


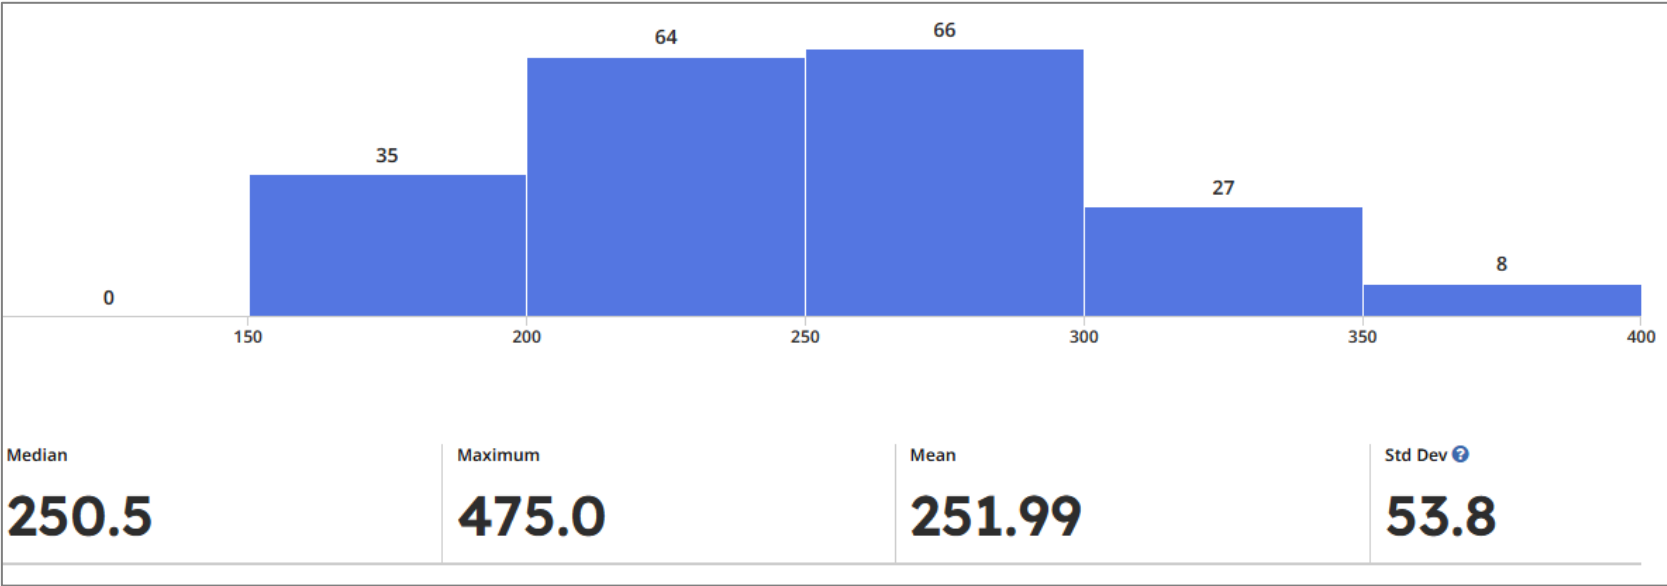
# **CS 211: Intro to Computer Architecture**

## ***9.1: The Hardware-Software Interface***

**Minesh Patel**

Spring 2025 – Tuesday 25 March

# Midterm Exam Statistics



# Announcements

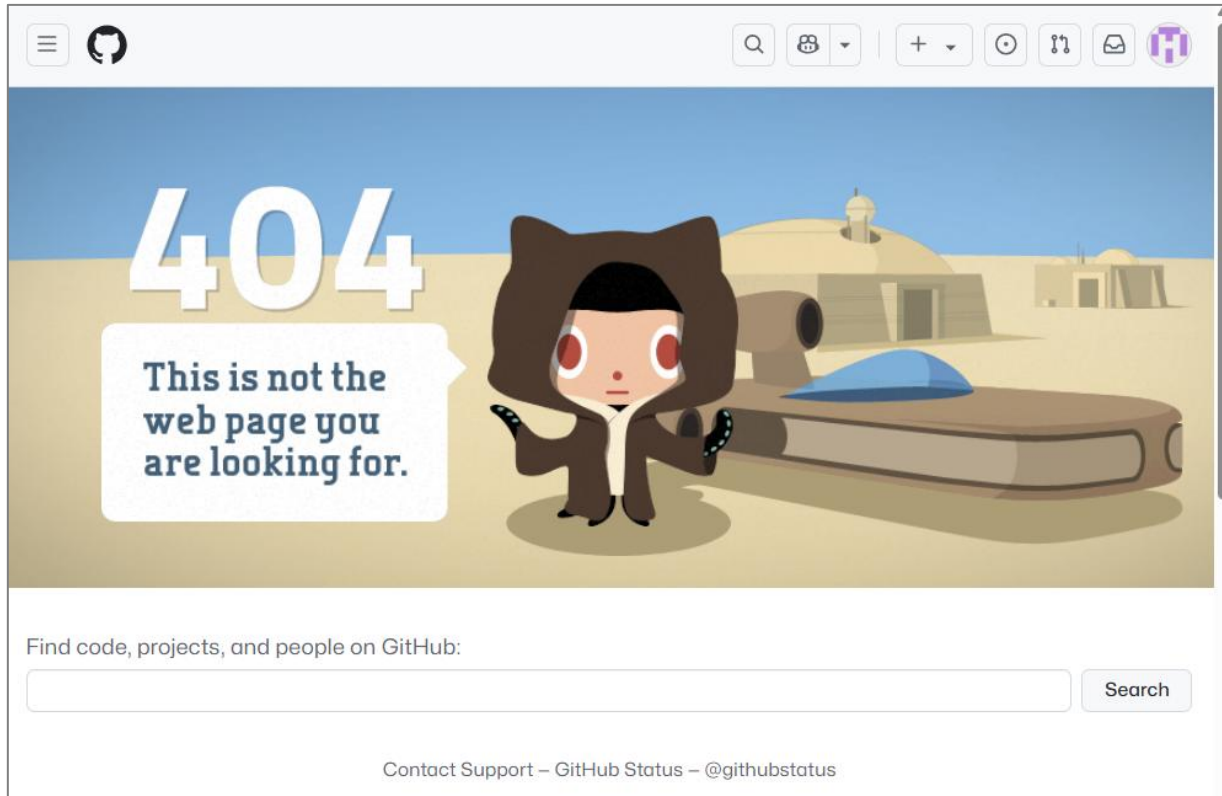
- Assignments
  - **PA3**: due **Friday @ 23:59**
  - **Extra Credit**: replaces WA6, due in **two weeks**

# Agenda

- **Aside:** PA3 Filesystem Permissions

# Filesystem Permissions

- Controls who can access a file



```
mp2099@ilab3:~/cs211/pa3$ ls -l
total 48
drwxr-x--- 7 mp2099 mp2099  39 Mar 14 18:15 filesystems
-rw-r----- 1 mp2099 mp2099 4417 Mar 19 13:31 Makefile
drwxr-x--- 2 mp2099 mp2099  11 Mar 19 13:29 src
----- 1 mp2099 mp2099    0 Mar 24 11:54 tmp
mp2099@ilab3:~/cs211/pa3$ cat tmp
cat: tmp: Permission denied
mp2099@ilab3:~/cs211/pa3$ |
```

- Fundamental part of most operating systems

# Filesystem Permissions on Linux

```
mp2099@ilab3: ~/cs211/pa3
mp2099@ilab3:~/cs211/pa3$ ls
filesystems  Makefile  src
mp2099@ilab3:~/cs211/pa3$ touch tmp
mp2099@ilab3:~/cs211/pa3$ ls -l
total 48
drwxr-x--- 7 mp2099 mp2099  39 Mar 14 18:15 filesystems
-rw-r----- 1 mp2099 mp2099 4417 Mar 19 13:31 Makefile
drwxr-x--- 2 mp2099 mp2099  11 Mar 19 13:29 src
-rw-r----- 1 mp2099 mp2099   0 Mar 24 11:54 tmp
mp2099@ilab3:~/cs211/pa3$ chmod -w tmp
mp2099@ilab3:~/cs211/pa3$ ls -l
total 48
drwxr-x--- 7 mp2099 mp2099  39 Mar 14 18:15 filesystems
-rw-r----- 1 mp2099 mp2099 4417 Mar 19 13:31 Makefile
drwxr-x--- 2 mp2099 mp2099  11 Mar 19 13:29 src
-r--r----- 1 mp2099 mp2099   0 Mar 24 11:54 tmp
mp2099@ilab3:~/cs211/pa3$ rm tmp
rm: remove write-protected regular empty file 'tmp'? n
mp2099@ilab3:~/cs211/pa3$ chmod -r tmp
mp2099@ilab3:~/cs211/pa3$ ls -l
total 48
drwxr-x--- 7 mp2099 mp2099  39 Mar 14 18:15 filesystems
-rw-r----- 1 mp2099 mp2099 4417 Mar 19 13:31 Makefile
drwxr-x--- 2 mp2099 mp2099  11 Mar 19 13:29 src
----- 1 mp2099 mp2099   0 Mar 24 11:54 tmp
mp2099@ilab3:~/cs211/pa3$ cat tmp
cat: tmp: Permission denied
mp2099@ilab3:~/cs211/pa3$
```

← **Creating a new file  
(read || write by default)**

← **chmod to read-only**

← **chmod to no  
permissions at all**

# Filesystem Permissions for PA3

**Think about whether you should allow:**

- creating a file...
  - In a read-only directory?
  - In a r+w subdirectory of a read-only directory?
- removing a file...
  - That has read-only permissions?
  - Inside a read-only directory?
- changing permissions on a file/directory...
  - That has no access permissions?
  - What about its children?
- ...

# Agenda


- **What's Next: Layers of Abstraction**
- Computer Organization Revisited
- The Representation of Code
  - Machine Language and ISAs
  - Choosing an ISA



# So Far: Representation of Data

64-bit address space  
( $2^{64}$  bytes)

- **Characters**  
    'a' '\n' '\0'
- **Numbers**  
    0b1011      -6½  
    0xae.14      $\pi^{-2}$
- **Arrays / strings**  
    [0, 1, 2]    "me"
- ...

