


Writing an operating system in 2.5 years

Yunhao Zhang

But first, writing an OS in 0.5 years

- P0: understand C and user-level instructions
- P1: understand context-switch and multi-threading
- P2: understand exception and privilege levels
- P3: understand file system abstractions
- P4: understand file system implementation
- P5: understand bus and I/O devices

Why **2.5 years**? An overview



In **June 2020**, we created egos



~20K lines of code


run on **Intel/Arm** CPU

run as a **Linux / MacOS** process




20K lines of code

Students read a **very small** portion



2K lines of code

Students read a **large** portion



Intel x86 (1987)



CPU document has
several thousands of pages



RISC-V (2010)

CPU document
has **<100** of pages

The motivation



~20K → ~2K

x86 → RISC-V


Linux / MacOS → hardware

Lesson

Good motivations should convince
non-experts why the work is **important**

non-experts like students, my friend doing ML theory, etc.

Yet, ideal \neq possible



Obstacles

Summer 2020

Fall 2020



Obstacles & Hope



The hardware had only **24KB memory**



There was **no disk**

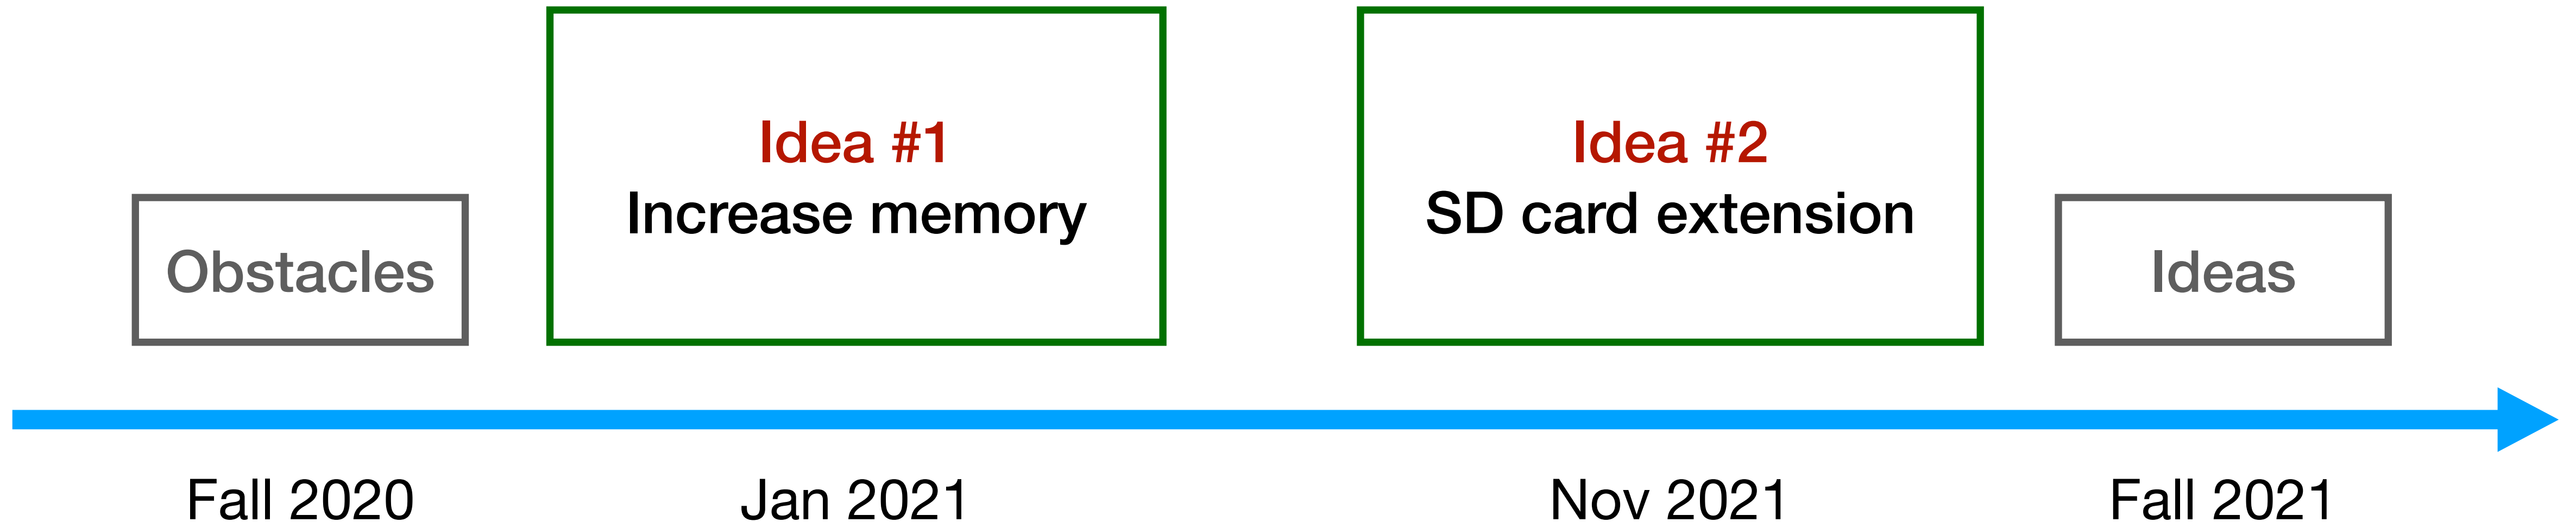


The hardware supports **timer interrupt**



The hardware supports **privilege levels** and **exceptions**

Overcome obstacles with ideas!



Open-source hardware

This repository has been archived by the owner before Nov 8, 2022. It is now read-only.

sifive / freedom Public archive Watch 187 Fork 269 Star 1k

[Code](#) [Issues 64](#) [Pull requests 6](#) [Actions](#) [Projects](#) [Security](#) [Insights](#)

