

HALucinator: Firmware Re-hosting Through Abstraction Layer Emulation


Abraham Clements*, **Eric Gustafson***, Tobias Scharnowski, Paul Grosen, David Fritz,
Christopher Kruegel, Giovanni Vigna, Saurabh Bagchi, and Mathias Payer





IoT and Operational Technology



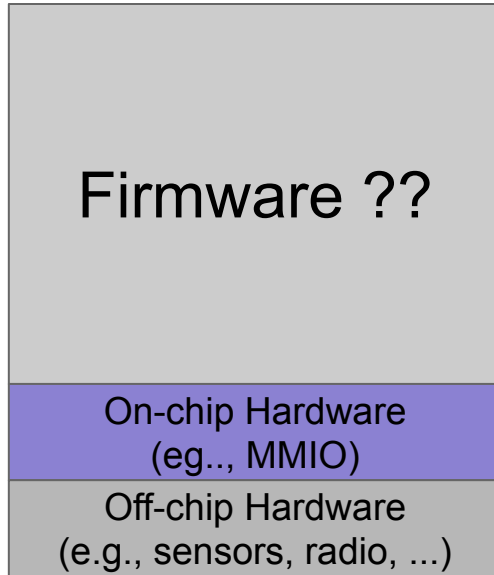
SECLAB



Device Internals

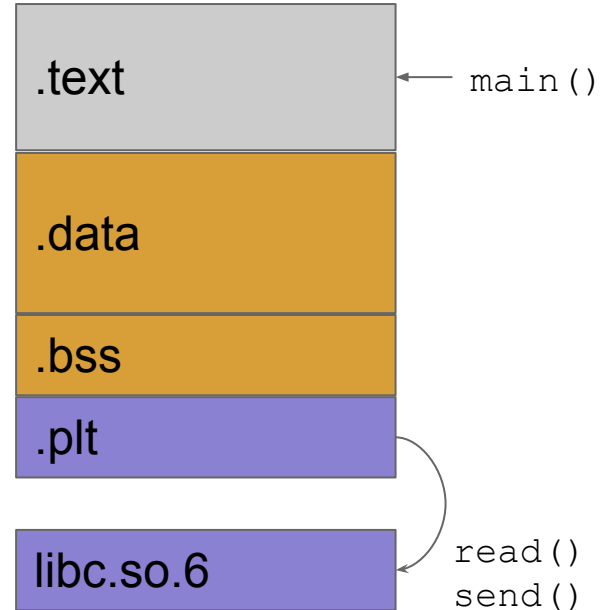


Baremetal



Raw hardware access

Linux ELF file



Kernel abstractions used for hardware interactions


Debug access

- Should be disabled
- If present, very limited

Limits parallelism

Other limitations



- Can be expensive (\$100 - \$10k)
- Brittle - easily bricked



Re-hosting to the Rescue?



SECLAB



**HALucinator's Goal:
Enable scalable firmware testing without requiring specialized hardware**

Peripherals



On chip

CPU
AES Accelerator
Hash
Coprocessor
Timers
Counters
Flash Controller
Clock Config
IAP
DMA

Off chip

Ethernet
SD-MMC
GPIO
Camera
LCD
Touch Screen
Wireless
EEPROM
Serial
CAN
Analog IO
USB

Peripherals Prevent Re-hosting



SECLAB

Peripherals



On chip	Off chip
CPU	Ethernet
AES Accelerator	SD-MMC
...	...

Mouser Lists
44,520 Microcontrollers
3,502 Datasheets
26 Manufacturers

...

Peripherals Prevent Re-hosting



SECLAB

Peripherals

On chip

CPU
AES Accelerator

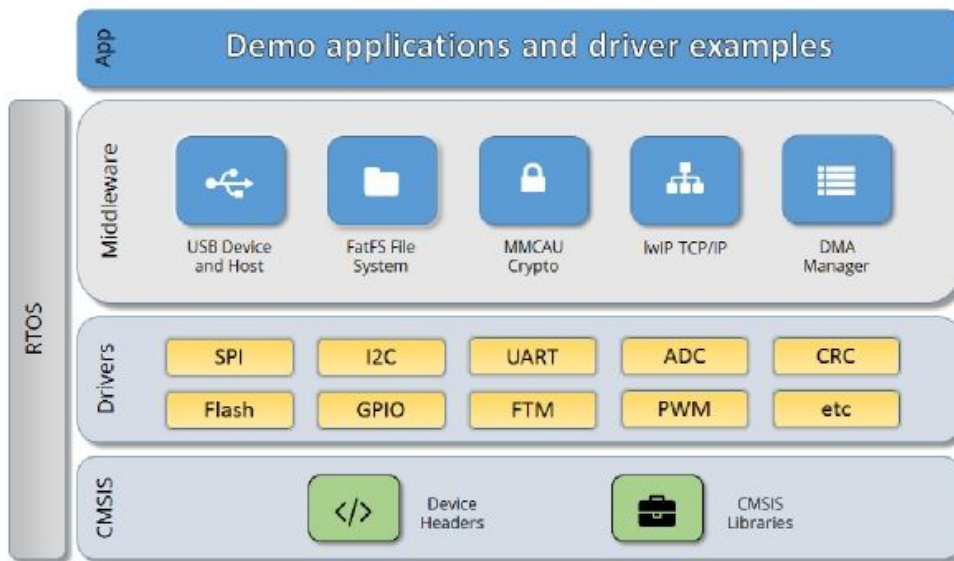
Off chip

Ethernet
SD-MMC



**Without support for peripherals baremetal firmware will not run!
There are 10,000's of peripherals and combinations there of!**

7. Analog IC
USB



HALs are Everywhere




SECLAB

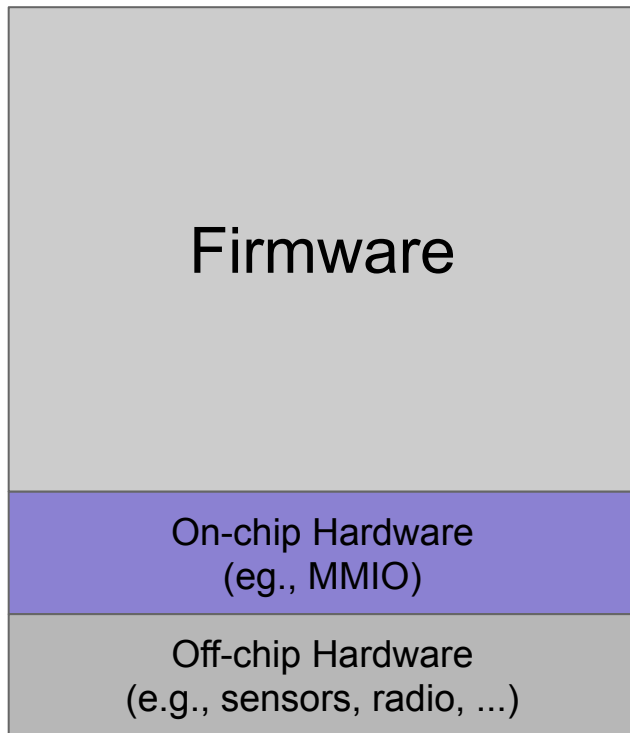




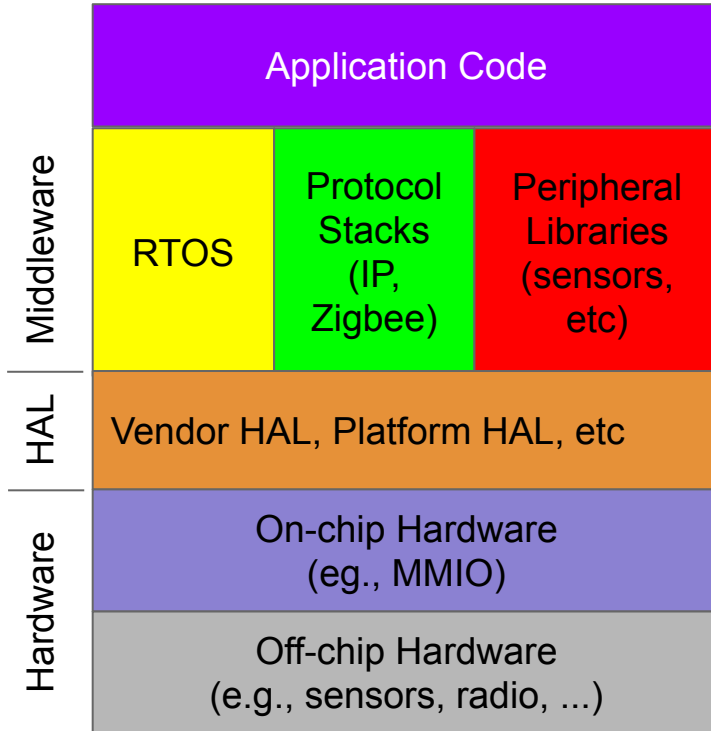
HALucinator

Enables replacing HALs and other libraries with high level implementations. Transforming the re-hosting scaling problem from supporting 10,000's of devices to dozens of HALS

