

# Reversing & Emulating Samsung's Shannon Baseband

Grant Hernandez & Marius Muench

Tyler Tucker, Hunter Searle, Weidong Zhu  
Patrick Traynor & Kevin Butler

**hardwear.io**

Hardware Security Conference and Training

October 1st, 2020



# # Intros

---

## Grant Hernandez

- Recent Ph.D. graduate (University of Florida)
- Works on Android security and firmware analysis
- Now with Qualcomm's product security team (QPSI)



## Marius Muench

- Ph.D. graduate (EURECOM)
- By now PostDoc @ VUSec
- Builds tooling for better dynamic analysis of embedded systems



# # Outline

---

- Intro to cellular and basebands
- Motivation and tooling overview
- Overview of the hardware platform
- Digging into the firmware
- ShannonOS overview
- Changing behavior with the ModKit
- Fuzzing with AFL
- Conclusions

# # Cellular & Baseband Overview

- Phones communicate wirelessly with cellular networks using the cellular protocols
  - Over 30 years of cellular standards across multiple generations (2G - 5G)
    - GSM (2G), WCDMA/UMTS (3G), LTE (4G), 5G-NR
  - Designed by committee (the 3GPP standards body)
- These protocols define the physical layer and up for cellular operators and devices
- **“Basebands” implement one half of the protocols**
  - Every phone today that has a SIM card has a baseband processor
- Protocol implementation is left up to individual manufacturers (below)
  - Shannon is-a particular implementation of these standards


The Qualcomm logo is displayed in a blue, sans-serif font.The HiSilicon logo features a stylized red and white 'H' symbol above the word 'HISILICON' in a bold, black, sans-serif font.The Mediatek logo is an orange pill-shaped box with the word 'MEDIATEK' in white. Below it are the Intel logo (blue and red) and the Infineon logo (blue and red).The Samsung Exynos Modem 5100 logo is a black square with a small '5G' icon in the top left, the text 'SAMSUNG Exynos Modem' in white, and the number '5100' at the bottom.

# # Getting Insight into Basebands

---

- Basebands are a large non-free component of mobile devices today
  - Can we replace them with FOSS equivalents?
  - How do they work and are they secure?
- Manual reverse engineering is an option, but basebands are huge (doesn't scale)
  - Shannon firmware is a binary blob of ~40MB in size
  - See [https://github.com/grant-h/shannon\\_s5000/](https://github.com/grant-h/shannon_s5000/) -> **At least 4,000,000 SLoC**
- Dynamic analysis would speed up our understanding
  - ✗ On device debugging requires exploitation (debug is fused out)
- **Skip on-target testing and let's build an emulator where we have control!**
  - Lot's of custom hardware to deal with

# # Tooling



- Great for static reverse engineering
- Reasonable collaboration features
- No pricy licensing scheme



- Dynamic reverse engineering framework based on QEMU
- Callbacks during emulation
- Recently released python wrappers

# avatar<sup>2</sup>


- Python framework with wrappers around various emulators (& more)
- Configurable Machine to quickly support targets
- Python based peripheral modeling

A red waveform graphic consisting of a horizontal line with several vertical spikes of varying heights extending above and below it. The spikes above the line are on the left side, and the spikes below the line are on the right side. The horizontal line extends across the width of the slide.

# Shannon Hardware Overview

# # Samsung EXYNOS & Shannon

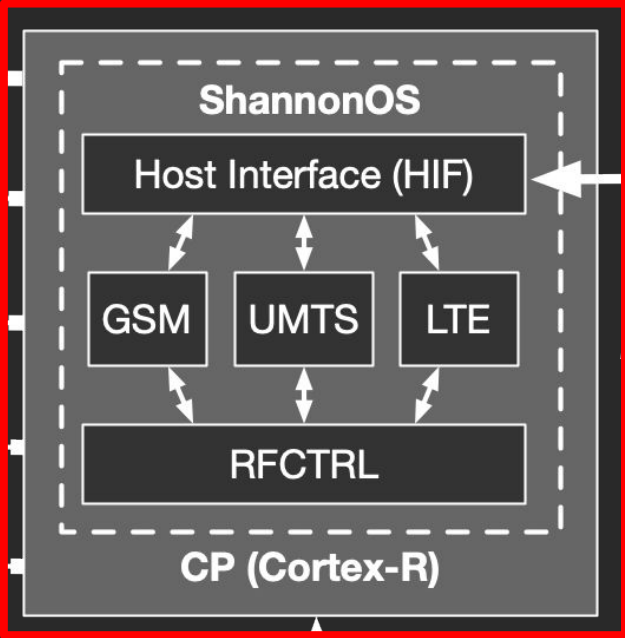
- Samsung U.S. devices prefer Qualcomm SoC's (and therefore their modem)
- Global edition devices instead use EXYNOS - Samsung LSI's in-house SoC
- Shannon co-exists in the SoC floorplan as an IP block
  - The modem is divided into digital and analog
  - The analog component (also known as RF Front-end) is on a separate chip
- We focused on Samsung S10 (SM-G973F) and S7 (SM-G935)





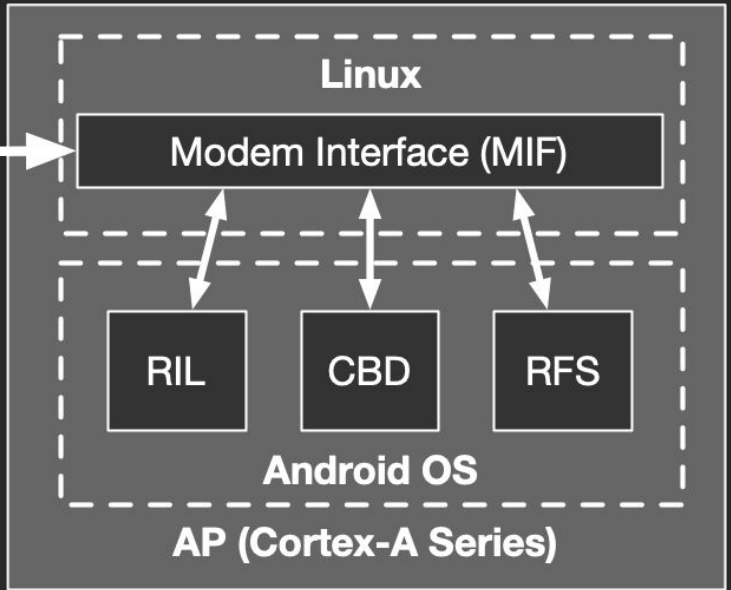
*Peripherals*

- MARCONI DSP
- RFIC
- USIM
- UART
- USI (I2C/SPI)