master Go to file Code **About**

Source files for SiFive's Freedom platforms

- Readme
- Apache-2.0 license
- 1k stars
- 187 watching
- 269 forks

Releases 1

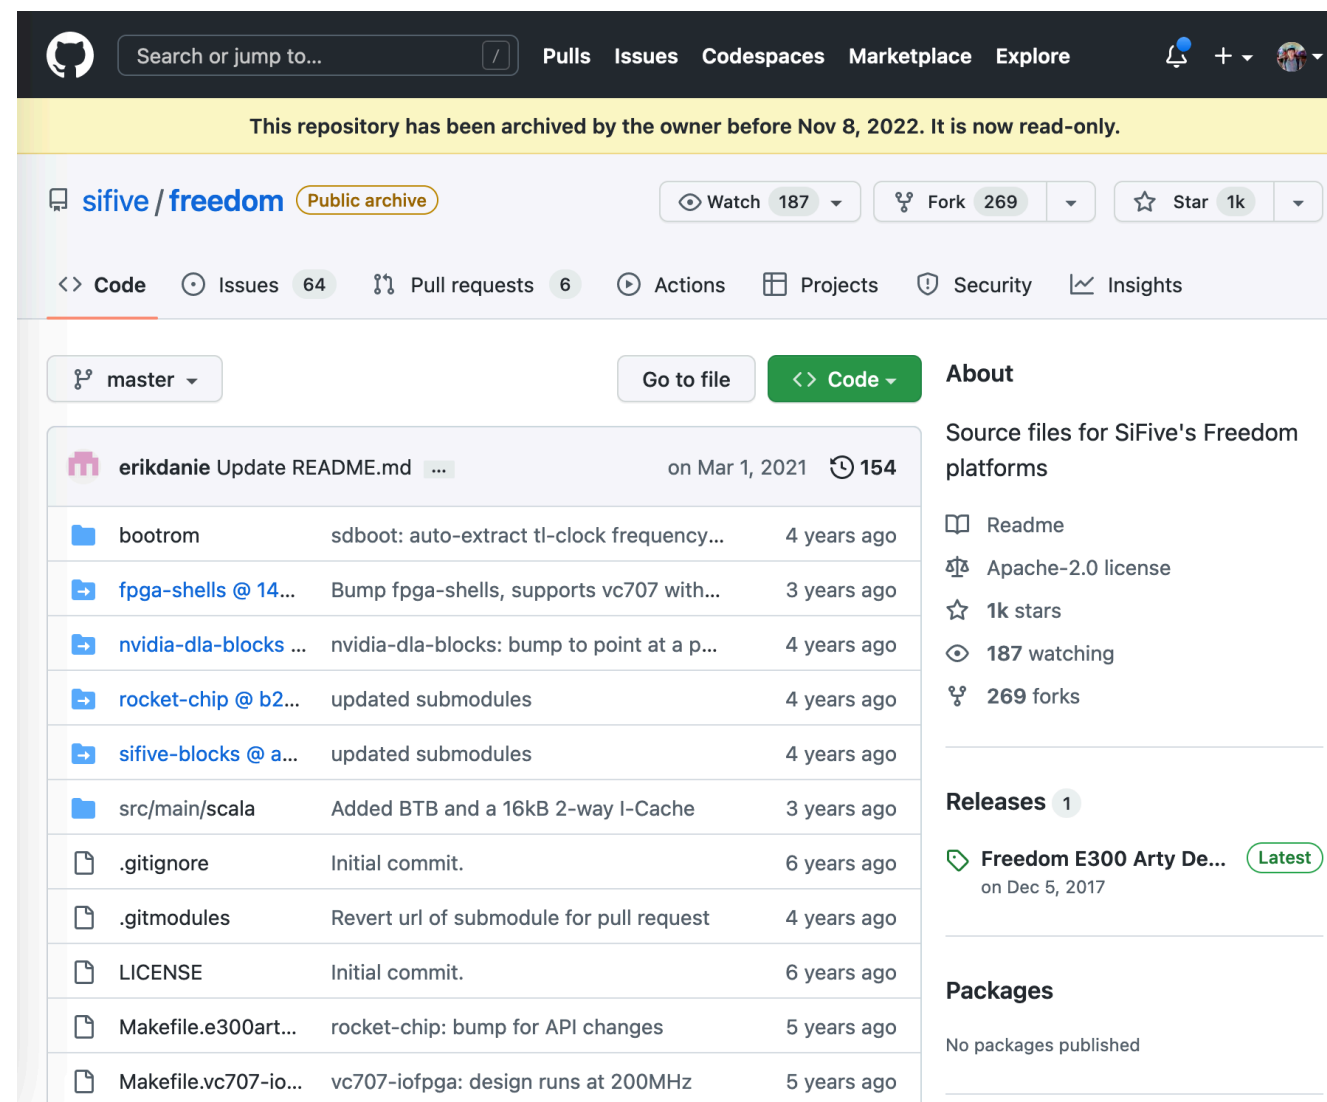
- Freedom E300 Arty De...** Latest
on Dec 5, 2017

Packages


No packages published

Author	Commit Message	Date
erikdanie	Update README.md	on Mar 1, 2021 154
	bootrom sdboot: auto-extract tl-clock frequency...	4 years ago
	fpga-shells @ 14... Bump fpga-shells, supports vc707 with...	3 years ago
	nvidia-dla-blocks ... nvidia-dla-blocks: bump to point at a p...	4 years ago
	rocket-chip @ b2... updated submodules	4 years ago
	sifive-blocks @ a... updated submodules	4 years ago
	src/main/scala Added BTB and a 16kB 2-way I-Cache	3 years ago
	.gitignore Initial commit.	6 years ago
	.gitmodules Revert url of submodule for pull request	4 years ago
	LICENSE Initial commit.	6 years ago
	Makefile.e300art... rocket-chip: bump for API changes	5 years ago
	Makefile.vc707-io... vc707-iofpga: design runs at 200MHz	5 years ago

Running open-source hardware




A **binary file**
encoding the hardware design
(**clocks, registers, circuits, etc.**)



FPGA: emulate the **hardware design**

Idea #1: Increase memory




```
86     dcache = Some(DCacheParams(  
87         rowBits = site(SystemBusKey).beatBits,  
88         → nSets = 256, // 16Kb scratchpad  
89         nWays = 1,  
90         nTLBEntries = 4,  
91         nMSHRs = 0,  
92         blockBytes = site(CacheBlockBytes),  
93         scratch = Some(0x80000000L))),
```

<https://github.com/chipsalliance/rocket-chip/blob/b21c7879b3ea22f69cb8457109561f37c225f8ea/src/main/scala/subsystem/Configs.scala#L78>

