

VST-A: A Foundationally Sound Annotation Verifier

LITAO ZHOU, Shanghai Jiao Tong University, China and University of Hong Kong, China

JIANXING QIN, Shanghai Jiao Tong University, China

QINSHI WANG, Princeton University, USA

ANDREW W. APPEL, Princeton University, USA

QINXIANG CAO*, Shanghai Jiao Tong University, China

Program verifiers for imperative languages such as C may be *annotation-based*, in which assertions and invariants are put into source files and then checked, or *tactic-based*, where proof scripts separate from programs are *interactively* developed in a proof assistant such as Coq. Annotation verifiers have been more automated and convenient, but some interactive verifiers have richer assertion languages and formal proofs of soundness. We present VST-A, an annotation verifier that uses the rich assertion language of VST, leverages the formal soundness proof of VST, but allows users to describe functional correctness proofs intuitively by inserting assertions.

VST-A analyzes control flow graphs, decomposes every C function into control flow paths between assertions, and reduces program verification problems into corresponding *straightline Hoare triples*. Compared to existing foundational program verification tools like VST and Iris, in VST-A such decompositions and reductions can be nonstructural, which makes VST-A more flexible to use.

VST-A's decomposition and reduction is defined in Coq, proved sound in Coq, and computed call-by-value in Coq. The soundness proof for reduction is totally logical, independent of the complicated semantic model (and soundness proof) of VST's Hoare triple. Because of the rich assertion language, not all reduced proof goals can be automatically checked, but the system allows users to prove residual proof goals using the full power of the Coq proof assistant.

CCS Concepts: • **Security and privacy** → **Logic and verification**; • **Theory of computation** → **Program verification**; *Hoare logic*; Automated reasoning; Parsing.

Additional Key Words and Phrases: Annotated Programs, Foundational Verification, Coq

ACM Reference Format:

Litao Zhou, Jianxing Qin, Qinshi Wang, Andrew W. Appel, and Qinxiang Cao. 2024. VST-A: A Foundationally Sound Annotation Verifier. *Proc. ACM Program. Lang.* 8, POPL, Article 69 (January 2024), 31 pages. <https://doi.org/10.1145/3632911>

1 INTRODUCTION

In the past 15 years, researchers have built several tools for program verification. These tools are used in different ways and have their own advantages.

Interactive program verification tools, such as Iris [Jung et al. 2018; Krebbers et al. 2017b] and VST [Appel 2011; Beringer 2021; Cao et al. 2018], are built in interactive theorem provers like Coq

*Corresponding author.

Authors' addresses: Litao Zhou, Shanghai Jiao Tong University, , China and University of Hong Kong, , China, tonyzhou0608@gmail.com; Jianxing Qin, Shanghai Jiao Tong University, , China, qdelta@sjtu.edu.cn; Qinshi Wang, Princeton University, , USA, wangqinshi1995@gmail.com; Andrew W. Appel, Princeton University, , USA, appel@princeton.edu; Qinxiang Cao, Shanghai Jiao Tong University, , China, caoqinxiang@gmail.com.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2024 Copyright held by the owner/author(s).

ACM 2475-1421/2024/1-ART69

<https://doi.org/10.1145/3632911>

or HOL4. Users write formal program correctness proofs in the same theorem prover, by using the lemmas and tactics of the program verification tool.

Some of these interactive tools themselves are *foundationally sound* (i.e., have a formal proof w.r.t. the language’s operational semantics in the proof assistant). This is especially meaningful for verifying real-world programs and higher-level properties such as functional correctness. First, real-world programming languages are complicated. For example, it is very subtle to determine what C programs may cause undefined behavior. Second, an advanced program logic for higher-order properties usually has a nontrivial soundness proof. For example, VST and Iris use step-indexed semantics to interpret impredicative assertion languages whose soundness proof is complicated.

Interactive program verification tools can also benefit from the rich logic language and the rich proof language of theorem provers (like Coq and HOL). These tools can easily shallow-embed higher-order functions and predicates in their assertion languages. They also make it convenient (in specifying programs and program assertions) to introduce additional logical connectives. Additionally, the tactic proof languages in proof assistants are very powerful when users need to describe proof strategies such as proof-by-induction and proof-by-contradiction.

A different strain of program verification tools requires programmers to write annotations in the source code. With sufficient annotations, tools like Dafny [Leino 2010], Hip/Sleek [Chin et al. 2012], VeriFast [Jacobs et al. 2011], Viper [Müller et al. 2017], Frama-C [Baudin et al. 2021] and CN [Pulte et al. 2023] can verify program correctness automatically. By restricting the assertion languages, tools like CBMC [Kroening and Tautschnig 2014], F-Soft [Ivančić et al. 2015, 2005], and Infer [Calcagno and Distefano 2011] can reduce the annotation overhead for programmers while preserving automation.

Compared with writing tactical proofs in a theorem prover, we believe that *writing annotations is a much more straightforward way of demonstrating that a program is correct*. Even proofs in theory papers and proofs written completely in an interactive prover are often presented in research papers as annotated programs, e.g. Reynolds’s first paper about separation logic uses annotations to describe separation logic proofs [Reynolds 2002, page 10], and recent verification papers like Jung et al. [2020]’s work extending Iris to support prophecy variables also uses annotations to describe their proofs [Jung et al. 2020, page 7]. Fig. 1 shows an implementation of an in-place linked list reversal and its functional correctness proof.¹ The annotations on lines 3-5 describe the specification this function should satisfy: for any list of integers l (the With clause on line 3), if l is stored in a linked list and this linked list’s head pointer is passed to the function by the program variable p (the Require clause on line 4), then the function reverses the linked list and returns the new head pointer (the Ensure clause on line 5). The assertion on line 9 describes the main idea of a functional correctness proof. It states the criteria that the program state should satisfy every time the program enters the loop body. Assertion-annotated programs can present proofs succinctly; by contrast, in interactive verifiers, key insights into program correctness easily get mixed with structural tactics and become lengthy proof scripts.

In this paper, we demonstrate how to combine the benefits of interactive tools and annotation verifiers. We present VST-A, a foundationally sound verification tool, that is implemented on top of VST. VST-A enjoys rich assertion languages and flexible proof strategies, and it allows users to write readable assertion annotations directly as comments exactly as in Fig. 1.

¹This is a separation logic [Reynolds 2002] proof. In other words, we use an assertion of form “ $P * Q$ ” to say that the memory can be split into two disjoint pieces of which one satisfies P and the other satisfies Q . $\text{ll}(p, l)$ is a separation logic predicate that asserts on the location referenced by variable p stores a linked list of l .

```

1 struct list {unsigned head; struct list *tail;};
2 struct list *reverse (struct list *p) {
3     /*@ With l,
4         Require ll(p, l)
5         Ensure ll(ret, rev(l)) */
6     struct list *w, *t, *v;
7     w = NULL; v = p;
8     while (v) {
9         /*@ Assert  $\exists l_1 u x l_2. l = \text{rev}(l_1) \ x \ l_2 \wedge v \mapsto (x, u) * ll(w, l_1) * ll(u, l_2) */$ 
10        t = v->tail; v->tail = w; w = v; v = t; }
11    return w; }

```

Fig. 1. Annotations for verifying linked-list reversal




Fig. 2. Verification workflow in VST-A

We illustrate the VST-A workflow in Fig. 2: (1) Users first provide a C program with assertion annotations. (2a) Our front-end parser then converts the source code into ClightA, the Coq representation of this annotated C language. (2b) Next, the C program's functional correctness is reduced to smaller proof goals using the annotations. Specifically, a *split function* accepts a ClightA program and its pre-/post-conditions as input, and returns a set of straightline Hoare triples, each of which consists of a sequence of primary statements (assignment statements, function calls) and/or **assume** commands. For example, Fig. 3 shows the split result of the reverse function in Fig. 1: four triples are returned as verification goals. As illustrated by the control flow graph in Fig. 4, the functionality of this *split function* is natural; it computes all of the control flow paths that are separated by assertion annotations in the source program. (3) Finally, users are left to prove each straightline Hoare triple in the split result. (4) The VST-A soundness theorem ensures the correctness of the original program if all of the paths have been verified.

In summary, the contributions of this paper are:

- (1) **A new framework for program verification** that combines the benefits of interactive provers and the readability of annotated programs.

(1) $\forall l,$
 $\{ \llbracket p, l \rrbracket \}$
 $w = \text{NULL}; v = p; \text{assume } v;$
 $\{ \exists l_1 c x l_2. l = \text{rev}(l_1) x l_2 \wedge$
 $v \mapsto (x, c) * \llbracket w, l_1 \rrbracket * \llbracket c, l_2 \rrbracket \}$

(2) $\forall l,$
 $\{ \llbracket p, l \rrbracket \}$
 $w = \text{NULL}; v = p; \text{assume } !v; \text{ret} = w;$
 $\{ \llbracket \text{ret}, \text{rev}(l) \rrbracket \}$

(3) $\forall l_1 c x l_2,$
 $\{ l = \text{rev}(l_1) x l_2 \wedge v \mapsto (x, c) * \}$
 $\{ \llbracket w, l_1 \rrbracket * \llbracket c, l_2 \rrbracket \}$
 $t = v \rightarrow \text{tail}; v \rightarrow \text{tail} = w;$
 $w = v; v = t; \text{assume } v;$
 $\{ \exists l_1 c x l_2. l = \text{rev}(l_1) x l_2 \wedge$
 $v \mapsto (x, c) * \llbracket w, l_1 \rrbracket * \llbracket c, l_2 \rrbracket \}$

(4) $\forall l_1 c x l_2,$
 $\{ l = \text{rev}(l_1) x l_2 \wedge v \mapsto (x, c) * \}$
 $\{ \llbracket w, l_1 \rrbracket * \llbracket c, l_2 \rrbracket \}$
 $t = v \rightarrow \text{tail}; v \rightarrow \text{tail} = w;$
 $w = v; v = t; \text{assume } !v;$
 $\text{ret} = w;$
 $\{ \llbracket \text{ret}, \text{rev}(l) \rrbracket \}$

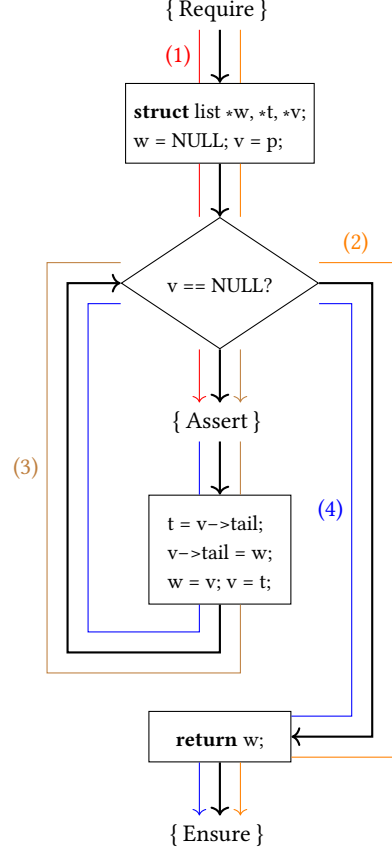


Fig. 3. Split results for verifying linked-list reversal Fig. 4. Control flow graph of linked-list reversal

- (2) **A formal language for annotated programs:** We define ClightA, a formal language for annotated C programs, as a method of describing how the functional correctness proof for a large program can be reduced.
- (3) **A control-flow-based verification splitting algorithm:** The VST-A proof reduction framework uses a split function that is implemented in Coq and proved sound w.r.t. the VST program logic. We believe this split algorithm and its soundness proof are general and can be applied to other Hoare-style imperative verification tools as well.

In the rest of this paper: We will introduce background information about VST-A in §2. We will present our annotation-based proof language and analyze its expressiveness in §3. We will define the split function and prove it sound in §4. We will discuss the connection between our soundness proof and Hoare logic’s conjunction rule in §5. We put statistics of VST-A verification examples in §6. Finally, we will discuss related works in §7 and conclude in §8.

2 BACKGROUND

We used Coq [Boutillier et al. 2014] and VST to implement our annotation verifier, VST-A. VST is an interactive program verification tool that is built in Coq. Its primary components are

can be interpreted as starting from a program state that satisfies P . After executing c , if the statement exits normally, the program state satisfies Q . Similarly, the program state should satisfy Q_{brk} , Q_{con} , and Q_{ret} when the statement exits with break, continue, or return, respectively. We use \vec{Q} as an abbreviation of the last three postconditions.

