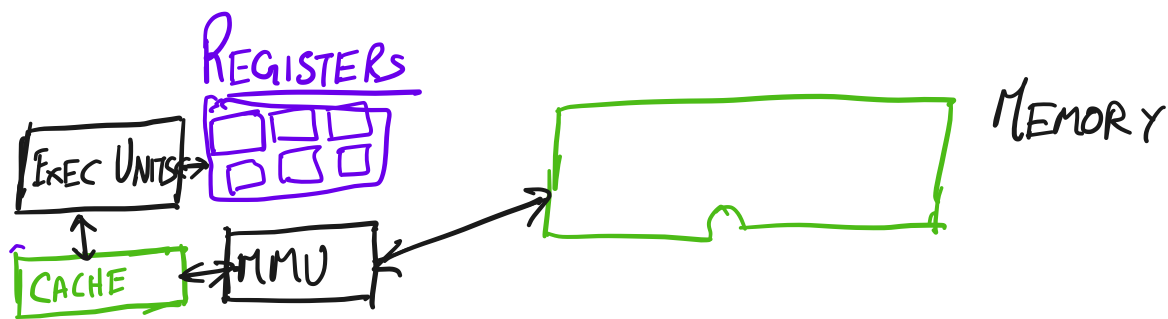
- EFFICIENCY



- Combining functionality

- Q1. (a) WHAT DOES THE PROCESSOR LOOK LIKE TO A PROCESS?
(b) WHAT DOES MEMORY LOOK LIKE TO A PROCESS?

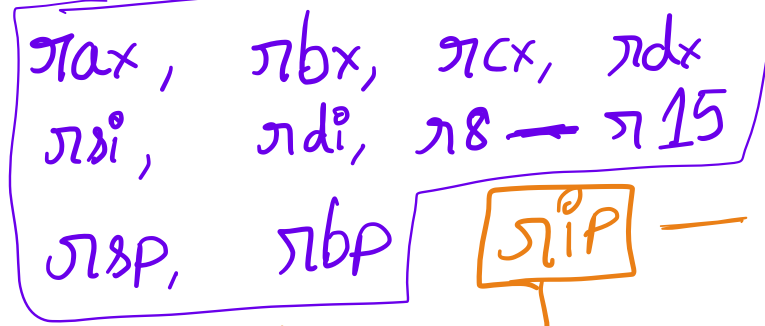
PROCESSOR ("CORE") & MEMORY: AN ABSTRACT VIEW



Exec units: Have no state.
Must read both **Instruction** & **data**
from elsewhere

Registers: A set (of known size, determined by processor designer) of fixed-size holders of state (in this class 64 bits).

x86-64



general purpose

special purpose

This one is special

CONSISTENT + FAST ACCESS

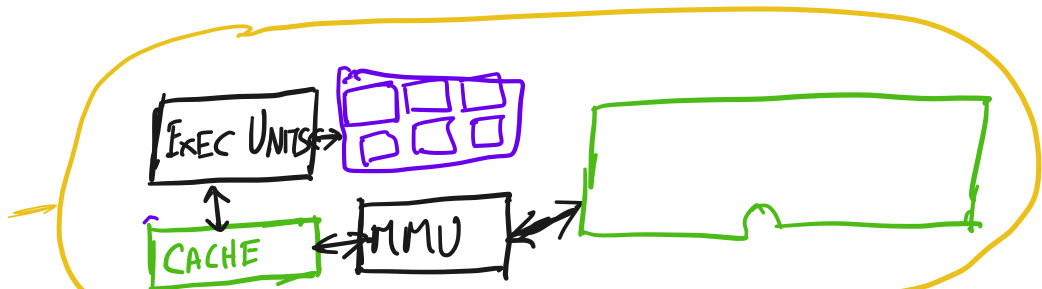
MEMORY: MORE PLentiful: AMOUNT DECIDED WHEN BUILDING A COMPUTER

BUT SLOWER ACCESS (0ns - 100ns) with variance.

```

...
-> int p = fork();
if (p != 0) {
-> printf("I'm yellow");
}

