



EMBEDDED HACKING

FIRST EDITION 1.0001 IN-DEVELOPMENT-ALPHA

Copyright © 2025 Kevin Thomas

Forward

I remember when I started learning programming to which my first language was 6502 Assembler to allow me to program a Commodore 64 and right from the beginning of my journey, I learned the lowest level development possible.

Literally every piece of the Commodore 64 was understood as it was a simple machine. There was absolutely no abstraction layer of any kind.

We had an absolute mastery of everything however it was a very simple architecture.

Microcontrollers are small systems without an operating system and are also very simple in their design. They are literally everywhere from your toaster to your fridge to your TV and billions of other electronics that you never think about.

Most microcontrollers are developed in the C programming language which has its roots to the 1970's however dominates the landscape.

We will take our time and learn the basics of C utilizing a Pico 2 microcontroller.

Below are items you will need for this course.

Raspberry Pi Pico 2 w/ Header

<https://www.pishop.us/product/raspberry-pi-pico-2-with-header>

USB A-Male to USB Micro-B Cable

<https://www.pishop.us/product/usb-a-male-to-usb-micro-b-cable-6-inches>

Raspberry Pi Pico Debug Probe

<https://www.pishop.us/product/raspberry-pi-debug-probe>

Complete Component Kit for Raspberry Pi

<https://www.pishop.us/product/complete-component-kit-for-raspberry-pi>

10pc 25v 1000uF Capacitor

<https://www.amazon.com/Cionyce-Capacitor-Electrolytic-Capacitors-Microwave/dp/B0B63CCQ2N?th=1>

NOTE: The item links may NOT be available, but the descriptions allow you to shop on any online or physical store of your choosing.

Let's begin...


Table Of Contents

Chapter 1:	hello, world
Chapter 2:	Debugging hello, world
Chapter 3:	Hacking hello, world
Chapter 4:	Embedded System Analysis
Chapter 5:	Intro To Variables
Chapter 6:	Debugging Intro To Variables
Chapter 7:	Hacking Intro To Variables
Chapter 8:	Uninitialized Variables
Chapter 9:	Debugging Uninitialized Variables
Chapter 10:	Hacking Uninitialized Variables
Chapter 11:	Integer Data Type
Chapter 12:	Debugging Integer Data Type
Chapter 13:	Hacking Integer Data Type
Chapter 14:	Floating-Point Data Type
Chapter 15:	Debugging Floating-Point Data Type
Chapter 16:	Hacking Floating-Point Data Type
Chapter 17:	Double Floating-Point Data Type
Chapter 18:	Debugging Double Floating-Point Data Type
Chapter 19:	Hacking Double Floating-Point Data Type
Chapter 20:	Static Variables
Chapter 21:	Debugging Static Variables
Chapter 22:	Hacking Static Variables
Chapter 23:	Constants
Chapter 24:	Debugging Constants
Chapter 25:	Hacking Constants
Chapter 26:	Operators
Chapter 27:	Debugging Operators
Chapter 28:	Hacking Operators
Chapter 29:	Static Conditionals
Chapter 30:	Debugging Static Conditionals
Chapter 31:	Hacking Static Conditionals
Chapter 32:	Dynamic Conditionals
Chapter 33:	Debugging Dynamic Conditionals
Chapter 34:	Hacking Dynamic Conditionals
Chapter 35:	Functions, w/o Param, w/o Return
Chapter 36:	Debugging Functions, w/o Param, w/o Return
Chapter 37:	Hacking Functions, w/o Param, w/o Return
Chapter 38:	Functions, w/ Param, w/ Return
Chapter 39:	Debugging Functions, w/ Param, w/ Return
Chapter 40:	Hacking Functions, w/ Param, w/ Return

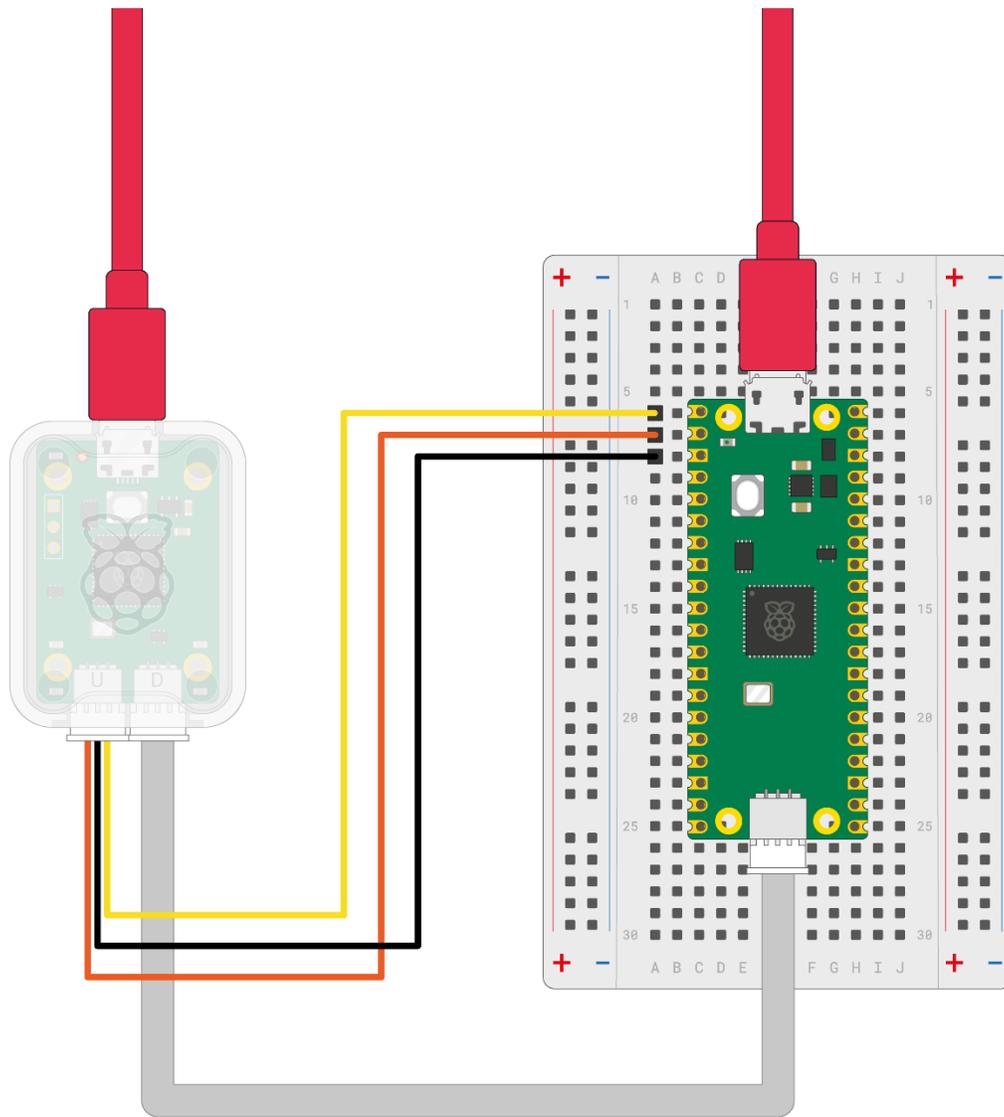
Chapter 1: hello, world

We begin our journey building the traditional *hello, world* example in Embedded C.

Below we see our diagrams for our breadboard schematic which includes our Pico 2 microcontroller and the Pico debug probe.



fritzing



To setup our development environment, we will download VS Code.

<https://code.visualstudio.com/download>

Once VS Code is installed, we will install the Raspberry Pi Pico VS Code extension.

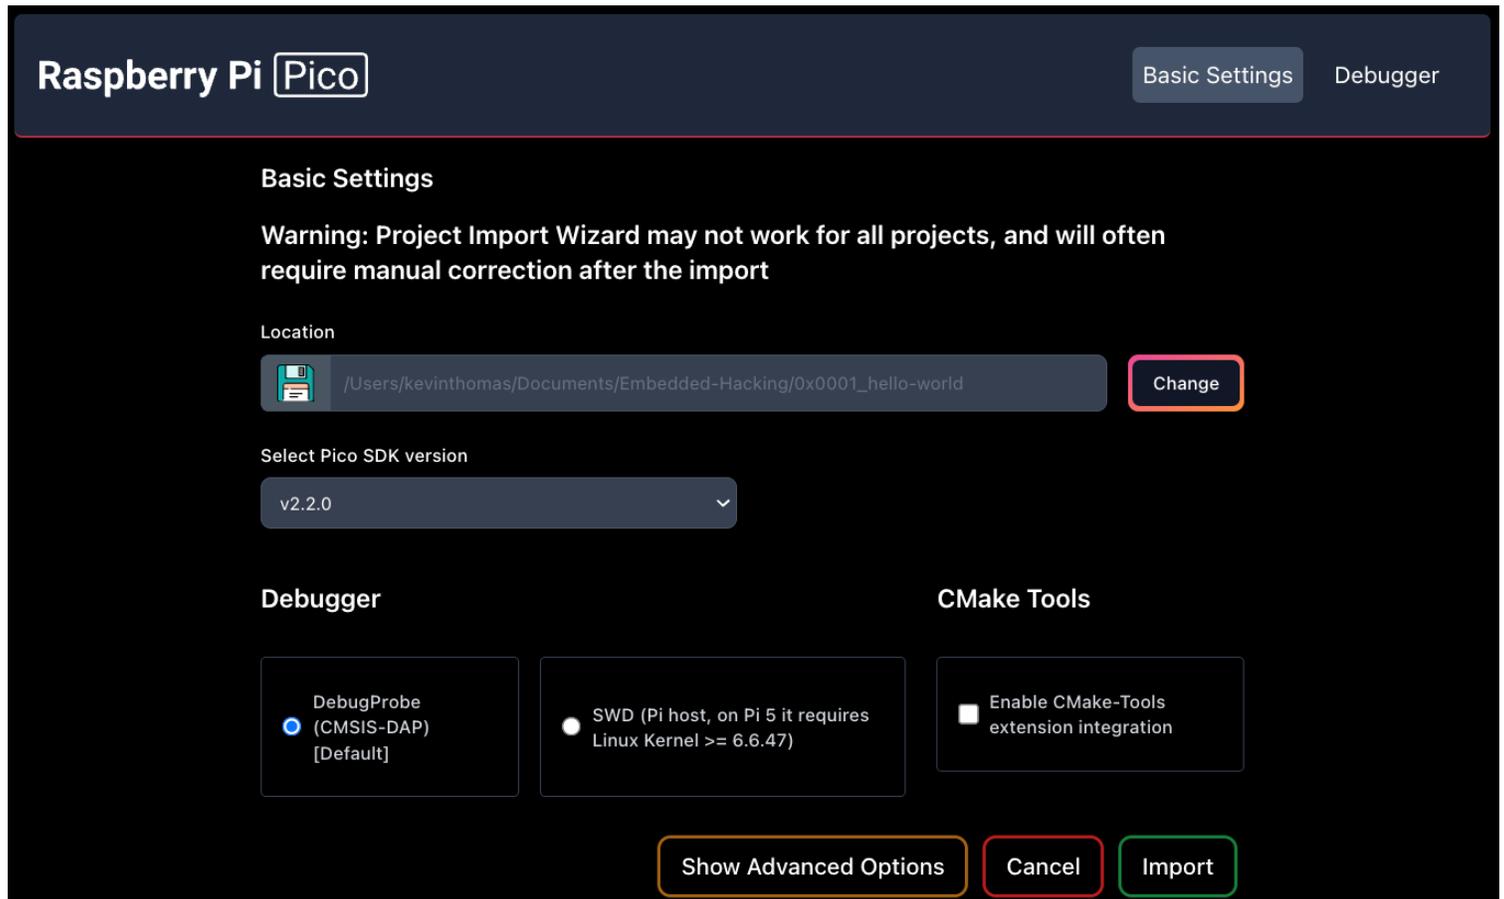
<https://marketplace.visualstudio.com/items?itemName=raspberry-pi.raspberry-pi-pico>

We will setup the Raspberry Pi Pico Debug Probe as there are detailed instructions below as well to get started.

<https://www.raspberrypi.com/documentation/microcontrollers/debug-probe.html>

<https://www.raspberrypi.com/documentation/microcontrollers/images/pico-2-r4-pinout.svg>

If you do not have Git installed, here is a link to install git on Windows, MAC and Linux.



<https://git-scm.com/book/en/v2/Getting-Started-Installing-Git>

We need to clone our course repo to whatever folder you prefer.

```
git clone https://github.com/mytechtalent/Embedded-Hacking.git
```

Open VS Code and click **File** then **Open Folder** then click on the **Embedded-Hacking** folder and then select **0x0001_hello-world**.

This may pop up a screen asking to import the project. Once visible, click **Import**, otherwise just continue.

Now we are ready to compile and flash our code onto the Pico.

You can click on **Compile** and then **Run** in the bottom right-hand side of the VS Code editor assuming you have your Pico 2 plugged in.

Press and hold the push button we attached to the breadboard while pressing the white BOOSEL button on

the Pico 2; then release the white BOOTSEL button on the Pico 2 and then release the push button we attached to the breadboard.

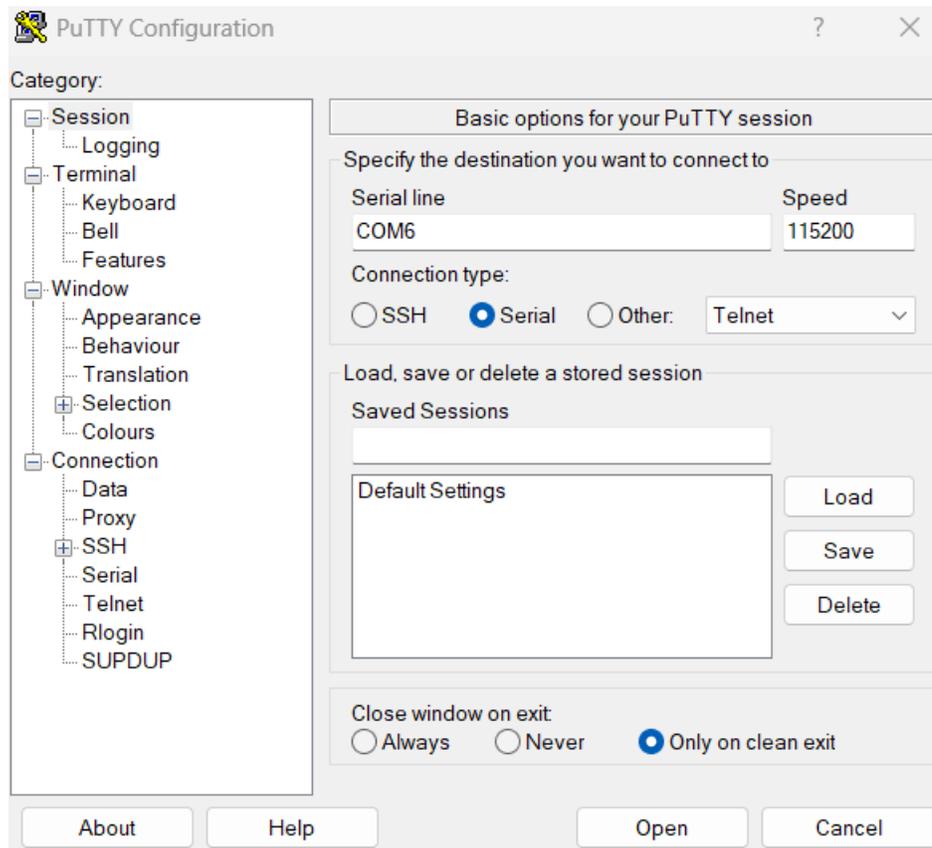
If the **Compile** and **Run** buttons within VS Code does not work, you can also open a file explorer window to copy our **0x0001_hello-world.uf2** firmware into the **RPI-RP2** drive.

We need to download a serial monitor to interact with our Pico. If you are on Windows download PuTTY as the link is below.

<https://www.putty.org>

If you are on Windows, you can open the Device Manager and look for the COM port that will be used to connect PuTTY to. There are at minimum two ports one for the Pico 2 UART and the other for the Pico Debug Probe. Try both and one of them will be UART that we are looking for.





The next step is to run PuTTY.

You want to type in your COM port, in my case COM6, and click the **Open** button.

If you are on MAC or Linux, you can use the screen program.

```
ls /dev/tty.  
screen /dev/tty.XXX 115200
```

Now let's review our **0x0001_hello-world.c** file as this is located within the main folder.

```
#include <stdio.h>  
#include "pico/stdlib.h"  
  
int main(void) {  
    stdio_init_all();  
  
    while (true)  
        printf("hello, world\r\n");  
}
```

Let's break down this code.

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

This line includes the `stdio.h` header file, which contains declarations for standard input and output functions.

```
#include "pico/stdlib.h"
```

This line includes the `pico/stdlib.h` header file, which contains declarations for various Raspberry Pi Pico standard library functions.

```
int main(void)
```

The above line declares the main function, which is the entry point for all C and Python programs.

```
stdio_init_all();
```

This line initializes the standard input and output system.

```
while (true)
```

This line starts a while loop that will run forever.

```
printf("hello, world\r\n");
```

This line prints the message, *hello, world*, to the console.

Open the terminal to see, *hello, world*, as expected being printed over and over again.

Chapter 2: Debugging hello, world

Today we debug!

There are two main types of reverse engineering: static and dynamic. Static reverse engineering involves examining the binary without executing it. Tools like Ghidra allow you to inspect raw assembly instructions, control flow, and code structure. Dynamic reverse engineering, on the other hand, involves running the binary and observing its behavior in real time. With tools like GDB, you can monitor memory changes, register values, and execution paths as the program runs.

We will download Ghidra, a free static disassembler from the NSA at the link below.

<https://github.com/NationalSecurityAgency/ghidra/releases>

If you are using Windows, we will move the Ghidra folder to the **C:** drive and make sure to update the path accordingly. If you are on MAC or Linux, move to the root of your drive as well and update version in path.

`C:\ghidra_11.4.2_PUBLIC`

Please download and install the proper Java version based on your system.

<https://adoptium.net/temurin/releases>

Once complete, a file called **ghidraRun** will be created. To launch Ghidra, execute this file. If you're on Windows, be sure to run the batch file version.

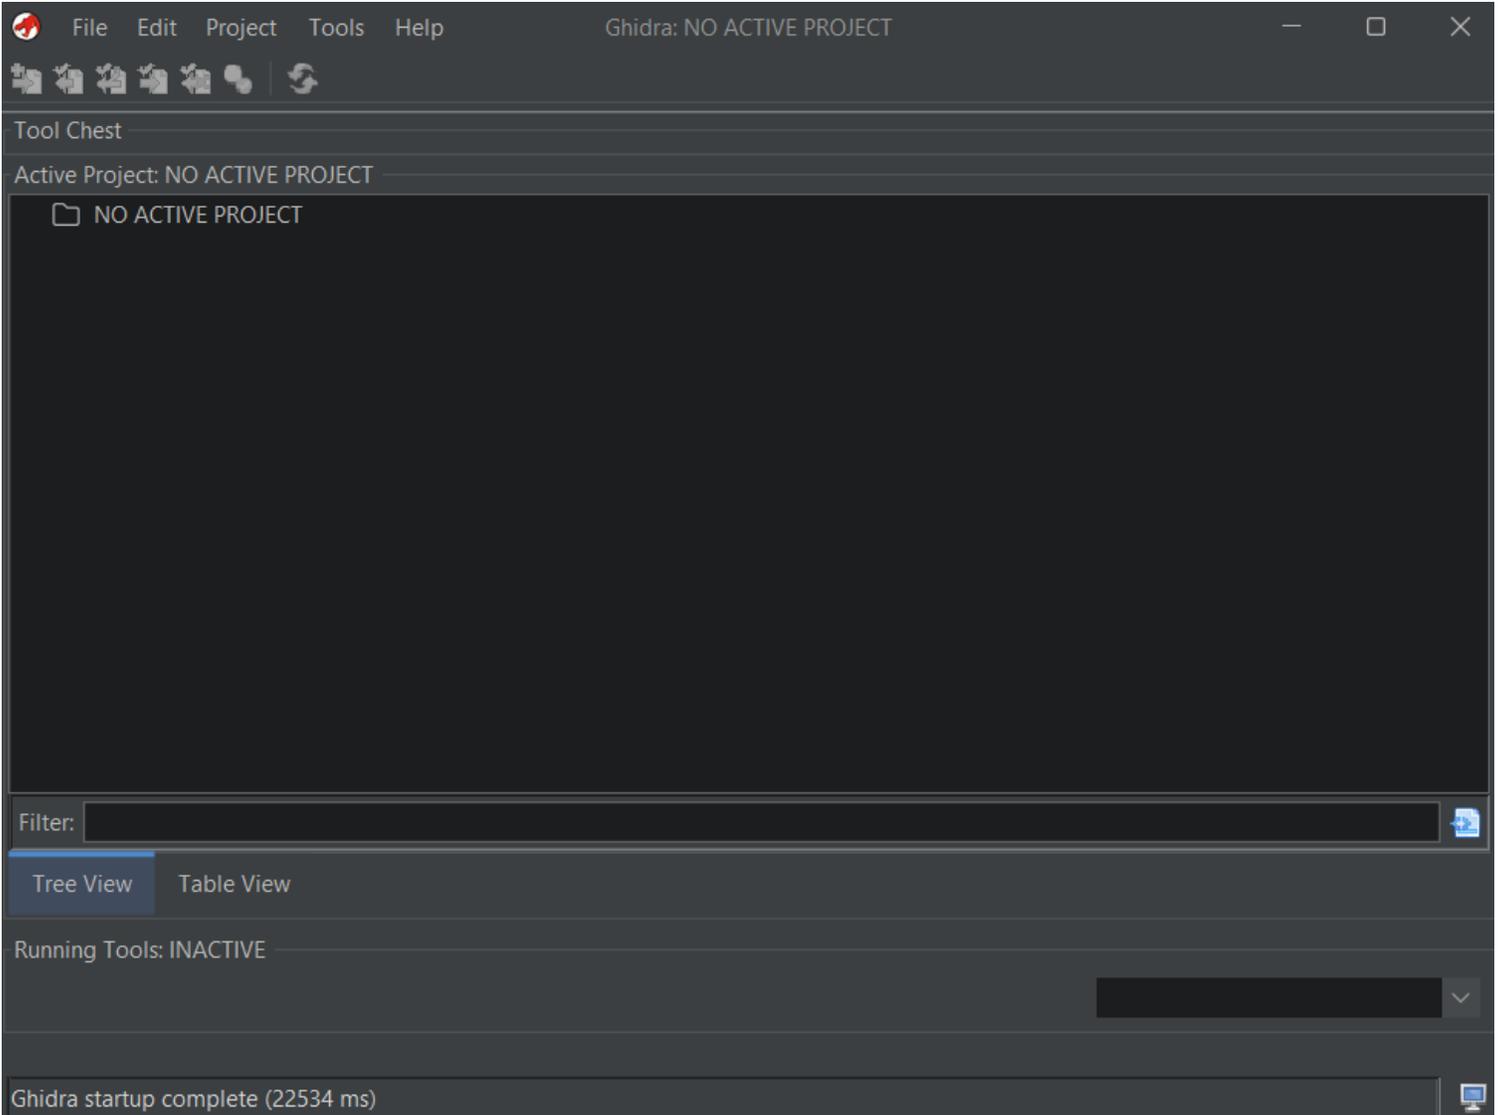
A window will appear where we will select **File, New Project, Non-Shared Project, Next**, and create a **Project Name**. Here we will call it **0x0001_hello-world** and press **Finish**.


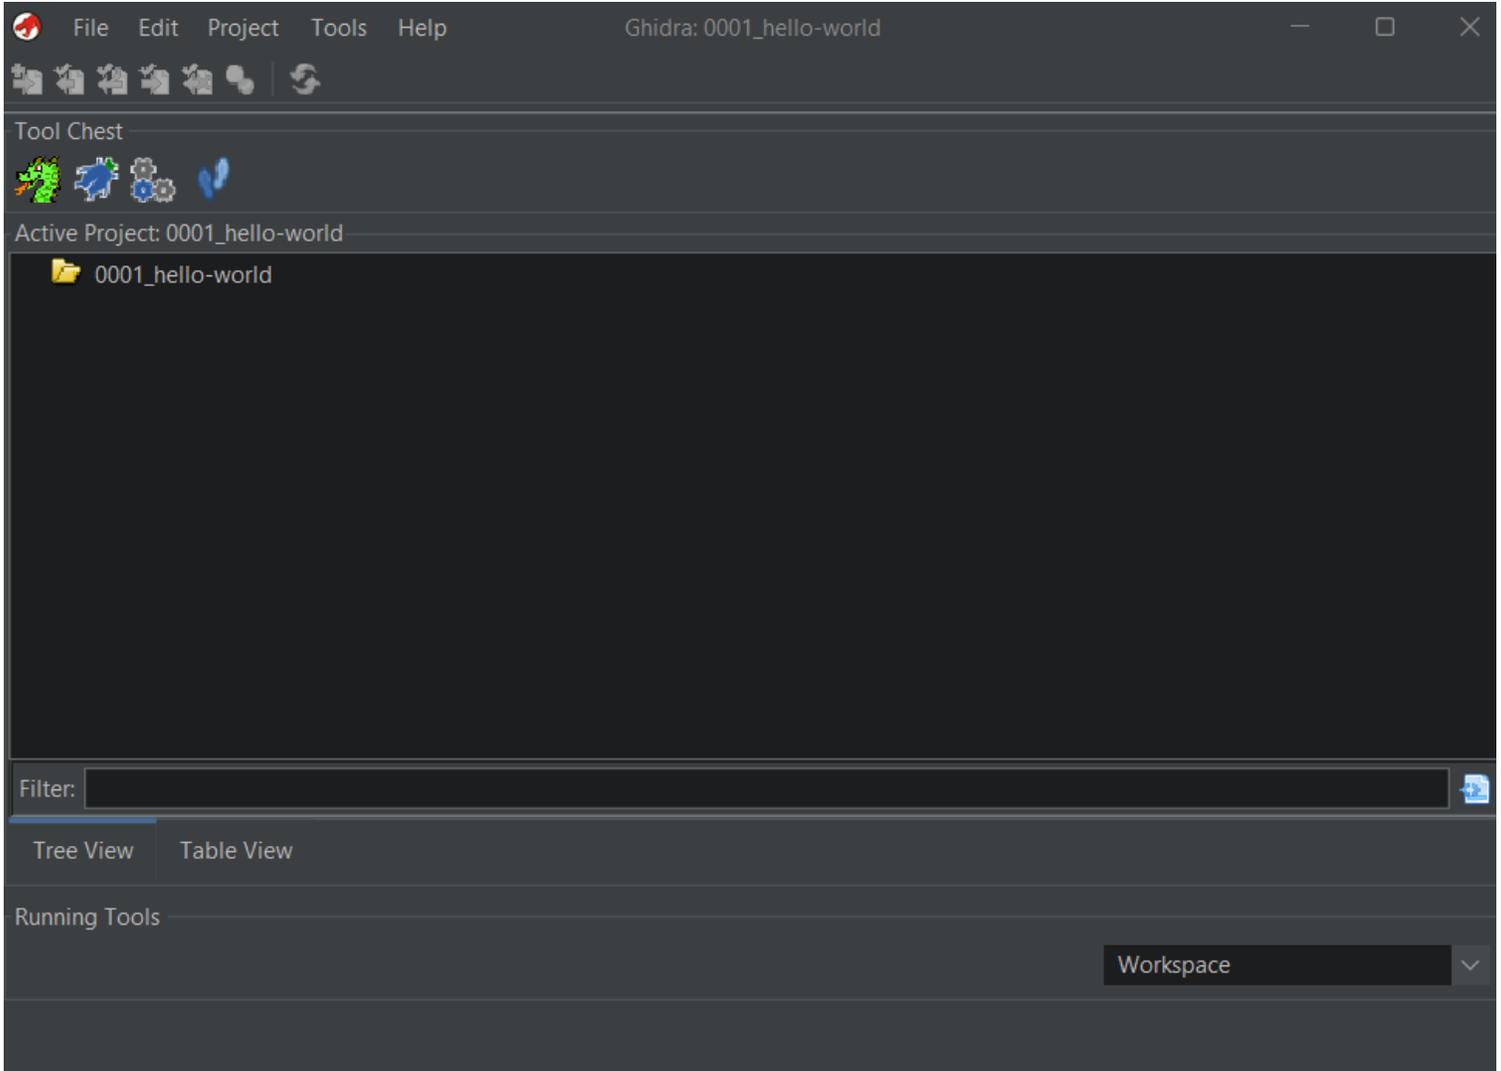
Open the file explorer and navigate to the **Embedded-Hacking** folder and drag-and-drop the **0x0001_hello-world.elf** file into the folder within the Ghidra application panel.

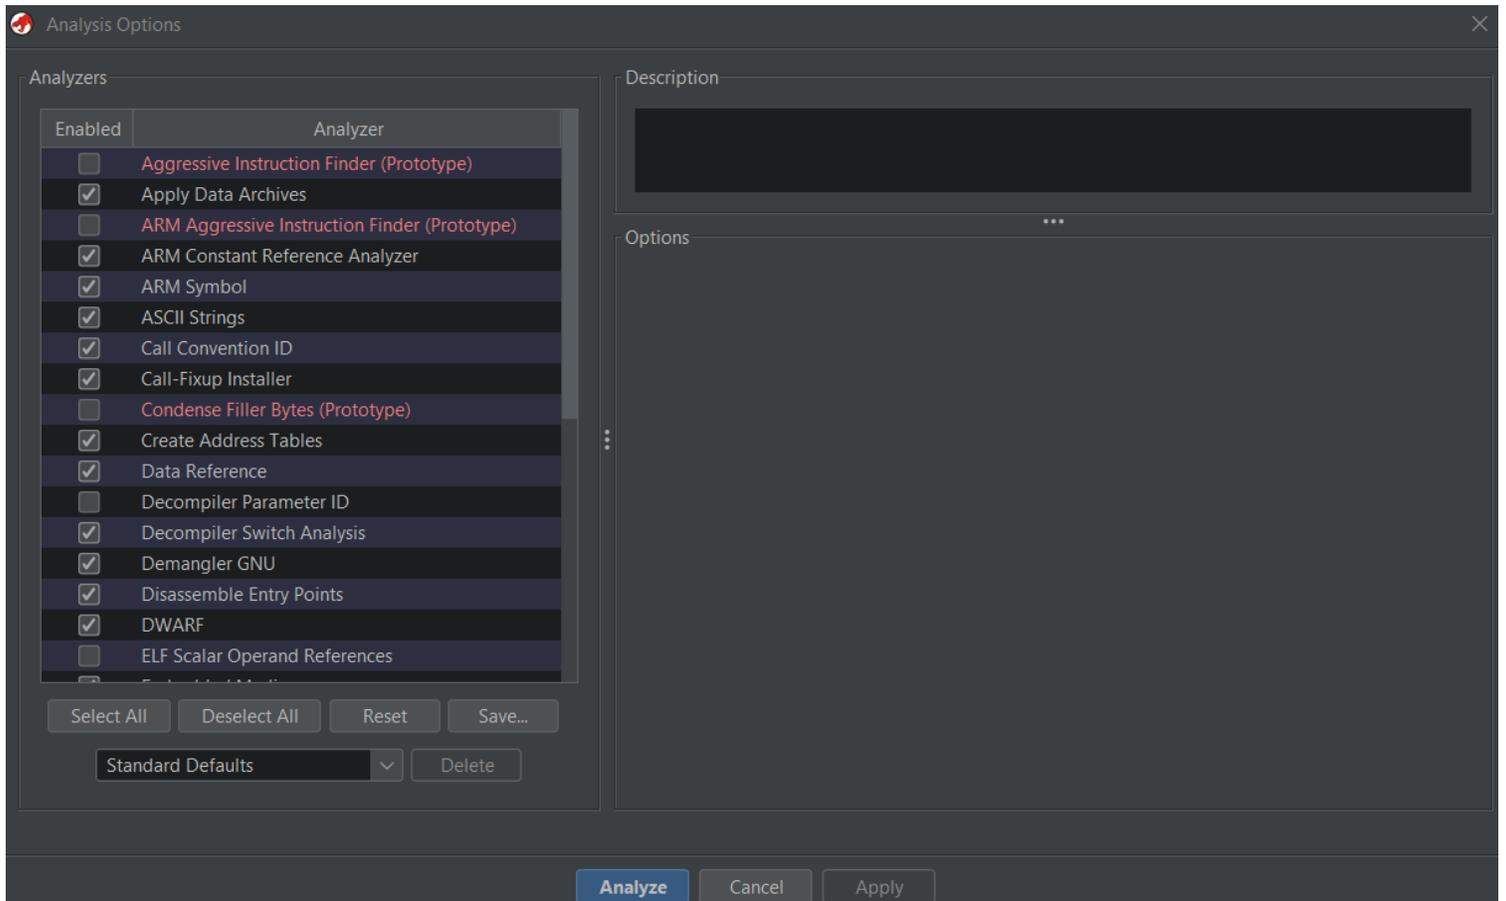
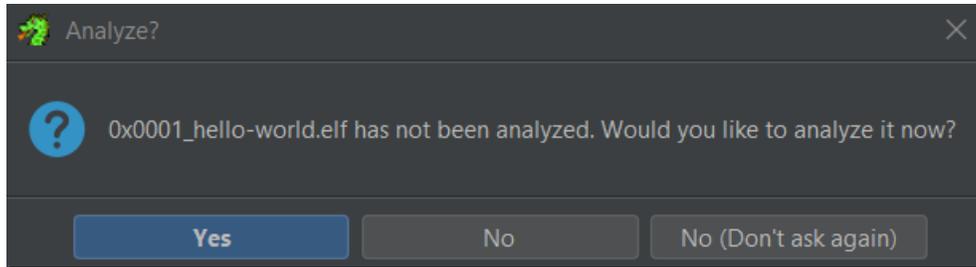
In the small window that appears, you will see the file identified as an ELF, which stands for Executable and Linkable Format. This format includes symbols that make reverse engineering easier. In future chapters, we will work with stripped binaries that do not contain these symbols.

At this point, click **Ok** and then double-click on the file within the window.

Finally click the auto-analyze and let's begin reviewing the binary.







```

*****
*
*                               FUNCTION                               ...
*****
int main(void)
    assume LRset = 0x0
    assume TMode = 0x1
int      r0:4      <RETURN>
main          XREF[3]:  Entry Point(*),
                _reset_handler:1000018c(c),
                .debug_frame::00000018(*)
0x0001_hello-world.c:4 (2)
0x0001_hello-world.c:5 (2)
10000234 08 b5      push      {r3,lr}
                0x0001_hello-world.c:5 (4)
10000236 01 f0 99 f9  bl      _Bool stdio_init_all(void)
                LAB_1000023a          XREF[1]:  10000240(j)
                0x0001_hello-world.c:7 (6)
                0x0001_hello-world.c:8 (6)
1000023a 02 48      ldr      r0=>__EH_FRAME_BEGIN__, [DAT_10000244] = "hello, world\r"
                = 100019CCh
1000023c 01 f0 de f9  bl      __wrap_puts      int __wrap_puts(char * s)
                0x0001_hello-world.c:7 (8)
10000240 fb e7      b      LAB_1000023a
10000242 00      ??      00h
10000243 bf      ??      BFh

```

```

1
2 /* WARNING: Unknown calling convention */
3
4 int main(void)
5
6 {
7     stdio_init_all();
8     do {
9         __wrap_puts("hello, world\r");
10    } while( true );
11 }
12

```

I have held off on exploring the deeper meaning behind all of this because our first goal is to establish a solid static reverse engineering workflow.

Now we can see our main function displayed in raw assembly, a decompiled view, and a pseudo source code window.

One of the first differences we notice is that our original source used a while true loop, but the decompiled output shows a do while loop. This is not a major issue, as the logic is still clear and we can see the code echoing *hello, world* to the terminal.

In our original source, we used the `printf` function. After compilation, the compiler optimized this and replaced it with the `puts` function, which is a common substitution for simple output.

