Verified functional programming of an IoT operating system's bootloader

Shenghao Yuan, Jean-Pierre Talpin

Verified functional programming of an IoT operating system's bootloader

SHENGHAO YUAN, INRIA, IRISA, France
JEAN-PIERRE TALPIN, INRIA, IRISA, France

The fault of one device on a grid may incur severe economical or physical damages. Among the many critical components in such IoT devices, the operating system's bootloader comes first to initiate the trusted function of the device on the network. However, a bootloader uses hardware-dependent features that make its functional correctness proof difficult. This paper uses verified programming to automate the verification of both the C libraries and assembly boot-sequence of such a, real-world, bootloader in an operating system for ARM-based IoT devices: RIoT. We first define the ARM ISA specification, semantics and properties in F★ to model its critical assembly code boot sequence. We then use Low★, a DSL rendering a C-like memory model in F★, to implement the complete bootloader library and verify its functional correctness and memory safety. Other than fixing potential faults and vulnerabilities in the source C and ASM bootloader, our evaluation provides an optimized and formally documented code structure, a reasonable specification/implementation ratio, a high degree of proof automation and an equally efficient generated code.

CCS Concepts: • **Software and its engineering** → *Embedded software*; *Correctness*; **Software verification**; *Functional languages*; *Domain specific languages*; **Formal software verification**; • **Security and privacy** → *Logic and verification.*

Additional Key Words and Phrases: Verified programming, IoT kernel, boot loader, case study

ACM Reference Format:
Shenghao Yuan and Jean-Pierre Talpin. 2021. Verified functional programming of an IoT operating system's bootloader. In *MEMOCODE '21: ACM-IEEE International Conference on Formal Methods and Models for System Design, Nov 20–22, 2021, Beijing, China.* ACM, New York, NY, USA, 24 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

Among the critical components in the operating system stack of an embedded device, the first whose reliability is put to trial is the bootloader to check, load and execute the image of the operating system or unikernel. Failure to boot renders an embedded device useless, leaving its possibly networked and mission-critical function unattended until maintenance. Well-known mechanisms, such as Trusted Boot [2] or Verified Boot[12], mainly focus on the validation of loaded images. However, a bootloader is itself a small but complex piece of software, tightly coupled to its hardware platform, making it quite difficult to be verified, especially when hand-written in an unsafe language like C and/or assembly code.

Program verification techniques have become popular to ensure the correctness of programs written in unsafe languages like C. Deductive programming tools would for instance allow one to verify a bootloader at its source C code-level, but probably require the use of another tool to check its necessary assembly boot sequence correct. As such, VCC[7], Verifast[16] and refinedC[28] allow to verify C or Java programs annotated with pre- and post-conditions. These conditions however introduce a heterogeneous syntax of annotations that rapidly scales in proportion to the source code to verify. For instance, a 14 lines long C program may require about 20 lines of annotations in refinedC [28, Sec. 2.2]. Another choice to formally verify a bootloader is to homogeneously express the implementation and verification conditions in a proof assistant, e.g. SABLE[8] in Isabelle/HOL[21] or the the first-stage bootloader of [29] in Coq[5]. Although these specifications may automatically be generated by using VST[1], verification conditions still

2021. Manuscript submitted to ACM

require to undergo a time-consuming process of mechanized proof. In fact, the first-stage bootloader only proves part of a Sanctum-like bootloader functionally correct.

In this paper, we adopt a verified programming methodology to implement and verify a real-world bootloader. Our approach gains both proof automation while maintaining homogeneous specifications and implementations in the F⋆[30] programming environment. We implement a verified *riotboot*[26], the bootloader of a friendly Operating System for IoT devices, RIoT[4], together with the ARM Cortex-M architecture of its target platform(s).

The source code of *riotboot* is a mixture of C and assembly code. The C code can be modeled in Low⋆[22], a low-level subset of F⋆ that renders this C-like memory model and enjoys a translation to C that doesn't require a runtime library using the KreMLin compiler[32]. The hardware-dependent part of *riotboot*, written in assembly code, may however not be modeled using existing F⋆ libraries. The most related project, VALE [6, 11], only supports the verification of x84/x64 architectures in F⋆. Hence, this paper presents the following contributions:

- *ARM-F⋆*: We define a complete F⋆ model of an instruction set available in most ARM platforms, including modes, conditionals and suffixes. Our model includes 1/ the formal syntax and operational semantics of the chosen ISA in F⋆, 2/ a formal specification of its critical properties (as explained in the ARM assembly user guide), and 3/ a series of lemmas in F⋆ to automate verification of programs.
- *Verified riotboot*: We use Low⋆ and our library ARM-F⋆ to model *riotboot* and verify its functional correctness and memory safety in the F⋆ environment. Our workflow contains: 1/ a model of *riotboot*'s C modules in Low⋆ and of its assembly code in ARM-F⋆, 2/ a functional correctness proof of *riotboot* in the F⋆ environment, and 3/ extracted C and assembly code from our F⋆ model.
- *Evaluation*: We show potential faults and vulnerabilities found with our F⋆/Low⋆ model, compare it with existing formally verified bootloaders by highlighting an optimized amount of required specification and, foremost, the high degree of proof automation gained.

As a result, we benefit from bare-metal executable code verified against all critical requirements at minimal specification cost and development time (i.e. one month). Our workflow provides a principled type-driven approach allowing IoT developers to specify and verify system- and application-specific properties in a way that maximizes proof automation while facilitating specification, that could further be minimized by using static analysis [10], that could further be extended to deal with physical and hardware constraints [31].

The rest of the article is organized as follows. Section 2 gives a short introduction to verified programming in F⋆. Section 3 briefly introduces the modules of Riotboot. Section 4 formalizes the ARM assembly documentation [20] in F⋆. Section 5 specifies *riotboot* in F⋆/Low⋆, verifies its functional correctness, and evaluates our model. Sections 6-7 conclude by discussing related and future works.

2 A BRIEF OVERVIEW OF F⋆ AND LOW⋆

F⋆ is a general-purpose functional programming language that, in the spirit of Liquid Haskell or Agda, is meant at verifying programs. In this aim, F⋆ supports a dependent-type system allowing to express type refinements of both pure and imperative functions with logical properties pertaining to their value domain, pre- and post-conditions. For instance, the type of the `tot`(al) function `abs` accepts any integer and returns its absolute value: `v:int -> Tot v:int{v>=0}` is its type. The `st`(ateful) function `get`, reading the value `v` of a reference `r` in the memory heap `h` has type `r:ref a -> ST v:a`, pre-condition `requires fun h -> (contains h r)`, saying that `h` musts contain `r`, and post-condition `ensures fun h v _ -> v = (sel h r)`, saying that the returned value `v` is exactly that of the sel(ected) heap location. Low★ can be seen as a domain-specific language embedded in F★ whose purpose is to render the computational model of imperative system languages like C. As a result, Low★ enjoys the powerful specification and proof capabilities of F★ while can generate verified C code readily usable without resource-hungry runtime library.

Subroutines used during the image validation process are typical Low★ programs. For instance, considering function rb_hdr_t2uint16_t in details, it marshals header struct(ure) into an uint16_t buffer for input to the fletcher32 image validation algorithm. It consists of a type and logical specification with a val declaration, and an implementation with a let declaration. It takes two arguments s and d whose types are specified between arrows: rb_hdr_t and d:B.buffer UInt16.t. The first one is just the data-type of a *riotboot* header data-structure. The second is a Low★ buffer (i.e. B.buffer) containing 16bits unsigned integers with type refinement behind brackets {} saying that the length of buffer B.length d should be larger than the value of the constant UInt16.v offset_chksum.

```
val rb_hdr_t2uint16_t: rb_hdr_t -> d:B.buffer UInt16.t{B.length d > UInt16.v offset_chksum} -> ST unit
  (requires (fun h0 -> B.live h0 d)) (ensures (fun h0 v h1 -> (M.modifies (M.loc_buffer d) h0 h1) /\ B.live h1 d))
let rb_hdr_t2uint16_t s d =
  d.(0ul) <- uint32_to_uint16(s.magic_number); d.(1ul) <- uint32_to_uint16(s.magic_number >>^ 16ul); .../...;
  d.(5ul) <- uint32_to_uint16(s.start_addr >>^ 16ul); ()
```

The function body behind the let sequentially marshals the raw header data from s into the buffer d by performing a series of assignments and shifts. The function returns nothing, but has side-effects: it populates buffer d. Its assumptions and guarantees are specified by predicates in the monad ST. The pre-condition is defined by a function with the initial memory state h0 as a parameter. By stating B.live h0 d, it requires d to be a live memory area in h0. The post-condition is stated as a function taking the result v and initial and final memory states h0 and h1 as parameters. It says that the function returns a modified and live memory buffer d. To obtain this guarantee, the effect of each statement in the sequence is collected from a sentence to the next one by using monadic binding. This propagated information is then checked against the declared post-condition.

3 RIOTBOOT OVERVIEW

The bootloader of RIoT: *riotboot*, expects flash memory to be supplied and formatted in slots to host operating system images. The core of *riotboot* consists of two modules: *choose_image* and *cpu_jump_to_image*.

```
1 void kernel_init(void){
2   uint32_t version = 0; int slot = -1;
3   for (unsigned i = 0; i < riotboot_slot_numof; i++){ //choose_image beginning
4     const riotboot_hdr_t *riot_hdr = riotboot_slot_get_hdr(i);
5     if (riotboot_slot_validate(i)) { continue; }
6     if (riot_hdr->start_addr != riotboot_slot_get_image_startaddr(i)) { continue; }
7     if (slot == -1 || riot_hdr->version > version) { version = riot_hdr->version; slot = i;} } //choose_image ending
8   if (slot != -1) { riotboot_slot_jump(slot); } //cpu_jump_to_image
9   while (1) {} }
```

Function choose_image consists of a for-loop that chooses a suitable image from a list of slots in flash memory. It first selects an image header in that list (lines 3-4) and validates its header (line 5) using the *fletcher32* checksum algorithm (below). If no valid image is present, *kernel_init* falls into an infinite loop (line 9) whose behavior may actually be reduced to *nop* by an optimizing compiler. If several valid images are present in the list, it chooses that with the latest version number (line 7).

