

## **Securing the Airwaves**

Emulation, Fuzzing, and Reverse Engineering of iPhone Baseband Firmware

Bruno, Luca, Rachna {bruno,glockow,rachna}@srlabs.de



Nice to meet you :)



**Bruno  
Produit**

- Background in low-level security and cryptography
- Experience in fuzzing, telco and hardware hacking
- GitHub @brunoproduit



**Luca  
Glockow**

- Background in application and device security
- Experience in hacking hard- and software in telco
- GitHub @luglo



**Rachna  
Shriwas**

- Background in device testing, fuzzing and code assurance
- GitHub @rachsr1

## We start where other research projects stopped

“*Unlike other baseband implementations, Qualcomm leverages a fully-custom architecture known as Hexagon [...]*  
*Unfortunately, tooling for this architecture is sparse, and especially full-system emulators are lacking.*”

Hernandez et al. 2022, “FirmWire: Transparent Dynamic Analysis for Cellular Baseband Firmware”

# We investigate the other chip everyone has in their pocket

## Our goals

---



### Create Transparency

Understand, document, and share  
Hexagon security insights



### Enable vulnerability research

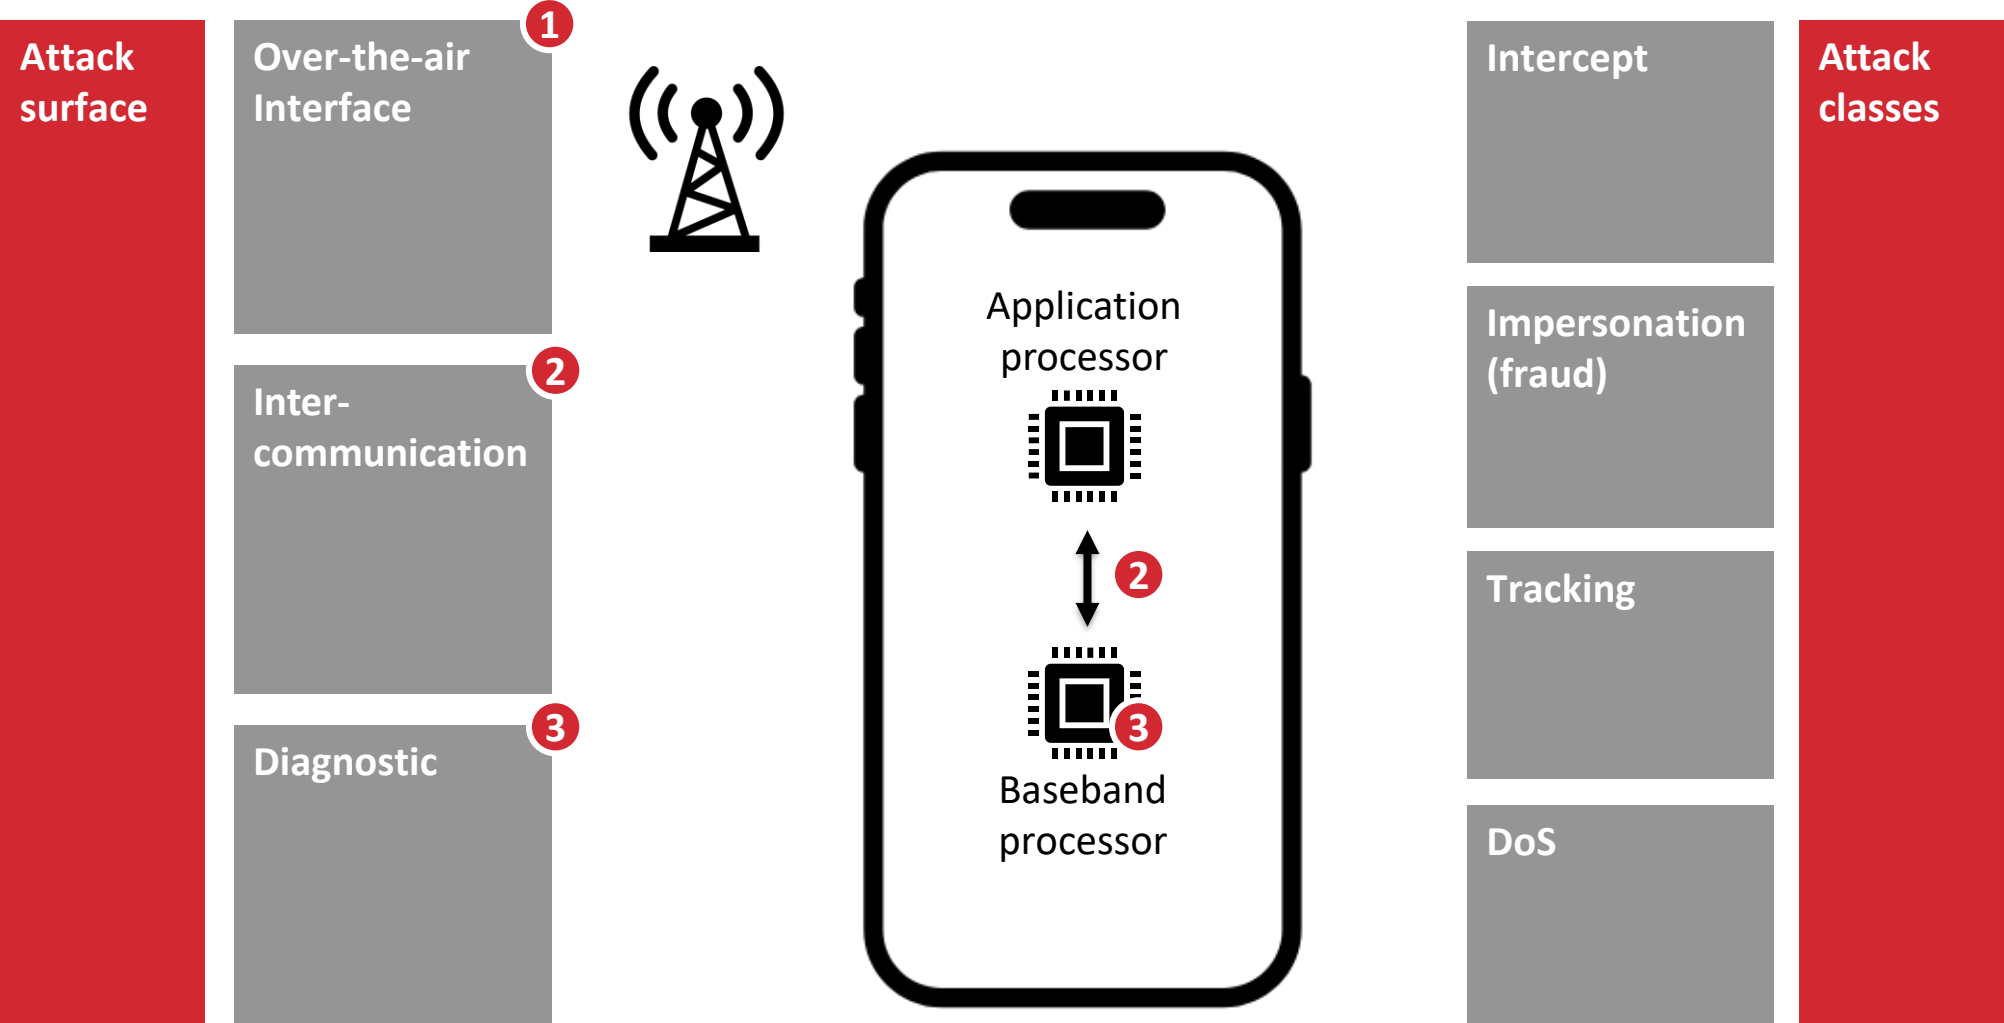
Help secure Qualcomm-based phones  
against exploitation














### Release Tooling

Open source fuzzing setup and tooling  
for Hexagon basebands

# The Qualcomm baseband exposes multiple interfaces, making it vulnerable to diverse attacks



# Existing baseband research does not cover Hexagon baseband full system emulation

	 Emulation	 Fuzzing	 Hexagon	
Reverse Engineer	<ul style="list-style-type: none"><li>▪ <b>Reversing Hexagon:</b> Burke 2018 [1]</li><li>▪ <b>Reversing DIAG:</b> Esage 2020 [2]</li></ul>			
Hardware Fuzz	<ul style="list-style-type: none"><li>▪ <b>Fuzzing on Hexagon hardware:</b> Gong &amp; Zhang 2021 [3]</li></ul>			 
Emulation Fuzz	<ul style="list-style-type: none"><li>▪ <b>Advanced rehosted baseband fuzzing:</b> Maier et al 2020 [4], Hernandez et al 2022 [5]</li></ul>			 
Emulation Hexagon Fuzz	The gap we want to fill			  

[1] Burke 2018, [A Journey into Hexagon](#)  
[2] Esage 2020, [Advanced Hexagon DIAG](#)  
[3] Gong & Zhang 2021, [In-Depth Analyzing and Fuzzing for Qualcomm Hexagon Processor](#)  
[4] Maier et al. 2020, [BaseSAFE: baseband sanitized fuzzing through emulation](#)  
[5] Hernandez et al. 2022, [FirmWire: Transparent Dynamic Analysis for Cellular Baseband Firmware](#)

---

## Introduction: Hexagon baseband

From research to tooling

Demo: Fuzzing Hexagon

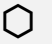


Opening up baseband security



---

# Hexagon is a whole new DSP-focussed CPU architecture, prompting research from the ground again

Challenges for researchers	
Custom CPU architecture	<ul style="list-style-type: none"><li>▪ <b>Lacks support</b> in most tooling</li><li>▪ <b>Tedious</b> to reverse</li></ul>
Custom OS and software stack	<ul style="list-style-type: none"><li>▪ <b>Proprietary</b> tooling</li><li>▪ <b>Adaptation</b> of common attack techniques required</li></ul>
Custom communication protocols	<ul style="list-style-type: none"><li>▪ <b>No free</b> tooling or documentation exists</li></ul>

## Design Highlights

-  **Optimized for parallel execution** (VLIW)
-  **DSP-specialized** (vs general CPU)
-  **Quirky registers** (chicken, duck, goose...)

```
6      #ifndef _HEXAGONTYPES_H_
473    typedef enum global_register_t
502    {
503        G_REG_ACC1,
504        G_REG_CHICKEN,
        G_REG_STFINST,
```
-  **Privilege mode separation:** monitor, guest, user
-  **Instruction packets**, group of parallel instructions

```
A2_addi
{
J2_call
A2_tfrsi
}
J2_cmpeq
```