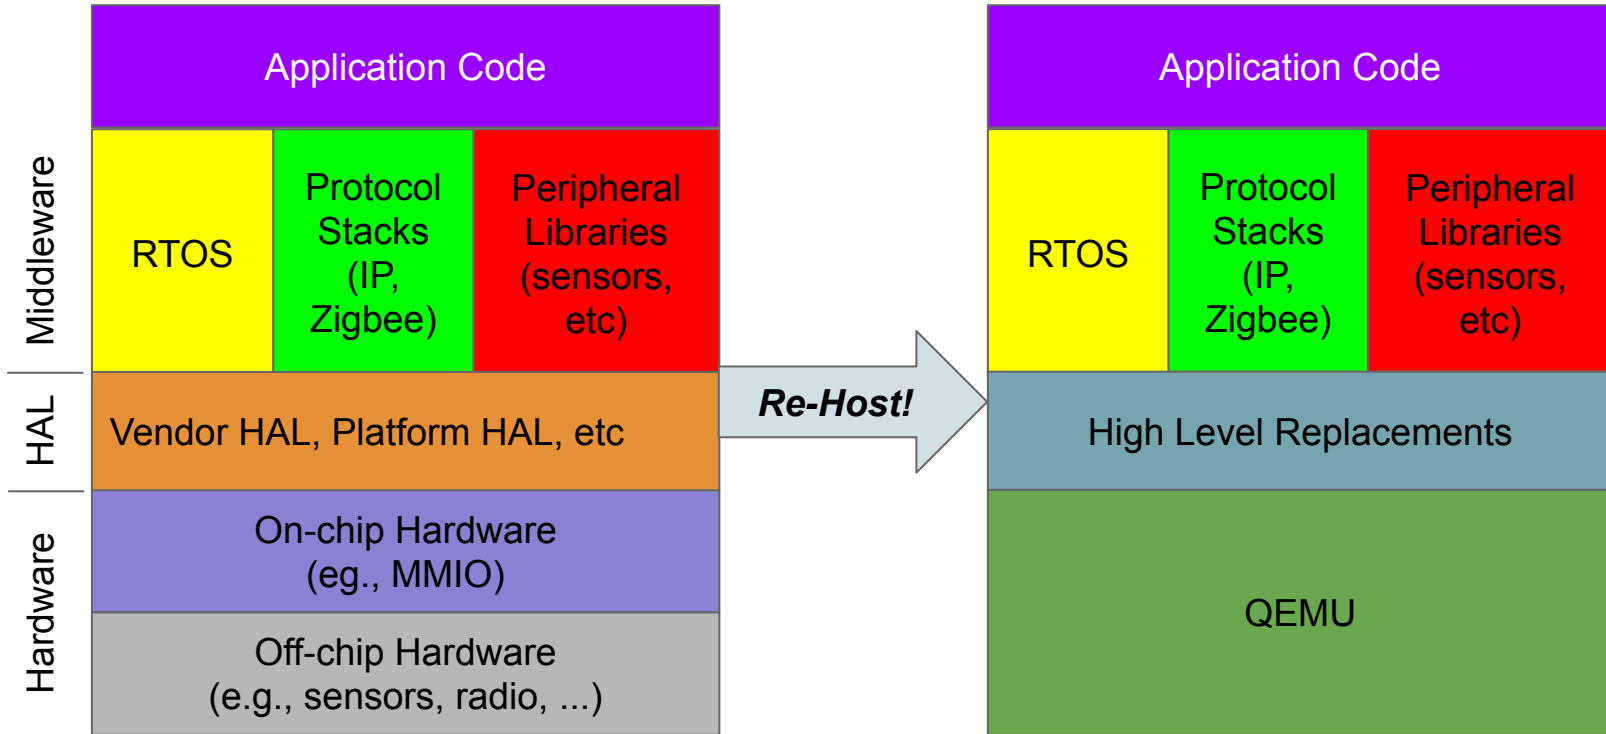




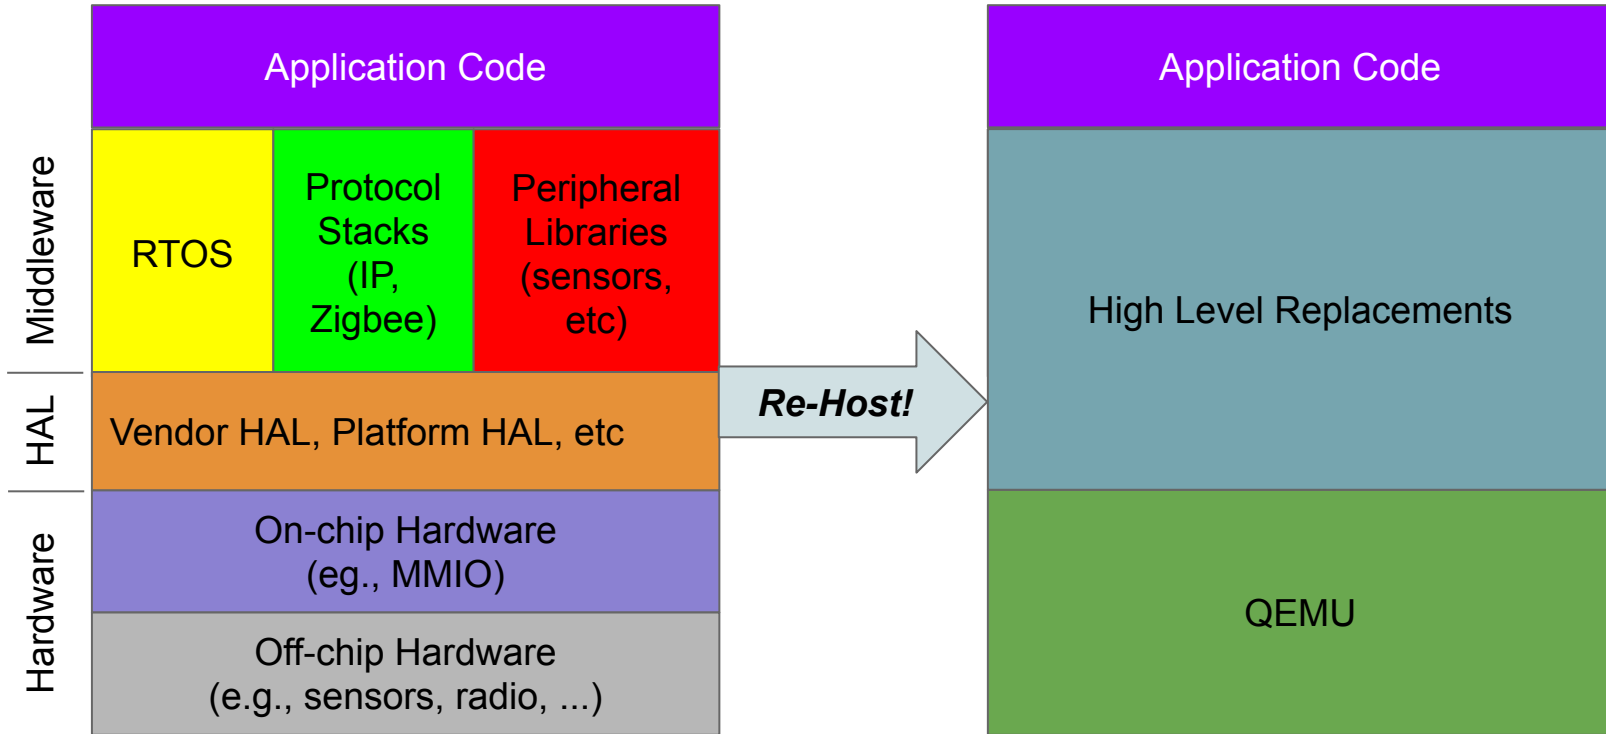
High Level Emulation



High Level Emulation



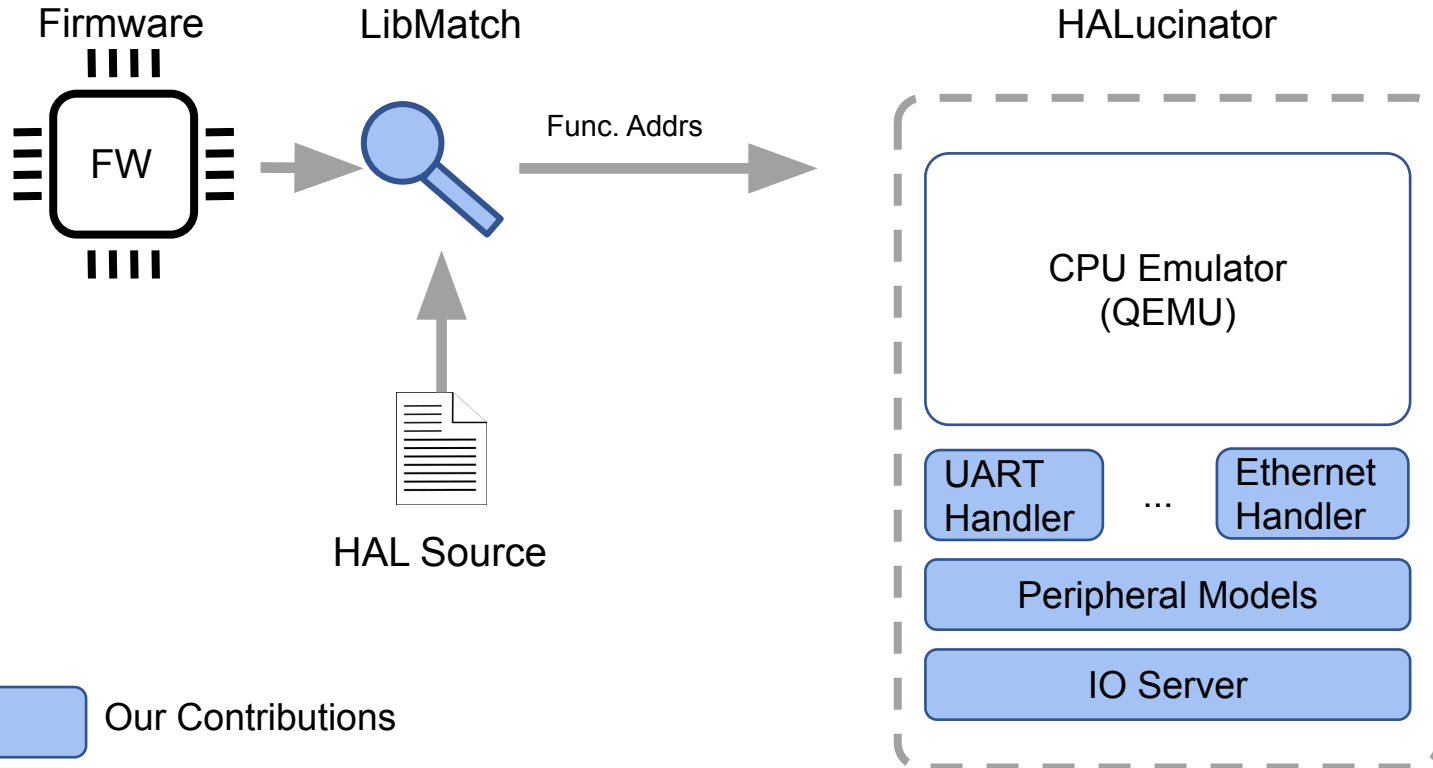
High Level Emulation