Background of the SPI bus

Instance	Flash Controller	Address	cs_width	div_width
QSPI 0	Y	0x10014000	1	12
SPI 1	N	0x10024000	4	12
SPI 2	N	0x10034000	1	12

Table 64: SPI Instances



Chapter 19 of Sifive FE310 manual, v19p04


<https://github.com/yhzhang0128/egos-2000/blob/main/references/sifive-fe310-v19p04.pdf>

Remap SPI1 to microSD card

Instance	Flash Controller	Address	cs_width	div_width
QSPI 0	Y	0x10014000	1	12
SPI 1	N	0x10024000	4	12
SPI 2	N	0x10034000	1	12


Table 64: SPI Instances

new SPI1




old SPI1

Remap SPI1 to microSD card




Old SPI1 mapping

6 pins
GND + VCC + SPI (4)



- Pin 1
- Pin 2
- Pin 3
- Pin 4
- Pin 5
- Pin 6

- ~CS
- MOSI
- MISO
- SCK
- GND
- VCC




New SPI1 mapping

Idea #2: Using microSD card as disk

ChipKit SPI

```
set_property -dict { PACKAGE_PIN G1 IOSTANDARD LVCMOS33 } [get_ports { ck_miso }];  
set_property -dict { PACKAGE_PIN H1 IOSTANDARD LVCMOS33 } [get_ports { ck_mosi }];  
set_property -dict { PACKAGE_PIN F1 IOSTANDARD LVCMOS33 } [get_ports { ck_sck }]; #I  
set_property -dict { PACKAGE_PIN C1 IOSTANDARD LVCMOS33 } [get_ports { ck_ss }]; #I
```



Find these 4 wires in the
repo and **replace** them

##Pmod Header JA

```
set_property -dict { PACKAGE_PIN G13 IOSTANDARD LVCMOS33 } [get_ports { ja_0 }]; #IO_0  
set_property -dict { PACKAGE_PIN B11 IOSTANDARD LVCMOS33 } [get_ports { ja_1 }]; #IO_1  
set_property -dict { PACKAGE_PIN A11 IOSTANDARD LVCMOS33 } [get_ports { ja_2 }]; #IO_2  
set_property -dict { PACKAGE_PIN D12 IOSTANDARD LVCMOS33 } [get_ports { ja_3 }]; #IO_L1
```



The road to ideas is **difficult**

No concrete **progress** for **>1 year**

Not sure whether this can succeed at all

Only person working on this project



Obstacles

Ideas * 2

Fall 2020

Fall 2021

Take-away:

Ideas are **difficult** to come up with
and there is **no guarantee of success**

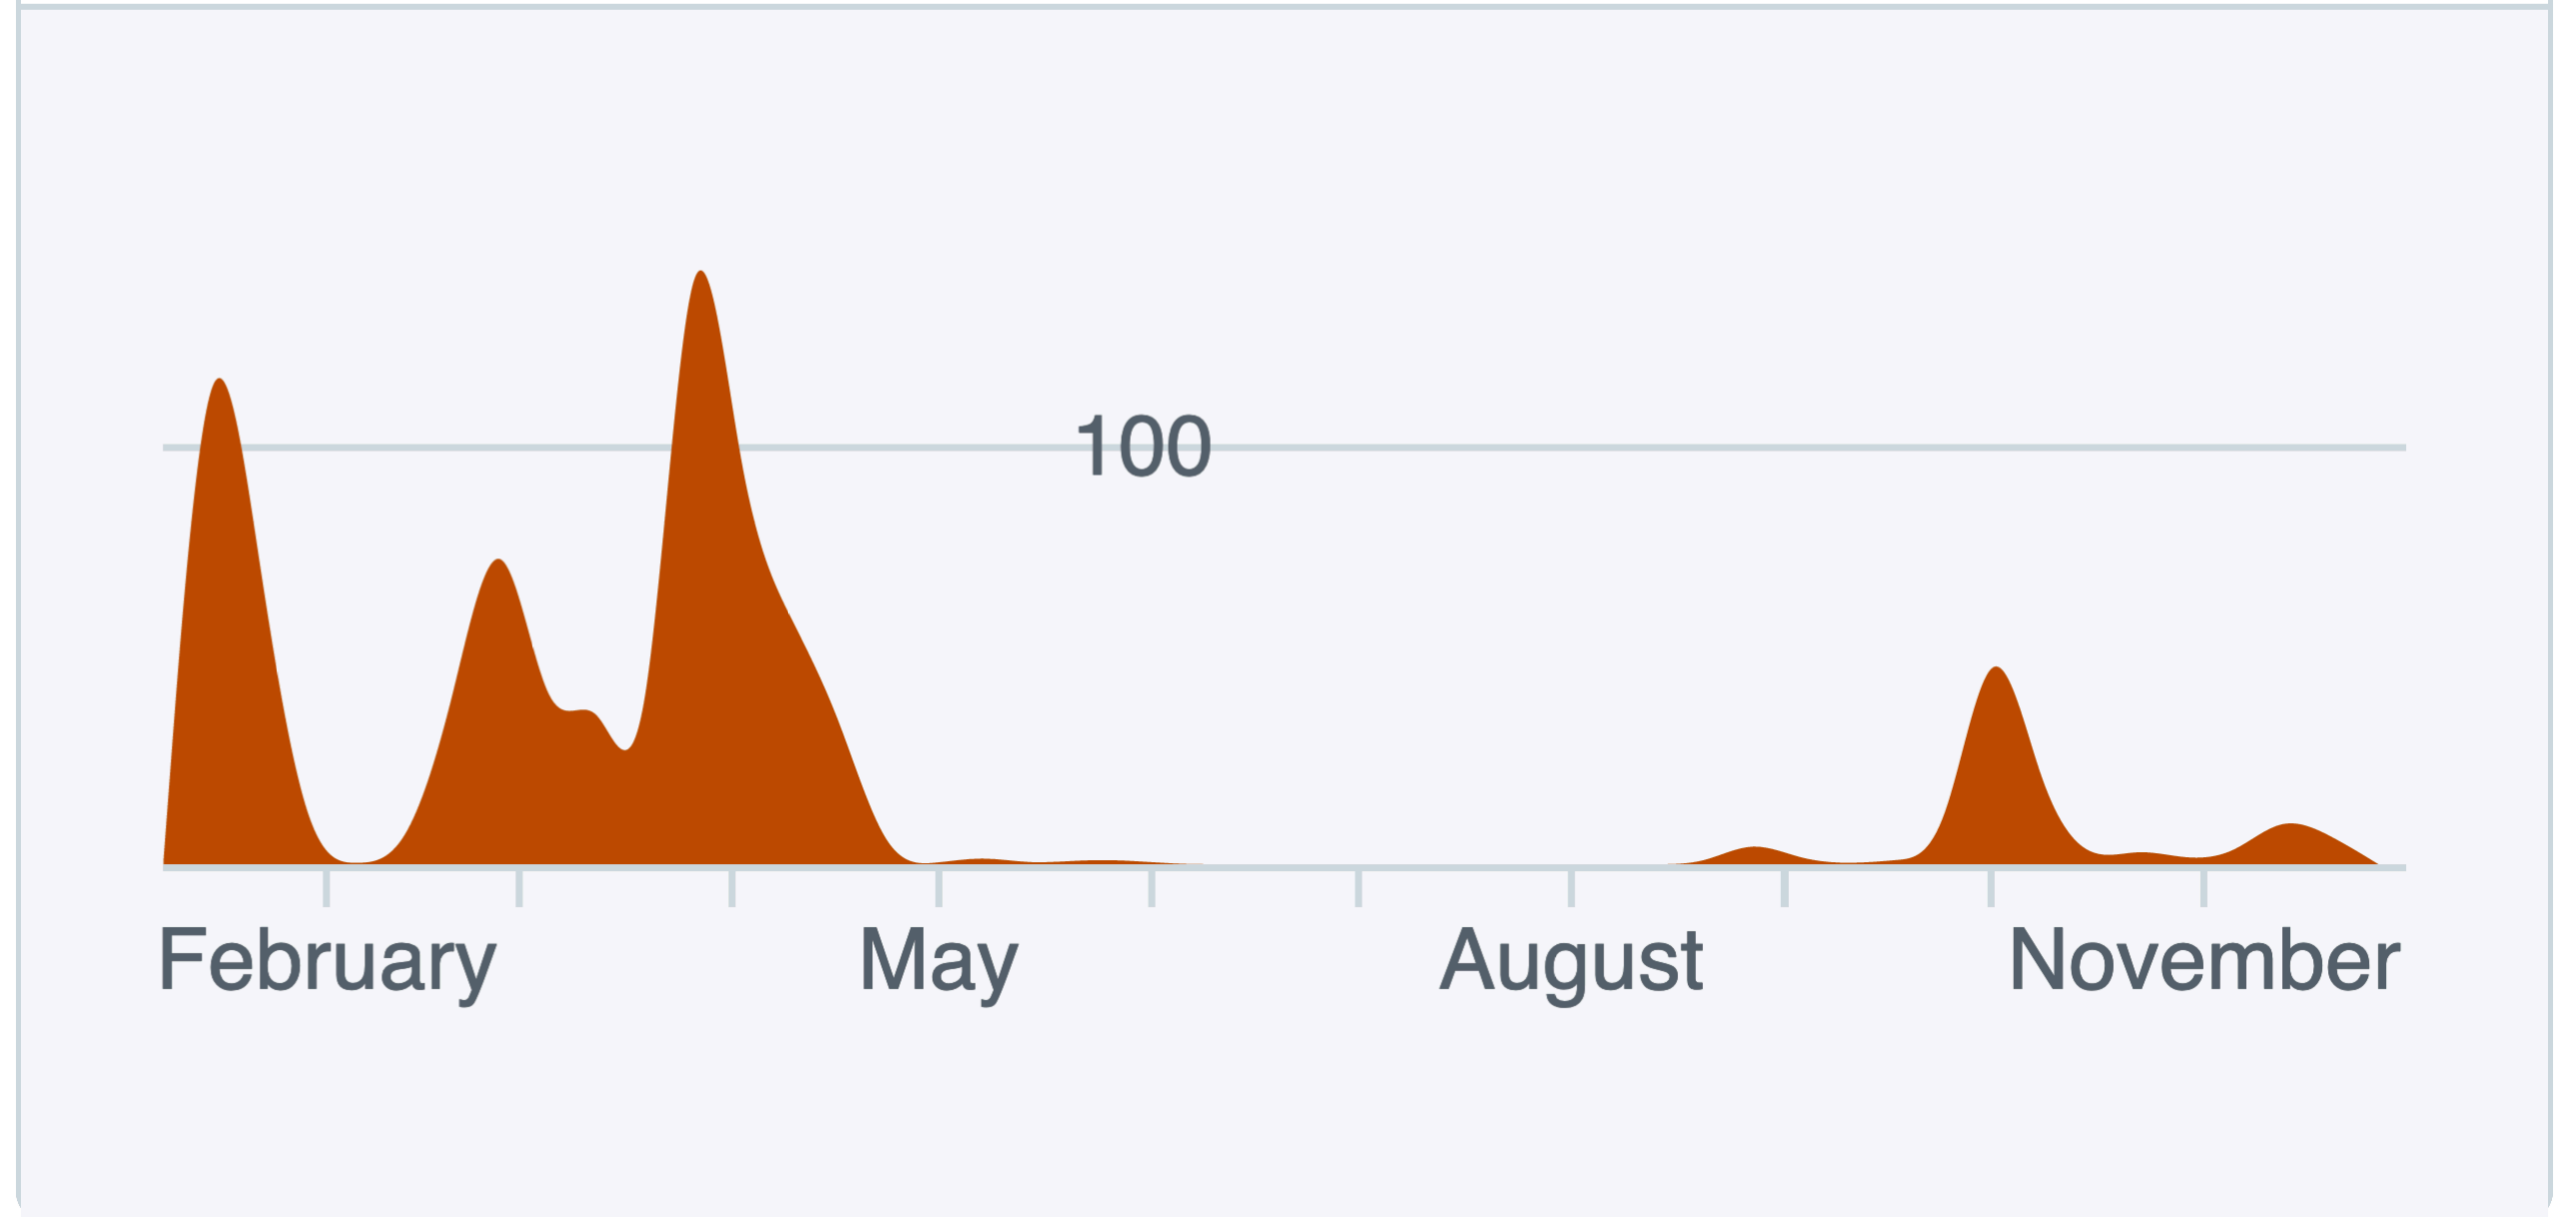




yhzhang0128

#1

914 commits 70,178 ++ 66,163 --



Motivation

Obstacles

Ideas

Implementation

Evaluation

Summer 2020

Fall 2020

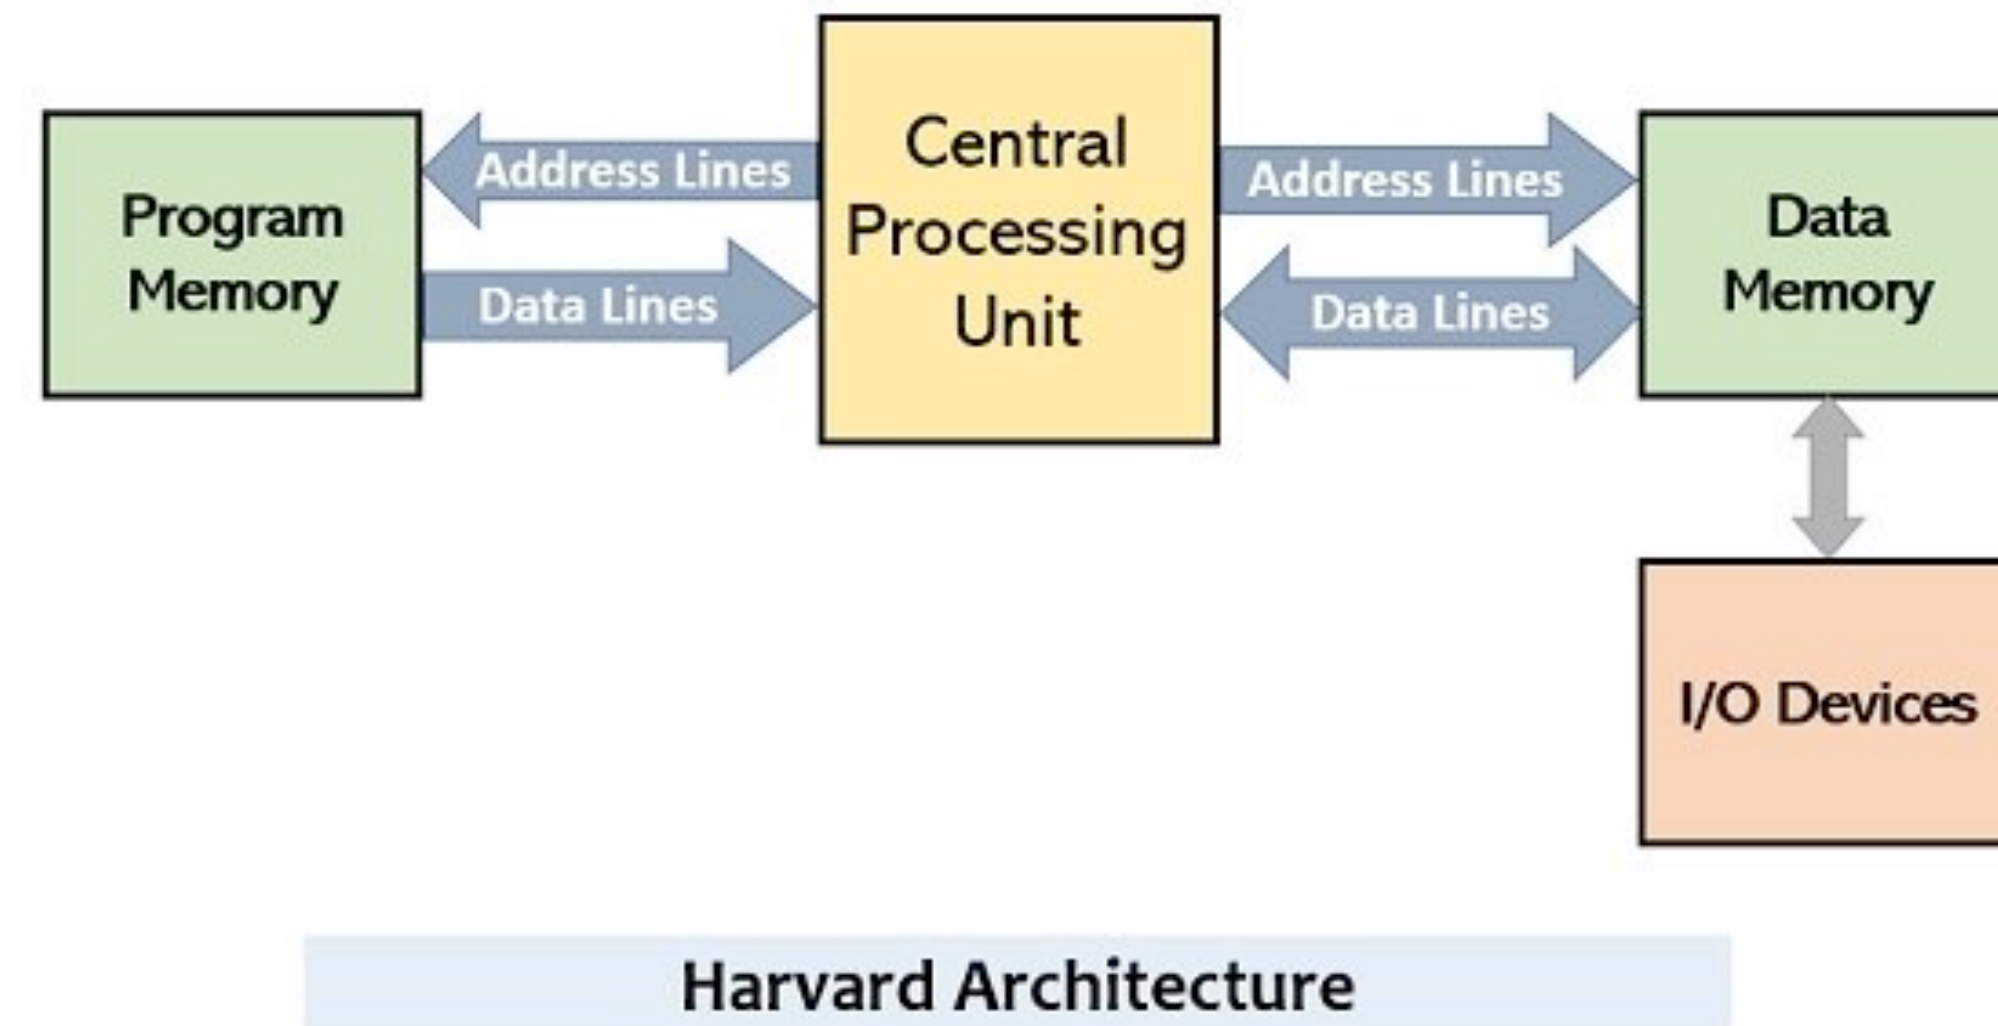
Fall 2021

Spring 2022

Fall 2022

The **bug** taking me **>1 day** to fix

```