Function `cpu_jump_to_image` is written in Cortex-M assembly code and performs a "long jump" to execute the imaged system. Line 2 sets the stack pointer (MSP) to the image address. Line 3 skips the image header. Lines 4-5 set the destination address and force the processor state to `Thumb` mode. Finally, line 6 branches execution at the destination. Such operations cannot be performed in a system language: a tempting `(*image_addr)()` in C would result in sharing the memory space of the bootloader with the image.

```
1 static inline void cpu_jump_to_image(uint32_t image_addr) {
2   __set_MSP(*(uint32_t*) image_addr);
3   image_addr += 4;
4   uint32_t destination = *(uint32_t*) image_addr;
5   destination |= 0x1;
6   __asm("BX %0" :: "r" (destination)); }
```

Our workflow starts with the definition of the ARM ISA in F★: its syntax, operational semantics and properties Then, the *choose_image* function is modelled in F★/Low★ and the *cpu_jump_to_image* function is expressed in ARM-F★. The model's functional correctness is automatically verified by F★ using the Z3 SMT-solver, and the verified model is used to extract executable C code by the Kremlin compiler and ASM assembly code using the ARM-F★ print module. The synthesis of extraction result finally produces a verified `riotboot`.

4 FORMALIZING THE ARM ISA IN F★

This section selects a general ARM instruction set, defines its syntax and semantics and proves its properties derived from the ARM ASM user guide to provide useful lemmas.

4.1 Syntax

The syntax of the ARM assembly language is shown in Fig.1. It comprises of three kinds of instructions[1]:

- Twelve arithmetic instructions from the 'Add with Carry' **adc** to the 'Store' instruction **str**.
- The logical instructions **mov** and four bitwise operations: conjunction **and**, disjunction **orr** and **orn**, exclusion **eor**.
- The shift instructions: Arithmetic Shift Right **asr**, Logical Shift Left **lsl**, Logical Shift Right **lsr** and Rotate Right **ror**.

$$
\begin{array}{rll}
ci & ::= & \{cond\}\ i \mid \{s\}\ i \mid \{s\}\ \{cond\}\ i \\
i & ::= & \textbf{adc}\ r_d\ r_n\ op_2 \mid \textbf{add}\ r_d\ r_n\ op_2 \mid \textbf{bx}\ r_d \mid \textbf{cmn}\ r_n\ op_2 \mid \textbf{cmp}\ r_n\ op_2 \mid \textbf{ldr}\ r_d\ r_n\ o \\
 & & \mid \textbf{mul}\ r_d\ r_n\ r_m \mid \textbf{neg}\ r_d\ r_n \mid \textbf{nop} \mid \textbf{sub}\ r_d\ r_n\ op_2 \mid \textbf{str}\ r_d\ r_n\ o \\
 & & \mid \textbf{mov}\ r_d\ op_2 \mid \textbf{and}\ r_d\ r_n\ op_2 \mid \textbf{eor}\ r_d\ r_n\ op_2 \mid \textbf{orn}\ r_d\ r_n\ op_2 \mid \textbf{orr}\ r_d\ r_n\ op_2 \\
 & & \mid \textbf{asr}\ r_d\ r_n\ r_s \mid \textbf{lsl}\ r_d\ r_n\ r_s \mid \textbf{lsr}\ r_d\ r_n\ r_s \mid \textbf{ror}\ r_d\ r_n\ r_s \\
cond & ::= & EQ \mid NE \mid CS \mid CC \mid MI \mid PL \mid VS \mid VC \mid LT \mid LE \mid GT \mid GE \mid AL \\
r & ::= & r_0 \mid r_1 \mid \ldots \mid r_{12} \mid sp \mid lr \mid pc \\
op_2 & ::= & c \mid r \mid r\ sop \\
sop & ::= & ASRshift\ sh_2 \mid LSLshift\ sh_1 \mid LSRshift\ sh_2 \mid RORshift\ sh_1 \\
sh_n & \in & [1, 30+n],\ o, c\ \in\ Int32,\ s\ \in\ String
\end{array}
$$

Fig. 1. Core syntax of the ARM assembly language

Conditional instructions execute when a condition flag is set by a prior instruction. A compound conditional instruction can be built by composing a simple one with a condition or a suffix or both condition and suffix.

Conditional code cond defines the condition that must be met for an instruction to execute. It can be equal (`EQ`), unequal (`NE`), negative (`MI`), positive or zero (`PL`), etc.

[1] The classification follows the same flags update principle: *str* and *adc* use the same function to update flags, while *mov* and bitwise instructions adopt another. Please refer to the ARM ASM user guide for details.

Optional suffix, if specified, sets the condition flag after the instruction is executed. Otherwise, the instruction has no effect on the condition flags.

General purpose registers are r_0-r_{12} and three special registers:the stack pointer register (sp), the link register(lr) and the program counter register (pc). The Application Program Status Register (APSR) holds the program status flags.

Operands and shifts are found as second operand op_2 of many ARM arithmetic and logical instructions. They can be a constant c, a register r or a register with a shift value.

4.2 Machine state

In the spirit of the VALE project [6, 11], we represent the machine state as a record (named arm_state) consisting of:

- *mem: addr -> int32*, as a map from physical addresses to bytes,
- *regs: reg -> int32*, as functions mapping register names to values,
- *flags: flag*, as the negative (N), zero (Z), carry (C), and overflow (V) condition flags of the APSR register.
- *isa_mode: mode*, three kinds of ISA modes: ARM, Thumb16 and Thumb32,
- *ok: bool*, a Boolean field ok representing the processor state.

A valid state (ok = true) indicates that the machine has safely executed until the current state, e.g. no segmentation fault occurred. While an invalid memory access or update would make the machine crash (ok = false).

4.3 Operational Semantics

We now define the operational semantics of key instructions from Fig. 1 in F⋆ by employing the methodology of VALE. The complete definition of the operational semantics of ARM instructions can be found in a GitLab repository[34]. Firstly some auxiliary functions are defined to check the validity of ARM instructions (as per the reference manual [20]). Then the rules of the operational semantics are defined, as shown in Figure 2.

Valid functions. Most ARM instructions have constraints regarding the usage of registers and operands, e.g., the destination register of most instructions can not be the program counter. This paper defines validity functions to constrain the parameters of each instruction. In F⋆ and VALE, they can be modeled as predicates and enforced as pre-conditions to using the instructions.

Semantics. We first introduce the key operational semantics rules of the simple ARM instructions used for the bootloader, i.e., **add**, **bx**, **mov**, **orr**. Then, we introduce a special rule: the *memory_unsafe* rule. Fig 2 exemplifies these rules for selected instructions. The predicate *valid(r)* defines the validity condition of the related argument r.

All rules in Fig 2 satisfy two premises: the memory flag *ok* is true and all operands are valid. Then,

- (add) adds the values in r_n and op_2, stores the result in r_d and updates the pc register.
- (bx*) causes a branch to the address stored in r_d and switches the instruction set: (bx1) If bit(0) is 0, then the processor changes to ARM state; (bx2) if bit(0) is 1, the processor remains in Thumb state.
- (mov) copies the value of op_2 into r_d and updates the pc register.
- (orr) performs a bit-wise OR operation on r_n and op_2, stores the result in r_d and updates the pc.

If an operand is invalid, the memory flag is cleared (*ok=false*). If the memory flag is false, the processor aborts.

$$\frac{\neg st.ok \vee \neg valid(op)}{(ins, st) \longrightarrow \textbf{abort}} \quad (memory\_unsafe)$$

$$\frac{st.ok \wedge valid(rd) \wedge valid(rn) \wedge valid(op_2)}{(ADD\ rd\ rn\ op_2, st) \longrightarrow st[rd \mapsto rn + op_2, pc \mapsto pc + 1]}(add)$$
$$\frac{st.ok \wedge rd.bit(0) = 0 \wedge valid(rd)}{(BX\ rd, st) \longrightarrow st[st.isa\_mode \mapsto Thumb_{16}, pc \mapsto rd]}\ (bx1)$$
$$\frac{st.ok \wedge rd.bit(0) = 1 \wedge valid(rd)}{(BX\ rd, st) \longrightarrow st[st.isa\_mode \mapsto ARM, pc \mapsto rd]}\ (bx2)$$
$$\frac{st.ok \wedge valid(rd) \wedge valid(op_2)}{(MOV\ rd\ op_2, st) \longrightarrow st[rd \mapsto op2, pc \mapsto pc + 1]}\ (mov)$$
$$\frac{st.ok \wedge valid(rd) \wedge valid(rn) \wedge valid(op_2)}{(ORR\ rd\ rn\ op_2, st) \longrightarrow st[rd \mapsto rn \mid op_2, pc \mapsto pc + 1]}(orr)$$
$$\cdots$$

Fig. 2. Semantics of the simple ARM instruction set

4.3.1 Conditional Instructions. All simple ARM instructions can be executed conditionally by relying on the condition code *c* of the instruction and the value of the condition flags in the APSR, i.e. the memory state *st*. The condition function *cond(c,st)* states the condition *c* that must be met for an instruction to execute in the state *st*, Fig. 3.

$$cond(c, st) \overset{\text{def}}{=} \begin{cases} (st.flags).z & if\ c = EQ\ (*Equal*) \\ not\ ((st.flags).z) & if\ c = NE\ (*Not\ equal*) \\ \ldots & \\ true & if\ c = AL\ (*Default*) \end{cases}$$

Fig. 3. The *cond* function for conditional instructions

For instance, code `EQ` expects condition equality, which corresponds to the flag `Z` set to true. Code `NE` expects condition inequality and flag `Z` to false. Hence, a conditional instruction *ins* demands two rules:

- The *cond_true* rule: if the conditional instruction satisfies both the premises of the simple instruction rule and $cond(c, st)$ is true, then the conditional instruction performs the expected operation, referred to as ins_{pre} for premises and ins_{post} for conclusion from, e.g., Fig.2.
- The *cond_false* rule: if the conditional instruction meets the premise of the simple instruction rule but $cond(c, st)$ is false, then the instruction only updates the value of *pc*.

$$\frac{st.ok \wedge cond(c, st) \wedge ins_{pre}}{(ins, st) \rightarrow st[ins_{post}]}(cond\_true) \quad \frac{st.ok \wedge \neg cond(c, st) \wedge ins_{pre}}{(ins, st) \rightarrow st[pc \mapsto pc + 1]}(cond\_false)$$

4.3.2 Instructions with Condition Suffix. When suffix *s* is specified, conditional flags are updated after performing the default action of the instruction *i*. The `N` and `Z` flags update according to the result of *i*, while the other two relate to specific instructions. Three update functions are defined to classify the scenarios:

- The `upd_arithmetic` function updates the four condition flags according to the result of an instruction, e.g. **adc**.
- The `upd_logical` function is used to update N, Z and C flags after performing **mov** or bitwise instructions.
- The `upd_shift` function updates the three flags when shift operations (i.e. **asr**, **lsl**, **lsr** and **ror**) are performed.