```



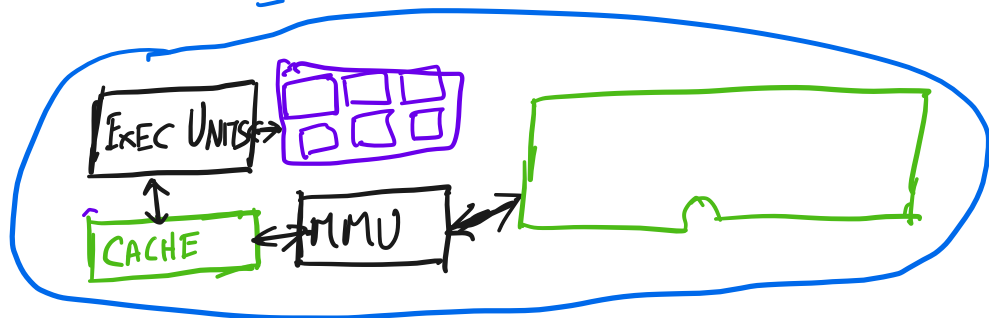
```

else {
  printf("I'm blue");
}

```

I am yellow

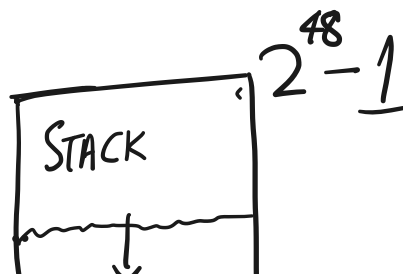
I'm blue.

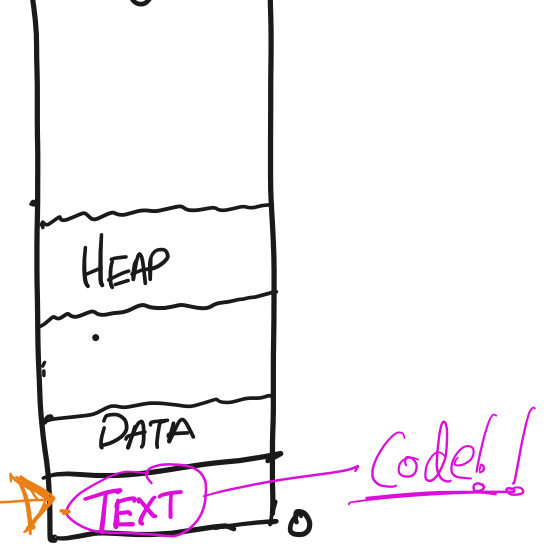
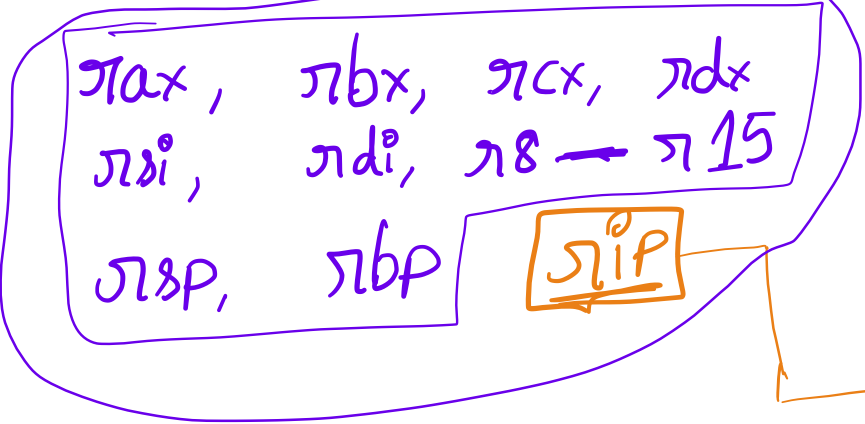


Each process gets

- Its own set of registers
- Its own view of memory
- Some other fiddly metadata.

