

# Abstract and Bio

A SIMPLE RISC-V MULTITASKING OS FOR LEARNING

2021-12-03 @ Ohio Linux Fest, Columbus, Ohio

## Abstract:

Explore the line between hardware and software by writing code with absolute control over the cpu and peripherals. We'll explore how to do this using a completely free and open source simulator (Renode), toolchain (GCC), and instruction set (RISC-V). Using assembly, we'll initialize parts of the system such as CPU interrupts and privilege levels. Finally we'll review the assembly code for doing a context switch, the key software to which enables multitasking via timesharing.

## Bio:

Zak Kohler is a Chemical Engineer by training but a hacker at heart. He started programming in 3rd grade and has never let up. His first foray in open source was in early high school, and he discovered Linux and Free Software at university. Electronics is his second love and he fuses the two by playing with early computer hardware, modern microcontrollers and FPGAs. When zak isn't messing with computers he can be found growing plants, drawing, and exploring the world on foot.

# A SIMPLE RISC-V MULTITASKING OS FOR LEARNING

2021-12-03 - Ohio Linuxfest - Zak Kohler



# Zak Kohler

(y2kbugger)

- Chemical Engineering, University of Akron 2010
- Materials engineer turned software developer
- Relevant Interests:
  - programming, electronics, retrocomputing
- Irrelevant Interests:
  - running, punk rock, cheesemaking

# Outline

Motivation

Explain baremetal, risc-v, renode

Operating system background

KohlerOS, my operating system

- Startup
- Shell
- Multitasking, the context switch
- System calls



# Motivation

# Learning from a toy Operating System

This talk seeks to explore the line **between hardware and software** using a minimal simulated environment to write and play with a toy operating system.