# Agenda

---

Introduction: Hexagon baseband

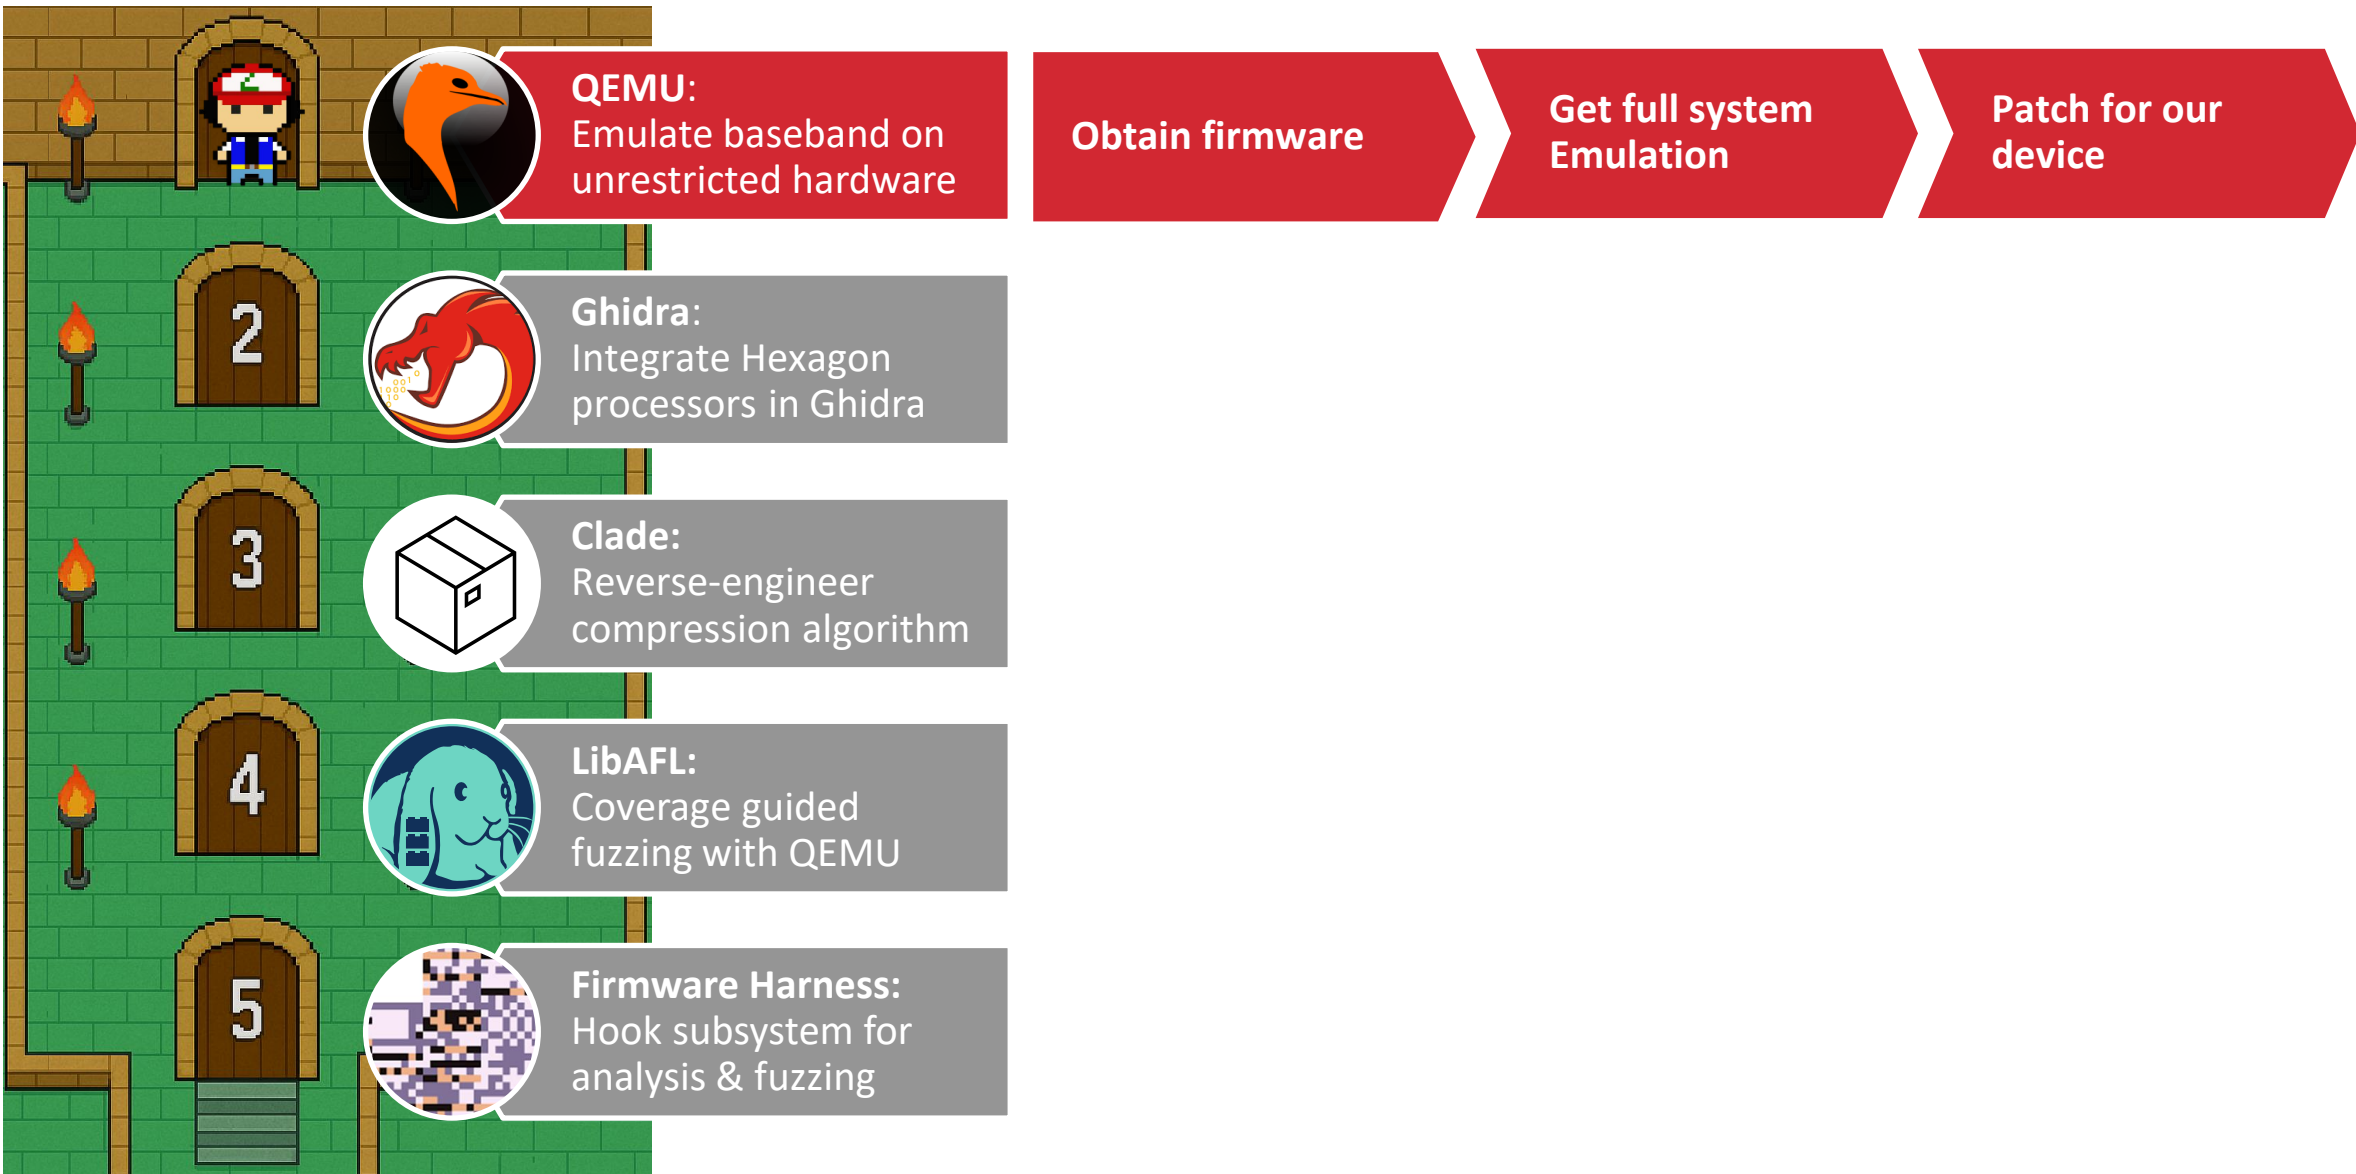
 **From research to tooling**

Demo: Fuzzing Hexagon

Opening up baseband security

---

# We need the firmware and a customized emulator to run Hexagon on a different machine



# The firmware is the initial step in our emulation effort



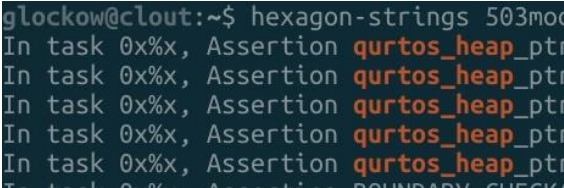
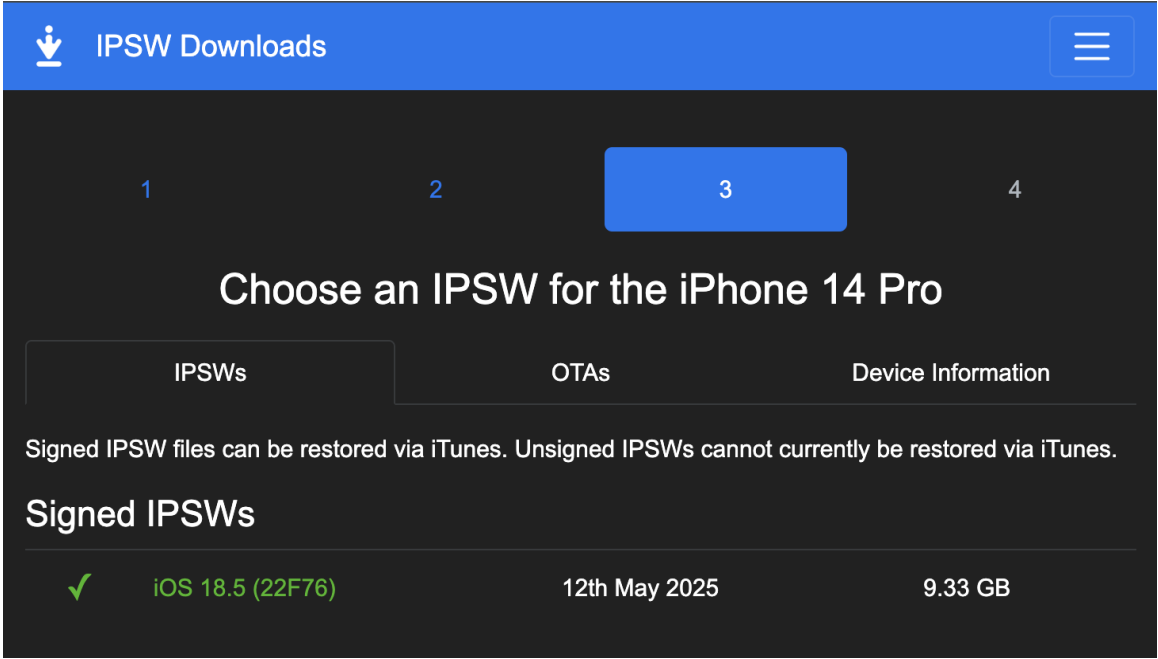
Obtain firmware

Get full system Emulation

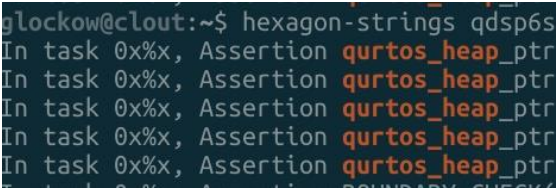
Patch for our device

Step	Details
Obtain firmware	<ul style="list-style-type: none"><li>▪ Download publicly accessible iOS IPSW firmware files (e.g. from ipsw.me)</li><li>▪ Extract baseband bundle from the IPSW archive and <b>identify QUALCOMM DSP6 executable (qdsp6sw.mbn)</b></li></ul>
Analyze firmware	<ul style="list-style-type: none"><li>▪ <b>Identical strings</b> confirm, firmware shares code</li><li>▪ <b>Build on the same Qualcomm Hexagon SDK</b></li></ul>

## Screenshots



Router firmware



iPhone firmware

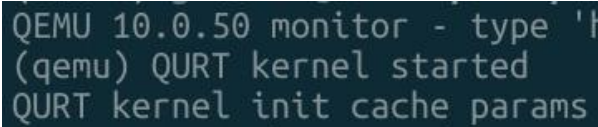
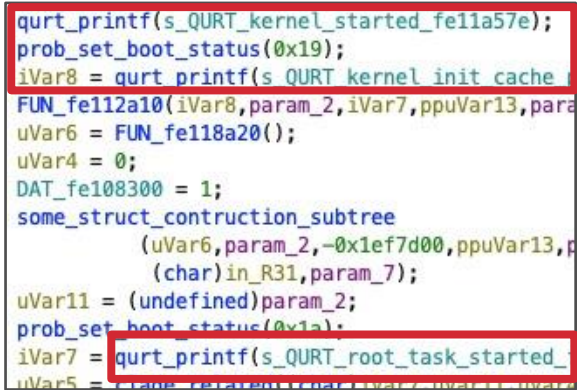
# Finding (Q)emu: Qualcomm works on Hexagon full system emulation



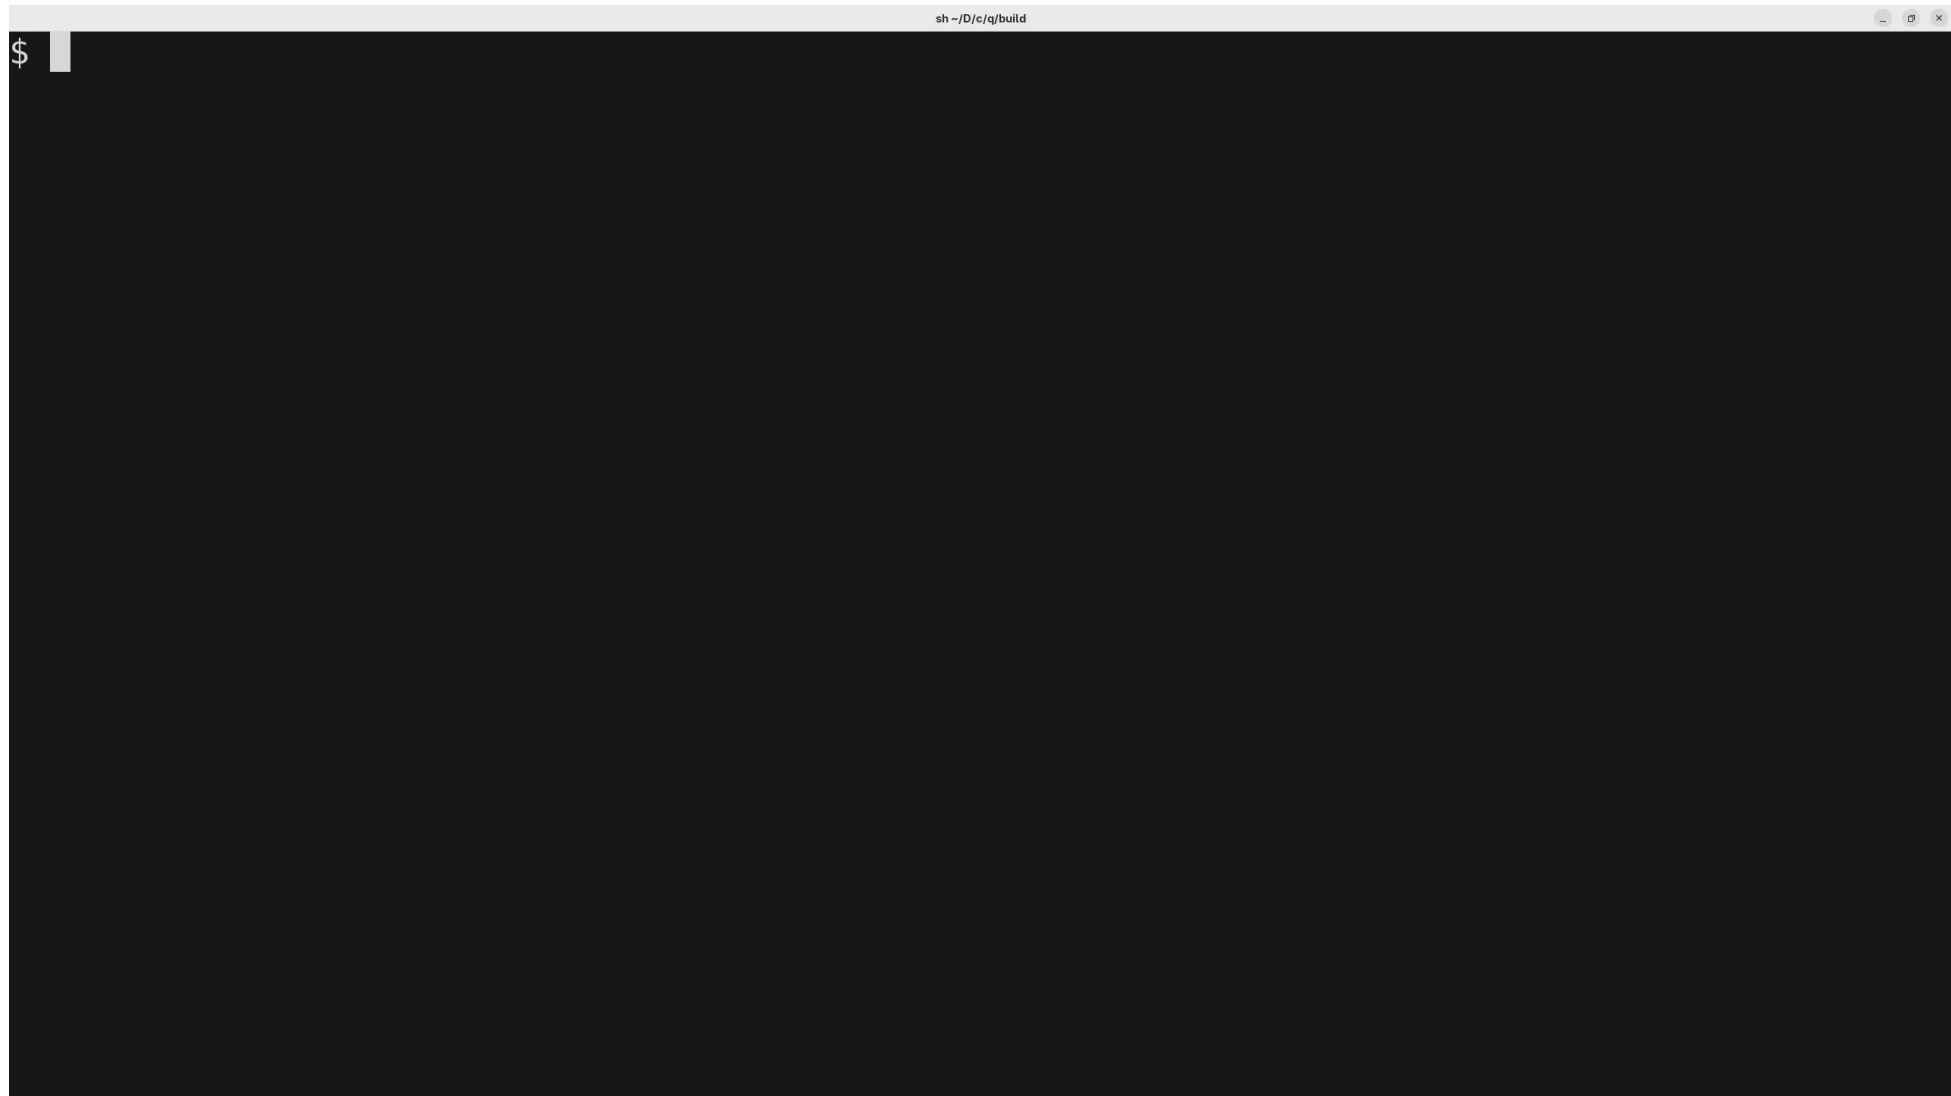
Obtain firmware

Get full system Emulation

Patch for our device

Step	Details	Screenshots
Find full system emulator	<ul style="list-style-type: none"><li>hexagon-softmmu not supported by QEMU</li><li>QEMU fork maintained by Qualcomm Innovation Center</li><li>hexagon-softmmu implementation on WIP branch</li></ul>	
Run emulator	<ul style="list-style-type: none"><li>Compile and run qemu-system-hexagon</li></ul> <pre>\$ qemu-system-hexagon -monitor stdio -display none -kernel qdsp6sw.mbn</pre>	<p>Boot logs in QEMU</p> 
Progress boot	<ul style="list-style-type: none"><li>More logs appear in reversed firmware code</li></ul>	<p>Boot logs in decompiled code</p>