HALucinator implementation




SECLAB

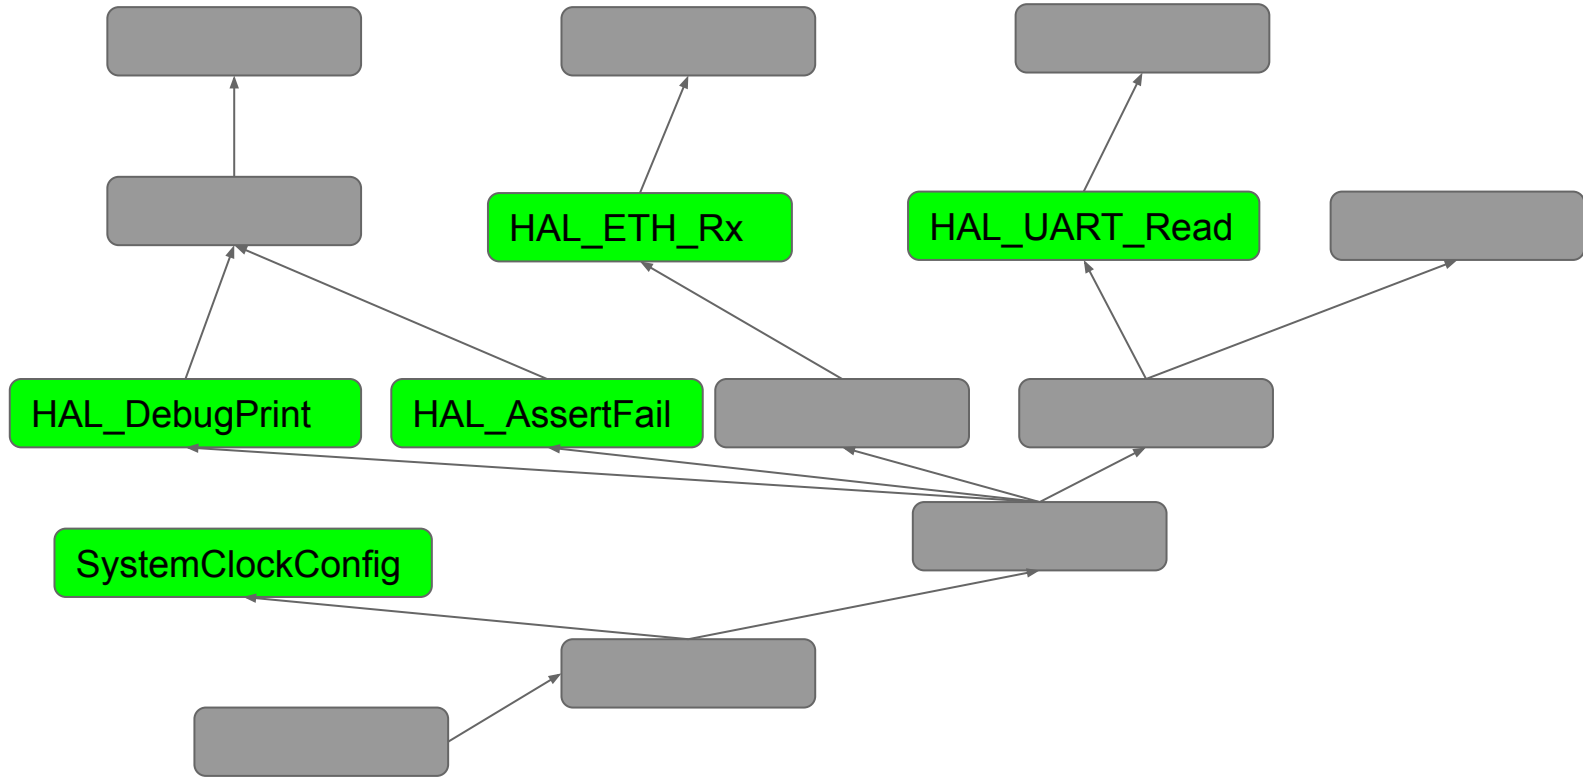


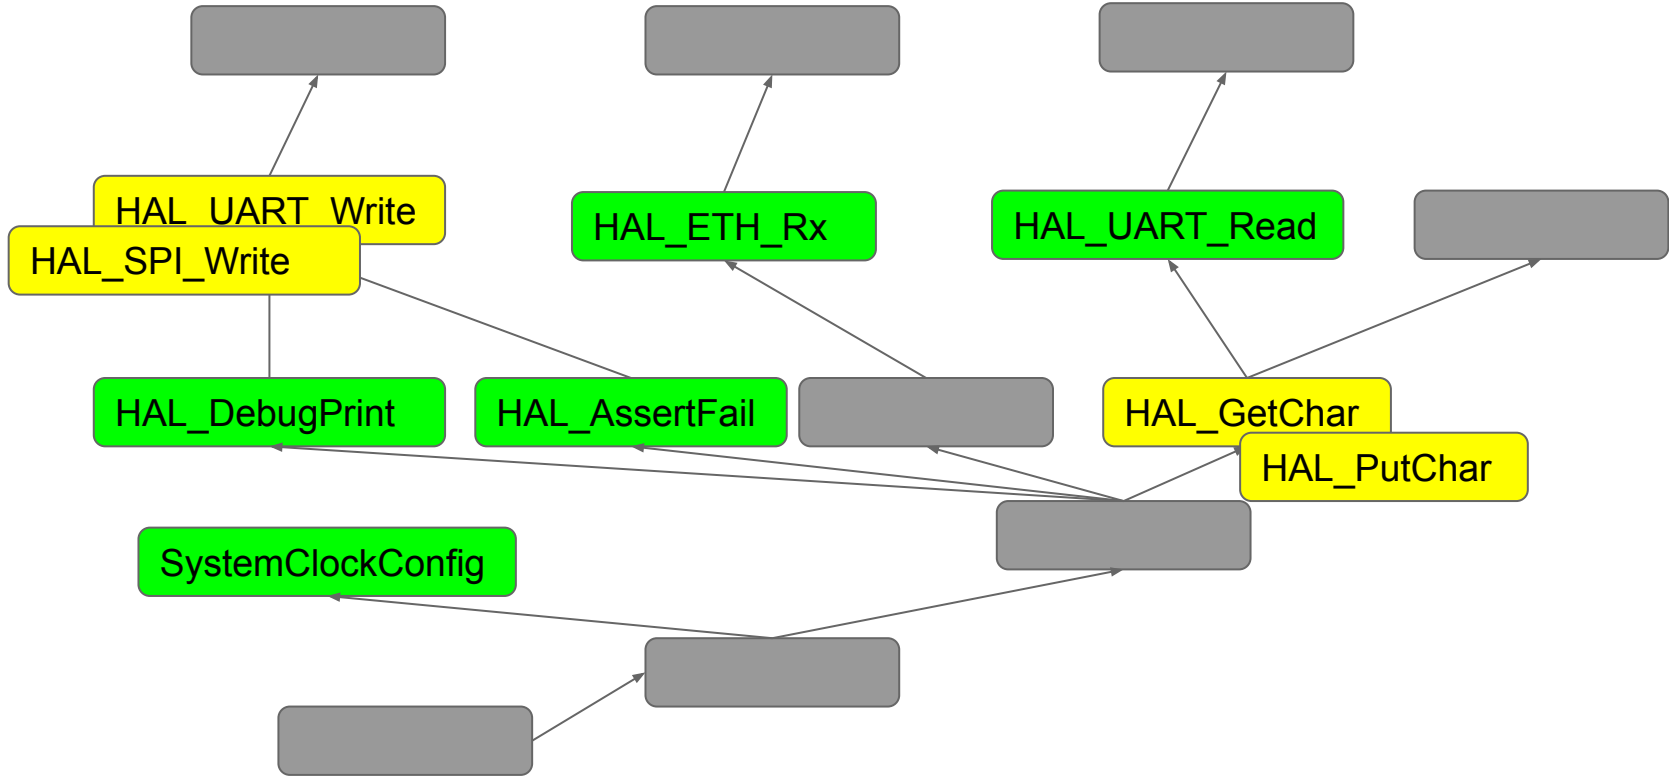
Handler Example

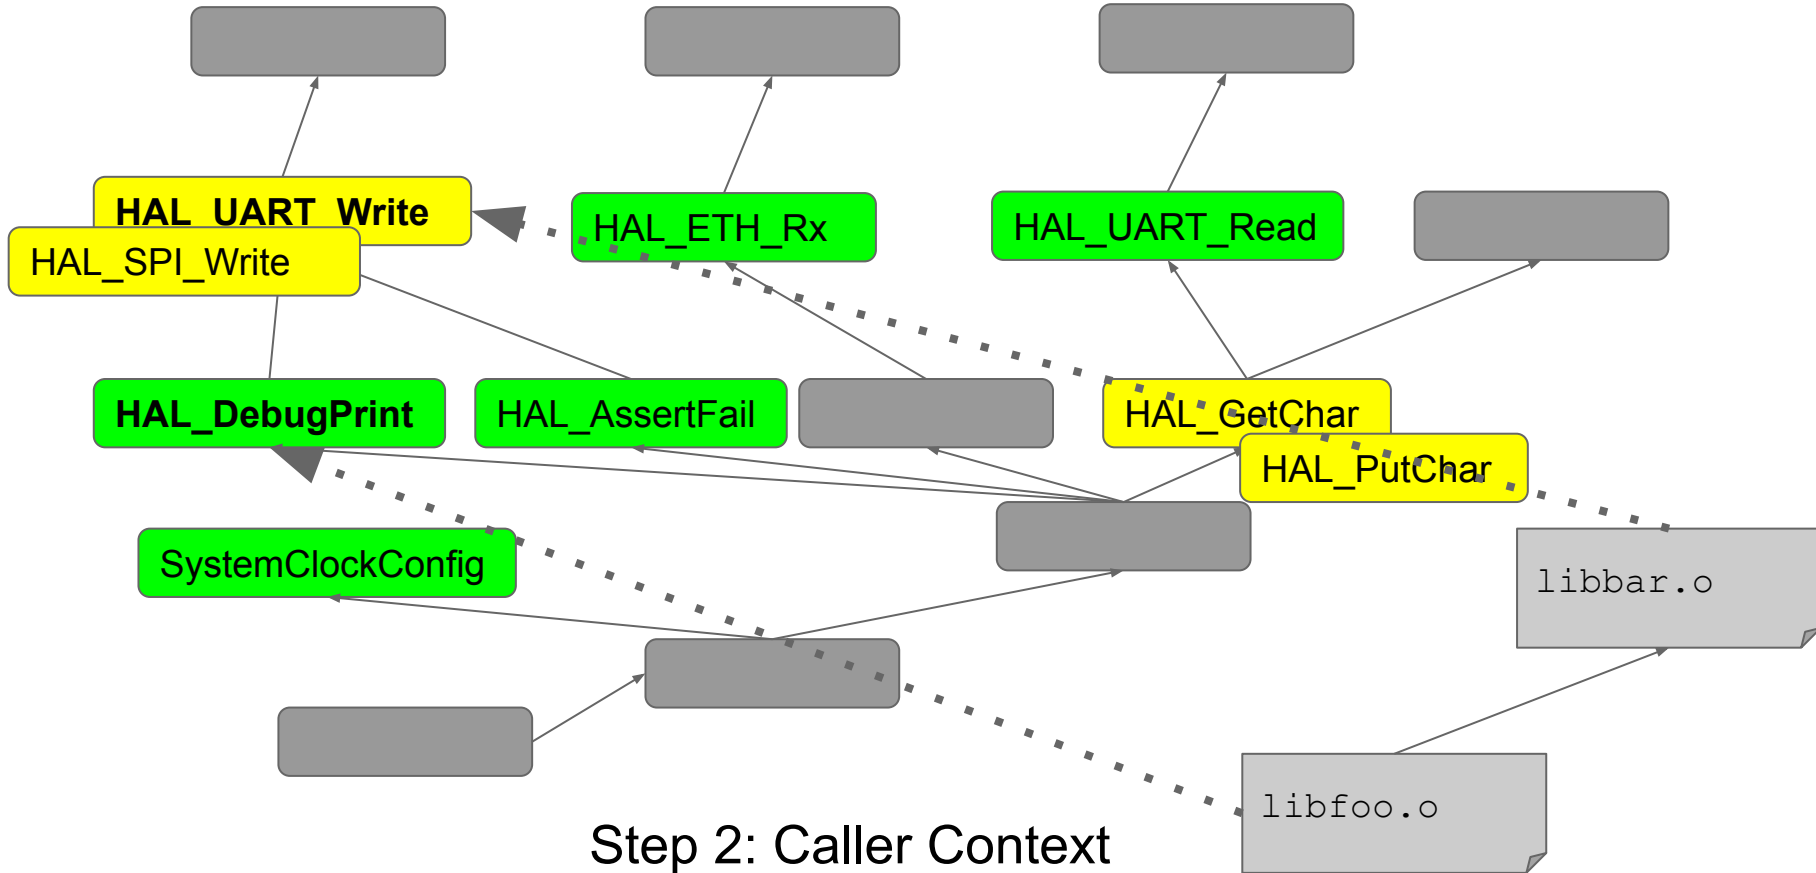


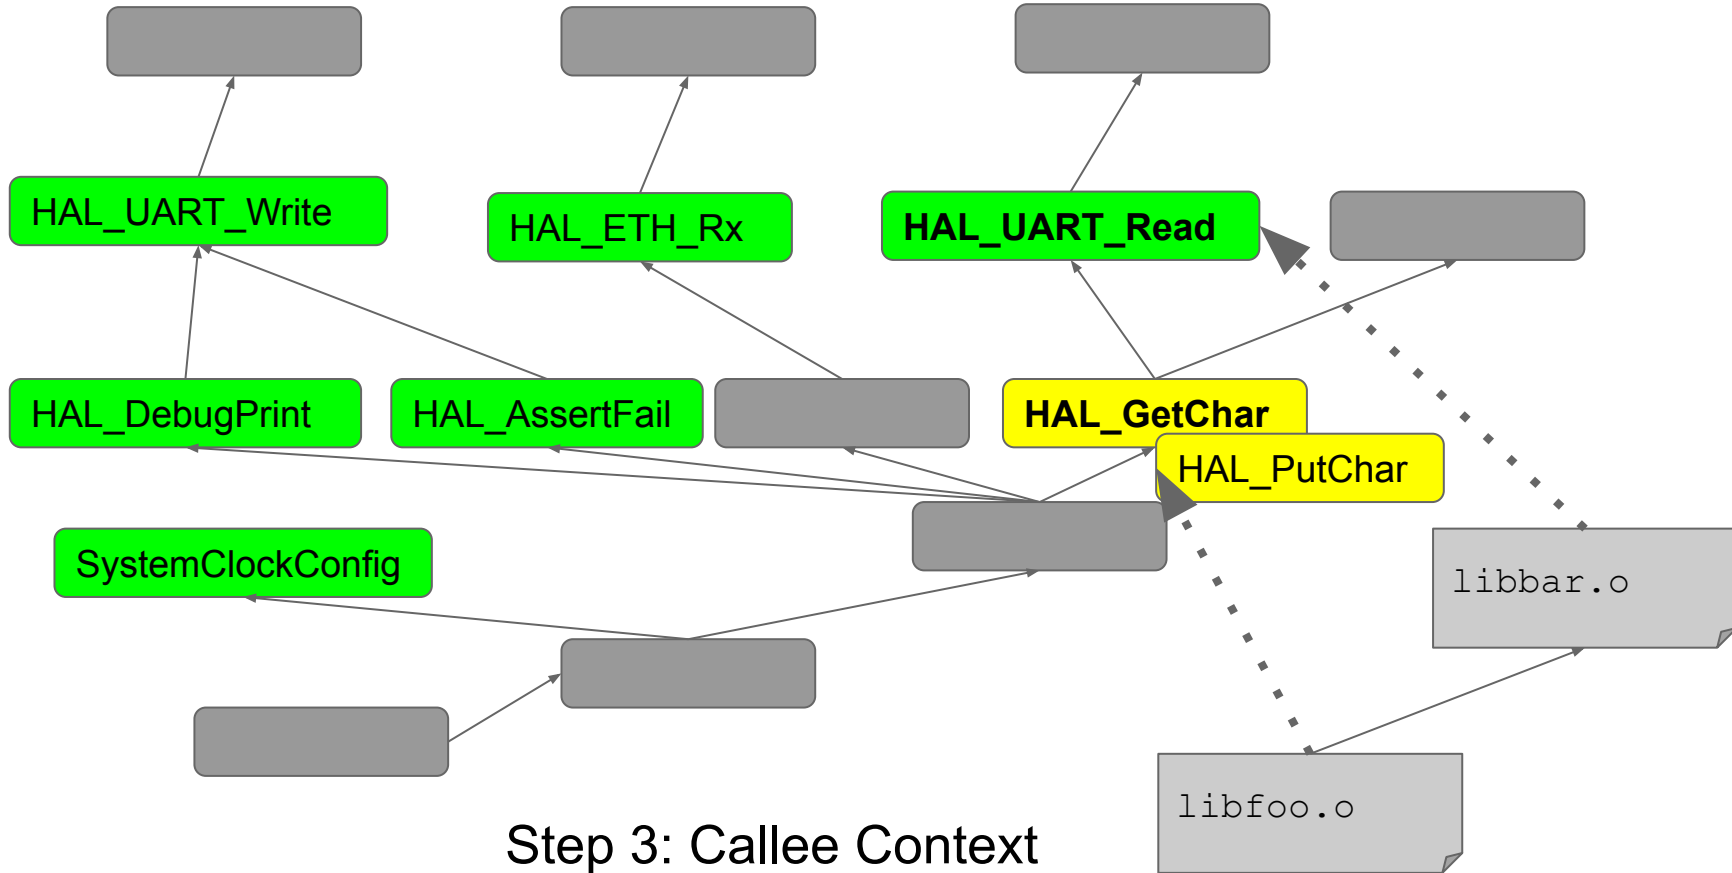
```
def i2c_read_buf(uc):  
    # i2c_read_buf(char* buf, int len);  
    buf = uc.regs.r1 # arg 0: The buffer  
    l = uc.regs.r2   # arg 1: Buffer length  
    assert(buf != 0) # Crash on bad arguments  
    assert(len > 0)  
    data = I2CModel.rx('i2c', 0, len) # Get the data  
                                       # from the virtual bus  
    uc.mem[buf] = data                 # Store it in the emulator
```

