

VeriFast's separation logic: a logic without lateres for modular verification of fine-grained concurrent programs

Bart Jacobs

KU Leuven, Department of Computer Science, DistriNet Research Group, Leuven, Belgium

ABSTRACT VeriFast is one of the leading tools for semi-automated modular formal program verification. A central feature of VeriFast is its support for *higher-order ghost code*, which enables its support for expressively specifying fine-grained concurrent modules, without the need for the *later* modality. We present the first formalization and soundness proof for this aspect of VeriFast's logic, and we compare it both to Iris, a state-of-the-art logic for fine-grained concurrency which features the later modality, as well as to some recent proposals for Iris-like reasoning without the later modality.

KEYWORDS Separation Logic, Fine-Grained Concurrency, Modular Verification.

1. Introduction

VeriFast (Vogels et al. 2015) is one of the leading tools for semi-automated modular formal verification of single-threaded and multithreaded C, Java, and Rust programs. It symbolically executes each function/method of the program, using a separation logic (O'Hearn et al. 2001; Reynolds 2002) representation of memory. It requires programs to be *annotated* with function/method preconditions and postconditions and loop invariants, as well as *ghost declarations*, such as definitions of separation logic predicates that specify the layout of data structures, and *ghost commands* for folding and unfolding predicates as well as invoking *lemma functions*, functions consisting entirely of ghost code. For expressive specification of fine-grained concurrent modules, it supports *higher-order ghost code*, in the form of *lemma function pointers* and *lemma function pointer type assertions*. While the general ideas underlying this specification approach have been described earlier (Jacobs & Piessens 2011), as have some examples of their use for solving verification challenges (Jacobs et al. 2015; Jacobs 2016), in this paper we present the first formalization and soundness proof for this aspect of VeriFast's logic. We define the programming language

JOT reference format:

Bart Jacobs. *VeriFast's separation logic: a logic without lateres for modular verification of fine-grained concurrent programs*. Journal of Object Technology. Vol. 25, No. 1, 2026. Licensed under Attribution 4.0 International (CC BY 4.0) <http://dx.doi.org/10.5381/jot.2026.25.1.a4>

```
let x = cons(0) in
( FAA(x, 1) || FAA(x, 1) );
let v = *x in
assert v = 2
```

Figure 1 An example program. `cons(0)` allocates a memory cell, initializes it to 0, and returns its address. The `FAA` command performs a sequentially consistent atomic fetch-and-add operation. `c1 || c2` is the parallel composition of commands `c1` and `c2`. `*x` returns the value stored at address `x`.

and introduce the running example in §2, define the syntax of annotations in §3, formalize the program logic implemented by VeriFast's symbolic execution algorithm in §4, and prove its soundness in §5. We discuss related work in §6.

2. Programming language

In order to focus on the complexities of the logic rather than those of the programming language, we present VeriFast's separation logic in the context of a trivial concurrent programming language whose syntax is given in Fig. 2 and whose small-step operational semantics is given in Fig. 3. An example program that allocates a memory cell, increments it twice in parallel, and then asserts that the cell's value equals two is shown in Fig. 1.

$$\begin{aligned}
z &\in \mathbb{Z}, x \in \mathcal{X} \\
e &::= z \mid x \\
i &::= \mathbf{cons}(e) \mid \mathbf{FAA}(e, e) \mid *e \mid \mathbf{assert} \ e = e \\
c &::= e \mid i \mid \mathbf{let} \ x = c \ \mathbf{in} \ c \mid (c \parallel c)
\end{aligned}$$

Figure 2 Syntax of the expressions e , instructions i , and commands c of the programming language. We assume a set \mathcal{X} of program variable names. $c; c'$ is a shorthand for $\mathbf{let} _ = c \ \mathbf{in} \ c'$, where $_$ is a designated element of \mathcal{X}

$$\begin{aligned}
&\frac{\ell \notin \text{dom } h}{(h, \mathbf{cons}(v)) \rightarrow (h[\ell := v], \ell)} \\
&\frac{\ell \in \text{dom } h}{(h, \mathbf{FAA}(\ell, z)) \rightarrow (h[\ell := h(\ell) + z], h(\ell))} \\
&\frac{\ell \in \text{dom } h}{(h, * \ell) \rightarrow (h, h(\ell))} \quad (h, \mathbf{assert} \ v = v) \rightarrow (h, 0) \\
&(h, \mathbf{let} \ x = v \ \mathbf{in} \ c) \rightarrow (h, c[v/x]) \\
&\frac{(h, c) \rightarrow (h', c')}{(h, \mathbf{let} \ x = c \ \mathbf{in} \ c'') \rightarrow (h', \mathbf{let} \ x = c' \ \mathbf{in} \ c'')} \\
&\frac{(h, c) \rightarrow (h', c')}{(h, (c \parallel c'')) \rightarrow (h', (c' \parallel c''))} \\
&\frac{(h, c) \rightarrow (h', c')}{(h, (c'' \parallel c)) \rightarrow (h', (c'' \parallel c'))} \quad (h, v \parallel v') \rightarrow (h, 0)
\end{aligned}$$

Figure 3 Small-step operational semantics of the programming language

We define the multiset of threads of a command c as follows:

$$\text{thrds}(c) = \begin{cases} \text{thrds}(c_1) & \text{if } c = \mathbf{let} \ x = c_1 \ \mathbf{in} \ c_2 \\ \text{thrds}(c_1) \uplus \text{thrds}(c_2) & \text{if } c = (c_1 \parallel c_2) \\ \{\{c\}\} & \text{otherwise} \end{cases}$$

We say a configuration (h, c) is *reducible* if it can make a step:

$$\frac{(h, c) \rightarrow (h', c')}{\text{red}(h, c)}$$

We say a configuration is *finished* if its command is a value.