core = RocketCoreParams(  
  useVM = false,  
  fpu = None,  
  mulDiv = Some(MulDivParams(mulUnroll = 8)),  
  btb = None,  
  dcache = Some(DCacheParams(  
    rowBits = site(SystemBusKey).beatBits,  
    nSets = 256, // 16Kb scratchpad  
    nWays = 1,  
    nTLBEntries = 4,  
    nMSHRs = 0,  
    blockBytes = site(CacheBlockBytes),  
    scratch = Some(0x80000000L)),  
  icache = Some(ICacheParams(  
    rowBits = site(SystemBusKey).beatBits,  
    nSets = 64,  
    nWays = 1,  
    nTLBEntries = 4,  
    blockBytes = site(CacheBlockBytes))))))
```




Lesson

Implementing a system is **non-trivial**
and requires **determination** and **hard work**

A 4.5-year research process



Then, **challenge** state-of-the-art



Next step: **Publish** the research

Fighting for a world where
every college student can
read **all** the code of an
operating system

Lines of Code	What?	Lines of Code	What?
199	Boot Loader & TTY Driver	336	File System
182	SD Card Driver	264	Applications & Daemons
32	Interrupt & Exception Handling	269	Library & Networking (TBA)
137	Page Table & Software Translation	64	Makefile
345	Timer, Scheduler & System Call	172	RISC-V Emulator & Board Tools