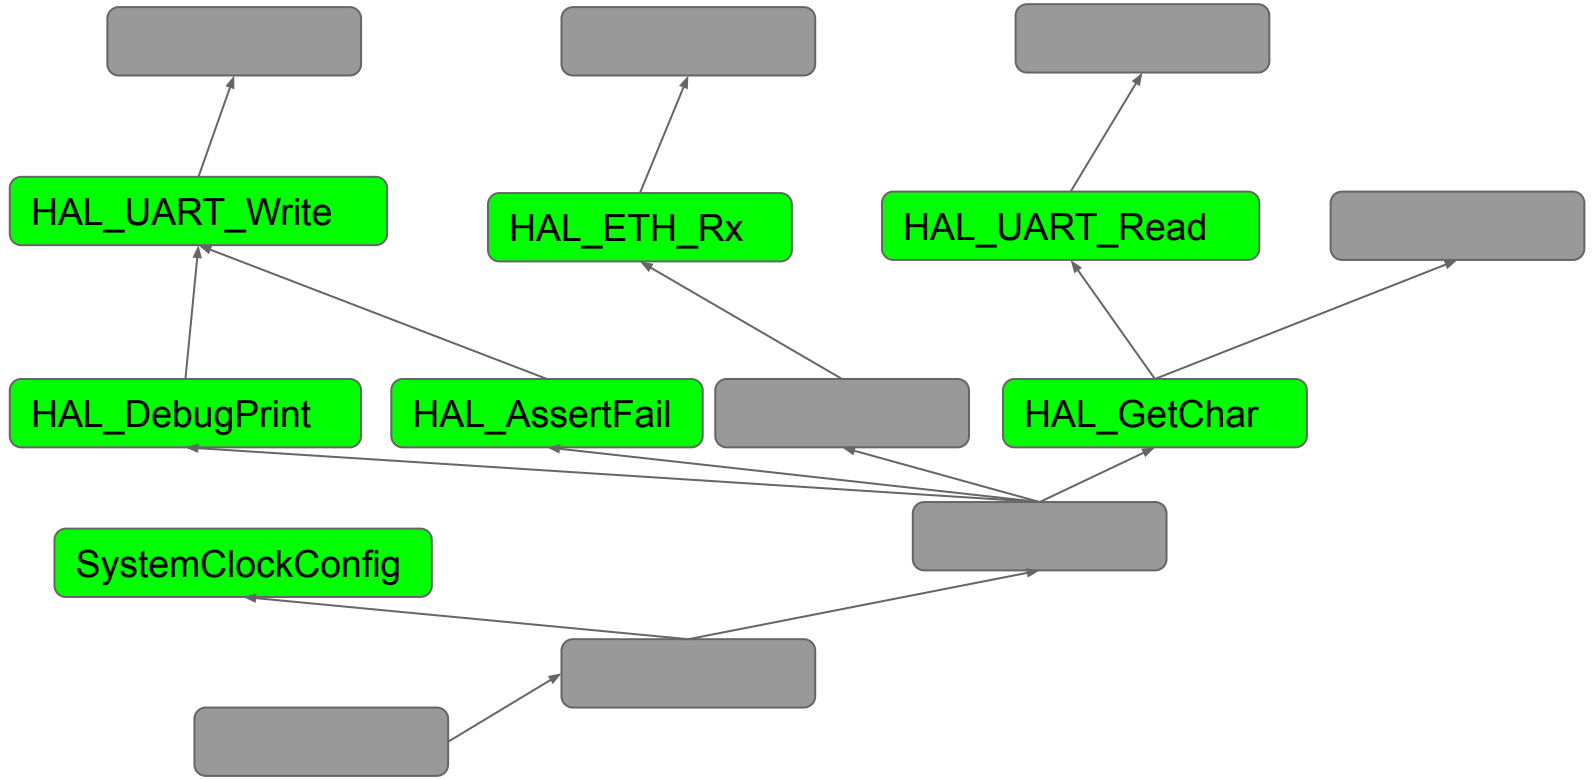


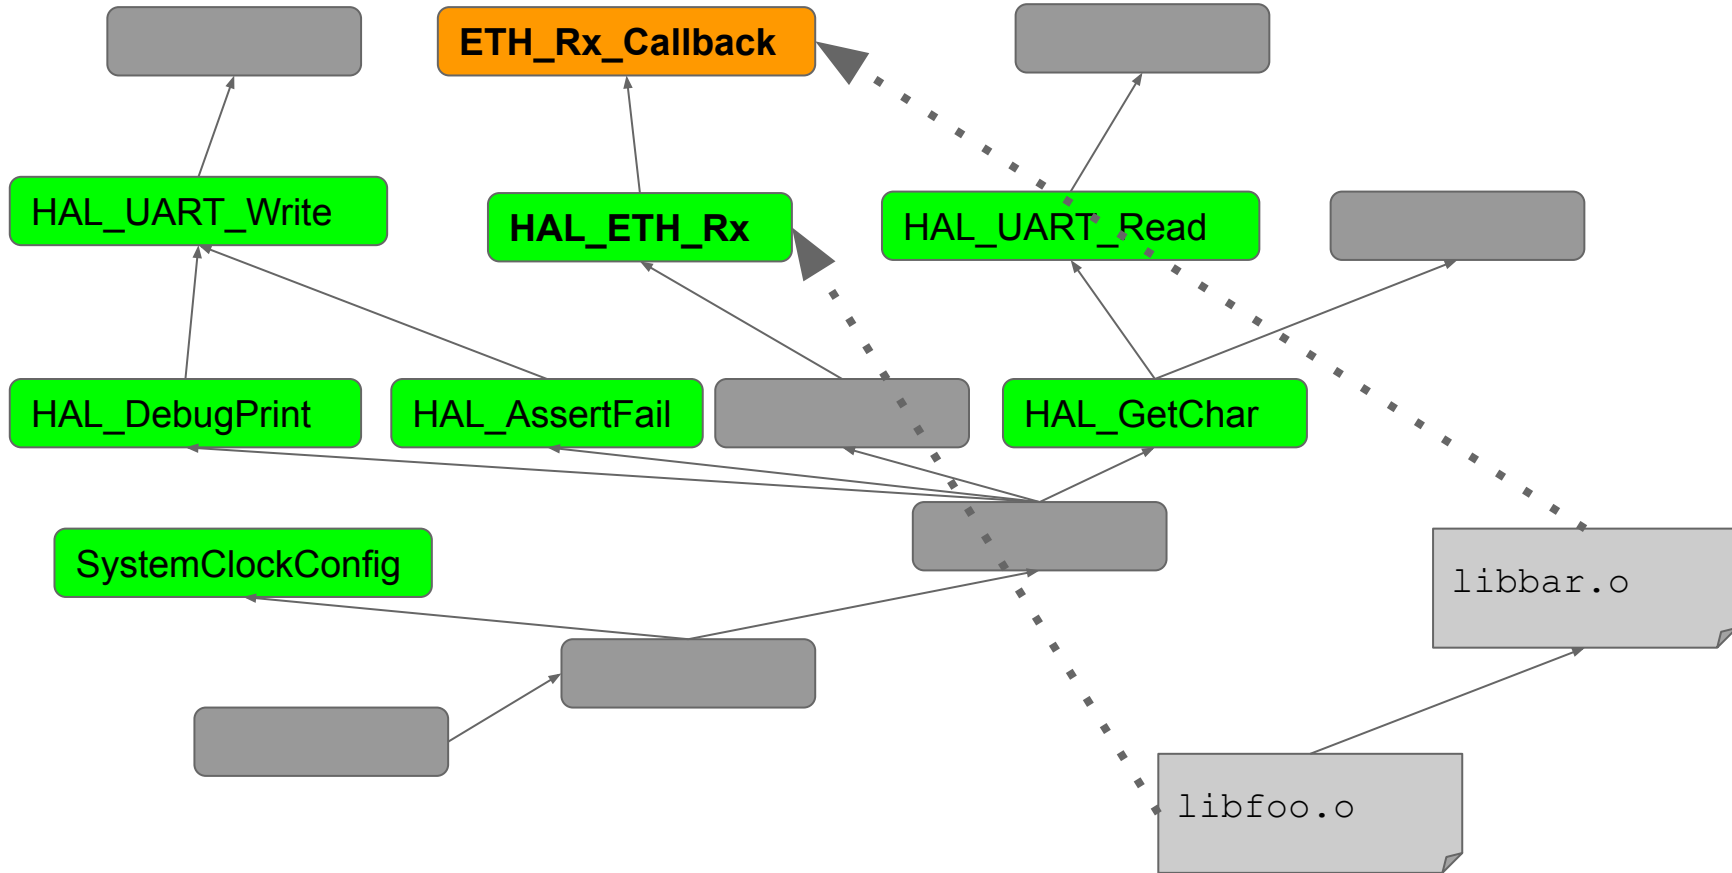


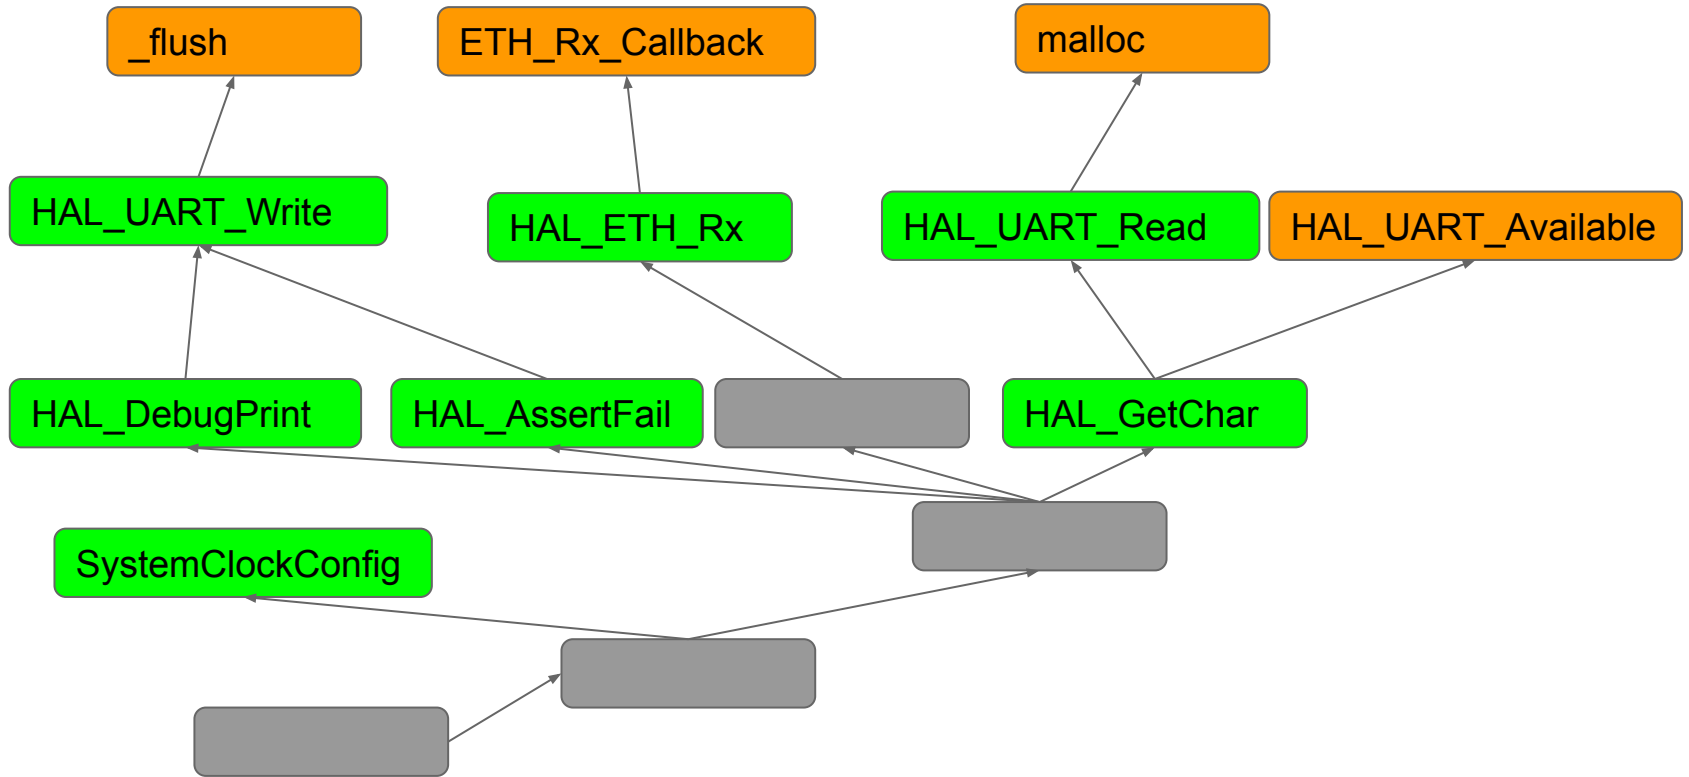












- Built on AFL-Unicorn
- Program exits when the input is exhausted
- Deterministic timers based on block counts
- Interrupt events also based on block counts
- Crashes detected via Unicorn's own error detector as well as handler assertions

16 Firmware Samples

- **ATMEL ASF**
 - USART
 - FAT32 on SD-Card
 - **HTTP Server**
 - **6LoWPAN Sender and Receiver**
- **STM32Cube**
 - UART
 - FAT32 on SD-Card
 - **UDP-Echo Server and Client**
 - **TCP-Echo Server and Client**
 - **PLC**
- **NXP -MCUXpresso**
 - UART
 - **UDP Echo Server**
 - **TCP Echo Server**
 - **HTTP Server**



LibMatch Results