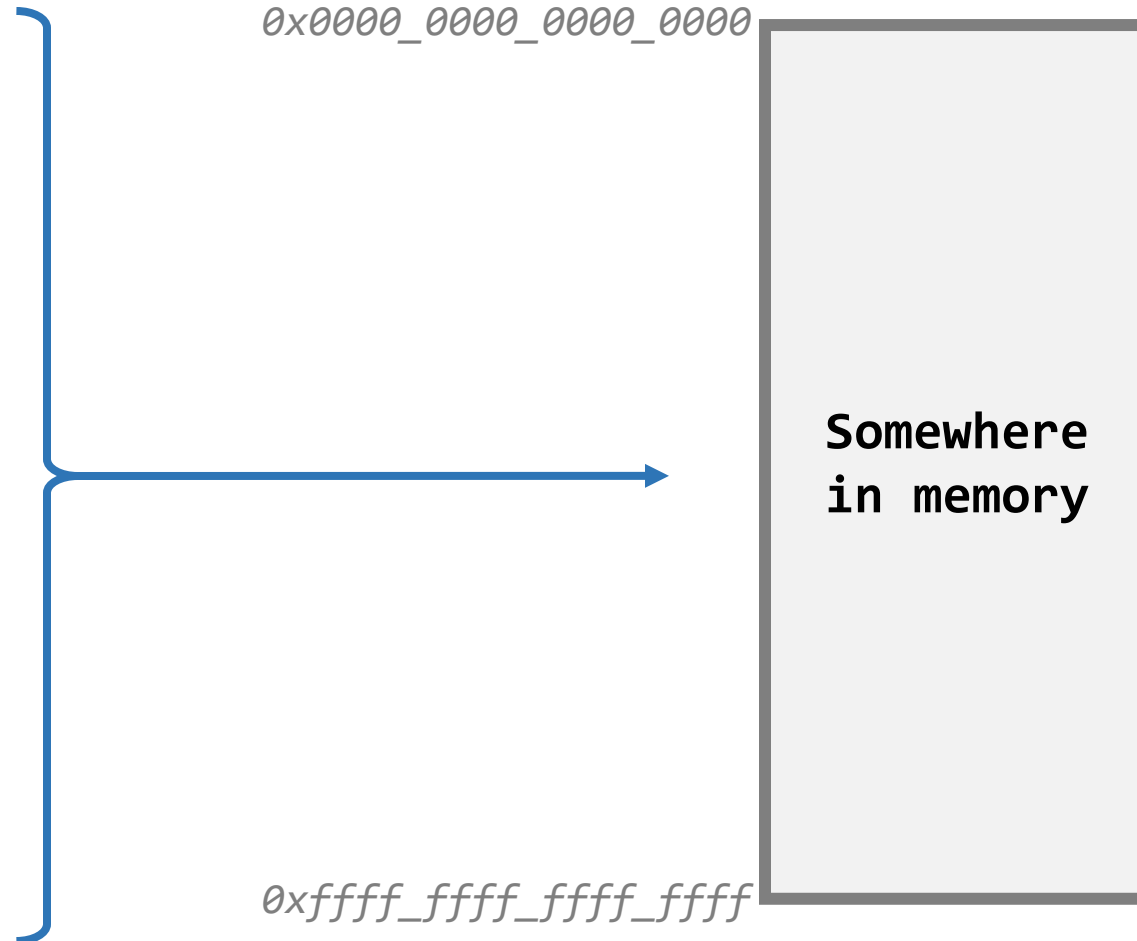


# Up Next: Representation of Code

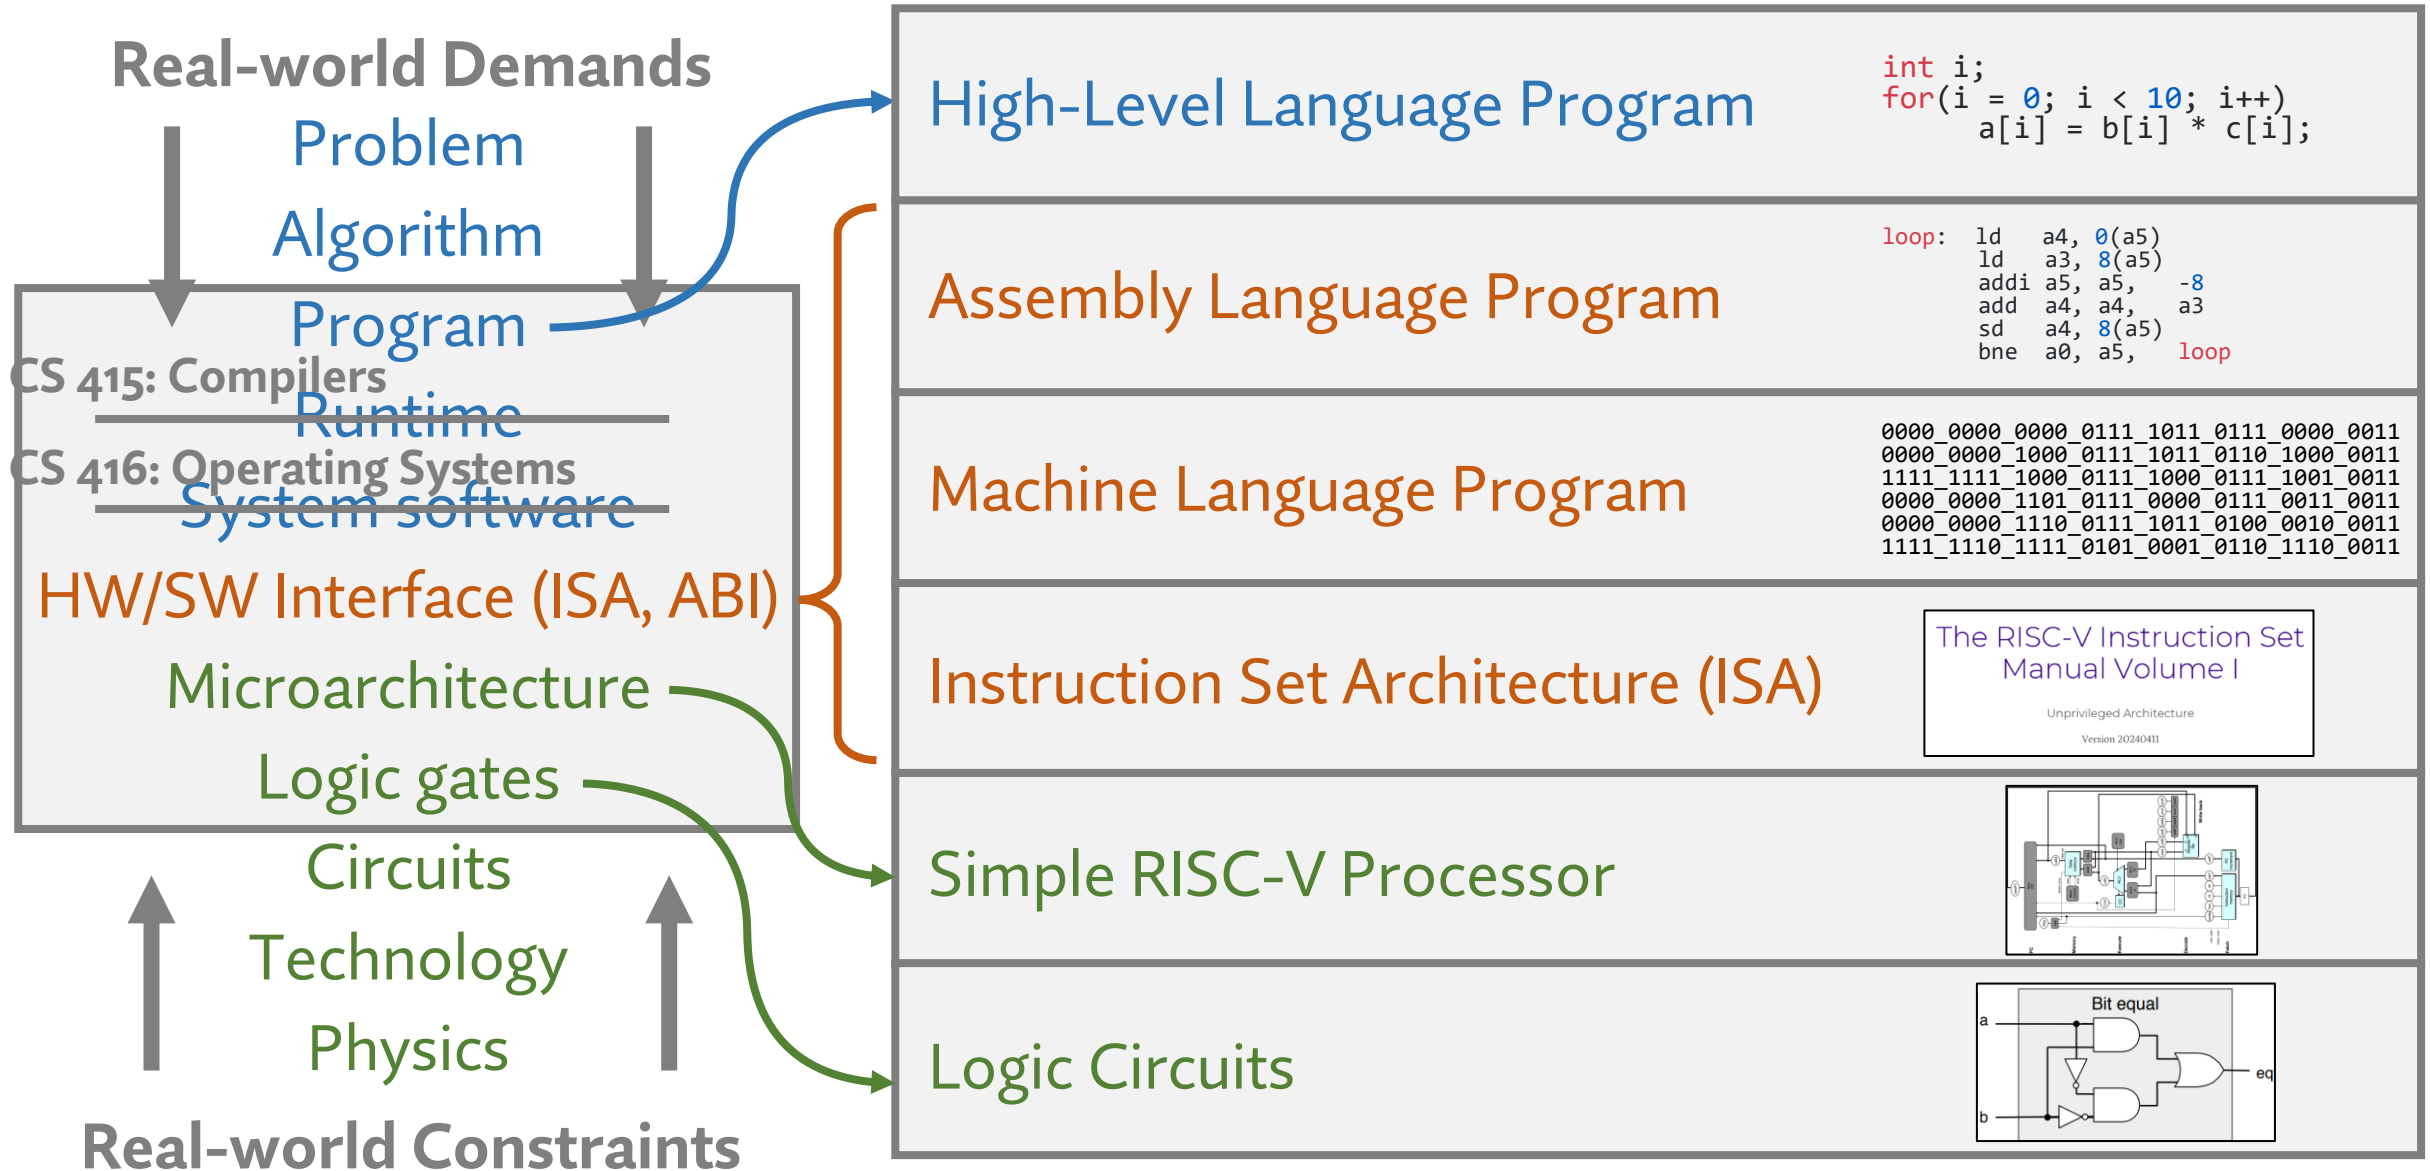
**64-bit address space**

( $2^{64}$  bytes)

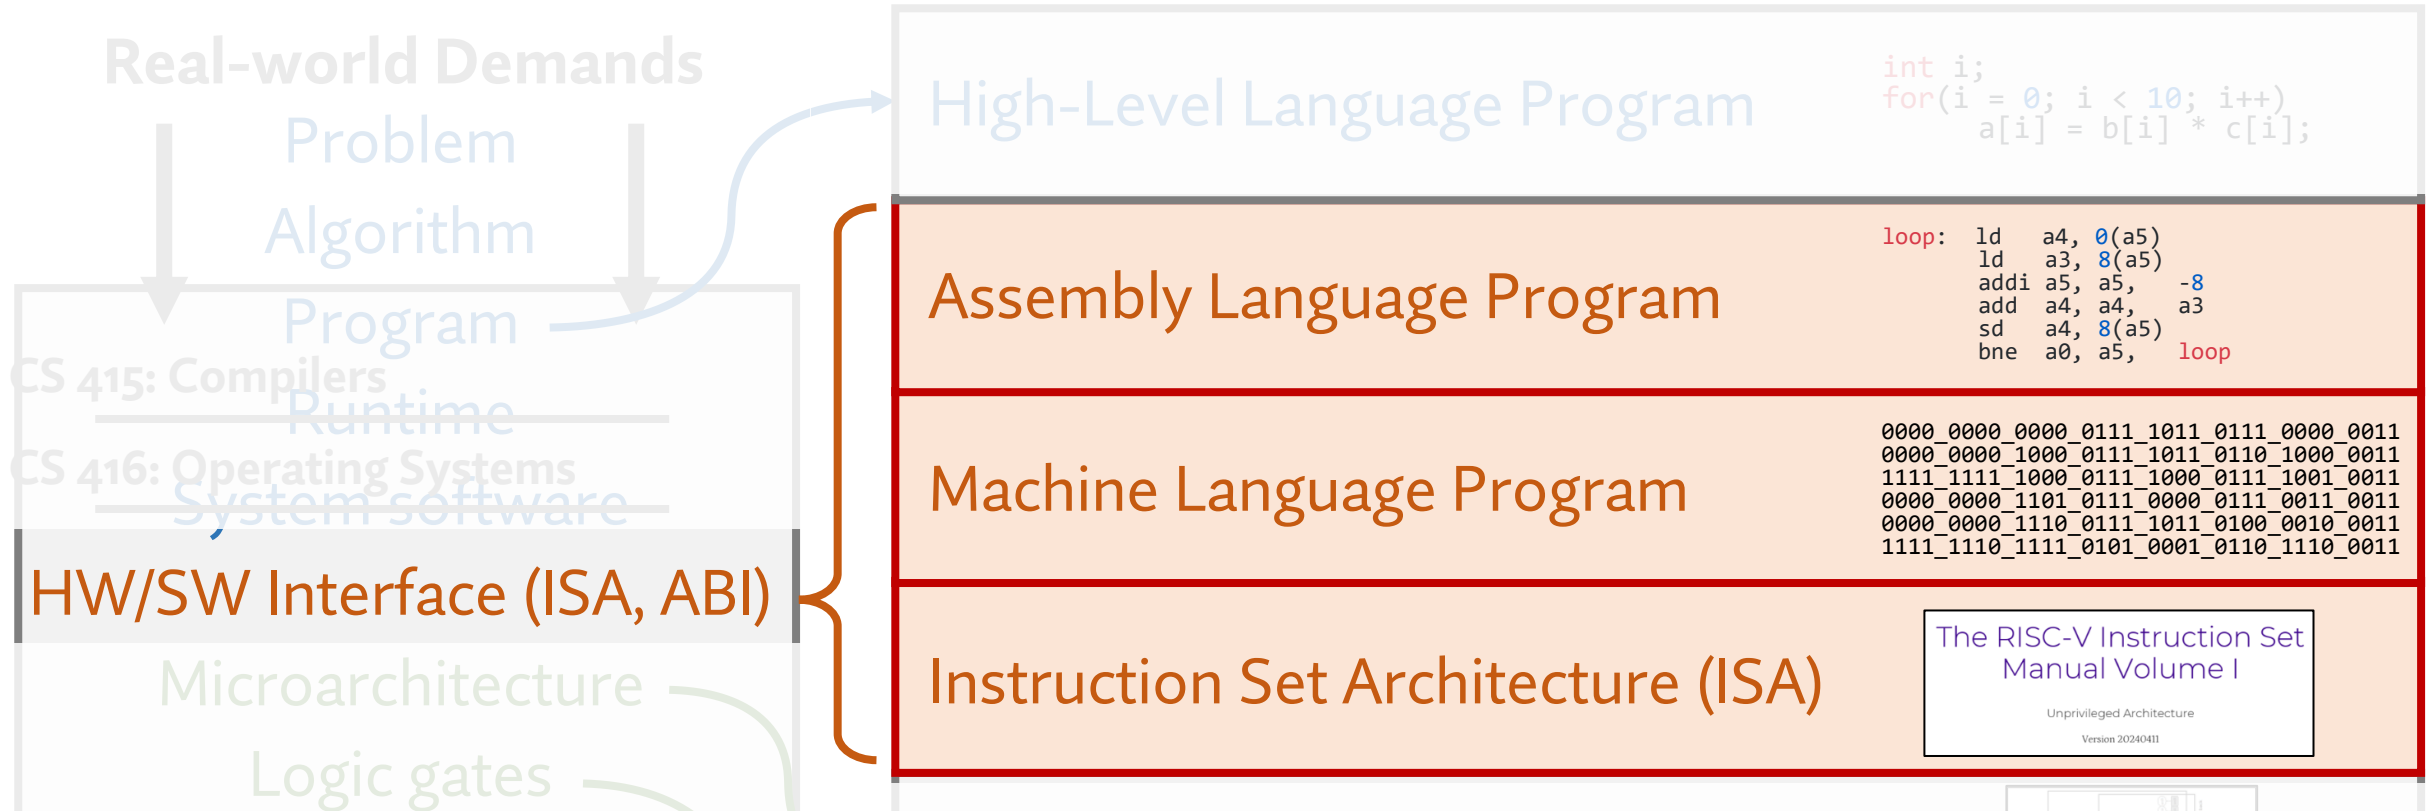
```
*a += 3;  
  
b = b ? 1 : 0;  
  
while(c) {...}  
  
func(a, b, c);  
  
...
```