**SIPC  
MBOX**

*Shared  
Memory  
link (SHM)*



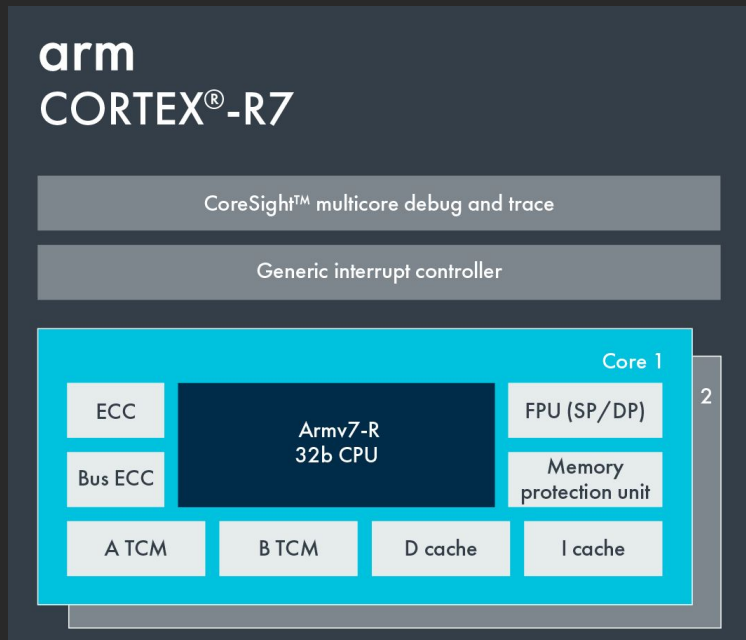
Interprocessor Interrupts

SMC (BP Secure Boot)

**Samsung EXYNOS Platform Services**

# # Shannon CPU Support

- ARM Cortex-R7/R8 (identified from strings and previous work)
- PANDA doesn't support these ARM variants
  - This seems like it could derail things before they even get started
- PANDA does support the R5
- What would be required to get the R7 and R8 working?



Taken from  
<https://developer.arm.com/ip-products/processors/cortex-r/cortex-r7>

```

diff --git a/target/arm/cpu.c b/target/arm/cpu.c
index 04b062c..9f6d05e 100644
--- a/target/arm/cpu.c
+++ b/target/arm/cpu.c
@@ -1132,6 +1132,15 @@ static void cortex_r5_initfn(Object *obj)
     define_arm_cp_regs(cpu, cortexr5_cp_reginfo);
 }

+static void cortex_r7_initfn(Object *obj)
+{
+    ARMCPU *cpu = ARM_CPU(obj);
+
+    cortex_r5_initfn(obj);
+    cpu->pmsav7_dregion = 32;
+}
+
static const ARMCPRegInfo cortexa8_cp_reginfo[] = {
    { .name = "L2LOCKDOWN", .cp = 15, .crn = 9, .crm = 0, .opc1 = 1, .opc2 = 0,
      .access = PL1_RW, .type = ARM_CP_CONST, .resetvalue = 0 },
@@ -1573,6 +1582,7 @@ static const ARMCPUInfo arm_cpus[] = {
    { .name = "cortex-m4",    .initfn = cortex_m4_initfn,
      .class_init = arm_v7m_class_init },
    { .name = "cortex-r5",    .initfn = cortex_r5_initfn },
+   { .name = "cortex-r7",    .initfn = cortex_r7_initfn },
    { .name = "cortex-a7",    .initfn = cortex_a7_initfn },

```

**Not too many  
changes needed!**

# # MPU setup and layout

```
0x00000000-0x00007fff id=00 perm=r-x
0x04000000-0x0401ffff id=01 perm=r-x
0x04800000-0x04803fff id=02 perm=rw-
0x40000000-0x47ffffff id=03 perm=rwx
0x40000000-0x40ffffff id=04 perm=r-x
0x48000000-0x4bffffff id=05 perm=rw-
0x4b800000-0x4bffffff id=06 perm=---
0x4b800000-0x4b8fffff id=07 perm=rw-
0x45600000-0x456fffff id=08 perm=rw-
0x46000000-0x467fffff id=09 perm=rw-
0x46000000-0x460fffff id=10 perm=rw-
0x46800000-0x46ffffff id=11 perm=rw-
0x47000000-0x47ffffff id=12 perm=rw-
0x47300000-0x473fffff id=13 perm=rw-
0x47400000-0x477fffff id=14 perm=r--
0x47780000-0x477fffff id=15 perm=rw-
0x80000000-0xffffffff id=16 perm=rw-
0xe0000000-0xe1ffffff id=17 perm=r--
0x50000000-0x50ffffff id=18 perm=rw-
0x51000000-0x513fffff id=19 perm=rw-
0x51400000-0x515fffff id=20 perm=rw-
0x4a800000-0x4affffff id=21 perm=rwx
0x4b000000-0x4b1fffff id=22 perm=rwx
0x00100000-0x0011ffff id=23 perm=rw-
```

Example MPU configuration


- Well reverse engineered already:
  - Breaking Band (N. Golde, D. Komaromy - REcon'16)
- But, firmware was extracted from RAM dump
  - How do we get there without it?



# The Firmware

# # Extracting modem.bin

- Firmware available from third-party websites like SamMobile



- Previous researchers and papers mentioned that modem.bin is encrypted
  - Apparently, Samsung opted out of this recently  $\backslash\_(\u263a)\_/$
- Firmware can also be extracted from the RADIO partition on a running device (if you have root)

```

00000000: 544f 4300 0000 0000 0000 0000 0000 0000  TOC.....
00000010: 0080 0040 1004 0000 0000 0000 0500 0000  ...@.....

00000020: 424f 4f54 0000 0000 0000 0000 2004 0000  BOOT.....
00000030: 0000 0040 401e 0000 d597 ad57 0100 0000  ...@@.....W....
00000040: 4d41 494e 0000 0000 0000 0000 6022 0000  MAIN.....`"..
00000050: 0000 0140 a079 5402 3fb1 20ef 0200 0000  ...@.yT.?. ....
00000060: 5653 5300 0000 0000 0000 0000 009c 5402  VSS.....T.
00000070: 0000 8047 60f6 5d00 04e5 2907 0300 0000  ...G`.]... )....
00000080: 4e56 0000 0000 0000 0000 0000 0000 0000  NV.....
00000090: 0000 6045 0000 1000 0000 0000 0400 0000  ..`E.....
000000a0: 4f46 4653 4554 0000 0000 0000 00aa 0700  OFFSET.....
000000b0: 0000 0000 0056 0800 0000 0000 0500 0000  .....V.....

```

Entry Name

File Offset

00000000:	544f 4300 0000 0000 0000 0000	0000 0000	TOC.....
00000010:	0080 0040 1004 0000 0000 0000	0500 0000	...@.....

Load Address          Size                  CRC                  Count/Entry ID

00000020:	424f 4f54 0000 0000 0000 0000	2004 0000	BOOT.....
00000030:	0000 0040 401e 0000 d597 ad57	0100 0000	...@@.....W....
00000040:	4d41 494e 0000 0000 0000 0000	6022 0000	MAIN.....`"..
00000050:	0000 0140 a079 5402 3fb1 20ef	0200 0000	...@.yT.?. ....
00000060:	5653 5300 0000 0000 0000 0000	009c 5402	VSS.....T.
00000070:	0000 8047 60f6 5d00 04e5 2907	0300 0000	...G`.]... ).....
00000080:	4e56 0000 0000 0000 0000 0000	0000 0000	NV.....
00000090:	0000 6045 0000 1000 0000 0000	0400 0000	..`E.....
000000a0:	4f46 4653 4554 0000 0000 0000	00aa 0700	OFFSET.....
000000b0:	0000 0000 0056 0800 0000 0000	0500 0000	.....V.....



