

FFXE


Dynamic Control Flow Graph Recovery for Embedded Firmware Binaries

**Ryan Tsang, Asmita, Doreen Joseph, Soheil Salehi,
Prasant Mohapatra, Houman Homayoun**

Motivation

- Auditing Embedded Systems Firmware
 - Software/Firmware supply chain
- Source code unavailable → Binary analysis
 - Human-in-the-loop process
 - Reliance on automated program analysis techniques
 - Quality depends on **Control Flow Graph** (CFG)
Recovery

```
000002b0 <Reset_Handler>:  
2b0: 4906      ldr r1, [pc, #24]; (2cc)  
2b2: 4a07      ldr r2, [pc, #28]; (2d0)  
2b4: 4b07      ldr r3, [pc, #28]; (2d4)  
2b6: 1a9b      subs r3, r3, r2  
2b8: dd03      ble.n 2c2  
2ba: 3b04      subs r3, #4  
2bc: 58c8      ldr r0, [r1, r3]  
2be: 50d0      str r0, [r2, r3]  
2c0: dcfb      bgt.n 2ba  
2c2: f000 f86b bl 39c <SystemInit>  
2c6: f7ff fbf7 bl 248 <_mainCRTStartup>  
2ca: 0000      .short 0x0000  
2cc: 00000990 .word 0x00000990  
2d0: 20000000 .word 0x20000000  
2d4: 20000010 .word 0x20000010
```

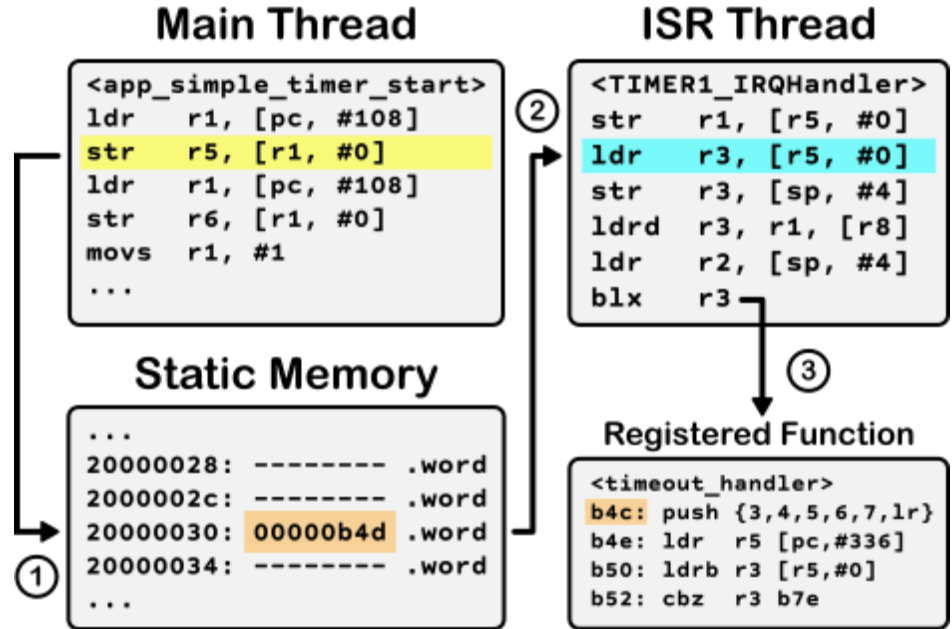


Motivation

- Registered Interrupt Handlers
 - Common pattern in embedded systems firmware
 - Branch targets computed and registered in different thread