In Verifiable C, most of the compositional rules are standard. Fig. 6 shows some representative ones.⁴ To fit the general loop syntax of Clight, the SEMAX-LOOP rule has two invariants, loop invariant I and continue invariant I_{con} . The loop invariant I is required to hold before each iteration, and I_{con} is required to hold at continue or before c_{incr} statements. Verifiable C's proof rules for primary statements are less important to end users – VST provides verified forward symbolic execution (§2.4), and thus VST's users do not need to use those rules directly. All proof rules in Verifiable C are proved sound foundationally w.r.t. the CompCert Clight semantics, and symbolic execution applies the proof rules. VST also provides some useful derived rules, a few of which are shown in Fig. 7. The SEQ-ASSOC rule reassociates sequential compositions, and rules EXTRACT-EXISTS and EXTRACT-PURE introduce variables and propositions from precondition to context, respectively.

$$\text{EXTRACT-PURE} \frac{\text{pure}(P_{\text{pure}}) \quad P_{\text{pure}} \Rightarrow \{P\} c \{Q, [\vec{Q}']\}}{\{P_{\text{pure}} \wedge P\} c \{Q, [\vec{Q}']\}}$$

$$\text{EXTRACT-EXISTS} \frac{\forall (x : A). \{P\} c \{Q, [\vec{Q}']\}}{\{\exists x : A. P\} c \{Q, [\vec{Q}']\}}$$

$$\text{SEQ-ASSOC} \frac{\{P\} c_1; (c_2; c_3) \{Q, [\vec{Q}']\}}{\{P\} (c_1; c_2); c_3 \{Q, [\vec{Q}']\}}$$

Fig. 7. Derived rules from C Hoare logic

2.3 Inversion Rules and Weakest Preconditions

The VST program logic is higher-order, i.e., assertions can quantify over assertions, so that one can state the following inversion lemmas, which have already been proven in VST.⁵

LEMMA 1 (INVERSION ON SEQUENCING). *If $\{P\} c_1; c_2 \{Q, [\vec{Q}']\}$, then*

$$\{P\} c_1 \{ \exists R : \text{assert. } R \wedge (\{R\} c_2 \{Q, [\vec{Q}']\}), [\vec{Q}'] \}$$

LEMMA 2 (INVERSION ON IF-BRANCHING). *If $\{P\} \text{if } (b) c_1 \text{ else } c_2 \{Q, [\vec{Q}']\}$, then*

$$P \models \exists P' : \text{assert. } P' \wedge (\{P' \wedge b \neq 0\} c_1 \{Q, [\vec{Q}']\}) \wedge (\{P' \wedge b = 0\} c_2 \{Q, [\vec{Q}']\})$$

It is worth mentioning that the normal postcondition appearing in the inversion on sequencing is VST's representation of weakest precondition, i.e.

$$\text{wp} \left(c, Q, [\vec{Q}'] \right) \triangleq \exists R : \text{assert. } R \wedge (\{R\} c \{Q, [\vec{Q}']\}).$$

⁴A full list of rules can be found in the extended version of the paper [Zhou et al. 2023a].

⁵Readers who knew that VST's separation Hoare logic was proved sound with a semantic proof in a shallow-embedded style may be surprised that it is possible to prove inversion lemmas. But in fact, several years ago the VST-Floyd developers layered a deep-embedded Hoare logic [Wang et al. 2023] over the shallow-embedded Hoare logic just so that useful lemmas of this kind can be supported.

This definition of weakest precondition satisfies basic properties like the following:

THEOREM 3. $P \vDash \text{wp}(c, Q, [\vec{Q}'])$ if and only if $\{P\} c \{Q, [\vec{Q}']\}$.

Thus, lemma 1 can be restated as:

$$\{P\} c_1; c_2 \{Q, [\vec{Q}']\} \text{ iff. } \{P\} c_1 \{\text{wp}(c_2, Q, [\vec{Q}']), [\vec{Q}']\}.$$

We utilize higher-order assertions in our VST-A development.

2.4 Forward Symbolic Execution

VST's forward verification tactics enable users to obtain the strongest postcondition of a sequence of primary statements automatically. For example, in the verification of path (3) in Fig. 3, the symbolic assignment executor $\text{ae}(P, c)$ for precondition P and statement c can compute the following (where ll is the linked-list predicate):

$$\begin{aligned} P &:= l = \text{rev}(l_1) \ x \ l_2 \ \wedge \ v \mapsto (x, u) * \text{ll}(w, l_1) * \text{ll}(u, l_2) \\ c &:= t = v \rightarrow \text{tail} \\ \text{ae}(P, c) &= l = \text{rev}(l_1) \ x \ l_2 \ \wedge \ t = u \ \wedge \ v \mapsto (x, u) * \text{ll}(w, l_1) * \text{ll}(u, l_2) \end{aligned}$$

The assignment executor may fail, if the precondition P cannot guarantee that c will run safely or P is not in a good form⁶ such that the symbolic executor can execute c , but we have found if users write assertions in their annotations, corresponding to correct proofs, the symbolic executor can run through the entire straightline Hoare triple. Users are left to prove the entailment from the inferred strongest postcondition (of the straight-line code) to the specified postcondition (arising from the “next” annotation, according to the split function). VST-Floyd also provides useful tactics for solving the entailment problem. Residual proof goals may be generated if some entailments cannot be automatically proven, and users can write their own flexible Coq proof scripts to address them. In summary, with the help of VST-Floyd, the back-end verification of the split results obtained by VST-A can be largely automated.

3 VST-A FRONT END

3.1 Annotated C Programs and Internal Representations

VST-A requires users to describe C function specifications and C programs' functional correctness proofs by writing annotations in C programs. Specifically, a function specification in VST-A is always in a `/*@ With ... Require ... Ensure ... */` form and

$$\text{With } [\vec{x} : \vec{A}] \text{ Require } P(\vec{x}) \text{ Ensure } Q(\vec{x})$$

represents a parameterized pre-/postcondition, i.e. it states that for any list of values \vec{x} of type \vec{A} , if the initial program state satisfies $P(\vec{x})$ then the C function can be safely executed (no C undefined behavior will happen). If it terminates, the ending state satisfies $Q(\vec{x})$.

For functional correctness proofs, users can describe the main proof skeleton by inserting assertions (including loop invariants) and “given” annotations in the source program. We formally define this annotated C language (Fig. 8), namely ClightA, and implement a front-end parser that converts annotated C programs into the ClightA abstract syntax. Compared with the Clight syntax, ClightA has two new components: assertions and ExGiven structures.

⁶Generally speaking, to ensure that an assignment executor $\text{ae}(P, c)$ always succeeds, P should be in the form of a symbolic heap assertion. When c is a load/store statement, there should be an explicit mapsto predicate for the manipulated variables in the separating conjunction clauses. Otherwise, one might need to apply the rule of consequence (and prove an entailment) to put P into that form.

$$\begin{array}{l}
\text{assertion : } P := \dots \\
\text{ClightA statements : } C := c_p \mid C_1; C_2 \mid \text{if } (e) C_1 \text{ else } C_2 \mid \text{loop } (C_2) C_1 \\
\quad \mid \text{skip} \mid \text{break} \mid \text{continue} \mid \text{return} \\
\quad \mid \text{assert } P \mid \text{ExGiven } x : A, \{P(x)\} C \\
\text{Annotation erasing : } C \Downarrow \in \text{Clight statement}
\end{array}$$

Fig. 8. ClightA: abstract C language for annotated programs

As for assertions, users can insert them anywhere by writing `/*@ Assert ... */` in the source program, which is directly converted into a leaf node in the ClightA syntax tree. Annotating loop structures with invariants is *not compulsory* in VST-A,⁷ but users can still write loop invariants as `/*@ Inv ... */` in C source files, to distinguish from assertions before loops. Our front-end automatically converts such annotations into the general syntax of ClightA.

As for the ExGiven structures, users can use a combination of `/*@ Assert $\exists x, \dots$ */` and `/*@ Given x */` to represent logical variable introduction in a Hoare logic proof. In a typical goal-directed proof strategy, one can extract existential variables from the precondition into the proof context (see EXTRACT-EXIST in Fig. 7). Then assertions that appear later in the focused proof can refer to the extracted variables as ordinary Coq assumptions. VST-A supports this proof method by defining the “ExGiven $x : A, \{P(x)\} C$ ” syntax. The syntax indicates that assertion P is existentially quantified by logical variable x . Moreover, the inner ClightA statement is also quantified by x , so within the (annotated) assertions of C , one can mention x .

Fig. 9 is a comparison between annotated C programs and ClightA syntax, showing how our front-end parser translates annotations into AST constructions.

<pre> /*@ Inv P₁ */ while (b) { c₁; /*@ Assert $\exists x, P_2(x)$ Given x */ c₂; /*@ Assert P₃(x) */ c₃; } </pre>	<pre> loop (skip) { assert(P₁); if (b) { skip; } else { break; } c₁; ExGiven x, P₂(x) { c₂; assert(P₃(x)); c₃; } } </pre>
---	---

Fig. 9. Annotated C program v.s. ClightA syntax

3.2 Expressiveness: Describing Hoare logic Proofs in VST-A

It is not surprising that the ClightA language is expressive enough to describe the main structures of Hoare logic proofs. Hoare logic proof rules decompose the verification target into smaller ones (e.g., as in Fig. 6 in §2). Describing such proofs in VST-A is natural. For example, to apply the sequence rule SEMAX-SEQ, we can insert the middle condition as an assertion annotation into the C program.

⁷That is, instead of an invariant at the beginning of the loop body, it may be more convenient to write assertion(s) elsewhere in the loop body, sufficient to break the control flow into straight-line segments; see §3.4.

3.3 Expressiveness: Describing Interactive Proofs in VST-A

ClightA is also expressive enough to describe interactive proofs used by existing verification tools like VST. In most cases, these tools try to find proof rules to apply using a *goal-directed strategy*. For example:

- In order to verify a Hoare triple of form $\{P\}x = e; c \left\{ Q, \left[\vec{Q}' \right] \right\}$ in VST, its forward symbolic execution tactic applies the sequence rule $\text{SEM}_{\text{MAX-SEQ}}$ and uses the strongest postcondition of P and $x = e$ as the middle condition. VST-A's users do not need to write any annotation to describe such proof strategy.
- In order to verify a Hoare triple of form

$$\{P\}\text{while}(e) \{c_1\}; c_2 \left\{ Q, \left[\vec{Q}' \right] \right\}$$

in VST (suppose no break statements appear in c_1), it asks users to provide a loop invariant I and then $I \wedge e = 0$ will be used as the middle condition between the loop and c_2 . In VST-A, the proof effort is similar. Users only need to provide such a loop invariant I in annotated C programs.

- In order to verify a Hoare triple whose precondition is existentially quantified, one will typically extract that existential variable into the proof context. The “given” annotation in VST-A describes this proof strategy.

Besides these goal-directed proof steps, users of interactive program-verification tools may use the rule of consequence anywhere in the middle of a proof, to replace a precondition with a weaker assertion (Fig. 10 shows an example). Correspondingly, users of VST-A can write an assertion annotation to “invoke” the rule of consequence.

$$\begin{array}{l} \{\ll(p, l) \wedge p \neq \text{NULL}\} \\ p \rightarrow \text{head} = 0 \\ \{\exists l'. \ll(p, l')\} \end{array} \quad \text{can be rewritten as} \quad \begin{array}{l} \{\exists q x l'. p \mapsto (x, q) * \ll(q, l')\} \\ p \rightarrow \text{head} = 0 \\ \{\exists l'. \ll(p, l')\} \end{array}$$

Fig. 10. Example of replacing the precondition with an existentially quantified assertion. The predicate $\ll(p, l)$ means that a linked list starts at address p representing (in the heads of its cons cells) the sequence l . The two preconditions above are provably equivalent.

In summary, typical interactive proofs are structural, following the syntax tree of the C statements (with local uses of existential variable extraction or the rule of consequence). The ClightA language is able to describe such structural proofs.

3.4 Expressiveness: Describing Nonstructural Proofs in VST-A

ClightA can also describe nonstructural proofs, those that do not follow the C syntax tree. Allowing users to write nonstructural proofs is a considerable convenience. For example, in order to verify a Hoare triple of form

$$\{P\}\text{if}(e) \{c_1\} \text{else} \{c_2\}; c_3 \left\{ Q, \left[\vec{Q}' \right] \right\}$$

in an interactive verification tool, users will probably be asked to provide a join condition after the if-statement (a precondition for c_3). In some cases this is appropriate and convenient, but sometimes it is both difficult and unnecessary. Here are some typical scenarios.

- The if-then branch c_1 is the break statement. In this case, it suffices to prove

$$\{P \wedge e \neq 0\}c_1 \{\perp, [Q_{\text{brk}}, \perp, \perp]\} \quad \text{and} \quad \{P \wedge e = 0\}c_2; c_3 \left\{ Q, \left[\vec{Q}' \right] \right\},$$

and users may hope to symbolically execute c_2 so that the middle condition between c_2 and c_3 can be generated instead of manually provided.⁸

- The statement c_3 is very short. In this case, verifying c_3 twice can be less work than writing down a join condition.
- Proving c_3 functionally correct needs very different proof strategy when e has a different boolean value. In this case, a join condition does not help to reduce the workload: in effect, one is case-splitting on e and then proving c_3 twice.

Similarly, in order to verify a Hoare triple of form

$$\{P\}\text{while } (e) \{ c_1 \}; c_2 \{ Q, [\vec{Q}'] \}$$

where break statements appear in some branches in the loop body c_1 , a join condition is needed (in addition to a loop invariant) in an interactive proof, since an execution may leave the loop by a break statement or by falsifying the loop condition e . In VST-A, users can choose to provide a join condition (the proof will be structural) or not to provide a join condition (the proof will be nonstructural).

Certain kinds of generalized loop invariants lead to nonstructural proofs. Consider a red-black tree (RBT) algorithm that reestablishes red-black invariants after insertion. Fig. 11 is a straightforward textbook implementation of this algorithm [Cormen et al. 2022]. When the rotation on line 15 is done, the loop exits immediately, since the assignment statement on line 14 ensures that the loop condition will be evaluated to false in the next iteration. If we were to write a loop invariant on line 3 before the loop condition is checked, we then need to both state the RBT bottom-up fixing invariant and describe the case in which the RBT has been fixed by the final rotation and should exit immediately. One may instead expect a single invariant on line 5, and reason separately about the control flow where the loop exits after a rotation. Such proof strategy is unavailable in common goal-directed verifiers, in which loop invariants are compulsory, but it is supported by VST-A.