Entry Name

File Offset

	Load Address	Size	CRC	Count/Entry ID	Entry Name
00000000:	544f 4300	0000 0000	0000 0000	0000 0000	TOC.....
00000010:	0080 0040	1004 0000	0000 0000	0500 0000	...@.....
00000020:	424f 4f54	0000 0000	0000 0000	2004 0000	BOOT.....
00000030:	0000 0040	401e 0000	d597 ad57	0100 0000	...@@...W...
00000040:	4d41 494e	0000 0000	0000 0000	6022 0000	MAIN....."
00000050:	0000 0140	a079 5402	3fb1 20ef	0200 0000	...@.yT?....
00000060:	5653 5300	0000 0000	0000 0000	009c 5402	VSS.....T.
00000070:	0000 8047	60f6 5d00	04e5 2907	0300 0000	...G`.]...).
00000080:	4e56 0000	0000 0000	0000 0000	0000 0000	NV.....
00000090:	0000 6045	0000 1000	0000 0000	0400 0000	..`E.....
000000a0:	4f46 4653	4554 0000	0000 0000	00aa 0700	OFFSET.....
000000b0:	0000 0000	0056 0800	0000 0000	0500 0000	...V.....

00000020:	424f 4f54 0000 0000 0000 0000	2004 0000	BOOT.....	...
00000030:	0000 0040 401e 0000 d597 ad57	0100 0000	...@ @... ..W	....

00000020:	424f 4f54 0000 0000 0000 0000	2004 0000	BOOT.....	...
00000030:	0000 0040 401e 0000 d597 ad57	0100 0000	...@ @... ..W	....
00000040:	4d41 494e 0000 0000 0000 0000	6022 0000	MAIN.....	`" ..
00000050:	0000 0140 a079 5402 3fb1 20ef	0200 0000	...@ .yT.? . .	....

```

00000020: 424f 4f54 0000 0000 0000 0000 2004 0000  BOOT.....
00000030: 0000 0040 401e 0000 d597 ad57 0100 0000  ...@@...W...


```

```

...
00000420: 3c00 00ea d8f1 9fe5 d8f1 9fe5 d8f1 9fe5 <.....
00000430: d8f1 9fe5 feff ffea d4f1 9fe5 d4f1 9fe5 .....
00000440: f81b 0000 fc1b 0000 241c 0000 0000 0000 .....$
00000450: 004c 4254 0000 0000 0000 0000 0000 0000 .LBT.....
00000460: 0000 0000 0000 0000 0000 0000 0000 0000 .....

```

Looks like ARM “always” conditionals!  
 Let’s disassemble from here :)



# # Exception Vectors

```
40000000 3c 00 00 ea    b    boot_RESET
-- Flow Override: CALL_RETURN (CALL_TERMINATOR)

LAB_40000004
40000004 d8 f1 9f e5    ldr    pc=>boot_UDI, [DAT_400001e4]
40000008 d8 f1 9f e5    ldr    pc=>boot_SWI, [DAT_400001e8]
4000000c d8 f1 9f e5    ldr    pc=>boot_PREFETCH, [DAT_400001ec]
40000010 d8 f1 9f e5    ldr    pc=>boot_DATA_ABORT, [DAT_400001f0]

boot_NA
40000014 fe ff ff ea    b    boot_NA
40000018 d4 f1 9f e5    ldr    pc=>boot_IRQ, [DAT_400001f4]
4000001c d4 f1 9f e5    ldr    pc=>boot_FIQ, [DAT_400001f8]
```

BOOT

# # Debugging Strings

- Over 150K debug messages found in modem.bin!
- Wrote Ghidra Python script to auto-type regions starting with DBT:

```
ty = fapi.getDataTypes("TraceEntry")
caddr = fapi.findBytes(caddr, "DBT:")

for caddr in caddr:
    te = fapi.createData(caddr, ty)
    filename = te.getComponent(...)
    message = te.getComponent(...)
    name = "TraceEntry::%s::%s" %
        (filename, message)
    fapi.createLabel(caddr, name, ... )
```

Address	Hex	Type	Value	Comment
41ca521c	44 42 54 3a 22 01 00 00 04 ...	TraceEntry		
41ca521c	44 42 54 3a	uint	3A544244h	magic
41ca5220	22 01 00 00	uint	122h	unk1
41ca5224	04 00 00 00	uint	4h	unk2
41ca5228	98 ba cd fe	uint	FECDBA98h	unk_magic
41ca522c	97 99 ca 41	char *	s_GetReconfirmAbortInt...	message
41ca5230	23 23 00 00	uint	2323h	linenum
41ca5234	e4 9b 4b 40	char *	s_.././././HEDGE/TL1/T...	file

44 42 54 3a 1e 01 00 00 02 00 00 00 98 ba cd fe	DBT:.....
08 96 ca 41 da 1d 00 00 e4 9b 4b 40 44 42 54 3a	...A.....K@DBT:
22 01 00 00 04 00 00 00 98 ba cd fe 3f 96 ca 41	".....?..A
b1 1e 00 00 e4 9b 4b 40 44 42 54 3a 22 01 00 00	.....K@DBT:"...
04 00 00 00 98 ba cd fe 86 96 ca 41 b3 1e 00 00	.....A....
e4 9b 4b 40 44 42 54 3a 19 01 00 00 01 00 00 00	..K@DBT:.....



# # First emulation steps

- We have everything needed for emulation!
  - CPU architecture
  - Memory map
  - Entry point
  - Bonus: debugging strings

```
from avatar2 import *
class ShannonMachine():
    self.avatar = Avatar(arch=ARM_CORTEX_R7)
    self.emu = avatar.add_target(PyPandaTarget, entry_address=0x40000000)

    for entry in self.modem_file.entries:
        self.add_memory_range(entry.load_address, entry.size,
                               file=section_path, name=entry.name)
    self.avatar.init_targets()
    self.avatar.cont()
```

# # Displaying the Boot UART

```
if (boot_mode == DUMP_MODE) {
    boot_unk_common_setup();
    uart_putc('#');
    boot_crash_or_dump();
    boot_unk_crash();
    uart_puts(s_Done_40000550);
    FUN_40000d84(&DAT_00002e00);
}
else {
    if (boot_mode == BOOT_MODE) {
        boot_unk_common_setup();
        uart_putc('#');
    }
}
```

```
self.create_peripheral(UARTPeripheral,
0x84000000, 0x1000, name='boot_uart')
```

```
boot_stage2(nextFnPointer);
goto LAB_400004f0;
}
r0 = s_Xxx!_40000570;
}
else {
    r0 = s_Unknown_4000055c;
}
uart_puts(r0);
}
```

```
class UARTPeripheral(ShannonPeripheral):
    def hw_read(self, offset, size):
        if offset == 0x18:
            return self.status

        return 0

    def hw_write(self, offset, size, value):
        if offset == 0:
            sys.stderr.write(chr(value & 0xff))
            sys.stderr.flush()
        else:
            self.log_write(value, size, "UART")

        return True

    def __init__(self, name, address, size, **kwargs):
        super(UARTPeripheral, self).__init__(name,
            address, size)
        self.status = 0
        self.write_handler[0:size] = self.hw_write
        self.read_handler[0:size] = self.hw_read
```