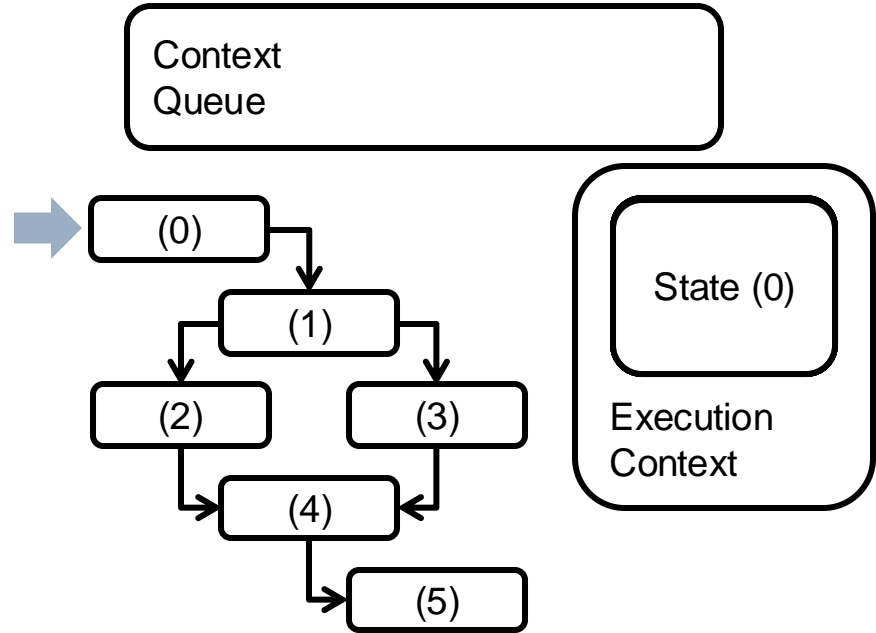
→ Asynchronous memory accesses

- Difficult to resolve indirection
- Runtime information useful



Background: Forced Execution

- Xu et al. 2009
- Concretely execute all conditional branch paths in DFS manner



Problem

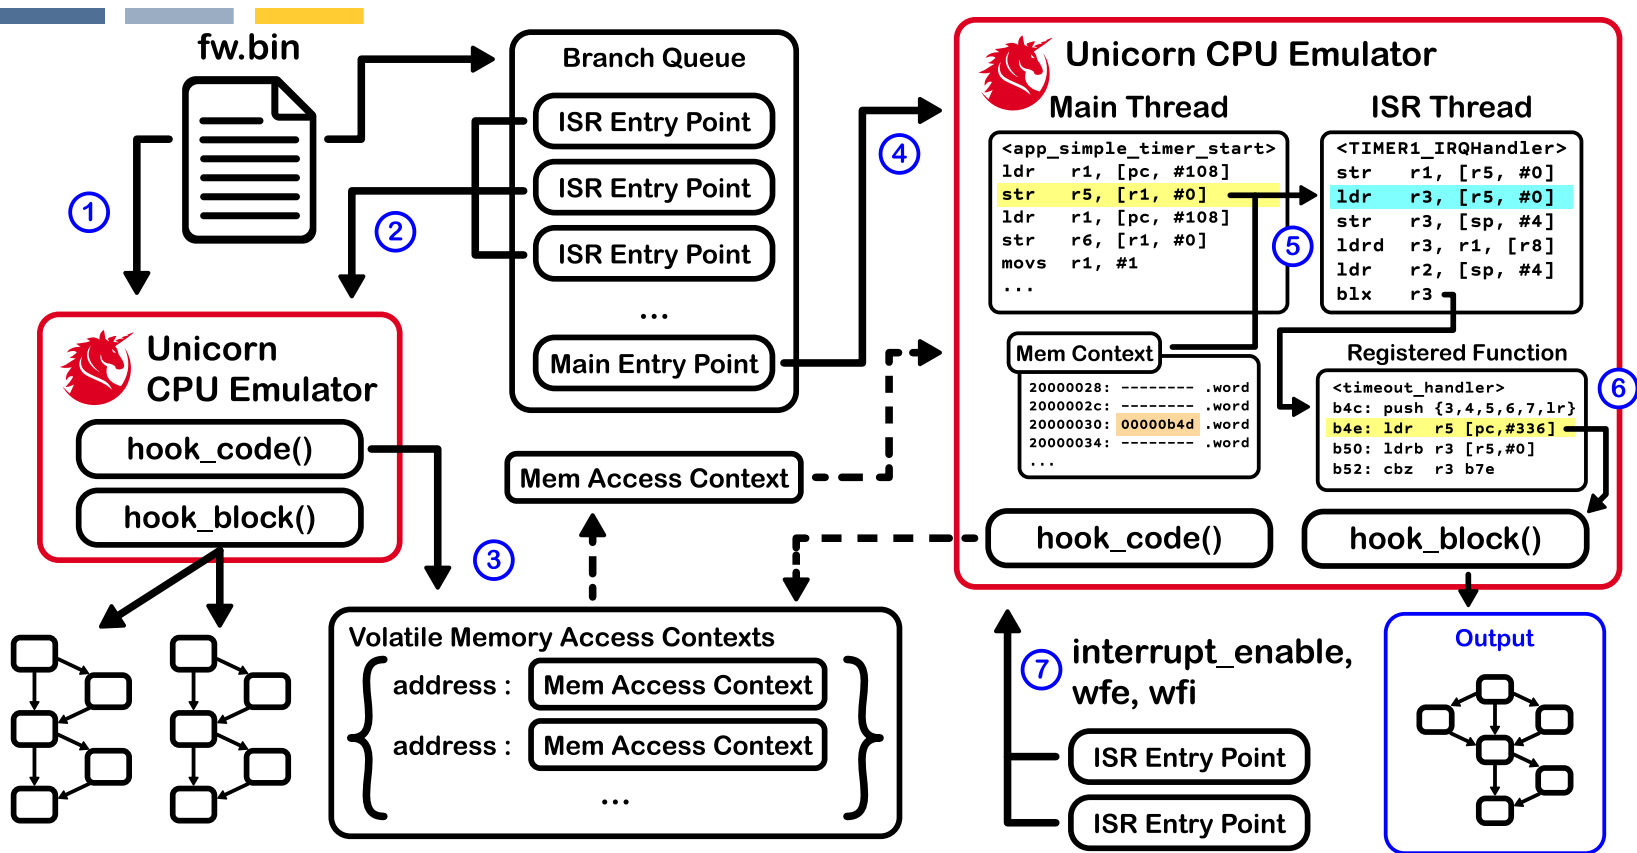


Forced Execution Assumption 2

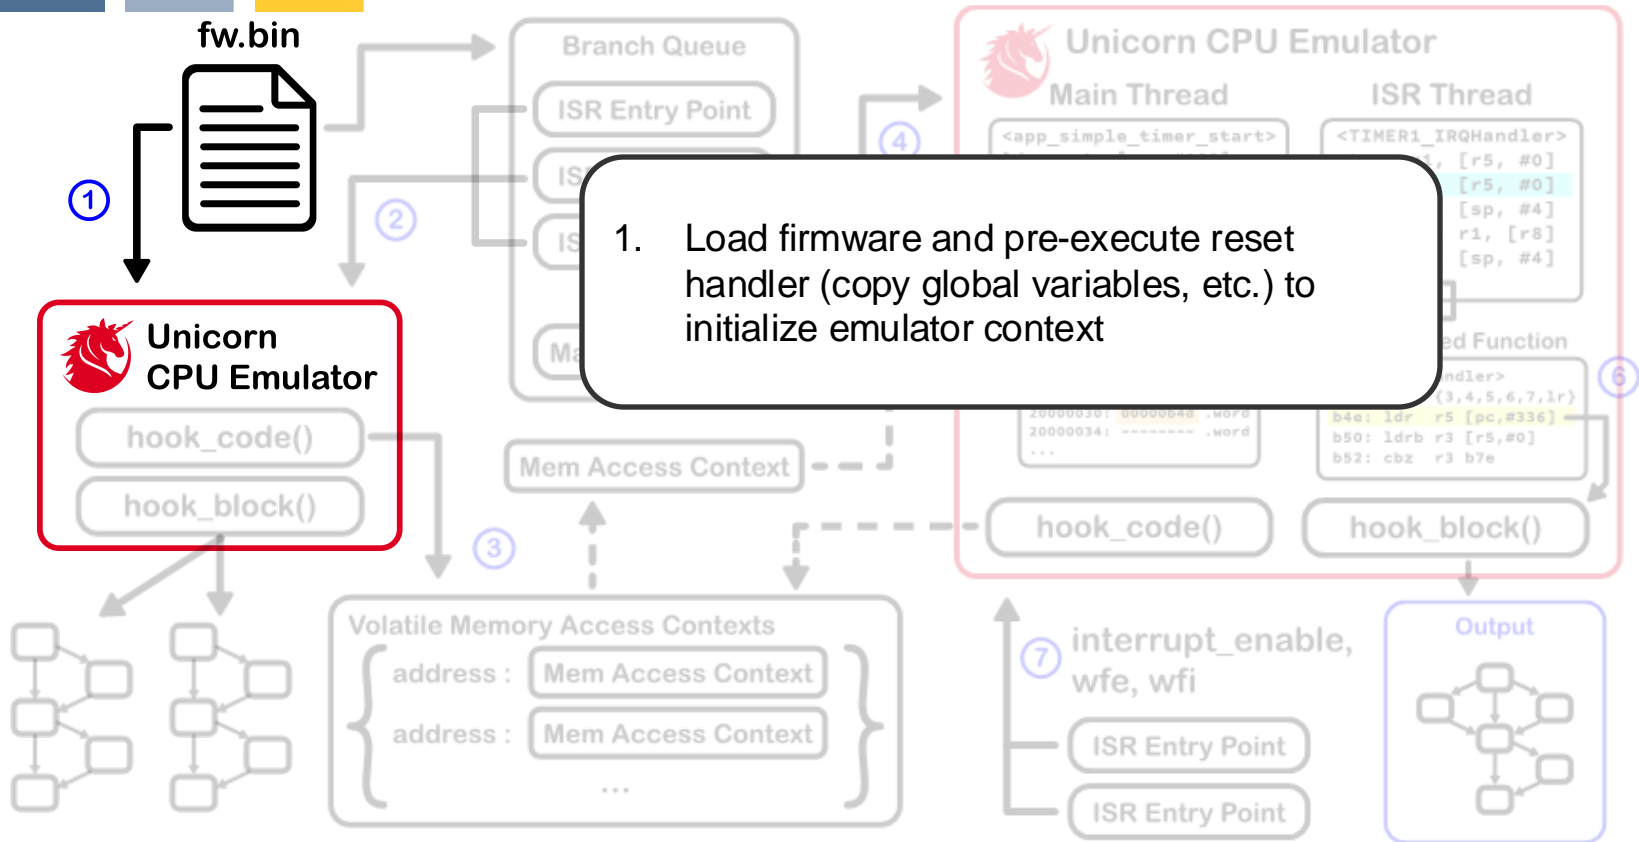
“The target of an indirect branch is completely determined by a control flow path to this indirect branch and is independent of intermediate program states.”