Without QEMU modification, firmware boot fails with only two lines of output



# We control the machine but not its software - reversed bug hunting

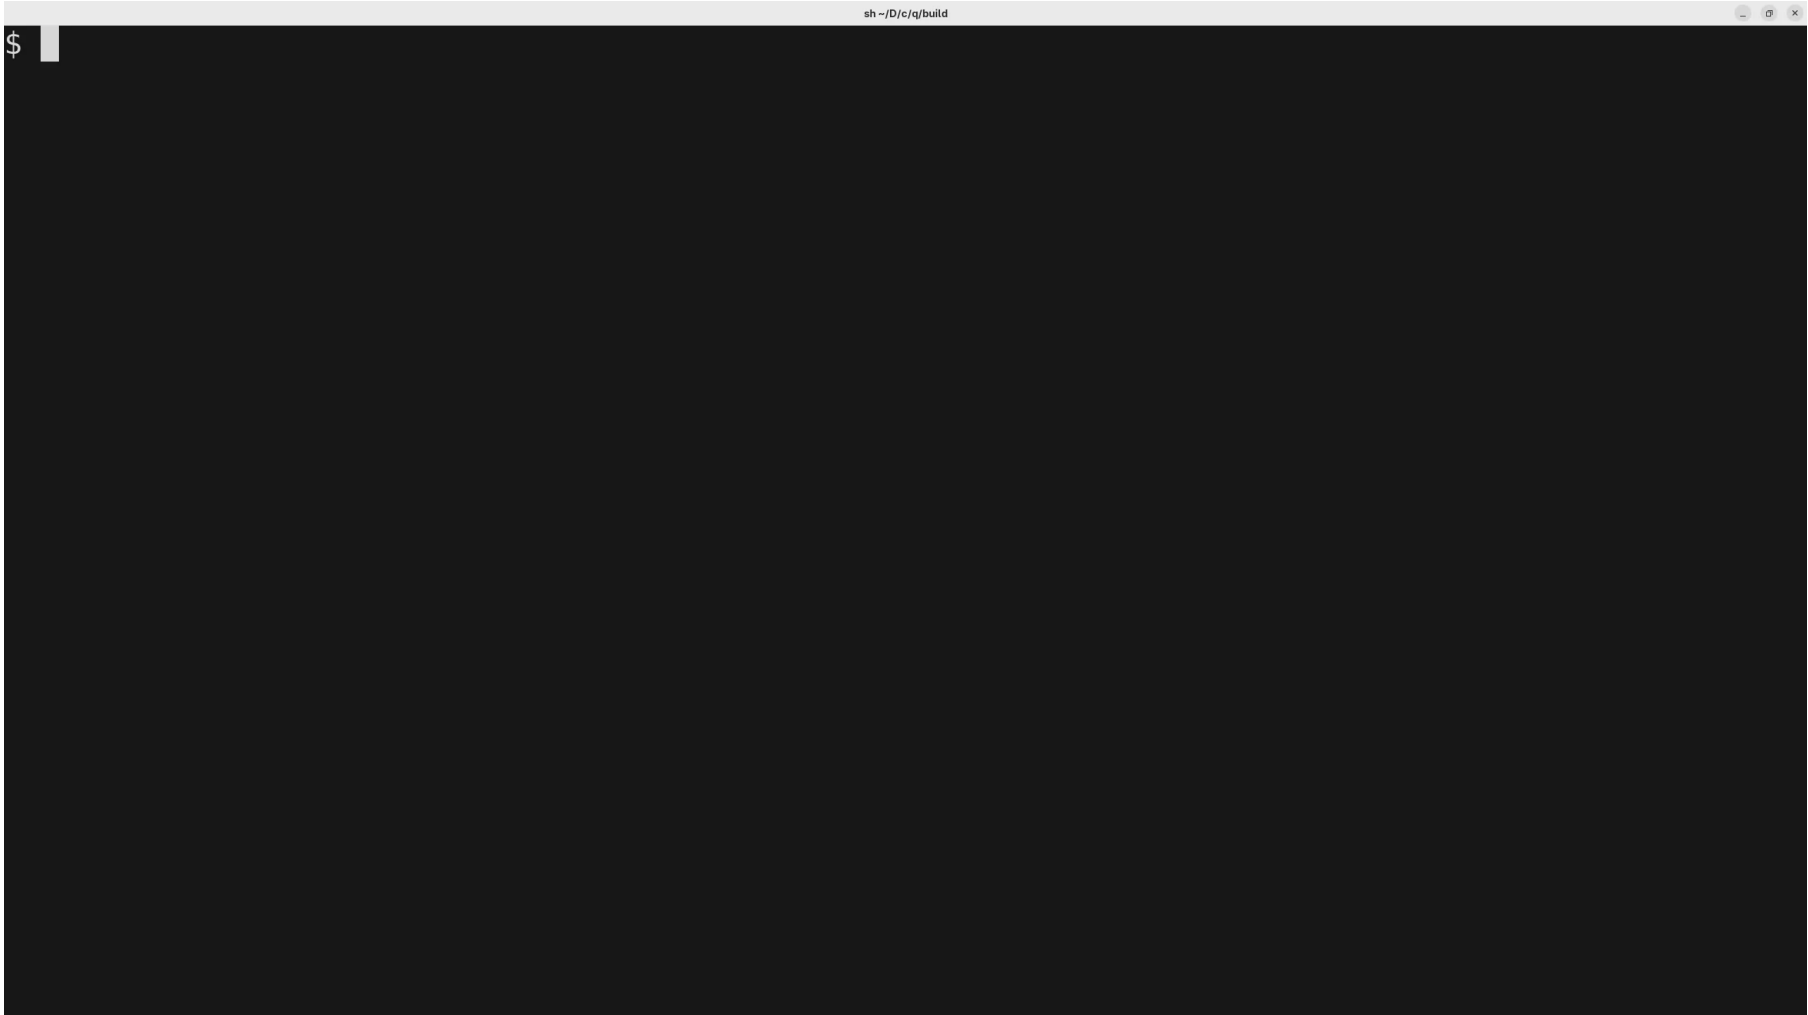
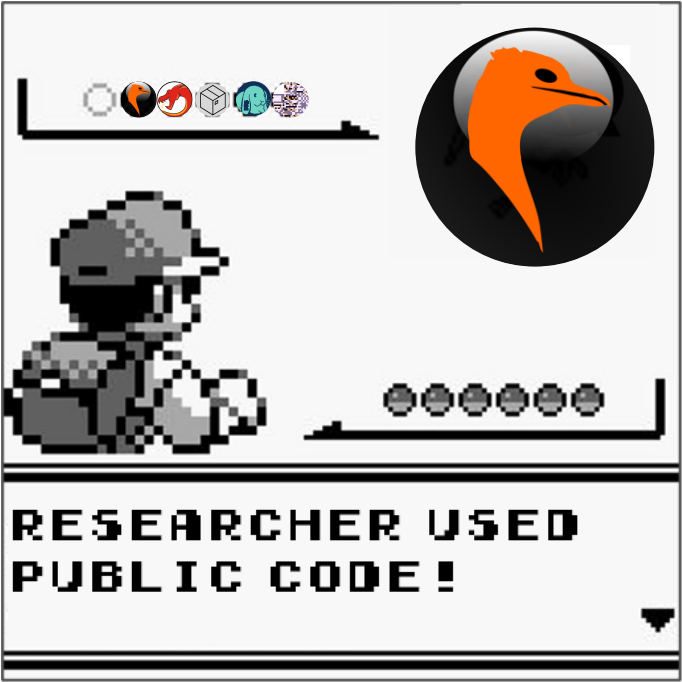


Step	Details	Screenshots
<div>Dynamic Debugging</div> <div>Example problem</div>	<ul style="list-style-type: none"><li>▪ <b>Reading the source</b> and QEMUs <b>logging infrastructure</b> is extremely helpful</li><li>▪ <b>Boot is stuck</b> without error</li></ul> <pre>\$ qemu-system-hexagon -display none -kernel qdsp6sw.mbn -d mmu</pre> <ul style="list-style-type: none"><li>▪ <b>TLB issue</b> and its location indicated by QEMU logs</li><li>▪ <b>TLB index</b> out of bounds</li><li>▪ <b>Unexpected behaviour</b> triggered</li></ul>	<pre>hexagon_tlb_fill: tid = 0x0, pc = 0xfe002f20, e = 4, MMU_DATA_LOAD, probe = 0, MMU_KERN TLB miss RW exception (0x6) caught: Cause code 0xfe002f20, BADVA = 0xfe104a88 hex_tlb_lock: 0</pre> <p><i>TLB logs during emulation</i></p>
<div>Fix</div>	<ul style="list-style-type: none"><li>▪ <b>Identify max</b> used TLB size</li><li>▪ <b>Modify</b> <code>target/hexagon/hex_mmu.h</code></li></ul>	<pre>#define NUM_TLB_ENTRIES 192</pre> <p><i>TLB size in QEMU machine definition</i></p>

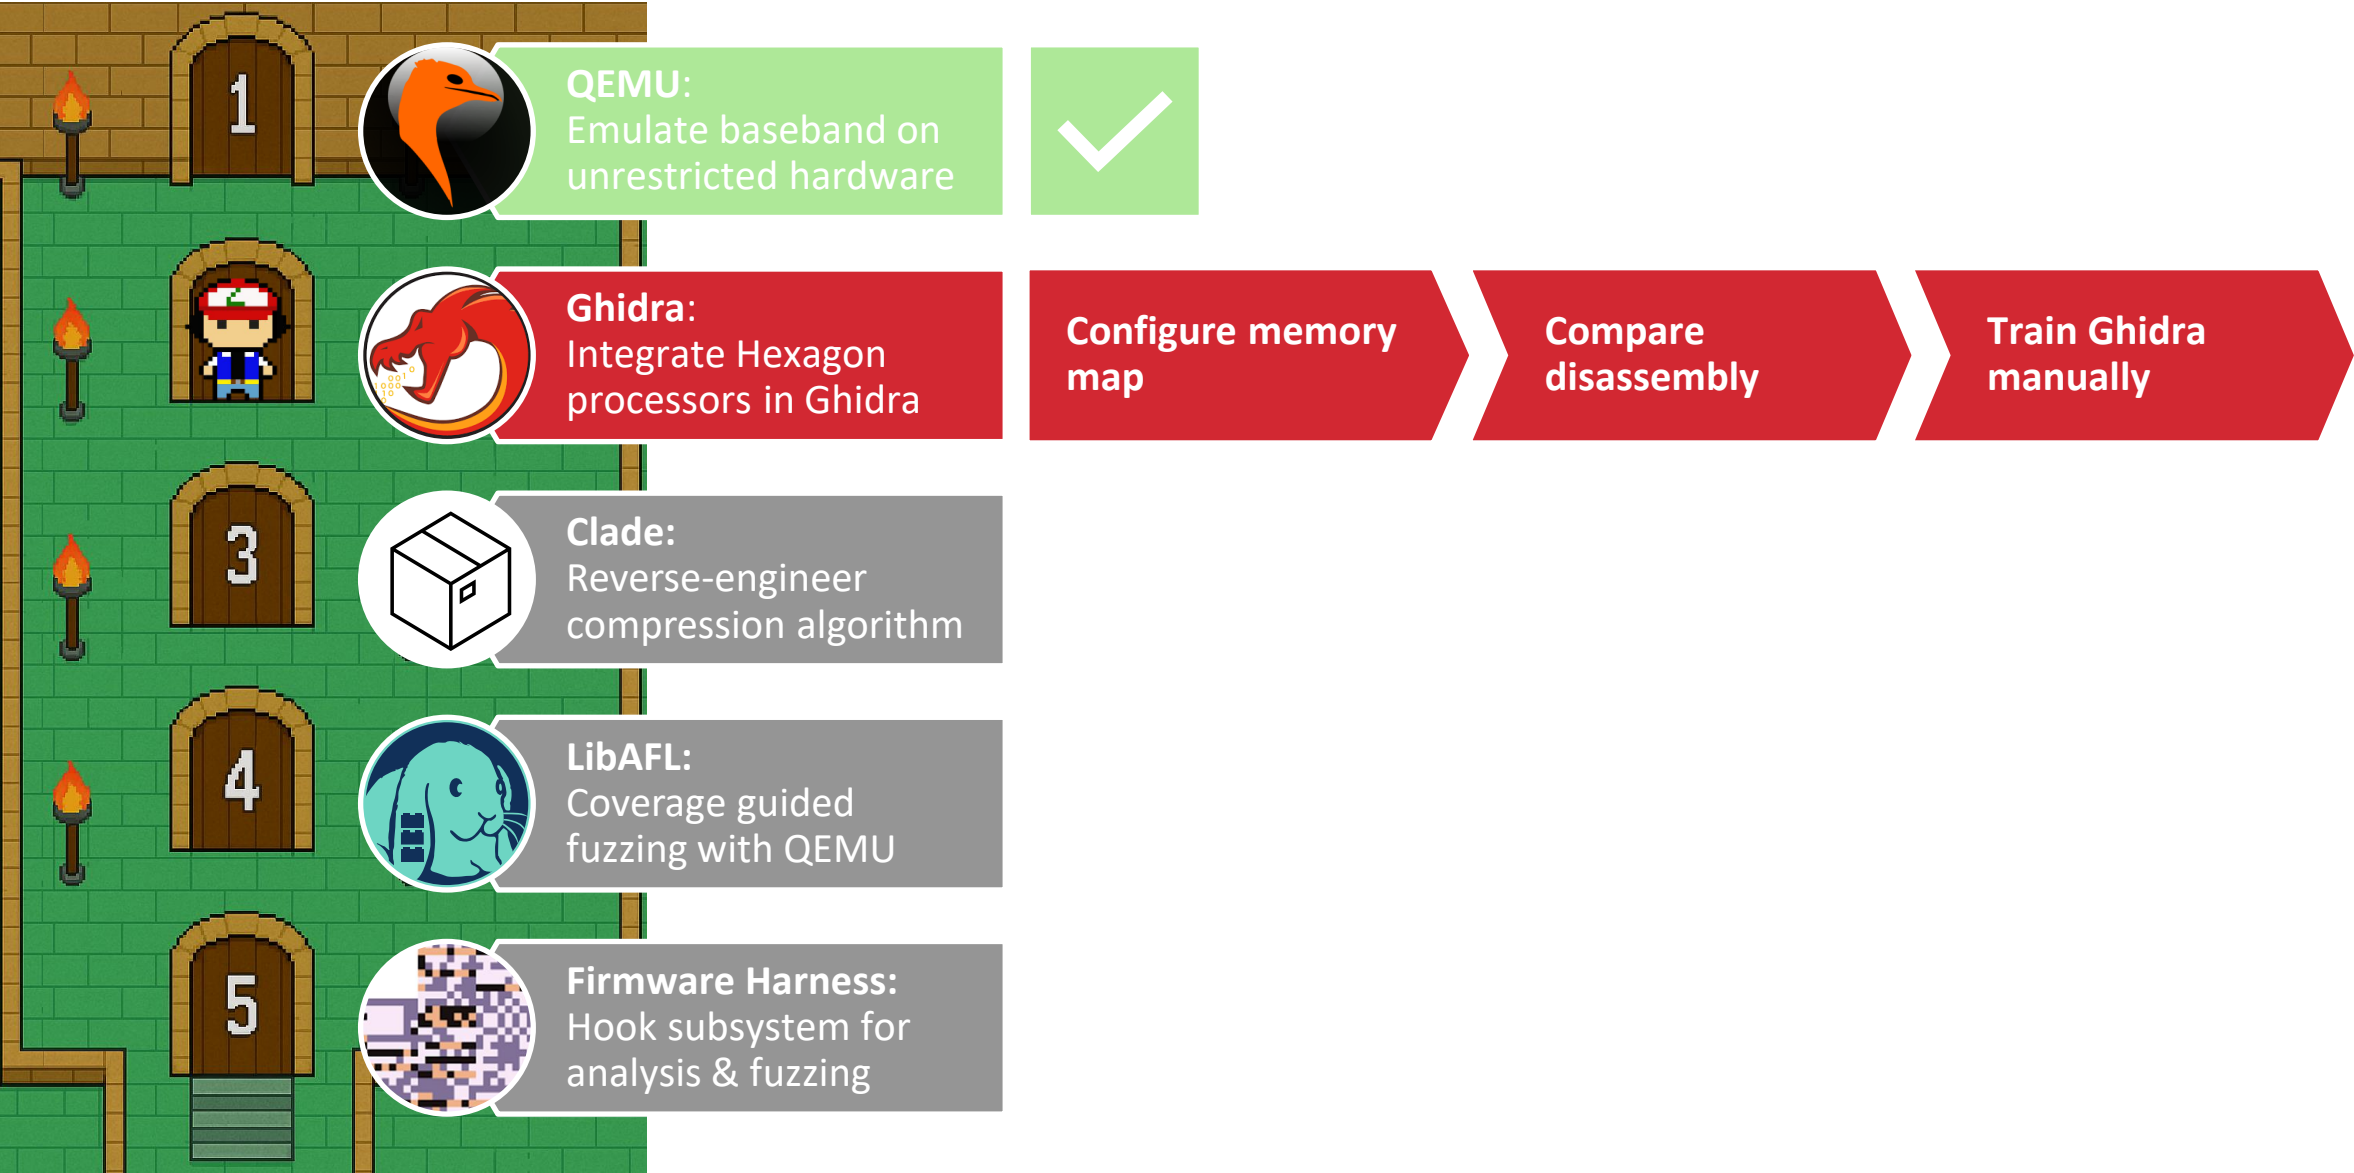
# Hacking around the TLB size yields one more line of boot log