Fig. 4 shows the semantics of the **add** instruction with condition suffix.Compared to rules in Fig. 2, instructions with condition suffixes mainly add flags to the updated memory state. For instance, the *adds* rule calls the *upd_arithmetic* to update the N, Z, C and V flags according to the result.

$$\frac{st.ok \wedge valid(rd) \wedge valid(rn) \wedge valid(op_2)}{(ADDS\ rd\ rn\ op_2, st) \rightarrow st[rd \mapsto rn + op_2, flags \mapsto upd\_arithmetic(rn +_i op_2), pc \mapsto pc + 1]} \quad (adds)$$

$$\frac{st.ok \wedge cond(c, st) \wedge valid(rd) \wedge valid(rn) \wedge valid(op_2)}{(ADDSC\ c\ \ rd\ rn\ op_2, st) \rightarrow st[rd \mapsto rn + op_2, flags \mapsto upd\_arithmetic(rn +_i op_2), pc \mapsto pc + 1]} \quad (addsc1)$$

$$\frac{st.ok \wedge \neg cond(c, st) \wedge valid(rd) \wedge valid(rn) \wedge valid(op_2)}{(ADDSC\ c\ \ rd\ rn\ op_2, st) \rightarrow st[pc \mapsto pc + 1]} \quad (addsc2)$$

Fig. 4. Semantics of instructions with condition suffix

In addition, if a conditional instruction has both condition and suffix, its semantic rule is composed with the aforementioned ones. For instance, the **add** instruction has two rules, Fig.4. *addsc1* is derived from the *adds* and *cond_true* rules. *addsc2* is an instance of the *cond_false* rule.

4.4 Properties of the ARM ISA

Based on the semantics of ARM instructions defined in *ARM-F★*, we specify the correctness requirements of isolation, branch, and no-effect listed in the ASM manual [20] and formalize them as Lemmas. Doing so allows us to prove the correctness of any ARM assembly code sequence modeled in F★. These lemmas are summarized in Tab 1.

Table 1. ASM properties from the ARM Compiler User Guide

Name	Specification
Isolation	A processor in one instruction set state cannot execute instructions from another instruction set.
Branch	Transformations between different modes only depend on the jump instruction [i.e. BX in the paper].
No-effect	If the condition test of a conditional instruction fails, the instruction has no effect.

Let i be the instruction that is to be executed next, st_0 be the current memory state, st_1 (i.e. $Eval(i, st_0)$) be the next memory state after executing i, and S_x be an instruction set in x mode. Theorem 1 says that the ARM operational semantics should ensure the processor never receives any instruction of the wrong instruction set in the current state.

Theorem 1 (Isolation) If $st_0.isa\_mode = x$, $st_0.ok$ and $st_1.ok$ then $i \in S_x$

Most instructions can execute in all modes, only **orn** is available in the Thumb32 mode. So the isolation property is equivalent to saying the validity holds if $st_0.ok$ and $st_1.ok$ are true.

Theorem 2 says that, if executing an instruction results in a memory state of a different mode as the current one, then the instruction can only be the jump instruction **bx**.

Theorem 2 (Branch) If $st_0.isa\_mode \neq st_1.isa\_mode$ then $i = bx$.

The formalization and proof of this property in F★ proceed by case analysis base on the definition of *mode*.

Theorem 3 says that, if the condition test of a conditional instruction fails, the instruction 1/ does not execute, 2/ does not write a value in the destination register, 3/ does not change any flag, and 4/ does not raise an exception.

Theorem 3 (No-effect) *Let* c, rd *be the condition code and the destination register (if present) of the instruction* i*. If* $Cond(c, st) = false$*, then* $st_1.pc = st_0.pc + 1$*,* $st_1.rd = st_0.rd$*,* $st_1.flags = st_0.flags$ *and* $st_1.ok = st_0.ok$*.*

We firstly say that **nop** and **nopc** have the same effect on a memory state. and then state the equivalence of a failed conditional instruction with **nop**. This way, the property is easily proved together with the lemmas of **nop**.

Along the way, we prove some auxiliary lemmas which will be useful for the verification of our case study. First, we can formalize the memory safety of an executed list of instructions L. Let st_0 be the initial memory state and st_1

(i.e.$EvalL(L, st_0)$) the final memory one, the $list\_ok\ L\ st_0$ function says that no instruction in L generates an exception if its current state is memory-safe.

Lemma 1 (List Memory Safety) If $st_0.ok$ and $(list\_ok\ L\ st_0)$ then $st_1.ok$.

Assuming a memory-safe initial state, the list memory safety lemma in F⋆ stipulates that no instruction in the list produces an unsafe state, and that the final state is memory-safe, by induction on the list L.

Additionally useful lemmas apply to specific instructions, such as nop_equiv_nopc and the 'load after store' lemma.

Lemma 2 (Load after Store) Let $st_1 = Eval((\mathbf{str}, r_d, r_n, o), st_0)$ and $st_2 = Eval((\mathbf{ldr}, r_d, r_n, o), st_1)$, then $st_0.r_d = st_1.r_d = st_2.r_d$.

It stipulates that the destination register always remains unchanged when the processor first executes the store instruction str with some registers and operands, and only executes the load instruction ldr with the same parameters.

5 EVALUATION CASE STUDY

Our goal is to specify the *riotboot* protocol and verify its correctness in F⋆. We first give the details of its implementation and verification. Next, we evaluate our formal model from the following perspectives: Bug-fixing/optimization, verification cost, and comparison with existing verified bootloaders.

5.1 Implementation & Verification

The *riotboot* in F⋆ has the same structure as its C version: *choose_image* and *cpu_jump_to_image*. For instance, *cpu_jump_to_image* corresponds to the ARM assembly code below. The input *image_address* is stored in $R0$. $i0$ copies the input to $R1$, $i1$ is to set MSP, $i2$ is to skip sp register (by 1 `int32` word instead of 4bits in ASM). $i3$ is to set thumb bit, i.e. bit[0] of $R0$ is 1. $i4$ causes a branch to the address contained in $R0$ and changes the instruction set to Thumb mode.

```
let i0: ins = LDR R1 R0 (OConst 0l)
let i1: ins = MOV SP (OReg R1)
let i2: ins = LDR R0 R0 (OConst 1l)
let i3: ins = ORR R0 R0 (OConst 1l)
let i4: ins = BX  R0
let cpu_jump_to_image_ins = [i0; i1; i2; i3; i4]
```

Theorem 4 (Functional Correctness) If *riotboot* finds a suitable image *i*, then 1/ *i* should be fletcher32-valid and be latest comparing with all valid images (*functional_correctness_aux0*) and 2/ the registers satisfy $sp = i.start\_addr$, $pc = i.start\_addr \mid 0x1$ and the processor mode is *Thumb* (*functional_correctness_aux1*).

Functional correctness is defined by two auxiliary lemmas according to the code structure of *riotboot*: *choose_image* requires the *images* is available and ensures the first lemma, while *cpu_jump_to_image* assumes the liveness of the selected image and guarantees the second lemma.

Theorem 5 (Memory Safety) *riotboot* requires an initially safe memory state and yields a safe final memory state.

Since Low⋆'s hyper-stack memory model guarantees memory safety of *choose_image*, we only need to prove that *cpu_jump_to_image* is memory-safe. The *list_memory_safety* lemma is useful for inductively proving this property.

5.2 Discussion

Building the *riotboot* case study in Low⋆ based on our ARM-F⋆ opens to interesting discussions regarding the memory models of Low⋆ and the ARM model, the validity of the booted image, and the extracted C and assembly code.

Memory Model. The *choose_image* module is encoded in Low⋆. It is based on its hyper-stack memory model while the *cpu_jump_to_image* function uses the ARM ISA memory model: a map from physical addresses to bytes. Hence a potential problem to compose the specification and factor the verification of two modules with different memory models. The technique used in VALE [11] can be reproduced to reconcile them by constraining the interface between the two modules: it defines a correct simulation relation that states that a Low★ memory is simulated by a VALE memory, and proves that store operations in VALE preserve this relation. Since *cpu_jump_to_image* does not contain any store instruction ('**str**'), it is hence trivial to prove that the assembly code sequence of riotboot never modifies memory.

Validity of the booted image. *riotboot* uses the fletcher32 algorithm to validate the checksum of the selected image. In the case study, a refinement type is introduced to prove the termination of fletcher32. To guarantee functional correctness of the algorithm, a solution is to add a predicate to the postcondition of the Low★ code relying on HACSPEC [14] to verify the functional correctness of cryptographic algorithms encoded in a Rust-like specification language from which F★ can be generated and used as the basis for proofs. In the present case study, we relied on the verified implementation of the fletcher32 algorithm provided by HACSPEC to trust image validation.

Extracted Code. *riotboot* is modeled in Low★ and *ARM-F★*. Code extraction relies on Low★'s KreMLin compiler to generate C code and on VALE to generate assembly code. They are composed as a standalone program.

5.3 Evaluation

Our F★/Low★ bootloader implementation relates to the *RIOT*[24] and *RIOT in Rust*[25] projects as part of Inria's *Future-Proof IoT* Challenge[15].

Monadic type checking improvements. We rapidly spotted an infinite loop in the original C&Rust versions of *riotboot* preventing code generation from Low★, as it would be given the `Div`(ergent) monad. Instead, we introduced an if statement (for the case no valid image is found). Strong typing in F★ also allowed us to spot and correct comparisons of header sequence numbers (i.e. versions) with header start addresses in the Rust version of *riotboot*[27]. Refinement types also allowed optimizations in *riotboot* while maintaining a verified equivalence with its unoptimized translation. For example, the if-statement on line 6 of *kernel_init* (Sec 3) has an unnecessary condition that can be omitted: the left part of the condition is equal to the statement `riotboot_slot_get_hdr(i)->start_addr` according to the definition of *riot_hdr* (line 4). But the right part is also equal to that statement according to the definition of `riotboot_slot_get_image_startaddr`.