# # Baseband Initialization: pal\_init1()

---

- A huge monolith function that starts all modem subsystems and tasks
  - Activates malloc heap
  - Loads NV items
  - Starts timers
  - Initializes DSP(s) and other peripherals
  - Starts all tasks
  
- Crucial for setting up the baseband state



# # Bas

- A

- Cr

```
1 void pal_init1(void)
2 {
3     int iVar1;
4     undefined4 uVar2;
5     LogCtx local_4110;
6     byte local_4108 [4];
7     undefined local_4104 [4];
8     undefined auStack16640 [16460];
9     undefined auStack180 [164];
10
11     local_4104[0] = 0;
12     local_4108[0] = 0;
13     memclear(auStack16640,&DAT_0000404c);
14     memclear(auStack180,0x4c);
15     FUN_4162c4e0(0xacc283f);
16     iVar1 = FUN_406d0408();
17     if (iVar1 != 0) {
18         iVar1 = 2;
19     }
20     FUN_406cf8f8(iVar1);
21     FUN_4072f5e2();
22     FUN_40fc8ea2();
23     FUN_406cfbfa();
24     FUN_406cfd34();
25     pal_MemDriverPmd();
26     pal_MemDriverRtk();
27     FUN_406d0ff4();
28     FUN_40fc09e4();
29     DAT_43549238._0_2_ = thunk_FUN_40735b1a();
30     FUN_40fc1e74();
31     pal_RegInit();
32     FUN_406cfae4();
33     gsm_unknown_complicated(&PTR_s_CP_BOOT_419bf1d0,&DAT_402e36e8,0,"RegInit completed");
34     UTIL_Math();
35     uVar2 = FUN_40fc8798("S355AP_DREAM");
36     FUN_408b6d26("S355AP_DREAM",uVar2);
37     DAT_4169abbc = 1;
38     FUN_40fc2e32(&DAT_000017be,auStack16640);
39     DAT_4169abb4 = 1;
40     iVar1 = thunk_FUN_040053b0(auStack16640,auStack180,0x4c);
41     if (iVar1 == 0) {
42         pal_Security_memory();
43         local_4110.te = &TraceEntry::pal_main::Protected_pal_RegInit_Secure_Block_Def;
44         local_4110.flags = &DAT_00001582;
45         log_printf(&local_4110,&DAT_fecdba98);
46     }
47     else {
48         local_4110.te = &TraceEntry::pal_main::Protected_pal_RegInit_Secure_Block_Ini;
49         local_4110.flags = &DAT_00001582;
50         log_printf(&local_4110,&DAT_fecdba98);
51     }
52     FUN_412e9810();
53     pal_Csc_defined();
54     TestConfigurationService_TCS_Init_NVtype_Features();
55 }
```

```
56 FUN_40fb005e();
57 Mti_Common_Define_Initialization();
58 FUN_40fc2e32(0x30,local_4104);
59 FUN_40fc2e32(0x4e,local_4108);
60 FUN_40fc8e74((uint)local_4108[0]);
61 FUN_406ce62e();
62 uart_main(2,&DAT_4310a644,0x400);
63 dbg_Core_Hello();
64 memReleaseNonSecureMemoryRx();
65 FUN_409cc58a();
66 DmTraceMsg_Default();
67 FUN_40735a6a();
68 modem_uart_printf("UpTimer Init: ");
69 FUN_408142ac();
70 modem_uart_printf("Complete.\r\n");
71 thunk_FUN_04001bfc();
72 thunk_FUN_04004584();
73 FUN_406cf82e();
74 do {
75     iVar1 = thunk_FUN_04015f06();
76 } while (iVar1 == 0);
77 hal_modem_InitModem();
78 ps_Controller_lips_2g_peri_client_id();
79 FUN_40fc22be();
80 pal_WatchdogKick();
81 FUN_406ef7fa();
82 FUN_406cf750();
83 thunk_FUN_04001d66();
84 tcs_rfcfg();
85 tcs_rfcfg();
86 local_4110.te = &TraceEntry::pal_main::BandSet_Done;
87 local_4110.flags = &DAT_00001582;
88 log_printf(&local_4110,&DAT_fecdba98);
89 FUN_40711d14();
90 FUN_406ce8bc();
91 hw_Device();
92 hw_Acpm_pal_SmCreateEventGroup();
93 FUN_408aade0();
94 FUN_40718528();
95 hw_ClkFindSysClkCofigInfoIndex(2,1);
96 FUN_40fb6380();
97 DSP_DM_CreateTimers();
98 FUN_406d0b82();
99 FUN_4077748a();
100 FUN_406d0eda();
101 DAT_416b525c = thunk_FUN_04015538(0x4e);
102 thunk_FUN_04003f9a();
103 thunk_FUN_0400400e(DAT_416b525c,&LAB_406cee78+1);
104 thunk_FUN_04003fa4(DAT_416b525c);
105 hw_Power_Information();
106 FUN_406cf9b8();
107 FUN_40fc2e32(0x5a,&DAT_416a0564);
108 gsm_unknown_complicated(&PTR_s_CP_BOOT_419bf1d0,&DAT_402e36f8,0,"pal_init1 completed");
109 FUN_40fc1200(&LAB_406cee34+1);
110 return;
```

# # pal\_init1() – peripherals

- Iteratively emulated peripherals by watching for crash strings or infinite loops
- MMIO monitoring used to see which peripherals needed more modeling
- Simple automated cyclic-bit pattern heuristic
- **Partially emulated: PMIC, CLK, DSP, SOC, SIPC/SHM, TIMER, GIC**

```
class ShannonPeripheral():
    def log_read(self, value, size, offset_name):
        log.log(self.log_level,
                "%s: %0" + str(size*2) + "x <- %s[%s]",
                self.format_address(self.pc), value,
                self.name, offset_name)

    def log_write(self, value, size, offset_name):
        log.log(self.log_level, "%s: %s[%s] <- %0" +
                str(size*2) + "x",
                self.format_address(self.pc), self.name,
                offset_name, value)

    def cyclic_bit(self, pattern=1, cycle_len=32):
        val = (pattern << self.cycle_idx) % 0xffffffff
        self.cycle_idx = (self.cycle_idx + 1) % cycle_len
        return val
```

# # We reached the banner!

---

```
[VARIANT/PALVar/Platform_EV/CHIPSET/S5000AP/device/User/src/pal_main.c] - BandSet: Done.
```

```
=====
```