# Layers of Abstraction for CS 211



# Layers of Abstraction for CS 211: What's Next



*Next ~2-3 weeks:*

The representation of code:

**assembly** and **machine** language programs

# Agenda


- What's Next: Layers of Abstraction
- **Computer Organization Revisited**
- The Representation of Code
  - Machine Language and ISAs
  - Choosing an ISA

# Recap: Computer Organization

**Input/Output**  
(network, USB, etc.)



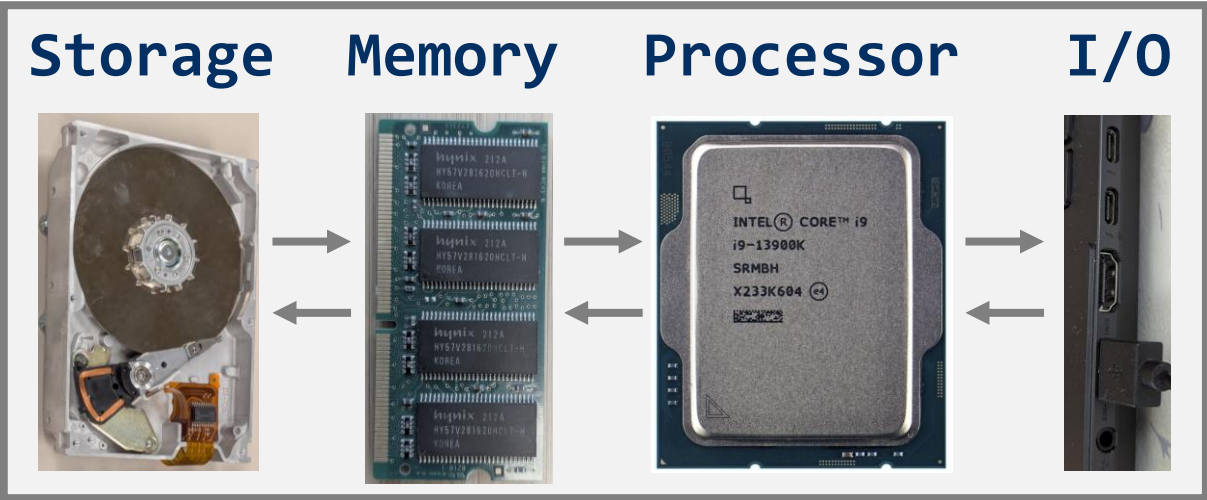




**Storage**  
(big, slow, cheap memory)



**Processor**  
(where code runs)

**Main Memory**  
(small, fast, expensive)

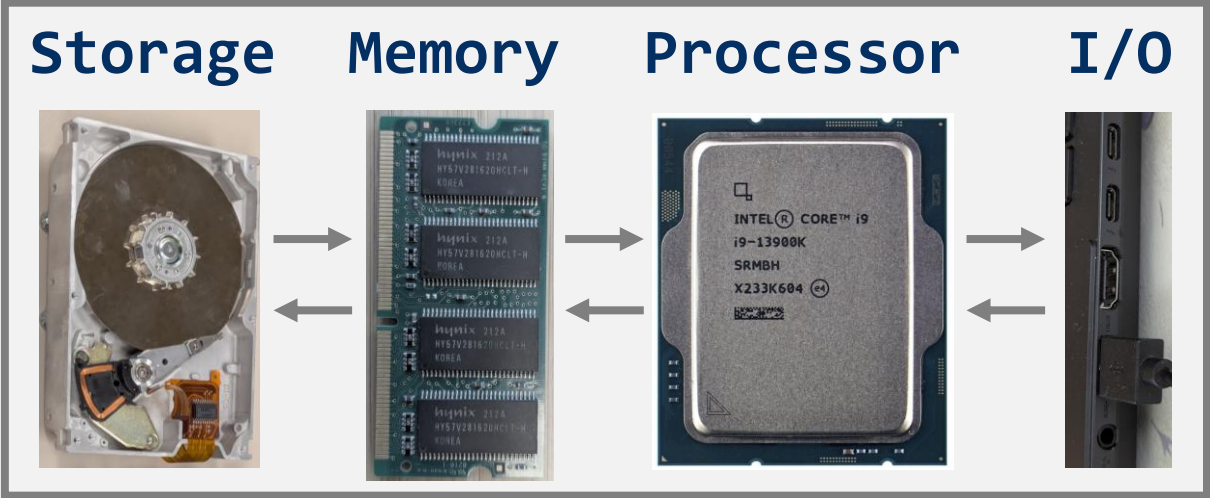
# Computer Organization: Mental Model




Computer

# Computer Organization: Simplified

## Computer



## Simplified Mental Model








# Abstract View: What Do We Want from a Computer?

- **Ideal:** runs **everything** quickly, efficiently, securely, reliably, etc.
  - Independent of your favorite programming language
- **Reality:** no single hardware design can achieve all of the above

## Theme 1: Tradeoffs - pick the right tools



<https://dlcdnwebimgs.asus.com/gain/d47ac467-7878-47b7-a29a-c9adfb80b9bd/>



<https://www.ebay.com/itm/154513065385?epid=1252327726&hash=item23f9b245a9:g:GoIAAOSwi7VcjULN>



<https://www.amazon.com/Google-Pixel-Unlocked-Smartphone-Ultrawide/dp/BogHJZPFDD>




<https://www.ultrabookreview.com/64501-asus-vivobook-pro-16x-oled-review/>

# Abstract View: What Do We Want from a Computer?

- **Ideal:** runs **everything** quickly, efficiently, securely, reliably, etc.
  - Independent of your favorite programming language
- **Reality:** no single hardware design can achieve all of the above

## Theme 2: Separation of Concerns


Software



Abstraction



Hardware

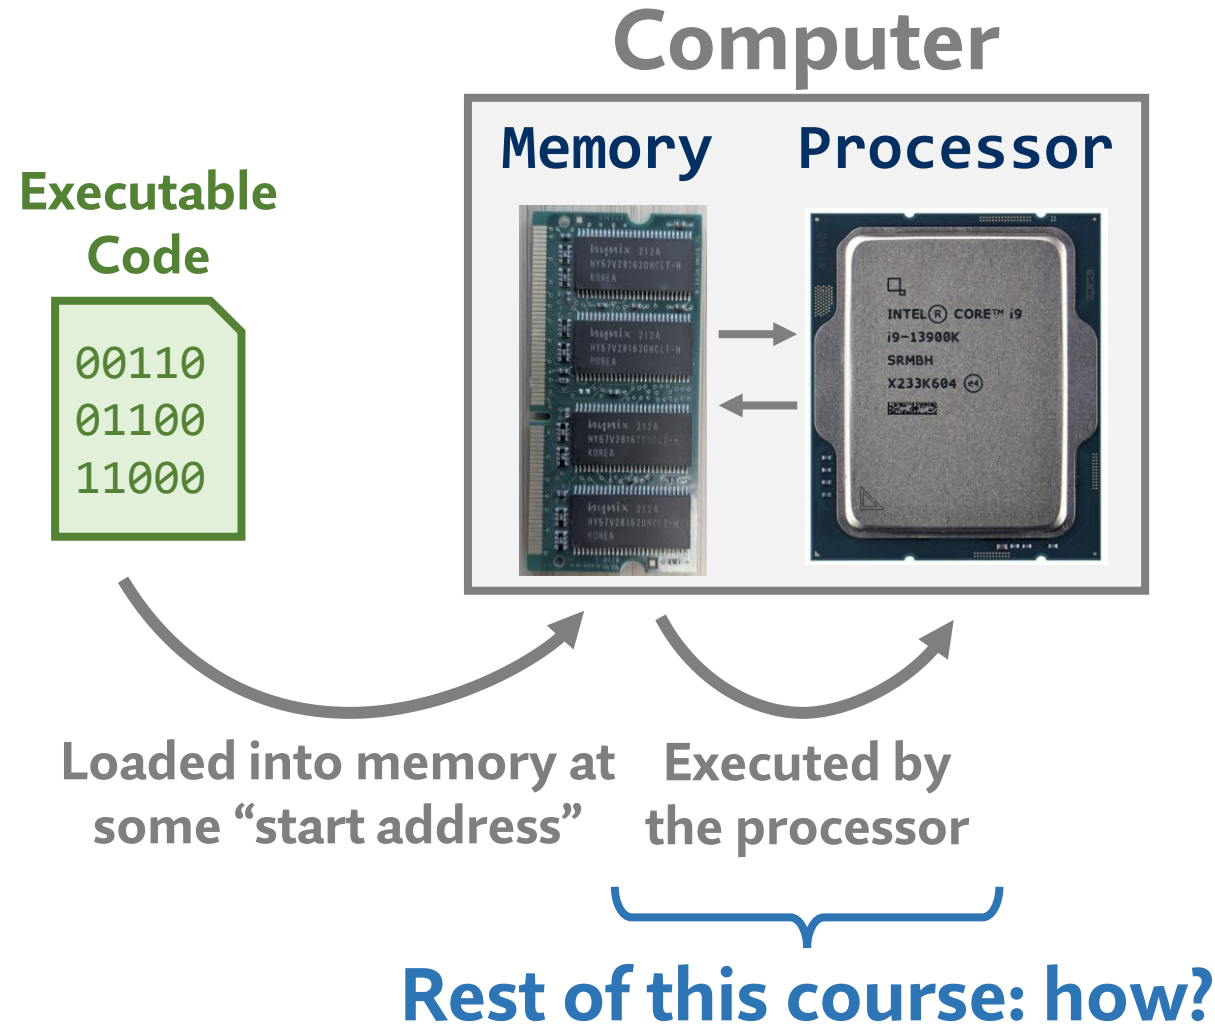


# Agenda

- What's Next: Layers of Abstraction
- Computer Organization Revisited
- **The Representation of Code**
  - Machine Language and ISAs
  - Choosing an ISA

# Running Code on the Processor

- Code lives somewhere **in memory**



# From Source Code to Execution




- The processor only understands **machine language code**
  - Anything that runs is **translated to machine language** (via compiler, interpreter, etc.)

## Various Source Codes

Rust

Golang




C++



Python

Java

Scala



The representation  
of code

Machine  
Language

```
00110  
01100  
11000
```

TRANSLATION



## Computer

Memory

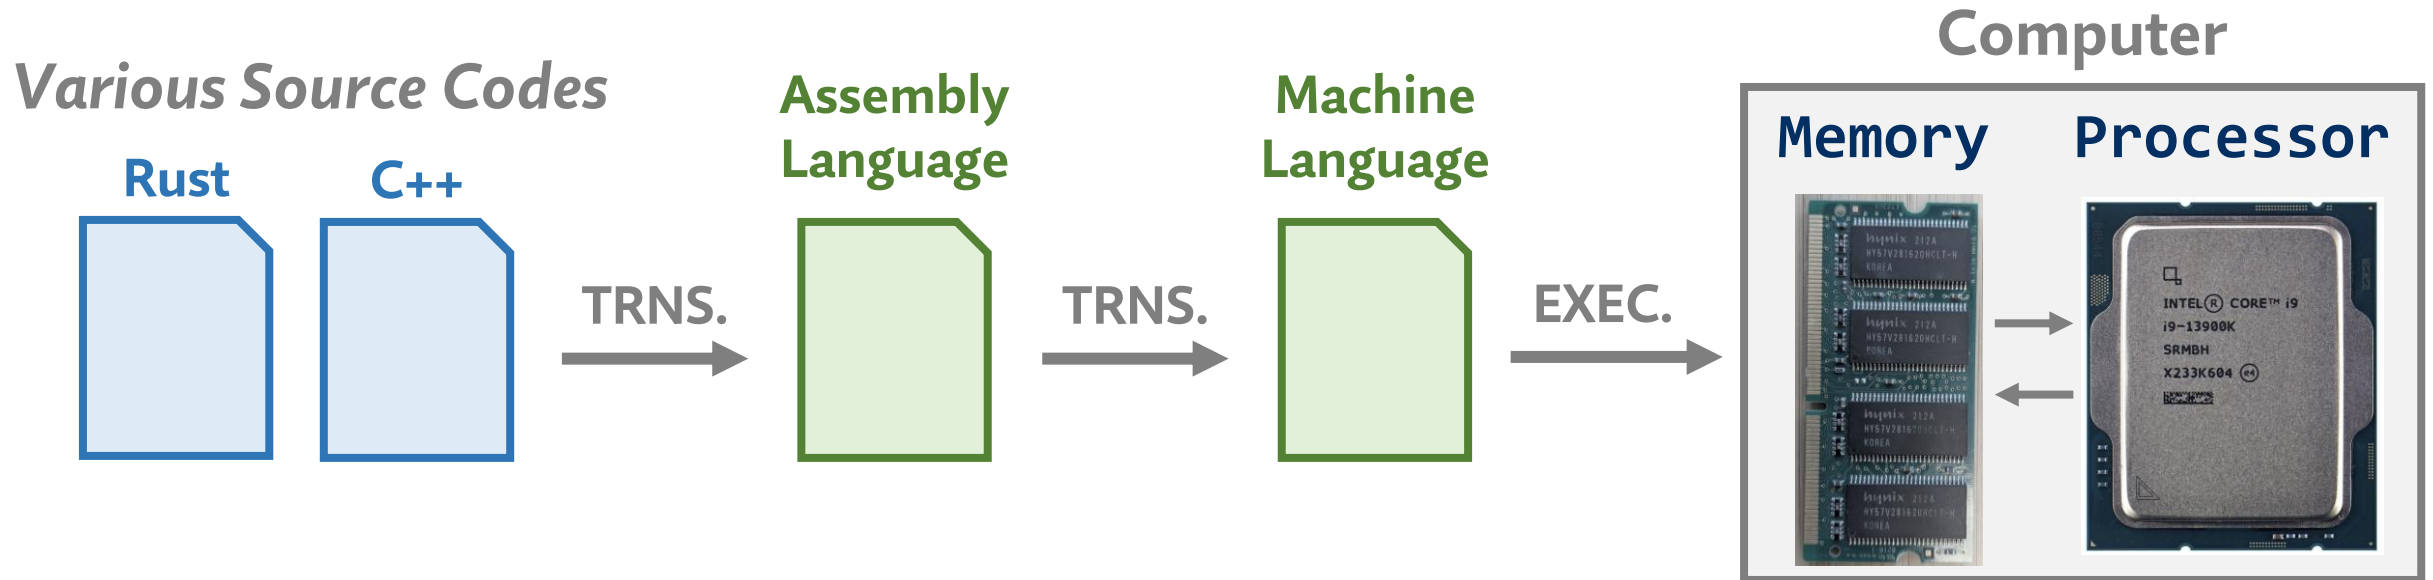
Processor



Loaded into memory at  
some "start address"