- Interrupt handler registration violates original assumptions
 - Branch target address is determined and written asynchronously
 - Forced execution will result in incorrect jumps
- **Forced Firmware Execution Engine (FFXE)**
 - Account for asynchronous memory accesses to resolve indirect branches
 - Resumes forced execution from volatile memory reads

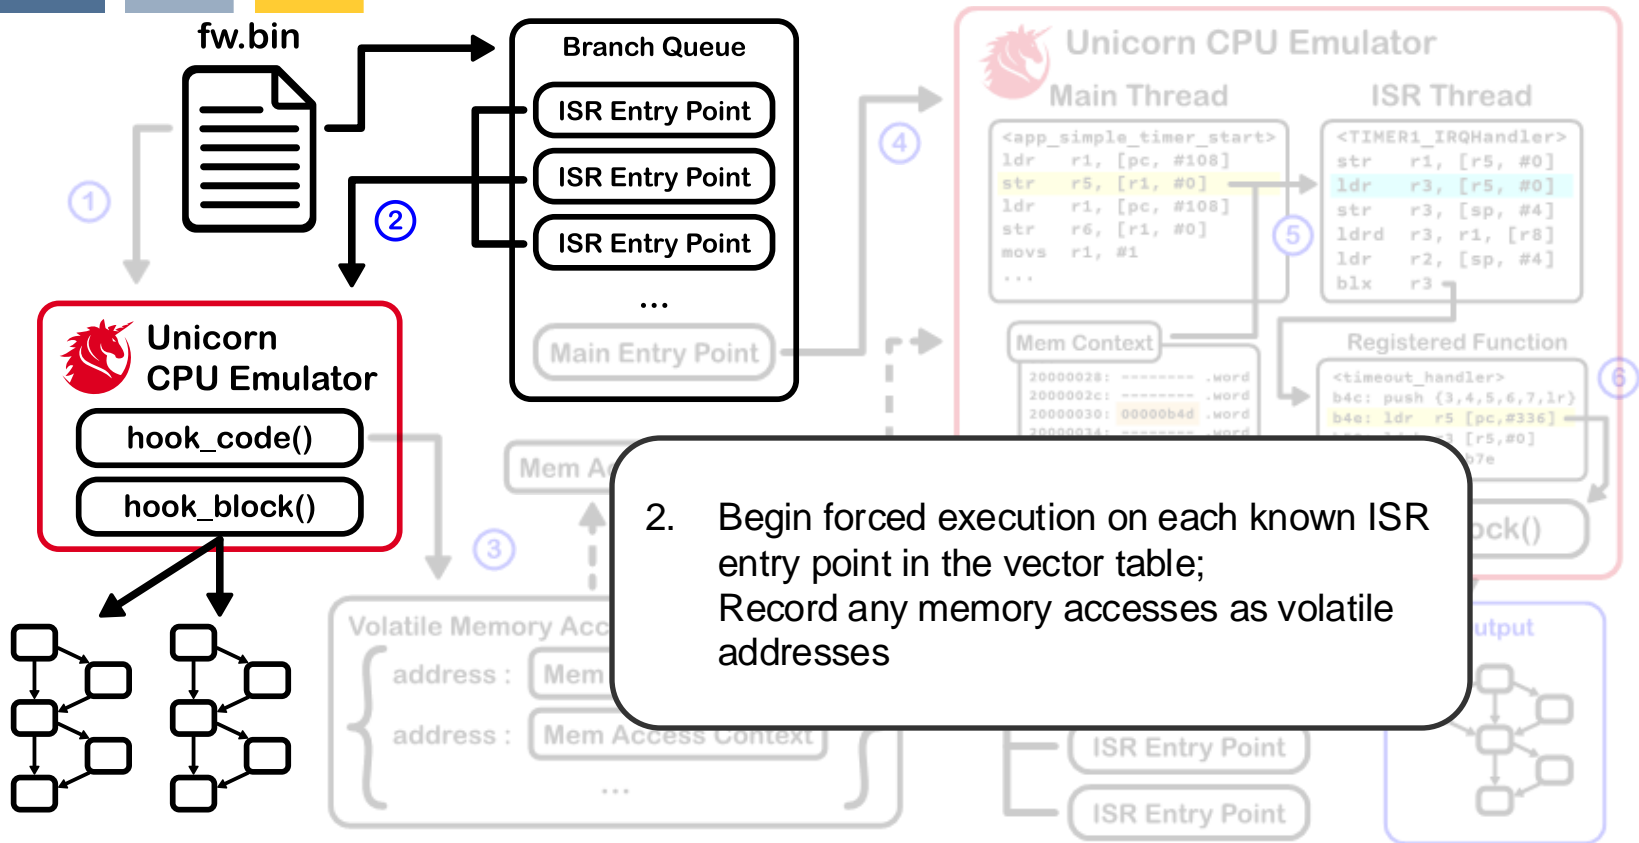
FFXE



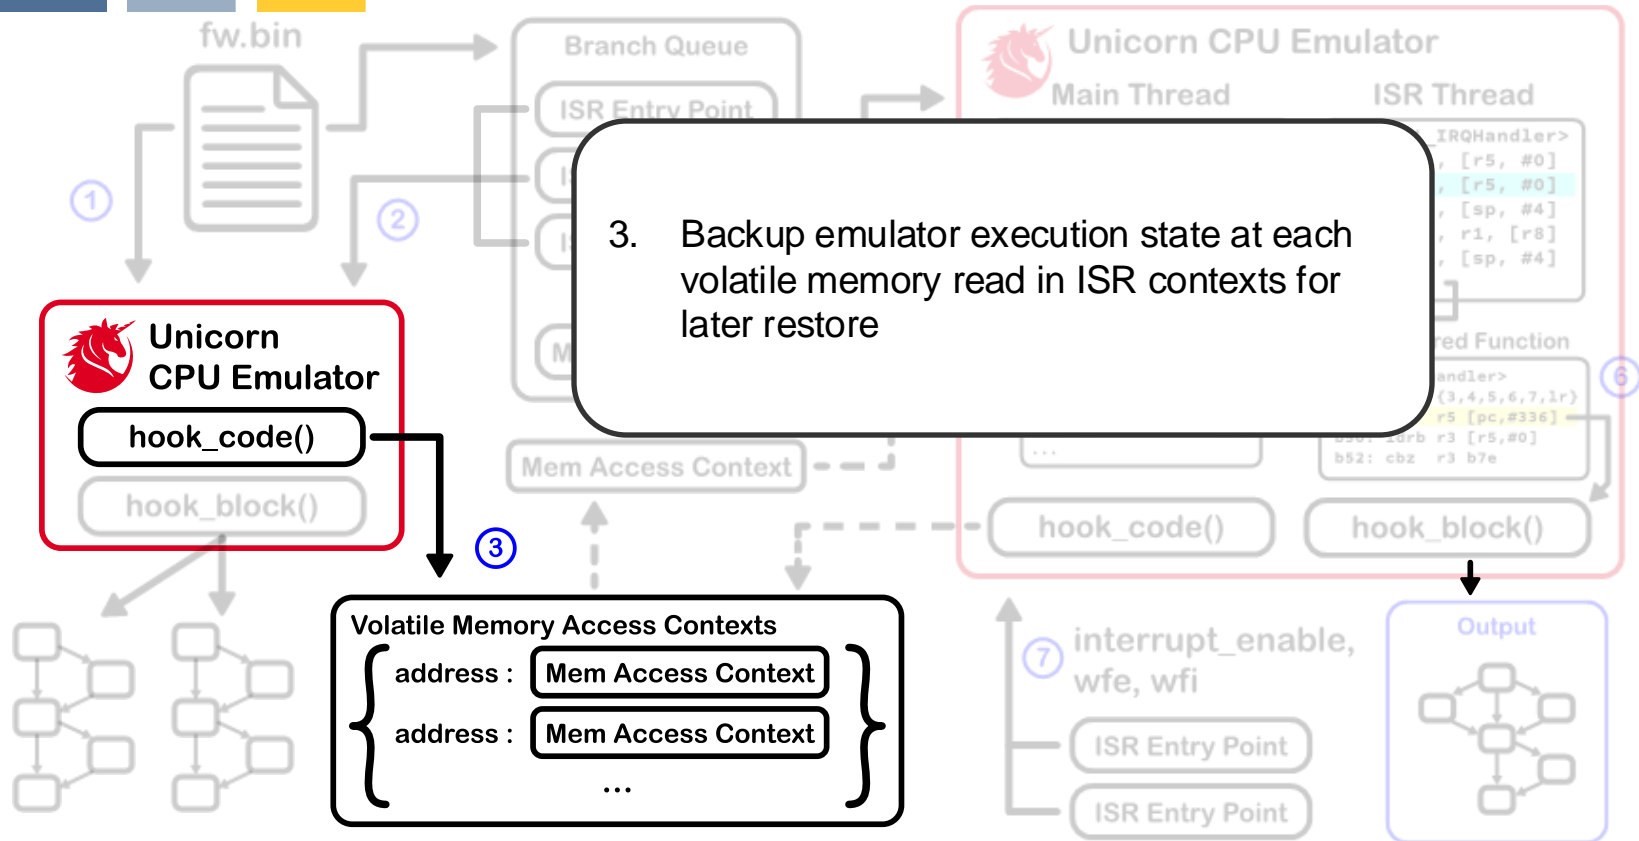
Method



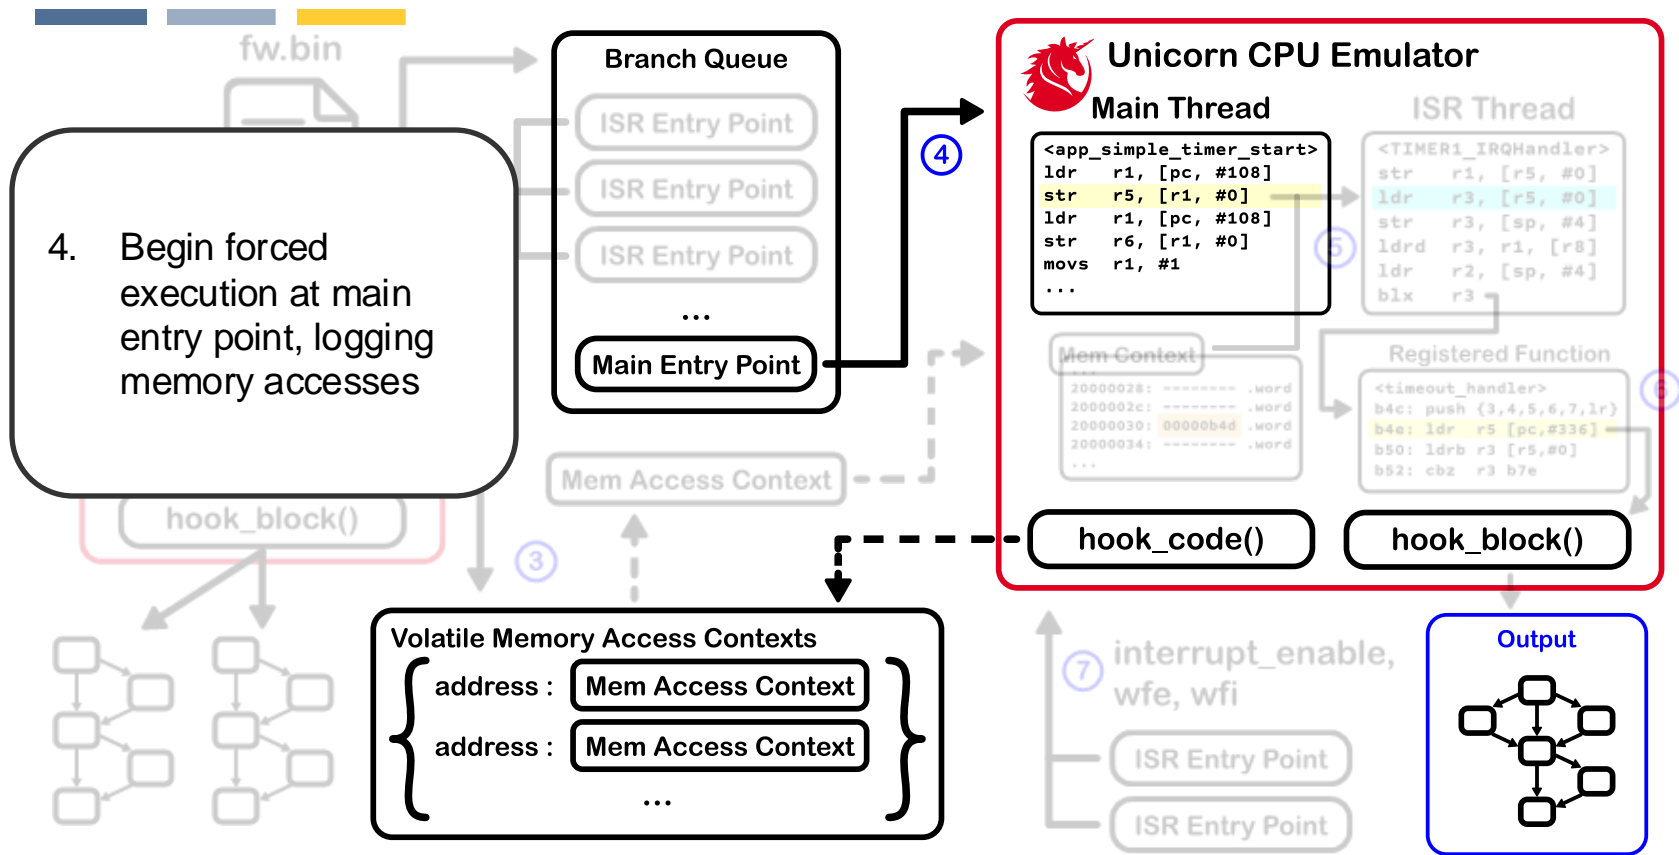