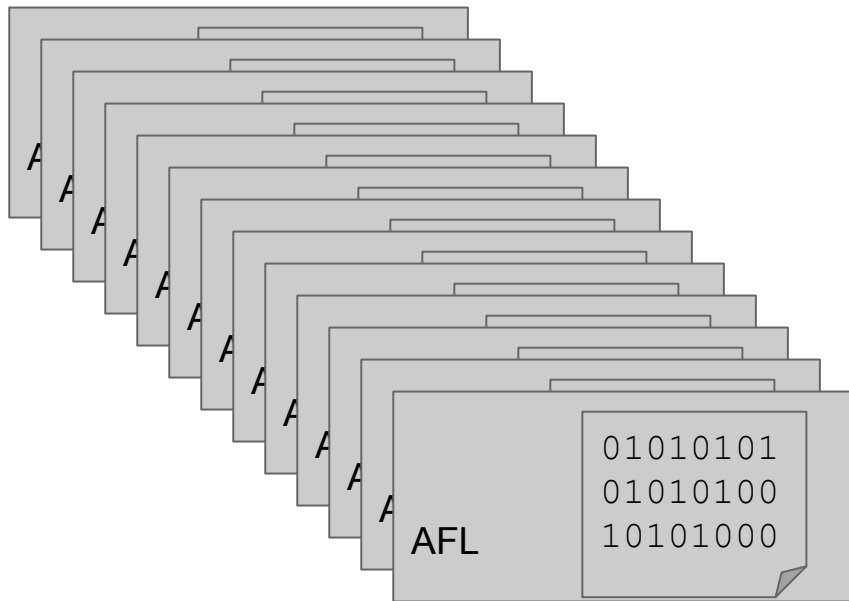
	“Naïve” LibMatch (Bindiff)	LibMatch w/ context
Correct	74.5%	87.4%
Missing	5.0%	3.2%
Collisions	18.8%	8.5%
Incorrect	2.5%	0.9%
External	--	9.96%

% matches across 16 test binaries

Three Handler categories:

- **Trivial:** Does nothing / returns a constant
- **Translating:** Collects arguments, interacts with a Model, returns a result
- **Internal Logic:** Needs to re-implement undocumented internal details

- Over 85% of handlers require little effort
 - 44.5% (37) are “trivial”
 - 42.2% (35) are “translating”
- Remainder (11): “Internal logic”
 - HAL behavior doesn’t abstract hardware well enough
 - HAL behavior makes assumptions not in the docs (e.g., uses its own heap allocator)