$$\text{finished}(h, z)$$

We say a configuration is *okay* if each thread is either reducible or finished.

$$\frac{\forall c_t \in \text{thrds}(c). \text{finished}(h, c_t) \vee \text{red}(h, c_t)}{\text{ok}(h, c)}$$

We say a configuration is *safe* if each configuration reachable from it is okay.

$$\frac{\forall h', c'. (h, c) \rightarrow^* (h', c') \Rightarrow \text{ok}(h', c')}{\text{safe}(h, c)}$$

We say a program c is *safe* if (\emptyset, c) is safe. The goal of the logic that we present here is to prove that a given program is safe. This implies that it does not access unallocated memory and that there are no assertion failures.¹

3. Annotated programs

When verifying a program with VeriFast, the user must first insert *annotations*, specifically *ghost declarations* and *ghost commands*, to obtain an *annotated program*. The syntax of ghost declarations and ghost commands is shown in Fig. 4. An annotated version of the example program is shown in Fig. 5. An annotated program may refer to ghost constructs declared in the *VeriFast prelude*, shown in Fig. 6.

There are two kinds of ghost declarations: *lemma type declarations* and *predicate constructor declarations*. These give meaning to *lemma type names* $t \in \mathcal{T}$ and *predicate constructor names* $p \in \mathcal{P}$. Conceptually, a lemma type is a predicate over a *lemma value* $\lambda \bar{g}. C$, a parameterized ghost command. A predicate constructor is a named, parameterized assertion. Applying a predicate constructor to an argument list produces a *predicate value* $p(\bar{V})$.

Besides integers, lemma values, and predicate values, ghost values may be pairs of ghost values, unit values $()$, and finite sets of ghost values.

Resources may be shared among threads using *atomic spaces* (analogous to Iris *invariants* (Jung et al. 2015, 2018)). An atomic space is (non-uniquely) identified by a *name* (any ghost value) and an *invariant* (a predicate value) (but there may be multiple atomic spaces with the same name and invariant at any

¹ In fact, the logic also proves that there are no data races, but for simplicity we do not consider data races here.

	$t \in \mathcal{T}$	lemma type names
	$p \in \mathcal{P}$	predicate constructor names
	$g \in \mathcal{G}$	ghost variable names
	$\pi \in \mathbb{R}^+$	fractions
ghost values $V ::=$	$z \mid (V, V) \mid () \mid \{\bar{V}\}$	
	$\mid p(\bar{V})$	predicate values
	$\mid \lambda \bar{g}. G$	lemma values
ghost expressions $E ::=$	$V \mid x \mid g \mid E + E$	
	$\mid p(\bar{E})$	predicate constructor applications
	$\mid (E, E) \mid ()$	pair expressions, empty tuple
	$\mid \emptyset \mid \{E\} \mid E \cup E \mid E \setminus E$	set expressions
assertions $a ::=$	$[\pi]E \mapsto E$	points-to assertions
	$[\pi]E \mapsto_g E$	ghost cell points-to assertions
	$E()$	predicate assertions
	$[\pi]\mathbf{atomic_space}(E, E)$	atomic space assertions
	$E : t(\bar{E})$	lemma type assertions
	$\exists g. a$	
	$\mathbf{atomic_spaces}(E)$	atomic spaces assertions
	$\mathbf{heap}(E)$	heap chunk assertions
	$a * a$	separating conjunctions
$gdecl ::=$	$\mathbf{lem_type} \ t(\bar{g}) = \mathbf{lem}(\bar{g}) \ \mathbf{forall} \ \bar{g} \ \mathbf{req} \ a \ \mathbf{ens} \ a$	
	$\mid \mathbf{pred_ctor} \ p(\bar{g})() = a$	
$I ::=$	$E(\bar{E})$	
	$\mid \mathbf{gcons}(E) \mid *E \leftarrow_g E$	
	$\mid \mathbf{open_atomic_space}(E, E) \mid \mathbf{close_atomic_space}(E, E)$	
	$\mid E \leftarrow_h E$	heap chunk update
$G ::=$	$I \mid \mathbf{glet}_i \ g = G \ \mathbf{in} \ G$	
$C ::=$	$G \mid \mathbf{produce_lem_ptr_chunk} \ t(\bar{E})(\bar{g}) \ \{ G \}$	
	$\mid \mathbf{create_atomic_space}(E, E) \mid \mathbf{destroy_atomic_space}(E, E)$	
$\hat{c} ::=$	$e \mid i \mid \mathbf{let} \ x = \hat{c} \ \mathbf{in} \ \hat{c} \mid \hat{c} \mid \hat{c} \mid \mathbf{glet} \ g = C \ \mathbf{in} \ \hat{c}$	