Method



Method

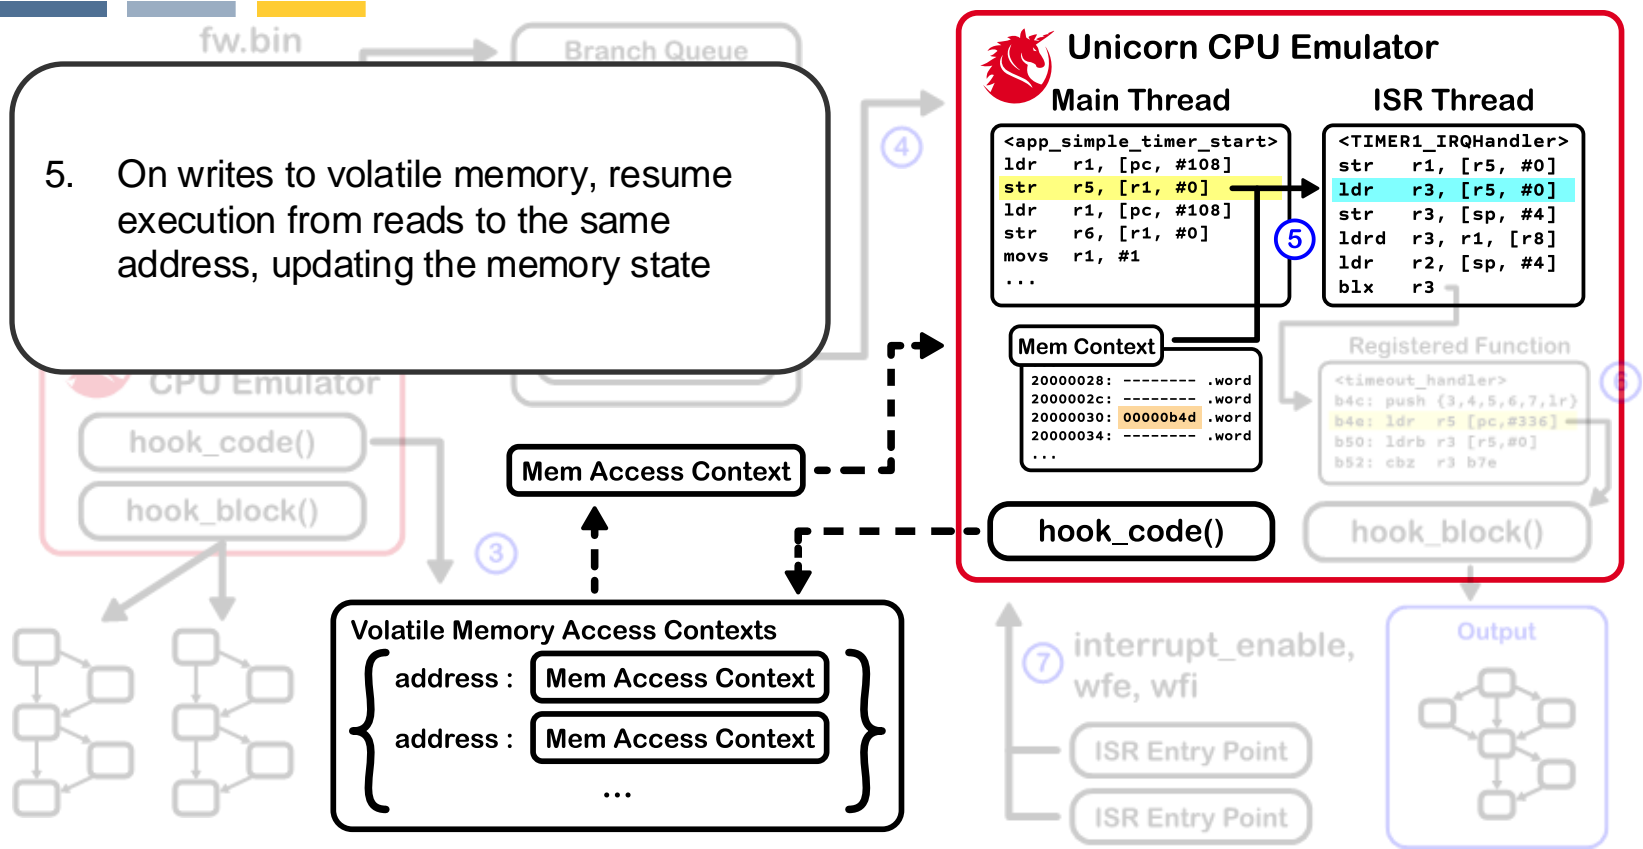


Method

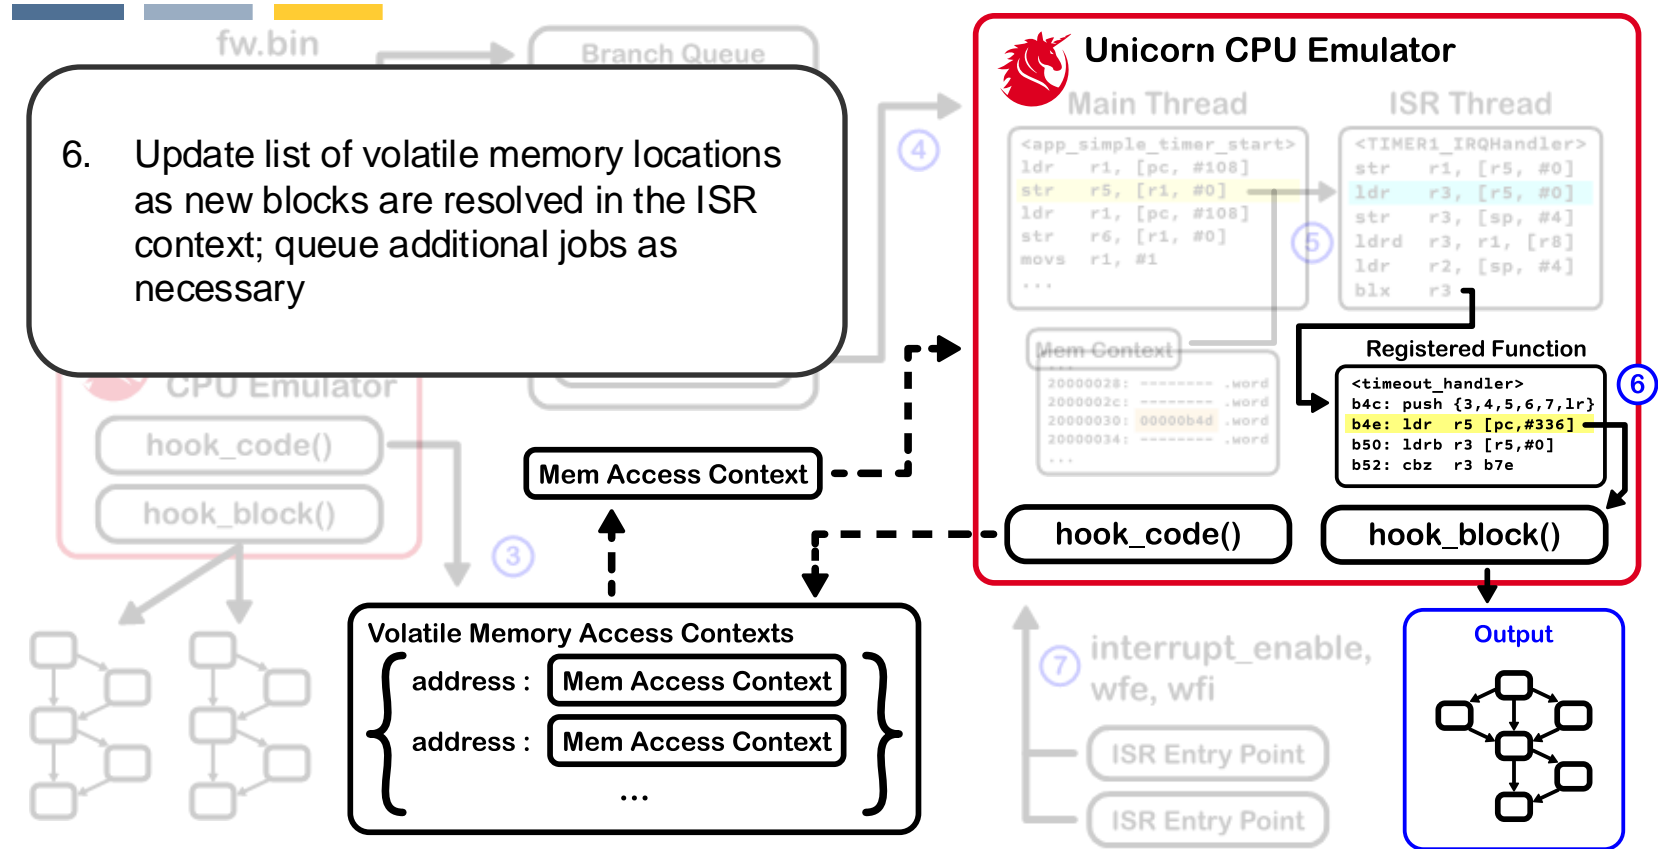


Method

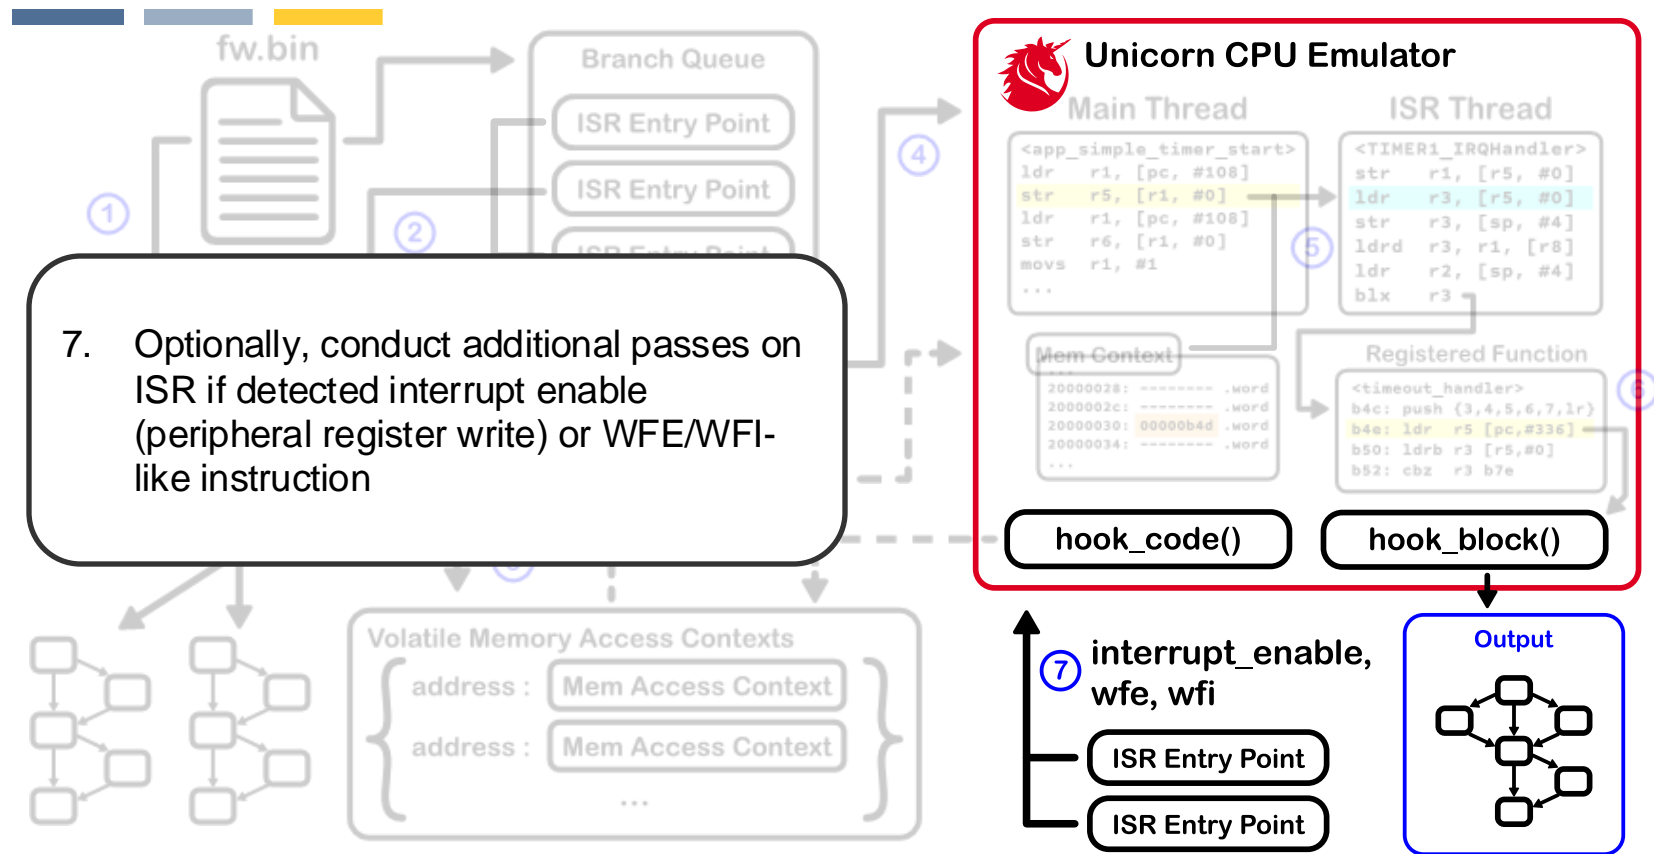
5. On writes to volatile memory, resume execution from reads to the same address, updating the memory state



Method



Method



Evaluation: Callback Resolution

- FFXE fully recovers callback edges to all handlers in our test set
 - Other tools typically fail to resolve these edges
- Nordic nRF5 SDK
 - 8 samples x 4 optimization levels = 32 binaries
