

Lemma Functions for Frama-C: C Programs as Proofs

Grigoriy Volkov

*Faculty of Computer Science
National Research University
Higher School of Economics
Moscow, Russia
gdvolkov@edu.hse.ru*

Mikhail Mandrykin

*Software Engineering Department
Ivannikov Institute for System Programming of the
Russian Academy of Sciences
Moscow, Russia
mandrykin@ispras.ru*

Denis Efremov

*Faculty of Computer Science
National Research University
Higher School of Economics
Moscow, Russia
defremov@hse.ru*

Abstract—This paper describes the development of an auto-active verification technique in the Frama-C framework. We outline the lemma functions method and present the corresponding ACSL extension, its implementation in Frama-C, and evaluation on a set of string-manipulating functions from the Linux kernel. We illustrate the benefits our approach can bring concerning the effort required to prove lemmas, compared to the approach based on interactive provers such as Coq. Current limitations of the method and its implementation are discussed.

Index Terms—formal verification, deductive verification, Frama-C, auto-active verification, lemma functions, Linux kernel

I. INTRODUCTION

Deductive verification is one of the most “heavyweight” static techniques of formal reasoning about software properties. It requires manual annotation of source code with formal specifications and manual or semi-automatic proof of conformance between code and specifications. Deductive verification is applied in areas where the cost of error can be too high. The main benefit of this approach is that it can mathematically guarantee that code is free from particular error classes.

In the toolsets based on the Hoare logic (such as Frama-C [1], VeriFast [2], Dafny [3], GNATprove [4]), in order to verify the correctness of a function, the verification engineer is required to define a precondition — a predicate that should hold in order for the following code to be functionally correct and contain no runtime errors (i. e., models of undefined behavior that are known to the verification tool), a postcondition — a predicate that should hold after a function execution and describe what this code should act like to be functionally correct. Pre- and postconditions form the function’s contract. Frame-conditions or effects are sometimes viewed as a separate entity in the contract, but typically they present a type of postcondition that restrict the memory state of a program. A formal contract is required to reason about the function’s correctness.

Given a formal contract, an engineer can try to prove that a function conforms to its specification. To do this one needs

to additionally specify loop invariants and termination expression (loop variants) in case the function under analysis contains loops. Then the verification instrument generates verification conditions (VCs), which can be separated into categories: safety and behavioral ones. Safety VCs are responsible for runtime error checks for known (modeled) types, while behavioral VCs represent the functional correctness checks. All VCs need to be discharged in order for the function to be considered fully proved (totally correct). This can be achieved either by manually proving all VCs discharged with an interactive prover (i. e., Coq, Isabelle/HOL or PVS) or with the help of automatic provers.

It is well known that most problems related to verifying the conformance between some code in basically any practical programming language and its formal specification in any sufficiently expressive specification language are undecidable in general. The situation when an automatic prover cannot discharge a verification condition is common. This may be due to various reasons: limited time or memory resources for the prover, an error in the code, an incorrect formal specification, an incomplete formal specification, an incapacity of the prover to finish the proof. The specifications debugging is hard. Despite the modern advances [5]–[7] in the field, the whole situation is still poor when specifications become complex.

As it can be seen the whole process of deductive verification requires a significant manual effort from a verification engineer. This implies that practically all verification tools have some capabilities for user interaction beyond simply supplying the program and its specification. This additional user interaction required to eventually prove the correspondence between the code and its specification is usually called the proof overhead. This overhead can significantly depend on the way the user interaction is accomplished in a particular verification tool. Even though these ways can generally be very diverse, for simplicity we dare to speculatively group them into the following three major categories:

- External provers. The vast majority of verification tools support this approach to user interaction at

some point, be it as a part of a standard workflow or as a last resort. The essence is that the verification conditions (VCs) produced by the tool are translated into the native language of some interactive proof assistant, e.g., Coq, Isabelle/HOL or PVS. Then the resulting theorems are proved using the capabilities of the prover, sometimes with some extensions provided by the verification tool, e.g., custom tactics. This approach is often both simple and efficient, but for some use cases, it has significant drawbacks such as the ones we describe in this paper.

- Interactive proof editors. Some verification tools such as WP plugin for Frama-C, Why3 verification platform or KIV [8] verification system go further and implement their own custom interactive proof management facilities such as tactics, transformations, and strategies. This is a hard yet powerful approach with its own advantages and shortcomings. One of its major disadvantages especially relevant to our particular case is a high implementation overhead, which prevents tentative adoption and quick preliminary evaluation of the results.
- Auto-active proofs. One of the possible techniques to reduce and automate the manual work and make deductive verification more suitable for such methods as continuous verification [9]–[11] is auto-active verification. The general idea is to express the proofs in the same input language as the program and its specification themselves relying on the Curry-Howard correspondence between the programs and the proofs. This approach is quite widely applied in particular in such verification tools as Dafny [3], VeriFast [2], GNATprove [12], and Why3 [13]. This is the approach that we pursue and evaluate in this paper.