Figure 4 Syntax of ghost declarations $gdecl$, ghost instructions I , inner ghost commands G , outer ghost commands C (collectively called *ghost commands*), and annotated commands \hat{c} . Heap chunk assertions and heap chunk update commands are *internal*; they are not accepted by VeriFast in source code and are introduced here only for the sake of the soundness proof.

```

pred_ctor Inv(x, g1, g2)() =
  ∃v1, v2. [1/2]g1 ↦g v1 * [1/2]g2 ↦g v2 * x ↦ v1 + v2
pred_ctor pre1(x, g1, g2)() =
  quad[1/2]atomic_space(Nx, Inv(x, g1, g2)) * [1/2]g1 ↦g 0
pred_ctor post1(x, g1, g2)() =
  [1/2]atomic_space(Nx, Inv(x, g1, g2)) * [1/2]g1 ↦g 1
pred_ctor pre2(x, g1, g2)() =
  [1/2]atomic_space(Nx, Inv(x, g1, g2)) * [1/2]g2 ↦g 0
pred_ctor post2(x, g1, g2)() =
  [1/2]atomic_space(Nx, Inv(x, g1, g2)) * [1/2]g2 ↦g 1

let x = cons(0) in
glet g1 = gcons(0) in
glet g2 = gcons(0) in
create_atomic_space(Nx, Inv(x, g1, g2));
(
  produce_lem_ptr_chunk
  FAA_ghop(x, 1, pre1(x, g1, g2), post1(x, g1, g2))(op) {
    open_atomic_space(Nx, Inv(x, g1, g2));i
    op();i
    *g1 ←g 1;i
    close_atomic_space(Nx, Inv(x, g1, g2))
  };
  FAA(x, 1)
||
  produce_lem_ptr_chunk
  FAA_ghop(x, 1, pre2(x, g1, g2), post2(x, g1, g2))(op) {
    open_atomic_space(Nx, Inv(x, g1, g2));i
    op();i
    *g2 ←g 1;i
    close_atomic_space(Nx, Inv(x, g1, g2))
  };
  FAA(x, 1)
);
destroy_atomic_space(Nx, Inv(x, g1, g2));
let v = *x in
assert v = 2

```

Figure 5 VeriFast proof of the example program. $Nx \triangleq ()$.

```

lem_type FAA_op(x, n, P, Q) = lem()
  forall v
  req x ↦ v * P()
  ens x ↦ v + n * Q()
lem_type FAA_ghop(x, n, pre, post) = lem(op)
  forall P, Q
  req atomic_spaces(∅) * op : FAA_op(x, n, P, Q) * P()
  * pre()
  ens atomic_spaces(∅) * op : FAA_op(x, n, P, Q) * Q()
  * post()
pred_ctor heap_(h)() = heap(h)

```

Figure 6 The ghost prelude (built-in ghost declarations). The declaration of `heap_` is *internal*. It is not meant to be used in annotated programs; it is introduced here only for the sake of the soundness proof.

given time). At any point in time, ownership of the stock of logical resources in the system is distributed among the threads and the atomic spaces. That is, at any point, each logical resource is owned either by exactly one thread or by exactly one atomic space, or has been leaked irrecoverably. (More precisely, given that fractional resources are supported, the bundles of resources owned by the threads and the atomic spaces sum up to a logical heap that contains each physical points-to chunk only once and each `atomic_space` chunk only as many times as there are atomic spaces with that name and invariant, etc.) Creating an atomic space transfers a bundle of resources satisfying the atomic space's invariant from the creating thread to the newly created atomic space. Opening an atomic space transfers the resources owned by the atomic space to the opening thread; closing an atomic space again transfers a bundle of resources satisfying the atomic space's invariant from the closing thread to the atomic space. Destroying an atomic space transfers ownership of the resources owned by the atomic space to the destroying thread. To destroy an atomic space, the destroying thread must have full ownership of the atomic space. To open it, only partial ownership is required. (To close it, no ownership is required. If no such atomic space exists, the resources are leaked.) To prevent the same atomic space from being opened when it is already open, the set of opened atomic spaces is tracked using an `atomic_spaces(S)` chunk, where S is a set of the name-invariant pairs of the atomic spaces that are currently open.²

Lemma type assertions $V : t(\bar{V})$ assert that a given lemma value V is of a given lemma type t , applied to a given lemma type argument list \bar{V} . Such assertions are *linear*. To call a lemma, a full lemma type chunk for that lemma must be available, and it becomes unavailable for the duration of the call. A lemma type chunk is produced by the `produce_lem_ptr_chunk` ghost command. Since that command is not allowed inside lemmas, the stock of lemma type

² This means it is not possible to open two atomic spaces with the same name-invariant pair at the same time, even if multiple such atomic spaces exist.

chunks in the system only decreases as the lemma call stack grows; absence of infinite lemma recursion follows trivially.³

Intermediate results produced by ghost commands can be stored in *ghost variables*, which are like program variables except that they are in a separate namespace and can therefore never hide a program variable.⁴ To facilitate reasoning about concurrent programs, annotated programs can furthermore allocate *ghost cells*; these are like physical memory locations except that they are allocated in a separate *ghost heap* and mutated using separate *ghost cell mutation commands*.

Points-to chunks, ghost points-to chunks, and atomic spaces can be owned *fractionally*, which allows them to be shared temporarily or permanently among multiple threads. A *fractional chunk* has a *coefficient* which is a positive real number.

4. Verification of annotated programs

In this section we formalize the program logic implemented by VeriFast's symbolic execution algorithm. We abstract over the mechanics of symbolic execution, the essence of which is described in Featherweight VeriFast (Vogels et al. 2015). In particular, the tool generally requires **open** and **close** ghost commands to unfold and fold predicates. Instead, here we use *semantic assertions*; predicates are fully unfolded during the interpretation of syntactic assertions as semantic assertions.

Core to VeriFast's verification approach is the concept of a *chunk* α :

$$\alpha ::= V \mapsto V \mid V \mapsto_g V \mid \mathbf{atomic_space}(V, V) \mid V : t(\bar{V}) \mid \mathbf{atomic_spaces}(V) \mid \mathbf{heap}(V)$$

A *logical heap* H is a function from chunks to nonnegative real numbers:

$$H \in \text{LogicalHeaps} = \text{Chunks} \rightarrow \mathbb{R}^+$$

We say a logical heap is *weakly consistent*, denoted $\text{wok } H$ if no points-to chunk or ghost points-to chunk is present with a coefficient greater than 1, and no two (fractions of) points-to chunks or two (fractions of) ghost points-to chunks are present with the same left-hand side (address) but a different right-hand side (stored value).