At this point, I am going to pause on reviewing the assembly and shift focus to setting up GDB. This will allow us to begin dynamic reverse engineering, along with a basic introduction to the ARM architecture we are working with.

To enable dynamic reverse engineering capabilities, we will download the GNU ARM toolchain tailored to our embedded architecture. Be sure to select the version appropriate for your system.

<https://developer.arm.com/downloads/-/arm-gnu-toolchain-downloads>

The next step is to download OpenOCD. If you are on Windows, there are pre-build binaries at the location below.

<https://gnutoolchains.com/arm-eabi/openocd>

If you are on Windows, the next step is to extract the folder to your `C:\` drive and update your path to include the following directories and keep in mind the version you downloaded as you may need to adjust the path.

```
C:\OpenOCD-20250710-0.12.0\bin
C:\OpenOCD-20250710-0.12.0\share\openocd\scripts\interface
C:\OpenOCD-20250710-0.12.0\share\openocd\scripts\target
```

For MAC, we first install Homebrew and the various dependencies and OpenOCD.

```
/bin/bash -c "$(curl -fsSL
https://raw.githubusercontent.com/Homebrew/install/HEAD/install.sh)"
brew install git libtool automake pkg-config libusb
brew install openocd
```

For Linux, we install the various dependencies and OpenOCD.

```
sudo apt update
sudo apt install git build-essential libtool autoconf pkg-config libusb-1.0-0-dev
libftd2rl-dev
sudo apt install openocd
```

Run OpenOCD with the below config.

```
openocd -f interface/cmsis-dap.cfg -f target/rp2350.cfg -c "adapter speed 5000"
```

Open a new terminal and then run the following to launch our dynamic debugger called GDB.

```
arm-none-eabi-gdb build/0x0001_hello-world.elf
```

Once it loads, we need to target our remote server.

```
target remote :3333
```

We need to halt the currently running binary.

```
monitor reset halt
```

```
C:\Users\assem.KEVINTHOMAS\Documents\Embedded-Hacking\0x0001_hello-world>arm-none-eabi-gdb build\0x0001_hello-world.elf
GNU gdb (Arm GNU Toolchain 14.3.Rel1 (Build arm-14.174)) 15.2.90.20241229-git
Copyright (C) 2024 Free Software Foundation, Inc.
License GPLv3+: GNU GPL version 3 or later <http://gnu.org/licenses/gpl.html>
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.
Type "show copying" and "show warranty" for details.
This GDB was configured as "--host=x86_64-w64-mingw32 --target=arm-none-eabi".
Type "show configuration" for configuration details.
For bug reporting instructions, please see:
<https://bugs.linaro.org/>.
Find the GDB manual and other documentation resources online at:
  <http://www.gnu.org/software/gdb/documentation/>.

For help, type "help".
Type "apropos word" to search for commands related to "word"...
Reading symbols from build\0x0001_hello-world.elf...
(gdb) target remote :3333
Remote debugging using :3333
uart_tx_wait_blocking (uart=warning: could not convert 'uart_inst' from the host encoding (CP1252) to UTF-32.
This normally should not happen, please file a bug report.
0x40070000)
   at C:/Users/assem.KEVINTHOMAS/.pico-sdk/sdk/2.2.0/src/rp2_common/hardware_uart/include/hardware/uart.h:432
432      while (uart_get_hw(uart)->fr & UART_UARTFR_BUSY_BITS) tight_loop_contents();
(gdb) monitor reset halt
[rp2350.cm0] halted due to debug-request, current mode: Thread
xPSR: 0xf9000000 pc: 0x00000088 msp: 0xf0000000
[rp2350.cm1] halted due to debug-request, current mode: Thread
xPSR: 0xf9000000 pc: 0x00000088 msp: 0xf0000000
(gdb)
```

Before we go any further, we need to turn to the RP2350 datasheet.

<https://datasheets.raspberrypi.com/rp2350/rp2350-datasheet.pdf>

2.2. Address map

The address map for the device is split into sections as shown in [Table 8](#). Details are shown in the following sections. Unmapped address ranges raise a bus error when accessed.

Each link in the left-hand column of [Table 8](#) goes to a detailed address map for that address range. The detailed address maps have a link for each address to the relevant documentation for that address.

Rough address decode is first performed on bits 31:28 of the address:

Table 8. Address Map Summary

Bus Segment	Base Address
ROM	0x00000000
XIP	0x10000000
SRAM	0x20000000
APB Peripherals	0x40000000
AHB Peripherals	0x50000000
Core-local Peripherals (SIO)	0xd0000000
Cortex-M33 private registers	0xe0000000

Above is page 30 where we see our address map.

XIP, a technique where firmware instructions are executed directly from non-volatile memory rather than being copied into RAM.

Table 10. Address map for XIP bus segment

Bus Endpoint	Base Address
XIP_BASE	0x10000000
XIP_NOCACHE_NOALLOC_BASE	0x14000000
XIP_MAINTENANCE_BASE	0x18000000
XIP_NOCACHE_NOALLOC_NOTRANSLATE_BASE	0x1c000000

At address 0x10000000, is where we will focus within GDB.

Before we dive into the assembler, we need to understand we are working with an RP2350 microcontroller that has a dual-core architecture.

This course will not focus on the RISC-V core however will focus on the ARM Cortex-M33 core as this is more prevalent in the industry today however a future course may cover the RISC-V core.

The ARM Cortex-M33 core is part of what we refer to as the Armv8-M Mainline family. We will review the Arm Cortex-M33 Processor Technical Reference Manual that is included in the course Github repo.

Name	Description
R0-R12	R0-R12 are general-purpose registers for data operations.
MSP (R13)	The <i>Stack Pointer</i> (SP) is register R13. In Thread mode, the CONTROL register indicates the stack pointer to use, <i>Main Stack Pointer</i> (MSP) or <i>Process Stack Pointer</i> (PSP).
PSP (R13)	When the Armv8-M Security Extension is included, there are two MSP registers in the Cortex-M33 processor: <ul style="list-style-type: none"> • MSP_NS for the Non-secure state. • MSP_S for the Secure state. When the Armv8-M Security Extension is included, there are two PSP registers in the Cortex-M33 processor: <ul style="list-style-type: none"> • PSP_NS for the Non-secure state. • PSP_S for the Secure state.
MSPLIM	The stack limit registers limit the extent to which the MSP and PSP registers can descend respectively.
PSPLIM	When the Armv8-M Security Extension is included, there are two MSPLIM registers in the Cortex-M33 processor: <ul style="list-style-type: none"> • MSPLIM_NS for the Non-secure state. • MSPLIM_S for the Secure state. When the Armv8-M Security Extension is included, there are two PSPLIM registers in the Cortex-M33 processor: <ul style="list-style-type: none"> • PSPLIM_NS for the Non-secure state. • PSPLIM_S for the Secure state.
LR (R14)	The <i>Link Register</i> (LR) is register R14. It stores the return information for subroutines, function calls, and exceptions.
PC (R15)	The <i>Program Counter</i> (PC) is register R15. It contains the current program address.
PSR	The <i>Program Status Register</i> (PSR) combines: <ul style="list-style-type: none"> • <i>Application Program Status Register</i> (APSR). • <i>Interrupt Program Status Register</i> (IPSR). • <i>Execution Program Status Register</i> (EPSR). These registers provide different views of the PSR.

On page B1-40, we see the above processor core register summary.

Our microcontroller has 13 general-purpose 32-bit wide registers called r_0-r_{12} . These registers will be used for storing intermediate values, passing function arguments, and performing arithmetic or logical operations during program execution. They form the core working set for most instructions and are essential for efficient data manipulation and control flow within the processor.

The r_{13} register is called the stack pointer. The stack pointer holds the address of the top of the stack, a region of memory used for temporary storage during function calls. When a function is called, local variables,

return addresses, and saved register states are pushed onto the stack. As the function exits, these values are popped off. The stack grows downward in memory on ARM Cortex-M systems, and the `sp` ensures that data is stored and retrieved in the correct order. It's critical for managing nested function calls and interrupt handling.

The `r14` register is called the link register. The link register stores the return address when a function or subroutine is called. In ARM assembly, instructions like `bl` (Branch with Link) automatically place the address of the next instruction into `lr` so the processor knows where to return after the function finishes. If `lr` is overwritten or mishandled, the program may jump to an unintended location, leading to crashes or undefined behavior. In exception handling, `lr` also plays a role in determining the return path after servicing an interrupt.

The `r15` register is called the program counter. The program counter holds the address of the next instruction to be executed. It's automatically updated as the processor steps through instructions, and can be manually modified during jumps, branches, or exceptions. The `pc` is central to control flow, whether you're executing sequential code, branching conditionally, or handling interrupts. In debugging or reverse engineering, tracking the `pc` helps you understand exactly where the processor is in its execution lifecycle.

We need to touch base on what XIP is within the RP2350 MCU microcontroller. This is the actual chip that powers the Pico 2.

As mentioned earlier, XIP is called, execute in place, and is capable of directly executing code from non-volatile storage (such as flash memory) without the need to copy the code to random-access memory (RAM) first. Instead of loading the entire program into RAM, XIP systems fetch instructions directly from their storage location and execute them on the fly.

Our goal is to find the main function within our binary to reverse engineer it. Before our main function there will be a large amount of setup code to include the vector table which will handle hardware interrupts and exceptions within our firmware which will be at the address close to the beginning of `0x10000000`.

Our XIP address starts at `0x10000000` so let's examine 1000 instructions and look for a `push {r3, lr}` followed by a call to `stdio_init_all` which would indicate our main stack frame being called.

```
(gdb) x/1000i 0x10000000
...
0x10000234 <main>:  push    {r3, lr}
...
```

This is our main program. If you are new to assembler, do not be discouraged as we will take this step-by-step!

To begin working effectively with the RP2350, it is important to understand how memory is organized within the microcontroller. The RP2350 features a dual-core ARM Cortex-M33 processor, which introduces more advanced memory management capabilities compared to earlier architectures. We start by examining the stack and heap, as these are essential concepts in embedded systems.

The stack is a region of memory used to manage function calls and local variables. It automatically grows and shrinks as functions are called and return. Each time a function is invoked, a stack frame is created to store its local variables and the return address. The stack pointer register keeps track of the current position in the stack and is updated automatically during function calls and returns.

Because the RP2350 has two cores, each core maintains its own dedicated stack. The size of each stack is typically defined in the linker script or project configuration and is constrained by the available RAM. When data is added to the stack, such as function parameters, it is referred to as a push operation. When data is removed, such as return values or saved registers, it is called a pop operation.

If the stack grows beyond its allocated space, it can result in a stack overflow. This may cause unpredictable behavior or system crashes. In contrast, the heap is a region of memory used for dynamic allocation. It is managed manually by the programmer, who must explicitly allocate and free memory as needed.

Dynamic memory allocation is performed using functions such as `malloc` in C or `new` in C++. This approach is useful for handling data structures whose size may vary during runtime. The heap in the RP2350 is typically located in the RAM region. Its size is flexible and can be adjusted based on the needs of the application.

Memory on the heap can be allocated to obtain a block of space and deallocated to return it for reuse. Over time, repeated allocation and deallocation can lead to fragmentation, which makes it harder to find large contiguous blocks of memory. The RP2350 uses standard C library functions such as `malloc` and `free` to manage heap memory. The size and location of the heap are usually defined in the linker script or project settings.

In this course, we will not necessarily focus on dynamic memory allocation. Instead, we will use safer and more predictable strategies for managing memory. The RP2350 has a limited amount of RAM, so careful planning is essential. Code is stored in Flash memory and is executed directly from that location. Understanding this memory layout is key to building reliable and efficient embedded applications.

Now let's examine our main function.

```
(gdb) x/5i 0x10000234
0x10000234 <main>:   push    {r3, lr}
0x10000236 <main+2>:   bl     0x1000156c <stdio_init_all>
0x1000023a <main+6>:   ldr    r0, [pc, #8]    @ (0x10000244 <main+16>)
0x1000023c <main+8>:   bl     0x100015fc <__wrap_puts>
0x10000240 <main+12>:  b.n    0x1000023a <main+6>
```

Let's set a breakpoint to our main function and continue.

```
(gdb) b *0x10000234
Breakpoint 1 at 0x10000234: file C:/Users/assem.KEVINTHOMAS/Documents/Embedded-
Hacking/0x0001_hello-world/0x0001_hello-world.c, line 5.
Note: automatically using hardware breakpoints for read-only addresses.
(gdb) c
Continuing.
Thread 1 "rp2350.cm0" hit Breakpoint 1, main ()
    at C:/Users/assem.KEVINTHOMAS/Documents/Embedded-Hacking/0x0001_hello-
world/0x0001_hello-world.c:5
```

```
warning: Source file is more recent than executable.
5         stdio_init_all();
```

Let's re-examine our main function and we will see an arrow pointing to the instruction we are about to execute. Keep in mind, we have NOT executed it yet.

```
(gdb) x/5i 0x10000234
=> 0x10000234 <main>:  push    {r3, lr}
0x10000236 <main+2>:  bl      0x1000156c <stdio_init_all>
0x1000023a <main+6>:  ldr    r0, [pc, #8]    @ (0x10000244 <main+16>)
0x1000023c <main+8>:  bl      0x100015fc <__wrap_puts>
0x10000240 <main+12>: b.n     0x1000023a <main+6>
```

We push the `r3` register and the `lr` register to the stack.

Keep in mind, the base pointer is not a register in the RP2350's ARM Cortex-M33 architecture. Unlike some other architectures such as x86, which use a dedicated base pointer for stack frame management, the Cortex-M33 relies on the stack pointer and the link register to handle function calls and returns.

In this architecture, the stack pointer, also known as `sp` or `r13`, points to the top of the stack and is automatically adjusted as functions are called and return. The link register, referred to as `lr` or `r14`, holds the return address when a function is invoked. These two registers work together to manage the stack and control program flow during subroutine execution.

The concept of a base pointer, as seen in x86/64 systems with the `RBP` register, is not part of the standard conventions used in the RP2350. Instead, stack frames are managed directly through `sp` and `lr` without a separate frame pointer.

It is important to note that in microcontroller environments like the RP2350, the main function typically runs in an infinite loop and does not return. As a result, the value stored in the link register after main begins execution is never used, but it remains part of the standard calling convention.

We have not executed our first main assembler function yet so let's first examine what our stack contains.

```
(gdb) x/10x $sp
0x20082000:  0x00000000    0x00000000    0x00000000    0x00000000
0x20082010:  0x00000000    0x00000000    0x00000000    0x00000000
0x20082020:  0x00000000    0x00000000
```

Now let's step-into which means take a single step in assembler.

```
(gdb) si
0x10000236    5          stdio_init_all();
(gdb) x/5i 0x10000234
0x10000234 <main>:  push    {r3, lr}
=> 0x10000236 <main+2>:  bl      0x1000156c <stdio_init_all>
0x1000023a <main+6>:  ldr    r0, [pc, #8]    @ (0x10000244 <main+16>)
0x1000023c <main+8>:  bl      0x100015fc <__wrap_puts>
```

```
0x10000240 <main+12>:      b.n      0x1000023a <main+6>
```

Let's review our stack.

```
(gdb) x/10x $sp
0x20081ff8:      0xe000ed08      0x1000018f      0x00000000      0x00000000
0x20082008:      0x00000000      0x00000000      0x00000000      0x00000000
0x20082018:      0x00000000      0x00000000
```

We can see that we have two new addresses that were pushed onto our stack.

To prove this, let's look at the values of `r3` and `lr`.

```
(gdb) x/x $r3
0xe000ed08:      Cannot access memory at address 0xe000ed08
(gdb) x/x $lr
0x1000018f <platform_entry+8>:  0x00478849
(gdb) x/x $sp
0x20081ff8:      0xe000ed08
```

The stack pointer is currently at address `0x20081ff8`, and the value at that location is `0xe000ed08`, which matches the value in `r3`. This suggests that `r3` was pushed onto the stack first.

```
(gdb) x/x $sp+4
0x20081ffc:      0x1000018f
```

We find the value `0x1000018f`, which matches the value in the link register. This confirms that the link register was pushed onto the stack after `r3`.

Because the stack grows downward in memory, each push operation moves the stack pointer to a lower address. The original stack pointer was at `0x20082000`, and after pushing two values, it moved down to `0x20081ff8`.

This behavior aligns with ARM calling conventions. During a function prologue, registers such as `lr` and any callee-saved registers are pushed onto the stack to preserve their values. The stack pointer is adjusted accordingly, and the return address stored in `lr` ensures that control can return to the correct location once the function completes.

I hope this helps you understand how the stack works. We will continue to examine the stack throughout this course.

Let's step-over the next instruction as it is a call to our below C- SDK function which is not of interest to as it simply sets up the MCU peripherals to communicate.

Our next step is to step-over the call to standard IO initialize all.

```
(gdb) x/5i 0x10000234
```

```

0x10000234 <main>:  push    {r3, lr}
=> 0x10000236 <main+2>: bl      0x1000156c <stdio_init_all>
0x1000023a <main+6>: ldr     r0, [pc, #8]    @ (0x10000244 <main+16>)
0x1000023c <main+8>: bl      0x100015fc <__wrap_puts>
0x10000240 <main+12>: b.n     0x1000023a <main+6>
(gdb) n
8          printf("hello, world\r\n");
(gdb) x/5i 0x10000234
0x10000234 <main>:  push    {r3, lr}
0x10000236 <main+2>: bl      0x1000156c <stdio_init_all>
=> 0x1000023a <main+6>: ldr     r0, [pc, #8]    @ (0x10000244 <main+16>)
0x1000023c <main+8>: bl      0x100015fc <__wrap_puts>
0x10000240 <main+12>: b.n     0x1000023a <main+6>

```

Now we are about to load the value **INSIDE** of a memory address at 0x10000244 into r0. The r0, [pc, #8] means take the value at the current program counter and add 8 to it and take that address's value and store it into r0. This is a pointer which means we are pointing to the value inside that address.

Let's *si* one step and examine what is inside r0 at this point.

```

(gdb) si
0x1000023c      8          printf("hello, world\r\n");
(gdb) x/x $r0
0x100019cc:    0x6c6c6568

```

Hmm... This does not look like an address however it does look like ascii chars to me. Let's look at an ascii table.

<https://www.asciitable.com>

We see 0x6c is l and we see it again so another l and 0x65 is e and 0x68 is h.

This is our *hello, world* string however it is backward! The reason is memory is stored in reverse byte order or little-endian order from memory to registers within the MCU.

We can see the full pointer to this char array or string by doing the below.

```

(gdb) x/s $r0
0x100019cc:    "hello, world\r"

```

In this chapter, we established a foundational reverse engineering workflow using both static and dynamic techniques. Through Ghidra, we examined the binary statically, observing the raw assembly and decompiled views to understand control flow and compiler optimizations. We noted subtle differences between our original source code and the decompiled output, such as the transformation of a `while (true)` loop into a `do-while` construct and the substitution of `printf` with `puts` for efficiency.

Using GDB, we transitioned into dynamic analysis, inspecting live register values and stack behavior during execution. We confirmed how the stack grows downward, how the link register is pushed to preserve return addresses, and how memory inspection reveals the inner workings of function calls. These observations aligned with ARM Cortex-M33 calling conventions and gave us a practical view of how the RP2350 handles execution at the instruction level.

Although the example was simple, it demonstrated the power of combining static and dynamic reverse engineering to gain insight into compiled binaries. With this workflow in place, we are now prepared to tackle more complex binaries, explore deeper architectural features of the RP2350, and refine our debugging strategies for embedded development.

In our next chapter we will hack this simple binary.


```
(gdb) x/5i 0x10000234
0x10000234 <main>:  push    {r3, lr}
0x10000236 <main+2>:  bl      0x1000156c <stdio_init_all>
0x1000023a <main+6>:  ldr     r0, [pc, #8]    @ (0x10000244 <main+16>)
0x1000023c <main+8>:  bl      0x100015fc <__wrap_puts>
0x10000240 <main+12>: b.n     0x1000023a <main+6>
```

The first thing we need to do to hack our system LIVE is to set a breakpoint to the address right before the call to puts and then continue.

```
(gdb) b *0x1000023c
Breakpoint 1 at 0x1000023c: file C:/Users/assem.KEVINTHOMAS/Documents/Embedded-
Hacking/0x0001_hello-world/0x0001_hello-world.c, line 8.
Note: automatically using hardware breakpoints for read-only addresses.
(gdb) c
Continuing.
```

```
Thread 1 "rp2350.cm0" hit Breakpoint 1, 0x1000023c in main ()
   at C:/Users/assem.KEVINTHOMAS/Documents/Embedded-Hacking/0x0001_hello-world/0x0001_hello-
world.c:8
warning: Source file is more recent than executable.
8      printf("hello, world\r\n");
```

```
(gdb) disas
Dump of assembler code for function main:
   0x10000234 <+0>:  push    {r3, lr}
   0x10000236 <+2>:  bl      0x1000156c <stdio_init_all>
   0x1000023a <+6>:  ldr     r0, [pc, #8]    @ (0x10000244 <main+16>)
=>  0x1000023c <+8>:  bl      0x100015fc <__wrap_puts>
   0x10000240 <+12>: b.n     0x1000023a <main+6>
   0x10000242 <+14>:  nop
   0x10000244 <+16>:  adds   r4, r1, r7
   0x10000246 <+18>:  asrs   r0, r0, #32
End of assembler dump.
(gdb)
```

The next thing we need to do is hijack the value of *hello, world* which is pointed to in *r0* and create our own data and fill it with a hacked malicious string.

```
(gdb) x/s $r0
0x100019cc:  "hello, world\r"
(gdb) set $r0 = "hacky, world\r"
evaluation of this expression requires the program to have a function "malloc".
```

```
(gdb) x/s $r0
0x100019cc:  "hello, world\r"
```

Oh no it did not work! Now what!

GDB interprets "hacky, world\r" as a string literal, and it tries to evaluate it as a pointer to a valid memory address where that string resides. But GDB itself does not allocate memory for that string unless the program being debugged has already loaded it somewhere, typically via the C runtime or a statically defined string in the binary.

The error isn't because GDB is trying to call `malloc`. It's because GDB is trying to resolve the string literal to a memory address, and it fails because that string doesn't exist in the program's memory space. If you're debugging a bare-metal binary or early startup code on the RP2350, there's no runtime environment to provide that string, and no global symbols or `.data` section initialized with it.

Therefore, we need to create our string in SRAM!

If we remember from the last chapter, the RP2350 datasheet states that the SRAM starts at `0x20000000`. With that we can create a new string in SRAM directly.

```
(gdb) set {char[14]} 0x20000000 = {'h','a','c','k','y',' ',' ',' ','w','o','r','l','d','\r','\0'}
(gdb) x/s 0x20000000
0x20000000 <ram_vector_table>: "hacky, world\r"
```

Now to need to hijack the address inside `r0` and change it to our hacked address in SRAM and verify our hack.

```
(gdb) set $r0 = 0x20000000
(gdb) x/x $r0
0x20000000 <ram_vector_table>: 0x68
(gdb) x/s $r0
0x20000000 <ram_vector_table>: "hacky, world\r"
```

Let's continue and execute our hack!

```
(gdb) c
Continuing.
```

```
Thread 1 "rp2350.cm0" hit Breakpoint 1, 0x1000023c in main ()
    at C:/Users/assem.KEVINTHOMAS/Documents/Embedded-Hacking/0x0001_hello-world/0x0001_hello-
world.c:8
8             printf("hello, world\r\n");
```


Chapter 4: Embedded System Analysis

We are working with a microcontroller so there is no operating system in use. This is what we refer to as bare-metal programming.

```
// ROOT ADDRESSES
#define BOOTROM_MAGIC_OFFSET 0x10
#define BOOTROM_FUNC_TABLE_OFFSET 0x14
#if PICO_RP2040
#define BOOTROM_DATA_TABLE_OFFSET 0x16
#endif

#if PICO_RP2040
#define BOOTROM_VTABLE_OFFSET 0x00
#define BOOTROM_TABLE_LOOKUP_OFFSET 0x18
#else
#define BOOTROM_WELL_KNOWN_PTR_SIZE 2
#if defined(__riscv)
#define BOOTROM_ENTRY_OFFSET 0x7dfc
#define BOOTROM_TABLE_LOOKUP_ENTRY_OFFSET (BOOTROM_ENTRY_OFFSET -
BOOTROM_WELL_KNOWN_PTR_SIZE)
#define BOOTROM_TABLE_LOOKUP_OFFSET (BOOTROM_ENTRY_OFFSET -
BOOTROM_WELL_KNOWN_PTR_SIZE*2)
#else
#define BOOTROM_VTABLE_OFFSET 0x00
```

We must start with what happens when the RP2350 gets power.

The RP2350 has an on-chip bootloader (bootrom) that executes immediately when the chip gets power.

```
src/rp2_common/boot_bootrom_headers/include/boot/bootrom_constants.h
```

The RP2350 bootrom is a mask ROM that contains the first-stage bootloader code. This bootrom provides various

```
static inline void rom_connect_internal_flash(void) {
    rom_connect_internal_flash_fn func = (rom_connect_internal_flash_fn)
rom_func_lookup_inline(ROM_FUNC_CONNECT_INTERNAL_FLASH);
    func();
}
```

functions including flash initialization, boot path selection, and hardware setup.

```
src/rp2_common/pico_bootrom/include/pico/bootrom.h
```

On RP2350, boot stage 2 is called as a regular function and must return normally, unlike RP2040, **boot2_generic_03h.S**. The second stage bootloaders are responsible for setting up external flash to enable XIP operation.

```

// The QMI is automatically configured for 03h XIP straight out of reset,
// but this code can't assume it's still in that state. Set up memory
// window 0 for 03h serial reads.

// Setup timing parameters: short sequential-access cooldown, configured
// CLKDIV and RXDELAY, and no constraints on CS max assertion, CS min
// deassertion, or page boundary burst breaks.

#define INIT_M0_TIMING (\
    1                << QMI_M0_TIMING_COOLDOWN_LSB |\
    PICO_FLASH_SPI_RXDELAY << QMI_M0_TIMING_RXDELAY_LSB |\
    PICO_FLASH_SPI_CLKDIV  << QMI_M0_TIMING_CLKDIV_LSB  |\
0)

// Set command constants
#define INIT_M0_RCMD (\
    CMD_READ          << QMI_M0_RCMD_PREFIX_LSB  |\
0)

// Set read format to all-serial with a command prefix
#define INIT_M0_RFMT (\
    QMI_M0_RFMT_PREFIX_WIDTH_VALUE_S << QMI_M0_RFMT_PREFIX_WIDTH_LSB  |\
    QMI_M0_RFMT_ADDR_WIDTH_VALUE_S   << QMI_M0_RFMT_ADDR_WIDTH_LSB   |\
    QMI_M0_RFMT_SUFFIX_WIDTH_VALUE_S << QMI_M0_RFMT_SUFFIX_WIDTH_LSB  |\
    QMI_M0_RFMT_DUMMY_WIDTH_VALUE_S  << QMI_M0_RFMT_DUMMY_WIDTH_LSB   |\
    QMI_M0_RFMT_DATA_WIDTH_VALUE_S   << QMI_M0_RFMT_DATA_WIDTH_LSB   |\
    QMI_M0_RFMT_PREFIX_LEN_VALUE_8   << QMI_M0_RFMT_PREFIX_LEN_LSB   |\
0)

```

The default `boot2_generic_03h` implementation configures the QMI for basic serial flash operation.
`src/rp2350/boot_stage2/boot2_generic_03h.S`

The configuration sets up timing parameters with a short cooldown, configurable clock divider and RX delay, and configures the QMI for 03h serial read commands with all-serial format.

After QMI configuration, boot stage 2 performs a dummy transfer to initialize the flash device and then configures continuous read mode.

```

// Dummy transfer
mov r1, #XIP_NOCACHE_NOALLOC_BASE
ldrb r1, [r1]

// Set prefix length to 0, as flash no longer expects to see commands
bic r0, #QMI_M0_RFMT_PREFIX_LEN_BITS
str r0, [r3, #QMI_M0_RFMT_OFFSET]

```

`src/rp2350/boot_stage2/boot2_w25q080.S`

The dummy transfer activates XIP mode, and the prefix length is set to 0 since the flash no longer expects command prefixes for subsequent reads.

Boot stage 2 returns control to the bootrom, which then jumps to the `reset_vector` as that value is the second entry in the vector table at `0x10000004` which in our case is `0x1000015d`.

```
// Pull in standard exit routine
#include "boot2_helpers/exit_from_boot2.S"
```

We will focus on execute in place, or XIP, a technique where firmware instructions are executed directly from non-volatile memory rather than being copied into RAM. In the context of the RP2350, this typically means that code is mapped from external or internal Flash memory into the processor's address space, allowing instructions to be fetched and executed without relocation.

This approach conserves RAM and simplifies startup, since the processor can begin executing code immediately after reset. The Flash region is memory-mapped, so the CPU treats it as part of its normal instruction space. While XIP is efficient for read-only code execution, it's important to note that Flash access times are generally slower than RAM, and write operations require special handling.

Understanding XIP is essential for debugging and reverse engineering, as it affects how code is laid out, how breakpoints behave, and how memory regions are protected or cached. Let me know if you'd like to walk through the RP2350's memory map or trace instruction fetches from Flash during startup.

When we examine the first few values at 0x10000000, we begin with the vector table.

```
(gdb) x/4x 0x10000000
0x10000000 <__vectors>: 0x20082000      0x1000015d      0x1000011b      0x1000011d
```

Address	Value	Meaning
0x10000000	0x20082000	Initial Stack Pointer (SP)
0x10000004	0x1000015d	Reset Handler (entry point after boot)
0x10000008	0x1000011b	NMI Handler
0x1000000C	0x1000011d	HardFault Handler

The reset handler is at 0x1000015d, so disassembling from there will show the actual startup logic.

```
(gdb) x/3i 0x1000015d
0x1000015d <_reset_handler>: mov.w    r0, #3489660928 @ 0xd0000000
0x10000161 <_reset_handler+4>: ldr    r0, [r0, #0]
0x10000163 <_reset_handler+6>: cbz   r0, 0x1000016a <hold_non_core0_in_bootrom+6>
```

On ARM Cortex-M chips, all code runs in Thumb mode, and the processor uses the least significant bit of an address to mark this: if bit 0 is set, it means "Thumb," if clear, it means "ARM." The actual instructions still live at even addresses, but debuggers and disassemblers handle this flag differently as GDB shows the address exactly as it appears in the vector table (with the Thumb bit set), while Ghidra strips that bit off and shows the true instruction address. So, both are correct, they're just presenting the same location in two slightly different ways.

Let's start from the reset handler and work our way to main.

At 0x1000015d: mov.w r0, #0xd0000000 - Load SIO base address.

At 0x10000161: ldr r0, [r0, #0x0] - Read the CPUID register.

At 0x10000163: cbz r0, LAB_1000016a - Branch if core 0 (r0 == 0).

```
*****
*
*                               FUNCTION
*****
undefined _reset_handler()
    assume LRset = 0x0
    assume TMode = 0x1
undefined  ▲ <UNASSIGNED> <RETURN>
    _reset_handler
    crt0.S:446 (4)
1000015c 4f f0 50 40    mov.w    r0,#0xd0000000
    crt0.S:447 (2)
10000160 00 68            ldr     r0,[r0,#0x0]=>DAT_d0000000
    crt0.S:452 (2)
10000162 10 b1            cbz     r0,LAB_1000016a
```

At 0x10000164-0x10000168, if not core 0, send back to bootrom.

```
    hold_non_core0_in_bootrom
    crt0.S:456 (4)
10000164 4f f0 00 00    mov.w    r0,#0x0
    crt0.S:457 (2)
10000168 f2 e7            b       _enter_vtable_in_r0
```

Data Copy Phase (0x1000016a-0x10000176)

- Copies initialized data from flash to RAM using the data_cpy_table.
- The loop at LAB_1000016c processes each entry in the copy table.

```

LAB_1000016a                                XREF[1]: 10000162(j)
crt0.S:481 (2)
1000016a 0d a4          adr          r4, [0x100001a0]

LAB_1000016c                                XREF[1]: 10000176(j)
crt0.S:485 (2)
1000016c 0e cc          ldmia     r4!, {r1,r2,r3}=>data_cpy_table
                                                    = 10003804h
                                                    = 20000110h
                                                    = 2000062Ch
                                                    = 10003D20h
                                                    = 20080000h
                                                    = D3h

crt0.S:486 (2)
1000016e 00 29          cmp          r1,
crt0.S:487 (2)
10000170 02 d0          beq          LAB byte 0h      0
                                                    Hex Decimal
crt0.S:488 (4)
10000172 00 f0 12 f8    bl          data_cpy          undefined data_cpy()
crt0.S:489 (2)
10000176 f9 e7          b          LAB_1000016c

```

BSS Clear Phase (0x10000178-0x10000184)

- Zeros out the BSS section in RAM.
- The loop clears memory from 0x2000062c to 0x20000858.

```

LAB_10000178                                XREF[1]: 10000170(j)
crt0.S:494 (2)
10000178 15 49          ldr          r1, [DAT_100001d0]
                                                    = 2000062Ch
crt0.S:495 (2)
1000017a 16 4a          ldr          r2, [DAT_100001d4]
                                                    = 20000858h
crt0.S:496 (2)
1000017c 00 20          movs        r0, #0x0
crt0.S:497 (2)
1000017e 00 e0          b          bss_fill_test

bss_fill_loop                                XREF[1]: 10000184(j)
crt0.S:499 (2)
10000180 01 c1          stmia     r1! => __TMC_END__, {r0}

bss_fill_test                                XREF[1]: 1000017e(j)
crt0.S:501 (2)
10000182 91 42          cmp          r1, r2
crt0.S:502 (2)
10000184 fc d1          bne     bss_fill_loop

```

Runtime Initialization (0x10000186-0x10000188)

- Calls `runtime_init`.
- This sets up the C runtime environment.

```
platform_entry
crt0.S:512 (2)
10000186 14 49      ldr      r1, [DAT_100001d8]          = 10002E7Dh
crt0.S:513 (2)
10000188 88 47      blx     r1=>runtime_init          void runtime_init(void)
```

Main Function Call (0x1000018a-0x1000018c)

- Finally calls `main` at `0x10000234`.

```
crt0.S:514 (2)
1000018a 14 49      ldr      r1, [DAT_100001dc]          = 10000235h
crt0.S:515 (2)
1000018c 88 47      blx     r1=>main                    int main(void)
```

But where does this all come from?

We setup VSCode with the Pico extension. In Windows you will see something like the following.

`C:\Users\assem.KEVINTHOMAS\.pico-sdk\sdk\2.2.0\src\rp2_common\pico crt0`

There is a file called `crt0.S` to which this all begins!

Below is a snippet from the file.

```
.section .vectors, "ax"
.align 2

.global __vectors, __VECTOR_TABLE, __vectors_end
__VECTOR_TABLE:
__vectors:
.word __StackTop
.word _reset_handler
```

These entries correspond to `0x20082000` which is the stack pointer and `0x1000015d` which is the reset handler.

The RP2350 vector table is a critical structure that defines how the microcontroller responds to exceptions and interrupts, but it's not the first thing the ARM Cortex-M33 core looks at when it powers up as the on-chip bootrom executes first, followed by boot stage 2 configuration of the flash interface, and only then does the bootrom read the vector table and jump to the application's reset handler.

The vector table lives at `0x10000000` in the RP2350's XIP Flash region with the stack pointer at offset `0x00` and the reset vector at offset `0x04`, but this location is determined by the application's linker script rather than being a fixed hardware requirement as the bootrom uses the Vector Table Offset Register (VTOR) to locate the table dynamically.

We will find the linker scripts specifically for our 2.2.0 sdk in a folder similar to this.

C:\Users\assem.KEVINTHOMAS\.pico-sdk\sdk\2.2.0\src\rp2_common\pico crt0\rp2350

There you will see **memmap_default.ld** which is the standard XIP configuration where code executes directly from Flash at 0x10000000.