In this paper, we present our experience with integration of support for auto-active verification in ACSL [14], Frama-C, and AstraVer deductive verification tool and applying this approach to our existing set of specified string-manipulating functions from the Linux kernel. We demonstrate that integration of auto-active proofs to ACSL and Frama-C is both beneficial and promising approach to automated deductive verification that significantly reduces manual proof overhead.

This paper is organized as follows. In Section II we give an overview of the VerKer project and its specifications and describe the problems we are solving by applying auto-active verification to this project. Next, in Section III we present the implementation of the technique in the Frama-C framework and the AstraVer plugin. In Sections IV and V we describe the obtained result and discuss future work on open issues.

II. VERKER PROJECT

The VerKer project [15] consists of standard library functions taken from the Linux kernel. The main idea of

the project is to develop the correct and detailed ACSL [14] specifications and fully prove these functions with the Frama-C framework. Thus, evaluating the toolset (Frama-C + AstraVer Plugin (a fork of Jessie) + Why3 + Solvers) on the real code (not artificial examples) and highlighting C idioms that are hard to describe with ACSL specification, not possible to model with the deductive verification plugin, or just giving the right direction for the tools development.

ACSL is a Behavioral Interface Specification Language (BISL) [16] implemented in Frama-C. At the same time, many higher-level formalisms for efficient specification and verification of heap-manipulating or parallel programs such as dynamic frames, separation logic, permissions, rely/guarantee reasoning (e.g., ownership methodology), specification refinement (e.g., function contract inclusion) or auto-active verification are not included in ACSL. Therefore ACSL is not tightly bound to any particular verification methodology and support for its features can significantly vary in tools.

A. Verification Approach

Specifications development is an iterative process. It cannot be broadly described in the general case. Every algorithm requires its own approach [17]–[20] with separate axiomatizations fitting each particular case and implementation. A common way to write formal specifications is to factor out the most complex logical statements into theorems and lemmas so that the provers can discharge the rest of the verification conditions automatically. The theorems and lemmas are then proved manually with an interactive prover. The main benefits of this approach are the following:

- Lemmas are formulated in a more general way so that they contain less implementation context specific details (e.g., the assertion in a particular function). This eases the process of their manual proving and allows one to reuse them in multiple proofs of VCs and other lemmas.
- In this case, lemmas are “aside from code”, thus making formal specifications more maintainable and tolerable to small code changes because they rarely provoke changes in general statements.

The VerKer project adheres to this approach in the ACSL specifications. The project currently contains 26 fully proved functions from Linux kernel library folder, 16 functions among them are string related (i.e., have the `str*` prefix and perform operations on strings), and 7 are memory related (i.e., have `mem*` prefix). Only 9 functions (see Table I) contain in its formal specification a logical function with lemmas formulated on it. The formal contracts are written in a way to prove correspondence between logical functions and C functions. The lemmas are excessive, and provers sometimes can automatically prove the correctness of a lemma based on other lemmas.

In the paper [15] it is claimed that the functions are fully-proved but lemmas are not proved (only automatic

Function Name	Number of Lemmas	Proved Automatically	Unproved
check_bytes8	3	3	0
strlen	10	4	6
skip_spaces	7	1	6
strchr	7	4	3
strchrnul	7	5	2
strspn	8	5	3
strcspn	5	2	3
strnlen	17	11	6
strpbrk	5	2	3
Total	69	37	32

TABLE I
LEMMAS IN THE VERKER PROJECT

checks for contradiction was performed). The proving of all lemmas in the project considered as a future work in the paper.

B. Shortcomings and limitations

Thus we can found the following limitations of the approach based on external provers:

- The use of external interactive proof assistants or proof editors requires the users to get accustomed to at least two different specification representations, corresponding languages and tools. This makes it harder for the users to employ their skills acquired during the process of specification within one verification environment (e. g., the deductive verification plugin) later in the process of proof within the other one (e. g., the Coq proof assistant). In particular, the range of supported language constructs can significantly differ. Many ACSL constructs such as pointer dereference do not have direct counterparts in the language of the interactive prover.
- The reproduction of external proofs requires the use of external tools that work on their own representations — yet the translations performed by the verification tool in order to obtain the external representations can cause various surprising instabilities, where the proof suddenly fails after small supposedly irrelevant changes are made to program source code or specifications. This is especially well demonstrated by the pervasive problem of generating unique, stable and predictable names for various program and specification entities such as variables and lemmas. For instance, when a lemma or variable is renamed, the proofs have to be adjusted accordingly, yet the relation between the names used in the original program and the names used in the proofs can be quite complex.
- The external proofs are kept separately from the specifications and are therefore not directly available to the reader of the verified program. This can make it harder to estimate the effort needed for modification of certain parts of the program or specification or to keep the proofs in sync with the program and specification.
- The use of external interactive provers prevents an application of various encodings optimized for automatic

provers, e. g., SMT-solvers since the VCs should remain manageable by different tools that apply different tactics or decision procedures. In particular, while SMT solvers efficiently handle large ground (quantifier-free) formulas modulo theories, most interactive tools are much more suitable to handle smaller formulas in higher-order logic.

- External proofs introduce additional problems of proof location and sharing when the same lemma is included or imported in several source code files.

III. AUTO-ACTIVE PROOFS

According to its original definition in [21], *auto-active verification* is any verification technique where user input is supplied before VC generation. The user input can be given in various forms including but not limited to: auxiliary annotations such as verifier pragmas, triggers, ad-hoc lemma instantiations, ghost code, and lemma functions. Among those, lemma functions are a primary tool intended to convey proofs of user-defined auxiliary lemmas. The proofs are expressed as ordinary pure (i. e., without side-effects) total (i. e., always successfully terminating) imperative functions. Though the execution of pure total imperative code does not produce any observable effects with respect to the programming language semantics, when interpreted by the verification tool the instructions of the code can produce a substantial effect on its internal state or the internal state of the automatic theorem prover. In particular, through the process of VCs generation, the proof-carrying code can determine the fragment of current proof context that is explicitly known to the verifier or the prover versus the fragment that is only known implicitly as a set of derivable inferences. As most of the logic relevant to deductive verification is generally undecidable, neither the verifier nor the prover can ever explore all propositions derivable in the current context. The code of the lemma functions thus serves as a guide for the verification tool that efficiently indicates the relevant parts of the search space. One of the most significant cases where verification tools need such a guide is with inductive proofs as they have very limited support in automatic provers.

Consider the following inductive proof expressed as a lemma function:

```

logic char * strchrnul(char * s, char c) =
  *s ≡ c ? s :
  *s ≡ '\0' ? s :
    strchrnul(s + 1, c);
  ...

```

```

/*@ ghost
  @ /@ lemma
  @ @ requires valid_str(s);
  @ @ decreases strlen(s);
  @ @ ensures s ≤ strchrnul(s, c) ≤ str + strlen(s);
  @ @/
  @ void strchrnul_in_range(char * s, char c)
  @ {
  @   if (*s != '\0' && *s != c)
  @     strchrnul_in_range(s + 1, c);
  @ }
  @*/

```

The lemma and its proof are expressed within a single recursive function definition `strchrnul_in_range`. The induction here is simple as it exactly follows the recursive definition of the logic function `strchrnul`. The pre-condition of the function requires the parameter string `s` to be a valid C string so that it has a well-defined finite length. Then the `if` statement in the function body splits between the base and inductive cases. In the base cases `*s == '\0'` and `*s == c` the post-condition trivially follows from the definition of the function `strchrnul`. In the inductive case, which corresponds to the body of the `if` statement, the post-condition follows from the inductive assumption for strings of a smaller length and an additional instance of a lemma about the length of the string `s + 1` that is implicitly present in the proof context (this lemma is actually also proved by a lemma function that is provided earlier in the code). The recursion is well-founded since in the only recursive call the value of the variant `strlen(s)` specified in the `decreases` clause is strictly smaller than that in the caller function. With the `AstraVer` plugin, the correctness of the function `strchrnul_in_range` can be verified automatically by discharging all the VCs using an SMT solver such as `Alt-Ergo` or `CVC4`. The automatic proofs require no induction and only a few quantifier instantiations, while all the necessary terms are explicitly present in the VCs. So in either solver, the verification takes far less than a second.

Now since the function `strchrnul_in_range` is proved correct for arbitrary values of its parameters, and since the function is pure and total, the statement about its correctness with respect to the contract can be expressed as the following lemma:

```

lemma strchrnul_in_range :
  ∀char * s, char c;
  valid_str(s) ⇒
    s ≤ strchrnul(s, c) ≤ str + strlen(s);

```

This lemma can be automatically generated from the function contract and used in the proofs later in the program code, both explicitly by the user by calling the function `strchrnul_in_range` with some concrete arguments or implicitly by the SMT solver by instantiating the lemma with some concrete terms.