- Obtain firmware
- Get full system Emulation
- Patch for our device



# After the emulator runs, we deeply inspect the firmware with Ghidra





# Reversing the firmware with Ghidra involves modifying the memory map



Configure memory map

Compare disassembly

Train Ghidra manually

Step                      Details                      Screenshots

Problems with reversing using Ghidra

- Use Ghidra with Hexagon plugin
- **Misidentified** memory segments
- **Requires** manual correction

Fix: Configure memory map manually

- **Memory segments defined by QEMU machine**
- **Extract memory segment details from the firmware (mtree)**
- **Adjust memory map to set correct start and end addresses for the segments**

```
(qemu) info mtree
address-space: cpu-memory-0
address-space: cpu-memory-1
address-space: cpu-memory-2
address-space: cpu-memory-3
address-space: cpu-memory-4
address-space: cpu-memory-5
address-space: memory
0000000000000000-ffffffffffffffff (prio 0, i/o): system
0000000000000000-00000000ffffffff (prio 0, ram): lpddr4.ram
00000000d8180000-00000000d81801ff (prio 0, rom): config_table.rom
00000000d81e0000-00000000d81e0fff (prio 0, i/o): fast
00000000d8400000-00000000d87fffff (prio 0, ram): vctm.ram
00000000fab20000-00000000fab20fff (prio 0, i/o): qutimer
00000000fc910000-00000000fc910fff (prio 0, i/o): l2vic
00000000fc921000-00000000fc922fff (prio 0, i/o): qutimer_views
0000000091000000-0000000091ffffff (prio 0, ram): cpz.ram
```

Memory segments information using mtree

segment_26	cde80000	cd091fff
segment_27	edfbc000	edfbcfff
config_table.rom	d8180000	d81801ff
fast	d81e0000	d81e0fff
vctm.ram	d8400000	d87fffff
qutimer	fab20000	fab20fff
l2vic	fc910000	fc910fff
qutimer_views	fc921000	fc922fff
segment_4	fe000000	fe003fff
segment_31	fe200000	ff7e1fff
_elfHeader	_elfHeader::0...	_elfHeader::0...
elfProgramHeaders	elfProgramHe...	elfProgramHe...
cpz.ram	cpz.ram::9100...	cpz.ram::91ff...
unallocated_0	unallocated_0...	unallocated_0...
lpddr4.ram	lpddr4.ram::0...	lpddr4.ram::f...

Memory mapping in Ghidra

# Comparing the disassembled code with Hexagon tools output refines the reversing efforts



Step	Details	Screenshots																				
Problem	<ul style="list-style-type: none"><li>▪ <b>Gap:</b> Ghidra misses registers, memory addresses and VLIW boundaries</li><li>▪ <b>Poor readability:</b> Opcodes are less readable compared to <code>objdump</code> output</li></ul>	<div><div>LAB_c0312040XREF[1]:</div><table><tr><td>c0312040</td><td>04 60 00 7c</td><td>{ A2_combineii</td><td>R5R4 0x0 0x1</td></tr><tr><td>c0312044</td><td>01 40 00 78</td><td>A2_tfrsi</td><td>R1 0x0</td></tr><tr><td>c0312048</td><td>81 41 01 3c</td><td>S4_storeirb_io</td><td>R1 0x3 0x1</td></tr><tr><td>c031204c</td><td>02 c8 01 3c</td><td>S4_storeirb_io</td><td>R1 DAT_00000010 0x2</td></tr><tr><td>c0312050</td><td>06 40 c2 91</td><td>{ L2_loadrd_io</td><td>R7R6 R2 0x0</td></tr></table></div> <div>Disassembled in Ghidra with hexagon-plugin</div>	c0312040	04 60 00 7c	{ A2_combineii	R5R4 0x0 0x1	c0312044	01 40 00 78	A2_tfrsi	R1 0x0	c0312048	81 41 01 3c	S4_storeirb_io	R1 0x3 0x1	c031204c	02 c8 01 3c	S4_storeirb_io	R1 DAT_00000010 0x2	c0312050	06 40 c2 91	{ L2_loadrd_io	R7R6 R2 0x0
	c0312040	04 60 00 7c	{ A2_combineii	R5R4 0x0 0x1																		
c0312044	01 40 00 78	A2_tfrsi	R1 0x0																			
c0312048	81 41 01 3c	S4_storeirb_io	R1 0x3 0x1																			
c031204c	02 c8 01 3c	S4_storeirb_io	R1 DAT_00000010 0x2																			
c0312050	06 40 c2 91	{ L2_loadrd_io	R7R6 R2 0x0																			
Fix	<ul style="list-style-type: none"><li>▪ <b>Validate:</b> Use Hexagon SDK’s tool <code>hexagon-llvm-objdump</code> to validate disassembled code in Ghidra</li><li>▪ <b>Update</b> the register names in Ghidra, if needed</li></ul>	<div><table><tr><td>c0312040:</td><td>04 60 00 7c</td><td>7c006004 {</td><td>r5:4 = combine(#0,#1)</td></tr><tr><td>c0312044:</td><td>01 40 00 78</td><td>78004001</td><td>r1 = #0</td></tr><tr><td>c0312048:</td><td>81 41 01 3c</td><td>3c014181</td><td>memb(r1+#3) = #1</td></tr><tr><td>c031204c:</td><td>02 c8 01 3c</td><td>3c01c802</td><td>memb(r1+#16) = #2 }</td></tr><tr><td>c0312050:</td><td>06 40 c2 91</td><td>91c24006 {</td><td>r7:6 = memd(r2+#0)</td></tr></table></div> <div>Disassembled with hexagon-llvm-objdump (official Hexagon-SDK tool)</div>	c0312040:	04 60 00 7c	7c006004 {	r5:4 = combine(#0,#1)	c0312044:	01 40 00 78	78004001	r1 = #0	c0312048:	81 41 01 3c	3c014181	memb(r1+#3) = #1	c031204c:	02 c8 01 3c	3c01c802	memb(r1+#16) = #2 }	c0312050:	06 40 c2 91	91c24006 {	r7:6 = memd(r2+#0)
c0312040:	04 60 00 7c	7c006004 {	r5:4 = combine(#0,#1)																			
c0312044:	01 40 00 78	78004001	r1 = #0																			
c0312048:	81 41 01 3c	3c014181	memb(r1+#3) = #1																			
c031204c:	02 c8 01 3c	3c01c802	memb(r1+#16) = #2 }																			
c0312050:	06 40 c2 91	91c24006 {	r7:6 = memd(r2+#0)																			

# Renaming structs and functions in Ghidra takes time but enhances code comprehension



Step	Details	Screenshots
Refine function naming	<ul style="list-style-type: none"><li>Manually rename functions based on our knowledge and the print statements in the code</li></ul>	<div><pre>&gt; f prob_set_boot_status &gt; f prob_set_clade2_cfg_base &gt; f prob_task_indexing &gt; f process_kill ✓ q &gt; f qmi_time_client_connect</pre></div> <div>Renamed functions in Ghidra</div>
Define data-structs	<ul style="list-style-type: none"><li>Configure data types and structs based on open-source code and header definitions</li></ul>	<div><div>Data Type Manager</div><div><div>qurt_thread.h</div><div><div>defines</div><div><div>qurt_cache_partition_t</div><div>qurt_thread_attr_t</div><div>qurt thread t</div></div></div></div></div> <div>Configured structs in Ghidra</div>
Analyze register states	<ul style="list-style-type: none"><li>Analyze runtime register values using QEMU to compare with the code logic</li></ul>	<div><pre>(qemu) info registers  CPU#0 TID 0 : General Purpose Registers = {   r00 = 0x00000000   r01 = 0x00000000   r02 = 0x001c0071   r03 = 0x0000000c</pre></div> <div>Runtime memory analysis</div>

# Bonus: Mapping boot flow with the decompiled code calls for automation



Configure memory map

Compare disassembly

Train Ghidra manually

Step	Details	Screenshots
Capture the boot trace	<ul style="list-style-type: none"><li>▪ <b>Capture full trace</b> with QEMU monitor and save to a file</li></ul> <pre>./qemu-system-hexagon -kernel qdsp6sw.mbn -monitor stdio -d exec 2&gt;trace.txt</pre>	
Create a Ghidra script	<ul style="list-style-type: none"><li>▪ <b>Parse</b> the trace using a Python script</li><li>▪ <b>Highlight:</b> Color each memory address reached during the boot process</li></ul>	
Optimize the script	<ul style="list-style-type: none"><li>▪ <b>Scale:</b> Update script to support multiple threads</li><li>▪ Different colors for different threads</li></ul>	