Executed by  
the processor

# Assembly Language: Human-Readable Machine Code



- Machine language **exposes the raw capabilities** of a processor
  - E.g., **primitive operations** such as arithmetic, memory access, etc.
  - **Not** meant for humans to read/write
- Assembly language provides an **(almost) 1:1 abstraction** of machine code

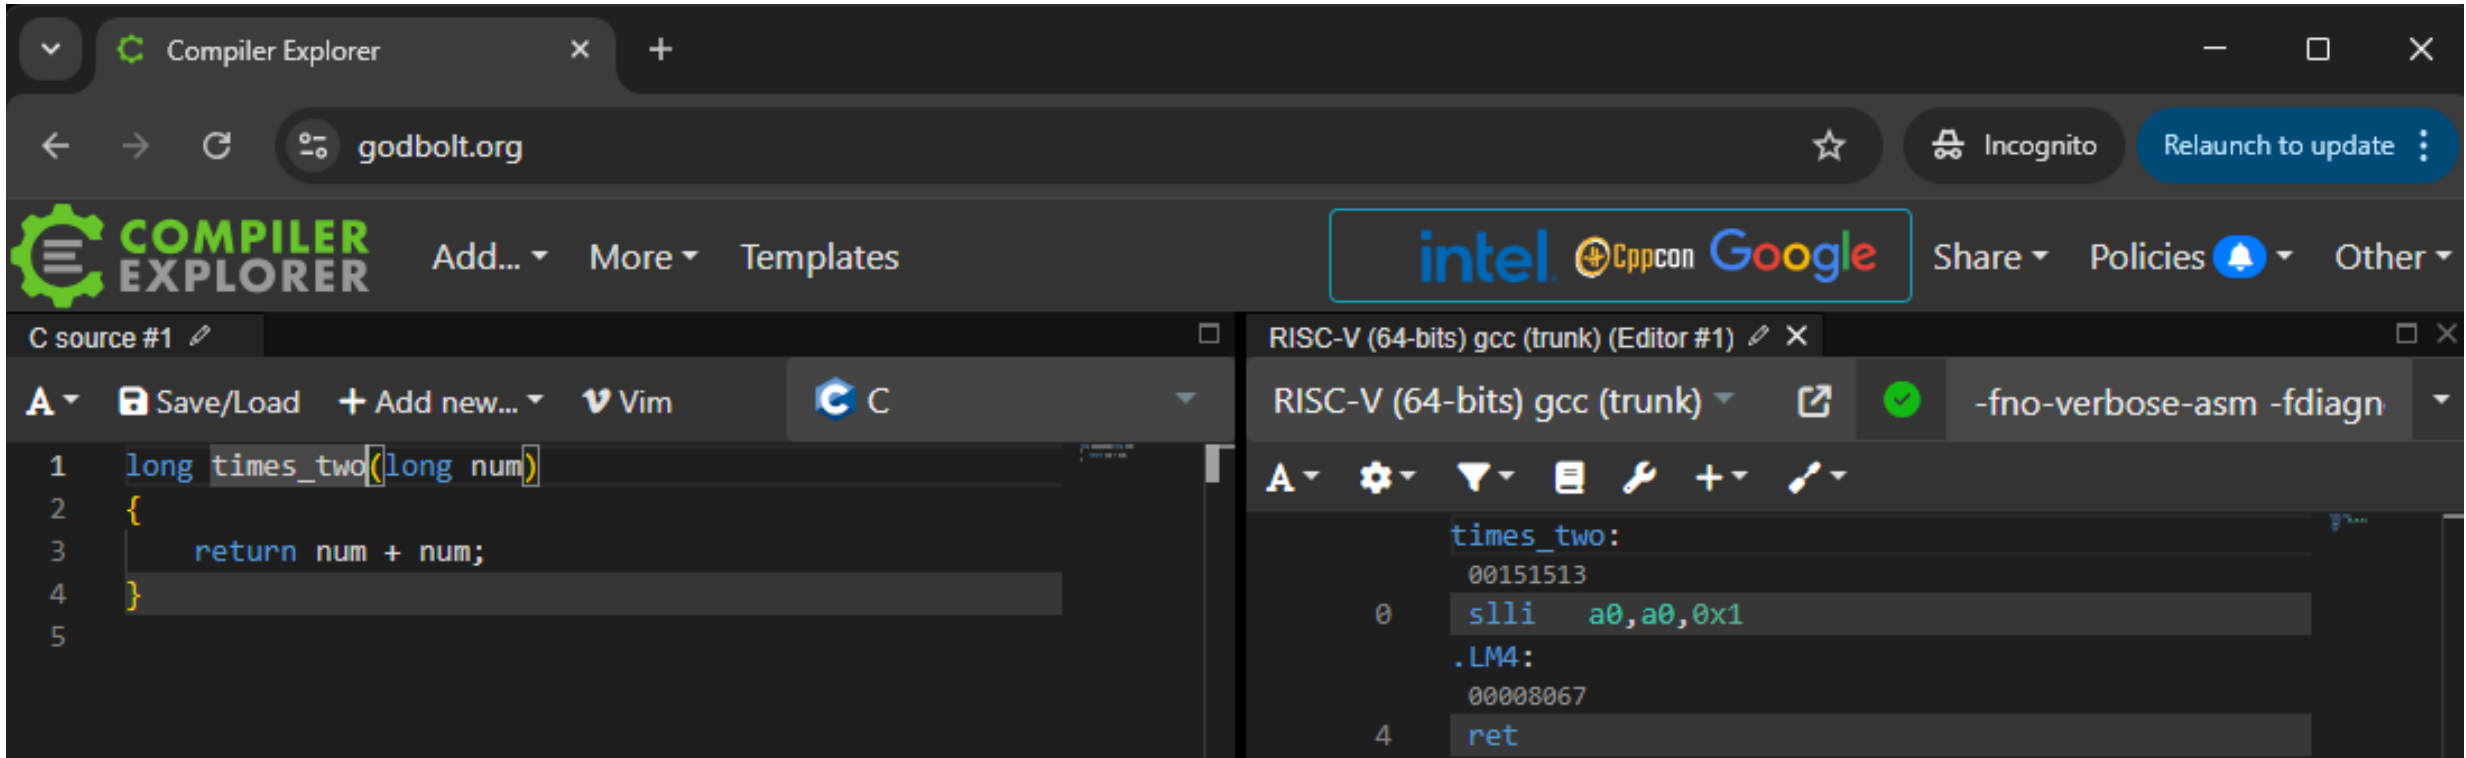
# The Terminator T-800 Runs MOS 6502 Assembly



```
8 .....
9
10          ORG  $4000
11 A1       =   $3C
12 A2       =   $3E
13 A4       =   $42
14 AUXMOVE  =   $C311
15
16 .....
17 • SETUP - move data for VTOC
18 • and catalog to auxmem at
19 • 0000-03FF (pseudo trk 11
20 • 0-3)
21 .....
22 SETUP   LDA  #<VTOC
23         STA  A1
24         LDA  #>VTOC
25         STA  A1+1
26         LDA  #<END
27         STA  A2
28         LDA  #>END
29         STA  A2+1
30         LDA  #$00
31         STA  A4
32         LDA  #$00
33         STA  A4+1
34         SEC
35         JMP  AUXMOVE
36
```

# Example: Godbolt Compiler Explorer

<https://godbolt.org/z/cqvME3xEn>



C Code

Assembly Code

Machine Code

```
long times_two(long num)
{
    return num + num;
}
```



```
.times_two
    slli a0, a0, 0x1
    ret
```

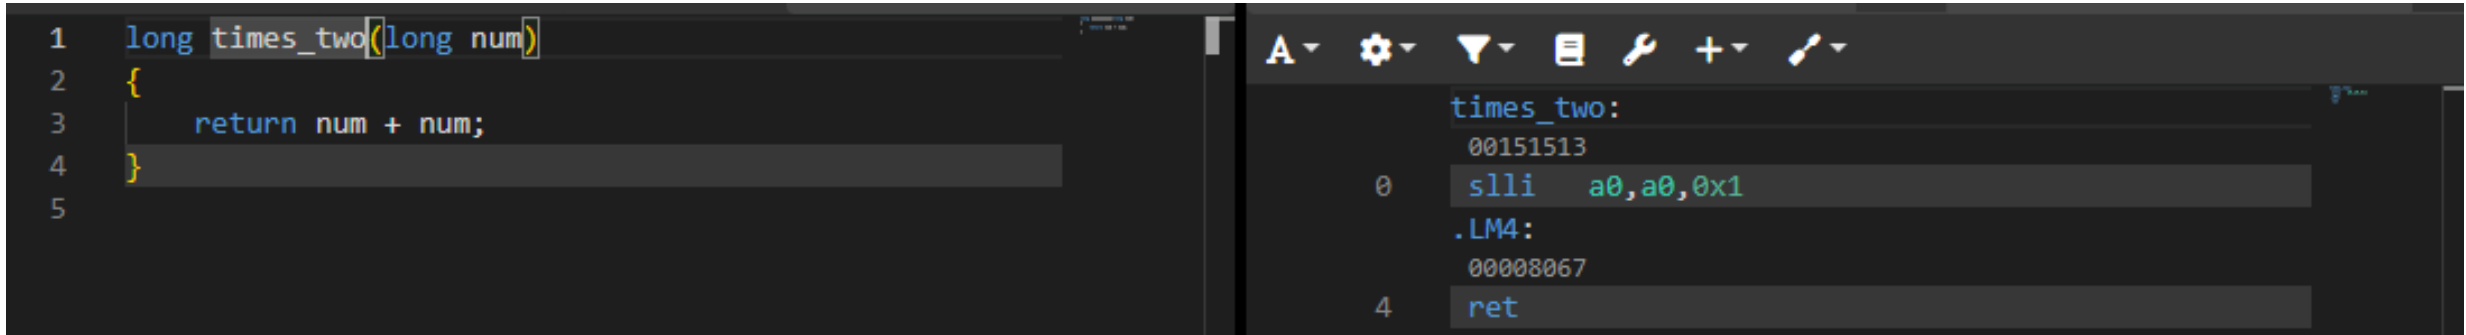


```
0x00151513
0x00008067
```



# Aside: Compiler Transformation

<https://godbolt.org/z/cqvME3xEn>



The screenshot shows a debugger interface with two panels. The left panel displays C code for a function named `times_two` that takes a `long num` and returns `num + num`. The right panel shows the corresponding assembly code for the function, which includes instructions for loading the argument, shifting it left by one bit, and returning the result.

```
1 long times_two(long num)
2 {
3     return num + num;
4 }
5
```

```
times_two:
00151513
0 slli a0,a0,0x1
.LM4:
00008067
4 ret
```

- Consider what I wrote in C vs. what processor will execute

## Supplied C Code

```
return num + num;
```

## “Decompiled” C Code

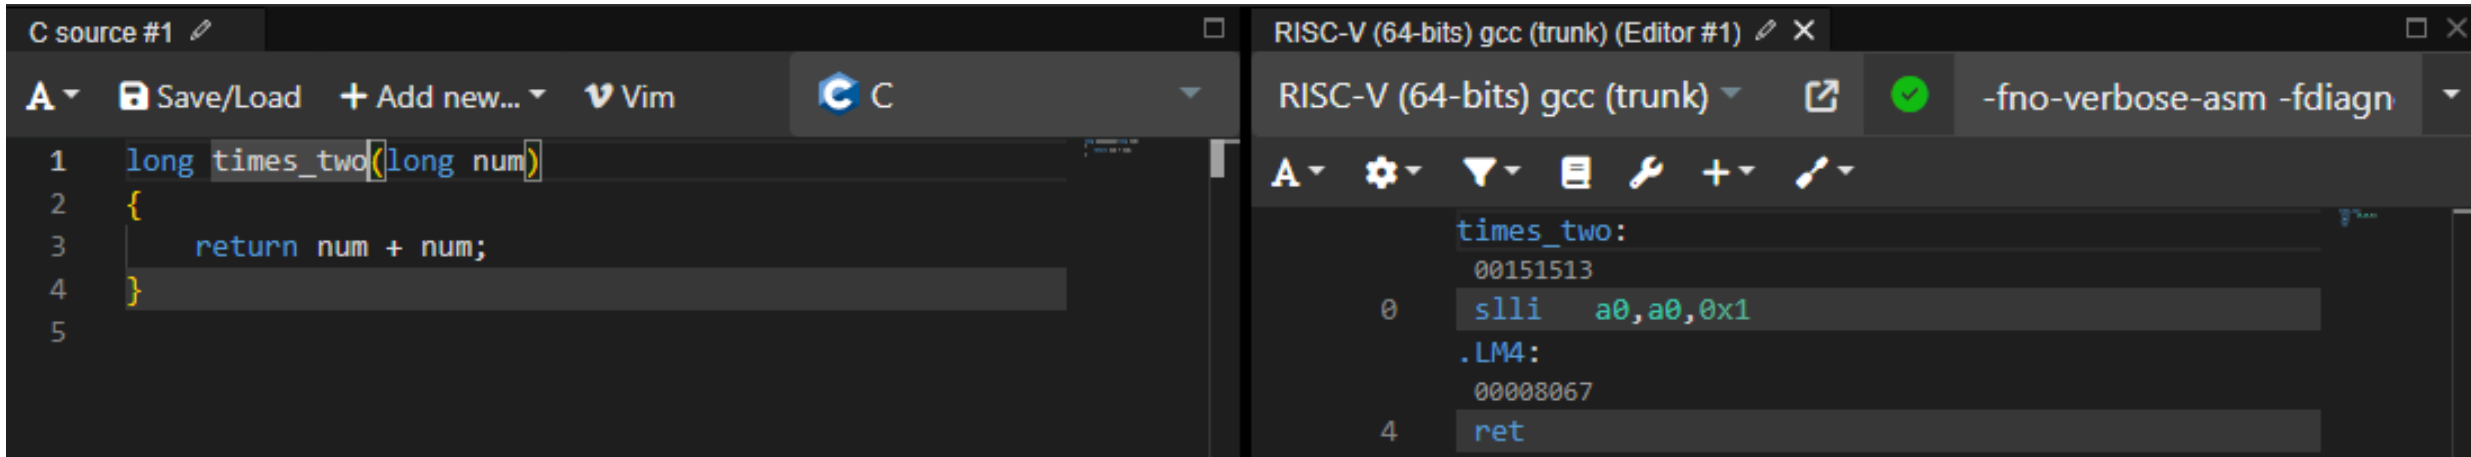
```
return num << 1;
```



Technically, this is undefined behavior in C!