This is the basic idea of lemma functions and their use. Now let us consider our approach to the integration of

lemma function into ACSL in particular. We first present our existing approach and current results that we obtained by using it, and then we discuss some of its limitations and possible directions of future developments.

A. Integrating Lemma Functions Into ACSL/Frama-C

In our fork of the `Frama-C` platform developed as part of the `AstraVer` toolset, we implemented relatively simple support for lemma functions as a special case of total pure ghost functions. Both ACSL specification and its implementation in mainline `Frama-C` provide support for ghost functions. Meanwhile, the `AstraVer` plugin is able to generate VCs to check ACSL `assigns`, `allocates`, `decreases` and `terminates` clauses of function contracts. So the only missing part to provide basic support for lemma functions was the automatic generation of lemmas and their implicit importing in all the code functions located further in the source code.

More concretely, in order to allow lemma functions to be used in a similar way to regular ACSL lemmas (which can only be written in logic), we introduced a new `lemma` keyword for C function specifications.

Our modified version of `Frama-C` processes it in the following way:

- For a function contract of the form


```

/@ requires R(p1, ..., pn, g1, ..., gn);
@ ensures E(p1, ..., pn, g1, ..., gn, \result);
@/,

```

where p_1, \dots, p_n are function parameters, g_1, \dots, g_n are global variables and the special variable `\result` denotes the return value of the function, a corresponding axiom is generated, that is, an axiom of the form

$$\begin{aligned}
 & \forall \text{typeof}(p_1) p_1, \dots, \text{typeof}(p_n) p_n, \\
 & \text{typeof}(g_1) g_1, \dots, \text{typeof}(g_n) g_n; \\
 & \exists \text{typeof}(\backslash\text{result}) \backslash\text{result}; \\
 & R(p_1, \dots, p_n, g_1, \dots, g_n) \implies \\
 & E(p_1, \dots, p_n, g_1, \dots, g_n, \backslash\text{result});
 \end{aligned}$$

- All logic definitions such as functions, predicates, axioms, and lemmas in ACSL are grouped into axiomatic blocks. The generated axiom is placed in a new axiomatic block. `AstraVer` implements on-demand import of all the definitions from axiomatic blocks based on occurrences of symbols defined in the axiomatic. As the axiom name cannot be directly used in an ACSL specification, the axiom is not included in any proof context until some other symbol in this new axiomatic is defined and later used in another specification (code function or another axiomatic block). In particular, this approach removes the axiom from the context of the lemma function itself, preventing the axiom from trivially proving the function's verification conditions.
- A dummy identically true predicate with a unique name is generated, and its definition is added into the new axiomatic block.
- A precondition requiring the dummy predicate from the new axiomatic block is inserted into the contracts

of all function located after the currently processed lemma function, which causes the definitions from the axiomatic block (including the generated axiom) to be imported into the proof context of these functions. As the dummy predicate is generated by the tool and has a unique name, the user cannot access it and thus create any extra dependencies disrupting the ordering of lemma proofs according to their code locations.

- The lemma function itself gets some additional clauses inserted into its contract specification:
 - **assigns** $\backslash nothing$; **allocates** $\backslash nothing$ — to ensure the function is pure (i. e., it does not touch any global state);
 - **terminates** $\backslash true$; – to prevent abrupt termination (i. e., calls to `exit()`, `abort()` etc.) — loop and recursion termination is already required by default in the AstraVer plugin.

An alternative approach enforcing the topological ordering of lemma functions has been considered: explicit removal of the axiom from the context of the functions located before the lemma body. This is possible to achieve using the `meta remove_prop` pragma in the generated Why3ML model. However, this would have required Frama-C to track the connection between the function and the axiom explicitly, and the AstraVer plugin to use that connection to generate the `meta` statement in the WhyML output. A minimal implementation that fully translates lemma functions into existing specification entities was chosen in the end.

IV. RESULTS

Even though at the previous stage of our project [15] we ended up with 32 unproved lemmas, we did not attempt to match our lemma function proofs at this stage exactly with the lemmas used at the previous stage. Instead, the goal was to ensure all the code can be verified without resorting to any additional assumptions by completely replacing all the lemmas involved with a minimum required set of lemma functions. Thus we ended up with 31 lemma functions required to fully prove the VCs for some code functions as presented in Table II.