4 CONTROL FLOW SPLITTING AND SOUNDNESS

One of the most important components in VST-A is the verified split function which reduces an entire C program’s functional correctness to a series of straightline Hoare triples, based on the annotations. Intuitively, this split function is a CFG-based computation (like our demonstration in Fig. 4), and its soundness must ultimately relate to C’s small-step semantics—any execution trace of the program statement can be decomposed into these separated paths. However, we choose to prove this soundness theorem directly using VST’s program logic (*Verifiable C*), instead of proving it indirectly by first showing the execution trace decomposition lemma that we mentioned above. Our main consideration is: it is nontrivial to formally establish a theoretical connection between a program logic and an operational semantics, especially when a complicated program logic for a realistic programming language with a lot of subtleties are considered. VST has already done that once—Verifiable C’s soundness proof takes 60K lines of Coq definitions and proofs. If we would choose to prove the split function sound using small-step semantics as intermediate proof steps, we might need to develop similar lengthy proofs.

Even in applications to other languages and other operational semantics, it will be useful to build annotation verifiers on top of program logics in the way that we present here, rather than try to prove soundness directly from operational semantics. Proved-sound verification tools tend to be based on similar program logics (at least in terms of the core rules targeted by our split function: the sequence rule, the consequence rule, etc.). But they may be based on quite different styles

⁸VST users do not need to provide a join condition in such cases since 2018. VST implemented this feature by adding a built-in program transformation to its verification tactics.

```

1 void insert_balance(struct tree *p, struct tree *root) {
2   ... // pre-processing code omitted
3   /*@ Inv: ... (RBT invariant)  $\vee$  ... (loop exit property) */
4   while (p != root && p->parent->color != RED) {
5     /*@ Assert: ... (RBT invariant) */
6     struct tree *p_par = p->parent, *p_gpar = p_par->parent;
7     if (p_par == p_gpar->left) { // p's parent is a left child
8       struct tree *p_uncle = p_gpar->right;
9       if (p_uncle->color == RED) {
10        p_par->color = BLACK; p_uncle->color = BLACK;
11        p_gpar->color = RED; p = p_gpar;
12      } else {
13        if (p == p_par->right) { p = p_par; left_rotate(p, root); }
14        p->parent->color = BLACK; p_gpar->color = RED;
15        right_rotate(p->parent->parent, root);
16      } else { ... } // dual case where p's parent is a right child, omitted
17    } }

```

Fig. 11. The fix-up function of red-black tree insertion

of operational semantics (e.g. imperative HOL [Bulwahn et al. 2008] uses big step semantics and CompCert Clight uses small step semantics), or they may (like VST [Appel et al. 2014] or Iris [Jung et al. 2018]) incorporate modal impredicativity.

Since most Hoare logic proof rules are syntax-oriented, we implement our split function through recursion on ClightA syntax tree (§4.2) and then we prove it sound by induction (§4.3). In the rest part of this section, we will start from defining the Coq type of split results (§4.1).

4.1 The Type of Split Result

As illustrated in Fig. 12, control flow paths between the assertions can be divided into four classes:

- (1) head paths – control flow paths from the precondition to an internal assertion;
- (2) tail paths – control flow paths from an internal assertion to the postcondition;
- (3) full paths – control flow paths between two internal assertions;
- (4) assertion-free paths – control flow paths from the precondition to the postcondition (we call them *assertion-free* since they pass through no assertions inside the annotated program).

Formally, the split result is a record that consists of “head/tail paths”, “full paths” and “assertion-free paths”, which are essentially a list of *basic program statements* \vec{c}_i , annotated with one single assertion, two assertions, and no assertions, respectively. A basic statement can either be a primary Clight statement, c_p , or a special statement, *assume e*, that represents an if-condition (positively or negatively) in the control flow.

Recall that in the VST program logic, a Hoare triple has multiple postconditions for different kinds of program exits (i.e., exit by break, by continue, by return, or normal fall-through). Correspondingly, the split result also makes distinctions among the different exits. Thus, in the definition of our intermediate split result, the record contains one set of “full paths” between the annotated assertions, one set of “head paths” from the entry point to the internal assertions, four sets of “tail paths” from the internal assertions to the four different kinds of exits, and four sets of “assertion-free paths” from the entry point to the four different kinds of exits. To handle existential variables in

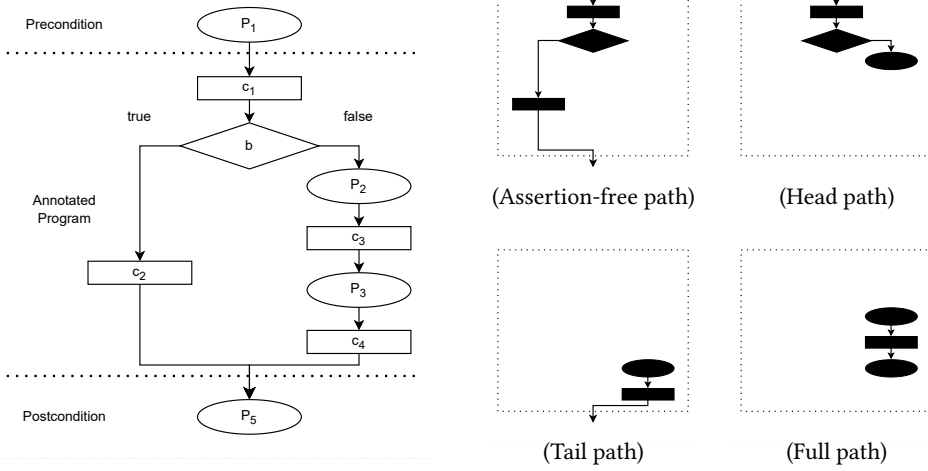


Fig. 12. Control flow graphs v.s. paths in split results

their scope, full paths can be universally quantified. With these fields, the split result record can sufficiently reveal all control flow information in a ClightA program. Fig. 13 shows the definition.

$$\begin{aligned}
 \text{Basic statement : } c_b &:= c_p \mid \text{assume } e \\
 \text{Assertion-free paths : } p_- &:= \vec{c}_b \\
 \text{Head paths : } p_{\dashv} &:= \vec{c}_b \{P\} \\
 \text{Tail paths : } p_{\vdash} &:= \{P\} \vec{c}_b \\
 \text{Full paths : } p_{\dashv\vdash} &:= \{P_1\} \vec{c}_b \{P_2\} \mid \forall (x : A). p_{\dashv\vdash}
 \end{aligned}$$

$$\text{split}(C) = \left\{ \begin{array}{l} p_-^{\text{nor}}, p_-^{\text{brk}}, p_-^{\text{con}}, p_-^{\text{ret}}, \\ p_{\vdash}^{\text{nor}}, p_{\vdash}^{\text{brk}}, p_{\vdash}^{\text{con}}, p_{\vdash}^{\text{ret}}, \\ p_{\dashv\vdash} \end{array} \right\}, \text{ where } \begin{array}{l} p_-^{\text{nor}}, p_-^{\text{brk}}, p_-^{\text{con}}, p_-^{\text{ret}} \subseteq \text{Assertion-free paths,} \\ p_{\vdash}^{\text{nor}}, p_{\vdash}^{\text{brk}}, p_{\vdash}^{\text{con}}, p_{\vdash}^{\text{ret}} \subseteq \text{Tail paths,} \\ p_{\dashv\vdash} \subseteq \text{Head paths, } p_{\dashv\vdash} \subseteq \text{Full paths} \end{array}$$

Fig. 13. The type of split results

By supplementing “head/tail paths” or “assertion-free paths” with the pre-/postconditions, we can interpret the split result into a collection of closed Hoare triples as hypotheses of split’s soundness theorem (these hypotheses are illustrated in Fig. 12 and formally defined in Fig. 14):

THEOREM 4 (SOUNDNESS). *For any ClightA program C and pre-/post-conditions $P, Q, Q_{\text{brk}}, Q_{\text{con}}$ and Q_{ret} , if $\text{split}(C) = \{p_-^{\text{nor}}, p_-^{\text{brk}}, p_-^{\text{con}}, p_-^{\text{ret}}, p_{\vdash}^{\text{nor}}, p_{\vdash}^{\text{brk}}, p_{\vdash}^{\text{con}}, p_{\vdash}^{\text{ret}}, p_{\dashv\vdash}\}$ and*

- all straightline Hoare triples from the precondition P to internal assertions are provable,*
 - all straightline Hoare triples from internal assertions to the postconditions \vec{Q} are provable,*
 - all straightline Hoare triples between internal assertions are provable,*
 - all straightline Hoare triples from the precondition P to the postconditions \vec{Q} are provable,*
- i.e. (defined in Fig. 14),*

- HeadHypo(P, p_{\dashv}),*
- TailHypo($Q, p_{\vdash}^{\text{nor}}$), TailHypo($Q_{\text{brk}}, p_{\vdash}^{\text{brk}}$), TailHypo($Q_{\text{con}}, p_{\vdash}^{\text{con}}$) and TailHypo($Q_{\text{ret}}, p_{\vdash}^{\text{ret}}$),*
- FullHypo($p_{\dashv\vdash}$),*

$$\begin{aligned}
\text{HeadHypo}(P, \vec{c}_b\{Q\}) &= \{P\} \vec{c}_b \{Q, [\bar{\top}]\} \\
\text{HeadHypo}(P, \mathbf{p}_\downarrow) &= \bigwedge_{p \in \mathbf{p}_\downarrow} \text{HeadHypo}(P, p) \\
\text{AssnFreeHypo}(P, Q, \vec{c}_b) &= \{P\} \vec{c}_b \{Q, [\bar{\top}]\} \\
\text{AssnFreeHypo}(P, Q, \mathbf{p}_-) &= \bigwedge_{p \in \mathbf{p}_-} \text{AssnFreeHypo}(P, Q, p) \\
\text{TailHypo}(Q, \{P\}\vec{c}_b) &= \{P\} \vec{c}_b \{Q, [\bar{\top}]\} \\
\text{TailHypo}(Q, \mathbf{p}_\uparrow) &= \bigwedge_{p \in \mathbf{p}_\uparrow} \text{TailHypo}(Q, p) \\
\text{FullHypo}(\{P\}\vec{c}_b\{Q\}) &= \{P\} \vec{c}_b \{Q, [\bar{\top}]\} \\
\text{FullHypo}(\forall x. \mathbf{p}_{\downarrow}) &= \forall x. \text{FullHypo}(\mathbf{p}_{\downarrow}) \\
\text{FullHypo}(\mathbf{p}_{\downarrow}) &= \bigwedge_{p \in \mathbf{p}_{\downarrow}} \text{FullHypo}(p)
\end{aligned}$$

Fig. 14. Hypotheses of the soundness theorem

(d) $\text{AssnFreeHypo}(P, Q, \mathbf{p}_-^{\text{nor}})$, $\text{AssnFreeHypo}(P, Q_{\text{brk}}, \mathbf{p}_-^{\text{brk}})$, $\text{AssnFreeHypo}(P, Q_{\text{con}}, \mathbf{p}_-^{\text{con}})$ and $\text{AssnFreeHypo}(P, Q_{\text{ret}}, \mathbf{p}_-^{\text{ret}})$,

then $\{P\} C \Downarrow \{Q, [Q_{\text{brk}}, Q_{\text{con}}, Q_{\text{ret}}]\}$, where $C \Downarrow$ represents the result of erasing all annotations from C .

Noticing that all of the control flows in a function body should end with a return statement, we directly use the following corollary in VST-A.

COROLLARY 5. For any *ClightA* program C and pre-/post-conditions P and Q , if

$$\text{split}(C) = \{\emptyset, \emptyset, \emptyset, \mathbf{p}_-^{\text{ret}}, \emptyset, \emptyset, \emptyset, \mathbf{p}_+^{\text{ret}}, \mathbf{p}_\downarrow, \mathbf{p}_{\downarrow}\}$$

and (a) $\text{HeadHypo}(P, \mathbf{p}_\downarrow)$, (b) $\text{TailHypo}(Q, \mathbf{p}_+^{\text{ret}})$, (c) $\text{FullHypo}(\mathbf{p}_{\downarrow})$, and (d) $\text{AssnFreeHypo}(P, Q, \mathbf{p}_-^{\text{ret}})$, then $\{P\} C \Downarrow \{Q\}$.

Remark. CompCert *Clight* does not have an assume statement. We choose to encode the assume statement into *Clight* AST, and encode straightline Hoare triples into VST Hoare triples, so that VST-A's users can directly use VST's tactics to prove those straightline triples. Specifically,

assume $e \triangleq$ **if** (e) **skip** **else** **break**;

$$\{P\} c_1; c_2; \dots; c_n \{Q\} \triangleq \{P\} c_1; c_2; \dots; c_n \{Q, [\top, \perp, \perp]\}.$$

We proved:

LEMMA 6. For straightline Hoare triples, $\{P\}$ **assume** e ; $c \{Q\}$ if and only if $\{P \wedge e \neq 0\} c \{Q\}$.


4.2 Split Function

The core split function is defined by recursion on the abstract syntax tree of the input *ClightA* program. Note that the split function is a partial function, since we will not try to compute the reduction result if there is an assertion-free loop in the CFG.