# Example: RISC-V vs. x86-64

<https://godbolt.org/z/cqvME3xEn>



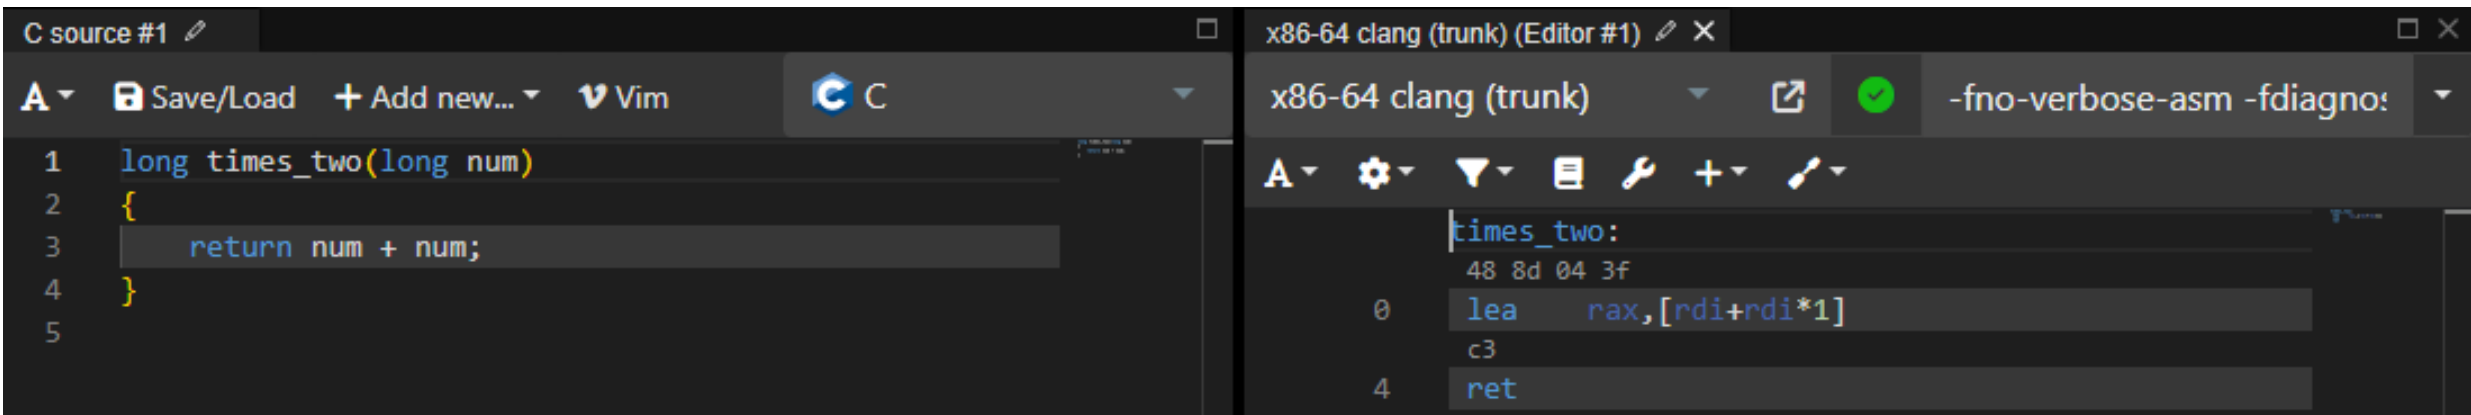
The screenshot shows a side-by-side comparison of C source code and its RISC-V assembly. The left pane, titled 'C source #1', contains the following code:

```
1 long times_two(long num)
2 {
3     return num + num;
4 }
5
```

The right pane, titled 'RISC-V (64-bits) gcc (trunk) (Editor #1)', shows the assembly output with the following instructions:

```
times_two:
00151513
0 slli a0,a0,0x1
.LM4:
00008067
4 ret
```

<https://godbolt.org/z/bxaxzEaKo>



The screenshot shows a side-by-side comparison of the same C source code and its x86-64 assembly. The left pane is identical to the one above. The right pane, titled 'x86-64 clang (trunk) (Editor #1)', shows the assembly output with the following instructions:

```
times_two:
48 8d 04 3f
0 lea rax,[rdi+rdi*1]
c3
4 ret
```

# Sections of an Executable File

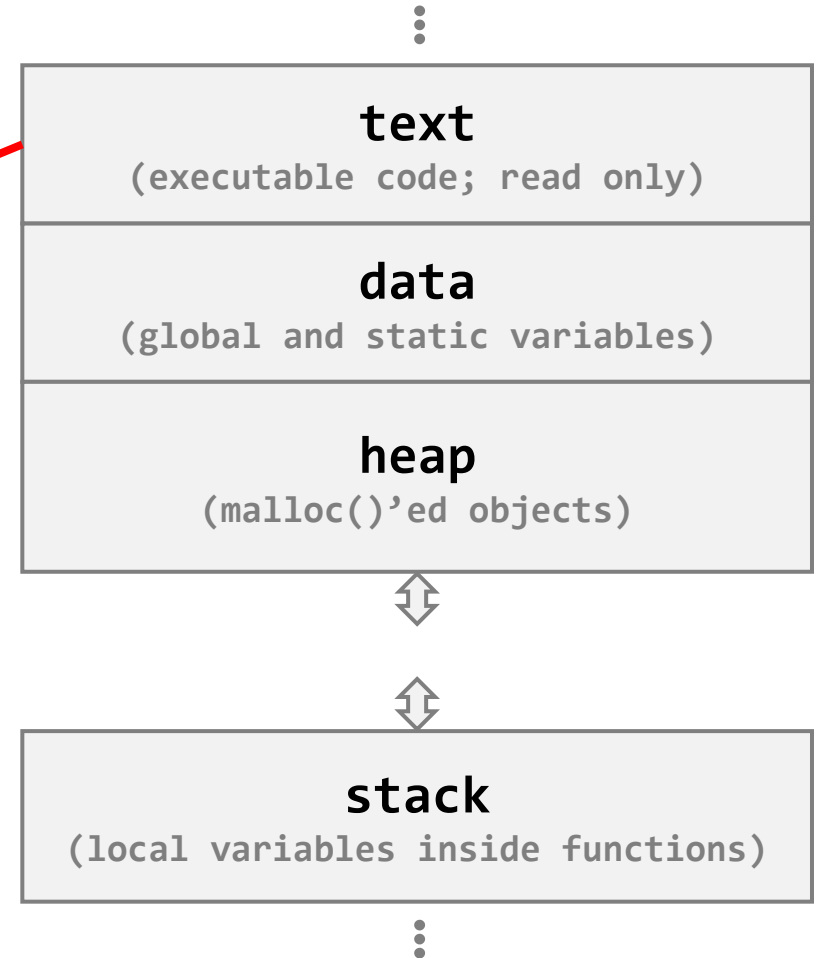
The tool “**objdump**” allows us to inspect sections of a program

```
mp2099@ilab4: ~/cs211/expe x + v - □ x
mp2099@ilab4:~/cs211/experiment$ /common/users/shared/cs211_s25_5678/toolchain
_glibc3/bin/riscv64-unknown-linux-gnu-objdump --section=.text hello.o -h

hello.o:      file format elf64-littleriscv

Sections:
Idx Name      Size      VMA           LMA           File off  Algn
 10 .text      000000c0  00000000000010460  00000000000010460  00000460  2**2
                CONTENTS, ALLOC, LOAD, READONLY, CODE
mp2099@ilab4:~/cs211/experiment$
```

## Typical C Program's Address Space



# Example: Hello World on ilab

```
#include <stdio.h>
#include <stdlib.h>

int main(int argc, char *argv[])
{
    print("Hello, World");
    return EXIT_SUCCESS;
}
```

hello.c

```
mp2099@ilab2: ~/cs211/experi x + v
mp2099@ilab2:~/cs211/experiment$ /common/users/shared/cs211_s25_5678/toolchain_glibc3/bin/riscv64-unknown-linux-gnu-gcc hello.c -o hello.o -mabi=lp64 -march=rv64i -Os -fno-inline -fomit-frame-pointer
mp2099@ilab2:~/cs211/experiment$ /common/users/shared/cs211_s25_5678/toolchain_glibc3/bin/riscv64-unknown-linux-gnu-objdump -L -S --section=.text hello.o

hello.o:          file format elf64-littleriscv

Disassembly of section .text:

000000000010460 <main>:
   10460:      00010537      lui      a0,0x10
   10464:      ff010113      addi     sp,sp,-16
   10468:      52850513      addi     a0,a0,1320 # 10528 <_IO_stdin_used+0x8>
   1046c:      00113423      sd      ra,8(sp)
   10470:      fe1ff0ef      jal     10450 <puts@plt>
   10474:      00813083      ld      ra,8(sp)
   10478:      00000513      li      a0,0
   1047c:      01010113      addi     sp,sp,16
   10480:      00008067      ret
```

**“disassembly”**

# Recap: C Compilation

Platform-Independent  
-----  
Platform-Dependent

Text File  
(C Source Code)

```
#include <stdio.h>
#include <stdlib.h>

int main(int argc, char *argv[])
{
    print("Hello, World");
    return EXIT_SUCCESS;
}
hello.c
```

C Preprocessor

Text File  
(Prep. C Source Code)

```
[...ENTIRE CONTENTS OF "stdio.h" and
stdlib.h...]

int main(int argc, char *argv[])
{
    print("Hello, World");
    return 0;
}
hello.i
```

C Compiler

Text File  
(ASM Source Code)

```
main:
    addi sp,sp,-32
    sd ra,24(sp)
    sd s0,16(sp)
    addi s0,sp,32
    mv a5,a0
    sd a1,-32(s0)
    sw a5,-20(s0)
    lla a0,.LC0
    call puts@plt
    li a5,0
    mv a0,a5
    ld ra,24(sp)
    ld s0,16(sp)
    addi sp,sp,32
    jr ra
hello.s
```

Assembler

Binary File  
(Object File)

```
00000070: 4865 6c6c 6f2c 2057 6f72 6c64 0000 4743 Hello, World..gc
00000080: 433a 2028 5562 756e 7475 2031 312e 342e c: (Ubuntu 11.4.
00000090: 302d 3175 6275 6e74 7531 7e32 322e 3034 0-lubuntu1-22.04
000000a0: 2920 3131 2e34 2e30 0041 3200 0000 7269 .) 11.4.0.A2...ri
000000b0: 7363 7600 0128 0000 0005 7276 3634 6932 scv,...rv64i2
000000c0: 7030 5f6d 3270 305f 6132 7030 5f66 3270 p0_m2p0_a2p0_f2p
000000d0: 305f 6432 7030 5f63 3270 3000 002e 7368 0_d2p0_c2p0...sh
000000e0: 7374 7274 6162 002e 7465 7874 002e 6461 strtab..text..da
000000f0: 7461 002e 6273 7300 2e72 6f64 6174 6100 ta..bss..rc
00000100: 2e63 6f6d 6d65 6e74 002e 6e6f 7465 2e47 .comment..r
00000110: 4e55 2d73 7461 636b 002e 7269 7363 762e nu-stack...r
00000120: 6174 7472 6962 7574 6573 0000 0000 0000 attributes..
```

Linker

Binary File  
(Object File)

Binary File  
(Object File)

Binary File  
(Object File)