Unfortunately, a thorough comparison of the proof overhead between the interactive and auto-active proofs turned out to be very hard. The current support for interactive proofs that existed in AstraVer plugin was quite basic, which resulted in considerably large manual proofs. The reduction of these proofs seems only possible with very significant implementation effort including implementation of relevant custom tactics and optimizing the VC generation for manual proofs. For the reference, we provide an illustrative comparison of an interactive auto-active proof of a single lemma in Figure 1 as it can be performed with the current implementation. This example shows the reduction in proof overhead of more than tenfold, yet a different more optimized manual proof augmented with appropriate

Function Name	Number of Ghost Functions	Number of Lemma Functions
strlen	4	2
skip_spaces	1	3
strchr	0	5
strchrnul	0	5
strspn	1	5
strcspn	0	3
strnlen	3	6
strpbrk	0	2
Total	9	31

TABLE II
LEMMA FUNCTIONS

custom tactics could potentially result in a much smaller number.

However, in the context of our project, the use of lemma functions became a decisive choice that enabled complete, practical, stable and future-proof verification of auxiliary lemmas with minimal implementation effort. The ability to instantiate lemmas by explicitly calling the corresponding lemma functions and the availability of all general C constructs in the proof code provided the verification engineer with sufficient control. By these means in our examples we were able to reduce the verification time of each lemma function (cumulative time of the solvers on all function’s VCs) to a few seconds.

V. OPEN ISSUES AND FUTURE WORK

A. Logic types and expressions in code

Unlike many specification languages originally developed with auto-active verification in mind, ACSL maintains a clear distinction between code and specification languages. In particular, C and logic types as well as C and logic expressions are represented entirely separately within the Frama-C kernel and are only allowed in their respective positions. Therefore lemma functions cannot be directly used to prove the lemmas with quantification over logic types. This limitation can be overcome in individual cases of quantified logic variables ranging over the domains bounded by C values.

For universal quantification of the form

$$\forall \mathbb{L} v; v_1 \leq v < v_2 \implies P(v),$$

where v_1 and v_2 are C values of type \mathbb{C} a proof can be accomplished through the use of an auxiliary lemma

$$\text{aux: } \forall \mathbb{C} v; v_1 \leq v < v_2 \implies P(v),$$

which can then be generalized using a dummy proof of the

```

1 Theorem Strchr_skipped :
2   ∀ voidP_str_0_30_alloc_table_at_L : alloc_table voidP,
3   ∀ charP_charM_str_0_30_at_L : map (pointer voidP) Int8.t,
4   ∀ str_0_0 : pointer voidP,
5   ∀ c_1_0 : Int8.t,
6   ∀ i_0_0 : Uint64.t,
7   valid_str
8     str_0_0
9     voidP_str_0_30_alloc_table_at_L
10    charP_charM_str_0_30_at_L ∧
11    strchr
12     str_0_0
13     c_1_0
14    charP_charM_str_0_30_at_L ≠ (null : pointer voidP) ∧
15    Uint64.infix_lseq (Uint64.of_int 0%Z) i_0_0 ∧
16    (Uint64.to_int i_0_0 < sub_pointer
17     (strchr
18      str_0_0
19      c_1_0
20      charP_charM_str_0_30_at_L)
21     str_0_0))%Z →
22
23    get
24     charP_charM_str_0_30_at_L
25     (shift str_0_0 (Uint64.to_int i_0_0)) ≠ c_1_0.
26 (* Why3 intros
27   voidP_str_0_30_alloc_table_at_L
28   charP_charM_str_0_30_at_L
29   str_0_0 c_1_0 i_0_0
30   (hl,(h2,(h3,(h4,h5))))). *)
31 intros
32 voidP_str_4_82_alloc_table_at_L
33 charP_charM_str_4_82_at_L
34 str_4_c_4_0_i_1
35 (hl,(h2,(h3,(h4,h5)))).
36
37 Definition P (i : int) :=
38   ∀ a : alloc_table voidP,
39   ∀ m : map (pointer voidP) Int8.t,
40   ∀ s : pointer voidP,
41   ∀ c : Int8.t,
42   valid_str s a m ∧
43   strchr s c m ≠ (null : pointer voidP) ∧
44   (0 ≤ i)%Z ∧
45   (i < sub_pointer (strchr s c m) s)%Z →
46   get m (shift s i) ≠ c.
47 specialize Wf_Z.natlike_rec3 with (P := P) as Ind.
48 assert (P 0) as Base.
49 unfold P.
50 intros a m s c Valid.
51 unfold not.
52 ...
53
54
55
56
57
58   rewrite First in Valid.
59   rewrite Sub_pointer_self in Valid.
60   omega.
61   assert (∀ z : int, (0 < z)%Z → P (Z.pred z) → P z) as Step.
62   intros z Gt0 Pred.
63   unfold P.
64   intros a m s c Valid.
65 ...
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114   apply Strchr_same_block with
115     (voidP_s_4_26_alloc_table_at_L := a).
116   tauto.
117   assert (∀ z : int, (0 ≤ z)%Z → P z) as Lemma.
118   apply Ind. 1,2: tauto.
119   unfold P in Lemma.
120   apply Lemma with (a := voidP_str_4_82_alloc_table_at_L).
121   unfold Uint64.infix_lseq in h3.
122   rewrite Uint64.Of_int in h3.
123 ...
124
125
126
127   unfold Uint64.in_bounds. omega.
128   tauto.
129 Qed.