³ This is a simplification with respect to the actual VeriFast tool, which does support production of lemma type chunks inside lemmas, using a variant of the **produce_lem_ptr_chunk** syntax that additionally takes a block of ghost code. The chunk is available only until the end of that block. Now, suppose there is an infinite lemma call stack. Since the program text contains only finitely many **produce_lem_ptr_chunk** commands, among the lemmas that appear infinitely often in that call stack, there is one that is syntactically maximal, i.e. that is not itself contained within another lemma that also appears infinitely often. It follows that from some point on, the call stack contains no lemmas bigger than this maximal one. Since a lemma type chunk for a given lemma can only be produced by a bigger lemma (since the latter's body must contain a **produce_lem_ptr_chunk** command producing the former's), the stock of lemma type chunks for this maximal lemma will, from that point on, only decrease, which leads to a contradiction. (Note: for measuring the size of a lemma, the size of contained lemma values is not taken into account. It follows that substitution of values for ghost variables never affects the size of a lemma.)

⁴ In the actual VeriFast tool, they are in the same namespace, but VeriFast checks that real code never uses a ghost variable.

We define *satisfaction* of an assertion a by a logical heap H , denoted $H \models a$, inductively as follows:

$$\frac{H(\alpha) \geq \pi}{H \models [\pi]\alpha}$$

$$\frac{\mathbf{pred_ctor } p(\bar{g})() = a \quad |\bar{V}| = |\bar{g}| \quad H \models a[\bar{V}/\bar{g}]}{H \models p(\bar{V})()}$$

$$\frac{H \models a[V/g] \quad H \models a \quad H' \models a'}{H + H' \models a * a'}$$

A *semantic assertion* is a set of logical heaps. We define the *interpretation* $\llbracket a \rrbracket$ of an assertion as a semantic assertion as $\llbracket a \rrbracket = \{H \mid H \models a\}$.

We define *correctness* of an annotated command or ghost command \hat{c} with respect to a precondition P and a postcondition Q (both semantic assertions), denoted $\{P\} \hat{c} \{Q\}$, inductively in Fig. 7. We define implication of semantic assertions as follows:

$$P \Rightarrow Q \triangleq \forall H \in P. \text{wok } H \Rightarrow H \in Q$$

Note: nesting **produce_lem_ptr_chunk** commands is not allowed.

A correctness proof outline for the example annotated program is shown in Fig. 8.

We say an annotated program \hat{c} is *correct* if $\{\text{True}\} \hat{c} \{\text{True}\}$.

We define the erasure of an annotated command \hat{c} to a command $c = \text{erasure}(\hat{c})$ as follows:

$$\begin{aligned} \text{erasure}(c) &= c \\ \text{erasure}(\mathbf{let } x = \hat{c} \mathbf{in } \hat{c}') &= \mathbf{let } x = \text{erasure}(\hat{c}) \mathbf{in } \text{erasure}(\hat{c}') \\ \text{erasure}(\hat{c} \parallel \hat{c}') &= \text{erasure}(\hat{c}) \parallel \text{erasure}(\hat{c}') \\ \text{erasure}(\mathbf{glet } g = C \mathbf{in } \hat{c}) &= \text{erasure}(\hat{c}) \end{aligned}$$

Theorem 1. *If an annotated program \hat{c} is correct, then its erasure $\text{erasure}(\hat{c})$ is safe.*

5. Soundness

We say a logical heap is *strongly consistent*, denoted $\text{sok } H$, if, for every $V : t(\bar{V})$ such that $H(V : t(\bar{V})) > 0$, we have that V semantically is of type $t(\bar{V})$, denoted $\models V : t(\bar{V})$, defined as follows:

$$\frac{\mathbf{lem_type } t(\bar{g}')(\bar{g}'') \mathbf{req } a \mathbf{ens } a'' \quad |\bar{V}| = |\bar{g}'| \quad |\bar{g}| = |\bar{g}''|}{\forall \bar{V}'. \frac{\{ \llbracket a[\bar{V}/\bar{g}', \bar{V}'/\bar{g}''] \rrbracket \} G[\bar{V}'/\bar{g}'] \{ \llbracket a'[\bar{V}/\bar{g}', \bar{V}'/\bar{g}''] \rrbracket \}}{H \models \lambda \bar{g}. G : t(\bar{V})}}$$

A *ghost heap* \hat{h} is a partial function from integers to ghost values.

An *atomic spaces bag* A is a multiset of pairs $((V, V), H)$ of name-invariant pairs and logical heaps, such that for each element $((_, V), H)$ we have $H \models V()$. We define the atomic