Verification Cost. Our case study demonstrates that the verified programming workflow presented in the paper has a major impact on validation costs as most verification conditions generated by its type checker can automatically be discharged by F★'s companion SMT solver Z3. Verification conditions in *riotboot* are easy to define and express in F★/Low★. The number of refinement types, pre- and post- conditions we specified are listed in Table 2:

Table 2. Data Statistics of verification conditions

Module	Refinement Type	Pre-/Post-condition
Choose Image	14	11 / 18
Cpu Jump to Image	0	11 / 26

Refinement types in *riotboot* are used to set the length or scope of some parameters and can be directly derived from the source code. e.g. an input variable *i*, representing an index in an image table, should be less than the table's length. Our F★/Low★ implementation expresses the requirement $i \in [0, length - 1]$ by a refinement type.

Most pre/post-conditions in *Choose Image* concern the liveness of buffer pointers holding image headers, because Low★ requires a pointer to reference a live memory buffer before operation. The preconditions of *Cpu Jump to Image* express this memory safety condition and the post-conditions enforce them for all intermediate states.

Comparison. Table 3 compares our F⋆ model with related verified bootloaders.

Table 3. Comparison of verified bootloaders

Name	SourceCode	Model	Proofs	Language
SABLE	600+	250+	400+	Isabelle/HOL
First-stage	200+	n.a.	n.a.	Coq
riotboot	150+	180+	12	F⋆/Low⋆

$^{\text{n.a.}}$no artifact or data available.

To our knowledge, SABLE is the first formally verified bootloader. It uses the methodology of seL4[17] and adopts the Isabelle/HOL proof assistant. Its source code is over 600 lines and its formal specification 250 lines. The verification effort of SABLE represents more than 400 lines of proof.

The first-stage bootloader, another verified bootloader, formally verifies Sanctum's secure boot[9] (more than 200 lines of C) down to its RISC-V instruction semantics in Coq. Currently, this project is carrying on the whole correctness proof and, at the time of writing, no data or artifact are available for comparison.

Compared with related works, our verified implementation of *riotboot* with about 150 lines of C code in F⋆/Low⋆. The formal specification has a similar code size, and verification benefits a high degree of proof automation using the Z3 SMT solver. To prove the *riotboot* functional correctness and memory safety, only 7 auxiliary lemmas needed to be defined and 12 lines of manual proof declared.

6 RELATED WORKS

6.1 Bootloaders

Secure boot and trusted boot are two well-known features of bootloaders not to conflate with the verified programming of a bootloader. Secure boot is a valuable feature to help maintain the integrity of a platform at runtime, for example *Android's Verified Boot*. Trusted boot, defined by Trusted Computing Group (TCG), is a process to let a running application check if the system has booted into a trusted environment, e.g., *ARM's Trusted Boot*. While designed with the highest engineering skills, neither secure or trusted boot have provers' verified implementations. In this paper, our goal is to additionally propose a method to guarantee the functional correctness of a bootloader at minor additional engineering costs. Although some tools, like BootStomp[23], allow to identify bootloader vulnerabilities, our method allows to formally verify the absence thereof (up to the considered memory model).

Coreboot[33] is an open-source firmware platform delivering a lightning fast and secure boot. Some libraries of Coreboot, e.g. libgfxinit, written in the SPARK language, can automatically be proved to have no runtime errors, but most of Coreboot, written in C and assembly, is unverified.

Instead, SABLE is a formally verified bootloader developed using Isabelle/HOL. SABLE's method proves that the formalized behavior of bootloader's implementation, in C, satisfies its abstract specification requirements. Compared to our approach, the proof scale is quite important (more than 400 lines of proof) and compilation from C to machine code still remains is unverified (which it could using, e.g., CompCert). It considers Sanctum system's bootloader deployed on the RISC-V architectures and verify the first stage of boot. One advantage of this approach is to reuse existing Coq libraries, for instance the riscv-coq project[13], which implements the RISC-V ISA specification in Coq. However, this method requires a fully mechanical proof in Coq, and has, at the time of writing and the best of our knowledge, not delivered a complete and available correctness theorem.

Our method improves related works by employing verified programming to enforce functional correctness properties at compile-time in a way that maximizes proof automation, as presented in Sec 5.3.

6.2 Verified assembly languages

Pioneering works such as CompCert[19], seL4 and Sail[3] have formalized many architecture specifications, such as the x86/x64, ARM and RISC-V ISAs, allowing embedded systems designers to verify the expected properties of low-level programs using the artifacts of these projects, and complete detailed manual proofs using Isabelle/HOL or Coq.

To the best of our knowledge, the closest and only related work to ARM-F★, presented in this paper is the verified assembly language environment VALE. VALE is a tool to formally verify high-performance applications written in assembly language by relying on existing verification frameworks, such as Dafny[18] and F★. Currently, it includes a limited ARM ISA for the Dafny verification framework and doesn't support the verification of ARM assembly code in F★. Our ARM-F★ model covers a complete ARM ISA as found in practical applications like *riotboot*, which comprises registers (e.g. *sp*), advanced instructions like **orr** and **bx**, and mode transformations between *ARM* and *Thumb* ISAs. The ARM-F★ model also formalizes the correctness requirements listed in the ASM manual and provides both a methodology and useful lemmas for reuse in practical applications.

7 CONCLUSION AND FUTURE WORKS

In this paper, we have formalized the ARM instruction set in F★, and developed a verified implementation of the RIoT bootloader. Our formalization of the ARM ISA supports a general instruction set available in most ARM platforms. We also specify the correctness requirements from the ARM ASM manual and prove them as `Lemmas` in F★. Next, we model the RIoT bootloader in Low★, and verify functional correctness and memory safety of its main components. Our evaluation shows that, not only strong typing in the verified *riotboot* fixes potential vulnerabilities, provides an optimized code structure, but most importantly gains from a high degree of proof automation.

Our next project is to verify RIoT's rBPF subsystem[35] using the same methodology as for *riotboot*. We expect that an F★-verified rBPF will provide a more industrial-size experience to highlight the effectiveness of our workflow. Our final goal is to build useful libraries for the F★/Low★ community to verify low-level embedded programs, and also provide a set of verified subsystems for the RIoT community.

ACKNOWLEDGMENT

This work is partially supported by the Inria challenge RIOT-fp. The authors are grateful to Kaspar Schleiser and Emmanuel Baccelli for valuable information and help with riotboot, and to members of the HACSPEC project for providing a verified model of riotboot's fletcher32 checksum function.

REFERENCES

[1] Andrew W. Appel. 2011. Verified Software Toolchain. In *European Symposium on Programming* (Saarbrücken, Germany) *(ESOP'11/ETAPS'11)*. Springer-Verlag, Berlin, Heidelberg, 1–17.

[2] ARM. 2021. Arm Trusted Firmware. https://github.com/ARM-software/arm-trusted-firmware

[3] Alasdair Armstrong, Thomas Bauereiss, Brian Campbell, Alastair Reid, Kathryn E. Gray, Robert M. Norton, Prashanth Mundkur, Mark Wassell, Jon French, Christopher Pulte, Shaked Flur, Ian Stark, Neel Krishnaswami, and Peter Sewell. 2019. ISA Semantics for ARMv8-a, RISC-v, and CHERI-MIPS. *Proc. ACM Program. Lang.* 3, POPL, Article 71 (Jan. 2019), 31 pages. https://doi.org/10.1145/3290384

[4] Emmanuel Baccelli, Cenk Gündoğan, Oliver Hahm, Peter Kietzmann, Martine S Lenders, Hauke Petersen, Kaspar Schleiser, Thomas C Schmidt, and Matthias Wählisch. 2018. RIOT: An open source operating system for low-end embedded devices in the IoT. *IEEE Internet of Things Journal* 5, 6 (2018), 4428–4440.

[5] Yves Bertot and Pierre Castéran. 2013. *Interactive theorem proving and program development: Coq'Art: the calculus of inductive constructions.* Springer Science & Business Media, Berlin, Heidelberg.

[6] Barry Bond, Chris Hawblitzel, Manos Kapritsos, K. Rustan M. Leino, Jacob R. Lorch, Bryan Parno, Ashay Rane, Srinath Setty, and Laure Thompson. 2017. Vale: Verifying High-Performance Cryptographic Assembly Code. In *26th USENIX Security Symposium (USENIX Security 17)*. USENIX

Association, Vancouver, BC, 917–934. https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/bond

[7] Ernie Cohen, Markus Dahlweid, Mark Hillebrand, Dirk Leinenbach, Michał Moskal, Thomas Santen, Wolfram Schulte, and Stephan Tobies. 2009. VCC: A practical system for verifying concurrent C. In *International Conference on Theorem Proving in Higher Order Logics*. Springer, Springer Berlin Heidelberg, Berlin, Heidelberg, 23–42.

[8] Scott D Constable, Rob Sutton, Arash Sahebolamri, and Steve Chapin. 2018. *Formal Verification of a Modern Boot Loader.* Technical Report. Electrical Engineering and Computer Science.

[9] Victor Costan, Ilia Lebedev, and Srinivas Devadas. 2016. Sanctum: Minimal hardware extensions for strong software isolation. In *25th USENIX Security Symposium USENIX Security 16)*. USENIX Association, Austin, TX, 857–874.

[10] Lucas Franceschino, David Pichardie, and Jean-Pierre Talpin. 2021. Verified functional programming of an abstract interpreter. In *Static Analysis Symposium*. ACM, Springer, Cham, Chicago, United States, 1–20.

[11] Aymeric Fromherz, Nick Giannarakis, Chris Hawblitzel, Bryan Parno, Aseem Rastogi, and Nikhil Swamy. 2019. A verified, efficient embedding of a verifiable assembly language. *Proceedings of the ACM on Programming Languages* 3, POPL (2019), 1–30.

[12] GOOGLE. 2021. Verifying Boot. https://source.android.com/security/verifiedboot/verified-boot.html

[13] MIT PLV Group. 2021. riscv-coq:RISC-V Specification in Coq. https://github.com/mit-plv/riscv-coq

[14] hacspec. 2021. A specification language for crypto primitives in Rust. https://github.com/hacspec/hacspec

[15] Inria. 2021. RIOT-fp. https://future-proof-iot.github.io/RIOT-fp/about

[16] Bart Jacobs and Frank Piessens. 2008. *The VeriFast program verifier.* Technical Report. Citeseer.

[17] Gerwin Klein, Kevin Elphinstone, Gernot Heiser, June Andronick, David Cock, Philip Derrin, Dhammika Elkaduwe, Kai Engelhardt, Rafal Kolanski, Michael Norrish, Thomas Sewell, Harvey Tuch, and Simon Winwood. 2009. SeL4: Formal Verification of an OS Kernel. In *Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles (SOSP '09)*. Association for Computing Machinery, New York, NY, USA, 207–220.