## DEVELOPMENT PLATFORM

- ARM Emulation Baseboard | Cortex-R7
- Software Build Date : Jul 19 2019 05:03:07
- Software Builder :
- Compiler Version : ARM RVCT 50.6 [Build 422]

Platform Abstraction Layer (PAL) Powered by  
CP Platform Part

```
=====
```

```
[INFO] OS_Create_Event_Group(0x4177b710, ACPM_PMIC_RX_EVENT)
```

```
[VARIANT/PALVar/Platform_EV/HAL/Common/driver/Acpm/src/hw_Acpm.c] -
```

```
[ACPM] Shannon OS [_ShannonOS_3.2_R8_AC5]
```



Shannon0S

# # Tasks

- All metadata about tasks are stored in TaskStructs
- A global array contains the TaskStructs for each task
  - Around 100 tasks present in the baseband!
- Each task has its own stack
  - The stackbase is guarded by the magic constant '\xef\xbe\xad\xde'
- Communicate with each other via Queues
  - Over 290 different ones!

```
0x00: Start of structure
[...]
0x0c: TaskID
0x10: Stackbase
[...]
0x24: Name Pointer
[...]
0x2c: Stacksize
0x30: Main function
0x34: Pre-main function
[...]
0x108: End of structure
```

Acpm	CC	LTEL2LRx	REG_SAP	IMS_CC
Default	MM	LTEL2LTx	AS_SAP	LPP
DM	SM	LTEL2TCM	SMS_SAP	SHM
DM_TX	SS	LTEL2IDLE	CC_SS_SAP	UL2CC
BDA	L1C	LTEL2HTx	SIM_SAP	UL2DL
CIQD	PPP	LTEL2HRx	DBG_SAP	UL2UL
CIQD_FE	GDA	LTE_TLP	DS_REG_SAP	UDATA
Background	CDH	LTE_MTM	DS_AS_SAP	UBMCTask
TpTest	VSUP	LTE_DM	DS_SMS_SAP	ephyFramework
TaskReg	VCG	EDFS	DS_CC_SS_SAP	syncTask
DBGUNS	VCE	URRC	DS_SIM_SAP	recMailTask
DBGCMD	SAEL3	HSPA_CALIBRATION	DS_DBG_SAP	sendMailTask
DBGCMD2	DS_SAE_L3	LLC	MMC	BTL
InitPacketHandler	PDNMGR	GRR	MMC_IF	CLM
PacketHandler	SIM	RLC	SR_IF	CLTCP
PBM	DS_SIM	GMAC	LTE_MMC_GL1	SecuCh
DS_PBM	LteRrm	GLAPD	USAT	SHUB_MSG
ATI	LTE_L1LC	SNDCP	DS_USAT	SSH
MTI	LteRrc	SRM	LTE_TCPIP	CPCOP
SMS	LteRrc_DS	LCSM	LTE_SISO_ASYNC	PROXIMITY

# # Scheduling & Interrupts

---

- Tasks block (yield) on receiving IPC messages
- Can be preempted by higher priority tasks or interrupts (if not disabled)
- Interrupt Flow
  - IRQ Received
  - Find and call IRQ handler OR set task to wakeup
  - Call OS\_irq\_schedule\_task
  - End of interrupt
- Over 260 IRQ handlers!
  - A sobering look at how hard it would be to fully emulate
  - <https://github.com/grant-h/ShannonBaseband/blob/master/firmware/extdata/G973FXXU3ASG8/interrupts.h>

# # Timers

- Monumental component for scheduler
  - Emulation hangs without them
  - Expected to raise IRQ on expiry
- Do NOT follow the specification of R7/R8 timers
- Prototyping in Python, then implementation in PANDA

```
static void shannon_timer_write(void *opaque, hwaddr offset,
                               uint64_t value, unsigned size) {
    [...]
    switch (offset) {
        case 0x00: /* TimerLoad */
            s->limit = value;
            ptimer_set_limit(s->timer, s->limit, 1);
            break;
        case 0x04: /* TimerControl */
            if (s->control & STIMER_CTRL_ENABLE)
                ptimer_stop(s->timer);
            s->control = value;
            freq = s->freq;
            ptimer_set_limit(s->timer, s->limit, 1);
            ptimer_set_freq(s->timer, freq);
            if (s->control & STIMER_CTRL_ENABLE)
                ptimer_run(s->timer, (s->control & STIMER_CTRL_PERIODIC) == 0);
            break;
        case 0x0c: //STOP?
            ptimer_stop(s->timer);
            break;
        case 0x10:
            s->int_level = 0;
            shannon_timer_update(s);
            break;
        case 0x14: /* TIM_IRQ_LEVEL */
            s->int_level = value;
            break;
    }
    [...]
```





# The ModKit

# # ModKit

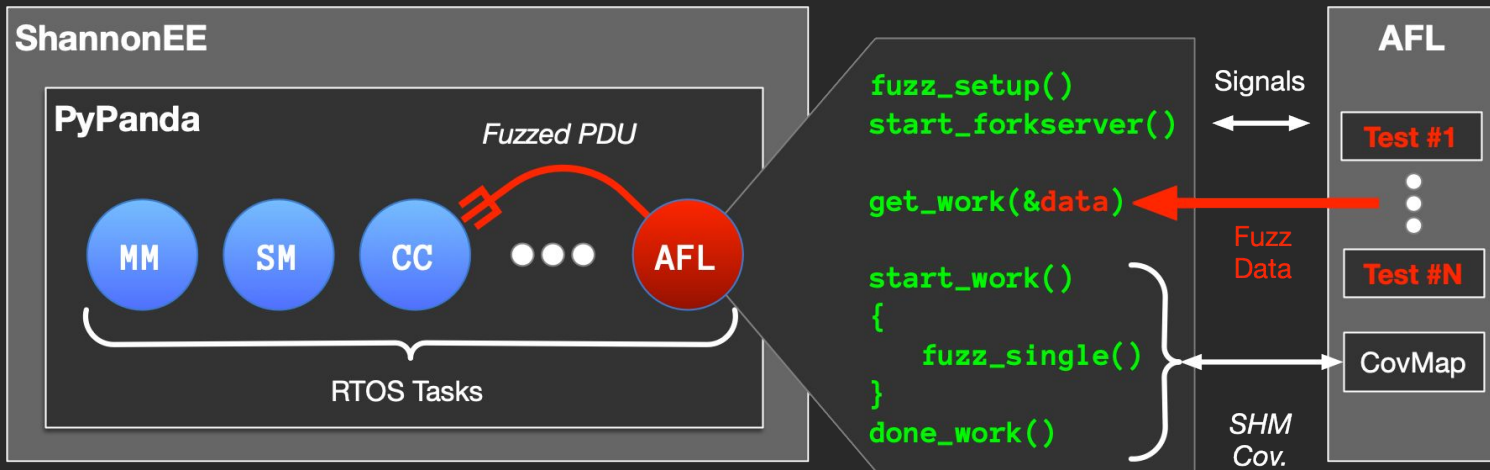
- Inject custom tasks inside the emulated baseband
  - Dynamically loaded during emulation
  - Allows for interaction with the baseband
- C-based compiled modules
  - Dynamic symbol resolution via patterns
  - Creative use of preprocessor macros:

```
DYNAMIC_SYMBOLS = {  
    [...]  
    "OS_fatal_error" : {  
        "pattern" : "70 b5 05 46 ???????? ?? 48  
?? 24",  
    },  
    "pal_MemAlloc" : {  
        "pattern" : "2d e9 f0 4f 0d 00 83 b0  
99 46 92 46 80 46",  
    },  
    [...]  
}
```

```
#define MODKIT_FUNCTION_SYMBOL(ret, name, ...) \  
ret (*__SYMREQ_FUNC_ ## name)(__VA_ARGS__) = (ret (*)(__VA_ARGS__))0xaaaaaaaa; \  
static ret (*name)(__VA_ARGS__) __attribute__((weakref, alias ("__SYMREQ_FUNC_" # name)));
```