```

Generated by Why3

```

1 /*@ ghost
2   @ /@ lemma
3   @ @ requires valid_str(str);
4   @ @ requires strchr(str,c) ≠ \null;
5   @ @ requires 0 ≤ i < strchr(str,c) - str;
6   @ @ decreases i;
7   @ @ ensures str[i] ≠ c;
8   @ @/
9   @ void strchr_skipped(char *str, char c,
10     size_t i)
11   @ {
12   @   if (i > 0 && *str != '\0' && *str != c)
13   @     strchr_skipped(str + 1, c, i - 1);
14   @ }
15   @*/

```

93 lines of user proof

13 lines

55 lines

8 lines

Fig. 1. Proofs of lemma strchr_skipped. Coq proof on the left. Auto-active proof with lemma function on the right.

```

form
  /@ lemma
  @ ensures  $\forall \mathbb{L} v; v_1 \leq v < v_2 \implies P(v);$ 
  @/
  void gen() {
    /@ loop invariant  $v_1 \leq i \leq v_2;$ 
    @ loop invariant  $\forall \mathbb{C} j; v_1 \leq j < i \implies P(j);$ 
    @/
    for ( $\mathbb{C} i = v_1; i < v_2; i++$ )
      aux(i);
  }

```

For logic expressions $l(\bar{u})$ ranging over the \mathbb{C} values $v_1 \leq l(\bar{u}) < v_2$, where \bar{u} are \mathbb{C} expressions, a proxy \mathbb{C} function can be introduced:

```

/*@ ensures \result  $\equiv l(\bar{u});$ 
  @ assigns \nothing;
  @*/
 $\mathbb{C}$  proxy( $\bar{u}$ );

```

Note that $l(\bar{u})$ may contain logic constructs that are thus encapsulated by the proxy. In general, proxies should be introduced carefully as their post-conditions are not checked and may introduce inconsistencies, but in some cases, the proxy can be given an implementation (definition) that satisfies the corresponding specification. In our proofs, we only introduced proxies with definitions.

Ultimately, the proper support for logic types and expressions in the ghost code should be added to ACSL specification and implemented in Frama-C to fully support auto-active proofs with lemma functions.

B. Lemma functions and axiomatics

Normally, lemmas in ACSL are grouped into axiomatic blocks. However, axiomatic blocks cannot include code. Therefore, lemma functions cannot be directly grouped into axiomatic blocks. Moreover, allowing lemma functions to be declared as lemmas in axiomatic blocks introduces a problem with ordering of proofs. To be well-founded, the lemma functions should be split into strongly connected components and topologically ordered. However, unlike code functions that can only be called explicitly, lemmas can be instantiated implicitly, provided they are preset in the proof context. When each lemma function has just a single definition, the definition order of lemma functions can serve as their required topological order, with an only additional restriction that mutually recursive functions should be somehow declared simultaneously. However, when lemma functions are allowed to have separate declarations and definitions, the ordering of lemma functions becomes generally unspecified. In the current implementation, the lemma functions are allowed to have separate declarations, but they are later grouped into strongly connected components, and the definition of the first function in the component is then used for topological ordering. This essentially simulates simultaneous definitions. However, this approach is completely ignorant of the lemma declarations

grouped into axiomatic blocks. Ultimately, instead of choosing between the declaration and the definition points to be used for topological ordering, we propose including the lemma function definitions (bodies) into the axiomatic blocks directly. The ordering then can be retained as in the current implementation. However, this would be a significant change to the ACSL specification.