In our linker script we see the following.

```
MEMORY
{
  INCLUDE "pico_flash_region.ld"
  RAM(rwx) : ORIGIN = 0x20000000, LENGTH = 512k
  SCRATCH_X(rwx) : ORIGIN = 0x20080000, LENGTH = 4k
  SCRATCH_Y(rwx) : ORIGIN = 0x20081000, LENGTH = 4k
}
```

Then as we look deeper, we see the following.

```
__StackTop = ORIGIN(SCRATCH_Y) + LENGTH(SCRATCH_Y);
```

We see above that the `ORIGIN(SCRATCH_Y)` is `0x20081000` and the length is `4k` therefore we get the following which we can verify in GDB.

```
__StackTop = 0x20081000 + 0x1000 = 0x20082000
```

```
(gdb) x/x 0x10000000
0x10000000 <__vectors>: 0x20082000
```

This value is emitted into the vector table at address `0x10000000` via the **crt0.S** file.

```
.section .vectors, "ax"
.align 2

.global __vectors, __VECTOR_TABLE, __vectors_end
__VECTOR_TABLE:
__vectors:
.word __StackTop
.word reset_handler
```

The Cortex-M33 core loads this into the stack pointer register and places the stack at the top of the `SCRATCH_Y` region, which is a small, dedicated RAM block reserved for the core 0 stack.

At the end of the vector table, we see the following.

```
(gdb) x/36i 0x10000110
0x10000110 <isr_usagefault>: mrs      r0, IPSR
0x10000114 <isr_usagefault+4>: subs   r0, #16
0x10000116 <unhandled_user_irq_num_in_r0>: bkpt   0x0000
0x10000118 <isr_invalid>: bkpt   0x0000
0x1000011a <isr_nmi>: bkpt   0x0000
0x1000011c <isr_hardfault>: bkpt   0x0000
0x1000011e <isr_svcall>: bkpt   0x0000
0x10000120 <isr_pendsv>: bkpt   0x0000
0x10000122 <isr_systick>: bkpt   0x0000
```

```

0x10000124 <__default_isrs_end>:                                     @ <UNDEFINED> instruction: 0xebf27188
0x10000128 <__default_isrs_end+4>: subs    r0, r4, r4
0x1000012a <__default_isrs_end+6>: asrs    r0, r0, #32
0x1000012c <__default_isrs_end+8>: subs    r4, r1, r5
0x1000012e <__default_isrs_end+10>: asrs   r0, r0, #32
0x10000130 <__default_isrs_end+12>: lsls   r0, r4, #6
0x10000132 <__default_isrs_end+14>: asrs   r0, r0, #32
0x10000134 <__default_isrs_end+16>: add    r3, pc, #576 @ (adr r3, 0x10000378
<runtime_init_per_core_irq_priorities+44>)
0x10000136 <__default_isrs_end+18>: b.n    0xfffff6e
0x10000138 <__binary_info_header_end>: udf    #211 @ 0xd3
0x1000013a <__binary_info_header_end+2>: @ <UNDEFINED> instruction:
0xffff0142
0x1000013e <__binary_info_header_end+6>: asrs   r1, r4, #32
0x10000140 <__binary_info_header_end+8>: lsls   r7, r7, #7
0x10000142 <__binary_info_header_end+10>: movs   r0, r0
0x10000144 <__binary_info_header_end+12>: subs   r0, r6, r6
0x10000146 <__binary_info_header_end+14>: movs   r0, r0
0x10000148 <__binary_info_header_end+16>: adds   r5, #121 @ 0x79
0x1000014a <__binary_info_header_end+18>: add    r3, sp, #72 @ 0x48
0x1000014c <__entry_point>: mov.w  r0, #0
0x10000150 <__enter_vtable_in_r0>: ldr    r1, [pc, #120] @ (0x100001cc
<data_cpy_table+44>)
0x10000152 <__enter_vtable_in_r0+2>: str    r0, [r1, #0]
0x10000154 <__enter_vtable_in_r0+4>: ldmia  r0!, {r1, r2}
0x10000156 <__enter_vtable_in_r0+6>: msr    MSP, r1
0x1000015a <__enter_vtable_in_r0+10>: bx     r2
0x1000015c <__reset_handler>: mov.w  r0, #3489660928 @ 0xd0000000
0x10000160 <__reset_handler+4>: ldr    r0, [r0, #0]
0x10000162 <__reset_handler+6>: cbz    r0, 0x1000016a <hold_non_core0_in_bootrom+6>

```

The first section is the `isr_usagefault` to which we will do a little digging.

```

arm-none-eabi-nm -C build\0x0001_hello-world.elf | findstr isr_usagefault
10000110 W isr_usagefault

```

This means this is weakly defined as **crt0.S** has only the stub but the code we see below is elsewhere.

```

0x10000110 <isr_usagefault>: mrs      r0, IPSR
0x10000114 <isr_usagefault+4>: subs   r0, #16

```

In **crt0.S** we see the following.

```

// Declare a weak symbol for each ISR.
// By default, they will fall through to the undefined IRQ handler below (breakpoint),
// but can be overridden by C functions with correct name.

.macro decl_isr_bkpt name
.weak \name
.type \name,%function
.thumb_func
\name:
    bkpt #0
.endm

```

We can try searching with PowerShell.

```

PS C:\Users\assem.KEVINTHOMAS> Get-ChildItem -Recurse -Include *.S -Path
"C:\Users\assem.KEVINTHOMAS\pico-sdk" | Select-String "mrs r0, IPSR"

```

Sadly, this returns no result. Let's look within our running GDB instance.

```

(gdb) list isr_usagefault
315     .global __unhandled_user_irq
316     .thumb_func
317     __unhandled_user_irq:
318     // if we include the implementation if there could be a valid IRQ hanler in the
vtable that uses it
319     #if !(PICO_NO_RAM_VECTOR_TABLE && PICO_MINIMAL_STORED_VECTOR_TABLE)
320         mrs r0, ipsr
321         subs r0, #16
322     .global unhandled_user_irq_num_in_r0
323     unhandled_user_irq_num_in_r0:
324     #endif

```

```

// All unhandled USER IRQs fall through to here.
// Additionally, if the Armv9-M MemManage/BusFault/UsageFault/SecureFault/DebugMonitor
exceptions
// are enabled, but the handlers are not defined, then unhandled_user_irq_num_in_r0 will
// also be reached, but with a negative exception number (e.g. MemManage == -12)
.global __unhandled_user_irq
.thumb_func
__unhandled_user_irq:
// if we include the implementation if there could be a valid IRQ hanler in the vtable
that uses it
#if !(PICO_NO_RAM_VECTOR_TABLE && PICO_MINIMAL_STORED_VECTOR_TABLE)
    mrs r0, ipsr
    subs r0, #16
#endif

```

Now when we look in **cr0.S**, we can see the following.

Let's now examine the next few lines of GDB.

```

(gdb) x/36i 0x10000110
..
0x10000116 <unhandled_user_irq_num_in_r0>:  bkpt      0x0000
0x10000118 <isr_invalid>:  bkpt      0x0000

```

```

0x1000011a <isr_nmi>:          bkpt    0x0000
0x1000011c <isr_hardfault>:   bkpt    0x0000
0x1000011e <isr_svcall>:     bkpt    0x0000
0x10000120 <isr_pendsv>:     bkpt    0x0000
0x10000122 <isr_systick>:    bkpt    0x0000
..

```

We can see in **crt0.S**, directly below our other code, we see the following.

```

.global unhandled_user_irq_num_in_r0
unhandled_user_irq_num_in_r0:
#endif
    // note the next instruction is a breakpoint too, however we have a 2 byte alignment
hole
    // and it is preferable to have distinct labels, to inform the user what has happened
in the debugger.
    bkpt #0

decl_isr_bkpt isr_invalid
#if !PICO_MINIMAL_STORED_VECTOR_TABLE
// these are separated out into individual BKPT instructions with label for clarity
decl_isr_bkpt isr_nmi
decl_isr_bkpt isr_hardfault
decl_isr_bkpt isr_svcall
decl_isr_bkpt isr_pendsv
decl_isr_bkpt isr_systick
#endif

```

Let's continue our analysis with the next few lines.

```

(gdb) x/36i 0x10000110
..
0x10000124 <__default_isrs_end>:                               @ <UNDEFINED> instruction: 0xebf27188
0x10000128 <__default_isrs_end+4>:   subs    r0, r4, r4
0x1000012a <__default_isrs_end+6>:   asrs    r0, r0, #32
0x1000012c <__default_isrs_end+8>:   subs    r4, r1, r5
0x1000012e <__default_isrs_end+10>:  asrs    r0, r0, #32
0x10000130 <__default_isrs_end+12>:  lsls    r0, r4, #6
0x10000132 <__default_isrs_end+14>:  asrs    r0, r0, #32
0x10000134 <__default_isrs_end+16>:  add     r3, pc, #576    @ (adr r3, 0x10000378)
<runtime_init_per_core_irq_priorities+44>
0x10000136 <__default_isrs_end+18>:  b.n     0xfffff6e
..

```

In our **crt0.S**, we see only the following.

```

.global __default_isrs_end
__default_isrs_end:

```

Where does this actual code come from?

It's not code it is a binary-info header emitted by the startup assembly, sitting immediately after the default ISR marker.

In PowerShell, let's do the following.

```

arm-none-eabi-objdump -d --source build\0x0001_hello-world.elf |
  Select-String '^\\s*1000012[4-9]|^\\s*1000013[0-6]' -Context 1,2 |
  ForEach-Object { $_.Context.PreContext + $_.Line + $_.Context.PostContext } |
  ForEach-Object { $_.Trim() } |
  Where-Object { $_ -ne "" } |
  Select-Object -Unique
10000124 <__default_isrs_end>:
10000124:      7188ebf2      .word    0x7188ebf2
10000128:      10001b20      .word    0x10001b20
1000012c:      10001b4c      .word    0x10001b4c
10000130:      100001a0      .word    0x100001a0
10000134:      e71aa390      .word    0xe71aa390
10000138 <__binary_info_header_end>:

```

Address	Value	Field / Symbol	Description
0x10000124	0x7188EBF2	Magic signature	A fixed identifier marking the start of the binary-info header. Used by tools/boot ROM to recognize this structure.
0x10000128	0x10001B20	Binary info start pointer	Address of the first entry in the .binary_info section. In this build, that's __bi_ptr84.
0x1000012C	0x10001B4C	Binary info end pointer	Address just past the last .binary_info entry. Here it's start + 0x2C bytes.
0x10000130	0x100001A0	Data copy table pointer	Address of data_cpy_table, used by the reset handler to copy initialised .data from flash to RAM.
0x10000134	0xE71AA390	Reserved / trailer constant	A fixed value defined in the SDK's startup assembly; may serve as a checksum, version marker, or reserved field.
0x10000138	(label)	__binary_info_header_end	Symbol marking the end of the binary-info header block.

```

PS C:\Users\assem.KEVINTHOMAS\Documents\Embedded-Hacking\0x0001_hello-world> arm-none-eabi-objdump -s -j .text build\0x0001_hello-world.elf | Select-String "10000120" -Context 0,6

```

```

10000120 00be00be f2eb8871 201b0010 4c1b0010 .....q ...L...
10000130 a0010010 90a31ae7 d3deffff 42012110 .....B.!.
10000140 ff010000 b01b0000 793512ab 4ff00000 .....y5..O...
10000150 1e490860 06c881f3 08881047 4ff05040 .I.`.....GO.P@
10000160 006810b1 4ff00000 f2e70da4 0ecc0029 .h..O.....)
10000170 02d000f0 12f8f9e7 1549164a 002000e0 .....I.J. ..
10000180 01c19142 fcd11449 88471449 88471449 ...B...I.G.I.G.I

```

Let's continue with our GDB analysis.

```
(gdb) x/36i 0x10000110
```

```
..
```

```

0x10000138 <__binary_info_header_end>:      udf      #211      @ 0xd3
0x1000013a <__binary_info_header_end+2>:    @ <UNDEFINED> instruction:
0xffff0142
0x1000013e <__binary_info_header_end+6>:    asrs     r1, r4, #32
0x10000140 <__binary_info_header_end+8>:    lsls     r7, r7, #7
0x10000142 <__binary_info_header_end+10>:   movs     r0, r0
0x10000144 <__binary_info_header_end+12>:   subs     r0, r6, r6
0x10000146 <__binary_info_header_end+14>:   movs     r0, r0
0x10000148 <__binary_info_header_end+16>:   adds     r5, #121      @ 0x79
0x1000014a <__binary_info_header_end+18>:   add      r3, sp, #72    @ 0x48
..

```

In **crt0.S** we see the following.

```

.section .binary_info_header, "a"

// Header must be in first 256 bytes of main image (i.e. excluding flash boot2).
// For flash builds we put it immediately after vector table; for NO_FLASH the
// vectors are at a +0x100 offset because the bootrom enters RAM images directly
// at their lowest address, so we put the header in the VTOR alignment hole.

#if !PICO_NO_BINARY_INFO
binary_info_header:
.word BINARY_INFO_MARKER_START
.word __binary_info_start
.word __binary_info_end
.word data_cpy_table // we may need to decode pointers that are in RAM at runtime.
.word BINARY_INFO_MARKER_END
#endif

#include "embedded_start_block.inc.S"

```

Let's dig in and see what we can find.

```

arm-none-eabi-objdump -d --source build\0x0001_hello-world.elf |
  Select-String '^\\s*1000013[8-9]|^\\s*1000014[0-9a-f]' -Context 1,2 |
  ForEach-Object { $_.Context.PreContext + $_.Line + $_.Context.PostContext } |
  ForEach-Object { $_.Trim() } |
  Where-Object { $_ -ne "" } |
  Select-Object -Unique
10000138 <__binary_info_header_end>:
10000138:      fffffded3      .word   0xffffded3
1000013c:      10210142      .word   0x10210142
10000140:      000001ff      .word   0x000001ff
10000144:      00001bb0      .word   0x00001bb0
10000148:      ab123579      .word   0xab123579

```

```

PS C:\Users\assem.KEVINTHOMAS\Documents\Embedded-Hacking\0x0001_hello-world> arm-none-eabi-
objdump -s --start-address=0x10000138 --stop-address=0x1000014c build\0x0001_hello-world.elf

```

```

build\0x0001_hello-world.elf:      file format elf32-littlearm

```

```

Contents of section .text:

```

```

10000138 d3deffff 42012110 ff010000 b01b0000  ....B.!.....
10000148 793512ab                          y5..

```

```

PS C:\Users\assem.KEVINTHOMAS\Documents\Embedded-Hacking\0x0001_hello-world> # Dump all

```

symbols and grep for our addresses

```
PS C:\Users\assem.KEVINTHOMAS\Documents\Embedded-Hacking\0x0001_hello-world> arm-none-eabi-nm
--numeric-sort build\0x0001_hello-world.elf |
>>      Select-String "10000138|1000013c|10000140|10000144|10000148"
```

```
10000138 T __binary_info_header_end
```

```
10000138 t embedded_block
```

Raw words and interpretations.

```
0x10000138: 0xFFFFDED3
```

- Marker start: Picobin block start marker (BlockMarkerStart).

```
0x1000013C: 0x10212142
```

- Four item-header bytes: This is not a pointer; it's the first item header packed into 4 bytes (1B head/type + 1B size + 2B typedata). In default RP2350 builds this is the IMAGE_TYPE item emitted by `embedded_start_block.inc.S3`.

```
0x10000140: 0x000001FF
```

- Next item header bytes or size field: Another 4 bytes belonging to the item sequence (depends on which items are compiled in; see decode steps below). For minimum metadata images, you'll see the LAST item header here1.

```
0x10000144: 0x00001BB0
```

- Link to next block (relative bytes) or continuation of item data: Picobin blocks store a 32-bit link
- "offset to next block from this header." If `END_BLOCK` is enabled in your build, this will be a positive offset to the end block; otherwise, it is 0 to loop to self2.

```
0x10000148: 0xAB123579
```

- Marker end: Picobin block end marker (BlockMarkerEnd).

As we continue our analysis.

```
(gdb) x/36i 0x10000110
```

```
..
0x1000014c <_entry_point>:  mov.w   r0, #0
..
```

In `crt0.S` we see the following.

```
#if !PICO_CRT0_NO_RESET_SECTION
.section .reset, "ax"

// On flash builds, the vector table comes first in the image (conventional).
// On NO_FLASH builds, the reset handler section comes first, as the entry
// point is at offset 0 (fixed due to bootrom), and VTOR is highly-aligned.
// Image is entered in various ways:
//
// - NO_FLASH builds are entered from beginning by UF2 bootloader
//
// - Flash builds vector through the table into _reset_handler from boot2
//
// - Either type can be entered via _entry_point by the debugger, and flash builds
//   must then be sent back round the boot sequence to properly initialise flash

// ELF entry point:
.type _entry_point,%function
.thumb_func
.global _entry_point
_entry_point:
```

```

#if PICO_NO_FLASH
    // on the NO_FLASH case, we do not do a rest thru bootrom below, so the RCP may or may
not have been initialized:
    //
    // in the normal (e.g. UF2 download etc. case) we will have passed thru bootrom
initialization, but if
    // a NO_FLASH binary is loaded by the debugger, and run directly after a reset, then
we won't have.
    //
    // we must therefore initialize the RCP if it hasn't already been

#if HAS_REDUNDANCY_COPROCESSOR
    // just enable the RCP which is fine if it already was (we assume no other co-
processors are enabled at this point to save space)
    ldr r0, = PPB_BASE + M33_CPACR_OFFSET
    movs r1, #ARM_CPU_PREFIXED(CPACR_CP7_BITS)
    str r1, [r0]
    // only initialize canary seeds if they haven't been (as to do so twice is a fault)
    mrc p7, #1, apsr_nzcv, c0, c0, #0
    bmi 1f
    // i dont think it much matters what we initialized to, as to have gotten here we must
have not
    // gone thru the bootrom (which a secure boot would have)
    mcrr p7, #8, r0, r0, c0
    mcrr p7, #8, r0, r0, c1
    sev
1:
#endif
#if !__ARM_ARCH_6M__
    // Make sure stack limit is 0 if we came in thru the debugger; we do not know what it
should be
    movs r0, #0
    msr msplim, r0
#endif

    ldr r0, =__vectors
    // Vector through our own table (SP, VTOR will not have been set up at
// this point). Same path for debugger entry and bootloader entry.
#else
    // Debugger tried to run code after loading, so SSI is in 03h-only mode.
    // Go back through bootrom + boot2 to properly initialise flash.
    ldr r0, =BOOTROM_VTABLE_OFFSET
#endif

```

What we see is `ldr r0, =BOOTROM_VTABLE_OFFSET` and this optimizes down to `mov.w r0, #0`.

The rest of the code up to `main` is here and directly translates in **crt0.S** nicely. Let's break this down piece by piece.

```
_enter_vtable_in_r0:
    ldr r1, =(PPB_BASE + ARM_CPU_PREFIXED(VTOR_OFFSET))
    str r0, [r1]
    ldmia r0!, {r1, r2}
    msr msp, r1
    bx r2
```

(gdb) x/36i 0x10000110

```
..
0x10000150 <_enter_vtable_in_r0>:    ldr    r1, [pc, #120] @ (0x100001cc
<data_cpy_table+44>)
0x10000152 <_enter_vtable_in_r0+2>:  str    r0, [r1, #0]
0x10000154 <_enter_vtable_in_r0+4>:  ldmia  r0!, {r1, r2}
0x10000156 <_enter_vtable_in_r0+6>:  msr    MSP, r1
0x1000015a <_enter_vtable_in_r0+10>: bx     r2
..
```

On the RP2350's Cortex-M33 core, `_enter_vtable_in_r0` is a tiny hand-off routine that takes a pointer to a new vector table in `r0`, writes it into the Vector Table Offset Register (VTOR) so all future exceptions and interrupts use it, then reads the first two words from that table so the initial Main Stack Pointer value and the `Reset_Handler` address and loads the MSP accordingly, and finally branches to the `Reset_Handler`, effectively transferring execution as if the CPU had just reset into the new firmware.

```
.type _reset_handler,%function
.thumb_func
_reset_handler:
    // Note if we entered thru here on core 0, then we should have gone thru bootrom, so
    // SP (and MSPLIM) on Armv8-M
    // should already be set

    // Only core 0 should run the C runtime startup code; core 1 is normally
    // sleeping in the bootrom at this point but check to be sure (e.g. if
    // debugger put core 1 at the ELF entry point for some reason)
    ldr r0, =(SIO_BASE + SIO_CPUID_OFFSET)
    ldr r0, [r0]
#ifdef __ARM_ARCH_6M__
    cmp r0, #0
    beq 1f
#else
    cbz r0, 1f
#endif
```

(gdb) x/36i 0x10000110

```
..
0x1000015c <_reset_handler>: mov.w   r0, #3489660928 @ 0xd0000000
0x10000160 <_reset_handler+4>: ldr    r0, [r0, #0]
0x10000162 <_reset_handler+6>: cbz    r0, 0x1000016a
..
```

This `_reset_handler` snippet is the very first C-runtime entry point after reset on the RP2350, and its opening instructions are checking which CPU core is running. The `mov.w r0, #0xd0000000 / ldr r0, [r0]` sequence reads the `SIO_CPUID` register in the RP2350's SIO block, which returns 0 for core 0 and 1 for core 1. The `cbz r0, 1f` means "if this is core 0, branch to label 1," allowing only core 0 to proceed into the full C runtime startup (stack already set by the boot ROM). Core 1 normally sits idle in the boot ROM until explicitly started, so this guard prevents both cores from running the same initialization code and avoiding double-init of data sections, clocks, and peripherals if, for example, a debugger dropped core 1 directly at the ELF entry point.

```

hold_non_core0_in_bootrom:
    // Send back to the ROM to wait for core 0 to launch it.
    ldr r0, =BOOTROM_VTABLE_OFFSET
    b _enter_vtable_in_r0
1:

#if !PICO_RP2040 && PICO_EMBED_XIP_SETUP && !PICO_NO_FLASH
    // Execute boot2 on the core 0 stack (it also gets copied into BOOTRAM due
    // to inclusion in the data copy table below). Note the reference
    // to __boot2_entry_point here is what prevents the .boot2 section from
    // being garbage-collected.
    _copy_xip_setup:
        ldr r1, =__boot2_entry_point
        mov r3, sp
        add sp, #-256
        mov r2, sp
        bl data_cpy
    _call_xip_setup:
        mov r0, sp
        adds r0, #1
        blx r0
        add sp, #256
#endif

    // In a NO_FLASH binary, don't perform .data etc copy, since it's loaded
    // in-place by the SRAM load. Still need to clear .bss
#if !PICO_NO_FLASH
    adr r4, data_cpy_table

    // assume there is at least one entry
1:
    ldmia r4!, {r1-r3}
    cmp r1, #0
    beq 2f
    bl data_cpy
    b 1b
2:
#endif

```

```
(gdb) x/36i 0x10000110
..
0x10000164 <hold_non_core0_in_bootrom>:   mov.w   r0, #0
0x10000168 <hold_non_core0_in_bootrom+4>:   b.n     0x10000150 <_enter_vtable_in_r0>
0x1000016a <hold_non_core0_in_bootrom+6>:   add     r4, pc, #52      @ (adr r4, 0x100001a0)
<data_cpy_table>
0x1000016c <hold_non_core0_in_bootrom+8>:   ldmia  r4!, {r1, r2, r3}
0x1000016e <hold_non_core0_in_bootrom+10>:  cmp     r1, #0
0x10000170 <hold_non_core0_in_bootrom+12>:  beq.n  0x10000178
<hold_non_core0_in_bootrom+20>
0x10000172 <hold_non_core0_in_bootrom+14>:  bl     0x1000019a <data_cpy>
0x10000176 <hold_non_core0_in_bootrom+18>:  b.n    0x1000016c
<hold_non_core0_in_bootrom+8>
0x10000178 <hold_non_core0_in_bootrom+20>:  ldr    r1, [pc, #84]    @ (0x100001d0)
<data_cpy_table+48>
0x1000017a <hold_non_core0_in_bootrom+22>:  ldr    r2, [pc, #88]    @ (0x100001d4)
<data_cpy_table+52>
0x1000017c <hold_non_core0_in_bootrom+24>:  movs   r0, #0
0x1000017e <hold_non_core0_in_bootrom+26>:  b.n    0x10000182 <bss_fill_test>
..
```

This block funnels non-core0 straight back into the Boot ROM and then performs core0's C-runtime staging: the label loads `r0` with `BOOTROM_VTABLE_OFFSET` (in the build you're disassembling it assembles to 0) and immediately branches to `_enter_vtable_in_r0`, which installs the Boot ROM's vector table and jumps into its reset handler so secondary cores wait there until launched by core0; if we're on core0, the code optionally stages and runs the boot2 XIP setup stub on core0's stack (copy via `data_cpy`, then `blx` into it) to bring external flash online, then iterates the `data_cpy_table` with `ldmia r4!, {r1-r3}` until a zero sentinel in `r1`, copying each region described by the triples, and finally loads the `.bss` start/end from the literal pool, sets `r0=0`, and falls through to the `bss` zeroing routine.

```
// Zero out the BSS
ldr r1, =__bss_start__
ldr r2, =__bss_end__
movs r0, #0
b bss_fill_test
bss_fill_loop:
stm r1!, {r0}
bss_fill_test:
cmp r1, r2
bne bss_fill_loop
```

```
(gdb) x/36i 0x10000110
..
0x10000180 <bss_fill_loop>: stmia  r1!, {r0}
0x10000182 <bss_fill_test>: cmp     r1, r2
0x10000184 <bss_fill_test+2>: bne.n  0x10000180 <bss_fill_loop>
..
```

This is the RP2350's standard `.bss` zero-fill loop that runs during C runtime startup to ensure all uninitialized global/static variables start at zero, as required by the C standard. It loads `__bss_start__` into `r1` and `__bss_end__` into `r2`, sets `r0` to zero, then repeatedly executes `stmia r1!, {r0}` to store that zero word into memory and post-increment `r1` to the next word. After each store, it compares `r1` to `r2`; if they're not equal, it branches back to `bss_fill_loop` and continues until the entire `.bss` region is cleared. Once `r1` reaches

__bss_end__, the loop exits and the system can safely enter main with all zero-initialized data in place.

```
platform_entry: // symbol for stack traces
#if PICO_CRT0_NEAR_CALLS && !PICO_COPY_TO_RAM
    bl runtime_init
#endif
```

(gdb) x/36i 0x10000110

```
..
0x10000186 <platform_entry>: ldr    r1, [pc, #80]    @ (0x100001d8 <data_cpy_table+56>)
0x10000188 <platform_entry+2>: blx   r1
0x1000018a <platform_entry+4>: ldr    r1, [pc, #80]    @ (0x100001dc
<data_cpy_table+60>)
0x1000018c <platform_entry+6>: blx   r1
0x1000018e <platform_entry+8>: ldr    r1, [pc, #80]    @ (0x100001e0
<data_cpy_table+64>)
0x10000190 <platform_entry+10>: blx   r1
0x10000192 <platform_entry+12>: bkpt  0x0000
0x10000194 <platform_entry+14>: b.n   0x10000192 <platform_entry+12>
0x10000196 <data_cpy_loop>: ldmia  r1!, {r0}
0x10000198 <data_cpy_loop+2>: stmia r2!, {r0}
0x1000019a <data_cpy>:      cmp    r2, r3
0x1000019c <data_cpy+2>:    bcc.n 0x10000196 <data_cpy_loop>
0x1000019e <data_cpy+4>:    bx     lr
0x100001a0 <data_cpy_table>: subs  r4, r1, r5
0x100001a2 <data_cpy_table+2>: asrs  r0, r0, #32
0x100001a4 <data_cpy_table+4>: lsls  r0, r2, #4
0x100001a6 <data_cpy_table+6>: movs  r0, #0
0x100001a8 <data_cpy_table+8>: lsls  r4, r5, #10
0x100001aa <data_cpy_table+10>: movs  r0, #0
0x100001ac <data_cpy_table+12>: adds  r0, r5, #3
0x100001ae <data_cpy_table+14>: asrs  r0, r0, #32
0x100001b0 <data_cpy_table+16>: movs  r0, r0
0x100001b2 <data_cpy_table+18>: movs  r0, #8
--Type <RET> for more, q to quit, c to continue without paging--
0x100001b4 <data_cpy_table+20>: movs  r0, r0
0x100001b6 <data_cpy_table+22>: movs  r0, #8
0x100001b8 <data_cpy_table+24>: adds  r0, r5, #3
0x100001ba <data_cpy_table+26>: asrs  r0, r0, #32
0x100001bc <data_cpy_table+28>: asrs  r0, r0, #32
0x100001be <data_cpy_table+30>: movs  r0, #8
0x100001c0 <data_cpy_table+32>: asrs  r0, r0, #32
0x100001c2 <data_cpy_table+34>: movs  r0, #8
0x100001c4 <data_cpy_table+36>: movs  r0, r0
0x100001c6 <data_cpy_table+38>: movs  r0, r0
0x100001c8 <data_cpy_table+40>: bx     lr
0x100001ca <data_cpy_table+42>: movs  r0, r0
0x100001cc <data_cpy_table+44>:      @ <UNDEFINED> instruction:
0xed08e000
0x100001d0 <data_cpy_table+48>: lsls  r4, r5, #10
0x100001d2 <data_cpy_table+50>: movs  r0, #0
0x100001d4 <data_cpy_table+52>: lsls  r0, r3, #19
0x100001d6 <data_cpy_table+54>: movs  r0, #0
0x100001d8 <data_cpy_table+56>: asrs  r5, r7, #13
0x100001da <data_cpy_table+58>: asrs  r0, r0, #32
0x100001dc <data_cpy_table+60>: lsls  r5, r6, #8
0x100001de <data_cpy_table+62>: asrs  r0, r0, #32
```

```

0x100001e0 <data_cpy_table+64>:      asrs    r5, r6, #13
0x100001e2 <data_cpy_table+66>:      asrs    r0, r0, #32
0x100001e4 <_init>: push      {r3, r4, r5, r6, r7, lr}
0x100001e6 <_init+2>:                nop
0x100001e8 <register_tm_clones>:      ldr     r3, [pc, #24]    @ (0x10000204
<register_tm_clones+28>)
0x100001ea <register_tm_clones+2>:    ldr     r1, [pc, #28]    @ (0x10000208
<register_tm_clones+32>)
0x100001ec <register_tm_clones+4>:    subs   r1, r1, r3
0x100001ee <register_tm_clones+6>:    asrs   r1, r1, #2
0x100001f0 <register_tm_clones+8>:    it     mi
0x100001f2 <register_tm_clones+10>:   addmi  r1, #1
0x100001f4 <register_tm_clones+12>:   asrs   r1, r1, #1
0x100001f6 <register_tm_clones+14>:   beq.n 0x10000200 <register_tm_clones+24>
0x100001f8 <register_tm_clones+16>:   ldr    r3, [pc, #16]    @ (0x1000020c
<register_tm_clones+36>)
0x100001fa <register_tm_clones+18>:   cbz    r3, 0x10000200 <register_tm_clones+24>
0x100001fc <register_tm_clones+20>:   ldr    r0, [pc, #4]     @ (0x10000204
<register_tm_clones+28>)
0x100001fe <register_tm_clones+22>:   bx     r3
--Type <RET> for more, q to quit, c to continue without paging--
0x10000200 <register_tm_clones+24>:   bx     lr
0x10000202 <register_tm_clones+26>:   nop
0x10000204 <register_tm_clones+28>:   lsls  r4, r5, #10
0x10000206 <register_tm_clones+30>:   movs  r0, #0
0x10000208 <register_tm_clones+32>:   lsls  r4, r5, #10
0x1000020a <register_tm_clones+34>:   movs  r0, #0
0x1000020c <register_tm_clones+36>:   movs  r0, r0
0x1000020e <register_tm_clones+38>:   movs  r0, r0
0x10000210 <frame_dummy>:      push   {r3, lr}
0x10000212 <frame_dummy+2>:      ldr    r3, [pc, #20]    @ (0x10000228 <frame_dummy+24>)
0x10000214 <frame_dummy+4>:      cbz    r3, 0x1000021e <frame_dummy+14>
0x10000216 <frame_dummy+6>:      ldr    r1, [pc, #20]    @ (0x1000022c <frame_dummy+28>)
0x10000218 <frame_dummy+8>:      ldr    r0, [pc, #20]    @ (0x10000230 <frame_dummy+32>)
0x1000021a <frame_dummy+10>:     nop.w
0x1000021e <frame_dummy+14>:     ldmia.w sp!, {r3, lr}
0x10000222 <frame_dummy+18>:     b.w   0x100001e8 <register_tm_clones>
0x10000226 <frame_dummy+22>:     nop
0x10000228 <frame_dummy+24>:     movs  r0, r0
0x1000022a <frame_dummy+26>:     movs  r0, r0
0x1000022c <frame_dummy+28>:     lsls  r0, r2, #18
0x1000022e <frame_dummy+30>:     movs  r0, #0
0x10000230 <frame_dummy+32>:     adds  r4, r1, r7
0x10000232 <frame_dummy+34>:     asrs  r0, r0, #32
..

```

In the final linked binary, `platform_entry` has been expanded far beyond the single `b1 runtime_init` you see in `cr0.S` as the compiler and linker have transformed that into a small call sequence that loads three function pointers from a nearby literal pool and calls them in turn. Those pointers, stored at `data_cpy_table+56`, `+60`, and `+64`, are filled in at link time with whatever initialization routines the Pico SDK and GCC's C runtime require. In a typical build, they correspond to the SDK's `runtime_init`, the standard `__libc_init_array` for running C++ constructors, and finally your application's `main` (or a wrapper). Using `ldr/blx` through a literal pool instead of a direct `b1` allows the linker to insert any combination of functions, handle long call distances, and keep the assembly source minimal.

Immediately after `platform_entry` is the `data_cpy` routine, a generic word-copy loop used earlier in startup to populate RAM sections from flash or other sources. It works by loading a word from the source pointer in `r1`, storing it to the destination in `r2`, and looping until `r2` reaches the end address in `r3`. The label `data_cpy_table` that follows is not actually executable code, it's a block of constants the startup code uses. The first part holds triples of (source, destination, end) addresses for each region that needs copying. Later entries include other constants such as the VTOR register address (`0xE000ED08` in little-endian form) and the `.bss` bounds, as well as the three function pointers used by `platform_entry`. GDB's disassembler shows these raw words as nonsensical Thumb instructions because it doesn't know they're data.

After this data region come a few standard GCC/EABI stubs: `_init`, `register_tm_clones`, and `frame_dummy`. These are pulled in automatically by the toolchain. `_init` is a hook for pre-main setup, often empty in embedded builds. The `register_tm_clones` and `frame_dummy`, are part of GCC's support for transactional memory and exception frame registration; on bare-metal targets they usually do nothing but are still linked in. Together, this sequence shows how a minimal assembly entry point in `crt0.S` grows into a fully linked startup chain, with the linker and runtime glue inserting the necessary initialization calls, memory setup routines, and housekeeping code before your program ever reaches `main`.

Chapter 5: Intro To Variables

In this chapter we are going to introduce the concept of a variable. If we have a series of boxes all laid out in a row and we numbered them from 0 to 9 (we start with 0 in Engineering) and then placed item 0 in box 0 and then item 1 in box 1 all the way to item 9 in box 9.

The boxes in this analogy represents our SRAM. The items are nothing more than variables of different types, which we will discuss later, that are stored in each of these addresses.

For the Developer, you simply provide a type and a name and the compiler will assign to the value to an actual address.

One of the most important considerations is that you have to declare variables before you use them in a program.

The process of declaration provides the compiler the size and name of the variable you are creating.

The process of definition allocates memory to a variable. These two processes are usually done at the same time.

Let's look at some code.

```
uint8_t age;
```

Here we have a data type which is `uint8_t` and the name of the variable which is `age`.

The data type determines how much space a variable is going to occupy in memory. This will signal the compiler to allocate space for it.

A semicolon signals to the compiler that a statement is complete. In our case the statement was the `uint8_t age`.

The `uint8_t` type takes up 1 byte of memory it is an unsigned integer type that can store a value between 0 and 255.

If you declare a value during declaration it is referred to as initialization.