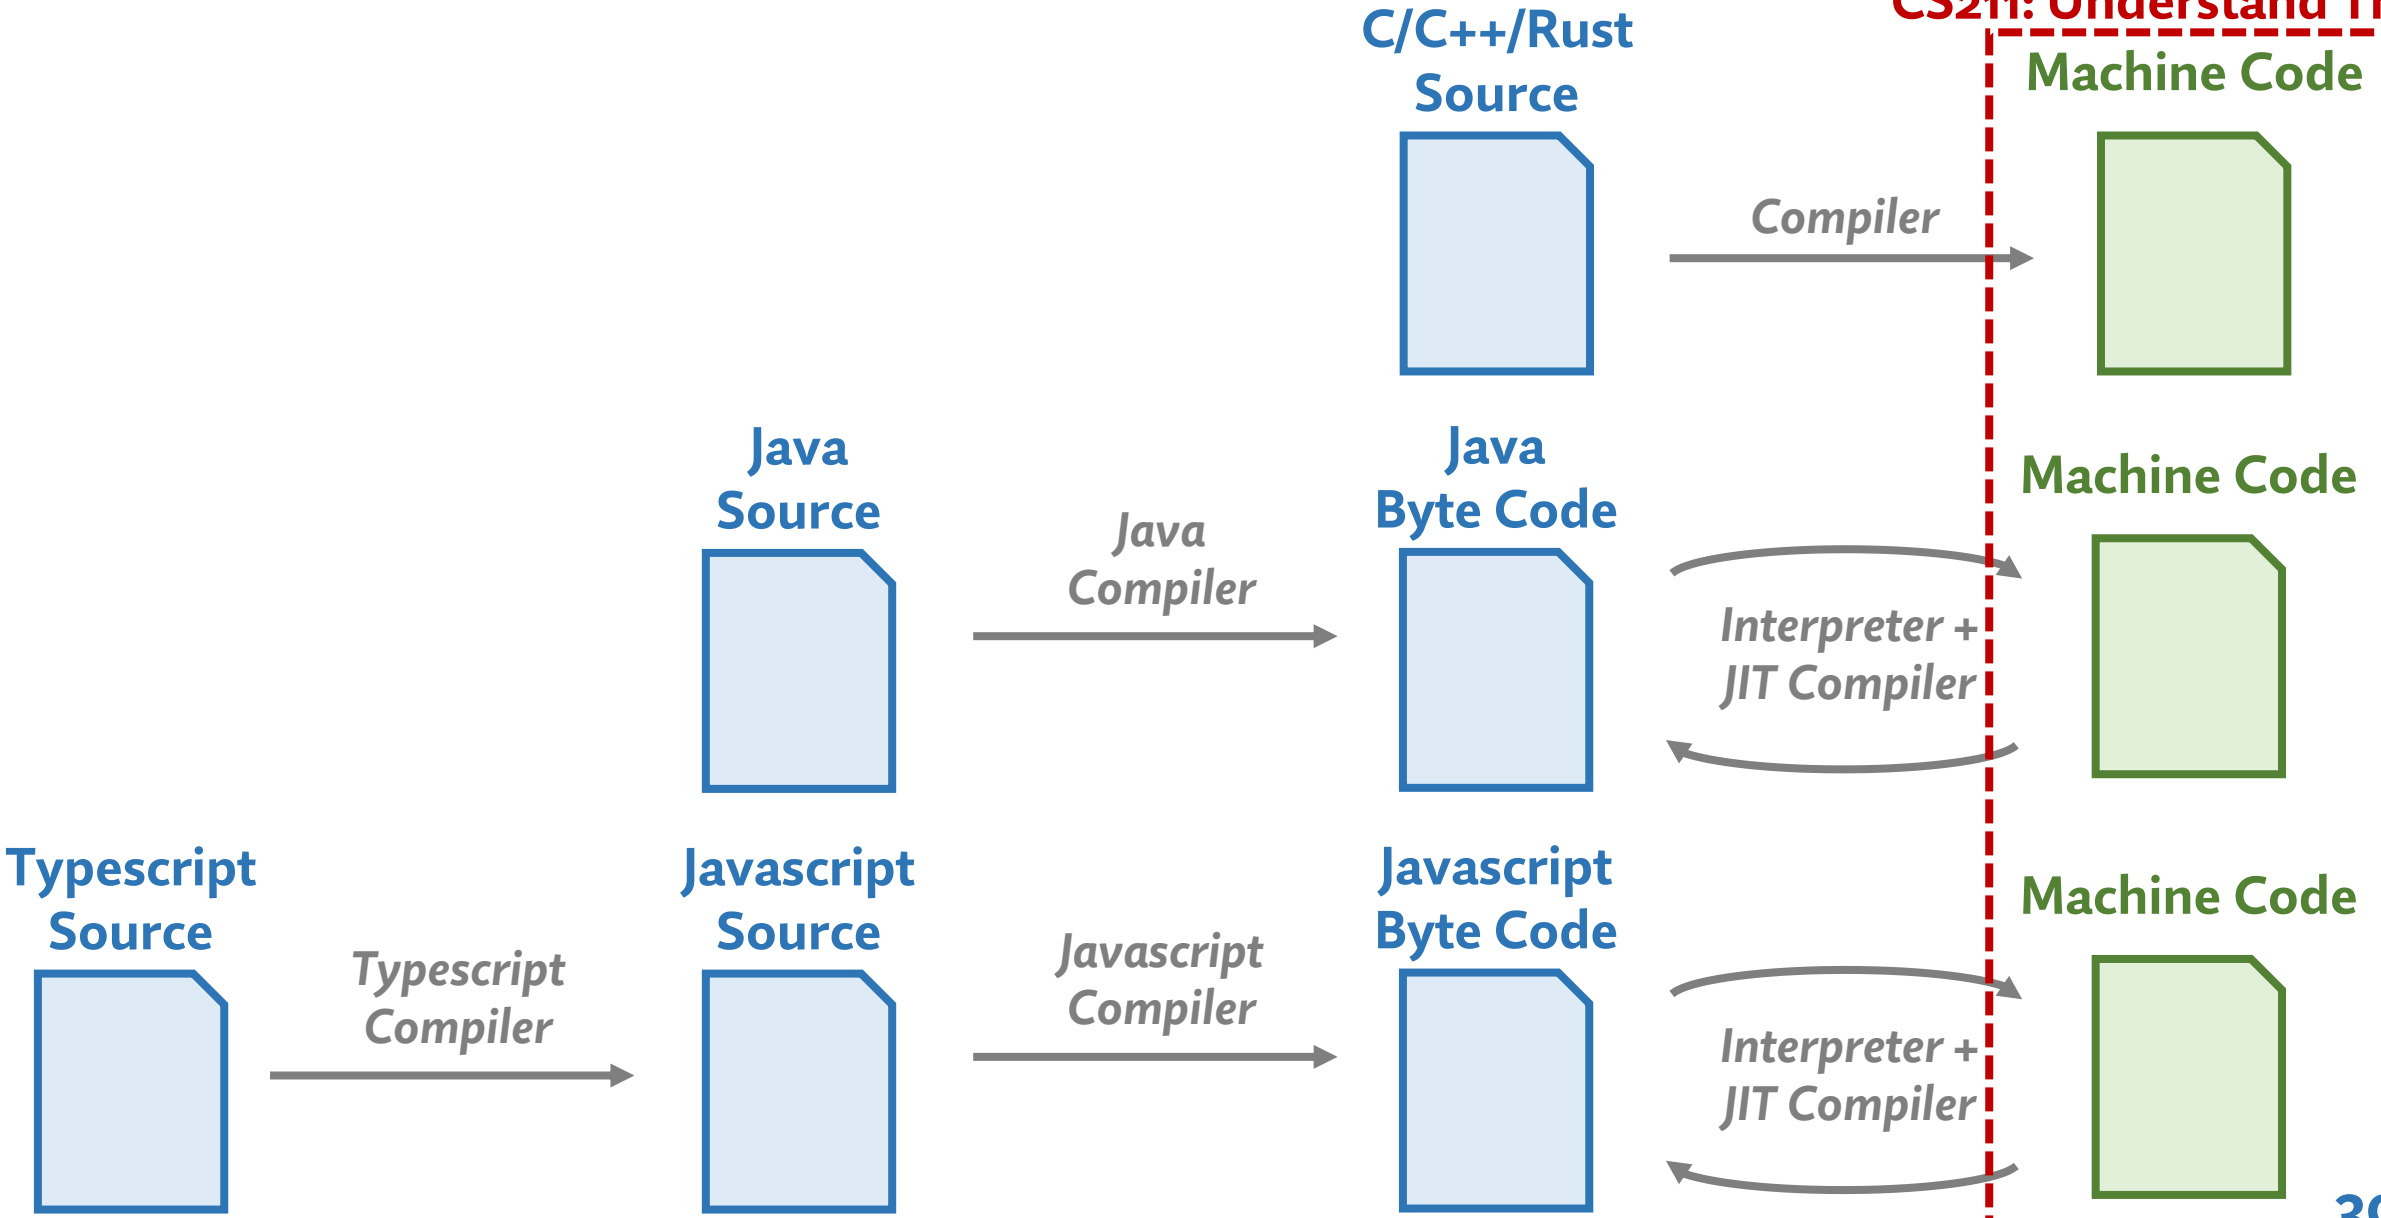
Binary File  
(Executable)

```
00001050: 4743 433a 2028 5562 756e 7475 2031 312e gcc: (Ubuntu 11.
00001060: 342e 302d 3175 6275 6e74 7531 7e32 322e 4.0-lubuntu1-22.
00001070: 3034 2920 3131 2e34 2e30 0041 3200 0000 04) 11.4.0.A2...
00001080: 7269 7363 7600 0128 0000 0005 7276 3634 riscv,...rv64
00001090: 6932 7030 5f6d 3270 305f 6132 7030 5f66 12p0_m2p0_a2p0_f
000010a0: 3270 305f 6432 7030 5f63 3270 3000 002e 2p0_d2p0_c2p0...
000010b0: 7368 7374 7274 6162 002e 696e 7465 7270 shstrtab..interp
000010c0: 002e 6e6f 7465 2e67 6e75 2e62 7369 6c64 ..note.gnu.build
000010d0: 2d69 6400 2e6e 6f74 652e 4142 492d 7461 -fd..note.ABI-ta
000010e0: 6700 2e67 6e75 2e68 6173 6800 2e64 796e g..gnu.hash..dyn
000010f0: 7379 6d00 2e64 796e 7374 7200 2e67 6e75 sym..dynstr..gnu
00001100: 2e76 6572 7369 6f6e 002e 676e 752e 7665 .version..gnu.ve
00001110: 7273 696f 6e5f 7200 2e72 656e 612e 6479 rston_f..rela.dy
00001120: 6e00 2e72 656c 612e 706c 7400 2e74 6578 n..rela.plt..tex
00001130: 7400 2e72 6f64 6174 6100 2e65 685f 6672 t..rodata..eh_fr
00001140: 616d 655f 6864 7200 2e65 685f 6672 616d ame_hdr..eh_fram
00001150: 6500 2e70 7265 696e 6974 5f61 7272 6179 e..preinit_array
00001160: 002e 696e 6974 5f61 7272 6179 002e 6669 ..init_array..fi
00001170: 6e69 5f61 7272 6179 002e 6479 6e61 6d69 n1_array..dynam
00001180: 6300 2e64 6174 6100 2e67 6f74 002e 6273 c..data..got..bs
00001190: 7300 2e63 6f6d 6d65 6e74 002e 7269 7363 s..comment..risc
000011a0: 762e 6174 7472 6962 7574 6573 0000 0000 v.attributes....
```

Machine  
Language

# Example: Translating High-Level Languages

CS211: Understand This



# Agenda

- What's Next: Layers of Abstraction
- Computer Organization Revisited
- The Representation of Code
  - **Machine Language and ISAs**
  - Choosing an ISA




# Designing a Machine Language

## Various Source Codes

Rust

Golang




C++



Python

Java


Scala



TRANSLATION



What should this be?



## Computer

Memory

Processor



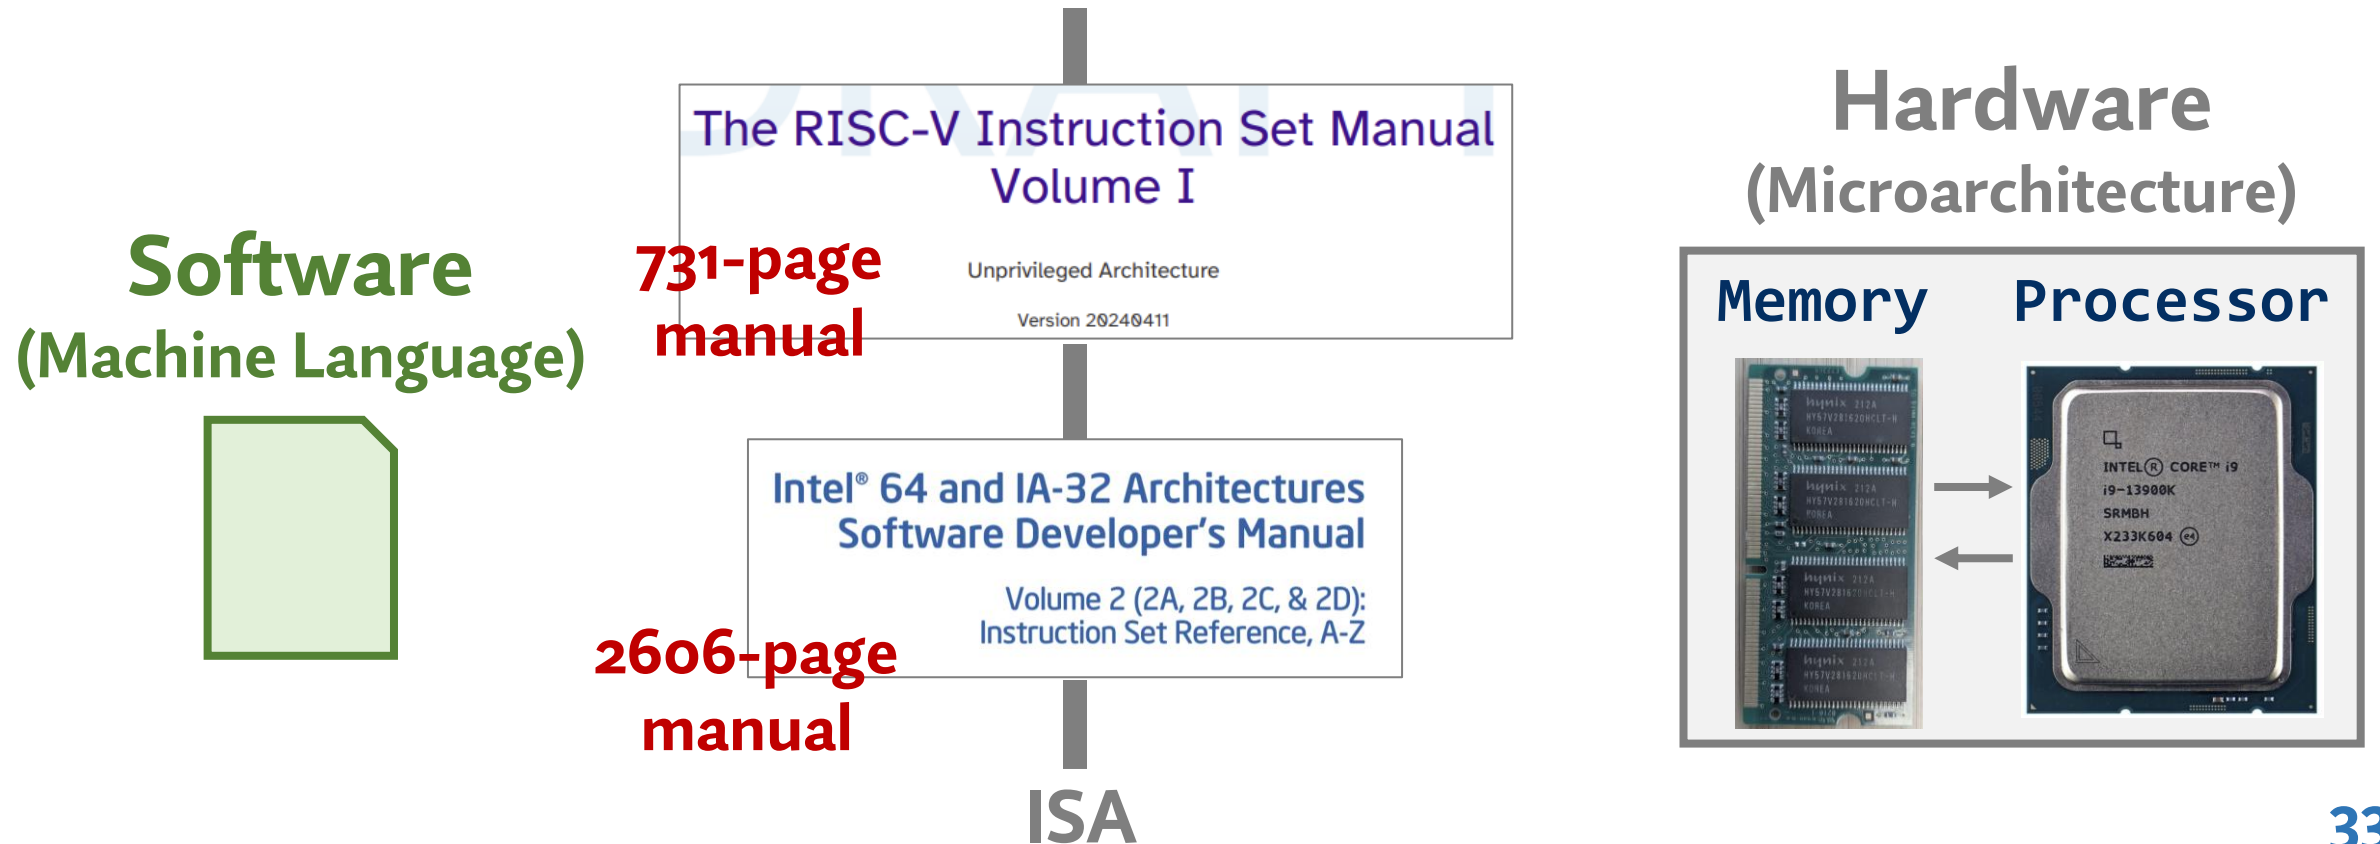
Loaded into memory at some "start address"