Colored opcodes in Ghidra

# We can now follow along emulator execution in our Ghidra setup



Configure  
memory map

Compare  
disassembly

Train Ghidra  
manually

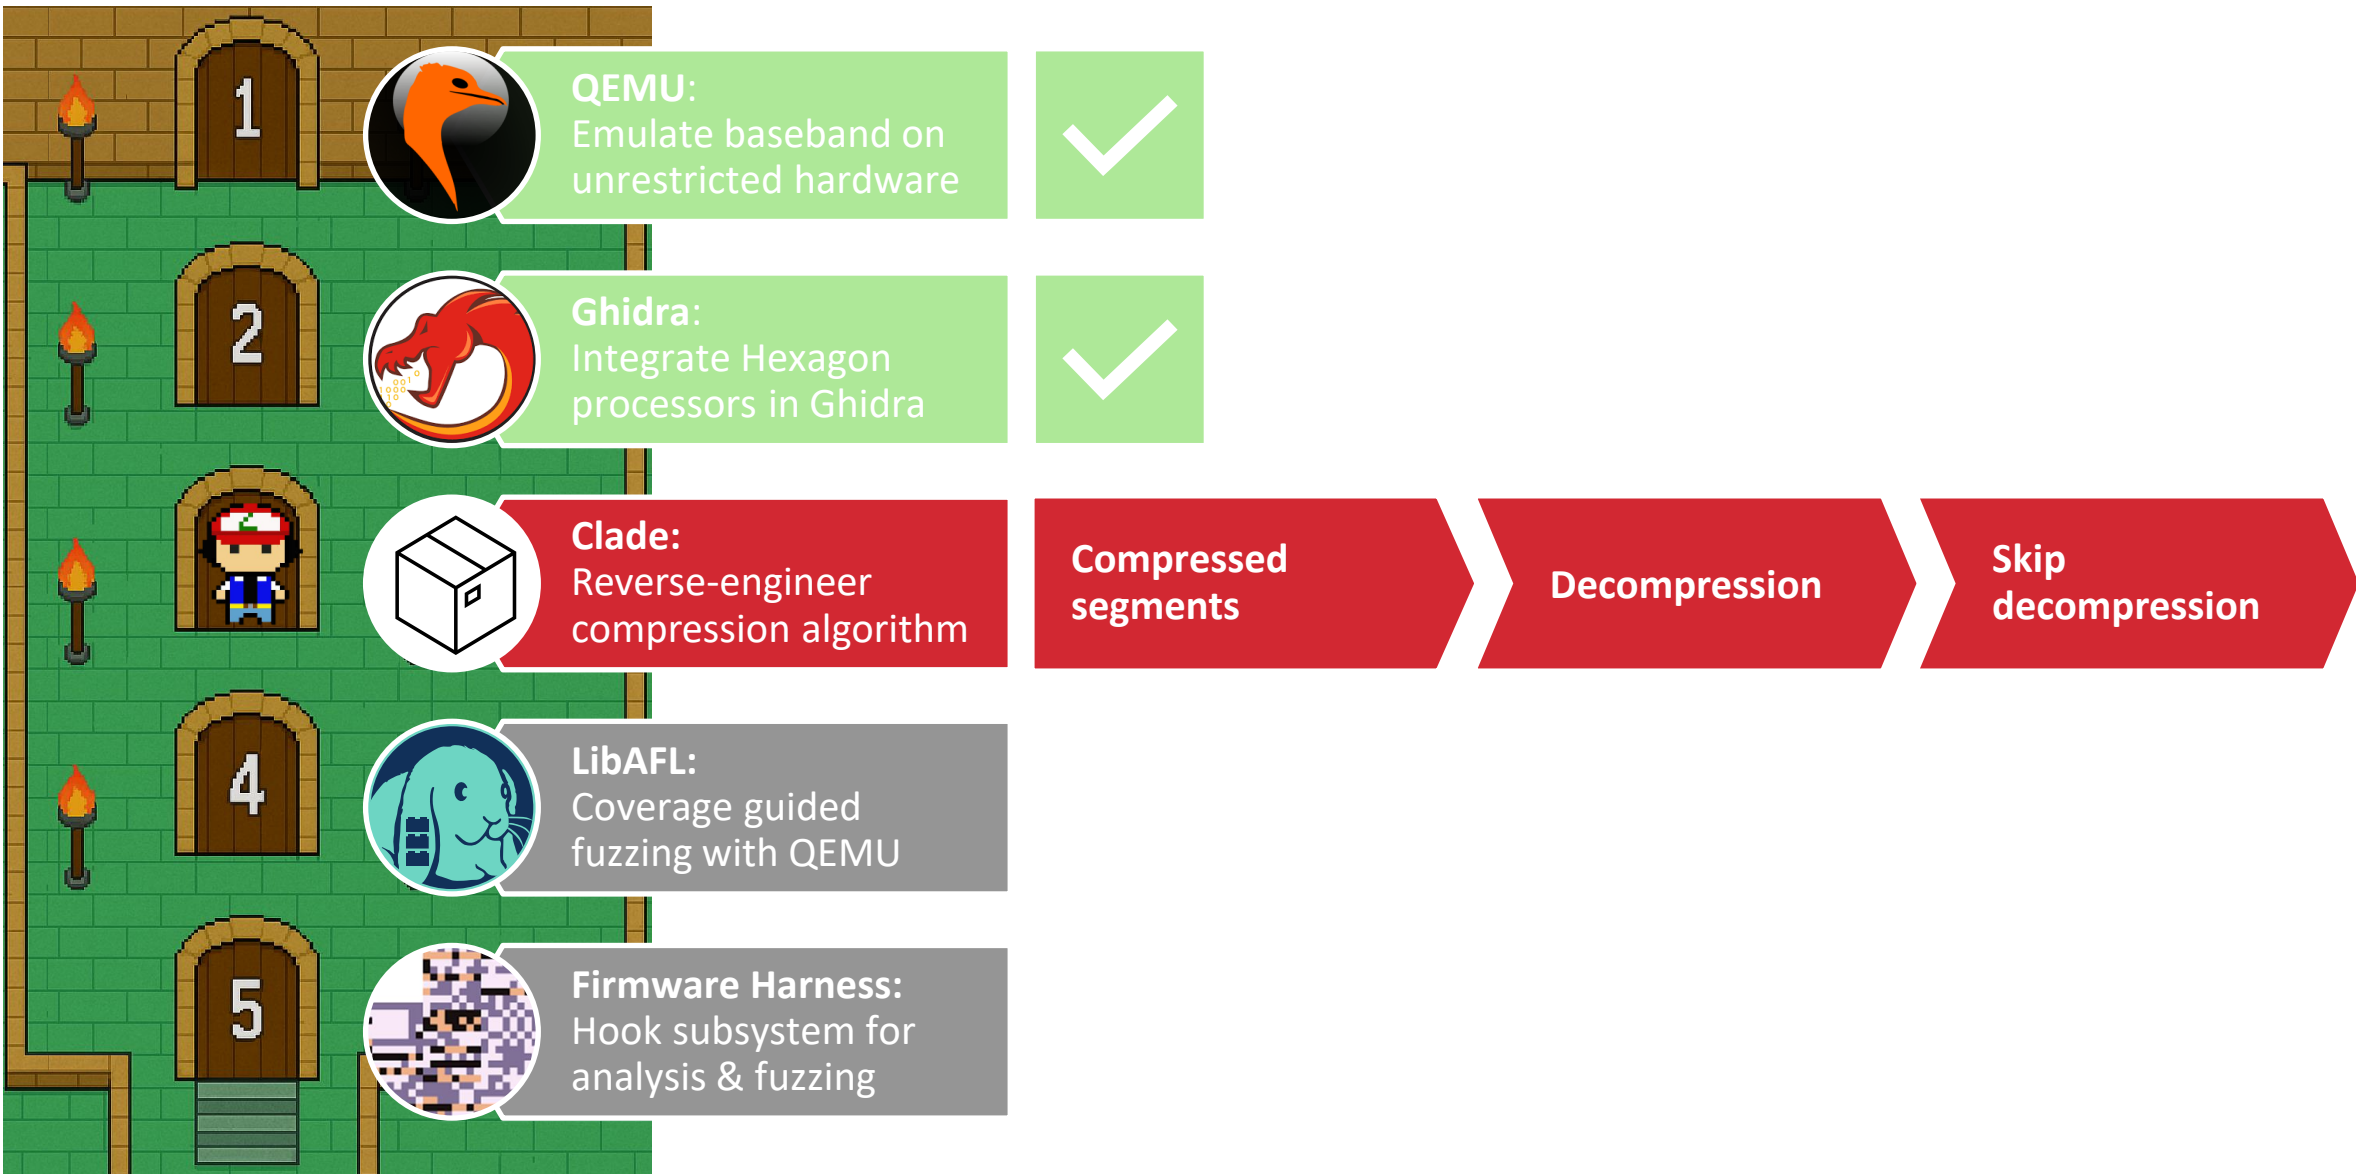
```
fe10c080 4a 75 fe 5b { J2_call      prob_set_boot_status
fe10c084 00 c3 00 78 A2_tfrsi      uVar4 0x18
fe10c088 14 59 00 5a { J2_call      qurt_printf
fe10c08c 95 46 e1 0f A4_ext        s_0x%x,_BADVA:_0x%x,_SSR:_0x%x,_SP_fe11a53b+5
fe10c090 c0 c7 00 78 A2_tfrsi      uVar4=>s_QURT_kernel_started_fe11a57e s_QURT_k...
fe10c094 40 75 fe 5b { J2_call      prob_set_boot_status
fe10c098 20 c3 00 78 A2_tfrsi      uVar4 0x19
fe10c09c 0a 59 00 5a { J2_call      qurt_printf
fe10c0a0 96 46 e1 0f A4_ext        s_RT_kernel_started_fe11a57e+2
fe10c0a4 60 c2 00 78 A2_tfrsi      iVar8=>s_QURT_kernel_init_cache_params_fe11a59...
fe10c0a8 b4 f4 00 5a { J2_call      FUN_fe112a10
fe10c0ac ba e4 01 5a { J2_call      FUN_fe118a20
fe10c0b0 0c 42 e1 0f { A4_ext        DAT_fe108300
fe10c0b4 02 c0 00 78 A2_tfrsi      uVar12=>DAT_fe108300 DAT_fe108300
fe10c0b8 5c 6b 00 5a { J2_call      some_struct_construction_subtree
fe10c0bc 01 c0 42 3c S4_storeiri_io uVar12=>DAT_fe108300 0x0 0x1
fe10c0c0 2a 75 fe 5b { J2_call      prob_set_boot_status
fe10c0c4 40 c3 00 78 A2_tfrsi      iVar8 0x1a
fe10c0c8 f4 58 00 5a { J2_call      qurt_printf
fe10c0cc 96 46 e1 0f A4_ext        s_RT_kernel_started_fe11a57e+2
fe10c0d0 40 c6 00 78 A2_tfrsi      iVar8=>s_QURT_root_task_started_fe11a5b2 s_QURT...
fe10c0d4 a6 ee ff 5b { J2_call      clade_related
fe10c0d8 4c f1 ff 5b { J2_call      clade_related2
```

*Colored opcodes in Ghidra reached during the boot flow*

```
prob_set_boot_status(0x18);
qurt_printf(s_QURT_kernel_started_fe11a57e);
prob_set_boot_status(0x19);
iVar8 = qurt_printf(s_QURT_kernel_init_cache_params_fe11a593);
FUN_fe112a10(iVar8,param_2,iVar7,ppuVar13,param_5,param_6);
uVar6 = FUN_fe118a20();
uVar4 = 0;
DAT_fe108300 = 1;
some_struct_construction_subtree
    (uVar6,param_2,-0x1ef7d00,ppuVar13,param_5,param_6,unaff_R16,unaff_R17,unaff_R30,
    (char)in_R31,param_7);
uVar11 = (undefined)param_2;
prob_set_boot_status(0x1a);
iVar7 = qurt_printf(s_QURT_root_task_started_fe11a5b2);
uVar5 = clade_related((char)iVar7,uVar11,uVar4,(char)ppuVar13,(char)param_5,(char)param_6,
    CONCAT44(unaff_R17,unaff_R16),CONCAT44(in_R31,unaff_R30));
clade_related2(uVar5,uVar11,uVar4,(char)ppuVar13,(char)param_5,(char)param_6,unaff_R16,unaff_R17,
    unaff_R30,(char)in_R31,param_7,param_8);
```

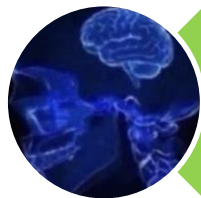
*Matching decompiled code*

# We can observe different compression mechanisms in the firmware, our next stage to tackle





# Various proprietary compression algorithms slow the pace of reverse engineering efforts



**Deltacompress:** Used to compress data in Qualcomm firmware; tools available for decompression



**Q6zip:** Used to compress code in Qualcomm firmware; tools available for decompression



**CLADE:** Replaces Q6zip; requires CLADE dict and config for decompression



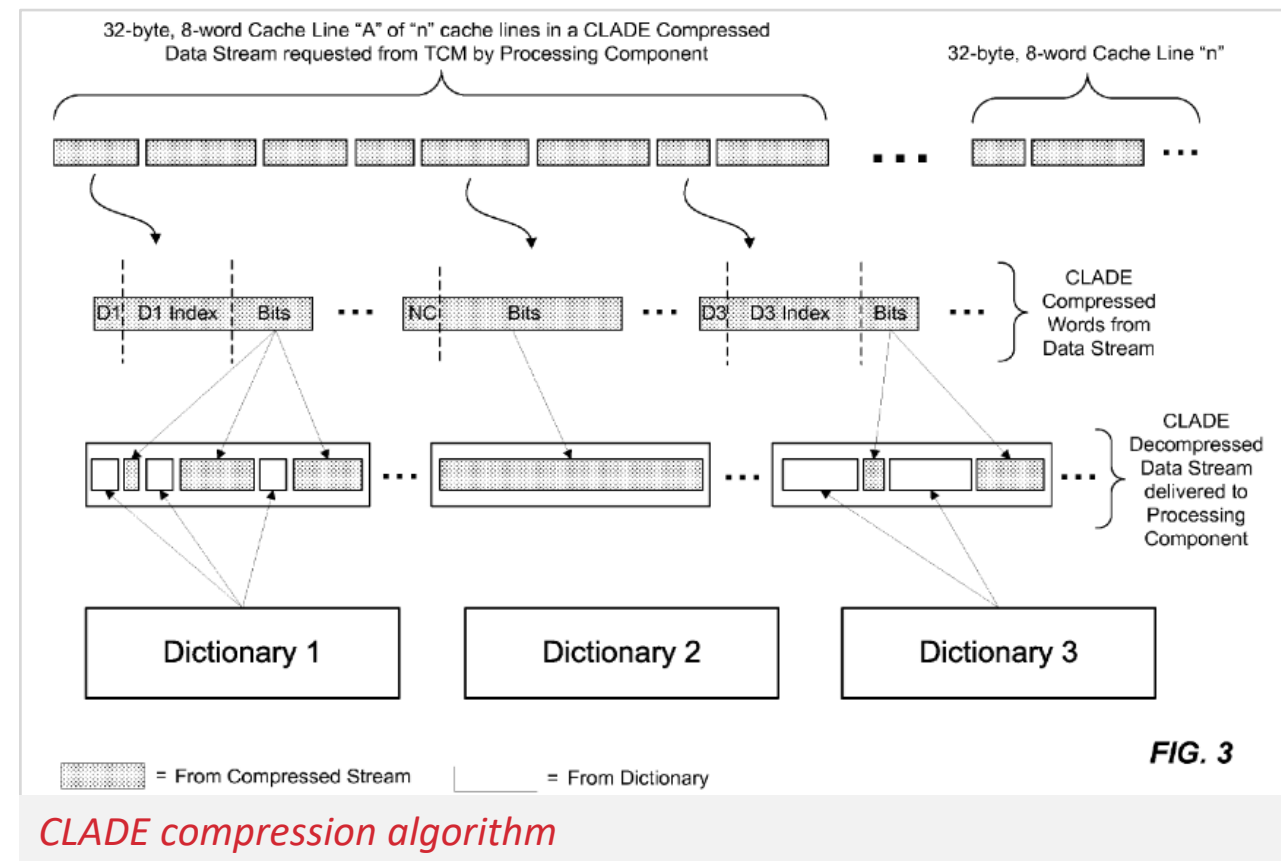
**CLADE2:** Enhanced version of clade; utilizes hardware for decompression



Compressed segments

Decompression

Skip decompression

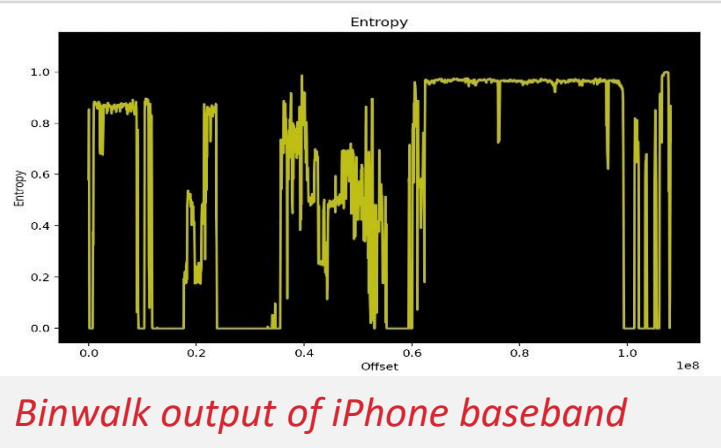


The manual decompression of clade section achieves partial progress, but full resolution requires additional analysis



Step	Details
Decompress firmware sections	<ul style="list-style-type: none"><li>Binwalk and binutils identify compressed sections</li><li>Public tooling decompresses sections successfully</li></ul>
Import sections with Ghidra	<ul style="list-style-type: none"><li>Add decompressed sections with Ghidra Memory Manager</li></ul>

Screenshots



```
...m::d7ffffff  ??  ??
...m::d8000000  ??  ??
...m::d8000001  ??  ??
...m::d8000002  ??  ??
...m::d8000003  ??  ??
...m::d8000004  ??  ??
...m::d8000005  ??  ??
...m::d8000006  ??  ??
...m::d8000007  ??  ??
...m::d8000008  ??  ??
...m::d8000009  ??  ??
...m::d800000a  ??  ??
...m::d800000b  ??  ??
```