Let's open up our folder **0x0005_intro-to-variables**.

Now let's review our **0x0005_intro-to-variables.c** file as this is located in the main folder.

```
#include <stdio.h>
#include "pico/stdlib.h"

int main(void) {
    uint8_t age = 42;

    age = 43;

    stdio_init_all();

    while (true)
        printf("age: %d\r\n", age);
}
```

Let's flash the uf2 file onto the Pico 2. If you are unsure about this step, please take a look at Chapter 1 to get re-familiar with this process.

The first lines you should be familiar with and if not again refer to Chapter 1 to get re-familiar with those lines.

Let's break down this code.

```
uint8_t age = 42;
```

We start by declaring and initializing the variable to hold a 1-byte unsigned integer and assign the value of 42 to it.

```
age = 43;
```

We then change the value stored in `age` to 43.

Then inside the while loop we have a `printf` where we print text to indicate that we are going to print the age and then use what we refer to as a format specifier which is `%d` to indicate we are using a decimal value and then our new line chars `\r\n` and then we have the value that will populate `%d` which is 43.

Let's open up PuTTY or your terminal editor of choice and we will see our values being printed in an infinite loop.

Chapter 6: Debugging Intro To Variables

Today we debug!


We will start with Ghidra.

Open up a terminal and run **ghidraRun** and when the window appears, we will select **File, New Project, Non-Shared Project, Next**, and create a **Project Name**. Here we will call it **0x0005_intro-to-variables** and press **Finish**.


Open the file explorer and navigate to the **Embedded-Hacking** folder and drag-and-drop the **0x0005_intro-to-variables.bin** file into the folder within the Ghidra application panel.

In the small window that appears, you will see the file identified as a BIN, which is a binary format without symbols. We will be using the BIN format going forward as this is what we would normally see in the wild so there will be additional setup required based on what we have learned so far.

The window will show a Raw Binary format. Here we click on the three dots to the right of Language and search for Cortex. We want to select Cortex little endian default and click **Ok**.



Click on the **Options...** button. Change the Block Name to **.text** and the base address to **XIP** which is **10000000** hex and click **Ok**.



Let's double-click on the file within the window.

Finally click the auto-analyze and let's begin reviewing the binary.

Let's look at the Functions in the Symbol Tree.



Remember back to Chapter 4, what function existed at 0x1000019a?

The answer is `data_cpy`, so now we can resolve this symbol in Ghidra.

Click on `FUN_1000019a`, in the Decompile view, click on the function name and right-click and select **Edit Function Signature**.

Edit Function at 1000019a

```
void FUN_1000019a (undefined4 param_1, undefined4 * param_2, undefined4 * param_3, undefined4 * param_4)
```

Function Name:

Calling Convention:

Function Attributes:

- Varargs
- In Line
- No Return
- Use Custom Storage

Function Return/Parameters

Index	Datatype	Name	Storage
	void	<RETURN>	<VOID>
1	undefined4	param_1	r0:4
2	undefined4 *	param_2	r1:4
3	undefined4 *	param_3	r2:4
4	undefined4 *	param_4	r3:4

Call Fixup:

Commit all return/parameter details

OK Cancel

Update this to `data_cpy` then click **Ok**.

Edit Function at 1000019a

```
void data_cpy (undefined4 param_1, undefined4 * param_2, undefined4 * param_3, undefined4 * param_4)
```

Function Name:

Calling Convention:

Function Attributes:

- Varargs
- In Line
- No Return
- Use Custom Storage

Function Return/Parameters

Index	Datatype	Name	Storage
	void	<RETURN>	<VOID>
1	undefined4	param_1	r0:4
2	undefined4 *	param_2	r1:4
3	undefined4 *	param_3	r2:4
4	undefined4 *	param_4	r3:4

Call Fixup:

Commit all return/parameter details

OK Cancel

In Chapter 4, what was the function at `FUN_10000210`.

The answer is `frame_dummy` so let's update that function as well then click **Ok**.

undefined4 `frame_dummy` (undefined4 `param_1`)

Function Name:

Calling Convention:

Function Attributes:

- Varargs
- In Line
- No Return
- Use Custom Storage

Function Return/Parameters

Index	Datatype	Name	Storage
	undefined4	<RETURN>	r0:4
1	undefined4	param_1	r0:4

Call Fixup:

Commit all return/parameter details

The final function we will resolve is main then click **Ok**.

Edit Function at 10000234

```
int main (void)
```

Function Name:

Calling Convention:

Function Attributes:

- Varargs
- In Line
- No Return
- Use Custom Storage


Function Return/Parameters

Index	Datatype	Name	Storage
	int	<RETURN>	r0:4

Call Fixup:

Commit all return/parameter details


Let's review the assembler and decompile views.



The screenshot shows a debugger interface with two panels. The left panel displays assembly code with addresses and instructions. The right panel shows the decompiled C code. The assembly code includes instructions like `push {r3,lr}`, `bl FUN_10002f54`, `movs r1,#0x2b`, `ldr r0=>s_age:_fd_100034a0,[DAT_10000244]`, `bl FUN_100030e4`, and `b LAB_1000023a`. The decompiled code shows a `main` function that calls `FUN_10002f54()`, `FUN_100030e4(DAT_10000244, 0x2b)`, and enters a `while(true)` loop.

We see two more functions that need to be resolved. The first one is the Pico C SDK `stdio_init_all`.

If we review our source code, we see that the function returns a `bool`.




The screenshot shows a code editor with the following C code:

```
1 #include <stdio.h>
2 #inc
3
4 int
5
6
7
8
9 stdio_init_all();
10
11 while (true)
12 |     printf("age: %d\r\n", age);
13 }
14
```

A tooltip is displayed over the `stdio_init_all()` call, containing the following information:

- bool stdio_init_all(void)**
- Initialize all of the present standard stdio types that are linked into the binary.
- Returns:** true if at least one output was successfully initialized, false otherwise. `stdio_uart`, `stdio_usb`, `stdio_semihosting`, `stdio_rtt`
- [Generate Copilot summary](#)


Therefore, we need to update accordingly and click **Ok**.



The other function we have to resolve is `printf`.

```
int printf(const char *__restrict__, ...)  
*****  
Global functions, printf  
*****  
*****  
printf()  
Function description  
print a formatted string using RTT and SEGGER RTT formatting.  
Generate Copilot summary
```

Here you want to select `Varargs` for variadic args as `printf` can take any number of args and click **Ok**.



Let's reexamine our assembler and de-compilation.

```
*****...
*                               FUNCTION                               ...
*****...
int __stdcall main(void)
int r0:4 <RETURN>
main+1 XREF[1,1]: 1000018c(c), 1000018a(*)
main
10000234 08 b5 push {r3,lr}
10000236 02 f0 8d fe bl stdio_init_all bool stdio_init_all(void)
LAB_1000023a XREF[1]: 10000242(j)
1000023a 2b 21 movs r1,#0x2b
1000023c 01 48 ldr r0=>s_age:_%d_100034a0,[DAT_10000244] = "age: %d\r\n"
= 100034A0h
1000023e 02 f0 51 ff bl printf int printf(...)
10000242 fa e7 b LAB_1000023a
```

```
1
2 /* WARNING: Heritage AFTER dead
3 /* WARNING: Restarted to delay
4
5 int main(void)
6
7 {
8     undefined4 in_r0;
9
10    stdio_init_all();
11    do {
12        printf(in_r0,0x2b);
13    } while( true );
14 }
15
```

We know that 0x2b in hex is 43. We can always double-check with the ascii table that we have worked with previously.

Take note that the initialization of `uint8_t age = 42` was optimized out by the compiler so we are only seeing 43 which the original code was `age = 43`.

In our next lesson we will hack this!

Chapter 7: Hacking Intro To Variables

Let's continue where we were in Ghidra from our last chapter.


```
*****
*                               FUNCTION                               ...
*****
int __stdcall main(void)
r0:4 <RETURN>
main+1 XREF[1,1]: 1000018e(c), 1000018a(*)
main
10000234 08 b5 push {r3,lr}
10000236 02 f0 8d fe bl stdio_init_all bool stdio_init_all(void)
LAB_1000023a XREF[1]: 10000242(j)
1000023a 2b 21 movs r1,#0x2b
1000023c 01 48 ldr r0=>s_age:_%d_100034a0,[DAT_10000244] = "age: %d\r\n"
1000023e 02 f0 51 ff bl printf = 100034A0h
10000242 fa e7 b int printf(...)
int printf(...)
```

Let's hack 0x2b to 0x46! Highlight 0x2b and right-click and select **Patch Instruction**.


```
LAB_1000023a XREF[1]: 10000242(j)
1000023a 46 21 movs r1,#0x46
```

Let's export the hacked bin.

```
Listing: 0x0005_intro-to-variables.bin
*****
*                               FUNCTION                               ...
*****
int __stdcall main(void)
r0:4 <RETURN>
main+1 XREF[1,1]: 1000018e(c), 1000018a(*)
main
10000234 08 b5 push {r3,lr}
10000236 02 f0 8d fe bl stdio_init_all bool stdio_init_all(void)
LAB_1000023a XREF[1]: 10000242(j)
1000023a 2b 21 movs r1,#0x2b
1000023c 01 48 ldr r0=>s_age:_%d_100034a0,[DAT_10000244] = "age: %d\r\n"
1000023e 02 f0 51 ff bl printf = 100034A0h
10000242 fa e7 b int printf(...)
int printf(...)
```



Select **Raw Bytes** as a **Format** and put the file in the **0x0005_intro-to-variables** bin directory and name the new bin **0x0005_intro-to-variables-h.bin** and click **Ok**.



We need to use a tool to convert this hacked binary into the UF2 format.

```
python ..\uf2conv.py build\0x0005_intro-to-variables-h.bin --base 0x10000000 --family 0xe48bff59 --output build\hacked.uf2
```

After flashing the **hacked.uf2** to the Pico 2, we see the following in the serial terminal.



Boom! We hacked it!

In the coming chapters we will see a great deal of repetition so to some of you this may be a bit boring however to the majority I hope this helps to reinforce techniques that will help you beyond this course as an embedded reverse engineer.

Chapter 8: Uninitialized Variables

In this chapter we are going to examine what happens in memory when we create variables that are not initialized.

We will also introduce the RP2350 GPIO or general-purpose input/output by toggling our red LED on GPIO16.

Let's open up our folder **0x0008_uninitialized-variables**.

Now let's review our **0x0008_uninitialized-variables.c** file as this is located in the main folder.

```
#include <stdio.h>
#include "pico/stdlib.h"

#define LED_PIN 16

int main(void) {
    uint8_t age;

    stdio_init_all();

    gpio_init(LED_PIN);
    gpio_set_dir(LED_PIN, GPIO_OUT);

    while (true) {
        printf("age: %d\r\n", age);

        gpio_put(LED_PIN, 1);
        sleep_ms(500);

        gpio_put(LED_PIN, 0);
        sleep_ms(500);
    }
}
```

Let's flash the uf2 file onto the Pico. If you are unsure about this step, please take a look at Chapter 1 to get re-familiar with this process.

The only difference is that we have no idea what the value will be inside of age or do we?

In other versions of C you would see garbage data if a value is uninitialized however what we see in the C Pico SDK is that like other modern compilers, if you have a value that is not initialized, it will get assigned to the `.bss` section of memory.

The entire `.bss` section is assigned an address in RAM via the linker and does not reside in the binary or flash.

When the Pico boots, behind the scenes `memset` which is a C standard lib function is zeroing out the entire `.bss` so this is why these values are in fact 0.

When you initialize a variable, it will go into the .data section.

When you initialize a constant it will go into the .rodata section.

Let's look at what `stdio_init_all` is behind the scenes.

```
bool stdio_init_all(void) {
    // todo add explicit custom, or registered although you can call stdio_enable_driver
    explicitly anyway
    // These are well known ones

    bool rc = false;
#ifdef LIB_PICO_STDIO_UART
    stdio_uart_init();
    rc = true;
#endif

#ifdef LIB_PICO_STDIO_SEMIHOSTING
    stdio_semihosting_init();
    rc = true;
#endif

#ifdef LIB_PICO_STDIO_RTT
    stdio_rtt_init();
    rc = true;
#endif

#ifdef LIB_PICO_STDIO_USB
    rc |= stdio_usb_init();
#endif
    return rc;
}
```

The `gpio_init` function prepares the chosen pin for use, and `gpio_set_dir` configures it as an output so it can drive the LED. Inside the main loop, `gpio_put` is called with a value of 1 to switch the LED on and with 0 to switch it off. A call to `sleep_ms` is added between these operations to create a visible delay, producing the familiar blink effect at a human-perceivable rate.

Let's review our other 4 functions within the Pico C SDK.

```
void gpio_init(uint gpio) {
    gpio_set_dir(gpio, GPIO_IN);
    gpio_put(gpio, 0);
    gpio_set_function(gpio, GPIO_FUNC_SIO);
}
```

```

static inline void gpio_set_dir(uint gpio, bool out) {
#if PICO_USE_GPIO_COPROCESSOR
    gpioc_bit_oe_put(gpio, out);
#elif PICO_RP2040 || NUM_BANK0_GPIOS <= 32
    uint32_t mask = 1ul << gpio;
    if (out)
        gpio_set_dir_out_masked(mask);
    else
        gpio_set_dir_in_masked(mask);
#else
    uint32_t mask = 1u << (gpio & 0x1fu);
    if (gpio < 32) {
        if (out) {
            sio_hw->gpio_oe_set = mask;
        } else {
            sio_hw->gpio_oe_clr = mask;
        }
    } else {
        if (out) {
            sio_hw->gpio_hi_oe_set = mask;
        } else {
            sio_hw->gpio_hi_oe_clr = mask;
        }
    }
}
#endif
}

```

```

static inline void gpio_put(uint gpio, bool value) {
#if PICO_USE_GPIO_COPROCESSOR
    gpioc_bit_out_put(gpio, value);
#elif NUM_BANK0_GPIOS <= 32
    uint32_t mask = 1ul << gpio;
    if (value)
        gpio_set_mask(mask);
    else
        gpio_clr_mask(mask);
#else
    uint32_t mask = 1ul << (gpio & 0x1fu);
    if (gpio < 32) {
        if (value) {
            sio_hw->gpio_set = mask;
        } else {
            sio_hw->gpio_clr = mask;
        }
    } else {
        if (value) {
            sio_hw->gpio_hi_set = mask;
        } else {
            sio_hw->gpio_hi_clr = mask;
        }
    }
#endif
}


```

```

void sleep_ms(uint32_t ms) {
    sleep_us(ms * 1000ull);
}

```

Let's flash and examine the binary. We also see the red LED blinking.



In our next lesson we will debug this.

Chapter 9: Debugging Uninitialized Variables

Today we debug!

We will start with Ghidra.

Open up a terminal and run **ghidraRun** and when the window appears, we will select **File, New Project, Non-Shared Project, Next**, and create a **Project Name**. Here we will call it **0x0008_uninitialized-variables** and press **Finish**.

Open the file explorer and navigate to the **Embedded-Hacking** folder and drag-and-drop the **0x0008_uninitialized-variables.bin** file into the folder within the Ghidra application panel.

In the small window that appears, you will see the file identified as a BIN, which is a binary format without symbols. We will be using the BIN format going forward as this is what we would normally see in the wild so there will be additional setup required based on what we have learned so far.

The window will show a Raw Binary format. Here we click on the three dots to the right of language and search for Cortex. We want to select Cortex little endian default and click **Ok**.

We will skip all of the Ghidra setup as these are detailed in Chapter 6.

First, we need to set up our Cortex little-endian and options to the `.text` section to `0x10000000`.

We then auto-analyze the binary and set up the memory map as well.

We then update our function signature of `int main(void)` at `FUN_10000234`.

We then update our function signature of `bool stdio_init_all(void)` at `FUN_100030cc`.

We then update our function signature of `void gpio_init(uint gpio)` at `FUN_100002b4`.

```
10000240 4f f0 01 05    mov.w    r5,#0x1
10000244 10 23            movs     r3,#0x10
10000246 45 ec 44 30    mcrr     p0,0x4,r3,r5,cr4
```

The `gpioc_bit_out_put` is a tiny, always-inlined helper that atomically sets or clears a single GPIO by emitting a coprocessor instruction: it calls `pico_default_asm_volatile("mcrr p0, #4, %0, %1, c0" :: "r"(pin), "r"(val))`, passing the pin number and the boolean value to the RP2 GPIO coprocessor; the effect is equivalent to `"if (val) gpioc_hilo_out_set(1ull << pin); else gpioc_hilo_out_clr(1ull << pin)"`, so a true value sets the pin, false clears it, and the operation happens in one atomic coprocessor-backed cycle.

```

static inline void gpio_set_dir(uint gpio, bool out) {
#if PICO_USE_GPIO_COPROCESSOR
    gpioc_bit_oe_put(gpio, out);
#elif PICO_RP2040 || NUM_BANK0_GPIOS <= 32
    uint32_t mask = 1ul << gpio;
    if (out)
        gpio_set_dir_out_masked(mask);
    else
        gpio_set_dir_in_masked(mask);
#else
    uint32_t mask = 1u << (gpio & 0x1fu);
    if (gpio < 32) {
        if (out) {
            sio_hw->gpio_oe_set = mask;
        } else {
            sio_hw->gpio_oe_clr = mask;
        }
    } else {
        if (out) {
            sio_hw->gpio_hi_oe_set = mask;
        } else {
            sio_hw->gpio_hi_oe_clr = mask;
        }
    }
}
#endif
}

```

We have worked with Ghidra and GDB. Let's take another perspective working with our GPIO and LED by first looking at that raw code. We have **0x0008_uninitialized-variables-a.c** so let's review.

```

#include <stdio.h>
#include "pico/stdlib.h"

#define LED_PIN 16

int main(void)
{
    gpio_init(LED_PIN);
    gpio_set_dir(LED_PIN, GPIO_OUT);

    while (true) {
        gpio_put(LED_PIN, 1);
        sleep_ms(500);

        gpio_put(LED_PIN, 0);
        sleep_ms(500);
    }
}

```

Let's take a step deeper.

We have **0x0008_uninitialized-variables-b.c** so let's review.

```
#include <stdio.h>
#include "pico/stdlib.h"

#define LED_PIN 16

int main(void)
{
    // gpio_init(LED_PIN);
    gpio_set_dir(LED_PIN, GPIO_IN);
    gpio_put(LED_PIN, 0);
    gpio_set_function(LED_PIN, GPIO_FUNC_SIO);

    // gpio_set_dir(LED_PIN, GPIO_OUT);
    gpioc_bit_oe_put(LED_PIN, GPIO_OUT);

    while (true) {
        // gpio_put(LED_PIN, 1);
        gpioc_bit_out_put(LED_PIN, 1);
        // sleep_ms(500);
        sleep_us(500 * 1000ull);

        // gpio_put(LED_PIN, 0);
        gpioc_bit_out_put(LED_PIN, 0);
        // sleep_ms(500);
        sleep_us(500 * 1000ull);
    }
}
```

Here we see some additional internal sdk functions. Let's dig deeper to have a better understanding of what is going on.

We have **0x0008_uninitialized-variables-c.c** so let's review.

```
#include <stdio.h>
#include "pico/stdlib.h"

#define LED_PIN 16

int main(void)
{
    // gpio_init(LED_PIN);
    // gpio_set_dir(LED_PIN, GPIO_IN);
    gpioc_bit_oe_put(LED_PIN, GPIO_OUT);
    // gpio_put(LED_PIN, 0);
    gpioc_bit_out_put(LED_PIN, 0);
    // gpio_set_function(LED_PIN, GPIO_FUNC_SIO);
    hw_write_masked(&pads_bank0_hw->io[LED_PIN],
                   PADS_BANK0_GPIO0_IE_BITS,
                   PADS_BANK0_GPIO0_IE_BITS | PADS_BANK0_GPIO0_OD_BITS
    );
    io_bank0_hw->io[LED_PIN].ctrl = GPIO_FUNC_SIO << IO_BANK0_GPIO0_CTRL_FUNCSEL_LSB;
    hw_clear_bits(&pads_bank0_hw->io[LED_PIN], PADS_BANK0_GPIO0_ISO_BITS);

    // gpio_set_dir(LED_PIN, GPIO_OUT);
    gpioc_bit_oe_put(LED_PIN, GPIO_OUT);

    while (true) {
        // gpio_put(LED_PIN, 1);
        gpioc_bit_out_put(LED_PIN, 1);
        // sleep_ms(500);
        sleep_us(500 * 1000ull);

        // gpio_put(LED_PIN, 0);
        gpioc_bit_out_put(LED_PIN, 0);
        // sleep_ms(500);
        sleep_us(500 * 1000ull);
    }
}
```

Here we see some more lower-level activity that we will review at the assembler level.

We have **0x0008_uninitialized-variables-b.d** so let's review.

```
#include <stdio.h>
#include "pico/stdlib.h"

#define LED_PIN 16

int main(void)
{
    // gpio_init(LED_PIN);
    // gpio_set_dir(LED_PIN, GPIO_IN);
    // gpio_set_dir(LED_PIN, GPIO_OUT);
    pico_default_asm_volatile ("mcr p0, #4, %0, %1, c4" : : "r" (LED_PIN), "r"
(GPIO_OUT));
    // gpio_put(LED_PIN, 0);
    // gpio_set_dir(LED_PIN, GPIO_OUT);
    pico_default_asm_volatile ("mcr p0, #4, %0, %1, c4" : : "r" (LED_PIN), "r"
(GPIO_OUT));
    // gpio_set_function(LED_PIN, GPIO_FUNC_SIO);
    // hw_write_masked(&pads_bank0_hw->io[LED_PIN],
    //                 PADS_BANK0_GPIO0_IE_BITS,
    //                 PADS_BANK0_GPIO0_IE_BITS | PADS_BANK0_GPIO0_OD_BITS
    //                 );
    // hw_xor_bits(addr, (*addr ^ values) & write_mask);
    pico_default_asm_volatile (
        "ldr r2, [%0]\n"           // load current pad register
        "eor r2, r2, %1\n"        // xor with IE bit
        "and r2, r2, %2\n"        // mask with (IE|OD)
        "eor r2, r2, %1\n"        // recombine (hw_xor_bits logic)
        "str r2, [%0]\n"          // write back
        :
        : "r" (&pads_bank0_hw->io[LED_PIN]),
        "r" (PADS_BANK0_GPIO0_IE_BITS),
        "r" (PADS_BANK0_GPIO0_IE_BITS | PADS_BANK0_GPIO0_OD_BITS)
        : "r2", "memory"
    );

    // io_bank0_hw->io[LED_PIN].ctrl = GPIO_FUNC_SIO << IO_BANK0_GPIO0_CTRL_FUNCSEL_LSB;
    pico_default_asm_volatile (
        "str %1, [%0]\n"
        :
        : "r" (&io_bank0_hw->io[LED_PIN].ctrl),
        "r" (GPIO_FUNC_SIO << IO_BANK0_GPIO0_CTRL_FUNCSEL_LSB)
        : "memory"
    );
    // hw_clear_bits(&pads_bank0_hw->io[LED_PIN], PADS_BANK0_GPIO0_ISO_BITS);
    pico_default_asm_volatile (
        "ldr r2, [%0]\n"           // load current register value
        "bic r2, r2, %1\n"        // clear the ISO bits (bit clear)
        "str r2, [%0]\n"          // write back
        :
        : "r" (&pads_bank0_hw->io[LED_PIN]),
        "r" (PADS_BANK0_GPIO0_ISO_BITS)
        : "r2", "memory"
    );
};
```

```

// gpio_set_dir(LED_PIN, GPIO_OUT);
/// gpioc_bit_oe_put(LED_PIN, GPIO_OUT);
pico_default_asm_volatile ("mcr p0, #4, %0, %1, c4" :: "r" (LED_PIN), "r"
(GPIO_OUT));

while (true) {
// gpio_put(LED_PIN, 1);
/// gpioc_bit_out_put(LED_PIN, 1);
pico_default_asm_volatile ("mcr p0, #4, %0, %1, c0" :: "r" (LED_PIN), "r" (1));
/// sleep_ms(500);
sleep_us(500 * 1000ull);

// gpio_put(LED_PIN, 0);
/// gpioc_bit_out_put(LED_PIN, 0);
pico_default_asm_volatile ("mcr p0, #4, %0, %1, c0" :: "r" (LED_PIN), "r" (0));
/// sleep_ms(500);
sleep_us(500 * 1000ull);
}
}

```

Here we start to dive into assembler, let's review in our next and final deeper dive.

We have **0x0008_uninitialized-variables-e.e** so let's review.

```
int main(void) {
    __asm__ volatile (
        // gpio_init(LED_PIN);
        /// gpio_set_dir(LED_PIN, GPIO_IN);
        //// gpioc_bit_oe_put(LED_PIN, GPIO_OUT);
        "movs r4, #0x10\n"           // GPIO16
        "movs r5, #0x01\n"         // bit 1; used for OUT/OE writes
        "mcr r p0, #4, r4, r5, c4\n" // gpioc_bit_oe_put(16, 1); p102

        // gpio_set_function(LED_PIN, GPIO_FUNC_SIO);
        /// hw_write_masked(&pads_bank0_hw->io[LED_PIN],
        ///                 PADS_BANK0_GPIO0_IE_BITS,
        ///                 PADS_BANK0_GPIO0_IE_BITS | PADS_BANK0_GPIO0_OD_BITS
        /// );
        //// hw_xor_bits(addr, (*addr ^ values) & write_mask);
        "ldr r3, =0x40038044\n"     // &pads_bank0_hw->io[16]; p785, p796
        "ldr r2, [r3]\n"           // load current config
        "bic r2, r2, #0x80\n"      // clear OD; output disable
        "orr r2, r2, #0x40\n"      // set IE; enable input buffer
        "str r2, [r3]\n"           // store updated config
        /// io_bank0_hw->io[LED_PIN].ctrl = GPIO_FUNC_SIO <<
        IO_BANK0_GPIO0_CTRL_FUNCSEL_LSB;
        "ldr r3, =0x40028084\n"     // &io_bank0_hw->io[16].ctrl; p603, p637
        "ldr r2, [r3]\n"           // load current config
        "bic r2, r2, #0x1f\n"      // clear FUNCSEL bits [4:0]
        "orr r2, r2, #5\n"         // set FUNCSEL = 5 (SIO)
        "str r2, [r3]\n"           // store updated config
        /// hw_clear_bits(&pads_bank0_hw->io[gpio], PADS_BANK0_GPIO0_ISO_BITS);
        "ldr r3, =0x40038044\n"     // &pads_bank0_hw->io[16]; p785, p796
        "ldr r2, [r3]\n"           // load current config
        "bic r2, r2, #0x100\n"     // clear ISO bit (bit 8) un-isolate pad
        "str r2, [r3]\n"           // store updated config

        // gpio_set_dir(LED_PIN, GPIO_OUT);
        /// gpioc_bit_oe_put(LED_PIN, GPIO_OUT);
        "movs r4, #0x10\n"           // GPIO16
        "movs r5, #0x01\n"         // bit 1; used for OUT/OE writes
        "mcr r p0, #4, r4, r5, c4\n" // gpioc_bit_oe_put(16, 1); p102

        // while (true)
        "1:\n"                       // loop start

        // gpio_put(LED_PIN, 1);
        /// gpioc_bit_out_put(LED_PIN, 1);
        "movs r4, #0x10\n"           // GPIO16
        "movs r5, #0x01\n"         // bit 1; used for OUT/OE writes
        "mcr r p0, #4, r4, r5, c0\n" // gpioc_bit_out_put(16, 1)
        // sleep_ms(500);
        /// sleep_us(500 * 1000ull);
        "ldr r2, =0x17D7840\n"     // r2 = ~8.4M cycles
        "2:\n"                       // delay loop
        "subs r2, r2, #1\n"         // decrement counter
        "bne 2b\n"                 // repeat until zero
    );
}
```

```

// gpio_put(LED_PIN, 1);
/// gpioc_bit_out_put(LED_PIN, 1);
"movs r4, #0x10\n"           // GPIO16
"movs r5, #0x00\n"           // bit 0; used for OUT/OE writes
"mcrp p0, #4, r4, r5, c0\n"  // gpioc_bit_out_put(16, 0)
// sleep_ms(500);
/// sleep_us(500 * 1000ull);
"ldr r2, =0x17D7840\n"       // r2 = ~8.4M cycles
"3:\n"                        // delay loop
"subs r2, r2, #1\n"          // decrement counter
"bne 3b\n"                    // repeat until zero

// jmp
"b 1b\n"                      // repeat forever
);
}

```

This code demonstrates a complete GPIO blink implementation for the RP2350 microcontroller using inline assembly with GPIO coprocessor instructions. The program initializes GPIO pin 16 as an output and creates an infinite loop that toggles the pin state every 500 milliseconds. The code leverages RP2350-specific features, particularly the GPIO coprocessor accessible through `mcrp p0, #4` instructions, which provides efficient hardware-accelerated GPIO control that wasn't available on the earlier RP2040 chip.

The initialization sequence follows the standard GPIO setup pattern but uses direct register manipulation for maximum performance. First, it configures the pad control register at address `0x40038044` (`PADS_BANK0_BASE + 0x44` for GPIO16) to enable input buffering and disable output disable functionality. Then it sets the GPIO function to SIO (Software I/O) by writing value 5 to the `IO_BANK0` control register at `0x40028084`, and removes pad isolation by clearing bit 8. Finally, it uses the GPIO coprocessor to enable output mode for the pin.

The main blink loop demonstrates precise timing control using a software delay loop calibrated for the RP2350's 150MHz system clock. The delay value `0x17D7840` (approximately 25 million iterations) is calculated to produce a 500ms delay, accounting for the loop overhead of roughly 3 CPU cycles per iteration. The loop alternates between setting GPIO16 high and low using the coprocessor instructions `mcrp p0, #4, r4, r5, c0`, where the coprocessor efficiently handles the bit manipulation without requiring traditional memory-mapped I/O operations. This approach provides deterministic timing and minimal CPU overhead compared to using the standard SDK GPIO functions.


In our next chapter we will hack this.

Chapter 10: Hacking Uninitialized Variables

Today we hack!

We will start with Ghidra.

Let's open our **0x0008_uninitialized-variables** project.



In our last chapter we resolved some of the functions such as `main`.

```
***** ...
*                FUNCTION                ...
***** ...

int __stdcall main(void)
int      r0:4    <RETURN>
main+1   XREF[1,1]: 1000018c (c), 1000018a (*)
main
10000234 38 b5      push   {r3,r4,r5,lr}
10000236 02 f0 49 ff   bl     stdio_init_all          bool stdio_init_all(void)
1000023a 10 20      movs  r0,#0x10
1000023c 00 f0 3a f8   bl     gpio_init             void gpio_init(uint gpio)
10000240 4f f0 01 05   mov.w  r5,#0x1
10000244 10 23      movs  r3,#0x10
10000246 45 ec 44 30   mcrr  p0,0x4,r3,r5,cr4

LAB_1000024a XREF[1]: 10000270 (j)
1000024a 00 21      movs  r1,#0x0
1000024c 09 48      ldr   r0=>s_age:_%d_10003618,[DAT_10000274] = "age: %d\r\n"
                                                = 10003618h
1000024e 03 f0 05 f8   bl     FUN_1000325c          undefined FUN_1000325c()
10000252 10 24      movs  r4,#0x10
10000254 45 ec 40 40   mcrr  p0,0x4,r4,r5,cr0
10000258 4f f4 fa 70   mov.w  r0,#0x1f4
1000025c 00 f0 58 fd   bl     FUN_10000d10          undefined FUN_10000d10()
10000260 4f f0 00 03   mov.w  r3,#0x0
10000264 43 ec 40 40   mcrr  p0,0x4,r4,r3,cr0
10000268 4f f4 fa 70   mov.w  r0,#0x1f4
1000026c 00 f0 50 fd   bl     FUN_10000d10          undefined FUN_10000d10()
10000270 eb e7      b     LAB_1000024a
```

We know that in our last chapter we enabled GPIO 16 and it blinked the red LED.

The first thing we will hack is instead of GPIO 16 we will make GPIO 17 blink.

The first thing we need to patch is the following.

```
1000023a 10 20          movs      r0,#0x10
```

Here we will patch this to 0x11 which is 17 in decimal.

In earlier chapters we went over this in detail so the first thing we will do is **right-click** on the instruction, select **Patch Instruction** and change 0x10 to 0x11.

That corresponds with us calling the `gpio_init(LED_PIN)` code from our source.

We also have to patch the following.

```
10000240 4f f0 01 05      mov.w     r5,#0x1
10000244 10 23          movs     r3,#0x10
10000246 45 ec 44 30      mcrr    p0,0x4 ,r3,r5,cr4
```

Here, we will **right-click** and **Patch Instruction** 0x10 to 0x11 as well.

These lines correspond to `gpio_set_dir(LED_PIN, GPIO_OUT)` code from our source.

Now we are going to enter into our loop.

Let's hack the 0x0 to say 0x42. We will do the same patch instruction as we did earlier.

```
1000024a 00 21          movs     r1,#0x0
```

We also have to patch our GPIO in the below line to 0x11.

```
10000252 10 24          movs     r4,#0x10
```

At this point we can click **File, Export Program, select Format: Raw Bytes** and update our output file to **0x0008_unitialized-variables-h.bin** and click **Ok**.

Finally, we have to convert to a UF2.

```
python ..\uf2conv.py build\0x0008_unitialized-variables-h.bin --base 0x10000000 --
family 0xe48bff59 --output build\hacked.uf2
```

Let's flash our Pico 2 and we will notice the green LED blinking so YAY!

If we open up PuTTY or another terminal program, we will see 0x42 or 66 decimal as well. BOOM!

A screenshot of a PuTTY terminal window titled "COM3 - PuTTY". The window has a black background and displays the text "age: 66" repeated 20 times, one line per iteration, in a white monospaced font. The window's title bar includes a small icon on the left and standard minimize, maximize, and close buttons on the right.

```
age: 66
```

In our next lesson we will cover the integer data type.

Chapter 11: Integer Data Type

In this chapter we are going to discuss the integer data type. We have already covered examples with this however the goal of this course is to continue to reinforce learning so that you have a mastery of the embedded process.

Let's open up our folder **0x000b_integer-data-type**.

Now let's review our **0x000b_integer-data-type.c** file as this is located in the main folder.

```
#include <stdio.h>
#include "pico/stdlib.h"

int main(void) {
    uint8_t age = 43;
    int8_t range = -42;

    stdio_init_all();

    __asm volatile (
        "ldr r3, =0x40038000\n"           // address of PADS_BANK0_BASE
        "ldr r2, =0x40028004\n"         // address of IO_BANK0 GPIO0.ctrl
        "movs r0, #16\n"                // GPIO16 (start pin)

        "init_loop:\n"                  // loop start
        "lsls r1, r0, #2\n"              // pin * 4 (pad offset)
        "adds r4, r3, r1\n"              // PADS base + offset
        "ldr r5, [r4]\n"                 // load current config
        "bic r5, r5, #0x180\n"           // clear OD+ISO
        "orr r5, r5, #0x40\n"            // set IE
        "str r5, [r4]\n"                 // store updated config

        "lsls r1, r0, #3\n"              // pin * 8 (ctrl offset)
        "adds r4, r2, r1\n"              // IO_BANK0 base + offset
        "ldr r5, [r4]\n"                 // load current config
        "bic r5, r5, #0x1f\n"            // clear FUNCSEL bits [4:0]
        "orr r5, r5, #5\n"               // set FUNCSEL = 5 (SIO)
        "str r5, [r4]\n"                 // store updated config

        "mov r4, r0\n"                  // pin
        "movs r5, #1\n"                  // bit 1; used for OUT/OE writes
        "mccrr p0, #4, r4, r5, c4\n"     // gpioc_bit_oe_put(pin,1)
        "adds r0, r0, #1\n"              // increment pin
        "cmp r0, #20\n"                  // stop after pin 18
        "blt init_loop\n"                // loop until r0 == 20
    );
```

```

uint8_t pin = 16;

while (1) {
    __asm volatile (
        "mov r4, %0\n"           // pin
        "movs r5, #0x01\n"      // bit 1; used for OUT/OE writes
        "mccrr p0, #4, r4, r5, c0\n" // gpioc_bit_out_put(16, 1)
        :
        : "r"(pin)
        : "r4", "r5"
    );
    sleep_ms(500);

    __asm volatile (
        "mov r4, %0\n"           // pin
        "movs r5, #0\n"         // bit 0; used for OUT/OE writes
        "mccrr p0, #4, r4, r5, c0\n" // gpioc_bit_out_put(16, 0)
        :
        : "r"(pin)
        : "r4", "r5"
    );
    sleep_ms(500);

    pin++;
    if (pin > 18) pin = 16;

    printf("age: %d\r\n", age);
    printf("range: %d\r\n", range);
}
}

```

Here we have a heavy mix of inline assembler and C. We start off with a `uint8_t` `age = 43` which is an unsigned 8-bit integer which is 43 and an `int8_t` `range = -42` which is 42.