Executed by the processor



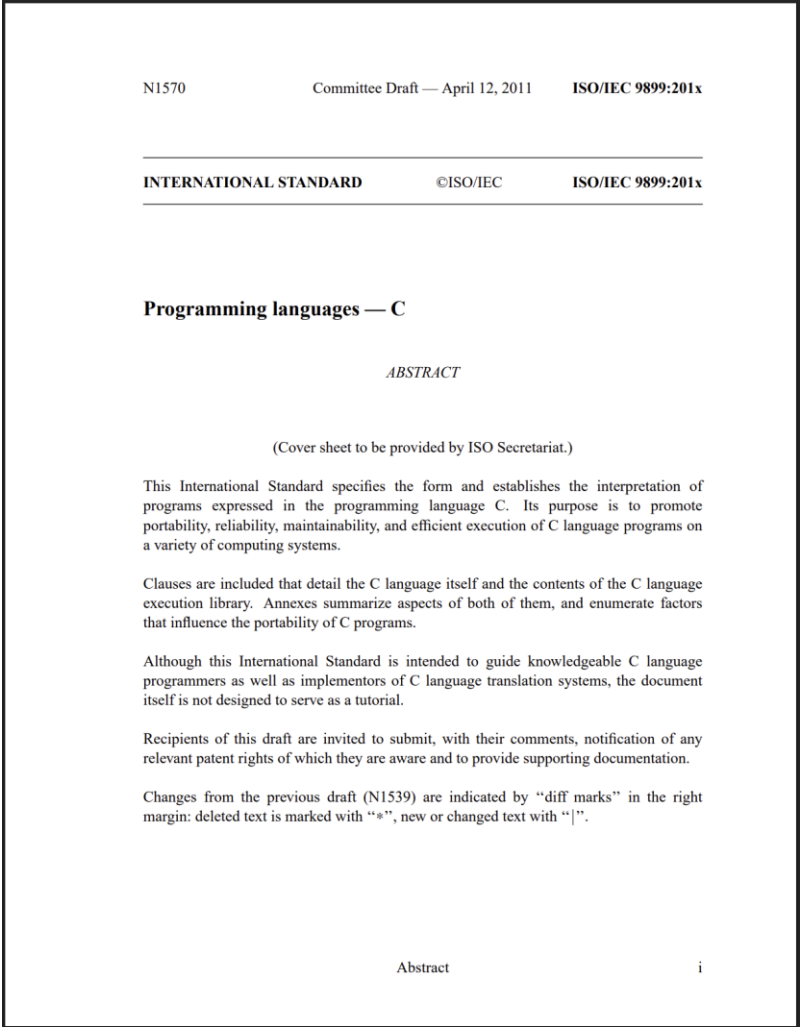
# Instruction Set Architecture: The HW-SW Interface

- **Instruction Set Architecture (ISA):** the contract between HW and SW
  - Precisely specifies the machine language
  - Abstracts away *how* the hardware implements the machine language




# Language Specifications

## C Programming Language Standard



## RISC-V ISA Manual



# What Does the ISA Specify?

- Describes:
  - All operations a given processor *can* perform
  - The *binary representations* of each operation
- Abstracts away *how* the processor does any of it
  
- Example operations:
  - **Scalar arithmetic:** Add, subtract, multiply, divide, etc.
  - **Vector arithmetic:** Dot product, element shuffling
  - **Memory access:** load/store to a given memory address
  - **Function calls:** Call and return
  - **Security primitives:** Hash functions, polynomial computations
  - **Configuration:** configure the hardware (e.g., enable/disable features)
  - ... and more

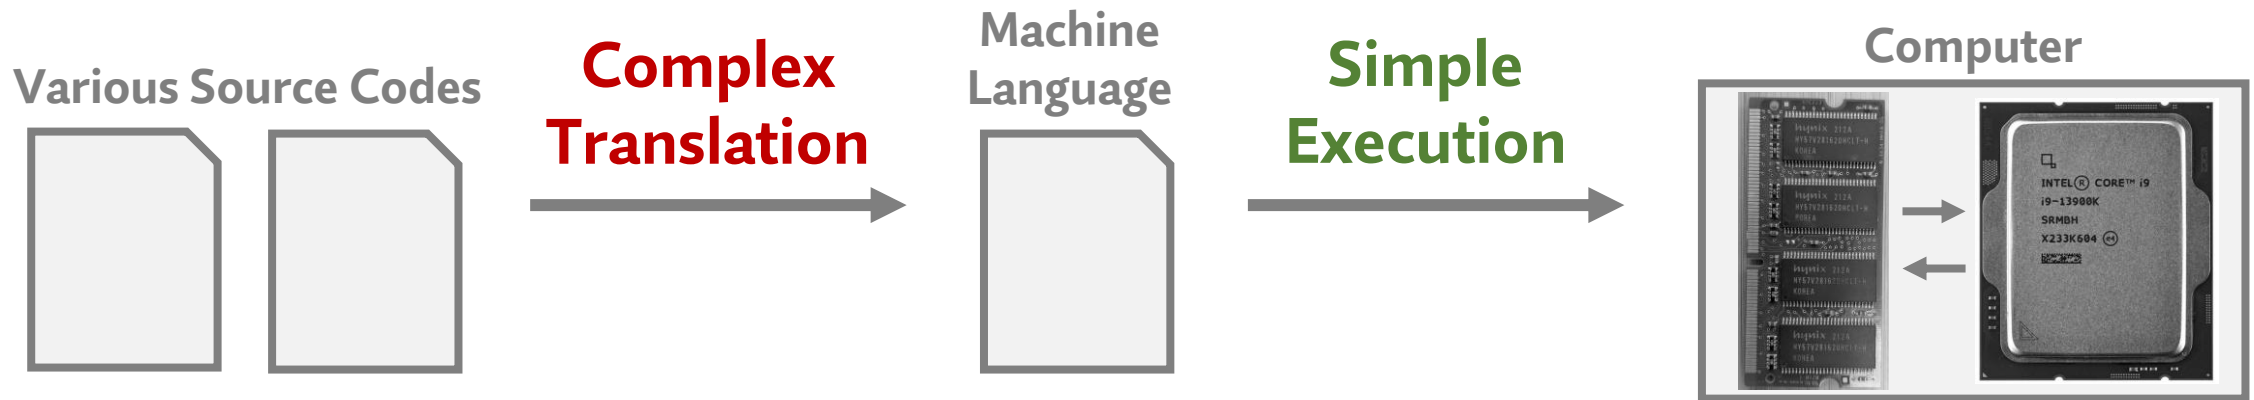
# Agenda

- What's Next: Layers of Abstraction
- Computer Organization Revisited
- The Representation of Code
  - Machine Language and ISAs
  - **Choosing an ISA**


# Choice of ISA Impacts HW/SW Design Complexity

- **There's no free lunch:** you can move complexity around but not remove it

## Simple Machine Language





## Complex Machine Language




# Choosing an ISA

Various Source Codes



TRANSLATION

Machine  
Language



EXECUTION

Computer

Memory



Processor



- **No single answer:** need to know your end-to-end system design goals
- **Complex tradeoffs** between different capabilities
  - Design complexity (software & hardware)
  - Performance optimizations
  - Debuggability
  - Ability to patch problems in the field
  - Security considerations
  - ...

# Choosing an ISA: High Level Machine Language

- **Option 1:** high-level machine language (e.g., C, Java, Fortran, etc.)

**tl;dr: it's been tried**

**Retrospective on High-Level Language Computer Architecture**

*David R. Ditzel*  
Bell Laboratories  
Computing Science Research Center  
Murray Hill, New Jersey

*David A. Patterson*  
Computer Science Division  
Department of Electrical Engineering and Computer Sciences  
University of California  
Berkeley, California

**Introduction**

High-level language computers (HLLC) have attracted interest in the architectural and programming community during the last 15 years; proposals have been made for machines directed towards the execution of various languages such as ALGOL,<sup>1,2</sup> APL,<sup>3,4,5</sup> BASIC,<sup>6,7</sup> COBOL,<sup>8,9</sup> FORTRAN,<sup>10,11</sup> LISP,<sup>12,13</sup> PASCAL,<sup>14</sup> PL/I,<sup>15,16,17</sup>

- Esoteric: Aesthetics or no stated advantages.

An almost universal justification for high-level language comput-

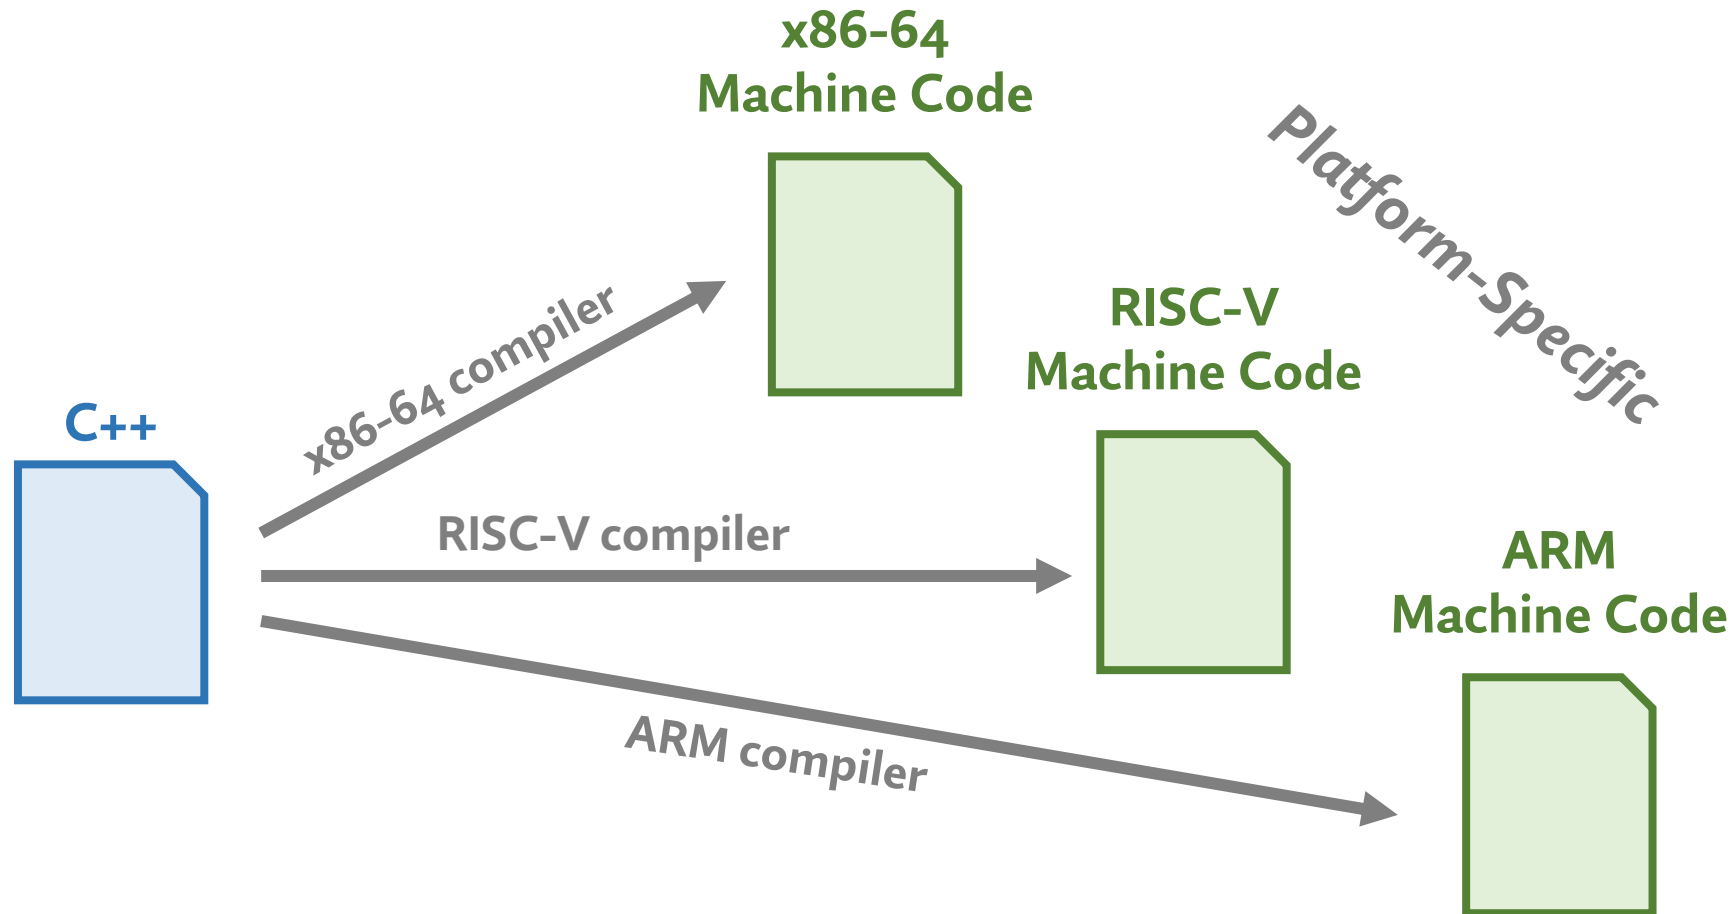
# High-Level Machine Languages

- **Older (1965 – 1980):** Fortran, Lisp, Pascal, Algol, and many others
- **ARMv5 (1998):** “Java mode” side-by-side with ARM mode
- Vision:
  - **Simple, cheap software** (e.g., simple compilers, software debugging tools)
- Reality:
  - **Moves complexity** to the hardware: not simpler or cheaper!
  - **No real difference** (or advantage) to application-level programmers
  - Often gives software **less control** over the processor and memory