[18] K Rustan M Leino. 2010. Dafny: An automatic program verifier for functional correctness. In *International Conference on Logic for Programming Artificial Intelligence and Reasoning*. Springer, Springer Berlin Heidelberg, Berlin, Heidelberg, 348–370.

[19] Xavier Leroy. 2020. *The CompCert C verified compiler: Documentation and user's manual.* Intern report. Inria. 1–78 pages.

[20] ARM Limited. 2016. ARM Compiler armasm User Guide v5.06. https://developer.arm.com/documentation/dui0473/m

[21] Tobias Nipkow, Lawrence C Paulson, and Markus Wenzel. 2002. *Isabelle/HOL: a proof assistant for higher-order logic*. Vol. 2283. Springer Science & Business Media, Berlin, Heidelberg.

[22] Jonathan Protzenko, Jean-Karim Zinzindohoué, Aseem Rastogi, Tahina Ramananandro, Peng Wang, Santiago Zanella-Béguelin, Antoine Delignat-Lavaud, Catalin Hritcu, Karthikeyan Bhargavan, Cédric Fournet, and Nikhil Swamy. 2017. Verified Low-Level Programming Embedded in F*. *PACMPL* 1, ICFP (Sept. 2017), 17:1–17:29. https://doi.org/10.1145/3110261

[23] Nilo Redini, Aravind Machiry, Dipanjan Das, Yanick Fratantonio, Antonio Bianchi, Eric Gustafson, Yan Shoshitaishvili, Christopher Kruegel, and Giovanni Vigna. 2017. BootStomp: On the Security of Bootloaders in Mobile Devices. In *26th USENIX Security Symposium (USENIX Security 17)*. USENIX Association, Vancouver, BC, 781–798. https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/redini

[24] RIoT. 2021. RIOT-OS. https://github.com/RIOT-OS/RIOT

[25] RIoT. 2021. RIOT-rs. https://github.com/future-proof-iot/RIOT-rs

[26] RIoT. 2021. riotboot. https://github.com/RIOT-OS/RIOT/tree/master/bootloaders/riotboot

[27] RIoT. 2021. riotboot-rs. https://github.com/kaspar030/riotboot-rs

[28] Michael Sammler, Rodolphe Lepigre, Robbert Krebbers, Kayvan Memarian, Derek Dreyer, and Deepak Garg. 2021. RefinedC: Automating the Foundational Verification of C Code with Refined Ownership Types. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*. Association for Computing Machinery, New York, NY, USA, 158–174. https://doi.org/10.1145/3453483.3454036

[29] Zygimantas Straznickas. 2020. *Towards a verified first-stage bootloader in Coq.* Ph. D. Dissertation. Massachusetts Institute of Technology.

[30] Nikhil Swamy, Juan Chen, Cédric Fournet, Pierre-Yves Strub, Karthikeyan Bhargavan, and Jean Yang. 2011. Secure distributed programming with value-dependent types. *ACM SIGPLAN Notices* 46, 9 (2011), 266–278.

[31] Jean-Pierre Talpin, Jean-Joseph Marty, Shravan Narayan, Deian Stefan, and Rajesh Gupta. 2019. Towards verified programming of embedded devices. In *DATE 2019 - 22nd IEEE/ACM Design, Automation and Test in Europe*. IEEE, Florence, Italy, 1445–1450. https://doi.org/10.23919/DATE.2019.8715067

[32] FStar Team. 2021. KreMLin. https://github.com/FStarLang/kremlin

[33] The Coreboot Development Team. 2021. Coreboot. https://www.coreboot.org/

[34] Shenghao YUAN and Jean-Pierre Talpin. 2021. verified riotboot in FStar. https://gitlab.inria.fr/syuan/verified-riotboot

[35] Koen Zandberg and Emmanuel Baccelli. 2020. Minimal Virtual Machines on IoT Microcontrollers: The Case of Berkeley Packet Filters with rBPF. In *9th IFIP/IEEE International Conference on Performance Evaluation and Modeling in Wired and Wireless Networks*. IEEE, Berlin / Virtual, Germany, 1–6.

APPENDIX

This appendix lists parts of our implementation of `riotboot` referenced in this article. As already mentioned, the complete implementation can be downloaded from a GitLab repository for evaluation purposes: https://gitlab.inria.fr/syuan/memocode-riotboot.

Section 4.3 defines valid functions to describe the constraints of most ARM instructions: The *exception_pc* function says that r_d can be *pc* only for a Thumb32 instruction and with a constant *c* in range 0-4095; The $valid(i, r_n, m)$ function says users are suggested to use *pc* or *sp* as the first operand in most ARM instructions;

$$exception\_pc(i, r_d) \overset{\text{def}}{=} \begin{cases} 0 \le n \le 4095 & if\ r_d = pc\ and\ i.op_2 = c \\ false & if\ r_d = pc \\ true & otherwise \end{cases}$$

$$valid(i, r_n, m) \overset{\text{def}}{=}$$

$$\begin{cases} r_n \ne pc\ \&\&\ r_n \ne sp & if\ i = \textbf{adc}\ and\ m = ARM \\ r_n \ne pc\ \&\&\ r_n \ne sp & if\ i = \textbf{add}\ and\ m = ARM \\ \ldots & \end{cases}$$

$$valid(i, op_2, m) \overset{\text{def}}{=}$$

$$\begin{cases} \quad reg\_not\_in\_operand(pc, op_2) & \\ \&\&\ reg\_not\_in\_operand(sp, op_2) & if\ i = \textbf{adc}|\textbf{add}\ \ldots \\ \quad reg\_not\_in\_operand(pc, op_2) & \\ \&\&\ reg\_not\_in\_operand(sp, op_2) & \\ \&\&\ no\_reg\_shift(op_2)\ if\ i = \textbf{and} & and\ m = ARM \\ \quad reg\_not\_in\_operand(pc, op_2) & \\ \&\&\ reg\_not\_in\_operand(sp, op_2) & \\ \&\&\ 1 \le i.sh \le 32\ if\ i = \textbf{asr} & and\ m = Thumb_i \\ \quad \ldots & \end{cases}$$

$$reg\_not\_in\_operand(reg, op_2) \overset{\text{def}}{=} \begin{cases} true & if\ op_2 = c \\ reg \ne r & if\ op_2 = r || r\ sop \end{cases}$$

$$no\_reg\_shift(op_2) \overset{\text{def}}{=} \begin{cases} r \ne pc & if\ op_2 = r\ sop \\ true & otherwise \end{cases}$$

$$valid(o, m) \overset{\text{def}}{=} \begin{cases} -4095 \le o \le 4095 & if\ m = ARM \\ -255 \le o \le 4095 & if\ m = Thumb32 \\ 0 \le o \le 124 & if\ m = Thumb16 \end{cases}$$

$$valid(i, m) \overset{\text{def}}{=} \begin{cases} m \ne ARM\ \&\&\ m \ne Thumb_{16} & if\ i = \textbf{orn} \\ true & otherwise \end{cases}$$

Fig. 5. The valid functions and related functions

The $valid(i, op_2, m)$ function says if the second operand is a register, it should not be *pc* or *sp* (i.e. *reg_not_in_operand* (reg, op_2)), and if it is a register with a shift, the shift register should not be *pc* (i.e. $no\_reg\_shift(op_2)$); The $valid(o, m)$ function says the offset in ARM mode should be in range [-4095,4095], in Thumb32 mode is [-255,4095] and in Thumb16 should be [0,124]; The $valid(i, m)$ says `ORN` is only available in the Thumb32 instruction set.

.1 Semantics of the ARM instruction set

This section details the complete operational semantics of the ARM instruction set as outline in Figures 7-9 in the style of a state transition system subject to the validity preconditions.

.1.1 Semantics of the simple ARM instruction set.

Mode. The processor must be in the correct *instruction set state* for the ARM instructions it is executing. ARM instructions are 32 bits wide. Thumb instructions are 16 or 32-bits wide. This paper models three kinds of modes in F★:

```
type mode = ARM | Thumb32 | Thumb16
```

Condition Flags. The APSR register is a record (`flag`) holding the negative (`N`), Zero (`Z`), Carry (`C`), and Overflow (`V`) condition flags. The processor uses them to determine whether or not to execute conditional instructions.

```
type flag = {   (*true => 1; false => 0*)
    n : bool;   (*Negative*)
    z : bool;   (*Zero*)
    c : bool;   (*Carry*)
    v : bool;   (*Overflow*)  }
```

Auxiliary Definitions. The memory model in ARM assembly is exposed by four operations declared as total functions in F★.

- `eval_mem`: reads from memory at given address.
- `upd_mem`: writes into memory at given address.
- `eval_reg`: reads from a given register (`[[reg]]`).
- `upd_reg`: writes into given register (`reg/val`).

'[[_]]' is also overloaded to get the value of condition flags, for instance `[[flags.c]]` returns the Carry value. In F★, these operations are encoded as follows:

```
unfold let eval_mem (addr: int32) (s:arm_state): Tot int32 =
  load_mem addr s.mem
let upd_mem (a:int32) (v:int32) (s:arm_state):Tot arm_state=
  {s with mem = store_mem a v s.mem}
unfold let eval_reg (r:reg) (s:arm_state) : Tot int32 =
  s.regs r
let upd_reg (r:reg) (v:int32) (s:arm_state): Tot arm_state =
  {s with regs = regs_make (fun (r':reg) ->
    if r = r' then v else s.regs r') }
```

Some symbols used in the Fig. 7 and Fig. 8 are explained below:

- $+, -, \times, \neg$ designate operations in range $[-2^{31}, 2^{31} - 1]$.
- &, |, ⊗, and ∼ are bitwise AND, OR, exclusive OR and NOT operations respectively.
- $\gg_a$ is the arithmetic right shift operation.
- $\ll$ is the logical left shift operation.
- $\gg_l$ is the logical right shift operation.
- $\gg_r$ is the rotate right shift operation.

The unit of a memory cell is a 32-bit integer, so the *pc* register usually increases by 1 (i.e. 4 bytes).

In order to better explain the shift operations, Fig. 6 shows four examples, where

- *ASR #n* moves the left-hand 32-n bits of a register to the right by n places, into the right-hand 32-n bits of the result. It copies the original bit(31) of the register into the left-hand n bits of the result.