# # Using the ModKit for the AFL Task

- Dynamic baseband symbol resolution for a nice FFI
  - Provides access to recovered Shannon symbols
- Mods dynamically linked into RWX page in baseband memory

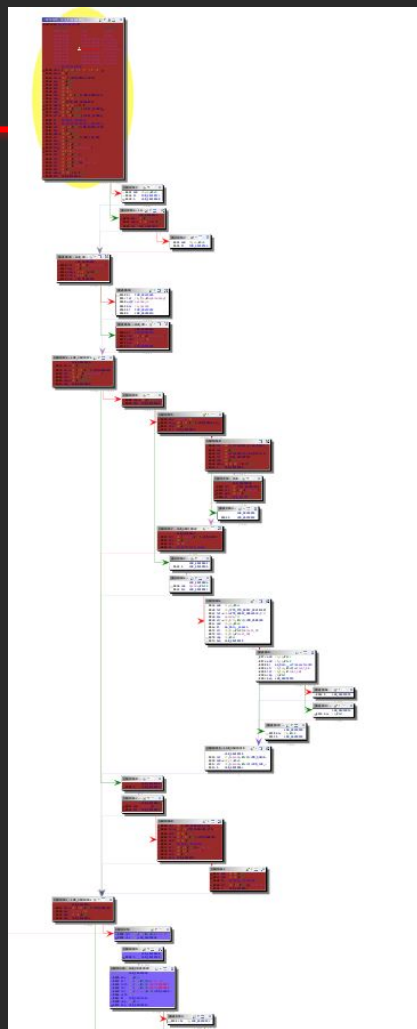


# # Fuzzing with AFL

- Campaigns
  - GSM SM (GPRS)
  - GSM MM
  - GSM CC
- Harnesses < 50 lines of C thanks to our ModKit
- Able to export coverage from emulator to debug fuzzer hangs in the DragonDance plugin

```
afl-fuzz -i inputs -o campaign -- ./shannon_emu.py  
--restore-snapshot g973 --fuzz gsm_mm --fuzz-input @@
```

```
# Get coverage and debug output of `input`  
./shannon_emu.py --restore-snapshot g973 --fuzz-triage  
gsm_mm --fuzz-input input
```



# # Fruits of our labor

---

- Could replicate several n-days
  - E.g., the bug discussed in “A Walk with Shannon” (Amat Cama - Infiltrate)
- Found 0day: SVE-2020-18098
  - using a fuzzing harness on the GSM CC task

## **SVE-2020-18098: Buffer overflow vulnerability in Baseband with abnormal SETUP message**

Severity: Critical

Affected versions: O(8.x), P(9.0), Q(10.0) devices with Exynos chipsets

Reported on: June 19, 2020

Disclosure status: Privately disclosed.

A possible buffer overflow vulnerability in baseband allows arbitrary code execution.

The patch adds the proper validation of the buffer length.



Let's wrap it up!

# # Summary

---

- Baseband firmware is huge and complicated
  - But reversing is still possible!
  
- Emulation is tedious to set up
  - But rewards with detailed insights
  - Allows for additional security testing!
  
- Choose your tooling wisely
  - Ghidra: collaborative RE
  - PANDA: full-system emulation capabilities
  - Avatar<sup>2</sup>: peripheral modeling and orchestration

# # Releases



<https://github.com/grant-h/ShannonBaseband>

- GHIDRA Tools
  - Shannon firmware loader
  - Debug annotation script
  - Auto-renamer
- Shannon reversing details
  - Export of 2017 Shannon image (reversed)
  - On-device log parser (.BTL) and file format info



[https://github.com/grant-h/shannon\\_s5000/](https://github.com/grant-h/shannon_s5000/)

- Code skeleton


<https://github.com/grant-h/ShannonFirmware>

- Firmware repository



<https://github.com/grant-h/ShannonEE>

- **ShannonEE** (will be released at a later time)





# Acknowledgements



®  
Semiconductor  
Research  
Corporation



NSF/SRC CNS-1815883

NWO 628.001.030 ("TROPICS")



AFOSR FA8560-18-S-1201

ONR OTA-N00014-20-1-2205

AFRL/AFOSR-2018-003

# Thanks !



@Digital\_Cold



@nSinusR





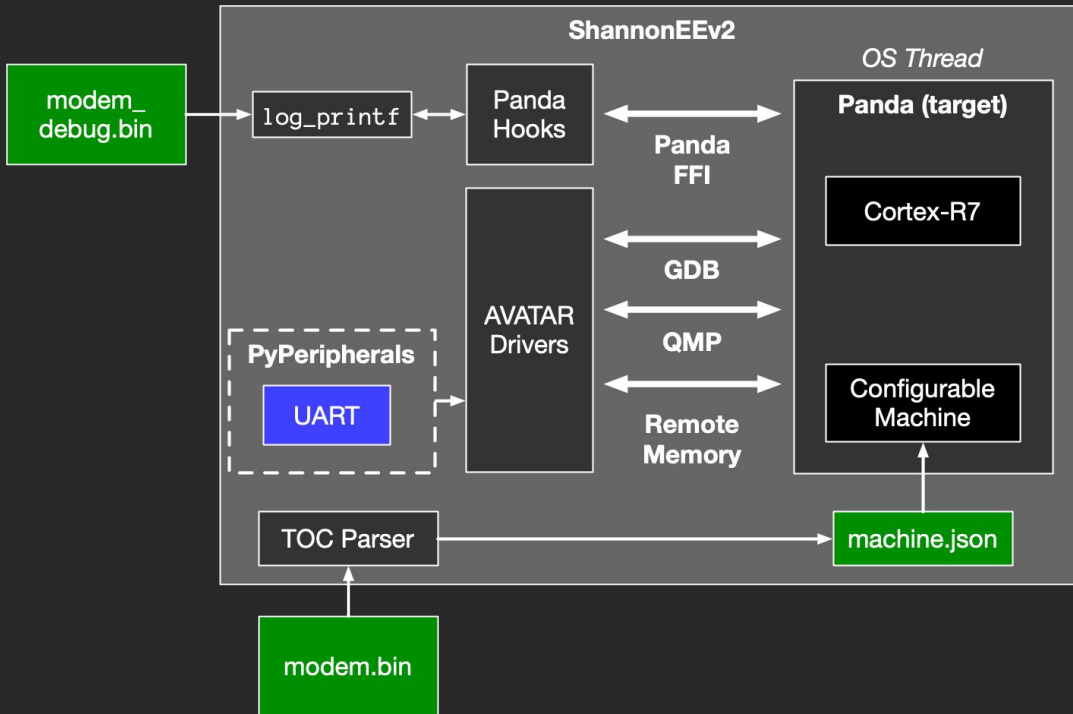
# Backup Slides

A red waveform is plotted on a dark gray background. The waveform starts with a series of positive pulses of varying heights, followed by a series of negative pulses of varying depths. A horizontal red line serves as the zero baseline. The word "Demo" is written in white text above the waveform.