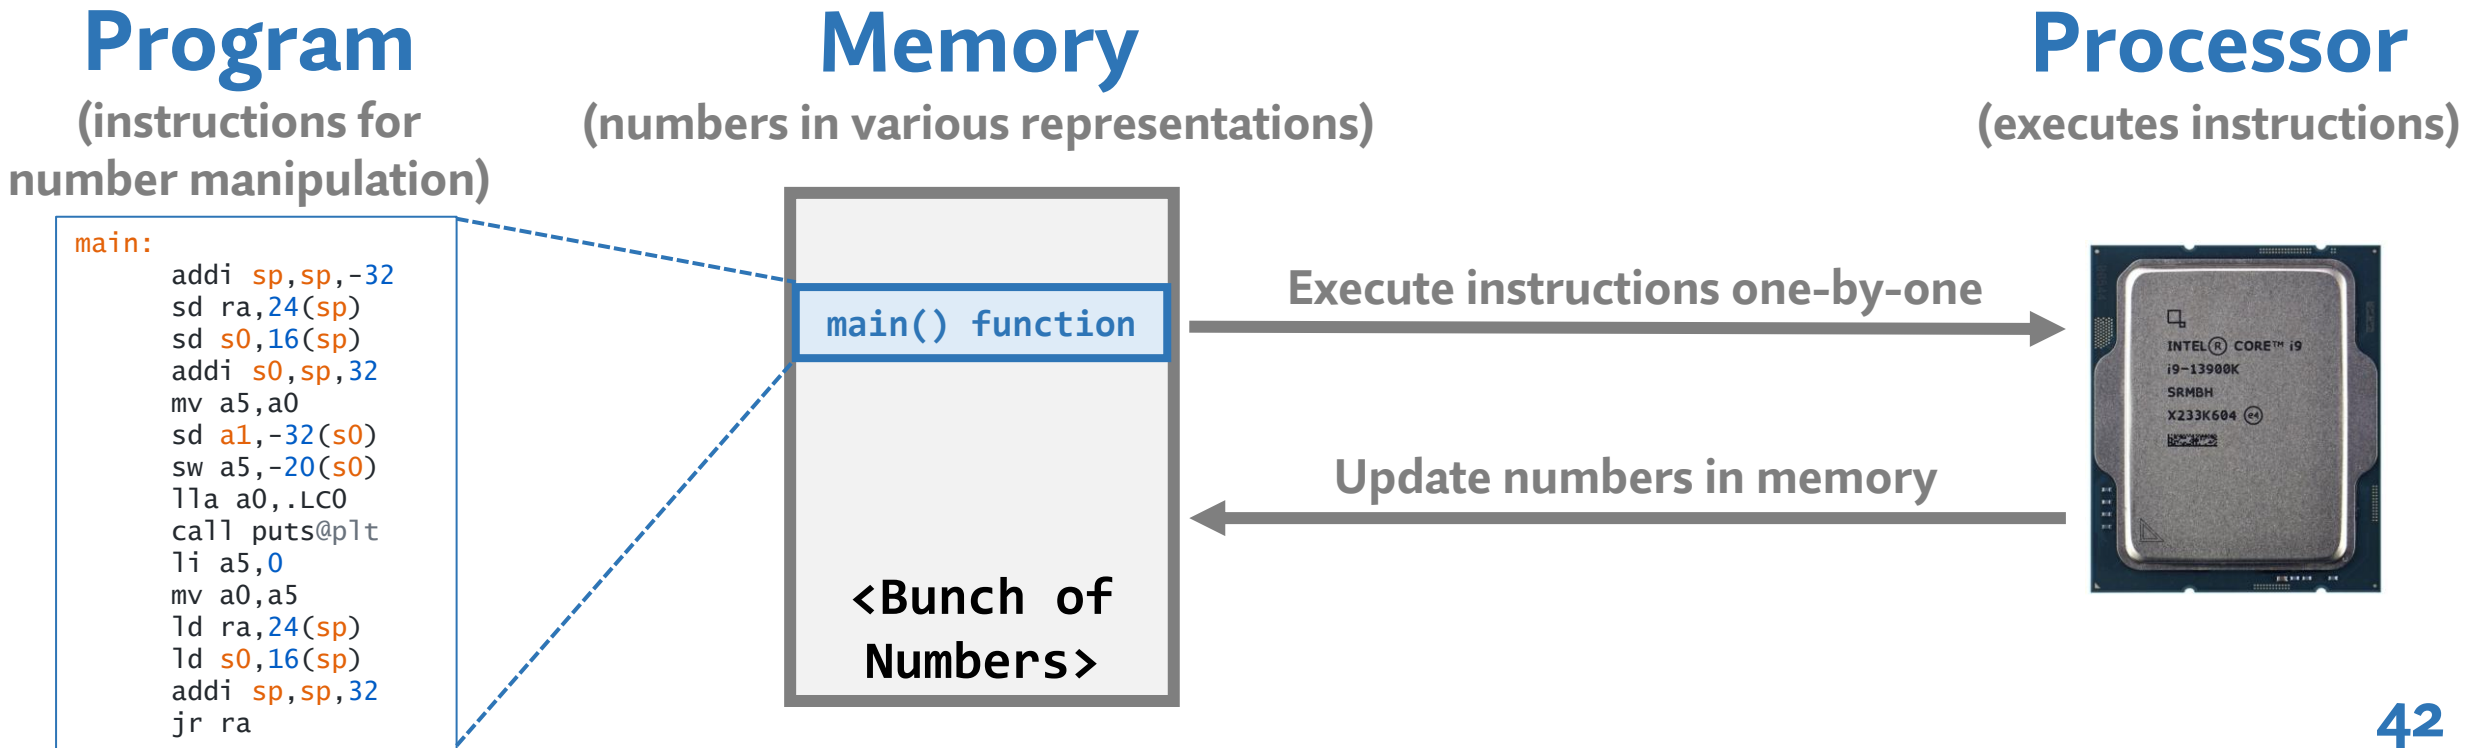
# Choosing an ISA: Low Level Machine Language

- Break things down to **small primitive operations** (i.e., “instructions”)
- Let compilers **translate** from HLLs to instructions



# Modern ISAs

- Today's systems almost all use some variant of low-level instructions
  - All data and instructions are represented as **numbers in memory**
  - The processor **executes instructions** out of memory one-by-one
  - Every instruction updates the **numbers in memory**



# Example Instruction: RISC-V NOP


- A “No Operation (NOP)” instruction does **nothing**
  - Essential for debugging, performance, alignment, etc.

## Functional Description of a NOP Instruction

The NOP instruction does not change any architecturally visible state, except for advancing the pc and incrementing any applicable performance counters. NOP is encoded as ADDI x0, x0, 0.

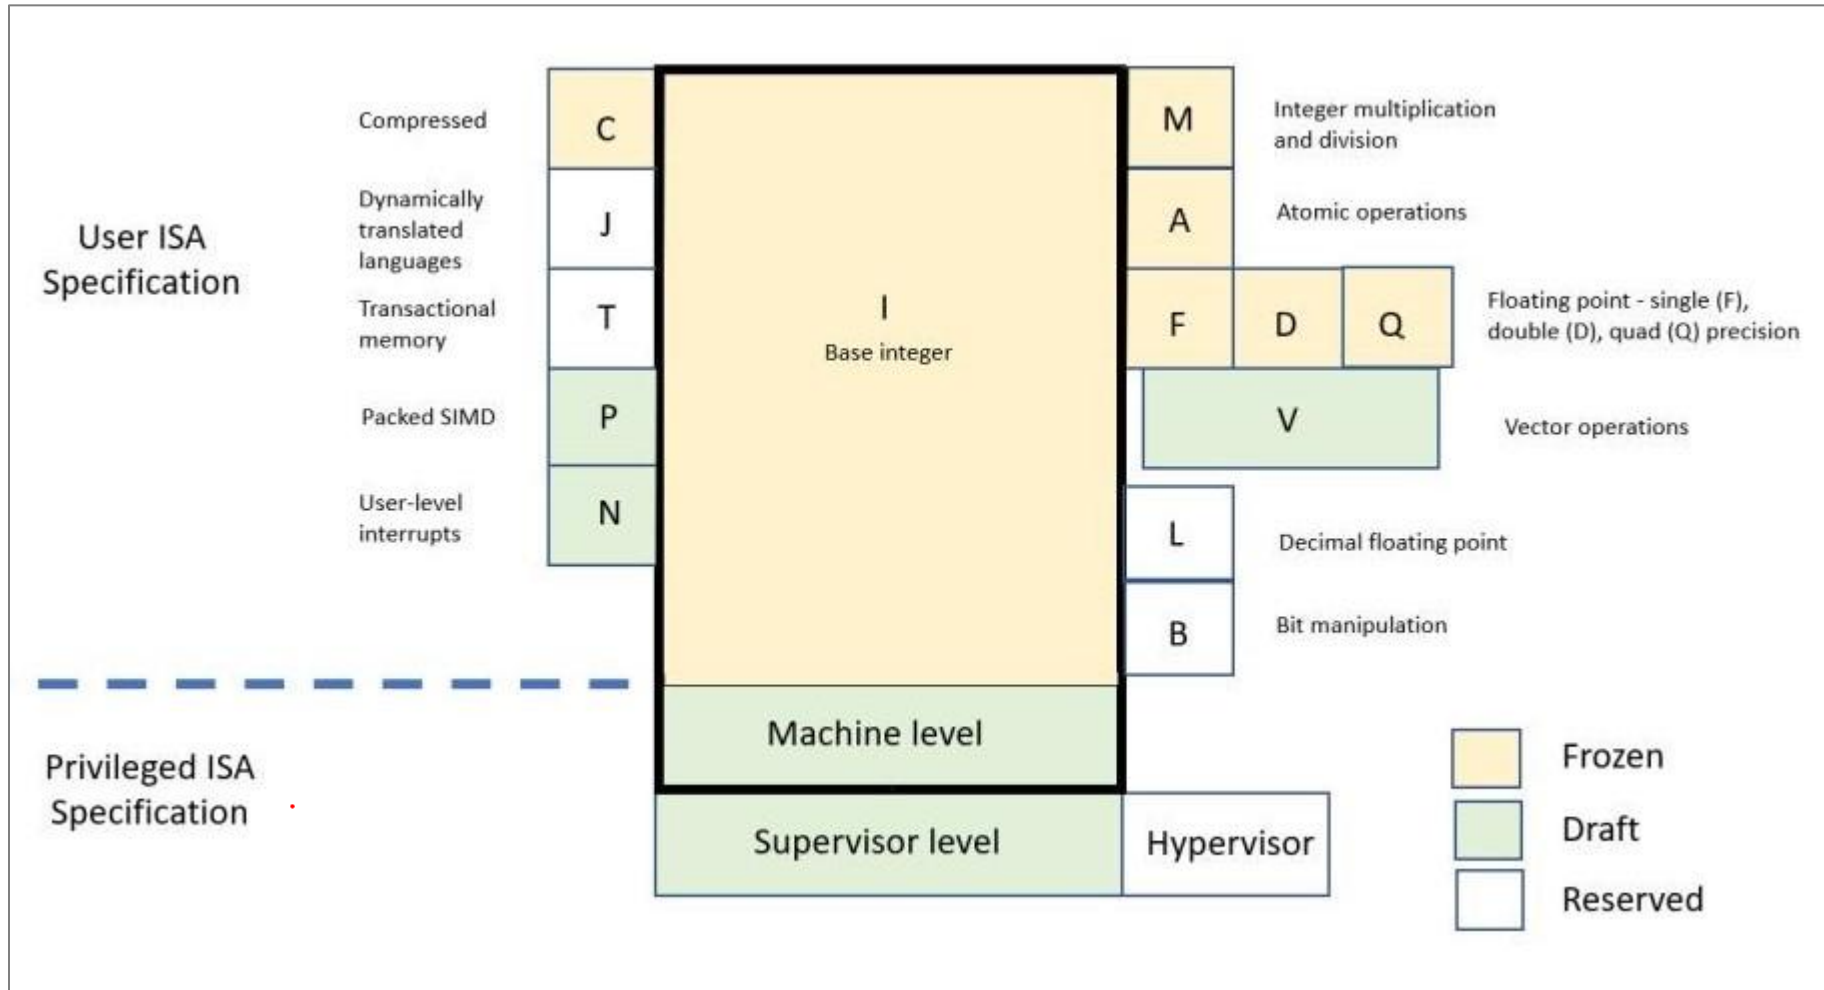
## Binary Representation of a NOP Instruction (32 Bits)

### 2.4.3. NOP Instruction



# RISC-V

- Good middle ground for this course: simple, yet realistic



<https://www.microcontrollertips.com/wp-content/uploads/2020/11/RISC-V-ISA-block-diagram.jpg>

# **CS 211: Intro to Computer Architecture**

## ***9.1: The Hardware-Software Interface***

**Minesh Patel**

Spring 2025 – Tuesday 25 March