We then init our `stdio_init_all` for the purposes of our UART terminal interface.

Let's take a moment and explain UART. Universal Asynchronous Receiver Transmitter is what we use to communicate with our terminal. Up to this point it has only been in a receive capacity where we are only receiving print statements rather than being interactive.

On the RP2350, the UART is one of the chip's flexible serial interfaces, designed for simple, low-pin-count communication between the microcontroller and external devices such as PCs, sensors, or other MCUs. Each UART block handles full-duplex communication with independent transmit (TX) and receive (RX) FIFOs for efficient data handling. Like other RP-series chips, the RP2350 uses a GPIO muxing system, so any eligible GPIO pin can be mapped to a UART function, giving developers freedom in board layout. The UART supports standard features such as configurable word length, stop bits, parity, and flow control for general-purpose serial communication in embedded applications.

We have GPIO 0 which on the board is our UART0 TX pin connected to our Pico Debug Probe's UART RX pin. In addition, we have our GPIO 1 pin on the board which is UART0 RX pin connected to our Debug Probe's UART TX pin so we can have communication.

We will explore UART in more depth in future chapters.

After we init our `stdio_init_all` which will in our case enable UART communications, we have a larger inline assembler block.

Here we load the hardware addresses of `0x4003800` which is the `PADS_BANK0_BASE` address we saw in the RP2350 datasheet from prior chapters.

```
"ldr r3, =0x40038000\n"           // address of PADS_BANK0_BASE
```

On page 32 of the datasheet, we see it clearly.

RP2350 Datasheet

Bus Endpoint	Base Address
CLOCKS_BASE	0x40010000
PSM_BASE	0x40018000
RESETS_BASE	0x40020000
IO_BANK0_BASE	0x40028000
IO_QSPI_BASE	0x40030000
PADS_BANK0_BASE	0x40038000


If we turn to page 595 of the datasheet, we see more info on pads.

Each GPIO is connected off-chip via a **pad**. Pads are the electrical interface between the chip's internal logic and external circuitry. They translate signal voltage levels, support higher currents and offer some protection against electrostatic discharge (ESD) events. You can adjust pad electrical behaviour to meet the requirements of external circuitry in the following ways:

- Output drive strength can be set to 2mA, 4mA, 8mA or 12mA.
- Output slew rate can be set to slow or fast.
- Input hysteresis (Schmitt trigger mode) can be enabled.
- A pull-up or pull-down can be enabled, to set the output signal level when the output driver is disabled.
- The input buffer can be disabled, to reduce current consumption when the pad is unused, unconnected or connected to an analogue signal.

An example pad is shown in [Figure 42](#).

Figure 42. Diagram of a single IO pad.



The next line of assembler is the address of `IO_BANK0`, `GPIO0_CTRL`.

```
"ldr r2, =0x40028004\n" // address of IO_BANK0 GPIO0.ctrl
```

On page 604-605 of the datasheet, we see this info.

9.11. List of registers

9.11.1. IO - User Bank

The User Bank IO registers start at a base address of `0x40028000` (defined as `IO_BANK0_BASE` in SDK).

Table 649. List of IO_BANK0 registers

Offset	Name	Info
0x000	GPIO0_STATUS	
0x004	GPIO0_CTRL	

Here we see at offset 0x4, the CTRL value we are looking for as we can now dig into this 32-bit wide register on page 610 of the datasheet.

IO_BANK0: GPIO0_CTRL Register

Offset: 0x004

Table 651.
GPIO0_CTRL Register

Bits	Description	Type	Reset
31:30	Reserved.	-	-
29:28	IRQOVER	RW	0x0
	Enumerated values:		
	0x0 → NORMAL: don't invert the interrupt		
	0x1 → INVERT: invert the interrupt		
	0x2 → LOW: drive interrupt low		
	0x3 → HIGH: drive interrupt high		
27:18	Reserved.	-	-
17:16	INOVER	RW	0x0
	Enumerated values:		
	0x0 → NORMAL: don't invert the peri input		
	0x1 → INVERT: invert the peri input		
	0x2 → LOW: drive peri input low		
	0x3 → HIGH: drive peri input high		

Bits	Description	Type	Reset
	0x2 → LOW: drive output low		
	0x3 → HIGH: drive output high		
11:5	Reserved.	-	-
4:0	FUNCSEL : 0-31 → selects pin function according to the gpio table 31 == NULL	RW	0x1f
	Enumerated values:		
	0x00 → JTAG_TCK		
	0x01 → SPI0_RX		
	0x02 → UART0_TX		
	0x03 → I2C0_SDA		
	0x04 → PWM_A_0		
	0x05 → SIO_0		
	0x06 → PIO0_0		
	0x07 → PIO1_0		
	0x08 → PIO2_0		
	0x09 → XIP_SS_N_1		
	0x0a → USB_MUXING_OVERCURRETECT		
	0x1f → NULL		

Here we see all of the bits within the 32-bit register that we can configure.

The next assembler code line is moving the value of 16 into the r0 register.

```
movs r0, #16\n" // GPIO16 (start pin)
```

This is the address of our start LED start pin.

Let's review our assembler.

```

__asm volatile (
    "ldr r3, =0x40038000\n"           // address of PADS_BANK0_BASE
    "ldr r2, =0x40028004\n"         // address of IO_BANK0 GPIO0.ctrl
    "movs r0, #16\n"                // GPIO16 (start pin)

    "init_loop:\n"                  // loop start
    "lsls r1, r0, #2\n"             // pin * 4 (pad offset)
    "adds r4, r3, r1\n"             // PADS base + offset
    "ldr r5, [r4]\n"                // load current config
    "bic r5, r5, #0x180\n"          // clear OD+ISO
    "orr r5, r5, #0x40\n"           // set IE
    "str r5, [r4]\n"                // store updated config

    "lsls r1, r0, #3\n"             // pin * 8 (ctrl offset)
    "adds r4, r2, r1\n"             // IO_BANK0 base + offset
    "ldr r5, [r4]\n"                // load current config
    "bic r5, r5, #0x1f\n"           // clear FUNCSEL bits [4:0]
    "orr r5, r5, #5\n"              // set FUNCSEL = 5 (SIO)
    "str r5, [r4]\n"                // store updated config

    "mov r4, r0\n"                  // pin
    "movs r5, #1\n"                 // bit 1; used for OUT/OE writes
    "mcrnr p0, #4, r4, r5, c4\n"    // gpioctl_oe_put(pin,1)
    "adds r0, r0, #1\n"             // increment pin
    "cmp r0, #20\n"                 // stop after pin 18
    "blt init_loop\n"              // loop until r0 == 20
);

```


This inline assembly routine is walking through GPIO pins 16–18 on our RP2350 and configuring them for direct software control via the SIO (single-cycle I/O) block. For each pin, it first accesses the PADS_BANK0 register set to clear the output-disable/isolation bits and enable the input buffer, ensuring the pad is electrically active. Then it moves to the IO_BANK0 control register for that pin, clears the function select field, and sets it to 5, which maps the pin to the SIO peripheral rather than an alternate function. Finally, it uses a coprocessor register write (`mcrnr`) to enable the output driver for that pin. The loop increments through pins 16–18, so by the end, those three GPIOs are initialized as standard digital I/O lines under SIO control, ready for bit-banging or direct register-driven toggling.

Why do we have a compare to 20 as we are only dealing with GPIO 16-18?

```
"cmp r0, #20\n" // stop after pin 18
```

This comes down to how the loop termination works. The instruction `cmp r0, #20` is paired with `blt init_loop`, which means “branch back while **r0** < 20.” Since you start at `r0 = 16`, the loop runs for 16, 17, 18, 19. That gives you four iterations, and GPIO18 gets configured on the third pass (when `r0 = 18`). If you instead compared against 19, the loop would only run while `r0 < 19`, so it would stop after finishing `r0 = 17` therefore never reaching 18. In other words, the compare value is always set to one past the last pin you want, because the branch condition checks for “less than,” not “less than or equal.”

We see the red, green and yellow LED’s toggling every 500ms and we see in our terminal the values of our integer values printing in the same timeframe.

A screenshot of a PuTTY terminal window titled "COM3 - PuTTY". The terminal output consists of a repeating sequence of two lines: "range: -42" followed by "age: 43". This sequence is repeated 12 times in total. The terminal background is black, and the text is white. The window title bar includes standard minimize, maximize, and close buttons.

```
range: -42
age: 43
range: -42
```

In our next lesson we will debug this.

Chapter 12: Debugging Integer Data Type

In this chapter we are going to discuss debugging the integer data type.

Run OpenOCD with the below config.

```
openocd -f interface/cmsis-dap.cfg -f target/rp2350.cfg -c "adapter speed 5000"
```

Open a new terminal and then run the following to launch our dynamic debugger called GDB.

```
arm-none-eabi-gdb build/0x000b_integer-data-type.bin
```

Once it loads, we need to target our remote server.

```
target remote :3333
```

We need to halt the currently running binary.

```
monitor reset halt
```

We need to review 46 instructions from main. We remember main is 0x10000234.

```
x/46i 0x10000234
0x10000234: push    {r4, r5, r6, lr}
0x10000236: bl     0x100030cc
0x1000023a: ldr   r3, [pc, #124] @ (0x100002b8)
0x1000023c: ldr   r2, [pc, #124] @ (0x100002bc)
0x1000023e: movs  r0, #16
0x10000240: lsls  r1, r0, #2
0x10000242: adds  r4, r3, r1
0x10000244: ldr   r5, [r4, #0]
0x10000246: bic.w r5, r5, #384 @ 0x180
0x1000024a: orr.w r5, r5, #64 @ 0x40
0x1000024e: str   r5, [r4, #0]
0x10000250: lsls  r1, r0, #3
0x10000252: adds  r4, r2, r1
0x10000254: ldr   r5, [r4, #0]
0x10000256: bic.w r5, r5, #31
0x1000025a: orr.w r5, r5, #5
0x1000025e: str   r5, [r4, #0]
0x10000260: mov   r4, r0
0x10000262: movs  r5, #1
0x10000264: mcrr  0, 4, r4, r5, cr4
0x10000268: adds  r0, #1
0x1000026a: cmp   r0, #20
0x1000026c: blt.n 0x10000240
0x1000026e: movs  r6, #16
0x10000270: mov   r4, r6
0x10000272: movs  r5, #1
0x10000274: mcrr  0, 4, r4, r5, cr0
0x10000278: mov.w r0, #500 @ 0x1f4
0x1000027c: bl   0x10000d10
0x10000280: mov   r4, r6
0x10000282: movs  r5, #0
0x10000284: mcrr  0, 4, r4, r5, cr0
0x10000288: adds  r6, #1
0x1000028a: uxtb  r6, r6
0x1000028c: mov.w r0, #500 @ 0x1f4
```

```
--Type <RET> for more, q to quit, c to continue without paging--
0x10000290: bl      0x10000d10
0x10000294: cmp     r6, #19
0x10000296: mov.w  r1, #43 @ 0x2b
0x1000029a: ldr    r0, [pc, #20] @ (0x100002b0)
0x1000029c: it     eq
0x1000029e: moveq  r6, #16
0x100002a0: bl     0x1000325c
0x100002a4: mvn.w  r1, #41 @ 0x29
0x100002a8: ldr    r0, [pc, #8] @ (0x100002b4)
0x100002aa: bl     0x1000325c
0x100002ae: b.n    0x10000270
```

I know the first thought might be one of feeling completely overwhelmed but let's take this step-by-step.

We first know that our source code is raw assembler so a good deal will match up however we are going to take this exact code and go step-by-step to better understand what is happening.

```
0x10000234: push   {r4, r5, r6, lr}
```

We begin by pushing these 4 registers onto the stack. We can break on 0x10000234 and review the stack before and after.

```
(gdb) b *0x10000234
Breakpoint 1 at 0x10000234
Note: automatically using hardware breakpoints for read-only addresses.
(gdb) c
Continuing.
Thread 1 "rp2350.cm0" hit Breakpoint 1, 0x10000234 in ?? ()
(gdb) x/4x $sp
0x20082000: 0x00000000 0x00000000 0x00000000 0x00000000
(gdb) x/i $pc
=> 0x10000234: push   {r4, r5, r6, lr}
```

We can see the stack is empty. We see we are about to execute the push statement so let's first review the values in the registers.

```
(gdb) i r
...
r4      0x100001d0      268435920
r5      0x88526891      -2007865199
r6      0x4f54710       83183376
r7      0x400e0014     1074659348
r8      0x43280035     1126694965
r9      0x0            0
r10     0x10000000     268435456
r11     0x62707361     1651536737
r12     0xa5024200     -1526578688
sp      0x20082000     0x20082000
lr      0x1000018f     268435855
...
```

Let's step into and review the stack.

```
(gdb) si
0x10000236 in ?? ()
(gdb) x/4x $sp
0x20081ff0: 0x100001d0 0x88526891 0x04f54710 0x1000018f
(gdb) x/x $sp
0x20081ff0: 0x100001d0
(gdb) x/x $r4
0x100001d0: 0x2000062c
(gdb) x/x $sp+0x4
0x20081ff4: 0x88526891
(gdb) x/x $r5
0x88526891: Cannot access memory at address 0x88526891
(gdb) x/x $sp+0x8
0x20081ff8: 0x04f54710
(gdb) x/x $r6
0x4f54710: 0x00000000
(gdb) x/x $sp+0xc
0x20081ffc: 0x1000018f
(gdb) x/x $lr
0x1000018f: 0x00478849
```

Here we can see clearly how the push works. Let's look at the next instruction is which will live within the program counter.

```
(gdb) x/i $pc
=> 0x10000236: bl 0x100030cc
```

We know this is our call to `stdio_init_all`. We can simply step over it.

```
(gdb) b *0x1000023a
Breakpoint 2 at 0x1000023a
(gdb) c
Continuing.
Thread 1 "rp2350.cm0" hit Breakpoint 2, 0x1000023a in ?? ()
(gdb) x/i $pc
=> 0x1000023a: ldr r3, [pc, #124] @ (0x100002b8)
```

Let's examine what is located at the memory address of `0x100002b8`.

```
(gdb) x/x 0x100002b8
0x100002b8: 0x40038000
```

This is the address of `PADS_BANK0_BASE` we saw in our source code.

As we step again, we see the value of the address of `IO_BANK0_GPIO0.ctrl` at the next instruction.

```
(gdb) si
0x1000023c in ?? ()
(gdb) x/x 0x100002bc
0x100002bc: 0x40028004
```

As we step again, we see GPIO16, our start pin, moved into r0.

```
(gdb) si
0x1000023e in ?? ()
(gdb) x/i $pc
=> 0x1000023e:  movs    r0, #16
```

The next set of instructions are part of our `init_loop`. This is an identical assembler match to our source code.

```
(gdb) si
0x10000240 in ?? ()
(gdb) x/18i $pc
=> 0x10000240:  lsls    r1, r0, #2
0x10000242:  adds   r4, r3, r1
0x10000244:  ldr    r5, [r4, #0]
0x10000246:  bic.w  r5, r5, #384    @ 0x180
0x1000024a:  orr.w  r5, r5, #64     @ 0x40
0x1000024e:  str    r5, [r4, #0]
0x10000250:  lsls   r1, r0, #3
0x10000252:  adds   r4, r2, r1
0x10000254:  ldr    r5, [r4, #0]
0x10000256:  bic.w  r5, r5, #31
0x1000025a:  orr.w  r5, r5, #5
0x1000025e:  str    r5, [r4, #0]
0x10000260:  mov    r4, r0
0x10000262:  movs   r5, #1
0x10000264:  mcrr   0, 4, r4, r5, cr4
0x10000268:  adds   r0, #1
0x1000026a:  cmp    r0, #20
0x1000026c:  blt.n  0x10000240
```

This loop is systematically initializing GPIO pins 16 through 19 on the RP2350 so they can be driven directly by the SIO block. For each pin, it first calculates the correct `PADS_BANK0` register offset and updates the pad configuration: clearing the output-disable and isolation bits, then enabling the input buffer. Next, it computes the `IO_BANK0` control register offset, clears the function-select field, and sets it to 5, which maps the pin to SIO rather than an alternate peripheral. With the pad and function configured, it then enables the pin's output driver using the `mcrr` instruction (a coprocessor write that acts like `gpio_set_oe(pin, 1)`). Finally, the loop increments the pin number and repeats until r0 reaches 20, which ensures pins 16, 17, 18, and 19 are all configured before exiting.

Let's review the remaining code.

```
(gdb) b *0x1000026e
Breakpoint 3 at 0x1000026e
(gdb) c
Continuing.
Thread 1 "rp2350.cm0" hit Breakpoint 3, 0x1000026e in ?? ()
(gdb) x/23i $pc
=> 0x1000026e:  movs    r6, #16
0x10000270:  mov    r4, r6
0x10000272:  movs   r5, #1
0x10000274:  mcrr   0, 4, r4, r5, cr0
0x10000278:  mov.w  r0, #500        @ 0x1f4
0x1000027c:  bl     0x10000d10
0x10000280:  mov    r4, r6
```

```
0x10000282:  movs    r5, #0
0x10000284:  mcrr    0, 4, r4, r5, cr0
0x10000288:  adds   r6, #1
0x1000028a:  uxtb   r6, r6
0x1000028c:  mov.w  r0, #500          @ 0x1f4
0x10000290:  bl     0x10000d10
0x10000294:  cmp    r6, #19
0x10000296:  mov.w  r1, #43 @ 0x2b
0x1000029a:  ldr    r0, [pc, #20]   @ (0x100002b0)
0x1000029c:  it     eq
0x1000029e:  moveq  r6, #16
0x100002a0:  bl     0x1000325c
0x100002a4:  mvn.w  r1, #41 @ 0x29
0x100002a8:  ldr    r0, [pc, #8]   @ (0x100002b4)
0x100002aa:  bl     0x1000325c
0x100002ae:  b.n   0x10000270
```

This corresponds to our `uint8_t pin = 16` and `while` loop.

```
uint8_t pin = 16;

while (1) {
    __asm volatile (
        "mov r4, %0\n"           // pin
        "movs r5, #0x01\n"      // bit 1; used for OUT/OE writes
        "mcr r0, #4, r4, r5, c0\n" // gpio_cbit_out_put(16, 1)
        :
        : "r"(pin)
        : "r4", "r5"
    );
    sleep_ms(500);

    __asm volatile (
        "mov r4, %0\n"           // pin
        "movs r5, #0\n"         // bit 0; used for OUT/OE writes
        "mcr r0, #4, r4, r5, c0\n" // gpio_cbit_out_put(16, 0)
        :
        : "r"(pin)
        : "r4", "r5"
    );
    sleep_ms(500);

    pin++;
    if (pin > 18) pin = 16;

    printf("age: %d\r\n", age);
    printf("range: %d\r\n", range);
}
}
```

Let's verify this by breaking at the following address and proving `age` is within the `r0` register.

```
(gdb) b *0x1000029a
Breakpoint 4 at 0x1000029a
(gdb) c
Continuing.
Thread 1 "rp2350.cm0" hit Breakpoint 4, 0x1000029a in ?? ()
(gdb) x/4i $pc
=> 0x1000029a: ldr    r0, [pc, #20] @ (0x100002b0)
   0x1000029c: it     eq
   0x1000029e: moveq  r6, #16
   0x100002a0: bl     0x1000325c
(gdb) si
0x1000029c in ?? ()
```

Here we can see within `r0` our string about to be passed to the `printf` function at `0x1000325c`.

```
(gdb) x/s $r0
0x10003618: "age: %d\r\n"
```

Let's open up our PuTTY terminal or screen and see what happens when we break and continue to the next `printf` statement.

```
(gdb) b *0x100002a8
Breakpoint 5 at 0x100002a8
(gdb) c
Continuing.
Thread 1 "rp2350.cm0" hit Breakpoint 5, 0x100002a8 in ?? ()
```

We noticed the red LED flashed and we see `age: 43` within our terminal.



Let's review the last 3 instructions.

```
(gdb) x/3i $pc
=> 0x100002a8: ldr    r0, [pc, #8]    @ (0x100002b4)
    0x100002aa: bl     0x1000325c
    0x100002ae: b.n   0x10000270
```

We know we have an unconditional break of `b.n 0x10000270` which will take us to the top of our loop and continue indefinitely. Let's step again and review `r0`.

```
(gdb) si
0x100002aa in ?? ()
(gdb) x/s $r0
0x10003624:    "range: %d\r\n"
```

We have verified our code now when we continue, we will see the green LED flash and see `range: -42` printed.



In our next chapter we will hack this!

Chapter 13: Hacking Integer Data Type

In this chapter we are going to discuss hacking the integer data type.

Run OpenOCD with the below config.

```
openocd -f interface/cmsis-dap.cfg -f target/rp2350.cfg -c "adapter speed 5000"
```

Open a new terminal and then run the following to launch our dynamic debugger called GDB.

```
arm-none-eabi-gdb build/0x000b_integer-data-type.bin
```

Once it loads, we need to target our remote server.

```
target remote :3333
```

We need to halt the currently running binary.

```
monitor reset halt
```

We need to break at main and review 46 instructions from main. We remember main is 0x10000234.

```
(gdb) b *0x10000234
Breakpoint 1 at 0x10000234
Note: automatically using hardware breakpoints for read-only addresses.
(gdb) c
Continuing.
Thread 1 "rp2350.cm0" hit Breakpoint 1, 0x10000234 in ?? ()
(gdb) x/46i $pc
=> 0x10000234:  push    {r4, r5, r6, lr}
0x10000236:  bl      0x100030cc
0x1000023a:  ldr     r3, [pc, #124] @ (0x100002b8)
0x1000023c:  ldr     r2, [pc, #124] @ (0x100002bc)
0x1000023e:  movs   r0, #16
0x10000240:  lsls   r1, r0, #2
0x10000242:  adds   r4, r3, r1
0x10000244:  ldr    r5, [r4, #0]
0x10000246:  bic.w  r5, r5, #384 @ 0x180
0x1000024a:  orr.w  r5, r5, #64 @ 0x40
0x1000024e:  str    r5, [r4, #0]
0x10000250:  lsls   r1, r0, #3
0x10000252:  adds   r4, r2, r1
0x10000254:  ldr    r5, [r4, #0]
0x10000256:  bic.w  r5, r5, #31
0x1000025a:  orr.w  r5, r5, #5
0x1000025e:  str    r5, [r4, #0]
0x10000260:  mov    r4, r0
0x10000262:  movs   r5, #1
0x10000264:  mcrr  0, 4, r4, r5, cr4
0x10000268:  adds   r0, #1
0x1000026a:  cmp    r0, #20
0x1000026c:  blt.n  0x10000240
0x1000026e:  movs   r6, #16
0x10000270:  mov    r4, r6
0x10000272:  movs   r5, #1
0x10000274:  mcrr  0, 4, r4, r5, cr0
0x10000278:  mov.w  r0, #500 @ 0x1f4
0x1000027c:  bl     0x10000d10
```

```

0x10000280:  mov    r4, r6
0x10000282:  movs   r5, #0
0x10000284:  mcrr   0, 4, r4, r5, cr0
0x10000288:  adds   r6, #1
0x1000028a:  uxtb   r6, r6
0x1000028c:  mov.w  r0, #500          @ 0x1f4
0x10000290:  bl     0x10000d10
0x10000294:  cmp    r6, #19
0x10000296:  mov.w  r1, #43 @ 0x2b
0x1000029a:  ldr    r0, [pc, #20]   @ (0x100002b0)
0x1000029c:  it     eq
0x1000029e:  moveq  r6, #16
0x100002a0:  bl     0x1000325c
0x100002a4:  mvn.w  r1, #41 @ 0x29
0x100002a8:  ldr    r0, [pc, #8]    @ (0x100002b4)
0x100002aa:  bl     0x1000325c
0x100002ae:  b.n    0x10000270

```

Here we can hack 3 things, let's hack the starting LED, and the two integer values.

First, let's set breakpoints on the below.

```

0x1000026e:  movs   r6, #16
0x10000296:  mov.w  r1, #43 @ 0x2b
0x100002a4:  mvn.w  r1, #41 @ 0x29
(gdb) b *0x1000026e
Breakpoint 2 at 0x1000026e
(gdb) b *0x10000296
Breakpoint 3 at 0x10000296
(gdb) b *0x100002a4
Breakpoint 4 at 0x100002a4
(gdb) c
Continuing.
Thread 1 "rp2350.cm0" hit Breakpoint 2, 0x1000026e in ?? ()

```

1st hack, you will hack the green LED to light up.

```

(gdb) x/i $pc
=> 0x1000026e:  movs   r6, #16
(gdb) si
0x10000270 in ?? ()
(gdb) set $r6 = 17
(gdb) c
Continuing.

```

2nd hack, you will change the age to 44.

```

(gdb) x/i $pc
=> 0x10000296:  mov.w  r1, #43 @ 0x2b
(gdb) si
0x1000029a in ?? ()
(gdb) set $r1 = 44
(gdb) c
Continuing.

```

3rd hack, you will change the range to 50.

```
(gdb) x/i $pc
=> 0x100002a4: mvn.w    r1, #41 @ 0x29
(gdb) si
0x100002a8 in ?? ()
(gdb) set $r1 = 50
(gdb) c
Continuing.
```



Boom! We hacked the LED to turn green when it should have been red, we hacked age to 44 when it should have been 43 and we hacked range to 50.

In our next chapter we will discuss the floating-point data type.

Chapter 14: Floating-Point Data Type

In this chapter we are going to discuss the floating-point data type.

Let's open up our folder **0x000e_floating-point-data-type**.

Now let's review our **0x000e_floating-point-data-type.c** file as this is located in the main folder.

```
#include <stdio.h>
#include "pico/stdlib.h"

int main(void) {
    float fav_num = 42.5;


    stdio_init_all();

    while (true)
        printf("fav_num: %f\r\n", fav_num);
}
```

We start off with a `float fav_num = 42.5` which is a 32-bit float.

We then init our `stdio_init_all` for the purposes of our UART terminal interface.

We then simply echo `fav_num: 42.5` in the terminal.



The screenshot shows a terminal window titled "COM3 - PuTTY". The output of the program is a continuous stream of the text "fav_num: 42.500000" printed on multiple lines, demonstrating the loop in the code.

In our next lesson we will debug this.

Chapter 15: Debugging Floating-Point Data Type

In this chapter we are going to discuss debugging the floating-point data type.

Run OpenOCD with the below config.

```
openocd -f interface/cmsis-dap.cfg -f target/rp2350.cfg -c "adapter speed 5000"
```

Open a new terminal and then run the following to launch our dynamic debugger called GDB.

```
arm-none-eabi-gdb build/0x000b_integer-data-type.bin
```

Once it loads, we need to target our remote server.

```
target remote :3333
```

We need to halt the currently running binary.

```
monitor reset halt
```

```
(gdb) b *0x10000234
```

```
Breakpoint 1 at 0x10000234
```

Note: automatically using hardware breakpoints for read-only addresses.

```
(gdb) c
```

Continuing.

```
Thread 1 "rp2350.cm0" hit Breakpoint 1, 0x10000234 in ?? ()
```

```
(gdb) x/9i $pc
```

```
=> 0x10000234: push    {r3, r4, r5, lr}
0x10000236: bl      0x10002f5c
0x1000023a: movs   r4, #0
0x1000023c: ldr    r5, [pc, #12] @ (0x1000024c)
0x1000023e: mov    r2, r4
0x10000240: mov    r3, r5
0x10000242: ldr    r0, [pc, #12] @ (0x10000250)
0x10000244: bl     0x100030ec
0x10000248: b.n    0x1000023e
```

```
(gdb) b *0x10000244
```

```
Breakpoint 2 at 0x10000244
```

```
(gdb) c
```

Continuing.

```
Thread 1 "rp2350.cm0" hit Breakpoint 2, 0x10000244 in ?? ()
```

```
(gdb) x/i $pc
```

```
=> 0x10000244: bl      0x100030ec
```

Let's review what is inside `r2` and `r3`.

```
(gdb) i r $r2 $r3
r2          0x0          0
r3          0x40454000  1078280192
```

We see `r2` has the low word and `r3` has the high word so together `0x4045400000000000 = IEEE 754 encoding of 42.5`.

```
(gdb) set $bits = ((long long)$r2 << 32) | $r3
(gdb) set {long long}0x20000000 = $bits
(gdb) x/gf 0x20000000
0x20000000: 42.5
```

When working with the RP2350 (Cortex M33), it is important to understand how floating point values are passed to functions like `printf`. In C, when you call `printf("%f", fav_num)` with a float, the compiler promotes that value to a double because `printf` is a variadic function. On ARM Cortex M, a double is 64 bits wide, and according to the procedure call standard, it is split across two 32-bit registers. In this case, the low 32 bits of the double go into one register and the high 32 bits go into another. For the value `42.5`, the IEEE 754 double encoding is `0x4045400000000000`. That means one register holds `0x00000000` and the other holds `0x40454000`. If you look at only one register, the value appears meaningless, but together they form the correct double.

Inside GDB, you can reconstruct this double by combining the two registers. First, you shift the high word left by 32 bits and OR it with the low word to form a 64-bit integer. This gives you the raw bit pattern of the double. However, if you simply cast that integer to a double in GDB, it will convert the number's value rather than reinterpret its bits, which produces the wrong result. To force GDB to reinterpret the bits, you must store the 64-bit integer into memory and then examine that memory as a double. For example, by writing the packed value into a safe RAM location and using `x/gf` to display it, GDB will decode the bytes as a floating-point number. Doing this with the registers from the RP2350 shows the correct result: `42.5`. This process demonstrates both how the ABI splits doubles across registers and how to use GDB to reassemble and verify them.

In our next chapter we will hack this!

Chapter 16: Hacking Floating-Point Data Type

In this chapter we are going to discuss hacking the floating-point data type.

Run OpenOCD with the below config.

```
openocd -f interface/cmsis-dap.cfg -f target/rp2350.cfg -c "adapter speed 5000"
```

Open a new terminal and then run the following to launch our dynamic debugger called GDB.

```
arm-none-eabi-gdb build/0x000b_integer-data-type.bin
```

Once it loads, we need to target our remote server.

```
target remote :3333
```

We need to halt the currently running binary.

```
monitor reset halt
(gdb) b *0x10000234
Breakpoint 1 at 0x10000234
Note: automatically using hardware breakpoints for read-only addresses.
(gdb) c
Continuing.
Thread 1 "rp2350.cm0" hit Breakpoint 1, 0x10000234 in ?? ()
(gdb) x/9i $pc
=> 0x10000234: push    {r3, r4, r5, lr}
    0x10000236: bl      0x10002f5c
    0x1000023a: movs   r4, #0
    0x1000023c: ldr    r5, [pc, #12] @ (0x1000024c)
    0x1000023e: mov    r2, r4
    0x10000240: mov    r3, r5
    0x10000242: ldr    r0, [pc, #12] @ (0x10000250)
    0x10000244: bl     0x100030ec
    0x10000248: b.n   0x1000023e
(gdb) b *0x10000244
Breakpoint 2 at 0x10000244
(gdb) c
Continuing.
Thread 1 "rp2350.cm0" hit Breakpoint 2, 0x10000244 in ?? ()
(gdb) x/i $pc
=> 0x10000244: bl      0x100030ec
```

Let's review what is inside r2 and r3.

```
(gdb) i r $r2 $r3
r2          0x0          0
r3          0x40454000   1078280192
```

Let's say we want to hack 42.5 to be 43.375 so we need to change r3 to be 0x4045B000.

```
(gdb) i r $r2 $r3
r2          0x0          0
r3          0x40454000   1078280192
(gdb) set $r3 = 0x4045B000
(gdb) c
```

Continuing.

```
Thread 1 "rp2350.cm0" hit Breakpoint 2, 0x10000244 in ?? ()
```

We see 43.375 in the terminal.

```
(gdb) i r $r2 $r3
r2          0x0          0
r3          0x40454000   1078280192
(gdb) set $r3 = 0x4045C000
(gdb) c
```

Continuing.

```
Thread 1 "rp2350.cm0" hit Breakpoint 2, 0x10000244 in ?? ()
```

We now see 43.5 in the terminal.

When you look at the hexadecimal encodings of doubles, the difference between `0x4045B00000000000` and `0x4045C00000000000` comes down to the mantissa bits in the IEEE-754 format. Both numbers share the same sign bit (0 for positive) and the same exponent field (10000000100), which corresponds to an exponent of 5 after subtracting the bias of 1023. That exponent means the binary point is shifted so the number is expressed as something times 2^5 . The only part that changes between 43.375 and 43.5 is the mantissa, which encodes the fractional part of the number after normalization.

For **43.375**, the binary representation is `101011.011`. Normalized, that becomes `1.01011011 × 25`. The mantissa begins with `01011011...`, and when packed into the 52-bit fraction field, those bits line up to produce the hex sequence `...5B...`. That is why the high word of the double is `0x4045B000`. For **43.5**, the binary is `101011.1`, which normalizes to `1.010111 × 25`. The mantissa here is `010111...`, slightly larger than the previous case. When encoded, those bits produce the hex sequence `...5C...`, giving the high word `0x4045C000`. So, the difference between “B” and “C” in the hex is simply the mantissa incrementing by one step, moving the represented value from 43.375 to 43.5. This illustrates how tightly the mantissa bits map to fractional steps in the IEEE-754 encoding.



In our next chapter we will discuss the double floating-point data type.

Chapter 17: Double Floating-Point Data Type

In this chapter we are going to discuss the double floating-point data type.

Let's open up our folder **0x0011_double-floating-point-data-type**.

Now let's review our **0x0011_double-floating-point-data-type.c** file as this is located in the main folder.

```
#include <stdio.h>
#include "pico/stdlib.h"

int main(void) {
    double fav_num = 42.52525;


    stdio_init_all();

    while (true)
        printf("fav_num: %lf\r\n", fav_num);
}
```

We start off with a `double fav_num = 42.52525` which is a 64-bit float.

We then init our `stdio_init_all` for the purposes of our UART terminal interface.

We then simply echo `fav_num: 42.52525` in the terminal.



The screenshot shows a PuTTY terminal window titled "COM3 - PuTTY". The terminal output consists of multiple lines of the text "fav_num: 42.525250", demonstrating the program's output in a loop.

In our next lesson we will debug this.

Chapter 18: Debugging Double Floating-Point Data Type

In this chapter we are going to discuss debugging the double floating-point data type.