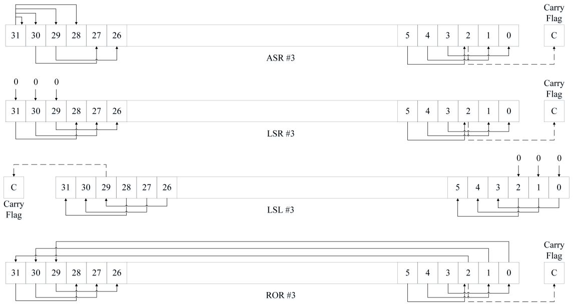

Fig. 6. Four shift operation: examples.

- *LSR #n* moves the left-hand 32-n bits of a register to the right by n places, into the right-hand 32-n bits of the result. It sets the left-hand n bits of the result to 0.
- *LSL #n* moves the right-hand 32-n bits of a register to the left by n places, into the left-hand 32-n bits of the result. It sets the right-hand n bits of the result to 0.
- *ROR #n* moves the left-hand 32-n bits of a register to the right by n places, into the right-hand 32-n bits of the result. It also moves the right-hand n bits of the register into the left-hand n bits of the result.

Note that the shift operations don't modify the Carry flag if the instruction lacks the condition flag suffix.

.1.2 Instructions with Condition Suffix. Most instructions can update the condition flags when the suffix *s* is specified. But there are two special cases: the instructions **cmp** and **cmn** always update this flag, while the instructions **bx, ldr, neg, nop** and **str** never do (they don't support that suffix). This section mainly discusses the semantics of instructions with condition suffix, i.e. of the form '$\{s\}$ i'.

Three update functions are defined to classify the scenarios:

- The `upd_arithmetic` function updates the four condition flags according to the result of an instruction.
  - `C` = 1 if an addition instruction (**adc**/**add**/**cmn**) produces a carry, or a subtraction instruction (**cmp**/**sub**) produce a borrow, otherwise `C` = 0.
  - `V` = 1 if the result of a signed add, subtract, or compare is greater than or equal to 2^{31}, or less than -2^{31}
- The `upd_logical` function is used to update N, Z and C flags after performing the **mov** instruction or bitwise instructions.
  - `C`: updates the flag during calculation of 2^{nd} operand.
  - `V`: does not affect the flag.
- The `upd_shift` function updates the three flags.
  - `C`: The flag is updated to the last bit shifted out.
  - `V`: does not affect the flag.

Fig. 9 shows the semantics rules of some instructions with condition suffix.

.2 Functional correctness of the assembly boot code

This section is the proof of functional correctness of the core assembly boot sequence code of the bootloader.

$$\frac{st.ok \wedge valid(rd) \wedge valid(rn) \wedge valid(op_2)}{(ADC\ rd\ rn\ op_2, st) \longrightarrow st[rd/_{[[rn]]+[[op_2]]+[[flags.c]]}, pc/_{[[pc]]+1}]} \qquad (adc)$$

$$\frac{st.ok \wedge valid(rd) \wedge valid(rn) \wedge valid(op_2)}{(ADD\ rd\ rn\ op_2, st) \longrightarrow st[rd/_{[[rn]]+[[op_2]]}, pc/_{[[pc]]+1}]} \qquad (add)$$

$$\frac{st.ok \wedge valid(rd) \wedge valid(rn) \wedge valid(op_2)}{(AND\ rd\ rn\ op_2, st) \longrightarrow st[rd/_{[[rn]]\&[[op_2]]}, pc/_{[[pc]]+1}]} \qquad (and)$$

$$\frac{st.ok \wedge valid(rd) \wedge valid(rn) \wedge valid(rs)}{(ASR\ rd\ rn\ rs, st) \longrightarrow st[rd/_{rn \gg_a rs}, pc/_{[[pc]]+1}]} \qquad (asr)$$

$$\frac{st.ok \wedge [[rd]].bit(0) = 0 \wedge valid(rd)}{(BX\ rd, st) \longrightarrow st[st.isa\_mode/_{Thumb_{16}}, pc/_{[[rd]]}]} \qquad (bx1)$$

$$\frac{st.ok \wedge [[rd]].bit(0) = 1 \wedge valid(rd)}{(BX\ rd, st) \longrightarrow st[st.isa\_mode/_{ARM}, pc/_{[[rd]]}]} \qquad (bx2)$$

$$\frac{st.ok \wedge valid(rn) \wedge valid(op_2)}{(CMN\ rn\ op_2, st) \longrightarrow st[flags/_{upd\_arith([[rn]]+[[op_2]])}, pc/_{[[pc]]+1}]} \qquad (cmn)$$

$$\frac{st.ok \wedge valid(rn) \wedge valid(op_2)}{(CMP\ rn\ op_2, st) \longrightarrow st[flags/_{upd\_arith([[rn]]-[[op_2]])}, pc/_{[[pc]]+1}]} \qquad (cmp)$$

$$\frac{st.ok \wedge valid(rd) \wedge valid(rn) \wedge valid(op_2)}{(EOR\ rd\ rn\ op_2, st) \longrightarrow st[rd/_{[[rn]] \otimes [[op_2]]}, pc/_{[[pc]]+1}]} \qquad (eor)$$

$$\frac{st.ok \wedge valid(rd) \wedge valid(rn) \wedge valid(o)}{(LDR\ rd\ rn\ o, st) \longrightarrow st[rd/_{\{\{[[rn]]+[[o]]\}\}}, pc/_{[[pc]]+1}]} \qquad (ldr)$$

$$\frac{st.ok \wedge valid(rd) \wedge valid(rn) \wedge valid(rs)}{(LSL\ rd\ rn\ rs, st) \longrightarrow st[rd/_{[[rn]] \ll [[rs]]}, pc/_{[[pc]]+1}]} \qquad (lsl)$$

$$\frac{st.ok \wedge valid(rd) \wedge valid(rn) \wedge valid(rs)}{(LSR\ rd\ rn\ rs, st) \longrightarrow st[rd/_{[[rn]] \gg_l [[rs]]}, pc/_{[[pc]]+1}]} \qquad (lsr)$$

$$\frac{st.ok \wedge valid(rd) \wedge valid(op_2)}{(MOV\ rd\ op_2, st) \longrightarrow st[rd/_{[[op2]]}, pc/_{[[pc]]+1}]} \qquad (mov)$$

$$\frac{st.ok \wedge valid(rd) \wedge valid(rn) \wedge valid(rm)}{(MUL\ rd\ rn\ rm, st) \longrightarrow st[rd/_{[[rd]] \times [[rm]]}, pc/_{[[pc]]+1}]} \qquad (mul)$$

$$\frac{st.ok \wedge valid(rd) \wedge valid(rm)}{(NEG\ rd\ rm, st) \longrightarrow st[rd/_{\neg[[rm]]}, pc/_{[[pc]]+1}]} \qquad (neg)$$

$$\frac{st.ok}{(NOP, st) \longrightarrow st[pc/_{[[pc]]+1}]} \qquad (nop)$$

$$\frac{st.ok \wedge valid(rd) \wedge valid(rn) \wedge valid(op_2) \wedge valid(st.isa\_mode)}{(ORN\ rd\ rn\ op_2, st) \longrightarrow st[rd/_{[[rn]]\ |\ (\sim[[op_2]])}, pc/_{[[pc]]+1}]} \qquad (orn)$$

$$\frac{st.ok \wedge valid(rd) \wedge valid(rn) \wedge valid(op_2)}{(ORR\ rd\ rn\ op_2, st) \longrightarrow st[rd/_{[[rn]]\ |\ [[op_2]]}, pc/_{[[pc]]+1}]} \qquad (orr)$$

$$\frac{st.ok \wedge valid(rd) \wedge valid(rn) \wedge valid(rs)}{(ROR\ rd\ rn\ op_2, st) \longrightarrow st[rd/_{[[rn]] \gg_r [[op_2]]}, pc/_{[[pc]]+1}]} \qquad (ror)$$

$$\frac{st.ok \wedge valid(rd) \wedge valid(rn) \wedge valid(o)}{(STR\ rd\ rn\ o, st) \longrightarrow st[\{[[rn]] + [[o]]\}/_{[[rd]]}, pc/_{[[pc]]+1}]} \qquad (str)$$

$$\frac{st.ok \wedge valid(rd) \wedge valid(rn) \wedge valid(op_2)}{(SUB\ rd\ rn\ op_2, st) \longrightarrow st[rd/_{[[rn]]-[[op_2]]}, pc/_{[[pc]]+1}]} \qquad (sub)$$

Fig. 7. Semantics of the simple ARM instruction set

$$\frac{st.ok \land cond(c,st) \land valid(rd) \land valid(rn) \land valid(op_2)}{(ADCC\ c\ rd\ rn\ op_2, st) \rightarrow st[rd/_{[[rn]]+[[op_2]]+[[flags.c]]}, pc/_{[[pc]]+1}]} \qquad (adcc)$$

$$\frac{st.ok \land cond(c,st) \land valid(rd) \land valid(rn) \land valid(op_2)}{(ADDC\ c\ rd\ rn\ op_2, st) \rightarrow st[rd/_{[[rn]]+[[op_2]]}, pc/_{[[pc]]+1}]} \qquad (addc)$$

$$\frac{st.ok \land cond(c,st) \land valid(rd) \land valid(rn) \land valid(op_2)}{(ANDC\ c\ rd\ rn\ op_2, st) \rightarrow st[rd/_{[[rn]]\&[[op_2]]}, pc/_{[[pc]]+1}]} \qquad (andc)$$

$$\frac{st.ok \land cond(c,st) \land valid(rd) \land valid(rn) \land valid(rs)}{(ASRC\ c\ rd\ rn\ rs, st) \rightarrow st[rd/_{rn \gg_a rs}, pc/_{[[pc]]+1}]} \qquad (asrc)$$

$$\frac{st.ok \land cond(c,st) \land [[rd]].bit(0) = 0 \land valid(rd)}{(BXC\ c\ rd, st) \rightarrow st[st.isa\_mode/_{Thumb_{16}}, pc/_{[[rd]]}]} \qquad (bxc1)$$

$$\frac{st.ok \land cond(c,st) \land [[rd]].bit(0) = 1 \land valid(rd)}{(BXC\ c\ rd, st) \rightarrow st[st.isa\_mode/_{ARM}, pc/_{[[rd]]}]} \qquad (bxc2)$$