 - 24 manually identified callback functions

Table 1: Registered Function Resolution Comparison of CFG Recovery Methods.

Firmware	angr_emu	angr_fast	ffxe	fxe	ghidra	
gpote	-00	0/0/1	0/1/1	1/1/1	0/0/1	0/1/1
	-01	0/0/1	0/1/1	1/1/1	0/0/1	0/0/1
	-02	0/0/1	0/1/1	1/1/1	0/0/1	0/0/1
	-03	0/0/1	0/1/1	1/1/1	0/0/1	0/0/1
i2s	-00	0/0/4	0/4/4	9/4/4	0/0/4	0/4/4
	-01	0/0/4	0/4/4	10/4/4	0/0/4	0/3/4
	-02	0/0/4	0/4/4	12/4/4	0/0/4	0/3/4
	-03	0/0/4	0/4/4	12/4/4	0/0/4	0/3/4
saadc	-00	0/0/2	0/2/2	4/2/2	0/0/2	0/2/2
	-01	0/0/2	0/2/2	4/2/2	0/0/2	0/2/2
	-02	0/0/2	1/2/2	4/2/2	0/0/2	0/2/2
	-03	0/0/2	1/2/2	4/2/2	0/0/2	0/2/2
simple_timer	-00	0/0/2	0/2/2	2/2/2	0/0/2	0/2/2
	-01	0/0/2	0/2/2	2/2/2	0/0/2	0/1/2
	-02	0/0/2	1/2/2	2/2/2	0/0/2	0/1/2
	-03	0/0/2	1/2/2	2/2/2	0/0/2	0/1/2
spi	-00	0/0/5	0/5/5	7/5/5	0/0/5	0/5/5
	-01	0/0/5	0/5/5	8/5/5	0/0/5	0/4/5
	-02	0/0/5	0/5/5	10/5/5	0/0/5	0/4/5
	-03	0/0/5	0/5/5	10/5/5	0/0/5	0/4/5
timer	-00	0/0/1	0/1/1	1/1/1	0/0/1	0/1/1
	-01	0/0/1	0/1/1	1/1/1	0/0/1	0/1/1
	-02	0/0/1	0/1/1	1/1/1	0/0/1	0/1/1
	-03	0/0/1	0/1/1	1/1/1	0/0/1	0/1/1
twi_sensor	-00	0/0/5	0/5/5	7/5/5	0/0/5	0/5/5
	-01	0/0/5	0/5/5	8/5/5	0/0/5	0/4/5
	-02	0/0/5	0/5/5	10/5/5	0/0/5	0/4/5
	-03	0/0/5	0/5/5	10/5/5	0/0/5	0/4/5
uart	-00	1/1/4	0/4/4	9/4/4	0/0/4	0/4/4
	-01	0/0/4	0/4/4	10/4/4	0/0/4	0/4/4
	-02	0/0/4	2/4/4	12/4/4	0/0/4	0/2/4
	-03	0/0/4	2/4/4	12/4/4	0/0/4	0/2/4

Table elements: (# of found edges to registered function)/(# of found registered function entry blocks)/(# of known registered functions).