→

```

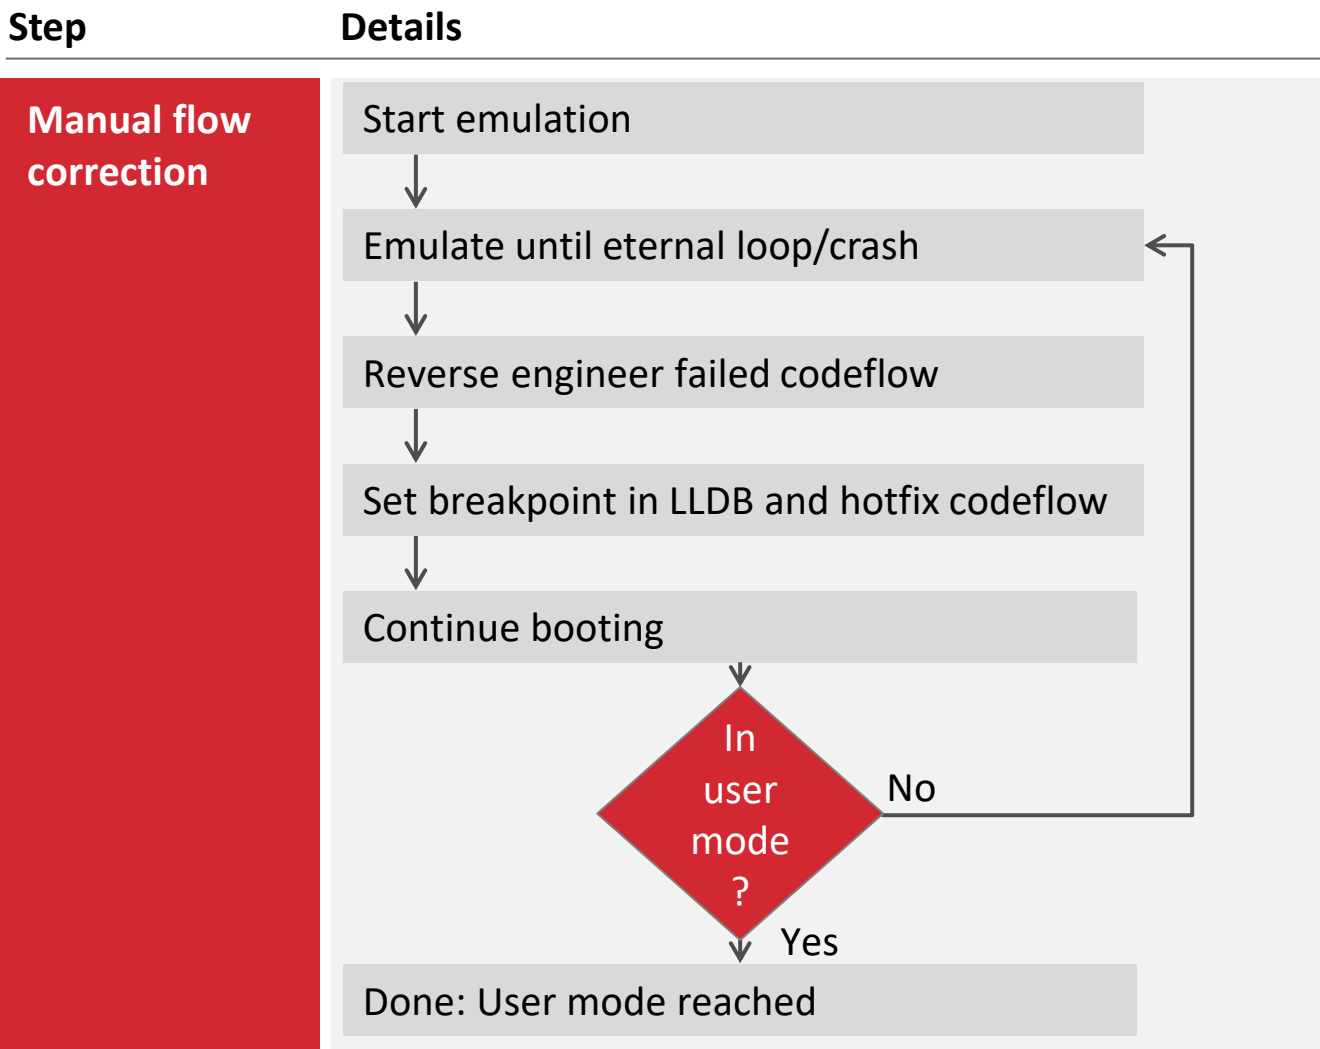
FUN_d8000000
d8000000 b8 58 01 5a { J2_call FUN_d800b170
d8000004 00 c0 00 78 A2_tfrsi R0 0x0
d8000008 14 c0 0c 10 { J4_cmpeq... R20 0x0 LAB_d8000030
d800000c 8a 4c 30 5b { J2_call SUB_d7981920
d8000010 21 40 15 b8 A2_addi R1 R21 -0x7fff
d8000014 00 40 71 70 A2_tfr R0 R17
d8000018 42 eb 12 78 A2_tfrsi R2 0x255a
d800001c 00 62 30 73 { A4_combi... R1R0 R16 0x10
d8000020 dc 52 00 0e A4_ext 0xe004b700
d8000024 a2 43 00 78 A2_tfrsi R2 0xe004b71d
d8000028 c3 eb 12 78 A2_tfrsi R3 0x255e
d800002c 60 e5 ff 5b { J2_call SUB_d7ffcaec
```

Before decompression

After decompression



# Using hexagon-lldb, enables dynamic analysis on the firmware



### Screenshots

```
init static mapping: ppn 0x00005780, vpn 0x00005780, pages 0x00000040
init static mapping: ppn 0x00004900, vpn 0x00004900, pages 0x00000010
finished static mem
App Images Init
```

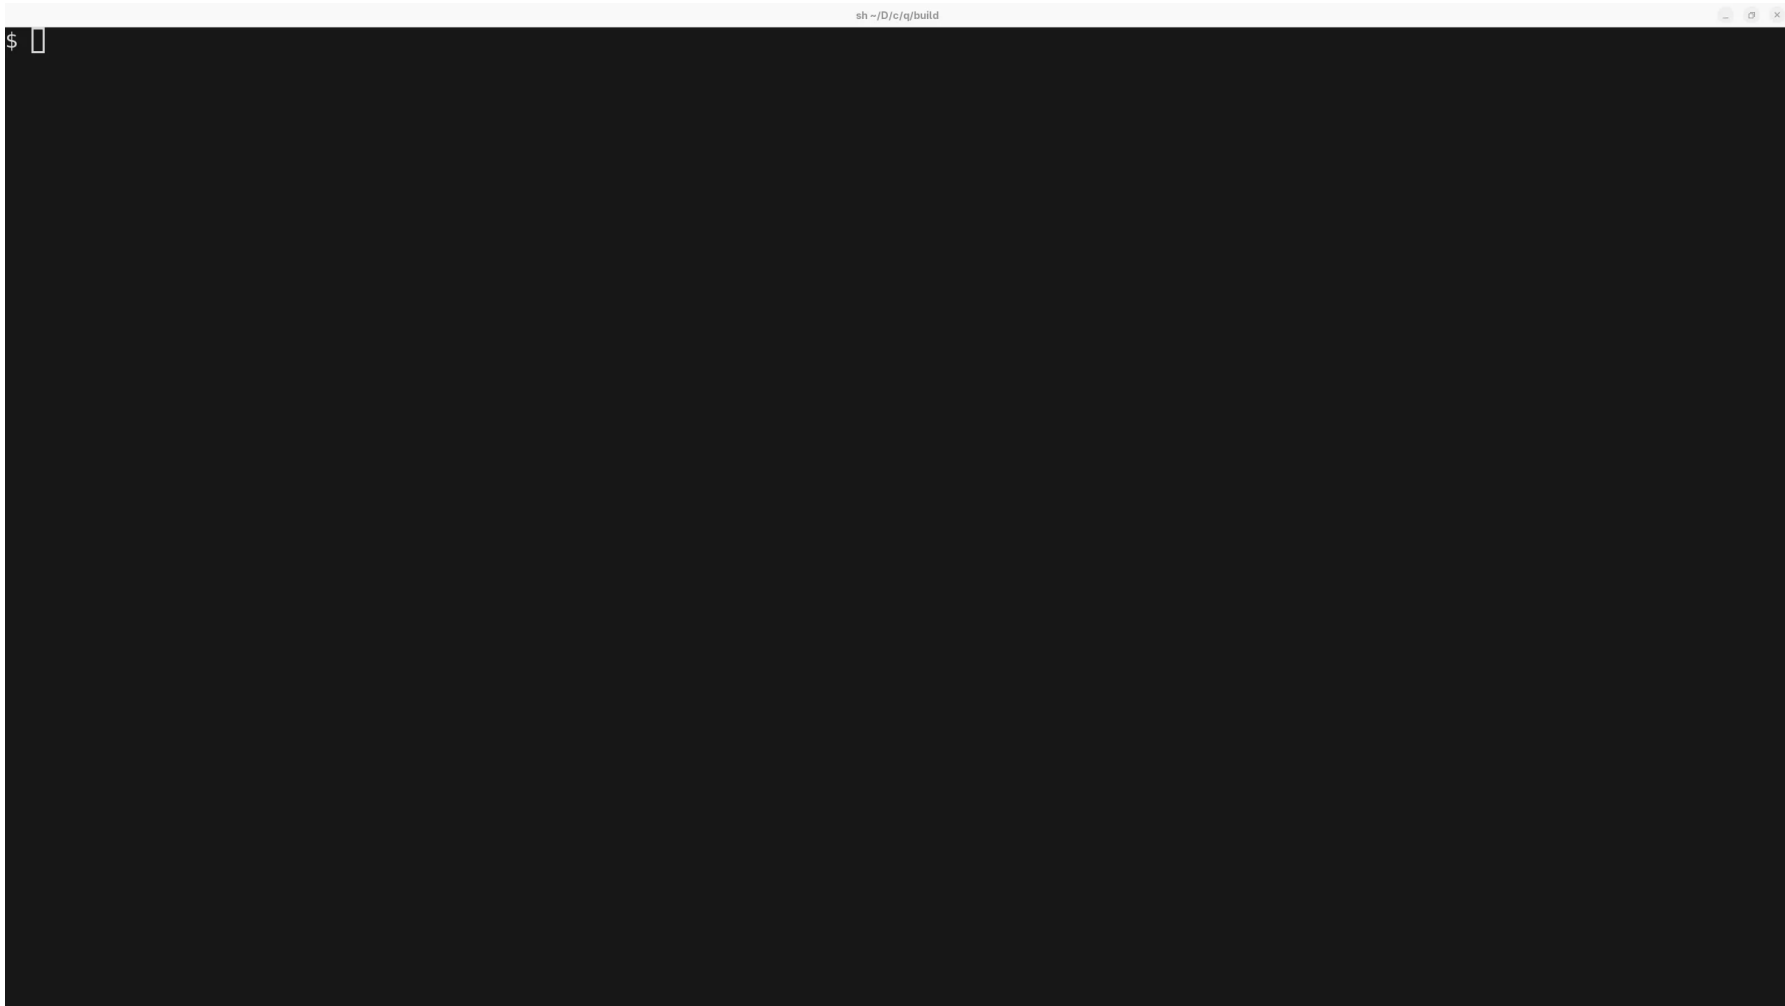
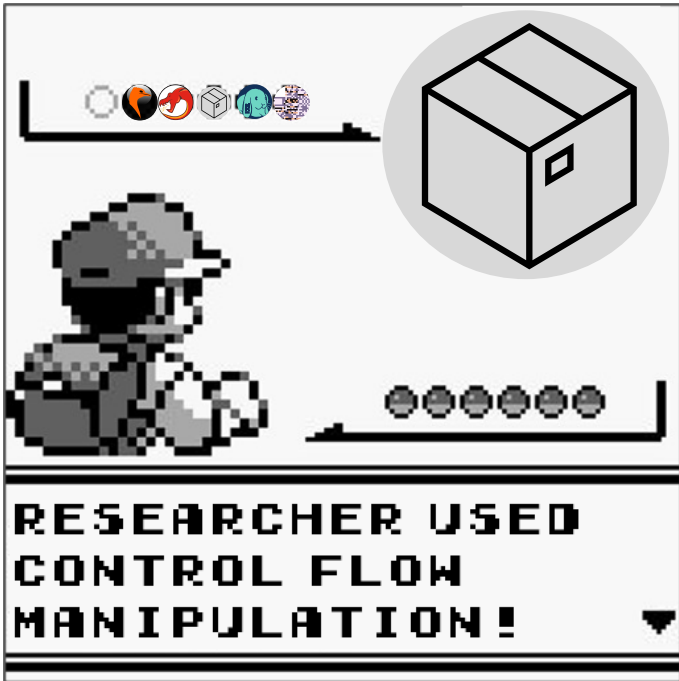
*Finished initializing memory areas*

```
[(lldb) sysstatus
modectl: 0x003d003f
bestwait: 0x1ff / 511 (dec)
schedcfg: 0x0000010f - int #0f / 15 (dec), EN:enabled
syscfg: 0x0095807f
```