Base cases (Fig. 15). For primary statements c_p , the normal assertion-free path set (the $\mathbf{p}_-^{\text{nor}}$ field) is a singleton of the statement itself, i.e. $\{[c_p]\}$, and all other path sets are empty. For a break statement, the assertion-free break-exit path set is a singleton of an empty list of basic statements, i.e. $\{[]\}$. The split results of continue and return are similar. For the assertion annotation **assert** P , the split result has only one head path $\{[]\{P\}\}$ and one normal tail path $\{\{P\}\{[]\}\}$.

$$\begin{aligned}
\text{split}(c_p) &= \left\{ \begin{array}{l} \{[c_p]\}, \emptyset, \emptyset, \emptyset \\ \emptyset, \emptyset, \emptyset, \emptyset \\ \emptyset, \emptyset \end{array} \right\} & \text{split}(\text{break}) &= \left\{ \begin{array}{l} \emptyset, \{[]\}, \emptyset, \emptyset \\ \emptyset, \emptyset, \emptyset, \emptyset \\ \emptyset, \emptyset \end{array} \right\} & \text{split}(\text{return}) &= \left\{ \begin{array}{l} \emptyset, \emptyset, \emptyset, \{[]\} \\ \emptyset, \emptyset, \emptyset, \emptyset \\ \emptyset, \emptyset \end{array} \right\} \\
\text{split}(\text{continue}) &= \left\{ \begin{array}{l} \emptyset, \emptyset, \{[]\}, \emptyset \\ \emptyset, \emptyset, \emptyset, \emptyset \\ \emptyset, \emptyset \end{array} \right\} & \text{split}(\text{assert } P) &= \left\{ \begin{array}{l} \emptyset, \emptyset, \emptyset, \emptyset \\ \{\{P\}\}[], \emptyset, \emptyset, \emptyset \\ \{[]\}\{P\}, \emptyset \end{array} \right\}
\end{aligned}$$

Fig. 15. Split function for basic statements

Fig. 16. Different full paths in $C_1; C_2$

If

$$\text{split}(C_1) = \left\{ \begin{array}{l} p_-^{\text{nor}}, p_-^{\text{brk}}, p_-^{\text{con}}, p_-^{\text{ret}}, \\ p_+^{\text{nor}}, p_+^{\text{brk}}, p_+^{\text{con}}, p_+^{\text{ret}}, \\ p_{\downarrow}, p_{\uparrow} \end{array} \right\}$$

$$\text{split}(C_2) = \left\{ \begin{array}{l} q_-^{\text{nor}}, q_-^{\text{brk}}, q_-^{\text{con}}, q_-^{\text{ret}}, \\ q_+^{\text{nor}}, q_+^{\text{brk}}, q_+^{\text{con}}, q_+^{\text{ret}}, \\ q_{\downarrow}, q_{\uparrow} \end{array} \right\},$$

then

$$\text{split}(C_1; C_2) = \left\{ \begin{array}{l} p_-^{\text{nor}} \cdot q_-^{\text{nor}}, \\ p_-^{\text{brk}} \cup p_-^{\text{nor}} \cdot q_-^{\text{brk}}, \\ p_-^{\text{con}} \cup p_-^{\text{nor}} \cdot q_-^{\text{con}}, \\ p_-^{\text{ret}} \cup p_-^{\text{nor}} \cdot q_-^{\text{ret}}, \\ p_+^{\text{nor}} \cup p_+^{\text{nor}} \cdot q_+^{\text{nor}}, \\ p_+^{\text{brk}} \cup q_+^{\text{brk}} \cup p_+^{\text{nor}} \cdot q_+^{\text{brk}}, \\ p_+^{\text{con}} \cup q_+^{\text{con}} \cup p_+^{\text{nor}} \cdot q_+^{\text{con}}, \\ p_+^{\text{ret}} \cup q_+^{\text{ret}} \cup p_+^{\text{nor}} \cdot q_+^{\text{ret}}, \\ p_{\downarrow} \cup p_-^{\text{nor}} \cdot q_{\downarrow}, \\ p_{\uparrow} \cup q_{\uparrow} \cup p_+^{\text{nor}} \cdot q_{\uparrow} \end{array} \right\}.$$

Fig. 17. The definition of $\text{split}(C_1; C_2)$

Recursion cases. Using sequential composition as an example, all “full paths” between the assertions in $C_1; C_2$ can be divided into three classes (Fig. 16): (1) paths that completely fall in the CFG of C_1 ; (2) paths that completely fall in the CFG of C_2 ; and (3) paths that combine two parts, one of which is a “tail path” of C_1 and the other of which is a “head path” of C_2 . Such “tail paths” and “head paths” of an assertion-annotated C program can also be recursively computed. We show our complete definition in Fig. 17.

In this definition, we use \cdot to represent the concatenation of two paths, and overload this notation to connect two sets of paths:

$$\begin{aligned}
p_+ \cdot q_{\downarrow} &= \{P\}(\vec{c}_b \uparrow \vec{c}_b')\{Q\} & \text{where } p_+ &= \{P\}\vec{c}_b & \text{and } q_{\downarrow} &= \vec{c}_b'\{Q\} \\
p_- \cdot q_{\downarrow} &= (\vec{c}_b \uparrow \vec{c}_b')\{Q\} & \text{where } p_- &= \vec{c}_b & \text{and } q_{\downarrow} &= \vec{c}_b'\{Q\} \\
p_+ \cdot q_- &= \{P\}(\vec{c}_b \uparrow \vec{c}_b') & \text{where } p_+ &= \{P\}\vec{c}_b & \text{and } q_- &= \vec{c}_b' \\
p_- \cdot q_- &= \vec{c}_b \uparrow \vec{c}_b' & \text{where } p_- &= \vec{c}_b & \text{and } q_- &= \vec{c}_b' \\
\mathbf{p} \cdot \mathbf{q} &= \{\mathbf{p} \cdot \mathbf{q} \mid p \in \mathbf{p}, q \in \mathbf{q}\}
\end{aligned}$$

Computing $\text{split}(\text{if } (b) C_1 \text{ else } C_2)$ simply adds assume statements to the head of all head paths and assertion-free paths in the two if-branches, and returns the union of the two split results.

$$\begin{aligned}
& \text{If } \text{split}(C_1) = \left\{ \begin{array}{l} \mathbf{p}_-^{\text{nor}}, \mathbf{p}_-^{\text{brk}}, \mathbf{p}_-^{\text{con}}, \mathbf{p}_-^{\text{ret}}, \\ \mathbf{p}_+^{\text{nor}}, \mathbf{p}_+^{\text{brk}}, \mathbf{p}_+^{\text{con}}, \mathbf{p}_+^{\text{ret}}, \\ \mathbf{p}_\perp, \mathbf{p}_\perp \end{array} \right\}, \\
& \text{then } \text{split}(\text{ExGiven } x : A, \{P(x)\} C_1) = \\
& \left(\begin{array}{l} \emptyset, \emptyset, \emptyset, \emptyset \\ \{ \{ \exists x : A. Q \} \vec{c}_b \mid \{ Q \} \vec{c}_b \in \mathbf{p}_+^{\text{nor}} \} \cup \{ \{ \exists x : A. P(x) \} \} \cdot \mathbf{p}_-^{\text{nor}} \\ \{ \{ \exists x : A. Q \} \vec{c}_b \mid \{ Q \} \vec{c}_b \in \mathbf{p}_+^{\text{brk}} \} \cup \{ \{ \exists x : A. P(x) \} \} \cdot \mathbf{p}_-^{\text{brk}} \\ \{ \{ \exists x : A. Q \} \vec{c}_b \mid \{ Q \} \vec{c}_b \in \mathbf{p}_+^{\text{con}} \} \cup \{ \{ \exists x : A. P(x) \} \} \cdot \mathbf{p}_-^{\text{con}} \\ \{ \{ \exists x : A. Q \} \vec{c}_b \mid \{ Q \} \vec{c}_b \in \mathbf{p}_+^{\text{ret}} \} \cup \{ \{ \exists x : A. P(x) \} \} \cdot \mathbf{p}_-^{\text{ret}} \\ \{ \{ \neg \exists x : A. P(x) \} \}, \\ \{ \{ \forall x : A. \{ P \} \vec{c}_b \{ Q \} \} \vec{c}_b \{ Q \} \in \mathbf{p}_\perp \} \cup \{ \forall x : A. \{ Q \} \vec{c}_b \{ R \} \mid \{ Q \} \vec{c}_b \{ R \} \in \mathbf{p}_\perp \} \end{array} \right)
\end{aligned}$$

Fig. 18. The definition of $\text{split}(\text{ExGiven } x : A, \{P(x)\} C_1)$

Computing $\text{split}(\text{loop}(C_2) C_1)$ is similar. Detailed definitions can be found in the extended version of the paper and our Coq development.

Handling logical variables (Fig. 18). The focus of computing $\text{split}(\text{ExGiven } x : A, \{P(x)\} C_1)$ is to handle the logical variable x . (1) The ExGiven structure has an existentially quantified precondition $P(x)$ in the head. Therefore, there are no assertion-free paths in the split result, and the result of the head paths is a singleton of $\{ \neg \exists x : A. P(x) \}$. (2) The tail paths in $\text{ExGiven } x : A, \{P(x)\} C_1$ can either be paths from the precondition $\exists x : A. P(x)$ to exits of C_1 or tail paths of C_1 itself. When these tail paths connect to head paths later, the postconditions of those head paths will not be in the scope of x . Thus, we existentially quantify over the variable x in all those tail paths' preconditions now, i.e. if $\{ Q(x) \} \vec{c}_b$ is a tail path of C_1 , then $\{ \exists x : A. Q(x) \} \vec{c}_b$ is a tail path of $\text{ExGiven } x : A, \{P(x)\} C_1$. (3) In full paths, we need to unify the existential variable x in P with those in the head paths that are split from C_1 , so that x can be shared among the pre-/post-conditions of each combined full path. Full paths in C_1 (\mathbf{p}_\perp) are also collected after adding a universal binder x to the result. In our Coq development, we implement this definition using Coq dependent types. We put technique details in the extended version of the paper.

4.3 Proof of Soundness

We prove theorem 4 by induction over ClightA syntax trees. For the convenience of presentation, we use $\text{AllHypo}(P, \vec{Q}, \text{split}(C))$ to represent all ten hypotheses of this soundness theorem: In order words, the soundness theorem says: $\text{AllHypo}(P, \vec{Q}, \text{split}(C))$ implies $\{P\} C \Downarrow \{ \vec{Q} \}$.

$$\begin{aligned}
& \text{AllHypo} \left(P, Q, [Q_{\text{brk}}, Q_{\text{con}}, Q_{\text{ret}}], \left\{ \begin{array}{l} \mathbf{p}_-^{\text{nor}}, \mathbf{p}_-^{\text{brk}}, \mathbf{p}_-^{\text{con}}, \mathbf{p}_-^{\text{ret}}, \\ \mathbf{p}_+^{\text{nor}}, \mathbf{p}_+^{\text{brk}}, \mathbf{p}_+^{\text{con}}, \mathbf{p}_+^{\text{ret}}, \\ \mathbf{p}_\perp, \mathbf{p}_\perp \end{array} \right\} \right) \\
& \triangleq \begin{array}{l} \text{AssnFreeHypo}(P, Q, \mathbf{p}_-^{\text{nor}}) \quad \wedge \quad \text{AssnFreeHypo}(P, Q_{\text{brk}}, \mathbf{p}_-^{\text{brk}}) \quad \wedge \\ \text{AssnFreeHypo}(P, Q_{\text{con}}, \mathbf{p}_-^{\text{con}}) \quad \wedge \quad \text{AssnFreeHypo}(P, Q_{\text{ret}}, \mathbf{p}_-^{\text{ret}}) \quad \wedge \\ \text{TailHypo}(Q, \mathbf{p}_+^{\text{nor}}) \quad \wedge \quad \text{TailHypo}(Q_{\text{brk}}, \mathbf{p}_+^{\text{brk}}) \quad \wedge \\ \text{TailHypo}(Q_{\text{con}}, \mathbf{p}_+^{\text{con}}) \quad \wedge \quad \text{TailHypo}(Q_{\text{ret}}, \mathbf{p}_+^{\text{ret}}) \quad \wedge \\ \text{HeadHypo}(P, \mathbf{p}_\perp) \quad \wedge \quad \text{FullHypo}(\mathbf{p}_\perp). \end{array}
\end{aligned}$$

In this proof, only induction steps about sequential compositions, if-statements and loops are interesting. We describe the main idea of proving $\text{split}(C_1; C_2)$ sound in this section. The proofs for if-statements and loops are similar and can be found in the extended version of the paper.




Fig. 19. Soundness proof example of sequential composition

Fig. 19 shows an example of sequential composition with precondition P , postcondition Q , and annotations P_1, P_2, Q_1, Q_2 . The split function will generate these straightline Hoare triples:

$$\begin{array}{l} \{P_1\} \quad c_{1a}; \text{assume } b; c_{2a} \quad \{Q_1\} \\ \{P_1\} \quad c_{1a}; \text{assume } !b; c_{2b} \quad \{Q_2\} \\ \{P_2\} \quad c_{1b}; \text{assume } b; c_{2a} \quad \{Q_1\} \\ \{P_2\} \quad c_{1b}; \text{assume } !b; c_{2b} \quad \{Q_2\} \\ \dots \end{array}$$

In order to prove $\{P\} C_1 \Downarrow; C_2 \Downarrow \{Q\}$, we need to find an intermediate assertion R such that $\{P\} C_1 \Downarrow \{R\}$ and $\{R\} C_2 \Downarrow \{Q\}$ (according to SEMAX-SEQ). These two judgments can be established by induction hypothesis and the following four Hoare triples:

$$\{P_1\} c_{1a} \{R\} \quad \{R\} \text{assume } b; c_{2a} \{Q_1\} \quad \{P_2\} c_{1b} \{R\} \quad \{R\} \text{assume } !b; c_{2b} \{Q_2\}$$

These requirements can be simply satisfied if we let R be

$$\text{wp}(\text{assume } b; c_{2a}, Q_1) \wedge \text{wp}(\text{assume } !b; c_{2b}, Q_2).$$

In the general case, we instantiate R to the conjunction of the weakest preconditions of all head paths and assertion-free paths in $\text{split}(C_2)$. Using VST's higher order logic, this middle condition can be written as:

$$R \triangleq \left\{ \begin{array}{l} \exists R. R \wedge \text{HeadHypo}(R, q_+) \\ \wedge \text{AssnFreeHypo}(R, Q, q_-^{\text{nor}}) \\ \wedge \text{AssnFreeHypo}(R, Q_{\text{brk}}, q_-^{\text{brk}}) \\ \wedge \text{AssnFreeHypo}(R, Q_{\text{con}}, q_-^{\text{con}}) \\ \wedge \text{AssnFreeHypo}(R, Q_{\text{ret}}, q_-^{\text{ret}}) \end{array} \right\}$$

According to the induction hypothesis, it suffices to prove the following two propositions.

$$\text{AllHypo}(P, R, [Q_{\text{brk}}, Q_{\text{con}}, Q_{\text{ret}}], \text{split}(C_1)) \quad (1)$$

$$\text{AllHypo}(R, Q, [Q_{\text{brk}}, Q_{\text{con}}, Q_{\text{ret}}], \text{split}(C_2)) \quad (2)$$

The proof of (2) is simple, which can be justified by the following lemma in VST.