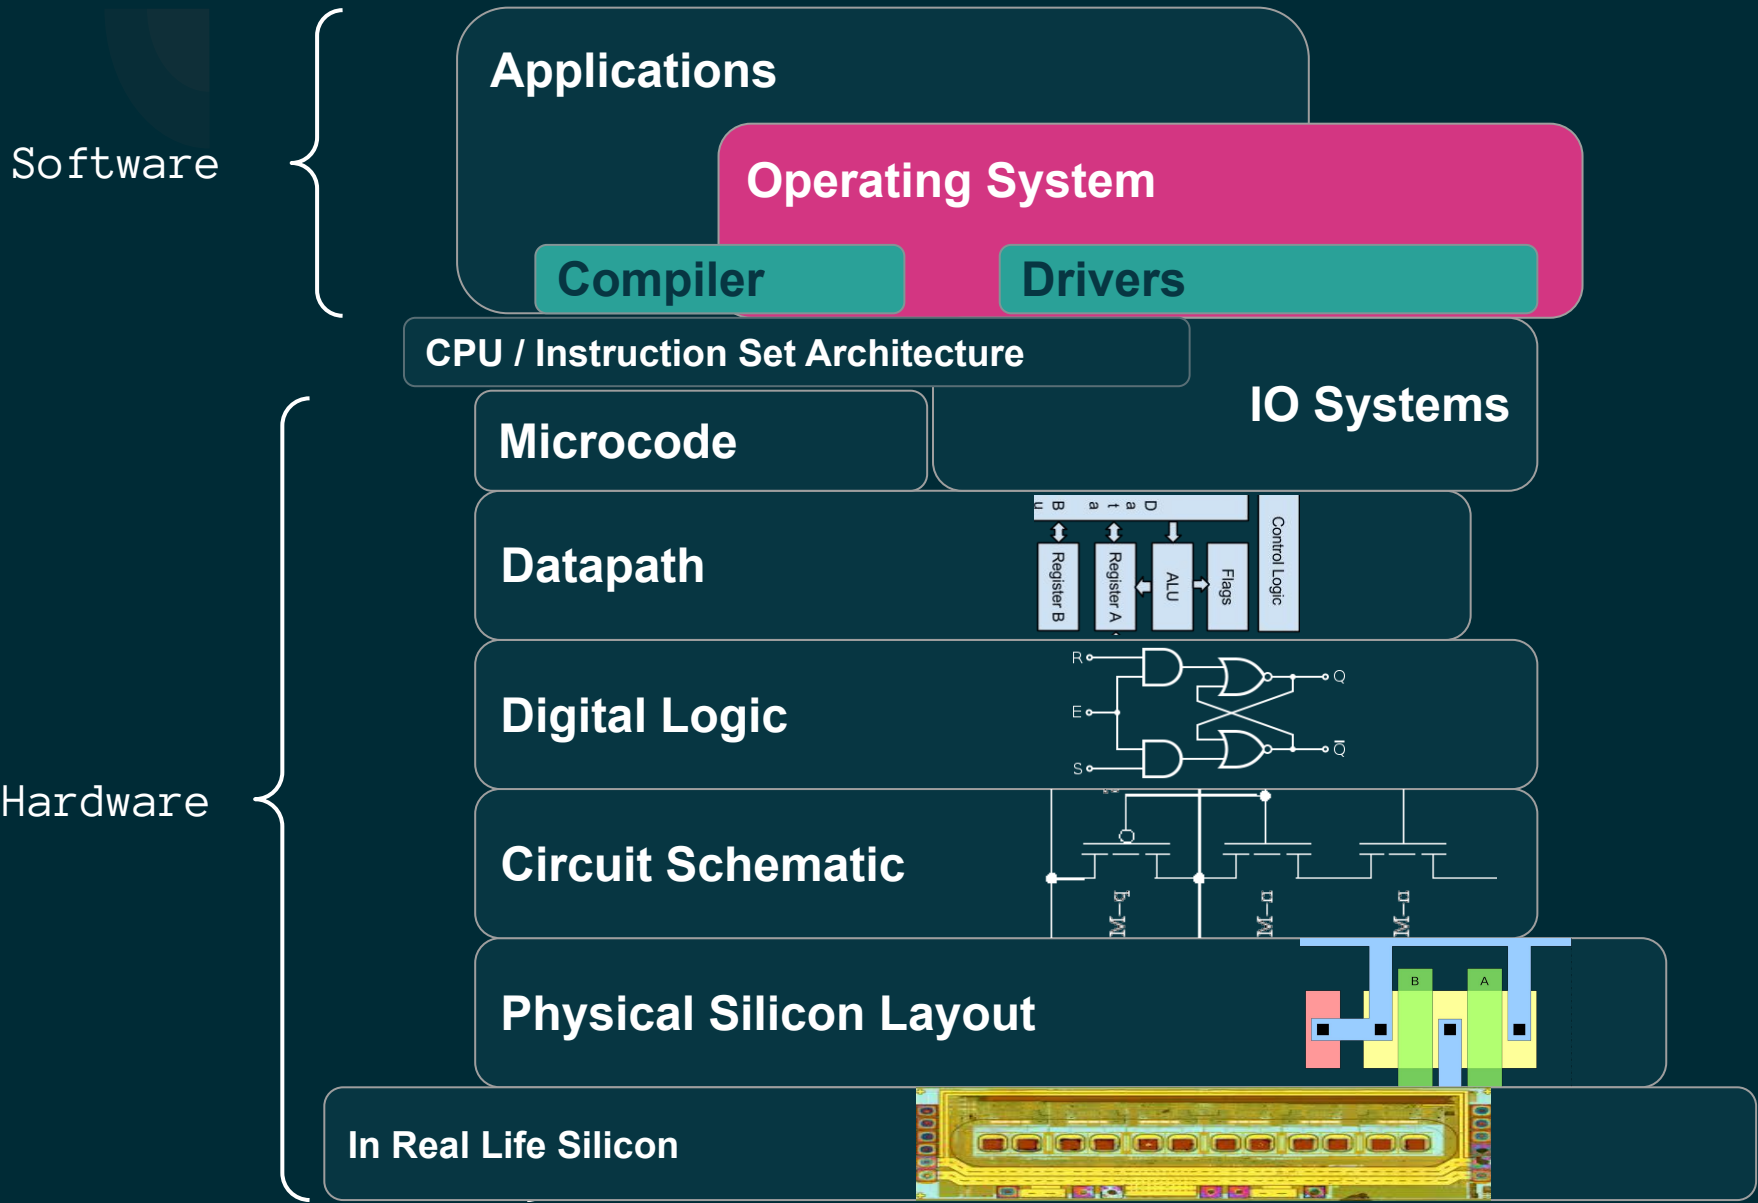
# Abstraction

Simple interfaces *allow* complex systems

You don't think about radios each time you send a text via sms, this is normally considered a good thing.




# Abstractions in Computers






In pursuit of killing the magic I've been playing with hardware and toy CPU emulators  
This talk is a(nother) step towards building a IRL CPU



```
In [61]: cpu = Computer()
cpu.reset()
cpu.switches.load_program(moveacross_bytes)
setup_bokeh_gui(cpu)

for i in range(400):
    #time.sleep(.01)
    cpu.step(instructionwise=True, debug=False)
```

Bitmap of Memory




# What is baremetal?

In the non-embedded world, when you compile and link a C program into an executable you are doing so with the intention of running it within a specific operating system.

In contrast when you compile baremetal or -freestanding you are telling the compiler that you intend to run this without relying on an operating system.

# What is baremetal?



This could be used, for example, to write an operating system. Alternatively it can be used to access the hardware of a system directly on an embedded system. Doing so sacrifices higher level luxuries such as memory management, standard IO, thread/process control, etc.

Because of this, sometimes it makes sense to run on a type of minimal OS optimized for embedded, e.g. a real time operating system (RTOS).

# What is RISC-V?

Wikipedia

RISC-V (pronounced "risk-five") is an open standard instruction set architecture (ISA) based on established reduced instruction set computer (RISC) principles. Unlike most other ISA designs, the RISC-V ISA is provided under open source licenses that do not require fees to use.