$$\frac{st.ok \land cond(c,st) \land valid(rd) \land valid(rn) \land valid(op_2)}{(EORC\ c\ rd\ rn\ op_2, st) \rightarrow st[rd/_{[[rn]] \otimes [[op_2]]}, pc/_{[[pc]]+1}]} \qquad (eorc)$$

$$\frac{st.ok \land cond(c,st) \land valid(rd) \land valid(rn) \land valid(o)}{(LDRC\ c\ rd\ rn\ o, st) \rightarrow st[rd/_{\{\{[[rn]]+[[o]]\}\}}, pc/_{[[pc]]+1}]} \qquad (ldrc)$$

$$\frac{st.ok \land cond(c,st) \land valid(rd) \land valid(rn) \land valid(rs)}{(LSLC\ c\ rd\ rn\ rs, st) \rightarrow st[rd/_{[[rn]] \ll [[rs]]}, pc/_{[[pc]]+1}]} \qquad (lslc)$$

$$\frac{st.ok \land cond(c,st) \land valid(rd) \land valid(rn) \land valid(rs)}{(LSRC\ c\ rd\ rn\ rs, st) \rightarrow st[rd/_{[[rn]] \gg_l [[rs]]}, pc/_{[[pc]]+1}]} \qquad (lsrc)$$

$$\frac{st.ok \land cond(c,st) \land valid(rd) \land valid(op_2)}{(MOVC\ c\ rd\ op_2, st) \rightarrow st[rd/_{[[op2]]}, pc/_{[[pc]]+1}]} \qquad (movc)$$

$$\frac{st.ok \land cond(c,st) \land valid(rd) \land valid(rn) \land valid(rm)}{(MULC\ c\ rd\ rn\ rm, st) \rightarrow st[rd/_{[[rd]] \times [[rm]]}, pc/_{[[pc]]+1}]} \qquad (mulc)$$

$$\frac{st.ok \land cond(c,st) \land valid(rd) \land valid(rm)}{(NEGC\ c\ rd\ rm, st) \rightarrow st[rd/_{\neg[[rm]]}, pc/_{[[pc]]+1}]} \qquad (negc)$$

$$\frac{st.ok \land cond(c,st)}{(NOPC\ c, st) \rightarrow st[pc/_{[[pc]]+1}]} \qquad (nopc)$$

$$\frac{st.ok \land cond(c,st) \land valid(rd) \land valid(rn) \land valid(op_2) \land valid(st.isa\_mode)}{(ORNC\ c\ rd\ rn\ op_2, st) \rightarrow st[rd/_{[[rn]] \mid (\sim[[op_2]])}, pc/_{[[pc]]+1}]} \qquad (ornc)$$

$$\frac{st.ok \land cond(c,st) \land valid(rd) \land valid(rn) \land valid(op_2)}{(ORRC\ c\ rd\ rn\ op_2, st) \rightarrow st[rd/_{[[rn]] \mid [[op_2]]}, pc/_{[[pc]]+1}]} \qquad (orrc)$$

$$\frac{st.ok \land cond(c,st) \land valid(rd) \land valid(rn) \land valid(rs)}{(RORC\ c\ rd\ rn\ op_2, st) \rightarrow st[rd/_{[[rn]] \gg_r [[op_2]]}, pc/_{[[pc]]+1}]} \qquad (rorc)$$

$$\frac{st.ok \land cond(c,st) \land valid(rd) \land valid(rn) \land valid(o)}{(STRC\ c\ rd\ rn\ o, st) \rightarrow st[\{[[rn]] + [[o]]\}/_{[[rd]]}, pc/_{[[pc]]+1}]} \qquad (strc)$$

$$\frac{st.ok \land cond(c,st) \land valid(rd) \land valid(rn) \land valid(op_2)}{(SUBC\ c\ rd\ rn\ op_2, st) \rightarrow st[rd/_{[[rn]]-[[op_2]]}, pc/_{[[pc]]+1}]} \qquad (subc)$$

Fig. 8. Semantics of conditional ARM instruction sets

```
val functional_connectness_aux2_0: st:arm_state -> Lemma
  (requires (st.ok=true))
  (ensures  (let st0 = eval_cond_ins i0 st in
             let r0' = eval_reg R0 st in
             let r0  = eval_reg R0 st0 in
             let r1_0 = eval_reg R1 st0 in
```

$$\frac{st.ok \wedge cond(c, st) \wedge valid(rd) \wedge valid(rn) \wedge valid(op_2)}{(ADCSC\ c\ rd\ rn\ op_2, st) \rightarrow st[rd/_{[[rn]]+[[op_2]]+[[flags.c]]}, pc/_{[[pc]]+1}, flags/upd\_arith([[rn]] +_i [[op_2]] +_i [[flags.c]])]} \qquad (adcsc)$$

$$\frac{st.ok \wedge cond(c, st) \wedge valid(rd) \wedge valid(rn) \wedge valid(op_2)}{(ADDSC\ c\ rd\ rn\ op_2, st) \rightarrow st[rd/_{[[rn]]+[[op_2]]}, pc/_{[[pc]]+1}, flags/upd\_arith([[rn]] +_i [[op_2]])]} \qquad (addsc)$$

$$\frac{st.ok \wedge cond(c, st) \wedge valid(rd) \wedge valid(rn) \wedge valid(op_2)}{(ANDSC\ c\ rd\ rn\ op_2, st) \rightarrow st[rd/_{[[rn]]\&[[op_2]]}, pc/_{[[pc]]+1}, upd\_logical([[rn]] +_i [[op_2]])]} \qquad (andsc)$$

$$\frac{st.ok \wedge cond(c, st) \wedge valid(rd) \wedge valid(rn) \wedge valid(rs)}{(ASRSC\ c\ rd\ rn\ rs, st) \rightarrow st[rd/_{rn \gg_a rs}, pc/_{[[pc]]+1}, upd\_logical([[rn]], [[rs]])]} \qquad (asrsc)$$

$$\frac{st.ok \wedge cond(c, st) \wedge valid(rd) \wedge valid(rn) \wedge valid(op_2)}{(EORSC\ c\ rd\ rn\ op_2, st) \rightarrow st[rd/_{[[rn]] \otimes [[op_2]]}, pc/_{[[pc]]+1}, upd\_logical([[rn]] +_i [[op_2]])]} \qquad (eorsc)$$

$$\frac{st.ok \wedge cond(c, st) \wedge valid(rd) \wedge valid(rn) \wedge valid(rs)}{(LSLSC\ c\ rd\ rn\ rs, st) \rightarrow st[rd/_{[[rn]] \ll [[rs]]}, pc/_{[[pc]]+1}, upd\_logical([[rn]], [[rs]])]} \qquad (lslsc)$$

$$\frac{st.ok \wedge cond(c, st) \wedge valid(rd) \wedge valid(rn) \wedge valid(rs)}{(LSRSC\ c\ rd\ rd\ rs, st) \rightarrow st[rd/_{[[rn]] \gg_l [[rs]]}, pc/_{[[pc]]+1}, upd\_logical([[rn]], [[rs]])]} \qquad (lsrsc)$$

$$\frac{st.ok \wedge cond(c, st) \wedge valid(rd) \wedge valid(op_2)}{(MOVSC\ c\ rd\ op_2, st) \rightarrow st[rd/_{[[op2]]}, pc/_{[[pc]]+1}, flags/upd\_arith([[op_2]])]} \qquad (movsc)$$

$$\frac{st.ok \wedge cond(c, st) \wedge valid(rd) \wedge valid(rn) \wedge valid(rm)}{(MULSC\ c\ rd\ rn\ rm, st) \rightarrow st[rd/_{[[rd]] \times [[rm]]}, pc/_{[[pc]]+1}, flags/upd\_arith([[rn]] +_i [[rm]])]} \qquad (mulsc)$$

$$\frac{st.ok \wedge cond(c, st) \wedge valid(rd) \wedge valid(rn) \wedge valid(op_2) \wedge valid(st.isa\_mode)}{(ORNSC\ c\ rd\ rn\ op_2, st) \rightarrow st[rd/_{[[rn]]\ |\ (\sim[[op_2]])}, pc/_{[[pc]]+1}, upd\_logical([[rn]] +_i [[op_2]])]} \qquad (ornsc)$$

$$\frac{st.ok \wedge cond(c, st) \wedge valid(rd) \wedge valid(rn) \wedge valid(op_2)}{(ORRSC\ c\ rd\ rn\ op_2, st) \rightarrow st[rd/_{[[rn]]\ |\ [[op_2]]}, pc/_{[[pc]]+1}, upd\_logical([[rn]] +_i [[op_2]])]} \qquad (orrsc)$$

$$\frac{st.ok \wedge cond(c, st) \wedge valid(rd) \wedge valid(rn) \wedge valid(rs)}{(RORSC\ c\ rd\ rn\ rs, st) \rightarrow st[rd/_{[[rn]] \gg_r [[rs]]}, pc/_{[[pc]]+1}, upd_logical([[rn]], [[rs]])]} \qquad (rorsc)$$

$$\frac{st.ok \wedge cond(c, st) \wedge valid(rd) \wedge valid(rn) \wedge valid(op_2)}{(SUBSC\ c\ rd\ rn\ op_2, st) \rightarrow st[rd/_{[[rn]]-[[op_2]]}, pc/_{[[pc]]+1}, flags/upd\_arith([[rn]] +_i [[op_2]])]} \qquad (subsc)$$

Fig. 9. Semantics of conditional ARM instructions with condition suffix

```
               r1_0 == (eval_mem r0' st) /\
               r0 == r0'
             ))
let functional_connectness_aux2_0 st = ()

val functional_connectness_aux2_1: st:arm_state -> Lemma
  (requires (st.ok=true))
  (ensures  (let st1 = eval_cond_ins i1 st in
             let r0' = eval_reg R0 st in
             let r0  = eval_reg R0 st1 in
             let r1' = eval_reg R1 st in
             let r1 = eval_reg R1 st1 in
             let sp = eval_reg SP st1 in
               sp == r1 /\
               r1 == r1' /\
               r0 == r0'
             ))
let functional_connectness_aux2_1 st = ()