LEMMA 7. For any program c and postcondition \vec{Q} , $\left\{ \exists P : \text{assert}. P \wedge \left(\{P\} c \left\{ \vec{Q} \right\} \right) \right\} c \left\{ \vec{Q} \right\}$ holds.

For proposition 1, the inversion lemma for sequencing (Lemma 1) has already shown that the weakest precondition of the second statement can serve as the intermediate assertion for the sequential composition. Based on Lemma 1, we can prove a corresponding inversion lemma on the \cdot operator for each type of path in the split results.

PROPOSITION 8 (INVERSION LEMMAS FOR SPLIT RESULTS).

- (1) If $\text{AssnFreeHypo}(P, Q, p_- \cdot q_-)$, then $\text{AssnFreeHypo}(P, \exists R. R \wedge \text{AssnFreeHypo}(R, Q, q_-), p_-)$
- (2) If $\text{HeadHypo}(P, p_- \cdot q_+)$, then $\text{AssnFreeHypo}(P, \exists R. R \wedge \text{HeadHypo}(R, q_+), p_-)$

- (3) If $\text{TailHypo}(Q, p_+ \cdot q_-)$, then $\text{TailHypo}(Q, \exists R. R \wedge \text{AssnFreeHypo}(R, Q, q_-), p_+)$
(4) If $\text{FullHypo}(p_+ \cdot q_+)$, then $\text{TailHypo}(\exists R. R \wedge \text{HeadHypo}(R, q_+), p_+)$

According to the conjunction rule, AssnFreeHypo and TailHypo from Proposition 8 still hold if we combine all of those weakest preconditions of C_2 's partial paths (i.e. those preconditions are such that $\text{AssnFreeHypo}(R, Q, q_-)$ and $\text{HeadHypo}(R, q_+)$). Formally, Proposition 8 can be extended into the following form by the conjunction rule.

PROPOSITION 9 (GROUPED INVERSION LEMMAS).

- (1) If $\text{AssnFreeHypo}(P, Q, p_- \cdot q_-)$, then $\text{AssnFreeHypo}(P, \exists R. R \wedge \text{AssnFreeHypo}(R, Q, q_-), p_-)$
(2) If $\text{HeadHypo}(P, p_- \cdot q_+)$, then $\text{AssnFreeHypo}(P, \exists R. R \wedge \text{HeadHypo}(R, q_+), p_-)$
(3) If $\text{TailHypo}(Q, p_+ \cdot q_-)$, then $\text{TailHypo}(Q, \exists R. R \wedge \text{AssnFreeHypo}(R, Q, q_-), p_+)$
(4) If $\text{FullHypo}(p_+ \cdot q_+)$, then $\text{TailHypo}(\exists R. R \wedge \text{HeadHypo}(R, q_+), p_+)$

THEOREM 10 (CONJUNCTION RULE). *If Hoare triples $\{P\} c \{Q_1, [\vec{Q}'_1]\}$ and $\{P\} c \{Q_2, [\vec{Q}'_2]\}$ are derivable, then $\{P\} c \{Q_1 \wedge Q_2, [\vec{Q}'_1 \wedge \vec{Q}'_2]\}$ is derivable.*

Next, we also use the conjunction rule to combine the weakest preconditions of the different kinds of paths and prove proposition 1, which completes the soundness proof of $\text{split}(C_1; C_2)$.

However, the conjunction rule is not ubiquitous among the Hoare logic variants proposed in the literature: for example, the current VST program logic cannot derive the conjunction rule. We will leave the discussion of the conjunction rule to §5. For now, we assume that the conjunction rule holds, so that Proposition 9 and the proof of $\text{split}(C_1; C_2)$ soundness can hold.

5 CONJUNCTION RULE AND PRECISENESS

The conjunction rule is natural in traditional Hoare logics and separation logics for sequential programs, but some extensions to the logics will make the conjunction rule inadmissible. In this section, we more extensively discuss why the conjunction rule is required by our soundness proof (§5.1, §5.2). To make the conjunction rule admissible in VST-A (§5.3, §5.4, §5.5), we identify a new notion of preciseness to restrict the function specifications being called during verification. We also discuss the trade-offs of using conjunction rules and precise function specifications, and we suggest some future directions for improvement (§5.6).

5.1 A Small Example

Suppose we would like to prove the Hoare triple

$$\{P\} c_1; \mathbf{if} (b) c_2 \mathbf{else} c_3 \{R\} \quad (3)$$

given that the following split results hold (here, we assume that c_1 , c_2 and c_3 are primary statements):

$$\{P\} c_1; \mathbf{assume} b; c_2 \{R\} \quad (4)$$

$$\{P\} c_1; \mathbf{assume} !b; c_3 \{R\} \quad (5)$$

By inversion on sequential composition (lemma 1), proposition (3), (4) and (5) are equivalent to:

$$\{P\} c_1 \{\text{wp}(\mathbf{if} (b) c_2 \mathbf{else} c_3, R)\}$$

$$\{P\} c_1 \{\text{wp}(\mathbf{assume} b; c_2, R)\}$$

$$\{P\} c_1 \{\text{wp}(\mathbf{assume} !b; c_3, R)\}$$

Furthermore, by inversion (Lemmas 1 and 2) and properties of \mathbf{assume} (Lemma 6), the postcondition of the first triple $\text{wp}(\mathbf{if} (b) c_2 \mathbf{else} c_3, R)$ is equivalent to

$$\text{wp}(\mathbf{assume} b; c_2, R) \wedge \text{wp}(\mathbf{assume} !b; c_3, R)$$

Thus, the split function’s soundness on the example above can be reduced to an instance of the conjunction rule.

5.2 Unsoundness When the Conjunction Rule Is Inadmissible

So far, we have seen a tight connection between the conjunction rule and the split function’s soundness — our soundness proof uses the conjunction rule (§4.3) and a very simple instance of this soundness theorem can be reduced to an instance of the conjunction rule (§5.1). But what will happen to split’s soundness if the conjunction rule is not admissible? Consider Hoare logic with ghost updates [Krebbers et al. 2017a] as an example. Ghost states are “logical states” that help with the program’s proof, and they particularly useful for verifying concurrent programs.⁹ When users prove programs with ghost states, they can apply a ghost update when they use the consequence rule, see SEMAX-CONSEQ-GHOST below. Here, $\models P$ says: there is at least one possible ghost update which makes the state satisfy P .

$$\text{SEMAX-CONSEQ-GHOST} \frac{P_1 \models P_2 \quad R_2 \models R_1 \quad \vec{R}'_2 \models \vec{R}'_1 \quad \{P_2\} c \{R_2, [\vec{R}'_2]\}}{\{P_1\} c \{R_1, [\vec{R}'_1]\}}$$

The conjunction rule is not admissible in this logic — if the proofs of $\{P\} c \{Q\}$ and $\{P\} c \{Q'\}$ use different and conflicting ghost updates, $\{P\} c \{Q \wedge Q'\}$ cannot be valid since two conflicting ghost updates cannot happen simultaneously.

In this logic, our split function is unsound. This loss of soundness is not determined by the way we prove soundness in §4.3 but by the framework we propose to first split the program into individual paths, which are then verified separately. Consider the following annotated program:

```

1   /*@ Assert g ↦ A */
2   f();
3   x = nondetermined_0_or_1();
4   if (x) {
5       /*@ Assert g ↦ B1 ∧ x = 1 */
6       f1();
7   } else {
8       /*@ Assert g ↦ B0 ∧ x = 0 */
9       f0();
10  }
11  /*@ Assert g ↦ C1 ∧ x = 1 ∨ g ↦ C0 ∧ x = 0 */

```

in which g is a ghost location for storing the status of the following STS (state transition system) [Turon et al. 2013]: $A \rightarrow A_0$, $A \rightarrow A_1$, $A_0 \rightarrow B_0$, $A_1 \rightarrow B_1$, $B_0 \rightarrow C_0$, $B_1 \rightarrow C_1$. We assume that $f()$,

⁹Ghost states are not the same as the “ghost variables” of traditional Hoare logics. Ghost variables are logical variables that were introduced to relate old values of variables to current values, and to relate current values to abstract mathematical values. VST and Iris support ghost variables using ordinary Coq variables; those ghost variables are fully compatible with our VST-A program decomposition, to support data abstraction and modular verification [Berlinger 2021]. Examples of such variables in Fig. 1 are l, l_1, u, x, l_2 .

$f0()$, and $f1()$'s specifications are:

$$\begin{aligned} &\text{for any } b \in \mathbf{bool}, \\ &\{g \mapsto (\text{if } b \text{ then } A_1 \text{ else } A_0)\} f() \{g \mapsto (\text{if } b \text{ then } B_1 \text{ else } B_0)\} \\ &\{g \mapsto B_0\} f0() \{g \mapsto C_0\} \\ &\{g \mapsto B_1\} f1() \{g \mapsto C_1\}. \end{aligned}$$

In this example, all straightline Hoare triples in `split`'s result are provable, especially the following two triples about `f()`:

$$\begin{aligned} &\{g \mapsto A\} f(); x = \text{nondetermined_0_or_1}(); \mathbf{assume} \ x; \{g \mapsto B_1 \wedge x = 1\} \\ &\{g \mapsto A\} f(); x = \text{nondetermined_0_or_1}(); \mathbf{assume} \ !x; \{g \mapsto B_0 \wedge x = 0\} \end{aligned}$$

For the first triple, we can choose to take the $A \rightarrow A_1$ step in this STS before calling `f()`. For the second triple, we can choose to take the $A \rightarrow A_0$ step in this STS before calling `f()`. However, the whole Hoare triple is not provable since we cannot determine the value of `x` before `x = nondetermined_0_or_1()` is executed. That breaks `split`'s soundness.

5.3 VST-A's Design Choice and Proof Strategy

In the current design of VST-A, we focus on sequential program verification and disallow all ghost updates. VST-A uses a more restricted variant of the VST program logic. This variant is still proved sound w.r.t. CompCert Clight semantics and the most significant change is that the ghost update operator is removed from the consequence rule.

Despite the removal of ghost updates, users are still able to write unrestricted higher-order predicates and prove many complex sequential programs in VST-A. We derive the conjunction rule (theorem 10) from this stronger logic by induction over Clight abstract syntax tree. Our inductive proof steps are all Hoare-logic-based, and we believe that such a proof strategy is (1) easier to formalize in Coq, and (2) in fact more general than a semantic-based proof, since it is independent of how the soundness of the Hoare logic was proved w.r.t. its semantic model.

Consider the induction step for $c = c_1; c_2$. By applying Lemma 1 to the premises we obtain the following:

$$\begin{aligned} &\{P\} c_1 \left\{ \exists R_1. R_1 \wedge \{R_1\} c_2 \left\{ Q_1, [\vec{Q}'_1] \right\}, [\vec{Q}'_1] \right\} \\ &\{P\} c_1 \left\{ \exists R_2. R_2 \wedge \{R_2\} c_2 \left\{ Q_2, [\vec{Q}'_2] \right\}, [\vec{Q}'_2] \right\} \end{aligned}$$

We can apply the induction hypothesis of c_1 and make use of SEMAX-CONSEQ to obtain the following:

$$\{P\} c_1 \left\{ \begin{array}{l} \exists R. R \wedge \{R\} c_2 \left\{ Q_1, [\vec{Q}'_1] \right\} \\ \wedge \{R\} c_2 \left\{ Q_2, [\vec{Q}'_2] \right\} \end{array} \right\}, [\vec{Q}'_1 \wedge \vec{Q}'_2]$$

where R can be instantiated as $R_1 \wedge R_2$. According to SEMAX-SEQ, we are left to prove:

$$\left\{ \begin{array}{l} \exists R. R \wedge \{R\} c_2 \left\{ Q_1, [\vec{Q}'_1] \right\} \\ \wedge \{R\} c_2 \left\{ Q_2, [\vec{Q}'_2] \right\} \end{array} \right\} c_2 \left\{ \begin{array}{l} Q_1 \wedge Q_2, \\ [\vec{Q}'_1 \wedge \vec{Q}'_2] \end{array} \right\}$$

Using EXTRACT-EXISTS and EXTRACT-PROP, it suffices to prove:

$$\{R\} c_2 \left\{ Q_1, [\vec{Q}'_1] \right\}, \{R\} c_2 \left\{ Q_2, [\vec{Q}'_2] \right\} \Rightarrow \{R\} c_2 \left\{ Q_1 \wedge Q_2, [\vec{Q}'_1 \wedge \vec{Q}'_2] \right\}$$

which immediately follows the induction hypothesis of c_2 .