C. Proof reuse and higher-order logic

Careful consideration of the resulting auto-active proofs from our study reveals that many of them are very similar to each other and represent instances of some more general proof patterns. In general, it is natural to express such more general lemmas and their proofs using higher-order logic (HOL). However, the support for higher-order reasoning in SMT-based deductive verification tools such as Jessie or AstraVer is usually very limited due to the inherent incompleteness of decision procedures for HOLs. An alternative solution to direct support for HOL would be a stratified solution, where the core specification language does not support any HOL constructs, but a separate higher-level language such as the language of modules does provide constructs that can be used to express and prove higher-order propositions. This approach is already used in Why3ML specification language. Let us illustrate this approach on two example lemmas from our study:

```

logic char * strchr(char * s, char c) =
  *s  $\equiv c$  ? s :
  *s  $\equiv '\backslash 0'$  ? (char*)\null : strchr(s + 1, c);
lemma strchr_skipped :
   $\forall$ char * s, c, size_t i;
  valid_str(s)  $\wedge$ 
  strchr(s, c)  $\not\equiv \backslash null$   $\wedge$ 
   $0 \leq i < strchr(s, c) - s \implies$ 
  s[i]  $\not\equiv c$ ;
  ...
logic char * strchrnul(char * s, char c) =
  *s  $\equiv c$  ? s :
  *s  $\equiv '\backslash 0'$  ? s : strchrnul(s + 1, c);
lemma strchrnul_skipped :
   $\forall$ char * s, c, size_t i;
  valid_str(s)  $\wedge$ 
   $0 \leq i < strchrnul(s, c) - s \implies$ 
  s[i]  $\not\equiv c$ ;

```

Currently, the lemmas have to be proved as two separate lemma functions. Though there is a single generic function that suits both cases, its definition makes use of higher-order features :

```

logic char * strchrcomb(boolean cond(char * s, char c),
  char * dft(char * s),
  char * s, char c) =
  cond(s, c) ? s : *s  $\equiv '\backslash 0'$  ? dft(s, c) : strchrcomb(s + 1, c);

```

Within the stratified approach, this function should be placed in a separate module alongside the functions corresponding to its parameters, auxiliary definitions, and

appropriate lemmas:

predicate $cond(\text{char} * s, \text{char} c)$ reads $* s$;
logic $\text{char} * dft(\text{char} * s)$;
predicate $dft_same = \forall \text{char} * s; dft(s) \equiv s$;
predicate $dft_null = \forall \text{char} * s; dft(s) \equiv \backslash null$;
logic $\text{char} * strchrcomb(\text{char} * s, \text{char} c) =$
 $cond(s, c) ? s :$
 $* s \equiv \backslash 0' ? dft(s) :$
 $strchrcomb(s + 1, c)$;
lemma $strchrcomb_skipped :$
 $\forall \text{char} * s, c, \text{size_t } i$;
 $valid_str(s) \wedge$
 $(dft_same \vee dft_null \wedge strchrcomb(s, c) \not\equiv \backslash null) \wedge$
 $0 \leq i < strchrcomb(s, c) - s \implies$
 $!cond(s + i, c)$;

Then the lemmas can be proved in the general case as usual, e. g., by using lemma-functions. The resulting module can be instantiated by providing concrete definitions for logic functions $cond$ and dft to obtain both logic functions $strchr$ and $strchrnul$. Upon instantiation, their corresponding lemmas do not have to be proven, and they can be safely instantiated directly into axioms.

VI. CONCLUSION

In this paper we propose the implementation of auto-active proofs technique in the Frama-C framework and AstraVer plugin, evaluate our implementation on VerKer project specifications by proving all string-related lemmas, and demonstrate the benefits of the approach by comparing it with manual lemma proofs in Coq.

The results of the work are publicly available. The source code with specifications and verification protocols are available as a part of the VerKer project. The proofs are easy reproducible since there is a configuration of a continuous integration system. The modified toolset is available on the AstraVer project page.

The obtained results show the approach is very promising and should be adopted in the mainline of AstraVer toolset. Auto-active verification in Frama-C ease the work of a verification engineer, make specifications more maintainable and potentially can lower the cost of deductive verification.

REFERENCES

- [1] F. Kirchner, N. Kosmatov, V. Prevosto, J. Signoles, and B. Yakobowski, “Frama-C: A software analysis perspective,” *Formal Aspects of Computing*, vol. 27, no. 3, pp. 573–609, May 2015. [Online]. Available: <https://doi.org/10.1007/s00165-014-0326-7>
- [2] B. Jacobs, J. Smans, P. Philippaerts, F. Vogels, W. Penninckx, and F. Piessens, “VeriFast: A powerful, sound, predictable, fast verifier for C and Java,” in *NASA Formal Methods Symposium*. Springer, 2011, pp. 41–55.
- [3] K. R. M. Leino, “Dafny: An automatic program verifier for functional correctness,” in *International Conference on Logic for Programming Artificial Intelligence and Reasoning*. Springer, 2010, pp. 348–370.
- [4] B. Carré and J. Garnsworthy, “SPARK — an annotated ada subset for safety-critical programming,” in *Proceedings of the conference on TRI-ADA’90*. ACM, 1990, pp. 392–402.