Run OpenOCD with the below config.

```
openocd -f interface/cmsis-dap.cfg -f target/rp2350.cfg -c "adapter speed 5000"
```

Open a new terminal and then run the following to launch our dynamic debugger called GDB.

```
arm-none-eabi-gdb build/0x0011_double-floating-point-data-type.bin
```

Once it loads, we need to target our remote server.

```
target remote :3333
```

We need to halt the currently running binary.

```
monitor reset halt
```

Let's debug.

```
(gdb) b *0x10000238
Breakpoint 1 at 0x10000238
Note: automatically using hardware breakpoints for read-only addresses.
(gdb) c
Continuing.
Thread 1 "rp2350.cm0" hit Breakpoint 1, 0x10000238 in ?? ()
(gdb) x/9i $pc
=> 0x10000238: push    {r3, r4, r5, lr}
0x1000023a: add     r5, pc, #24      @ (adr r5, 0x10000254)
0x1000023c: ldrd   r4, r5, [r5]
0x10000240: bl     0x10002f64
0x10000244: mov    r2, r4
0x10000246: mov    r3, r5
0x10000248: ldr   r0, [pc, #4]     @ (0x10000250)
0x1000024a: bl     0x100030f4
0x1000024e: b.n   0x10000244
```

We literally have to follow the exact procedure in our last few chapters. If the below is confusing, please review the chapters related to the floating-point data type.

```
(gdb) b *0x1000024a
Breakpoint 2 at 0x1000024a
(gdb) c
Continuing.
Thread 1 "rp2350.cm0" hit Breakpoint 2, 0x1000024a in ?? ()
(gdb) i r $r2 $r3
r2          0x645a1cac          1683627180
r3          0x4045433b          1078281019
(gdb) set $bits = ((long long)$r2 << 32) | $r3
(gdb) set {long long}0x20000000 = $bits
(gdb) x/gf 0x20000000
0x20000000:    42.52525
(gdb) x/s *0x10000250
0x100034b0:    "fav_num: %lf\r\n"
```

Here we see the same `r2` and `r3`, both being 32-bit wide, each share the total of a single 64-bit value that when formatted with `printf`, we see returns `42.52525`.

In our next chapter we will hack this!

Chapter 19: Hacking Double Floating-Point Data Type

In this chapter we are going to discuss hacking the double floating-point data type.

Run OpenOCD with the below config.

```
openocd -f interface/cmsis-dap.cfg -f target/rp2350.cfg -c "adapter speed 5000"
```

Open a new terminal and then run the following to launch our dynamic debugger called GDB.

```
arm-none-eabi-gdb build/0x0011_double-floating-point-data-type.bin
```

Once it loads, we need to target our remote server.

```
target remote :3333
```

We need to halt the currently running binary.

```
monitor reset halt
```

Let's hack.

```
(gdb) b *0x10000238
Breakpoint 1 at 0x10000238
Note: automatically using hardware breakpoints for read-only addresses.
(gdb) c
Continuing.
Thread 1 "rp2350.cm0" hit Breakpoint 1, 0x10000238 in ?? ()
(gdb) x/9i $pc
=> 0x10000238: push    {r3, r4, r5, lr}
    0x1000023a: add     r5, pc, #24      @ (adr r5, 0x10000254)
    0x1000023c: ldrd   r4, r5, [r5]
    0x10000240: bl     0x10002f64
    0x10000244: mov    r2, r4
    0x10000246: mov    r3, r5
    0x10000248: ldr   r0, [pc, #4]     @ (0x10000250)
    0x1000024a: bl    0x100030f4
    0x1000024e: b.n   0x10000244
```

We literally have to follow the exact procedure in our last few chapters. If the below is confusing, please review the chapters related to the floating-point data type.

```
(gdb) b *0x1000024a
Breakpoint 2 at 0x1000024a
(gdb) c
Continuing.
Thread 1 "rp2350.cm0" hit Breakpoint 2, 0x1000024a in ?? ()
(gdb) i r $r2 $r3
r2                0x645a1cac                1683627180
r3                0x4045433b                1078281019
```

Let's have a little fun! We can figure out how to change our number to say, 43.52525. We can figure this out by doing the following.

```
(gdb) set {double}0x20000000 = 43.52525
(gdb) x/2wx 0x20000000
```

```
0x20000000:      0x4045c33b      0x645a1cac
```

So, we know r2 has the proper values for the values right of the decimal so we need to hack r3.

```
(gdb) set $r3 = 0x4045c33b
```

```
(gdb) c
```

Continuing.

```
Thread 1 "rp2350.cm0" hit Breakpoint 2, 0x1000024a in ?? ()
```

Let's review our PuTTY!



Boom! We hacked it!

In our next chapter we will discuss static variables.

Chapter 20: Static Variables

In this chapter we are going to discuss static variables and GPIO inputs.

Let's open up our folder **0x0014_static-variables**.

Now let's review our **0x0014_static-variables.c** file as this is located in the main folder.

```
#include <stdio.h>
#include "pico/stdlib.h"

int main(void) {
    stdio_init_all();

    const uint BUTTON_GPIO = 15;
    const uint LED_GPIO = 16;
    bool pressed = 0;

    gpio_init(BUTTON_GPIO);
    gpio_set_dir(BUTTON_GPIO, GPIO_IN);
    gpio_pull_up(BUTTON_GPIO);

    gpio_init(LED_GPIO);
    gpio_set_dir(LED_GPIO, GPIO_OUT);

    while (true) {
        uint8_t regular_fav_num = 42;
        static uint8_t static_fav_num = 42;

        printf("regular_fav_num: %d\r\n", regular_fav_num);
        printf("static_fav_num: %d\r\n", static_fav_num);

        regular_fav_num++;
        static_fav_num++;

        pressed = gpio_get(BUTTON_GPIO);
        gpio_put(LED_GPIO, pressed ? 0 : 1);
    }
}
```

We start off with two constants. The first being `BUTTON_GPIO` which is assigned to GPIO15 and the second is `LED_GPIO` which is assigned to GPIO16.

Additionally, we have a boolean called `pressed` which will store the value of our button press.

We have our `gpio_init` which we have seen before in great detail where it gets our `BUTTON_GPIO` setup for usage. We also set the direction as input and we have something new here where we use a `gpio_pull_up` turns on the internal pull-up so the pin is held at a logic HIGH when nothing is driving it. Without a pull resistor, the input is floating and will read random values.

We should use a pull-down resistor but for the sake of explanation, I wanted to drive it high so this forces us to

adjust our logic later in the code to account for this. The pull-up will normally keep the value high unless you press the button.

We then init our output which we have seen in great depth in earlier chapters.

The interesting thing is within our while loop as we have a `regular_fav_num` which is created as a `uint8_t`, unsigned 8-bit integer each time throughout the loop. This will literally redefine itself to 42 every pass through the loop so it will stay consistently 42 even though we increment it with `regular_fav_num++`.


The static variable however will live on the heap rather than the stack so its value will live through every reiteration of the while loop and will in-fact increment until it overflows to where it will start at 0 and work its way up to 255 and back to 0.

So... What is the heap vs the stack?

In our loop the regular variable is a stack variable where each time the loop iteration starts, a fresh `regular_fav_num` is created on the stack and initialized to 42, so incrementing it only affects that single instance before it gets destroyed at the end of the iteration and re-created again on the next pass and that's why it always appears to stay 42. The static variable is not allocated on the stack or the heap; it lives in static storage (a fixed data area in the program image or RAM reserved for globals/static variables) and is initialized once before the program runs. Because it persists across function calls and loop iterations, `static_fav_num++` updates the same memory location each time and the value accumulates until it overflows (for an 8-bit unsigned type it wraps from 255 back to 0). Static storage differs from heap allocation (`malloc/new`): heap objects are dynamically allocated and freed at runtime, while static storage is fixed for the program lifetime.

The final pair of lines reads the raw GPIO value for the button into `pressed` and then writes a corresponding logic level to the LED pin using a ternary expression. Because we enabled an internal pull-up, `gpio_get(BUTTON_GPIO)` returns 1 when the button is released and 0 when it is pressed, so the variable name `pressed` is a bit misleading (it holds the raw input, not a true boolean "pressed" flag). The ternary `pressed ? 0 : 1` maps that raw value to the LED output: if the input is non-zero (button released) it writes 0 to `LED_GPIO`, otherwise (button pressed) it writes 1 so the LED will be driven high when the button is actually pressed (assuming the LED is wired active-high). A clearer equivalent is `gpio_put(LED_GPIO, !gpio_get(BUTTON_GPIO))` and you can also invert the mapping or change wiring if you prefer the opposite behavior.

Let's open PuTTY and observe the behavior!



```
static_fav_num: 115
regular_fav_num: 42
static_fav_num: 116
regular_fav_num: 42
static_fav_num: 117
regular_fav_num: 42
static_fav_num: 118
regular_fav_num: 42
static_fav_num: 119
regular_fav_num: 42
static_fav_num: 120
regular_fav_num: 42
static_fav_num: 121
regular_fav_num: 42
static_fav_num: 122
regular_fav_num: 42
static_fav_num: 123
regular_fav_num: 42
static_fav_num: 124
regular_fav_num: 42
static_fav_num: 125
```

In our next lesson, we will debug this!

Chapter 21: Debugging Static Variables

In this chapter we are going to discuss debugging static variables and GPIO inputs.

Let's open up Ghidra and create a project for our **0x0014_static-variables**.

We have several chapters explaining how to create a project in Ghidra. At this stage you should be comfortable to do such.

```
***** ...
* FUNCTION ...
***** ...
undefined FUN_10000234 ()
undefined <UNASSIGNED> <RETURN>
FUN_10000234+1 XREF[1,1]: 1000018c (c), 1000018a (*)
FUN_10000234
10000234 10 b5 push {r4,lr}
10000236 02 f0 ed fe bl FUN_10003014 undefined FUN_10003014()
1000023a 0f 20 movs r0,#0xf
1000023c 00 f0 60 f8 bl FUN_10000300 undefined FUN_10000300()
10000240 0f 20 movs r0,#0xf
10000242 4f f0 00 03 mov.w r3,#0x0
10000246 43 ec 44 00 mcrr p0,0x4,r0,r3,cr4
1000024a 00 22 movs r2,#0x0
1000024c 01 21 movs r1,#0x1
1000024e 00 f0 43 f8 bl FUN_100002d8 undefined FUN_100002d8()
10000252 10 20 movs r0,#0x10
10000254 00 f0 54 f8 bl FUN_10000300 undefined FUN_10000300()
10000258 10 23 movs r3,#0x10
1000025a 4f f0 01 02 mov.w r2,#0x1
1000025e 42 ec 44 30 mcrr p0,0x4,r3,r2,cr4
10000262 0b 4c ldr r4,[DAT_10000290] = 200005A8h
LAB_10000264 XREF[1]: 1000028e (j)
10000264 2a 21 movs r1,#0x2a
10000266 0b 48 ldr r0=>s_regular_fav_num:_%d_10003560,[DAT_100002...= "regular_fav_num: %d\r\n"
= 10003560h
10000268 02 f0 9c ff bl FUN_100031a4 undefined FUN_100031a4()
1000026c 21 78 ldrb r1,[r4,#0x0]=>DAT_200005a8
1000026e 0a 48 ldr r0=>s_static_fav_num:_%d_10003578,[DAT_1000029...= "static_fav_num: %d\r\n"
= 10003578h
```

```
10000270 02 f0 98 ff bl FUN_100031a4 undefined FUN_100031a4()
10000274 4f f0 50 41 mov.w r1,#0xd0000000
10000278 23 78 ldrb r3,[r4,#0x0]=>DAT_200005a8
1000027a 10 22 movs r2,#0x10
1000027c 01 33 adds r3,#0x1
1000027e 23 70 strb r3,[r4,#0x0]=>DAT_200005a8
10000280 4b 68 ldr r3,[r1,#offset DAT_d0000004]
10000282 c3 f3 c0 33 ubfx r3,r3,#0xf,#0x1
10000286 83 f0 01 03 eor r3,r3,#0x1
1000028a 43 ec 40 20 mcrr p0,0x4,r2,r3,cr0
1000028e e9 e7 b LAB_10000264
```

Let's take our time and update these functions properly. If you are unclear on how to update function signatures please refer to earlier chapters.

Let's identify `main` which will be at `0x10000234`. This will be `int main(void)`.

Let's identify `stdio_init_all` which will be at `0x10003014`. This will be `bool stdio_init_all(void)`.

Let's identify `gpio_init` which will be at `0x1000023c` and `0x10000254`. This will be `void gpio_init(uint gpio)`.

For our next function, we need to understand about the concept of optimization. We programmed using a function called `gpio_pull_up` however it does not exist in our binary. When we drill-down into our binary, `gpio_pull_up` calls a function called `gpio_set_pulls` instead.

Let's identify `gpio_set_pulls` which will be at `0x1000024e`. This will be `void gpio_set_pulls(uint gpio, bool up, bool down)` instead.

Let's identify `printf` which will be at `0x10000268` and `0x10000270`. This will be `int printf(char *format, ...)` as `printf` is a variadic function which mean it can take an unlimited amount of arguments.

Within the while loop, the instruction `movs r1, #0x2a` loads the immediate value 42 into register `r1`. This corresponds to the initialization of the local variable `regular_fav_num = 42`. Immediately after, the `ldr r0, =s_regular_fav_num...` pulls in the address of the format string `"regular_fav_num: %d\r\n"`. With `r0` holding the format string and `r1` holding the integer value, the `bl printf` call matches the C statement `printf("regular_fav_num: %d\r\n", regular_fav_num)`.

Next, the compiler handles the static variable. Unlike the automatic local, which is reinitialized each loop iteration, the `static_fav_num` lives in the `.data` section at a fixed RAM address (`DAT_200005a8`). The instruction `ldrb r1, [r4, #0x0]` fetches its current value into `r1`. Then `ldr r0, =s_static_fav_num...` loads the format string `"static_fav_num: %d\r\n"`. Again, `bl printf` prints it. This matches the C line `printf("static_fav_num: %d\r\n", static_fav_num)`.

After printing, the code increments both counters. For the static variable, you can see `ldrb r3, [r4, #0x0]` to load it, adds `r3, #0x1` to increment, and `strb r3, [r4, #0x0]` to store it back. This is the compiled form of `static_fav_num++`. The automatic `regular_fav_num++` is optimized away in this loop because it's reinitialized to 42 every iteration, so its increment has no lasting effect, hence you don't see a store back to memory.

The bottom half of the loop corresponds to the GPIO logic. The instruction `mov.w r1, #0xd0000000` sets up a base address for a memory-mapped peripheral. Then `ldr r3, [r1, #offset DAT_d0000004]` reads from a register (likely the button input). The `ubfx r3, r3, #0xf, #0x1` extracts a single bit (bit 15), which is the button state. The `eor r3, r3, #0x1` flips it, implementing the ternary `pressed ? 0 : 1`. Finally, `mcrn p0, 0x4, r2, r3, cr0` is a coprocessor register write, which in this context is the compiler's way of emitting a store to the GPIO output register, effectively toggling the LED.

The unconditional branch `b LAB_10000264` at the end loops execution back to the start, reproducing the `while (true)` infinite loop. So, in summary: the assembly faithfully implements the C code by reinitializing a local variable, maintaining a static counter across iterations, printing both, and then reading a GPIO input to drive an LED output, all wrapped in an endless loop. The differences you notice, like the missing increment of the automatic variable, are the compiler's optimizations, since that increment has no observable effect.

Let's review our updated Ghidra.

```

*****|***** ...
*          FUNCTION          ...
*****|***** ...

int __stdcall main(void)
int      r0:4      <RETURN>
main+1                                XREF[1,1]:  1000018c (c), 1000018a (*)
main
10000234 10 b5      push      {r4,lr}
10000236 02 f0 ed fe  bl      stdio_init_all      bool stdio_init_all(void)
1000023a 0f 20      movs     r0,#0xf
1000023c 00 f0 60 f8  bl      gpio_init      void gpio_init(uint gpio)
10000240 0f 20      movs     r0,#0xf
10000242 4f f0 00 03  mov.w   r3,#0x0
10000246 43 ec 44 00  mcrr    p0,0x4,r0,r3,cr4
1000024a 00 22      movs     r2,#0x0
1000024c 01 21      movs     r1,#0x1
1000024e 00 f0 43 f8  bl      gpio_set_pulls  void gpio_set_pulls(uint gpio, b ...
10000252 10 20      movs     r0,#0x10
10000254 00 f0 54 f8  bl      gpio_init      void gpio_init(uint gpio)
10000258 10 23      movs     r3,#0x10
1000025a 4f f0 01 02  mov.w   r2,#0x1
1000025e 42 ec 44 30  mcrr    p0,0x4,r3,r2,cr4
10000262 0b 4c      ldr      r4,[DAT_10000290]      = 200005A8h

LAB_10000264                                XREF[1]:  1000028e (j)
10000264 2a 21      movs     r1,#0x2a
10000266 0b 48      ldr      r0=>s_regular_fav_num:_%d_10003560 ,[DAT_100002...= "regular_fav_num: %d\r\n"
                                                = 10003560h
10000268 02 f0 9c ff  bl      printf      int printf(char * format, ...)
1000026c 21 78      ldrb     r1,[r4,#0x0]=>DAT_200005a8
1000026e 0a 48      ldr      r0=>s_static_fav_num:_%d_10003578 ,[DAT_1000029...= "static_fav_num: %d\r\n"
                                                = 10003578h

10000270 02 f0 98 ff  bl      printf      int printf(char * format, ...)
10000274 4f f0 50 41  mov.w   r1,#0xd0000000
10000278 23 78      ldrb     r3,[r4,#0x0]=>DAT_200005a8
1000027a 10 22      movs     r2,#0x10
1000027c 01 33      adds    r3,#0x1
1000027e 23 70      strb     r3,[r4,#0x0]=>DAT_200005a8
10000280 4b 68      ldr      r3,[r1,#offset DAT_d0000004]
10000282 c3 f3 c0 33  ubfx    r3,r3,#0xf,#0x1
10000286 83 f0 01 03  eor     r3,r3,#0x1
1000028a 43 ec 40 20  mcrr    p0,0x4,r2,r3,cr0
1000028e e9 e7      b       LAB_10000264

```

In our next lesson, we will hack this!

Chapter 22: Hacking Static Variables

In this chapter we are going to discuss hacking static variables and GPIO inputs.

Let's open up Ghidra and hack the project for our **0x0014_static-variables**.

Let's say we want to simply hack 0x2a or 42 decimal to 43.

```
LAB_10000264 XREF[1]: 1000028e (j)
10000264 2a 21      movs     r1,#0x2a
```

Let's patch the instruction to 0x2b. We have covered this in detail in chapter 7. If this process is confusing, please review.


```
LAB_10000264 XREF[1]: 1000028e (j)
10000264 2b 21      movs     r1,#0x2b
```

In the past we have hacked the GPIO output so let's try something new with the GPIO input!

As we know we have a pull-up on the button GPIO so let's hack the default ternary operator from 0x1 to 0x0 so the button will be on by default not off!

```
10000286 83 f0 01 03      eor     r3,r3,#0x1
10000286 83 f0 00 03      eor     r3,r3,#0x0
```


Now let's patch the binary.



We need to use a tool to convert this hacked binary into the UF2 format.

```
python ..\uf2conv.py build\0x0014_static-variables-h.bin --base 0x10000000 --family 0xe48bff59 --output build\hacked.uf2
```

After flashing the **hacked.uf2** to the Pico 2, we see the following in the serial terminal.



```
static_fav_num: 37
regular_fav_num: 43
static_fav_num: 38
regular_fav_num: 43
static_fav_num: 39
regular_fav_num: 43
static_fav_num: 40
regular_fav_num: 43
static_fav_num: 41
regular_fav_num: 43
static_fav_num: 42
regular_fav_num: 43
static_fav_num: 43
regular_fav_num: 43
static_fav_num: 44
regular_fav_num: 43
static_fav_num: 45
regular_fav_num: 43
static_fav_num: 46
regular_fav_num: 43
static_fav_num: 47
regular_fav_num: 43
static_fav_num: 48
```

Boom! We see 43 instead of 42 and on our breadboard we see GPIO16 lit up by default and when you press the button it goes off so we hacked that behavior as well!

In our next lesson we will cover constants.

Chapter 23: Constants

In this chapter we are going to discuss constants and I²C w/ a 1602 LCD display.

Let's open up our folder **0x0017_constants**.

Now let's review our **0x0017_constants.c** file as this is located in the main folder.

```
#include <stdio.h>
#include <string.h>
#include "pico/stdlib.h"
#include "hardware/i2c.h"
#include "lcd_1602.h"

#define FAV_NUM 42

#define I2C_PORT i2c1
#define I2C_SDA_PIN 2
#define I2C_SCL_PIN 3

const int OTHER_FAV_NUM = 1337;

int main(void) {
    stdio_init_all();

    i2c_init(I2C_PORT, 100000);
    gpio_set_function(I2C_SDA_PIN, GPIO_FUNC_I2C);
    gpio_set_function(I2C_SCL_PIN, GPIO_FUNC_I2C);
    gpio_pull_up(I2C_SDA_PIN);
    gpio_pull_up(I2C_SCL_PIN);

    lcd_i2c_init(I2C_PORT, 0x27, 4, 0x08);

    lcd_set_cursor(0, 0);
    lcd_puts("Reverse");
    lcd_set_cursor(1, 0);
    lcd_puts("Engineering");

    while (true) {
        printf("FAV_NUM: %d\r\n", FAV_NUM);
        printf("OTHER_FAV_NUM: %d\r\n", OTHER_FAV_NUM);
    }
}
```

The Inter-Integrated Circuit, or I²C, is a two-wire serial communication protocol that allows a single controller to communicate with multiple peripheral devices over a shared bus. It uses one line for data, called SDA, and one line for the clock, called SCL. Because the lines are open-drain, devices only pull them low, and pull-up resistors are required to restore the high state. On the RP2350 Pico 2, the I²C controllers are flexible and can be mapped to different GPIO pins, which makes it possible to adapt the bus to a variety of hardware layouts. Each device on the bus has a unique address, and the controller selects which device to talk to by sending this address before transmitting or receiving data. This makes I²C especially efficient for connecting displays, sensors, and other small

In our next chapter, we will debug this.

Chapter 24: Debugging Constants

In this chapter we are going to discuss debugging constants and I²C with our LCD 1602 module.

We are at the stage where our debugging will be significantly more involved than prior chapters so if this appears to be moving at a pace you are not comfortable, please take your time and review the last 10 or so chapters.

Let's open up Ghidra and create a project for our **0x0017_constants**.

We start with `main`. In the past we know `main` traditionally starts at `0x10000234`.

```
*****
*                               FUNCTION
*****
undefined FUN_10000234 ()
undefined  ▲ <UNASSIGNED>  <RETURN>
FUN_10000234+1      XREF[1,1]:  1000018c (c), 1000018a (*)
FUN_10000234
10000234 08 b5      push      {r3,lr}
10000236 03 f0 e1 fa  bl      FUN_100037fc      undefined FUN_100037fc()
1000023a 1a 49      ldr      r1,[DAT_100002a4]  = 000186A0h
1000023c 1a 48      ldr      r0,[DAT_100002a8]  = 2000062Ch
1000023e 03 f0 4d fd  bl      FUN_10003cdc      undefined FUN_10003cdc()
10000242 03 21      movs     r1,#0x3
10000244 02 20      movs     r0,#0x2
10000246 00 f0 53 fb  bl      FUN_100008f0      undefined FUN_100008f0()
1000024a 03 21      movs     r1,#0x3
1000024c 08 46      mov      r0,r1
1000024e 00 f0 4f fb  bl      FUN_100008f0      undefined FUN_100008f0()
10000252 00 22      movs     r2,#0x0
10000254 01 21      movs     r1,#0x1
10000256 02 20      movs     r0,#0x2
10000258 00 f0 68 fb  bl      FUN_1000092c      undefined FUN_1000092c()
1000025c 00 22      movs     r2,#0x0
1000025e 01 21      movs     r1,#0x1
10000260 03 20      movs     r0,#0x3
10000262 00 f0 63 fb  bl      FUN_1000092c      undefined FUN_1000092c()
10000266 08 23      movs     r3,#0x8
10000268 04 22      movs     r2,#0x4
1000026a 27 21      movs     r1,#0x27
1000026c 0e 48      ldr      r0,[DAT_100002a8]  = 2000062Ch
1000026e 00 f0 25 f8  bl      FUN_100002bc      undefined FUN_100002bc()
10000272 00 21      movs     r1,#0x0
10000274 08 46      mov      r0,r1
10000276 00 f0 3d fa  bl      FUN_100006f4      undefined FUN_100006f4()
1000027a 0c 48      ldr      r0=>s_Reverse_10003ee8,[DAT_100002ac]  = "Reverse"
                                                = 10003EE8h
```

```

1000027c 00 f0 b8 fa  b1      FUN_100007f0      undefined FUN_100007f0()
10000280 01 20          movs      r0,#0x1
10000282 00 21          movs      r1,#0x0
10000284 00 f0 36 fa  b1      FUN_100006f4      undefined FUN_100006f4()
10000288 09 48          ldr      r0=>s_Engineering_10003ef0 , [DAT_100002b0]
                                           = "Engineering"
                                           = 10003EF0h
1000028a 00 f0 b1 fa  b1      FUN_100007f0      undefined FUN_100007f0()

LAB_1000028e                                XREF[1]: 100002a0 (j)
1000028e 2a 21          movs      r1,#0x2a
10000290 08 48          ldr      r0=>s_FAV_NUM: _%d_10003efc , [DAT_100002b4]
                                           = "FAV_NUM: %d\r\n"
                                           = 10003EFCh
10000292 03 f0 7b fb  b1      FUN_1000398c      undefined FUN_1000398c()
10000296 40 f2 39 51  movw     r1,#0x539
1000029a 07 48          ldr      r0=>s_OTHER_FAV_NUM: _%d_10003f0c , [DAT_100002b8] = "OTHER_FAV_NUM: %d\r\n"
                                           = 10003F0Ch
1000029c 03 f0 76 fb  b1      FUN_1000398c      undefined FUN_1000398c()
100002a0 f5 e7          b        LAB_1000028e
100002a2 00          ??      00h
100002a3 bf          ??      BFh

```

I know this can be overwhelming to see all these stripped symbols however we will take our time and take them one step at a time.

The first think we want to do is look at the comments on the right-hand side as this will help us understand what the functionality or subroutines/functions might be.

Let's first update main as this will be at 0x10000234 which will be `int main(void)`.

In the past we know the first function called within main is our `bool stdio_init_all(void)` as this will be at 0x10000236.

In the last chapter, we glossed over the concept of I²C. As we start to understand this protocol, we will first know that our next function has to arguments that are being passed which would be `i2c_init`.

The first argument is what we refer to as a struct pointer. A struct is a complex datatype made up of other primitive datatypes or other structs with primitive data types.

If we dive into the pico-c-sdk library, we see the following.

```

/** \file hardware/i2c.h
 * \defgroup hardware_i2c hardware_i2c
 *
 * \brief I2C Controller API
 *
 * The I2C bus is a two-wire serial interface, consisting of a serial data line SDA and a
serial clock SCL. These wires carry
 * information between the devices connected to the bus. Each device is recognized by a
unique 7-bit address and can operate as
 * either a “transmitter” or “receiver”, depending on the function of the device. Devices
can also be considered as masters or
 * slaves when performing data transfers. A master is a device that initiates a data
transfer on the bus and generates the
 * clock signals to permit that transfer. The first byte in the data transfer always
contains the 7-bit address and
 * a read/write bit in the LSB position. This API takes care of toggling the read/write
bit. After this, any device addressed
 * is considered a slave.
 *
 * This API allows the controller to be set up as a master or a slave using the \ref
i2c_set_slave_mode function.
 *
 * The external pins of each controller are connected to GPIO pins as defined in the GPIO
muxing table in the datasheet. The muxing options
 * give some IO flexibility, but each controller external pin should be connected to only
one GPIO.
 *
 * Note that the controller does NOT support High speed mode or Ultra-fast speed mode, the
fastest operation being fast mode plus
 * at up to 1000Kb/s.
 *
 * See the datasheet for more information on the I2C controller and its usage.
 *
 * \subsection i2c_example Example
 * \addtogroup hardware_i2c
 * \include bus_scan.c
 */

typedef struct i2c_inst i2c_inst_t;

```

The above is the **i2c.h** file in the sdk.

Here is the **i2c.c** file portion in the sdk.

```

// -----
// Generic input/output

struct i2c_inst {
    i2c_hw_t *hw;
    bool restart_on_next;
};

```

In C, a struct (short for structure) is a user-defined data type that groups together related variables under one name. These variables, called members or fields, can be of different primitive types (like int, char, bool) or even other structs. Unlike arrays, which hold multiple values of the same type, structs let you combine heterogeneous data into a single logical unit. This makes them perfect for modeling hardware registers, configuration blocks, or higher-level abstractions like an I²C controller instance.

In the Pico SDK, you'll often see a pattern like this.

```
typedef struct i2c_inst i2c_inst_t;
```

This line does two things:

1. It forward-declares a struct called `i2c_inst`. At this point, the compiler knows such a struct exists but doesn't yet know its contents.
2. It creates an alias `i2c_inst_t` for `struct i2c_inst`. This is purely for readability and convention as the SDK code prefers the `_t` suffix for typedef'd types.

Later, the struct is fully defined.

```
struct i2c_inst {  
    i2c_hw_t *hw;  
    bool restart_on_next;  
};
```

Now the compiler knows the struct has two members.

`i2c_hw_t *hw` - a pointer to the actual hardware registers for the I²C peripheral. This is the low-level connection to the silicon, where each bit in memory maps to a control or status register.

`bool restart_on_next` - a software flag used by the SDK to track whether the next I²C transaction should issue a re-start condition instead of a full stop/start cycle.

Circling back, at **main.c**, we see the following.

```
#define I2C_PORT i2c1
```

This is a preprocessor macro that tells the compiler: whenever you write `I2C_PORT`, replace it with `i2c1`. It's a convenience alias so your application code can be written in terms of `I2C_PORT` without hard-coding which controller you're using. If you later want to switch to `i2c0`, you only change this one line.

Now, what is `i2c1`? In `i2c.h`, it's defined as the following.

```
#define i2c1 (&i2c1_inst) ///< Identifier for I2C HW Block 1
```

Here, `i2c1` is not a variable—it's another macro. It expands to the address of a global instance called `i2c1_inst`. This means that whenever you pass `i2c1` into an SDK function, you're really passing a pointer to a struct that represents the I²C1 controller.

That struct instance is created in `i2c.c` is as follows.

```
i2c_inst_t i2c1_inst = {i2c1_hw, false};
```

This line defines a global variable `i2c1_inst` of type `i2c_inst_t` (the struct we discussed earlier). It is initialized with two values:

1. `i2c1_hw` - a pointer to the actual hardware registers for I²C1.
2. `false` - the initial value of the `restart_on_next` flag.

So `i2c1_inst` is the software object that bundles together the hardware pointer and the SDK's state tracking for I²C1.

But what is `i2c1_hw`? Back in `i2c.h`, it's defined as the following.

```
#define i2c1_hw ((i2c_hw_t *)I2C1_BASE)
```

This macro casts the constant `I2C1_BASE` into a pointer of type `i2c_hw_t *`. In other words, it says: "treat the memory starting at `I2C1_BASE` as if it were a struct of type `i2c_hw_t`." That struct (`i2c_hw_t`) is a register map that matches the layout of the I²C peripheral in silicon. By dereferencing this pointer, the SDK can directly read and write the controller's registers.

Finally, in `addressmap.h`, we see the following.

```
#define I2C1_BASE _u(0x40098000)
```

This is the literal base address of the I²C1 peripheral in the RP2040's memory map. The `_u()` macro just ensures the constant is treated as an unsigned integer. This address is fixed by the chip's design: the hardware engineers wired the I²C1 block to live at `0x40098000` in the system bus.

Putting it all together...

- `I2C1_BASE` is the raw memory address of the I²C1 hardware block.
- `i2c1_hw` casts that address into a pointer to a register map struct (`i2c_hw_t *`).
- `i2c1_inst` is a higher-level struct (`i2c_inst_t`) that stores this hardware pointer plus SDK state.
- `i2c1` is a macro that expands to `&i2c1_inst`, giving you a pointer to that struct.
- `I2C_PORT` is an application-level alias that maps to `i2c1`, so your code can stay generic.

I know this was a lot to go through but as we progress, we get a better understanding of how this all works together.

As we work our way back to Ghidra we see the following.

```

1000023a 1a 49      ldr    r1,[DAT_100002a4]      = 000186A0h
1000023c 1a 48      ldr    r0,[DAT_100002a8]      = 2000062Ch
1000023e 03 f0 4d fd  bl     FUN_10003cdc            undefined FUN_10003cdc()

```

On ARM Cortex-M devices, the standard procedure call convention (AAPCS) dictates that the first four function arguments are passed in registers `r0` through `r3`, with any additional arguments placed on the stack. The return value is also delivered back to the caller in `r0`.

In the disassembly, this convention is followed exactly as the compiler prepares the call to `i2c_init` by loading `r0` with the first argument and `r1` with the second. The instruction `ldr r0, [DAT_100002a8]` loads the address `0x2000062C` into `r0`, which is the SRAM location of the global `i2c1_inst` structure. This pointer is the first argument, corresponding to `I2C_PORT` in our source code, which expands to `i2c1` and then to `&i2c1_inst`. The instruction `ldr r1, [DAT_100002a4]` loads the immediate value `0x000186A0` into `r1`, which is decimal 100,000, the baudrate we passed as the second argument.

Finally, the `bl FUN_10003cdc` instruction branches to the function body of `i2c_init`, with `r0` and `r1` already holding the two arguments. At runtime, `i2c_init` interprets `r0` as a pointer to the `i2c_inst_t` struct, which itself contains a pointer to the hardware register block at `0x40098000`, and `r1` as the baudrate value to configure. This is a clean demonstration of how the C source, preprocessor macros, and ABI rules collapse into a simple register-based calling sequence that directly connects your high-level function call to the underlying silicon.


We know our struct is the following.


```

struct i2c_inst {
    i2c_hw_t *hw;
    bool restart_on_next;
};


```

Let's right click on our bin in **Data Type Manager** and create a new **Structure**.






We need to now create our `i2c_inst_t` struct.




Now we need to add this struct into our original one in addition to adding in a bool. We also need to update our names and comments as shown in the image below.




Let's double-click on our r0 address. Then we will add the datatype of our i2c_inst.

```
1000023c 1a 48      ldr     r0, [DAT_100002a8] = 2000062Ch
```



Right-click and select **Data** then **Pointer**.




Now we will have the following.



Right-click again and select **Data** and **Choose Data Type...**



Now we will update the PTR_DATA_100002a8 to i2c_inst.



We have successfully created our own nested struct!

Therefore, the function at 0x1000023e is `i2c_init` and we will update that in Ghidra.

```
uint i2c_init(i2c_inst_t *i2c, uint baudrate)
```

The next function is the `gpio_set_function` at 0x10000246. This sets up our I²C and we can edit the function signature to be `void gpio_set_function(uint gpio, int fn)`. The `fn` is a function pointer meaning it is going to take an address to a function to be called. This function has other locations which will be set by this.

Next, we have our pull up which we have as `gpio_pull_up` however the compiler optimizes it down to its internal function called `gpio_set_pulls`. We can set that at 0x10000258 as `void gpio_set_pulls(uint gpio, bool up, bool down)`. This function has other locations which will be set by this.

We now come across our `lcd_i2c_init` function at 0x1000026e. We will edit the function signature to `void lcd_i2c_init(i2c_inst_t * i2c, uint pcf_addr, int nibble_shift, uint backlight_mask)`.

Following the init, we have the `lcd_set_cursor` function at 0x10000276. We will edit the function signature to `void lcd_set_cursor(int line, int position)`. This function has other locations which will be set by this.

Once we set our cursor, we have our `lcd_puts` at 0x1000027c. We will edit the function signature to `void lcd_puts(char *s)`. This function has other locations which will be set by this.

Finally, we have our `printf` function at 0x10000292. We will edit the function signature to `int __wrap_printf(char *format, ...)`.

I realize this was a lot to get our head wrapped around however it takes practice. In the real-world, we won't have the source code so a good deal of functions will remain unresolved however as you start to go through the reversing process, you will identify critical pieces necessary for proper binary modification.

In our next chapter we will hack this!

Chapter 25: Hacking Constants

In this chapter we are going to discuss hacking constants and PC with our LCD 1602 module.

We are at the stage where our debugging will be significantly more involved than prior chapters so if this appears to be moving at a pace you are not comfortable, please take your time and review the last 10 or so chapters.

Let's open up Ghidra and create a project for our **0x0017_constants**.

At 0x1000028e, we have our 42 const and at 0x10000296, we have our 1337 const.

We can patch these to be 0x02b or 43 and 0x540 or 1344.