Akin to x86, arm, mips, alpha, 6502, powerpc

# What is Renode

Renode is a simulator designed for embedded firmware, including networks of devices.

Traditional alternatives such as QEMU aren't as optimized for the embedded space.

Video game console emulators target a specific platform.



# Operating Systems

# First operating system: Human operators

- Batch Processed
- Each program has full control of the entire mainframe.



# The role of operating systems

The job of an operating system is to abstract away hardware. They do this by implementing standard interface for tasks that talk to real world.

- Input/Output
- Timesharing
- Memory Management
- File systems
- User/Machine Mode



# Timesharing

Imagine if this guy had to manually set the program counter each time a human interrupted with some input...or change out one job for another thousands of times per second.



# My OS

The job of an operating system is to abstract away hardware. They do this by implementing standard interface for tasks that talk to real world.

- ~~Input/Output~~
- Timesharing
- ~~Memory Management~~
- ~~File systems~~
- ~~User/Machine Mode~~

# Monitor/Shell

The shell's special job:

```
kos> ?
```

Programs:

```
h: hello world
e: hello ecall world
l: laugh
f: laugh forever
c: count forever
```

Shell commands:

```
?: show this help
@: list stopped processes
^: restart all stopped processes in background
!: stop all background processes.
C-Z: stop foreground process
C-C: kill foreground process
```

launch and control other processes.

If the shell launches programs, how does the shell get launched?