For other induction steps (e.g. for if-statements and for loops), the proof idea is somewhat similar in that their corresponding inversion lemmas are first applied, then the induction hypotheses can

be used to complete the proof. For control flow statements (break, continue and return), the proof is trivial. Thus, we are only left to derive the conjunction rules for primary statements.

5.4 The Conjunction Rule for Memory Stores

Consider a simple store statement and corresponding proof rule in VST.

$$\forall v. \{ p \mapsto v * (p \mapsto v' \text{ } \ast R) \} \ast p = v' \{R\} \quad (6)$$

The precondition states that the heap can be split into two parts. One part is a singleton heap, in which the nonaddressable variable p is a pointer to the old value v , and it is described by the mapsto predicate $x \mapsto v$. The other part should satisfy the postcondition R when joined with a singleton heap that stores the new value v' .

Note that in this specification, the old value v is universally quantified. If we want to conjoin the postconditions of two Hoare triples about $\ast p = v'$, we need to unify the two different instantiations of v into one. To be specific, the following property needs to hold:

PROPOSITION 11 (PRECISENESS OF STORE).

$$\begin{aligned} & (\exists v_1. p \mapsto v_1 * (p \mapsto v' \text{ } \ast R_1)) \wedge \\ & (\exists v_2. p \mapsto v_2 * (p \mapsto v' \text{ } \ast R_2)) \\ \models & \exists v. p \mapsto v * (p \mapsto v' \text{ } \ast R_1 \wedge R_2) \end{aligned}$$

In VST-A, we prove this through semantic-model-level reasoning based on its underlying memory model. Then, the conjunction rule for such simplified store statements can be logically derived. We leave the detailed proof for the extended version of the paper.

In VST, assignment statements $e_1 := e_2$ include (1) assigning a value to a nonaddressable variable, (2) loading a value from memory to a nonaddressable variable and (3) storing a value in memory. VST's higher-order assertion language also supports separation logic predicates with fractional permissions. In VST-A, we have proved that the conjunction rule holds for all primary set/load/store operations.

5.5 The Conjunction Rule for Function Calls

The store rule mentioned above can be treated as a specification for a function call that contains only a single store instruction. We have shown that Proposition 11 is an important property for deriving the conjunction rule. In this section, we generalize this property, which we refer to as *precise function specifications*, to derive the conjunction rule for function calls.

As previously mentioned, VST-A function specification has a form of

$$\text{With } [\vec{a} : \vec{A}]. \text{ Require } \{\lambda \vec{y}. P\} \text{ Ensure } \{\lambda r. Q\},$$

where P is a precondition parameterized by a list of formal parameters \vec{y} , Q is a postcondition parameterized by the return value r , and types \vec{A} represents the type of logical values to be shared between P and Q . Therefore, both P and Q will also be abstracted over a set of logical variables \vec{a} typed \vec{A} . We define the precise function specifications as follows:

Definition 12 (Precise function specification). With $[\vec{a} : \vec{A}]. \text{ Require } \{\lambda \vec{y}. P\} \text{ Ensure } \{\lambda r. Q\}$ is precise, if for any formal parameters \vec{b} , return value r and assertions R_1, R_2 , it holds that

$$\begin{aligned} & \left(\exists \vec{x}_1 : \vec{A}. P \vec{b} \vec{x}_1 * (Q r \vec{x}_1 \text{ } \ast R_1) \right) \wedge \left(\exists \vec{x}_2 : \vec{A}. P \vec{b} \vec{x}_2 * (Q r \vec{x}_2 \text{ } \ast R_2) \right) \\ \models & \exists \vec{x} : \vec{A}. P \vec{b} \vec{x} * (Q r \vec{x} \text{ } \ast R_1 \wedge R_2) \end{aligned}$$

In VST-A’s program logic (a variant of VST’s program logic), the function call rule, `SEMAX-CALL`, requires the callee function specification to be precise.

$$\text{SEMAX-CALL} \frac{\begin{array}{l} \phi = \text{With } [\vec{x} : \vec{A}]. \text{Require } \{\lambda \vec{y}. P\} \text{Ensure } \{\lambda r. Q\} \\ f \text{ satisfies } \phi \quad \phi \text{ is a precise specification} \end{array}}{\left\{ \vec{e} \Downarrow \vec{b} \wedge \exists \vec{x} : \vec{A}. P \vec{b} \vec{x} * (Q \text{ a } \vec{x} * R) \right\} a := f(\vec{e}) \{R, [\vec{\perp}]\}}$$

The notion of “precise function specification” that we propose is defined with respect to an operation (e.g. a store statement or a function call), while traditionally, “preciseness” is defined for a predicate. A typical example of a precise predicate is the $p \mapsto v$ predicate used in `SEMAX-STORE`.

Definition 13 (Precise predicate). A predicate P is precise if for any Q_1, Q_2 ,

$$(P * Q_1) \wedge (P * Q_2) = P * (Q_1 \wedge Q_2)$$

In fact, our definition of precise function specifications includes a common use case of precise predicates. As proved by Gotsman et al. [2011] and Vafeiadis [2011], if the conjunction rule is expected in concurrent separation logic with locks, the resource invariant for locks should be a precise predicate. Consider the following specification for a release operation for lock l :

$$\forall l \pi. \{ \pi \text{ is readable} \wedge \text{locked}_\pi(l, R) * R \} \text{release}(l) \{ \text{unlocked}_\pi(l, R) \}$$

where R is the predicate that describes the locked memory. We use locked_π and unlocked_π to denote whether this memory is owned by the current thread or not; π is a fractional permission. It is clear that if the resource invariant R is a precise predicate, then one can show that this specification is a precise specification. The precise function specification also makes a difference in the sequential setting we are considering. One example of a non-precise function specification is:

```

1 void foo (int * p)
2 /*@ With q Require p ↦ 0 * F(q) Ensure p ↦ 1 * F(q) */
3 { * p = 1; return; }

```

Here F is a predicate dependent on a logical variable q describing the frames that the function `foo` does not modify. It is possible to find such F that holds for two different instantiations of q , and breaks the precise function specification requirement.

In our Coq development, we find that, for most function specifications, we can derive that for any formal parameters \vec{b} , logical variables \vec{x}_1, \vec{x}_2 :

$$\left(P \vec{b} \vec{x}_1 * \text{True} \right) \wedge \left(P \vec{b} \vec{x}_2 * \text{True} \right) \vDash \vec{x}_1 = \vec{x}_2 \quad (7)$$

Then, to prove a specification precise, it is sufficient to show that P is a precise predicate. A typical example is the store predicate “ \mapsto ”. Clearly:

$$\forall p v_1 v_2. ((p \mapsto v_1 * \text{True}) \wedge (p \mapsto v_2 * \text{True})) \vDash v_1 = v_2 \quad (8)$$

For linked lists’ representation predicate ll , we also have:

$$\forall p l_1 l_2. ((\text{ll}(p, l_1) * \text{True}) \wedge (\text{ll}(p, l_2) * \text{True})) \vDash l_1 = l_2 \quad (9)$$

Moreover, property (7) is composable. For example, in order to prove:

$$((p \mapsto x_1 * \text{ll}(x_1, l_1)) * \text{True}) \wedge ((p \mapsto x_2 * \text{ll}(x_2, l_2)) * \text{True}) \vDash x_1 = x_2 \wedge l_1 = l_2$$

we can first use (8) to derive $x_1 = x_2$, then use after (9) to derive $l_1 = l_2$.

It is well-known that precise predicates are also composable, thus we developed several extensible automatic tactics in VST-A to help users prove the precise specification by combining property (7) and predicates' preciseness.¹⁰

We believe that most C functions have specifications that are naturally precise. Non-precise specifications are usually caused by separation logic conjuncts that describe inaccessible memory slices to the function, such as the $F(q)$ proposition in the example above. In practical verification tasks, one can usually remove these conjuncts using the "frame rule" and obtain a precise specification. The pruned part of the specification can be expressed elsewhere in the program, where the memory slice is really used.

5.6 Discussion and Future Work

In our current design of VST-A, we require the conjunction rule to be derivable in the logic to ensure the splitting algorithm's soundness. We do not consider our current design choice as a fatal problem for future extensions. Here are some potential research directions that may support ghost updates or remove the restriction of only using precise specifications in the future.

Verify the sequential fragments of a concurrent program. As was discussed in §5.2, the existence of the conjunction rule forbids the use of ghost updates. However, there are ways to verify the sequential fragment of a concurrent program with VST-A. We have proved the following property:

THEOREM 14. *If $\{P\} c_1; c_2 \models \{R, [\vec{R}']\}$, where \models indicates that the triple is derivable from the original VST logic that supports ghost state updates, then there exists Q, \vec{Q}' such that $\{P\} c_1 \{Q, [\vec{Q}']\}$ and $\{Q\} c_2 \models \{R, [\vec{R}']\}$.*

The triple about c_1 is derivable from the stronger program logic in VST-A. We can apply the current VST-A framework to the sequential fragment c_1 , and leave the rest of the program to be verified by the full-power VST interactive verifier.

New annotations for ghost actions. Another possible direction is to improve the capability of annotations in terms of expressing proofs. For example, in future versions of VST-A, we could allow users to write either explicit ghost commands in their programs, or write two consecutive assertions where the latter can be derived from the former with a ghost update. As a result, we do not need to worry about the possibility of having two conflicting ghost updates simultaneously.

Additional annotations for function calls. As was discussed in §5.5, in order to derive the conjunction rule for function calls, the callee should meet the requirement of precise function specification. However, this requirement could be removed if the logical variables of the specification are explicitly instantiated. This avoids the possibility that proofs of multiple straightline Hoare triples containing the same function call could actually instantiate those logical variables differently, causing issues in proving the conjunction rule. Intuitively, this is like putting two assertions before and after the function call statement as pseudo "join points" in the control flow, ensuring that every function call would only appear in one unique path in the split result. Then, split's soundness would be trivial for function calls.

Prophecy variables. In our counterexample in §5.2, the verification target is unprovable but will become provable if we are allowed to insert a prophecy variable [Abadi and Lamport 1988; Jung

¹⁰In several cases only weaker equivalence relations can be derived from the conjunction. Our tactics also support proof automation for that, if this relation implies the equality of the specification

et al. 2020] into the program. In general, we may be able to prove split’s soundness with the help of prophecy variables even if ghost updates are permitted in the logic.

6 EVALUATION

In this section we evaluate VST-A in various aspects, including verification effort, verification time, and statistics of our development. We also draw a comparison between VST-A and VST.

Verification effort using VST-A. We test VST-A using several sets of programs and present the statistics about verification effort in Table 1. VST-A proofs are divided into two parts: the annotated programs (which describe the main idea of the proof in a way easier to read than VST proofs) and Coq verification of straightline Hoare triples (similar to corresponding parts of VST proofs). Therefore, we count the lines of annotations and proofs separately, in the Specification and Assertion columns and the Proof column, respectively.

Our benchmarks include some measured by Sammler et al. [2021], including linked lists and binary search trees. We also include a slightly larger example—a small-step interpreter of a toy imperative language, which includes a number of if-branches. It can be seen from our evaluation that without providing additional assertions in the program code, VST-A is still able to split and verify the program.

Table 1. Number of C functions, lines of C code (without annotations), VST-A specification lines annotated in the C comments, assertion annotations in the C comments, Coq proofs in VST-A, versus Coq specifications and proofs in VST.

Program	Functions	Code	VST-A			VST
			Spec	Assert	Proof	Proof
Basics	8	78	56	5	84	156
Singly linked list	18	350	85	55	969	1212
Doubly linked list	4	95	16	16	171	213
Binary search tree	4	115	44	23	202	364
Interpreter	5	337	47	0	423	653

We also conduct a comparison in the proof effort between VST-A and VST. The last column in Table 1 shows the lines of proofs for verifying the same function in VST; compare this to the “Specification+Assertion+VST-A Proof” columns. The proof lines do not count auxiliary predicate definitions and lemmas since they are the same in VST and VST-A.

Not manifest in the line-count, but still important, is that the VST user must learn to use several different tactics for different forms of control flow, and each of these has several options depending on which assertions are supplied (e.g., if-postcondition, for-loop continue assertion, for-loop break assertion, etc.). In contrast, VST-A’s control-flow splitting takes care of all of this, leaving the user to learn only the straight-line *forward* and *forward_call* tactics. In some cases, splitting makes the Coq proof easier to automate. For example, verifying the interpreter example in Table 1 using VST-A only needs to use *forward* to handle assumes repeatedly, but this proof strategy cannot be used in the corresponding VST proof.

To summarize, we believe that to verify the same program, the manual effort of VST-A is no more than that of VST, while in the meantime VST-A can provide more intuitive and readable proofs with annotations in the C source program.

Verification time for different phases in VST-A. The verification time is shown in Table 2. Since the splitting process in VST-A is a computationally proven sound function, it is not surprising that the

Table 2. Verification time for different phases (in seconds). Reduction: time used for parsing C and generating straightline paths in Coq. Common: time used for compiling common definitions and theorems. Avg. compile: the average time used for compiling the straightline Hoare triple definitions of all generated paths \times the number of paths. Avg. verify: the average time used for verifying each path \times the number of paths. Max. verify: the maximum time used for verifying each path. VST verify: the time used for verifying the same programs in VST.