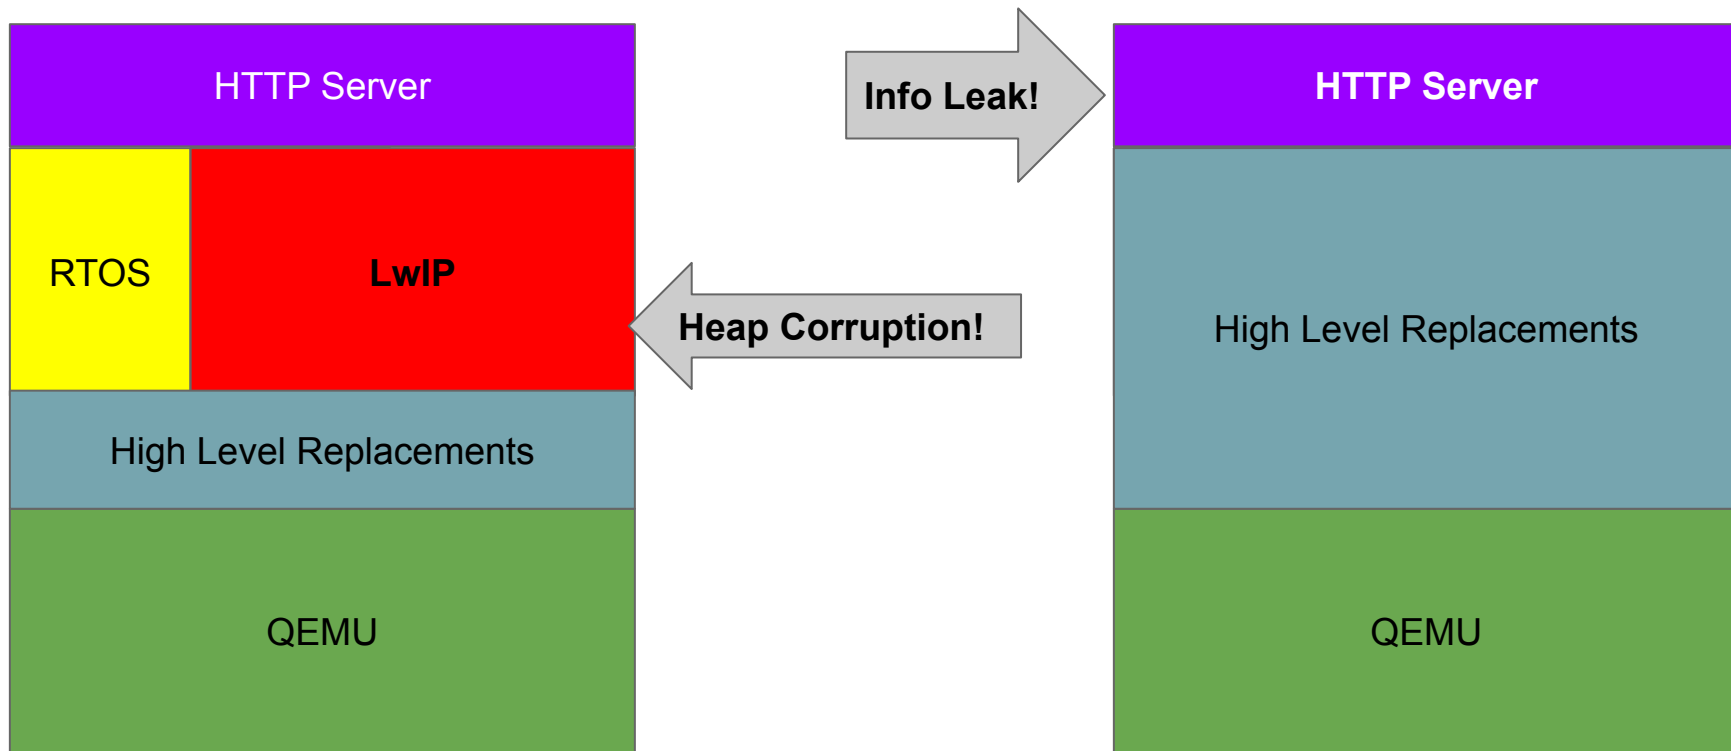


New crashes in:

- STM's ST-PLC Kit
- Atmel's HTTP Server example
- Atmel's Contiki 6LowPAN examples


**Hundreds of millions of executions with
real parallel AFL**

Multi-layer Fuzzing!

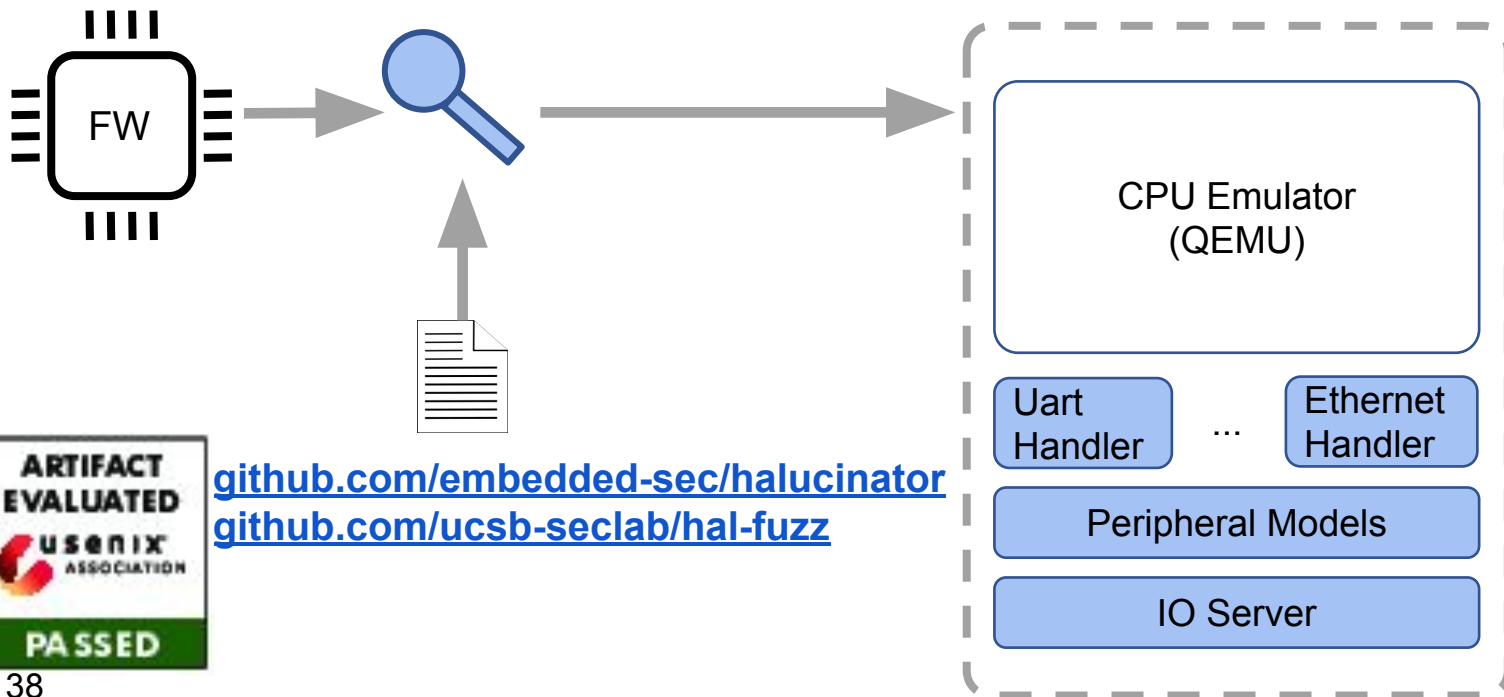


- **CVE-2019-8359**: Remote code execution via buffer overflow in packet reassembly of Contiki OS
- **CVE-2019-9183**: Remote Denial-of-Service via Integer underflow in packet reassembly of Contiki OS

- Re-hosted ARM portion of all challenge sets
- Solved 18/19 challenges
- Verified 17/18 solutions w/ just the emulator
- Solved 3 challenges automatically using fuzzing
- Won first place!



HALucinator eliminates implementing 10,000s of peripherals by using HALs



github.com/embedded-sec/halucinator
github.com/ucsb-seclab/hal-fuzz