# Early Startup

## Setup interrupts

ASM baremetal.s

```
start:
.....# setup machine trap vector
1:.....auipc  t0, %pcrel_hi(mtvec_interrupt_handler) # load mtvec_interrupt_handler(hi)
.....addi  t0, t0, %pcrel_lo(1b).....# load mtvec_interrupt_handler(lo)
.....csrw  mtvec, t0

.....# set mstatus.MIE=1 (enable M mode interrupts in general)
.....li    t0, 0b00000000000001000
.....csrrs zero, mstatus, t0
```

# Early Startup

Initialize stack and process pointers

```
26 .....# setup a stack pointer
27 .....la sp, memtop
28
29 .....# no process is running by default
30 .....# squat on tp to hold which process is running
31 .....# linux kinda does this soo..
32 .....li tp, 0
33
```

# Early Startup

Jump to C to initial rest of kernel

```
39  
40     ..... call init_kernel  
41  
42     forever:  
43     ..... j forever
```

```
void init_kernel()  
{  
    ..... init_uart();  
    ..... register_all_programs();  
    ..... current_process = init_process(lookup_program('s'));  
    ..... init_timer();  
    ..... asm volatile("ecall");  
}
```

# Monitor/Shell

kos> ?

Programs:

h: hello world

e: hello ecall world

l: laugh

f: laugh forever

c: count forever

Shell commands:

?: show this help

@: list stopped processes

^: restart all stopped processes in background

!: stop all background processes.

C-Z: stop foreground process

C-C: kill foreground process



# Programs and Processes

Programs are simply C functions registered with a name.

A running program is a Process.

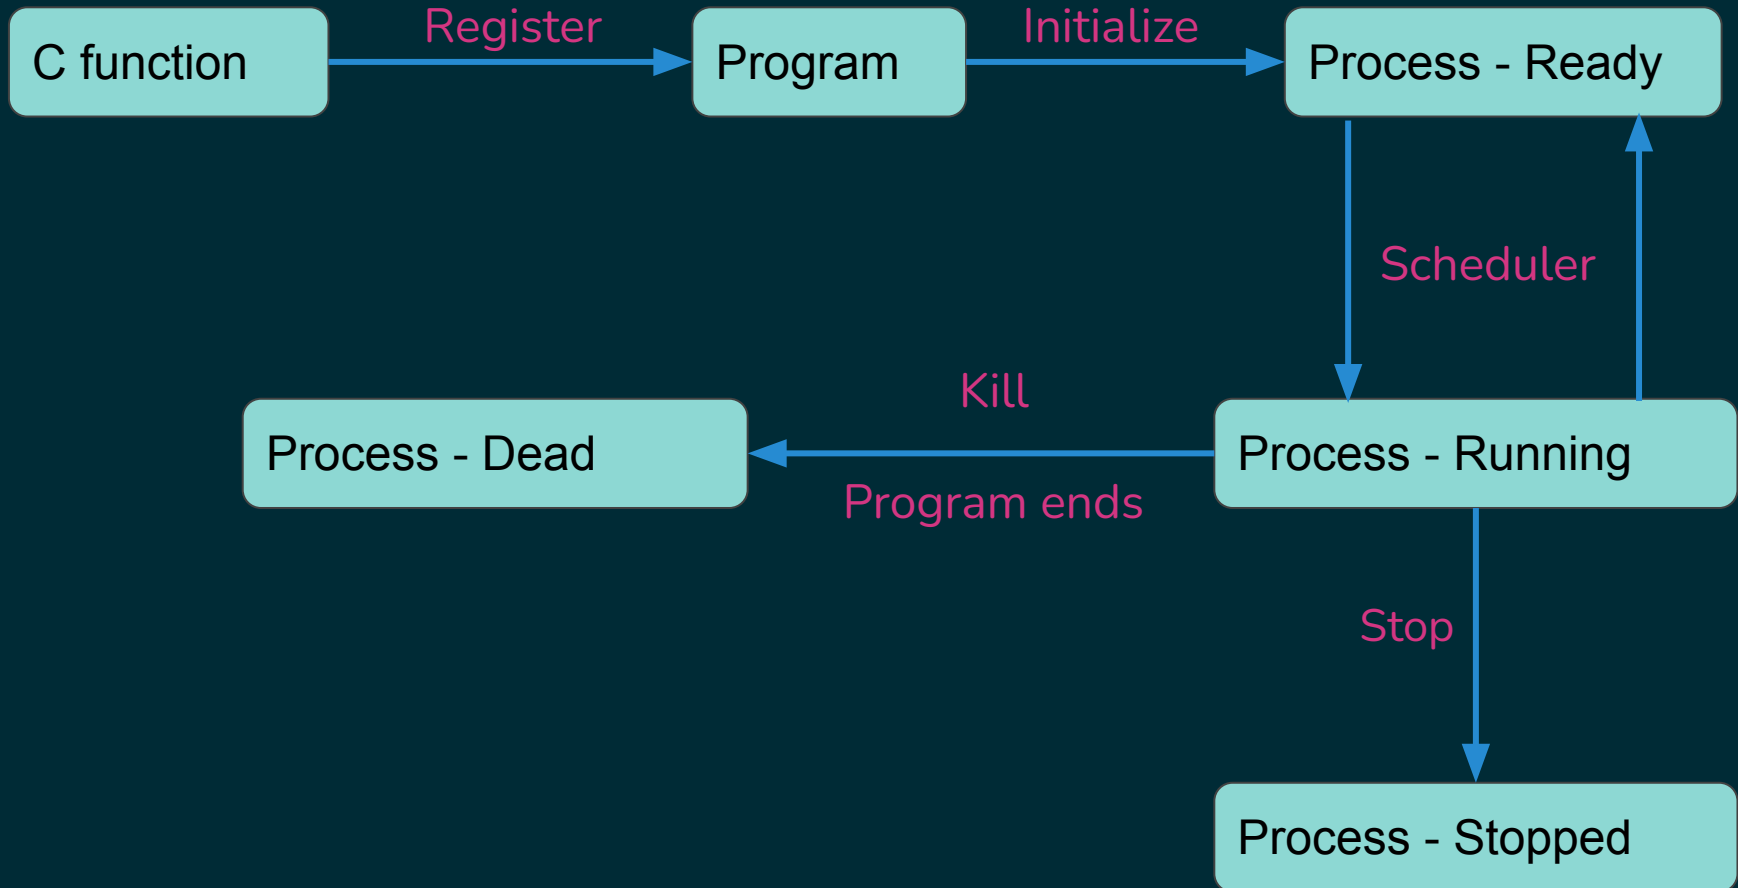
Scheduling is just a naive round-robin for all "Ready" Processes.

```
struct Program
{
    void (*function)();
    unsigned char name;
};

struct Process
{
    size_t *sp;
    struct Program *program;
    enum ProcessStatus status;
    size_t stack[PROC_STACK_SIZE];
};
```

```
enum ProcessStatus
{
    Uninitialized,
    Ready,
    Running,
    Dead,
    Stopping,
    Stopped,
};
```