Let's patch these in Ghidra as we have done this in prior chapters.

```
1000028e 2b 21      movs      r1,#0x2b
10000290 08 48      ldr       r0=>s_FAV_NUM:_$d_10003efc,[DAT_100002b4]    = "FAV_NUM: %d\r\n"
                                                    = 10003EFCh
10000292 03 f0 7b fb  bl      __wrap_printf                                int __wrap_printf(char * format, ...
10000296 40 f2 40 51  movw     r1,#0x540
```


In addition, we will hack our LCD, "Reverse Engineering", text.

```
1000027a 0c 48      ldr       r0=>s_Reverse_10003ee8,[DAT_100002ac]      = "Reverse"
                                                    = E8h
1000027c 00 f0 b8 fa  bl      lcd_puts                                    void lcd_puts(char * s)
10000280 01 20      movs     r0,#0x1
10000282 00 21      movs     r1,#0x0
10000284 00 f0 36 fa  bl      lcd_set_cursor                               void lcd_set_cursor(int line, in ...
10000288 09 48      ldr       r0=>s_Engineering_10003ef0,[DAT_100002b0]  = "Engineering"
                                                    = 10003EF0h
1000028a 00 f0 b1 fa  bl      lcd_puts                                    void lcd_puts(char * s)
```

If we double-click on DAT_100002ac, we see 4 bytes which are a memory address pointing to our string.

```
DAT_100002ac | XREF[1]: main:1000027a (R)
100002ac e8    ??    E8h    ? -> 10003ee8
100002ad 3e    ??    3Eh    >
100002ae 00    ??    00h
100002af 10    ??    10h
```

We have the proper address in little endian form so we need to look at 0x10003ee8. Let's press g and type in this address.



```
10003ee8 52 65 76      s_Reverse_10003ee8      XREF[1]:  main:1000027a (*)
          65 72 73      ds      "Reverse"
          65 00


10003ef0 45 6e 67      s_Engineering_10003ef0  XREF[1]:  main:10000288 (*)
          69 6e 65      ds      "Engineering"
          65 72 69 ...
```

Simply right-click and change the word *Reverse* to *Exploit*. Right-click and select **Patch Data** on *Reverse* and change it to *Exploit*. It **MUST** be the same number of bytes.

```
10003ee8 45 78 70      s_Exploit_10003ee8     XREF[1]:  main:1000027a (*)
          6c 6f 69      ds      "Exploit"
          74 00

10003ef0 45 6e 67      s_Engineering_10003ef0 XREF[1]:  main:10000288 (*)
          69 6e 65      ds      "Engineering"
          65 72 69 ...
```

Now we will **File, Export** and select **Raw Bytes**, and name it **0x0017_constants-h.bin** and press **Ok**.



Chapter 26: Operators

In this chapter we are going to discuss operators with a DHT11 humidity and temperature sensor.

Let's open up our folder **0x001a_operators**.

Now let's review our **0x001a_operators.c** file as this is located in the main folder.

```
#include <stdio.h>
#include "pico/stdlib.h"
#include "dht11.h"

int main(void) {
    stdio_init_all();

    dht11_init(4);

    int x = 5;
    int y = 10;
    int arithmetic_operator = (x * y);
    int increment_operator = x++;
    bool relational_operator = (x > y);
    bool logical_operator = (x > y) && (y > x);
    int bitwise_operator = (x<<1); // x is now 6 because of x++ or 0b00000110 and (x<<1)
    is 0b00001100 or 12
    int assignment_operator = (x += 5);

    while (true) {
        printf("arithmetic_operator: %d\r\n", arithmetic_operator);
        printf("increment_operator: %d\r\n", increment_operator);
        printf("relational_operator: %d\r\n", relational_operator);
        printf("logical_operator: %d\r\n", logical_operator);
        printf("bitwise_operator: %d\r\n", bitwise_operator);
        printf("assignment_operator: %d\r\n", assignment_operator);

        float hum, temp;
        if (dht11_read(&hum, &temp)) {
            printf("Humidity: %.1f%, Temperature: %.1f°C\r\n", hum, temp);
        } else {
            printf("DHT11 read failed\r\n");
        }

        sleep_ms(2000);
    }
}
```

Here we introduce the DHT11 temperature and humidity sensor. Each loop iteration prints the current humidity and temperature from the DHT11 sensor, followed by the results of several operator evaluations. These values remain constant across iterations because the operator expressions are evaluated once before the loop begins and are not recalculated dynamically.

The humidity and temperature readings are 51.0% and 23.8°C respectively which indicate successful communication with the DHT11 sensor. These values are typical for indoor environments and confirm that the sensor is functioning correctly. The `dht11_read()` function returns true, triggering the formatted output.


The operator results are derived from the initial values of `x = 5` and `y = 10`. The `arithmetic_operator` result is 50, which comes from `x * y`. This is a straightforward multiplication and confirms that both variables are initialized correctly. The `increment_operator` is 5, which might seem confusing at first. This is because the post-increment `x++` returns the original value of `x` before incrementing it. After this line, `x` becomes 6, but the stored result is still 5.

The `relational_operator` evaluates `(x > y)` after `x` has been incremented to 6. Since 6 is not greater than 10, the result is 0 (false). Similarly, the `logical_operator` checks whether `(x > y) && (y > x)`. One condition is true and the other is false, so the result is also 0. These results reinforce how logical and relational operators depend on the current state of variables and how post-increment can subtly affect outcomes.

The `bitwise_operator` result is 12, which comes from shifting `x` (now 6) left by one bit. In binary, 6 is 00000110, and shifting left by one yields 00001100, which is 12 in decimal. This demonstrates how bitwise operations manipulate individual bits and can be used for efficient arithmetic or hardware control.

Finally, the `assignment_operator` result is 11, which reflects `x += 5`. Since `x` was 6 after the increment, adding 5 yields 11. This shows how compound assignment operators both modify and return the updated value.

We can see the result in PuTTY.



```
COM4 - PuTTY
assignment_operator: 11
Humidity: 51.0%, Temperature: 23.8°C
arithmetic_operator: 50
increment_operator: 5
relational_operator: 0
logical_operator: 0
bitwise_operator: 12
assignment_operator: 11
Humidity: 51.0%, Temperature: 23.8°C
arithmetic_operator: 50
increment_operator: 5
relational_operator: 0
logical_operator: 0
bitwise_operator: 12
assignment_operator: 11
Humidity: 51.0%, Temperature: 23.8°C
arithmetic_operator: 50
increment_operator: 5
relational_operator: 0
logical_operator: 0
bitwise_operator: 12
assignment_operator: 11
Humidity: 51.0%, Temperature: 23.8°C
```

In our next chapter we will debug this.

Chapter 27: Debugging Operators

In this chapter we are going to discuss debugging operators with a DHT11 humidity and temperature sensor.

Let's create a new project in Ghidra called **0x001a_operators**.

As in prior chapters, it will be a Cortex ARM:LE:32 little endian and we have to set the options flash to `0x10000000` which is the base of XIP as we have discussed.

A little review shall we! We know that the `Reset_Handler` leads us to `main`. We know working through multiple examples that `main` is on or about `0x10000234` however we also know that in ARM the address of offset 4 from the base of XIP is the address of our `Reset_Handler`.

```
//
// flash
// ram:10000000-ram:1000443b
//

assume spsr = 0x0 (Default)
DAT_10000000 XREF[2]: 10001624 (*), 10001678 (*)
10000000 00 ?? 00h
10000001 20 ?? 20h
10000002 08 ?? 08h
10000003 20 ?? 20h
10000004 5d 01 00 10 addr LAB_1000015c+1
```

Ok here we see `5d 01` however we know that we have to reverse the byte order as we are dealing with little endian.

In addition, we know that the offset is `+1` because of the thumb bit so our real `Reset_Handler` is at the offset of `15c`.

```

LAB_1000015c+1                                XREF[0,1]: 10000004 (*)
1000015c 4f f0 50 40    mov.w    r0,#0xd0000000
10000160 00 68        ldr     r0,[r0,#0x0]=>DAT_d0000000
10000162 10 b1        cbz    r0,LAB_1000016a
10000164 4f f0 00 00    mov.w    r0,#0x0
10000168 f2 e7        b     LAB_10000150

LAB_1000016a                                XREF[1]: 10000162 (j)
1000016a 0d a4        adr     r4,[0x100001a0]

LAB_1000016c                                XREF[1]: 10000176 (j)
1000016c 0e cc        ldmia  r4!,{r1,r2,r3}=>DAT_100001a0
                                           = 10003B94h
                                           = 20000110h
                                           = 200009A4h
                                           = 10004428h
                                           = 20080000h
                                           = 4Fh 0

1000016e 00 29        cmp    r1,#0x0
10000170 02 d0        beq    LAB_10000178
10000172 00 f0 12 f8    bl     FUN_1000019a      undefined FUN_1000019a()
10000176 f9 e7        b     LAB_1000016c

LAB_10000178                                XREF[1]: 10000170 (j)
10000178 15 49        ldr    r1,[DAT_100001d0]
                                           = 200009A4h
1000017a 16 4a        ldr    r2,[DAT_100001d4]
                                           = 20000BD4h
1000017c 00 20        movs  r0,#0x0
1000017e 00 e0        b     LAB_10000182

LAB_10000180                                XREF[1]: 10000184 (j)
10000180 01 c1        stmia  r1! =>DAT_200009a4,{r0}

LAB_10000182                                XREF[1]: 1000017e (j)
10000182 91 42        cmp    r1,r2
10000184 fc d1        bne   LAB_10000180
10000186 14 49        ldr    r1=>FUN_10003134+1,[DAT_100001d8]
                                           = 10003135h
10000188 88 47        blx   r1=>FUN_10003134      undefined FUN_10003134()

1000018a 14 49        ldr    r1=>FUN_10000234+1,[DAT_100001dc]
                                           = 10000235h
1000018c 88 47        blx   r1=>FUN_10000234      undefined FUN_10000234()
1000018e 14 49        ldr    r1=>FUN_1000312c+1,[DAT_100001e0]
                                           = 1000312Dh
10000190 88 47        blx   r1=>FUN_1000312c      undefined FUN_1000312c()

LAB_10000192                                XREF[1]: 10000194 (j)
10000192 00 be        bkpt  0x0
10000194 fd e7        b     LAB_10000192



LAB_10000196                                XREF[1]: 1000019c (j)
10000196 01 c9        ldmia  r1!,{r0}
10000198 01 c2        stmia  r2!,{r0}

```


Well, that does not look like the beginning of a function so what do we do!

This is valuable to know as imagine you are working on another ARM device that is not a RP2350 and you need to locate main. We know this is the proper address and we KNOW that this is the `Reset_Handler` so let's create a function here in Ghidra.

Highlight the address and **right-click** and select **F** for function.



Let's go ahead and update the signature to `Reset_Handler`.



At the bottom we will find three functions, the first will be some sort of `init` and the last will be `exit` so the middle is `main`. This is our `platform_entry` and is part of the `cr0.S` we discussed in chapter 4.

```

10000188 88 47      blx      r1=>FUN_10003134      undefined FUN_10003134()
1000018a 14 49      ldr      r1=>FUN_10000234+1,[DAT_100001dc] = 10000235h
1000018c 88 47      blx      r1=>FUN_10000234      undefined FUN_10000234()
1000018e 14 49      ldr      r1=>FUN_1000312c+1,[DAT_100001e0] = 1000312Dh
10000190 88 47      blx      r1=>FUN_1000312c      undefined FUN_1000312c()

```

And there it is! The FUN_10000234 is our main!

Let's double-click it and go there.

```

*****
*                               FUNCTION                               *
*****
undefined FUN_10000234 ()
undefined  ▲ <UNASSIGNED> <RETURN>
undefined4  Stack[-0x1c]:4 local_1c      XREF[1]: 1000024c (R)
undefined4  Stack[-0x20]:4 local_20      XREF[1]: 10000244 (R)
undefined4  Stack[-0x28]:4 local_28      XREF[1]: 1000025c (W)
FUN_10000234+1      XREF[1,1]: Reset_Handler:1000018c (c),
FUN_10000234      Reset_Handler:1000018a (*)

```

We can also see the XREF to Reset_Handler as well. Let's update our signature to main which will be int main (void).

```

*****
|
|                               FUNCTION
|
*****
int __stdcall main(void)
int      r0:4      <RETURN>
undefined4      Stack[-0x1c]:4 local_1c      XREF[1]:      1000024c (R)
undefined4      Stack[-0x20]:4 local_20      XREF[1]:      10000244 (R)
undefined4      Stack[-0x28]:4 local_28      XREF[1]:      1000025c (W)
main+1
main      XREF[1,1]:  Reset_Handler:1000018c (c),
          Reset_Handler:1000018a (*)
10000234 f0 b5      push      {r4,r5,r6,r7,lr}
10000236 85 b0      sub      sp,#0x14
10000238 03 f0 a4 f8    bl      FUN_10003384      undefined FUN_10003384()
1000023c 04 20      movs    r0,#0x4
1000023e 00 f0 49 f8    bl      FUN_100002d4      undefined FUN_100002d4()
10000242 14 e0      b       LAB_1000026e
LAB_10000244      XREF[1]:      100002a8 (j)
10000244 02 98      ldr     r0,[sp,#local_20]
10000246 03 f0 13 fb    bl      <EXTERNAL>::thunk_EXT_FUN_200004e8      undefined thunk_EXT_FUN_2
1000024a 04 46      mov     r4,r0
1000024c 03 98      ldr     r0,[sp,#local_1c]
1000024e 0d 46      mov     r5,r1
10000250 03 f0 0e fb    bl      <EXTERNAL>::thunk_EXT_FUN_200004e8      undefined thunk_EXT_FUN_2
10000254 06 46      mov     r6,r0
10000256 0f 46      mov     r7,r1
10000258 22 46      mov     r2,r4
1000025a 2b 46      mov     r3,r5
1000025c cd e9 00 67    strd   r6,r7,[sp,#0x0]=>local_28
10000260 14 48      ldr     r0=>DAT_10003978,[DAT_100002b4]      = 48h      H
          = 10003978h
10000262 03 f0 57 f9    bl      FUN_10003514      undefined FUN_10003514()
LAB_10000266      XREF[1]:      100002b0 (j)
10000266 4f f4 fa 60    mov.w  r0,#0x7d0
1000026a 00 f0 ad fe    bl      FUN_10000fc8      undefined FUN_10000fc8()

```

```

LAB_1000026e                                XREF[1]: 10000242 (j)
1000026e 32 21      movs      r1,#0x32
10000270 11 48      ldr       r0=>s_arithmetic_operator:_%d_100038d8 ,[DAT_1... = "arithmetic_operator: %d\r\n"
                                                = 100038D8h
10000272 03 f0 4f f9  bl       FUN_10003514                          undefined FUN_10003514()
10000276 05 21      movs      r1,#0x5
10000278 10 48      ldr       r0=>s_increment_operator:_%d_100038f4 ,[DAT_10... = "increment_operator: %d\r\n"
                                                = 100038F4h
1000027a 03 f0 4b f9  bl       FUN_10003514                          undefined FUN_10003514()
1000027e 00 21      movs      r1,#0x0
10000280 0f 48      ldr       r0=>s_relational_operator:_%d_10003910 ,[DAT_1... = "relational_operator: %d\r\n"
                                                = 10003910h
10000282 03 f0 47 f9  bl       FUN_10003514                          undefined FUN_10003514()
10000286 00 21      movs      r1,#0x0
10000288 0e 48      ldr       r0=>s_logical_operator:_%d_1000392c ,[DAT_10000...= "logical_operator: %d\r\n"
                                                = 1000392Ch
1000028a 03 f0 43 f9  bl       FUN_10003514                          undefined FUN_10003514()
1000028e 0c 21      movs      r1,#0xc
10000290 0d 48      ldr       r0=>s_bitwise_operator:_%d_10003944 ,[DAT_10000...= "bitwise_operator: %d\r\n"
                                                = 10003944h
10000292 03 f0 3f f9  bl       FUN_10003514                          undefined FUN_10003514()
10000296 0b 21      movs      r1,#0xb
10000298 0c 48      ldr       r0=>s_assignment_operator:_%d_1000395c ,[DAT_1... = "assignment_operator: %d\r\n"
                                                = 1000395Ch
1000029a 03 f0 3b f9  bl       FUN_10003514                          undefined FUN_10003514()
1000029e 03 a9      add      r1,sp,#0xc
100002a0 02 a8      add      r0,sp,#0x8
100002a2 00 f0 27 f8  bl       FUN_100002f4                          undefined FUN_100002f4()
100002a6 00 28      cmp      r0,#0x0
100002a8 cc d1      bne     LAB_10000244
100002aa 09 48      ldr       r0=>s_DHT11_read_failed_100039a4 ,[DAT_100002d0] = "DHT11 read failed\r"
                                                = 100039A4h
100002ac 03 f0 b2 f8  bl       FUN_10003414                          undefined FUN_10003414()
100002b0 d9 e7      b       LAB_10000266
100002b2 00      ??      00h
100002b3 bf      ??      BFh

```

```

          DAT_100002b4          XREF[1]:  main:10000260 (R)
100002b4 78 39 00 10  undefine... 10003978h      ? -> 10003978

          DAT_100002b8          XREF[1]:  main:10000270 (R)
100002b8 d8 38 00 10  undefine... 100038D8h      ? -> 100038d8

          DAT_100002bc          XREF[1]:  main:10000278 (R)
100002bc f4 38 00 10  undefine... 100038F4h      ? -> 100038f4

          DAT_100002c0          XREF[1]:  main:10000280 (R)
100002c0 10 39 00 10  undefine... 10003910h      ? -> 10003910

          DAT_100002c4          XREF[1]:  main:10000288 (R)
100002c4 2c 39 00 10  undefine... 1000392Ch      ? -> 1000392c

          DAT_100002c8          XREF[1]:  main:10000290 (R)
100002c8 44 39 00 10  undefine... 10003944h      ? -> 10003944

          DAT_100002cc          XREF[1]:  main:10000298 (R)
100002cc 5c 39 00 10  undefine... 1000395Ch      ? -> 1000395c

          DAT_100002d0          XREF[1]:  main:100002aa (R)
100002d0 a4 39 00 10  undefine... 100039A4h      ? -> 100039a4

```

Ok so it's a big function, never fear we know exactly what to do! Let's update all of the signatures as follows.

The first function we need to update is at 0x10000238 which is `stdio_init_all` so it will be `bool stdio_init_all (void)`.

The next function we need to update is at 0x1000023e which is `dht11_init`. At this point you may say, how the hell would we know that? The answer is the following...

We discussed before that function arguments will be passed in registers and if there are more than 4 they go on the stack.

Here we see `R0` having `0x4` moved into it. Since we have the device, what is pin 4 attached to? We can literally trace the wire to the DHT11 so we know this is the `dht11_init`. Let's update our signature to `void dht11_init (uint pin)`.

```

1000023c 04 20          movs      r0,#0x4
1000023e 00 f0 49 f8      bl       FUN_100002d4          undefined FUN_100002d4()

```

There are a ton of `printf` so all we have to do is update one and we are set. Let's grab the first one at 0x10000262 and update to `int printf(char *format, ...)` as we know we have `*format` which is a variadic function which can take an unlimited amount of arguments so we have `*` which is a pointer to the address of each argument.

We see a function at 0x1000026a which is `FUN_10000fc8`. What can this be???

Again, what is the argument going into `R0`? We see `0x7d0`. In decimal is 2000.

MCU's typically have sleep functions that work in milliseconds and we know by observation in the last chapter that we saw prints in the terminal every 2 seconds so therefore this must be our `sleep_ms` function.

Let's update the function signature to `void sleep_ms (uint ms)`.

```
10000266 4f f4 fa 60    mov.w    r0,#0x7d0
1000026a 00 f0 ad fe    bl      sleep_ms                                void sleep_ms(uint ms)
```

Next, we see a function with two args. We see `r0`, the first arg, having `0x8` moved into it and `r1`, the second arg, having `0xc` moved into it.

Here we need to think about what is actually going on a bit.

Let's review our source code.

```
float hum, temp;
if (dht11_read(&hum, &temp)) {
    printf("Humidity: %.1f%, Temperature: %.1f°C\r\n", hum, temp);
} else {
    printf("DHT11 read failed\r\n");
}
```

We have not discussed pointers in any detail yet so now is the time.

We start by creating two variables which are `hum` for humidity and `temp` for temperature which are floats. We have discussed floats in the past as you should be familiar with them. The important thing here is we see `&hum` and `&temp`. The `&` operator here represents the "address of" `hum` and the "address of" `temp`.

In Ghidra, we see the following.

```
1000029e 03 a9          add     r1,sp,#0xc
100002a0 02 a8          add     r0,sp,#0x8
100002a2 00 f0 27 f8    bl     FUN_100002f4                undefined FUN_100002f4()
100002a6 00 28          cmp     r0,#0x0
100002a8 cc d1          bne    LAB_10000244
100002aa 09 48          ldr     r0=>s_DHT11_read_failed_100039a4 ,[DAT_100002d0] = "DHT11 read failed\r"
                                                    = 100039A4h
100002ac 03 f0 b2 f8    bl     FUN_10003414                undefined FUN_10003414()
100002b0 d9 e7          b      LAB_10000266
```

The instructions `add r0, sp, #0x8` and `add r1, sp, #0xc` are setting up arguments for the `dht11_read` function by calculating the addresses of two local variables: `hum` and `temp`. These variables are declared as `float hum, temp` and in C, and each occupies 4 bytes. The compiler places them on the stack at offsets relative to the stack pointer (`sp`), with `hum` at `sp + 0x8` and `temp` at `sp + 0xc`. As mentioned, the `&` operator in C retrieves the memory address of a variable, and these assembly instructions replicate that behavior by computing the addresses and storing them in registers `r0` and `r1`, which are then passed to the function.

This layout reflects how local variables are managed in stack frames during function calls. By passing the addresses of `hum` and `temp`, the `dht11_read` function can directly modify their values, allowing the calling function to access

the updated humidity and temperature readings. The use of stack-relative addressing ensures that each variable is properly aligned and accessible, and the offsets (0x8 and 0xc) are chosen based on the size and order of the variables in memory. This is a standard technique in embedded C and assembly programming for managing local data and interfacing with functions that operate on pointers.

Now the question is, looking at this does not clearly identify what this function is as the args are not giving us an idea of what this function is doing so we must look for clues after the function.

We can see clearly that DHT11 is mentioned with a potential read failure. Here our string saved us!

```
100002aa 09 48      ldr      r0=>s_DHT11_read_failed_100039a4 , [DAT_100002d0] = "DHT11 read failed\r"
                                                = 100039A4h
```

Let's update the function at 0x100002a2 to `bool dht11_read(float *humidity, float *temperature)` as we know these are the values.

Finally, we have our last function which is taking the above argument in r0 into a FUN_10003414 which we know the value in r0 is a string as we see in the image above. This is not printf so it must be puts.

Let's update the function signature to `int puts (char *s)` as the arg is a pointer to a string address which will contain the beginning of the string.

```
100002aa 09 48      ldr      r0=>s_DHT11_read_failed_100039a4 , [DAT_100002d0] = "DHT11 read failed\r"
                                                = 100039A4h
100002ac 03 f0 b2 f8  bl      puts      int puts(char * s)
```

In our next lesson we will hack this!

Chapter 28: Hacking Operators

In this chapter we are going to discuss hacking operators with a DHT11 humidity and temperature sensor.

Let's open our project in Ghidra called **0x001a_operators**.


Several chapters have showed you how to manipulate values going into `printf` or `puts` this time we will actually hack the values that will affect humidity and temperature!

```
DAT_1000042c                                XREF[1]:    dht11_read:10000408 (R)
1000042c cd cc cc 3d    undefine... 3DCCCCDh
```

At the very end of our `dht_read` function we will find the above. Let's double-click on the function and go there.

We see `cd cc cc 3d` bytes.

Let's open out our bytes editor. Click **Window** and then **Bytes**.



Here we see our first byte `cd` highlighted. There is a little pencil and paper icon as well. Click that once to enable editing.

A screenshot of the Ghidra Bytes Editor window. The window title is 'Bytes: 0x001a_operators.bin'. It displays a table with two columns: 'Addresses' and 'Hex'. The first byte at address 10000420 is highlighted in blue and contains the hex value 'cd cc cc 3d'.

Addresses	Hex
10000330	42 f2 10 72 4f f0 50 40 08 fa 03 f3 01 e0 01 3a
10000340	50 d0 41 68 19 42 fa d1 42 f2 10 72 4f f0 50 40
10000350	01 e0 01 3a 46 d0 41 68 0b 42 fa d0 42 f2 10 71
10000360	4f f0 50 40 01 e0 01 39 3c d0 42 68 1a 40 fa d1
10000370	10 46 4f f0 50 4c 2c 4e 42 f2 10 72 01 e0 01 3a
10000380	30 d0 dc f8 04 10 0b 42 f9 d0 42 f2 10 72 b7 6a
10000390	01 e0 01 3a 26 d0 dc f8 04 10 0b 42 f9 d1 02 aa
100003a0	02 eb e0 0e b1 6a 1e f8 08 2c c9 1b 52 00 d2 b2
100003b0	28 29 00 f1 01 00 88 bf 42 f0 01 02 28 28 0e f8
100003c0	08 2c d9 d1 9d f8 00 10 9d f8 01 60 9d f8 02 20
100003d0	9d f8 03 00 8b 19 13 44 9d f8 04 70 03 44 db b2
100003e0	9f 42 03 d0 00 20 03 b0 bd e8 f0 83 07 ee 90 6a
100003f0	b8 ee e7 6a 07 ee 90 1a b8 ee e7 7a 07 ee 90 0a
10000400	f8 ee e7 6a 07 ee 90 2a df ed 08 5a f8 ee e7 7a
10000410	a6 ee 25 7a e6 ee a5 7a 85 ed 00 7a 01 20 c4 ed
10000420	00 7a e0 e7 7c 0b 00 20 00 00 0b 40 cd cc cc 3d

The picture shows it not selected so you know what to look for. Once you click it, change `cd cc cc 3d` to `00 00 a0 40` which will increase our temp by 25% 😊.

Let's open a Python shell.

```
>>> import struct
>>> struct.unpack('<f', bytes.fromhex('cdcccc3d'))
(0.10000000149011612,)
>>> struct.unpack('<f', bytes.fromhex('0000a040'))
(5.0,)
```

When computers represent real numbers, they typically use the IEEE 754 floating-point standard. In this format, a 32-bit single-precision float is divided into three fields: one sign bit, 8 exponent bits, and 23 fraction (or mantissa) bits. The value of the float is computed as:


$$\text{Value} = (-1 \text{ to the power of sign}) \times (1 + \text{fraction}) \times 2^{(\text{exponent} - 127)}$$

The sign bit determines whether the number is positive or negative. The exponent controls the scale of the number by powers of two, while the mantissa encodes the significant digits between 1.0 and just under 2.0. Because of this structure, even a single-bit change in the exponent field can dramatically alter the magnitude of the number, while changes in the mantissa adjust the precision within that scale.


To see this in practice, consider the float encodings of 23.8 and 63. The number 23.8 is stored as the hexadecimal value 0x41BE147B. Breaking this down, the sign bit is 0 (positive), the exponent field is 10000011 (131 in decimal), and the mantissa encodes approximately 1.4875. Applying the formula gives $(1.4875 \times 2^4) = 23.8$. By contrast, 63.0 is stored as 0x427C0000. Here the sign bit is still 0, but the exponent has increased to 10000100 (132 in decimal), which doubles the scale factor from 16 to 32. The mantissa is now 1.96875, so the value becomes $(1.96875 \times 32) = 63$. The key observation is that the exponent incremented by one, which doubled the scale, and the mantissa re-normalized to keep the number in the valid range. This combination caused the value to jump from 23.8 to 63. In percentage terms, 63 is about 2.65 times larger than 23.8, which corresponds to a 164.7% increase. This example illustrates how floating-point numbers are not linear in their bit patterns: changing a few bits in the exponent can cause large jumps in value, while changes in the mantissa produce finer adjustments within a given scale.

- The difference is $63 - 23.8 = 39.2$
- Divide by the original value: $39.2 / 23.8 \approx 1.64739$
- Multiply by 100: **164.7% increase**.


Our current sensor readings are as follows.

A screenshot of a PuTTY terminal window titled 'COM4 - PuTTY'. The terminal displays a sequence of sensor readings: 'assignment_operator: 11', 'Humidity: 51.0%, Temperature: 23.8°C', 'arithmetic_operator: 50', 'increment_operator: 5', 'relational_operator: 0', 'logical_operator: 0', 'bitwise_operator: 12', 'assignment_operator: 11', and 'Humidity: 51.0%, Temperature: 23.8°C'. The text is white on a black background.

Now our bytes editor will look like the following.



Let's save out our binary.



0x001a_operators-h.bin

As in the past, we need to use a tool to convert this hacked binary into the UF2 format.

```
python ..\uf2conv.py build\0x001a_operators-h.bin --base 0x10000000 --family 0xe48bff59 --output build\hacked.uf2
```

After flashing the **hacked.uf2** to the Pico 2, we see the following in the serial terminal.

```

COM4 - PuTTY
assignment_operator: 11
Humidity: 51.0%, Temperature: 63.0°C
arithmetic_operator: 50
increment_operator: 5
relational_operator: 0
logical_operator: 0
bitwise_operator: 12
assignment_operator: 11
Humidity: 51.0%, Temperature: 63.0°C
arithmetic_operator: 50
increment_operator: 5
relational_operator: 0
logical_operator: 0
bitwise_operator: 12
assignment_operator: 11
Humidity: 51.0%, Temperature: 63.0°C

```

That is one way to do it lets put back our original values of `cd cc cc 3d` and see what other things we can do!

Addresses	Hex
10000210	7c 0b 00 20 20 e9 10 43 00 20 4a 41 03 b0 03 40
10000300	0c 46 4f f0 01 08 3b 68 8d f8 04 60 00 96 48 ec
10000310	44 30 46 ec 40 30 12 20 00 f0 56 fe 3b 68 48 ec
10000320	40 30 28 20 00 21 00 f0 e3 fd 3b 68 46 ec 44 30
10000330	42 f2 10 72 4f f0 50 40 08 fa 03 f3 01 e0 01 3a
10000340	50 d0 41 68 19 42 fa d1 42 f2 10 72 4f f0 50 40
10000350	01 e0 01 3a 46 d0 41 68 0b 42 fa d0 42 f2 10 71
10000360	4f f0 50 40 01 e0 01 39 3c d0 42 68 1a 40 fa d1
10000370	10 46 4f f0 50 4c 2c 4e 42 f2 10 72 01 e0 01 3a
10000380	30 d0 dc f8 04 10 0b 42 f9 d0 42 f2 10 72 b7 6a
10000390	01 e0 01 3a 26 d0 dc f8 04 10 0b 42 f9 d1 02 aa
100003a0	02 eb e0 0e b1 6a 1e f8 08 2c c9 1b 52 00 d2 b2
100003b0	28 29 00 f1 01 00 88 bf 42 f0 01 02 28 28 0e f8
100003c0	08 2c d9 d1 9d f8 00 10 9d f8 01 60 9d f8 02 20
100003d0	9d f8 03 00 8b 19 13 44 9d f8 04 70 03 44 db b2
100003e0	9f 42 03 d0 00 20 03 b0 bd e8 f0 83 07 ee 90 6a
100003f0	b8 ee e7 6a 07 ee 90 1a b8 ee e7 7a 07 ee 90 0a
10000400	f8 ee e7 6a 07 ee 90 2a df ed 08 5a f8 ee e7 7a
10000410	a6 ee 25 7a e6 ee a5 7a 85 ed 00 7a 01 20 c4 ed
10000420	00 7a e0 e7 7c 0b 00 20 00 00 0b 40 cc cc cc 3d

```

100003ec 07 ee 90 6a vmov s15,r6
100003f0 b8 ee e7 6a vcvf.f32... s12,s15
100003f4 07 ee 90 1a vmov s15,temperature
100003f8 b8 ee e7 7a vcvf.f32... s14,s15
100003fc 07 ee 90 0a vmov s15,humidity
10000400 f8 ee e7 6a vcvf.f32... s13,s15
10000404 07 ee 90 2a vmov s15,r2
10000408 df ed 08 5a vldr.32 s11,[pc,#0x20]=>DAT_1000042c = 3DCCCCCH
1000040c f8 ee e7 7a vcvf.f32... s15,s15
10000410 a6 ee 25 7a vfmf.f32 s14,s12,s11
10000414 e6 ee a5 7a vfmf.f32 s15,s13,s11
10000418 85 ed 00 7a vstr.32 s14,[r5]
1000041c 01 20 movs humidity,#0x1
1000041e c4 ed 00 7a vstr.32 s15,[r4]
10000422 e0 e7 b LAB 100003e6

```

Let's look at 0x10000410 and 0x10000414.

Now that we patched back our binary to its original value, we again read 23.8c. Let's say we want to take 10c from the measurement so that as it gets hotter it will actually report back being significantly cooler when in reality that is simply not true as the normal reading must be 15c otherwise it will trigger an alarm and stop our hack from happening and nuclear missiles will be launched at Washington, D.C.; a bit dramatic an example but you get it.

When working at the binary level, even the smallest change can have dramatic consequences. The patcher we will develop below illustrates this principle by showing how to surgically alter a compiled firmware image at precise offsets. Instead of recompiling the entire project, we identify the exact machine instructions and literal pool constants that control the behavior of the program, then overwrite them with new encodings. This approach is common in embedded forensics and reverse engineering, where source code may not be available but behavior still needs to be modified or corrected.

The patcher is designed with incremental safety in mind. It offers three modes: `pool_only`, `temp_only_vadd`, and `both_vadd`. The first mode is the least invasive, changing only the literal pool constant at offset 0x42C. This constant is loaded into a floating-point register and used in subsequent arithmetic. By altering it from 0.1f to a small negative value, we can observe how the scaling of results changes without touching the instruction stream. This mode is useful for confirming that the patcher is working correctly and that the firmware responds predictably to a single-word modification.

The second mode, `temp_only_vadd`, demonstrates a more targeted intervention. At offset 0x414, the original instruction is a fused multiply-add (`vfmf.f32`) that combines a sensor reading with the pool constant. By rewriting this instruction to a simple floating-point add (`vadd.f32`), and pairing it with a negative constant, we effectively bias the temperature output by a fixed amount. Importantly, humidity remains governed by its original instruction at 0x410, so only temperature is shifted. This illustrates the power of selective patching: by changing a single four-byte opcode, we alter the semantics of one computation while leaving the rest of the system untouched.

The third mode, `both_vadd`, applies the same transformation to both humidity and temperature paths. This is the most invasive option, since it changes two instructions and the shared pool constant. It is included to show how the same technique can be scaled up when a uniform bias is desired across multiple outputs. However, the progression from pool-only to temp-only to both is deliberate: it teaches the importance of incremental testing. Each step builds

confidence that the patcher is functioning as intended, and that the firmware continues to boot and run correctly after modification.

From a pedagogical perspective, this exercise reinforces several key lessons. First, it shows how to map virtual addresses in disassembly to file offsets in a binary image, a critical skill for anyone working with ELF or UF2 formats. Second, it demonstrates the relationship between high-level operations (like “subtract 2.0 from temperature”) and their low-level encodings in ARM Thumb-2 floating-point instructions. Finally, it emphasizes reproducibility and safety: every patch is verified by reading back the modified bytes and printing both their hexadecimal and floating-point interpretations. This ensures that learners not only see the effect of their changes at runtime, but also understand exactly what was written into the binary.