Demo



# # Emulation infrastructure






# On-Device Replication

# # Over-the-air Confirmation

- SDR: bladeRF x40 (\$420)
- Base station: YateBTS (modified with exploit payload)
- Target device: Samsung S10 / S7
- SIM Card: Osmocom USIM (for easy phone connection)
- Thanks to Tyler Tucker for recording the demo and managing the test setup!







<https://youtu.be/Bhs054ma5OQ>



# Boot Mode & IPC

# # Shannon Boot Mode

- Bootloader debug prints using UART
  - Saved to early boot log, can be dumped from kernel as well
- DUMP mode activated during crash dump
- BOOT mode for normal startup
- **Who signals these boot modes?**

```
if (boot_mode == DUMP_MODE) {
    boot_unk_common_setup();
    uart_putc('#');
    boot_crash_or_dump();
    boot_unk_crash();
    uart_puts(s_Done_40000550);
    FUN_40000d84(&DAT_00002e00);
}
else {
    if (boot_mode == BOOT_MODE) {
        boot_unk_common_setup();
        uart_putc('#');
        boot_prepare_mpu_next_ram();
        uart_puts(s_Boot_40000568);
        nextFnPointer = boot_comm_ap();
        if ((void *)0x10000000 < nextFnPointer) {
            boot_stage2(nextFnPointer);
            goto LAB_400004f0;
        }
        r0 = s_XxX!_40000570;
    }
    else {
        r0 = s_Unknown_4000055c;
    }
    uart_puts(r0);
}
}
```

# # A side-quest to the Samsung Kernel

---

- Examined the Samsung S10's kernel sources to better understand the modem's boot
  - Samsung S10 Kernel - <https://github.com/grant-h/SM-G973F-Kernel>
- `drivers/misc/modem_v1`
  - `modem_io_device.c` - dev node ioctl handler (for booting), read/write for commands
  - `link_device_memory.h` - **Contains structs on shared memory regions**
- Cellular Boot Daemon (CBD) communicates with `/dev/umts_boot0` (`modem_v1`)
  - `IOCTL_MODEM_RESET` - Reset and await new boot
  - `IOCTL_SECURITY_REQUEST` - Establish modem bootloader secure boot
  - `IOCTL_MODEM_ON`, `IOCTL_MODEM_BOOT_ON` - Start boot process
  - `IOCTL_MODEM_DL_START` - Download code to modem memory
  - `IOCTL_MODEM_BOOT_OFF` - Finalizes boot process

# # Cross-core Comms

- Modem driver in Linux maps shared memory region between CP and AP for DMA-based IPC (ring buffers)
- Commands are queued in a full-duplex fashion
- Commands come from Samsung RIL

```
#define MEM_IPC_MAGIC      0xAA
#define MEM_CRASH_MAGIC   0xDEADDEAD
#define MEM_BOOT_MAGIC    0x424F4F54
#define MEM_DUMP_MAGIC    0x44554D50
```

link\_device\_memory.h

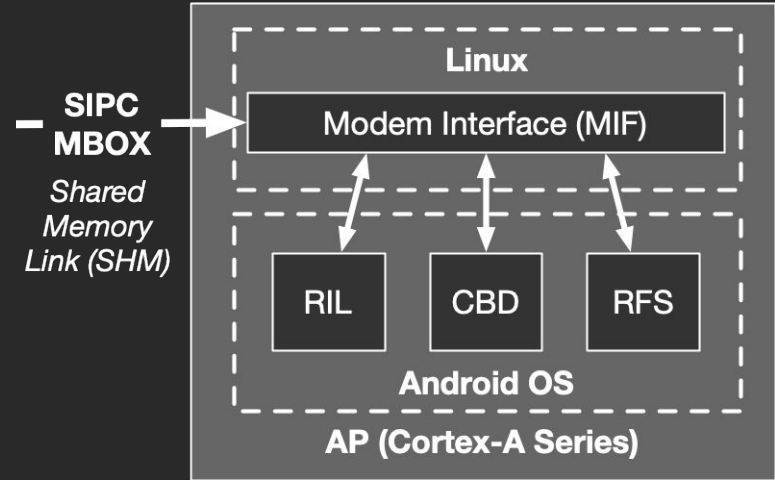
```
struct __packed shmem_4mb_phys_map {
    u32 magic;
    u32 access;

    u32 fmt_tx_head, fmt_tx_tail;
    u32 fmt_rx_head, fmt_rx_tail;

    u32 raw_tx_head, raw_tx_tail;
    u32 raw_rx_head, raw_rx_tail;
    ...
    char fmt_tx_buff[SHM_4M_FMT_TX_BUFF_SZ];
    char fmt_rx_buff[SHM_4M_FMT_RX_BUFF_SZ];
    ...
    char raw_tx_buff[SHM_4M_RAW_TX_BUFF_SZ];
    char raw_rx_buff[SHM_4M_RAW_RX_BUFF_SZ];
};
```

# # Samsung IPC (SIPC) Clients

- SIPC flows bidirectionally over this SHM link
- Cellular Boot Daemon (CBD)
  - Maps modem from local Android partition into baseband memory
  - Sends configuration / debugging commands
- Remote File System (RFS)
  - Provides open/read/write to Android filesystem (a.k.a. EFS) from modem
  - Stores modem config (PLMN, bandset, etc.)
- Radio Interface Layer (RIL)
  - Provides a HAL for Android telephony services
  - Significant configuration API for modem
- FOSS version of RIL -  
<https://redmine.replicant.us/projects/replicant/wiki/Libsamsunq-ril>



Replicant

A decorative graphic consisting of a horizontal red line with several vertical red lines of varying heights extending above and below it, each ending in a small red dot. The graphic is positioned on the left side of the slide, partially overlapping the title text.