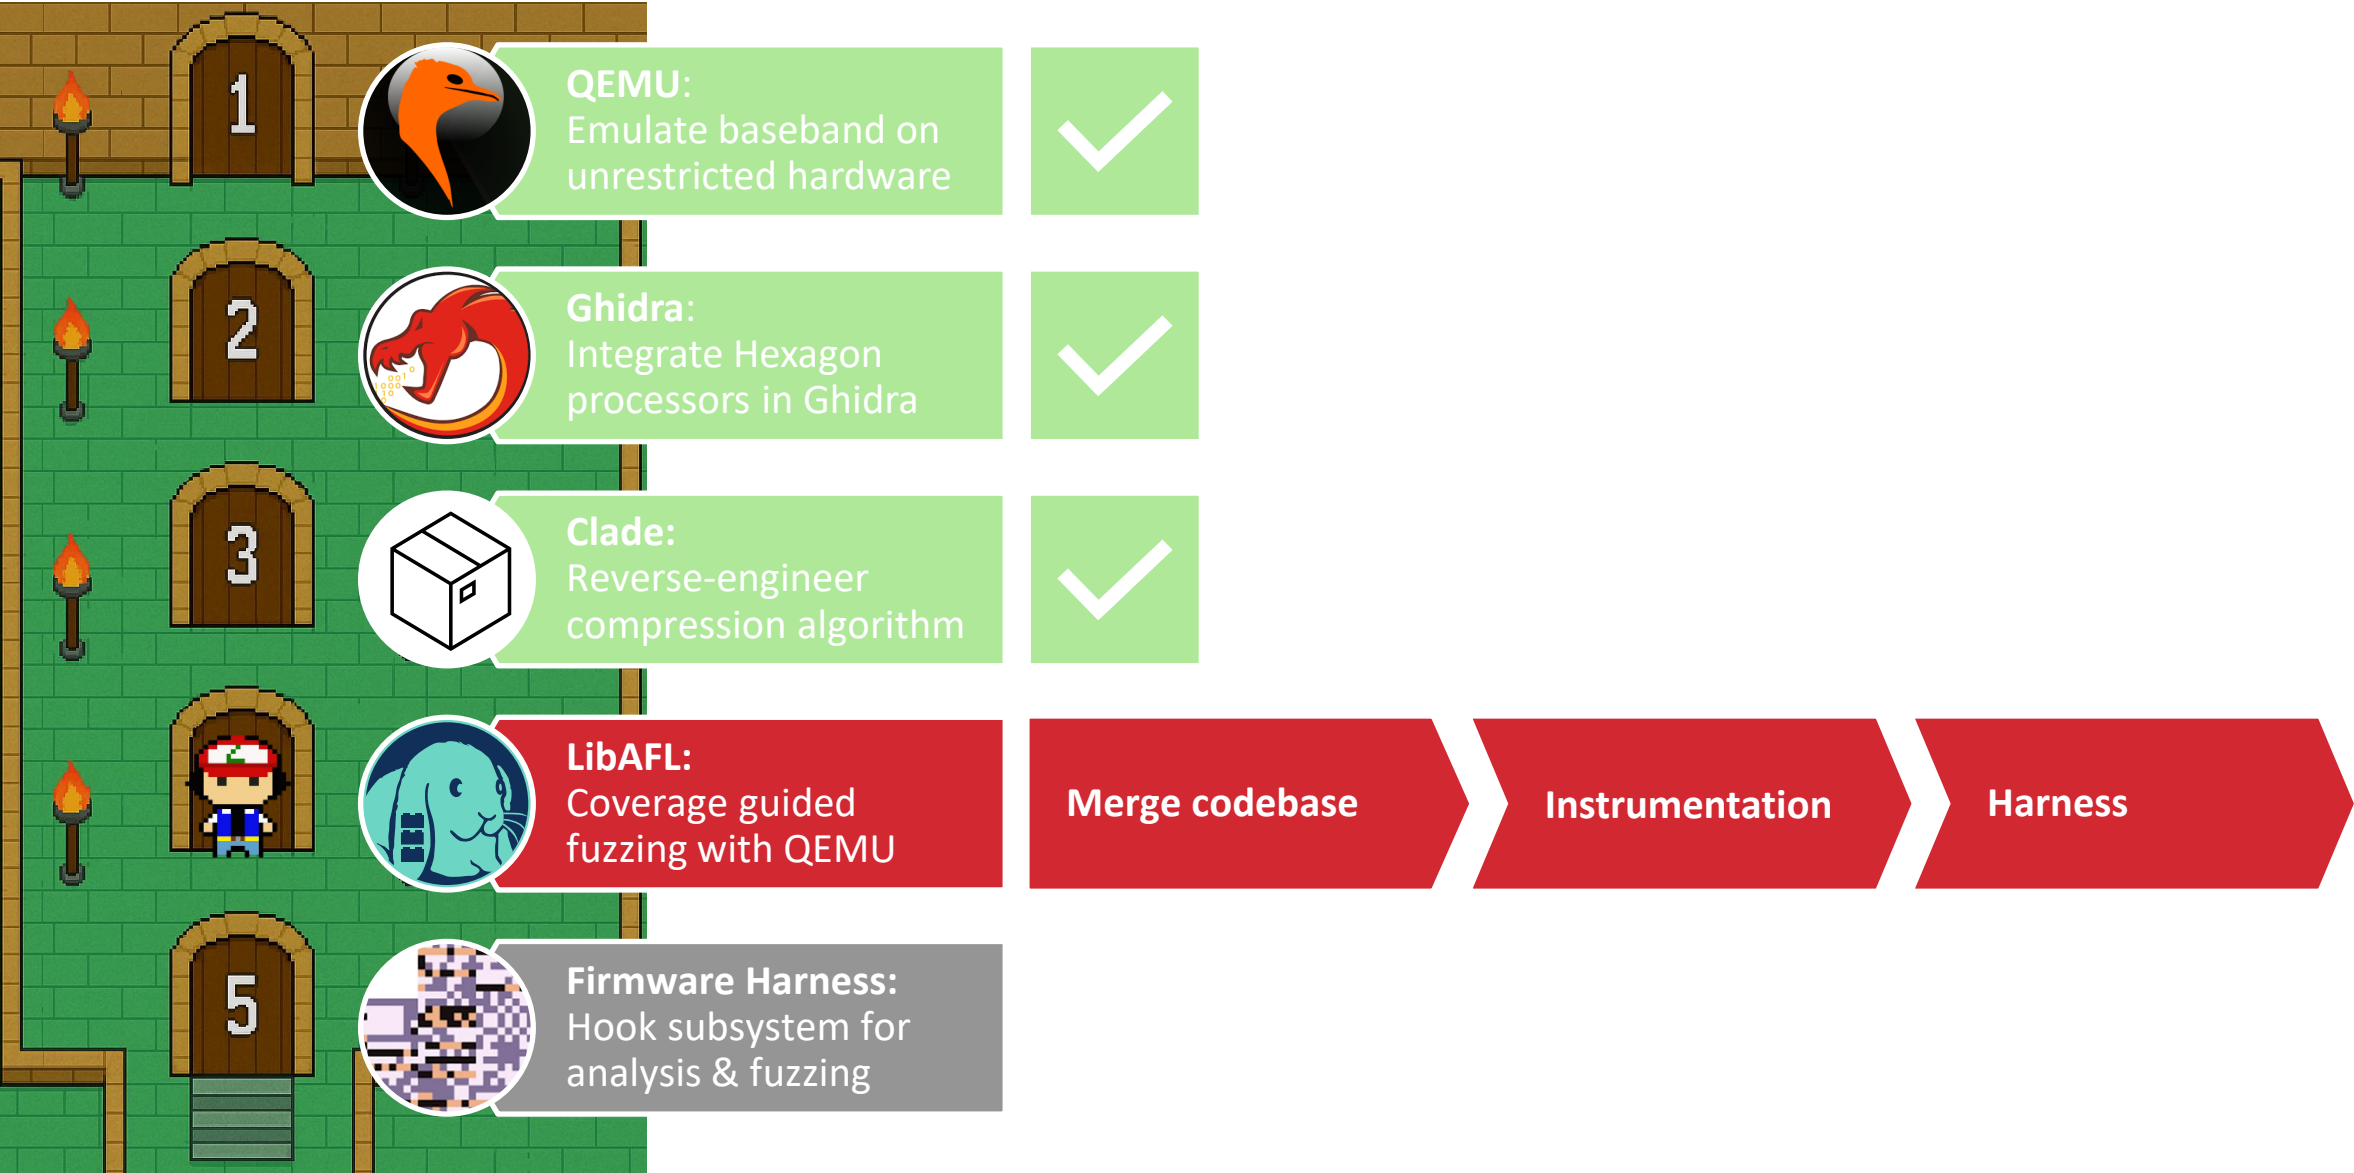
TID	Prio	Mode	Priv	Cause	Set	Unset
0	255	WAIT	Monitor	0x43	EX,GM,IE	SS,UM,XE
1	1	RUN	Monitor	0x1d	EX,GM,IE,UM	SS,XE
2	255	WAIT	Monitor	0x43	EX,GM,IE	SS,UM,XE
3	255	WAIT	Monitor	0x43	EX,GM,IE	SS,UM,XE
4	255	WAIT	Monitor	0x43	EX,GM,IE	SS,UM,XE

*Custom LLDB status commands for Hexagon*

Connecting LLDB to QEMU, we can solve code flow errors and skip parts where needed



After controlling the flow manually, we want full control and fuzzing with LibAFL

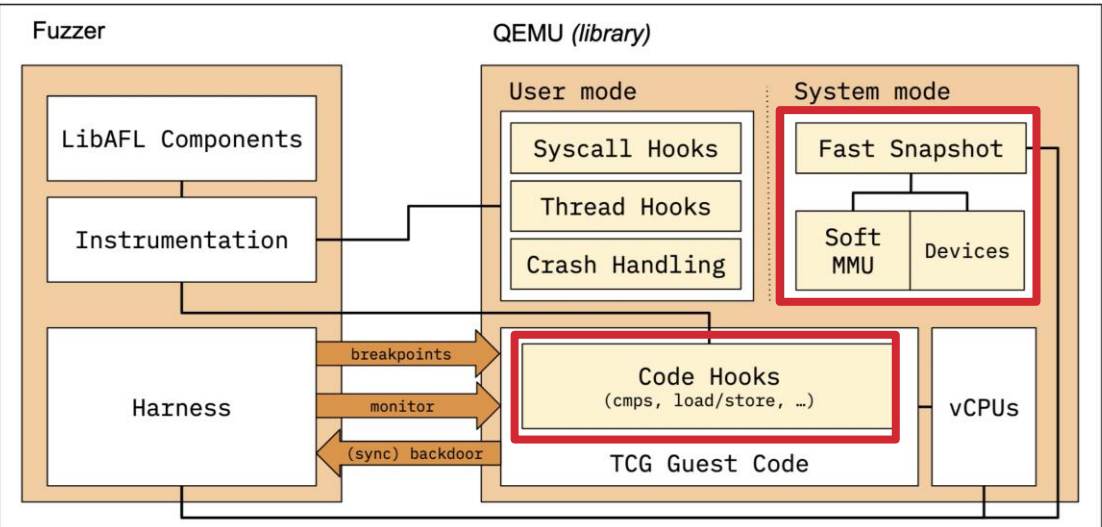


# LibAFL's QEMU provides exactly what we need but merging it with Hexagon QEMU is hard



Step	Details	Screenshots
libafl-qemu-bridge	<ul style="list-style-type: none"><li>▪ <b>Hooks</b> for control flow manipulation</li><li>▪ <b>Snapshots</b> of the process</li><li>▪ <b>User mode emulation</b> implemented for hexagon</li><li>▪ <b>No common ancestor</b> with Hexagon QEMU</li></ul>	
Merge branches manually	<ul style="list-style-type: none"><li>▪ <b>Diff bridge</b> against QEMU base version</li><li>▪ <b>Apply diff</b> on Hexagon QEMU</li><li>▪ <b>Fix</b> compile errors</li></ul>	
Feature integration	<ul style="list-style-type: none"><li>▪ <b>Patch</b> in support for Hexagon system mode support</li></ul>	

LibAFL QEMU Process



LibAFL QEMU architecture diagram

# LibAFL provides hooks, which we use to instrument the baseband boot procedure



Merge codebase

Instrumentation

Harness

Step	Details	Screenshots
------	---------	-------------

## Define breakpoints

- **Rust function hook** in the fuzzer to handle breakpoints
- **Introspection** in common functions like printf
- **JSON configuration** for breakpoints

```
"breakpoints": [  
  {  
    "name": "qurt_println",  
    "address": "0xfe10f2b0",  
    "handler": "HandlePrintln"  
  },  
]
```

*JSON config*

## Compare with disassembled code

- **Colored trace analysis** guides flow manipulation

## Control flow manipulation

- **Boot optimizations** by skipping memory zeroing and device initialization functions
- **Boot progression**

```
QURT kernel started  
QURT kernel init cache params  
QURT root task started  
DLPager_rw_swap_vaddr a0000000, DLPager_rw_swap_size 50000  
QuRTOS heap init: start cb3bc000 size 0x80000  
QuRTOS ISLAND heap init: start bfea0200 size 0x2800  
Init memory pools...  
Found Virtual Pool in overrides. size 124, pool id 4096  
Adding pages 0x00001000->0x000bfe00 to Default Virtual pool.  
0 name DEFAULT_PHYSP00L pool cb3c4490 island 0, collapse_type 0  
Found requested Pool ID in overrides. size 88, pool id 0  
Using over-ride ranges for pool index 0; number of ranges 10  
0: 0x900dc000->0x90100000  
1 name HWIO_POOL pool cb3c4570 island 0, collapse_type 0  
Using configured ranges for pool index 1; number of ranges 1  
0: 0x09000000->0x0a000000  
11 name CLADE_DICT pool cb3c4e30 island 0, collapse_type 0  
Found requested Pool ID in overrides. size 8, pool id 11
```

*Boot logs after control flow manipulation*

# Creating a harness requires taming a 108MB proprietary codebase



Merge codebase

Instrumentation

Harness

Step

Details

Screenshots

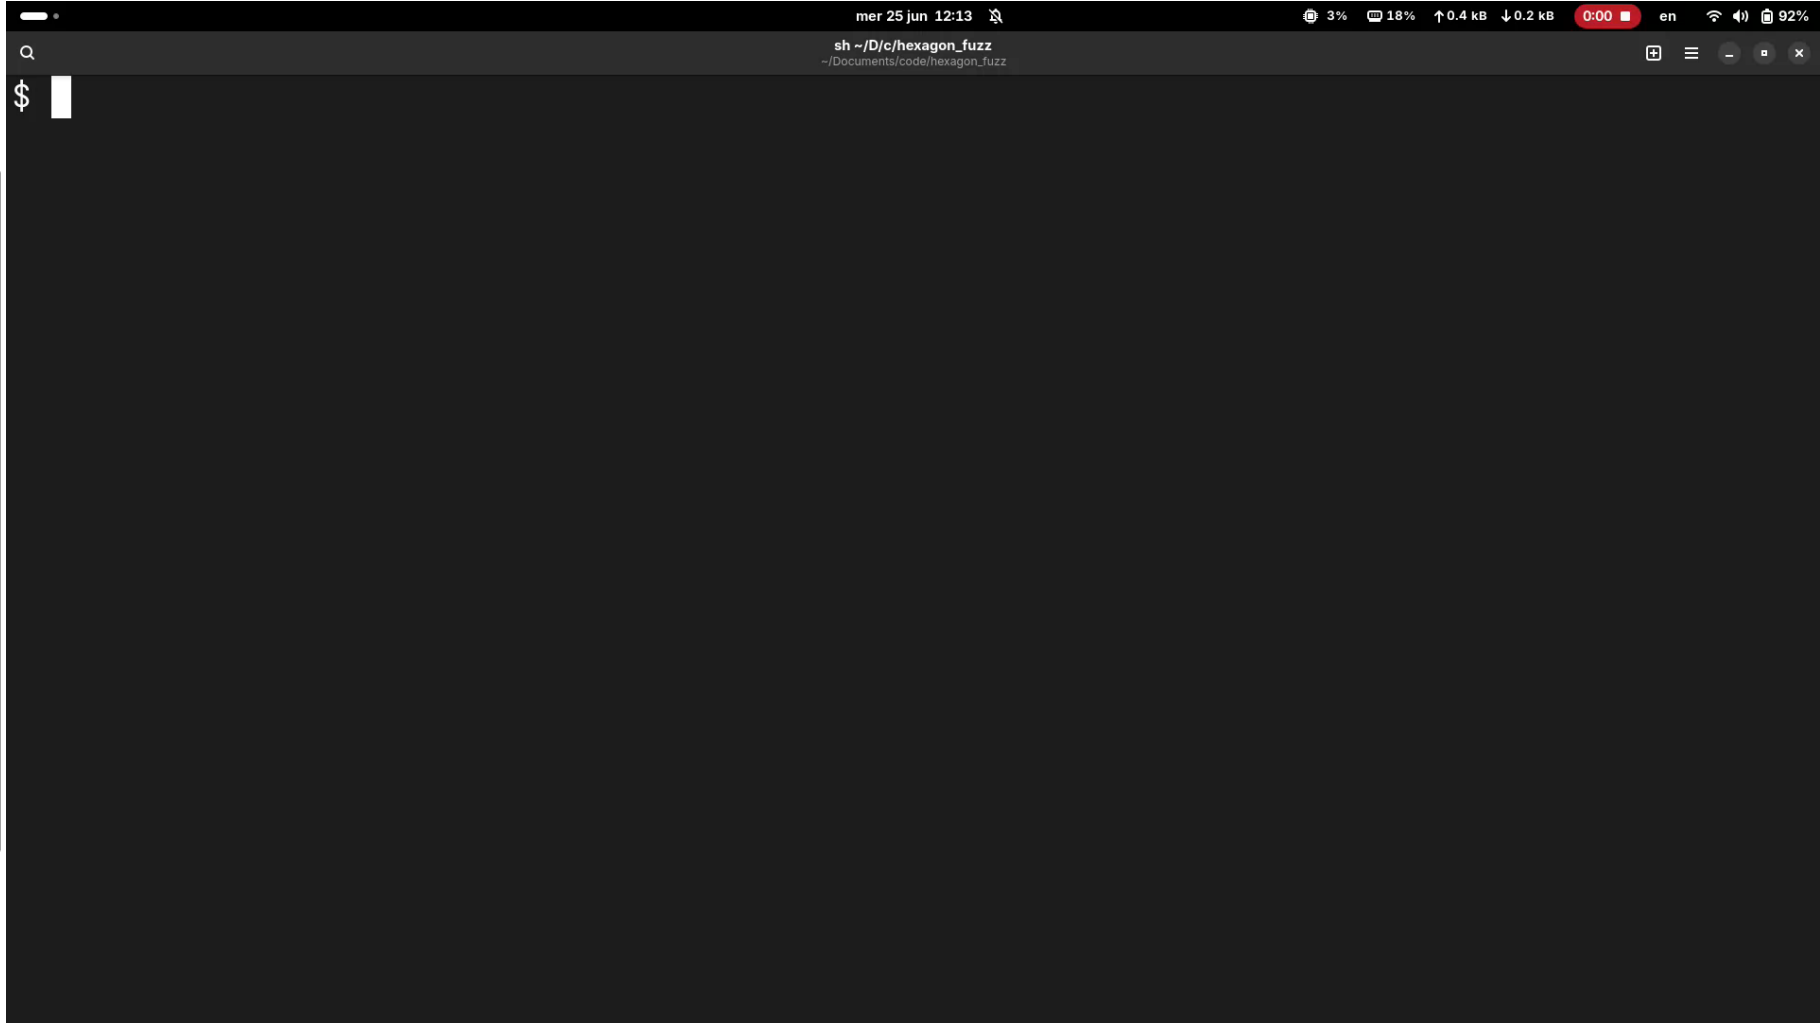
Codebase navigation

- **70k strings, 30k functions** most without references
- **Cross reference** with source code from SDK and internet
- **Utilize Ghidra scripting** to recover function names
- **Compare with Qualcomm tooling** e.g. QXDM monitoring