Now let’s review our **hack-temp.py** file as this is located in the main folder.

```
#!/usr/bin/env python3

"""
FILE: hack-temp.py

DESCRIPTION:
Incremental firmware patcher with minimal, conservative steps to bias readings.

BRIEF:
Provides three modes:
- pool_only: change only the literal pool constant (scaling effect; least invasive)
- temp_only_vadd: convert temperature path to vadd.f32 and set a small negative pool
- both_vadd: convert both paths to vadd.f32 and set a small negative pool

Intended to let you test changes incrementally, keeping humidity stable while
you evaluate temperature bias first.

AUTHOR: Kevin Thomas
CREATION DATE: November 1, 2025
UPDATE DATE: November 1, 2025
"""

import sys
import os
import struct

BIN_PATH = "build/0x001a_operators.bin"

# Offsets
OFF_410 = 0x410
OFF_414 = 0x414
OFF_POOL = 0x42C

# Original encodings (for visibility only)
ORIG_410 = bytes.fromhex("a6ee257a") # vfma.f32 s14, s12, s11
ORIG_414 = bytes.fromhex("e6eea57a") # vfma.f32 s15, s13, s11
ORIG_POOL = bytes.fromhex("cccccc3d") # 0.1f
```

```

def read_at(path, offset, n):
    """Read bytes from a binary file at a given offset.

    Parameters
    -----
    path : str
        Filesystem path to the binary file.
    offset : int
        Byte offset from the start of the file.
    n : int
        Number of bytes to read.

    Returns
    -----
    bytes
        The raw bytes read from the file.

    Raises
    -----
    IOError
        If the file cannot be opened or read.
    """
    with open(path, "rb") as f:
        f.seek(offset)
        return f.read(n)

def write_at(path, offset, data):
    """Write bytes into a binary file at a given offset.

    Parameters
    -----
    path : str
        Filesystem path to the binary file.
    offset : int
        Byte offset from the start of the file.
    data : bytes
        Bytes to write into the file.

    Raises
    -----
    IOError
        If the file cannot be opened or written.
    """
    with open(path, "rb+") as f:
        f.seek(offset)
        f.write(data)

```

```

def fmt_float(b):
    """Format a 4-byte sequence as hex and float if possible.

    Parameters
    -----
    b : bytes
        A 4-byte sequence.

    Returns
    -----
    str
        Hexadecimal string with float interpretation if valid.
    """
    hx = b.hex()
    try:
        val = struct.unpack("<f", b)[0]
        return f"{hx} (float {val:.6f})"
    except Exception:
        return hx

def patch_pool_only(new_pool_float=-1.0):
    """Patch only the pool constant to a new float value.

    Parameters
    -----
    new_pool_float : float, optional
        New float value to write into the pool constant (default -1.0).

    Returns
    -----
    None
    """
    print("[mode] pool_only")
    curp = read_at(BIN_PATH, OFF_POOL, 4)
    print(f"Current @0x1000042C: {fmt_float(curp)} (expected {ORIG_POOL.hex()} == 0.1f)")
    newb = struct.pack("<f", new_pool_float)
    write_at(BIN_PATH, OFF_POOL, newb)
    chkp = read_at(BIN_PATH, OFF_POOL, 4)
    print(f"Patched @0x1000042C: {fmt_float(chkp)} (target {newb.hex()} ==
{new_pool_float}f)")

```

```

def patch_temp_only_vadd(new_pool_float=-2.0):
    """Patch temperature path to vadd and set pool to a small negative.

    Parameters
    -----
    new_pool_float : float, optional
        New float value to write into the pool constant (default -2.0).

    Returns
    -----
    None
    """
    print("[mode] temp_only_vadd")
    cur414 = read_at(BIN_PATH, OFF_414, 4)
    curp = read_at(BIN_PATH, OFF_POOL, 4)
    print(f"Current @0x10000414: {cur414.hex()} (expected {ORIG_414.hex()})")
    print(f"Current @0x1000042C: {fmt_float(curp)}")
    write_at(BIN_PATH, OFF_414, VADD_S15_S11)
    write_at(BIN_PATH, OFF_POOL, struct.pack("<f", new_pool_float))
    chk414 = read_at(BIN_PATH, OFF_414, 4)
    chkp = read_at(BIN_PATH, OFF_POOL, 4)
    print(f"Patched @0x10000414: {chk414.hex()} (should be {VADD_S15_S11.hex()})")
    print(f"Patched @0x1000042C: {fmt_float(chkp)} (target {new_pool_float}f)")

def patch_both_vadd(new_pool_float=-2.0):
    """Patch both humidity and temperature paths to vadd and set pool.

    Parameters
    -----
    new_pool_float : float, optional
        New float value to write into the pool constant (default -2.0).

    Returns
    -----
    None
    """
    print("[mode] both_vadd")
    cur410 = read_at(BIN_PATH, OFF_410, 4)
    cur414 = read_at(BIN_PATH, OFF_414, 4)
    curp = read_at(BIN_PATH, OFF_POOL, 4)
    print(f"Current @0x10000410: {cur410.hex()} (expected {ORIG_410.hex()})")
    print(f"Current @0x10000414: {cur414.hex()} (expected {ORIG_414.hex()})")
    print(f"Current @0x1000042C: {fmt_float(curp)}")
    write_at(BIN_PATH, OFF_410, VADD_S14_S11)
    write_at(BIN_PATH, OFF_414, VADD_S15_S11)
    write_at(BIN_PATH, OFF_POOL, struct.pack("<f", new_pool_float))
    chk410 = read_at(BIN_PATH, OFF_410, 4)
    chk414 = read_at(BIN_PATH, OFF_414, 4)
    chkp = read_at(BIN_PATH, OFF_POOL, 4)
    print(f"Patched @0x10000410: {chk410.hex()} (should be {VADD_S14_S11.hex()})")
    print(f"Patched @0x10000414: {chk414.hex()} (should be {VADD_S15_S11.hex()})")
    print(f"Patched @0x1000042C: {fmt_float(chkp)} (target {new_pool_float}f)")

```

```

def main():
    """Dispatch patch mode and perform incremental modifications.

    Returns
    -----
    None

    Raises
    -----
    FileNotFoundError
        If the binary file specified by BIN_PATH does not exist.
    """
    if not os.path.exists(BIN_PATH):
        raise FileNotFoundError(f"Binary not found: {BIN_PATH}")
    mode = (sys.argv[1] if len(sys.argv) > 1 else "pool_only").strip().lower()
    if mode == "pool_only":
        patch_pool_only(new_pool_float=-1.0)
    elif mode == "temp_only_vadd":
        patch_temp_only_vadd(new_pool_float=-2.0)
    elif mode == "both_vadd":
        patch_both_vadd(new_pool_float=-2.0)
    else:
        print(f"Unknown mode: {mode}")
        print("Use: pool_only | temp_only_vadd | both_vadd")
        sys.exit(2)

if __name__ == "__main__":
    main()
    print("Patch complete. Convert to UF2 and flash. Test each mode incrementally.")

```

The very first line, `#!/usr/bin/env python3`, is known as a shebang. It tells Unix-like systems to execute this file using the Python 3 interpreter found in the user's environment. This makes the script directly runnable from the command line without explicitly typing `python3`.

The triple-quoted string that follows is a module-level docstring. It documents the file name, a description of its purpose, a brief summary of its functionality, and metadata such as author and creation date. This is a professional practice that makes the script self-describing and suitable for inclusion in a larger codebase or textbook.

The `import` statements bring in standard Python modules. `sys` is used for accessing command-line arguments and exiting with error codes. `os` provides filesystem utilities such as checking if a file exists. `struct` is essential here because it allows packing and unpacking of binary data into Python types, which is exactly what we need when manipulating machine instructions and floating-point constants.

The constant `BIN_PATH` defines the default path to the binary file we want to patch. This makes the script specific to your build output, but also easy to change if the file is moved. The next three constants, `OFF_410`, `OFF_414`, and `OFF_POOL`, are the file offsets where the relevant instructions and literal pool constant reside. These offsets were determined by analyzing the disassembly and mapping virtual addresses to file positions.

The `ORIG_410`, `ORIG_414`, and `ORIG_POOL` variables hold the expected original bytes at those offsets. They are expressed as hexadecimal strings converted to raw bytes. These serve as sanity checks: when you read the file, you can confirm that you are indeed looking at the unmodified instructions (`vfma.f32`) and the original constant (`0.1f`). The `VADD_S14_S11` and `VADD_S15_S11` variables hold the replacement encodings for `vadd.f32` instructions, packed into little-endian 32-bit words using `struct.pack`.

The function `read_at` encapsulates the logic for reading a sequence of bytes from a file at a given offset. Its docstring follows a structured format, listing parameters, return values, and possible exceptions. Inside, it opens the file in binary mode, seeks to the requested offset, and returns the bytes read. This is the fundamental primitive for inspecting the binary.

The function `write_at` is the complement: it writes a sequence of bytes into the file at a given offset. It opens the file in read-write binary mode, seeks to the offset, and writes the provided data. This is the primitive that actually performs the patching.

The helper `fmt_float` takes a 4-byte sequence and formats it as a hexadecimal string, and if possible, interprets it as a 32-bit IEEE-754 float. This is useful for verifying that the pool constant has the expected numeric value, not just the correct hex encoding.

The function `patch_pool_only` demonstrates the least invasive modification. It prints the current pool value, writes a new float (default `-1.0`) into the pool, and then reads it back to confirm the change. This mode does not touch any instructions, so it only affects scaling behavior.

The function `patch_temp_only_vadd` is more targeted. It prints the current instruction at offset `0x414` and the pool constant, then overwrites the instruction with the encoding for `vadd.f32 s15, s15, s11`. It also writes a small negative float (default `-2.0`) into the pool. Finally, it reads back both the instruction and the pool to verify the patch. This biases the temperature path while leaving humidity unchanged.

The function `patch_both_vadd` applies the same transformation to both humidity and temperature paths. It prints the current instructions and pool, overwrites both instructions with `vadd.f32` encodings, and writes the new pool constant. It then verifies all three modifications. This is the most invasive mode, altering both outputs simultaneously.


The `main` function orchestrates the script. It first checks that the binary file exists, raising a `FileNotFoundError` if not. It then parses the first command-line argument to determine which mode to run, defaulting to `pool_only` if none is provided. Depending on the mode, it calls the appropriate patch function with default float values. If the mode is unrecognized, it prints usage instructions and exits with an error code.

Finally, the `if __name__ == "__main__":` block ensures that `main` is executed only when the script is run directly, not when it is imported as a module. After running `main`, it prints a reminder to convert the patched binary to UF2 format and flash it, and to test each mode incrementally. This reinforces the pedagogical theme of safe, stepwise experimentation.

Run the following...

```
python .\hack-temp.py
```

```
python ..\uf2conv.py build\0x001a_operators.bin --base 0x10000000 --family 0xe48bff59 --  
output build\hacked.uf2
```



```
COM4 - PuTTY
assignment_operator: 11
Humidity: 51.0%, Temperature: 15.0°C
arithmetic_operator: 50
increment_operator: 5
relational_operator: 0
logical_operator: 0
bitwise_operator: 12
assignment_operator: 11
Humidity: 51.0%, Temperature: 15.0°C
arithmetic_operator: 50
increment_operator: 5
relational_operator: 0
logical_operator: 0
bitwise_operator: 12
assignment_operator: 11
Humidity: 51.0%, Temperature: 15.0°C
```

Now if we wanted to do this manually in Ghidra we would go back to our offset and change the bytes from `cc cc cc 3d` to `00 00 80 bf` as we will go through this one more time as this can be complicated.

The IEEE-754 single-precision floating-point format uses 32 bits divided into three fields: one bit for the sign, eight bits for the exponent, and twenty-three bits for the fraction, also called the mantissa. The exponent is stored with a bias of 127, which means the actual exponent is the stored value minus 127. The value of the number is determined by combining the sign, the normalized mantissa, and the adjusted exponent.

In our binary, the constant at address `0x1000042c` is stored as the four bytes `cc cc cc 3d`. Because the system is little-endian, the actual word is `0x3dcccccc`. Written in binary, this is `00111101 11001100 11001100 11001100`. Breaking that down, the sign bit is zero, so the number is positive. The exponent field is `01111011`, which equals 123 in decimal. Subtracting the bias of 127 gives an effective exponent of -4 . The mantissa bits are `10011001100110011001100`, which represent the fractional part. Putting this together, the value is approximately 1.6 multiplied by 2 to the power of -4 , which comes out to about 0.1. That matches the intended constant.


To represent -1.0 , the IEEE-754 rules are straightforward. The sign bit must be one, indicating a negative number. The exponent must be 127, which corresponds to 2 raised to the power of 0. The mantissa must be all zeros, because the value is exactly one with no fractional part. This produces the binary pattern `10111111 10000000 00000000 00000000`, which in hexadecimal is `0xbf800000`. In little-endian byte order, the four bytes are stored as `00 00 80 bf`.

So, at address `0x1000042c`, the firmware originally holds `cc cc cc 3d`, which encodes 0.1. To change this to -1.0 , you overwrite those four bytes with `00 00 80 bf`. The change is entirely explained by flipping the sign bit to one and setting the exponent and mantissa to represent exactly 1. This is why the transformation from `cc cc cc 3d` to `00 00 80 bf` achieves the desired constant.

Addresses	Hex
100002f0	7c 0b 00 20 2d e9 f0 43 00 26 4a 4f 83 b0 05 46
10000300	0c 46 4f f0 01 08 3b 68 8d f8 04 60 00 96 48 ec
10000310	44 30 46 ec 40 30 12 20 00 f0 56 fe 3b 68 48 ec
10000320	40 30 28 20 00 21 00 f0 e3 fd 3b 68 46 ec 44 30
10000330	42 f2 10 72 4f f0 50 40 08 fa 03 f3 01 e0 01 3a
10000340	50 d0 41 68 19 42 fa d1 42 f2 10 72 4f f0 50 40
10000350	01 e0 01 3a 46 d0 41 68 0b 42 fa d0 42 f2 10 71
10000360	4f f0 50 40 01 e0 01 39 3c d0 42 68 1a 40 fa d1
10000370	10 46 4f f0 50 4c 2c 4e 42 f2 10 72 01 e0 01 3a
10000380	30 d0 dc f8 04 10 0b 42 f9 d0 42 f2 10 72 b7 6a
10000390	01 e0 01 3a 26 d0 dc f8 04 10 0b 42 f9 d1 02 aa
100003a0	02 eb e0 0e b1 6a 1e f8 08 2c c9 1b 52 00 d2 b2
100003b0	28 29 00 f1 01 00 88 bf 42 f0 01 02 28 28 0e f8
100003c0	08 2c d9 d1 9d f8 00 10 9d f8 01 60 9d f8 02 20
100003d0	9d f8 03 00 8b 19 13 44 9d f8 04 70 03 44 db b2
100003e0	9f 42 03 d0 00 20 03 b0 bd e8 f0 83 07 ee 90 6a
100003f0	b8 ee e7 6a 07 ee 90 1a b8 ee e7 7a 07 ee 90 0a
10000400	f8 ee e7 6a 07 ee 90 2a df ed 08 5a f8 ee e7 7a
10000410	a6 ee 25 7a e6 ee a5 7a 85 ed 00 7a 01 20 c4 ed
10000420	00 7a e0 e7 7c 0b 00 20 00 00 0b 40 00 00 80 bf

After applying the patch in Ghidra, run the following and we will see the same PuTTY output.

```
python ..\uf2conv.py build\0x001a_operators-h.bin --base 0x10000000 --family 0xe48bff59 --output build\hacked.uf2
```



```
COM4 - PuTTY
assignment_operator: 11
Humidity: 53.0%, Temperature: 15.0°C
arithmetic_operator: 50
increment_operator: 5
relational_operator: 0
logical_operator: 0
bitwise_operator: 12
assignment_operator: 11
Humidity: 53.0%, Temperature: 15.0°C
arithmetic_operator: 50
increment_operator: 5
relational_operator: 0
logical_operator: 0
bitwise_operator: 12
assignment_operator: 11
Humidity: 53.0%, Temperature: 15.0°C
```

Hurray! We did it!

Imagine if we discovered an island off the coast of Eastern Russia that was not on any map and that that no one knew about which housed what is called Dark Eyes which is the equivalent of the Five Eyes but for other nations.

Imagine it was secretly manufacturing nuclear missiles that were 1 month from maturing to which will be launched at Washington, D.C.

Imagine if we needed to task someone to hack the temperature sensor to report a different reading while spinning up their centrifuges but reporting a normal temperature 😊.

Is Embedded Reverse Engineering important? IT'S THAT F*****G IMPORTANT!

Chapter 29: Static Conditionals

In this chapter we are going to discuss static conditionals as well as an intro to PWM with a SG90 servo motor.

Let's open up our folder **0x001d_static-conditionals**.

Now let's review our **0x001d_static-conditionals.c** file as this is located in the main folder.

```
#include <stdio.h>
#include "pico/stdlib.h"
#include "servo.h"

#define SERVO_GPIO 6

int main(void) {
    stdio_init_all();

    int choice = 1;

    servo_init(SERVO_GPIO);

    while (true) {
        if (choice == 1) {
            printf("1\r\n");
        } else if (choice == 2) {
            printf("2\r\n");
        } else {
            printf("?\r\n");
        }

        switch (choice) {
            case 1:
                printf("one\r\n");
                break;
            case 2:
                printf("two\r\n");
                break;
            default:
                printf("??\r\n");
        }

        servo_set_angle(0.0f);
        sleep_ms(500);
        servo_set_angle(180.0f);
        sleep_ms(500);
    }
}
```

In this program we see two intertwined lessons: how static conditionals work in C, and how servo motors are controlled using Pulse Width Modulation (PWM) on the Raspberry Pi Pico 2 (RP2350). The first part of the code sets up a variable called `choice` and evaluates it using both `if/else` statements and a `switch` block. Because `choice` is initialized to 1 and never changes inside the loop, the program will always print "1" and "one" on each iteration. This is an example of a static conditional: the branching logic is present, but the outcome is predetermined because the condition never varies. It is useful for teaching because it shows how different conditional structures behave, but it also highlights that conditionals only become powerful when they respond to changing inputs such as sensor readings, user commands, or external events.

The more substantial lesson comes from the servo driver. A hobby servo such as the SG90 expects a control signal at a fixed frequency of 50 Hz, which means one complete cycle every 20 ms. Within each 20 ms frame, the width of the high pulse encodes the desired angle of the servo shaft. A pulse of about 1000 μ s corresponds to 0°, while a pulse of about 2000 μ s corresponds to 180°. Intermediate values map linearly to intermediate angles. This is why the driver defines `SERVO_DEFAULT_MIN_US` and `SERVO_DEFAULT_MAX_US` as 1000 and 2000 microseconds. The servo does not care about the absolute frequency of the microcontroller's system clock; it only cares that the control signal repeats every 20 ms and that the pulse width falls within its expected range.

On the RP2350, the system clock runs at 150 MHz by default. This is far faster than the 50 Hz signal we need. The PWM hardware bridges this gap by dividing the system clock down to a manageable tick rate and then counting up to a wrap value. In the driver, `servo_wrap` is set to 19,999, meaning the PWM counter will count 20,000 ticks before wrapping back to zero. The clock divider is then calculated so that the counter increments at 1 MHz. With a 1 MHz tick rate, 20,000 ticks take exactly 20 ms, which produces the required 50 Hz frame. In other words, the 150 MHz system clock is divided by 150, yielding 1 MHz, and the wrap value ensures that each PWM cycle lasts 20 ms. This is the critical translation: the microcontroller's fast clock is scaled down to the slow, precise timing that the servo expects.

The function `pulse_us_to_level` performs the conversion from microseconds to PWM counts. Since the counter ticks at 1 MHz, each tick represents 1 μ s. A desired pulse width of 1000 μ s therefore corresponds to 1000 ticks, while 2000 μ s corresponds to 2000 ticks. This value is written to the PWM channel level register, which determines how long the signal stays high during each 20 ms frame. The hardware then automatically generates the waveform, repeating the frame continuously, while the software only needs to adjust the duty cycle to change the pulse width. This separation of concerns is important: the hardware guarantees stable timing, while the software provides the desired control values.

To make the interface more intuitive, the driver provides `servo_set_angle`. Instead of thinking in terms of microseconds, the programmer specifies an angle in degrees. The function clamps the input to the valid range of 0–180°, calculates the corresponding pulse width using a linear mapping, and then calls `servo_set_pulse_us`. This allows the programmer to think in physical terms, angles of rotation, while the driver handles the electrical details. In the main loop, alternating calls to `servo_set_angle(0.0f)` and `servo_set_angle(180.0f)` demonstrate how the servo can be driven back and forth between its extremes. The `sleep_ms(500)` calls provide a half-second delay between movements, making the motion visible and controlled.

Taken together, this example shows how embedded software connects abstract logic to physical hardware. The static conditionals illustrate branching structures in C, while the servo driver demonstrates how PWM translates microcontroller clock cycles into precise timing signals for real-world devices. The RP2350's 150 MHz system clock is far too fast for a servo directly, but by using a clock divider and wrap value, the PWM hardware produces a stable 50 Hz signal with adjustable pulse widths. The programmer can then command angles in degrees, and the

servo responds with mechanical motion. This is the essence of embedded systems: software logic determines what should happen, and hardware peripherals carry out the exact timing needed to make it happen in the physical world.

To make the timing relationship concrete, consider what happens when we command the servo to move to 90 degrees, the midpoint of its range. The driver maps angles linearly between 1000 microseconds (0°) and 2000 microseconds (180°). At 90°, the ratio is halfway, so the pulse width is:

$$\text{Pulse width} = 1000 \mu\text{s} + 0.5 \times (2000 - 1000) \mu\text{s} = 1500 \mu\text{s}$$

This means the control signal should stay high for 1500 microseconds within each 20-millisecond frame.

Now let's see how this translates into PWM counts on the RP2350. As mentioned, the system clock runs at 150 MHz, but the PWM hardware divides it down by 150, giving a tick rate of 1 MHz. At 1 MHz, each tick represents 1 microsecond. The wrap value is set to 19,999, so the counter runs from 0 to 19,999, producing a full cycle length of 20,000 ticks, which equals 20 milliseconds. To generate a 1500 μs pulse, we need the signal to stay high for 1500 ticks before going low. The conversion function `pulse_us_to_level` calculates this directly:

$$\text{Level} = 1500 \mu\text{s} \times 1 \text{ tick}/\mu\text{s} = 1500 \text{ ticks}$$

So, the PWM channel level register is set to 1500. The hardware then ensures that in each 20 ms frame, the signal is high for 1500 ticks (1.5 ms) and low for the remaining 18.5 ms. The servo interprets this as a command to move to its midpoint, 90 degrees.

This worked example shows the complete chain of reasoning: the programmer specifies an angle, the driver converts it to a pulse width, the hardware translates that into counter ticks, and the servo responds with physical motion. By walking through the numbers, students can see how the abstract concept of "90 degrees" becomes a precise electrical signal timed against the RP2350's 150 MHz system clock. It is this bridge between software logic, hardware timing, and mechanical action that makes embedded systems both challenging and rewarding to study.

To reinforce the concept, let's walk through the extremes as well as the midpoint.

Example 1: 0° (minimum angle): At 0°, the driver maps directly to the minimum pulse width of 1000 microseconds. Since the PWM counter ticks at 1 MHz, each tick equals 1 μs . Therefore, the level value is 1000 ticks. In each 20 ms frame, the signal is high for 1000 μs (1 ms) and low for the remaining 19 ms. The servo interprets this as a command to rotate fully to its zero position.

Example 2: 90° (midpoint angle): At 90°, the driver calculates a pulse width halfway between 1000 μs and 2000 μs , which is 1500 μs . This corresponds to 1500 ticks in the PWM counter. The signal is high for 1.5 ms and low for 18.5 ms in each 20 ms frame. The servo interprets this as the midpoint of its travel, holding the shaft at 90°.

Example 3: 180° (maximum angle): At 180°, the driver maps to the maximum pulse width of 2000 μs . This equals 2000 ticks in the PWM counter. The signal is high for 2 ms and low for 18 ms in each 20 ms frame. The servo interprets this as a command to rotate fully to its maximum position.

By comparing these three cases, students can see the linear relationship between angle, pulse width, and PWM counts:


- $0^\circ \rightarrow 1000 \mu\text{s} \rightarrow 1000$ ticks
- $90^\circ \rightarrow 1500 \mu\text{s} \rightarrow 1500$ ticks
- $180^\circ \rightarrow 2000 \mu\text{s} \rightarrow 2000$ ticks

This makes it clear that the driver is simply scaling the angle into the correct pulse width, and the PWM hardware translates that into precise timing against the RP2350's 150 MHz system clock. The servo then responds with mechanical motion. This chain of logic — angle to pulse width to ticks to motion — is the essence of servo control with PWM, and it demonstrates how embedded software bridges abstract logic and physical behavior.

IMPORTANT: 25v 1000uF Capacitor

MAKE ABSOLUTELY SURE YOU CONNECT VBUS (USB 5V) TO A SEPARATE POWER RAIL AND TIE ALL GROUNDS TOGETHER, BECAUSE YOUR SERVO MUST BE POWERED FROM 5V ALONG WITH THE BULK CAPACITOR. DO NOT POWER THE SERVO FROM THE PICO'S 3.3V RAIL. IF YOU DRAW SERVO CURRENT DIRECTLY FROM THE USB PORT, YOU CAN EXCEED 500 MA AND CAUSE BROWNOUTS OR DAMAGE, SINCE SERVO SPIKES CAN GO OVER 1000 MA. USE A STRONG 5V SUPPLY AND KEEP GROUNDS COMMON.

In PuTTY, we see the following along with our servo moving.



```
COM4 - PuTTY
1
one
1
one
1
one
1
one
1
one
```

In our next chapter we will debug this.

Chapter 30: Debugging Static Conditionals

In this chapter we are going to discuss debugging static conditionals as well as an intro to PWM with a SG90 servo motor.

Let's create a new project in Ghidra called **0x001d_static-conditionals**.

As in prior chapters, it will be a Cortex ARM:LE:32 little endian and we have to set the options flash to `0x10000000` which is the base of XIP as we have discussed.

As always, we know that the `Reset_Handler` leads us to `main`. We know working through multiple examples that `main` is on or about `0x10000234` however we also know that in ARM the address of offset 4 from the base of XIP is the address of our `Reset_Handler`.

```

*****
*                               FUNCTION                               *
*****
undefined FUN_10000234 ()
undefined  ▲<UNASSIGNED> <RETURN>
FUN_10000234+1          XREF[1,1]:  1000018c (c) , 1000018a (*)
FUN_10000234
10000234 38 b5          push    {r3,r4,r5,lr}
10000236 01 f0 dd fa    bl      FUN_100017f4          undefined FUN_100017f4()
1000023a 06 20          movs   r0,#0x6
1000023c 00 f0 le f8        bl      FUN_1000027c          undefined FUN_1000027c()
10000240 00 25          movs   r5,#0x0
10000242 0b 4c          ldr    r4,[DAT_10000270]
                                     = 43340000h

LAB_10000244          XREF[1]:  1000026c (j)
10000244 0b 48          ldr    r0=>DAT_10001c54,[DAT_10000274]
                                     = 00000D31h
                                     = 10001C54h
10000246 01 f0 ld fb        bl      FUN_10001884          undefined FUN_10001884()
1000024a 0b 48          ldr    r0=>DAT_10001c5c,[DAT_10000278]
                                     = 6Fh   o
                                     = 10001C5Ch
1000024c 01 f0 la fb        bl      FUN_10001884          undefined FUN_10001884()
10000250 28 46          mov    r0,r5
10000252 00 f0 5d f8        bl      FUN_10000310          undefined FUN_10000310()
10000256 4f f4 fa 70        mov.w  r0,#0x1f4
1000025a 00 f0 e1 fd        bl      FUN_10000e20          undefined FUN_10000e20()
1000025e 20 46          mov    r0,r4
10000260 00 f0 56 f8        bl      FUN_10000310          undefined FUN_10000310()
10000264 4f f4 fa 70        mov.w  r0,#0x1f4
10000268 00 f0 da fd        bl      FUN_10000e20          undefined FUN_10000e20()
1000026c ea e7          b      LAB_10000244
1000026e 00          ??    00h
1000026f bf          ??    BFh

DAT_10000270          XREF[1]:  FUN_10000234:10000242 (R)
10000270 00 00 34 43      undefine... 43340000h

DAT_10000274          XREF[1]:  FUN_10000234:10000244 (R)
10000274 54 1c 00 10      undefine... 10001C54h
                                     ? -> 10001c54

DAT_10000278          XREF[1]:  FUN_10000234:1000024a (R)
10000278 5c 1c 00 10      undefine... 10001C5Ch
                                     ? -> 10001c5c

```

Let's update our main function to `int main(void)` at `0x10000234`.

We know `FUN_100017f4` is `stdio_init_all` so let's update the function signature to `bool stdio_init_all(void)`.

Next, we see the following.

```
1000023a 06 20      movs    r0,#0x6
1000023c 00 f0 1e f8    bl     FUN_1000027c                undefined FUN_1000027c()
```

As we know, we see `r0` with an immediate value so this is a parameter. We have our breadboard and know the servo GPIO is 6 so this must be the `servo_init` function. Let's update it to `void servo_init(uint pin)`.

Next, we see the following.

```
10000244 0b 48      ldr     r0=>DAT_10001c54,[DAT_10000274]    = 00000d31h
                                           = 10001c54h
10000246 01 f0 1d fb    bl     FUN_10001884                undefined FUN_10001884()
1000024a 0b 48      ldr     r0=>DAT_10001c5c,[DAT_10000278]    = 6fh      o
                                           = 10001c5ch
1000024c 01 f0 1a fb    bl     FUN_10001884                undefined FUN_10001884()
```

Going into `r0`, we see a pointer to an address and the same function called twice. Here this is clearly our `puts` function. Let's update it to `int puts(char *s)`.

Now how do we know this is `puts`? We see a value `0x0000d31h` so in the `ascii` table we know that in little endian `0x31` is "1" and `0x0d` is "\r" and well `0x00` is "\0" the null terminator. In addition, in PuTTY we saw "1" echoed so this must be `puts`. Finally, in PuTTY, we saw "one" and we can clearly see `6fh` or the letter "o" as this is a pointer to the string.

Next, we see the following.

```
10000250 28 46      mov     r0,r5
10000252 00 f0 5d f8    bl     FUN_10000310                undefined FUN_10000310()
```

Ok this is not helpful as we have no idea what is in `r5`! Let's double-click on the function and see what is inside to get a better idea.

```
10000358 4f f4 fa 63    mov.cs.w    r3,#0x7d0
1000035c b3 f5 7a 7f    cmp.w      r3,#0x3e8
```

As we look around, we see two hard-coded hex values of `0x7D0` and `0x3E8` which are 2000 and 1000 respectively, so this gives it away as we know the basics of our PWM program and these numbers represent the pulse-width limits used in `servo_set_angle`.

In other words, 1000 corresponds to a 1.0 ms pulse (the minimum duty cycle that drives the servo to one extreme, typically 0°), while 2000 corresponds to a 2.0 ms pulse (the maximum duty cycle that drives the servo to the opposite extreme, typically 180°). The function maps the requested angle into this range, scaling linearly between 1000 and 2000 microseconds, so every intermediate value produces the correct position of the servo arm.

Let's update this function to `void servo_set_angle(float degrees)`.

Finally, we have the following.

```
10000256 4f f4 fa 70    mov.w      r0,#0x1f4
1000025a 00 f0 e1 fd    bl        FUN_10000e20      undefined FUN_10000e20()
```

This is easy as we see it repeated again after the next `servo_set_angle` so this is obviously our `sleep_ms` as we know `0x1f4` is 500 in decimal our delay. Let's update our function signature to `void sleep_ms(uint ms)`.

In our next lesson we will hack this!

Chapter 31: Hacking Static Conditionals

In this chapter we are going to discuss hacking debugging static conditionals as well as an intro to PWM with a SG90 servo motor.


Let's open our project in Ghidra called **0x001d_static-conditionals**.

We will start with the trivial hacking of “1” and “one”.

Let's double-click on 10001c54h.

```
DAT_10000274                                     XREF[1]:    main:10000244 (R)
10000274 54 1c 00 10    undefine... 10001C54h                                ? -> 10001c54
```

Let's open our bytes editor which we used in prior chapters.



Addresses	Hex
10001b30	40 f8 04 3b 04 3a f9 d2 04 32 08 d0 d2 07 1c bf
10001b40	11 f8 01 3b 00 f8 01 3b 01 d3 0b 88 03 80 60 46
10001b50	70 47 00 bf 08 2a 13 d3 8b 07 b1 d0 10 f0 03 03
10001b60	ae d0 c3 f1 04 03 d2 1a db 07 1c bf 11 f8 01 3b
10001b70	00 f8 01 3b a4 d3 31 f8 02 3b 20 f8 02 3b 9f e7
10001b80	04 3a d9 d3 01 3a 11 f8 01 3b 00 f8 01 3b f9 d2
10001b90	0b 78 03 70 4b 78 43 70 8b 78 83 70 60 46 70 47
10001ba0	20 f0 03 01 10 f0 03 00 c0 f1 00 00 51 f8 04 3b
10001bb0	00 f1 04 0c 4f ea cc 0c 6f f0 00 02 1c bf 22 fa
10001bc0	0c f2 13 43 4f f0 01 0c 4c ea 0c 2c 4c ea 0c 4c
10001bd0	a3 eb 0c 02 22 ea 03 02 12 ea cc 12 04 bf 51 f8
10001be0	04 3b 04 30 f4 d0 c2 f1 00 01 02 ea 01 02 b2 fa
10001bf0	82 f2 c2 f1 1f 02 00 eb d2 00 70 47 f8 b5 00 bf
10001c00	5f f8 00 f0 e5 01 00 20 5f f8 00 f0 7d 01 00 20
10001c10	51 14 00 10 c5 13 00 10 f9 13 00 10 e5 14 00 10
10001c20	31 14 00 10 49 13 00 10 79 14 00 10 35 13 00 10
10001c30	5d 08 00 10 7d 14 00 10 d9 0b 00 10 95 13 00 10
10001c40	65 14 00 10 15 14 00 10 ad 04 00 10 d5 04 00 10
10001c50	11 02 00 10 31 0d 00 00 00 00 00 00 6f 6e 65 0d

We know 0x31 is 1 so let's change it to 2. Click on the pencil icon and update it to 0x32.

Let's change “one” to “fun” shall we! 😊

```
DAT_10001c5c                                     XREF[1]:    main:1000024a (*)
10001c5c 6f    ??    6Fh    o
10001c5d 6e    ??    6Eh    n
10001c5e 65    ??    65h    e
```

Click on 6fh and it will highlight in the bytes editor and overwrite 6f 6e 65 to 66 75 6e.

Now in Ghidra we see it updated.

```
          DAT_10001c5c          XREF[1]:  main:1000024a (*)
10001c5c 66          ??          66h    f
10001c5d 75          ??          75h    u
10001c5e 6e          ??          6Eh    n
```


Now let's have some REAL fun and speed up this servo to something silly!

```
10000252 00 f0 5d f8  b1      servo_set_angle      void servo_set_angle(float
10000256 4f f4 fa 70  mov.w    r0,#0x1f4
1000025a 00 f0 e1 fd  b1      sleep_ms             void sleep_ms(uint ms)
1000025e 20 46          mov      r0,r4
10000260 00 f0 56 f8  b1      servo_set_angle      void servo_set_angle(float
10000264 4f f4 fa 70  mov.w    r0,#0x1f4
10000268 00 f0 da fd  b1      sleep_ms             void sleep_ms(uint ms)
```

We know 0x1f4 is 500 so... What if we wanted to reduce the delay to say 100! Let's patch both values from 0x1f4 to 0x64.

```
10000252 00 f0 5d f8  b1      servo_set_angle      void servo_set_angle(float
10000256 4f f0 64 00  mov.w    r0,#0x64
1000025a 00 f0 e1 fd  b1      sleep_ms             void sleep_ms(uint ms)
1000025e 20 46          mov      r0,r4
10000260 00 f0 56 f8  b1      servo_set_angle      void servo_set_angle(float
10000264 4f f0 64 00  mov.w    r0,#0x64
10000268 00 f0 da fd  b1      sleep_ms             void sleep_ms(uint ms)
```

Let's save out our binary.



0x001d_static-conditionals-h.bin

As always, let's convert this hacked binary into the UF2 format.

```
python ..\uf2conv.py build\0x001d_static-conditionals-h.bin --base 0x10000000 --family 0xe48bff59 --output build\hacked.uf2
```

After flashing the **hacked.uf2** to the Pico 2, we see the following in the serial terminal.