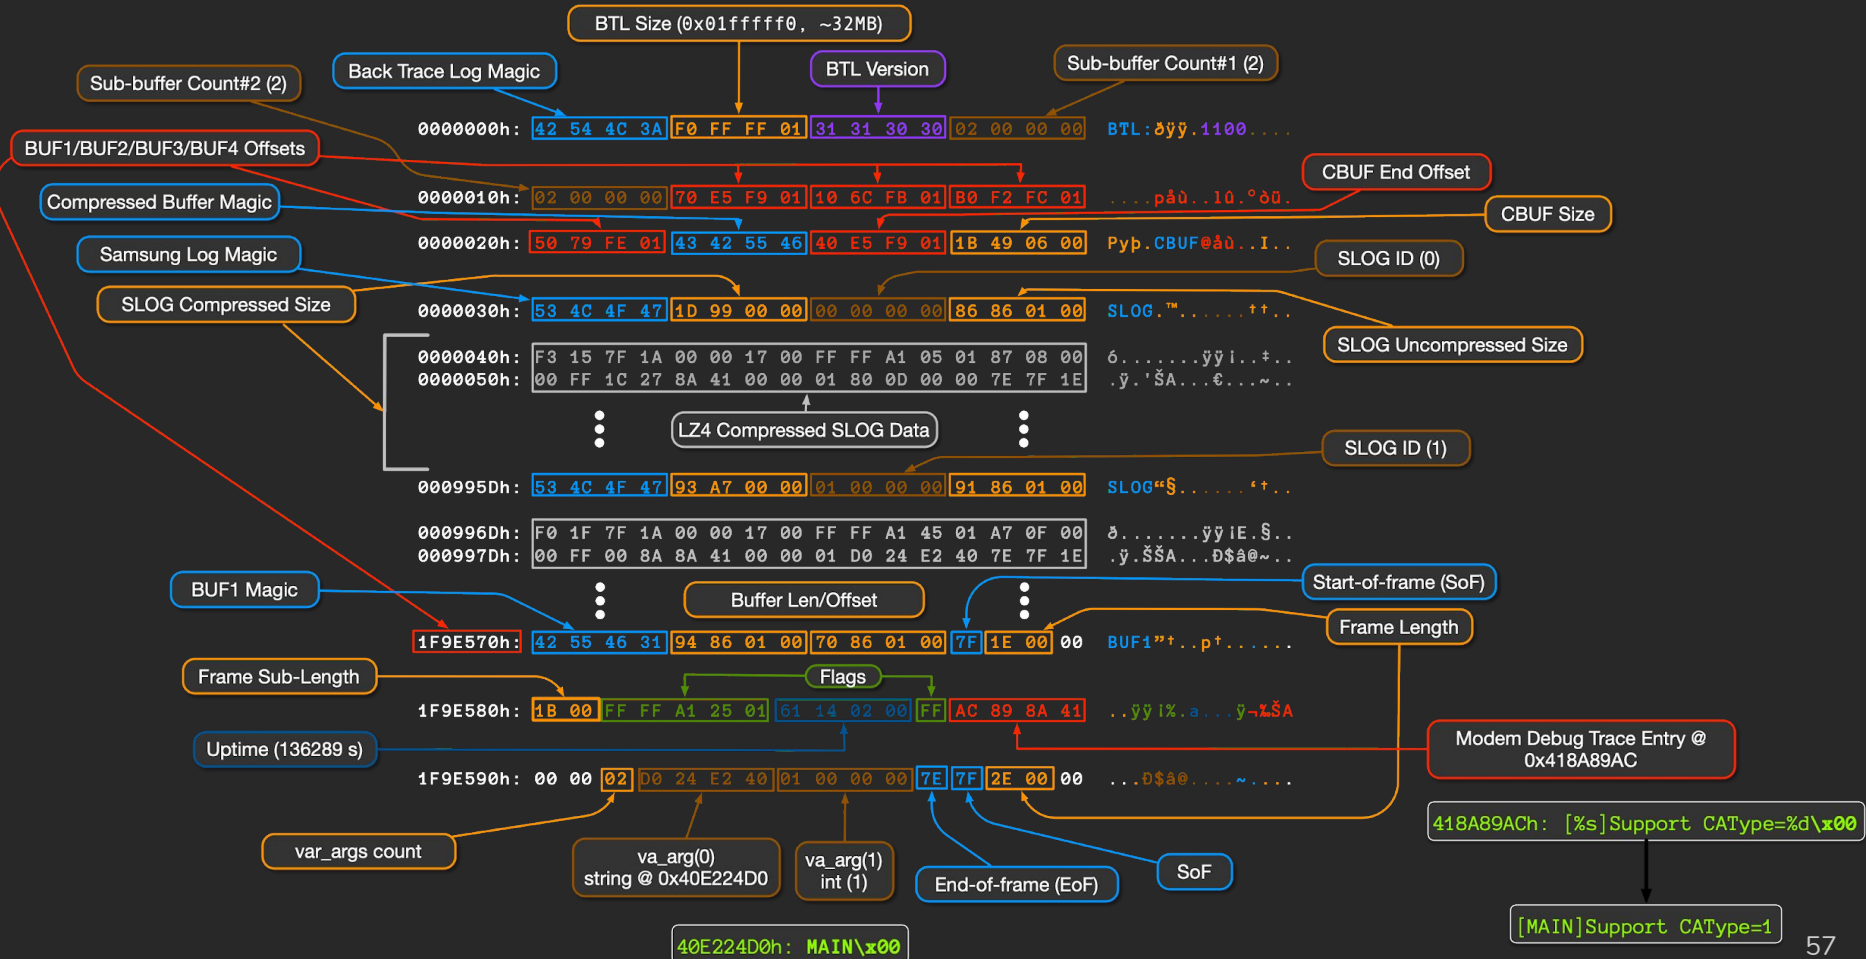
# Backtrace Logging Format

# # Modem Backtrace Logging (BTL)

---

- Dumped on crash, but can be manually extracted (.BTL files)
- Reverse engineered BTL task
- LZ4 encoded log ring buffer in baseband memory
  - Log entries are pointers to format strings in modem\_debug.bin
  - Varargs are pointers or literals, depending on format specifier
- Not 100% reverse engineered, but good enough to compare ShannonEE with real modem logging







# OS Internals: Queues

# # Queues

- IPC is realized via Queues
- Queue messages allocated by sender and freed by called
- Over 290 queue destinations!
  - [https://github.com/grant-h/ShannonB aseband/blob/master/firmware/extdata/G973FXXU3ASG8/pal\\_queues.h](https://github.com/grant-h/ShannonB aseband/blob/master/firmware/extdata/G973FXXU3ASG8/pal_queues.h)
- Message payload variable depending on msgID

```
struct qitem_header {
    union {
        struct {
            uint16_t src; // source queue
            uint16_t dst; // destination queue
        };
        uint32_t dir;
    };
    uint16_t size; // size of the payload
    uint16_t msgId; // [msgGroup:8][msgNumber:8]
} PACKED;
```



5G

# # What about 5G?

---

- Shannon also present in Samsung's Next-gen 5G Modem (51XX)
  - Present for instance in Galaxy S20
- Radical change or processor profile!
  - Cortex-R instead of Cortex-A
  - Full-blown MMU
- Lot's of new code for handling 5G:
  - At least 500k lines of new code
  - But not everything changed!