PROCESS MEMORY



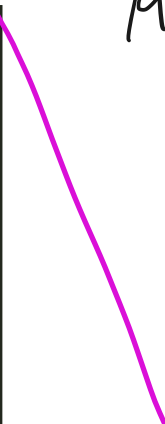


IMPORTANT: ALL INFORMATION USED BY THE PROCESSOR IS EITHER IN REGISTERS OR MEMORY

```

1 #include <stdio.h>
2 #include <stdint.h>
3
4 uint64_t f(uint64_t* ptr);
5 uint64_t* q;
6
7 int main(void)
8 {
9     uint64_t x = 0;
10    uint64_t arg = 8;
11
12    x = f(&arg);
13
14    printf("x: %lu\n", x);
15    printf("dereference q: %lu\n", *q);
16
17    return 0;

```



```

18 }
19
20 uint64_t f(uint64_t* ptr)
21 {
22     uint64_t x = 0;
23     x = g(*ptr);
24     return x + 1;
25 }
26
27 uint64_t g(uint64_t a)
28 {
29     uint64_t x = 2*a;
30     q = &x; // <-- THIS IS AN ERROR (AKA BUG)
31     return x;
32 }

```

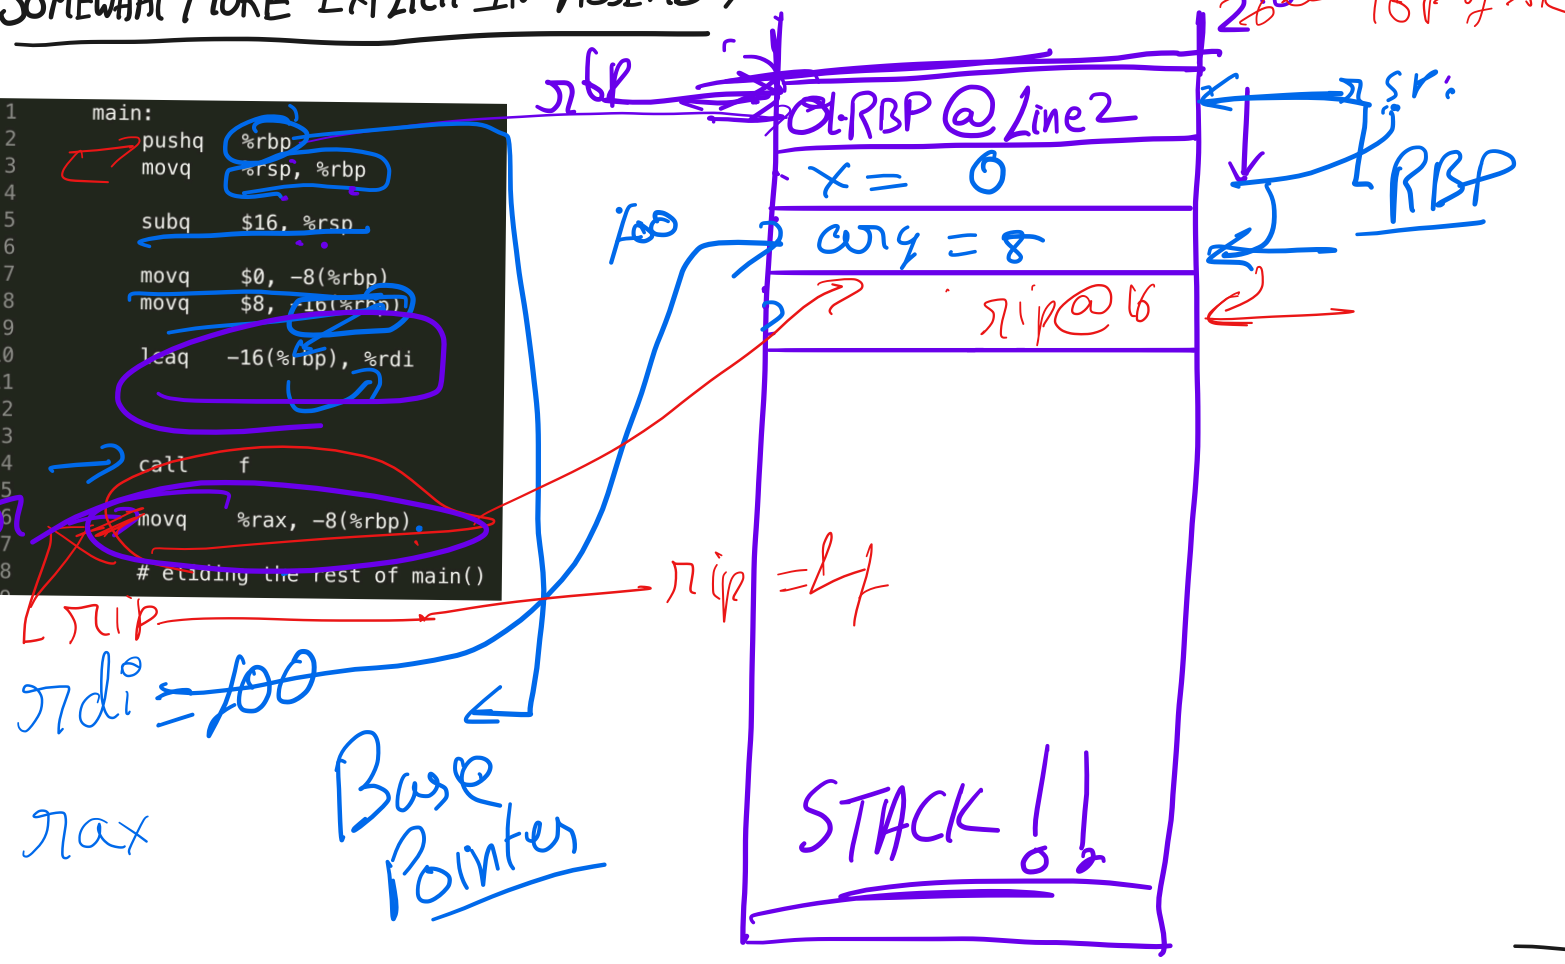


SOMEWHAT MORE EXPLICIT IN ASSEMBLY

```

1 main:
2     pushq %rbp
3     movq  %rsp, %rbp
4
5     subq  $16, %rsp
6
7     movq  $0, -8(%rbp)
8     movq  $8, -16(%rbp)
9
10    leaq  -16(%rbp), %rdi
11
12
13
14    call  f
15
16    movq  %rax, -8(%rbp)
17
18    # ending the rest of main()
19

```



Simple assembly

%rax ← Registers written with %
 \$22 ← Numbers (immediates) written with \$

mov(%rbp) ← pointer with metric

-8(%rbp)

%rbp - 8

movq source, destination

subq x, y => y = y - x

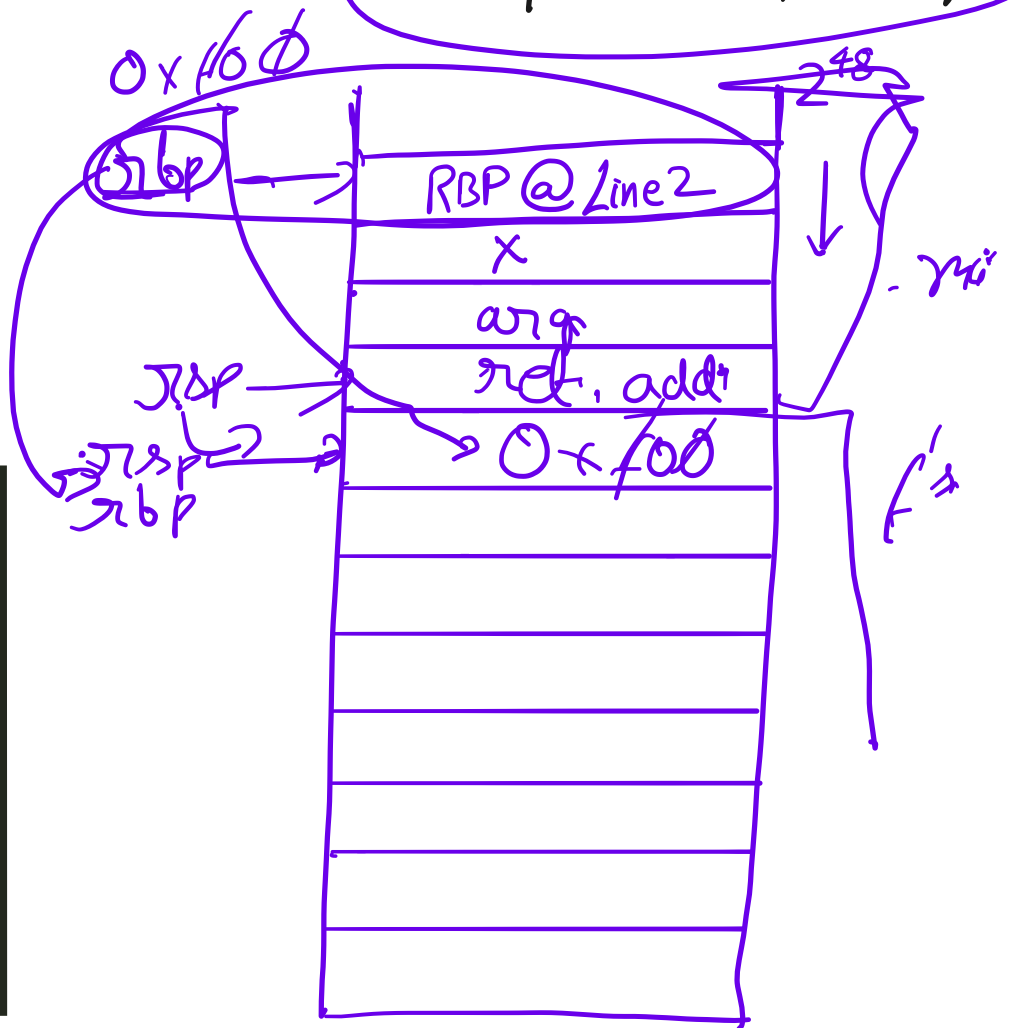
pushq %rax ≡ %rax = %rax - 8
 movq %rax, %rax

```

1 main:
2   pushq %rbp
3   movq  %rsp, %rbp
4
5   subq  $16, %rsp
6
7   movq  $0, -8(%rbp)
8   movq  $8, -16(%rbp)
9
10  leaq  -16(%rbp), %rdi
11
12
13
14  call  f
15
16  movq  %rax, -8(%rbp)
17
18  # eliding the rest of main()