```

```
val functional_connectness_aux2_2: st:arm_state -> Lemma
  (requires (st.ok=true))
  (ensures  (let st2 = eval_cond_ins i2 st in
             let r0' = eval_reg R0 st in
             let r0 = eval_reg R0 st2 in
             let r1' = eval_reg R1 st in
             let r1 = eval_reg R1 st2 in
             let addr = Int32.int_to_t (add_mod (Int32.v r0') (Int32.v 1l)) in
             let sp' = eval_reg SP st in
             let sp = eval_reg SP st2 in
              r0 == eval_mem addr st /\
              r1 == r1' /\
              sp' == sp
             ))
let functional_connectness_aux2_2 st = ()

val functional_connectness_aux2_3: st:arm_state -> Lemma
  (requires (st.ok=true))
  (ensures  (let st3 = eval_cond_ins i3 st in
             let r0' = eval_reg R0 st in
             let r0 = eval_reg R0 st3 in
             let r1' = eval_reg R1 st in
             let r1 = eval_reg R1 st3 in
             let sp' = eval_reg SP st in
             let sp = eval_reg SP st3 in
               bit_n (Int32.v r0) 31 == true /\
               r0  == Int32.int_to_t (logor (Int32.v r0') (Int32.v 1l)) /\
               r1 == r1' /\
               sp' == sp
             ))

#push-options "--ifuel 50 --fuel 50 --z3rlimit 320"
let functional_connectness_aux2_3 st = ()
#pop-options

val functional_connectness_aux2_4: st:arm_state -> Lemma
  (requires (st.ok=true /\
             (let r0 = eval_reg R0 st in
              bit_n (Int32.v r0) 31 == true)
              ))
  (ensures  (let st4 = eval_cond_ins i4 st in
             let pc  = eval_reg PC st4 in
             let r0'  = eval_reg R0 st in
             let r0  = eval_reg R0 st4 in
             let r1' = eval_reg R1 st in
             let r1 = eval_reg R1 st4 in
             let sp' = eval_reg SP st in
             let sp = eval_reg SP st4 in
              st4.isa_mode == Thumb16 /\
              sp == sp' /\
              r0 == r0' /\
              r1 == r1' /\
              pc == r0
             ))
let functional_connectness_aux2_4 st = ()

val functional_connectness_aux1: st:arm_state -> Lemma
  (requires (st.ok = true))
  (ensures  (let st' = eval_list_ins cplist st in
             let st0 = eval_cond_ins i0 st in
             let st1 = eval_cond_ins i1 st0 in
             let st2 = eval_cond_ins i2 st1 in
             let st3 = eval_cond_ins i3 st2 in
             let st4 = eval_cond_ins i4 st3 in
               st' == st4
              ))
let functional_connectness_aux1 st = ()

val functional_connectness_aux2: st:arm_state -> Lemma
  (requires (st.ok = true))
  (ensures  (let st0 = eval_cond_ins i0 st in
             let st1 = eval_cond_ins i1 st0 in
             let st2 = eval_cond_ins i2 st1 in
             let st3 = eval_cond_ins i3 st2 in
             let st4 = eval_cond_ins i4 st3 in
             let r0' = eval_reg R0 st in
             let addr = Int32.int_to_t (add_mod (Int32.v r0') (Int32.v 1l)) in
             let sp = eval_reg SP st4 in
             let r1 = eval_mem r0' st in
             let pc  = eval_reg PC st4 in
             let r0  = eval_reg R0 st4 in
              st4.isa_mode == Thumb16 /\
              sp == r1 /\
              r0 == Int32.int_to_t (logor (Int32.v (eval_mem addr st)) (Int32.v 1l)) /\
              pc == r0
              ))

#push-options "--ifuel 50 --fuel 50 --z3rlimit 320"
let functional_connectness_aux2 st =
  let st0 = eval_cond_ins i0 st in
  let st1 = eval_cond_ins i1 st0 in
  let st2 = eval_cond_ins i2 st1 in
  let st3 = eval_cond_ins i3 st2 in
    functional_connectness_aux2_0 st;
    functional_connectness_aux2_1 st0;
    functional_connectness_aux2_2 st1;
    functional_connectness_aux2_3 st2;
    functional_connectness_aux2_4 st3
#pop-options

val functional_connectness: st:arm_state -> Lemma
  (requires (st.ok=true))
  (ensures  (let st1 = eval_list_ins cplist st in
             let r0' = eval_reg R0 st in
             let sp = eval_reg SP st1 in
             let r0 = eval_reg R0 st1 in
             let addr = Int32.int_to_t (add_mod (Int32.v r0') (Int32.v 1l)) in
             let pc = eval_reg PC st1 in
             let r1 = eval_mem r0' st in
               r0 == Int32.int_to_t (logor (Int32.v (eval_mem addr st)) (Int32.v 1l)) /\
               sp == r1 /\
               pc == r0
               ))

#push-options "--ifuel 50 --fuel 50 --z3rlimit 320"
let functional_connectness st =
   functional_connectness_aux1 st; functional_connectness_aux2 st
#pop-options
```

.3 Validated choose_image function

Finally, this section lists the verified fletcher32 and choose_image functions of the bootloader.

```
type fletcher = (pub_uint32 & pub_uint32)
type header = (pub_uint32 & pub_uint32 & pub_uint32 & pub_uint32)

let riotboot_magic : pub_uint32 =
  pub_u32 0x544f4952
let new_fletcher () : fletcher =
  (pub_u32 0x0, pub_u32 0x0)
let max_chunk_size () : uint_size =
  usize 360

let reduce_u32 (x_0 : pub_uint32) : pub_uint32 =
  ((x_0) &. (pub_u32 0xffff)) +. ((x_0) `shift_right` (pub_u32 0x10))

let combine (lower_1 : pub_uint32) (upper_2 : pub_uint32) : pub_uint32 =
  (lower_1) |. ((upper_2) `shift_left` (pub_u32 0x10))

let update_fletcher (f_3 : fletcher) (data_4 : seq pub_uint16) : fletcher =
  let max_chunk_size_5 = max_chunk_size () in
  let (a_6, b_7) = f_3 in
  let (a_6, b_7) =
    foldi (usize 0) (seq_num_chunks (data_4) (max_chunk_size_5)) (fun i_8 (
        a_6,
        b_7
      ) ->
      let (chunk_len_9, chunk_10) =
        seq_get_chunk (data_4) (i_8) (max_chunk_size_5)
      in
      let intermediate_a_11 = a_6 in
      let intermediate_b_12 = b_7 in
      let (intermediate_a_11, intermediate_b_12) =
        foldi (usize 0) (chunk_len_9) (fun j_13 (
            intermediate_a_11,
            intermediate_b_12
          ) ->
          let intermediate_a_11 =
            (intermediate_a_11) +. (
              cast U32 PUB (array_index
                 (**) #pub_uint16 #chunk_len_9
                 (chunk_10) (j_13)))
          in
          let intermediate_b_12 = (intermediate_b_12) +. (intermediate_a_11) in
          (intermediate_a_11, intermediate_b_12))
        (intermediate_a_11, intermediate_b_12)
      in
      let a_6 = reduce_u32 (intermediate_a_11) in
      let b_7 = reduce_u32 (intermediate_b_12) in
      (a_6, b_7))
    (a_6, b_7)
  in
  let a_6 = reduce_u32 (a_6) in
  let b_7 = reduce_u32 (b_7) in
  (a_6, b_7)

let value (x_14 : fletcher) : pub_uint32 =
  let (a_15, b_16) = x_14 in
  combine (a_15) (b_16)

let header_as_u16_slice (h_17 : header) : seq pub_uint16 =
  let (magic_18, seq_number_19, start_addr_20, _) = h_17 in
  let magic_21 = u32_to_be_bytes (magic_18) in
  let seq_number_22 = u32_to_be_bytes (seq_number_19) in
  let start_addr_23 = u32_to_be_bytes (start_addr_20) in
  let u8_seq_24 = seq_new_ (pub_u8 0x0) (usize 12) in
  let u8_seq_25 =
    seq_update_slice (u8_seq_24) (usize 0) (magic_21) (usize 0) (usize 4)
  in
  let u8_seq_26 =
    seq_update_slice (u8_seq_25) (usize 4) (seq_number_22) (usize 0) (usize 4)
  in
  let u8_seq_27 =
    seq_update_slice (u8_seq_26) (usize 8) (start_addr_23) (usize 0) (usize 4)
  in
  let u16_seq_28 = seq_new_ (pub_u16 0x0) (usize 6) in
  let (u16_seq_28) =
    foldi (usize 0) (usize 6) (fun i_29 (u16_seq_28) ->
      let u16_word_30 =
        array_from_seq (2) (
          seq_slice (u8_seq_27) ((i_29) * (usize 2)) (usize 2))
      in
      let u16_value_31 = u16_from_be_bytes (u16_word_30) in
      let u16_seq_28 = array_upd u16_seq_28 (i_29) (u16_value_31) in
      (u16_seq_28))
    (u16_seq_28)
  in
  u16_seq_28

let is_valid_header (h_32 : header) : bool =
  let (magic_number_33, seq_number_34, start_addr_35, checksum_36) = h_32 in
  let slice_37 =
    header_as_u16_slice (
      (magic_number_33, seq_number_34, start_addr_35, checksum_36))
  in
  let result_38 = false in
  let (result_38) =
    if (magic_number_33) = (riotboot_magic) then begin
      let fletcher_39 = new_fletcher () in
      let fletcher_40 = update_fletcher (fletcher_39) (slice_37) in
      let sum_41 = value (fletcher_40) in
      let result_38 = (sum_41) = (checksum_36) in
      (result_38)
    end else begin (result_38)
    end
  in
  result_38

let choose_image (images_42 : seq header) : (bool & pub_uint32) =
  let image_43 = pub_u32 0x0 in
  let image_found_44 = false in
  let (image_43, image_found_44) =
    foldi (usize 0) (seq_len (images_42)) (fun i_45 (image_43, image_found_44
      ) ->
      let header_46 = array_index
        (**) #header #(seq_len images_42)
        (images_42) (i_45)
      in
      let (magic_number_47, seq_number_48, start_addr_49, checksum_50) =
        header_46
      in
      let (image_43, image_found_44) =
        if is_valid_header (
          (magic_number_47, seq_number_48, start_addr_49, checksum_50
          )) then begin
          let change_image_51 =
            not ((image_found_44) && ((seq_number_48) <=. (image_43)))
          in
          let (image_43, image_found_44) =
            if change_image_51 then begin
              let image_43 = start_addr_49 in
              let image_found_44 = true in
              (image_43, image_found_44)
            end else begin (image_43, image_found_44)
          end
        in
        (image_43, image_found_44)
      end else begin (image_43, image_found_44)
      end
    in
    (image_43, image_found_44))
  (image_43, image_found_44)
in
(image_found_44, image_43)
```