# Lifecycle of a kOS process



Launch a simple program using the shell.

Processes are written as C functions, but this is out of convenience. We could add an arbitrary binary loader to the shell.

DEMO

Ok, that's not very  
interesting

It just looks as if the shell is dispatching to a C function based on user input.


Next I will give a demo to prove that this is not the case.

DEMO



# Context switch

enable preemptive  
multitasking



# The context switch

A software construct that

- Cannot be written in C
- Is CPU architecture dependent
- Is triggered by a hardware timer interrupt

The basic stages are:

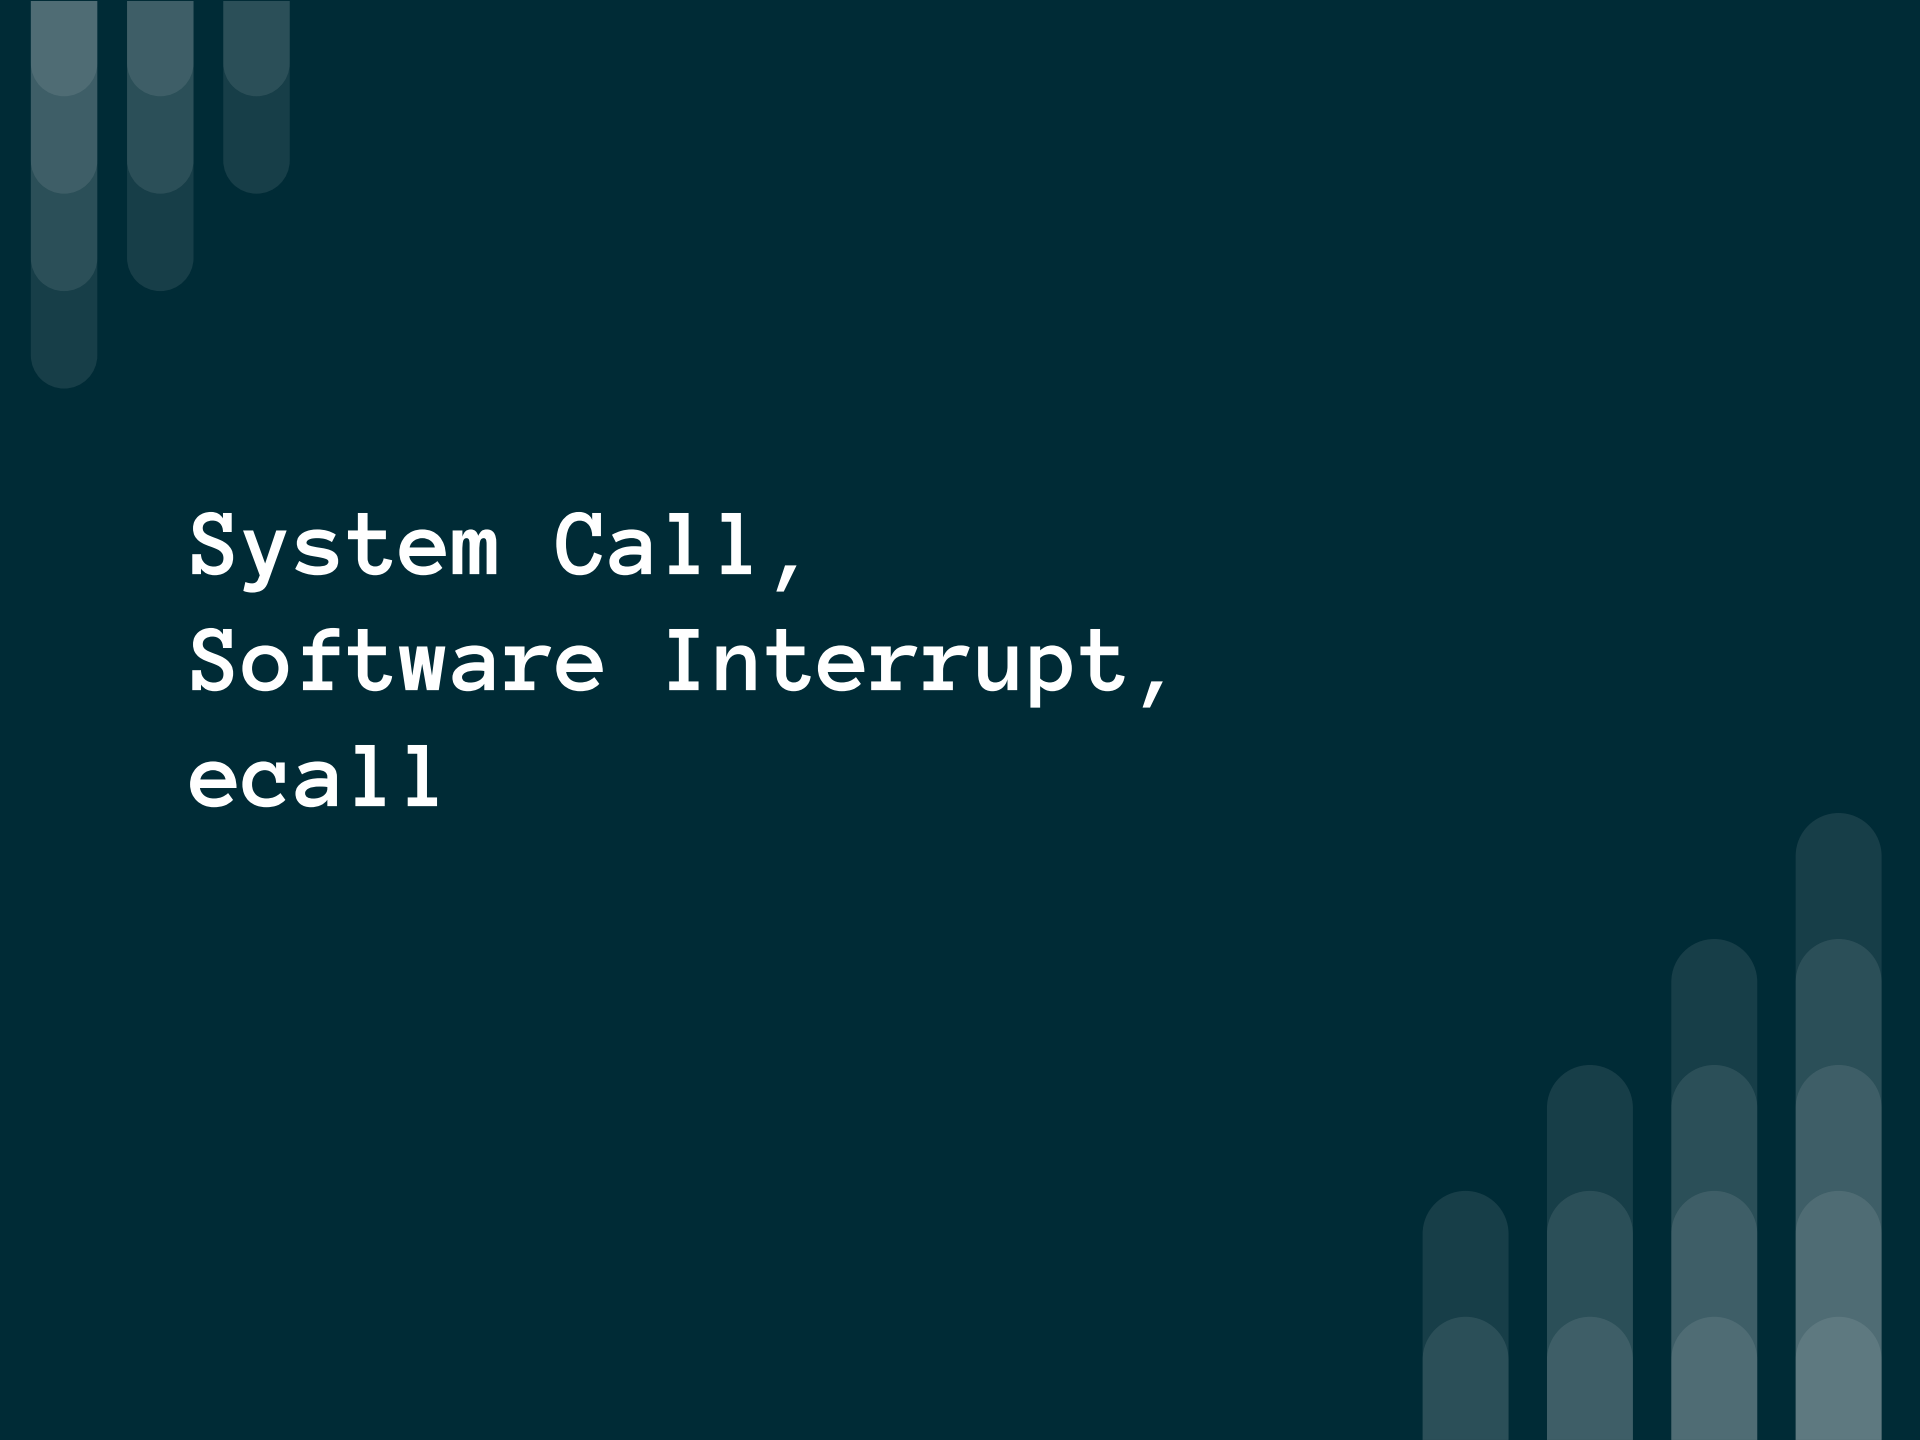
- Backs up state of CPU registers
- Restores a different state
- and then continues execution.