$$\begin{array}{c}
\{\text{True}\} \mathbf{cons}(V) \{\text{res} \mapsto V\} \quad \{[\pi]\ell \mapsto V\} * \ell \{[\pi]\ell \mapsto V \wedge \text{res} = V\} \quad \{\ell \mapsto V\} \ell \leftarrow V' \{\ell \mapsto V'\} \\
\\
\frac{\{P\} \hat{c} \{R\} \quad \forall v. \{R[v/\text{res}]\} \hat{c}'[v/x] \{Q\}}{\{P\} \mathbf{let} x = \hat{c} \mathbf{in} \hat{c}' \{Q\}} \quad \frac{\{V : \text{FAA_ghop}(\ell, z, V', V'') * \llbracket V'() \rrbracket\} \quad \mathbf{FAA}(\ell, z)}{\{V : \text{FAA_ghop}(\ell, z, V', V'') * \llbracket V''() \rrbracket\}} \quad \frac{\{P\} \hat{c} \{Q\} \quad \{P'\} \hat{c}' \{Q'\}}{\{P * P'\} \hat{c} \parallel \hat{c}' \{Q * Q'\}} \\
\\
\{\text{True}\} \mathbf{gcons}(V) \{\text{res} \mapsto_g V\} \quad \{\ell \mapsto_g V\} \ell \leftarrow_g V' \{\ell \mapsto_g V'\} \\
\\
\llbracket V'() \rrbracket \mathbf{create_atomic_space}(V, V') \{\mathbf{atomic_space}(V, V')\} \\
\\
\frac{(V, V') \notin S \quad \{\mathbf{atomic_spaces}(S) * \llbracket V'() \rrbracket\} \quad \mathbf{close_atomic_space}(V, V') \quad \{\mathbf{atomic_spaces}(S \setminus \{(V, V')\})\}}{\{\mathbf{atomic_spaces}(S) * [\pi]\mathbf{atomic_space}(V, V')\} \quad \mathbf{open_atomic_space}(V, V') \quad \{\mathbf{atomic_spaces}(S \cup \{(V, V')\}) * [\pi]\mathbf{atomic_space}(V, V') * \llbracket V'() \rrbracket\}} \\
\\
\{\mathbf{atomic_space}(V, V')\} \mathbf{destroy_atomic_space}(V, V') \{\llbracket V'() \rrbracket\} \\
\\
\frac{\mathbf{lem_type} t(\bar{g}) = \mathbf{lem}(\bar{g}') \mathbf{req} a \mathbf{ens} a' \quad |\bar{V}| = |\bar{g}| \quad |\bar{g}''| = |\bar{g}'| \quad \forall \bar{V}'. |\bar{V}'| = |\bar{g}'| \Rightarrow \{\llbracket a[\bar{V}/\bar{g}, \bar{V}'/\bar{g}'] \rrbracket\} G[\bar{V}'/\bar{g}''] \{\llbracket a'[\bar{V}/\bar{g}, \bar{V}'/\bar{g}'] \rrbracket\}}{\{\text{True}\} \mathbf{produce_lem_ptr_chunk} t(\bar{V})(\bar{g}'') \{G\} \{\text{res} : t(\bar{V})\}} \\
\\
\frac{\mathbf{lem_type} t(\bar{g}) = \mathbf{lem}(\bar{g}') \mathbf{req} a \mathbf{ens} a' \quad |\bar{V}'| = |\bar{g}'|}{\{V : t(\bar{V}) * \llbracket a[\bar{V}/\bar{g}, \bar{V}'/\bar{g}'] \rrbracket\} V(\bar{V}') \{V : t(\bar{V}) * \llbracket a'[\bar{V}/\bar{g}, \bar{V}'/\bar{g}'] \rrbracket\}} \\
\\
\{\mathbf{heap}(h) * \ell \mapsto _ \} \ell \leftarrow_h v \{\mathbf{heap}(h[\ell := v]) * \ell \mapsto v\} \quad \frac{\{P\} \hat{c} \{Q\}}{\{P * R\} \hat{c} \{Q * R\}} \quad \frac{\forall V. \{P[V/g]\} \hat{c} \{Q\}}{\{\exists g. P\} \hat{c} \{Q\}} \\
\\
\frac{P \Rightarrow P' \quad \{P'\} \hat{c} \{Q\} \quad Q \Rightarrow Q'}{\{P\} \hat{c} \{Q'\}}
\end{array}$$

Figure 7 Correctness of annotated commands and ghost commands. We use \hat{c} to range over both annotated commands and ghost commands.

```

{emp}
let x = cons(0) in glet g1 = gcons(0) in glet g2 = gcons(0) in
{x ↦ 0 * g1 ↦g 0 * g2 ↦g 0}
close Inv(x, g1, g2)();
{Inv(x, g1, g2)() * [1/2]g1 ↦g 0 * [1/2]g2 ↦g 0}
create_atomic_space(Nx, Inv(x, g1, g2));
{atomic_space(Nx, Inv(x, g1, g2)) * [1/2]g1 ↦g 0 * [1/2]g2 ↦g 0}
(
  {[1/2]atomic_space(Nx, Inv(x, g1, g2)) * [1/2]g1 ↦g 0}
  glet lem = produce_lem_ptr_chunk FAA_ghop(x, 1, pre1(x, g1, g2), post1(x, g1, g2))(op) {
    For all P, Q,
    {atomic_spaces(∅) * op : FAA_op(x, 1, P, Q) * P() * pre1(x, g1, g2)()}
    open pre1(x, g1, g2)();
    {
      atomic_spaces(∅) * op : FAA_op(x, 1, P, Q) * P() *
      [1/2]atomic_space(Nx, Inv(x, g1, g2)) * [1/2]g1 ↦g 0
    }
    open_atomic_space(Nx, Inv(x, g1, g2)); open Inv(x, g1, g2)();
    {
      ∃v2. atomic_spaces({(Nx, Inv(x, g1, g2))}) * op : FAA_op(x, 1, P, Q) * P() *
      [1/2]atomic_space(Nx, Inv(x, g1, g2)) * g1 ↦g 0 * [1/2]g2 ↦g v2 * x ↦ v2
    }
    For all v2,
    {
      atomic_spaces({(Nx, Inv(x, g1, g2))}) * op : FAA_op(x, 1, P, Q) * P() *
      [1/2]atomic_space(Nx, Inv(x, g1, g2)) * g1 ↦g 0 * [1/2]g2 ↦g v2 * x ↦ v2
    }
    op();
    {
      atomic_spaces({(Nx, Inv(x, g1, g2))}) * op : FAA_op(x, 1, P, Q) * Q() *
      [1/2]atomic_space(Nx, Inv(x, g1, g2)) * g1 ↦g 0 * [1/2]g2 ↦g v2 * x ↦ 1 + v2
    }
    *g1 ←g 1;
    {
      atomic_spaces({(Nx, Inv(x, g1, g2))}) * op : FAA_op(x, 1, P, Q) * Q() *
      [1/2]atomic_space(Nx, Inv(x, g1, g2)) * g1 ↦g 1 * [1/2]g2 ↦g v2 * x ↦ 1 + v2
    }
    close Inv(x, g1, g2)(); close_atomic_space(Nx, Inv(x, g1, g2));
    {
      atomic_spaces(∅) * op : FAA_op(x, 1, P, Q) * Q() *
      [1/2]atomic_space(Nx, Inv(x, g1, g2)) * [1/2]g1 ↦g 1
    }
    close post1(x, g1, g2)()
    {atomic_spaces(∅) * op : FAA_op(x, 1, P, Q) * Q() * post1(x, g1, g2)()}
  } in
  {[1/2]atomic_space(Nx, Inv(x, g1, g2)) * [1/2]g1 ↦g 0 * lem : FAA_ghop(x, 1, pre1(x, g1, g2), post1(x, g1, g2))}
  close pre1(x, g1, g2)();
  {pre1(x, g1, g2)() * lem : FAA_ghop(x, 1, pre1(x, g1, g2), post1(x, g1, g2))}
  FAA(x, 1);
  {post1(x, g1, g2)() * lem : FAA_ghop(x, 1, pre1(x, g1, g2), post1(x, g1, g2))}
  open post1(x, g1, g2)()
  {[1/2]atomic_space(Nx, Inv(x, g1, g2)) * [1/2]g1 ↦g 1}
  ||
  ...
);
{atomic_space(Nx, Inv(x, g1, g2)) * [1/2]g1 ↦g 1 * [1/2]g2 ↦g 1}
destroy_atomic_space(Nx, Inv(x, g1, g2)); open Inv(x, g1, g2)();
{g1 ↦g 1 * g2 ↦g 1 * x ↦ 2}
let v = *x in
assert v = 2