```

rax
rdi
rbp
rsp



```

f:
pushq %rbp
movq  %rsp, %rbp

subq  $32, %rsp
movq  %rdi, -24(%rbp)

movq  $0, -8(%rbp)

movq  -24(%rbp), %r8
movq  (%r8), %r9
movq  %r9, %rdi

```

"PROLOG"


```
movq    %rax, -8(%rbp)
movq    -8(%rbp), %r10
addq    $1, %r10
movq    %r10, %rax
```

```
movq    %rbp, %rsp
popq    %rbp
ret
```

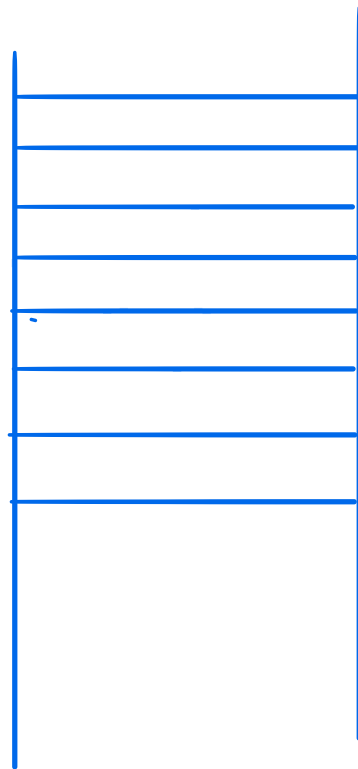
"EPILOG"

OK! Now let us understand the bug in 'g'

uint64_t * q;

```
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 }
```

...
x = g(*ptr);
...



Calling Convention

① How to pass arguments

② Where does the return value go

③ Who saves what registers?

Caller saved / volatile

%rax, %rcx, %rdx, %r8, %r9, %r10, %r11

Callee saved

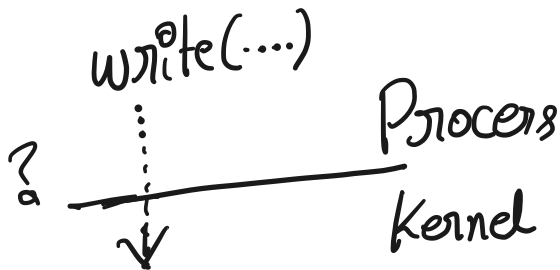
%rbx, %rbp, %rdi, %rsi, %r12-15

```
movq    %rax, -8(%rbp)
movq    -8(%rbp), %r10
addq    $1, %r10
movq    %r10, %rax

movq    %rpb, %rsp
popq   %rbp
ret
```

f epilog

Syscalls



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);
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;
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
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

```