Thread Name	Thread ID	Proirity	PID	MCPS	Delta CPU %	Total CPU %	Stack Size (bytes)
ml1_mgr	148	98	0	0.9	0.16	0.24	9216
RFLM_QLNK_OFD1	106	51	0	0.79	0.18	0.21	4088
LFW_SCHD_CMN_1	3052	27	0	0.51	0.09	0.13	8112
VSTMR_irq40	20d2	77	0	0.48	0.12	0.13	8192
RFLM_CCS	f7	33	0	0.4	0.1	0.1	4016
rf_fe	176	99	0	0.36	0.11	0.09	32768
gsm_msggr_t1	142	99	0	0.29	0.07	0.08	4608
slpc_worker	17d	131	0	0.32	0.07	0.08	8192
Prof TP	25d04a	124	0	0.25	0.07	0.07	4096
GFW_GENERIC_1	106e	40	0	0.23	0.15	0.06	8192
GFWRF_TRIG2_A0	1070	27	0	0.23	0.1	0.06	4096
GFWRF_TRIG1_A0	1071	27	0	0.23	0.06	0.06	4096
DPC_Task	1000	126	0	0.47	0.05	0.06	4096
RFLM_QLNK	107	33	0	0.18	0.05	0.05	4088
gsm_rr	120	164	0	0.18	0.04	0.05	384

Running processes in QXDM

With integration of LibAFL we successfully boot the Hexagon baseband and are ready to fuzz





With this setup we are ready to explore the unknown

1		<b>QEMU:</b> Emulate baseband on unrestricted hardware	✓
2		<b>Ghidra:</b> Integrate Hexagon processors in Ghidra	✓
3		<b>Clade:</b> Reverse-engineer compression algorithm	✓
4		<b>LibAFL:</b> Coverage guided fuzzing with QEMU	✓
5		<b>Firmware Harness:</b> Hook subsystem for analysis & fuzzing	



# Agenda

---

Introduction: Hexagon baseband

From research to tooling

 **Demo: Fuzzing Hexagon**

Opening up baseband security

---

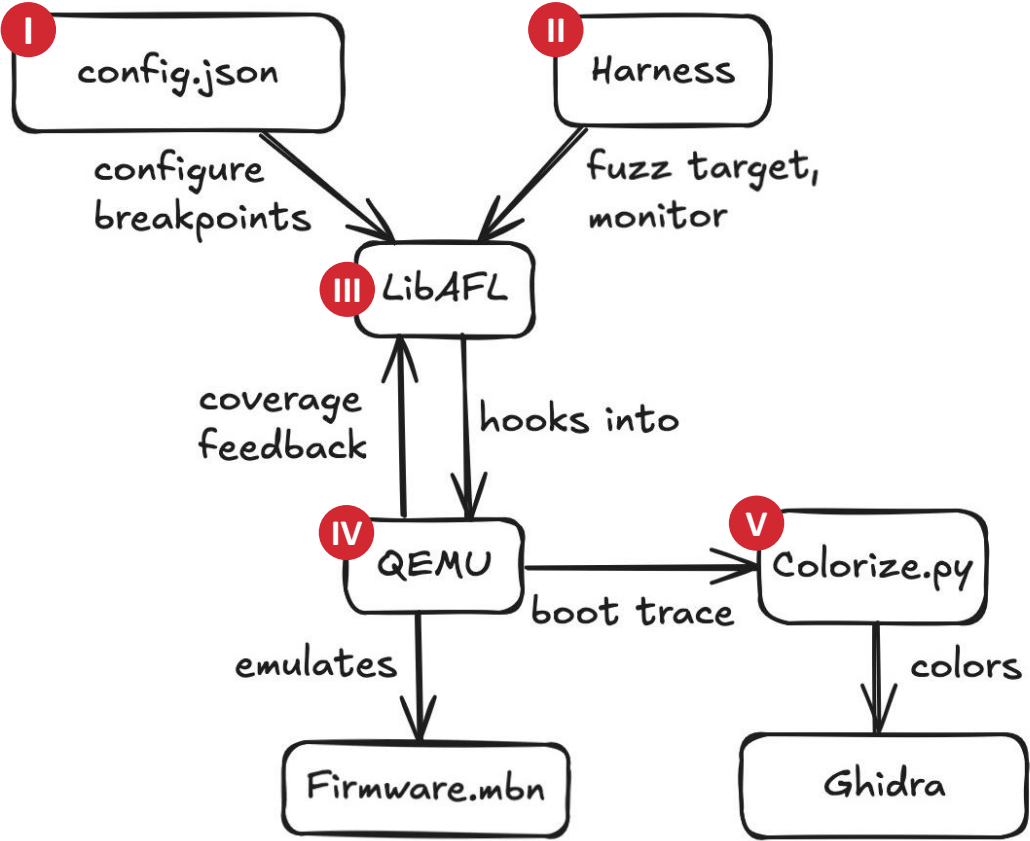
# Fuzzing the Hexagon baseband: we share our tooling

## Architecture Components

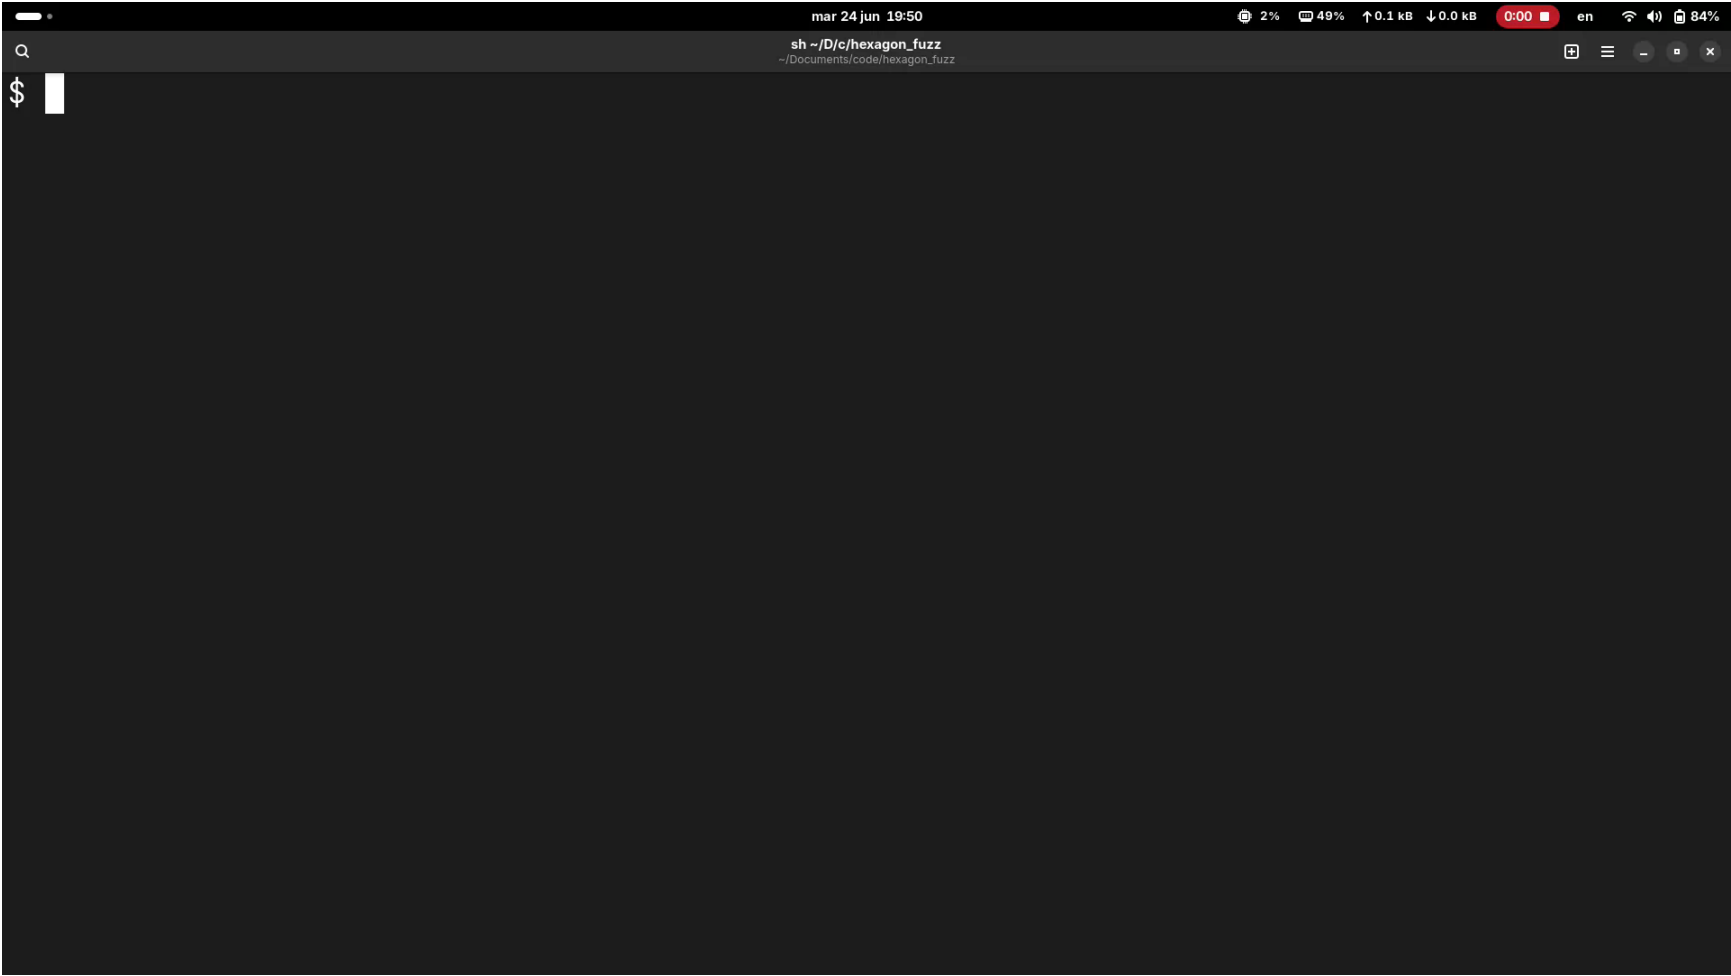
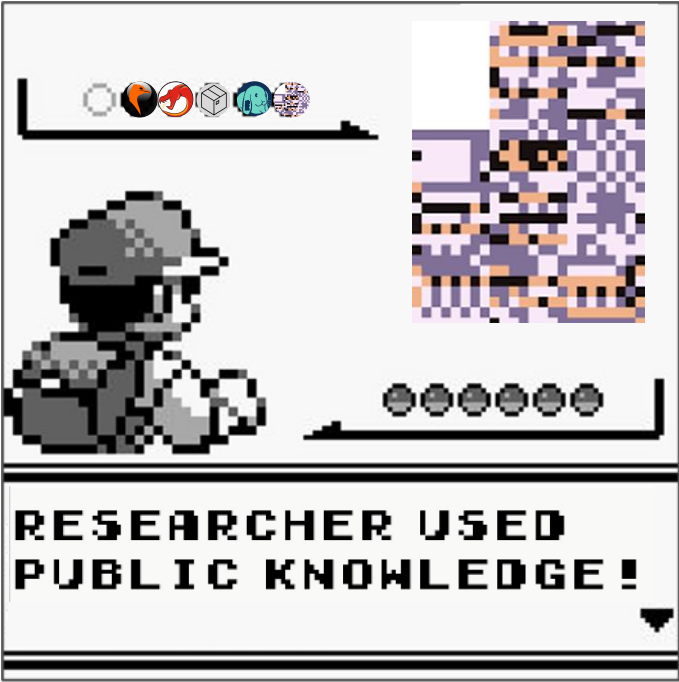
Our Hexagon baseband fuzzer consists of following components:

- I Config.** Set fuzzing parameters and breakpoint addresses
- II Harness.** Rust harness to interact with LibAFL to set breakpoints, fuzz target, monitor process and process coverage
- III LibAFL.** Hook into firmware and introspect
- IV QEMU.** Emulate the Hexagon firmware and get coverage feedback
- V Ghidra scripts.** Improve understanding of decompiled code and progress boot

## Architecture Diagram



# Demo



# Agenda

---

Introduction: Hexagon baseband

From research to tooling

Demo: Fuzzing Hexagon

 **Opening up baseband security**

---

# We refine and integrate our tool to enable successful vulnerability research

Goal	Approach
Targeted fuzzing	<ol style="list-style-type: none"><li>1. Map critical functions</li><li>2. Map exposure: Identify functions that can be triggered from phone/SIM</li><li>3. Optimize fuzzing harness to dynamically adapt to the target function</li></ol>
Optimize CLADE2 decompression	<ol style="list-style-type: none"><li>1. Decompress CLADE2 segments to enable proper task initialization</li><li>2. Ghidra integration: Integrate decompressed output with Ghidra to refine decompiled code</li></ol>
On-device verification	Dynamic testing on iPhone: Trigger crashes, eg. using an SDR
Integrate with existing tools	Integrate Hexagon fuzzer with FirmWire



❤️ Thank you ❤️

- Androm3da
- CUB3D
- domenukk
- Janne
- nlitsme
- nsr
- Mzakocs
- toshipiazza
- ...