```

Figure 8 Proof outline for the example proof

space chunks $\text{chunks}(A)$ and the atomic spaces total owned heap $\text{heap}(A)$ as follows:

$$\begin{aligned}\text{chunks}(A) &= \{\{\mathbf{atomic_space}(V, V') \mid ((V, V'), _) \in A\}\} \\ \text{heap}(A) &= \biguplus_{(_, H) \in A} H\end{aligned}$$

A stock of lemma type chunks Σ is a multiset of (V, t, \bar{V}) tuples. We say such a stock is *consistent* if for each (V, t, \bar{V}) in Σ , V is semantically of type $t(\bar{V})$.

We say a heap h and logical heap H are *consistent*, denoted $h \sim H$, if there exists a ghost heap \hat{h} , an atomic spaces bag A , and a consistent stock of lemma type chunks Σ such that $h + \hat{h} + \text{chunks}(A) + \Sigma \geq \text{heap}(A) + H$, where a heap is interpreted as a set of \mapsto chunks and a ghost heap is interpreted as a set of \mapsto_g chunks. Notice: if $h \sim H$, it follows that H is strongly consistent.

We define the *weakest precondition* for n steps of a command c with respect to postcondition Q , denoted $\text{wp}_n(c, Q)$, as the semantic assertion that is true for a logical heap H if either c is a value and $H \in Q$ or $n = 0$ or for each heap h and frame H' such that $h \sim H + H'$, all threads of c are either finished or reducible and for each step that (h, c) can make to some configuration (h', c') , there exists a logical heap H'' such that $h' \sim H'' + H'$ and H'' satisfies the weakest precondition of c' with respect to Q for $n - 1$ steps:

$$\begin{aligned}H \in \text{wp}_n(c, Q) &\Leftrightarrow \\ &\text{finished}(\emptyset, c) \wedge H \in Q \vee n = 0 \vee \\ &\forall h, H'. h \sim H + H' \Rightarrow (h, c) \text{ ok} \wedge \\ &\quad \forall h', c'. (h, c) \rightarrow (h', c') \Rightarrow \\ &\quad \exists H''. h' \sim H'' + H' \wedge H'' \in \text{wp}_{n-1}(c', Q)\end{aligned}$$

We define the atomic space chunks $\text{chunks}(S)$ for a set S of opened atomic spaces as follows:

$$\text{chunks}(S) = \{\{\mathbf{atomic_space}(V, V') \mid (V, V') \in S\}\}$$

We say a logical heap H is *self-consistent with depth bound k* , denoted $H \text{ ok}_k$, if there exists a heap h , a ghost heap \hat{h} , an atomic spaces bag A , a set of opened atomic spaces S , and a consistent stock of lemma type chunks Σ of size at most k such that $\{\{\mathbf{heap}(h)\}\} + h + \hat{h} + \text{chunks}(A) + \text{chunks}(S) + \{\{\mathbf{atomic_spaces}(S)\}\} + \Sigma \geq \text{heap}(A) + H$, where a heap is interpreted as a set of \mapsto chunks and a ghost heap is interpreted as a set of \mapsto_g chunks. Notice: if $H \text{ ok}_k$, it follows that H is strongly consistent.

Notice that $h \sim H$ if and only if $\exists k, (H + \{\{\mathbf{heap}(h), \mathbf{atomic_spaces}(\emptyset)\}\}) \text{ ok}_k$.