- [5] F. Rothenberger, “Integration and analysis of alternative SMT solvers for software verification,” Master’s thesis, ETH-Zürich, 2016.
- [6] G. Petiot, N. Kosmatov, B. Botella, A. Giorgetti, and J. Julliand, “How testing helps to diagnose proof failures,” *Formal Aspects of Computing*, Jun 2018. [Online]. Available: <https://doi.org/10.1007/s00165-018-0456-4>
- [7] M. Hentschel, R. Bubel, and R. Hähnle, “The symbolic execution debugger (SED): a platform for interactive symbolic execution, debugging, verification and more,” *International Journal on Software Tools for Technology Transfer*, 2018. [Online]. Available: <https://doi.org/10.1007/s10009-018-0490-9>
- [8] W. Reif, “The KIV-approach to software verification,” in *KO-RSO: Methods, Languages, and Tools for the Construction of Correct Software*. Springer, 1995, pp. 339–368.
- [9] C. Calcagno, D. Distefano, J. Dubreil, D. Gabi, P. Hooimeijer, M. Luca, P. O’Hearn, I. Papakonstantinou, J. Purbrick, and D. Rodriguez, “Moving fast with software verification,” in *NASA Formal Methods*, K. Havelund, G. Holzmann, and R. Joshi, Eds. Cham: Springer International Publishing, 2015, pp. 3–11.
- [10] A. Chudnov, N. Collins, B. Cook, J. Dodds, B. Huffman, C. MacCárthaigh, S. Magill, E. Mertens, E. Mullen, S. Tasiran, A. Tomb, and E. Westbrook, “Continuous formal verification of Amazon s2n,” in *Computer Aided Verification*, H. Chockler and G. Weissenbacher, Eds. Cham: Springer International Publishing, 2018, pp. 430–446.
- [11] P. W. O’Hearn, “Continuous reasoning: Scaling the impact of formal methods,” in *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science*, ser. LICS ’18. New York, NY, USA: ACM, 2018, pp. 13–25. [Online]. Available: <http://doi.acm.org/10.1145/3209108.3209109>
- [12] C. Dross and Y. Moy, “Auto-active proof of red-black trees in SPARK,” in *NASA Formal Methods*, C. Barrett, M. Davies, and T. Kahsai, Eds. Cham: Springer International Publishing, 2017, pp. 68–83.
- [13] J.-C. Filliâtre and A. Paskevich, “Why3 — where programs meet provers,” in *Proceedings of the 22nd European Symposium on Programming*, ser. Lecture Notes in Computer Science, M. Felleisen and P. Gardner, Eds., vol. 7792. Springer, Mar. 2013, pp. 125–128.
- [14] P. Baudin, P. Cuoq, J.-C. Filliâtre, C. Marché, B. Monate, Y. Moy, and V. Prevosto, “ACSL: ANSI/ISO C Specification Language,” CEA LIST and INRIA, Tech. Rep. 1.12, March 2017.
- [15] D. Efremov, M. Mandrykin, and A. Khoroshilov, “Deductive Verification of Unmodified Linux Kernel Library Functions,” in *Leveraging Applications of Formal Methods, Verification and Validation. Verification*, T. Margaria and B. Steffen, Eds. Cham: Springer International Publishing, 2018, pp. 216–234.
- [16] J. Hatcliff, G. T. Leavens, K. R. M. Leino, P. Müller, and M. Parkinson, “Behavioral interface specification languages,” *ACM Comput. Surv.*, vol. 44, no. 3, pp. 16:1–16:58, Jun. 2012. [Online]. Available: <http://doi.acm.org/10.1145/2187671.2187678>
- [17] J.-C. Filliâtre, “Verifying two lines of C with Why3: an exercise in program verification,” in *Proceedings of the 4th International Conference on Verified Software: Theories, Tools, Experiments*, ser. VSTTE’12. Berlin, Heidelberg: Springer-Verlag, 2012, pp. 83–97. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-27705-4_8
- [18] A. Tafat and C. Marché, “Binary heaps formally verified in Why3,” INRIA, Research Report 7780, Oct. 2011, <http://hal.inria.fr/inria-00636083/en/>.
- [19] A. Blanchard, N. Kosmatov, and F. Loulergue, “Ghosts for Lists: A Critical Module of Contiki Verified in Frama-C,” in *Tenth NASA Formal Methods Symposium - NFM 2018*, Newport News, United States, Apr. 2018. [Online]. Available: <https://hal.inria.fr/hal-01720401>
- [20] J. Burghardt, R. Clausecker, J. Gerlach, and H. Pohl, “ACSL By Example,” Fraunhofer Institute for Open Communication Systems, Tech. Rep., 2017.
- [21] K. R. M. Leino and M. Moskal, “Usable auto-active verification,” 2010.