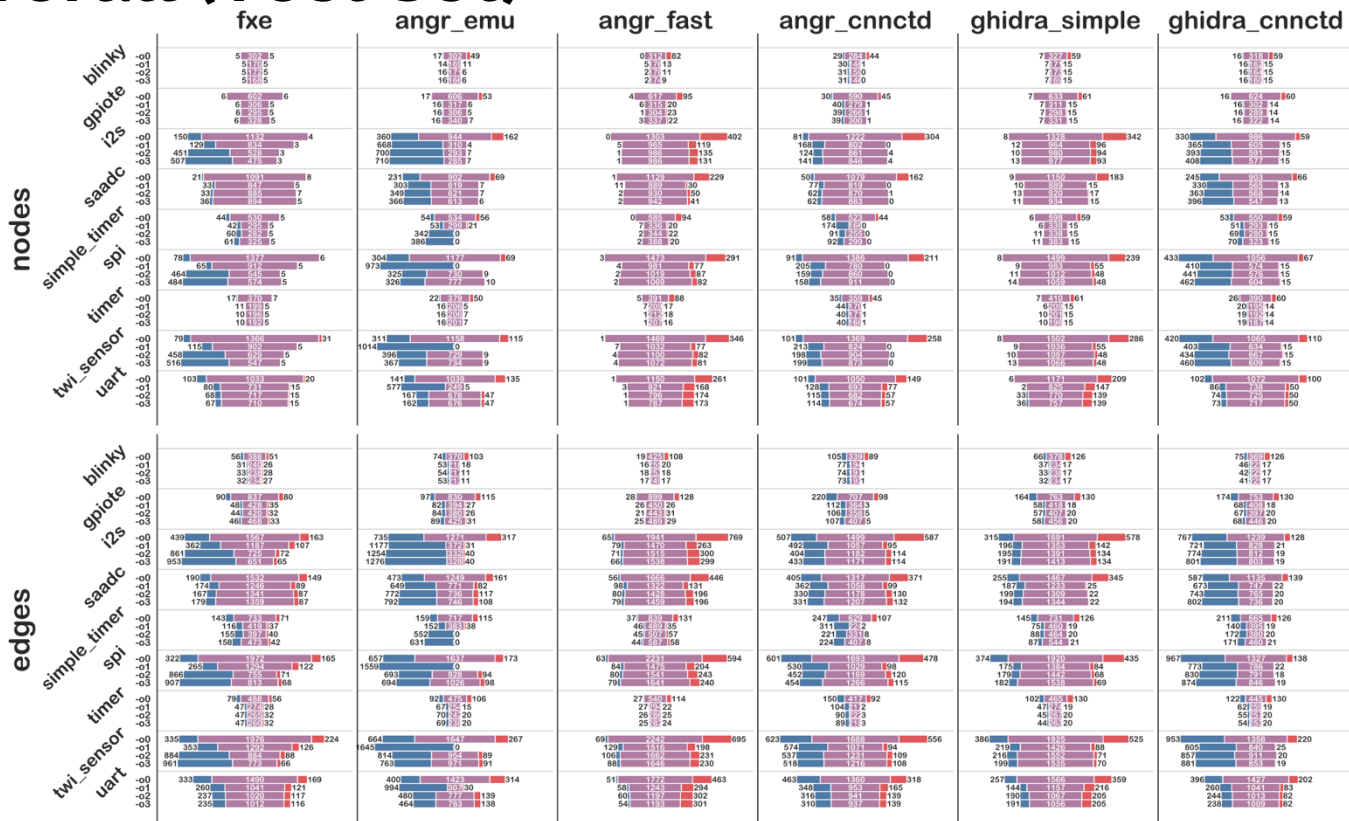
Evaluation: Overall (Test Set)



FFXE-only (blue)

Other engine (red)

Overlap (purple)



Evaluation: Overall (Test Set)

Complementary coverage for reachable graphs



Evaluation: Overall (Real-World)

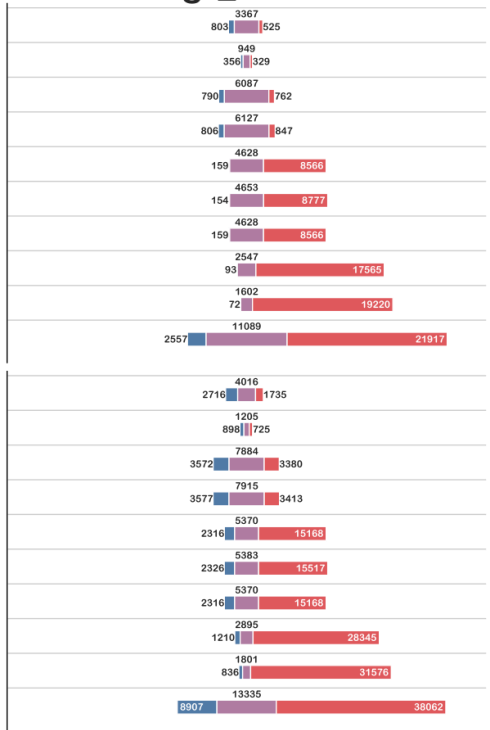


block coverage


edges

switchmate_light_2_99_16
 switchmate_light_2_21
 switchmate_bright_2_9_11
 switchmate_bright_1_46
 flex_7_88_flash
 flex_7_81_flash
 flex_7_64_flash
 chargehr_18_128
 chargehr_18_128
 BCM20702A1


angr_cnnctd



angr_emu



ghidra_cnnctd



Conclusion



- FFXE designed to resolve edges to interrupt callback functions
 - Common occurrence in bare-metal firmware SDKs
- Evaluated FFXE's ability to resolve these edges against existing tools
 - All other tools failed to locate such edges in our test set
- Demonstrated FFXE provides complementary coverage to existing tools
 - Ghidra, angr's static and dynamic methods, original FXE

GitHub: <https://github.com/rchtsang/ffxe>