Lemma 1 (Soundness of inner ghost command correctness).

$$\{P\} G \{Q\} \wedge H \in P \wedge (H + H') \text{ ok}_k \Rightarrow \exists H'' \in Q. (H'' + H') \text{ ok}_k$$

Proof. By induction on k and nested induction on the size of G . The outer induction hypothesis is used to deal with lemma calls. \square

Lemma 2. *If an annotated command \hat{c} is correct with respect to precondition P and postcondition Q , then, for all n , P implies*

the weakest precondition of the erasure of \hat{c} with respect to Q for n steps:

$$\{P\} \hat{c} \{Q\} \Rightarrow \forall n. P \Rightarrow \text{wp}_n(\text{erasure}(\hat{c}), Q)$$

Proof. By induction on the derivation of the correctness judgment. The most interesting case is $\hat{c} = \mathbf{FAA}(\ell, z)$. Fix an n and a logical heap $H \in P$. Fix a heap h , a ghost heap \hat{h} , an atomic spaces bag A , a consistent stock of lemma type chunks Σ , and a frame H_F such that $h + \hat{h} + \text{chunks}(A) + \Sigma = \text{heap}(A) + H + H_F$. By $H \in P$ and H strongly consistent we can fix a g , a G and an H' such that $H = \{\{\lambda g. G : \mathbf{FAA_ghop}(\ell, z, V_{\text{pre}}, V_{\text{post}})\}\} + H'$ and $H' \models V_{\text{pre}}()$. By strong consistency of H , we have $\forall op, V_P, V_Q. \{\{\mathbf{atomic_spaces}(\emptyset) * op : \mathbf{FAA_op}(\ell, z, V_P, V_Q) * \llbracket V_P() \rrbracket * \llbracket V_{\text{pre}}() \rrbracket\}\} G[op/g] \{\{\mathbf{atomic_spaces}(\emptyset) * op : \mathbf{FAA_op}(\ell, z, V_P, V_Q) * \llbracket V_Q() \rrbracket * \llbracket V_{\text{post}}() \rrbracket\}\}$. We take $op = \lambda. \ell \leftarrow_h h(\ell) + z, V_P = \text{heap}_-(h)$, and $V_Q = \text{heap}_-(h[\ell := h(\ell) + z])$. We have that semantically, op is of type $\mathbf{FAA_op}(\ell, z, V_P, V_Q)$, so $\Sigma' = \Sigma - \{\{\lambda g. G : \mathbf{FAA_ghop}(\ell, z, V_{\text{pre}}, V_{\text{post}})\}\} + \{\{op : \mathbf{FAA_op}(\ell, z, V_P, V_Q)\}\}$ is consistent. We apply Lemma 1 to G using $H' + \{\{\mathbf{atomic_spaces}(\emptyset), op : \mathbf{FAA_op}(\ell, z, V_P, V_Q), \mathbf{heap}(h)\}\}$ for H , H_F for H' and the size of Σ' for k to obtain that there exists an $H'' \in \llbracket V_{\text{post}}() \rrbracket$ such that $(H'' + \{\{\mathbf{atomic_spaces}(\emptyset), op : \mathbf{FAA_op}(\ell, z, V_P, V_Q), \mathbf{heap}(h[\ell := h(\ell) + z])\}\} + H_F) \text{ ok}_k$ and therefore $\{\{\lambda g. G : \mathbf{FAA_ghop}(\ell, z, V_{\text{pre}}, V_{\text{post}})\}\} + H'' \in Q$ and $h[\ell := h(\ell) + z] \sim H'' + \{\{\lambda g. G : \mathbf{FAA_ghop}(\ell, z, V_{\text{pre}}, V_{\text{post}})\}\} + H_F$. \square

Lemma 3. *If $h \sim H$ and $H \in \text{wp}_n(c, \text{True})$, then any configuration reached by (h, c) in at most n steps is okay.*

Proof. By induction on n . \square

Theorem 1. *If an annotated program \hat{c} is correct, then its erasure $\text{erasure}(\hat{c})$ is safe.*

Proof. We first apply Lemma 2. Then, since the empty heap is consistent with the empty logical heap, we can finish the proof by applying Lemma 3 with $H = \emptyset$. \square

6. Related work and discussion

Iris In contrast to the state-of-the-art logic for fine-grained concurrency verification Iris (Jung et al. 2015, 2018), the presented logic does not require the *later* modality. This is because atomic space invariants are stored in the logical heap in a *syntactic* form, rather than as propositions over logical heaps. As a result, no recursive domain equations are involved.

A downside of our approach compared to Iris, however, is that our logic does not directly support separating implications (a.k.a. magic wands), viewshifts, or other logical connectives in which some operand assertions appear in *non-positive* positions, i.e. whose truth is not monotonic in the truth of some of the operand assertions. This is because we define the meaning of predicate values using a least fixpoint construction.

We recover the functionality of separating implications and viewshifts to some extent by means of lemma values, with the major limitation that lemma type assertions are *linear*, which makes them more awkward to work with than the Iris constructs, although in practice this has not hindered us significantly so far; in fact, while we do vaguely remember encountering cases where this was inconvenient (or worse), we have trouble recalling the specific circumstances.

Having said that, we use VeriFast as a tool for verifying particular programs, not for metatheory development. It is very likely that the limitations of our logic would become prohibitive if we attempted to replicate deep metatheory developments such as RustBelt’s lifetime logic (Jung 2020) in VeriFast. We do, however, make use of the *results* of such developments in VeriFast, through axiomatisation. The soundness of such axiomatisations, however, is a nontrivial question. While our axiomatisation of the lifetime logic *appears* sound, it is future work to build a formal argument of that, perhaps by connecting a Rocq mechanisation of the development of the present paper with that of the lifetime logic. Very important related work in this regard is Nola, which proposes a version of the lifetime logic, machine-checked in Rocq, without the later modality (see below).