Program	Reduction	Common	Avg. compile	Avg. verify	Max. verify	VST verify
Basics	0.239	12.0	1.7×14	4.8×14	8.1	29.1
SLL	0.087	11.0	1.7×57	6.9×57	12.2	159.2
DLL	0.064	16.0	1.7×13	7.1×13	13.5	48.5
BST	0.069	11.0	1.7×18	7.4×18	12.2	54.8
Interpreter	0.090	25.0	1.7×45	20.8×45	41.9	222.9

reduction time is short. For different verification tasks, users need to provide different separation logic predicates and lemmas to specify and prove the program. This constitutes a majority of the compilation time, which is also needed by a VST proof.

A significant difference between VST-A and VST is that in VST-A each straightline Hoare triple's proof can be developed as a separate lemma, residing in distinct files, facilitating parallel checking and compilation. Therefore, we present the average time for these phases and the number of paths. In VST, the correctness theorem of a function needs to be proved as a whole, which is more difficult to parallelize.

In addition to parallelized compilation, the separation of straightline Hoare triple proofs also makes it easier to maintain the correctness of the proofs when the program is modified. Only those paths that are affected by the modification need to be re-verified. We conducted several experiments and the results are shown in Table 3.

Table 3. Case study of the number of paths that need to be re-verified when the program is modified. $4 \rightarrow 3$ indicates that the program modification changed the number of paths.

Program	Function	Paths	Changes	Changed paths
Singly linked list	append	5	modify the pre-condition	1
Singly linked list	rev_append	4	change the loop invariant	3
Singly linked list	reverse	$4 \rightarrow 3$	remove an assertion in the loop	$2 \rightarrow 1$
Binary search tree	lookup	5	modify the post-condition	2
Interpreter	eval	21	change the code in one branch	1

Statistics of the development of VST-A. We present the line of codes statistics for our development in VST-A below. Our development in VST-A includes the following parts:

- Coq formalization of the restricted fragment of VST program logic, including the logical rules, auxiliary lemmas, and the conjunction rule proof: 6,236 lines
- Coq formalization of the VST-A framework, including the split algorithm, and its soundness proof: 6,287 lines
- Modification to the CompCert C parser to parse annotated C programs: 1.87% of change
- Modification to the VST-Floyd lemmas and automation tactics to support forward symbolic execution in the restricted fragment of VST program logic: 6.74% of change

- OCaml development that provides an efficient implementation of the split algorithm: 3,176 lines in addition to original CompCert

7 RELATED WORK

7.1 Traditional Annotation Verifiers

Many annotation verifiers work by reducing annotated programs to SMT assertion entailments: these include Frama-C, Dafny, VeriFast, Viper, Hip/Sleek. Some of those systems use a specialized intermediate language for verification, connectable to several SMT back-ends and to several programming-language front-ends; for example Frama-C uses Why3 [Filliâtre and Paskevich 2013] and Dafny uses Boogie [Barnett et al. 2006]. Modern SMT solvers effectively and efficiently solve many of the resulting entailments. These annotation verifiers are, in practice, “interactive:” one starts by annotating the program with function specifications and some loop invariants, and the verifier inevitably points out several places where the proof fails—with sufficiently good error messages that the user can adjust the assertions, add new assertions and invariants, and try again, and again. This works well in practice, and it is what we wanted to emulate.

The disadvantage of those systems is in the poverty of their assertion languages. Because SMT (or Why3 or Boogie) accommodates only first-order or near-first-order logics, the rich specification languages of VST or of logics built in Iris cannot be used. In VST one often proves high-level properties about the behavior of programs, in application-specific mathematics that would be difficult to fit into SMT. In Frama-C, Dafny and CN (etc.), some authors work around this by writing higher-level proofs in Coq and program-logic proofs in the annotation verifier, and stapling together the two verifications (in different logics without a common foundation) [Boldo et al. 2014]. That approach could be made more foundational by embedding a semantics of Why3 in Coq [Cohen and Johnson-Freyd 2024], but the assertion language would still be near-first-order.

Another disadvantage of those annotation verifiers is that none has a machine-checked proof of soundness (e.g., w.r.t. an operational semantics). This lack is not inherent, as VST-A demonstrates.

When proving programs in VST or in an Iris-based logic, one generally uses automated solvers to prove entailments, or at least to prove the easy parts and leave residuals for the user. These solvers may be programmed in Coq (using tactics or computational reflection) or may be external (such as SMT). VST-A does not choose a specific solver. Users can choose their own solver to prove the split result correct.

BRiCk [Malecha et al. 2022], used by BedRock Systems Inc. to verify its microkernel/hypervisor, is a program logic for C++, built in Iris, on principles inspired by VST. We expect that our VST-A method would work well in such a C++ program logic.

Tools like F* [Martínez et al. 2019], LiquidHaskell [Vazou et al. 2014, 2018], and ATS [Chen and Xi 2005] have managed to combine higher-order programming with theorem proving in a dependent type system so that rich higher-order properties can be automatically verified in the style of the program’s annotations. However, they require users to either write programs in a new domain-specific language or construct the proof as a term in the program. By contrast, VST-A works on the standard (and practical) C programming language while also enabling reasoning with higher-order properties.

We also note the work of sledgehammer [Böhme and Nipkow 2010] and auto2 [Zhan 2016] for proof automation in Isabelle. Sledgehammer relies on SMT solvers, while auto2 builds compositional proof automation using a saturation-based proof automation system in which goal-directed proof strategies can be encoded. Users of auto2 can easily extend auto2 with their own domain-specific proof strategy. Zhan [2018] built an auto2 instance that supports separation logic reasoning for verifying sequential programs. Although auto2 supports flexible saturation-based proof strategies,

this specific instance of sequential program verification is mainly goal-directed. VST-A is not goal-directed, and is open to any solver or proof style when verifying entailments, regardless of whether it is, based on an interactive proof a tactic-based solver, or a model checking based one.

7.2 Interactive Prover-Based Program Verification

There is no deep reason why annotation verifiers should lack soundness proofs (e.g., Frama-C, Dafny, Verifast) and tactic-based verifiers should have machine-checked soundness proofs (e.g., VST, Iris). We guess that the reason is: such soundness proofs are naturally higher-order, more easily accommodated in the kinds of higher-order logics implemented in proof assistants such as Coq and Isabelle, so it is natural that designers of VST and Iris *also* have their *users* operate in the same proof assistants.

Soundness proofs are important for real-world programming languages, which have many subtle features in their semantics and compilers. Users want what they prove about a program to be consistent with the compiled machine code semantics, so the foundational soundness of VST-A is a real benefit. In this section, we compare VST-A with other foundational tools.

VST and various Iris-based verifiers have invested significantly in increasing proof automation so that users can verify their programs conveniently. However, they all require their users to complete correctness proofs for the entire program in an interactive theorem prover, which is not easy for an ordinary software engineer to learn.

There are also works that build annotation verification into interactive theorem provers and have achieved foundational soundness. RefinedC [Sammler et al. 2021] is an automated and foundational annotation verifier for C programs, that defines a restricted fragment of the Iris logic, Lithium, so that proof searches can be guided by translating the assertion annotations into a Lithium program. DiaFrame [Mulder et al. 2022] is also an automated and foundational tool. It employs a similar structural approach to RefinedC, but is more targeted at proving fine-grained concurrent programs. These tools use tactic-based proof strategy design and achieve some reasonable automation. In other words, a Hoare triple will be reduced to smaller proof goals (and even directly solved) by automatically applying a series of proof tactics, which use proved-sound logic rules or verified single-step symbolic execution. In comparison, VST-A is based on one computational proved-sound reduction function. Thus, developers of VST-A do not need to *decompose* this reduction step into multiple proof tactics, which in the end allows users to describe more flexible proofs using annotated C programs. Here is an example of how reduction is decomposed into tactics. Given an annotated program of form (here, we use a pair of braces to emphasize that sequential composition is right associative in CompCert Clight and in our ClightA syntax):

$$/*@ P */ \text{if } (b) \ c_1 \ \text{else } c_2; \{ c_3; /*@ Q */ c_4 \} /*@ R */$$

VST-A will generate 3 straightline Hoare triples:

$$\{P\} \text{assume } b; c_1; c_3 \{Q\} \quad \{P\} \text{assume } !b; c_2; c_3 \{Q\} \quad \text{and} \quad \{Q\} c_4 \{R\}.$$

In order to achieve this in RefinedC's or DiaFrame's tactic-based proof automation system, the system needs to apply the SEQ-ASSOC rule (see §2.2) first, turning the proof goal into:

$$\{P\} \{ \text{if } (b) \ c_1 \ \text{else } c_2; c_3 \}; c_4 \{R\}$$

and then apply the sequence rule with middle condition Q . After that, one more proof rule¹¹ is needed for turning $\text{if } (b) \ c_1 \ \text{else } c_2; c_3$ into $\text{if } (b) \{ c_1; c_3 \} \ \text{else} \{ c_2; c_3 \}$ so that the SEMAX-IF rule can be used to complete the reduction. However, such tactic-based decomposition is not always easy to design, and it can even be impossible. Especially, it is not obvious how to design tactic-based

¹¹In most Iris-based verifiers, this last step is not needed since Iris's symbolic execution can do that implicitly.

proof automation for handling nontraditional loop invariants supported by VST-A. That is, our split function (in effect) performs some nontrivial static analysis and is verified by a nontrivial soundness proof. Besides supporting more flexible proofs, the core split function in VST-A is computation-based so that VST-A can first complete the reduction step very efficiently, and users can then prove straightline Hoare triples manually or using their own domain-specific proof automations. Also, this design of VST-A can better support incremental development.

7.3 Conjunction Rule and Preciseness

The conjunction rule is naturally sound in traditional Hoare logic [Floyd 1993]. However, for concurrent separation logic with locks, a counterexample that leads to unsoundness [O’Hearn 2004] can be found. As a workaround, De Vilhena et al. [2020] proved a restricted version called the candidate rule, which requires postconditions to be pure (independent of resources, especially ghost resources) to solve their verification problem. However, this rule cannot be applied in our setting. In VST-A, we do not propose alternative rules but prove the conjunction rule on top of a VST logic without ghost updates. As for supporting concurrency, we proposed several possible directions in §5.6.

We are not aware of any similar notions of precise function specifications in the literature as we have defined in this paper. Traditionally, preciseness restrictions are placed on assertion predicates. For example, in concurrent separation logic, the resource invariant should be a precise predicate to make the conjunction rule sound [Gotsman et al. 2011; Vafeiadis 2011]. Our aim is to define a notion of preciseness for specifications, so that the conjunction rule is derivable from the existing logical rules. Compared with traditional preciseness, we showed in §5.5 that our notion of preciseness is more expressive, as it accounts for a pair of pre-/post-conditions for an operation and allows the specification to be quantified by logical variables.

8 CONCLUSION

We have presented VST-A, an annotation verifier that is foundationally verified. VST-A targets a widely used real-world language, C, and supports higher-order assertions in the very rich specification language of VST that includes the full expressive power of Coq. VST-A splits the verification of a large program into verifications of straightline control flow paths separated by assertions. The soundness of this approach requires the conjunction rule to be derivable in the program logic. We have identified a novel notion of precise specifications in the proof of the conjunction rule. Currently, VST-A only supports sequential C program verification, but we have proposed ways to extend VST-A to support concurrency in the future. Our formal annotation language and other major designs are not C-specific, nor are they separation-logic-specific, nor VST-specific. A similar development can be used to design other Hoare-style annotation verifiers for imperative languages.

Comparing to existing foundational program verification tools built in interactive theorem provers, VST-A has the following advantages:

- Annotation-based proof is a more readable way to explain why a program is correct.
- Our annotation-based proof language ClightA is expressive enough to describe nonstructural proofs, which cannot be supported systematically using goal-directed tactic-based proof automation;
- VST-A is easier to use – users only need to write assertions in annotations, and use forward symbolic execution to prove straightline Hoare triples. In comparison, users of other tools like VST and Iris need to use different tactics to handle different program structures like if-conditions, recursions and different loops.

- VST-A reduces proof recompilation time. When a verified program is updated slightly, its correctness proof also needs corresponding updates. In existing interactive verifiers, users must rerun all tactical proof scripts, even though only a small portion needs to be updated. Now, only the part of the program that has been changed and the corresponding proof need to be recompiled, since other parts in split's result are unchanged.

We aim to enhance the verification process of VST-A further. Future work includes the introduction of domain-specific heuristics for automatically manipulating separation logic predicates during symbolic execution and proving separation logic entailments on straightline Hoare triples. We also plan to allow user to specify partial assertions, so that users do not need to write assertions for the entire state of the function throughout the whole program.

ACKNOWLEDGMENTS

This material is based upon work supported (in part) by NSF China 61902240, and the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR001120C0160.

DATA AVAILABILITY STATEMENT

The VST-A formalization, implementation and evaluation examples are openly available in Zenodo [Zhou et al. 2023b].