[Show Code](#)

# Programs and Processes

```
struct Program
{
    void (*function)();
    unsigned char name;
};

struct Process
{
    size_t *sp;
    struct Program *program;
    enum ProcessStatus status;
    size_t stack[PROC_STACK_SIZE];
};
```

The image features a dark teal background with decorative elements. In the top-left corner, there are three vertical bars of varying heights, each composed of several overlapping rounded rectangular segments. In the bottom-right corner, there are four vertical bars of increasing height from left to right, also composed of overlapping rounded rectangular segments. The text is centered in the middle-left area of the slide.

System Call,  
Software Interrupt,  
ecall



# Software Interrupts

- In an operating system using privilege modes, processes are prevented from accessing the hardware directly.
- In order to talk to hardware a process must talk to the kernel.

Q: How can a process talk to the kernel?

A: Software interrupts

Show Code

# How a real standard library uses system calls (musl)

```
master | musl-libc / src / stdio / fopen.c
```

Rich Felker support kernels with no SYS\_open syscall, only SYS\_openat ...

0 contributors

31 lines (24 sloc) | 535 Bytes

```
1 #include "stdio_impl.h"
2 #include <fcntl.h>
3 #include <string.h>
4 #include <errno.h>
5
6 FILE *fopen(const char *restrict filename, const char *restrict mode)
7 {
8     FILE *f;
9     int fd;
10    int flags;
11
12    /* Check for valid initial mode character */
13    if (!strchr("rwa", *mode)) {
14        errno = EINVAL;
15        return 0;
16    }
17
18    /* Compute the flags to pass to open() */
19    flags = __fmodeflags(mode);
20
21    fd = sys_open_cp(filename, flags, 0666);
22    if (fd < 0) return 0; musl-libc / src / internal / syscall.h
23
24    f = __fdopen(fd, mode);
25    if (f) return f;
26
27    __syscall(SYS_close, fd);
28    return 0;
29 }
```

```
master | musl-libc / src / internal / mips / syscall.s
```

Rich Felker add 7-arg syscall support for mips ...

0 contributors

25 lines (24 sloc) | 442 Bytes

```
1 .set    noreorder
2
3 .global __syscall
4 .type   __syscall,@function
5 __syscall:
6     move    $2, $4
7     move    $4, $5
8     move    $5, $6
9     move    $6, $7
10
11    lw      $7, 16($sp)
12    lw      $8, 20($sp)
13    lw      $9, 24($sp)
14    lw      $10, 28($sp)
15    subu    $sp, $sp, 32
16    sw      $8, 16($sp)
17    sw      $9, 20($sp)
18    sw      $2, 28($sp)
19    lw      $2, 28($sp)
20    syscall
21    beq     $7, $0, 1f
22    addu    $sp, $sp, 32
23    subu    $2, $0, $2
24 1:      jr      $ra
25        nop
```

# Wrap Up

Motivation

Explain baremetal, risc-v, renode

Operating system background

My operating system

- Startup
- Shell
- Multitasking, the context switch
- System calls



# Further Learning



# Topics Not Covered

- Memory management
- Race conditions
- CPU modes/rings
- File Systems
- Drivers
  - UART
  - Timer
- RTOS
- Deploying to hardware

# Want to try this out for yourself?

Last year I covered the basics of using Renode to run baremetal risc-v code.

<https://blog.y2kbugger.com/baremetal-riscv-renode-1.html>

# Thank You!

[@y2kbugger](#) on twitter

<https://blog.y2kbugger.com>

## Questions??

<https://blog.y2kbugger.com/baremetal-riscv-renode-1.html>