Nola, Lilo VeriFast’s logic is by no means the only later-less logic for fine-grained concurrency. Nola (Matsushita 2023; Matsushita & Tsukada 2025) is an Iris library that generalizes Iris’s support for invariants (the Iris construct analogous to our atomic spaces) by parameterizing it over the type Fml of *formulas* that describe the contents of invariants, and the *semantics* $\llbracket \cdot \rrbracket$ of *formulas*, that maps a formula to an Iris proposition. Classical Iris invariants are obtained by taking $Fml = \blacktriangleright iProp$, the type of Iris propositions but guarded by the later functor, and $\llbracket next(P) \rrbracket = \triangleright P$. However, an alternative is to take as Fml a type of *syntactic* separation logic formulae. For most syntactic constructs, the semantics can be defined without the need for the later modality, thus enabling later-less invariant reasoning for a wide class of invariants. Matsushita & Tsukada (2025) were even able to implement a later-less version of RustBelt’s lifetime logic this way. Furthermore, they show how to handle an even wider class of invariants without later by applying *stratification*, where formulae at a layer $k + 1$ can quantify over formulae at layer k . Lilo (Lee et al. 2025) applies Nola’s idea of stratification to enable later-less invariant reasoning in a logic for verifying termination of busy-waiting programs under fair scheduling. Lilo was used to build the first modular total correctness proof of an elimination stack. Nola and Lilo are fully mechanized in Rocq.

Acknowledgments

We thank Justus Fasse for proofreading. This research is partially funded by the Research Fund KU Leuven, and by the Cybersecurity Research Program Flanders.

References

Jacobs, B. (2016). Partial solutions to VerifyThis 2016 challenges 2 and 3 with VeriFast. In V. Klebanov (Ed.), *Pro-*

ceedings of the 18th Workshop on Formal Techniques for Java-like Programs, FTfJP@ECOOP 2016, Rome, Italy, July 17-22, 2016 (p. 7). ACM. Retrieved from <http://dl.acm.org/citation.cfm?id=2955818>

Jacobs, B., & Piessens, F. (2011). Expressive modular fine-grained concurrency specification. In T. Ball & M. Sagiv (Eds.), *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011* (pp. 271–282). ACM. Retrieved from <https://doi.org/10.1145/1926385.1926417> doi: 10.1145/1926385.1926417

Jacobs, B., Smans, J., & Piessens, F. (2015). Solving the VerifyThis 2012 challenges with VeriFast. *Int. J. Softw. Tools Technol. Transf.*, 17(6), 659–676. Retrieved from <https://doi.org/10.1007/s10009-014-0310-9> doi: 10.1007/S10009-014-0310-9

Jung, R. (2020). *Understanding and evolving the Rust programming language* (Doctoral dissertation, Saarland University, Saarbrücken, Germany). Retrieved from <https://publikationen.sulb.uni-saarland.de/handle/20.500.11880/29647>

Jung, R., Krebbers, R., Jourdan, J., Bizjak, A., Birkedal, L., & Dreyer, D. (2018). Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *J. Funct. Program.*, 28, e20. Retrieved from <https://doi.org/10.1017/S0956796818000151> doi: 10.1017/S0956796818000151

Jung, R., Swasey, D., Sieczkowski, F., Svendsen, K., Turon, A., Birkedal, L., & Dreyer, D. (2015). Iris: Monoids and invariants as an orthogonal basis for concurrent reasoning. In S. K. Rajamani & D. Walker (Eds.), *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015* (pp. 637–650). ACM. Retrieved from <https://doi.org/10.1145/2676726.2676980> doi: 10.1145/2676726.2676980

Lee, D., Lee, J., Yoon, T., Cho, M., Kang, J., & Hur, C.-K. (2025, April). Lilo: A higher-order, relational concurrent separation logic for liveness. *Proc. ACM Program. Lang.*, 9(OOPSLA1). Retrieved from <https://doi.org/10.1145/3720525> doi: 10.1145/3720525

Matsushita, Y. (2023). *Non-step-indexed separation logic with invariants and Rust-style borrows* (Doctoral dissertation, University of Tokyo). doi: 10.15083/0002013242

Matsushita, Y., & Tsukada, T. (2025, June). Nola: Later-free ghost state for verifying termination in Iris. *Proc. ACM Program. Lang.*, 9(PLDI). Retrieved from <https://doi.org/10.1145/3729250> doi: 10.1145/3729250

O’Hearn, P. W., Reynolds, J. C., & Yang, H. (2001). Local reasoning about programs that alter data structures. In L. Fribourg (Ed.), *Computer Science Logic, 15th International Workshop, CSL 2001. 10th Annual Conference of the EACSL, Paris, France, September 10-13, 2001, Proceedings* (Vol. 2142, pp. 1–19). Springer. Retrieved from https://doi.org/10.1007/3-540-44802-0_1 doi: 10.1007/3-540-44802-0_1

Reynolds, J. C. (2002). Separation logic: A logic for shared mutable data structures. In *17th IEEE Symposium on Logic in Computer Science (LICS 2002), 22-25 July 2002, Copenhagen, Denmark, Proceedings* (pp. 55–74). IEEE Computer

Society. Retrieved from <https://doi.org/10.1109/LICS.2002.1029817> doi: 10.1109/LICS.2002.1029817
Vogels, F., Jacobs, B., & Piessens, F. (2015). Featherweight VeriFast. *Log. Methods Comput. Sci.*, 11(3). Retrieved from [https://doi.org/10.2168/LMCS-11\(3:19\)2015](https://doi.org/10.2168/LMCS-11(3:19)2015) doi: 10.2168/LMCS-11(3:19)2015

About the author

Bart Jacobs is an associate professor at the DistriNet research group at the department of Computer Science at KU Leuven (Belgium). His main research interest is in modular formal verification of concurrent programs. You can contact the author at bart.jacobs@kuleuven.be or visit <https://distrinet.cs.kuleuven.be/people/BartJacobs/>.