REFERENCES

- Martín Abadi and Leslie Lamport. 1988. The Existence of Refinement Mappings. In *Proceedings of the Third Annual Symposium on Logic in Computer Science (LICS '88), Edinburgh, Scotland, UK, July 5-8, 1988*. IEEE Computer Society, 165–175. <https://doi.org/10.1109/LICS.1988.5115>
- Andrew W. Appel. 2011. Verified software toolchain (Invited talk). In *Lecture Notes in Computer Science*, Vol. 6602 LNCS. Springer, Berlin, Heidelberg, 1–17. https://doi.org/10.1007/978-3-642-19718-5_1
- Andrew W. Appel, Robert Dockins, Aquinas Hobor, Lennart Beringer, Josiah Dodds, Gordon Stewart, Sandrine Blazy, and Xavier Leroy. 2014. *Program Logics for Certified Compilers*. Cambridge. <https://doi.org/10.1017/CBO9781107256552>
- Mike Barnett, Bor-Yuh Evan Chang, Robert DeLine, Bart Jacobs, and K. Rustan M. Leino. 2006. Boogie: A Modular Reusable Verifier for Object-Oriented Programs. In *Formal Methods for Components and Objects*, Frank S. de Boer, Marcello M. Bonsangue, Susanne Graf, and Willem-Paul de Roeper (Eds.). Springer, Berlin, 364–387. https://doi.org/10.1007/11804192_17
- Patrick Baudin, François Bobot, David Bühler, Loïc Correnson, Florent Kirchner, Nikolai Kosmatov, André Maroneze, Valentin Perrelle, Virgile Prevosto, Julien Signoles, and Nicky Williams. 2021. The dogged pursuit of bug-free C programs: The Frama-C Software Analysis Platform. , 56–67 pages. <https://doi.org/10.1145/3470569>
- Lennart Beringer. 2021. Verified Software Units. In *Lecture Notes in Computer Science*, Vol. 12648 LNCS. Springer, Cham, 118–147. https://doi.org/10.1007/978-3-030-72019-3_5
- Sascha Böhme and Tobias Nipkow. 2010. Sledgehammer: Judgement day. In *Lecture Notes in Computer Science*, Vol. 6173 LNAI. Springer, Berlin, Heidelberg, 107–121. https://doi.org/10.1007/978-3-642-14203-1_9
- Sylvie Boldo, François Clément, Jean-Christophe Filliâtre, Micaela Mayero, Guillaume Melquiond, and Pierre Weis. 2014. Trusting computations: A mechanized proof from partial differential equations to actual program. *Computers and Mathematics with Applications* 68, 3 (2014), 325–352. <https://doi.org/10.1016/j.camwa.2014.06.004>
- Pierre Boutillier, Stéphane Gloudu, Benjamin Grégoire, Hugo Herbelin, Pierre Letouzey, Pierre-Marie Pédrot, Yann Régis-Gianas, Matthieu Sozeau, Arnaud Spiwack, and Enrico Tassi. 2014. Coq 8.4 Reference Manual. (jul 2014). <https://hal.inria.fr/hal-01114602>
- Lukas Bulwahn, Alexander Krauss, Florian Haftmann, Levent Erkök, and John Matthews. 2008. Imperative Functional Programming with Isabelle/HOL. In *Theorem Proving in Higher Order Logics, 21st International Conference, TPHOLs 2008, Montreal, Canada, August 18-21, 2008. Proceedings (Lecture Notes in Computer Science, Vol. 5170)*, Otmame Ait Mohamed, César A. Muñoz, and Sofiène Tahar (Eds.). Springer, 134–149. https://doi.org/10.1007/978-3-540-71067-7_14
- Cristiano Calcagno and Dino Distefano. 2011. Infer: An automatic program verifier for memory safety of C programs. In *Lecture Notes in Computer Science*, Vol. 6617 LNCS. Springer, Berlin, Heidelberg, 459–465. https://doi.org/10.1007/978-3-642-20398-5_33

- Qinxiang Cao, Lennart Beringer, Samuel Gruetter, Josiah Dodds, and Andrew W. Appel. 2018. VST-Floyd: A Separation Logic Tool to Verify Correctness of C Programs. *Journal of Automated Reasoning* 61, 1-4 (jun 2018), 367–422. <https://doi.org/10.1007/S10817-018-9457-5>
- Chiyan Chen and Hongwei Xi. 2005. Combining programming with theorem proving. In *Proceedings of the ACM SIGPLAN International Conference on Functional Programming, ICFP*. 66–77. <https://doi.org/10.1145/1086365.1086375>
- Wei Ngan Chin, Cristina David, Huu Hai Nguyen, and Shengchao Qin. 2012. Automated verification of shape, size and bag properties via user-defined predicates in separation logic. *Science of Computer Programming* 77, 9 (aug 2012), 1006–1036. <https://doi.org/10.1016/j.scico.2010.07.004>
- Joshua M. Cohen and Philip Johnson-Freyd. 2024. A Formalization of Core Why3 in Coq. *Proceedings of the ACM on Programming Languages* 8, POPL (2024).
- Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. 2022. *Introduction to Algorithms*. MIT press.
- Paulo Emilio De Vilhena, François Pottier, and Jacques Henri Jourdan. 2020. Spy game: Verifying a local generic solver in iris. *Proceedings of the ACM on Programming Languages* 4, POPL (2020). <https://doi.org/10.1145/3371101>
- Jean-Christophe Filliâtre and Andrei Paskevich. 2013. Why3—where programs meet provers. In *Programming Languages and Systems: 22nd European Symposium on Programming (ESOP'13)*. Springer, 125–128. https://doi.org/10.1007/978-3-642-37036-6_8
- Robert W. Floyd. 1993. *Assigning Meanings to Programs*. Springer Netherlands, Dordrecht, 65–81. https://doi.org/10.1007/978-94-011-1793-7_4 (reprint of a 1967 paper).
- Alexey Gotsman, Josh Berdine, and Byron Cook. 2011. Precision and the conjunction rule in concurrent separation logic. In *Electronic Notes in Theoretical Computer Science*, Vol. 276. 171–190. <https://doi.org/10.1016/j.entcs.2011.09.021>
- Franjo Ivančić, Gogul Balakrishnan, Aarti Gupta, Sriram Sankaranarayanan, Naoto Maeda, Takashi Imoto, Rakesh Potthengil, and Mustafa Hussain. 2015. Scalable and scope-bounded software verification in Varvel. *Automated Software Engineering* 22, 4 (dec 2015), 517–559. <https://doi.org/10.1007/s10515-014-0164-0>
- F. Ivančić, Z. Yang, M. K. Ganai, A. Gupta, I. Shlyakhter, and P. Ashar. 2005. F-SOFT: Software verification platform. In *Lecture Notes in Computer Science*, Vol. 3576. Springer, Berlin, Heidelberg, 301–306. https://doi.org/10.1007/11513988_31
- Bart Jacobs, Jan Smans, Pieter Philippaerts, Frédéric Vogels, Willem Penninckx, and Frank Piessens. 2011. VeriFast: A powerful, sound, predictable, fast verifier for C and Java. In *Lecture Notes in Computer Science*, Vol. 6617 LNCS. 41–55. https://doi.org/10.1007/978-3-642-20398-5_4
- Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Aleš Bizjak, Lars Birkedal, and Derek Dreyer. 2018. Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *Journal of Functional Programming* 28 (2018). <https://doi.org/10.1017/S0956796818000151>
- Ralf Jung, Rodolphe Lepigre, Gaurav Parthasarathy, Marianna Rapoport, Amin Timany, Derek Dreyer, and Bart Jacobs. 2020. The future is ours: prophecy variables in separation logic. *Proc. ACM Program. Lang.* 4, POPL (2020), 45:1–45:32. <https://doi.org/10.1145/3371113>
- Robbert Krebbers, Ralf Jung, Ales Bizjak, Jacques-Henri Jourdan, Derek Dreyer, and Lars Birkedal. 2017a. The Essence of Higher-Order Concurrent Separation Logic. In *Programming Languages and Systems - 26th European Symposium on Programming, ESOP 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings (Lecture Notes in Computer Science, Vol. 10201)*, Hongseok Yang (Ed.). Springer, 696–723. https://doi.org/10.1007/978-3-662-54434-1_26
- Robbert Krebbers, Amin Timany, and Lars Birkedal. 2017b. Interactive proofs in higher-order concurrent separation logic. *ACM SIGPLAN Notices* 52, 1 (jan 2017), 205–217. <https://doi.org/10.1145/3009837.3009855>
- Daniel Kroening and Michael Tautschnig. 2014. CBMC - C Bounded Model Checker. In *Lecture Notes in Computer Science*, Vol. 8413 LNCS. Springer, Berlin, Heidelberg, 389–391. https://doi.org/10.1007/978-3-642-54862-8_26
- K. Rustan M. Leino. 2010. Dafny: An automatic program verifier for functional correctness. In *Lecture Notes in Computer Science*, Vol. 6355 LNAI. 348–370. https://doi.org/10.1007/978-3-642-17511-4_20
- Xavier Leroy. 2009. Formal verification of a realistic compiler. *Commun. ACM* 52, 7 (jul 2009), 107–115. <https://doi.org/10.1145/1538788.1538814>
- Gregory Malecha, Gordon Stewart, František Farka, Jasper Haag, and Yoichi Hirai. 2022. Developing With Formal Methods at BedRock Systems, Inc. *IEEE Security & Privacy* 20, 3 (2022), 33–42. <https://doi.org/10.1109/MSEC.2022.3158196>
- Guido Martínez, Danel Ahman, Victor Dumitrescu, Nick Giannarakis, Chris Hawblitzel, Cătălin Hrițcu, Monal Narasimhamurthy, Zoe Paraskevopoulou, Clément Pit-Claudel, Jonathan Protzenko, Tahina Ramananandro, Aseem Rastogi, and Nikhil Swamy. 2019. Meta-F*: Proof Automation with SMT, Tactics, and Metaprograms. In *Lecture Notes in Computer Science*, Vol. 11423 LNCS. Springer Verlag, 30–59. https://doi.org/10.1007/978-3-030-17184-1_2 arXiv:1803.06547
- Ike Mulder, Robbert Krebbers, and Herman Geuvers. 2022. Diaframe: Automated Verification of Fine-Grained Concurrent Programs in Iris. (2022), 16. <https://doi.org/10.1145/3519939.3523432>
- Peter Müller, Malte Schwerhoff, and Alexander J. Summers. 2017. Viper: A verification infrastructure for permission-based reasoning. In *Dependable Software Systems Engineering*. Vol. 9583. Springer Verlag, 104–125. <https://doi.org/10.3233/978->

1-61499-810-5-104

- Peter W. O'Hearn. 2004. Resources, concurrency, and local reasoning. *Lecture Notes in Computer Science* 2986 (2004), 1–2. https://doi.org/10.1007/978-3-540-24725-8_1
- Christopher Pulte, Dhruv C. Makwana, Thomas Sewell, Kayvan Memarian, Peter Sewell, and Neel Krishnaswami. 2023. CN: Verifying Systems C Code with Separation-Logic Refinement Types. *Proc. ACM Program. Lang.* 7, POPL, Article 1 (jan 2023), 32 pages. <https://doi.org/10.1145/3571194>
- John C. Reynolds. 2002. Separation logic: A logic for shared mutable data structures. In *Proceedings - Symposium on Logic in Computer Science*. 55–74. <https://doi.org/10.1109/lics.2002.1029817>
- 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 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*. ACM, 158–174. <https://doi.org/10.1145/3453483.3454036>
- Aaron Joseph Turon, Jacob Thamsborg, Amal Ahmed, Lars Birkedal, and Derek Dreyer. 2013. Logical relations for fine-grained concurrency. In *The 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '13, Rome, Italy - January 23 - 25, 2013*, Roberto Giacobazzi and Radhia Cousot (Eds.). ACM, 343–356. <https://doi.org/10.1145/2429069.2429111>
- Viktor Vafeiadis. 2011. Concurrent separation logic and operational semantics. In *Electronic Notes in Theoretical Computer Science*, Vol. 276. Elsevier, 335–351. <https://doi.org/10.1016/j.entcs.2011.09.029>
- Niki Vazou, Eric L Seidel, Ranjit Jhala, Dimitrios Vytiniotis, and Simon Peyton-Jones. 2014. Refinement types for Haskell. In *ACM SIGPLAN Notices*, Vol. 49. 269–282. <https://doi.org/10.1145/2628136.2628161>
- Niki Vazou, Anish Tondwalkar, Vikraman Choudhury, Ryan G Scott, Ryan R Newton, Philip Wadler, and Ranjit Jhala. 2018. Refinement reflection: complete verification with SMT. *Proceedings of the ACM on Programming Languages* 2, POPL (2018), 1–31. <https://doi.org/10.1145/3158141> arXiv:1711.03842
- Zhongye Wang, Qinxiang Cao, and Yichen Tao. 2023. Verifying Programs with Logic and Extended Proof Rules: Deep Embedding v.s. Shallow Embedding. <https://doi.org/10.48550/arXiv.2310.17616> arXiv:2310.17616 [cs.PL]
- Bohua Zhan. 2016. AUTO2, a saturation-based heuristic prover for higher-order logic. In *International Conference on Interactive Theorem Proving*. Springer, 441–456. https://doi.org/10.1007/978-3-319-43144-4_27
- Bohua Zhan. 2018. Efficient Verification of Imperative Programs Using Auto2. In *Tools and Algorithms for the Construction and Analysis of Systems - 24th International Conference, TACAS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 10805)*, Dirk Beyer and Marieke Huisman (Eds.). Springer, 23–40. https://doi.org/10.1007/978-3-319-89960-2_2
- Litao Zhou, Jianxing Qin, Qinshi Wang, Andrew W. Appel, and Qinxiang Cao. 2023a. VST-A: A Foundationally Sound Annotation Verifier. *CoRR* abs/1909.00097 (2023). <https://doi.org/10.48550/arXiv.1909.00097> arXiv:1909.00097
- Litao Zhou, Jianxing Qin, Qinshi Wang, Andrew W. Appel, and Qinxiang Cao. 2023b. VST-A: A Foundationally Sound Annotation Verifier. <https://doi.org/10.5281/zenodo.10125129>

Received 2023-07-11; accepted 2023-11-07